N° d'ordre: 1234

THÈSE

Présentée pour obtenir

LE GRADE DE DOCTEUR EN SCIENCES DE
L'UNIVERSITÉ PARIS-SUD XI

Spécialité: Informatique

par

Pierre YGER

Chaos control and plasticity in large scale neuronal networks with ongoing activity

Soutenue le 13 Octobre 2010 devant la Commission d'examen:

M.	Andrew DAVISON	
Mme.	Hélène PAUGAM-MOISY	
M.	Yves FRÉGNAC	(Directeur de thèse)
M.	Philippe TARROUX	(Président du jury)
M.	Gérard BERRY	

Rapporteurs:

M.	Frédéric Alexandre
M.	Wulfram Gerstner

Thèse préparée au laboratoire
UNIC
Unité de Neurosciences Intégratives et Computationnelles (UPR 3293)
CNRS, 1 Avenue de la Terrasse
91 198 Gif sur Yvette

CHAOS CONTROL AND PLASTICITY IN LARGE SCALE NEURONAL NETWORKS WITH ONGOING ACTIVITY

Résumé

Comprendre comment s'organisent et fonctionnent les aires sensorielles primaires du neocortex est une étape cruciale pour l'analyse des mécanismes sous tendant le fonctionnement, d'un point de vue algorithmique, de l'activité cérébrale. Cette compréhension de la dynamique sensorielle à grande échelle passe par l'utilisation de modèles de neurones simplifiés, du type "intègre et tire", et par un cadre de travail particulier, celui du réseau "balancé", permettant de se placer dans un régime proche de ceux observés dans l'activité spontanée *in vivo*, où les neurones déchargent de façon irrégulière et à relativement basse fréquence.

La première partie de cette thèse montre, par des simulations grande échelle et l'utilisation de patterns de stimulations ayant des statistiques proches de l'activité spontanée passée, que des réseaux de neurones dans des régimes irréguliers et asynchrones comme le cortex visuel *in vivo* pourrait fonctionner au bord d'un régime dynamique particulier, le chaos déterministe. Alors que l'activité spontanée de ces réseaux est souvent considérée comme du bruit, la structure particulière de cette activité issue de la connectivité récurrente peut avoir un rôle particulier pour la dynamique fonctionnelle du réseau. Cette étude permet de mieux comprendre la propagation d'information dans le contexte de l'activité persistante des réseaux récurrents. Par ailleurs, une étude menée en collaboration sur des données intracellulaires *in vitro* et *in vivo* permet de mieux comprendre comment le niveau global de corrélation au sein du réseau peut être observé via l'analyse du spectre de puissance de la trace du potentiel membranaire d'un neurone. La similarité entre le niveau de corrélation au repos (en activité spontanée), et lorsque le neurone traite une scène visuelle avec des statistiques naturelles nous permettra de conforter le lien entre activité spontanée et activité évoquée, et la possibilité que l'activité spontanée rejoue des patterns sensoriels mémorisés au cours de l'interaction avec l'environnement.

L'analyse analytique et exacte de la dynamique de tels réseaux de neurones, dits balancés, dans des régimes asynchrones irréguliers se révèle impraticable pour des réseaux construits avec des paramètres biophysiques : lorsque la connectivité devient plus structurée, que des vitesses de propagation finies sont prises en considération, la zoologie des dynamique observables s'apprécie principalement via des explorations numériques. Une partie de mon travail s'est donc naturellement portée sur l'étude et la compréhension des paramètres clés gouvernants la dynamique de réseaux topographiques, plus particulièrement en ce qui concerne la distribution des corrélations inter-neurones en fonction de la distance. Plusieurs évidences montrent en effet que ces corrélations entre décharges neuronales peuvent refléter la nature du codage de l'information au sein du cerveau, rendant important de comprendre comment ces dernières s'organisent et quels sont les paramètres les contrôlant. En s'intéressant plus particulièrement au cas des régimes synchrones réguliers, où ces corrélations sont plus marquées, nous montrerons que ces corrélations sont principalement imposées par la balance excitation/inhibition, et comment elles peuvent être modulées via une stimulation extérieure.

Pour essayer d'enseigner des statistiques ou motifs particuliers au réseau qui seront reflétées ultérieurement dans son activité spontanée, dans le cadre d'un apprentissage non supervisé, la connectivité au sein du réseau doit s'auto-organiser en fonction de l'activité. Un candidat idéal pour assurer ces modifications est une règle de plasticité dépendante du temps précis d'arrivée des potentiels d'action, appelée Spike Timing Dependent Plasticity (STDP). Cette règle (Markram and Tsodyks, 1996, Bi and Poo, 1998), observée principalement *in vitro*, montre que si deux neurones ont une décharge causale (le neurone pré-synaptique étant actif juste avant le neurone post-synaptique), la synapse liant ces deux neurones est renforcée : l'information du neurone pré-synaptique a été utile et pertinente a la décharge du neurone post. Inversement, si les décharges sont acausales, a savoir si pre est actif après post, la synapse se trouve diminuée. Cette règle, intéressante dans sa formulation la plus simple, souffre néanmoins de nombreux problèmes lorsque utilisée dans des réseaux large échelle. Une nouvelle règle de plasticité synaptique sera donc explorée, incluant des contraintes d'homéostasie et de méta plasticité. Cette hypothèse permet de réconcilier différents schémas de plasticités théoriques par des mécanismes biophysiques réalistes. Son intérêt fonctionnel est de permettre au même réseau corti-

cal de traiter de manière fiable l'information sensorielle tout en engrammant les associations causales nouvelles crées par l'interaction avec l'environnement sensoriel.

Toutes les simulations de cette thèse ont été faites en prenant en compte l'hétérogénéité des différents outils de simulation disponibles. Dans le but d'unifier tous ces outils, un travail d'homogénéisation a été effectué pour simplifier la confrontation des résultats obtenus via différents outils de simulation. Plus spécifiquement, une interface de programmation en Python, PyNN, a été écrite et développée pour permettre la définition d'un langage commun à différents logiciels de simulations. Cette unification est importante pour la structuration du domaine des neurosciences computationnelles, pour augmenter la confiance dans des simulations complexes où la certitude que le résultat ne dépend pas des méthodes de simulations considérées doit être écartée.

Mots-clefs : Réseaux de neurone, traitement de l'information, plasticité, outils de simulations.

Abstract

Understanding how the primary sensory areas of the neocortex are structured in order to process sensory inputs is a crucial step in analysing the mechanisms underlying the functional role, from an algorithmic point of view, of cerebral activity. This understanding of the sensory dynamics, at a large scale level, implies using simplified models of neurons, such as the "integrate-and-fire models", and a particular framework, the "balanced" network, which allows the recreation of dynamical regimes of conductances close to those observed *in vivo*, in which neurons spike at low rates and with an irregular discharge.

The first part of this thesis shows, using large scale simulations and particular patterns of stimulation whose statistics are close to those of the spontaneous activity, that such neuronal networks in asynchronous and irregular firing states, such as the primary visual cortex *in vivo* can operate at the border of a particular dynamical regime, deterministic chaos. While the ongoing activity of those networks is often considered as noise, the particular structure of this activity, emerging from the recurrent connectivity, could have a functional role in information processing. This study allows to better understand information transmission within the context of persistent activity, inherent to those networks. In the meanwhile, a study led in collaboration on intra-cellular recordings *in vivo* and *in vitro* allows to have a better insight on how the global level of correlation received by a neuron can be observed by analysis of the power spectrum of the membrane potential recorded intracelullarly. The similarity between the level of correlation in the resting state (in spontaneous activity, while no inputs are presented to the system) and when the network processes information coming from a visual scene with natural statistics will enforce the link between spontaneous and evoked activities, and the possibility that ongoing activity replays sensory patterns stored during the interaction with the environment.

An exact mathematical analysis of the dynamics in such balanced neuronal networks, in asynchronous and irregular regimes, is hardly tractable when networks are built with biophysical parameters. Particularly, when connectivity is not random but structured and when finite axonal propagation speeds are taken into account, the diversity of the dynamics that can be observed has to be studied with numerical simulations. Part of my work was then focused on the study and the understanding of the key parameters governing the behaviour of such topographical neuronal networks, more particularly regarding the distribution of the correlations between pairs of neurons as a function of distance. Since several pieces of biological evidence show that these pairwise correlations may be an important part of the neural code in the brain, it is important to know how they emerge, and what parameters can influence their distributions. By focusing more particularly on the synchronous regular regimes, where correlations are more pronounced, we have shown that they are mainly driven by the balance between excitation and inhibition, and how they can be modulated by external stimulations.

To try to store some particular statistics or patterns to a neuronal network, and to see if a trace of this learning is kept in the ongoing activity in the framework of unsupervised learning, the connectivity within the neuronal network must evolve as a function of the activity. An recent candidate to achieve such unsupervised changes is a plasticity rule depending on the precise arrival time of the actions

potentials, called spike timing dependent plasticity (STDP). This rule (Markram and Tsodyks, 1996, Bi and Poo, 1998), observed mainly *in vitro*, shows that if two neurons have a causal discharge (the pre-synaptic neuron firing just before the post-synaptic one), the synapse between them is reinforced: the information of the pre-synaptic neuron has been useful and pertinent for the post-synaptic one. On the contrary, if the discharges are not causal, meaning if pre is active after post, the synapse is weakened. This rule, interesting in its general formulation, suffers from a lot of problems when used in recurrent large-scale networks. A new rule of plasticity was therefore explored, incorporating some homeostatic constraints in the framework of metaplasticity. This hypothesis can link different theoretical schemes of plasticity by the definition of a biophysically realistic mechanism. The functional interest is to be able to process incoming sensory informations in a cortical network, and at the same time, to store the causal associations generated by the interaction with the sensory environment.

All the simulations in this thesis were made while taking into account the heterogeneity of the simulation tools available. With the target goal of unifying all these tools, a work has been performed to simplify the comparison between results obtained with various simulators. More precisely, an application programming interface in Python, PyNN, has been designed and implemented to allow the definition of a common language for large-scale neuronal network simulations. This unification process is important and valuable for the consolidation of the neuroscience community, to gain confidence in the results of complex simulations, where the certainty that the final results do not depend on the simulations strategies considered should be banished.

Keywords : Neuronal networks, information processing, plasticity, simulation tools.

Remerciements

Des mots seront peu de choses pour toutes les personnes qui m'ont aidé durant ces quatre années de transition, d'immersion dans les neurosciences computationelles et dans le monde fascinant de la recherche et du cerveau. L'abandon de la carte étudiant et du diplôme d'ingénieur pour un aller simple dans l'inconnu dont je ne suis à priori encore sorti. Même si ce sera forcément imparfait et incomplet, je voudrais quand même dire un petit mot pour toutes les personnes qui m'ont aidé, guidé, soutenu. Pardon d'avance pour les oublis et les imperfections.

Tout d'abord, merci Yves, pour ton énergie : celle avec laquelle tu portes le laboratoire et tes idées. Ta solide confiance m'a été précieuse, de même que tes conseils scientifiques, ta connaissance pointue et rigoureuse de la bibliographie et tes conseils avisés.
Andrew, pour ton professionnalisme, ta dévotion contagieuse au culte de Python, ta confiance dans les projets que j'ai menés. Dire que j'ai progressé en anglais à tes cotés serait surement mentir, tu as surement plus progressé en français avec moi.
Alain, pour ta constance et ton enthousiasme communicatif sur les projets. Un certain don pour positiver les choses, donner envie de creuser et de porter les idées à maturation, quand dehors parfois il fait doute.
Markus et Abigail, pour m'avoir enseigné NEST et les joies de la plasticité synaptique, pour les bons moments passés au Japon.
Merci à Philippe Tarroux de m'avoir donné un pieds d'appui au LIMSI, de bien avoir voulu accepter d'être mon directeur de thèse.
Un petit mot aussi à Luc Bougé, pour m'avoir enrolé à l'ENS et fait confiance, et sans qui finalement tout aurait été plus difficile.
Merci enfin à mes rapporteurs, Wulfram Gerstner et Frédéric Alexandre d'avoir accepté de prendre le temps de relire et de critiquer ce manuscrit.
Merci aussi a Gérard Berry d'avoir relu ce manuscrit et accepté de participer à mon jury.

Pour continuer sur mes collègues de bureau, merci Olivier, maintenant en train de saccager des rétines dans des lointaines contrées, pour ton énergie, ta soif de questions, ton enthousiasme. Travailler avec toi fut un plaisir, et la meilleure entrée en matière pour moi.
Sami, pour ta rigueur, ta motivation communicative, et pour tous les bons fous rires passés ensemble alors que des codes ne marchaient pas, que des équations patinaient, voire qu'il pleuvait (!) sur les clusters de calcul, un grand merci.
Charlotte, pour avoir toujours été là, disponible (sauf sur Gmail), pour un café ou pour parler, pour ta patience et ton écoute, pour tes "c'est quelle heure" que j'entendrai encore longtemps, merci.
Julie, pour m'avoir supporté dans toutes les circonstances, pour avoir été là quand il fallait que tu épargnes quelques rats, pour m'avoir donné des conseils que je n'ai jamais suivi.
Michelle, pour des discussions intarissable et sans cesse surprenante, Luc, qui est tout sauf prétentieux et arrogant, n'en déplaise à d'autres, Pierre-Jean, qui a parfois tort de penser avoir raison, David, qui m'a appris de nuit à jongler avec les balles et les idées et Romain, fidèle d'Oberkampf et du curry, et puis merci à tous les autres : Nazied, Manuel, Dan, Pedro, Valérie, Guillaume, Paul, la liste est imparfaite.

Et comme il y a une vie en dehors du travail, parfois, je tenais à remercier ceux qui m'ont soutenu et encouragé dans cette démarche et dans ces recherches, parfois sans le savoir. Une

question, une curiosité, un rien suffit pour se remettre à l'ouvrage. Mon père et ses maths, ma mère et sa confiance, ma soeur et sa constance. Mes amis de toujours, qu'ils soient du collège, du lycée, de prépa ou d'école, toujours partants et présents. Et puis les amies, celles qui ont comptées, qui m'ont toujours soutenu, et qui parfois sont parties.

Contents

I Materials and Methods **17**

1 Diving into a cortical network: the biology **18**

1.1 Overview of a neuron 18

1.2 The synapse 19

1.3 Neuronal networks: a canonical structure 20

2 Computational models: the building blocks **23**

2.1 Compartmental models 23

2.2 The Integrate-and-Fire model 23

2.3 Conductance based synapses 25

2.4 Mean field models 25

3 Network activity **26**

3.1 Neuronal networks 26

3.2 The balanced random network 26

4 Coding in neuronal networks **27**

4.1 The quest for the code 27

4.2 An overview of some temporal codes 29

II On the ongoing activity *in vivo* **32**

5 Introduction **33**

6 The spontaneous activity **34**

6.1 Origin of the ongoing activity 34

6.2 Nature and structure *in vivo* 35

7 Similarity between evoked and ongoing activity **37**

7.1 Evidence from the literature 37

7.2 Main results 38

8 Information transmission under ongoing activity **58**

8.1 Ongoing activity and information propagation 58

8.2 Reliability as a function of the stimulus 59

8.3 Main results 61

9 Discussion **73**

9.1 A general paradigm 73

9.2 On the role of the recurrent connections 76

III Structure of the ongoing activity 81

10 Introduction **82**

10.1 Toward more realistic networks 82

10.2 A link with macroscopic measures 82

11 Macroscopic invariants in topological networks **84**

11.1 Main results 84

11.2 Supplementary figures 106

12 Discussion **110**

12.1 General comments 110

12.2 Toward a mean-field-based description of topographical networks 112

IV Plasticity in neuronal networks 116

13 Introduction **117**

13.1 Supervised learning 118

13.2 Unsupervised learning 119

14 Hebbian rules and spike based plasticity **122**

14.1 Refinement of the Hebbian rule 122

14.2 Spike based plasticity and spike timing dependent plasticity 123

14.3 Classical model of STDP 125

15 A memory retention problem **127**

15.1 Materials and Methods 128

15.2 Results on the dynamics 130

15.3 Results on axonal versus dendritic delays 134

15.4 Emergence of structures 136

16 Discussion **137**

V Towards a new STDP rule **140**

17 A more complex view of STDP **141**
17.1 STDP as an epiphenomenon ? 141
17.2 Metaplasticity 145

18 STDP and metaplasticity **146**
18.1 Main results 146
18.2 Supplementary materials 162

19 Discussion **169**

VI The simulation workflow **172**

20 Introduction **173**
20.1 Neuronal network simulators 174
20.2 The simulations tools 174
20.3 A complex trade-off 176

21 PyNN: A Common Interface for Neuronal Network Simulators. **181**
21.1 Main results 181

22 The analysis workflow in neuroscience **192**
22.1 Heterogeneity of the data 192
22.2 On the analysis tools 192
22.3 The NeuroTools project 192

23 Discussion **194**

VII Discussion & Further work **196**

References **201**

Introduction

"Les hommes se croient libres parce qu'ils ont conscience de leurs volitions et de leur appétit, et qu'ils ne pensent pas, même en rêve, aux causes qui les disposent à désirer et à vouloir, parce qu'ils les ignorent."
Spinoza, Ethique, I, Appendice, 1677

A living organism can be seen as a response to the environment, provided through a complex and uncontrolled evolutionary process. Accorded to the theory of Darwin (1859), every organism is fighting for its own survival, and this permanent struggle progressively shapes organisms able to react, adapt, and reproduce themselves, taking into account all their surroundings. Nothing is constant in the world, except these changes.

Whether behaviours are hard-wired or resulting from internal and computations performed by the brain, organisms need to have fast and appropriate responses to well-defined situations, and the slow evolutionary forces gradually selecting species push for organisms more and more suited to such fast responses, through complex and hidden constraints imposed by the external world. Only those able to adapt are kept in the next generation, and adaptation requires and imposes an understanding and an integration of the environment. In that evolutionary spirit of better capturing the essence of the surrounding world, the central nervous system has been turned into a prominent integration center, gathering every inputs signal received from the sensory modalities in order to, provide appropriate responses, as fast as possible. It can be a spoken word, a quick jump, an uncontrolled feeling, but it is always the result of an ongoing computation, performed by the brain constantly impinged upon by sensory inputs.

Indeed, the brain is an efficient computational device for integrating multi-modal information coming from the sensory world, and turning them into behaviour. Computational neuroscience is the attempt to understand, with equations, how the brain processes and deals with information. Without aiming to answer the philosophical question "Who am I ?", it is a pragmatic attempt to reveal the basis of brain computations, and to understand the basic principles of neuronal dynamics.

Computational neuroscience

The goal of computational neuroscience is to understand the underlying properties of the information processing tasks performed by the brain. The term "computational", which comes from the Latin word *computare*, refers to logically and mathematically well-defined computations, and therefore sets the formal background of the field, at the intersection of mathematics, computer science, theoretical physics and biology. Computational neuroscience aims to dissect and observe, from an algorithmic point of view and based on biological observations and recordings, the activity of the brain. How the brain encodes information and how a group of electrical impulses is converted into an action or a thought is a fascinating question, leading and embracing various fields such as philosophy, artificial intelligence, robotics, computer science and ethics.

The motivations of computational neuroscience are numerous. Setting aside philosophical questions and the question of consciousness, a better comprehension of the neuronal code and dynamics can tremendously help and guide certain medical treatments and prostheses. Popular success, which are direct consequences of the understanding of the brain, are: cochlear implants, neuronal prostheses that can recover the loss of a member such as an arm, deep brain

stimulation to treat chronic pain, Parkinson's disease or preventing epileptic crises, brain machine interfaces (BMI) allowing the to control, by on-line processing of brain activity recorded through various sensors, a third apparatus that can be either an electronic arm or a mouse cursor on a screen. In the case of paraplegia or lock-in syndrome, brain can be fully functional and operational, but the nerves may not be able to transmit motor commands. Moreover, gaining a clear understanding of the brain is fundamental to design more efficient drugs and treatments. A lot of untreatable diseases nowadays involve the brain, nerves or neurons, since the brain is the most complex organ in the human body. Treatment of autism, schizophrenia, depression and many more, all those neuronal disorder may gain from a better understanding of brain activity.

The classical method commonly used to analyse the brain, in order to deal with such a complex system, is to try to isolate all its key elements and, knowing how every single small building block works, to try to infer the global activity by putting together all these individual behaviours. This generic method is not restricted to computational neuroscience, but in our case, for the brain, one needs to focus on elements such as neurons, axons, dendrites, synapses, glia, neurotransmitters, neuromodulators, ion channels, and so on, and all of these need to be understood and properly modelled in order to have a clear insight about their roles and their effects. Capturing their fundamental characteristics is achieved by the design of biophysical models, reproducing experimental results obtained with experiments *in vivo* or *in vitro*. Those models are an attempt to turn a physical element or a physical system into variables and equations reflecting biologically observed quantities. The main biophysical models that underlie this manuscript will be briefly reviewed later in Part I. These will be our essential building blocks. The only problem is that such a bottom up approach is not possible in practice. In theory, it could be achieved, but theory is too often far from reality. Being able to have a clear understanding of all these individual subunits would require to be able to separate them in a clean manner, to suppose that they can be understood independently, that interactions between them are linear, and so on. Knowing what is a sand grain and what happens with a small pile of sand grains does not explain the complex behaviour of a sand dune. Non linear interactions and complex properties may emerge when such individual building blocks are gathered, and it is a challenge to get a clear idea about where the key interactions lie. From this comes the need for theory and predictions.

The converse approach, top to bottom, can be seen as a black-box approach, also widely used in computational neuroscience. Since having access to detailed information about all the individual sub-blocks of the system is near impossible, and even if it were not, since knowing the non-linearities of their interactions is certainly impossible, one can present inputs to the system, measure its responses, and try to infer what the mechanisms and the interactions between the sub-blocks should be, in order to explain its general response. One drawback of this method is that the state of the system itself can bias the observations that are performed. If one thinks about *in vivo* experiments, when brain activity is recorded, these recordings are often made under anaesthesia, or at least in experimental conditions that may be far from the real working conditions of the system. This unknown bias is the price to pay, and what needs to always be kept in mind is the distinct hypothesis that has been made to get an intuition about how this tremendously complex system works. This black-box approach led, in psychology, to the notion of behaviourism: the idea that organisms should be studied only through their behaviours and responses to external stimuli and the environment. Nevertheless, behaviourism has been turned nowadays into cognitive science, with the combined idea that even if the observation of the behaviour is important, one can still try to make testable infer-

	Human Brain	Blue Gene computer
Neurons/Processors	10^{12}	10^{4}
Synapses/Memory	10^{16} bytes	10^{11} bytes
Operations	10^{17} flops	10^{14} flops
Power Consumption	20 W	505000 W + cooling
Frequency	$\simeq$ 5 Hz	2 GHz
Parallel scheme	Asynchronous	Synchronous

Figure 1: Comparison between some key functional properties of the brain and a modern supercomputer, such as the Blue Gene.

ences about brain processes underlying it. An infinite set of responses can still be generated by a finite number or combinations of mental processes and computational neuroscience aims at revealing them by a reverse engineering approach.

The brain as a computing device

The brain is a computational device made of neurons able to efficiently process the huge flow of sensory information that is constantly impinging upon organisms, in order to extract relevant information and produce an adequate behavioural action as a response. The parallel between binary signals that have been conceived in binary logic, at the source of modern computer science, and the all-or-nothing behaviour of a neuronal spike is appealing for theoretical consideration, and since we are now entering a world where computer science and information technology is more and more pervasive, every optimisation that could be gained from the observation of the brain's behaviour is valuable. To focus only on the visual system, if one considers a rather simple and low-level task, performed by almost every mammal, such as image categorization, the brain is able, in less than 200 ms, to discriminate a particular item in an image. This is the so-called go or no-go experiment reported in Thorpe et al. (1996) in monkeys. Detecting a face in a crowd, segmenting an image, detecting a salient point in a visual scene, all of these algorithmic tasks are performed almost instantaneously, in an effortless manner, by the visual system while it is still a challenge for real computers. Edge detection under several luminosity levels, or dealing with partial or complete masking objects, requires complex algorithms and a lot of external knowledge to have a result which is still slower and less efficient that the one given by the brain. One has only to consider in Table 1 some similarities and differences between a Blue Gene computer and a human brain to measure the gap between these two philosophies of computation.

The brain is a massively asynchronous and low-frequency device, while computer science has had the tendency to consider sequential computations, and to work at higher and higher frequencies in order to make them faster. This second strategy is currently showing its limitations, both in Moore's law Moore (1998) and in the energy consumption and the cooling problem. Computer science is more and more interested in reducing power consumption, and in designing multi-core computations that can be launched in parallel, in order to break the sequential framework of the von Neumann machine. The new design of what will be the future of computer science should be made by being as close as possible to what has already been conceived by evolution. Moving to less energy consuming cores working at lower frequencies, but exchanging information in a more efficient way, is an attempt to go toward neuromorphic

architectures of computation, closer to the brain.

Since the first cartography of brain functions performed by Brodman in 1909, who noticed that the brain cortex was organized in distinct areas, each of them responsible for some specialized coding functions, a particular interest has been devoted to the so called "primary" sensory areas. Dissecting the brain and trying to link its shape and/or activity to behaviour is an active field of research and after a lot of progress achieved in anatomy and microscopy, two world wars and a better understanding of mental diseases (based on a lot of case studies with lesions and brain pathology), we have nowadays a fairly good picture of these primary sensory cortical areas which may be sufficient to pretend that a clear understanding of their functional role is possible. These areas (vision, somatosensory, hearing, taste, olfaction) are the first entrance stages of the sensory inputs in the neocortex. Schematically, sensory inputs gathered by sensors and converted, by transduction mechanisms, from the physical world to electrical impulsions are sent, after a relay in the thalamus to the cortex, in dedicated areas that process and broadcast their results to higher areas. These primary sensory areas are therefore very good candidates to understand the basis operations performed by neurons, and how the transduction of an input signal affects the dynamics of a population of neurons. As a first processing stage, they are supposed to form a simple basis where information is decomposed and encoded, before being reliably transmitted to higher order areas that will, based on those low-level computations, gather, merge, decide and act. Throughout this manuscript, results will be mainly explored by having in mind the visual system, but are intended to be as generic as possible.

Dynamics of the primary sensory cortical areas

By focussing on generic models of primary sensory areas, that will be explained more in detail, we will try to understand and dissect some key mechanisms governing the dynamics of these neuronal networks. To understand how they can produce computations and encode information, one need to understand the nature of their activity, and how and why, in return, activity may shape the structure of the network itself. Organization of sensory cortical areas is indeed linked to the structure of the stimuli they process. In the visual system, for example, information is retinotopically organized, and neurons in the primary visual area (V1) are sensitive to the primitives of the visual scene: oriented lines (Hubel and Wiesel, 1959). Simple cells sharing similar orientation preferences tend, for example to be more interconnected than others, and the wiring principles establishing those connections are thought to be the result of the joint activity of the cells. In this manuscript, we will try to study some current views about the ongoing activity in generic networks. As we will see, even without any external input, neuronal networks within the brain are constantly active and, unlike in digital computer science, information processing is performed in the presence of noise and variability. The robustness of the brain in low level processing tasks is striking when one takes into account this "noise", and so is the learning capacity in this stochastic state. The brain's capacity to produce reliable and reproducible responses despite its ongoing and fluctuating activity, which will be developed in Part II, is a challenging and interesting problem. In computer science, computers are as fault tolerant as possible, because errors in bits exchanged would automatically lead to errors in algorithms and to bugs. Error correction codes are an efficient way to deal with reliable transmission of information over unreliable communication channels, but they only ensure that the transmitted information content is the same. The resting state of a computer being well defined, if the message is the same, then the output result of the internal algorithm

will be the same. But how does the brain do this, in the presence of a stochastic resting state ?

Structure of the manuscript

In this thesis I aim to show how the similarity between the statistics of the ongoing and the evoked activity could be an advantage for the system, from an information transmission point of view, and how this similarity could be achieved by plasticity in large-scale networks of integrate-and-fire neurons.

The structure of this thesis is as follows: in the first part, I briefly review and discuss the basis and the key properties of the models that will be used in all the different sections of the manuscript. Biophysical models such as neurons or synapses and the key concepts needed to understand the main results obtained in the manuscript will be developed. In the second part, I will present a study on the similarities between spontaneous and evoked activity. After having defined the background of the spontaneous activity in the brain, and how intracellular recordings can be used to confirm the idea that evoked activity (under external stimulation) and spontaneous activity share a lot of similarities, I will present a model design to address the question of information transmission and reliability of the cortical responses in the presence of this ongoing activity. With the help of large scale simulations and the use of stimulation patterns with statistics close to those of ongoing activity, I show why asynchronous irregular networks such as the visual cortex may operate on the edge of a particular dynamical state, deterministic chaos. This study helps to better understand information propagation in the presence of ongoing activity (often considered as "noise"), a key element of recurrent networks. This particular model, compared to experimental results and which I call the Frozen Paradigm explains the difference in reliability that can be observed in the responses of neuronal networks facing several distinct stimuli.

In the third part, I present an exhaustive simulation study of two dimensional layered networks of integrate-and-fire neurons, in order to better understand what are the key parameters influencing the dynamical states of the system. How do the correlations between the spiking activities of pairs of neurons organize according to some key parameters in the model is an open and interesting, question since correlations in neuronal networks have been shown to be a crucial part of the code. We will see that macroscopic invariants can be extracted, whatever the fine details of the connectivity are.

In the fourth part, I tackle the question of unsupervised learning in recurrent neuronal networks. What is the state of the art, what are the problems that one needs to solve to obtain a stable system that can retain memory of its inputs? After having demonstrated the memory retention problems that can appear when a biological learning rule such as spike timing dependent plasticity (STDP) is implemented in large scale recurrent networks, even in topographical networks such as the one analysed in the earlier part of this manuscript, the need for a new plasticity model will be developed. In order to observe how some particular statistics can be stored by the networks and may be later replayed in the structure of the ongoing activity, I study a new rule for synaptic plasticity, based on STDP models but incorporating some homeostatic constraints in the framework of meta-plasticity. This hypothesis allows a link to be established between different theoretical schemas of plasticity that will be developed with the help of some biophysically realistic mechanisms.

Finally, I discuss the heterogeneity of the simulation tools that may be present in computational neuroscience. Since this is a growing field, without the expertise and the history

of well-defined topics such a mathematics and theoretical physics, neuronal simulation are still trying to settle their foundations. By keeping this in mind, one can observe that simulation are made with software that should be compared and cross-validated in order to be sure simulations are correct and give reproducible results. To that aim, I will present PyNN, a meta-language written in Python allowing to write code once, to run it without conversion, and to compare the results given by various neuronal network simulators.

Part I
Materials and Methods

Summary

1 Diving into a cortical network: the biology **18**
1.1 Overview of a neuron . . . 18
1.2 The synapse . . . 19
1.3 Neuronal networks: a canonical structure . . . 20

2 Computational models: the building blocks **23**
2.1 Compartmental models . . . 23
2.2 The Integrate-and-Fire model . . . 23
2.3 Conductance based synapses . . . 25
2.4 Mean field models . . . 25

3 Network activity **26**
3.1 Neuronal networks . . . 26
3.2 The balanced random network . . . 26

4 Coding in neuronal networks **27**
4.1 The quest for the code . . . 27
4.2 An overview of some temporal codes . . . 29

1 Diving into a cortical network: the biology

1.1 Overview of a neuron

Neurons are the basis of the nervous system, and the brain is made of $\simeq 10^{12}$ neurons, forming a highly connected network of very complex processing units (see Figure 2). Neurons are excitable cells able to create and propagate electrical signals (called spikes or action potentials) in an all-or-none manner, and this ability lies in a difference of electrical potential between the intra and extra-cellular media. A neuron is made of a cell body, called the soma, and two main extensions: the dendritic arbour and the axon (see Figure 4 top). Dendrites receive signals and information from other neurons, while the axon sends information to other neurons. The connection between an axon and a dendrite is a synapse, which we will describe later. Like every cell in a living organism, a neuron is a very complex machine, the details of whose function are out of the scope of this manuscript. We consider here the neuron only from the point of view of its transfer function, i.e. how it can convert a stream of electrical inputs into an output signal.

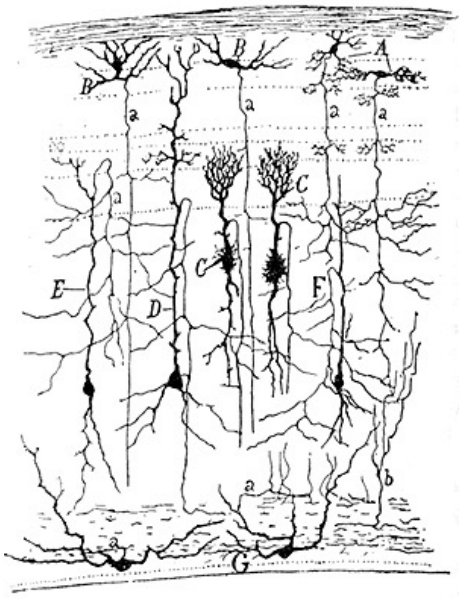

Figure 2: Hand-made drawings made around 1900 by Ramon y Cajal. Various shapes can be observed for neurons, reflecting the huge diversity that can be found. Neurons are sorted according to their shapes, their location in the brain areas, their electrical properties, and so on.

To have a better insight about this input/output transfer function, we will describe briefly how and why action potentials (also called *spikes*) emerge and propagate in neurons. For a more comprehensive overview, the reader is refered to Kandel et al. (2000). The initiation and creation of the action potential has been studied in depth in the axon of the giant squid, in the seminal work of Hodgkin and Huxley (1952), and this stereotyped electrical activity exchanged by neurons through the synapse is due to the fact that neurons are polarized cells. At rest, there is a difference of potential between the interior and the exterior of the neuron, delimited by its membrane and following the Nernst equation. This difference is due to a difference in ionic concentrations, between the two sides of the membrane, and is approximately -70mV. The genesis of the action potential is a complex sequence of ionic movements in

and out the intracellular region, which is possible because the membrane is a lipid bilayer with many diverse protein assemblages embedded in it, especially ion pumps able to let ions cross the membrane and to flow against their concentration gradient.

This sequence can be decomposed into three main steps, governed by a particular activation sequence of voltage-gated ionic channels (see Figure 3). First, as the membrane potential increases, sodium channels open and are responsible for an sudden influx of sodium in the neuron which in turn is responsible for an increase of the intracellular potential. This is the so-called depolarization step. Subsequently, potassium ions are pumped out of the neuron, and this phenomenon reduces the membrane potential. This is the so called hyper-polarization phase. Finally, the third phase is a recovery one, where pumps help to return the membrane potential to its resting state. The whole sequence, i.e. the generation of the action potential, lasts only a few milliseconds. The electrical impulse, once generated at the axon hillock, propagates in the axon and influences other neurons through synapses.

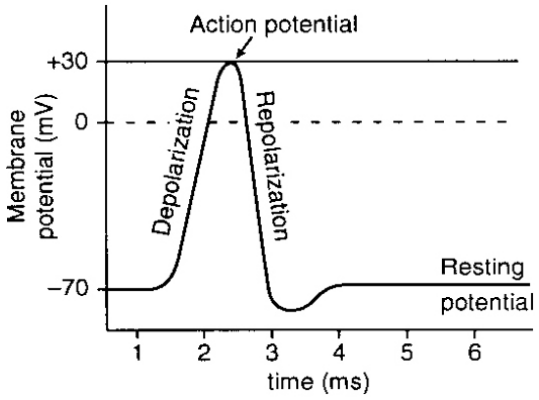

Figure 3: Illustration of an action potential. The membrane potential of the neuron increases from its resting value during a depolarization phase, reaches a peak, and then goes back during a hyperpolarization phase before recovering its resting value. This stereotyped and fast electrical waveform is an action potential.

1.2 The synapse

The synapse is a key element where the axon of a *pre-synaptic* neuron *A* connects with the dendritic arbour of a *post-synaptic* neuron *B* (see Figure 4 top). It transmit the electrical influx emitted by neuron *A* to *B*. Every cortical neuron is connected to approximately 10^4 others, so the total number of synapses in the human brain is $\simeq 10^{16}$. Synapses are crucial in shaping a network's structure, and their ability to modify their efficacy according to the activity of the pre and the post-synaptic neuron is at the origin of synaptic plasticity and memory retention in neuronal networks: see chapter IV for more details about synaptic plasticity and learning in neuronal networks. Synapses can be either chemical or electrical, but again, for a more exhaustive description, the reader should refer to Kandel et al. (2000). To focus only on the chemical synapses, the only ones that will be considered in the rest of this manuscript, the pre-synaptic neuron *A* releases a neurotransmitter into the synaptic cleft which then binds to receptors located on the surface of the post-synaptic neuron *B*, embedded in the plasma membrane. These neurotransmitters are stored in vesicles, regenerated continuously, but a too strong stimulation of the synapse may lead to a temporary lack of neurotransmitter, or

to a saturation of the post-synaptic receptors on *B*. This short-term plasticity phenomenon is called synaptic adaptation (Tsodyks et al., 2000).

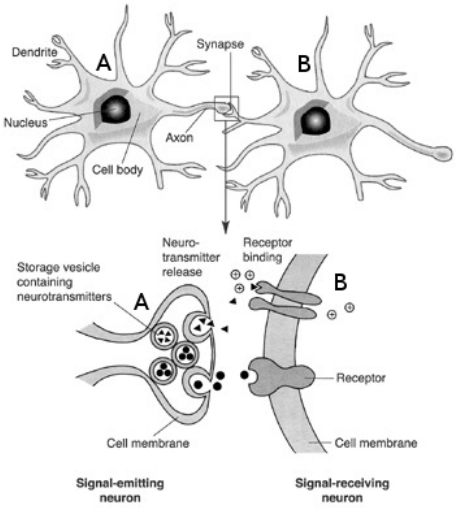

Figure 4: Top: schematic illustration of a synaptic contact between two neurons. The axon of pre-synaptic neuron *A* establishes a synapse with a dendrite of post-synaptic neuron *B*. Bottom: detail of the synaptic cleft. Neurotransmitters stored in vesicles are liberated when the pre-synaptic membrane is depolarized, and then docked onto receptors of *B*

The type of neurotransmitter which is received to the post-synaptic neuron influences its activity. The synaptic current is cancelled for a given inversion potential: if this inversion potential is below V_{thresh} (the voltage threshold for triggering an action potential), the net synaptic effect inhibits the neuron, and if it is below, it excits the cell. Classical neurotransmitter such as glutamate leads to a depolarization (i.e. an increase of the membrane potential), and the synapse is said to be *excitatory*. In contrast, gamma-aminobutyric acid (GABA) leads to an hyper-polarization (a decrease of the membrane potential), and the synapse is said to be *inhibitory*. In general, a given neuron produces only one type of neurotransmitter, being either only excitatory or only inhibitory. This principle is known as the *Dale's principle*, and is a common assumption made in the models of neuronal networks. Nevertheless, neurons in particular areas may partially contradict this principle (Sulzer and Rayport, 2000), and it has to be kept in mind that several neurotransmitters exist and impact post-synaptic electrical activity in different ways.

1.3 Neuronal networks: a canonical structure

Most of the simulations in this manuscript focus on the dynamics of generic cortical networks, as can be found in primary sensory areas, such as V1: these areas are the first entrance stages of sensory inputs in the neuronal networks that will integrate and send responses back to the thalamus. In this classical feed-forward view, where information proceeds in a hierarchical manner (from area to area), sensory signals are turned, by a transduction mechanism, into electrical signals projected into the thalamus and sub-cortical structures, before being transmitted to higher and higher cortical areas, with the underlying idea that the further we go in the cortical areas, the more complex and elaborate is the processing performed on the inputs. Detecting a particular face moving in a crowd, for example, will be achieved by neurons in the infero-temporal areas (Desimone et al., 1984), almost at the end of the visual information

processing flow (in the so called "*what*" pathway). This task is indeed very complex: it needs a raw and low level segmentation of the scene (in V1), a more complex detection of movement (in IT), and a progression along the what pathway in the infero-temporal structures (in V2, V4, IT) (see Figure 5 Left).

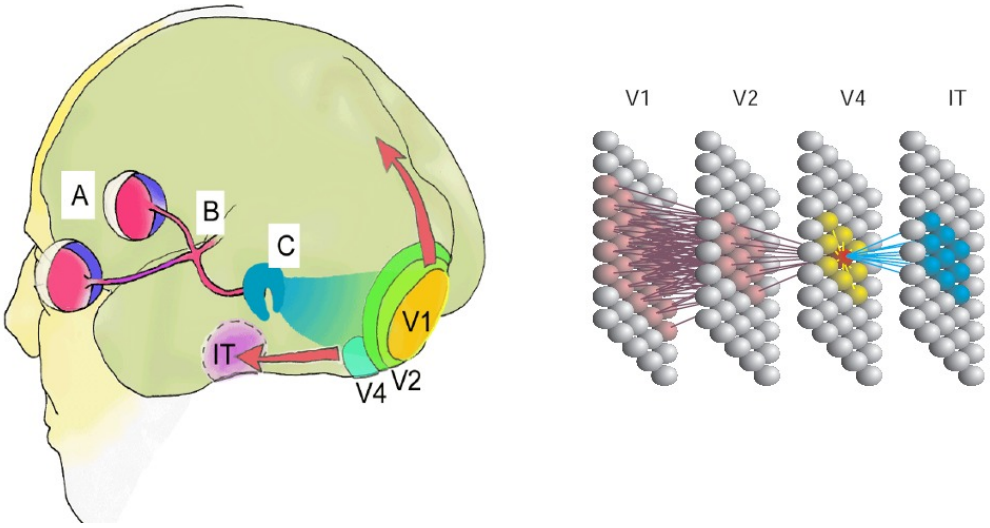

Figure 5: Adapted with permission from Lamme and Roelfsema (2000) and Perrinet (2003). Left: The visual pathway, with the succession of cortical areas where visual information is projected. Right: Representation of feedforward (pink), horizontal (yellow), and feedback (blue) connections to a hypothetical neuron in V4 (red cell).

The integration field of a neuron is composed of the synaptic afferents targeting its dendrites. Neurons integrate all this information, and emit action potentials if particular spatio-temporal patterns of activity are present or absent. As one can see in Figure 5 (Right), the integration fields, or receptive fields, of the neurons are larger and larger as we consider higher areas. Neurons in V1 have receptive fields that allow them to sample only a subpart of the visual scene, but then, by a cascade and a feed-forward structure, receptive fields in superior areas are more and more broad and span the whole visual field. This classical feed-forward view will be challenged in the following, because it has to be stressed that every layer within this straigth-forward information processing scheme (V1, V2, ...) is a higly inter-connected network of neurons. Areas are strongly coupled with higher ones sending many feedback connections to primary sensory cortical areas. For example, the connections from cortico-thalamical projections originating from V1 are ten times more numerous than the thalamo-cortical ones to V1.

Apart from this structure in areas, from an anatomical point of view, the neocortex is organized in columns and in a laminar structure formed with six distinct layers, described and drawn originally by Ramon y Cajal around 1900, see Figure 2 (for a good review of his work, see the work of Garcia-Lopez et al. (2010)) and Figure 6 for a schematic view of a cortical column. The distinction between the layers is based on anatomical observations, neuron types and densities.

In a cubic millimetre of cortex, the number of neurons is estimated to be $\simeq 100000$ (Braitenberg and Schüz, 1991, Beaulieu and Colonnier, 1983, Briones et al., 2004), and the probability of connection between any pair of neurons around $\simeq 10\%$, depending on the area and on the species. The functional role of this layered organization is far from being understood, and

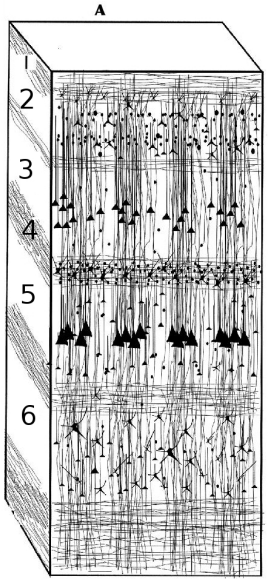

Figure 6: Schematic view of the cortical layers. Layer 4 receives the sensory inputs from the thalamus. Large pyramidal neurons have their soma in layer 4/5, and a huge dendritic arborization spawning into superficial layers (2/3). Interneurons (inhibitory cells), are more local.

a clear picture of this would be a major breakthrough in computational neuroscience. In the following sections, models are as generic as possible, and except if stated otherwise, they should be considered as models of layer 4, which is the first layer where sensory information is projected from the thalamus. Layer 4 is the entrance gate into the cortex, and, to be schematic, information is then sent to superficial layers 2/3 before going back to deep layers 5/6 and being sent back to the thalamus. Several hypotheses have been made, trying to link the columnar organization with functional units and small canonical microcircuits that could be replicated (Binzegger et al., 2004, Thomson et al., 2002), but none of them dominates. It is important to understand that the results that are obtained in the following sections are intended to be generic and easily transposable to other areas and/or cortical structures. The neuronal networks that are used, even if not biologically plausible according to their structures and their parameters, aim to capture the key dynamics of randomly or topographically connected networks of neurons. This target goal seems to be plausible, since all mammals have about the same density of neurons, and similar layered organization (Jehee and Murre, 2008). This simple observation is sufficient to think that neuronal networks may perform particular functions, from a computational point of view, that may be captured and understood without the need of diving too much into the exact details of the anatomy.

The ratio of glial cells to neurons is not constant across species, and some evidence suggests that it may be related to the wiring structure of the network. The more dense and complex is the network, the more glial cells are necessary to keep neurons operational, to ensure the stability of the system. Glial cells will be ignored in the rest of the manuscript, but more and more studies suggest that they may play an important role in neuronal network dynamics. There are indeed growing evidences that glial cells can influence neuronal activity, and, as such, may directly participate in information processing in the brain. They respond to external stimuli (Kelly and Essen, 1974, Schummers et al., 2008), are closely linked with calcium dynamics, and their functional role is still debated (Kirchhoff, 2010). An exhaustive description

of their effects is out of the scope of this manuscript, but it is important to stress that the next generation of neuronal network models will have to deal more carefully with their dynamics.

2 Computational models: the building blocks

From a computational point of view, neurons are active processing units, collecting and integrating over space and time along their dendrites the electrical activities incoming from approximatively 10^4 other neurons. They may be seen as very complex transistors that will choose whether to propagate electrical activity according to some key properties of the input signals that are gathered, turning them either into coincidence detectors in their inputs (either in space and/or time) or into integrators. As the results of biological processes, they display a huge variety of types and shapes, and even though in this thesis we will mostly ignore this aspect, for computational reasons that will be described later, it has to be stressed that neuronal morphologies by themselves could be important factors in the local computations performed by neurons: a pyramidal neuron may not be able to do the same calculations as a Purkinje cell, simply because of the strong differences in their morphologies and biophysical properties such as channel types and/or density, which constrain their functional responses.

2.1 Compartmental models

If we want to model properly the electrical activity within a single neuron, the situation is already problematic. When considering one neuron, such as those shown in Figure 2, designing an exhaustive set of equations that will perfectly describe its activity is not easy. One needs to model its morphology, and approximate it with what is called a compartmental model: a succession of small elements, or sections, being a discretization of the continuous structure of the neuron. In each of these compartments, one can solve differential equations coming from the discretization of the cable equation (Rall, 1957), and by measuring the appropriate density of ionic channels within all of them obtain a calculation of the electrical activity within the discrete approximated neuron. The spike generation mechanism mainly obeys the Hodgkin and Huxley (1952) system of equations. The system can easily end up with a neuron made of thousands of compartments, each of them being a computational node where differential equations should be solved: limits on computational power make it difficult to model a large assembly of neurons with such a level of precision. This is a computational challenge in itself, faced by the Blue Brain project nowadays (Markram, 2006) (to model only 10000 such neurons requires a 8000 nodes Blue Gene supercomputer). The drawback of this approach is that the risk of being overwhelmed by parameters is large. The advantage is that they give insights about active processing and the dendrites (Shepherd et al., 1985). The positions of the synapses along the dendritic tree and the distribution of ionic channels can shape the neuronal response of the cell along its dendrites. Moreover, it has been shown that spikes can back-propagate along the dendritic tree (Frégnac, 1999, Larkum et al., 2001, Nevian et al., 2007), and influence the input/output transfer function. They may also be linked with plasticity, as we will see in Part IV.

2.2 The Integrate-and-Fire model

In contraste to compartmental models, the integrate-and-fire model is a tractable oversimplification of a neuron. It does not take the structure into account, and turns all the cell into a single point in space, where one or a set of differential equations are solved. This is a point process model, and was introduced by Lapicque (1907). Inputs to the neuron are described as ionic currents flowing through the cell membrane when neurotransmitters are released. Their sum is seen as a physical time-dependent current $I(t)$ and the membrane is described as an *RC* circuit, charged by $I(t)$. When the membrane potential reaches a threshold value V_{thresh}, a spike is emitted and the membrane potential is reset (see Figure 7). In its basic form, the equation of the integrate and fire model is:

$$\tau_{\text{m}} \frac{dV(t)}{dt} = -V(t) + RI(t) \tag{1}$$

where V is the membrane potential, and R the resistance of the membrane, with $\tau_{\text{m}} = RC$.

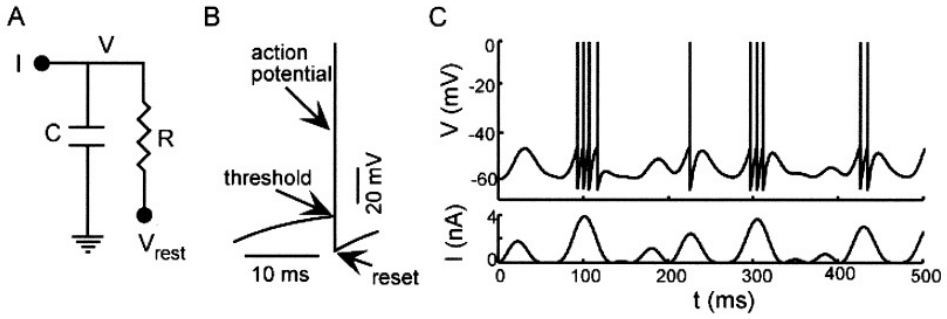

Figure 7: Taken from Abbott (1999) The integrate-and-fire model of Lapicque. (A) The equivalent circuit with membrane capacitance C and membrane resistance R. V is the membrane potential, V_{rest} is the resting membrane potential, and I is an injected current. (B) The voltage trajectory of the model. When V reaches a threshold value, an action potential is generated and V is reset to a subthreshold value. (C) An integrate-and-fire model neuron driven by a time varying current. The upper trace is the membrane potential and the bottom trace is the input current.

In general, the leaky term can be developed and the equation is written as follows, V_{rest} being the resting potential of the neuron:

$$\tau_{\text{m}} \frac{dV(t)}{dt} = g_{\text{leak}}(V_{\text{rest}} - V(t)) + RI(t) \tag{2}$$

where $I(t)$ is the current input coming from the synapses, and if we model the synaptic inputs as current injections, this gives:

$$I(t) = \sum_i w_i \sum_k syn(t - t_k - d_j) \tag{3}$$

where d_i is the conduction delay from neuron i to the considered neuron, w_i the synaptic weight, and t_k the times of the spikes. *syn* is a function which describes the time course of the current induced by one synaptic input, which can be modelled as a Dirac delta function, a

decaying exponential function, or an alpha function. A lot of variations on this integrate-and-fire model have been made: among those, one can make the threshold not constant (Brette and Gerstner, 2005), make the internal dynamic more close to the Hudgin-Huxley formalism (Fitzhugh, 1961), or refine internal dynamics (Izhikevich, 2003).

Capturing in a single point in space the essence of the neuronal dynamics has the big advantage of being much more tractable, on a computational level, and hence allowing large-scale simulations of hundreds of thousands and up to millions of cells (Ananthanarayanan et al., 2009, Mehring et al., 2003, Morrison et al., 2007a). Moreover, depending on the variables of interest, it is important to stress that these point models, even if they represent a huge reduction of complexity, are not necessarily less accurate than more biological and detailed models. For example, if one compares the ability of a leaky integrate and fire model with adaptation and an exponential threshold (Brette and Gerstner, 2005) to predict the spiking responses of a real neuron receiving a fluctuating current injected *in vitro* in its soma, the performance is better than for a full and more complex Hodgkin Huxley model. This is probably because it has many fewer parameters and therefore is easier to fit. In addition, exploration of the parameter space has been able to reproduce the huge variety of the firing behaviour observed in real neuron (Izhikevich, 2003, Naud et al., 2008, Jolivet et al., 2004). Nevertheless, integrate-and-fire models can not give any insight about dendritic integration and non linear interactions between voltage gated channels in the membrane.

2.3 Conductance based synapses

A cortical neuron in an active network receives a massive synaptic bombardment. As explained in the previous section, this neuronal input is commonly approximated as a fluctuating current $I(t)$ but synaptic drives are better modelled by fluctuating conductances: the amplitudes of the post synaptic potentials (PSP) evoked by neurotransmitter release from presynaptic neuron depend on the post-synaptic depolarization level. A lot of study focuses now on this integrate-and-fire model with conductance-based synapses (Destexhe et al., 2001, Tiesinga et al., 2000, Cessac and Viéville, 2008, Vogels and Abbott, 2005). The equation of the membrane potential dynamic is then:

$$\tau_{\mathrm{m}}\frac{dV(t)}{dt} = (V_{\mathrm{rest}} - V(t)) + g_{\mathrm{exc}}(t)(E_{\mathrm{exc}} - V(t)) + g_{\mathrm{inh}}(t)(E_{\mathrm{inh}} - V(t)) \tag{4}$$

When V_{m} reaches the spiking threshold V_{thresh}, a spike is generated and the membrane potential is held at the resting potential for a refractory period of duration τ_{ref}. Synaptic connections are modelled as conductance changes: when a spike is emitted $g \rightarrow g + \delta g$ followed by exponential decay with time constants τ_{exc} and τ_{inh} for excitatory and inhibitory post-synaptic potentials, respectively. The shape of the PSP may not be exponential. Other shapes for the PSP can be used, such as alpha synapses $(t/\tau_{\mathrm{syn}})\exp(1 - t/\tau_{syn})$, or double shaped exponentials synapses $(1/(\tau^1_{\mathrm{syn}} - \tau^2_{\mathrm{syn}}))(\exp(-t/\tau^1_{\mathrm{syn}}) - \exp(-t/\tau^2_{\mathrm{syn}}))$. E_{exc} and E_{inh} are the reversal potentials for excitation and inhibition.

The conductance-based model offers richer dynamics than the current-based one, being also more biologically realistic. The only problem is that analytical analyse are more complex, because the transfer function from inputs to output rate is more complicated. More and more analytical studies have tried to capture the non-linear properties of such conductance models (Burkitt et al., 2003, Kovacic et al., 2009, Cessac and Viéville, 2008). In contrast to current-

based networks, neuronal networks with conductance-based synapses are able to display self-sustained activity, without the need for external noise (Vogels and Abbott, 2005, Marre et al., 2009b). Initial bumps of activity can reverberate and be sustained by the network, while this is not the case with current-based synapses. Moreover, synaptic bombardment *in vivo* can lead to four-fold conductance increases as compared to the quiescent case. Such an increase can have dramatic effects on the integrative properties of the neuron - effects that are neglected in current-based models.

2.4 Mean field models

While the integrate-and-fire and compartmental models consider that the exact times of spike occurrence are important and may play a role in the coding strategies used in the cortex, other models consider that the pertinent information is in the instantaneous firing rate of the neuron. Since the discharge of the cell can be noisy and irregular (see Part II), the spikes are not modelled and the only relevant information used by those models is the firing rate of the neuron. At each time, the neuron can emit spikes with a certain probability $r(t)$, directly related to the activities $r_i(t)$ of its pre-synaptic sources, weighted by some factors w_i:

$$\frac{dr(t)}{dt} = \sum_i w_i f(r_i(t)) \tag{5}$$

where f is a positive, monotonic and increasing function, inducing a non linear relationship between the summed inputs and the instantaneous firing rate $r(t)$. Usually, f is a sigmoidal or a hyperbolic tangent function. An alternative viewpoint is to say that r represent the average firing rate over a population of identical neurons, rather than an instantaneous frequency, and that is why these equations are called mean-field or rate-based models. Solid mathematical results can be obtained with these mean field models, concerning either their dynamics or their learning properties. Some of these results will be mentioned in the following parts, even though the main model used in this manuscript will be the integrate-and-fire.

3 Network activity

3.1 Neuronal networks

When we have a satisfactory model one neuron, capturing its non-linear dynamics with either an integrate-and-fire or a mean field model, we can think about modelling a network of neurons. By connecting them with synapses, network models are a powerful tool to understand brain dynamics and the origin of the electrical activity observed *in vivo*. The type of the neuron models used and the connectivity scheme influence the kind of dynamics that can be observed: in the following, we will focus mainly on recurrent networks of integrate and fire neurons. Among alternatives approaches, networks of binary neurons are easier to analyse analytically (van Vreeswijk and Sompolinsky, 1996, 1998, Sompolinsky et al., 1988). Those models treat the neuron like the spin of an elementary particle, its activity being a binary variable (spiking or silent), and established a link between theoretical physic and neuroscience. Models such as the Ising model (Schneidman et al., 2006, Marre et al., 2009a) are used to infer correlations and structure of the neuronal code. Behind all these models, the key point is

to explain the irregularity of the neuronal discharges observed *in vivo*. To explain it without stochastic inputs, one needs to have large fluctuations of the neuronal dynamics, counterbalanced by weak synaptic weights and by the fact that the activity should be balanced: an excitatory and and inhibitory population should act such that the average activity stays below a certain threshold, while fluctuations may be high and let the neurons cross a threshold, in order to emit spikes.

3.2 The balanced random network

The balanced random network (van Vreeswijk and Sompolinsky, 1996, 1998, Brunel, 2000, Vogels et al., 2005, Kumar et al., 2008b, El Boustani and Destexhe, 2009b, Amit and Brunel, 1997, Renart et al., 2010) is a common and convenient framework for studying the dynamics of large-scale populations of sparsely-connected integrate-and-fire neurons. In these networks, two generic populations of excitatory and inhibitory neurons are reciprocally coupled with weights J_e and J_i (see Figure 8) to generate a balanced regime where the average depolarization of the neurons is roughly constant, subthreshold, and irregular spiking is the result of fluctuations. There is a classical ratio of 4 excitatory neurons for 1 inhibitory neurons, based on the measured ratio in cortex (Braitenberg and Schüz, 1991) and a sparse, random connectivity. Every neuron is typically connected to $1-10\%$ of the others, and depending on certain key parameters, mainly the amount of external noise injected into the system and the balance between excitatory and inhibitory weights, several regimes of activity can be observed. Those regime have been described and classified in Brunel (2000), and can be asynchronous/synchronous (from a population viewpoint) and regular/irregular (from a neuron viewpoint) (see raster plots in Figure 9).

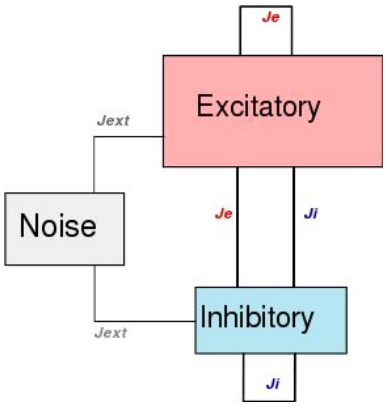

Figure 8: Schema of a balanced random network. Two populations, one excitatory and one inhibitory, are reciprocally coupled with excitatory and inhibitory weights J_e and J_i, receiving extra noise with weight J_ext.

The average firing rate of all the neurons within the network can be constant (asynchronous) or display oscillations (synchronous). The individual discharge of one neuron can be regular (the inter spike intervals (ISIs) are almost all equal), or irregular (the ISIs follow a Poisson distribution). This irregularity is often quantified by the coefficient of variation (CV), given by $CV = \frac{\overline{ISI}}{\sigma(ISI)}$ where denotes the average and σ the standard deviation. A pure Poisson process has a CV equal to 1. The more the discharge is regular, the more the CV tends to 0. Classical values observed *in vivo* in the spontaneous regime are usually close to 1 (Nawrot et al.,

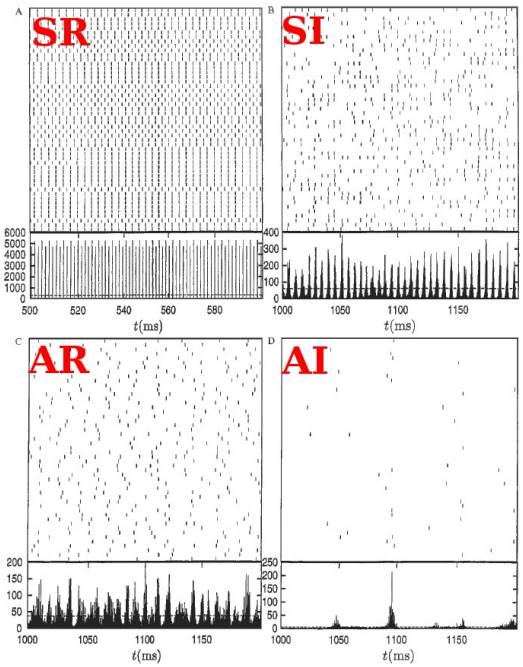

Figure 9: The main activity regimes observed in neuronal networks. Top Left: Synchronous regular (SR): the global activity of the network is oscillatory, and all the neurons fire regularly at intervals of their refractory period. Top Right: Synchronous irregular (SI): the global activity is oscillatory, but neurons fire irregularly as Poisson-like sources. Bottom Left: Asynchronous regular (AR), the global activity is constant, but neurons fire regularly. Bottom Left: Asynchronous irregular (AI), both the individual discharges of the neurons and the global firing rate are irregular. Adapted from Brunel (2000)

2008), so simulations tend to focus on the irregular regime. In such a regime, neurons fire in an irregular manner, behaving almost like Poisson processes, and the average pairwise cross-correlation is modulated by the internal balance or the external input. This regime is also well suited to produce slow oscillations comparable with oscillations observed *in vivo* under anaesthesia (Han et al., 2008, Arieli et al., 1996). A full numerical study of these dynamics will be the subject of Part III. Mean field models are a common tool used to establish and predict, analytically, the stationary average response of homogeneous networks of integrate-and-fire neurons under certain assumptions (neuronal discharges should be independent and Poissonian, in the irregular regime) (El Boustani and Destexhe, 2009b, Brunel, 2000). However, they are much harder to use when complex models of neurons are used, or when inhomogeneities, such as delays, are taken into account.

4 Coding in neuronal networks

4.1 The quest for the code

The question of how information is encoded in primary sensory cortical areas is a fundamental one in computational neuroscience. Based on the assumption that the information content is contained in the spiking activity of the neurons, several coding schemes are debated. On the two opposite sides of a wide field, one can consider information as being either rate coded,

meaning that the exact times of the action potentials are not the relevant information, but that information content is encoded in the average discharge of the cell across time, or time coded, meaning that information is precisely encoded in the exact spike times of the neuron. Both schemes have pros and cons, and it may be that the brain uses one or the other depending on the context or on the sensory area which is considered (Shadlen and Newsome, 1998, Mehta et al., 2002).

If we consider the rate code, this coding scheme is known to be present in sensorimotor areas, where population vectors have been shown to be a very good estimate of the direction of movement (Georgopoulos et al., 1986). Individual motor neurons are poorly tuned and their individual firing patterns contain only little information about the direction and velocity of limb movements, but when their activities are summed and averaged, precise predictions and estimations can be achieved, as if the brain was integrating the relevant information over a population of cells. With such a code, rate-based models are good candidates for understanding the generic properties of the code. The brain can use Bayesian framework to encode/decode information, average over time and/or populations to gain a robust estimate of ongoing processing in a cortical network. The drawback of this rate-based scheme, is that to have reliable and accurate estimates of the firing rates, one need to integrate information over a certain time window, certain population size: since firing rates can be rather low, even *in vivo*, and since neurons are limited in their frequency of discharge ($\simeq$ 200 Hz) by the refractory period, rate-based coding needs receptive fields of neurons to be rather large and information coded in a redundant manner by a large number of cells.

A more recent theory is that information is precisely encoded in the relative spike times of neurons (Gray et al., 1989). This theory has the advantage of circumventing the problem of the time and/or population integration needed by rate-based codes, and several observations are in line with the idea that codes where the spike times do not play any role may be wrong. For a good review on the evidences and the advantages of time codes, one can refer to (Tiesinga et al., 2008). Several theories have been built on this idea, and the time-coding framework has led to successful approaches such as latency coding (Thorpe et al., 1996, Butts et al., 2007), synchrony coding (Singer and Gray, 1995), synfire chains (Abeles, 1991), and phase coding in subcortical structures such as the hippocampus (O'Keefe and Recce, 1993). The time-based codes are built on the fact that precise temporal correlations exist between neurons (Abeles et al., 1994), and that they could be the signature of precise interactions valuable for the system. Since neurons spatially and temporally integrate synaptic inputs, the co-activation, either in space or in time, of a pre-synaptic group of neurons will more reliably drive the post-synaptic one, and propagate information. Simulations and experiments show that synchrony arises naturally in a classical multi-layered feed-forward network of spiking neurons, even when stimulated by uncorrelated inputs (Reyes, 2003). Nevertheless, as for rate coding, time coding suffers from several limitations. First, if exact times are important for the system, then the system must be very robust and insensitive to noise. A response, given twice to the same stimuli, should be twice the same, whatever the state of the brain, and neuronal integration should be reliable. As we will see in the following, the extent to which this is true is not clear from experimental data, and the subject is an ongoing debate (Jacobs et al., 2009). Moreover, in the case of the latency coding, when information is encoded as time differences between spikes, one needs to know what is the "first" spike, and if relative or absolute time differences should be considered. This raises the question of a clock, and could be partially solved by linking the spike discharges to general oscillations/rhythms of the background, such as in hippocampal structures. In addition, the robustness of the code should also be questioned.

What happens if one neuron does not? Spike codes require reliable responses, and raise the concern of reproducibility, from a trial to trial basis, in the responses. To explore these time codes, the integrate-and-fires neurons are essential, and the rest of this manuscript, based on large scale simulations with such neurons, will focus on these codes. More precisely, to understand temporal codes, and to understand how they can work with the ongoing activity of the brain, one needs to understand how correlations between spiking activities emerge in recurrent networks, since it is well known that correlations are a crucial component of the neuronal assembly code (Singer and Gray, 1995, Nirenberg and Latham, 2003) also linked to behaviour (Zohary et al., 1994).

Among the temporal schemes that have been proposed, we will review more in depth the synchrony theory and the synfire chain theory. They both rely on the idea that correlations among neurons carry information about the stimulus. Formally, this means that the mutual information between the stimulus s and the responses r_1, r_2 of two neurons $I_{\text{mut}}(r_1, r_2, s)$ is higher than $I_{\text{mut}}^{\text{indep}}(r_1, r_2, s)$, the mutual information when assuming that $P(r_1, r_2|s) = P(r_1|s)P(r_2|s)$. Note that this issue differs from the issue of the temporal structure of the activity. Here we are interested in the link between the structure of the neural code and the stimulus.

4.2 An overview of some temporal codes

Synchrony The synchrony theory relies on correlations and was first proposed by von der Malsburg (1995) to solve the binding problem. Binding is the fusion, by the brain, of different pieces of local information into a single and coherent higher-order percept. In the visual system, binding helps for example to see not only a group of independent lines close by, but shapes and contour lines that will be interpreted, in higher areas, as particular geometric shapes. In this coding scheme hypothesis, information is encoded by synchronous assemblies of neurons that are co-active. This co-activation has several advantages. First, this avoids the problem of having one neuron coding for each particular feature. If we take for example a visual scene, one can consider that a dedicated neuron can code for a cube, one for a triangle, and so on. We could have as many neuron as shapes, and this hypothesis is often referred to as the grand-mother or pontifical cell concept established by Barlow (1972). In the literature we can also find the term pontifical or grandmother neuron: you may have, somewhere deep in your associative areas, one neuron coding for your grandmother, and firing only when something related to her appears as a stimuli. Evidence for such cells is sparse (Quiroga et al., 2005), and this is not computationally tractable, because having one neuron per concept leads to a combinatorial problem. To avoid this, one can say than a shape is the co-activation of some neurons within a synchronous assembly. For example, four neurons coding for four lines would encode the concept of a square. The theory has been widely debated, and its opponents have raised the argument that synchronous assemblies are not a way to solve the binding problem by themselves, since they do not compute the binding, they just signal it. Many theoretical criticisms therefore focus on the way the synchrony can be detected (Shadlen and Movshon, 1999). In particular, it might not be easy to detect synchronous assemblies among the irregular activity of cortical networks, where the sustained firing rate of all the neurons makes synchronous assemblies emerge by chance. However, recent work has shown that the detection of synchronous assemblies can be done with realistic mechanisms (Shamir and Sompolinsky, 2004, Gütig and Sompolinsky, 2006). The experimental evidence behind the synchrony theory has also been widely discussed, and is still controversial.

Synfire chains Synfire chains can be seen as an extension of the synchrony theory to delayed correlations. Why should only synchronous correlations be considered, when it is known that correlations with delay can also be detected by post-synaptic neurons, provided that the delay in the correlation is compensated for by different conduction delays. The synfire chain theory proposed that the stimulus, decision or actions could be encoded in particular and precise spiking patterns, rather than in the firing rate. This idea is close to the idea of neuronal assemblies originally proposed by Hebb (1949), and is also sometimes referred to "cortical songs", even if slight differences between the two concepts exist. Synfire chains occur as a response to a stimulus, and within a relatively short time scale (a few hundred milliseconds), while cortical songs are more general and could last over longer time scales (up to seconds). A synfire chain is precisely defined by a travelling packet of synchronous spiking events across the cortical network. This would be the signature of an assembly which codes for a particular feature, or action. The feasibility of the propagation of synfire chains in cortical networks has been shown for feed-forward network models (Diesmann et al., 1999), and has been discussed in recurrent network models (Kumar et al., 2008a). It has been shown experimentally by Vaadia et al. (1995) that the correlations between neuron, in the frontal cortex of monkeys performing a behavioural task are dynamically modulated with time, and that this modulation can be linked to behavioural events. The authors proposed that this fast modulation is due to the selective emergence of spiking patterns in the cortex, each of them being associated with a behavioural event. More recently, Ikegaya et al. (2004) demonstrated the reactivation of some particular patterns in the ongoing activity, that may be interpreted as synfire chains. Nevertheless, there are several criticisms of the synfire chain theory, mainly focusing on the methods used to demonstrate their existence (Mokeichev et al., 2007). Establishing the significant occurrence of a given pattern is difficult to assess with finite-size recordings and a restricted number of electrodes. Several studies have shown that at least part of the results which are used to demonstrate significant pattern occurrences could result from stochastic activity (Oram et al., 1999), but careful examinations of the statistics show the presence of higher order correlations in the activity.

Putting all these results together, the existence of correlation based codes, whether synchrony or synfire chains, is still a matter of debate. If correlation codes are acting in cortical areas, the demonstration of their existence is hampered by two limitations. First, the distributed nature of these codes makes them difficult to demonstrate without a large number of simultaneous recordings, and their significance may increase with the number of recorded neurons (see Schneidman et al. (2006) for an example of correlations appearing at larger scales). Conversely, this may be an explanation of why a neural code based on mean firing rate seems to prevail with single cell recording techniques. Furthermore, even if correlation codes exist, they may appear clearly only in complex behavioural tasks. Intuitively, we could say that, when faced with a task of low complexity, the cortex is not required to use its full computational capacity, and will use only mean firing rates. There is an interesting piece of evidence for this is the olfactory system, where abolishing synchrony through suppressing lateral inhibition has no effect on the discrimination of dissimilar odours (easy task), but severely impairs discrimination of similar odours (hard task) (Yokoi et al., 1995).

Part II

On the ongoing activity *in vivo*

Summary

5 Introduction **33**

6 The spontaneous activity **34**
6.1 Origin of the ongoing activity 34
6.2 Nature and structure *in vivo* 35

7 Similarity between evoked and ongoing activity **37**
7.1 Evidence from the literature 37
7.2 Main results 38

8 Information transmission under ongoing activity **58**
8.1 Ongoing activity and information propagation 58
8.2 Reliability as a function of the stimulus 59
8.3 Main results 61

9 Discussion **73**
9.1 A general paradigm 73
9.2 On the role of the recurrent connections 76

5 Introduction

Close your eyes, cover your ears, and do not touch anything. Even without any active afferent stimuli coming to the sensory areas, the brain is known to display a spontaneous activity whose nature and origin is still a matter of debate. Whatever the time is, your brain is constantly active and millions of neurons are emitting action potentials every millisecond. If not, this is a clear clinical sign that you are dead. This spontaneous activity, also called ongoing activity, is by definition the running activity of the brain when not "facing" or processing particular stimuli. This is somehow the resting state of the brain, when no particular or at least no known actions are performed. The nature of this activity is correlated with behavioural states, being different during awake states and sleep. Its nature can be captured at different scales, depending on the device used to record it. For example with electro-or magneto encephalography (EEG or MEG), the frequency content of large scale brain electrical activity can be used to asses and characterize different cortical states of activity. It is usually divided into frequency bands (see Table 1): α rhythm, with frequencies around 8-12 Hz, can be seen as an awake resting state, when eyes are closed. β rhythm (10-30 Hz) is a kind of normal working state. γ oscillations, at more than 30 Hz appear during cross-modal tasks, active and intense concentration and memory tasks (in addition to slow oscillations observed in the EEG signal, such as δ ($<$4Hz) and θ (4-10 Hz)).

Name	Frequency range
δ	1-4 Hz
θ	4-10 Hz
α	8-12 Hz
β	10-30 Hz
γ	30-80 Hz

Table 1: Table of the different frequency ranges often mentioned in the literature. Note that the precise borders of these ranges are often modulated from one study to the other.

The exact nature of those oscillations is irrelevant for the rest of the manuscript, and this enumeration was just intended to underline and stress the fact that ongoing activity is constantly present in the brain. The naive vision of a brain waiting for a particular stimulus needs to be challenged to achieve a major breakthrough in computational neuroscience. What is the nature of this ongoing activity, how does it affect, shape and filter the incoming sensory streams?

The ongoing activity is a major difference between biological systems and silicon ones. In modern computers, transistor properties are calibrated and designed such that performing twice the same operations will give twice exactly the same results. Variability, noise and fluctuations are enemies of reproducibility so that error correction codes, controlled design and other solutions are used to be sure that the output of the system is deterministically linked to its input. This uncorrelated and residual activity is often considered as "noise" in experimental protocols, that needs to be eliminated by averaging over many trials to obtain a clear evoked response to a particular stimulus. This averaging, if possible in a well designed and controlled setup where the exact same stimulus can be presented several times to the system, is far from what the brain is used to dealing with. This "noise" is an intrinsic property the system is, at the best able to cope with, perhaps able to exploit. This intriguing capacity will be the subject of this Part. Since variability *in vivo* is the result of the recurrent interactions between

neurons and since neurons are reliable when receiving a fluctuating input current (Mainen and Sejnowski, 1995), the statistical structure of the noise may be less random than expected and directly linked to the underlying topology of the network. Its particular structure can even play a role in information transmission and signal processing, and this is the question that will be tackled in this Part. After having discussed the origin and the nature of the spontaneous activity in neuronal networks, results obtained with large scale models of integrate and fire neurons are presented to show how recurrent interactions can interact with an incoming signal to provide a new conceptual framework for information transmission.

6 The spontaneous activity

6.1 Origin of the ongoing activity

The origin of this spontaneous activity in the cortex is far from being clearly understood. Isolated neurons by themselves seem to have some homeostatic processes making them able to be spontaneously active, even if no incoming activity is present. This phenomenon has also been reported in cortical neurons (Llinás et al., 1991, Mazzoni et al., 2007), either taken *in vitro* brain slices (Timofeev et al., 2000) or in cultures (Gross et al., 1982). Such activity, once created, can easily reverberate and be amplified in the microcircuits made by the neuronal connections, and lead to a self-sustained activity regime. How long this ongoing activity will last depends on the preparation, on the size of the network, and some other unknown parameters. Plasticity, the fact that the efficacy transmission of the synapses is not constant over time, the fact that neurons may be dying during the time of the *in vitro* preparation: all affect this activity, and progressively silence the network. Nevertheless, even an isolated brain (Andjus et al., 1967) displays such a spontaneous activity, i.e. the brain does not require external inputs to generate its own recurrent activity.

To focus on the sensory systems *in vivo*, the origins of ongoing activity are already found in the transduction layers, where the external world is transformed into electrical activities that are relayed to the cortex. It is the case for example in the auditory cortex (Tritsch et al., 2007), and also in the visual system: the retina generates a spontaneous activity, even in dark conditions, of approximatively 5 to 30 Hz (Wong et al., 1998). During its development, travelling waves of activity spontaneously pop up and propagate, being integrated by bipolar and amacrine cells in the retina before being sent to the lateral geniculate nucleus (LGN) of the thalamus, and then to V1 in the cortex. Noise in the sensors is then constantly impacting the activity of the neurons downstream. The resting state of the brain is then an active one, and relevant information need to be extracted from these continuous and ongoing flows of sensory inputs. To illustrate this, one can look at recent data from techniques that allow to record more and more accurately the detailed and exact ongoing activity in awake animals. While this can be done intra-cellularly, at the single neuron level (Lee et al., 2006) and revealed by inverting inhibitory events when recording with KCL electrodes (Figure 3B in Bringuier et al. (1997)), the ongoing activity is much easier to observe at the population level, with multi-electrode arrays. In Figure 10, taken from the work of Lin et al. (2006), one can see the spiking activity in an awake mouse, and the subtle difference when the animal receives a clear external stimulation (red line). More than 200 neurons are recorded in parallel with a multi-electrode array. Understanding how pertinent information is extracted from this ongoing activity is a crucial step.

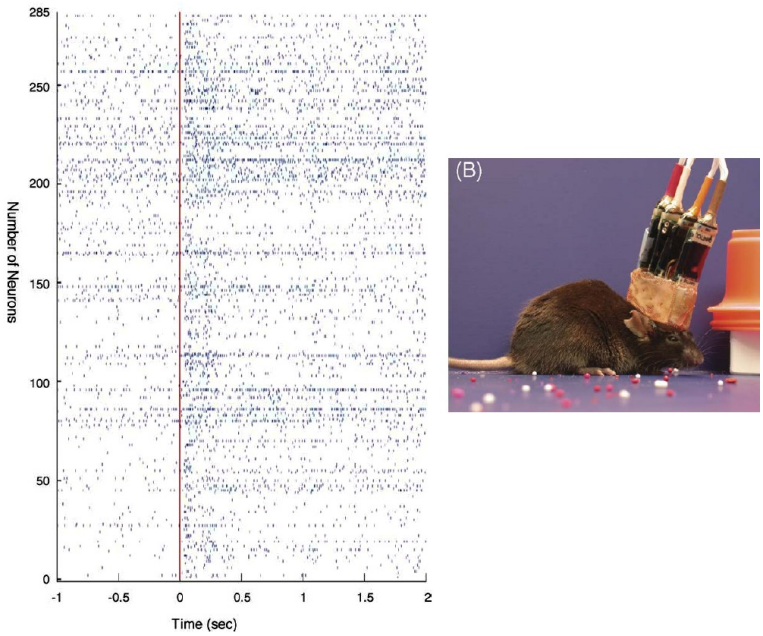

Figure 10: Extracted from Lin et al. (2006). Left: spontaneous spiking activity, as a raster plot, for 200 neurons recorded in awake behaving mouse. Stimulation time (a puff of air is made on the leg of the animal) is indicated by a red line. Right: picture of the device.

6.2 Nature and structure *in vivo*

The exact nature of the spontaneous activity is hard to capture, for several reasons. The first one is that the dynamical nature of this resting state is not that clear, and the question of its stationarity remains open. Imaging tools available nowadays (from intracellular recordings to local field potentials (LFP), voltage-sensitive dyes (VSD), two-photons) can give an insight about its nature, but not a full and exhaustive view of the electrical activity over large portions of the brain. However, we will try here to summarize the principal characteristics of this ongoing activity, and the main observations that can made on its dynamics.

Age, wakefulness and anaesthetic dependence First, the statistics of this spontaneous activity recorded *in vivo* depend on the age and the state of the animal. Most of the *in vivo* data that have been acquired in experiments over the past 50 years have came from anaesthetized animals, and the brain activity is far more oscillatory than what is nowadays reported in awake animals. It has also been acquired in various species, and across different ages. As one can see in Figure 12 (below), ongoing activity in the visual cortex of the ferret at three distinct stages of development is different: neuronal networks are shaped extra-utero during development by electrical activity (Spitzer, 2006, Katz and Shatz, 1996). But even if we record spontaneous activity in a "stable" and mature network, anaesthetics are known to perturb the balance between excitation and inhibition (Winters, 1976), leading to pathological activity that may be sometimes far from the awake regime. More and more efforts are devoted to developing awake recordings that will be as precise and controlled as in the anaesthetized context (Lin et al., 2006, Ferezou et al., 2006, Greenberg et al., 2008), to draw a better picture of the brain's activity. Recent techniques allowing the recording of precise neuronal activity over a large cortical surface are multi-electrode arrays and voltage sensitive dyes (VSD), but in both cases there is a trade-off between spatial and temporal resolutions and having a clear

snapshot of the instantaneous dynamics in a particular volume of the cortex is not yet feasible. An observation that can be made from these awake recordings is that while activity in the anaesthetized regime may be oscillatory and synchronous, activity in the awake regime is far more irregular and asynchronous. Ongoing activity is also rather sparse and neurons fire spontaneously at relatively low firing rates. Again the situation varies as a function of the age, the area, the species, and is used the anaesthetic, but in freely moving awake rats, sparse activity was observed intra-cellularly (Lee et al., 2006). An exhaustive review of anaesthesia and age effects is out of the scope of this manuscript: its aim is more to demonstrate to the reader that ongoing activity, if clearly observed, has a complex nature which is hard to capture.

Slow oscillations In spontaneous activity under anaesthesia, slow oscillations observed (in the EEG but also in the membrane potential of individual neurons (Steriade et al., 1993a,b)) have been considered as reflecting a switch between "up" and "down" states. In the membrane potential trace recorded *in vivo*, one observes some silent periods, where the membrane stays close to its resting potential ("down" state), and some very active periods where the membrane is strongly depolarized and the neuron sustains a strong irregular spiking activity ("up" state). The similarity of structure between the "up" and the awake state, observed at the intracellular level and at the local EEG level, has led some to consider "up" states as "fragments of wakefulness" (Destexhe et al., 2007), and to link it with ongoing activity in awake regime. Under different anaesthesia, with the notable exception of barbiturate, Steriade et al. (1993b) reported *in vivo* similar slow oscillations (the most striking binary switch between two states is obtained is with ketamine-xylazine Steriade et al. (1993a), and this anaesthesia triggers an artificially high level of synchrony (Amzica and Steriade, 1995)). Dual recordings in the cortex and in thalamo-cortical neurons reveal that these oscillations are synchronized between thalamus and cortex, but only over small time windows (Contreras and Steriade, 1995).

Irregularity of the ongoing activity Oscillatory or not, it has been observed that the spiking activity of neurons *in vivo* is rather sparse and highly irregular. Most V1 neurons for example display Poissonian or supra-Poisson spike count variability in response to low dimensional stimuli such as bars and gratings (Dean, 1981). This and other experimental data are in favour of the synchronous or asynchronous irregular regimes explained in Part I: neurons fire as Poisson sources, irregularly, with a coefficient of variation for their inter-spike intervals close to 1 (Nawrot et al., 2008). The origin of this irregular activity observed in the sub-threshold voltage and/or in spiking activity is linked to synaptic activity. To check that is is not the result of stochastic opening of intrinsic channels, Pare et al. (1998) injected and anaesthetic, the tretrodoxin (TTX), *in vivo*, which drastically reduced the synaptic activity. This provides a direct demonstration of the synaptic origin of the background activity. Later computational studies estimated that voltage gated channels contribute less than 10% of the background activity (Destexhe and Paré, 1999). To illustrate the irregularity, Figure 11 shows the membrane potential of a cat V1 neuron in absence of stimulation (data taken from Frégnac's lab). As we can see, the cell is spontaneously firing action potentials, and its membrane potential is fluctuating, as indirect evidence of the synaptic bombardment received by the neuron.

Ongoing activity and chaos Such an irregular and ongoing activity has led several authors to raise the question of whether it is chaotic, and several studies stress the hypothesis that the dynamical properties of ongoing activity are close to those of a chaotic system (Faure and Korn, 2001, Korn and Faure, 2003, El Boustani and Destexhe, 2009a, London et al., 2010).

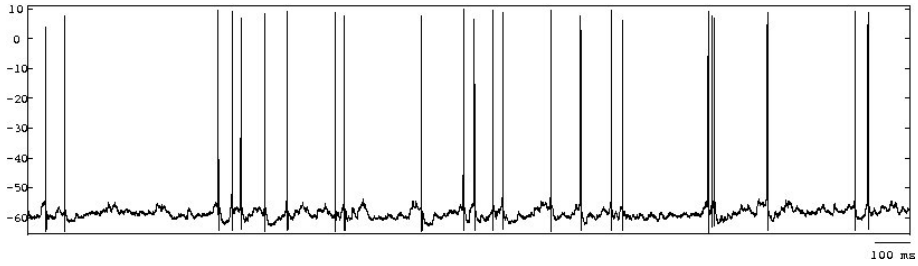

Figure 11: Intracellular recording of the membrane potential of a cat V1 neuron in spontaneous activity, under anesthesia. Internal lab data. The irregularity of the discharge can be observed, and the small fluctuations of the membrane potential reflect the ongoing synaptic bombardment.

Chaotic dynamical systems exhibit behaviours in which two slightly different initial conditions can lead to two very distinct evolutions of the system, leading to two distinct trajectories diverging exponentially. This is the so-called "butterfly effect": even the change of the activity in one neuron can drastically impact all the remaining dynamics. The sensitivity of the ongoing activity to initial conditions is hard to control, and therefore the Lyapunov exponents of the brain as a chaotic dynamical system are hard to estimate. In the cortex, the influence of one particular spike on the evolution of the dynamics is not understood, but *in vivo* interesting experiments have been performed to show that injection of an electrical pulse in a single neuron of the thalamus can drastically affect and change the dynamical state of the cortex (Li et al., 2009).

Combined with other results from broadband recordings, such as EEG, there is therefore accumulating evidence that ongoing activity in the brain is chaotic, or close to a chaotic state. They transition between order and chaos is call the edge of chaos, and the brain may operate close to a transition point (Bertschinger and Natschlager, 2004, Kitzbichler et al., 2009). The so-called balanced random network framework, developed in Part I, where neurons are randomly and sparsely connected, provides a good model for understanding neuronal properties in such regimes: with such a system in the AI regime, the response to a given input might not be reproducible (Banerjee et al., 2008), and it is therefore important to understand how reliability and reproducibility of neuronal responses can emerge in this context. This question has been partially addressed by several theoretical works. First, the averaged response of a network of binary neuron model to a uniform stimulation was shown to answer to the stimulation reliably and in a very fast manner, with a time constant below the time constant of a single neuron (van Vreeswijk and Sompolinsky (1996, 1998), see also Gerstner (2000) for the same property in networks of integrate-and-fire neurons). The response amplitude increases linearly with the stimulation strength. However, this concerns only the "macroscopic" response: the information is extracted by averaging over the whole population. Secondly, some work studied Lyapunov exponents of pulse-coupled network of oscillators neurons, in response to a common fluctuating noise injected into all the cells (Teramae and Fukai, 2008): when the ratio between the coupling strength and the variance of the external inputs is high enough, the system can be reliable.

7 Similarity between evoked and ongoing activity

7.1 Evidence from the literature

Single cell level At the level of the cortical area, the structure of the ongoing activity seems to contain correlations. A striking example is the work of Lampl et al. (1999): recording pairs of cells intra-cellularly in the cat primary visual cortex, they showed that the correlation between the membrane potentials of the cells is high during spontaneous activity. In another work, Lampl and colleagues, this time in the rat cortex, showed that excitation and inhibition impinging upon a cortical cell are also very synchronised (or phase-locked with a delay of a few milliseconds), during both spontaneous and evoked activity (Okun and Lampl, 2008). In the visual system, other groups of cells are particularly strongly correlated. Several works have shown that the correlations between cells sharing similar orientation preference is higher than for the others, by using multiple extracellular recordings (Ts'o et al., 1986, Nelson et al., 1992). This structure of ongoing activity has been also studied by Tsodyks et al. (1999), using imaging in the deeply anaesthetized cat with voltage sensitive dyes (VSD). They showed that the ongoing activity in the cat visual cortex measured with this technique has a spatial structure similar to that evoked by a drifting grating presented to the animal. To obtain this result, they recorded the activity of a cell extracellularly simultaneously with the VSD activity map. The spike trigger average (STA) of the VSD map based on the single cell activity gives similar maps for both preferred evoked and ongoing activities, which means that triggering the VSD acquisition on spontaneous spikes reveals the orientation map domain to which the cell belongs to. In auditory cortex, similar results were found by Luczak et al. (2009).

Network level To illustrate the similarity between evoked and ongoing activity at a larger scale, Figure 12 shows some results from Fiser et al. (2004). Ongoing activity is recorded with a single row multi-electrode array (MEA) in the visual cortex of awake ferrets, at three distinct periods of the development (three different ages), and each time in three different conditions: when the animal is in the dark, watching dense noise, or watching a natural movie. The key point of this is the similarity, at the spiking level, between the three distinct conditions. This can be also observed in the similarity between the temporal correlation functions computed for the three stimulus conditions. From an external point of view, it is hard to tell when the system is processing an input or not. Activity is rather constant from dark to the natural scene, and even if the spatio-temporal profile of the spiking correlations is affected, it is not drastically different.

Evidence gathered with spike-triggered LFP or multi-electrode arrays (MEA) (Nauhaus et al., 2009, Smith and Kohn, 2008) confirms largely inferences made from the spatio-temporal study of synaptic echoes by Bringuier et al. (1999) and reveals that correlations spread in space and in time: due to the local connectivity of the brain (Hellwig, 2000), spikes can trigger waves that pop up and propagate in the ongoing activity, creating correlated activity in the recurrent networks (Swadlow and Alonso, 2009). These waves are also triggered by evoked stimulations, spreading over several cortical areas, as can be observed in many VSD imaging studies (Han et al., 2008, Contreras, 2007, Mohajerani et al., 2010). The waves can be bottom up and go from primary sensory areas to higher areas, but evidence for lateral horizontal connections (Bringuier et al., 1999) or top to bottom feedback needs also to be taken into account (Roland et al., 2006). The similarity between ongoing and evoked activity in visual cortex can be observed at such a network scale (Arieli et al., 1996, Kenet et al., 2003,

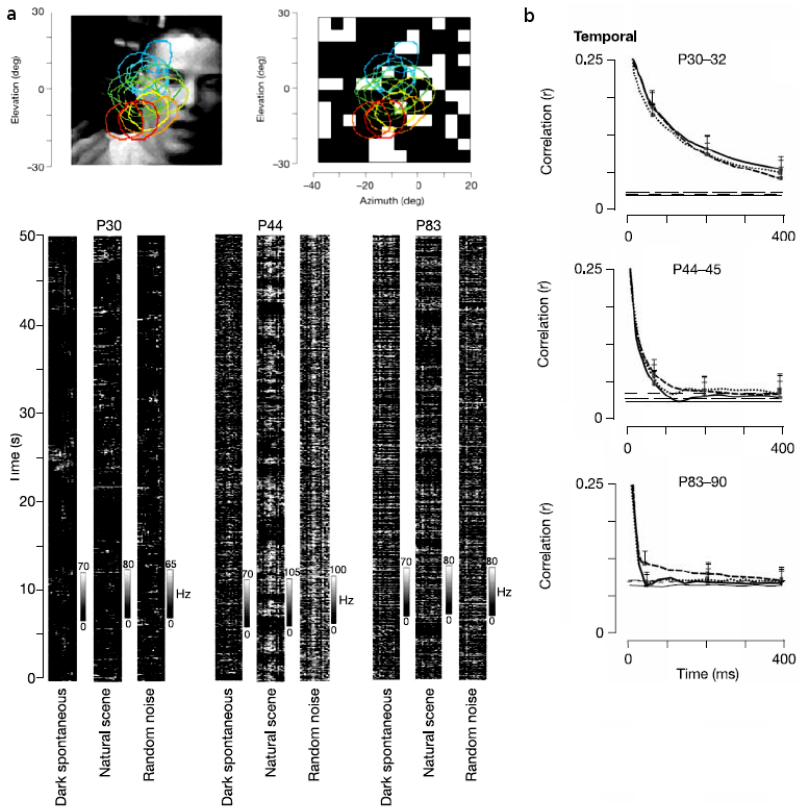

Figure 12: Extracted from Fiser et al. (2004). a) Time series plots of neural activity recorded under three interleaved stimulus conditions at three different ages. The time series graphs were obtained from a single animal in each age group. At each age, visual stimulation modulates the spatio-temporal pattern of spontaneous activity, but does not significantly alter its basic correlation structure. Bin width, 20 ms. b) Temporal correlation functions computed for the three stimulus conditions. Thin horizontal lines show plots of correlation functions computed at each age for each condition using randomly shuffled binned spikes. Random temporal shuffling of spike trains abolished all correlations in all three age groups, demonstrating that the observed shifts in correlated activity were not simply a result of the developmental increase in cell firing rates. Bin width, 20 ms. Error bars represent s.e.m.

Han et al., 2008), but also with two-photons imaging at a smaller scale (Greenberg et al., 2008).

7.2 Main results

In a collaborative study, myself and several colleagues showed intra-cellularly how the similarity between evoked and ongoing activity can be assessed *in vivo* by considering the power spectra of the membrane potential traces (after adequate filtering from the spikes). The membrane potential of a neuron reflects the combined impact of the ongoing synaptic bombardment it is subject to, and therefore could be used to asses the state of this incoming activity. A novel analysis, based on the frequency content of recorded membrane potentials, during evoked or ongoing activity, can provide direct evidence that the average amount of correlations during both states is similar. This study was published in the following article. My contributions to this study were mainly achieved on the simulation part and on the *in computo*

results.

OPEN ACCESS Freely available online

PLoS COMPUTATIONAL BIOLOGY

Network-State Modulation of Power-Law Frequency-Scaling in Visual Cortical Neurons

Sami El Boustani[•], Olivier Marre[•], Sébastien Béhuret, Pierre Baudot, Pierre Yger, Thierry Bal, Alain Destexhe, Yves Frégnac*

Unité de Neurosciences Intégratives et Computationnelles (UNIC), UPR CNRS 2191, Gif-sur-Yvette, France

Abstract

Various types of neural-based signals, such as EEG, local field potentials and intracellular synaptic potentials, integrate multiple sources of activity distributed across large assemblies. They have in common a power-law frequency-scaling structure at high frequencies, but it is still unclear whether this scaling property is dominated by intrinsic neuronal properties or by network activity. The latter case is particularly interesting because if frequency-scaling reflects the network state it could be used to characterize the functional impact of the connectivity. In intracellularly recorded neurons of cat primary visual cortex *in vivo*, the power spectral density of V_m activity displays a power-law structure at high frequencies with a fractional scaling exponent. We show that this exponent is not constant, but depends on the visual statistics used to drive the network. To investigate the determinants of this frequency-scaling, we considered a generic recurrent model of cortex receiving a retinotopically organized external input. Similarly to the *in vivo* case, our *in computo* simulations show that the scaling exponent reflects the correlation level imposed in the input. This systematic dependence was also replicated at the single cell level, by controlling independently, in a parametric way, the strength and the temporal decay of the pairwise correlation between presynaptic inputs. This last model was implemented *in vitro* by imposing the correlation control in artificial presynaptic spike trains through dynamic-clamp techniques. These *in vitro* manipulations induced a modulation of the scaling exponent, similar to that observed *in vivo* and predicted *in computo*. We conclude that the frequency-scaling exponent of the V_m reflects stimulus-driven correlations in the cortical network activity. Therefore, we propose that the scaling exponent could be used to read-out the "effective" connectivity responsible for the dynamical signature of the population signals measured at different integration levels, from Vm to LFP, EEG and fMRI.

Citation: El Boustani S, Marre O, Béhuret S, Baudot P, Yger P, et al. (2009) Network-State Modulation of Power-Law Frequency-Scaling in Visual Cortical Neurons. PLoS Comput Biol 5(9): e1000519. doi:10.1371/journal.pcbi.1000519

Editor: Karl J. Friston, University College London, United Kingdom

Received May 28, 2009; **Accepted** August 25, 2009; **Published** September 25, 2009

Funding: OM was supported by Délégation Générale pour l'Armement and Fondation de la Recherche médicale. SB was supported by Délégation Générale pour l'Armement. This work was supported by CNRS, ANR (Natstats, HR-Cortex) and European Commission-funded (FACETS project, FP6-2004-IST-FETPI 15879) grants. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: fregnac@unic.cnrs-gif.fr

• These authors contributed equally to this work.

Introduction

Assigning a functional role to the correlations in network activity is still controversial. While many studies have proposed that the mean firing rate of the neuron contains much of the information about the sensorimotor interaction with the environment, or the behavioral task being performed [1,2], other studies have suggested a specific role of higher-order interactions in cortical processing [3–5].

Here, we explore another way to extract correlations, through the scaling properties of the power spectrum (hereby called "power spectral density" or PSD) of the membrane potential of single neurons. A particularly common form of frequency scaling is the power-law, according to which the PSD scales as $1/f^{\alpha}$ at high frequencies, with some exponent α which may be integer or fractional (fractal). Power-law frequency-scaling is ubiquitous in electrophysiological measurements of neuronal population activity, from spiking activity [6] to fMRI signals [7], but its function and origin are still controversial. Some studies consider it as the manifestation of neural "avalanches", a special form of cell assembly dynamics which would appear when the cortical network is in a critical state [8,9] and which would be optimal for information processing. Power-law decay functions may also provide the basis for long-lasting interactions in adaptation [10,11] or memory storage [12].

Several explanations for the origin of power-law scaling have been proposed. At the intracellular level the membrane potential activity was shown to have power-law scaling at high frequencies, with exponent values around $\alpha = 2.5$ for synaptic background activity *in vivo* [13,14] and channel noise [15–17]. Cable equations predict values of α between 3 and 4 for inputs distributed in soma and dendrites, and the non-ideality of the membrane capacitance was proposed to account quantitatively for these values [18]. However, it is unclear whether this exponent can also be modulated by extrinsic factors *in vivo*, and in particular by the synaptic bombardment evoked by sensory input.

As we report in this paper, we decided to approach this issue by analyzing the V_m activity of neurons recorded intracellularly in cat primary visual cortex *in vivo*, when the network is driven so as to be in an irregular activity regime. We found that the power-law scaling observed in the intracellular activity PSD at high frequencies is modulated by the stimulus. We examined whether

Author Summary

Intracellular recording of neocortical neurons provides an opportunity of characterizing the statistical signature of the synaptic bombardment to which it is submitted. Indeed the membrane potential displays intense fluctuations which reflect the cumulative activity of thousands of input neurons. In sensory cortical areas, this measure could be used to estimate the correlational structure of the external drive. We show that changes in the statistical properties of network activity, namely the local correlation between neurons, can be detected by analyzing the power spectrum density (PSD) of the subthreshold membrane potential. These PSD can be fitted by a power-law function $1/f^{\alpha}$ in the upper temporal frequency range. *In vivo* recordings in primary visual cortex show that the α exponent varies with the statistics of the sensory input. Most remarkably, the exponent observed in the ongoing activity is indistinguishable from that evoked by natural visual statistics. These results are emulated by models which demonstrate that the exponent α is determined by the local level of correlation imposed in the recurrent network activity. Similar relationships are also reproduced in cortical neurons recorded *in vitro* with artificial synaptic inputs by controlling *in computo* the level of correlation in real time.

the scaling exponent variations observed *in vivo* can be accounted for by theoretical models *in computo*, using paradigms where the correlation among inputs can be modulated. First, we designed a recurrent network model composed of a thalamic and a cortical layer and showed that when varying the correlation of the thalamic input to the cortical layer power-law exponent modulations were consistent with the *in vivo* results. The scaling exponent thus reflects in this model a specific correlational state of the network imposed by the input. We then dissected out those aspects in the activity impinging on the recorded neuron that can modulate the scaling exponent, and also explored the alternative hypothesis that intrinsic properties of the individual neuron are sufficient to explain the observed modulation. For this purpose, we applied different correlated synaptic inputs to neuron models. This confirmed that a change in the correlation of the synaptic input can modify the power-law exponent. Finally, we investigated this paradigm in cortical neurons *in vitro* using the dynamic-clamp technique and confirmed the results obtained with computational models. We discuss how these results are consistent with the theory that the power-law exponent modulation reflects changes in the correlation state of the network activity.

Results

Stimulus Dependence of Frequency Scaling in V1

15 neurons were recorded intracellularly in the primary visual cortex of the anesthetized and paralyzed cat (see Materials and Methods). Each neuron was recorded while presenting four full-field stimuli through the dominant eye (Fig. 1): a drifting grating at the cell's optimal orientation and spatial frequency (DG), a high spatial definition dense noise (DN), a natural image animated with a simulated eye movement sequence (NI), and a grating animated with the same eye movement sequence (GEM). After removing the spikes from the V_m signals by interpolation, we computed their PSDs (see Materials and Methods). These PSDs systematically exhibit a scaling behaviour in a broad, high-frequency band. To extract the scaling exponent, we fitted a linear function to the log-log representation of the PSD, for a range of frequencies from 75 to 200 Hz (Fig. 2B), where the quality of the linear fit is high (mean correlation coefficient $r = 0.95$). Note that this chosen band is also above the frequencies at which synaptic and membrane filtering cut-off appear [19].

Figure 2A shows the PSDs of the intracellular responses to the four stimuli for the same cell. In the log-log scale representation we observed a dependence of the slope, and hence the frequency-scaling exponent, on the stimulus. To confirm these effects at the population level, we compared for each cell the values of the exponent between pairs of stimuli. Figure 2C shows the comparison between stimuli DG and NI for each cell, and averaged over trials. Although the absolute value of the exponent was highly variable from cell to cell (ranging from 2.0 to 3.5), it was systematically lower, for the same cell, for NI than for DG (paired Wilcoxon test, $p \leq 0.003$). The magnitude of this difference was much larger than the standard error of the mean (SEM) among the different trials for the same protocol.

We checked whether the value of the exponent could be correlated with the recorded cell's averaged V_m or firing rate. The corresponding correlation coefficients were computed for each stimulus and then averaged together. We found that neither the firing rate ($r = 0.13$) nor the averaged V_m ($r = 0.2$) were correlated enough to explain the variations of scaling exponent (although these weak correlations were marginally significant ($p \leq 0.07$), except for the NI protocol where no correlation was found).

We also estimated whether these systematic modulations were visible at the spiking level, or present only at the V_m level. We computed the Fano factor exponent (see Materials and Methods) for the *in vivo* spiking responses. In contrast to the frequency-scaling of the V_m, we did not observe any consistent variation of the spiking scaling-exponent with the visual stimulus. Moreover, there was no significant correlation between the V_m and the spiking scaling exponents ($r = 0.2$, $p \geq 0.1$).

In some cells, the same protocol was repeated consecutively, interleaved with 2–3 s of spontaneous activity. We could not see any consistent difference between the power law exponents of the first trial and the others. This means that the dynamics reflected by the power law exponent appear in less than 10 seconds.

These results indicate that the changes of frequency-scaling for the same cell as a function of the stimulus context are mainly determined by the differences in the visual stimulus statistics. Based on the comparison of the frequency-scaling exponents between all possible pairs of stimuli, we divided the stimuli into 2 groups. The exponents obtained from the intracellular responses to DG and GEM were not significantly different from each other but differed significantly from those obtained with NI and DN. We summarized these results by computing the relative changes from DN to the other protocols (Fig. 2D).

For a subset of cells, we also presented three additional stimuli designed as surrogates of the natural stimulus. The "Spatial Random" stimulus is composed of the natural image "scrambled" by randomizing the phases of its Fourier coefficients and animated with the same sequence of eye movements. The "Time Random" stimulus is composed of the natural image animated with a similarly "phase-scrambled" version of the eye movement trajectory. Finally, the "space and time random" stimulus is composed of the scrambled image animated with the scrambled eye movements (plotted as Natural Image Surrogate or NIS in Fig. 2D). These three stimuli evoke power-law exponents similar to the DN protocol (no significant difference, Wilcoxon paired test, $p \geq 0.32$, $p \geq 0.014$, $p \geq 0.13$ respectively, and see Fig. 2D for the third surrogate). Even though we did not see a significant difference between NI and DN or between DN and NIS, there is a significant difference between NI and NIS, the latter being the

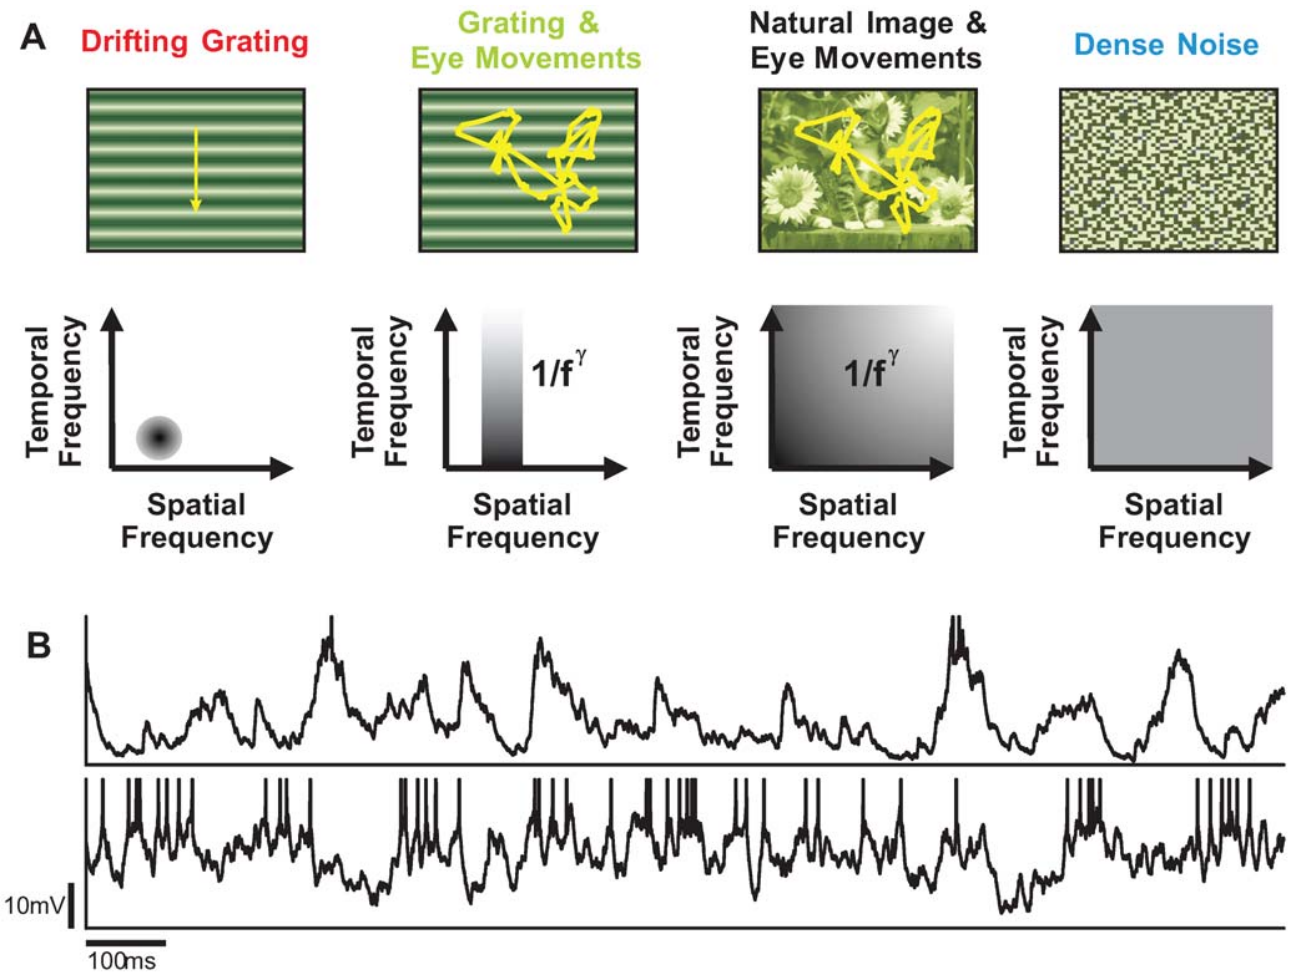

Figure 1. Protocols of visual context dependence. A: Stimuli used in the *in vivo* experiments. From left to right: Drifting Grating (DG): a sinusoidal grating with optimal spatial frequency and orientation, drifting at optimal frequency; Grating & Eye Movements (GEM): the same grating animated by a trajectory simulating the dynamics of eye movements; Natural Image & Eye Movements (NI): a natural image animated by the same trajectory mimicking eye movements; Dense Noise (DN): a dense noise of high spatial and temporal definition. All these stimuli were full-field and presented monocularly in the dominant eye. **B**: examples of intracellular responses of the same cell to the NI (top trace) and the DG (bottom trace) stimuli (data from Baudot, Marre, Levy, Monier and Frégnac, submitted; Baudot et al., 2004; Frégnac et al., 2005).
doi:10.1371/journal.pcbi.1000519.g001

same stimulus with reduced phase coherence (Wilcoxon paired test, $p \leq 0.003$, $p \leq 0.003$, $p \leq 0.006$ respectively for the three surrogate stimuli).

From this study, we concluded that the value of the frequency-scaling exponent of the intracellular signal is strongly dependent on the visual input. It is interesting to note that the scaling exponent always seems to be smaller when the stimulus is less correlated (DN being the extreme case where there is no correlation in the stimulus).

Spontaneous Activity

We applied the frequency-scaling analysis to periods of spontaneous activity recorded in the same cells. Comparison between the frequency-scaling exponent of Spontaneous Activity (SA) and those in response to the five different stimuli was also performed at the population level. We observed a systematic increase from SA to the DG and GEM stimuli (Fig. 2D and Fig. 2F; paired rank Wilcoxon test, $p \leq 0.0003$; the average difference between paired data SA-DG or SA-GEM is significantly different from zero, t-test, $p \leq 0.0001$). In contrast, the SA frequency-scaling exponents are similar to those for DN, NIS and NI (Fig. 2E; for NI $r = 0.81$, paired rank Wilcoxon test, $p \geq 0.5$; slope = 0.82; the average difference between paired data SA-NI or SA-DN is not significantly different from zero, t-test, $p \geq 0.1$).

Multifractal Analysis

To estimate how much the frequency-scaling exponent tells us about the multiscale statistics of the intracellular signal, we performed a multifractal analysis (see Materials and Methods). We therefore computed the two first moments of the singularity spectrum over the different cells and protocols.

The first moment is linearly related to the frequency-scaling exponent measured on the PSD [20]. The respective values were indeed correlated over the population. The second moment is slightly above 0 for the four protocols (DG: 0.0757 ± 0.1035, GEM: 0.0816 ± 0.1062, NI: 0.1018 ± 0.1244, NIS: 0.0680 ± 0.0909 and DN: 0.0755 ± 0.1015), while no significant differences were found between protocols. The intracellular signal is thus very close to a monofractal process, exhibiting self-similar behaviour. Furthermore, the first-order part of the singularity spectrum is the only one which varies with the visual stimulation. The functional sensitivity of our

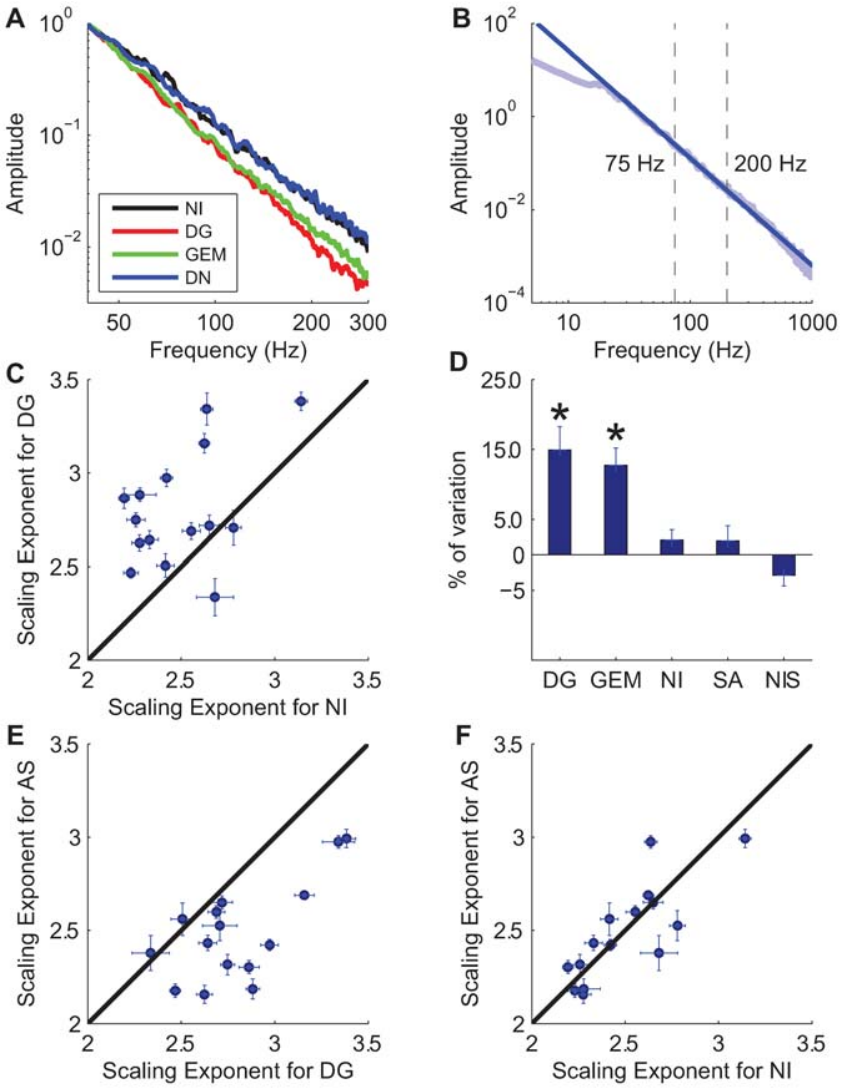

Figure 2. Change of frequency-scaling according to visual context. A Power spectral density (PSD) for a given cell in response to the four different stimuli presented in Fig. 1. The traces have been normalized so as to obtain the same value at 40 Hz, for the sake of clarity. **B** Illustration of the linear fit between 75 and 200 Hz for the dense noise protocol. The power-law scaling region extends beyond those frequencies but is affected by synaptic filtering at low frequencies and by noise artefacts at high frequencies. **C** Frequency scaling exponent comparison between DG and NI stimuli for each cell. The error bars represent the standard error of the mean (SEM) on the estimation of the frequency-scaling exponent across the 10 repetitions for each stimulus. The black abscissa line indicates equality between the DG and NI condition. **D** Population analysis relative to the DN case. Each bar indicates the percentage of variation from the DN frequency-scaling exponent. The asterisks (*) indicate a significant difference over the population of cells between the frequency-scaling exponents in response to DN and a given stimulus (paired Wilcoxon test, $p < 0.005$). The fourth bar represents the relative change between the spontaneous activity (SA) and the DN condition. **E** Comparison between the frequency-scaling exponent measured during NI stimulation and spontaneous activity (SA) for each cell. The black line indicates equality. **F** Same comparison than **E** between DG and SA.
doi:10.1371/journal.pcbi.1000519.g002

multiscale statistics can be reduced to the power-law behaviour of the V_m trace.

Frequency Scaling in a Simple Retinotopic Cortical Model

To study the effect of correlated input, we considered a simple model of a cortical network fed by an input with a controlled level of synchrony. This model was shown to be sufficient to reproduce the frequency-scaling exponent modulation measured above. In order to mimic the cortical network and the retinotopy of the input, we simulated topographically-connected networks of excitatory and inhibitory neurons using integrate-and-fire models and conductance-based synapses (see Materials and Methods). We considered networks with topographically organized connectivity where each neuron is connected to its neighbours according to a Gaussian distribution (Fig. 3A).

The stimuli used during *in vivo* experiments have different levels of correlation (Fig. 1A): the DG stimulus is highly correlated across space and time (one Dirac impulse in the spatio-temporal spectral plane), while the DN is, by definition, fully uncorrelated (flat spectrum in space and time). We chose to stimulate the recurrent network model with inputs having different levels of synchrony. The visually driven thalamic inputs project in a local region of space (Fig. 3A), and the cortical response is thus the product of both the thalamic input and the recurrently mediated activity. The different levels of

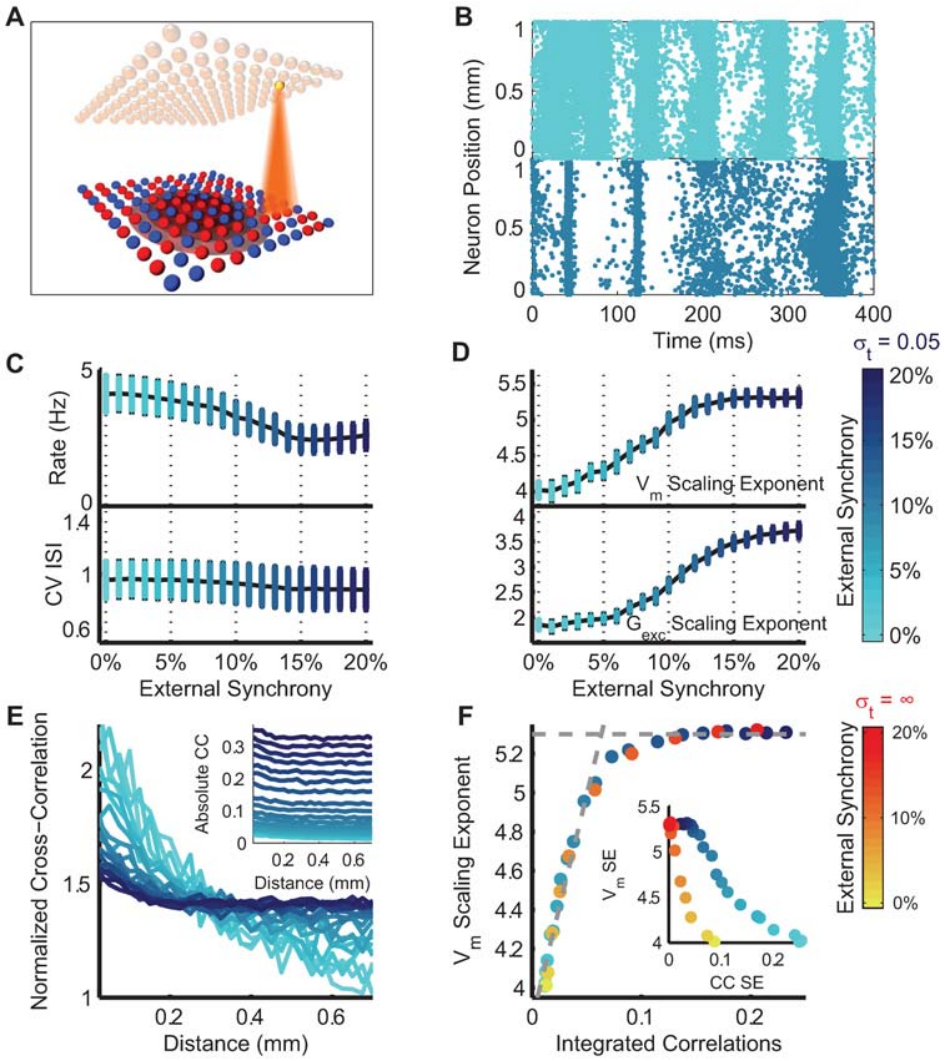

Figure 3. Modulation of the frequency-scaling in a recurrent network model with inputs of variable synchrony and spread. A Schematic representation of the network structure and connectivity. The cortical (lower sheet, blue and red neurons) and thalamic input (upper sheet, yellow neurons) layer-like networks ($1\ mm^2$) face each other. The cortical neurons are locally connected together, according to a Gaussian distribution ($\sigma_c = 0.15\ mm$) and the retino-thalamic input projects its synaptic connections on the cortical layer through a narrower Gaussian distribution ($\sigma_t = 0.05\ mm$). **B** Example of raster plots in the cortical layer in response to two thalamic input synchrony levels (top: synchrony of 0%; bottom: synchrony of 10%). **C** Mean firing rate (top) and coefficient of variation (bottom) of the cortical layer response to thalamic inputs of different synchrony levels. For each simulation, twenty neurons were randomly chosen among the network population to estimate error bars. **D** V_m (top) and G_{exc} (bottom) frequency-scaling exponents as functions of the input synchrony. Bars indicate standard deviations of the scaling exponent values. **E** Averaged spatial cross-correlation between neuronal activities as a function of the distance between pairs of neurons, for different input synchrony levels, normalized by the total area of the distant-dependent cross-correlation function. Inset: same graph without the normalisation. **F** values of the V_m frequency-scaling exponent as a function of the coefficient of correlation integrated over distance. Inset: values of the frequency-scaling exponent as a function of the correlation extent in the network activity (see text). The same results are shown in red for an infinite spread of the thalamic input.

doi:10.1371/journal.pcbi.1000519.g003

synchrony give rise to responses in the cortical area with different structures (Fig. 3B), although the mean firing rate and the coefficient of variation of the cortical activity remain roughly constant over the different levels of input synchrony (Fig. 3C). In particular, the cortical layer displays spontaneous waves of activity with an irregular and low frequency firing regime (rate $\simeq 4$ Hz and ISI CV $\simeq 1$) when there is no synchrony within the thalamic discharge. The presence of correlation in the external input disrupts these waves and creates synchronous firing in the cortical layer (Fig. 3B).

The frequency-scaling exponent in the model was estimated from the V_m traces of twenty cells (see Materials and Methods). The values of the V_m and G_{exc} frequency-scaling exponents both increased when the input synchrony increased (Fig. 3D). This also held for the inhibitory conductance G_{inh} which behaved as its excitatory counterpart (data not shown). This is consistent with the *in vivo* results where stimuli with more correlation (DG, GEM) evoke higher values of the scaling-exponent than the "decorrelated" stimuli (NI, NIS and DN).

Determinants of the Scaling Exponent

We next examined which features of the network activity structure could be related to this modulation of the scaling

exponent. Fig. 3E shows the spatial pairwise cross-correlation between pairs of neuron as a function of the interneuronal distance, for different levels of the input synchrony. The increase in input synchrony resulted in two simultaneous changes: a global increase of the cross-correlation values (Fig. 3E, inset) as well as a flatter spread profile over larger distances; when normalizing by the integral of the correlation over distance, it appears that the fall-off of the cross-correlation function (CC) is steeper for lower levels of synchrony (Fig. 3E). In summary, the different levels of input synchrony modulate not only the global level of the correlation in the cortical network, but also its topographic extent and distance dependence.

We next quantified the two features of the network activity that are modulated by the input synchrony and compared their modulation to that of the V_m exponent. We first compared the V_m exponent values to the *integrated correlation*, defined as the normalised cross-correlation integrated over distance. The frequency-scaling exponent increased linearly with the integrated correlation (from 0.0 to 0.05) and saturated around 5.25, for an integrated correlation of approximately 0.1 (Fig. 3F).

We also observed that the pairwise correlation between neurons scales with distance when expressed in logarithmic coordinates, which could be related to the V_{-m} frequency-scaling exponent. The corresponding cross-correlation scaling exponent (CC SE), which reflects the fall-off gradient of the spatial correlation, decreases linearly when the V_m exponent increases (Fig. 3F, inset).

To disentangle the influence of these two factors, we tested the effect of the spread of the thalamic projection to the cortical layer, which parameterizes the extent of the spatial correlation of the inputs. We ran the same simulations with an infinite spread (i.e., the thalamo-cortical connections were random). This condition might be related to the effect of a decorrelated background noise. While the relation between the cross-correlation scaling exponent and the V_m exponent was shifted, the relation between the integrated correlation and the V_m exponent remained unchanged. We found similar results by varying the spread between these two extreme values (data not shown): the spread had no direct influence on the V_m exponent value but shifted the baseline cross-correlation scaling exponent. Thus the variation of the spread, which determines the spatial structure of the input, did not alter the relation between the integrated cross-correlation and the V_m exponent.

This important relationship shows that, in this model, the integrated correlation is detected at the single-cell level through the membrane potential power spectrum scaling property for any stimulus. This measure thus provides a reliable hint about the actual functional state of the network. It also appears that, even if the spatial structure of the correlation is varied, the exponent value remains unchanged. This latter observation could explain why stimuli differing in their spatial structure can produce similar exponents *in vivo*.

As in the previous *in vivo* study, we estimated the Fano Factor scaling exponent. Even when averaging over a population of randomly assigned neurons, the mean Fano Factor did not exhibit any systematic variation with the input synchrony, the integrated correlation or the cross-correlation scaling exponent. This is in accordance with the *in vivo* results.

Finally, it is interesting to note that this network model can reproduce the changes in the frequency-scaling of the V_m observed *in vivo*, despite its simplicity and the absence of any form of power-law in the spatial rules of connectivity: the thalamo-cortical and the cortico-cortical connectivities are drawn in our simulations from Gaussian distributions. Therefore it is not necessary to implement a scale-free connectivity to observe a frequency-scaling exponent emerging in the synaptic bombardment.

Frequency-Scaling in Single-Cell Models

We have shown that the V_m scaling exponent is related to the integrated cross-correlation of the network activity. This integrated correlation depends on at least two factors: the global correlation level of the activity (*correlation strength*) and the spatial extent of the network correlation (*correlation extent*). In our recurrent network model, both are modified simultaneously when varying the input, which makes the isolation of the precise feature modulating the scaling exponent difficult. We thus turned to the modeling of a single neuron receiving parameterized correlated synaptic noise in order to dissect out the influence of the different parameters of this correlated noise on the postsynaptic V_m scaling exponent.

Furthermore, although the network model provides a possible explanation for the V_m frequency-scaling exponent modulation, this does not exclude a possible alternative hypothesis for our *in vivo* observations : due to the non-linearity in the neuronal transfer function, the V_m frequency-scaling exponent variation *in vivo* could be due to the variation of the input firing rate or the different depolarisation levels from one protocol to the other.

For these reasons we measured the V_m frequency-scaling exponent in isolated neuronal models in response to several correlated synaptic inputs, where all these parameters can be varied independently. We also injected the same correlated synaptic patterns into biological neurons *in vitro* through dynamic clamp. This allowed us to test independently the effect of the correlation strength and extent, and to test the simpler hypothesis aforementioned.

To further understand the relationship between the presynaptic activity and the V_m frequency-scaling, we designed a model assuming that the irregular activity originates in the synaptic activity impinging on the recorded cell. Indeed, since the frequency-scaling exponent varies for the same cell and different visual stimuli, it must be linked to the activity of the network surrounding the observed neuron. Note that, being interested only by these relative changes, we did not search for the mechanisms shaping the absolute value of the V_m PSD scaling, which may include intrinsic mechanisms [18,21,22]. For this reason we show the *relative* modulation of the values of the frequency-scaling exponent in different models and *in vitro* experiments, the baseline being the exponent in response to Poisson stimulation, unless otherwise noted.

In the retinotopic model discussed in the previous section, synchronous input in the thalamic layer evoked synchronous firing in the cortical layer at random positions. These firing assemblies affect the recorded neuron through lateral connections with different propagation delays, which depend on the distance from the presynaptic neuron. The temporal correlations in the presynaptic spike train impinging on the recorded cell thus reflect both the direct thalamic input and the spatial correlations observed in the intracortical distance-dependent cross-correlation. Our aim was to determine how these temporal correlations present in the afferent pattern are conveyed from the presynaptic bombardment to the subthreshold activity through cell integration. Note that the propagation delays play a crucial role in the translation of spatial correlations into temporal correlations. Indeed, if the presynaptic population could interact instantaneously with the postsynaptic cell (no propagation delay), synchronous firing would only increase the membrane potential variance.

The model is composed of N presynaptic neurons (Poisson processes) that all fire at the same mean rate ν, with a constant synchrony fraction r. This means that each emission of a spike occurs simultaneously in $k+1=rN$ neurons (Fig. 4). These presynaptic neurons then project with different conduction delays

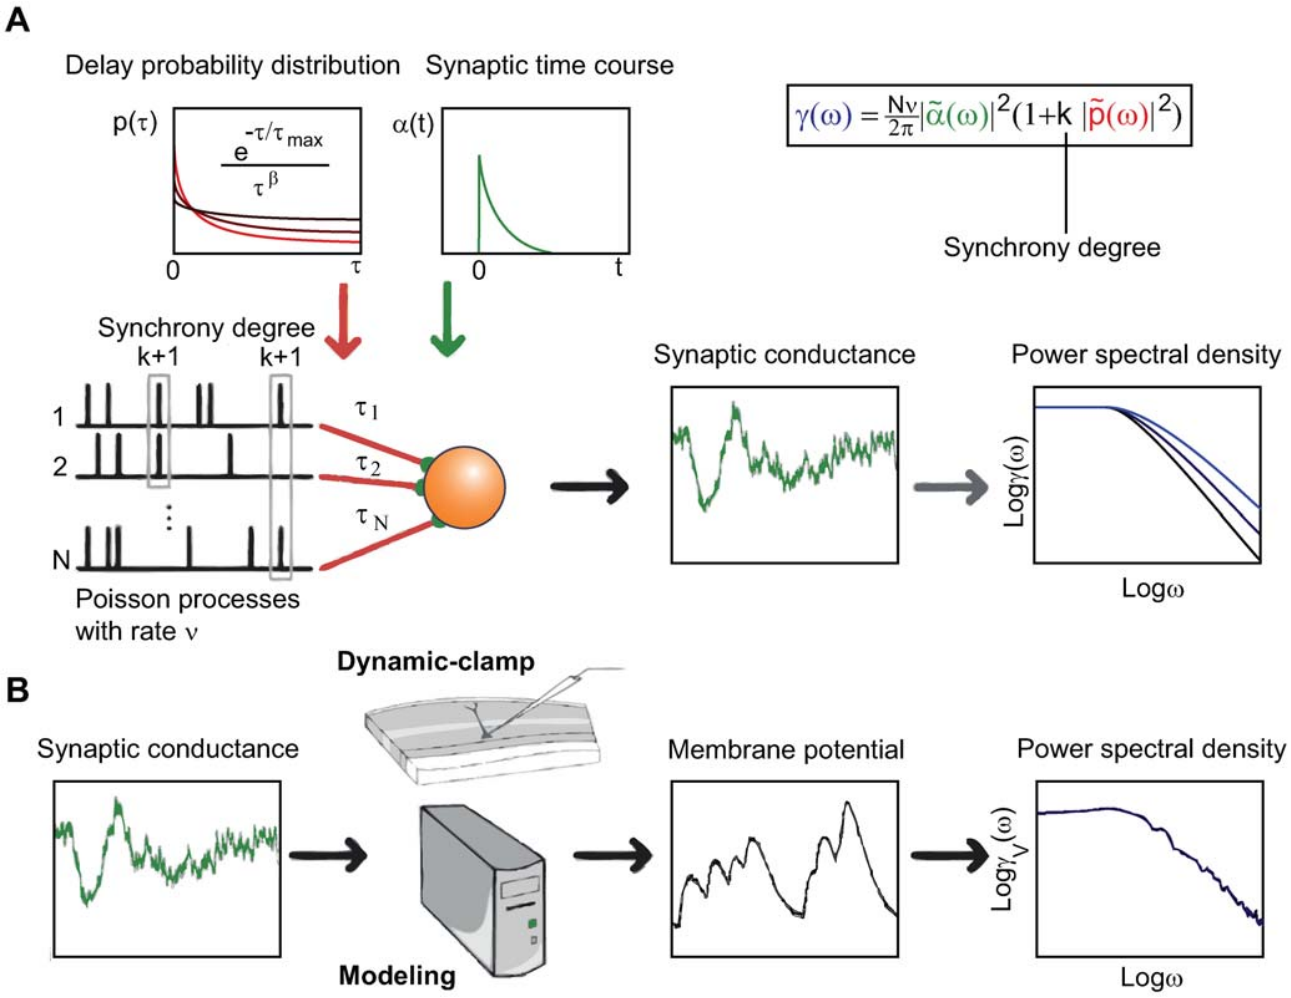

Figure 4. Conceptual scheme of the synchrony generator model and the corresponding conductance injection in model and *in vitro* neurons. A Simple representation of the conductance generator. At each time step dt, with a probability proportional to the firing rate νdt, k+1 neurons emit a spike synchronously. These spikes are then conveyed to the postsynaptic neuron with different delays, distributed according to a power-law probability density function (red curves). The arriving spikes then trigger post-synaptic conductances of exponential form (green curve, synaptic time course). The resulting conductance trace G_{exc} (green trace) has a PSD (blue curve) with a frequency power-law scaling behaviour. The analytical relation between the Fourier transform of the delay distribution and the PSD is given above the graphs. **B** The resulting synaptic conductance is then injected either in a model of single neuron or in a biological neuron through dynamic-clamp (see Materials and Methods). In both cases, the resulting membrane potential is measured and the corresponding PSD is estimated.
doi:10.1371/journal.pcbi.1000519.g004

to the same postsynaptic neuron, which represents the recorded cell. This means that spikes emitted simultaneously by various presynaptic sources will arrive with different delays at the postsynaptic neuron, thus creating a high-order structured temporal correlation pattern. The delays are chosen randomly according to a distribution $p(\tau)$ (Fig. 4).

We emphasize that this model is not biologically realistic: it is a correlated spike train generator parameterized by the synchrony level r and the delay distribution $p(\tau)$. To give more intuition about what these parameters represent, and to make a link with the recurrent model, we can interpret r as the strength of the correlations in presynaptic activity, and $p(\tau)$ as the way these correlations are temporally distributed. Note that both of these parameters would influence the integrated correlation measured previously in the recurrent model (the spatial correlation in the recurrent model becomes a temporal correlation when considering the delays between distant neurons).

In this model, it can be shown [23,24] that the analytical expression for the conductance PSD resulting from the synaptic integration of all these inputs is given by Eq. 5

$$\gamma(\omega)=\frac{Nv}{2\pi}|\alpha(\omega)|^2(1+k|p(\omega)|^2)$$

where $\alpha(\omega)$ is the Fourier transform of the synaptic time course (when the synapse is exponential, this is a Lorentzian curve), and $p(\omega)$ is the Fourier transform of the delay probability distribution.

From this expression, we find that a controlled way to impose an activity-dependent frequency scaling behaviour in this model is to impose a temporal delay distribution having itself a power-law form. Furthermore, this form of correlation is what we found in the recurrent model, although it was not implemented in the connectivity. For this reason the delay distribution will have the

form

$$p(\tau) \propto \frac{\exp\left(\frac{-\tau}{\tau_{max}}\right)}{\tau^{\beta}} \quad (1)$$

The β parameter parameterizes the extent of the delay distribution: the higher is β, the narrower will be the delay distribution. An infinite value of β would correspond to all delays equal to 0. We emphasize that this choice of delay distribution is not *ad hoc*, but rather is imposed in order to control the V_m frequency-scaling exponent. Other forms of delay distribution might produce more realistic presynaptic patterns, but we focus here on the part of the correlations that will exert a direct control over the postsynaptic frequency scaling.

The power spectral density of this delay distribution is [6]:

$$|p(\omega)|^2 \propto \frac{1}{(1+(\omega\tau_{max})^2)^{(1-\beta)}} \quad (2)$$

The synaptic conductance G_{syn} frequency-scaling exponent is thus equal to $2+2(1-\beta)$ for frequencies beyond the synaptic filtering and the delay cut-offs. Note that, as already shown at the population level in Fig. 3F, the synchrony level detected in the presynaptic train has a "gating" role according to (Equ. 5): no synchrony at all would give a G_{syn} frequency-scaling exponent of 4 whatever the value of β. Moreover, the relationship between the exponent and β is here uncovered as soon as a minimal level of synchrony is present in the presynaptic population (theoretically, any $k>0$).

Excitatory-Only Simulations

We numerically simulated this model to check the previous analytical expression. We took a population of $N=5000$ neurons and first fixed the presynaptic firing rate to $\nu=10$ Hz. For different values of the delay distribution parameter β_{exc}, and synchrony r, we simulated the model to produce G_{exc} and V_m traces. Figure 5A shows the resulting G_{exc} and V_m PSDs, for a fixed synchrony level $r=6\%$, and β_{exc} ranging from 0 to 1. The PSD frequency scaling decreases when β_{exc} increases for frequencies above 20 Hz.

We then measured the frequency-scaling exponents in these traces to quantify this result (see Materials and Methods) and plotted them as a function of the synchrony level r and β_{exc} (relative to the Poisson exponent). As predicted, the exponent decreases when the parameter β_{exc} increases (Fig. 5B). This inverse relation between the G_{exc} frequency-scaling exponent and β_{exc} appears more and more clearly as the synchrony r increases, and saturates for $r>4\%$ (Fig. 5B). Nevertheless, even with an amount of synchrony as small as $r=0.5\%$, the dependence of the power-law on β is already monotonic. We obtained a linear relation between β and the output frequency-scaling exponent, although the absolute values are not exactly those predicted by the analytical relation, most probably due to a finite-size bias of the estimation.

To illustrate this "gating" effect of the synchrony, we plotted the frequency-scaling exponent against the synchrony level r, for fixed β_{exc} (Fig. 5C). When increasing r, the exponent first increases and then saturates to a plateau which depends on β_{exc}.

Identical results were obtained for V_m but with a systematic shift of 2 corresponding to the membrane integration (absolute exponent values were between 2 and 4 for the conductance, and between 4 and 6 for V_m). This is what we would expect for a current-based model for which the effect of membrane integration results in a shift of 2 in the frequency-scaling exponent. This shows numerically that the non-linearity induced by the use of conductance-based synapses does not alter this relationship. Therefore, as long as few neuron assemblies are firing simultaneously in the presynaptic population, their correlations are made visible through the postsynaptic membrane potential PSD. Note that the results displayed in panels B and C of Fig. 5 are reminiscent of those obtained for the retinotopic cortical network in Fig. 3F. Indeed, increasing the synchrony or decreasing the β parameter would both increase the integrated cross-correlation, which in turn increases the V_m scaling exponent.

Excitatory-Inhibitory Simulations

The synaptic bombardment received by a cortical neuron is composed of both excitatory and inhibitory inputs. We extended our model by adding a population of presynaptic inhibitory neurons which has the same organization as the excitatory population described earlier, parameterized by the synchrony r and the delay distribution parameter β_{inh}. While independently varying the inhibitory and excitatory exponents (β_{exc},β_{inh}, we measured the corresponding V_m frequency-scaling exponent. We first performed this analysis with the two presynaptic populations having a fixed amount of synchrony ($r=6\%$), to ensure the impact on the G_{exc} and G_{inh} frequency-scaling exponents, and being completely uncorrelated. Fig. 6A shows how the V_m frequency-scaling exponent varies with β_{exc} and β_{inh}. The V_m frequency-scaling exponent seemed to be dominated by the β_{exc} parameter, while the influence of the inhibitory inputs remained marginal. Since the firing rate is similar for excitatory and inhibitory neurons, this dominance was due to the excitatory-inhibitory ratio ($\frac{N_{exc}}{N_{inh}}=4$). We checked that it was not due to the closer inhibitory reversal potential in additional simulations where we changed the reversal potential (data not shown). Note that when $\beta_{exc}=\beta_{inh}$, the V_m frequency-scaling exponent behaves as in the excitatory-only case (Fig. 6D).

We then examined the case where excitatory and inhibitory inputs are correlated, which is more realistic in view of most of the *in vivo* studies [25–27]. The functional relationship between conductance correlations and the V_m frequency-scaling exponent is conserved for stronger excitatory-inhibitory correlation, although it is slightly affected, especially for small β_{exc} values (Fig. 6B–C). To illustrate this effect, we plotted the variation of the V_m frequency-scaling exponent for $\beta_{inh}=\beta_{exc}$ and different levels of correlation (Fig. 6D).

For a sufficient amount of synchrony, the final V_m frequency-scaling exponent will thus be mainly influenced by the frequency-scaling exponent of the delay distribution β_{exc}, and, to a lesser extent, influenced by the correlation between excitatory and inhibitory conductances, and β_{inh}. We found that adding a constant delay between the excitation and inhibition as often observed experimentally does not change the V_m PSD slope value.

To conclude, our model shows how changes in the parameters which determine the correlation in the presynaptic bombardment affect the frequency-scaling exponent of the V_m signal. These changes are of the same order of magnitude as that which was observed *in vivo*. Increasing synchrony increases the V_m frequency-scaling exponent up to a limit which depends on the β parameters. Increasing β_{exc} or β_{inh}, or the correlation between excitation and inhibition, decreases the V_m exponent. However, it is much more affected by the correlations present in the excitatory neurons than in the inhibitory ones, since there are many more excitatory neurons.

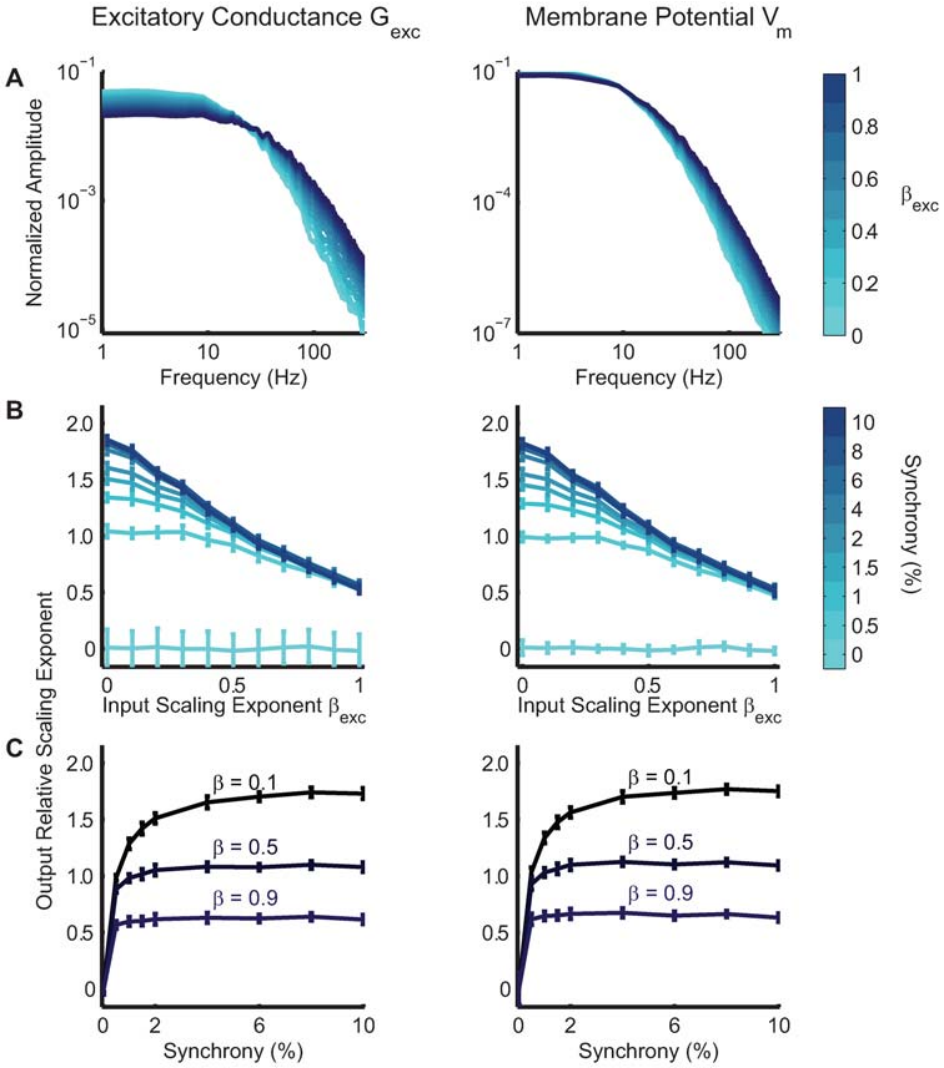

Figure 5. Variation of the value of the frequency-scaling exponent at the conductance and membrane potential levels for excitatory input only as a function of the parameters β_{exc} and r (synchrony percentage). Excitatory conductance G_{exc} and membrane potential V_m are plotted in the left and right column respectively. **A** Illustration of the PSD modulation on a log-log scale for different values of the parameter β_{exc} ranging from 0 (light blue) to 1 (dark blue). **B** Variation of the output frequency-scaling exponent with the β_{exc} parameter, for different levels of synchrony. When 4% of the presynaptic neurons are synchronous, the relation is almost saturated. **C** The gating effect of synchrony. For three fixed values of $\beta = 0.1$, 0.5 and 0.9, the curves represent the modulation of the output frequency-scaling exponent according to percent synchrony.

doi:10.1371/journal.pcbi.1000519.g005

Spike and V_m Power Law Relationships

Previous work on power-law frequency-scaling has mainly been based on extracellular recordings, either to characterize single-cell spiking correlations [6] or self-organized avalanche dynamics in networks [8]. Intracellular recordings, as used in the present study, offer a larger sampling of the network dynamics. Indeed, we can ask whether correlations in the synaptic input visible at the V_m level are still present in the spiking output. We estimated the Fano Factor (FF) for the numerical model to better understand the V_m-spike frequency-scaling exponent relation.

We measured the frequency-scaling exponent in the spiking activity in response to different correlated synaptic input patterns, built by varying the parameters β_{exc} and β_{inh}. Figure 7A illustrates the Fano factor scaling behaviour for $\beta_{exc} = \beta_{inh}$ ranging from 0 to 1, and shows a linear increase of the spiking frequency-scaling exponent with β_{exc} for time bins between 10 and 100 milliseconds. However, we then tested whether the same relationship holds for different resting potentials V_{rest} of the postsynaptic neuron (Fig. 7B). It appears that the relation between the spiking and the V_m frequency-scaling exponents is strongly dependent on the depolarization level.

This dependency is confirmed when varying β_{exc} and β_{inh} independently. Other parameters can drastically influence the spiking frequency-scaling exponent. As an illustrative example, figure 7C–D show the corresponding spiking frequency-scaling exponents for two different depolarization levels and excitation-inhibition correlation levels; in 7C the postsynaptic $V_{rest} = -65mV$ and there is no correlation, whereas in 7D $V_{rest} = -62.5mV$ and the correlation is set to 0.4%.

In light of these results, the lack of correlation between V_m and spiking frequency-scaling exponents, and the absence of systematic modulations for the spiking exponent *in vivo* and in the recurrent model can be explained. This is likely due to the sensitivity of the

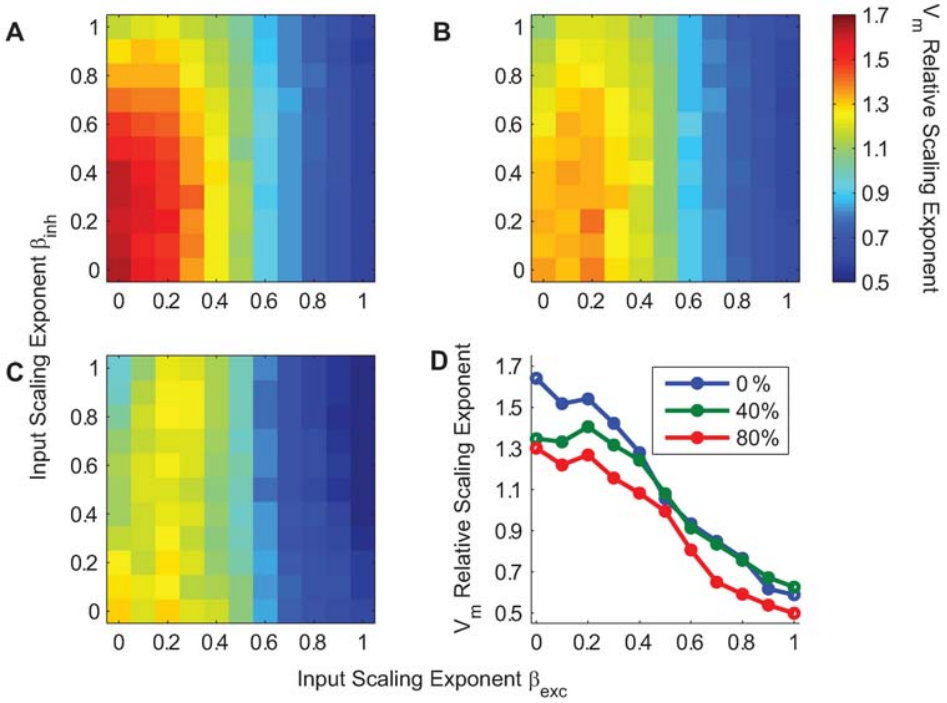

Figure 6. V_m Relative values of the frequency-scaling exponent for different excitatory and inhibitory parameters β_{exc} and β_{inh}. The synchrony percentage r has been fixed to 6% in each simulation. **A** The relative V_m frequency-scaling exponent (color-coded) for β_{exc} and β_{inh} ranging from 0 to 1 without any correlation between excitatory and inhibitory inputs. **B,C** Same graph but with 40% (panel B) and 80% (panel C) correlation between excitatory and inhibitory inputs. In each graph, the excitatory input has a stronger influence on the output frequency-scaling exponent than the inhibitory input. **D** For $\beta_{inh} = \beta_{exc}$, the output frequency-scaling exponent modulation is represented according to different correlation levels.
doi:10.1371/journal.pcbi.1000519.g006

latter to other parameters that also vary with the stimulus, such as the depolarization level. The spiking frequency-scaling exponent for single-cell study is thus hardly sufficient to characterize the self-similar behaviour of the neural activity. In the *in vivo* data, the FF is measured across a high heterogeneity of depolarization levels, and is thus not reliably linked with the presynaptic correlation. In contrast, the subthreshold activity has shown its robustness to changes in depolarisation, and thus provides a much better insight into the network correlation state, being averaged over a large number of presynaptic spiking neural elements.

Controls for Different Firing Rates and Resting Potentials

So far our model has shown how the frequency-scaling exponent can be modulated by the correlations present in the presynaptic activity pattern. However, we had to control for a simpler alternative hypothesis. In *in vivo* data the evoked neuronal mean activity was modulated by the different stimuli (on average 160% decrease from DG to NI), implying that the presynaptic firing rate of the recorded cell varies from one visual stimulus to the other. It is possible that this increase of firing rate induces a change in the frequency power-law scaling. In the following, we call this hypothesis the "first-order hypothesis". The weak correlation between the cell firing rate and the frequency-scaling exponent observed in the *in vivo* section makes such an hypothesis rather unlikely. However, to directly test this hypothesis on our model, we changed the input mean firing rate from 2.5 Hz to 10 Hz for both excitatory and inhibitory synaptic inputs. For each condition, we computed the V_m frequency-scaling exponent. Figure 8B (left panel) shows that it is almost unaffected by the input firing rate. Although we observed a small decrease in the frequency-scaling exponent when increasing firing rate, this could still not explain the *in vivo* results. Indeed, in the latter case, even though the correlation is weak, the frequency-scaling exponent increase is concurrent with an increase of the cell firing rate.

We also checked whether the membrane potential level V_{rest} can influence the frequency-scaling exponent. To do so, we varied the recorded cell membrane potential level by adjusting the synaptic strengths (see Materials and Methods). As for the firing rate, no significant influence in the frequency-scaling exponent can be attributed to the depolarization level (Fig. 8C, left panel), confirming the weak correlation observed *in vivo*.

Despite the lack of evidence for the "first-order hypothesis", our model does not incorporate biologically realisitic integrative features. It has been shown in previous studies [15,21,22] that the cell's intrinsic properties, shaped by its ionic channels, could have an impact on the V_m PSD form when the cell is submitted to noisy inputs. We performed the same analysis by replacing the integrate-and-fire model with a Hodgkin-Huxley model. The Na^+ and K^+ ionic channels could have an influence on the variation of the frequency-scaling exponent. However, adding these mechanisms did not alter the V_m frequency-scaling exponent's dependence on the input firing rate, nor on the mean postsynaptic membrane potential (Fig. 8B–C, middle panel). The results are identical to those obtained with the integrate-and-fire model. Controls were also performed with normally distributed synaptic weights for various standard deviations and gave identical results (Fig. S1A–B). On another set of controls, we changed the synaptic waveform by using synapses with a rise time on the order of 1 ms (β-synapse). The controls with this new type of synapse gave identical results to previous cases (Fig. S1C–D).

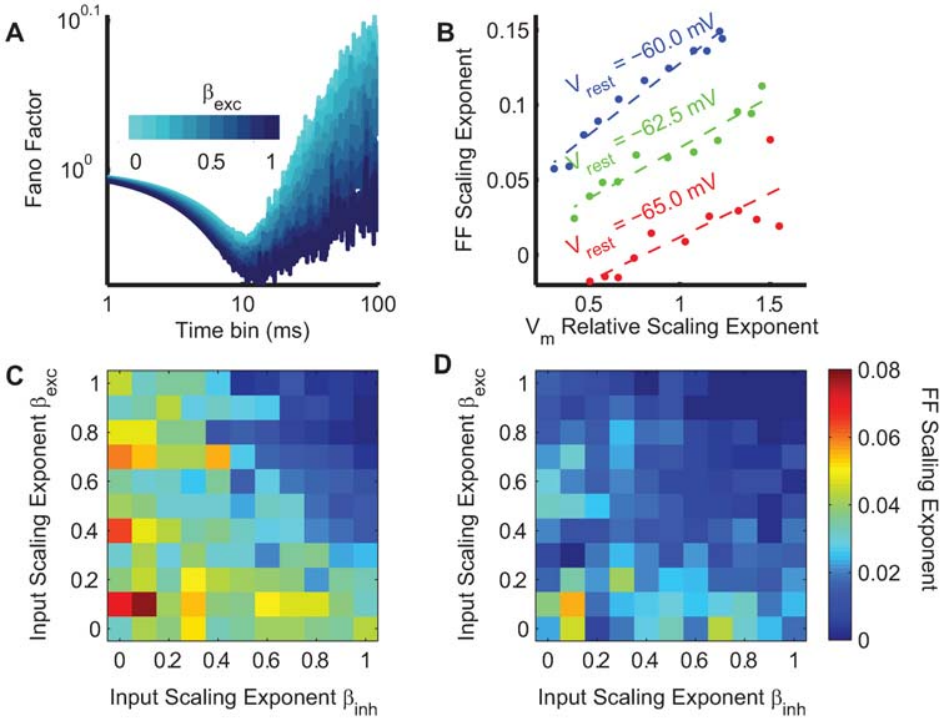

Figure 7. Relation between the V_m frequency-scaling exponent and that measured from the Fano Factor (FF) of the output spike train. A Example of the FF changes as a function of time bin, for different input parameters β_{exc}. The resting potential V_{rest} has been set to −60 mV to ensure a large enough number of spikes. The synchrony parameter is fixed at 6%. **B** Relation between spiking and relative V_m frequency-scaling exponents for different resting potentials (V_{rest} = −65 mV, −62.5 mV and −60 mV). **C,D** Fano Factor frequency-scaling exponents as a bivariate function of excitatory and inhibitory β_{exc} and β_{inh} parameters, in the absence of excitatory-inhibitory correlation and for and V_{rest} = −65 mV (C), and in the case of 40% of correlation and V_{rest} = −62.5 mV (D). In this latter case, V_{rest} has been increased by a few mV to ensure a reasonable level of spiking activity.
doi:10.1371/journal.pcbi.1000519.g007

Apart from the intrinsic mechanisms present in the somatic membrane, a possible source of modulation of the absolute value of the frequency-scaling exponent is the integrative property of the dendritic tree. To test how the dendritic arborization might impact the somatic subthreshold activity, we simulated synaptic input distributed in the dendrites of reconstructed pyramidal neurons. As shown in Table S1, the relative modulations of the exponent are well captured by correlation changes in the model, while global conductance changes had a negligible effect. However, it is important to note that these simulations were done using standard simulation tools (NEURON in this case), and thus used the standard cable equations. It has previously been shown that the standard cable equations cannot reproduce the correct frequency-scaling of the V_m PSD, and that taking into account the non-ideal character of the membrane capacitance could yield the correct frequency-scaling [18]. This could explain why the *in vivo absolute* values of the scaling exponent are not well reproduced here. However, the *relative* modulations of the exponent are well captured by correlation changes in the model, while global conductance changes had a negligible effect.

Dynamic-clamp experiments *in vitro*

Numerical simulations gave important insights about the role of intrinsic properties in the effects we see, but no computational model can guarantee an exhaustive exploration of such mechanisms. Indeed, even though the first-order hypothesis was invalidated for Hodgkin-Huxley models, we cannot exclude the influence of other ionic currents. Therefore, we performed the same test on real biological neurons through dynamic-clamp *in vitro*.

The correlated conductance traces generated by our model were directly injected into rat visual cortex neurons recorded *in vitro* (n = 9) using the dynamic-clamp technique (see Materials and Methods and Fig. 4B). We performed the same control as above changing the mean input firing rate. The frequency-scaling exponent barely changed (Fig. 8B, right panel; $r = -0.09$, $p \geq 0.3$), confirming that the overall presynaptic activity level has a negligible effect compared to the conductance correlations (characterized by the β parameter). Even the weak correlation observed between the mean input firing rate and the frequency-scaling exponent has the opposite sign to what is observed *in vivo*. The relative variation for different β has the same magnitude as the numerical models ($r = 0.92$).

The previous results were obtained for different resting membrane potentials and did not show any noticeable effect regarding the mean depolarization (Fig. 8B, right panel, $r = -0.002$, $p \geq 0.9$).

In order to measure the influence of the depolarization level on the frequency-scaling exponent, we systematically varied the conductance strength to change the mean V_m of the recorded cell. The frequency-scaling exponent did not exhibit significant variation (Fig. 8C, right panel). *In vitro* experiments thus confirm our previously observed results from numerical models.

In summary, the correlation in the activity impinging on the recorded cell plays a major role in determining the frequency-scaling exponent of the V_m. Other parameters, such as the total conductance (see also Fig. S3) and the balance between excitatory

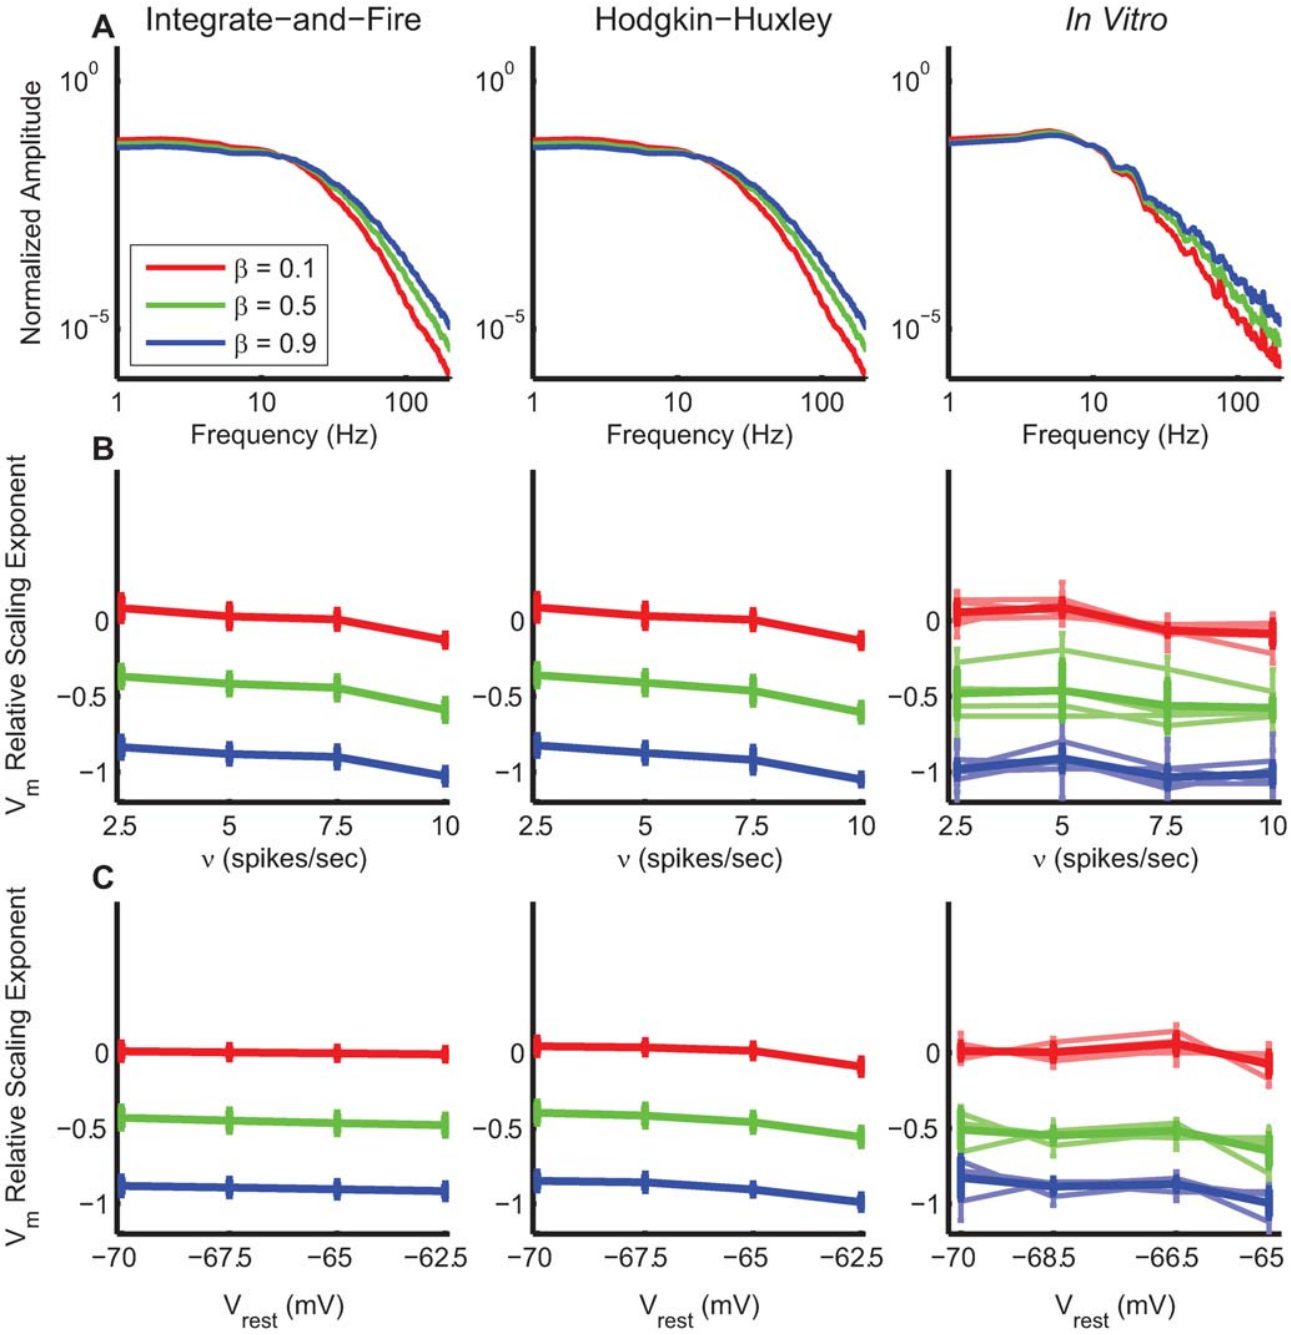

Figure 8. V_m frequency-scaling exponent changes for different input frequencies ν and for different resting membrane potential V_{rest}. These controls were performed with integrate-and-fire neurons (left column), Hodgkin-Huxley neurons (middle column) and with biological neurons during *in vitro* experiments (right column). The synchrony percentage was kept at 6% and there was no correlation between excitatory and inhibitory synaptic inputs. For the *in vitro* experiments, each light line represents a cell, for which ten trials have been repeated with the same parameters. Error bars are the standard deviation over the trials. The bold line represents the average across cells and trials. Note that the reference value subtracted to each measured exponent is the one obtained when the input parameter $\beta = 0.1$ to allow a direct comparison between models and *in vitro* data. **A** PSDs obtained for three values of $\beta = \beta_{exc} = \beta_{inh} \in \{0.1, 0.5, 0.9\}$. The modulation of the PSD slope is apparent. The absolute slope values are respectively (see Materials and Methods): -3.35 -3.82 and -4.4 (integrate and fire, left); -3.35 -3.82 and -4.4 (Hodgkin-Huxley, middle); -3.28, -3.7 and -3.92 (*in vitro*, right). **B** For three values of $\beta = \beta_{exc} = \beta_{inh} \in \{0.1, 0.5, 0.9\}$, the modulation of the V_m output frequency-scaling exponent according to the mean input firing rate per presynaptic neuron. **C** Same measures according to the postsynaptic resting membrane potential V_{rest}.

doi:10.1371/journal.pcbi.1000519.g008

and inhibitory conductances, have negligible effects. These results support the idea that changes in the frequency-scaling exponent observed *in vivo* reflect changes in correlations in the external stimulus-driven activity.

Discussion

In this paper we have analyzed the factors affecting power-law frequency-scaling in the membrane potential of cortical neurons. Our main findings are that (1) intracellular recordings of cat primary visual cortex neurons *in vivo* display power-law frequency-scaling at high frequencies, with a fractional exponent which depends on the spatio-temporal statistics of the visual stimuli; (2) this effect was reproduced in computational models of a recurrent network, and of single neurons of various degrees of complexity; the main determinant of the exponent was the correlation waveform in the presynaptic activity correlation; (3) other factors such as the conductance state had no effect on this measure. These findings were confirmed in cortical neurons *in vitro* using dynamic-clamp injection of random synaptic conductances with controlled degrees of correlation. We discuss below the implications of these findings and how they relate to previous work.

Influence of Network Correlations and Intrinsic Properties

Our central finding *in vivo* is that the frequency-scaling exponent in V1 is modulated by the visual stimulus statistics. Because such changes are detected in the same cells, they must necessarily reflect changes in the spatio-temporal structure of presynaptic activity. Guided by the fact that intracellular activity in sensory and prefontal cortex shows long lasting temporal correlations, we hypothesized that the main factor affecting frequency-scaling exponents is the correlation in presynaptic activity. This hypothesis was supported by numerical simulations. A similar modulation of the V_m frequency-scaling exponent was also found in a recurrent network for which the input correlation was varied : the scaling exponent increased when the input correlation increased above a certain threshold (required to be detectable). This threshold was not reached during decorrelated states, such as those imposed by surrogate natural scenes.

In the recurrent model, the input correlation modulated both the the absolute strength and temporal structure of correlations. To investigate separate modulations of these two factors, we chose a model of presynaptic inputs with a temporal power-law structure. This choice was made for two reasons: first, because this temporal structure was observed in our network model, without implementing any scaling in the connectivity; second, because it provided an operational way to control the form of the correlations in the input, and isolate which factors influence the output frequency-scaling exponent. The input is thus characterized by its frequency-scaling exponent, and we found that the V_m frequency-scaling exponent of the subthreshold output is linearly related to this input exponent. However, this relationship is present only if the correlation strength is large enough. According to these results, the V_m frequency-scaling exponent increase observed *in vivo* could plausibly be due to a global correlation strengthening in the surrounding network and/or by a narrowing of the spatial spread of correlation.

The hypothesis for a determinant role of correlations is also consistent with *in vitro* experiments, where we recreated artificial and controllable synaptic activity by dynamic-clamp. The fact that correlation changes are reflected by changes in the frequency-scaling exponent of the V_m frequency-scaling means that intrinsic cellular properties do not have major dynamical influences on this scaling, and that it mostly reflects synaptic activity controlled by the visual stimulation context. In particular, we showed that neither the mean level of synaptic bombardement nor the postsynaptic depolarization level could significantly modulate the V_m frequency-scaling exponent, even though the cell integrative properties shape its static absolute value [15–17,21].

A Signature of Avalanche Dynamics?

The finding that V_m activity presents power-law frequency-scaling is reminiscent of the power-law relationships of self-organized critical states, similar to those found from multi-site local-field potential recordings *in vitro* [8,28]. In the latter case, self-organized critical states are characterized by the production of "avalanches" of activity, whose size distribution follows a power-law. However, the power-law relations were found there in the frequency domain, which is very different from the distribution of event sizes detected in our study, so our results should not be taken as evidence for avalanche dynamics. We have performed an avalanche analysis on the recurrent network model, and as was reported in a previous study [29], we did not find evidence for avalanche type dynamics in the network during AI states.

Moreover, it has to be noted that the power-law relations found here depend on the stimulus, which means that the frequency-scaling exponent does not represent a unique signature of cortical network activity, but rather reflects a measure of the dynamic interplay between the sensory evoked activity and the ongoing recurrent network activity.

Relationship between the Subthreshold and Spiking Frequency-Scaling Exponents

Power-law frequency-scaling was reported previously in extracellularly-recorded spiking activity [6,30,31]. We observed that the V_m and spiking frequency-scaling exponents are linearly related. However, the exact value of the frequency-scaling of spiking activity critically depends on the V_m depolarisation level, and thus does not reliably reflect network correlation state. Our study shows that the V_m frequency-scaling exponent, which reflects the integration of thousands of input sources, can uncover features of the population activity that were not visible at the single cell spiking level or when assigning a limited number of cells at random.

Correlation States in Evoked and Spontaneous Activities

Our results imply that tracking the relative changes of the V_m frequency-scaling exponent could be a way to characterize dynamic changes in the correlations hidden in the global connectivity network, but read out at the subthreshold level by each member cell of these overlaid functional assemblies. Having interpreted the relative variations of the frequency-scaling exponent, we can now link these variations with the type of visual stimulus presented.

In order to emphasize the role of dynamic cortical non-linearities in the stimulus-dependency of the power-scaling, we checked whether or not these exponent changes were already apparent in the linear prediction of the V_m responses. To do so, we used the first-order kernel of the receptive field obtained by dense noise mapping to reconstruct linear predictions of the subthreshold dynamics for the different classes of stimuli and tested the contextual dependency of the spectral scaling properties of the linear predictor. The modulatory effects were not retrieved, which was expected since the estimation of the frequency-scaling exponent is performed on high frequencies (between 75 Hz and 200 Hz) that are not accounted for by the linear kernel (data not shown). We conclude that the exponent variations are not a linear

read-out of the scaling behaviour of the stimulus but rather the product of the non-linearities in the input-output relationship imposed by the cortical network.

According to our recurrent network study, the frequency-scaling exponent decreases when switching from DG stimuli to NI or DN stimuli should correspond to a decrease in the correlation strength. Following this interpretation, it could appear surprising that stimuli with very different structures, such as NI and DN stimuli, evoke similar values of the V_m scaling exponent. However, our study showed that the V_m scaling exponent is invariant to changes in the spatial structure of the input. As a consequence, stimuli with different spatial structures can evoke similar scaling exponents provided their global correlation levels are all low.

On the one hand, although it has not been demonstrated directly, natural movie stimuli probably induce decorrelation, for several reasons. First, our natural image is animated most of the time by fixational eye movements, which may already decorrelate activity at the LGN [32]. Second, the decorrelation theory [33] predicts that cortical responses to natural scenes should be decorrelated in order to maximize the transmitted information, and this prediction has been confirmed in V1 studies [34]. On the other hand, dense noise, as a fully uncorrelated stimulus, also evokes a very decorrelated response.

These low correlation levels for both stimuli are probably what make them indistinguishable from the perspective of the scaling exponent. In short, even if the structures of these inputs are very different, thalamic and cortical processing may reduce the initial correlations down to a similar level. Furthermore the scaling exponent captures neither the difference in the spatial structure of these resulting activities nor the difference in the low frequency band dominated by the stimulus spectrum. Taken together these arguments can explain why we observed similar scaling exponents. The same remark holds for DG and GEM stimuli: despite their difference in temporal structure, they might evoke similar levels of correlation, and thus similar scaling exponents, despite the difference in input spatial structure and low frequency content.

Finally, the same argument may explain why we found similar exponents for the spontaneous activity and the natural stimulus: for high frequencies, both exponents likely correspond to a very decorrelated activity, even if there might be a residual synchrony. Note however that this striking correlation between NI and AS is not necessarily present at lower frequencies.

Several studies have compared the structure of the spontaneous activity to that of the evoked activity. The spatial structure of the spontaneous activity measured with voltage-sensitive dye (VSD) imaging has been found to be similar to the DG-evoked activity [35,36], although this result could not be replicated in awake animals [37]. On the other hand, [38] found that the temporal correlations measured in multi-unit recordings seems to be similar for dense noise, natural scenes and spontaneous activity. Our results and a recent theoretical study [39] seem to be compatible with the latter observations. However, they are not necessarily in total contradiction with the VSD results since our measures concern different frequency bands: while we measured frequency-scaling exponents between 75 and 200 Hz, the VSD measures mostly concerned dye signal fluctuations at frequencies below 20 Hz. It thus appears most likely that V1 responses to natural scenes and spontaneous activity share similar correlation features in the high-frequency band.

We have shown that the frequency-scaling exponents measured in the intracellular activity can vary under the influence of the visual context for the same cell. Our model relates this modulation to a dynamic change in the network correlation state and could be associated to the underlying dynamic dimensionality [40]. Further studies need to address at the population level (LFP or VSD) how the frequency-scaling exponents of the network activity may vary with the stimulus context [41], and if such changes could be indicative of the detection of specific sensory statistics in the external drive or their spontaneous recall by the recurrent structure of the network.

Materials and Methods

Animal Experimentation

All *in vitro* and *in vivo* research procedures concerning the experimental animals and their care adhered to the American Physiological Society's Guiding Principles in the Care and Use of Animals, to European Council Directive 86/609/EEC and to European Treaties Series 123 and were also approved by the regional ethics committee "Ile-de-France Sud" (Certificate 05-003).

In vivo Preparation

Cells in the primary visual cortex of anaesthetized (Althesin) and paralyzed adult cats were recorded *in vivo* using sharp electrode (potassium methylsulfate 3 M, 70–100 MΩ) recordings (average $V_{rest} = -67$ mV, 0 nA) as described elsewhere [25,42]. Data processing and visual stimulation protocols used in-house software (G. Sadoc, Elphy, CNRS-UNIC).

Visual Stimulation

The analyzed data come from *in vivo* experiments to be presented in full in a companion paper (Baudot, Marre, Levy, Monier and Frégnac, submitted). Preliminary accounts have been given elsewhere [43,44]. Stimuli were displayed on a 21" CRT monitor with a 1024×768 pixel resolution and a 150 Hz refresh rate, with a background luminance of 12 cd/m^2. Receptive fields were mapped using sparse noise and classical tunings were determined by automated exploration. Intracellular responses were compared for four full-field visual stimuli of 10 s duration and increasing complexity (see Fig. 1): a) a drifting grating of optimal orientation, direction, and spatial and temporal frequencies (DG), b) the same optimal grating animated by a modeled eye-movement sequence (GEM), c) a natural image animated by the same virtual scanpath (NI), and d) dense binary white noise (DN). The mean luminance and contrast of each movie were equalized. Each movie was presented 10 times. For the NI condition, we used a high definition natural image (2048×1536 pixels) animated with a virtual eye movement sequence [43,44] (note that the size of the image is larger than the size of the screen, so that no blank region appears when the image is moved along the oculomotor trajectory). White noise consisted of a dynamic sequence (13.3 ms refresh period) of high spatial definition (50×50 pixels of side length 0.39°) binary dense noise.

Numerical Models

All the simulations (including dynamic-clamp experiments) were performed with the NEURON software [http://www.neuron.yale.edu] except for the recurrent model which was been run under NEST [45] using the PyNN interface [http://neuralensemble.org/PyNN]. A time step of $dt = 0.1$ ms was used systematically. We ran some simulations with $dt = 0.01$ ms to verify that our results were not dependent on the integration time step (data not shown).

The postsynaptic neuron follows an integrate-and-fire equation with conductance-based synapses whose time evolution is given by

$$\tau_m \frac{dV(t)}{dt} = (V_{leak} - V(t)) + g_{exc}(t)(E_{exc} - V(t)) + g_{inh}(t)(E_{inh} - V(t)) \quad (3)$$

with the resting membrane time constant $\tau_m = 20$ ms, the leak membrane potential $V_{leak} = -80$ mV and the excitatory and inhibitory conductances given in units of leak conductance $G_{leak} = 10$ nS. When $V(t)$ reaches the spiking threshold $V_{thresh} = -50$ mV, a spike is generated and the membrane potential is reset to $V_{reset} = -60$ mV for a refractory period of duration $\tau_{ref} = 5$ ms. $E_{exc} = 0$ mV and $E_{inh} = -70$ mV are the reversal potentials for the excitatory and inhibitory exponential synapses $syn = \{exc, inh\}$ whose dynamics follow

$$\tau_{syn} \frac{dg_{syn}(t)}{dt} = -g_{syn}(t) + \Delta g_{syn} S_{syn}(t) \quad (4)$$

where τ_{syn} is the synaptic time constant with $\tau_{exc} = 3$ ms and $\tau_{inh} = 7$ ms. Δg_{exc} and Δg_{inh} are the quantal synaptic strengths elicited by each presynaptic spike and $S_{syn}(t)$ is the point process modelling the incoming spike train. Δg_{exc} and Δg_{inh} are chosen in order to satisfy the ratio $< g_{exc} > + < g_{inh} > = 3$ where the brackets signify an average according to $S_{syn}(t)$, and so that the effective resting potential is $V_{rest} = -65$ mV on average. Identical results were been obtained for synapses with a finite rise time (β-synapses). Parameters for the Hodgkin-Huxley model were taken from [46].

The recurrent network is composed of 10000 excitatory and 2500 inhibitory neurons, sparsely connected, with a connection probability of 2% within each population and between the two populations. The synaptic weights are $\Delta g_{exc} = 4.0$ nS and $\Delta g_{inh} = 85.0$ nS. Each neuron has a topographic position on a cortical layer-like surface of 1 mm^2, and connects to its neighbours according to a Gaussian distribution of standard deviation $\sigma_c = 0.15$ mm. Periodic boundary conditions are used. Conduction delays d are distant-dependent with $d(x) = 0.5 + 5x$ (ms) where x is the distance between the two neurons expressed in millimetres. The slope value of $d(x)$ (giving a propagation speed of 0.2 mm/ms) is taken from a previous *in vivo* study showing a lateral propagation speed ranging dominantly between 0.1 and 0.3 mm/ms [42]. The retinotopic drive was modelled as another thalamic layer-like network facing the previous one where each neuron acts as a Poisson process with a controlled amount of synchrony between the firing. To mimic a retinotopic mapping, each cell in the thalamic layer projects to the recurrent network in a topographically organized manner following a Gaussian distribution of standard deviation $\sigma_t = 0.05$ mm (Fig. 3). The connection probability from the thalamic layer to the cortical layer is also 2%.

In some simulations, we used models based on morphologically-reconstructed neurons from cat cortex, obtained from two published reference studies (layer II–III of cat primary visual cortex Douglas et al. [47]; layer VI of cat somatosensory cortex Contreras et al. [48]), where biological details were given. The three-dimensional morphology of the reconstructed neurons was incorporated into the NEURON simulation environment, which enables simulating cable equations in complex three-dimensional structures [49]. *In vivo*-like activity was simulated in passive models using a previously published model of synaptic bombardment at excitatory and inhibitory synapses [50] (see this paper for details about the parameters and numerical simulations). The density of synapses was constant per unit membrane area according to published morphological studies, and was (per 100 μm^2): 60 for dendritic AMPA synapses, 10 for dendritic $GABA_A$ and 20 for somatic $GABA_A$ synapses. This gives 9947 AMPA and 2461 $GABA_A$ synapses for the layer II–III cell, and 16563 and 3376, respectively, for the layer VI cell. The release rates, chosen to yield synaptic bombardment consistent with *in vivo* measurements, were $\nu_{exc} = 1$ Hz and $\nu_{inh} = 5.5$ Hz for AMPA and GABAergic synapses, respectively (see details in [50]).

Correlation Generator

In order to produce spike trains with arbitrary temporal correlations, we used the theory of cluster point processes [23,51]. The presynaptic activity can be characterized by two main features: on the one hand, the specific temporal structure given by the spike train temporal auto-correlation form, and on the other hand, the correlation strength which measures the temporal coherence between individual presynaptic spike trains (see [52] for a similar distinction). These two features can be controlled separately in the spike train generator composed of a population of presynaptic neurons following Poisson processes, and firing together with a certain amount of synchrony. They project to the postsynaptic neuron through different time delays, randomly chosen from a specific distribution (Fig. 4). The temporal structure is given by the delay distribution whereas the global synchrony in the presynaptic neuronal discharge gives the correlation strength. In our implementation, the presynaptic population is assumed to contain N neurons ($N_{exc} = 4000$ for the excitatory population and $N_{inh} = 1000$ for the inhibitory population, except stated otherwise); at each time step it was decided randomly whether or not some neurons will fire. The probability was adjusted to give a mean firing rate ν of the inputs. If so, $k+1$ neurons were chosen randomly to fire among the N constituting the population. This method allows to have always $k+1$ synchronous neurons, and still an apparent Poisson discharge at rate ν for each presynaptic neuron taken individually. Note that this gives back independent Poisson spike trains when $k = 0$. Correlation between excitatory and inhibitory neurons is implemented in the same manner. The delays are then attributed to each presynaptic spike train according to the chosen delay distribution.

From point process theory, this can be seen as two nested point processes. The first point process follows a Poisson process which determines the cluster positions and the second one determines randomly the position of $k+1$ points within each cluster according to an arbitrary density probability function. The correspondance between both representations is straightforward and the power spectrum density can be computed analytically with the Neyman-Scott equation [23,24,51]

$$\gamma(\omega) = \frac{N\nu}{2\pi} |\alpha(\omega)|^2 (1 + k|p(\omega)|^2) \quad (5)$$

where $p(\omega)$ is the Fourier transform of the delay distribution, $k+1$ is the number of synchronous neurons and $\alpha(\omega)$ is the Fourier transform of the synaptic filtering. In Eq. 5, the factor k can also be written $k = rN - 1$ where r is the ratio of synchronous neurons which does not depend anymore on N.

In this paper, we are interested in the power-law frequency-scaling in the temporal power spectrum density (PSD). Eq. 5 relates the delay distribution to the PSD so that a power-law behaviour at the conductance level needs a power-law scaling in the delay distribution. Therefore, the delay associated with each synapse was randomly chosen from a distribution proportional to $\frac{1}{t^\beta} \exp(-\frac{t}{\tau_{max}})$. The exponential term is added to avoid oscillations in the PSD due to an abrupt cut-off [6] with $\tau_{max} = 10$ ms. The parameter β is varied over the simulations and modulates the spread of temporal correlations. The presynaptic neurons are synchronously active according to the parameter k. The output frequency-scaling exponent (to be defined below) measured in the PSD (Eq.5) is thus equal to $2(1-\beta)$.

In vitro Preparation

In vitro experiments were performed on 350 μm-thick sagittal slices from the lateral portions of rat occipital cortex. Wistar Rats, 4–6 weeks old (CNRS, Gif-sur-Yvette), were anesthetized with sodium pentobarbital (30 mg/kg) before craniectomy and cortex removal. The slices were maintained in an interface style recording chamber at 34–35°C. Slices were prepared on a DSK microslicer (Ted Pella, Redding, CA) in a slice solution in which the NaCl was replaced with sucrose while maintaining an osmolarity of 314 mosM. During recording, the slices were incubated in slice solution containing (in mM) 126 NaCl, 2.5 KCl, 1.2 $MgSO_4$, 1.25 $NaHPO_4$, 2 $CaCl_2$, 26 $NaHCO_3$, and 25 dextrose and aerated with 95% O_2-5% CO_2 to a final pH of 7.4. After 30 minutes to 2 hours of recovery, intracellular recordings were performed in deep layers (layer IV–VI) in electrophysiologically identified regular spiking and intrinsically bursting cells. Micropipettes were filled with 1.2–2 M potassium acetate and 4 mM KCl and had resistances of 80–100 M after bevelling. The dynamic-clamp technique [53,54] coupled with an Active Electrode Compensation (AEC) method that we developed and validated recently *in vivo* and *in vitro* [24] was used to inject computer-generated conductances in real neurons. The AEC method allows the removal in real time of electrode noise from intracellular voltage recordings. Dynamic-clamp experiments were run using the Real Time-NEURON environment [55], which is a modified version of NEURON 6.0 [49].

The dynamic-clamp protocol was used to insert the fluctuating conductances underlying synaptic noise in cortical neurons using the previous model, the post-synaptic neuron being now the recorded neuron, similar to a previous study [56]. The injected current is determined from the fluctuating excitatory and inhibitory conductances as well as from the difference of the membrane voltage from the respective reversal potentials.

Power Spectrum Analysis

Spikes were removed from the original traces and replaced by a low-pass filtered version of the trace. To control the validity of this procedure, we compared whenever possible the power spectrum obtained from the interpolated trace with an identical trace generated without threshold. In all cases we observed that injecting a given conductance trace into a neuronal model and then removing the spikes gave the same power spectrum as injecting the same conductance in a neuronal model without spike threshold (Fig. S2). The spectra were computed with the multi-taper method [57], which allows a better estimation of the power-laws than the standard periodogram methods. Results were similar when using the Welch method and the Goertzel algorithm [58].

We then determined the frequency-scaling exponent by linear regression on a log-log representation of the PSD, for the range 75–200 Hz. Similar results were obtained for lower bounds above 50 Hz, and higher bounds below 200 Hz. Estimation of the scaling exponent from multifractal methods gave similar values. For the *in vitro* data, we also estimated the frequency-scaling exponent by fitting a generalized Lorentzian function [59], which gave equivalent relative values.

We chose to use the linear fit for its simplicity, and because it is easy to quantify the goodness of fit, and thus to assess the power-law scaling over the frequency band chosen. In comparison, the Lorentzian fit is very accurate when considering controlled models where the cut-off frequencies can be easily found or computed, but this model gave inaccurate results when applied to *in vivo* data because it can not account for the low frequency regime, which is strongly modulated by the stimulus. Finally, the multifractal analysis gave us no control over the goodness of fit. In the case of the recurrent network, the fit was performed between 75 and 200 Hz. Using narrower bands gives similar results. In the *in vitro* measurements, the absolute values of the frequency-scaling exponent displayed significant variations because of the available scaling region. Our study focused on the *modulation* of the frequency-scaling, rather than on absolute values, the relative values of the frequency-scaling exponent are shown for *in vitro* experiments and the corresponding models for each linear region of the PSD. For the model studies, unless otherwise mentioned, we systematically subtracted the value obtained for a classical Poisson input. For the *in vitro* study, the reference was the frequency-scaling exponent obtained with the input parameter $\beta = 0.1$, averaged over the different conditions tested. In this case, measuring the relative values also removed the cell-to-cell variability of the absolute values.

The total input conductance is reported to be about three times the leak conductance G_{leak} in the anaesthetized cat [26]. This is also what we used in our model and in the conductance injection *in vitro*. As a consequence, the cut-off frequency of the synaptic and membrane filtering are below the frequency band used for our fitting (they did not exceed 75 Hz), and could not affect our estimates (this point is futher discussed in the Results section).

Multifractal Analysis

The multifractal analysis characterizes the scaling behavior of a signal $x(t)$ [60]. For each point t_0, the Hölder exponent $H(t_0)$ is defined as the maximal value α such that there exists a polynomial $P(t)$, with $Deg(P) \leq [\alpha]$, a positive constant C, and an interval around t_0 where for any t

$$|x(t) - P(t - t_0)| \leq C|t - t_0|^{\alpha} \tag{6}$$

This coefficient $H(t_0)$ reflects the scaling behaviour around the point t_0. The singularity spectrum $D(h)$ is the Haussdorf dimension of $\{t : H(t) = h\}$. It thus describes how the singularities are distributed in the signal. A particular example is the self-similar process (also called monofractal), where $D(h) \neq 0$ only at one point H, where $D(H) = 1$. The practical estimation of the singularity spectrum is made difficult by the finite size of the signal, and by its discrete nature. However, the wavelet formalism allows a robust estimation of $\tau(q)$, which is the Legendre transform of the singularity spectrum:

$$D(h) = \min_q \{qh - \tau(q)\} \tag{7}$$

In the case of a monofractal/scale-invariant process, $\tau(q) = qH - 1$, H being its unique Hölder exponent. This corresponds to a fractional Brownian process. Note that H is related to the PSD slope which is equal to $-2H - 1$. The curvature of $\tau(q)$ quantifies the deviation from monofractality. The slope and the curvature are respectively the first and second moments of the singularity spectrum. We used an algorithm based on wavelet leaders [20,61] which directly estimates these two values.

Fano Factor and Power-Law in the Spiking Activity

Fano factors and power-laws on these Fano factors were measured as in [6]. To compute the Fano Factor for a given time bin, we counted the number of spikes in each time bin and took the ratio of the spike-count variance to the mean spike-count. The power-law was estimated by computing this Fano Factor over a large range of time bins. This function was then represented in a log-log scale, and the slope of the curve was estimated by linear regression. This gives the frequency-scaling exponent of the

spiking activity through the Fano Factor $F(T) \propto T^{\alpha}$ where T is the time bin and α the scaling exponent.

Supporting Information

Figure S1 Effect of heterogeneous synaptic weights and synaptic waveform on the power law frequency scaling exponent. (A–B) V_m frequency-scaling exponent changes for different input frequencies ν and for heterogeneous synaptic strengths. The synaptic strengths are randomly distributed for each incoming synaptic spike train according to a Gaussian distribution whose standard deviation is half the mean value in this case. These controls were performed with integrate-and-fire neurons (panel A) and Hodgkin-Huxley neurons (panel B). The synchrony percentage was kept a 6% and there was no correlation between excitatory and inhibitory synaptic inputs. Error bars are the standard deviation over the trials. The bold line represents the average across cells and trials. (C–D) Variation of the value of the frequency-scaling exponent at the membrane potential level for excitatory input only as a function of the parameters β_{exc} and for β-synapses ($r = 3\%$). (C) Illustration of the PSD modulation on a log-log scale for different values of the parameter β_{exc} ranging from 0 (light blue) to 1 (dark blue). In the inset, a stereotypic synaptic time course is represented (with a time rise of 1 ms). (D) Variation of the output frequency-scaling exponent with the β_{exc} parameter.
Found at: doi:10.1371/journal.pcbi.1000519.s001 (0.41 MB EPS)

Figure S2 Illustration of the spike filtering algorithm for neuron models with and without spiking mechanism. (A) Injection of correlated synaptic input to a HH model. Blue: raw trace; Red: after spike filtering. (B) Power spectra density corresponding to panel A. (C) injection of the same synaptic input in a COBA model without threshold (green), superimposed to the HH-spike-filtered trace plotted in panel A. (D) Power spectra density of the the two traces displayed in panel C: COBA without threshold and HH with spike filtered.
Found at: doi:10.1371/journal.pcbi.1000519.s002 (0.97 MB EPS)

Figure S3 Influence of the different integrative time constants on the PSD frequency scaling. (A) V_m power spectra for different levels of correlation in the input (blue: Poisson input; red: correlated input with $k = 6\%$ and $\beta = 0$). The level of conductance is low in this condition ($G_{tot} = 0.23G_{leak}$). The dotted coloured lines indicate the linear fits over the high frequency region delimited by the vertical dashed gray line. (B) same PSD, but for a very high conductance state ($G_{tot} = 12G_{leak}$). The four fits correspond to fit in different frequency bands, for the two PSDs. To illustrate more precisely the differential effect of the conductance state and of the input correlations on the frequency-scaling exponent, we show several examples of V_m power spectra for two different levels of global conductance regime, and two different $\beta_{inh} = \beta_{exc}$ parameters. In the low conductance state (panel A), the power spectrum is composed of two linear regions separated by a unique cut-off, which is determined by the time constants of the synaptic and membrane filtering. In the very high conductance state (panel B), these two time constants are clearly different, so the power spectrum shows three linear regions separated by two cut-offs. Very large (and surely not plausible in biological conditions) changes of the conductance state thus displaced the second frequency cut-off, but still did not affect the relative slope in the linear regions. Decreasing the β parameter increases the slope over both frequency bands and relative changes of the frequency-scaling exponent have the same magnitude in these different regions. This shows that the relative modulation observed is not dependent on the specific frequency band chosen to estimate the PSD slope, since it can be observed over a large range of frequencies. Furthermore, this figure illustrates the differential effect of the conductance state, and of the correlation state, on the power spectrum. Opposite to the latter, the former does no affect the scaling exponent.
Found at: doi:10.1371/journal.pcbi.1000519.s003 (0.52 MB EPS)

Table S1 Frequency-scaling exponents for detailed neuron models. Neuron models were obtained from neuronal morphologies reconstructed from a layer III cell (upper table) and a layer VI cell (lower table) of the cat cerebral cortex (see methods). The frequency-scaling exponent is computed for different synaptic input firing rates and different levels of synchrony. Three levels of incoming synaptic activity have been considered, following (Destexhe & Paré, 1999) : a high-conductance state (HC) with $\nu_{exc} = 1$ Hz, $\nu_{inh} = 5.5$ Hz; a low-conductance state (LC) with $\nu_{exc} = \nu_{inh} = 0.5$ Hz and a very low-conductance state (VLC) with $\nu_{exc} = \nu_{inh} = 0.1$ Hz. Each condition was performed with two levels of synchrony between synaptic spike trains, $r = 0\%$ and $r = 1.5\%$ respectively. Frequency-scaling exponents barely changed with increasing firing rate for both uncorrelated and correlated inputs, for both cells. However, the frequency-scaling exponent was affected by the level of synchrony, as expected from our previous results. These simulations show that the relative modulations of the scaling exponent are mostly due to correlation changes, while conductance changes have a negligible effect.
Found at: doi:10.1371/journal.pcbi.1000519.s004 (0.01 MB PDF)

Acknowledgments

We thank Gabriel Peyré and Philippe Ciuciu for help with the multifractal analysis, Cyril Monier for helpful discussions, Andrew Davison and Emily Caporello for comments on the manuscript and Axel Corjon for graphic assistance.

Author Contributions

Conceived and designed the experiments: SEB OM PB TB YF. Performed the experiments: OM SB PB TB. Analyzed the data: SEB OM PY AD. Contributed reagents/materials/analysis tools: SEB OM PY TB AD YF. Wrote the paper: SEB OM AD YF.

References

1. Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233: 1416–1419.
2. Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci 18: 3870–3896.
3. Abeles M (1991) Corticonics: neural circuits of the cerebral cortex. Cambridge: Cambridge University Press.
4. Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, et al. (1995) Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. Nature 373: 515–518.
5. Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. Annu Rev Neurosci 18: 555–586.
6. Teich MC, Heneghan C, Lowen SB, Ozaki T, Kaplan E (1997) Fractal character of the neural spike train in the visual system of the cat. J Opt Soc Am A Opt Image Sci Vis 14: 529–546.
7. Maxim V, Sendur L, Fadili J, Suckling J, Gould R, et al. (2005) Fractional gaussian noise, functional mri and alzheimer's disease. Neuroimage 25: 141–158.
8. Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. J Neurosci 23: 11167–11177.
9. Plenz D, Thiagarajan TC (2007) The organizing principles of neuronal avalanches: cell assemblies in the cortex? Trends Neurosci 30: 101–110.
10. Drew PJ, Abbott LF (2006) Models and properties of power-law adaptation in neural systems. J Neurophysiol 96: 826–833.

11. Lundstrom BN, Higgs MH, Spain WJ, Fairhall AL (2008) Fractional differentiation by neocortical pyramidal neurons. Nat Neurosci 11: 1335–1342.
12. Fusi S, Drew PJ, Abbott LF (2005) Cascade models of synaptically stored memories. Neuron 45: 599–611.
13. Destexhe A, Rudolph M, Paré D (2003) The high-conductance state of neocortical neurons in vivo. Nat Rev Neurosci 4: 739–751.
14. Rudolph M, Pelletier JG, Pare D, Destexhe A (2005) Characterization of synaptic conductances and integrative properties during electrically induced eeg-activated states in neocortical neurons in vivo. J Neurophysiol 94: 2805–2821.
15. Diba K, Lester HA, Koch C (2004) Intrinsic noise in cultured hippocampal neurons: experiment and modeling. J Neurosci 24: 9723–9733.
16. Jacobson GA, Diba K, Yaron-Jakoubovitch A, Oz Y, Koch C, et al. (2005) Subthreshold voltage noise of rat neocortical pyramidal neurones. J Physiol 564: 145–160.
17. Yaron-Jakoubovitch A, Jacobson GA, Koch C, Segev I, Yarom Y (2008) A paradoxical isopotentiality: a spatially uniform noise spectrum in neocortical pyramidal cells. Front Cell Neurosci 2: 3.
18. Bedard C, Destexhe A (2008) A modified cable formalism for modeling neuronal membranes at high frequencies. Biophys J 94: 1133–1143.
19. Destexhe A, Rudolph M (2004) Extracting information from the power spectrum of synaptic noise. J Comput Neurosci 17: 327–345.
20. Lashermes B, Jaffard S, Abry P (2005) Wavelet leader based multifractal analysis. In Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '05). volume 4. pp iv/161–iv/164.
21. Manwani A, Koch C (1999) Detecting and estimating signals in noisy cable structure, i: neuronal noise sources. Neural Comput 11: 1797–1829.
22. Steinmetz PN, Manwani A, Koch C, London M, Segev I (2000) Subthreshold voltage noise due to channel fluctuations in active neuronal membranes. J Comput Neurosci 9: 133–148.
23. Daley DJ, Vere-Jones D (2008) An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods, 2nd Edition Springer.
24. Brette R (2009) Generation of correlated spike trains. Neural Comput 21: in press.
25. Monier C, Chavane F, Baudot P, Graham LJ, Frégnac Y (2003) Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. Neuron 37: 663–680.
26. Monier C, Fournier J, Frégnac Y (2008) In vitro and in vivo measures of evoked excitatory and inhibitory conductance dynamics in sensory cortices. J Neurosci Methods 169: 323–365.
27. Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. Nat Neurosci 11: 535–537.
28. Beggs JM, Plenz D (2004) Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. J Neurosci 24: 5216–5229.
29. El Boustani S, Pospischil M, Rudolph-Lilith M, Destexhe A (2007) Activated cortical states: experiments, analyses and models. J Physiol Paris 101: 99–109.
30. Kim SI, Jeong J, Kwak Y, Kim YI, Jung SH, et al. (2005) Fractal stochastic modeling of spiking activity in suprachiasmatic nucleus neurons. J Comput Neurosci 19: 39–51.
31. Gebber GL, Orer HS, Barman SM (2006) Fractal noises and motions in time series of presympathetic and sympathetic neural activities. J Neurophysiol 95: 1176–1184.
32. Desbordes G, Rucci M (2007) A model of the dynamics of retinal activity during natural visual fixation. Vis Neurosci 24: 217–230. doi:10.1017/S0952523807070460.
33. Atick JJ (1992) Could information theory provide an ecological theory of sensory processing? Network 3 (2): 213–251.
34. Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. Science 287: 1273–1276.
35. Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. Science 286: 1943–1946.
36. Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. Nature 425: 954–956.
37. Omer D, Rom L, Ultchin Y, Grinvald A (2008) The dynamics of ongoing activity in awake and anesthetized monkey are significantly different. Soc Neurosci Abstracts 769.9.
38. Fiser J, Chiu C, Weliky M (2004) Small modulation of ongoing cortical dynamics by sensory input during natural vision. Nature 431: 573–578.
39. Marre O, Yger P, ADavison, Frégnac Y (2009) Reliable recall of spontaneous activity patterns in cortical networks. Journal of Neuroscience in press.
40. Krakovská A, Stolc S (2008) Spectral decay vs. correlation dimension of eeg. Neurocomputing 71: 2978–2985.
41. Ciuciu P, Abry P, Rabrait C, Wendt H, Roche A (2007) Leader-based multifractal analysis for evi fmri time series: Ongoing vs task-related brain activity. In Proc. 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro ISBI 2007. pp 404–407.
42. Bringuier V, Chavane F, Glaeser L, Frégnac Y (1999) Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. Science 283: 695–699.
43. Baudot P, Levy M, Monier C, Chavane F, René A, et al. (2004) Time-coding, low noise vm attractors, and trial-to-trial spiking reproducibility during natural scene viewing in v1 cortex. Society for Neuroscience Abstracts 948: 12.
44. Frégnac Y, Baudot P, Levy M, Marre O (2005) An intracellular view of time coding and sparseness in v1 during virtual oculomotor exploration of natural scenes. 2nd International Cosyne Conference in Computational and Systems Neuroscience Salt Lake City, Utah.17 p.
45. Diesmann M, Gewaltig M (2001) NEST: An environment for neural systems simulations. Forschung und wisschenschaftliches Rechnen, Beitrage zum Heinz-Biling-Preis 58: 43–70.
46. Destexhe A (1997) Conductance-based integrate-and-fire models. Neural Comput 9: 503–514.
47. Douglas R, Martin K, Â Whitteridge D (1991) An intracellular analysis of the visual responses of neurones in cat visual cortex. J Physiol 440: 659–696.
48. Contreras D, Destexhe A, Steriade M (1997) Intracellular and computational characterization of the intracortical inhibitory control of synchronized thalamic inputs *in vivo*. J Neurophysiol 78: 335–350.
49. Hines M, Carnevale N (1997) The neuron simulation environment. Neural Computation 9: 1179–1209.
50. Destexhe A, Paré D (1999) Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. J Neurophysiol 81: 1531–1547.
51. Hohn N, Veitch D, Abry P (2003) Cluster processes, a natural language for network traffic. IEEE Transactions on Signal Processing, Special Issue on Signal Processing in Networking 51: 2229–2244.
52. Moreno-Bote R, Renart A, Parga N (2008) Theory of input spike auto- and cross-correlations and their effect on the response of spiking neurons. Neural Comput 20: 1651–1705.
53. Robinson HP, Kawai N (1993) Injection of digitally synthesized synaptic conductance transients to measure the integrative properties of neurons. J Neurosci Methods 49: 157–165.
54. Sharp AA, O'Neil MB, Abbott LF, Marder E (1993) Dynamic clamp: computer-generated conductances in real neurons. J Neurophysiol 69: 992–995.
55. Sadoc G, Le Masson G, Foutry B, Le Franc Y, Piwkowska Z, et al. (2009) Recreating in vivo-like activity and investigating the signal transfer capabilities of neurons: dynamic-clamp applications using real-time NEURON. New York: Springer, in press.
56. Destexhe A, Rudolph M, Fellous JM, Sejnowski TJ (2001) Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. Neuroscience 107: 13–24.
57. McCoy E, Walden A, Percival D (1998) Multitaper spectral estimation of power law processes. IEEE Trans Sig Proc 46: 655–668.
58. Welch PD (1967) The use of fast fourier transforms for the estimation of power spectra: A method based on time averaging over short modified periodograms. IEEE Transactions on Audio and Electroacoustics 15: 70–73.
59. Piwkowska Z, Pospischil M, Brette R, Sliwa J, Rudolph-Lilith M, et al. (2008) Characterizing synaptic conductance fluctuations in cortical neurons and their influence on spike generation. J Neurosci Methods 169: 302–322.
60. Stanley H, Meakin P (1988) Multifractal phenomena in physics and chemistry. Nature 335: 405–409.
61. Wendt H, Abry P (2007) Multifractality tests using bootstrapped wavelet leaders. EEE Trans Signal Processing 55(10): 4811–4820.

8 Information transmission under ongoing activity

Ongoing activity is problematic for understanding how information is encoded and transmitted in the brain. Its spontaneous, irregular and chaotic nature observed *in vivo* appears incompatible with response reliability and time coding (Banerjee et al., 2008). In a seminal experiment, Salzman et al. (1990) showed how the response of a monkey to a perceptual judgement of motion direction can be biased by cortical microstimulations *in vivo*. Small pulses of current were injected during the task to perturb the decision of the animal. The decrease of its performance shows how sensitive the system is. Similarly, the work of Li et al. (2009) showed how a small pulse of current injected in the thamalus could drastically change the dynamical regime of the cortex. This over-sensitivity of the system to perturbations is striking, and supports the chaotic nature of the spontaneous activity.

To transmit information, the classical way to cope with this situation is to consider the chaotic background as noise, uncorrelated with the simulation, and average over many responses to get rid of this "artefact" in a rate-based coding strategy. This make sense if the cortex is using population codes to encode information: a neuron itself is unreliable and may not be efficient at encoding information, so averaging over a population of cells doing the same computations may help to extract the informative part of the signal, enhancing the reliability of the code, and dealing with the noise. This linear model of a deterministic signal corrupted by a noise component justifies, in this context, a code based on averaging across time and/or neuronal assemblies (Shadlen and Newsome, 1998) and has been applied at different scales of integration ranging from intracellular recordings (Azouz and Gray, 1999) through optical imaging (Arieli et al., 1996), to LFP recordings (Deweese and Zador, 2004). Recent experimental results suggest that sensitivity of the system, due to the chaotic nature of the ongoing activity, should be an argument in favour of more rate-based codes, able to cope with the situation (London et al., 2010).

8.1 Ongoing activity and information propagation

Information propagation is too often considered from a feed-forward point of view, explained in Part I: in spite of the dominance of recurrent connections at the anatomical level, most models of cortical dynamics assume that layers of neurons are connected and arranged, and one usually studies the propagation of external inputs across the layers (Reyes, 2003, Diesmann et al., 1999). Nevertheless, ongoing activity challenges the classical feed-forward view of a silent brain waiting for incoming sensory inputs. This activity should more be seen as an internally generated filter selecting some particular inputs, by creating a competition between feed-forward sensory input streams and recurrent interactions. The idea of the ongoing activity as a filter comes from the fact that, for example, 94% of the synapses bombarding a neuron in V1 come from the cortex, and only 6% from the feed-forward thalamic pathway. Despite the fact that thalamo-cortical synapses are more effective and stronger than cortico-cortical ones (Gil et al., 1999), the cortex is a very recurrent and self-sustained dynamical system, wandering on the high-dimensional attractor of its own, self-generated ongoing activity. Feed-forward sensory inputs perturb this dynamic, but the fine interplay between both should be more carefully understood. In this spirit, the theory of predictive coding by Rao and Ballard (1999) can be applied to the relationship and the recurrence between cortex and thalamus. The cortex could build internal predictions and by gating the thalamic inputs with the numerous projections sent back from layer 6, may encode only the differences between

information which is expected and the sensory inputs. For evidence of the role of cortical feedback on the thalamus, see the work of Andolina et al. (2007), which shows how reliability in the thalamus is driven by the feedback sent by the cortex. Indeed, from an information transmission point of view, this would be more efficient to code only the differences between what is expected by the system, and what is coming from the sensory modalities. Such a coding scheme has been found in the electro sensory lobe of the Electric Fish (Sawtell and Bell, 2008), and in the retina (Hosoya et al., 2005), but evidences in cortex are still unclear.

Information propagation in recurrent networks is highly problematic. In balanced random networks (see Part I), one can observe that the ongoing "noise" generated by the recurrent interactions tends to disrupt the structure of the external signal. For example, Figure 13, taken from Vogels and Abbott (2005), shows that when an increase of firing rate is presented to a sub-population of cells in a random balanced network of conductance based integrate-and-fire neurons, with an ongoing activity near 10 Hz, only the cells that are directly stimulated keep a trace of the input. The "second" layer, i.e. neurons receiving inputs from these directly stimulated cells, is almost insensitive to the primary input. The only way the authors could achieve signal propagation was to artificially increase the strength of the synapses along the pathway, and build a strong feed-forward structure within the recurrent network. However, this led the network into pathological states, such as the Synchronous Regular regime. More and more evidence suggests that ongoing activity should be taken into account in thinking about new computational paradigms (see Ringach (2009) for a review).

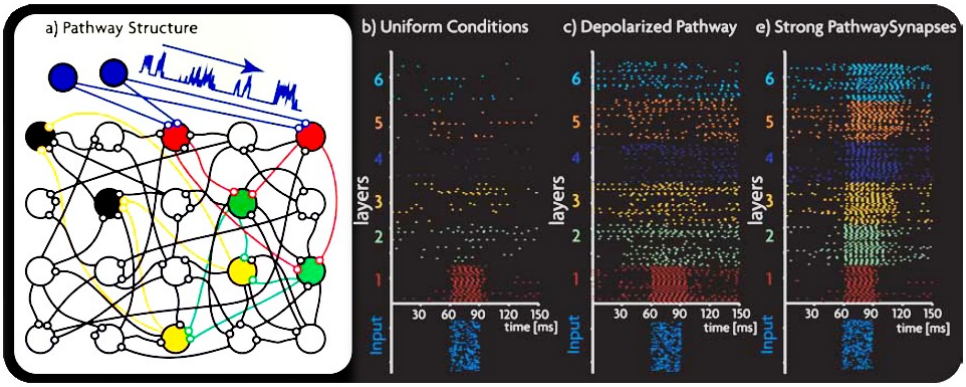

Figure 13: Adapted from Vogels and Abbott (2005). Signal propagation. a) Network diagram showing the layers of a candidate pathway. Input (blue) is fed into the network through strong synapses onto layer 1 neurons (red). In this and the following diagrams, layers 1– 6 are indicated by the colours green, yellow, dark blue, orange, and light blue, respectively. The white-filled circles denote non pathway neurons of the network. For this figure, layer 0 activity consists of a 30 ms pulse of activity at $\simeq$ 180 Hz. b) In a network with uniform excitatory and inhibitory synaptic strengths and neuronal parameters, no propagation occurs. c) Depolarization of pathway neurons by 15mV fails to induce propagation, although firing rates in all affected cells increase. e) Strengthening of pathway synapses by $\simeq$ 10-fold results in signal propagation and spreading.

8.2 Reliability as a function of the stimulus

Plenty of evidence shows that brain responses are, from trial to trial, rather "noisy" and unreliable. If the exact same stimulus is presented twice to the animal, spiking responses can

with difficulty be compared, and cross-correlations within trials are rather low. Nevertheless, unpublished work of P. Baudot *et al* (submitted to J NeuroPhysiol) in Frégnac's lab at UNIC, has demonstrated that the response variability of the same neuron in V1 depends on the stimulus. In Figure 14, one can see the spiking and subthreshold responses of a V1 neuron when stimulated in its receptive field by four distinct stimuli: a drifting grating, a drifting grating animated with a model of eye movement reproducing drifts, tremors and saccades, a natural image animated with eye movement, and dense noise. From trial to trial, variability of the neuron depends on the stimulus. For a drifting grating, responses are rather constant in terms of firing rate, but the exact times of spike occurrences are not reliable. On the contrary, for a natural image with eye movements, huge and very reproducible spike times can be observed. Similar results have been replicated in Haider et al. (2010), which shows that combined stimulation of the classical receptive field of the neuron and its surround increases the reliability and the sparseness of the responses of V1 pyramidal neurons.

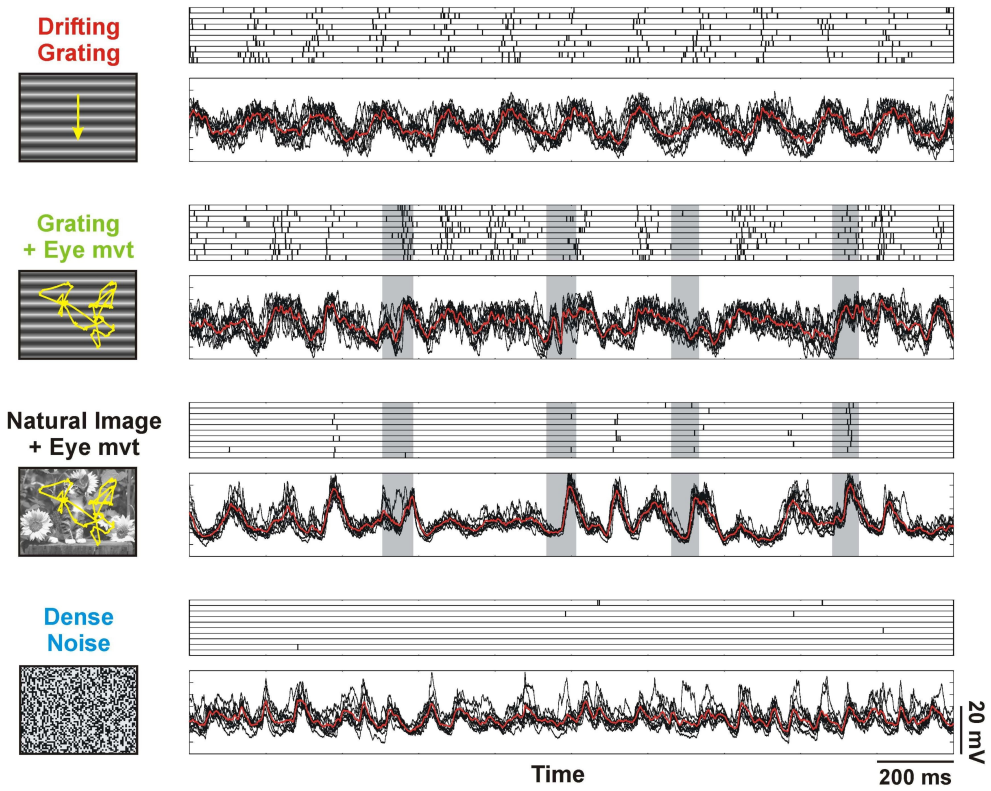

Figure 14: Intracellular recordings from a V1 neuron in an anaesthetized cat, presented with four different stimuli: from top to bottom; a sinusoidal grating, the same animated with simulated eye movements, a natural image with simulated eye movement, and a dense noise. Each time, the cell response is recorded 10 times, and spiking responses are shown for the 10 distinct trials. Membrane potentials for the ten trials are superimposed and the average, for each condition, is plotted in red. From Baudot et al, (2010).

These experimental results indicate that what is called "noise", i.e. the trial-to-trial variability of neuronal responses in the visual cortex, is modulated by the general context of full field input statistics sampled by the subthreshold synaptic integration field of the neuron. Different inputs to the same system can trigger different kinds of responses and dynamics and, in conjunction with the information transmission problems in recurrent networks, it raises the question of how do we need to stimulate such balanced random networks in order to maximize

information transmission. Are some inputs more appropriate than others?

8.3 Main results

The study presented in the following article is a theoretical attempt to better understand how stimulus-dependent reliability can emerge in simple and generic large-scale recurrent networks of spiking neurons. Does the reliability depend on the stimuli, and if so, what are the key properties of the stimuli triggering the reliability? Since several studies have shown how ongoing activity, i.e. "noise" can emerge from the recurrent cortical connectivity (van Vreeswijk and Sompolinsky, 1996, 1998), it appears logical to model the trial-to-trial variability as a product of the recurrent connectivity. To relate these models to our experimental results, we will raise two issues concerning the transmission of information within these networks:

– Is the reliability of the information transmission stimulus-dependent?
– What kind of stimulus maximizes the efficiency of information transmission?

I designed the experiments and the paradigm in collaboration with O. Marre, and we worked together equally on the project.

Behavioral/Systems/Cognitive

Reliable Recall of Spontaneous Activity Patterns in Cortical Networks

Olivier Marre*, Pierre Yger*, Andrew P. Davison, and Yves Frégnac
Unité de Neurosciences Intégratives et Computationnelles, UPR CNRS 2191, Centre National de la Recherche Scientifique, 91198 Gif-sur-Yvette, France

Irregular ongoing activity in cortical networks is often modeled as arising from recurrent connectivity. Yet it remains unclear to what extent its presence corrupts sensory signal transmission and network computational capabilities. In a recurrent cortical-like network, we have determined the activity patterns that are better transmitted and self-sustained by the network. We show that reproducible spiking and subthreshold dynamics can be triggered if the statistics of the imposed external drive are consistent with patterns previously seen in the ongoing activity. A subset of neurons in the network, constrained to replay temporal pattern segments extracted from the recorded ongoing activity of the same network, reliably drives the remaining, free-running neurons to call the rest of the pattern. Comparison with surrogate Poisson patterns indicates that the efficiency of the recall and completion process depends on the similarity between the statistical properties of the input with previous ongoing activity The reliability of evoked dynamics in recurrent networks is thus dependent on the stimulus used, and we propose that the similarity between spontaneous and evoked activity in sensory cortical areas could be a signature of efficient transmission and propagation across cortical networks.

Introduction

When injecting fluctuating current inputs into the soma of a neuron recorded *in vitro*, the spiking response is highly reliable (Mainen and Sejnowski, 1995). Such results, however, have been difficult to reproduce *in vivo* (Holt et al., 1996). The reason for this discrepancy is that much of the response variability observed *in vivo* seems to originate from the background activity (Arieli et al., 1996; DeWeese et al., 2005). Even in the absence of external drive, this ongoing neuronal activity is highly irregular (Timofeev et al., 2000), although the discharge statistics are still a matter of debate (Kenet et al., 2003; Fiser et al., 2004; Goldberg et al., 2004). Generic recurrent networks are a good model for understanding the possible interactions between ongoing and evoked activity in neocortical networks. Characterized by large, sparsely connected, excitatory and inhibitory populations, they can display a stable, self-generated regime called asynchronous irregular (AI) (van Vreeswijk and Sompolinsky, 1996; Brunel, 2000; Vogels et al., 2005; El Boustani and Destexhe, 2009a), which resembles the spontaneous activity observed *in vivo*. The central functional issue, as yet unsolved, is to characterize the sensitivity of these networks to external inputs (Destexhe and Contreras, 2006; Banerjee et al., 2008).

These recurrent AI networks are highly sensitive to small perturbations. As a consequence, propagation of either an increase in firing rate (Vogels et al., 2005) or a pulse of synchronized activity (Aviel et al., 2003; Mehring et al., 2003) is severely impaired by the ongoing activity, and in general this high variability and instability presents a severe challenge to information transmission or processing in recurrent networks. To achieve reliable signal propagation, these earlier studies introduced specific constraints in the network structure, either by selectively and substantially increasing synaptic weights (Vogels et al., 2005) or by adding connections (Mehring et al., 2003; Kumar et al., 2008) along a predetermined propagation path. Even in this latter case, the synfire chain stimulation can induce "synfire explosions," which can subsequently silence the network activity.

While previous studies adapted the network connectivity to improve the transmission of the chosen (a priori) input patterns, we explored the converse approach: leave the network structure unchanged and find the activity patterns that are better transmitted and sustained by the network. Several experimental studies have shown a similarity between spontaneous and evoked cortical activity (Tsodyks et al., 1999; Kenet et al., 2003; Fiser et al., 2004). From the theoretical point of view, the irregular and sustained patterns found in ongoing activity are by definition highly compatible with the recurrent architecture of the neocortical network. We have designed a new stimulation paradigm, in which we drive part of the network with temporal pattern segments extracted from the recorded ongoing activity of the exact same network. This gives inputs which mimic the spontaneous activity of the network model. We refer to this as the

Received Feb. 13, 2009; revised June 17, 2009; accepted July 15, 2009.

This work was supported by Centre National de la Recherche Scientifique, Agence Nationale de la Recherche (Natstats), and European Commission-funded (Bio-I3: Facets FP6-2004-IST-FETPI 15879) grants to Y.F. O.M. was supported by bursaries from the French Ministry of Defence and the Fondation pour la Recherche Médicale. P.Y. was supported by a Ministère de l'Education Nationale de la Recherche et de la Technologie bursary from the University of Paris XI. We thank Sami El Boustani and Drs. Alain Destexhe, Olivier Faugeras, and Cyril Monier for helpful discussions and critical comments on the manuscript.

*O.M. and P.Y. contributed equally to this work.

Correspondence should be addressed to Yves Frégnac, Unité de Neurosciences Intégratives et Computationnelles (UNIC), UPR CNRS 2191, Centre National de la Recherche Scientifique, Bât. 32-33, 1 Avenue de la Terrasse, 91198 Gif-sur-Yvette, France. E-mail: fregnac@unic.cnrs-gif.fr.

O. Marre's present address: Department of Molecular Biology, Princeton University, Washington Road, Princeton, NJ 08544.

DOI:10.1523/JNEUROSCI.0753-09.2009

 0270-6474/09/2914596-11$15.00/0

"frozen paradigm." In this article, we show using simulations that our paradigm produces efficient transmission, which is preserved over a broad range of parameters. To investigate the factors affecting transmission efficiency in the network, we performed further simulations with control, surrogate stimuli. Finally, we discuss the biological relevance of this paradigm.

Materials and Methods

Spiking network model

Neuron model. Networks are composed of 10,000 leaky integrate-and-fire neurons. Each neuron has membrane time constant $\tau_m = 20$ ms, and resting membrane potential $V_m = -60$ mV. When V_m reaches the spiking threshold $V_{thresh} = -50$ mV, a spike is generated and the membrane potential is held at the resting potential for a refractory period of duration $\tau_{ref} = 5$ ms. Synaptic connections are modeled as conductance changes, resulting in a model similar to the Conductance-Based, Integrate and Fire models case in the study by Vogels et al. (2005):

$$\tau_m \frac{dV(t)}{dt} = (V_{rest} - V(t)) + g_{exc}(t)(E_{exc} - V(t)) + g_{inh}(t)(E_{inh} - V(t)). \quad (1)$$

Reversal potentials are $E_{exc} = 0$ mV and $E_{inh} = -80$ mV, and synaptic activation is modeled as a conductance step $g \rightarrow g + \delta g$ followed by exponential decay with time constants $\tau_{exc} = 5$ ms and $\tau_{inh} = 10$ ms. The integration time step of our simulations is 0.1 ms (reducing the time step to 0.01 ms was found to produce no qualitative change in the network behavior), and synaptic delays are set to 0.1 ms.

Network connectivity. The network is composed of 8000 excitatory and 2000 inhibitory neurons, sparsely and randomly connected, with a connection probability of 2% (Vogels et al., 2005), independent of the identity of the target. While this connectivity ratio is at odds with anatomical data on neocortical connectivity [which shows a connectivity ratio of around 10%, and a dependence on the target cell type (Binzegger et al., 2004)], it allows a realistic pattern of activity to be generated and sustained, while higher connectivity ratios fail to do so (Vogels et al., 2005). We also note that for models with dynamic synapses (Markram and Tsodyks, 1996), due to the depression of synaptic strength of intracortical excitatory synapses which is dominantly reported *in vitro* (Thomson and West, 1993), there are fewer effective synapses than in the anatomical data, making the effective connectivity closer to that used in our simulations.

Synaptic parameters and network regime. Network global states are defined as either synchronous or asynchronous (population viewpoint) and as either regular or irregular (neuron viewpoint). For Figures 1–4, the network is set to an asynchronous irregular state (Brunel, 2000) at 13 Hz with a mean interspike interval (ISI) coefficient of variation (CV) of 1.57. Synaptic parameters are as follows: $\delta g_{exc} = 6$ nS and $\delta g_{inh} = 61$ nS. Weights are drawn from Gaussian distributions $N(g_{exc}, g_{exc}/3)$ and $N(g_{inh}, g_{inh}/3)$, with negative values discarded and then redrawn. To initiate the self-sustained activity, 5% of the cells receive an initial 50 ms burst of activity at 100 Hz.

Measures. To quantify the similarity between the "free-running" and "target" activity patterns, we calculated the recall index defined as the normalized cross-correlation of the target and actual spike trains (Aertsen et al., 1989) for a random sample of 500 cells, for zero phase delay and a time bin of 5 ms. We consider two spike trains of $N = 500$ neurons, binned with a time bin δt equal to the refractory period of the neuron (here, $\delta t = 5$ ms), obtained during two different trials. $\mathbf{S}_1$ and $\mathbf{S}_2$ are two matrices of size $N \times N_{bins}$, with $N_{bins} = T/\delta t$, T being the length of the spike trains. Note that $\mathbf{S}$ are binary matrices: for each $(i, j) \in N \times N_{bins}$, $S_{ij} \in \{0, 1\}$. The recall index ρ_{12} is given by the following formula:

$$\rho_{12} = \frac{\langle S_1 S_2 \rangle - \langle S_1 \rangle \langle S_2 \rangle}{\sqrt{\langle S_1 \rangle (1 - \langle S_1 \rangle) \langle S_2 \rangle (1 - \langle S_2 \rangle)}}, \quad (2)$$

where $\langle\rangle$ denotes the average number of filled bins per matrix. Since $\mathbf{S}$ is a binary matrix, we are sure that $\langle \mathbf{S} \rangle < 1$.

To estimate the reliability (reproducibility) of the neural responses for different kinds of stimuli, we computed the mean of the paired normalized cross-correlations between the trials ($n = 10$) (Schreiber et al., 2003), with the same time bin. Error bars were computed by using different stimulation patterns ($n = 10$).

Signal-to-noise ratio for the binned spiking activity. We chose a measure which can be directly related to the reliability and mean firing rates previously estimated. The signal S was given by the mean firing rate m multiplied by the time bin Δt: $S = m\Delta t$. The noise N was given by the standard deviation of the responses to the same stimulation. These values were estimated with the same time bin as the reliability to ensure a straightforward relationship between the two. If $X_i^t \in \{0; 1\}$ represents the binned response at time t, for the ith trial, the noise N is such that the following is true:

$$N^2 = \left\langle \frac{1}{n} \sum_i \left(X_i - \frac{1}{n} \sum_j X_j \right)^2 \right\rangle_t, \quad (3)$$

where n is the number of trials. Remarking that, since the bin size is small, $X_i = X_i^2$, and that $\langle X_i \rangle_t = S$, we obtain the following:

$$N^2 = \left(1 - \frac{1}{n}\right)(S - C), \quad (4)$$

where $C = \langle X_i X_j \rangle_t$. According to Palm et al. (1988), $C = rS(1 - S) + S^2$, where r is the reliability estimated above as the normalized cross-correlation between trials. The signal-to-noise ratio (SNR) is then the mean over the standard deviation, such that the following is true:

$$\mathrm{SNR}^2 = \frac{S}{1 - r(1 - S) - S}. \quad (5)$$

Note that we drop the $(1 - 1/n)$ factor, which will affect all the estimations in the same manner, and will reach one for an infinite number of trials. This is not problematic for our comparative study.

Distance between spike trains. to test the noise resistance, we used the Victor–Purpura distance (Victor and Purpura, 1996), which has the advantage of avoiding any binning artifact. Briefly, costs are assigned to three elementary operations over spike trains: changing the timing of a spike from t_a to t_b (cost $= q\|t_a - t_b\|$), adding a spike (cost $= 1$), and deleting a spike (cost $= 1$). The distance between two spike trains S_A and S_B is then defined as the least costly way of combining these three elementary operations to change spike train S_A into spike train S_B. This distance is dependent only on q, and this parameter is the cost for shifting the spikes when matching two spike trains. If $q = 0$, shifting spikes has no effect on the measure, so only the number of spikes will influence the distance between the two spike trains. Increasing q gives more importance to the precise spike times relative to their number. In our study, we estimated the distance between the two spike trains emitted by a same cell in response two different stimuli (the real pattern vs the jittered version of it), and averaged this estimation over $n = 500$ cells.

Simulator. All simulations were performed using the NEST simulator (Diesmann and Gewaltig, 2001) version 1.9 (http://www.nest-initiative.uni-freiburg.de), using the PyNN interface (Davison et al., 2008) (http://neuralensemble.org/PyNN). The code for the model is freely available from ModelDB (http://senselab.med.yale.edu/ModelDB/) and on the UNIC website (http://www.unic.cnrs-gif.fr).

Neural field model

The model consists of N local "neural fields" defined by inputs h_i and outputs $S_i = \tanh(gh_i)$. Their dynamics are given by Equation 1, above. All results given in this paper used $N = 2000$ neural fields. The frozen paradigm was applied by imposing the values of the h_i of a subset of neural fields.

Results

Convergence of the network activity to a target activity pattern

The frozen paradigm is implemented as follows: We divide our recurrent network into two populations. The selection of which neuron belongs to which population is made at random and the

connectivity of two neurons is independent of which populations they belong to. We first record a spontaneous pattern across the whole network (both populations) and then, while the spontaneous activity is ongoing, we force the neurons of one population to replay the sequence of spontaneous activity previously recorded (see Fig. 1*a*). We then measure the extent to which these clamped, or frozen, cells influences the free-running neurons to replay the spontaneous pattern previously recorded (see Fig. 1*a*). Replay of the recorded pattern by the free-running neurons is then equivalent to the recall of that pattern in the network.

We applied our paradigm to a sparsely connected, recurrent network of 8000 excitatory and 2000 inhibitory integrate and fire neurons (see Materials and Methods). The set of parameter values was that used by Vogels et al. (2005), except that the weights are drawn from Gaussian distributions instead of being all equal. The network settles in an asynchronous irregular regime in which it generates self-sustained activity with a mean rate of 13 Hz and a mean ISI CV of 1.57. It has been shown that, although the activity of this network appears very irregular on small time scales, where it cannot be distinguished from stochastic behavior, it can exhibit more coherent behavior at large time scales, where the dimensionality of the attractor can be reduced (El Boustani and Destexhe, 2009b). Figure 2 shows an example of the dependency on initial conditions in this network model: two initially identical trajectories quickly diverge after a perturbation as small as a single spike elicited in the same neuron. This divergence is characteristic of a chaotic network, and has already been observed and studied in several recurrent network models [see also Sompolinsky et al. (1988) and van Vreeswijk and Sompolinsky (1998, 2005)].

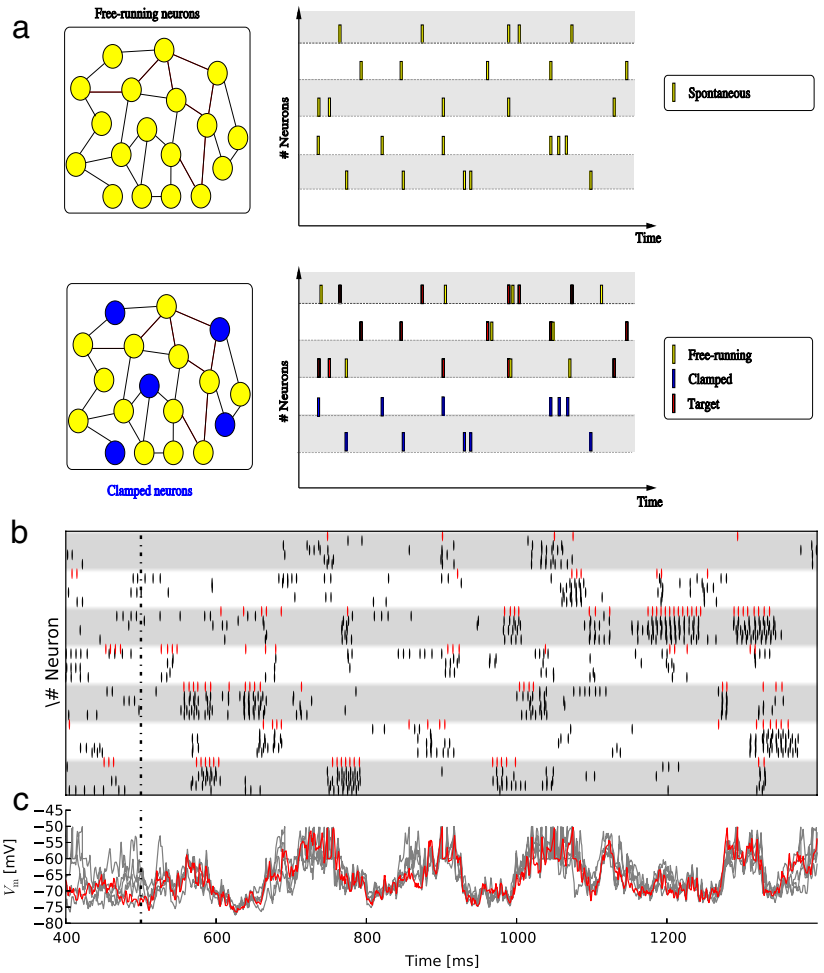

Figure 1. ***a***, Conceptual schema of the "frozen paradigm." A spontaneous pattern is recorded (top panel), and then a subset of neurons (labeled "frozen"; blue neurons and blue spikes) is forced to replay part of the pattern. We then examined whether the remaining, free-running neurons (yellow spikes) reliably reproduce the other part of the spontaneous pattern, which we label the "target" pattern (red spikes) (see Results for details). ***b***, Raster plot of responses of free-running neurons to the frozen stimulation. Each white or gray band represents the activity of one free-running neuron. Short vertical lines represent spikes. The red spikes are from the target pattern, and each row of black spikes represents a different trial with the same stimulation pattern but a different initial state of the network. ***c***, Superimposed V_m traces of responses to the same frozen stimulation, for one free-running neuron. The red trace indicates the target activity.

When 50% of the neurons are forced to replay a spontaneous pattern recorded previously (Fig. 1*a*, blue), the spiking activity of the free-running neurons (yellow) converges reliably to the target activity (red): repetitions of the same stimulation elicit temporally precise and reproducible spikes characterized by their temporal alignment across repetitions (Fig. 1*b*). The subthreshold membrane potentials V_m (Fig. 1*c*) of the free-running neurons closely follow the target activity waveforms (red) as soon as the input-recipient population is frozen. Note also the immediate reduction in the stimulus-locked variance following the "freeze" onset.

We quantified these observations with two measures. The recall index is a measure of how closely the free-running neurons of the network reproduce the target pattern (the previously recorded segment of spontaneous activity). It is defined as the cross-correlation between the target activity pattern and the response of the free-running neurons. The reliability is a measure of the variability of responses during frozen stimulation: the lower the variability the higher the reliability. It is defined as the mean cross-correlation between pairs of responses to the same frozen stimulus (see Materials and Methods for full details). Following the onset of stimulation, the recall index increases, within a time of ~50 ms that is independent of the proportion of frozen neurons, to a steady-state value (Fig. 3), reflecting a rapid convergence of the free-running activity to the target pattern. The recall index decays equally rapidly following the end of the stimulation. The firing rate averaged over the neuronal population (Fig. 3*a*, bottom panel) remains constant over time. The frozen stimulation thus preserves the statistics of the ongoing activity while eliciting a reproducible and faithful replay, extended over the full network, of the target pattern.

The steady-state value of the recall index measure is dependent on the proportion of frozen neurons (see Fig. 3*b*). As expected, the recall index increases when the proportion of frozen neurons increases, whether these are excitatory or inhibitory, except in a region where the low proportion of inhibitory free-running neurons evokes a transition to a "synchronous regular" regime or a quasisilent state, which are insensitive to the stimulus. The proportion of frozen inhibitory neurons affects the recall index more than does that of frozen excitatory neurons, which is

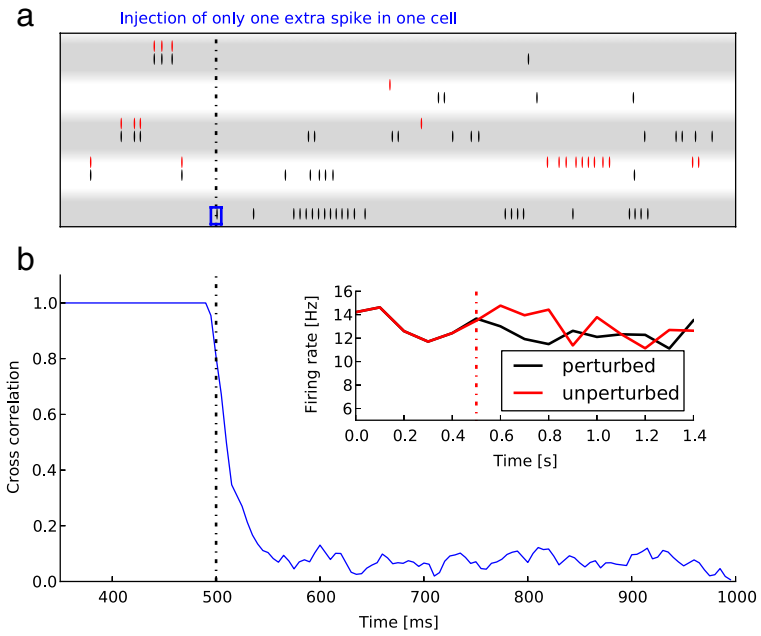

Figure 2. ***a***, Raster plot of five neurons during two runs of activity (red and black) in the same network, starting with the same initial condition. The black run is perturbed by the addition of one extra spike in one neuron at 500 ms (dotted line, blue inset in the bottom row). ***b***, Time course of the normalized cross-correlation between the two runs of activity. Correlation is computed with a time bin equal to the refractory period of the neurons, and measured on a sliding window. Note that this induces a smoothing of the estimated correlation. The vertical dotted lines indicate the time at which a perturbation (one extra spike) was artificially added to one of the neurons in one of the runs, inducing a fast drop of the correlation. The inset shows the mean firing rate for the two conditions as a function of time.

probably related to the higher weights of inhibitory synapses in this model ($\delta g_{exc} = 6$ nS, $\delta g_{inh} = 61$ nS; see Materials and Methods). The recall index is insensitive to the particular pattern chosen for stimulation and to the initial conditions (ANOVA, $p > 0.5$ for both pattern and initial condition dependencies), with a standard deviation approximately constant over the whole bidimensional plot shown in Figure 3*b* (0.006 on average). This absence of sensitivity to the timing of the stimulation relative to the ongoing activity is probably a consequence of the irregularity of the network activity: there is no clear oscillatory behavior which could induce a phase dependency, and our stimulation is not related to a particular frequency to which the response would lock.

To check whether the frozen stimulation always makes the network converge to the target pattern, we compared the values of the recall index and of the reliability for the same stimulation repeated several times. Convergence to a different pattern would lead to a reliability significantly higher than recall index, but we did not find any significant difference ($p > 0.4$, *t* test, $n = 10$), proving that the activity converges in all cases to the target pattern. Our results thus demonstrate that the frozen stimulation induces robust convergence for a large enough frozen population.

The importance of the recurrent architecture in this behavior was confirmed by a control experiment in which the connections between free-running neurons were cut and a stochastic current of equivalent mean amplitude injected into the neurons (Fig. 3*c*). The injected current is independent over all neurons. The mean firing rate received by each neuron is thus preserved, while the interactions between stimulation and recurrent architecture are suppressed. For a frozen proportion of 50%, the reliability drops from 0.47 ± 0.007 (SD) to 0.09 ± 0.005, indicating that the free-running activity in the intact network during frozen stimulation is not dominated by the frozen drive, but results from a dynamic cooperation between the frozen units and the recurrent connections.

Influence of input temporal structure on reliability

It could be argued that the reliability of the responses could have been obtained with any other imposed stimulation, whatever its statistics. We thus compared the frozen stimulation with surrogate stimulations having the same number of spikes, but a shuffled temporal structure. Several different surrogates were compared: a first series consisting of temporally "jittered" patterns with standard deviation of the jitter ranging from 5 to 25 ms; and a second series consisting of Poisson spike trains. In "local Poisson" (LP) stimulation, the mean firing rate of each neuron equals that of the same neuron during the reference, frozen pattern (equivalent to an infinite jitter). In "synchronous Poisson" (SP) stimulation, all neurons have the same firing rate, equal to the mean rate of the spontaneous activity, but the spike trains are correlated (Kuhn et al., 2003), to match the synchrony level of spontaneous activity. Finally, "global Poisson" (GP) stimulation matches only the mean firing rate of the spontaneous activity, but not the synchrony. Note that all these controls have approximately the same number of spikes as the reference spontaneous pattern. These different surrogate patterns represent a progressively increasing deviation from spontaneous activity statistics, while keeping first-order statistics unchanged.

For this comparison we chose a level of frozen neurons (50%) at which each free neuron receives equal numbers of connections from frozen and free neurons. As shown in Figure 4*a*, the more the stimulation statistics deviate from those of spontaneous activity the lower is the reliability. This difference is significant for jitter of 15 ms and greater ($p < 0.001$) and for LP and GP ($p < 0.0001$), being largest for the global Poisson process. For the global Poisson, the reliability drops by 17%. However, the level of reliability is the same for synchronous Poisson as for the frozen stimulation. The reliability modulation is accompanied by a modulation of the mean evoked firing rate in the free running network, as can be seen in Figure 4*b*: the mean firing rate drops significantly ($p < 0.0001$) for local (18%), synchronous (43%), and global (34%) Poisson surrogate patterns.

To synthesize these results, we measured the SNR of the spiking responses, defined as the mean firing rate divided by the standard deviation of the firing rate responses over trials (see Materials and Methods). This takes into account both the reliability and the mean firing rate of the output. This SNR is significantly higher for the spontaneous stimulation than for the surrogate Poisson stimulations (see Fig. 4*c*) [decrease of 27% (SP), 14% (LP), and 26% (GP)]. In the case of the local and global Poisson stimulations, the difference in SNR is produced by both the reliability and the mean firing rate drops. But in the case of the synchronous Poisson stimulation, the reliability is comparable to that of the spontaneous pattern, and the lower, but still significant, difference originates in the lower firing rate. Together, these

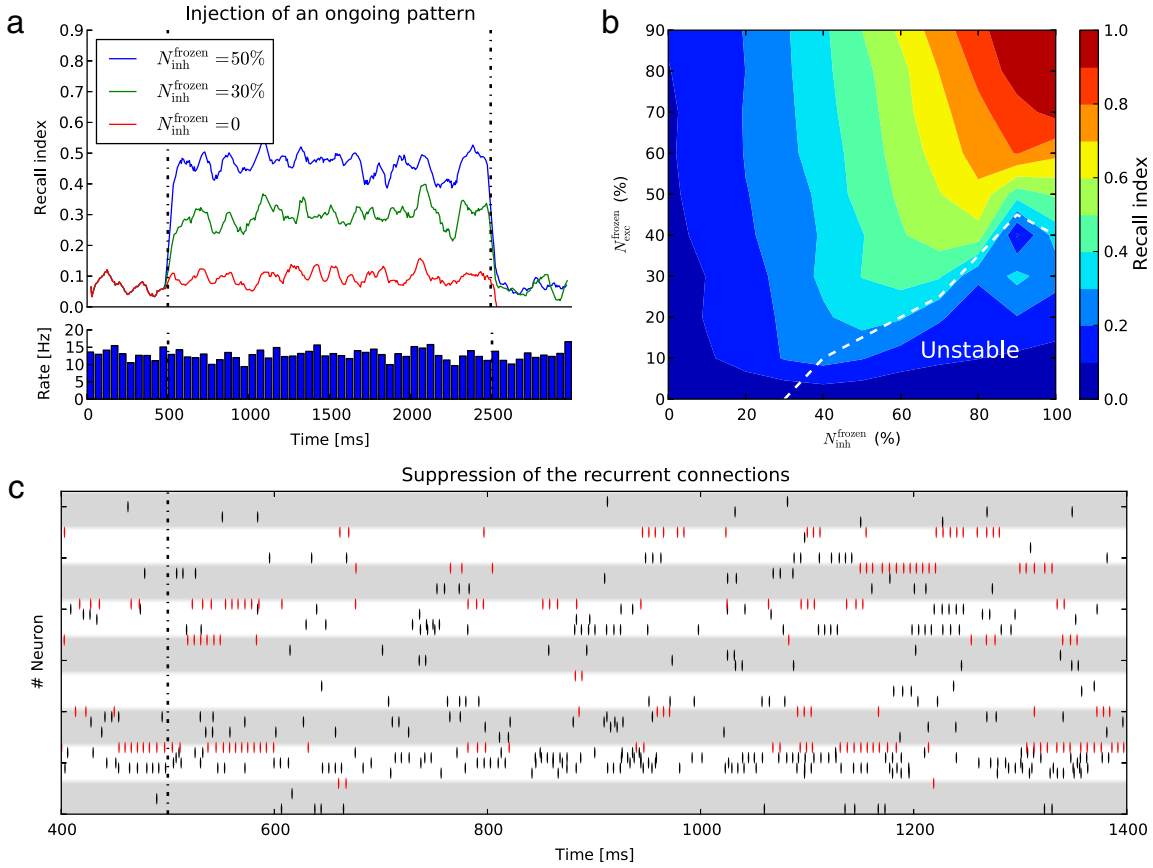

Figure 3. ***a***, Kinetics of the recall index (the normalized cross-correlation between the target activity and the response) as a function of time, before, during, and after the frozen stimulation. ***b***, Steady-state recall index color coded, as a function of the proportion of excitatory and inhibitory neurons that are frozen. The white dotted line delineates a zone of instability of the network in which the AI regime is not sustained. ***c***, Raster plot of responses, as in Figure 1*b*, when recurrent connections between free-running neurons are cut and replaced by independent Poisson spike trains producing an equivalent mean level of input to the cells.

measures imply a larger degradation in signal transmission, the more the input statistics deviate from the spontaneous statistics. Thus, for the same number of spikes, the forced replay of spontaneous ongoing patterns seems to induce a better signal transmission and recall.

We have explored how the input structure influences the output reliability and mean firing rate. To compare this influence to that of the mean input firing rate, we explored the effects on mean activity (Fig. 5*a*) and reliability (Fig. 5*b*) of independently varying the mean excitatory and inhibitory firing rates of the global Poisson stimulation. The combination of high excitatory and low inhibitory stimulation rates produces a high mean firing rate in the free-running units, but a low trial-to-trial reliability. In contrast, low excitatory but high inhibitory stimulation rates induce a highly reliable response, but with a very low mean firing rate. We empirically fitted this inverse relationship between reliability and mean activity of the free running units by the following power law:

$$\text{Rate}^{\alpha} \cdot \text{Reliability}^{\beta} = K \cdot (\text{Spike count}) \qquad (6)$$

($r^2 = 0.95$), where the spike count is the total number of spikes in the imposed stimulation pattern, summed over all the frozen neurons, and K is a constant ($\alpha = 0.46$, $\beta = 1.48$, $K = 11.59$).

Figure 5*c* represents the "iso-spike count" curves obtained when plotting reliability against mean firing rate. When compared with the Poisson surrogate stimulations, the particular structure of the imposed spontaneous pattern seems to be better adapted to the network connectivity: to reach similar levels of both reliability and response strength as observed for spontaneous statistics (Fig. 5*c*, black cross), the use of global Poisson stimulation would require a large increase in the spike frequency imposed on the frozen neurons (175% of the spontaneous rate, given in Fig. 5*c* by the iso-spike count curve intercepting the reference cross).

Together, these results indicate that, although the first determinant of the response reliability is the mean input firing rate, the structure of ongoing activity enables a more efficient signal transmission than uncorrelated stimulation. The input firing rate and synchrony levels explain a major part of this increase, but are not sufficient to reach similar levels of SNR.

Noise resistance

Our simulations allow us to replay precisely each time the same pattern in the frozen population. In more realistic situations, however, the frozen units, even if deterministic, could be corrupted by independent noise sources, which could interfere with signal transmission to the free-running units. To test the noise resistance of the pattern recall, we generated a set of degraded patterns from an original spontaneous pattern by randomly jittering each spike by a time drawn from a Gaussian distribution. The standard deviation of the distribution was varied from 1 to 50 ms. Each resulting pattern was then used to clamp the frozen units, and we compared the distance [measured as by Victor and Purpura (1996)] (see also Materials and Methods) between the spontaneous input and each jittered input against the distance between their corresponding outputs. We found that the modu-

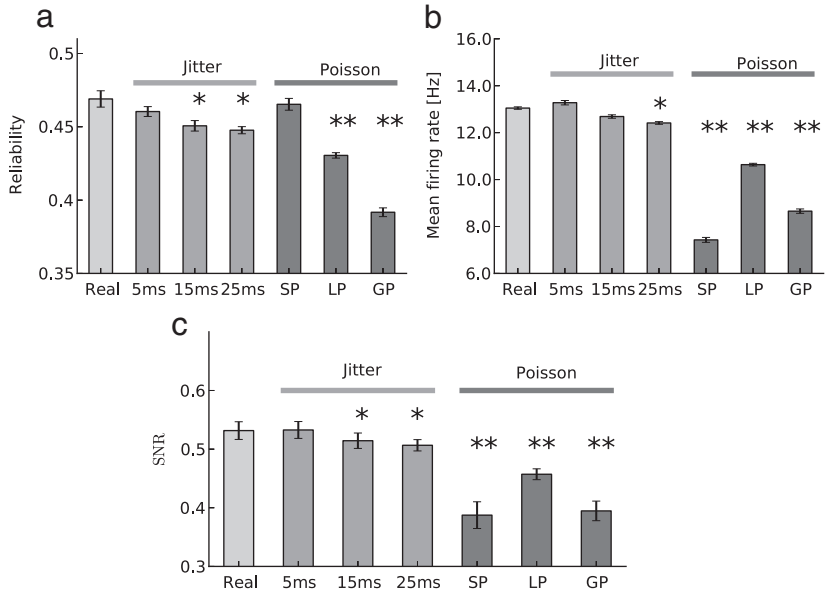

Figure 4. ***a***, ***b***, Comparison of the reliability (***a***) and mean firing rate (***b***) values obtained in response to different stimulation patterns, for 50% of excitatory and inhibitory frozen neurons. Stimulation patterns are, from left to right: a "real" pattern (Real) taken from the spontaneous activity, jittered versions of this pattern (5, 15, and 25 ms of jitter; see Materials and Methods), a Poisson stimulation (SP) reproducing the rate and the synchrony level observed in the real pattern, a Poisson stimulation in which each neuron has the same mean firing rate as in the spontaneous pattern (LP), and a Poisson stimulation in which all the neurons have the same firing rate, which is the global mean firing rate of the spontaneous activity (GP). ***c***, Same comparison for the SNR (see Materials and Methods) across all the surrogates.

lus of the output distance increases linearly with input distance (Fig. 6*a*), showing that the response to our stimulation is robust to noise. The slope of the input–output relationship is 0.57. This sensitivity to noise well below 1 shows that our stimulation paradigm is compatible with a transmission of more "biological" (i.e., noisy) inputs.

Noise could also be added in the model as external inputs that will perturb the recurrent activity of the free-running neurons. To check how robust is the preference for the ongoing statistics, while the network was asked to replay a particular pattern, all the free cells received, in addition to the pattern played by the frozen population, noisy inputs composed of uncorrelated excitatory and inhibitory Poisson spike trains at a certain frequency. This frequency is expressed relatively to ν_{thresh}, the input rate needed to reach the threshold in our particular neuron model with conductance-based synapses, in the absence of recurrent inputs:

$$\nu_{\text{thresh}} = \frac{g_{\text{leak}} V_{\text{rest}}}{\tau_{\text{exc}}\, g_{\text{exc}} V_{\text{thresh}}}. \quad (7)$$

As can be seen in Figure 6, *b* and *c*, the reliability and the recall index both decrease to a steady-state value when noise is increased. However, the important point is that the relative differences between stimulation types, observed in the purely deterministic case (see Fig. 4), are unchanged in the presence of noise (Fig. 6*b*). This result cannot be explained by a change in the firing rates, which remain approximately constant despite the external noise applied (Fig. 6*c*): since this external noise is balanced (excitatory and inhibitory), it has only a small impact on the resulting firing rates. Figure 6*e* shows the differences in the SNR according to the noise levels and for several surrogates. A linear decrease can be observed, but even for high noise levels (ν up to $3\nu_{\text{thresh}}$), the SNR remains significantly ($p < 0.001$ for LP and GP) higher for the frozen stimulation with real patterns. These results confirm the robustness to noise of the paradigm.

Effect of network parameters

To further test the generality of our observations, we varied the excitatory synaptic weights δg_{exc} between 1 and 10 nS and the inhibitory weights δg_{inh} between 1 and 91 nS, and used three different connection probabilities ϵ: 0.5%, 1%, and 2%. Over all the regions in which the network can generate an asynchronous irregular regime, the frozen paradigm with 50% frozen neurons produces qualitatively similar recall performance (Fig. 7*a–c*) (see also Vogels and Abbott, 2005). Surprisingly, the recall index is not directly related to the strength of the synaptic weights: an increase in synaptic weight does not necessarily increase the recall index. Over all conditions of connectivity and synaptic weights, the recall index increases with both the mean firing rate (Fig. 7*d*) and the ISI CV of the corresponding spontaneous activity. The structure of spontaneous activity, as characterized by the mean rate and ISI CV values, appears to be a better predictor of the convergence performance than is the connectivity structure of the network. This illustrates that our paradigm avoids focusing on the nonlinear relationship between the network structure and the way activity is transmitted inside this network by directly using those activity statistics that are proven to be sustained by the network. We also examined the level of reliability obtained with the GP stimulation over the same range of parameters. Figure 8 shows the normalized difference between the levels of reliability obtained for the spontaneous and GP stimulations. This difference could reach up to 25%. We did not find any clear correlation between the amount of synchrony in the spontaneous activity and this normalized difference (data not shown).

Frozen paradigm in a chaotic neural field model

Although it generates an irregular activity pattern and exhibits rapidly diverging responses to small perturbations, the recurrent model has not been proven mathematically to be chaotic. To have a better understanding of the effect of stimulation statistics in a well defined chaotic system, we apply the frozen paradigm to a neural field model (Sompolinsky et al., 1988) (see Material and Methods), defined by the following:

$$\frac{dh_i}{dt} = -h_i + \sum_j J_{ij} \tanh(g h_j). \quad (8)$$

The network is composed of N units, each characterized by its activity h_i. The synaptic weights J_{ij} are drawn from a Gaussian distribution with mean 0 and variance J^2/N. It has been shown for large N that when $gJ < 1$, the only stable state is the silent state. When $gJ > 1$, the network enters a chaotic regime (Sompolinsky et al., 1988).

Using the same strategy as for the spiking network, we froze the chaotic system by clamping part of the network to a previ-

ously recorded trajectory. During this frozen stimulation, we measured the convergence by the normalized cross-correlation between the free-running activity (values of h_i) and the target activity, equivalent to the previously defined recall index. As for the spiking network simulations, the cross-correlation (CC) rapidly increases to a plateau (Fig. 9*a*), whose value depends on the proportion of frozen units and on the network parameter gJ (Fig. 9*b*). Full convergence (CC $>$ 0.99) occurs if the proportion of frozen units is greater than a threshold value (Fig. 9*b*, dotted black line), which increases with gJ.

To test the resistance to noise of this convergence, we also injected random noise in addition to the frozen pattern. The convergence performance decreased linearly when the amplitude of this noise increased, the slope depending on the percentage of frozen units (see Fig. 9*c*). The noise resistance is thus similar to what we found for the spiking network.

To test the origin of this reliability, we again used a surrogate stimulation. We stimulated the subnetwork with random noise, of mean 0 and standard deviation σ equal to the spontaneous activity standard deviation σ_{SA}. The level of reliability was similar to that obtained with the real pattern. We then varied this standard deviation σ relative to σ_{SA} (see Fig. 9*d*). When σ decreased below σ_{SA}, the reliability decreased almost linearly. This confirms that the standard deviation of the activity is the main factor in explaining the reliability. Interestingly, when we increased σ above σ_{SA}, the reliability barely increased, whatever the absolute value of σ_{SA}. This result holds when the network parameters g and J and the ratio of frozen units do not give rise to full convergence (i.e., below the black line in Fig. 9). For a higher frozen fraction, stimulation with lower standard deviation can also lead to a saturated full convergence (data not shown). Thus, over a broad region of the network parameter space, σ_{SA} appears as a reference point in the relationship between σ and the output reliability.

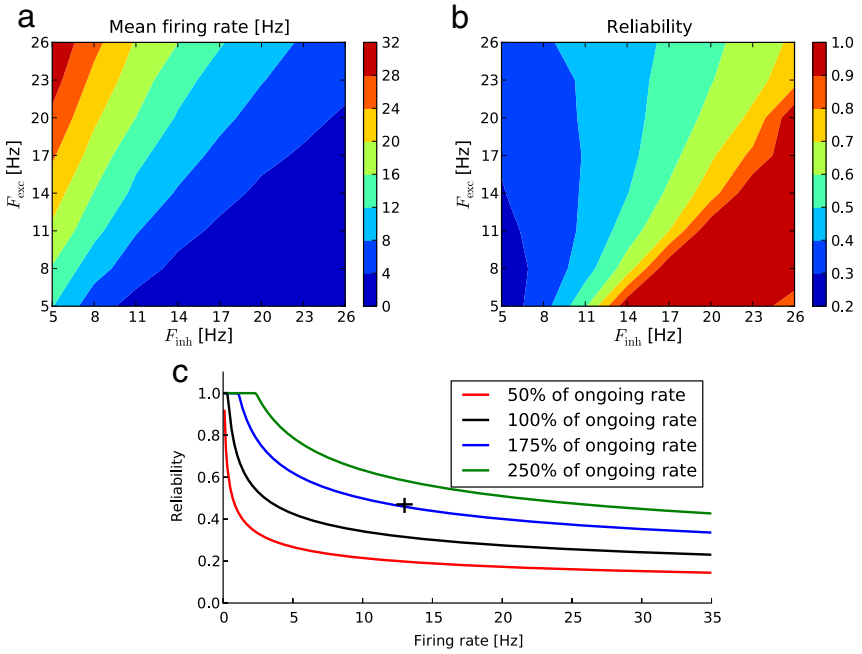

Figure 5. ***a***, ***b***, Mean firing rate (***a***) and reliability (***b***), color coded, in response to global Poisson stimulation, for different excitatory and inhibitory firing rates. ***c***, Reliability versus mean firing rate for different total spike counts. Each "iso-spike count" curve corresponds to the stimulations of ***a*** and ***b*** with the same total number of spikes as in the Poisson stimulation, but differently partitioned between excitatory and inhibitory populations. The black cross indicates the reliability and mean firing rate for the spontaneous pattern.

Discussion

In this paper, we analyzed the response of recurrent network models to a stimulation which mimics episodes of spontaneous activity. Our main findings are as follows: (1) Stimulating with spontaneous patterns of activity induces a predictable and noise-resistant recall of the full ongoing pattern, despite the highly irregular background activity and the context sensitivity of the network. (2) Despite the fact that the network is deterministic, the response reliability is modulated by the stimulus type, and is higher (up to 25%) for the spontaneous stimulation than for uncorrelated Poisson stimulation. (3) This reliability increase is mainly explained by the higher synchrony of the input. However, mimicking the input firing rate and synchrony is not sufficient to fully reproduce the signal-to-noise ratio obtained with the spontaneous stimulation.

On the basis of these findings, we make two main experimental predictions: first, the reliability of the cortical responses should be modulated by the stimulus statistics *in vivo*; second, evoked sensory responses with a structure similar to the spontaneous activity could be the signature of an efficient transmission of information.

Reliable response/completion despite irregular background activity

Our model demonstrates the feasibility of a reliable encoding of stimuli within irregular background activity. Several studies have shown that a chaotic dynamical system can be reliably driven by appropriate stimulation, a phenomenon termed "chaos control" (Garfinkel et al., 1992). Our frozen paradigm could be seen as a high-dimensional application of "chaos synchronization" (Pecora and Carroll, 1990). While these previous studies did not explore the effect of the input statistics, we are able to demonstrate that choosing a driving stimulation that respects the statistics of spontaneous activity has the advantage of allowing a prediction of the responses of the free-running neurons. Our investigation extends these earlier studies by applying these concepts to large-scale neuronal networks. It is important to note that we do not focus on any global fixed point reached by the system under the stimulation, but rather on the evoked "succession of transients," which are reliable and noise resistant. This concept has already been the subject of several studies (Rabinovich et al., 2008), and our results are a possible example of this concept in large-scale network models.

The neural field study illustrates the dual role of the stimulation. In this case, it has to be noted that the free subnetwork alone, without any stimulation, is chaotic (below the white line of Fig. 9*b*). Nevertheless, the stimulation due to the frozen population makes a reliable and noise-resistant propagation possible. The stimulation has two roles. First, it transmits the pattern, and second, it changes the context (background activity), so that the transmission is possible, and resistant to noise. These two complementary roles of the stimulation are due to the nature of recurrent networks, where the stimulation and the background activity interact with each other.

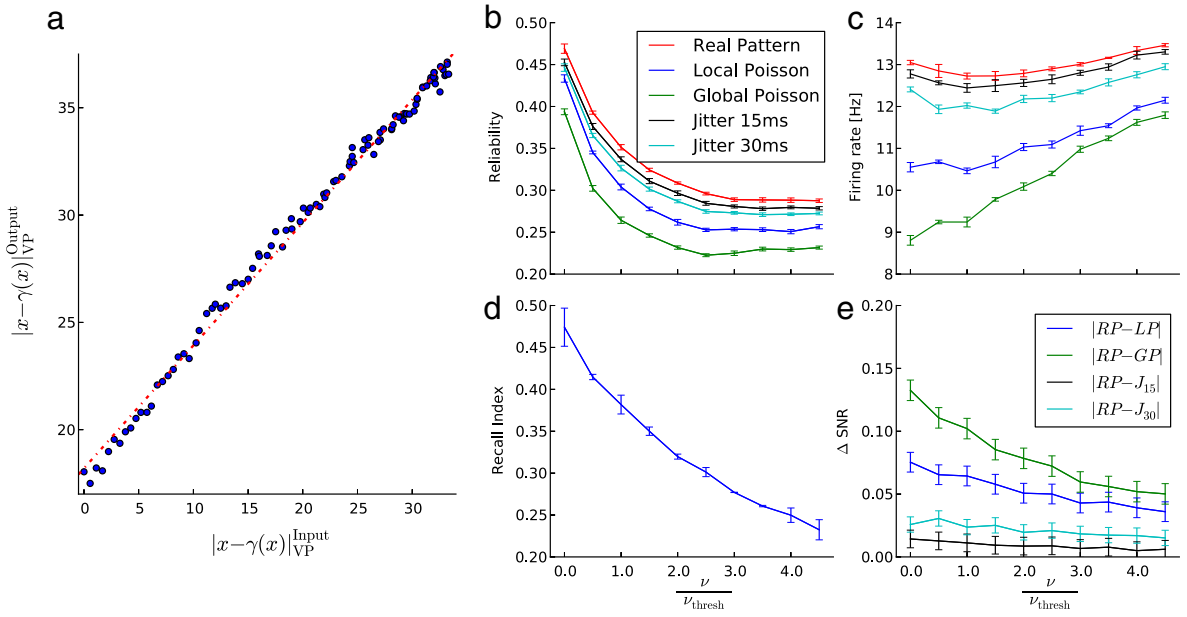

Figure 6. Noise resistance of the paradigm. ***a***, The coordinates of each point are the Victor–Purpura (VP) distance $\|x - \gamma(x)\|_{VP}$ between an input pattern x and the jittered version of it $\gamma(x)$ (horizontal axis), and the VP distance between the corresponding output patterns (vertical axis). ***b***, ***c***, Evolution of the reliability (***b***) and the firing rate (***c***) in the free-running neurons for several level of external noise (relatively to ν_{thresh}). Errors bars are obtained on 10 run per surrogate. ***d***, ***e***, Linear decrease of the reproducibility for one particular pattern (***d***), and difference between the SNRs obtained with a real pattern and those obtained for surrogates (***e***), as a function of the external noise.

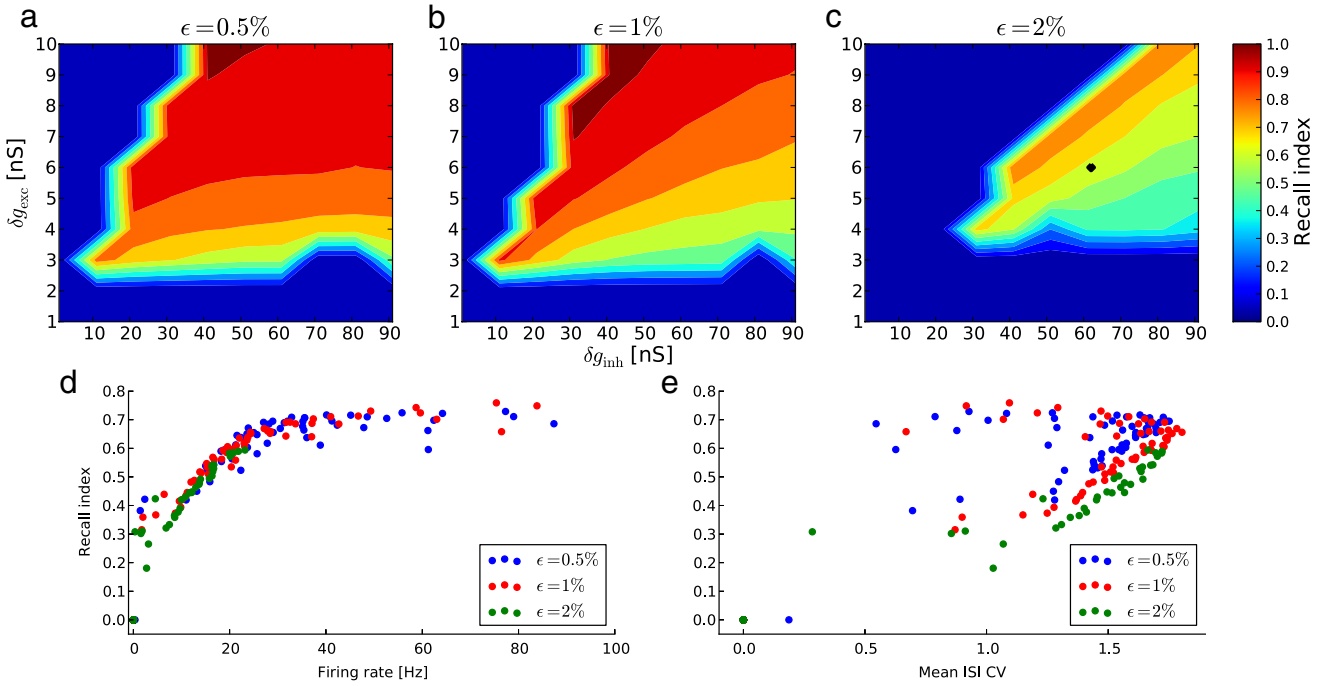

Figure 7. ***a–c***, Recall index for 50% of neurons frozen as a function of the excitatory and inhibitory synaptic weights, and for different connection probabilities ϵ: 0.5% (***a***), 1% (***b***), and 2% (***c***). The dark blue zones indicates regions in which the network either cannot sustain spontaneous activity or enters a synchronous regular state. ***d***, ***e***, Recall index values obtained in ***a–c*** plotted against the mean firing rate (***d***) and ISI CV (***e***) of the spontaneous activity of the same network.

Origin and stimulus dependence of the reliability

Our reliability study uncovers two findings. First, even if the network model used in this study is entirely deterministic, we observed a large trial to trial variability in response to the repetitions of the same stimulus. All the observed variability must originate in the surrounding network activity. As a consequence, experimentally, a large part of the trial to trial variability may come from the context sensitivity of a cortical network, and could thus be predicted by the preceding ongoing activity. This is in line with previous experimental reports where the preceding ongoing activity was a major source of variability (Arieli et al., 1996; Azouz and Gray, 1999; Deweese and Zador, 2004).

Second, our study also demonstrates that the response reliability can be modulated by the stimulus statistics in a nontrivial manner. Although the first determinant of the reliability level is the mean firing of the input, the higher-order structure of the stimulation has also a non-negligible effect (up to 25%), and this over the whole range of model parameters explored.

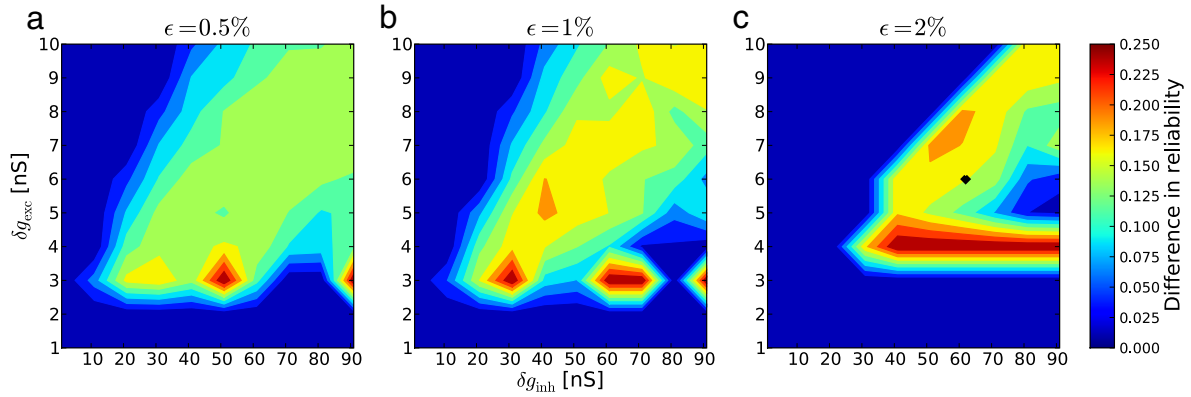

Figure 8. Normalized difference between the reliability obtained with global Poisson stimulation and that obtained with a real pattern, as a function of the synaptic weights, for different connectivities, ϵ.

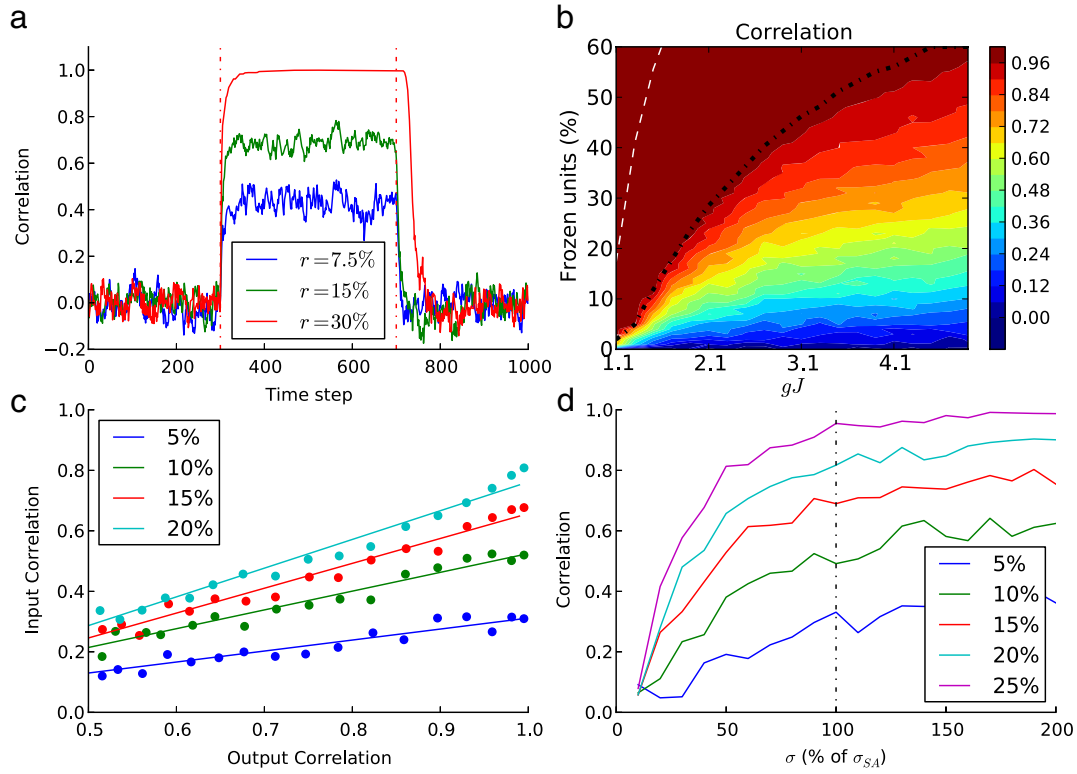

Figure 9. ***a***, Kinetics of the normalized cross-correlation between the target activity and the response of neural fields as a function of time, before, during, and after the frozen stimulation, for three different ratios of frozen neural fields. In this example, $gJ = 2.1$. ***b***, Steady-state values of the cross-correlation between target and free-running activities measured as a function of the ratio of frozen neural fields (vertical axis) and the network parameter gJ (horizontal axis). The black dotted line is the full convergence limit (CC > 0.99) and the white dotted line the limit under which the free subnetwork, without any stimulation, is still chaotic. ***c***, Normalized cross-correlation for several percentages of frozen units and for several amounts of noise added (in percentage of σ_{SA}) to an ongoing pattern. ***d***, Normalized cross-correlation when a surrogate Gaussian noise with a standard deviation expressed in percentage of σ_{SA} is used as a stimulation instead of a real pattern. In both cases, the network parameter was $gJ = 2.1$.

Previous theoretical results could be reinterpreted in light of this result. Vogels et al. (2005) noted that their modified network, which enhanced firing rate but not synfire chain propagation, occasionally transmitted "ghost signals." We interpret this reminiscence of the reliably evoked activity in spontaneous dynamics as another example of a network model that better transmits signals whose statistics match those of the spontaneous activity.

Beyond rate and synchrony

Our surrogate stimulation mimicking both the level of input firing rate and synchrony reached the same output reliability as the spontaneous stimulation. Nevertheless, since the output firing rate is lower, the resulting SNR is still lower than for the spontaneous stimulation. Previous studies have focused on the impact of rate (Vogels and Abbott, 2005) and synchrony (Mehring et al., 2003; Kumar et al., 2008) on the propagation of activity in recurrent networks. Our study showed that these two factors, though important, are not entirely sufficient to explain the SNR of the response to the spontaneous stimulation. Further studies will aim at identifying additional contributions to this SNR modulation. Among others, using a stimulation with heterogeneous firing rates could be one of these factors (see the difference between local and global Poisson stimulation).

Relation to experimental data

In light of our results, we can propose an explanation for the discrepancy mentioned in the Introduction between the reliable spiking activity observed *in vitro* (Mainen and Sejnowski, 1995) and the variability often observed *in vivo*. Our study indeed predicts that the reliability should be modulated by the stimulus statistics, and this has been observed experimentally: intracellular studies demonstrate that the reliability of subthreshold activity is increased during natural scene viewing (Baudot et al., 2004; Frégnac et al., 2005) or when the stimulus evokes strong shunting inhibition (Monier et al., 2008).

Some reports point out that the reliability of responses could depend on the interaction between internally generated and sensory-evoked activities (see above). In particular, the precision of the spiking responses may depend on the global reverberation rhythm generated by the recurrent network connectivity, which could preserve or destroy the sensory information, depending on the relative phase and the amplitude of the fast oscillations (for review, see Tiesinga et al., 2008). We propose to extend this prediction for nonoscillatory, irregular stimuli and background activity. Reliability should be increased when a cooperative mode (Heidmann et al., 1984) between the structure of the input and the recurrent connectivity of the network is uncovered. The frozen stimulation might be an example of such a mode.

If our stimulation, mimicking the spontaneous statistics, is indeed relevant in an experimental context, a similarity between the structure of spontaneous and evoked activities should be observed, and would correspond to an increased reliability of cortical responses. Similar levels of activity during spontaneous and sensory evoked regimes have been reported in primary auditory cortex, area A1 (deCharms and Merzenich, 1996), but the relationship between the structures of spontaneous and sensory-evoked activities is still a matter of debate in primary sensory areas. In the visual cortex, voltage-sensitive dye imaging in the anesthetized cat has shown similar cortical activity maps for spontaneous and grating-evoked responses (Tsodyks et al., 1999; Kenet et al., 2003). Additionally, in awake ferrets, multiunit recordings have shown that the temporal correlations of the activity remain unchanged when switching from ongoing activity to natural stimulation (Fiser et al., 2004). Finally, Han et al. (2008) have shown that a repeated stimulation modifies the structure of the spontaneous activity over several minutes, such this latter becomes more similar to the previously imposed evoked activity. In A1, it seems that the sets of responses to different types of stimuli are all included in the phase space delimited by the spontaneous activity (Luczak et al., 2009). It would be interesting to compare the reliability of these responses to other types of stimuli. Our results make us hypothesize that the reliability will be higher when the spontaneous and evoked activities are similar.

Conclusion

We have shown that, even in a deterministic network, the reliability of the responses can be modulated by the type of stimuli. As a consequence, to a given recurrent network would correspond a set of stimuli which are more efficiently transmitted than others. According to our results, the spontaneous activity could be a reference point in this set. This supports the hypothesis that similarity between evoked and spontaneous activities is the signature of an efficient mode of transmission across recurrent networks. However, further research will need to delimit this set by searching for other types of stimuli that are as efficiently transmitted.

Having stated that a match between spontaneous activity and input statistics evokes a better recall, we can hypothesize that the connectivity could have been shaped by a learning process so that the spontaneous activity matches the natural input statistics. In the case of a Boltzmann machine with binary neurons, where the inputs are also transmitted by "freezing" some neurons, the learning of the input statistics does induce such a match (Ackley et al., 1985). However, an equivalent learning process for a network of integrate and fire neurons is currently unknown. Nevertheless, the connectivity in our network model can be viewed as the result of an unseen learning process, where the network has learned to transmit more efficiently a particular set of inputs. We can hypothesize that at least part of these inputs are replayed by spontaneous activity. Experimentally, a possible consequence could be that the spontaneous ongoing activity replays the learned "neuronal songs" (Han et al., 2008). A complete implementation of this hypothesis would require a better knowledge of unsupervised plasticity mechanisms during the learning phase. According to our interpretation, the cortical network would efficiently transmit learned, i.e., predicted, patterns, without departing from ongoing activity. This would allow a robust mapping for some features selected by the spontaneous activity.

A more speculative extrapolation is to view the spontaneous activity as what the cortical network "expects" to transmit efficiently. The dynamic changes in the temporal structure of the spontaneous activity could be interpreted as transient switches between different sets of "expectations" or hypotheses made on the basis of the continuously updated incoming sensory flow. More extensive experimental research on the conditions under which ongoing activity in recurrent networks recapitulates fragments of previously learned memories (or "songs") (Louie and Wilson, 2001; Ikegaya et al., 2004) (but see Mokeichev et al., 2007) is needed to consolidate this view.

References

Ackley DH, Hinton GE, Sejnowski TJ (1985) A learning algorithm for Boltzmann machines. Cogn Sci 9:147–169.

Aertsen AM, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of "effective connectivity." J Neurophysiol 61:900–917.

Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science 273:1868–1871.

Aviel Y, Mehring C, Abeles M, Horn D (2003) On embedding synfire chains in a balanced network. Neural Comput 15:1321–1340.

Azouz R, Gray CM (1999) Cellular mechanisms contributing to response variability of cortical neurons *in vivo*. J Neurosci 19:2209–2223.

Banerjee A, Seriès P, Pouget A (2008) Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. Neural Comput 20:974–993.

Baudot P, Levy M, Monier C, Chavane F, Renè A, Huguet N, Marre O, Pananceau M, Kopysova I, Frégnac Y (2004) Time-coding, low noise Vm attractors, and trial-to-trial spiking reproducibility during natural scene viewing in V1 cortex. Soc Neurosci Abstr 30:948.12.

Binzegger T, Douglas RJ, Martin KA (2004) A quantitative map of the circuit of cat primary visual cortex. J Neurosci 24:8441–8453.

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J Comput Neurosci 8:183–208.

Davison AP, Brüderle D, Eppler J, Kremkow J, Muller E, Pecevski D, Perrinet L, Yger P (2008) PyNN: a common interface for neuronal network simulators. Front Neuroinformatics 2:11.

deCharms RC, Merzenich MM (1996) Primary cortical representation of sounds by the coordination of action-potential timing. Nature 381:610–613.

Destexhe A, Contreras D (2006) Neuronal computations with stochastic network states. Science 314:85–90.

Deweese MR, Zador AM (2004) Shared and private variability in the auditory cortex. J Neurophysiol 92:1840–1855.

DeWeese MR, Hromádka T, Zador AM (2005) Reliability and representational bandwidth in the auditory cortex. Neuron 48:479–488.

Diesmann M, Gewaltig M (2001) NEST: an environment for neural systems simulations. Forschung und wisschenschaftliches Rechnen, Beitrage zum Heinz-Biling-Preis 58:43–70.

El Boustani S, Destexhe A (2009a) A master equation formalism for macroscopic modeling of asynchronous irregular activity states. Neural Comput 21:46–100.

El Boustani S, Destexhe A (2009b) Does brain activity stem from high-dimensional chaotic dynamics? Evidence from the human electroencephalogram, cat cerebral cortex and artificial neuronal networks. Int J Bifurcat Chaos, in press.

Fiser J, Chiu C, Weliky M (2004) Small modulation of ongoing cortical dynamics by sensory input during natural vision. Nature 431:573–578.

Frégnac Y, Baudot P, Levy M, Marre O (2005) An intracellular view of time coding and sparseness in V1 during virtual oculomotor exploration of natural scenes. Paper presented at the 2nd International Cosyne Conference in Computational and Systems Neuroscience. Salt Lake City, Utah, March.

Garfinkel A, Spano ML, Ditto WL, Weiss JN (1992) Controlling cardiac chaos. Science 257:1230–1235.

Goldberg JA, Rokni U, Sompolinsky H (2004) Patterns of ongoing activity and the functional architecture of the primary visual cortex. Neuron 42:489–500.

Han F, Caporale N, Dan Y (2008) Reverberation of recent visual experience in spontaneous cortical waves. Neuron 60:321–327.

Heidmann A, Heidmann T, Changeux JP (1984) Selective stabilization of neuronal representations by resonance between spontaneous prerepresentations of the cerebral network and percepts evoked by interaction with the outside world. C R Acad Sci III 299:839–844.

Holt GR, Softky WR, Koch C, Douglas RJ (1996) Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. J Neurophysiol 75:1806–1814.

Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R (2004) Synfire chains and cortical songs: temporal modules of cortical activity. Science 304:559–564.

Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. Nature 425:954–956.

Kuhn A, Aertsen A, Rotter S (2003) Higher-order statistics of input ensembles and the response of simple model neurons. Neural Comput 15:67–101.

Kumar A, Rotter S, Aertsen A (2008) Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. J Neurosci 28:5268–5280.

Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. Neuron 29:145–156.

Luczak A, Barthó P, Harris KD (2009) Spontaneous events outline the realm of possible sensory responses in neocortical populations. Neuron 62:413–425.

Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. Science 268:1503–1506.

Markram H, Tsodyks M (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurons. Nature 382:807–810.

Mehring C, Hehl U, Kubo M, Diesmann M, Aertsen A (2003) Activity dynamics and propagation of synchronous spiking in locally connected random networks. Biol Cybern 88:395–408.

Mokeichev A, Okun M, Barak O, Katz Y, Ben-Shahar O, Lampl I (2007) Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. Neuron 53:413–425.

Monier C, Fournier J, Frégnac Y (2008) In vitro and in vivo measures of evoked excitatory and inhibitory conductance dynamics in sensory cortices. J Neurosci Methods 169:323–365.

Palm G, Aertsen AM, Gerstein GL (1988) On the significance of correlations among neuronal spike trains. Biol Cybern 59:1–11.

Pecora LM, Carroll TL (1990) Synchronization in chaotic systems. Phys Rev Lett 64:821–824.

Rabinovich M, Huerta R, Laurent G (2008) Neuroscience. Transient dynamics for neural processing. Science 321:48–50.

Schreiber S, Fellous J, Whitmera D, Tiesinga P, Sejnowski T (2003) A new correlation-based measure of spike timing reliability. Neurocomputing 52–54:925–931.

Sompolinsky H, Crisanti A, Sommers HJ (1988) Chaos in random neural networks. Phys Rev Lett 61:259–262.

Thomson AM, West DC (1993) Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. Neuroscience 54:329–346.

Tiesinga P, Fellous JM, Sejnowski TJ (2008) Regulation of spike timing in visual cortical circuits. Nat Rev Neurosci 9:97–107.

Timofeev I, Grenier F, Bazhenov M, Sejnowski TJ, Steriade M (2000) Origin of slow cortical oscillations in deafferented cortical slabs. Cereb Cortex 10:1185–1199.

Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. Science 286:1943–1946.

van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory inhibitory activity. Science 274:1724–1726.

van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. Neural Comput 10:1321–1371.

van Vreeswijk C, Sompolinsky H (2005) Irregular activity in large networks of neurons. In: "Les Houches Lectures LXXX" on methods and models in neurophysics, Chap 9 (Chow C, Gutkin B, Hansel D, Meunier C, Dalibard J, eds), pp 341–402. London: Elsevier.

Victor JD, Purpura KP (1996) Nature and precision of temporal coding in visual cortex: a metric-space analysis. J Neurophysiol 76:1310–1326.

Vogels TP, Abbott LF (2005) Signal propagation and logic gating in networks of integrate-and-fire neurons. J Neurosci 25:10786–10795.

Vogels TP, Rajan K, Abbott LF (2005) Neural network dynamics. Annu Rev Neurosci 28:357–376.

9 Discussion

9.1 A general paradigm

Generalization of the Frozen Paradigm Theoretically, the Frozen Paradigm is fairly general, and can be applied to many kinds of neural network model. Neuron models and connectivity can be changed in order to explore the performance of the recall according to several connectivity schemes. Apart from the neural field model, we have tested it on a collection of recurrent spiking network models taken from the literature. In a recurrent network similar to that of Vogels and Abbott (2005), when the synaptic weights are not constant but distributed according to a normal distribution, convergence is the same. Since this gave similar results to the spiking network described earlier, it shows that heterogeneous synaptic weights and delays do not hamper the convergence process. This is a crucial point, because heterogeneous networks are important in avoiding artificial synchronies.

The network described in Brunel (2000) is composed of spiking current based integrate-and-fire neurons, with a connectivity ratio of about 10%, and stochastic noise injected to sustain the ongoing activity. The Frozen Paradigm was applied to this model (data not shown) and results in a completion index reaching 0.1, a weak value but significantly above the baseline. The decrease of the performance compared to the Vogels and Abbott case should be attributed to the high level of stochastic noise injected, a non deterministic component which constitutes an obvious cause for this lower convergence: opposite to the irregular activity produced by the recurrent network, the stochastic noise injected in the model of Brunel (2000) cannot be modulated by any stimulation. Another point which could largely affect the convergence in our frozen stimulation is the addition of "dynamical components" in the model which could, at best, delay the fast convergence. We tested our paradigm on a recurrent model of current based integrate-and-fire neurons including so-called dynamical synapses (Tsodyks et al., 2000) with standard parameters for adaptation and depression. Observing the evolution of the completion index in Figure 15, it appears that the level of convergence is low, but the kinetics is as fast as before. Denser networks than the one considered in the article (connectivity $\varepsilon = 2\%$), but with dynamical synapses can lead to an effective connectivity rather low, and the Frozen Paradigm seems to be still valid.

The "inhibitory gate" Note that to keep the number of parameters as low as possible, the networks used in the previous article for the Frozen Paradigm are completely homogeneous. Results show that in such balanced networks, where inhibition is very strong to counter balance the excitatory inputs, clamping the inhibitory neurons is the key factor influencing the quality of the recall. But inhomogeneities could also affect it. If we consider for example the thalamo-cortical system as a single network, cortical neurons are much more numerous than thalamical ones ($\simeq 50$ fold more). Nevertheless, the thalamical neurons have much more projections onto the cortex, with stronger synapses. A scenario in which the thalamus is the clamped sub-population driving the rest of the cortex should be considered. External inputs from the retina to the LGN are known to be strong and could efficiently drive the thalamic neurons, making them almost-clamped and insensitive to the cortical feedback. This feedback, indeed, is mainly inhibitory and also is known to enhance the precision of the LGN responses, from a trial-to-trial basis (Andolina et al., 2007). Hence the ensemble formed by the thalamus and the layer 4 of cortex, including the feedforwad inhibitory activation of small inteneurons (the so-called "inhibitory gate" described by Bear, Dudek and Friedlander) could

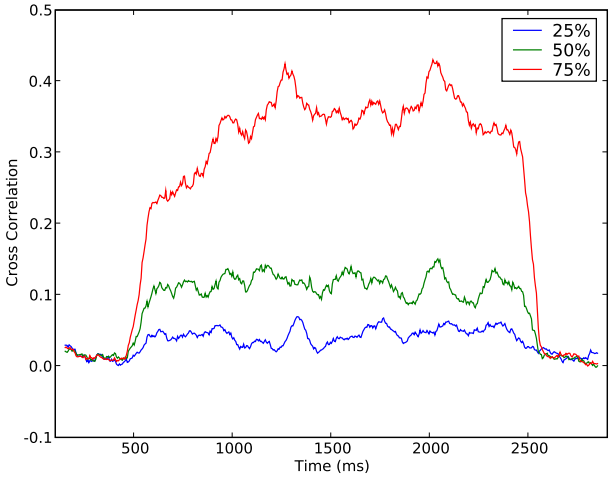

Figure 15: **a)** Time course of the normalized cross correlation between the target and the induced activity for different ratios of frozen neurons, but for a different network model (Tsodyks et al., 2000).

act as a frozen layer, reformatting the drive to cortex and inducing particular trajectories and global recall in the attractor of the ongoing activity of the reamining free cortical units (mostly outside layer 4) (see Figure 16).

Preliminary simulations have also shown that the frozen stimulation would also make networks with more local (Gaussian) connectivity converge, like those that will be studied in further Part III. Since it has been already reported by El Boustani and Destexhe (2009b) that changing the locality of the connections in a sparsely connected recurrent network does not change the macroscopic behaviour of the network, we can hypothesize that similar effects will be produced by the frozen stimulation. Nevertheless, a clear understanding of the dynamical states is a first step to better quantify the convergence in such networks. This will be the subject of the next Part.

Dependence of the convergence on the mean firing rate Our study shows that the number of spikes in the stimulation is a major factor in efficiently driving the free sub-network. The study on parameter dependency demonstrates how the mean firing rate and the recall performance are linked together. However, above 20 Hz, the recall performance seems to reach a plateau and is no longer sensitive to the firing rate. From our results, we can hypothesize that for any neural network there is a spontaneous firing rate which corresponds to a compromise between the "metabolic" cost of spiking and the transmission performance: if the firing rate is too low, the patterns of the ongoing activity, despite their optimality, will be poorly transmitted. The mean firing rate of the irregular activity measured *in vivo* may illustrate this compromise. In particular, the increase of firing rate when switching from sleep to the awake state observed in several cortical areas (see for example Steriade et al. (2001)) may correspond to a need to increase the transmission performance in the awake state.

Chaos control Theoretically, our paradigm is not limited to neural systems. Any dynamical system with several variables could be the subject of such a partial constraint, with the aim of propagate a stored pattern to the rest of the network. The first method that attempted to control

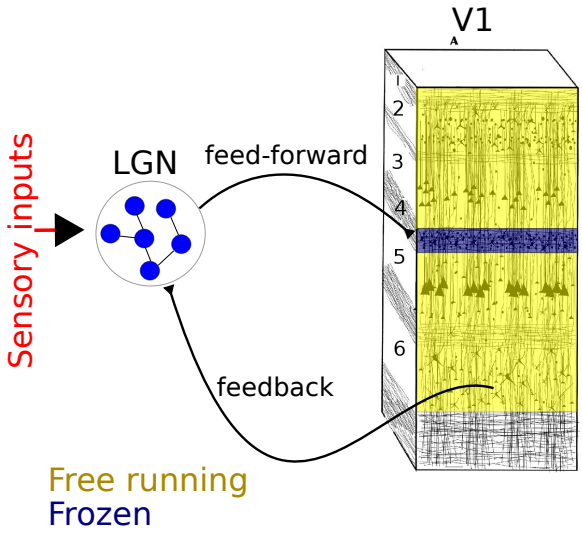

Figure 16: Schematic drawing of the Frozen Paradigm, applied to recurrent network thalamus & cortex. The ensemble formed by the thalamus and the layer 4 of cortex, including the feedforwad inhibitory activation of small inteneurons could act as a frozen layer, constraining the dynamical regime outside layer 4.

a chaotic system was termed the "chaos control" method. The idea was to benefit from the perturbation sensitivity of the chaotic system: a series of small perturbations is designed to drive the system into a determined orbit. This orbit becomes stable under these stimulations. However, this method is designed to drive the system in a very low dimensional attractor. A strategy more similar to ours has been partially used in Pecora and Carroll (1990). In this seminal paper, the authors proposed a method of controlling a chaotic system, and applied it to simple examples. For any dynamical system composed of N variables X_i, they chose a subset of variables and imposed random noise on them, while leaving the rest of the system freely running. Then, they measured whether this stimulation induces a convergence to the same trajectory at each trial. They applied their method to the Lorentz attractor and showed that imposing a random stimulation on the x variable is sufficient to control the y and z variables, since the different trials converge to the same trajectory. Their analytical study links this convergence to the negativity of the Lyapunov exponent of the sub-system.

Both the differences and the common points with our frozen paradigm appear clearly. On the one hand, both methods drive a system by freezing a subpart of it, and compute the reliability of the convergence in the same way. On the other hand, the frozen paradigm method also constrains the stimulation to belong to the attractor originally defining the dynamical system. This constraint allows to predict towards which trajectory the response should converge, even for complex systems with many dimensions, for which an analytical prediction would be impossible (see the work of Teramae and Fukai (2008)). Furthermore, we applied our paradigm to models having attractors of large dimensions, while the chaos control examples were restricted to low dimensional attractors, and the neural field study shows that our paradigm allows a quasi-convergence even in regions where the Lyapunov exponents stay positive. The model consists of N local "neural fields" defined by their activities h_i and with dynamics governed by equation:

$$\frac{dh_i}{dt} = -h_i + \sum_{j=1}^{N} J_{ij} \tan(g h_j) \tag{6}$$

where the synaptic weights J_{ij} are randomly taken from a gaussian distribution $N(0, \frac{J^2}{N})$. As explained in the article, for $N \to \infty$, if $gJ < 1$, the network activity decays to 0, and if $gJ > 1$, the system is in a chaotic regime (Sompolinsky et al., 1988).

The free subnetwork alone has $N_0 = (1-r)N$ units, where r is the ratio of frozen neural-field neurons. The variance of the synaptic weights is the free subnetwork is:

$$\text{var}(J_{ij}) = \frac{J^2}{N} = \frac{(\sqrt{1-r}J)^2}{N_0} \tag{7}$$

The free subnetwork can thus be related to a neural field network of parameters $g_0 = g$ and $J_0 = \sqrt{1-r}J$. It is chaotic when $g_0 J_0 > 1$, i.e. when:

$$g\sqrt{1-r}J > 1 \tag{8}$$

$$r < 1 - \frac{1}{(gJ)^2} \tag{9}$$

The equation of the limit of the chaotic zone (white dashed line) in Figure 9b of the article is thus:

$$r_{\text{lim}} = 1 - \frac{1}{(gJ)^2} \tag{10}$$

This result shows that despite the chaotic nature of the free subnetwork itself, the drive can still be achieved. It has also been shown recently (Rajan et al., 2010) that external inputs can suppress the chaotic nature of the ongoing activity in a network, bringing the system to a well defined state which could be linked to the convergence we are observing.

9.2 On the role of the recurrent connections

Role of the recurrent connections The recurrent connections are an important part of the network dynamics, and their functional role in the coding strategy used by the brain is still unclear. Shadlen and Newsome (1998) proposed that recurrent connections provide an input which maintains the network activity in a delimited range. To demonstrate this, they modelled the input coming from recurrent connections with a Poisson input, and searched for the necessary conditions for the neuron to produce a graded output. It appears that balancing excitation with inhibition is required for this purpose, and this generates as a consequence an irregular spiking output. From this study they deduced that the mean firing rate, averaged over hundreds of neurons, is the only measure which carries information, and does not require precise spike patterns. This implies that the computations carried out by cortical areas can only result from the convergence of feed-forward projections from one area to another in a rate based coding scheme. However, their reasoning is entirely based on the assumption that the irregular activity generated by intra-cortical connections can be modelled by a random walk. Since we know that the neuronal integration of a highly fluctuating current is reliable (Mainen and Sejnowski, 1995), this irregular background can be much more structured than expected and replicated from trial-to-trial even in a sensory drive condition. If the random walk can indeed by a good model for the irregularity of the sub-threshold activity, it is not appropriate to reproduce the trial-to-trial reliability that we have observed in our experimental results. It is also not the only way of generating an irregular activity, since van Vreeswijk and Sompolinsky (1996) have proven that even a deterministic model can generate this irregular

activity. Furthermore, the stochasticity assumption of Shadlen and Newsome is equivalent to assume that there is no cooperation between the input and the recurrent connections. Our model proves the opposite, while still being stable and having irregular activity.

Computational paradigm In his seminal study, Hopfield (1982) proposed a model where a network performs a completion of a pattern imposed on a sub population of units. In this case, the model is a purely recurrent network, fully connected, and the connectivity determines which patterns the network is able to complete. The philosophy of our paradigm is quite similar: the recurrent connections select which patterns will be efficiently transmitted. But we also apply our paradigm to more complex (and more realistic) models than the spin glass model of Hopfield. Furthermore, the Hopfield network, in its original implementation, only stores purely spatial patterns. In comparison, our paradigm can deal with spatio-temporal patterns. Brody and Hopfield (2003) have extended the Hopfield network to encode the stimulus information into spatio-temporal patterns. The principle of their model is that the stimulus is encoded by the latency of the spikes relative to an oscillation which is common to all the neurons. This oscillation constitutes a reference frame for the temporal aspect of the patterns, and is also present in the sub-threshold fluctuations of the activity, even in the spontaneous regime. Their work thus presents many points in common with our paradigm, the main difference being that the reference frame of the oscillatory drive is replaced in our case by the irregular activity sustained by the network. In view of the numerous experimental results showing the irregularity of the cortical activity, we hypothesize that our paradigm is better adapted to understand cortical processing, whereas, as Brody and Hopfield themselves remark, their model presents many similarities with the sensory processing performed in olfactory and hippocampal systems. These differences illustrate how strongly the spontaneous activity can impact on the way information is processed by recurrent networks.

The role played by the recurrent connections in our paradigm is thus close to the idea of Hopfield and the memory paradigm of Ackley et al. (1985): in his case, they participate in selecting some patterns that are better completed and transmitted than others. In the original Hopfield implementation, this was an all-or-none effect: either a pattern is stored, and will be efficiently completed, or it will be brought back to the pattern which includes it in its attraction basin. In our case, the efficiency of transmission decreases continuously when the stimulation pattern deviates from the spontaneous statistics. We can thus see our paradigm as a statistical and temporal version of the same idea: it better transmits the stimulus which is close to the statistics of the spontaneous activity. We can define an "optimal stimulus ensemble" (Machens et al., 2005) as the statistical set of stimuli which are efficiently transmitted by the network. The relationship between the recurrent connectivity and this ensemble is not trivial, but our work shows the identification of this ensemble with the spontaneous activity statistics.

What kind of computations can such a recurrent network model perform? Inspired by Hopfield networks and the Boltzmann machine (see Part IV), we propose that such a network is best suited to map a specific probabilistic input-output relationship between the frozen neurons and the free neurons. To detail this relation, we define the vector $\vec{u}$ describing the possible spiking activities of the frozen neurons ($\vec{u}$ is thus a vector of the spatio-temporal activity), and $\vec{v}$ the vector describing the activity of the free neurons. Any probabilistic input-output relationship can be defined by a conditional distribution $P(\vec{v}|\vec{u})$. Our proposition is that the recurrent network model is best suited to map the relation $P_s(\vec{v}|\vec{u})$, defined by

$$P_s(\vec{v}|\vec{u}) = \frac{P_s(\vec{v},\vec{u})}{P_s(\vec{u})} \tag{11}$$

where $P_s(\vec{v},\vec{u})$ is the distribution of the spontaneous activity over the whole network, and $P_s(\vec{u})$ is the distribution of the spontaneous activity over the frozen neurons. This latter is the optimal stimulus ensemble that we defined above. The recurrent network is best adapted to this probabilistic relation since the reliability of the response is maximal when the input respects the spontaneous activity statistics. More formally, we can say that, with the constraints of keeping the mean firing rate equal to the mean spontaneous activity, this conditional distribution $P_s(\vec{v},\vec{u})$ is probably the one with the lowest conditional entropy. This means that it gives the most reliable relation between $\vec{u}$ and $\vec{v}$. This is reminiscent of the theoretical foundations of the Boltzman machine, and many other representational models. The difference is that the Boltzmann machine aims at finding the stochastic binary recurrent network that best maps a given probabilistic input-output relationship. It learns the statistics of an input ensemble. Here, we start from a given recurrent network, and we try to find the input-output relationship that is best mapped by this network. Both approaches converge to the conclusion that the chosen input-output relation is found in the spontaneous activity statistics. As mentioned before, the symmetry of these approaches makes us consider the recurrent network model as the product of an unknown learning process of $P_s(\vec{v}|\vec{u})$.

Following this idea of an associative memory, presentation of the subpart $\vec{u}$, played only through the clamped neurons, can trigger the activation of a more complex pattern $\vec{v}$. In theory, the dimensions of $\vec{v}$ can be much higher than those of $\vec{u}$, if the network is built appropriately, with some heterogeneities. Therefore the network could be used for information compression: suppose you have data that you can convert into temporal streams of information, for example by temporal latencies within the spike trains. By inverting the problem and clamping the free neurons, one could generate a key $\vec{v}$, signature of the data. With some learning rules, like those that will be explained later (see Part IV), one could also try to shape the network connectivity in order to better store the pattern in its dynamics, since it has been shown to be feasible to embed synfire chains in asynchronous irregular regimes (Kumar et al., 2008a). This key $\vec{v}$ can then be used to replay the data, and the dimensionality reduction between $\vec{u}$ and $\vec{v}$ would be the compression factor. The missing part of these numerical simulations is about the capacity of such a system. Since the recall is not perfect, what are the key parameters constraining the size of the attraction basins for each pattern? Further studies need to be carried out to answer this question, in order to have a tractable and working system.

A learned match between evoked and ongoing activity The model we used is a generic one, and cannot be considered as a realistic model of V1. More realistic models will be explored in the following Part, mainly in the interest of understanding their dynamics. Nevertheless, the concept of the Frozen Paradigm leads to interesting experimental predictions. The Frozen Paradigm implies that the statistics of the spontaneous activity are efficient stimuli for driving the network and enhancing information transmission for a constant energy level in the stimulation. This should be linked with the experimental results shown previously in V1 neurons (see Figure 14). Neuronal responses in response to natural images, are more sparse and time locked than those in response to surrogate inputs, such as drifting gratings. Therefore the natural stimuli seem to correspond to an efficient stimulation, in the sense that they are able to evoke reproducible and irregular responses despite the low mean firing rate. One can hypothetize that the particular statistics of the visual world have been learned by the visual cortex during epigenesis (Spinelli and Jensen, 1979, Sur et al., 1999, Frégnac, 1999, Wörgötter et al., 1998, Kenet et al., 2003, Tsodyks et al., 1999) and that this learning has established a match between ongoing activity and external inputs. In the absence of any inputs, the system wan-

ders on the attractor of the ongoing activity, and therefore replay part of statistics it has been trained to. The whole conceptual picture then becomes that the cortical network learns statistics to which it can then respond in a reproducible manner, and the consequence would be that these responses are then replayed during spontaneous activity. This argument is favoured by the previous result based on the analysis of the power spectrum density of cortical neurons *in vivo*. The fact that the scaling exponents are very similar during spontaneous activity and natural scene stimulation indicates that both conditions trigger a similar amount of correlations within the network.

With the Frozen Paradigm, we implicitly make the hypothesis that the patterns played spontaneously by the network can be seen as the result of an unknown learning process: we do not control it, but a random network has a particular connectivity matrix, constraining the attractor of its ongoing activity, and defining a subspace of all possible dynamics. A good way to test it would be to design two random networks with slightly different structures, in order to have distinct dynamics, and to use the activity generated by the first one to drive the other. One should expect a better reliability if self-generated statistics are used, compared to the case where the activity generated by the other network is used.

If we now come back to the results observed by Fiser et al. (2004) (see Figure 12), we can reinterpret it under the light of this match with ongoing activity. Our theoretical work provides a reason why the activity structure does not change too much when imposing complex sensory inputs (natural images). According to our interpretation, the absence of changes would not mean that the sensory input does not drive the cortical activity. Rather, it shows that this drive is achieved without altering the global structure of the cortical activity. Taken together, these results show that our hypothesis is compatible with recent work that studies the properties of the spontaneous activity in the visual cortex. The issue thus remains to be examined more carefully in experimental studies.

Link with macroscopic fMRI observations Despite the success of our approach in accounting for the replay of particular statistics, the kind of networks we are using is more adapted to the scale of only one cortical area. At this level, a first order prediction of our model is that the change of mean activity should be minimal for certain types of stimuli, for which the information would be efficiently transmitted. For more macroscopic structures, functional Magnetic Resonance Imaging (fMRI) recordings offer a clean and nice view, at the population level, of this ongoing activity and of the fluctuations in this resting state (see Fox and Raichle (2007) for a review). Some fMRI studies have shown the absence of differential responses of primary visual areas to visual stimulation (Murray et al., 2002) interpret this absence of response as the result of a suppressive feedback, but the Frozen Paradigm hypothesis suggests that this absence of responses could just be due to the fact that these signals match the input statistics learned by the V1 network. When stimulating the area with patterns sharing statistical similarities with the ongoing activity, the response of the network is only slightly affected and this is hard to observe from an activity point of view. Most of the responses described in both fMRI and electro-encephalography signals correspond to transient responses to static stimuli. These transient variations have often been reported as corresponding to "unexpected" signals, with examples such as the mismatch negativity in MEEG. fMRI signals have also been interpreted as expectations about the environment (Pouget et al., 2003, Fox and Raichle, 2007). The fMRI response often decreases with the number of successive repetitions. It is often assumed that this decrease of activity corresponds to an adaptation effect at the neuronal level. Our theoretical results may be the future basis for an alternative interpreta-

tion, which would need further work, but can be tested. The network models we have used are "fixed": they do not have any capacity of adaptation or learning. But if we hypothesize now that for models incorporating adaptation, our paradigm still works, we can then picture that the "optimal stimulus ensemble" defined above will be dynamically changed with adaptation. It means that this kind of network will progressively learn the repeated test input statistics during successive trials. Consequently, the network response signalling deviation from the initial ongoing statistics will progressively disappear by plasticity (see Part IV). Ultimately, when this inclusion will be achieved, there will not be any more macroscopic responses to this stimulation. We will address this question of learning and adaptation in Part IV of this manuscript.

Part III

Structure of the ongoing activity

Summary

10 Introduction **82**
10.1 Toward more realistic networks 82
10.2 A link with macroscopic measures 82

11 Macroscopic invariants in topological networks **84**
11.1 Main results 84
11.2 Supplementary figures 106

12 Discussion **110**
12.1 General comments 110
12.2 Toward a mean-field-based description of topographical networks 112

10 Introduction

10.1 Toward more realistic networks

In the case of the Frozen Paradigm, we showed that the statistics of the spontaneous activity (and therefore its correlations) have a particular role when spontaneous patterns were used as external inputs. Comparison with surrogates patterns, reproducing first- and second-order statistics of the ongoing activity (Synchronous Poisson) produced a lower signal to noise ratio (SNR). This result implies that the higher-order correlations of the ongoing activity should recruit more non-linear interactions and may help the system to enhance its reliability. We also showed that these correlations, present in the ongoing activity, were similar to those evoked *in vivo* and in V1 neurons when the system was stimulated with complex stimuli, such as natural images. To be in line with the literature on balanced random networks, the Frozen Paradigm was implemented in a network as generic and as homogeneous as possible, with well known dynamics. Parameters were well controlled and studied, to centre the attention on information transmission and the design of the paradigm.

A natural step would be to extend this to topographical and other more realistic networks. This would allow the study of how the convergence of the recall depends on the detailed topology (defined here in terms of probability of connections), and how the frozen activity spreads into the rest of the network. Even if homogeneous and random balanced networks are interesting because they capture properties of the neuronal dynamics, they lack more realistic characteristics that can be found *in vivo*. Considering generic and simple wiring schemes can lead to analytical frameworks where key coding properties such as correlations can be described and predicted. This is for example the work of Kriener et al. (2009), analysing the distribution of pairwise cross-correlations in several random networks in the Asynchronous Irregular regime. As we saw in Part I, correlations are crucial because they are thought to be a major element in neuronal interactions and the basis of the temporal code. Therefore an important step, before extending the Frozen Paradigm to more biologically realistic networks, is to perform a clear and exhaustive study of topographical networks, and especially on how correlations organize in those networks. As already noted in Part I, the balanced random network is one of the simplest kinds of network able to generate an irregular activity. But cortical organization seems to be more organized, and anatomical data do not suggest that the connectivity is random. Without knowing if this connectivity is the result of a learning process or hard-wired, we can simply observe that, for example, in V1 neurons are connected in a patchy manner, according to the underlying orientation maps (see Figure 17). The exact structure of the canonical microcircuits within layers can be assessed by dual recordings *in vitro* (Binzegger et al., 2004, Thomson and Bannister, 2003), but such complex organizations, where connections are made not only as a function of distance, but also as a function of other features (preferred orientation, direction selectivity, ...) introduce more complex correlation patterns, and it is important to know what are the relevant quantities that may control such correlation levels.

10.2 A link with macroscopic measures

The aim of this Part is therefore to provide networks and simulations that will give a better insight about correlations in structured neuronal networks of spiking neurons. To target generic results, the key point is to keep the network as simple as possible, and to try to dissect its re-

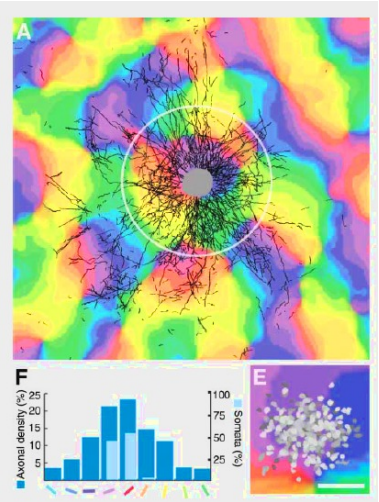

Figure 17: Taken from Stettler et al. (2002). Patchy lateral connectivity in macaque primary visual cortex: axons from an injection labelling 320 cells in the superficial layers of V1 are super-imposed upon the optical imaging orientation map for that portion of cortex. The white ring is 1 mm in diameter. While the axons proximal to the injection site display little orientation specificity, patches that form $\simeq$ 500 μm from the injection are targeted to orientation domains of similar preference to their cell bodies.

sponse without exploring the huge parameter space of all the parameters. Understanding the dynamics of such networks is important while recording techniques are nowadays able to give a glimpse of the macroscopic activity of a large cortical surface. For example voltage sensitive dye imaging (VSD) now allows recording of surfaces up to several square millimetres, and on that scale, the approximation of the balanced random network without any propagation delay cannot stand any more. Even if this approximation can be considered as valid in a small volume, understanding large scale phenomena such as synchronization within areas, or travelling waves (Han et al., 2008, Contreras, 2007, Mohajerani et al., 2010), is an important step. So if one wants to be able to know what are the relevant parameters influencing the dynamical states of topographical networks, one need to simulate them: simulations are important in order to be able to build a topographical description, with rate-based units, or neural fields, of large cortical networks. Since theoretical studies on balanced random network can accurately describe the dynamics without topology, with a master equation formalism (Brunel and Hakim, 1999, El Boustani et al., 2009), the next step is to have equations and predictions that can be matched to simulations with a structured connectivity, and according to some inhomogeneity in the coupling scheme (Jradeh, 2010). This work is also done with continuous neural field units (Wilson and Cowan, 1972, Amari, 1977, Bressloff, 2002). An encouraging step in the direction of a topographical description with a master equation formalism, without struggling with the number of parameters, is the demonstration by El Boustani et al. (2009), that some invariant in the connectivity may be relevant to accurately describe the dynamics: for balanced random networks with integrate-and-fire neurons and current based synapses, the number of incoming synapses per neuron N_{in} is the only quantity of the connectivity relevant for describing the stationary state of the system, up to its second order statistics (averaged firing rate and variance). Building a topographically organized network with a Gaussian probability of connection (see following article for more details) does not alter these estimates, as long as N_{in} is kept constant. Motivation of this Part is then to study the existence of such invariants, with a particular attention on the cross-correlations profile within the network. This is a im-

portant step for the understanding of the parameters driving the dynamical regimes observed with these particular layered networks and for establishing a link, through mean-fields models, with macroscopic measures now recorded *in vivo* (VSD, local fields potentials (LFP)). Since one cannot exclude that the mean-field assumptions themselves cancel the sensitivity of the model to microscopic irregularities, invariants help to fill the gap between micro and macroscopic models.

The other interesting question, already raised in the previous Part, with the Frozen Paradigm and the debate about the similarities between ongoing and evoked activities, is how the correlation structure is affected by the presentation of a stimulus. In recent experimental data, such as spike-triggered LFP data (Nauhaus et al., 2009), one can see how changing the input (increasing the contrast) can change the profile of the spatio-temporal correlations in V1. The more the network is stimulated, the more the spatial spread of the correlations is reduced. Inhibition is more active and tends to decorrelate the cells over large distances. This results was already observed in the PLoS article presented in the previous Part (see Figure 3, panel E).

11 Macroscopic invariants in topological networks

In the following article, we studied the spatial profile of pairwise cross-correlation coefficients in topographically organized networks of integrate and fire neurons, connected according to Gaussian profiles and including linear propagation delays with a finite velocity. The main regime observed was the Synchronous Regular regime, displaying oscillations and waves of activity close to those reported in voltage sensitive dyes imaging studies. I design the network and the experiment in collaboration with one of my colleague, Sami El Boustani, and we both worked equally on the project.

11.1 Main results

Journal of Computational Neuroscience manuscript No.
(will be inserted by the editor)

Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons

Pierre Yger* · **Sami El Boustani*** · **Alain Destexhe** · **Yves Frégnac**

Received: date / Accepted: date

Abstract The relation between the dynamics of neural networks as a function of the patterns of connectivity is far from clear, despite its importance for understanding functional properties. Here, we have studied sparsely-connected networks of conductance-based integrate-and-fire (IF) neurons with balanced excitatory and inhibitory connections, and with finite axonal propagation speed. We focused on the genesis of states with highly irregular spiking activity and synchronous firing patterns at low rate, called slow Synchronous Irregular (SI) states. In such balanced networks, we examined the "macroscopic" properties of the spiking activity, such as ensemble correlations and mean firing rates, for different intracortical connectivity profiles ranging from randomly connected networks to networks with Gaussian-distributed local connectivity. We systematically computed the distance-dependent correlations at the extracellular (spiking) and intracellular (membrane potential) levels between randomly assigned pairs of neurons. The main finding is that such properties, when they are averaged at a macroscopic scale, become invariant in relation with the different connectivity patterns, provided the excitatory-inhibitory balance is the same. In particular, the same correlation structure holds for different connectivity profiles. In addition, we examined the response of such networks to external input, and found that the correlation landscape can be modulated as a function of the mean level of synchrony imposed by the external drive. This modulation was found again independent of the external connectivity profile. We conclude that the first and second-order "mean-field" statistics of such networks do not depend on the details of the connectivity at a more microscopic scale. This study is an encouraging step toward a mean-field description of topological neuronal networks.

Keywords Topological networks · Mean-field · Correlation · Synchronous · - Intracortical connectivity

[1] co first-authors

Unité de Neuroscience Intégratives et Computationnelles (UNIC), UPR 2191 CNRS
Bat. 32-33, 1 Avenue de la Terrasse, 91198 Gif-sur-Yvette, France
Tel: 33-1-69-82-34-15, Fax: 33-1-69-82-34-27
Corresponding author: Alain Destexhe, E-mail: destexhe@unic.cnrs-gif.fr

1 Introduction

Spatio-temporal correlations are a key signature of the cortical population spiking discharge, measured in the recurrent spontaneous ongoing activity or from sensory-driven activity. Correlations have been classically considered as a crucial component of the neuronal assembly code (Singer and Gray, 1995; Nirenberg and Latham, 2003) also linked to behavior (Zohary et al, 1994). Today, a more precise insight about their spatiotemporal structure, and the supra- and/or sub-threshold level, is given by analysis techniques using intracellular recordings and modeling (Destexhe and Paré, 1999; El Boustani et al, 2009), multi electrode arrays (Smith and Kohn, 2008) or 2-photon imaging (Göbel et al, 2007; Greenberg et al, 2008). Understanding how these correlations emerge in recurrent neuronal networks and how their structure could be related to some generic network properties can help assessing their functional role and their relation or independence with the local built-in microscopic anatomical connectivity.

One important issue about neuronal correlations lies also in the way they are modulated by external stimulation. It is indeed well known that neuronal pairwise correlations can be affected by the presentation of a stimulus (Kohn and Smith, 2005; Nauhaus et al, 2009; Smith and Kohn, 2008; Mitchell et al, 2009), but the way they can change as a function of the stimulus statistics is poorly understood. Identifying the generic properties of the input that can influence the pairwise cross-correlation profile within sensory areas could provide an estimation of the sensory input properties knowing some experimental functional measurements.

In the superficial layers of the primary visual cortex, correlations are clustered as a function of the underlying orientation maps constraining the correlated inputs that are more frequently seen by a particular neuron (Berger et al, 2007; Nauhaus et al, 2009), but they can also span a large cortical surface (Smith and Kohn, 2008; Schwarz and Bolz, 1991), decaying with distance. In this context, the distance-dependent profile of the pairwise cross-correlations could be used to gain some knowledge about the underlying hard-wired connectivity.

The balanced random network (van Vreeswijk and Sompolinsky, 1996, 1998; Brunel, 2000; Vogels and Abbott, 2005; Kumar et al, 2008; El Boustani and Destexhe, 2009; Amit and Brunel, 1997; Renart et al, 2010) is a common and convenient framework to study the dynamic of large-scale populations of sparsely-connected integrate-and-fire neurons, and to reproduce the so-called slow Synchronous Irregular (SI) states observed *in vivo* (Brunel, 2000). In such regimes, neurons are firing in an irregular manner, behaving almost like Poisson processes, and the average pairwise cross-correlations value is modulated by the internal balance or the external input. This regime is also well suited to produce slow oscillations comparable with observed oscillations *in vivo* (Han et al, 2008; Arieli et al, 1996). and such models, analytical techniques can be used to study the distributions of the pairwise cross-correlations for some topological profiles and network regimes (Kriener et al, 2009).

However, despite its generality, this classical model with random connectivity lacks several important biological features, which complexifies the analytical approach of the problem. In this paper, we chose to study a more realistic two-dimensional network of integrate-and-fire neurons, which is more relevant biologically since it include propagation delays (Bringuier et al, 1999; Benucci et al, 2007) and conductance-based synapses (Vogels and Abbott, 2005; Cessac and Viville, 2008; Kumar et al, 2008; Marre et al, 2009). We provided a detailed numerical study of its spatio-temporal correlations for Gaussian connectivity profiles, previously introduced in the context of informa-

tion processing (Mehring et al, 2003). Such a model would be a minimal model to capture propagation phenomena which can be directly observed *in vivo* with large scale recordings (Voltage-Sensitive Dye Imaging, multi electrode recordings, 2-photon Imaging). Note that previous studies considered topological networks with irregular firing (Usher et al, 1994; Kitano and Fukai, 2007), but with different modeling and connectivity paradigms than here.

In the first part, we study the organization of the pairwise cross-correlations as a function of the distance in generic 2D networks of integrate-and-fire neurons subject to unstructured input, a case that will be referred as the spontaneous activity. This irregular but tonic bombardment is supposed to simulate the effect of the retinal "dark discharge" in thalamocortical visual networks which is detected in absence of any visual drive. We characterize the correlation profiles as a function of distance between pairs of neurons and their sensitivity when varying key parameters of the microscopic network structure. In the second part of the paper, we study the behavior of the same network when driven with synchronous inputs to study how the profile of its pairwise cross-correlations is affected.

2 Materials and Methods

Neuron model : The simulated networks were composed of 12500 (10000 excitatory and 2500 inhibitory) conductance-based leaky integrate-and-fire neurons with a membrane time constant $\tau_{\rm m} = 20$ms, a leak conductance of $G_{\rm leak} = 10$nS, and a resting membrane potential $V_{\rm rest} = -80$mV. When the membrane potential $V_{\rm m}$ reaches the spiking threshold $V_{\rm thresh} = -50$mV, a spike is generated and the membrane potential is clamped to the reset potential $V_{\rm reset} = -60$mV during a refractory period of duration $\tau_{\rm ref} = 5\,$ms. These parameters were kept fixed and were chosen as biologically plausible and in line with previous studies. Only the refractory period $\tau_{\rm ref}$ was varied in additional simulations to be sure that the results found, based on 0 time-lag correlations, were not qualitatively affected by this value. Others simulations were also performed to check the validity of the result with larger networks, up to 100,000 neurons.

Synapse model : The synaptic interactions between these neurons were modeled as transient conductance changes. The synaptic time course was modeled as an instantaneous rise followed by an exponential decay. The synaptic time constants were chosen to be $\tau_{\rm exc} = 3$ms and $\tau_{\rm inh} = 7$ms for excitation and inhibition respectively. The reversal potentials were $E_{\rm exc} = 0$mV and $E_{\rm inh} = -70$mV.
The complete set of equations describing the dynamic of a neuron is thus given by

$$\begin{aligned}\tau_{\rm m}\frac{dV(t)}{dt} &= (V_{\rm rest} - V(t)) + g_{\rm exc}(t)(E_{\rm exc} - V(t)) + g_{\rm inh}(t)(E_{\rm inh} - V(t))\\ \tau_{\rm syn}\frac{dg_{\rm syn}(t)}{dt} &= -g_{\rm syn}(t) + S_{\rm syn}(t)\end{aligned}$$

where syn $\in \{\text{exc}, \text{inh}\}$, $S_{syn}(t) = \sum_{i,k}\delta(t - t_i^k)$ are the incoming synaptic spike trains where $i \in \{1, .., N\}$ refers to presynaptic neurons and k to the different spike times of these neurons. Here $g_{\rm syn}(t)$ is expressed in units of the leak conductance. In this paper, we used 4nS for the excitatory conductance, and a balance of $g = 16$ unless stated otherwise. The main parameters are summarized in Table 2

Spatial organization : A cortical area of 1 mm^2 was simulated as a 2D-layer-like network with periodic boundary conditions and an excitatory/inhibitory neuron number ratio of 4:1. Note that since the density of neurons is arbitrarily selected, and since the density varies between species and cortical areas (Braitenberg and Schüz, 1998), this value of 1 mm^2 should not be taken as realistic. This scale is more in the order of a V1 hypercolumn where local circuits preveil and long-range horizontal connections are not included, being beyond the network size. It is therefore distinct from the larger one usually used in neural mass models. Neurons were arranged on a grid, and even if such a regular structure may bias slightly their connectivity (Voges et al, 2007), we checked with additional simulations that the results remain the same if neurons had been drawn uniformly across the layer. Every neuron was sparsely connected with the rest of the network with a connection probability that depends on the distance l_{ij} between two neurons in the network through a Gaussian profile

$$p_{ij} = e^{-\frac{l_{ij}^2}{2\sigma_c^2}} \tag{1}$$

where σ_c^2 is the variance of the connectivity profile, i.e. the spatial spread of the Gaussian profile. For each neuron, K incoming connections are drawn by randomly picking other neurons in the network that will or not create a projection according to a rejection method based on the Gaussian profile. The total number K of synapses per neuron was fixed, so whatever the σ_c value, each neuron keeps the same number of incoming synapses for the sake of comparability. Two connection densities were mainly studied: a highly sparse one (with a connection density $\epsilon = 0.5\%$) and a denser one (with $\epsilon = 5\%$). The network was considered to be in a spontaneous state when an unstructured and stationary external input was fed into it. In the case where the local network was stimulated, an other layer-like network projected onto the cortical network in a topological manner described by another Gaussian distribution with σ_{ext}^2 as a variance.

Delays : We used non-homogeneous delays, which depend linearly on the distances l_{ij} through

$$d_{ij} = d_{\text{syn}} + \frac{l_{ij}}{v} \tag{2}$$

where v is taken from the literature (Bringuier et al, 1999; Gonzlez-Burgos et al, 2000). A classical value of 0.1-0.5 m/s is usually reported, and in all simulations, we used $v = 0.2$m/s, and $d_{syn} = 0.2$ms.

Simulator : All simulations were performed using the **NEST** simulator (Diesmann and Gewaltig, 2001) and the **PyNN** interface (**?**). Correlated input in the external layer was built by combining Poisson processes and the Multiple Interaction Process (MIP) algorithm (Kuhn et al, 2003).

Data Analysis : Since the maximal distance between two neurons in our network is $\frac{\sqrt{2}}{2} \simeq 0.7$mm, we divided the spatial domain in slices of 50μm width and computed the distance-dependent spiking correlation by selecting, for each distance slice, random pairs of neurons with a distance between them falling into the range imposed by the slice boundaries and we estimated the Pearson coefficients on the spike counts. To be more precise, spike trains were digitized with a time bin equal to the refractory period of the

τ_{m}	20ms
τ_{ref}	5 ms
τ_{exc}	3 ms
τ_{inh}	7 ms
V_{rest}	-80 mV
V_{thresh}	-50 mV
V_{reset}	-60 mV
E_{exc}	0 mV
E_{inh}	-70 mV
G_{leak}	10 nS
Δg_{exc}	4nS
Δg_{inh}	$16\Delta g_{exc}$=64nS
d_{syn}	0.2 ms
v	0.2 m.s^{-1}
simtime	5500 s
ϵ	$0.5\%, 5\%$
σ_c	$[50, 1000]$ μm

Table 1 Parameter table summarizing all the cells and network parameters used in the simulations.

neurons, i.e 5 ms, and for each slice, we selected 2000 pairs of neurons before averaging the Pearson correlation coefficients computed over all these selected pairs. Cells that remained silent during the simulations were discarded from the analysis, and spiking data were gathered during 5s of stationary simulation. For the subthreshold activity, we selected a row of neurons in the network and we computed the cross-correlation at time lag 0 for each pair as a function of the interconnection distance. The characterization of this function was done by computing its integral value over distances ('Integrated correlation') and its linear slope in a log-log representation ('Scaling exponent of the correlation') (see Figure 5). The integrated correlations can be directly related to the average synchrony in the network while offering a more intuitive representation in term of distance-dependent correlations. For the network activity spike-triggered average, we sampled randomly 100 neurons for which average was computed over the whole spike train ensemble.

3 Topological Network Model

Cortical connectivity is still poorly understood, but is definitely not as random as it is usually modeled in previous studies. Whether the connectivity graph is small-world (the definition is ambiguous when considering propagation delays), clustered, or Gaussian is still unclear, but biological evidence shows that neurons in the cortex project mainly to their surrounding (Hellwig, 2000; Bienenstock, 1996). As a first approximation, neurons can be considered as being connected with a distance-dependent probability following a Gaussian profile. Even if it is well known that realistic connectivity is less isotropic and homogeneous (see for example in V1 the orientation maps and the patchy horizontal connectivity (Gilbert and Wiesel, 1983)), the Gaussian profile is a good description of a small cortical area where long-range interactions are ignored. Therefore, every neuron in our model is connected with the rest of the network with a 2D Gaussian probability function and a fixed number of incoming synapses, while periodic boundary conditions are used throughout the study to avoid any boundary effects. In order to obtain *in vivo*-like states, we adopted the usual integrate-and-fire balanced network

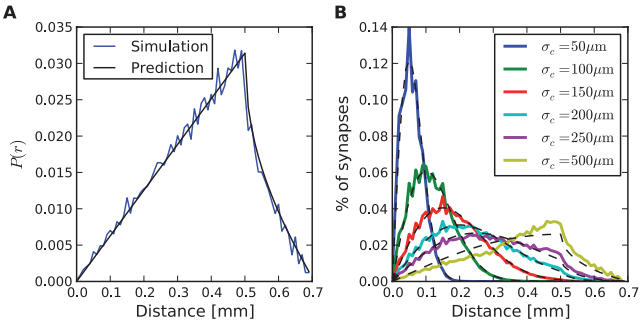

Fig. 1 Profile of the connections. **A**: Number of synapses as a function of the delay (and therefore distance) between neurons in the network, for several values of the connection spread, σ_c. The corresponding mean values are represented as dotted lines. **B**: Distribution of interneuron distances from one excitatory cell to all the other cells in the network

.

configuration comprising a ratio 4:1 between excitatory and inhibitory neurons (Brunel, 2000). The synaptic weights were chosen accordingly to obtain a balanced sub-threshold fluctuating dynamics responsible for the irregular firing.

Propagation delays are known to lead to a large diversity of states in large-scale neuronal networks (Roxin et al, 2005; Izhikevich et al, 2004). While they are often discarded in large-scale models, under the assumption that they could be neglected in a small cortical area, biological studies (Bringuier et al, 1999; Gonzlez-Burgos et al, 2000) have reported typical values of 0.1-0.5 m/s for conduction delays, and comparable values can be observed in Voltage Sensitive Dye Imaging, where activity waves propagate at a similar speed (Grinvald et al, 1994; Benucci et al, 2007; Nauhaus et al, 2009). Patch recordings *in vitro* also confirm that these delays scale linearly as a function of distance (Larkum et al, 2001) when considering the propagation from dendrites to the soma. Thus, even for a small patch of cortex of 1mm^2, with a synaptic delay of 0.2ms (due to neurotransmitter release), conduction delays are broadly distributed and should not be neglected. Moreover some artificial oscillations could arise in network where delays are homogeneous (Brunel, 2000). Our network was therefore built as an artificial square lattice of 1 mm^2 and we chose a propagation speed of $v = 0.2\text{m/s}$.

To have a clear picture of the network structure, one can have a look at Figure 1A where the distribution of the delays in the network as a function of the Gaussian spread σ_c is plotted. By construction, the distributions are continuously affected by σ_c, and are not Gaussian. Indeed, these functions are the product of the Gaussian profile of connectivity and the probability to find a pair of neuron for a given distance (see Figure 1B).

A recent study showed that the spike-triggered LFP in the awake animals could reflect the cross-correlation between LFP the LFP and the membrane potential of the very same cell. This result show that in term of correlations, the spiking activity of a cell convey the same information than the subthreshold response which is n coherence with our result for low rate regime

During the so-called spontaneous activity, every neuron in the network is stimulated with decorrelated Poisson input. Even if in terms of numbers of synapses the synaptic drive of cortical neurons in V1 originates mainly from the recurrent network, the efficacy

of the feedforward thalamocortical synapses is the largest (Gil et al, 1999). To simulate the functional balance between recurrent and feedforward input, each neuron in the cortical layer received the same number of external synapses than the recurrent ones. When we considered correlated input, an external layer was added on top of the network where external units produced synchronous Poisson spike trains projecting on a subset of the 2D network with Gaussian probability distributions. There was no delays from the stimulation layer to the recurrent network layer.

4 Response under an unstructured noise

In the spontaneous activity regime, i.e when uncorrelated Poisson external noise was applied to all the synapses at a mean frequency of $\nu_{\text{ext}} = 5$ spikes/s., the network displayed waves emerging at random places which tended to propagate all over the surface. It should be noted that such networks are not able to maintain self-sustained activity by themselves. Several studies previously reported that these ongoing and reverberating regimes could be observed in networks with conductance-based synapses (Vogels and Abbott, 2005; Kumar et al, 2008; El Boustani and Destexhe, 2009; Marre et al, 2009), but they were all achieved in networks without any propagation delays. The linear propagation time taken into account here (see Materials & Methods) increases the average synaptic delay and therefore the neuron density that would have been necessary to observe such a spontaneous regime. The average delay within the present network is close to 0.75 ms for $\sigma_c = 100\mu$m, and depends on σ_c (see Fig. 1A). Nevertheless, it has been shown that a weak uncorrelated external input does not alter the main statistical features of these models (Brunel, 2000; Vogels and Abbott, 2005) and their irregular activity could not be completely explained by the stochastic nature of this background activity, such that we can still study the interplay between spontaneous and evoked correlated activity.

As can be seen in Figure 2, large sub-threshold waves developed at the conductance level. Waves of excitation, popping up at random places, were immediately followed by an increase of inhibition, traveling simultaneously across the network due to the connection rule and the delays. These waves are reflected in the supra- and sub-threshold activity usually experimentally recorded (Han et al, 2008; Nauhaus et al, 2009).

To gain some insights on the role of the network structure, we varied in a systematic manner two main parameters : the spatial extent of the Gaussian profile used for the recurrent connections σ_c and the balance between excitatory and inhibitory synaptic strength g. The results are shown in Figure 3, for two different connectivity densities, $\epsilon \in \{0.5\%, 5\%\}$ to show the generality of the results. The first striking observation is that for both connectivity densities, averaged quantities such as the mean firing rate and the mean coefficient of variation of the inter-spike interval (ISI CV) do not depend on the connectivity spread σ_c but are only controlled by the balance g between excitation and inhibition (see Fig. 3A,B,E,F). A similar result has been previously observed in (El Boustani and Destexhe, 2009). At the population level, the only relevant parameters for these macroscopic quantities are therefore the average number of synapses received per neuron and their respective strength, not the precise lay-out of the recurrent connectivity. There is somehow a match between the level at which simple anatomical details should be taken into account and the scale of measurement. Of course, this is true as long as the connections are sparse enough: in the limit case of very small σ_c, we are almost in a first-neighbors situation and the local correlations are

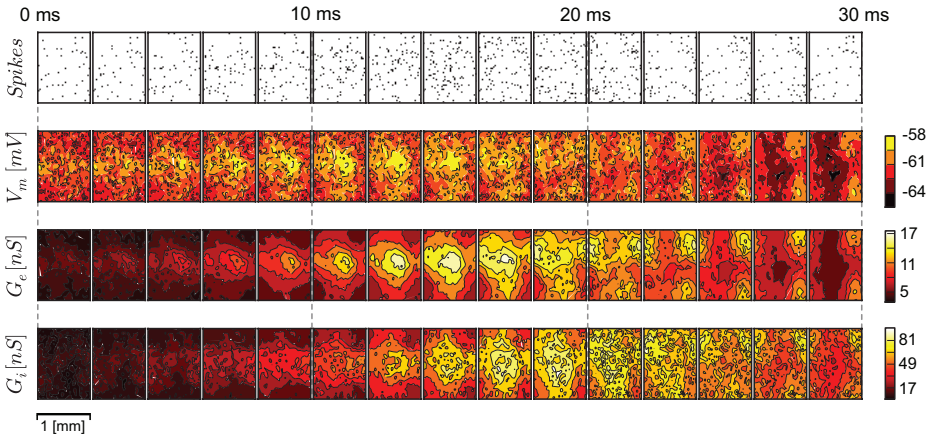

Fig. 2 Snapshots of the spontaneous ongoing activity in the 2D network. Neuronal responses at the supra-thresholds (spikes) and sub-threshold (V_{m}) levels for every neuron in the 2D network are shown respectively in the top and second rows. Snapshots are taken every 2 ms for a total duration of 28 ms. The last two rows shows instantaneous input conductance maps. Excitatory conductances are represented on third row and the inhibitory conductances in the last row.

too strong to keep averaged quantities invariant. This is an extreme situation where usual mean-field models are not valid anymore.

The second result concerns the irregular and oscillatory nature of the dynamics. Typical raster plot of the observed activity regime (Fig. 3D,H, $g = 16.$ and $\sigma_c = 100\mu$m) shows a low-rate irregular firing, with an oscillatory activity made of spontaneous waves (see also Fig. 2). Within these waves, neurons fire irregularly with a mean ISI CV close to one, while the frequency of these oscillations is only slightly affected by σ_c (Fig. 3C,G). Nevertheless, one can observe that by increasing the connectivity density, the influence of σ_c on the frequency tends to increase. So for highly connected network, one would expect a more significant impact of the connectivity on the network oscillations.

The network state depends on the balance g between excitation and inhibition, and on the frequency of the external noise ν_{ext}. Several regimes can be observed, among those reported in (Brunel, 2000; Mehring et al, 2003) for current-based synapses. We mainly focus on the states displayed in the raster plots of Fig. 3D,H where the network is in a slow Synchronous Irregular regime (SI) because we were interested in low firing rates and irregular activity. In such slow SI regimes, the network could display distinct waves of activity based on underlying topology and delays. Indeed, the spontaneous activity in sensory areas such as V1 is irregular, but it is also known, with Voltage Sensitive Dye Imaging studies (Han et al, 2008; Contreras, 2007; Arieli et al, 1996), that traveling waves appear and propagate. The slow SI regime seemed in that respect to be a good compromise to keep the irregularity and to promote the emergence of waves that mimic what is observed *in vivo*. To have a better insight, Fig. 4 shows distinct spatio-temporal profiles of the correlations in the case $\epsilon \in \{0.5\%, 5\%\}$ for two "extreme" connection spreads (Top : $\sigma_c = 50\mu$m, Bottom : $\sigma_c = 500\mu$m). One can clearly see the oscillations in the temporal domain due to the slow SI regime. Nevertheless, additional simulations shows that even in more asynchronous regimes, where correlations are less influenced by the structure and propagation waves are disrupted, conclusions are still

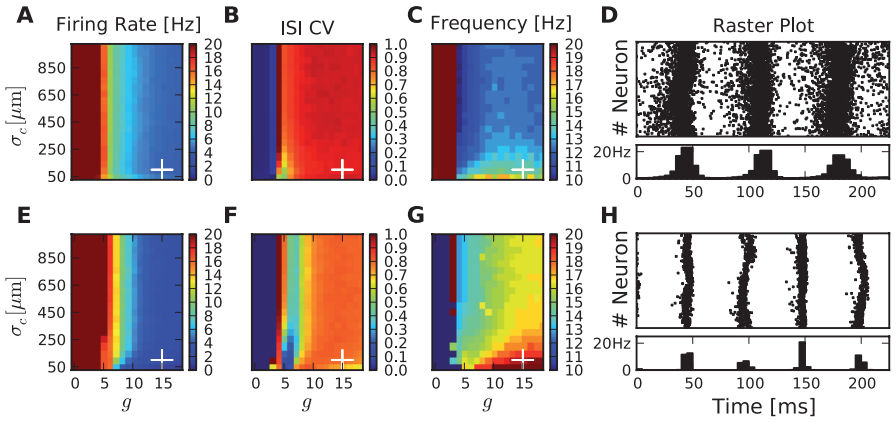

Fig. 3 Phase diagrams of the network statistics as a function of the local connectivity extent σ_c and the excitatory-inhibitory synaptic strength ratio g. **A-C** : Phase diagram of the network with a connection density $\epsilon = 0.5\%$. Mean firing rate (**A**), mean ISI CV (**B**) over the whole network, as a function of σ_c and g. **C**: Mean frequency of the spontaneous oscillation generated in the population dynamic, as can be seen in the raster plot (all excitatory neurons, i.e 10000 cells, **D**) of the activity for a particular regime ($g = 16$ and $\sigma_c = 100\mu$m, white cross), and average firing rate with a 5ms time bin. Bottom: **E-H** : Same as in **A-D**, but with a connection density $\epsilon = 5\%$.

valid (see Supplementary Figure 1). The same applies if we just increase the number of neurons, up to 100,000, without changing the connectivity: the density does not affect the result (see Supplementary Figure 2).

To quantify the distance-dependent correlation profile in the network at the spiking level, we used two measures to distinguish the global amount of synchrony and the decrease as a function of distance. Figure 5A shows a typical profile of the pairwise cross-correlations as a function of distance. For each distance, we selected 2000 pairs of neurons and we averaged the Pearson correlation coefficient computed over all these pairs (see Materials & Methods section). These coefficients were computed between the corresponding spike trains and digitized with a time bin equals to the refractory period of the neurons, i.e 5 ms. Nevertheless, the importance of that bin size has been checked and Supplementary Figure 3 shows that for a larger time bin, the main results are qualitatively similar. The integrated correlation is defined as the integral over distances, and it reflects the global amount of synchrony present in the network. Figure 5B shows that in a log-log scale, the decay of these pairwise correlations as a function of distance is approximatively linear. The slope of this linear region will be referred as the "Correlation Scaling Exponent" obtained by a least square fit. Similar analysis can be performed at the membrane potential level : correlation coefficients at 0 time-lag are used to assess the correlation between two membrane potentials, and the two measures described in Figure 5 can also be applied.

The exhaustive analysis of these correlation profiles in the phase space previously explored in Figure 3 is summarized in Figure 6, for a connection density of $\epsilon = 0.5\%$. Qualitatively similar results can be obtained for higher connection density (see Supplementary Figure 4). The correlations have been analyzed both at the spiking level and at the V_{m} level. In Figure 6, panels A and D show the integrated correlation and

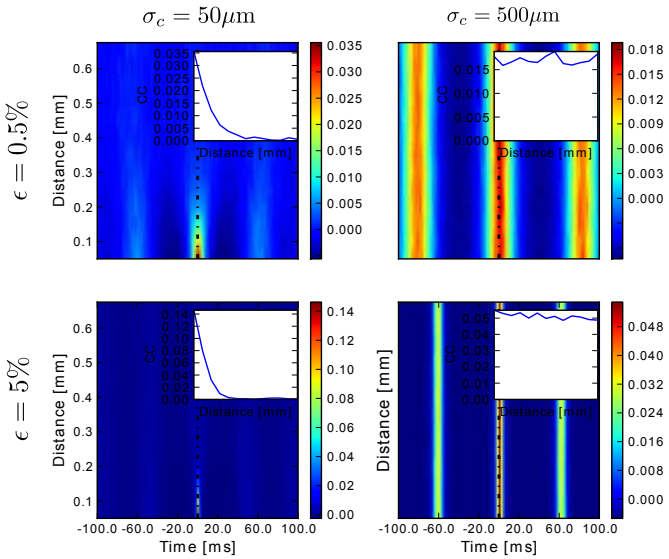

Fig. 4 Spatio-temporal profile of the spiking correlations in the networks, for $g = 16$ and for two values of σ_c: 50μm (Upper panels), and 500μm (Lower panels). Insets shows the spatial decay for the 0-time lag correlations, the curve used to compute the 'integrated correlation' and the 'correlation scaling exponent'.

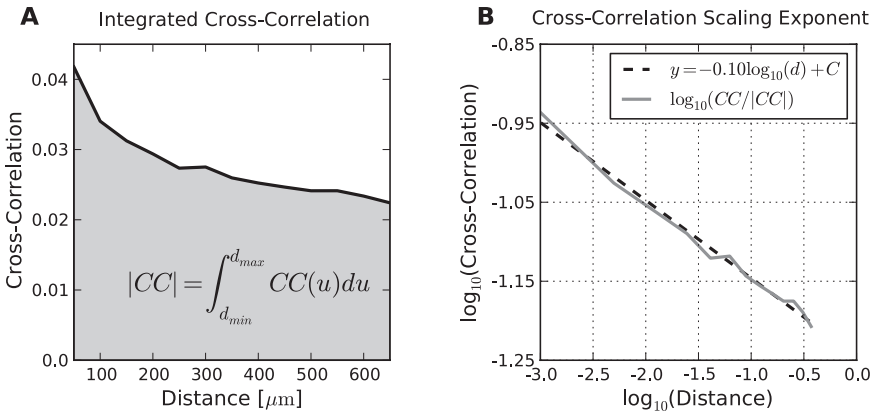

Fig. 5 Quantification of the distance-dependent correlation profile within the topological network. **A**: Typical profile of spiking pairwise cross-correlations as a function of distance. The integrated correlation is the integral over all distances. **B**: The correlation slope is fitted by a line in a log-log scale.

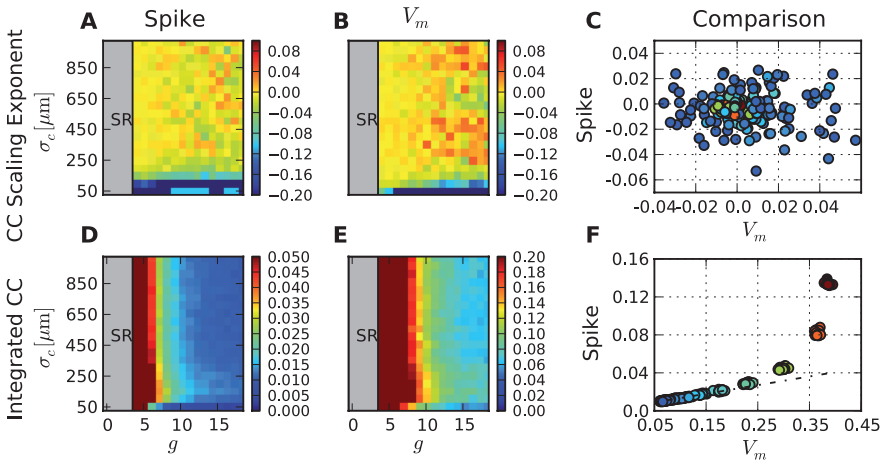

Fig. 6 Comparison of the distance-dependent correlations at the spiking and subthreshold (V_{m}) levels, for $\epsilon = 0.5\%$. **A-B** : Cross-correlation scaling exponent, analyzed either at the spiking level (A) or at the V_{m} level (B). **C** : Cross-correlations scaling exponents at the V_{m} level (observed in panel A) plotted against values at the spiking level (observed in panel B). Same color-code as in Figure 3, illustrating the firing rate of these particular points. **D-E** : Integrated correlations of the distance-dependent correlation profile, analyzed either at the spiking level (D) or at the V_{m} level (E). **F** : Integrated correlations obtained at the V_{m} level (observed in panel D) plotted against the values obtained at the spiking level (observed in panel E). Same color-code.

correlation scaling exponent for correlations measured on spike trains, and panels B and E show these measures on membrane potentials. In both cases, one observes that the correlation scaling exponent does not depend on the connectivity parameters (see Fig. 6A-B). Except in the Synchronous Regular (SR) regime where large oscillations corrupt these measures, the scaling exponent is almost independent of σ_c and of the balance g and tends to 0 on average. Figure 6C shows the values of the cross-correlation scaling exponent as measured in the spike trains against values measured on the basis of V_{m}. Since both values are invariants, a uniform cloud of points is found without any particular statistical bias. In Figure 6D-E, one can see the integrated correlations, again measured on spike trains and on V_{m}. In both cases, the balance g dictates the amount of synchrony which is present in the network in line with our previous results on averaged quantities. The more g is increased, the more dominant the inhibition is and the less synchronous the network activity is. Nevertheless, by plotting the integrated cross-correlations measured at the spiking level compared to those recorded at the V_{m} level (Fig. 6F), interestingly, for low rate regimes, the relation between the integrated spiking correlation and the sub-threshold correlations is almost linear and then increases in a nonlinear and monotonic manner for higher firing rates. The clusters are isolated according to the network firing rate shown in color code identical to the one used in Fig. 3. Therefore, in these networks configurations, the integrated correlation measured at the subthreshold level is uniquely determined by the spiking correlation.

To study the influence of the heterogeneity in the connection scheme, the ratio between the spread of the Gaussian profile used to connect the excitatory and the inhibitory neurons within the network was varied. Figure 7A illustrates how these two

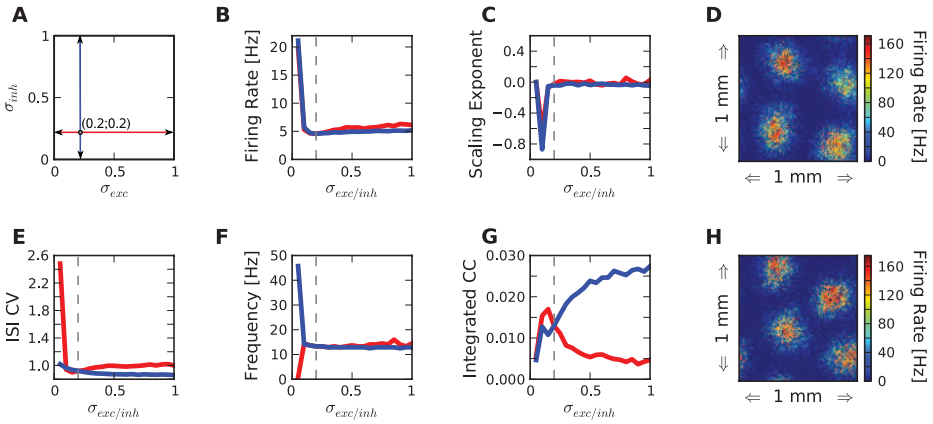

Fig. 7 Changing the spatial spread σ_c of the excitatory and inhibitory connections independently. **A** : Schematic illustration explaining the explored parametric region in other panels. In red, σ_{inh} is held constant while σ_{exc} is varied, while in blue, it is the opposite. In all subsequent panels (B,C,E-G), the intersection point is represented by the dashed gray line. **B, C, E, F, G** : Mean firing rate, cross-correlation scaling exponent, population activity peak frequency, mean ISI CV and integrated correlations as a function of σ_{exc} or σ_{inh}. **D, H** : Time-averaged activity maps of the two pathological cases that emerged for low σ_{exc} value (D) or σ_{inh} value (H).

parameters σ_{exc} and σ_{inh} were changed. Instead of exploring the whole parameter space, only two lines were explored: one with σ_{exc} fixed to 200 μm while σ_{inh} was varied in the range [0 to 1 mm], and another one where σ_{exc} was varied for a fixed σ_{inh} (respectively red and blue curves in Figure 7). As one can see in the Figure 7B-E-F, averaged quantities such as mean firing rate, population peak oscillatory frequency, or mean ISI CV are hardly impacted by these parameters, as soon as σ_{exc} or σ_{inh} are not too small. If they are, one can observe a symmetry breaking pushing the network into a pathological states (Fig. 7D-H) with very localized bumps of activity spontaneously jumping from a symmetric state to another. This scenario is reminiscent of the "hotspots"-like patterns obtained by (Usher et al, 1994) with Mexican-hat connectivity profile and hybrid neurons. Our finding is thus compatible with their model when the network connectivity is local enough to generate strong topological correlations. Otherwise, one can again notice that the cross-correlation scaling exponent is rather insensitive to the connectivity spatial spread and close to 0 (Fig. 7C), while the integrated cross-correlation is affected by the spread (Fig. 7G). Increasing the spread of inhibitory projections while keeping that of excitatory neurons constant increases the overall amount of synchrony within the network, by diluting the inhibition. On the contrary, increasing the spread of the excitation by keeping fixed that of inhibition decreases the amount of synchrony by diluting excitation.

Finally, we explored the role of propagation delays and their influence on the spatial spread of distance-dependent cross-correlations within the network. In particular, we studied the impact of interaction velocity for different network structure parameters: the connection spatial spread and the connection density. For highly localized connectivity ($\sigma_c = 50\mu$m , Fig. 8A,C) changing the velocity has no significant effect on the correlation profile, for both low and high connectivity densities. Indeed, when only nearby neurons are connected, the effective delay for various velocities are still

small enough to leave the profiles unaffected. Similarly, when a larger Gaussian profile is used in a network with high connectivity density ($\epsilon = 5\%$, $\sigma_c = 200\mu$m, see Fig. 8D), we found that changing the velocity induces a higher variability in the correlation profiles, without any consistent variation. However, when considering large connectivity spread in very diluted networks, where the propagation delays become crucial ($\epsilon = 0.5\%$, $\sigma_c = 200\mu$m, see Fig. 8C), the integrated correlation increases with velocity. It has to be stressed that the case $\sigma_c = 50\mu$m in Fig. 8A,C is a limit case, because the network is then close to a pathological state where the neuron are almost in a all-to-all nearest-neighbor connection scheme. The fact that the bottom rows of Fig. 6. A,B,C,D are different is consistent with the observation that in Fig. 8 there is a difference in the integrated correlations for the two values of σ_c used. In fact, one has to keep in mind that the shapes of the curves in Fig 8A,C are valid only for very local network ($\sigma_c = 50\mu$m), and that in all the other configurations, curves would look like those in Fig 8B,D). Altogether, we conclude that propagation delays have a significant effect on the spatial correlation profile only when long-range interactions are as important as local interactions. The linear relationship between delays and distances used in the model can be considered as too strong. Indeed, for very dense and intricate circuit, as the one studied in (Oswald and Reyes, 2008), this relationship is not that obvious. Indeed, if there are evidences that conduction times within dendrites and/or axons are linear, this linearity due to the wiring scheme may be more noisy. Nevertheless, we checked that the invariance and the results are still valid when delays are only correlated with distances (see Supplementary Figure 5), such as $d_{ij} = d_{\text{syn}} + \frac{l_{ij}}{v}(1 + \mathcal{N}(0, \sigma))$, where $\mathcal{N}(0, \sigma)$ is a Gaussian noise of variance $\sigma = 0.25$.

5 Effect of Structured Stimulation

Having studied the response of the network under unstructured stimulation (Poissonian input, mimicking the spontaneous ongoing activity coming from the thalamus), we were interested in adding an additional layer to inject spatial correlations (representing the sensory drive) in the network. More precisely, we simulated a layer of Poisson sources arranged also in a 2D plane (Fig. 9A), connected to the recurrent network with a divergent Gaussian spread with a variance σ^2_{ext}, and acting as compound processes made with Poisson sources and Multiple Interaction Processes (Kuhn et al, 2003). The combination of these two processes allowed us to control, with a continuous parameter $c \in [0, 1]$ (Fig. 9B), the amount of synchrony send to the network while keeping the input mean output firing rate ν_{ext} constant. To be more precise, c is the percentage of neurons that will emit simultaneous spikes during some volleys, appearing at a frequency $c\nu_{\text{ext}}$. Each external cells acts therefore as an independent Poissonian source, at a frequency $(1-c)\nu_{\text{ext}}$, and shares global inputs imposed by the MIP at a frequency $c\nu_{\text{ext}}$.

Figure 10 shows the response of the network for four external input synchrony levels $c \in \{0, 0.05, 0.1, 0.2\}$, and for different spreads of the external input divergence σ_{ext}. In Figure 10A, four typical raster plots are represented for these four levels of synchrony with an identical external input rate fixed at 5 spikes/s (as in the unstructured case). The more c is increased, the more efficiently the synchrony will trigger strong responses in the recurrent network, while, at the same time, decreasing the general firing rate by forcing every neuron in their refractory period (Fig. 10 B). For averaged quantities such as the mean firing rate and ISI CV (Fig. 10B,E), one can notice that

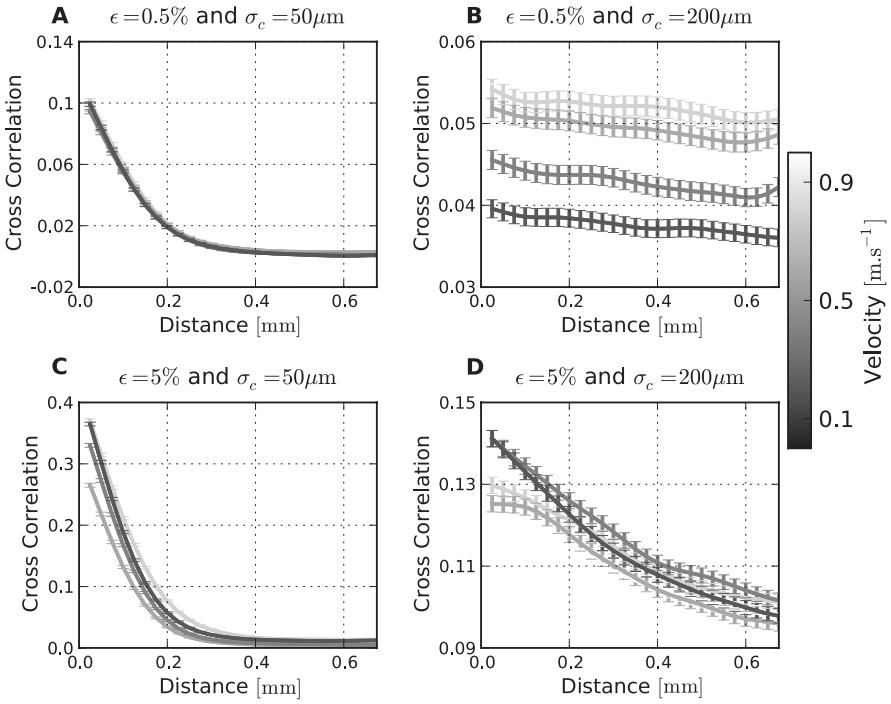

Fig. 8 Influence of the delays on the cross-correlation profiles. The average pairwise Pearson correlation coefficient is plotted in a network with a connectivity density $\epsilon = 0.5\%$, as a function of the distance within neuronal pairs, for four distinct velocity values. Error bars show the standard error of the mean. Two Gaussian profiles are considered. A: a very local one ($\sigma_c = 50\mu$m). B: a broader one ($\sigma_c = 200\mu$m). **C-D** : Same as **A-B**, but in a denser network with $\epsilon = 5\%$.

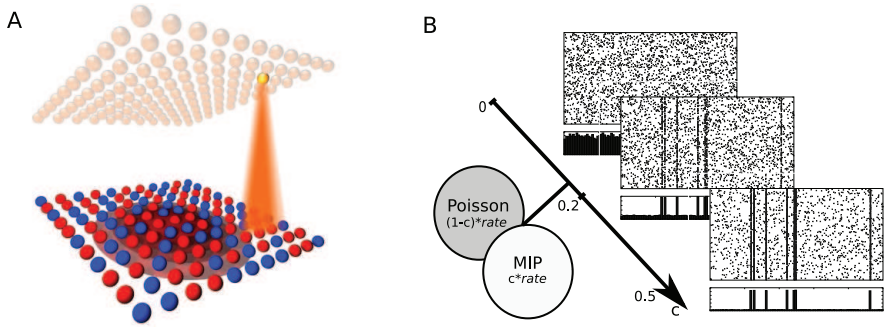

Fig. 9 Schematic view of the input layer used to inject spatial correlations. **A** : [Taken from (El Boustani et al, 2009)] An additional 2D layer of sources is added, where each source connects the recurrent network with a Gaussian profile of standard deviation σ_{ext} in a divergent manner. **B** : Illustration of the compound process made with a Multiple Interaction and Poisson processes, shown as raster plots of the activity for 2500 cells in the external layer, for several values of c. The parameter c controls the percentage of co-active neurons into synchronous volleys.

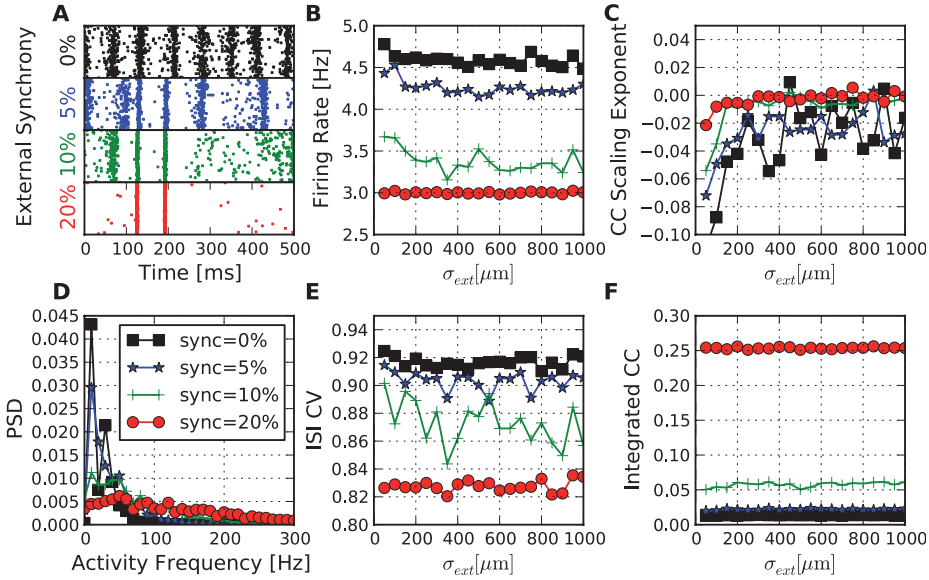

Fig. 10 Evoked activity under spatially correlated stimulations for $\epsilon = 0.5\%$. **A**: Three raster plots for 2500 neurons and different synchrony levels in the input. **B, C, E, F**: Mean firing rate, cross-correlation scaling exponent, mean ISI CV and integrated correlation in the recurrent network as a function of the input divergence σ_{ext}, for the four levels of synchrony. **D** : Power spectra density for a fixed $\sigma_{ext} = 200\mu$m, and the four levels of synchrony.

the external divergence σ_{ext} of the feedforward projection does not have any influence, a phenomena already observed in the presence of unstructured inputs. Increasing the external synchrony will increase the frequency content of the oscillatory activity in the population dynamics because of the stochastic nature of the input (Fig. 10D). Regarding the correlation profile as a function of distance, increasing the synchrony c induces an increase in the integrated correlation of the recurrent network (Fig. 10F), but more importantly, we observed that the external correlation is now able to change the correlation scaling exponent of the distance-dependent correlation profile, especially when the external connectivity spread is narrow (Fig. 10C).

These changes are summarized in Figure 11, where we chose $\sigma_{ext} = 50\mu$m for the sake of clarity. In Figure 11A, one can observe the increase in the integrated correlations following an increase of the external synchrony. The more the network receives strong synchronous spiking volleys, the stronger the pairwise correlation coefficients are. As one can see in Figure 11B, increasing the external synchrony also affects the cross-correlation scaling exponent of the distance-dependent correlation profile. Adding synchrony favors synchronous volleys that are strong enough to trigger spiking activity in the recurrent network at any position simultaneously. This global activation pattern creates in turn long-range correlations in the network. Conversely, for narrow stimulation spatial spread without any additional spatial correlations (no synchrony), small assemblies of neurons are becoming strongly correlated in a spatial range shorter than the one observed in the spontaneous activity.

The genesis of these processes can be better understood by studying the presynaptic pattern leading to spike, called Network Activity Spike-Triggered Average (NASTA).

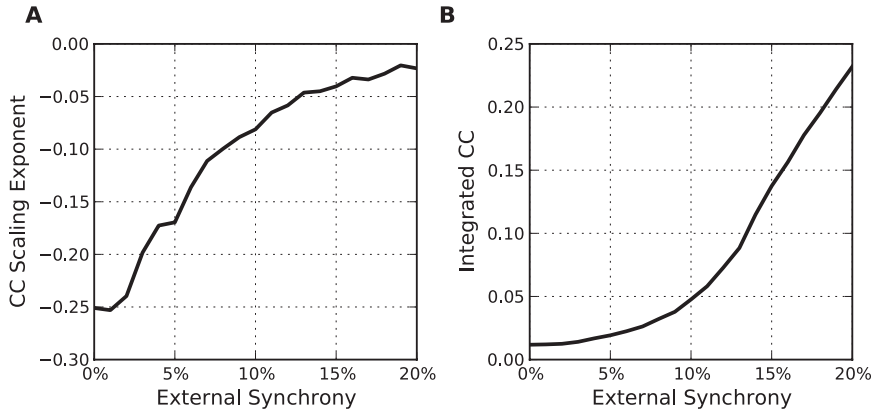

Fig. 11 Change in the distance-dependent correlation profile as a function of the level of synchrony in the input for $\epsilon = 0.5\%$. **A** : Correlation scaling exponent in the recurrent network as a function of the external synchrony. **B** : Integrated correlation in the stimulated layer as a function of the external synchrony.

.

This is shown in Figure 12A, where the recurrent inputs are mainly responsible for the emission of a spike for low levels of external synchrony. The fact that a localized PSTH is obtained for $c < 10\%$ means that the network activity of the surrounding neurons 2 ms earlier will be driving the post-synaptic neuron to spike. The more c is increased, the less important is the recurrent connectivity contribution and the stronger are the external spiking volleys to trigger by themselves a spike. In the NASTA spatio-temporal domain (Figure 12B), one can see that the broad temporal spread around the postsynaptic neuron account for the propagation delay with the presynaptic neurons in the network.Here again, the external spike volleys have a role in flattening the spatial correlation and sharpening the temporal spread of cross-correlation when increasing the external synchrony.

6 Discussion

Invariant Macroscopic Statistics In this paper, we have studied balanced network models with conductance-based synaptic interactions and different spatial profiles of connectivity. Our main finding is that such balanced networks own connectivity-invariant quantities as long as each neuron is sparsely connected to its neighbors. Surprisingly, this result holds even for very narrow Gaussian connectivity. This is a very encouraging result which shows that mean-field approaches, where no topology is taken into account, can offer a reliable description of the network at a macroscopic scales as long as the dynamic regime of conductances remains stationary and homogenous across the network.

For instance, when modeling data such as Voltage-Sensitive Dyes Imaging (VSDI) recordings, there is *a priori* no need to know the fine details of the connectivity within each pixel. In particular, we know that for the spontaneous activity regime, the mean firing rate, the mean ISI CV and the overall synchrony depend mainly on the synaptic E/I balance ratio (except for the first-neighbor extreme case). We also observed that

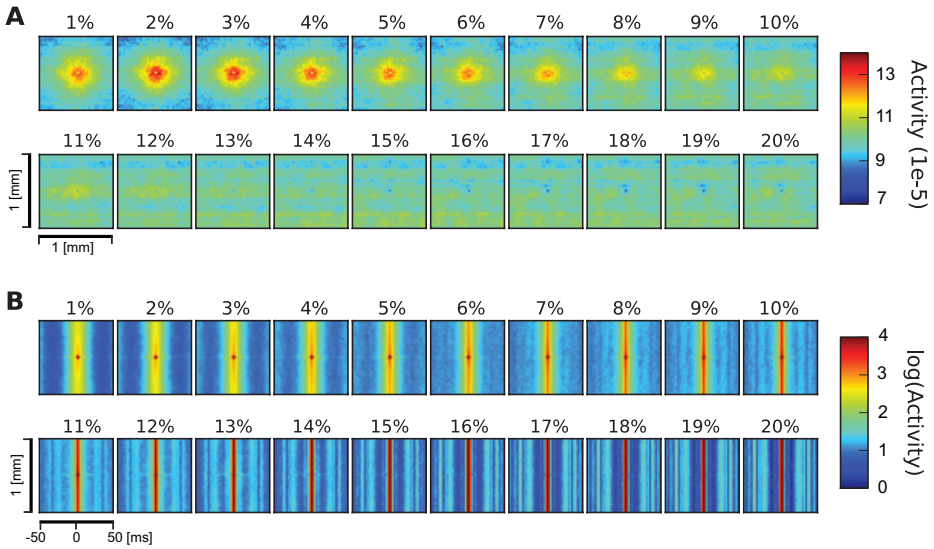

Fig. 12 Network activity spike-triggered average (NASTA) as a function of external synchrony for $\epsilon = 0.5\%$. **A** : NASTA for one neuron. The maps show the spatial network activity averaged over all its spikes, 2 ms before, for several level of external synchrony. **B** : Same as in **A**, but the activity is now plotted only for the central spatial line of the map and as a function of time.

the spatial correlation decay in this regime does not depend either on the connectivity extent nor on the conductance balance in the irregular regime. However, we do see a dependency of the correlation decay on the connectivity density (which is anyway a structural averaged quantity).

Similar conclusions apply to the case of networks subject to structured input (representing external correlation patterns imposed by the sensory drive). We showed that the fine details of the connectivity are less influential than averaged quantities such as the overall synchrony. It should be thus possible to find a simple relation describing the spatial correlation decay only in terms of macroscopic quantities such as the mean synaptic input per neuron or the mean synchrony in the external drive.

Thus, we observed numerically that most of the first- and second-order statistics in these networks are ruled by averaged macroscopic structural quantities. Therefore, it seems that for these models, structural (synaptic weights, mean synaptic input) and dynamical (mean firing rate, correlations) statistics are related to each other in a hierarchic manner as already observed in a simpler setting (Liu and Nykamp, 2009). We therefore do not conclude that connectivity is completely "decoupled" from correlation, but rather that this detail of description is irrelevant at a large-scale level of observation. The underlying mechanistic explanation is directly linked to the way these balanced activities are generated in sparse networks, and we show how these dynamics break in the extreme case of dense local connectivity.

Supra- and Sub-threshold Correlations Several studies have focused recently on the second-order transfer function of spiking neurons. More precisely, knowing the correlation structure in the presynaptic activity of one or two neurons sharing an input, one can ask which are the spiking auto-correlation of each neuron and/or their cross-

correlation. It has been shown recently that for the low-correlation regime, the supra- and sub-threshold activities are linearly related with a proportionality factor which mainly depends on the firing rate of these neurons (de la Rocha et al, 2007; Shea-Brown et al, 2008). In our situation, conclusion are harder to reach from an analytical point of view. However, we found a monotonic relation between the supra- and sub-threshold signal correlation proving that both levels offer a similar description of the correlation state in the network. For spontaneous activity (uncorrelated inputs), these correlations are generated by the recurrent connections within the network, so that this monotonic relation must satisfy a self-consistent relation in order to be stable. Describing this relation through close analytical equations for simpler models would bring us a step further in the understanding of recurrent network dynamics.

Comparison with other studies Previous studies intended to describe the relationship between network structure and dynamics. For example in (Usher et al, 1994), the authors studied networks displaying irregular spiking activity and long-range temporal correlations. They reported that in a small network with "hybrid" neurons violating Dale principle (Kriener et al, 2008), and connected in a Mexican hat manner (local excitation and long-range inhibition), they were able to generate network dynamics with a high degree of variability. However, their network regime was not balanced, and the irregular activity in the network was generated by "hotspot"-type activity patterns (similar to those described in Figure 7D-H), instead of waves. In (Kitano and Fukai, 2007), the authors report that in small and highly clustered small-world networks, the spiking irregularity is strongly dependent on the synaptic weight balance and the small-world rewiring parameter. However, this approach is far from the sparse connectivity we are considering here, and it does not match the mean-field requirements for a suitable prediction of macroscopic dynamical quantities. Nevertheless, in their computation of distance-dependent correlation and beyond the local region, there is a large region where the correlation scaling exponent does not depend as strongly on the connectivity scheme (see Fig. 5 of (Kitano and Fukai, 2007)). Even if these results are hardly comparable to the present ones, due to huge differences in the neurons models, network sizes and wiring scheme, several conclusions regarding the invariance of mean quantities hold in both studies.

Shaping the Correlation Landscape with Correlated Input When the network is fed with uncorrelated input, it was not possible to change the distance-dependent correlation scaling exponent by changing the connectivity extent or the synaptic weight ratio. However, changing the connectivity density or the correlations in the input either broadens or shortens this distance-dependent correlations. This could be also easily seen through the network activity spike-triggered average where the correlation landscape becomes uniform in space and narrow in time by increasing synchrony. Taken together with the fact that spiking and sub-threshold correlations have a monotonic relationship, this could be a way to estimate the impact of sensory input on ongoing activity using measurements such as VSD imaging, single-cell unit activity, or local field potentials.

Acknowledgements We thank Olivier Marre for helpful discussions and critical comments on the manuscript. P.Y. was supported by a MENRT bursary from the University of Paris XI, S.B was supported by the FRM. This work was supported by the CNRS, ANR (NATSTATS and HR-CORTEX) and the European Commission (FACETS FP6-2004-IST-FETPI 15879).

References

Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb Cortex 7(3):237–252

Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science 273(5283):1868–1871

Benucci A, Frazor RA, Carandini M (2007) Standing waves and traveling waves distinguish two circuits in visual cortex. Neuron 55(1):103–117

Berger D, Warren D, Normann R, Arieli A, Grün S (2007) Spatially organized spike correlation in cat visual cortex. Neurocomputing 70(10-12):2112–2116

Bienenstock E (1996) On the dimensionality of cortical graphs. J Physiol Paris 90(3-4):251–256

Braitenberg V, Schüz A (1998) Cortex: Statistics and Geometry of Neuronal Connectivity. Springer

Bringuier V, Chavane F, Glaeser L, Frégnac Y (1999) Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. Science 283(5402):695–699

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J Comput Neurosci 8(3):183–208

Cessac B, Viville T (2008) On dynamics of integrate-and-fire neural networks with conductance based synapses. Front Comput Neurosci 2:2, DOI 10.3389/neuro.10.002.2008

Contreras D (2007) Propagating waves in visual cortex. Neuron 55(1):3–5

Destexhe A, Paré D (1999) Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. J Neurophysiol 81(4):1531–1547

Diesmann M, Gewaltig M (2001) NEST: An environment for neural systems simulations. Forschung und wisschenschaftliches Rechnen, Beitrage zum Heinz-Biling-Preis 58:43–70

El Boustani S, Destexhe A (2009) A master equation formalism for macroscopic modeling of asynchronous irregular activity states. Neural Comput 21(1):46–100

El Boustani S, Marre O, Béhuret S, Baudot P, Yger P, Bal T, Destexhe A, Frégnac Y (2009) Network-state modulation of power-law frequency-scaling in visual cortical neurons. PLoS Comput Biol 5(9):e1000,519

Gil Z, Connors BW, Amitai Y (1999) Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability. Neuron 23(2):385–397

Gilbert CD, Wiesel TN (1983) Clustered intrinsic connections in cat visual cortex. J Neurosci 3(5):1116–1133

Göbel W, Kampa BM, Helmchen F (2007) Imaging cellular network dynamics in three dimensions using fast 3d laser scanning. Nat Methods 4(1):73–79

Gonzlez-Burgos G, Barrionuevo G, Lewis DA (2000) Horizontal synaptic connections in monkey prefrontal cortex: an in vitro electrophysiological study. Cereb Cortex 10(1):82–92

Greenberg DS, Houweling AR, Kerr JND (2008) Population imaging of ongoing neuronal activity in the visual cortex of awake rats. Nat Neurosci 11(7):749–751

Grinvald A, Lieke EE, Frostig RD, Hildesheim R (1994) Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. J Neurosci 14(5 Pt 1):2545–2568

Han F, Caporale N, Dan Y (2008) Reverberation of recent visual experience in spontaneous cortical waves. Neuron 60(2):321–327

Hellwig B (2000) A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. Biol Cybern 82(2):111–121

Izhikevich EM, Gally JA, Edelman GM (2004) Spike-timing dynamics of neuronal groups. Cereb Cortex 14(8):933–944

Kitano K, Fukai T (2007) Variability v.s. synchronicity of neuronal activity in local cortical network models with different wiring topologies. J Comput Neurosci 23(2):237–250

Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. J Neurosci 25(14):3661–3673

Kriener B, Tetzlaff T, Aertsen A, Diesmann M, Rotter S (2008) Correlations and population dynamics in cortical networks. Neural Comput 20(9):2185–2226

Kriener B, Helias M, Aertsen A, Rotter S (2009) Correlations in spiking neuronal networks with distance dependent connections. J Comput Neurosci

Kuhn A, Aertsen A, Rotter S (2003) Higher-order statistics of input ensembles and the response of simple model neurons. Neural Comput 15(1):67–101

Kumar A, Schrader S, Aertsen A, Rotter S (2008) The high-conductance state of cortical networks. Neural Comput 20(1):1–43

Larkum ME, Zhu JJ, Sakmann B (2001) Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. J Physiol 533(Pt 2):447–466

Liu CY, Nykamp DQ (2009) A kinetic theory approach to capturing interneuronal correlation: The feedforward case. Journal of Computational Neuroscience 26:339–368

Marre O, Yger P, Davison A, Frégnac Y (2009) Reliable recall of spontaneous activity patterns in cortical networks. Journal of Neuroscience 29(46):14,596–14,606

Mehring C, Hehl U, Kubo M, Diesmann M, Aertsen A (2003) Activity dynamics and propagation of synchronous spiking in locally connected random networks. Biol Cybern 88(5):395–408

Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. Neuron 63(6):879–888

Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. Nat Neurosci 12(1):70–76

Nirenberg S, Latham PE (2003) Decoding neuronal spike trains: how important are correlations? Proc Natl Acad Sci U S A 100(12):7348–7353

Oswald AMM, Reyes AD (2008) Maturation of intrinsic and synaptic properties of layer 2/3 pyramidal neurons in mouse auditory cortex. J Neurophysiol 99(6):2998–3008

Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) The asynchronous state in cortical circuits. Science 327(5965):587–590

de la Rocha J, Doiron B, Shea-Brown E, Josi K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. Nature 448(7155):802–806

Roxin A, Brunel N, Hansel D (2005) Role of delays in shaping spatiotemporal dynamics of neuronal activity in large networks. Phys Rev Lett 94(23):238,103

Schwarz C, Bolz J (1991) Functional specificity of a long-range horizontal connection in cat visual cortex: a cross-correlation study. J Neurosci 11(10):2995–3007

Shea-Brown E, Josi K, de la Rocha J, Doiron B (2008) Correlation and synchrony transfer in integrate-and-fire neurons: basic properties and consequences for coding. Phys Rev Lett 100(10):108,102

Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. Annu Rev Neurosci 18:555–586

Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci 28(48):12,591–12,603

Usher M, Stemmler M, Koch C, Olami Z (1994) Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. Neural Computation 6(5):795–836

Vogels TP, Abbott LF (2005) Signal propagation and logic gating in networks of integrate-and-fire neurons. J Neurosci 25(46):10,786–10,795

Voges N, Aertsen A, Rotter S (2007) Statistical analysis of spatially embedded networks: From grid to random node positions. Neurocomputing 70(10-12):1833–1837

van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274(5293):1724–1726

van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. Neural Comput 10(6):1321–1371

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370(6485):140–143

11.2 Supplementary figures

Journal of Computational Neuroscience manuscript No.
(will be inserted by the editor)

Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons: Supplementary Figures

Pierre Yger* · **Sami El Boustani*** · **Alain Destexhe** · **Yves Frégnac**

Received: date / Accepted: date

[1] co first-authors
Unité de Neuroscience Intégratives et Computationnelles (UNIC), UPR 2191 CNRS
Bat. 32-33, 1 Avenue de la Terrasse, 91198 Gif-sur-Yvette, France
Tel: 33-1-69-82-34-15, Fax: 33-1-69-82-34-27
Corresponding author: Alain Destexhe, E-mail: destexhe@unic.cnrs-gif.fr

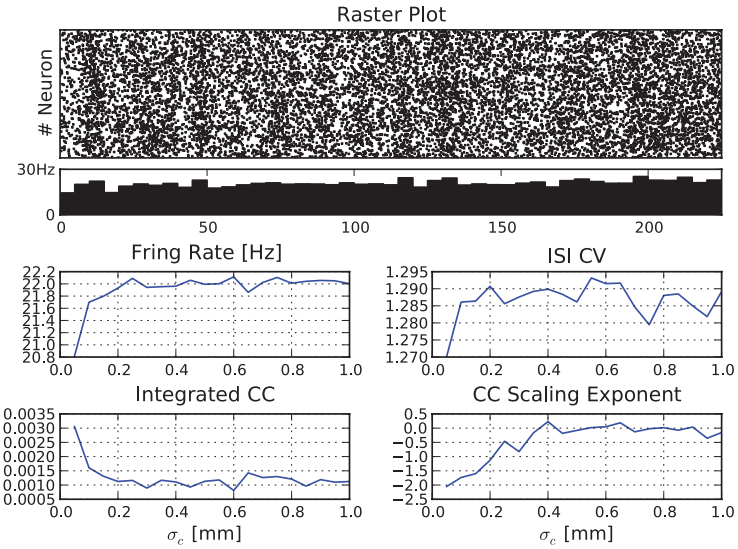

Supplementary Figure 1 Influence of the network regime. A diluted network ($\epsilon = 0.5\%$) with $g = 16$ is set into an AI state (see raster plot for 2500 excitatory cells). Subpanels shows mean firing rate (with a time bin of 10 ms), ISI CV, and the two parameters computed from the correlation profile as functions of σ_c.

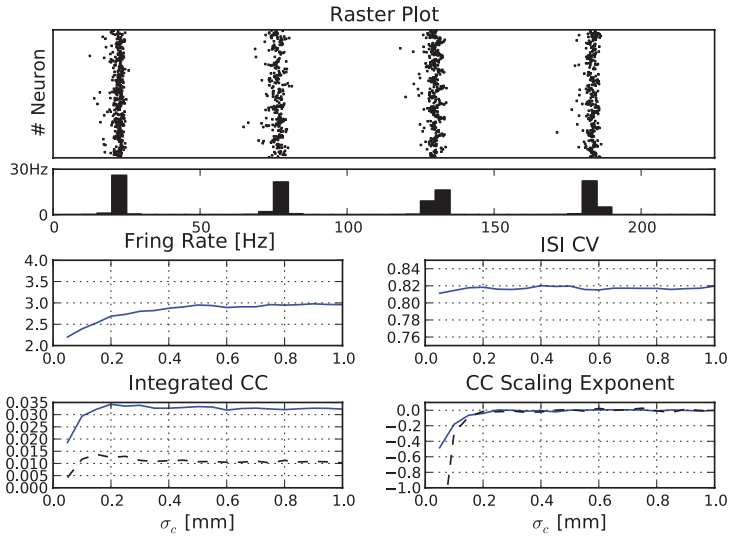

Supplementary Figure 2 Results obtained in a dense network with 100000 neurons, $\epsilon = 5\%$, $g = 16$, $\nu_{\text{ext}} = 15$ spike.s^{-1}. Raster plots of the activity for 2500 excitatory cells, with $\sigma_c = 0.1$mm. Subpanels shows mean firing rate, ISI CV, and the parameters computed from the correlation profile as a function of distance, for several σ_c. Dashed dot lines shows the curves in the network with 12500 cells.

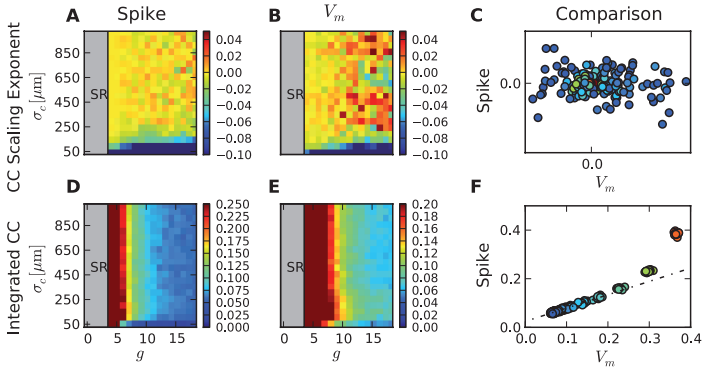

Supplementary Figure 3 Comparison of the distance-dependent correlations at the spiking and subthreshold (V_{m}) levels, for $\epsilon = 5\%$ and a larger bin to gather correlations ($T = 20$ms). **A-B** : Cross-correlation scaling exponent, analyzed either at the spiking level (A) or at the V_{m} level (B). **C** : Cross-correlations scaling exponents at the V_{m} level plotted against values at the spiking level. Same color-code as in Figure 2, illustrating the firing rate of these particular points. **D-E** : Integrated correlations of the distance-dependent correlation profile, analyzed either at the spiking level (D) or at the V_{m} level (E). **F** : Integrated correlations obtained at the V_{m} level plotted against the values obtained at the spiking level. Same color-code.

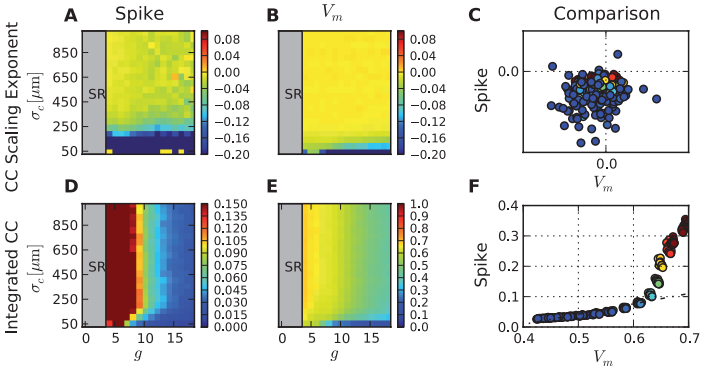

Supplementary Figure 4 Comparison of the distance-dependent correlations at the spiking and subthreshold (V_{m}) levels, for $\epsilon = 5\%$. **A-B** : Cross-correlation scaling exponent, analyzed either at the spiking level (A) or at the V_{m} level (B). **C** : Cross-correlations scaling exponents at the V_{m} level plotted against values at the spiking level. Same color-code as in Figure 2, illustrating the firing rate of these particular points. **D-E** : Integrated correlations of the distance-dependent correlation profile, analyzed either at the spiking level (D) or at the V_{m} level (E). **F** : Integrated correlations obtained at the V_{m} level plotted against the values obtained at the spiking level. Same color-code.

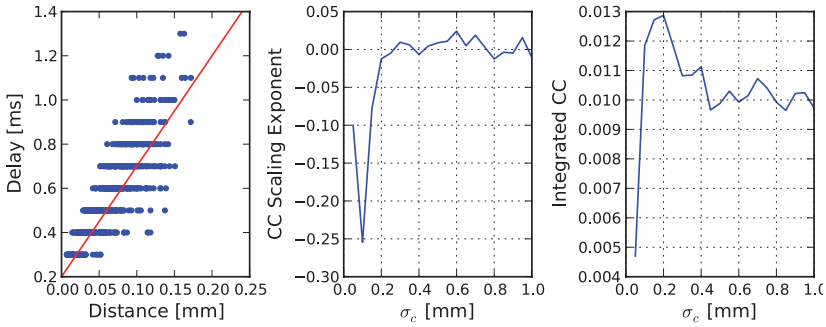

Supplementary Figure 5 Correlation Profile in a network with delays correlated with distances. Left: Distribution of the delays vs. distances for $\sigma_c = 100\mu$m. Red line shows the linear relationship between delays and distances $d_{ij} = d_{\mathrm{syn}} + \frac{l_{ij}}{v}$. Middle and Right: Variation of the Integrated CC and the CC Scaling exponent in a sparse network $\epsilon = 0.5\%$ with $g = 16$, as functions of σ_c.

12 Discussion

12.1 General comments

Comparison with previous results in literature It should be noted that several studies did not find the invariants we found here in topological networks. Of these studies, Usher et al. (1994) is hardly comparable with our model for several reasons. First, in their study they were looking for networks displaying irregular spiking activity and long-range temporal correlations. To do so, they did not used the balanced network paradigm but instead a topological network with very specific (and sometimes unrealistic) features. Neurons in their network violated Dale's principle by which a neuron should be either excitatory or inhibitory. These "dual-nature" neurons are known to be appropriate for generating AI regimes in sparsely connected networks (Kriener et al., 2008) although there existence has never been reported in neocortex (one example exists in Hermissenda). These neurons are also connect to their neighbours through a Mexican hat connectivity profile, which is also different from the connectivity schema we chose to study based on experimental and review studies (Hellwig, 2000, Bienenstock, 1996). They could succeed in generating network dynamics with a high degree of variability which was mainly produced by the "hotspot"-type activity patterns. Recent VSD imaging experiments that offer observation of superficial cortical dynamics with high temporal resolution support the idea of travelling waves such as those we observe (Han et al., 2008, Contreras, 2007, Arieli et al., 1996) whereas attractors of "hotspots" have rarely been reported in ongoing activity. However these authors report only two connectivity configurations in the paper, namely the Mexican hat and the random connectivity that displays oscillatory synchronous activity because the network state is unbalanced. However, electrophysiological experiments tend to support the idea of balanced conductance inputs (Shu et al., 2003, Haider et al., 2006, Hasenstaub et al., 2005, Monier et al., 2008) with local connectivity patterns that can be roughly approximated by a Gaussian profile, as in our study.

Another type of connectivity pattern that has been studied is the small-world network (SWN) type. In the work of Kitano and Fukai (2007), they chose to focus on networks of biophysical neurons following this connectivity pattern and in particular for a very small SWN p parameter (highly clustered). In these situations, very few synapses are randomly redistributed (on average 3.1 for most of their figures) leaving a network with essentially local and extremely dense connections. Even if this situation is far from the one we are considering (sparse connectivity), their results are comparable to ours.

Ongoing activity In contrast to the Frozen Paradigm network, we do not use here a self-sustained network with reverberating activity, and an external noise was injected into all neurons to settle the network into the Synchronous Irregular regime. Indeed, self sustained states in random networks are a minor fraction of all the possible states, and their existence is subject to several parameters. The network must be very diluted, in order to prevent too strong synchronous inputs (see Kumar et al. (2008b), Marre et al. (2009b)), and the conjunction of the propagation delay and the refractoriness must allow a certain amount of activity to be propagated to the network. In the Frozen Paradigm network, since we wanted to have a fully deterministic and chaotic dynamical system, we used a network set into such a self-sustained state. A brief transient pulse of activity is sent to the network, then activity is reverberated and maintained for a long transient time. Axonal velocities were infinite, with instantaneous synaptic transmission, and therefore ongoing activity was easy to sustain. Here, nevertheless,

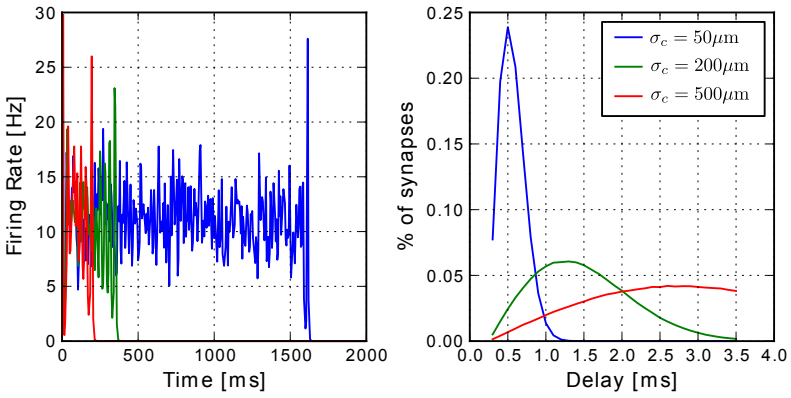

Figure 18: Left: Firing rate in a random network of 10000 neurons, without any structure, stimulated with an initial bump of activity for 100 ms, for several distributions of delays. Right: Distributions of the delays

.

obtaining such sustained regimes was more difficult, and to avoid extinctions of the network, subject to correlations and to a rather dense connectivity, we decided to add external noise to the system, to mimic the afferent inputs from the surroundings that may be received by the topographical area. The amount of external noise added is rather low, and as found in Brunel (2000), it can shift the activity from a synchronous irregular regime (the most common here) to an asynchronous irregular one, mainly noise driven (see Supplementary Figure 1).

It has to be noted that the propagation delay is the main parameter preventing sustained activity. To show this, we performed additional simulations in a self-sustained network where the distribution of the delays is similar, but with a completely randomized network structure. Figure 18 shows that increasing the spread of the delays (and therefore the mean delay in the network) reduces the lifetime of the self-sustained activity (the network is stimulated with an initial bump of activity during the first 100 ms). In such a 10000 cell network, activity is self-sustained for $d_{\text{syn}} = 0.2$ms (d_{syn} is the synaptic transmission time), but as soon as one adds a non-zero velocity (see previous article), this activity vanishes. We do not argue that the structure does not play a role, but delays and especially the mean delay seem to be here the crucial parameters.

External Stimulation As in the PLoS paper, presented in Part II, we used a multiple interaction process (MIP) (Kuhn et al., 2003) to stimulate the network and inject spatial correlations in the network. As we saw in the Frozen Paradigm, choosing an appropriate stimulus to stimulate neuronal networks is not straightforward. The MIP is an easy way to generate instantaneous correlations and to avoid strong volleys of synchrony, but as the output firing rate of the MIP process depends on the correlation coefficient, we compensate with a Poisson process to keep the total external input fixed (see Equation 12).

$$F = \text{MIP}(cF_{\text{ext}}) + \text{POISSON}((1-c)F_{\text{ext}}) \tag{12}$$

Only low coefficients of synchrony $\in [0, 0.2]$ were used. In order to get a better insight and compare our results with those of Kuhn and colleagues (see Fig 1C of Kuhn et al. (2003)), we have displayed in Figure 19 a raster plot of the activity of 500 neurons, with $c = 0.1$, for

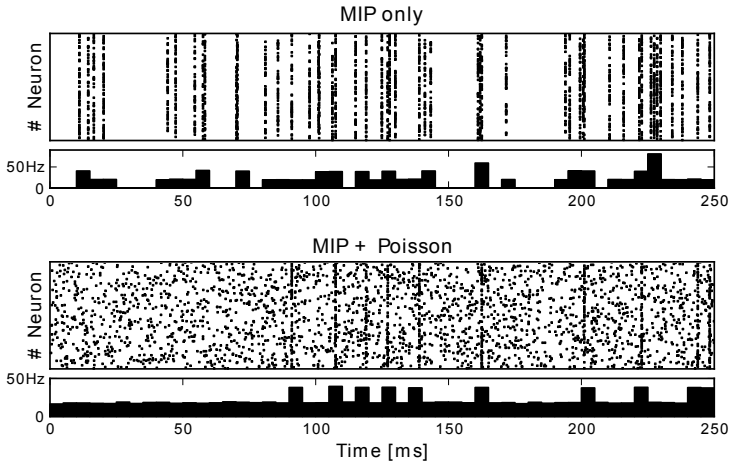

Figure 19: Illustration of the stimulation used in the 'thalamic layer', and comparison with a pure MIP stimulation. Up: raster plots of 500 cells and the firing rate during 250 ms, generated with a pure multiple interaction process (MIP) with rate 200 Hz and $c = 0.1$ (leading to spikes trains at a frequency $cF = 20$ Hz). Down: mixture of a MIP process at 20Hz and a Poisson source (see Equation 12), again with $c = 0.1$

.

both protocols: MIP only (with a mother spike train at 200 Hz), and the MIP/Poisson mixture. Of course, the parameter c is not directly related to the amount of synchrony in the input, but we could obtain a whole range of possible correlations with the same firing rate. It has to be stressed that the MIP process does not avoid bumps of synchrony: by pruning spikes from a mother process with a probability (1-c), when considering large populations, MIP generates large volleys (see Supplementary Figure 19)

The MIP is clearly not the only way to generate correlated spike trains in order to stimulate networks. Others techniques have been developed, for example in Brette (2008), and generative models such as the Glauber model could also be used to generate raster plots with controlled temporal correlations (Marre et al., 2009a).

12.2 Toward a mean-field-based description of topographical networks

The results of this topographical study are encouraging, because they show that depending on the scale you are looking at, the fine details of the underlying connectivity may not be necessary to explain the activity of the network. If you are interested in macroscopic and averaged measures, such as averaged firing rates, or global correlations over a whole population, then the fine structure in a local area is irrelevant, and in sparsely connected balanced networks, mean dynamical variables depend on mean structural variables. As shown in the main results section, it is mainly the state of the network (here the Synchronous Irregular regime, and the balance g between excitation and inhibition) that controls the profile of the spatio-temporal cross-correlations. Nowadays, recording techniques such as voltage sensitive dye give access to the averaged activity of the superficial layers over a rather large area of the cortex. In each pixel of the CCD camera, capturing the light emission which is proportional to the depolarization level of the neurons, VSD imaging record the averaged depolarization of a small group of neurons. Under the assumption that this group of neurons can be modelled as a random balanced network, which is, we agree, a "strong" assumption, its averaged activity (firing rate,

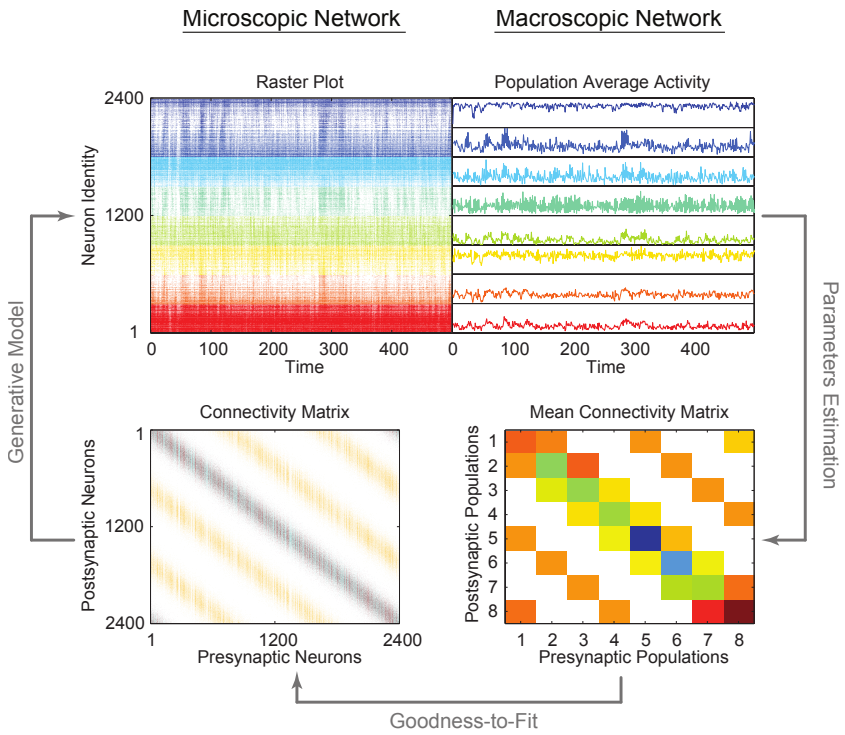

Figure 20: Schematic drawing of a mean field approach to infering anatomical connectivity from macroscopic dynamics. Microscopic activity generated by a particular connectivity matrix is recorded at a macroscopic scale, as population averaged activity, and one can try to infer a mean connectivity matrix at this macroscopic scale by finding the optimal coupling between a network of mean fields that gave rise to the observed activity.

variance) given by the VSD pixel can be used, knowing the existence of generic invariants, to infer some key parameters within this group of neurons. For example, the average excitatory/inhibitory couplings, the amount of external noise. Under this approximation, that each pixel can be considered as a random balanced network, whose dynamics could be estimated by mean-field equations, the full problem is therefore to study the coupled system made of several of those mean-fields, interacting one with the other. A schematic example of such an idea is represented on Figure 20. This is an inverse problem: knowing the averaged dynamics (mean and variance), given by the VSD data, try to find the optimal parameters/couplings with a network of mean-fields (biophysical models) to fit the data.

Preliminary results, obtained with simple binary neurons, show that the resolution of such an inverse problem may be feasible. Binary neurons are described by variables $\sigma_i \in \{0,1\}$. The state of the neuron i is given by its transfer function which in this case is an Heaviside function:

$$\sigma_i(\tau) = \Theta(I(\tau-1)) = \begin{cases} 1 \text{ if } I(\tau-1) \geq 0 \\ 0 \text{ if } I(\tau-1) < 0 \end{cases}$$

For mean-field models in irregular regimes, a Markovian description can be adopted to describe the average activity of mesoscopic populations (El Boustani et al., 2009). The corresponding transfer function for $i \in \{1,...,N\}$ populations is given by:

$$P(m_i^\tau | \{m_k\}^{\tau-1}) = \sqrt{\frac{N_i}{2\pi\nu_i(1-\nu_i)}} \exp\left(-\frac{N_i}{2}\frac{(m_i^\tau-\nu_i)^2}{\nu_i(1-\nu_i)}\right) \quad (13)$$

where m^{τ} is the activity at time τ, $m^{\tau-1}$ the activity at time $\tau-1$, N_i the number of neurons and ν_i the firing rate of population i.

The transfer function $\nu_i(m_k^{\tau-1})$ gives the mean activity of the population i knowing the previous state of all populations. These functions also depend on the coupling between every populations and we can use a maximum likelihood estimation to find these couplings from the activity. In the case of populations made with binary neurons, we have a sigmoidal transfer function to define the transfer function of a single population:

$$\nu_i(\{m_k\}^{\tau-1}) = \frac{1}{1+\exp(-\sum_{j=1}^{N} J_{ij} m_j^{\tau-1} - I_i^{\text{ext}})} \tag{14}$$

where J_{ji} are the coupling between population i and all the other, I^{ext} is an external noise term. With integrate-and-fire neurons and current based synapses, the transfer function is (Brunel and Hakim, 1999):

$$\nu_i(\{m_k\}^{\tau-1}) = \left(\tau^{\text{ref}} + \tau_{\text{m}}\sqrt{\pi} \int_{\frac{V_i^{\text{rest}} - <V_i>}{\sqrt{2}\sigma(V_i)}}^{\frac{V_i^{\text{thresh}} - <V_i>}{\sqrt{2}\sigma(V_i)}} du\, e^{u^2}(1+erf(u)) \right)^{-1} \tag{15}$$

with

$$< V_i >= V_{\text{rest}} + \tau_{\text{syn}}\left(\sum_{j=1}^{n} J_{ij} m_j^{\tau-1} + J_i^{\text{ext}} \nu_i^{\text{ext}}\right)$$

and

$$< \sigma^2(V_i) >= \frac{\tau_{\text{syn}}^2}{2(\tau_{\text{syn}}+\tau_{\text{m}})}\left(\sum_{j=1}^{n} J_{ij}^2 m_j^{\tau-1} + (J_i^{\text{ext}})^2 \nu_i^{\text{ext}}\right)$$

where, for each population i, V_{rest} is the resting potential of the integrate-and-fire neurons, V^{thresh} their spiking threshold, τ_{ref} their refractory period, τ_{m} their membrane time constant and ν_{ext} the frequency of the external noise.

And the problem is then to find:

$$\hat{J_{ij}}, \hat{I_i^{\text{ext}}} = \max_{J_{ij}, I_i^{\text{ext}}} \log\left(\prod_{\tau,i} P(m_i^{\tau} | \{m_k\}^{\tau-1}, J_{ij}, I_i^{\text{ext}}) \right)$$

Knowing the averaged activities of each pixel, one can try to infer the couplings that are more likely to have led to this activity. The constraints are given by the VSD recordings, and a biophysical model is used to invert the problem and, having the dynamics of a network of mean fields, guess the couplings. This idea should be developed further, because such a tool could reveal hidden anatomical structure that may not appear with simple cross-correlation measures. The case of correlated activity between populations should be more carefully studied, because the mean field approach described does not take those correlations into account, and assumes independence within populations. The resolution of the inverse problem depend on the observed activities. In asynchronous irregular regimes, the one mainly observed *in vivo*, this is rich enough to imprint, into the dynamical variables and their fluctuations, the influence of the connectivity. But in epileptic states (Synchronous Regular), where all the network is oscillating at a fixed frequency, the dimension of the activity is too low to constrain enough

the inverse problem. Further works should be continued in that direction. Such a conceptual framework may also be applied to other data type. As the Generalized Linear Models (Pillow et al., 2008), it can be a tool to unravel connectivity based on spatio-temporal correlations not within single neurons, but within populations of neurons. It may also be applied to calcium signals obtained by 2-photons imaging.

Part IV
Plasticity in neuronal networks

Summary

13 Introduction **117**
13.1 Supervised learning 118
13.2 Unsupervised learning 119

14 Hebbian rules and spike based plasticity **122**
14.1 Refinement of the Hebbian rule 122
14.2 Spike based plasticity and spike timing dependent plasticity 123
14.3 Classical model of STDP 125

15 A memory retention problem **127**
15.1 Materials and Methods 128
15.2 Results on the dynamics 130
15.2.1 Static properties of the network 130
15.2.2 Convergence and weights Distribution 132
15.2.3 Differences between excitatory and inhibitory synapses 133
15.3 Results on axonal versus dendritic delays 134
15.4 Emergence of structures 136

16 Discussion **137**

13 Introduction

Each millisecond, millions of electric signals are travelling in parallel from neuron to neuron, from cortical area to cortical area, converting almost instantaneously sensory inputs into active behaviour. Each millisecond, trillions of synapses are dynamically adapted, their weights being modified according to the past and present electrical activity in order to store and to enhance information transmission and reliability in this conversion procedure. Modification of these connection strengths between neurons in response to a particular external stimulation is supposed to be one of the key processes responsible for memory and information storage.

Several hypothesis have been made regarding the influence of the neuronal activity on the synapses, during the epigenetic development. Since the particular and exhaustive wiring scheme of the brain is not encoded genetically, it is commonly assumed that the topological structure of the neuronal network is dynamically rewired or refined during its development. Following the seminal works of Wiesel and Hubel (1965), Wiesel (1982) proposed a theory of plasticity for the development of the visual system (retina-LGN-cortex) with the "functional verification" hypothesis: functional connectivity is validated by the activity, stabilizing pre-existing connections and allowing the emergence of invariant feed-forward projections. The chemoaffinity theory, formulated by Sperry (1963), explains the genesis of mamalian visual systems in the retinotectal system with the idea that axon guidance and synaptogenesis proceed according to restrictive chemical markers between pre- and post-synaptic neurons. Both could be linked to the selective stabilisation theory of Changeux and Danchin (1976), where connections in the visual cortex or the neuro-muscular junction are supposed to be genetically specified between classes of cells, and gradually refined with the activity.

Many biological processes can act, on several time scales, to form and modify synaptic connections. Finding the exact mechanism by which the synapses between neurons are strengthened or weakened is a fundamental goal in the comprehension of the learning properties of the brain. Among all these plasticity mechanisms, associative plasticity is the one responsible for synaptic changes resulting from the association of two events. The classical view of associative memory in the brain is to view memories as cell assemblies or synfire chains that may have been co-activated once, and since this co-activation led to reinforcement of their coupling synapses, a memory may be recalled with a partial reactivation of this cell assembly, such as what was shown in Part II. Memories can be broadly distributed among cortical networks, spanning several sensory areas. In the neocortex, which is a cortex able to establish associations, this is particularly true, and multi-modal memories can be seen as an illustration of these cells assemblies. Having a particular picture in mind about you riding a bicycle on a windy road while smelling the odour of the freshly cut grass is a multi-modal association which can be recalled just by hearing again a particular sound first heard during this experience. Understanding how and why some neurons in distinct sensory areas succeed in establishing such associations is crucial to elucidate the learning problem faced by the brain.

With the Frozen Paradigm (Part II), we showed that statistics of the ongoing activity can be valuable, from a reliability point of view, when used as inputs to the system. With the topological network (Part III), we gained a better understanding of the dynamics and the structure of the activity in more realistic networks. In this Part, I would like to define the background and current views about the different kind of learning achieved in neuronal networks, in order to close the story of the Frozen Paradigm. If ongoing activity shares some statistical property with external inputs, in such a way that they are transmitted more efficiently into the network, then plasticity would be the missing piece allowing the storage of particular statistics and the

shaping of network activity. By establishing a match between ongoing and external activity, learning would tune the network to enhance information transmission within it, and this could be a guiding optimization principle constraining epigenetic development.

After having briefly reviewed the two main philosophies in machine learning, supervised and unsupervised learning, used to store information and memories in neuronal networks, I will focus more deeply on unsupervised learning theories and some experimental findings or observations made recently that have promoted the emergence of the concept of spike timing dependent plasticity (STDP). We will see how such a computational framework is attractive for storing information in networks of spiking neurons, but will also raise some key limitations and problems concerning the concept. The problem of memory retention in large scale neuronal networks is crucial, and based results, obtained in topographical networks such as the one explored in the previous Part, we will try to shape ongoing activity by external stimulations.

13.1 Supervised learning

Supervised learning, the more commonly used approach in artificial intelligence, is a machine learning framework in which where an error signal (or a reward signal, if we speak about reward learning) is present to instruct the system what should be learned, and when. Positive actions and/or behaviours and/or patterns produced in front of a set of learning stimuli are stored in order to produce a correct response to a new set of stimuli, by generalization. The system learns associations based on a list of training examples under the form $(\hat{x}, \hat{y})$, where $\hat{x}$ is typically a vector of input variables fed to the system, and $\hat{y}$ the desired output of the system: it can be a continuous value or a class if the goal is to classify the inputs. In the context of robotics, classification in order to produce decisions when confronted with new, unknown inputs is a challenging problem. There are several pieces of evidence that supervised learning exists in the brain, where reward signals may take the form of neurotransmitter or neuromodulator release. Each time a positive action is produced by the organism, neurotransmitters may validate or consolidate the synapses that lead to particular behaviours useful for the organism. Direct and indirect pieces of evidence are numerous to show how for example noradrenaline, dopamine and acetylcholine can influence the behaviour and therefore be a key factor in memory retention (Kasamatsu and Pettigrew, 1976, Romo and Schultz, 1990, Zhang et al., 2009). It has also been suggested that reward signals originating from the postsynaptic cell could back-propagate along the presynaptic axons and consolidate synaptic changes (Harris, 2008), providing a substrate for reinforcement of synaptic weights according to some feedback error signal. We will come to these points later.

Rate-based approaches In artificial neuronal networks, supervised learning can be achieved with widely used algorithms such as Kohonen learning rules, back-propagation algorithms or other learning procedures. Pioneering work was performed with the Perceptron by Rosenblatt (1958), which can be seen as the simplest kind of feed-forward neural network. The Perceptron is a linear classifier, biologically inspired with what can be seen as rate-based neurons and a neuronal structure. Nevertheless, despite being attractive, it suffered from severe limitations (unable to learn an XOR function), and more than 10 years were needed to find that several layers of such networks could do the job. Those frameworks are usually not applied to spiking neurons, because ongoing activity may disrupt the learning process and because the learning rules governing evolution of the synaptic weights would be too sensitive

to spike interactions. Instantaneous firing rates are more robust and rate-based models, neural mass or mean fields are more suited to implement, in a clear and theoretically tractable manner, analysis of convergence, stability, and so on. The key point of supervised learning is that a teacher is required to instruct the system what to do, so a training phase is always necessary to teach the correct actions and/or classifications to the system. Once examples have been learned, new stimuli can be presented.

Spike-based approaches The Liquid State Machine (LSM), explained in Maass et al. (2002) (or similarly the echo state network (Jaeger and Haas, 2004)) is a direct application of the learning possibilities given by chaotic neuronal networks made with spiking neurons. If you consider a random network of spiking neurons, it will spontaneously generate a very complex and unpredictable activity (chaotic): this is a complex dynamical system wandering on a high-dimensional attractor (see Part II). By correctly adjusting the weights of an "output" neuron (a reader), you can distinguish almost any particular pattern in the diversity produced by this reservoir. Appropriate linear combination of the non-linear interactions occurring within the reservoir can lead to a universal function approximator, close to support vector machines. Recent work on artificial intelligence and learning provide an extension of this LSM framework. On-line rules, based on averaged activities, can be used to learn any kind of pattern and to teach transformations to a neuronal network (Sussillo and Abbott, 2009). Other supervised paradigms are made with spiking neurons: recently, interesting techniques using spiking neurons have been developed, such as the tempotron (Gütig and Sompolinsky, 2006). Conceived as a temporal extension of the perceptron, it uses the integrative properties of the neuron to classify temporal patterns of spikes in a robust manner. The learning rule is supervised, and synaptic modification are made according to a home-made gradient descent. Last, by mimicking actions of neurotransmitters, such as dopamine, shown to have an effect *in vivo* on plasticity, one can validate synaptic changes at the synapses, and therefore create a learning window during which precise spike interactions will have an impact on the weight (Izhikevich, 2007, Legenstein et al., 2008, Baras and Meir, 2007).

13.2 Unsupervised learning

Unsupervised learning is the ability of the brain, to self-organize its connections within networks without any teacher or reward signal. This is the kind of learning which is thought to occur in all the primary sensory areas. In the visual cortex, for example, neurons are organized in functional maps of orientation preference, direction preference, ocular dominance and so on. No feedback signal is present to tell the brain what should be learned in the input signals. This spontaneous organization occurs even before any visual experience (Sur et al., 1999, Crair et al., 1998, Desai et al., 2002), and therefore the neurons spontaneously organize themselves into an efficient basis in which inputs can be easily decomposed and processed. Because of the presence of ongoing patterns of activity at a prenatal stage, numerous work in the vertebrate brain has shown that the travelling waves due "dark discharge" in the retina are sufficient to influence laminar segregation and ON & OFF specialization in the thalamus (Frégnac and Imbert, 1984, Shatz, 1996). The orientation columns in V1 are considered as crucial to allow fast and robust image segmentation, binding, and so on. This plasticity is mainly activity-driven: close-by neurons tend to receive correlated inputs, due to the retinotopic organization, and therefore a particular mechanism should consolidate the fact that receiving correlated activity implies a similar coding scheme, or common inputs. For a

good review of such activity-dependent changes, see Abbott and Nelson (2000). The basic principles underlying unsupervised learning theories are described in the following:

The Hebbian rule Most of the work on unsupervised learning in the brain has been based, since more than fifty years, on a well-known postulate made by Hebb (1949):

When an axon of cell A is near enough to excite cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.

Here can be found the basis of causal associations in the unsupervised learning framework. This postulate is a *local* and *associative* rule, meaning that causal activities between two connected neurons should lead to a synaptic reinforcement between them. Neurons do not need to integrate activity from other neurons: the evolution of the synapses between them results only from their own activities ("homosynaptic" plasticity, according to the terminology of Eccles). If $\langle x_i \rangle$ and $\langle x_j \rangle$ are the averaged firing rates of two neurons i and j, then if w_{ij} is a synapse between i and j, we have according to Hebb's postulate:

$$\frac{dw_{ij}}{dt} = F(w_{ij}, \langle x_i \rangle, \langle x_j \rangle) \tag{16}$$

with F being an unknown function. The rule is *cooperative*, the postulate stating that both pre- and post-synaptic neurons have to be active together in order to trigger synaptic modifications. Therefore, in the simplest form, the rule is used as:

$$\frac{dw_{ij}}{dt} = \eta \langle x_i \rangle \langle x_j \rangle \tag{17}$$

where η is a constant to scale the evolution of the weights. From a biological point of view, since synapses are well isolated and protected by glia cells, the synaptic clefts are rather isolated areas. Therefore, if we ignore the possible role of glia cells in memory consolidation, this Hebbian hypothesis that synaptic efficiency can be regulated only by pre-post electrical activity is tempting. More details about how exactly those changes in the synaptic weight may happen will be given in Part V. Nevertheless, biological evidence also suggest that synaptic modification may not only depend on the firing rates of the neurons. This framework of Hebbian learning has been very successful in reproducing the development of functional maps in V1, with firing rate models (Bednar et al., 2004). A network of mean-field models of excitatory and inhibitory populations will develop, according to Hebbian rules, and structure itself according to its inputs in such a way that can be closely compared to the architecture of the primary visual cortex. Nevertheless, some features of this Hebbian process are not biological and the rule, in its basic definition (see eq 17), is unstable. First, a normalization procedure is always assumed to keep weights in a certain range and avoid instability problems. This normalization procedure is far from being biologically explained. Then, as we will see later, weights can only be increased in this framework, leading to huge instabilities.

The Hopfield network One of the most famous examples of unsupervised learning in networks is made with binary neurons forming an associative memory and relies on this Hebbian

learning rule. This is the so called Hopfield model (Hopfield, 1982). Its simplicity and its efficiency has made it widely popular, and it has been used for a lot of learning problems, like pattern recognition or classification. The network is composed of N units x_i, each of them being connected to all the others with weights J_{ij} governed by equations:

$$x_i(t+1) = \text{sgn}(\sum_j J_{ij}x_j(t) - \theta_i) \tag{18}$$

θ_i is a threshold to determine the spin of the unit x_i: -1 or 1. The model is able to learn some repeatedly imposed patterns through a Hebbian algorithm, which will update the synaptic weight J_{ij} by adding ΔJ_{ij}, which will be proportional to the averaged correlation between the two activities:

$$\Delta J_{ij} \simeq \langle x_i \rangle \langle x_j \rangle \tag{19}$$

The convergence to imposed patterns will depend on their number, and their relative orthogonality (Hopfield, 1982). Theoretical studies have been performed on the ratio of number of patterns that can be stored in the network to the number of neurons (Amit et al., 1985), and the ratio is rather low, around 0.15.

The Boltzmann machine We introduce another learning model, very similar to the Hopfield network, which will be referred to in our following work. Roughly speaking, the Boltzmann machine is a stochastic version of the Hopfield network. If we keep the same notation as for the Hopfield network, we have a global energy value E, defined by the following equation:

$$E = -\sum_{1 \leq i,j \leq N} J_{ij}x_i x_j + \sum_i \theta_i x_i \tag{20}$$

Switching the spin of the unit i from -1 to 1 results in an energy difference of:

$$\Delta E_i = \sum_j J_{ij}x_j - \theta_i \tag{21}$$

In the Boltzmann machine, transitions of the units are stochastic. The probability of the unit i being in state 1 is:

$$p_i = \frac{1}{1 + \exp(-\beta \Delta E_i)} \tag{22}$$

This model can be used to find the global minimum of an energy function, by slowly decreasing the "temperature" $T = 1/\beta$ while always leaving network time to reach thermal equilibrium (Ackley et al., 1985). For learning, the Boltzmann machine is separated into two layers, a visible layer, and a hidden layer. An input from the environment corresponds to fixing the visible layer units to the values of the input. This is the "clamped" phase. Alternatively, the network can run freely. The Boltzmann machine is able to learn and internally represent a given distribution of inputs by the following algorithm (Ackley et al., 1985):

- The visible layer is clamped to some inputs sampled from the distribution which has to be learned, and the network runs until it reaches a thermal equilibrium. The probability of synchronous activation p_{ij} of the units i and j is estimated.
- The network runs freely. The probability of synchronous activation p'_{ij} of the units i and j is estimated.
- The weights are updated according to the rule: $\Delta J_{ij} = \varepsilon(p_{ij} - p'_{ij})$.
- Repeat these 3 steps until the weights converge to a steady state

The strength of the Boltzmann machine is that this algorithm is proven to work for any input distribution. The Kullback-Leilber distance between the environment-constrained distribution and the free-running distribution of visible units is denoted G, and by definition:

$$G = \sum_{\alpha} P_{\alpha} \ln(\frac{P_{\alpha}}{P'_{\alpha}}) \tag{23}$$

where P_{α} is the probability of the state α of the visible units when constrained by the environment, and P'_{α} the probability during free-running. The nice property of Boltzmann machine is that:

$$\frac{dG}{dJ_{ij}} = -\beta(p_{ij} - p'_{ij}) \tag{24}$$

The gradient thus depends only on local variables, and the preceding algorithm is thus a simple gradient descent. In the case where there are only visible units, the function G is convex, and the algorithm is thus guaranteed to converge to a local minimum.
The Boltzmann machine has been sometimes considered as a general learning machine. In practice, several problems arise when increasing the number of units. Notably, the learning time increases exponentially with the size of the machine, making it intractable for real-size problems.

14 Hebbian rules and spike based plasticity

14.1 Refinement of the Hebbian rule

The covariance rule In the very naive form, given Eq. 17, the Hebb rule does not really account for depression, i.e. for the fact that a synapse can also be weakened. Hebb's postulate has been extended by Stent (1973), taking into account the fact that synapses could also be depressed:

When the pre-synaptic axon of cell A repeatedly and persistently fails to excite the post-synaptic cell B while cell B is firing under the influence of other pre-synaptic axons, metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is decreased.

We also can then talk of a covariance rule (Stanton and Sejnowski, 1989). When activities covary, the synaptic link between them should be enforced or weakened. A detailed analysis of all these rules can be found in Gerstner and Kistler (2002). The covariance rule takes into account the averaged activity $\langle x_{i/j} \rangle_T$ over a long time window T, and leads to the following equation:

$$\frac{dw_{ij}}{dt} = \eta \langle x_i - \langle x_i \rangle_T \rangle \langle x_j - \langle x_j \rangle_T \rangle \tag{25}$$

η is again a scaling constant, and $\langle\rangle$ is the instantaneous firing rate. An illustration of the synaptic modification expected by the different plasticity scheme can be found in Figure 21. *In vivo*, these covariances rules have been explored by Frégnac et al. (1988, 1992, 1994), Stanton and Sejnowski (1989), Debanne et al. (1995, 1998). The fact that the rule varies according to $\langle x_{i/j} \rangle_T$ implies that it depends on the past activities of the pre- and post-synaptic neurons,

averaged over a certain time window T. This phenomenon is a form of metaplasticity: the plasticity rule itself is plastic, and this will be the subject of Part V.

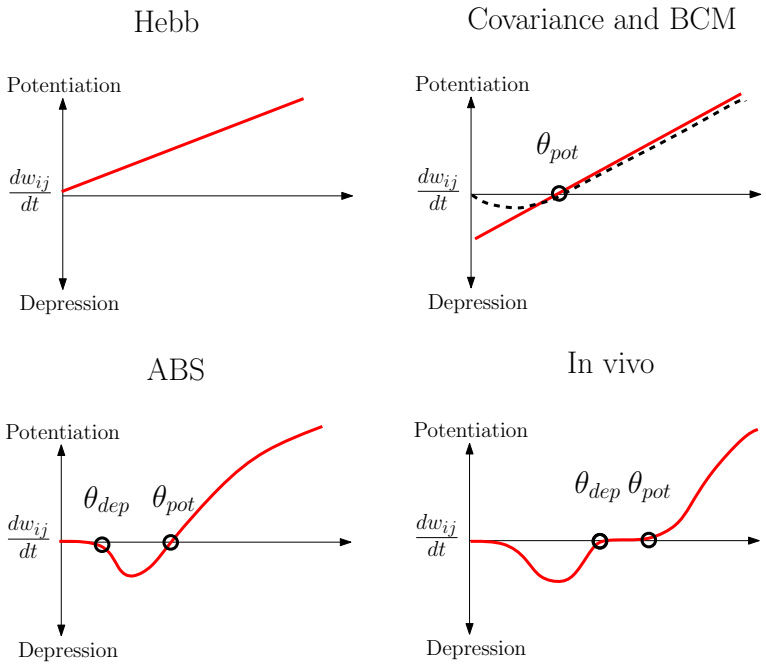

Figure 21: Illustration of different schemes of plasticity: the Hebbian rule, the covariance rule, the the BCM rule (black dash line), the ABS rule, and some *in vivo* observations. For a fixed pre-synaptic firing rate, the graphs show modification of the synaptic weights as a function of the post-synaptic firing rate, for Poisson firing statistics. Adapted from Frégnac et al. (2010).

The BCM rule In order to extend the covariance rule and approach the problem of weight normalization, Bienenstock et al. (1982) designed a model of synaptic plasticity that was able to reproduce phenomenologically several observations made *in vivo* (Clothiaux et al., 1991, Kirkwood et al., 1996). The BCM rule (BCM stands for Bienenstock, Cooper, and Munro), is a physical theory of learning in the visual cortex. The formalism was an efficient way to balance and regulate by a heterosynaptic process, the amount of plasticity according to past activity. It is close to the covariance rule, with a sliding threshold mechanism that regulates the amount of plasticity, according to the past activity of the synapse. Under the BCM rule, we have, if $\langle x_i \rangle$ is the firing rate of the pre-synaptic neuron and $\langle x_j \rangle$ is that of the post-synaptic one:

$$\langle \frac{dw_{ij}}{dt} \rangle = \langle x_i \rangle \Psi(\langle x_j \rangle, \theta(\langle x_j \rangle)) \tag{26}$$

with $\Psi(\langle x_j \rangle < \theta, \theta) < 0$ and $\Psi(\langle x_j \rangle > \theta, \theta) > 0$. θ is a sliding threshold and depends on $\langle x_j \rangle$. The relationship between θ and $\langle x_i \rangle$ should be supralinear: quadratic or more.

The ABS rule The ABS rule (Artola, Brocher and Singer) (Artola et al., 1990), linked with calcium-based models of plasticity (Lisman, 1989), adds one other sliding threshold for depression. Not only potentiation but also depression depends on the past activity. As one can see in Figure 21, no modification of the synapses occur if the post-synaptic activity is below a certain threshold θ_{dep}. If crossed, then depression takes places, before an other threshold θ_{pot} is crossed, triggering potentiation. The direction of the synaptic gain change depends on the

membrane potential of the post-synaptic cell, or on the amplitude of the surge of Ca^{2+}. The exact details of the machinery underlying these changes, from a biological point of view, will be discussed more in depth in Part V.

14.2 Spike based plasticity and spike timing dependent plasticity

Theta burst stimulation Bliss and Lomo (1973) discovered long-term potentiation (a long-lasting form of synaptic plasticity) in the hippocampus by high-frequency stimulation (HFS) of pre-synaptic afferents. Strong burst of pre-synaptic activity affected the amplitude of the synaptic EPSP in a stable manner, with changes lasting for more hours or more. This pioneering work showed how activity could shape synaptic efficacy.

Spike timing dependent plasticity The best experimental setup for exploring plasticity in a controlled manner is the *in vitro* setup. By using pairs of neurons clearly isolated and connected (using either brain slices or cultured neurons), one can patch the pre- and the post-synaptic neuron and observe the synaptic modifications between them according to their discharge. Based on those controlled experiments, the most recent and promising candidate to support unsupervised learning algorithms in the brain, based on neuronal activity, is the spike timing dependent plasticity (STDP).

There is indeed several evidence (Bi and Poo, 1998, Markram and Tsodyks, 1996, Gerstner et al., 1996) in neocortex that the efficiency of a synaptic connection between two neurons may be regulated by the precise timing of the joint activity of the neurons. This postulate, originally made by Hebb (1949), has been demonstrated in a lot of *in vitro* experimental studies in the form of the STDP rule. This is an *associative* rule that needs to be distinguished from short term plasticity or homeostasis phenomena, involving only integration of pre-synaptic activity (Tsodyks et al., 2000, Turrigiano and Nelson, 2004). As one can see in Figure 22, taken from Dan and Poo (2004), when pre-post pairings are made repeatedly at a fixed frequency of 1 Hz, with a particular time difference δt between pre and post spikes, synaptic modifications are observed whose magnitude depends on δt. For positive values of δt, when pre-synaptic spike occurs before the post, the synapse is potentiated. Oppositely, if δt is negative, the synapse is depressed. Both mechanisms occur in relatively short time windows of $\simeq 20$ ms, and a double exponential fit made on the data is the classical shape everybody has in mind when talking about STDP. That 20 ms time scale is the time window for triggering a change, but the actual change happens much more slowly. While some recent evidence may suggests that STDP can also be found *in vivo* (Crochet et al., 2006, Zhang et al., 1998, Young et al., 2007, Jacob et al., 2007), the impact of such a rule on a network level is still misunderstood, and part of the problem comes from the fact that there is a lack of data on the properties and the relevance of STDP *in vivo*. For review, see Caporale and Dan (2008). The STDP phenomenon as seen *in vitro* is appealing from a theoretical point of view. If a pre-synaptic spike occurs just before a post-synaptic one, the strength of the synapse between the two neurons tends to be increased. Conversely, if the pre-synaptic spike comes just after a post-synaptic one, the synaptic strength tends to be decreased. This rule establishes a link with Hebb's postulate and could allow neurons to learn causal chains of information: if pre-synaptic information is important in the discharge of the post-synaptic neuron, then synapse is strengthened, otherwise it is weakened. Interestingly, rules symmetric in sign have been observed in the electro senseory lobe of the electric fish by Bell et al. (1997) and have been used in models to decorrelate the sensory stream from expected inputs linked with the motor-induced reaffernce.

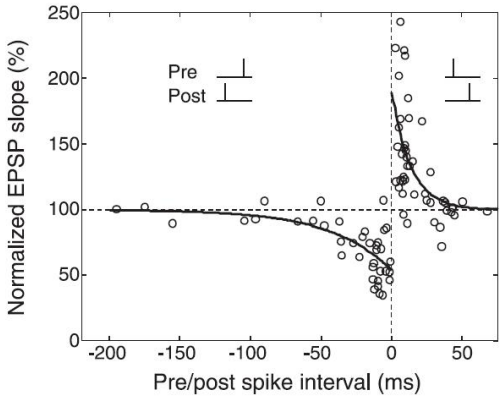

Figure 22: Illustration of spike timing dependent plasticity time windows, taken from Bi and Poo (1998). Depending on the precise time difference δt between a post- and a pre-synaptic spike, the synaptic weight can be either depressed or potentiated.

Since this seminal work, several theories have been proposed for a conceptual explanation of these STDP curves. The promising link between STDP and the Hebbian rule has led several authors to try to find a more generic optimization principle behind this canonical shape. The quest is "can STDP be seen as a biological response to an optimization problem?" with a goal function like $\Phi(x_{\text{input}}, y_{\text{output}})$, if x_{input} are the inputs to the neuron, and y_{output} its responses. Are the shapes of those curves telling us something about the learning strategies performed by the neuron? According to Toyoizumi et al. (2007), Chechik (2003), STDP could be seen as an attempt, by the neurons, to maximize the transmission of information and therefore the mutual information between inputs and outputs, $MI(y_{\text{output}}, x_{\text{input}})$. For Bohte and Mozer (2005), STDP is more a way to reduce the variability of the output knowing the input: $H(y_{\text{output}}/x_{\text{input}})$ (with H being the entropy). We can cite other examples such as slow feature analysis (Sprekeler et al., 2006), where STDP aims to decompose the signals into a basis of signals, slowly varying in time, or the predictive coding (Rao and Sejnowski, 2001) theory, where STDP is used to encode only time differences. Nevertheless, as we will see later, since STDP is still, from a biological point of view, a phenomenon which is not understood, all these theories, even if conceptually promising, can not pretend to understand STDP in its globality.

14.3 Classical model of STDP

From a modeller's point of view, the rule is ill defined. A good review on all the important aspects of such modelling is given in Morrison et al. (2008). In its most widely used formulation, one can model STDP with the following system of equations:

$$\delta w = \lambda \begin{cases} a_{\text{pot}} w^{\mu_{\text{pot}}} e^{-\frac{\delta t}{\tau_{\text{pot}}}} & \text{if } \delta t = t_{\text{post}} - t_{\text{pre}} > 0 \\ a_{\text{dep}} w^{\mu_{\text{dep}}} e^{-\frac{\delta t}{\tau_{\text{dep}}}} & \text{if } \delta t = t_{\text{post}} - t_{\text{pre}} < 0 \end{cases} \tag{27}$$

λ is the learning rate, a_{pot} and a_{dep} the scaling increments of the synaptic weights performed at each pairing, for potentiation and depression. Each time a pre or a post-synaptic event appears, weights are updated accordingly. τ_{pot} and τ_{dep} are the time constants of the double exponential shape observed in biological data, such as the one that can be seen in Figure 22.

Typical values are in the range 10-30 ms. μ_{pot} and μ_{dep} are generic exponents to model the fact that weight increments can depend on the current values of the weights. Taking advantage of the exponential, the most efficient way to implement this systems, at the synapse level, is to define two local variables $\theta_{\text{pot}}(t)$ and $\theta_{\text{dep}}(t)$, such that:

$$\begin{aligned} \frac{d\theta_{\text{pot}}(t)}{dt} &= -\frac{\theta_{\text{pot}}(t)}{\tau_{\text{pot}}} \\ \frac{d\theta_{\text{dep}}(t)}{dt} &= -\frac{\theta_{\text{dep}}(t)}{\tau_{\text{dep}}} \end{aligned} \tag{28}$$

Each time a pre-synaptic spike occurs, $\theta_{\text{pot}} \rightarrow \theta_{\text{pot}} + 1$, and each time a post synaptic spike occurs, $\theta_{\text{dep}} \rightarrow \theta_{\text{dep}} + 1$. In this case, if θ_{pot} and θ_{dep} are not bounded, the integration scheme of the STDP is said to be all-to-all. All previous pre- or post-synaptic spikes contribute to the modification of the weight at time t, since they have an impact on θ_{pot} and θ_{dep}.

In the so-called nearest-neighbour interaction scheme, θ_{pot} and θ_{dep} are bounded by 1, and only the nearest either pre or post synaptic spike is considered for potentiation or depression. This difference is important, because STDP in its basic form with an all-to-all interaction scheme is not compatible with the BCM theory, as shown in Izhikevich and Desai (2003). Only the nearest-neighbour scheme can provide a BCM behaviour with the rule. The values of λ, a_{pot}, a_{dep}, μ_{pot}, μ_{dep}, τ_{pot} and τ_{dep} are selected according to the STDP desired rule. To simplify the following notations, we set $a_{dep} = \lambda a_{pot}$ (and thus we should have $\lambda < 0$, because depression decreases the weight). The pairing scheme used during all the the simulations is all-to-all, meaning that all the interactions between pre and post synaptic spikes are taken into account.

Weight-dependent model of STDP Regarding the weight modifications performed by such plasticity rules, there are two main classes of STDP rules that are commonly used in modelling studies of neuronal networks. They are categorized as either "additive" or "weight-dependent", depending on how current synaptic weight impacts the change in the weight of the synapse (Gütig et al., 2003). These classes are established with the exponents μ_{pot} and μ_{dep}. If both are set to 0, then the STDP is "additive". Each time a weight modification is made, increments are only determined by a_{pot} and a_{dep}, without taking into account the current weight of the synapse. The rule needs a hard bound thresholding to constrain the weights between $[w_{\min}, w_{\max}]$. Oppositely, in the weight-dependent rule, the weight increments are a function of the current weight of the synapse. This is the case if μ_{pot} and μ_{dep} are positive. The biological evidence for "additive"-only rules is quite thin. The original data for STDP, and especially the synaptic modification observed as a function of the initial amplitude of the EPSP show (Figure 23) that the relative changes are not similar for potentiation and depression. Modifications for potentiation seem to be independent of the initial amplitude of the EPSP, while this is not the case for depression. Bi and Poo (1998) proposed, initially, a log-linear relationship for depression, while potentiation is much more additive. For a precise fit, see Morrison et al. (2008), Standage et al. (2007), but the exact values for μ_{pot} and μ_{dep} are not crucial, as long as they are not zero.

As pointed out in Gütig et al. (2003), the additive case, often used in models, is a very particular case with particular dynamics. It has been shown in van Rossum et al. (2000), Billings and van Rossum (2009), through a Fokker Plank approach, that an additive STDP rule always drive the weight distribution to a bimodal one, with all weights being clipped either at $w_{\min}$ or at $w_{\max}$. Nevertheless, they encourage synaptic competition and allow a better storage of

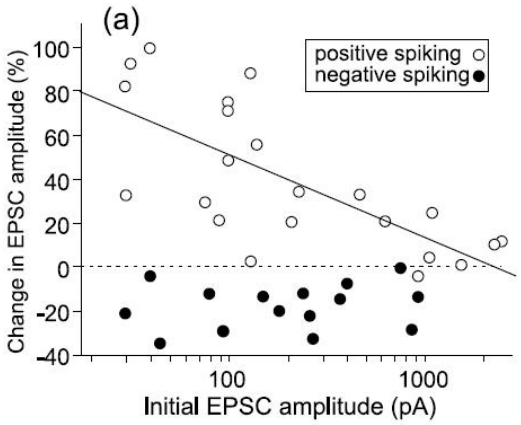

Figure 23: Relative weight modifications for potentiation (positive spiking) and depression (negative spiking), according to the initial amplitude of the EPSP. Taken from Bi and Poo (1998).

patterns (Fusi and Abbott, 2007). As we will see in the following, they are less sensitive to the memory retention problem occurring in recurrent networks. Moreover, it is also known that in cortex, and also in the cerebellum a lot of the synapses are considered as almost silent. One could see here the evidence for bimodal distribution resulting from the additive rules. On the contrary, weight-dependent rule leads to a unimodal distribution of the weights, more biologically plausible, but does not allow the emergence and the survival of neuronal structures in balanced random networks (Billings and van Rossum, 2009, Morrison et al., 2007a).

Plasticity of the inhibitory synapses Although inhibitory interneurons modulate many neuronal processes, the evidence for plasticity at inhibitory synapses remains scarce. Some studies report strengthening of inhibitory synapses in negative rate covariance regimes (Komatsu and Iwakiri, 1993), and spike timing dependent plasticity of inhibitory synapses has also been reported (Haas et al., 2006) as well as spike timing dependent depression of excitatory synapses on fast spiking inhibitory interneurons. Almost all models of plastic networks consider that only excitatory synapses are plastic, because most of the biological evidence for STDP has been gathered for synapse between excitatory neurons, far more numerous and easy to patch than inhibitory ones. Nevertheless, the question of plasticity at inhibitory synapses remains open, and could greatly help stabilization in recurrent networks. Haas et al. (2006) found an anti-Hebbian rule for inhibitory synapses. Pre-post pairing led to reinforcement of the synapse, meaning to an increase in the amplitude of the post-synaptic inhibitory post synaptic potential (IPSP), while post-pre led to a decrease. This anti-Hebbian rule, from a conceptual point of view, offers nice theoretical possibilities. In artificial neural networks, anti-Hebbian rules for inhibition are important to balance the changes at the excitatory synapses and allow the network to perform robust principal or independent component analysis (Plumbley, 1993). In all the following, we decided to have only plastic excitatory synapses, to be in line with previous literature and also to have a clear insight into the effects induced by plasticity.

15 A memory retention problem

To determine if external statistics learned through STDP can reverberate in the ongoing activity of a recurrent network, we create a large recurrent network with plastic synapses. After having converged to an equilibrium, the idea is to see if external inputs can establish a long lasting trace in the spontaneous activity, after stimulation. Such reverberations have been observed *in vivo*, with voltage sensitive dye imaging by Han et al. (2008). In this paper, the authors showed that the presentation of a visual stimulus in the rat barrel cortex elicits an evoked response in the form of a travelling waves that pop up and propagates in the visual cortex V1, and that this evoked response tends to be spontaneously replayed in the ongoing activity during several minutes after the and of the stimulation. When recording the ongoing activity in the visual cortex minutes after the stimulation, they were still able to see some kind of reverberation of the input. Such replay of particular patterns in ongoing activity is also tackled in a lot of other studies, such as Ikegaya et al. (2004) and Mokeichev et al. (2007). This replay would establish a link with the theoretical work performed with the Frozen Paradigm (see Part II): since the ongoing activity has a particular impact when used as an external stimulation, it would be interesting to see how its statistics, constrained by the wiring of the network, can be shaped by synaptic plasticity.

To address this question, we begin by studying the STDP rule equilibrium in a topographical network, similar to the one previously explored in Part III. Since we use weight-dependent rules, as in Morrison et al. (2007a), we already know that stable connections do not developed in random network. In this work, the authors showed that ongoing activity in random balanced networks was disrupting the formation of a stable connectivity, and demonstrated in large-scale recurrent networks the memory retention problem of weight-dependent STDP rules. The Frozen Paradigm requires an efficient way to shape, in a stable manner, the statistics of the ongoing activity. Stimulus statistics, as in the Han et al. (2008) VSD study, should be captured and integrated into the attractor of the ongoing activity, in order to enhance the reliability of the triggered responses. Since we know these statistics are driven by the connectivity, we need an efficient way to store it, in an unsupervised manner, into the synaptics weights of the network. The idea of this Part is to see whether the connectivity structure can promote or inhibit the development of such stable weights structures, and how stimulations could affect the equilibrium of a plastic reccurent nework.

15.1 Materials and Methods

Neuron model Neurons are modelled as leaky integrate-and-fire neurons with alpha-function shaped current-based synapses. The detail of their parameters is given in Table 2.

τ_m	τ_{ref}	τ_{exc}	τ_{inh}	V_{leak}	V_{th}	V_{reset}	c_m
20 ms	2 ms	0.33 ms	0.33 ms	0 mV	20 mV	0 mV	250 pF

Table 2: The cells parameters used in all the simulations. The cell model is a classical Integrate and Fire neuron with an alpha-shaped decay for its current-based synapses.

Network Structure We consider a topographically organized balanced network of N^2 neurons, arranged on a 2D grid as in Mehring et al. (2003), with a classical ratio of 4:1 between

the number of excitatory and inhibitory neurons: $N_{\text{inh}}^2 = \frac{N_{\text{exc}}^2}{4}$

A layered organization These excitatory and inhibitory neurons are placed on two grids of sizes $N_{\text{exc}} \times N_{\text{exc}}$ and $N_{\text{inh}} \times N_{\text{inh}}$ respectively, both corresponding to a cortical area of size $\alpha \times \alpha$ mm^2. In the following, we consider this area as a toroidal surface with periodic boundary conditions to avoid any border effects and grid artefacts will be neglected.

Probabilities of connections The probability p_{ij} of connection between two cells i and j is drawn from a random distribution dependent on the distance l_{ij} between the two cells, expressed in mm. More precisely, we have

$$p_{ij} = p_0 e^{-\frac{l_{ij}^2}{2\sigma^2}} \tag{29}$$

where p_0 is the amplitude of this Gaussian shaped curve, and not the probability for a neuron to be connected with himself: autapses (a neuron connecting to itself) are not allowed in the network for stability reasons. The values of p_0 and σ are extracted from fitting such a Gaussian curve to the data of Hellwig (2000) concerning the observed connections within layer III, which gives $p_0 = 0.7$ and $\sigma = 0.175$ mm (ensuring a connectivity of approximatively 13% in a square millimetre). Figure 24 summarizes the general layout of the network and the wiring scheme.

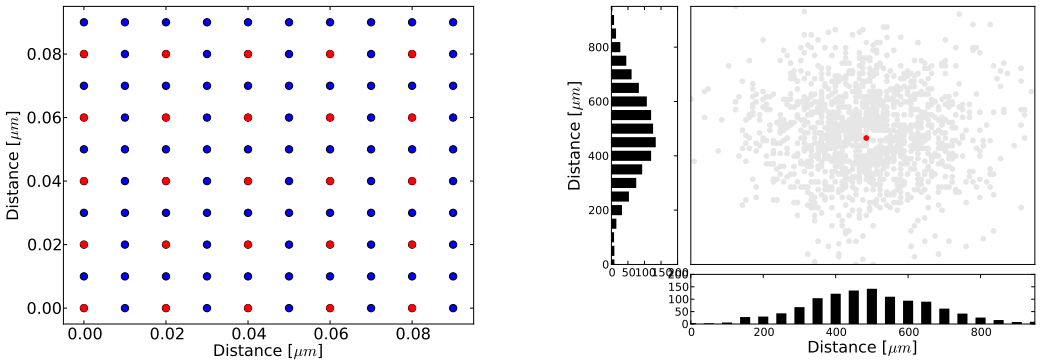

Figure 24: Connectivity of the topographical network. Left: cells are arranged on 2D grids with distinct spacing for excitatory (blue) and inhibitory (red) cells. Right: Each cell is connected to its local neighbourhood according to a Gaussian probability profile.

Linear delays Similarly to the model previously described in Part III, we used linear delays:

$$d_{ij} = d_{\text{syn}} + \frac{d}{v} \tag{30}$$

where d_{syn} is the synaptic transmission delay, set to 0.1 ms, d the Euclidean distance, on the torus, between neurons i and j, and v the axonal velocity. In all simulations, $v = 0.3$ m.s^{-1} (value taken from (Bringuier et al., 1999, Benucci et al., 2007)).

Control Network To check what are the effects due in particular to the topographical structure, we designed a control network with the same numbers of cells, the same averaged connectivity (in term of probability of connections) and delay distribution, but without any topographical organization. This shuffled version of the 2D network is a pure sparsely connected random network.

Simulation software For all the simulations, we used a parallel implementation of the NEST software (Gewaltig and Diesmann, 2007), with a time step of $dt = 0.1$ ms

Measures The fano factor (FF), like the coefficient of variation (CV), is a measure of the dispersion of a probability distribution. The fano factor is defined as: $FF = \frac{\text{var}(X)}{\langle X \rangle}$ where var is the variance and $\langle\rangle$ the mean of a random process X in some time window T. In our case, the fano factor is computed on the average spike count over a population of 1000 randomly sampled neurons within the network, binned with $T = 3$ ms.

Cross-correlations within the network are computed by averaging correlation coefficients over N pairs of cells. For such pairs, if x and y are the time series of the spiking activity histograms, binned with a time bin T, we have:

$$r = \frac{\sum_{i=1}^{n}(x_i - \langle x \rangle)(y_i - \langle y \rangle)}{(n-1)\text{std}(x)\text{std}(y)} \tag{31}$$

where $\langle x \rangle$ and $\langle y \rangle$ are the mean of the time series and $\text{std}(x)$ and $\text{std}(y)$ their standard deviations.

Equilibrium reached with STDP In the case of a neuron receiving uncorrelated Poisson spike trains, the equilibrium weight w^* achieved by STDP can be calculated analytically either in a nearest neighbour interaction scheme or in an all-to-all scheme (Standage et al., 2007). In the case of a multiplicative rule such as ours, and an all-to-all scheme, with $\delta w = \lambda a_{\text{pot}} w^{\mu_{\text{pot}}} e^{-\frac{\delta_t}{\tau_{\text{pot}}}}$ for potentiation, and $\delta w = \lambda a_{\text{dep}} w^{\mu_{\text{dep}}} e^{-\frac{\delta_t}{\tau_{\text{dep}}}}$ for depression we have, where r is the firing rate of the post-synaptic neuron :

$$w^* = -\left(\frac{1}{\lambda}\frac{\tau_{\text{pot}} + r}{\tau_{\text{dep}} + r}\right)^{\frac{1}{\mu_{\text{dep}} - \mu_{\text{pot}}}} \tag{32}$$

So, knowing the stationary firing rate r of the static recurrent network, without plasticity, if we want to have an equilibrium at w^*, then we have to fix:

$$\lambda = -\frac{1}{w^{*(\mu_{\text{dep}} - \mu_{\text{pot}})}}\frac{\tau_{\text{pot}} + r}{\tau_{\text{dep}} + r} \tag{33}$$

15.2 Results on the dynamics

All the following results are obtained in a 12500 neuron network (10000 excitatory cells, 2500 inhibitory cells), simulating a diluted cortical area of 1 mm^2. We summarize here briefly the parameters used to obtain these first results. We use $p_0 = 0.7, \sigma = 0.175$, the same parameters for the leaky integrate and fire neurons with alpha-function shaped current-based synapses as

in Morrison et al. (2007a), except that we add a refractory period of $\tau_{ref} = 2$ms. Each neuron receives 1600 external Poisson inputs at excitatory synapses, each at a frequency $\nu_{ext} = 6$ Hz. We use the rule described in van Rossum et al. (2000) for STDP at the excitatory synapses, with $\tau_{pot} = 14$ ms, $\tau_{dep} = 34$ ms, and λ is set in order to have an equilibrium at 100 pA. The parameters are summarized in Table 3.

λ	τ_{pot}	τ_{dep}	μ_{pot}	μ_{dep}
see Eq 33	14 ms	34 ms	0	1

Table 3: Spike timing dependent plasticity parameters. We used the rule proposed in van Rossum et al. (2000) with additive potentation and multiplicative depression. λ is defined according to the weight w^* we want to reach, and the time constant of the depression window is higher than the one for excitation.

15.2.1 Static properties of the network

Description of the Asynchronous Irregular regime The weights were initially set to put the network into an Asynchronous Irregular (AI) regime (see Brunel (2000)). We use $w_{exc} = 100$ pA and $w_{inh} = -gw_{exc}$ with $g = 6$. This AI state can be described as follows: in the static case, the network has an average firing rate of 10.2 Hz and a mean coefficient of variation of the inter-spike interval of 0.88. The targeted weights w^* is therefore only a rough approximation since these numbers shows us that the behaviour of the network can not be reduced to a pure Poisson process. The fano factor, computed with a bin size of 3 ms, is 5.85 (see Figure 25)

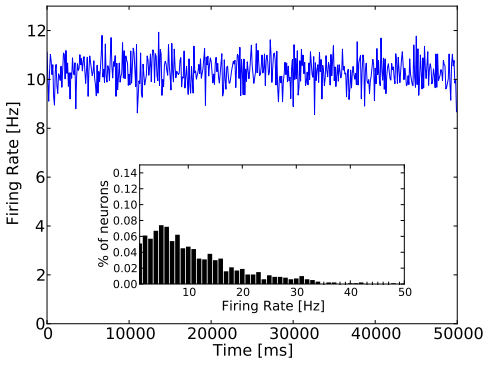

Figure 25: Firing rate of the network averaged over 1000 neurons for 50 s with a bin size of 100 ms. The inset shows the distribution of the firing rates for those recorded neurons. The heterogeneity of the number of synapses per neuron is responsible for the Poisson distribution (El Boustani and Destexhe, 2009b)

It is important to notice that, as can be seen in the inset of Figure 25, the heterogeneity in the mean number of efferens synapses per neuron introduces a variability in the firing rates. For a network where all neurons establish a fixed number of connections (Mehring et al., 2003), the distribution of the firing rate is much more closer to a Gaussian curve and the network suffers from a lack of heterogeneity. We check with the control network that these properties are generic and not due to the structure. In a randomly connected network with the same

macroscopic properties, we can observe similar values for the first-order statistics: the firing rate settles down at 11.2 Hz, and the ISI CV is 0.9.

Cross correlations within the Network The differences between the topographical and the randomly connected control network appear in the structure of the cross correlations. As one may expect, they depend on the distances in the former network, and not in the latter, similar to results shown in Part III. We computed the Pearson coefficient for the cross correlations between pairs of cells sorted according to the distances between them (see Figure 26, left). There is a strong dependence in the topographical network which is not present in the random one, and correlations roughly follow an exponential decay. The influence of the transmission velocity on this decay can be observed for several values of v. As a comparison, we plot the Pearson correlation coefficient in the random network.

Without taking the positions into account, we computed the cross correlogram of the activities between the excitatory cells. In the following, all the cross-correlogram are averaged over 500 pairs of neurons and 50 s of simulation using a bin size of 0.1 ms. Figure 26 (Right) shows that the cross-correlogram of the overall activity is symmetric, with fast oscillations, reflecting the balanced state of the network. Oscillations do not spread into the temporal domain.

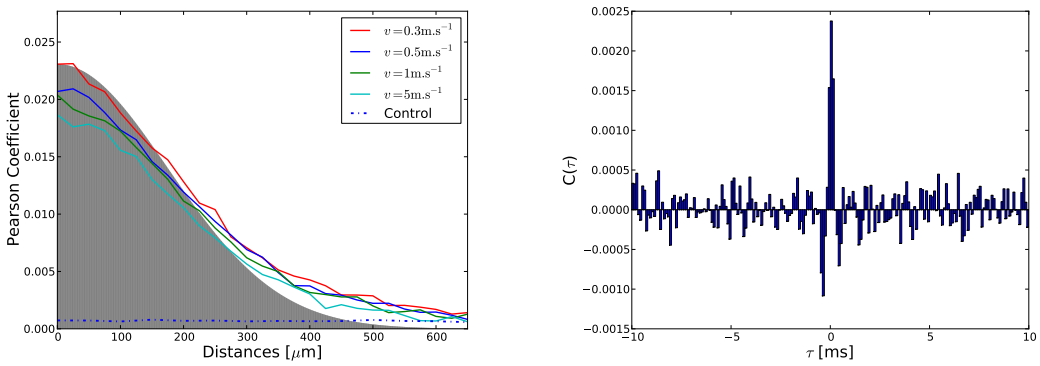

Figure 26: Cross correlations in the static network. Left: the Pearson correlation coefficient averaged over 500 pairs of neurons recorded for 50 s, and whose activities have been gathered with a bin size of 0.1 ms. In shaded gray, the Gaussian profile of the connections. Right: The averaged time course of all the pairwise cross-correlations.

For small distances (between 0 and 0.1 mm), we have a strong cross-correlation. For two populations separated by 0.3-0.4 mm, the peak of the correlation tends to decrease, and for two populations far apart, separated by 0.6-0.7 mm (0.7 mm being the maximal distance on the toroidal surface), we have only a small residual cross-correlation, implying that the major part is due to local activity. Those data, even if obtained in a slightly different network compared to the one simulated in Part III, and especially in a different regime (AI vs SI), are in line with previous results.

15.2.2 Convergence and weights Distribution

We next turned on STDP at all the excitatory synapses, and recorded the evolution of the efferent weights of 1000 neurons, simulating 250 s of biological time for the network. As we could predict, the weights for the excitatory to excitatory and excitatory to inhibitory

connections settled down into two Gaussian distributions (called respectively $\mathcal{N}(\mu_{EE}, \sigma_{EE})$ and $\mathcal{N}(\mu_{EI}, \sigma_{EI})$) with means close to the targeted equilibrium w^*: $\mu_{EE} = 106.3$ pA and $\mu_{EI} = 107.2$ pA ($\sigma_{EE} = 25.6$ pA and $\sigma_{EI} = 26$ pA). The statistics of the network are roughly invariant, with a global firing rate during the last 50 s, averaged over 1000 neurons, of 11 Hz. The mean CV of the inter-spike interval is 0.86 and the fano factor 6.15.

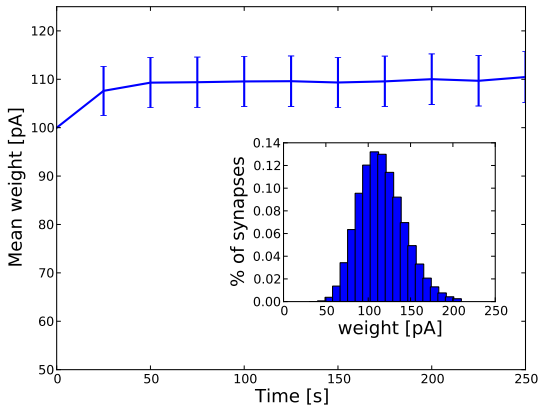

Figure 27: Convergence of the mean excitatory-excitatory weight in a plastic network. The error bars show the standard deviation of the distribution, and the inset represents the final distribution of the weights, after convergence has been reached.

As shown in Figure 27, the convergence to the distribution of weights $\mathcal{N}(\mu_{EE}, \sigma_{EE})$ is rapid, with a variance and a mean that is almost constant over time. Looking closer at the mean weight averaged according to the delays between two cells in Figure 30, we can observe a clear dependence of these delays, and therefore of the distance, on the mean weight.

Control Network We implemented the same STDP rule in the control network, without structure, and compared the results obtained after 250 s of biological time for the system. The control network settled down into a regime close to that of the structured network: mean firing rate of 10.5 Hz, mean ISI CV of 0.82 and fano factor of 7.93. By looking more closely at the convergence of the weights in this control network, we can notice than the convergence is similar, with a convergence to the distributions $\mathcal{N}(\mu_{EE} = 103.8 \text{ pA}, \sigma_{EE} = 24.7 \text{ pA})$ and $\mathcal{N}(\mu_{EI} = 104 \text{ pA}, \sigma_{EI} = 24.8 \text{ pA})$, but we can also notice that the weights do not depend on the delays. The balanced activity is kept, but we lose the topographical structure that could be seen via this influence of distance on the weights. Nevertheless, the macroscopic quantities such as firing rates, average level of correlations, ISI CV are identical in both networks, in line with the idea that some macroscopic invariants may constrain the dynamical properties, as shown in Part III.

15.2.3 Differences between excitatory and inhibitory synapses

In all the results, a slight difference between excitatory to excitatory (EE) and excitatory to inhibitory (EI) synapses is observed. To test the idea that the grid could be responsible for these differences, we plotted in Figure 28 how the way the delay are truncated due to the discretization of the grid introduces a small bias between the two distributions. More precisely, the inhibitory to inhibitory (II) and inhibitory to excitatory (IE) connections have

slightly different distance distribution from the EI, and EE connections, due to the different cell densities on the grid, and the network is not fully homogeneous. The inhibitory grid being less dense than the excitatory one, the discretization impacts the total number of inhibitory synapses. This bias is responsible for the slight difference between EI and EE synapses. Nevertheless, this difference does not drastically affect the results. For an in-depth study of the artefacts that may be induced by grid structure, see Voges et al. (2007). Additional simulations were made, drawing cells with random positions in a square of 1 mm^2 and without the grid, both with this network or the one previously built in Part III, and this difference vanish (data not shown).

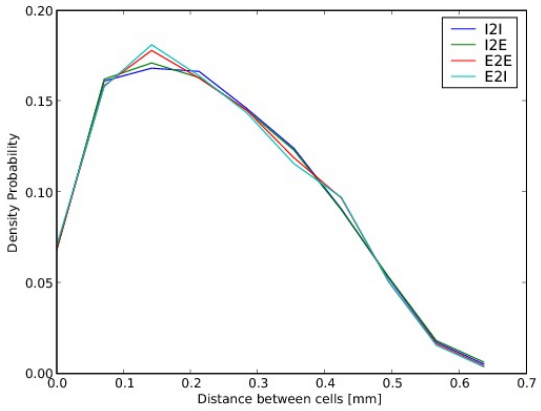

Figure 28: Distribution of the distances, and hence delays, observed in the topographical network, for all types of connections after the grid discretization.

15.3 Results on axonal versus dendritic delays

Next, we modified the delays between the cells to take into account more precisely the axonal and the dendritic parts of the delays, to be closer to biology. The interval in the STDP rule, $\delta t = t_{\text{post}} - t_{\text{pre}}$, may be more precisely written, from a synaptic point of view, as $\delta t = (t_{\text{post}} + t_{\text{den}}^{\text{back}}) - (t_{\text{pre}} + t_{\text{axon}} + t_{\text{syn}})$, where $t_{\text{den}}^{\text{back}}$ represents the time for an action potential that was elicited at time t_{post} at the soma of the post synaptic neuron to back-propagate within the dendrite to the synapse, t_{axon} the time for a pre-synaptic spike emitted at time t_{pre} to be propagated along the axon to the synapse due to the axonal velocity, and t_{syn} the time needed by the synapse for neurotransmitter release and diffusion. (see Figure 29, Left).

For STDP, back propagating action potentials are the best candidates for informing the synapse that a post-synaptic spike has been emitted (Kampa and Stuart, 2006, Golding et al., 2002, Letzkus et al., 2006, Frégnac, 1999). They have been observed *in vitro*, and even though their back-propagation is subject to several conditions, such as depolarization of the dendrite and so on, they could provide a key mechanism for computing the time difference, at the synapse level, which is needed by the STDP framework. To simplify the notation, for the sake of clarity, we merge t_{syn} into t_{axon} in order to deal only with two terms. Then, we have:

$$\delta t = t_{\text{post}} - t_{\text{pre}} + (t_{\text{back}}^{\text{den}} - t_{\text{axon}}) \tag{34}$$

Since the global transmission delay between the two cell somas, d, is also partly axonal and partly dendritic, we can say that if γ is the percentage of the delay which is considered as

being axonal, we have $t_{axon} = \gamma d$ and $t_{den} = (1-\gamma)d$. Under the assumption that the dendritic excitatory post synaptic potential (EPSP) forward propagation time is the same as the action potential back propagation time ($t_{\text{den}} = t_{\text{den}}^{\text{back}}$), we have in the previous equation:

$$\begin{aligned} \delta t &= t_{\text{post}} - t_{\text{pre}} + ((1-\gamma)d - \gamma d) \\ &= t_{\text{post}} - t_{\text{pre}} + (1-2\gamma)d \end{aligned} \tag{35}$$

We can therefore see that introducing an asymmetry between $t_{\text{back}}^{\text{den}}$ and t_{axon} will induce a shift of the STDP rule. When two connected cells fire together $t_{\text{pre}} = t_{\text{post}}$, the way we are modelling the delay and the value of γ will lead either to potentiation if $\gamma < 0.5$ or to depression if $\gamma > 0.5$ (see Figure 29, Right). Since STDP is highly sensitive to high temporal frequencies in the spike pairings, we could hypothetize that this shift will strongly impact the convergence of the weights. In such a topographical network, most of the correlations are localized (see Figure 26), occurring in a rather short temporal window, so if the delay is large, the net effect on the weights could be either a potentiation, for neurons that are close by, or a depression.

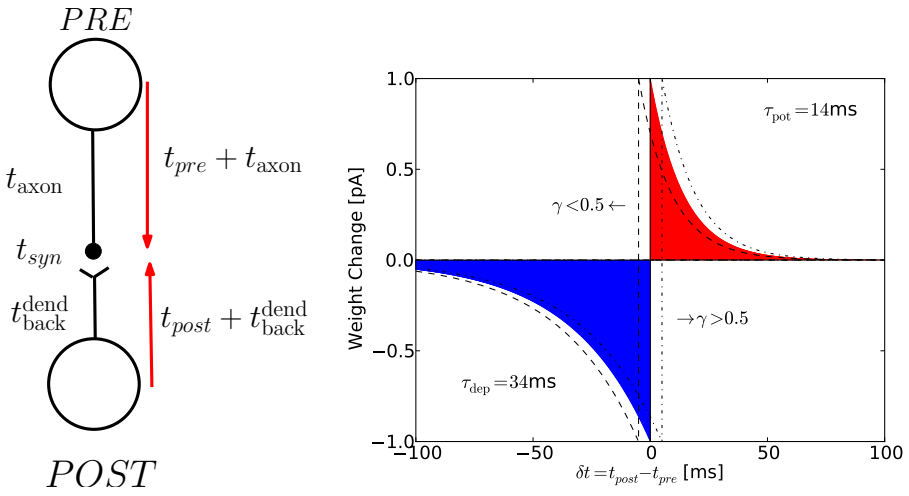

Figure 29: Left: Schematic view of the influence of γ, the proportion of axonal delay in the propagation delay between the pre and post synaptic neuron. Right: Influence of the axonal and dendritic delay in the STDP curve. The variation of γ can shift the curve.

To check how this might affect the behaviour of the network, we launched three simulations: in a first protocol, delay is as before, purely dendritic. This is the default situation with the simulator used, NEST, since the synapses are always hosted by the pre-synaptic neurons. In a second protocol, we split the delays between neurons into two part: 30% of the delay was considered as being axonal, and 70% as being dendritic. We observed on the 2D network how this repartition affected the dynamics and the steady state of the network (Currently, NEST does not allowed to have more than half of the delay to be axonal).

The network settled in a stable state, with a mean ISI CV of 0.82, a mean firing rate of 10.9 Hz, and a fano factor of 7.52. The weights for the convergence are $\mu_{EE} = 104.8$ pA and $\mu_{EI} = 105.8$ pA. Without tweaking the value of λ, the balance between depression and potentiation, the convergence of the macroscopic quantities within the network seems to be rather insensitive to this axonal delay. So the introduction of a partly axonal delay reduces the asymmetry of the total dendritic delays, without altering too much the properties of the network. Nevertheless, as we can see in Figure 30, the dependence of the weights on the distance is drastically reduced compared to a full dendritic delay.

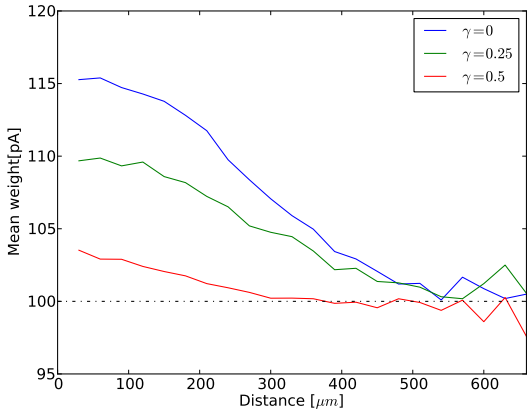

Figure 30: Influence of the delay modelling. Differences in the equilibrium reached by the plastic network according to the percentage of the delays which is considered as being axonal.

In a third simulation, we split all the delays into two equals parts, 50% is considered as being axonal and 50% is considered as being dendritic. The final statistics of the network are again still very similar (firing rate of 10.6 Hz, Fano Factor of 6.7, mean ISI CV of 0.81). The weights for the convergence are $\mu_{EE} = 102.4$ pA and $\mu_{EI} = 103.4$ pA, and as can be seen in Figure 30, this almost suppresses the dependence of the weights on distance. Most of the modification of the weights is due to small differences in δt. The closer δt is to 0, the higher is the amplitude of the modification of the weights with STDP. This shift induced by γ is therefore crucial for promoting the development of structures according to the topology. The prediction of these simulations is that for a simulator that implements delay as entirely axonal, with post-synaptic neuron hosting the synapse, then results shown in Figure 30 should be inverted. Weights for synapses connecting nearby neurons, and therefore correlated regions would be depressed. If the delay is partly axonal and partly dendritic, then when the two neurons fire simultaneously, we have potentiation, but we are reducing its effect by shifting the STDP curve toward negative values.

15.4 Emergence of structures

It was previously reported, in random networks, that even if weight-dependent STDP is compatible with the Asynchronous Irregular regime, no stable structures were able to be maintained. More precisely, Morrison et al. (2007a) have shown that for a given set of synapses, those that are strong at time t will tend to decay after a while, subject to fluctuations within the unimodal distribution of the weights achieved in the equilibrium. The more the network is dense, the more the time of survival of such structures is short, following a power law evolution. The theoretical analysis of recurrent random networks under additive rules was performed in Burkitt et al. (2007), who whowed how weights in this case converge to a stable bimodal distribution. Weights, once they have reached a bound of this bimodal distribution, stay constant, but the activity regime is no longer Asynchronous Irregular.

As in the random case, no stable structures developed spontaneously in the topographical or in the control network, either in a purely dendritic delay scheme, or in a mixed axonal and dendritic delay one. Figure 31 shows the distribution of the standard deviations σ for all the

weights time series, considered after 50s of simulation, when the network has converged to an equilibrium (with $\gamma = 0$). As one can see, the distribution of the temporal variances is broad, and the amplitude of the fluctuations has the same order of magnitude compared to the one of the global weight distribution (with a mean of $\mu \simeq 34$ compared to 26 for the standard deviation of the stable weight distribution when equilibrium has been reached). This result implies that, during the equilibrium, even if the global weight distribution is kept constant, individual weights are always fluctuating within this distribution. This is confirmed with the inset, showing the weight autocorrelation function, averaged over 5000 individual weight trajectories sampled randomly.

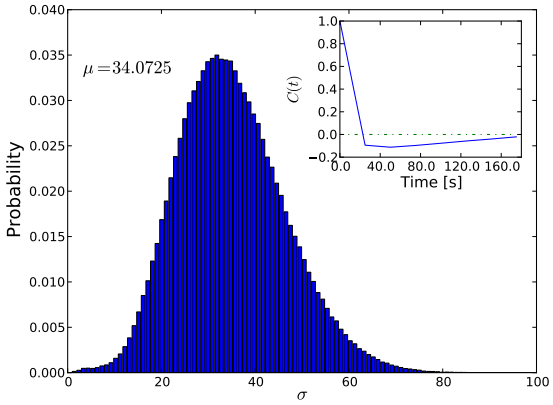

Figure 31: Distribution of the standard deviations σ for all the weights time series, considered after 50s of simulation, when the network has converged to an equilibrium (with $\gamma = 0$). The mean of the distribution is $\mu = 34$. Inset show the weight autocorrelation function, averaged over 5000 individual weight trajectories sampled randomly.

Unsupervised learning achieved by STDP To see how a stimulus can affect the development of the weights, we stimulated the centre of the network with a Poisson process emitting synchronous volleys of spikes at the rate of 5 Hz in order to force a local area in the network to fire synchronously. One example of such stimulation is given in Figure 32. The number of stimulated neurons is set to approximatively 4% of each population (exc/inh). The stimulating source is also connected with a Gaussian probability profile.

The stimulation was applied after convergence of the weights into a stable state over 150 s, and then we observe the relaxation process. During the stimulation, we observe synfired explosions with echo, a phenomenon that has been previously reported in Mehring et al. (2003). As we can see in Figure 33, after 150 s of stimulation, the distribution of the weights is drastically changed for the weights from stimulated cells to non stimulated cells, compared to the distribution of the weights between non-stimulated cells that stays roughly constant. On the top, one can see the average afferent synaptic weights per neuron. Projections within the centre of the stimulated area are modified during the stimulations, but return quickly to their equilibrium values as soon as the stimulation is stopped. The average firing rate of the whole network is only slightly affected by the external stimulation and this wash-out of the memory.

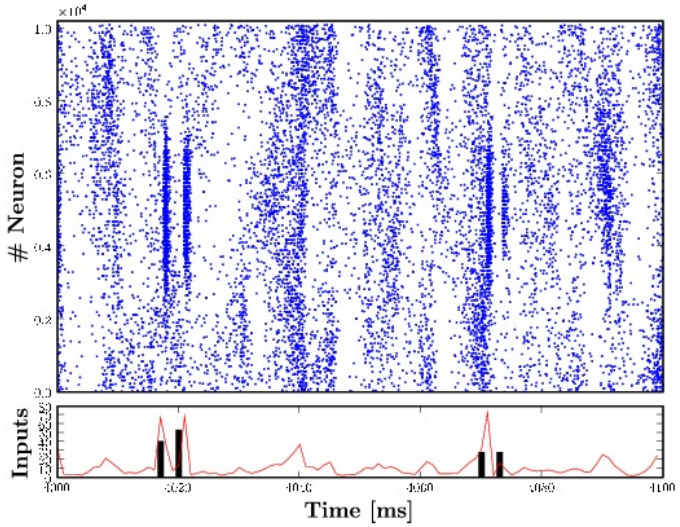

Figure 32: Stimulation of a localized area within the network by a synchonous volley of spikes. Top: Raster plots of the neurons, labelled according to their id in the 2D network. Bottom: Average firing rate in the network, with black line indicating the time of the external stimulation

16 Discussion

Over-sensitivity to ongoing activity Adding a topographical structure to the network does not lead to the development of stable connectivity with weight-dependent rules, as was already reported in random networks (Morrison et al., 2007a). The ongoing bombardment, inherent to the balanced network regime, is constantly impacting synapses and therefore, the weights are constantly fluctuating. In this recurrent network, ongoing Poisson activity disrupts any kind of structure that can be learned by the system, as long as weights stay in a unimodal distribution, and do not push the system away from the attractor of its ongoing activity. During stimulation with external inputs, the statistics of the weights can be changed, but as soon as the stimulation is relaxed, neurons tends to immediately re-learn the incoming "noise" they are subject to, and go back to their equilibrium. This is a severe problem for memory retention and emergence of structures in recurrent networks. If weight increments are modelled as additive, then structures develop (data not shown, but see analytical work of Burkitt et al. (2007) and Siri et al. (2008)), but the Asynchronous Irregular state is not maintained. The network forms into co active clusters and progressively transitions into a Synchronous Regular regime, ending in a stable crystalline state that can hardly be modified. This sensitivity of the STDP rule to its implementation, and also for example to the way delays are modelled between pre- and post-synaptic neurons, makes large-scale results hard to interpret.

STDP, since its discovery, has proven to be successful in reproducing generic results, such as receptive field development (Abbott, 2003, Song and Abbott, 2001), orientation preference maps (Wenisch et al., 2005), or learning input/output relations such as coordinate transformation (Davison and Frégnac, 2006), synfire chain formation (Hosaka et al., 2008) and much more. Nevertheless, almost all those studies ignored the recurrent connections and the ongoing activity, and for storing dynamical patterns and sculpting the spatio-temporal profile of the correlations within neuronal networks, the rule seems to be inappropriate. The question of stability over long time scales is often disregarded, even if this is a crucial concern if the

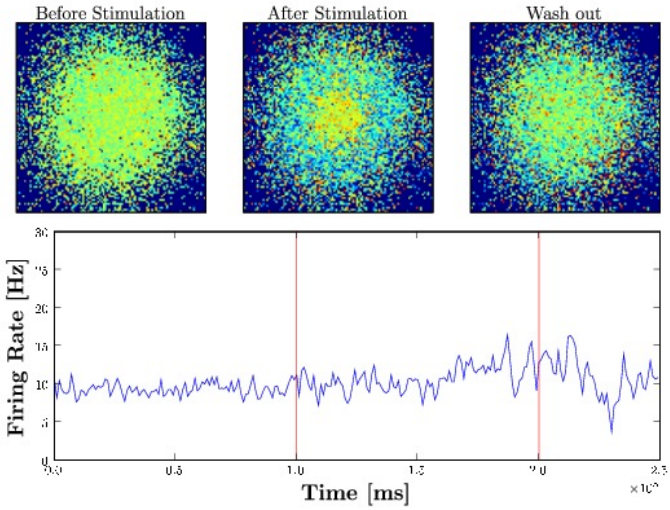

Figure 33: Stimulation of a localized area within the network by a synchronous volley of spikes. Top: Mean afferent synaptic weights per neuron, just before, after 50 s of stimulation, and 100 s after the stimulation. Bottom, average firing rate over the whole network during the whole simulation. The red lines show when stimulation starts and stops.

system is supposed to be plastic all the time. Memory retention, the stability of the irregular regime and of the macroscopic quantities is problematic with STDP: as shown in Morrison et al. (2007a), explosions of the network and convergence to a Synchronous Regular states can happen if no homeostasis is included to prevent divergences. It is well known that neurons *in vitro* have the intrinsic property of adapting their synaptic weights according to the background activity. This homeostatic scaling (Turrigiano and Nelson, 2004) could counter balance the effect of Hebbian learning, and this could be included in STDP rules.

STDP and BCM The missing link between STDP and BCM is also problematic: since the BCM theory is a serious candidate *in vivo* to explain ocular dominance or receptive field development, a general theory of plasticity should be able to conciliate both mechanisms. As already said, the theoretical shape of the BCM curve (see Figure 21) can be retrieved if nearest-neighbour spike interactions are considered for the STDP model (Izhikevich and Desai, 2003). Nevertheless, no biological evidence for this has been reported so far. In addition, even if the general shape of the curve is similar, the sliding threshold θ, acting as an homeostatic process in the BCM theory, regulating the balance between potentiation and depression and supposed to vary according to the post-synaptic firing rate history, can not be retrieved. This homeostatic balance, if incorporated into a more general model of plasticity, may help to enhance the stability of the system.

Modifications of the STDP rule In the context of supervised learning, some authors (Legenstein et al., 2008, Izhikevich, 2007) have started to propose that STDP could be used combined with an external reward signal. Changes can be accumulated at the synapse level, and validated only when particular feedback is sent to the system. Such an idea has the advantage that it can turn on or off the rule in some particular time window, during which the system is able to learn or not. This is a complementary approach to the homeostatic one. An other

argument in favour of the fact that STDP may be over-sensitive to the ongoing activity is that, from an energetic point of view, it may be costly for the system to constantly modify the synaptic weight, each time a pre/post pairing is performed. Constraining the time windows during which incoming patterns and/or statistics should be stored may save useless energy. In that spirit, a solution could be to consider that synapses do not vary under a continuum, but are more like binary switches that could take discrete values. This is the idea followed by Fusi and Abbott (2007), and supported by recent biological evidence (Montgomery and Madison, 2004). Transitions between such stable discrete states would be triggered by activity, and make the synapse more robust to noise. An other possibility is a recently found mechanism known as "synaptic tagging" (Young et al., 2006). Without going into the biological details, LTP or LTD can be pre-activated at the synapse level (the synapse is "tagged" during an early phase), and changes are stored only if confirmed during a late phase (Clopath et al., 2008). The distinction between early and late phase is again a kind of mechanism that could help in making the synapse more robust to noise. Nevertheless, these theories are more conceived based on an Hebbian framework, i.e. with the idea that exact spike times are not the crucial part of the information, which is more based on the average firing rates.

As we will see in the following Part, and as already suggested by this Discussion, the concept of STDP needs to be revised in a more general manner. By examining the biological data gathered since it has been discovered, one can appreciate how general the phenomenon is, and how the simple double-exponential shape used in this Part reduces its complexity (Lisman and Spruston, 2005). Instead of trying to play with the network's structure, and/or the parameters of the external stimulation, a better understanding of spike timing dependent plasticity need to be gained. As we will see in the following Part, a lot of biological evidence suggests that STDP, in this naive and basic form, is not able to reproduce all the biological results observed. Even if suited for simple problems, its robustness in the face of ongoing activity is severely impaired, with weights fluctuating without being hard to stabilize. If we consider that memory is indeed stored in the precise connections strengths between neurons, and if those connections are supposed to vary by plasticity in an analogue manner, then weight-dependent rules, more generic and plausible, are problematic from a memory point of view.

Part V

Towards a new STDP rule

Summary

17 A more complex view of STDP **141**
17.1 STDP as an epiphenomenon ? 141
17.2 Metaplasticity 145

18 STDP and metaplasticity **146**
18.1 Main results 146
18.2 Supplementary materials 162
18.2.1 Materials and Methods 162
18.2.2 Deriving the BCM rule from the STDP 165

19 Discussion **169**

17 A more complex view of STDP

In this Part, I will try to explain why a new STDP rule is necessary, to circumvent several incoherencies found when gathering biological data on associative plasticity, and how this rule can be mapped, in a metaplasticity framework, to the BCM theory while basic STDP can not (Izhikevich and Desai, 2003). After having discussed several pieces of evidence from biology, showing the limitations of STDP, we will explain in detail a plausible biological mechanism that could explain weight changes in a synapse. The rule, incorporating some homeostatic constraints in the framework of metaplasticity, will establish a link between different theoretical schema of plasticity (STDP, BCM, ABS) with the help of some biophysically realistic processes. Non-linear interactions between the long term potentation (LTP) and the long term depression (LTD) molecular transcription pathways will naturally emerge, and we will study how such a rule behaves in feed-forward and in recurrent networks, similar to those of the previous Part.

17.1 STDP as an epiphenomenon ?

It has recently been discovered that STDP is a much more complex phenomenon than it may appears at a simple glance. In this Part, we will review some of the key phenomena that have been reported on spike timing dependent plasticity, and that are not captured by the classical view of STDP defined by Eq. 28 in the previous Part.

Heterogeneity of the rule First of all, the shape of the rule seems to be dependent on the neuronal structures in which recordings were made, or on the cell types considered. As shown in Figure 34, the same protocol (repeated pre-post pairings at a fixed frequency $f_{\text{pairing}} = 1$ Hz), in different species and/or anatomical areas or neuron types can lead to a huge variety of rules. If STDP was the result of a generic optimization principle, it should take this heterogeneity into account, or be linked with a functional explanation correlated with the structure. Bell et al. (1997), for example, discovered in the electro sensory lobe of the electric fish an anti-Hebbian STDP rule, that may be useful from a functional point of view in achieving the perceptual filtering of expected changes in sensory input produced by motor action (here, electric discharge). The electric fish builds a sensory representation of its environment by sending an electrical discharge and establishing a difference between an efferent copy and the electrical signals gathered after interaction with its surroundings. Therefore, from a coding point of view, it encodes only the difference between what is expected and what is received, and for that purpose, an anti-Hebbian rule is well-suited. This coding scheme is in agreement with the predictive coding theory (Rao and Ballard, 1999).

Induction and dependence on the pairing frequency Spike timing dependency plasticity is usually observed in cortical neurons *in vitro* with a stereotyped protocol: approximatively 60 pairings are performed at a frequency of $f_{\text{pairing}} = 1$ Hz, with $\delta t = t_{\text{post}} - t_{\text{pre}}$ between pre- and post-synaptic spikes which is varied. The fact that 60 pairings need to be used is because plasticity needs to be induced. In Froemke et al. (2006), one can see that the amount of plasticity which is triggered does not depend linearly on the number of pairings. After an induction phase, it changes non-linearly up to a saturation plateau, around 60-100 pairings, the classical values considered in the protocols. This is something which is not taken into account

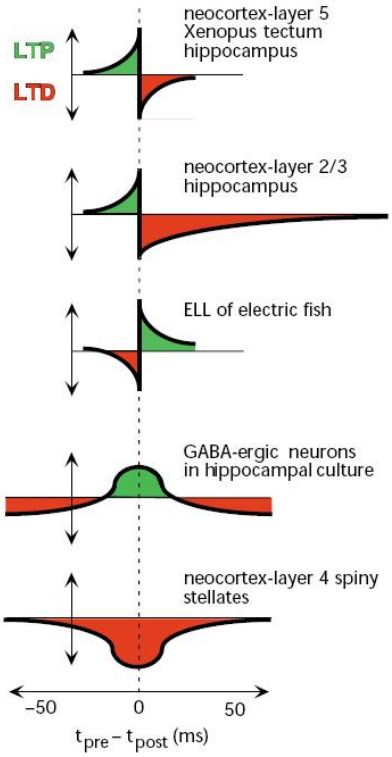

Figure 34: Taken from Abbott and Nelson (2000). Illustration of the variability of spike timing dependent plasticity rules. According to the precise time difference δt between a pre- and a post-synaptic spike, the synaptic weight can be either depressed or potentiated in many ways, according to species and/or area.

in models, but performing fewer pairings may not trigger any plastic changes, and performing more pairings have no further effect on the synapse. 60 pairings seems to be a compromise to reach the saturation in the plastic changes triggered. In addition, it has been also shown, in Sjöström et al. (2001), that if the frequency f_{pairing} of the pairing is changed, then the STDP curve shown in previous Part is not valid any more. Depression is only visible for low frequency pairings, when 60 pre-post pairings are performed with $\delta t < 0$ and $f_{\text{pairing}} < 20$ Hz. For $f_{\text{pairing}} > 20$ Hz, the synapse undergoes only potentiation, whatever δt. This crucial point led some authors to consider new models of STDP, taking not only the pairwise interactions between pre and post, but also higher order interactions such as triplet (pre-post-pre, or post-pre-post). This is the triplet model, developed in Pfister and Gerstner (2006), which is able to reproduce the observation that the shape of the STDP rule depends on the pairing frequency.

One other point, not considered so far in models of plasticity, is the fact that STDP changes are not instantaneous. Protocols are performed, then the final value of the weight is recorded up to 30 minutes later, divided by the number of pairings, and the final value is considered as being the instantaneous weight modification after each pairing (for an additive model of STDP). By examining the weight evolution curves found in the STDP literature (Froemke and Dan, 2002, Froemke et al., 2006, Bi and Poo, 1998, Sjöström et al., 2001) (see for example Figure 35), it can be seen that the weights seem to evolve continuously, at least for depression, after pairing. Plasticity is more a synaptic modification which is triggered by a transient stimulation, and slowly evolves toward a new equilibrium. Understanding the biological mechanisms responsible for those changes, at the molecular level, is necessary to gain an insight on how this

evolution would behave *in vivo*, when synapses are constantly impacted by random spikes.

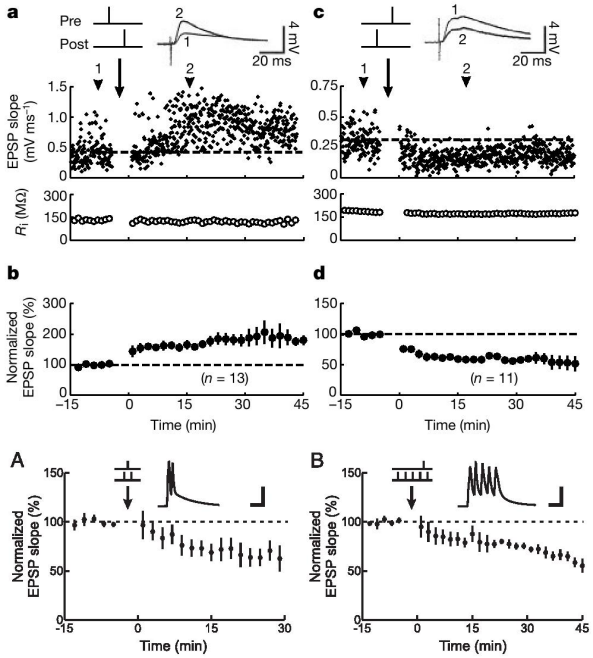

Figure 35: Taken from Froemke and Dan (2002) (4 upper panels) and Froemke et al. (2006) (2 lower panels). Evolution of the synaptic weights after plasticity induction. Details of the protocols can be found in the articles. An interesting point is the evolution, as a function of time, of the EPSP slope after induction. Variations can be observed during the whole window of observations, up to 30 minutes after the induction.

Asymmetry in the LTP/LTD molecular cascades Recent evidence has shown that LTP and LTD are not mediated by the same mechanisms, and are therefore not symmetric. Wang et al. (2005) induced plasticity by STDP protocols using triplets of spikes, such as pre-post-pre or post-pre-post. If δt_1 is the difference between the two first spike times, and δt_2 the difference between the two last one, then one would expect with STDP that if $\delta t_1 = \delta t_2$, LTP and LTD should cancel each other and the net result, at the synapse level, should be almost no modifications, or with the time constant τ_{dep} for depression being usually larger than for potentiation, one might expect LTD to dominate. Nevertheless, authors found a clear preference for LTP. This difference can be explained by the fact that LTP and LTD are not mediated by the same mechanisms. To detail the LTP mechanism, notice that at the synapse level, neurotransmitters (glutamate) can bind to the *N*-methyl *D*-aspartate (NMDA) receptors. If at the same time the post-synaptic cell is depolarized, the magnesium block is relieved and the NMDA channel opens: free calcium enters the cell, inducing a molecular cascade that phosphorylates the Ca^{2+}/calmodulin-dependent protein kinase (CaMKII), which in turn acts on the AMPA receptors activation and/or density. This match between the depolarization of the post-synaptic cell and the incoming calcium influx is supposed to be achieved, in the STDP framework, with the action potential that back-propagates from the soma to the synapse within the dendrite, establishing a temporal window during which LTP can be induced. The modification of the AMPA receptors sensitivity or density is responsible for LTP, since this is the channel responsible for the entrance of sodium influx, at the origin of the excitatory post synaptic potentials (see Part I). Protein kinases, and especially CaMKII are key molecules in

LTP induction, whereas phosphatase, such as calcineurin, are used in LTD induction (Lisman et al., 2002, Malinow and Malenka, 2002), through particular L-type Calcium channels, used by LTD but not by LTP (Bi and Poo, 1998). Distinct mechanisms are taking place at the synapse level for LTP and LTD, and such evidence suggest that they may interact one with each other in a highly non linear fashion.

STDP along the dendritic tree Several *in vitro* studies (Kampa and Stuart, 2006, Letzkus et al., 2006, Froemke et al., 2005) showed that the shape of the STDP curve obtained by pre-post pairings in cortical pyramidal neurons depends of the position of the synapse along the dendritic tree. They found evidence for a plasticity gradient from proximal to distal synapses: synapses close to the soma are regulated by Hebbian rules, while those on the distal part of the dendritic tuft are more anti-Hebbian. Since plasticity is regulated by calcium dynamics and NMDA receptors, the attenuation of the back-propagating spikes within the dendrite and the dendritic spikes in pyramidal neuron (Larkum et al., 2001) may be valide candidates to explain this phenomenon. It is interesting to observe that conceptually, this plasticity gradient could have a functional role in primary sensory cortical areas. We consider a cortical pyramidal neuron in layer 5, integrating information from the cortex and sending axons back to the thalamus. Since the thalamo-cortical synapses are established mainly in layer 4, close to its soma, Hebbian learning rules with STDP will consolidate feed-forward signals coming directly from the thalamus in parallel with the recurrent cortico-cortical synapses originating from layers 4/5. In the meanwhile, this pyramidal neuron forms distal synapses in superficial layers 2/3, so will also integrate, with anti-Hebbian rules, previously processed information: indeed, since pyramidal neurons of layer 4 project to layers 2/3, information is delayed in these layers, compared to incoming inputs in layer 4. So this pyramidal neuron in layer 5 may try, with Hebbian rule close to its soma (layer 4) and anti-Hebbian in its distal dendritic tuft (layers 2/3), to separate the new information from an expected background, based on past electrical activity reverberated. The gradient of plasticity and the small delay between information impinging layer 4 and the previously received and integrated information arriving in layers 2/3 (the background) can provide a theoretical substrate to correlate coherent information while, at the same time decorrelating it from the background. It could be an algorithmic way to build saliency maps, amplifying the difference between new incoming inputs and the background, and recent models are now trying to take advantage of these multiple forms of plasticity along the dendritic tree (Kaneki et al., 2009).

Sensitivity to the post-synaptic neuron As already mentioned in the previous Part, STDP is often not considered for inhibitory synapses. However, several evidences pointed out that depending on the celltype of the post-synaptic neuron, the shape of the rule could be different (see (Caporale and Dan, 2008, Davison and Frégnac, 2006)). For excitatory to inhibitory connections, the rule is inverted (Bell et al., 1997, Tzounopoulos et al., 2004). If $\delta t = t_{\text{post}} - t_{\text{pre}}$ is positive, then synapse is weakned, and is $\delta t < 0$, it is strenghtned. Nevertheless, since the post-synaptic neuron is inhibitory, the functional consequences, at the network level, are the same. Rule is also variable at the GABAergic synapses (Haas et al., 2006). This sensitivity to the post-synaptic celltype can be due to retrograde signals sent from the post-synaptic cell. It has indeed been shown that endocannabinoids (eCB) are important molecules in both short- and long-term depression of many synapses in the central nervous system (Chevaleyre et al., 2006). Hashimotodani et al. (2007) showed how an increase of the calcium concentration can promote the generation of eCB by the post-synaptic neuron, captured by the pre-synaptic

neuron through CB1 endocannabinoid receptors. These receptors are known to have an impact on the LTD molecular pathway, presumably by inhibiting presynaptic transmitter release (Sjöström et al., 2003).

17.2 Metaplasticity

STDP itself should be seen as a dynamical process that may change and evolve according to the state of the system, and this plasticity of the synaptic plasticity itself is referred to as metaplasticity (Abraham and Bear, 1996, Bear, 1995). A lot of evidence supports the idea that STDP should adapt and depend on other variables, such as, for example, the post-synaptic membrane potential (Ngezahayo et al., 2000), neuromodulation (Zhang et al., 2009), amount and kinetics of the calcium that may enter into the synapse (Yang et al., 1999). The idea of metaplasticity is that the synapse's previous history of activity influences its current plasticity.

The interactions over long time scales between the activity of the synapse and its plasticity can be observed with so-called priming experiments. Pre-activation of the synapse, either with LTP or LTD, has an impact on the reactivation of these LTP or LTD pathways for up to an hour. Experimental studies, especially in the Schäffer collateral pathways (SC) in hippocampal slices, have shown some "priming" effects: Huang et al. (1992) show that pre-activation, with weak stimulation of some collateral pathways projecting to a post-synaptic neuron can lead to potentiation. Reactivation of the same pathways, up to one hour later, led to a smaller potentiation than in a control experiment where the weak tetanic pre-stimulation was not applied. This means that activity dependent changes can trigger mechanisms in the synapses that last over long time scales and influence plasticity over time. Similar priming can be observed with depression (Mockett et al., 2002). This is closely related to all the "tagging" experiments of Frey and Morris (1997), and these results lead to the idea of an early and late phase for both changes (see Clopath et al. (2008) for a modelling study) to distinguish between plasticity induction and consolidation. STDP, in its simple form, should be explained in a framework that takes these results into account. More details on these protocols and on their results will be developed in the following article, but the idea is that extra time constants may be missing in the simple STDP framework. If the 20 ms time windows of the double exponential shape (τ_{pot} and τ_{dep}, see Figure 23) could be linked with the time course of the calcium influx within the synapse, biological internal mechanisms validating or consolidating the changes may have extra time constants that are much longer, and they should certainly be taken into account in order to understand the role of plasticity. The internal phosphorylation of the CaMKII molecule (LTP), or the evolution of the calcineurin concentration (LTD) may filter fast and transient changes induced by the instantaneous spike pairings. The synaptic changes at a given time t at the synapse are influenced by an additional variable which is the past activity dependent plasticity of the synapse.

Sensitivity to neuromodulators Associative plasticity is very sensitive to the neuromodulation, such as the endocannabinoids impact on the LTD, previously shown. Indeed, STDP is sensitive to the neuromodulators that are released by numerous varicosities (do not require a postsynaptic cleft) and diffuse in the extracellular medium. For instance, catecholaminergic ascending axons from subcortical origin run along the cortex from the frontal to occipital cortex and one axon can influence at the same time plasticity processes at synapses located in different cortical areas. It has been recently shown in Zhang et al. (2009) that in hippocampal synapses, STDP can be strongly modulated by dopamine. If dopamine is present, then the

depression part of the STDP curve vanishes and the rule is turned into a simple covariance rule, almost symmetric, where both pre-post and post-pre pairings lead to potentiation. This regulation of plasticity by self- or externally generated neuromodulators affects the dynamical aspect of plasticity and contributes to the metaplasticity.

Towards new models of STDP More and more models are now trying to understand STDP in the more general framework of metaplasticity. An attempt was made by Froemke and Dan (2002), with the idea to weights the spikes within the nearest-neighbours interaction scheme of the STDP pairings, to promote the influence of the closest by adding a suppression term. The triplet model of STDP (Pfister and Gerstner, 2006), which reproduces the asymmetry in the LTP and LTD pathways observed by Wang et al. (2005) considered not only pairs of spikes in the framework of STDP, but also triplet interaction. Interactions between spikes are not only made between pre-post and post-pre pairs, but also with all the possible triplets (pre-post-pre and post-pre-post). The model comes up with two time constants, either for potentiation or depression (so four in total): one responsible for the short STDP time window (the one observed in the biological data of Bi and Poo (1998)), and a longer one to capture the non-linearities imposed by the triplets. The triplet rule is able to reproduce the results from Sjöström et al. (2001) and the dependence on the pairing frequency f_{pairing}, but also to establish a link with the BCM framework, under several constraints. The triplet rule can promote rate-based competition, a feature which is missing with classical STDP. The BCM or ABS rules, introduced in Part IV, are metaplastic: the modification of the synapse depend on the sliding thresholds, which depend on the history of the post-synaptic spiking activity. The rules adapt smoothly as a function of the level of activity. The link between STDP and BCM theory is important to conciliate this two views of associative learning.

Similarly, Clopath et al. (2010) designed a rule sensitive to the post-synaptic membrane potential. This abstract model is also able to tackle the issue of frequency dependence, and also to explain some *in vitro* results with clamped membrane potentials. It also establishes a link with the BCM theory and has been shown to reproduce the development of receptive fields in a small recurrent network, as in the visual cortex. Both models are phenomenological, i.e. their internal quantities do not intend to be biological: they only focus on the final behaviour, and do not pretend to give a biophysical explanation of STDP. In the same spirit, the work of Savin et al. (2010) shows that STDP, when combined with internal homeostatic constraints, can perform independent components analysis and be used in blind source separation problems. This choice of implementation allows to keep a rule simple enough to be simulated in networks of neurons, nevertheless, none of them has been applied to large-scale networks of neurons.

In contrast, more bio-realistic models, based mainly on calcium dynamics, aim at explaining STDP in terms of biochemical signalling within the membrane. This is the case for example in Lisman (1994), Shouval et al. (2002), Graupner and Brunel (2007), Rubin et al. (2005), Badoual et al. (2006) and Saudargiene et al. (2005). All those biophysical models rely either on the Ca^{2+}/CAMKII phosphorylation and bistability, on calcium concentration, or on the dynamics of the AMPA receptors. Drawback of such models is that their complexity prevent them from being implemented in large networks.

18 STDP and metaplasticity

18.1 Main results

The study presented in the following draft is a theoretical attempt to better understand how stable synaptic modifications could occur in recurrent random networks. It provides a metaplastic rule of plasticity based on a putative biophysical mechanism based able to reproduce some classical results obtained in literature and establishing a link between STDP and the BCM theory. The rule is tested in feed-forward and recurrent networks of integrate-and-fire neurons. I designed the experiments and the paradigm in collaboration with S. El Boustani, and we worked together equally on the project. Materials & Methods are described as supplementary materials.

Synaptic learning in stochastic network states

Sami El Boustani [1], Pierre Yger [1], Yves Frégnac[CA] and Alain Destexhe[CA]

Unité de Neurosciences, Information et Complexité (UNIC), UPR CNRS 3293,
1 Avenue de la Terrasse,
91198 Gif-sur-Yvette, France.

Abstract

The mammalian cerebral cortex is characterized by irregular spontaneous activity, but how relevant information is processed and learned in such noisy states remains unknown. The usual plasticity rules are sensitive to spontaneous activity and therefore cannot operate in realistic network states. Here, we present a new class of learning rules which is based on metaplasticity. Synaptic metaplasticity can solve the problem of stable learning in noisy states, and at the same time, it naturally encompasses other known types of learning rule, placing them into a coherent framework. Finally, the metaplastic learning rule is shown to be consistent with a number of known molecular pathways, which leads to experimentally testable predictions.

Introduction

Various forms of synaptic plasticity have been discovered and characterized experimentally as well as theoretically. An identified type of associative plasticity is based on the relative spike timing between pre- and post-synaptic cells (Bell et al., 1997, Markram et al., 1997, Bi and Poo, 1998). This so-called spike timing dependent plasticity (STDP) has been shown to produce a form of long-term potentiation (LTP) and depression (LTD) in an asymmetric way depending on the temporal order of pre- and post-synaptic cell firing. Models have been proposed to capture such plasticity mechanisms and have been very successful in explaining receptive field emergence, synaptic competition and stability (van Rossum et al., 2000, Song et al., 2000, Song and Abbott, 2001, Gütig et al., 2003, Wenisch et al., 2005). However, previous STDP models suffer from an important limitation, their sensitivity to spontaneous activity. Because cerebral cortex is characterized by irregular spontaneous activity (Softky and Koch, 1993, Shadlen and Newsome, 1998), classical STDP models cannot be used to learn in the context of such stochastic-like activity states.

This problem was identified in computational models of realistic activity states (Morrison et al., 2007, Billings and van Rossum, 2009). For appropriate parameters, balanced networks of integrate-and-fire neurons can display asynchronous irregular (AI) states (Brunel, 2000, Vogels and Abbott, 2005, Kumar et al., 2008, Marre et al., 2009), which produce spike discharge patterns similar to those in awake animals (El Boustani et al., 2007). If endowed with a weight-dependent STDP rule, such networks can sustain AI regimes with stable synaptic weight distributions (Morrison et al., 2007). In this case, however, individual synaptic weights are subject to strong fluctuations due to the ongoing spiking activity. Consequently, changes in the weight distributions are always transient because synaptic weights relax to their stable stationary distribution. Due to this "catastrophic forgetting" effect, memory retention is impossible in these models (Morrison et al., 2007, Billings and van Rossum, 2009).

We investigate here a natural way to circumvent this problem through metaplasticity. Since early work on LTP and LTD, the need for higher-order regulation of plasticity rules has been evident. This process, termed metaplasticity (Deisseroth et al., 1995, Abraham and Bear, 1996), has been shown to take place through various molecular pathways and to be critical for the stabilization of learned statistics (see Abraham (2008) and references therein). In its more specific definition, metaplasticity is a change in the plasticity rule that depends on the past history of the synapse. One major role of metaplasticity is to promote in a conditional way (dependent on the prior history of the network) persistent changes in synaptic weights during learning. A cascade-based synaptic model that displays metaplasticity has been introduced previously (Fusi et al., 2005). However, this model did not incorporate STDP. Although knowledge on the biomolecular processes underlying STDP has tremendously progressed, no simple or biophysical model has been able to display at the same time these important features.

We suggest a new class of model of synaptic plasticity that includes metaplasticity based on a number of experimental observations. First, important non-linear interactions were observed if spike pairings occur closely in time (Sjöström et al., 2001, Wang et al., 2005). Second, experiments were conducted by including a priming protocol in classical LTP/LTD experiments with high-frequency and low-frequency stimulation (Huang et al., 1992, Christie and Abraham, 1992, Wang et al., 1998, Mockett et al., 2002), showing how pre-activation of LTP or LTD mechanisms could impact their re-activation. Based on these results, it appears that plastic processes occurring through the same molecular pathways exhibit regulation on three distinct time scales that account respectively for the nature of synaptic weight changes (STDP rule), the synaptic competition (short-term interactions) and metaplasticity (long-term interactions). We propose

here a model in which the functional diversity of the synaptic effects across these times scales naturally emerge. This metaplastic STDP (mSTDP) model is based on a putative biomolecular interaction between kinases and phosphatases acting on the synaptic weights, and is inspired by experimental data on the effect of priming stimulations on LTP and LTD induction in the hippocampus. The mSTDP model is shown to produce metaplasticity and to reproduce non-linear interactions observed in STDP experiments, thus making a link between experimental results which seemed difficult to integrate in a common framework.

Results

The mSTDP model is based on electrophysiological *in vitro* results reported previously (Wang et al., 2005) showing evidence that LTP and LTD are elicited by calcium influx through dedicated channels, NMDA and L-type respectively. For LTP, and for every pre-post pairing, calcium enters the cell through NMDA channels and binds to calmodulin. We consider that the STDP window for LTP is defined by an exponential curve, which follows the kinetics of the amount of Ca^{2+}/calmodulin that will be formed at the time of the post-synaptic spike. This variable will be labelled x_{LTP}. The action of this Ca^{2+}/calmodulin protein is two-fold. On a very short time-scale, it will activate a kinase K that will phosphorylate AMPA receptors thus increasing the excitatory synaptic conductance, proving a straightforward biophysical substrate of LTP. As a second stage, this protein will activate another pathway resulting in a non-linear negative feedback on the kinase activation, occurring on a slower time scale. This additional molecular pathway is critical in this model and has not been taken into account in previous models of calcium-dependent kinase kinetics (Lisman, 1989, 1994, Lisman et al., 2002, Graupner and Brunel, 2007). The slow and fast actions of Ca^{2+}/calmodulin are justified by the transient higher threshold for LTP induction following a priming stimulation of the test pathway (Huang et al., 1992).

In the model, these two antagonistic actions will be implemented through two different operators. The rapid Ca^{2+}/calmodulin action is modelled as an instantaneous update of synaptic weight given by the value of x_{LTP} at the time of the post-synaptic spike. We chose a linear relationship for the sake of simplicity, but a monotonic non-linear function involving detailed biophysical processes could be used as well. The long lasting negative regulation will be modelled as an average of $\exp(x_{\mathrm{LTP}})$ over the past history. This function accounts for the strong non-linear increase in the LTP threshold when weak tetanic priming stimulations are used prior to a strong tetanic induction (Huang et al., 1992). The net effect on the synaptic weight is proportional to the kinase concentration thus resulting in the rectified difference between these two terms (see Supplementary Material). This LTP mechanism is depicted in Fig. 1A and obeys the following equation for instantaneous weight changes: induced by LTP δw_{LTP}

$$\delta w_{\mathrm{LTP}}(t) = [\sum_{t_{\mathrm{post}}} x_{\mathrm{LTP}}(t)\delta(t - t_{\mathrm{post}}) - \frac{\alpha_{\mathrm{LTP}}}{T}\int_{-\infty}^{t} d\tau e^{(\tau-t)/T} e^{\beta x_{\mathrm{LTP}}(\tau)}]_{+} \quad (1)$$

where α_{LTP} measures the magnitude of the negative feedback, T is the slow time constant and β dictates the increased rate of suppression feedback.

Regarding LTD, it has been shown that priming stimulation can facilitate LTD induction with low frequency stimulation (Christie and Abraham, 1992, Wang et al., 1998, Mockett et al., 2002). Although the direct action of LTD is assumed to be conveyed by calcium influx through dedicated L-type channels (Graef et al., 1999, Wang et al., 2005), the priming effect has been shown to occur through NMDA channels and to covary with the metaplasticity observed for LTP

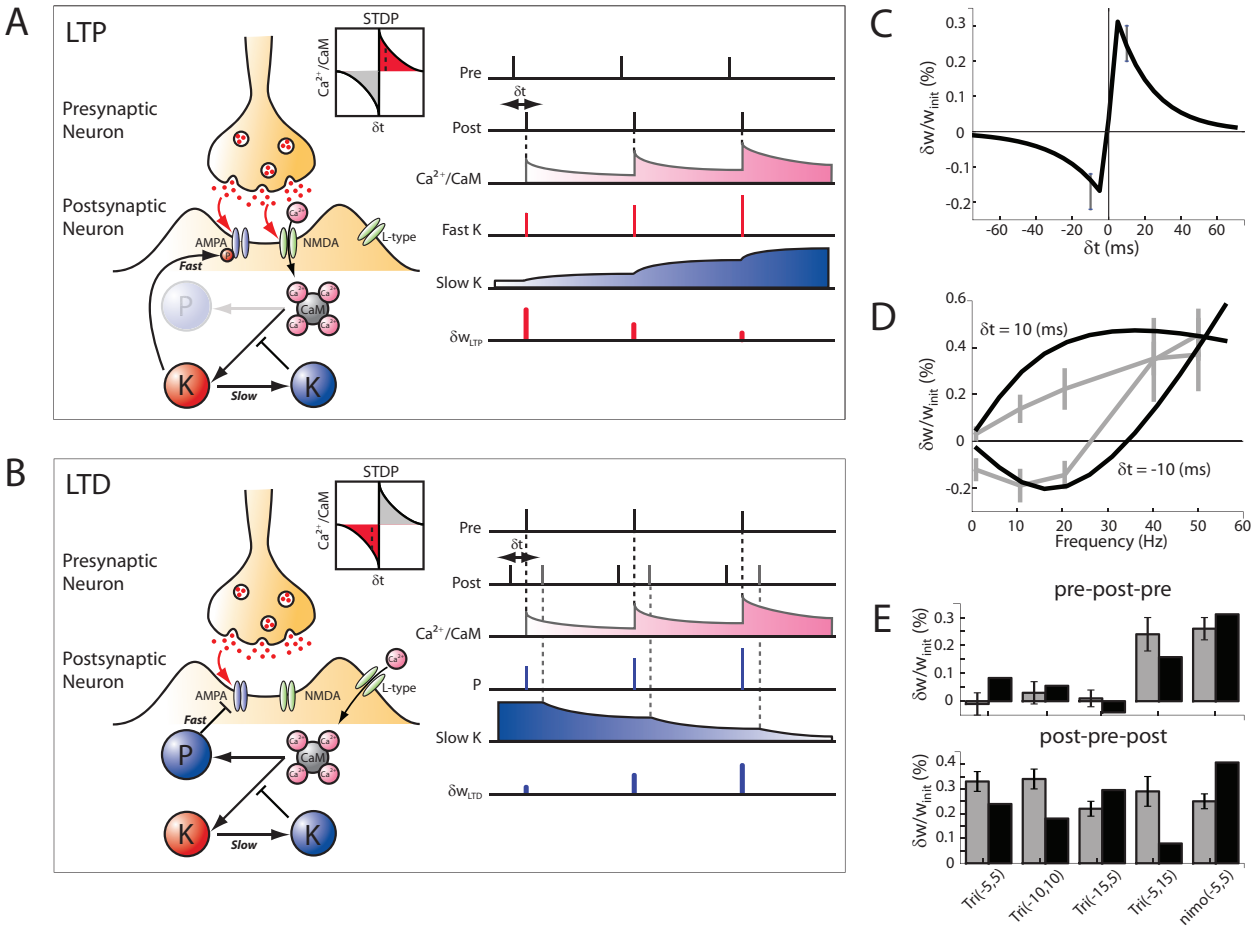

Figure 1: A model for spike-timing dependent plasticity with metaplasticity (mSTDP). A. Mechanism of long-term potentiation with a fast potentiation action and slow negative feedback: a pairing protocol is depicted on the right with the time evolution of each variable. The slowly increasing threshold progressively decreases the net effect on the synaptic weight. The pairing time difference δt is relative to the pre-synaptic spike time $\delta t = t_{\text{post}} - t_{\text{pre}}$. B. Mechanism of long-term depression with a fast depression action and facilitation produced by LTP negative feedback: on the right, a triplet protocol illustrates how post-pre pairings can increase the net effect of LTD on the synaptic weight. C. Typical STDP relation obtained with a simulated pairing protocol. Gray error-bars in the graph are taken from Wang et al. (2005) recorded in hippocampal neurons for δt equals 10 ms and -10 ms respectively. D. LTP and LTD dependency on the pairing frequency for 10 ms and -10 ms respectively. In gray, data obtained for visual cortex neurons by Sjöström et al. (2001) for the same protocol. E. LTP and LTD for triplet (denoted Tri) protocols compared to experimental data obtained by Wang et al. (2005) (upper panel: pre-post-pre; lower panel: post-pre-post). The shaded columns correspond to the control recordings for each protocol in the abscissa (with time intervals between each spike). The filled columns correspond to the model predictions. The last set of columns represents the condition where L-type calcium channels are blocked with nimo, an specific L-type calcium channel antagonist.

(Mockett et al., 2002). Indeed, LTD facilitation and LTP suppression are both elicited in priming experiments that activate NMDA receptors. We will assume that the direct action of "pre after post" pairing results in the binding of Ca^{2+}/calmodulin (variable x_{LTD}) to a phosphatase that will be responsible for rapid AMPA receptor de-phosphorylation and endocytosis. Moreover, we hypothesize that this process can be facilitated by the Ca^{2+}/calmodulin-dependent negative feedback generated through NMDA channels that will make more Ca^{2+}/calmodulin molecules available for the phosphatase. This is modelled as a buffering term inversely proportional to the slow LTP action $\exp(x_{\mathrm{LTP}})$, averaged over a slow time constant. This LTD mechanism is depicted in Fig. 1B and obeys the following equation:

$$\delta w_{\mathrm{LTD}}(t) = [\sum_{t_{\mathrm{pre}}} x_{\mathrm{LTD}}(t)\delta(t - t_{\mathrm{pre}}) - \frac{\alpha_{\mathrm{LTD}}}{T}\int_{-\infty}^{t} d\tau e^{(\tau - t)/T} e^{-\beta x_{\mathrm{LTP}}(\tau)}]_{+} \tag{2}$$

where α_{LTD} measures the magnitude of the facilitation of LTD. The time evolution of the synaptic weights is given by the integral over the synaptic changes :

$$w(t) = \lambda \int_{-\infty}^{t} ds\,(\delta w_{LTP}(s) - \delta w_{LTD}(s)) \tag{3}$$

with λ being the learning rate.

To determine if this model is consistent with STDP experiments, we first show that its predictions reproduce the plasticity effects reported for first and second order inter-spike interactions. In Fig. 1C, the bi-exponential STDP curve was obtained using a pairing protocol at 1 Hz. Triplet interactions were obtained in Fig. 1D-E for a given set of parameters chosen to qualitatively match the data reported in Sjöström et al. (2001) and Wang et al. (2005). Previous modelling efforts have also succeeded in fitting these data with similar short-term interactions (Pfister and Gerstner, 2006, Clopath et al., 2010). Here, the short-term interaction is inherited from the assumption that the Ca^2+-bound calmodulin concentration decays slowly compared to the STDP time constant. In Fig. 1D, plasticity curves are plotted for a fixed pairing interval occurring at various frequencies. In particular, for a post-pre pairing, the synapse first experiences depression followed by potentiation at higher frequencies, as reported in Sjöström et al. (2001). Figure 1E displays the synaptic changes for different patterns of pre-synaptic and post-synaptic spike triplet patterns (Wang et al., 2005). In a pre-post-pre firing pattern, the synapse displays virtually no plasticity if the second pre-synaptic spike is close to the post-synaptic spike. Whereas significant potentiation is obtained for longer intervals or post-pre-post patterns that would produce depression if short-term interactions were neglected. Moreover, when simulating L-type channel blockade by forcing x_{LTD} to zero, we obtained a comparable potentiation between pre-post-pre and post-pre-post patterns, a prediction in accordance with experimental data.

To further show the generality of this mechanism, we studied the relationship between mSTDP and the BCM plasticity rule (Bienenstock et al., 1982). Previous theoretical studies have succeeded in reproducing the BCM rule by considering non-linear interactions between pre-synaptic and post-synaptic spike patterns in STDP (Izhikevich and Desai, 2003, Burkitt et al., 2004, Pfister and Gerstner, 2006, Clopath et al., 2010). This fit was achieved either by assuming first-neighbour interactions between spikes (Izhikevich and Desai, 2003, Burkitt et al., 2004) or short-term interactions involving slower variables (Pfister and Gerstner, 2006, Clopath et al., 2010). However, these models were unable to account for one eminent feature of the BCM rule, the sliding threshold, which was either included as an ad-hoc mechanism or was absent (Bush et al., 2010). The sliding threshold is hyposits a supra-linear function of the past post-synaptic firing rate and has been introduced to ensure non trivial convergence in the plasticity algorithm (Bienenstock et al., 1982). Moreover, experimental evidence has pointed

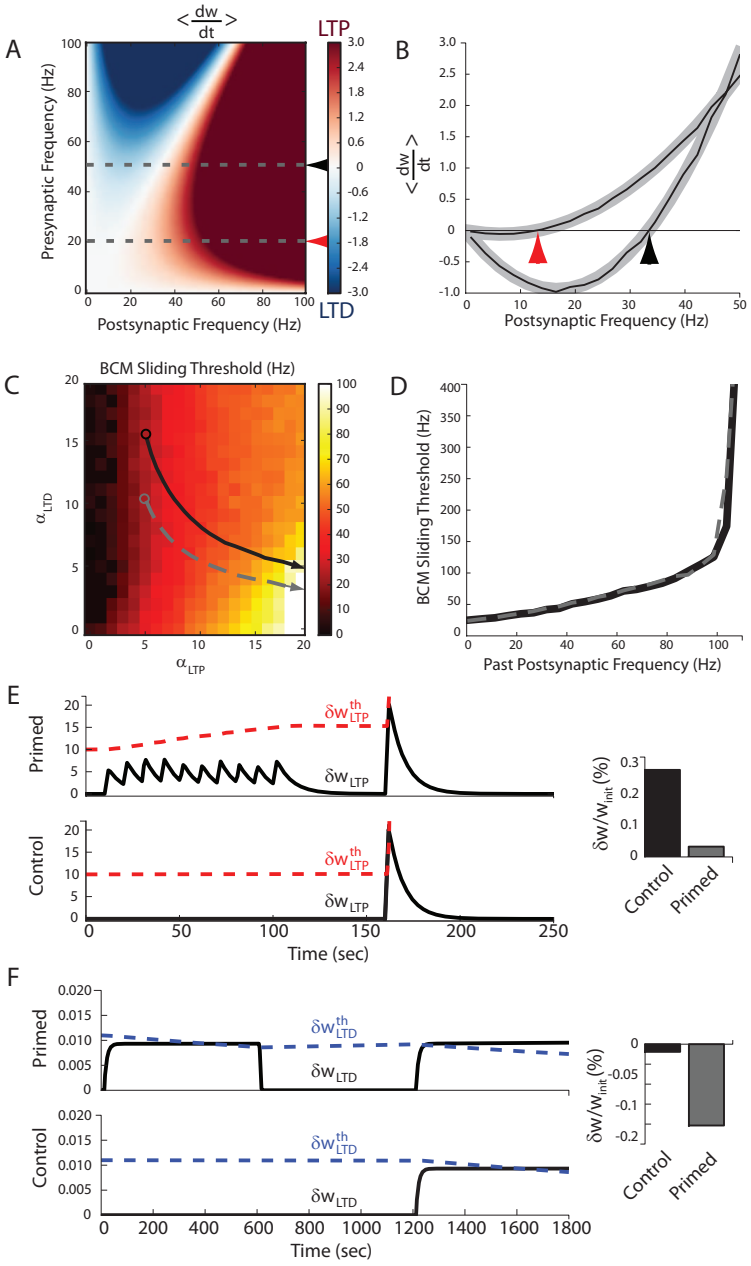

Figure 2: Relationship between mSTDP and the BCM rule. A. The net synaptic change ($\langle dw/dt \rangle$ (LTP or LTD) according to the pre-synaptic and post-synaptic firing rate in a two-neurons model with Poisson-process firing statistics and in the absence of slow variables. B. Two BCM curves corresponding to the gray dashed lines in A for a pre-synaptic firing rate of 20 Hz (top curve) and 50 Hz (lower curve). C. The BCM threshold as a function of the coefficients α_{LTP} and α_{LTD} and for an infinite T. Black lines represent the trajectory obtained for different postsynaptic firing rates for $T = 10$ sec and two coefficient pairs (α_{LTP},α_{LTD}) = (5,10) and (5,15). D. The supra-linear dependency of the BCM sliding threshold on the past post-synaptic firing rate history for various coefficients α_{LTP} and α_{LTD}. E. Simulation of an LTP priming experiment inspired by Huang et al. (1992), where weak tetanus stimulation can elicit LTP suppression. The top panel show the evolution of each variable with a priming stimulation and the lower panel show the control condition without priming. The histograms on the right are the resulting synaptic weight changes. F. Same as E but with a protocol inspired by Mockett et al. (2002) where low-frequency stimulation can elicit LTD facilitation.

out the necessity of additional mechanisms regulating plasticity induction – a process termed metaplasticity (Huang et al., 1992, Christie and Abraham, 1992, Mayford et al., 1995, Deisseroth et al., 1995, Abraham and Bear, 1996). In Fig. 2A, we first plotted the amplitude and sign of plasticity for various pre-synaptic and post-synaptic firing in a model consisting of two neurons firing as Poisson processes. Figure 2B shows two BCM curves that correspond to selected lines in the diagram. These curves were computed in absence of slow variables and are compared to theoretical predictions in bold. The depression domain and thus the BCM threshold increases with pre-synaptic firing as expected from the model definition (see Supplementary Material). This dependency creates competition between synapses originating from neurons with different firing rates, a property that classical STDP with all-to-all interactions lacks (Burkitt et al., 2004). To further explore this model, we then studied the role of the slow variables $\delta w_{\mathrm{LTP}}^{\mathrm{th}}$ and $\delta w_{\mathrm{LTD}}^{\mathrm{th}}$ in modulating the BCM threshold with the past post-synaptic firing rate history. We assumed that the slow time constant T is slow enough to be considered constant during the BCM curve measurement. For $T \to \infty$, the slow action on LTP and LTD are contained in the coefficients α_{LTP} and α_{LTD} of Eqs. 1,2 that act as constant thresholds defined by the past activity regime. Figure 2C shows the BCM thresholds obtained for different values of these coefficients. For a given parameter set and for a finite time constant T=10 sec, successive post-synaptic firing rates resulted in effective stationary variables $\delta w_{\mathrm{LTP}}^{\mathrm{th}}$ and $\delta w_{\mathrm{LTD}}^{\mathrm{th}}$ that are plotted as trajectories in the figure. Figure 2D, shows the corresponding BCM threshold as a function of the past post-synaptic firing rate assumed constant during the BCM protocol. This relationship is supra-linear as required by the theory for stabilizing the learning rule (Bienenstock et al., 1982).

To confirm that the mSTDP model faithfully reproduces the dynamics of metaplasticity as reported *in vitro*, we replicated priming experiments that were shown to either inhibit LTP or facilitate LTD. In Fig. 2E, the protocol used in Huang et al. (1992) was reproduced in a two-neuron model. When weak tetanus priming stimulations were used before subsequent strong tetanus stimulation (see Supplementary Material) the resulting synaptic weight change was considerably lower than the change elicited by the strong tetanus stimulation alone. Pre-activation of LTP increases its induction threshold and then decreases the amount of LTP obtained during the strong tetanus stimulation. Conversely, LTD induction by low-frequency stimulation was significantly facilitated when a previous low-frequency priming stimulation was used (Fig. 2F). This result is in accordance with LTD facilitation reported by (Christie and Abraham, 1992, Wang et al., 1998, Mockett et al., 2002). Low frequency stimulation of the LTD pathway decreases the induction threshold for LTD, and therefore increases the amount of depression compared to a control situation without priming (See Supplementary Material). Moreover, in Mockett et al. (2002), it was shown that this facilitation occurs through NMDA channels and is concurrent with LTP inhibition, which is in accordance with our definition of the slow variables both depending in x_{LTP}.

To show that the mSTDP model is robust to ongoing activity, we investigated several stimulation paradigms, from single-cell to network models. At the single-cell level, we considered a pre-synaptic population of 1,000 excitatory neurons and 250 inhibitory neurons projecting to a post-synaptic neuron, where only the excitatory synapses were subject to plasticity. To mimic the stochastic ongoing activity observed *in vivo*, each pre-synaptic neuron followed a Poisson process with a mean rate of 10 spikes/sec. The parameters α_{LTP} and α_{LTD} were chosen such that no synaptic changes were observed in this regime - considered as spontaneous activity. We separately considered the effect of heterogeneous pre-synaptic firing rate and synchrony within the pre-synaptic neurons. The non-linear inter-spike interactions responsible for the BCM plasticity can create competition between groups of neurons with different firing rates (Fig. 3A). We thus increased the firing rate of half the pre-synaptic excitatory neurons and indeed observed a rapid separation of their synaptic weights compared to those of the unstimulated neurons

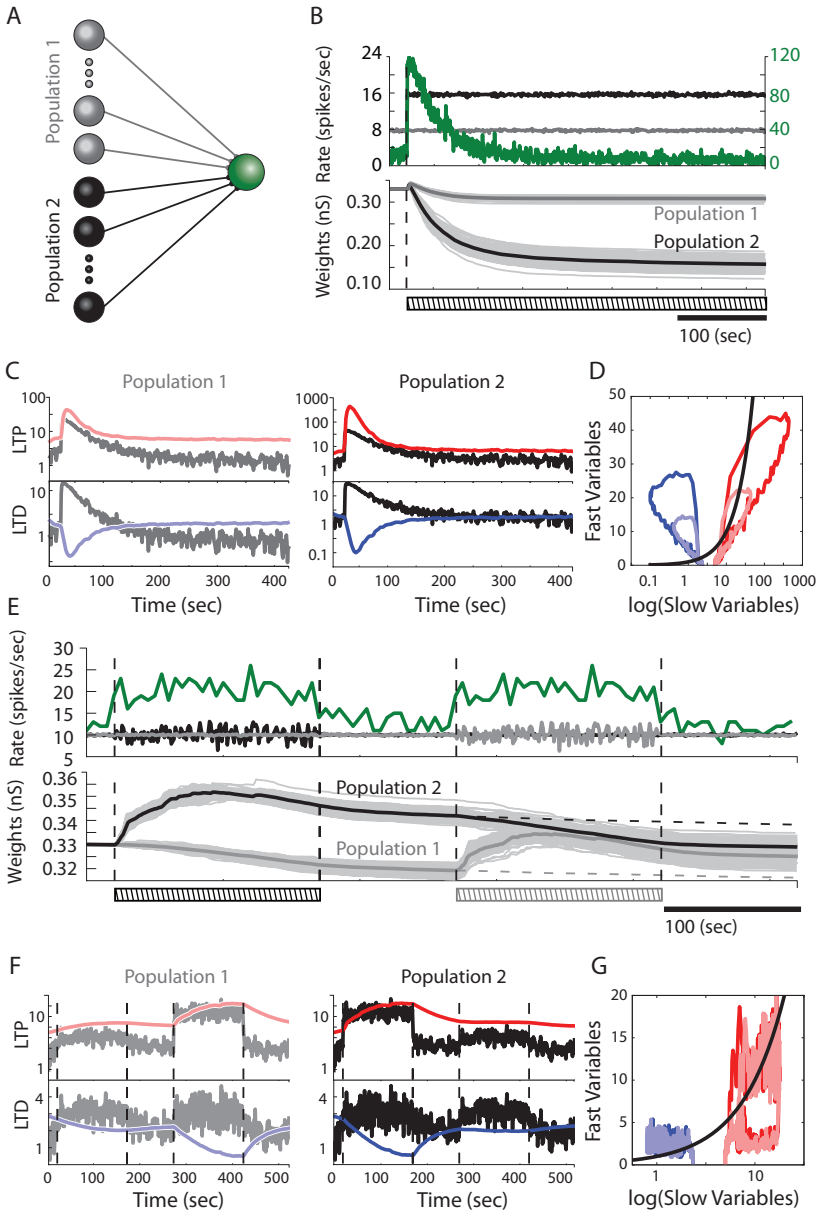

Figure 3: The BCM sliding threshold is predicted by the mSTDP model. A. Illustration of the feed-forward model. B. Time evolution of the firing rates (top) and the synaptic weights (bottom) in a feed-forward model with pre-synaptic Poisson processes. In this protocol, the post-synaptic neuron receives Poisson input from two pre-synaptic Populations. After 20 sec, the pre-synaptic neurons in population 2 double their firing rates from 8 to 16 spikes/sec (black curve) whereas the other population remain unchanged (gray curve). The post-synaptic neuron's firing rate is shown in green and individual synaptic weights in light gray (bottom). The dashed bar indicates the stimulation duration. C. Time evolution of the fast variables $\delta w_{\mathrm{LTP/LTD}}$ (black) and the corresponding slow variables $\delta w^{\mathrm{th}}_{\mathrm{LTP}}$ (red) and $\delta w^{\mathrm{th}}_{\mathrm{LTD}}$ (blue). These quantities are depicted in light color on the left for Population 1 and in dark color on the right for Population 2. The y-axis is plotted with a logarithmic scale. D. The trajectory in the space of slow and fast variables for Population 1 (light colours) and population 2 (dark colors). The x-axis is plotted with a logarithmic scale. E. Same as in B but with a different protocol where Population 2 neurons are synchronous Poisson processes (c=25%) between 20 sec and 170 sec (black cross-hatched bar) and Population 1 neurons are synchronous Poisson processes (c=25%) between 270 sec and 420 sec (gray cross-hatched bar). In dotted lines, the time evolution of the population mean synaptic weights are shown when only the first correlated stimulation is applied. F,G. Same as in C and D.

(Fig. 3B). This rapid plasticity phase was then followed by a phase where LTP was inactivated whereas LTD act on all synapses to scale the overall incoming synaptic activity until the variables δw_{LTP} and δw_{LTP} for both populations returned under their respective thresholds (Fig. 3C). The resulting post-synaptic firing rate was slightly diminished and the synaptic weights were perfectly stable in this new ongoing regime (Fig. 3B). In the phase space of both rapid and slow variables, the system trajectory is a loop that crosses the equality line and goes back to its origin point (Fig. 3D) thus effectively producing the desired BCM sliding threshold. The phase space trajectory for the high-firing rate population is exposed to stronger depression above the threshold compared to the unchanged stimulation. Therefore, heterogeneous firing rates elicit competitive synaptic weight scaling that will ensure that the fast variables return under their respective thresholds thus resulting in an overall diminution and stabilization of individual synaptic weights.

We next consider the case of correlated Poisson processes with a fixed mean firing rate. As reported in Song and Abbott (2001), if half the pre-synaptic neurons are driven by correlated synchronous inputs, inter-synaptic competition occurs and produces two separate synaptic weight distributions corresponding to each population. In Fig. 3E, in the first stage half the pre-synaptic neurons were driven by correlated Poisson processes (correlation coefficient of 0.25). The resulting synaptic weight distribution after the rapid learning phase is two separate unimodal distributions. This phase was then followed by a normalization phase that brings back the fast variables back below their induction thresholds. To show that these changes are not irreversible and that homogeneous distributions can be recovered with appropriate synaptic inputs, we then reversed the two population spiking patterns for the same stimulation duration. The final distribution is a unimodal distribution where both population synaptic weights are again mixed up. This second stimulation will bring back the fast variables near the identity line as can be seen in the time evolution of all variables (Fig. 3F). The full evolution of the system is shown in Fig. 3G where quasi-reversibility is illustrated by the overlapping population curves. From this result, we conclude that uncorrelated firing can segregate subpopulations of pre-synaptic neurons, but mainly results in down-scaling the overall synaptic weight distribution, whereas correlated inputs can elevate a subset of synaptic weights, thus making them more efficient in driving the post-synaptic neuron.

Finally, we considered the difficult problem of synaptic plasticity in recurrent networks displaying *in vivo*-like spontaneous activity. In this case, the synaptic weights directly change the network dynamics and conversely. We used balanced neuron networks, which have been shown to produce asynchronous irregular firing regimes (van Vreeswijk and Sompolinsky, 1996, 1998, Brunel, 2000) that are similar to the cortical activity observed *in vivo* (El Boustani et al., 2007). We compared a classic weight-dependent STDP rule (van Rossum et al., 2000) and the mSTDP rule in a topological network displaying an AI regime with a stable synaptic weight distribution (see Materials and Methods). A circular region of the network was stimulated with a different external pattern, while recurrent synaptic connections within and between stimulated and unstimulated regions were monitored (Fig. 4A). In these simulations, plasticity was authorized only within recurrent loops and externally driven synapses were kept constant. When the network was stimulated with uncorrelated Poisson inputs which increased the mean input rate, the recurrent synaptic weights strongly depressed during the stimulation (Fig. 4B-C, mSTDP). These changes were limited around the stimulated area (blue patch) since recurrent connection distribution was small compared to the size of the externally driven patch. These modifications did not last in the classical STDP model (Fig. 4B-C, STDP), because of the spontaneous network activity. This "catastrophic forgetting" effect was cancelled by the mSTDP model, which stabilized the induced synaptic changes (Fig. 4B-C, mSTDP). Only synaptic weights involving a direct connection with the stimulated region were affected. As a direct consequence, neuron firing rates in the stimulated region significantly diminished compared to the surrounding region

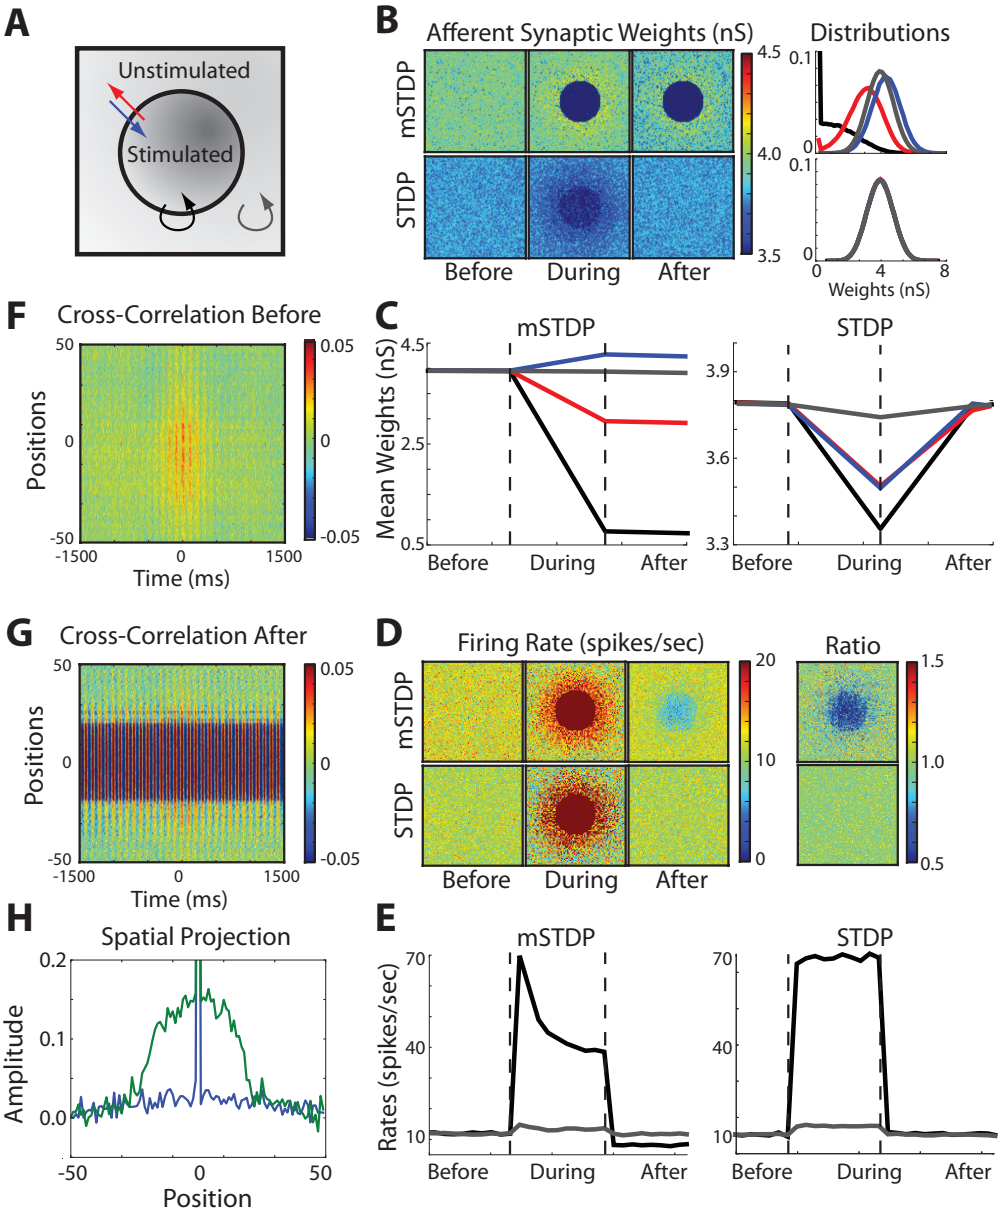

Figure 4: Stable learning in recurrent neuron networks with *in vivo*-like spontaneous activity. A. Stimulation protocol: a circular patch of the network received uncorrelated spike trains for a duration of 10 seconds. The color-coded arrows refer to different intrinsic recurrent network connections. Only the recurrent connections are plastic. B. Snapshots of the mean afferent synaptic weight per neuron is plotted for each neuron for the weight-dependent STDP and the mSTDP rules. The final population distributions are plotted of the right. C. Time evolution of mean synaptic weights for each connection group shown in A before, during and after the stimulation. D. Snapshots of neuron firing rates are plotted for both STDP rules. The ratio between the initial and final condition are shown on the right. E. Temporal evolution of both populations mean activity before, during and after the stimulation. F-H. Spatio-temporal membrane potential cross-correlation between a neuron in the centre of the stimulated region and a row spanning the network. This function is plotted before the stimulation (F) and after (G). H. The spatial profiles of the instantaneous cross-correlation function before and after the stimulation.

(Fig. 4D-E).

We also computed the spatio-temporal membrane potential cross-correlation between a neuron in the centre of the circular region and a row of neurons spanning the network. Figure 4F-G shows the spatio-temporal correlation landscapes, before and after stimulation respectively. The temporal behaviour is affected and the oscillatory stripes were enhanced and extended further in time after the stimulation. The spatial profiles dramatically broaded over time. They initial and final patterns can be compared in Fig 4H. Therefore, when mSTDP is acting on synaptic interactions, spatio-temporal statistics in the ongoing activity can be modified by the evoked activity in a stable way without disrupting the AI regime.

To conclude, we have shown that mSTDP can generalize a number of previous models and experiments on synaptic plasticity: it provides a plausible biophysical substrate for a sliding plasticity threshold and metaplasticity, and it can lead to stable learning in the presence of spontaneous activity without disturbing the network activity state. Moreover, the model ingredients are consistent with known molecular pathways involved in associative synaptic plasticity. Previous work has suggested that a likely candidate to mediate metaplasticity processes could be an auto-phosphorylated state of calcium/calmodulin kinase II (Mayford et al., 1995, Bear, 1995, Giese et al., 1998, Elgersma et al., 2002, Zhang et al., 2005). In particular, auto-phosphorylation at site Thr305/306 has an suppresive effect on the calmodulin kinase binding, acting on a very slow time scale as described in our model (Elgersma et al., 2002, Zhang et al., 2005). The CaMKII mechanism has been studied in theoretical models (Lisman, 1989, 1994, Lisman et al., 2002, Graupner and Brunel, 2007), showing that it is well suited to induce multi-stability states However, these models did not incorporate the slow negative feedback necessary to produce the BCM sliding threshold. Following Wang et al. (2005), LTD could be mediated through calmodulin-dependent activation of calcineurin following calcium influx trough L-type channels. We made an additional assumption regarding the availability of putative calcineurin for LTD when priming stimulations elicit LTP negative feedback. This hypothesis provides the facilitation term for LTD, which is an important ingredient of the mSTDP mechanism. This critical assumption can be tested experimentally to directly verify the present molecular model.

1 Acknowledgements

This work has been supported by the CNRS, the Agence Nationale de la Recherche (ANR) and EC contracts, BrainScales and Brain-i-nets (FP7-2009-ICT-FET 243914).

References

Abraham, W. C. (2008). Metaplasticity: tuning synapses and networks for plasticity. *Nat Rev Neurosci*, 9(5):387.

Abraham, W. C. and Bear, M. F. (1996). Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci*, 19(4):126–130.

Bear, M. F. (1995). Mechanism for a sliding synaptic modification threshold. *Neuron*, 15(1):1–4.

Bell, C. C., Han, V. Z., Sugawara, Y., and Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, 387(6630):278–281.

Bi, G. Q. and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 18(24):10464–10472.

Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci*, 2(1):32–48.

Billings, G. and van Rossum, M. C. W. (2009). Memory retention and spike-timing-dependent plasticity. *J Neurophysiol*, 101(6):2775–2788.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*, 8(3):183–208.

Burkitt, A. N., Meffin, H., and Grayden, D. B. (2004). Spike-timing-dependent plasticity: the relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Comput*, 16(5):885–940.

Bush, D., Philippides, A., Husbands, P., and O'Shea, M. (2010). Reconciling the STDP and BCM models of synaptic plasticity in a spiking recurrent neural network. *Neural Comput*.

Christie, B. R. and Abraham, W. C. (1992). Priming of associative long-term depression in the dentate gyrus by theta frequency synaptic activity. *Neuron*, 9(1):79–84.

Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci*, 13(3):344–352.

Deisseroth, K., Bito, H., Schulman, H., and Tsien, R. W. (1995). Synaptic plasticity: A molecular mechanism for metaplasticity. *Curr Biol*, 5(12):1334–1338.

El Boustani, S., Pospischil, M., Rudolph-Lilith, M., and Destexhe, A. (2007). Activated cortical states: experiments, analyses and models. *J Physiol Paris*, 101(1-3):99–109.

Elgersma, Y., Fedorov, N. B., Ikonen, S., Choi, E. S., Elgersma, M., Carvalho, O. M., Giese, K. P., and Silva, A. J. (2002). Inhibitory autophosphorylation of CaMKII controls PSD association, plasticity, and learning. *Neuron*, 36(3):493–505.

Fusi, S., Drew, P. J., and Abbott, L. F. (2005). Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611.

Giese, K. P., Fedorov, N. B., Filipkowski, R. K., and Silva, A. J. (1998). Autophosphorylation at thr286 of the alpha calcium-calmodulin kinase II in LTP and learning. *Science*, 279(5352):870–873.

Graef, I. A., Mermelstein, P. G., Stankunas, K., Neilson, J. R., Deisseroth, K., Tsien, R. W., and Crabtree, G. R. (1999). L-type calcium channels and GSK-3 regulate the activity of NF-ATc4 in hippocampal neurons. *Nature*, 401(6754):703–708.

Graupner, M. and Brunel, N. (2007). STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLoS Comput Biol*, 3(11):e221.

Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J Neurosci*, 23(9):3697–3714.

Huang, Y. Y., Colino, A., Selig, D. K., and Malenka, R. C. (1992). The influence of prior synaptic activity on the induction of long-term potentiation. *Science*, 255(5045):730–733.

Izhikevich, E. M. and Desai, N. S. (2003). Relating STDP to BCM. *Neural Comput*, 15(7):1511–1523.

Kumar, A., Schrader, S., Aertsen, A., and Rotter, S. (2008). The high-conductance state of cortical networks. *Neural Comput*, 20(1):1–43.

Lisman, J. (1989). A mechanism for the hebb and the anti-hebb processes underlying learning and memory. *Proc Natl Acad Sci U S A*, 86(23):9574–9578.

Lisman, J. (1994). The CaM kinase II hypothesis for the storage of synaptic memory. *Trends Neurosci*, 17(10):406–412.

Lisman, J., Schulman, H., and Cline, H. (2002). The molecular basis of CaMKII function in synaptic and behavioural memory. *Nat Rev Neurosci*, 3(3):175–190.

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275(5297):213–215.

Marre, O., Yger, P., Davison, A. P., and Frégnac, Y. (2009). Reliable recall of spontaneous activity patterns in cortical networks. *J Neurosci*, 29(46):14596–14606.

Mayford, M., Wang, J., Kandel, E. R., and O'Dell, T. J. (1995). CaMKII regulates the frequency-response function of hippocampal synapses for the production of both LTD and LTP. *Cell*, 81(6):891–904.

Mockett, B., Coussens, C., and Abraham, W. C. (2002). NMDA receptor-mediated metaplasticity during the induction of long-term depression by low-frequency stimulation. *Eur J Neurosci*, 15(11):1819–1826.

Morrison, A., Aertsen, A., and Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Comput*, 19(6):1437–1467.

Pfister, J.-P. and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci*, 26(38):9673–9682.

Shadlen, M. N. and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*, 18(10):3870–3896.

Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164.

Softky, W. R. and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci*, 13(1):334–350.

Song, S. and Abbott, L. F. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, 32(2):339–350.

Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*, 3(9):919–926.

van Rossum, M. C., Bi, G. Q., and Turrigiano, G. G. (2000). Stable hebbian learning from spike timing-dependent plasticity. *J Neurosci*, 20(23):8812–8821.

van Vreeswijk, C. and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726.

van Vreeswijk, C. and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput*, 10(6):1321–1371.

Vogels, T. P. and Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci*, 25(46):10786–10795.

Wang, H.-X., Gerkin, R. C., Nauen, D. W., and Bi, G.-Q. (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat Neurosci*, 8(2):187–193.

Wang, Y., Wu, J., Rowan, M. J., and Anwyl, R. (1998). Role of protein kinase c in the induction of homosynaptic long-term depression by brief low frequency stimulation in the dentate gyrus of the rat hippocampus in vitro. *J Physiol*, 513 (Pt 2):467–475.

Wenisch, O. G., Noll, J., and van Hemmen, J. L. (2005). Spontaneously emerging direction selectivity maps in visual cortex through stdp. *Biol Cybern*, 93(4):239–247.

Zhang, L., Kirschstein, T., Sommersberg, B., Merkens, M., Manahan-Vaughan, D., Elgersma, Y., and Beck, H. (2005). Hippocampal synaptic metaplasticity requires inhibitory autophosphorylation of ca2+/calmodulin-dependent kinase II. *J Neurosci*, 25(33):7697–7707.

18.2 Supplementary materials

18.2.1 Materials and Methods

The mSTDP rule In this Part, we describe in detail the implementation of the mSTDP rule and the corresponding equations. We will first describe the long-term potentiation (LTP). For a given connection, every pre-synaptic spike elicits an instantaneous increment of a variable r_{LTP} followed by exponential decay with a time constant τ_{LTP}. The ordinary differential equation corresponding to this eligibility trace is given by:

$$\tau_{\text{LTP}} \frac{dr_{\text{LTP}}}{dt} = -r_{\text{LTP}} + \sum_{t_{\text{pre}}} \delta(t - t_{\text{pre}}) \tag{36}$$

Once a post-synaptic spike is triggered, the current value of the LTP eligibility trace is used to determine the current increment of the variable x_{LTP} mimicking the amount of bound Ca^{2+}/calmodulin produced, which decays with a slower time constant T_{LTP}

$$T_{\text{LTP}} \frac{dx_{\text{LTP}}}{dt} = -x_{\text{LTP}} + \sum_{t_{\text{post}}} r_{\text{LTP}}(t)\delta(t - t_{\text{post}}) \tag{37}$$

x_{LTP} will be used to define the two actions on the kinase and their subsequent effect on the synaptic weight. The instantaneous action of LTP takes place at the post-synaptic spike time and results in an increase given by the current value of x_{LTP}. This term corresponds to the calcium-dependent kinase phosphorylation of AMPARs and can be written $\sum_{t_{\text{post}}} x_{\text{LTP}}(t)\delta(t - t_{\text{post}})$. The second component describes the negative feedback with a extremely slow time constant and thus can monitor the past history of a non-linear function of x_{LTP}. We hypothesize that this inhibitory term can be modelled as an exponential operator on x_{LTP} averaged over a long time interval T and can thus be written $\frac{\alpha_{\text{LTP}}}{T}\int_{-\infty}^{t} d\tau e^{(\tau - t)/T} e^{\beta x_{\text{LTP}}(\tau)}$. The coefficient α_{LTP} defines the impact of negative feedback relative to the instantaneous potentiation effect and β determines the rate of increase of the second term. The complete synaptic increment at time t is given by the integration over discrete events preceding t of the rectified difference between these two terms

$$\begin{aligned} w_{\text{LTP}}(t) &= \lambda \int_{-\infty}^{t} ds \left[\sum_{t_{\text{post}}} x_{\text{LTP}}(s)\delta(s - t_{\text{post}}) - \frac{\alpha_{\text{LTP}}}{T} \int_{-\infty}^{t} d\tau e^{(\tau - t)/T} e^{\beta x_{\text{LTP}}(\tau)} \right]_+ && (38) \\ &= \lambda \int_{-\infty}^{t} ds \left[\delta w_{\text{LTP}}(s) - \delta w_{\text{LTP}}^{th}(t)) \right]_+ && (39) \end{aligned}$$

where λ is the learning rate of the rule. We next discuss the equations describing long-term depression (LTD). Similarly to LTP, each post-synaptic spike elicits an instantaneous increment of a variable r_{LTD} followed by exponential decay with time constant τ_{LTD}

$$\tau_{\text{LTD}} \frac{dr_{\text{LTD}}}{dt} = -r_{\text{LTD}} + \alpha \sum_{t_{\text{post}}} \delta(t - t_{\text{post}}) \tag{40}$$

where α accounts for the asymmetric impact of the LTD and LTP sides of STDP. This variable is used to determine the increment of the slower variable x_{LTP}, representing the amount of available Ca^{2+}/calmodulin produced through calcium influx through L-type channels which decays with a time constant T_{LTD}. This update occurs at each pre-synaptic spike:

$$T_{\text{LTD}} \frac{dx_{\text{LTD}}}{dt} = -x_{\text{LTD}} + \sum_{t_{\text{pre}}} r_{\text{LTD}}(t)\delta(t - t_{\text{pre}}) \tag{41}$$

This molecule is hypothesized to bind to a phosphatase which has a rapid effect on the AMPAR endocytosis that will act to reduce the net synaptic weight of the synapse. We will assume that this effect is instantaneous and linear in the synaptic weight although any monotonic non-linear function could work as well. This is described by a Dirac delta operator on the variable x_{LTD} at the time the post-synaptic neuron has fired $\sum_{t_{\mathrm{pre}}} x_{\mathrm{LTD}}(t)\delta(t-t_{\mathrm{pre}})$. Moreover, as the LTD induction depends on Ca^{2+}/calmodulin which also binds with the kinase through NMDA channels, the negative feedback on the kinase will make more Ca^{2+}/calmodulin available for the phosphatase. This can be accounted for by introducing an additional term that is inversely proportional to this inhibitory term averaged over the past history $\frac{\alpha_{\mathrm{LTD}}}{T}\int_{-\infty}^{t} d\tau e^{(\tau-t)/T} e^{-\beta x_{\mathrm{LTP}}(\tau)}$. The coefficient α_{LTD} is the magnitude of facilitation of LTD. The net depressive effect on the synaptic weight at time t is given by the integration over the rectified difference between discrete events and the facilitation term

$$w_{\mathrm{LTD}}(t) = \lambda \int_{-\infty}^{t} ds \left[\sum_{t_{\mathrm{pre}}} x_{\mathrm{LTD}}(s)\delta(s-t_{\mathrm{pre}}) - \frac{\alpha_{\mathrm{LTD}}}{T} \int_{-\infty}^{t} d\tau e^{(\tau-t)/T} e^{-\beta x_{\mathrm{LTP}}(\tau)} \right]_{+} \quad (42)$$

$$= \lambda \int_{-\infty}^{t} ds \left[\delta w_{\mathrm{LTD}}(s) - \delta w_{\mathrm{LTD}}^{\mathrm{th}}(t)) \right]_{+} \quad (43)$$

Altogether, the synaptic weight fluctuations are dictated by the difference between LTP and LTD

$$w(t) = w_{\mathrm{LTP}}(t) - w_{\mathrm{LTD}}(t) \quad (44)$$

Neuron Model We consider leaky conductance-based integrate-and-fire neurons, with membrane time constant $\tau_{\mathrm{m}} = 20$ ms, and resting membrane potential $V_{\mathrm{m}} = -70$ mV. When V_{m} reaches the spiking threshold $V_{\mathrm{thresh}} = -54$ mV, a spike is generated and the membrane potential is held at the resting potential for a refractory period of duration $\tau_{\mathrm{ref}} = 5$ ms. Synaptic connections are modeled as conductance changes

$$\tau_{\mathrm{m}} \frac{dV(t)}{dt} = (V_{\mathrm{rest}} - V(t)) + g_{\mathrm{exc}}(t)(E_{\mathrm{exc}} - V(t)) + g_{\mathrm{inh}}(t)(E_{\mathrm{inh}} - V(t)) \quad (45)$$

where the reversal potentials are $E_{\mathrm{exc}} = 0$ mV and $E_{\mathrm{inh}} = -70$ mV. The synaptic activation is modelled as an instantaneous conductance increase followed by exponential decay:

$$\tau_{\mathrm{exc}} \frac{dg_{\mathrm{exc}}(t)}{dt} = -g_{\mathrm{exc}}(t) + S_{\mathrm{exc}}(t) \quad (46)$$

$$\tau_{\mathrm{inh}} \frac{dg_{\mathrm{inh}}(t)}{dt} = -g_{\mathrm{inh}}(t) + S_{\mathrm{inh}}(t) \quad (47)$$

with time constants $\tau_{\mathrm{exc}} = 5$ ms and $\tau_{\mathrm{inh}} = 5$ ms. $S_{\mathrm{exc/inh}}(t)$ are the synaptic spike trains -point processes- coming from the excitatory and inhibitory populations respectively. The integration time step of our simulations was 0.1 ms.

Feed-forward simulations In all feed-forward simulations, parameters for the plasticity were $\beta = 0.1$, $\alpha_{\mathrm{LTP}} = 5$, $\alpha_{\mathrm{LTD}} = 2.5$, $T = 50$ s. Initial weights for the pre-synaptic population were $w_{\mathrm{init}} = 0.33$ nS, and $w_{\mathrm{max}} = 1.5$ nS. The ratio between excitatory and inhibitory weights is such that $w_{\mathrm{inh}} = 15 w_{\mathrm{exc}}$. Every neuron in the pre-synaptic population had a constant firing rate of 10 Hz. In the rate protocol stimulation, half of the pre-synaptic neurons had their firing

rate increased to 20 Hz, while in the correlation protocol, half of the neurons kept the same firing rate but were turned, with the help of a multiple interaction process (Kuhn et al., 2003), into sources with a correlation coefficient of 0.25.

Priming simulations For the LTP priming protocol, we used a two-neurons model. The pre-synaptic neuron received 10 weak tetanic bursts (30 Hz for 200 ms each), with an inter-burst interval of 1 second. Then, 5 seconds later, it received a strong burst at 100 Hz for 500 ms. Synaptic weights were chosen such that pre-synaptic weak tetanic bursts were able to elicit some spikes in the post-synaptic neuron ($w_{\text{init}} = 0.12$ nS, and $w_{\max} = 1$ nS). Parameters for the plasticity rule were $\beta = 0.375$, $\alpha_{\text{LTP}} = 10$, $\alpha_{\text{LTD}} = 80$, and $T = 100$ seconds. The learning rate was fixed to $\alpha = 0.001$.

For the LTD priming protocol, the pre-synaptic neuron receives for 60 seconds a very low frequency stimulation at 5 Hz. Then, 60 seconds later, the same stimulation occurs. Synaptic weights are chosen such that pre-synaptic weak tetanic bursts are able to elicit some spikes in the post synaptic neuron ($w_{\text{init}} = 50$ nS, and $w_{\max} = 100$ nS). The strong values of those weights could be reduced if more than two neurons were considered. Parameters for the plasticity rule were $\beta = 0.5$, $\alpha_{\text{LTD}} = 0.001$, $\alpha_{\text{LTP}} = 80$, and $T = 100$ seconds. The learning rate was fixed to 2.5.

These parameters sets were chosen to reproduce the results obtained in Huang et al. (1992) and Mockett et al. (2002). These phenomenological results, which inspired the model, can also be obtained with different parameter sets.

Network model The network is composed of 10,000 excitatory and 2,500 inhibitory neurons, arranged on a grid in a 2D layer. The grid has periodic boundary conditions to avoid any border effects. Neurons are sparsely connected with a small-world connection scheme: every neuron is connected to all its neighbours in a circle of radius r with a connection probability of $\varepsilon = 2\%$. A fraction $p = 0.25\%$ of these local connections are uniformly redrawn and re-assigned to randomly selected neurons.

Network dynamical states are defined as either Synchronous or Asynchronous (population viewpoint) and as either Regular or Irregular (neuron viewpoint). The network was set to an Asynchronous Irregular state (Brunel, 2000) of approximatively 12 Hz with a mean ISI CV of 1.6. Synaptic parameters are drawn for Gaussian distributions such that their means are: $\delta g_{\text{exc}} = 4$ nS and $\delta g_{\text{inh}} = 40$ nS. The standard deviations of the Gaussian are one third of their means. Every neuron receives an additional Poisson input at 1000 spikes/sec. Synaptic delays are considered as being linearly dependent on distance, i.e if Δ is the distance between two neurons, we have $d_{ij} = d_{\text{syn}} + \frac{\Delta}{v}$ with d_{ij} the delay between the two neurons, $d_{\text{syn}} = 0.2$ ms the minimal delay due to synaptic transmission and v the velocity of axonal conduction. In all simulations, $v = 0.2$ m.s^{-1}, i.e. delays in the network are in the range $\left[0.2, 0.2 + \sqrt{2/2v} \simeq 3.7\right]$

Stimulation in the rate and correlation paradigms For 100 seconds, the network is plastic and converges to its equilibrium. Then, a circular area with a radius $r = 0.4$ is stimulated for 10 seconds, followed by 100 seconds for recovery. In the rate stimulation protocol, an increase of firing rate is sent during the stimulation time: every neuron in this stimulated area receive an additional external input of 1000 spikes/sec. In the case of correlations, neurons in the stimulated area receives synchronous spike volleys at a frequency of 50 Hz for 10s.

Spikes are recorded from the whole network, and weights are recorded every 10 s for all the synapses in the network. Moreover, to compute cross-correlations between neuron activities, we recorded during 1 second the membrane potentials of all the neurons in the network with a time resolution of 2 ms before, during and after the stimulation period.

Simulator All simulations were performed using the NEST simulator (Gewaltig and Diesmann, 2007), using the PyNN interface (Davison et al., 2008).

18.2.2 Deriving the BCM rule from the STDP

In this Part, we derive the Bienenstock-Cooper-Munro rule (BCM) from a simple model where two neurons are connected through a synapse and follow independent Poisson processes. We first solve the equations corresponding to the STDP eligibility traces for LTP and LTD Eqs. 36 and 40:

$$r_{\text{LTP}}(t) = \int_{-\infty}^{t} ds e^{-(t-s)/\tau_{\text{LTP}}} \sum_{t_{\text{pre}}} \delta(s - t_{\text{pre}}) \tag{48}$$

$$r_{\text{LTD}}(t) = \alpha \int_{-\infty}^{t} ds e^{-(t-s)/\tau_{\text{LTD}}} \sum_{t_{\text{post}}} \delta(s - t_{\text{post}}) \tag{49}$$

These solutions are then injected into Eqs. 37 and 41 that can be solved the same way:

$$x_{\text{LTP}}(t) = \int_{-\infty}^{t} ds' e^{-(t-s')/T_{\text{LTP}}} \int_{-\infty}^{s'} ds e^{-(s'-s)/\tau_{\text{LTP}}} \sum_{t_{\text{pre}}} \delta(s - t_{\text{pre}}) \sum_{t_{\text{post}}} \delta(s' - t_{\text{post}}) \tag{50}$$

$$x_{\text{LTD}}(t) = \alpha \int_{-\infty}^{t} ds' e^{-(t-s')/T_{\text{LTD}}} \int_{-\infty}^{s'} ds e^{-(s'-s)/\tau_{\text{LTD}}} \sum_{t_{\text{post}}} \delta(s - t_{\text{post}}) \sum_{t_{\text{pre}}} \delta(s' - t_{\text{pre}}) \tag{51}$$

Based on these equations, we can write the instantaneous action of LTP and LTD corresponding to the first terms of Eqs. 39 and 43:

$$\delta w_{\text{LTP}}(t) = \sum_{t_{\text{post}}} \delta(t - t_{\text{post}}) x_{\text{LTP}}(t) \tag{52}$$

$$\delta w_{\text{LTD}}(t) = \sum_{t_{\text{pre}}} \delta(t - t_{\text{pre}}) x_{\text{LTD}}(t) \tag{53}$$

Before considering the equations for the slow contribution of CaMKII (second terms in Eqs. 39 and 43), we will treat the simpler case where only the fast variables are taken into account. We will then treat the complete case.

BCM with fast variables only We will show that a static BCM curve can be obtained directly from the model with fast variables only. Because each variable δw_{LTP} and δw_{LTD} is always positive, the rectification is not necessary. We consider the time evolution of the mean synaptic weight so that we can write Eq. 44 with an average over point processes:

$$\langle w(t) \rangle = \langle w_{\text{LTP}}(t) \rangle - \langle w_{\text{LTD}}(t) \rangle \tag{54}$$

$$= \lambda \int_{-\infty}^{t} dt' \left(\langle \delta w_{\text{LTP}}(t') \rangle - \langle \delta w_{\text{LTD}}(t') \rangle \right) \tag{55}$$

By using Eq. 53 as well as the Poisson auto-correlation function we can write the corresponding LTP and LTD terms

$$\begin{aligned}
\langle \delta w_{\text{LTP}}(t') \rangle &= \tau_{\text{LTP}} \rho_{\text{pre}} \int_{-\infty}^{t'} ds' e^{-(t'-s')/T_{\text{LTP}}} \left(\rho_{\text{post}} \delta(s'-t') + \rho_{\text{post}}^2 \right) && (56)\\
&= \tau_{\text{LTP}} T_{\text{LTP}} \rho_{\text{pre}} \rho_{\text{post}}^2 + \tau_{\text{LTP}} \rho_{\text{pre}} \rho_{\text{post}} && (57)\\
\langle \delta w_{\text{LTD}}(t') \rangle &= \alpha \tau_{\text{LTD}} \rho_{\text{post}} \int_{-\infty}^{t'} ds' e^{-(t-s')/T_{\text{LTD}}} (\rho_{\text{pre}} \delta(s'-t) + \rho_{\text{pre}}^2) && (58)\\
&= \alpha \tau_{\text{LTD}} T_{\text{LTD}} \rho_{\text{post}} \rho_{\text{pre}}^2 + \tau_{\text{LTD}} \rho_{\text{pre}} \rho_{\text{post}} && (59)\\
& && (60)
\end{aligned}$$

where ρ_{pre} and ρ_{post} are respectively the pre-synaptic and post-synaptic firing rates. These terms do not depend on time anymore so that we can directly consider the time derivative of Eq. 55

$$\begin{aligned}
\frac{d\langle w(t) \rangle}{dt} &= \lambda \langle \delta w_{\text{LTP}} \rangle - \lambda \langle \delta w_{\text{LTD}} \rangle && (61)\\
&= \rho_{\text{pre}} \left(\lambda \tau_{\text{LTP}} T_{\text{LTP}} \rho_{\text{post}}^2 + \lambda \rho_{\text{post}} (\tau_{\text{LTP}} - \alpha \tau_{\text{LTD}} T_{\text{LTD}} \rho_{\text{pre}} - \alpha \tau_{\text{LTD}}) \right) && (62)\\
&= \rho_{\text{pre}} \Phi(\rho_{\text{post}}, \rho_{\text{pre}}) && (63)
\end{aligned}$$

where $\Phi(\rho_{\text{post}}, \rho_{\text{pre}})$ is a quadratic function in ρ_{post} which crosses the abscissa at 0 and possibly another point. To obtain the BCM curve, we asked that the second crossing point is positive as well as the second derivative with respect to ρ_{post} to make sure the curve is convex. The latter requirement is always true and the primer provides the following constrain

$$\theta_{\text{LTP}}^0 = \frac{\alpha \tau_{\text{LTD}} T_{\text{LTD}}}{\tau_{\text{LTP}} T_{\text{LTP}}} \rho_{\text{pre}} + \frac{\alpha \tau_{\text{LTD}} - \tau_{\text{LTP}}}{\tau_{\text{LTP}} T_{\text{LTP}}} > 0 \tag{64}$$

so that:

$$\frac{1}{T_{\text{LTD}} \rho_{\text{pre}} + 1} < \alpha \frac{\tau_{\text{LTD}}}{\tau_{\text{LTP}}} \tag{65}$$

which is true for any ρ_{pre} if $\tau_{\text{LTP}} < \alpha \tau_{\text{LTD}}$. Otherwise the existence of a crossing point will depend on the pre-synaptic firing rate.

BCM with fast and slow variables In the previous Part, we derived the BCM curve and the conditions to ensure the existence of a positive threshold. However, one of the main ingredient of the BCM theory that provides metaplasticity is the sliding threshold that depend on the past history of the post-synaptic firing through a supra-linear function. We will show in this Part that the slow variables linked to the negative feedback on the kinase fulfil this role. The variation of the slow components in LTP and LTD can be considered constant during measurement of the BCM curve because the time scales are several order of magnitude apart. As we can compute at least numerically the expected values for $\delta w_{\text{LTP}}^{\text{th}}(t)$ and $\delta w_{\text{LTD}}^{\text{th}}(t)$ for a given pre- and post-synaptic firing rate, we can derive the relationship between the BCM threshold and the post-synaptic firing averaged over the past history. We can thus write the

synaptic weight evolution Eq. 44 as

$$
\begin{aligned}
w(t) &= \lambda \int_{-\infty}^{t} ds \left(\left[\delta w_{\mathrm{LTD}}(s) - \delta w_{\mathrm{LTD}}^{\mathrm{th}}(t)) \right]_+ - \left[\delta w_{\mathrm{LTD}}(s) - \delta w_{\mathrm{LTD}}^{\mathrm{th}}(t)) \right]_+ \right) && (66) \\
&\simeq \lambda \int_{-\infty}^{t} ds \left(\left[\delta w_{\mathrm{LTD}}(s) - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}) \right]_+ - \left[\delta w_{\mathrm{LTD}}(s) - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}) \right]_+ \right) && (67) \\
& && (68)
\end{aligned}
$$

We first derive this equation with respect to time and compute the average over the synaptic weight:

$$
\frac{d\langle w(t) \rangle}{dt} = \lambda \left\langle \left[\delta w_{\mathrm{LTD}}(t) - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}) \right]_+ \right\rangle - \lambda \left\langle \left[\delta w_{\mathrm{LTD}}(t) - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}) \right]_+ \right\rangle \quad (69)
$$

The average on the Poisson processes can be written directly as an average over the possible values of δw_{LTD} described by a stationary distribution. We can thus write

$$
\begin{aligned}
\frac{1}{\lambda} \frac{d\langle w(t) \rangle}{dt} &= \int d\delta w_{\mathrm{LTP}} P(\delta w_{\mathrm{LTP}}) \left[\delta w_{\mathrm{LTP}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}}) \right]_+ && (70) \\
&\quad - \int d\delta w_{\mathrm{LTD}} P(\delta w_{\mathrm{LTD}}) \left[\delta w_{\mathrm{LTD}} - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}) \right]_+ && (71) \\
&= \int_{\delta w_{\mathrm{LTP}}^{\mathrm{th0}}}^{\infty} d\delta w_{\mathrm{LTP}} P(\delta w_{\mathrm{LTP}}) \left(\delta w_{\mathrm{LTP}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}} \right) && (72) \\
&\quad - \int_{\delta w_{\mathrm{LTD}}^{\mathrm{th0}}}^{\infty} d\delta w_{\mathrm{LTD}} P(\delta w_{\mathrm{LTD}}) \left(\delta w_{\mathrm{LTD}} - \delta w_{\mathrm{LTD}}^{\mathrm{th0}} \right) && (73)
\end{aligned}
$$

We will consider three scenarios to see how the slow variables affect the BCM curve for low post-synaptic firing rates. We will first consider the case in which the pre-synaptic and post-synaptic firing rate past history are low so that past x_{LTP} and x_{LTD} are themselves very small. In this case and if α_{LTD} is chosen to be large enough, the second term of Eq. 73 vanishes because the baseline threshold is considered very high for small post-synaptic firing rates of the BCM curve. The first term, however, can be approximated by:

$$
\frac{1}{\lambda} \frac{d\langle w(t) \rangle}{dt} = < \delta w_{\mathrm{LTP}} > - \delta w_{\mathrm{LTP}}^{\mathrm{th0}} = \tau_{\mathrm{LTP}} T_{\mathrm{LTP}} \rho_{\mathrm{pre}} \rho_{\mathrm{post}}^2 + \tau_{\mathrm{LTP}} \rho_{\mathrm{pre}} \rho_{\mathrm{post}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}} \quad (74)
$$

so that the LTP sliding threshold can be written:

$$
\theta_{\mathrm{LTP}}(\delta w_{\mathrm{LTP}}^{\mathrm{th0}}) = \frac{-1 + \sqrt{1 + \frac{4 T_{\mathrm{LTP}} \delta w_{\mathrm{LTP}}^{\mathrm{th0}}}{\tau_{\mathrm{LTP}} \rho_{\mathrm{pre}}}}}{2 T_{\mathrm{LTP}}} \quad (75)
$$

which is a supra-linear function of the post-synaptic function as required by the BCM theory. In other words, if the past synaptic activity was very low, the BCM curve will move upward and favours the impact of LTP. Interestingly, $\delta w_{\mathrm{LTP}}^{\mathrm{th0}}$ depends also on the pre-synaptic past activity thus making this metaplasticity both hetero-synaptic and homo-synaptic. Now consider the opposite case where the past activity was very high. In this situation, $\delta w_{\mathrm{LTP}}^{\mathrm{th0}}$ becomes extremely high so that the corresponding term in Eq. 73 vanishes whereas $\delta w_{\mathrm{LTD}}^{\mathrm{th0}}$ is close to 0 and we can write Eq. 73 as

$$
\frac{1}{\lambda} \frac{d\langle w(t) \rangle}{dt} = - < \delta w_{\mathrm{LTD}} > + \delta w_{\mathrm{LTD}}^{\mathrm{th0}} = -\alpha \tau_{\mathrm{LTD}} T_{\mathrm{LTD}} \rho_{\mathrm{post}} \rho_{\mathrm{pre}}^2 - \tau_{\mathrm{LTD}} \rho_{\mathrm{pre}} \rho_{\mathrm{post}} + \delta w_{\mathrm{LTD}}^{\mathrm{th0}} \quad (76)
$$

such that the LTD sliding threshold can be written:

$$\theta_{\mathrm{LTD}}(\delta w_{\mathrm{LTD}}^{\mathrm{th0}}) = \frac{\delta w_{\mathrm{LTD}}^{\mathrm{th0}}}{\alpha\tau_{\mathrm{LTD}}T_{\mathrm{LTD}}\rho_{\mathrm{pre}}^2+\tau_{\mathrm{LTD}}\rho_{\mathrm{pre}}} \tag{77}$$

In this situation, the LTD sliding threshold can slightly increase above 0 whereas the LTP sliding threshold is infinite and the synaptic weight experience depression. In other words, if the past synaptic activity is very high, the synaptic weight tends to decrease. The last situation we will consider is in-between the two previous ones. We will assume that we are close to an equilibrium point where LTP and LTD are close to their respective thresholds so that synaptic weights are almost stable. We can rewrite Eq. 73 as

$$\begin{aligned}
\frac{d\langle w(t)\rangle}{dt} &= \int_{-\infty}^{\infty} d\delta w_{\mathrm{LTP}} P(\delta w_{\mathrm{LTP}}) \left(\delta w_{\mathrm{LTP}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}}\right) \qquad (78)\\
&- \int_{-\infty}^{\infty} d\delta w_{\mathrm{LTD}} P(\delta w_{\mathrm{LTD}}) \left(\delta w_{\mathrm{LTD}} - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}\right)\\
&- \int_{-\infty}^{\delta w_{\mathrm{LTP}}^{\mathrm{th0}}} d\delta w_{\mathrm{LTP}} P(\delta w_{\mathrm{LTP}}) \left(\delta w_{\mathrm{LTP}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}}\right)\\
&+ \int_{-\infty}^{\delta w_{\mathrm{LTD}}^{\mathrm{th0}}} d\delta w_{\mathrm{LTD}} P(\delta w_{\mathrm{LTD}}) \left(\delta w_{\mathrm{LTD}} - \delta w_{\mathrm{LTD}}^{\mathrm{th0}}\right)\\
&= \Phi(\rho_{\mathrm{post}},\rho_{\mathrm{pre}})\rho_{\mathrm{pre}} + \lambda(\delta w_{\mathrm{LTD}}^{\mathrm{th0}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}}) - \lambda F_{\mathrm{LTP}}\left(\delta w_{\mathrm{LTP}}^{\mathrm{th0}}\right) + \lambda F_{\mathrm{LTD}}\left(\delta w_{\mathrm{LTD}}^{\mathrm{th0}}\right)
\end{aligned}$$

where $\Phi(\rho_{\mathrm{post}},\rho_{\mathrm{pre}})$ is the BCM curve found in the case where the slow variables are absent and $F_{\mathrm{LTP}}\left(\delta w_{\mathrm{LTP}}^{\mathrm{th0}}\right)$ and $F_{\mathrm{LTD}}\left(\delta w_{\mathrm{LTD}}^{\mathrm{th0}}\right)$ are two functions of the LTP and LTD slow variables respectively. These functions can be expanded around the mean value of each distribution. We will also approximate each probability distribution by a Gaussian. The general form of such an expansion is given by

$$\begin{aligned}
F(c) &= \int_{-\infty}^{c} dx P(x)(x-c) \qquad (79)\\
&= F(m) + \left.\frac{dF(c)}{dc}\right|_m (c-m) + O(\delta c^2) \qquad (80)\\
&\simeq \int_{-\infty}^{m} dx P(x)(x-m) - \int_{-\infty}^{m} dx P(x)(c-m) \qquad (81)\\
&= -\frac{\sigma}{\sqrt{2\pi}} - \frac{1}{2}(c-m) \qquad (82)
\end{aligned}$$

Once Eq. 82 is used in Eq. 79, we obtain

$$\frac{d\langle w(t)\rangle}{dt} = \frac{3}{2}\Phi(\rho_{\mathrm{post}},\rho_{\mathrm{pre}})\rho_{\mathrm{pre}} + \frac{3}{2}\lambda(\delta w_{\mathrm{LTD}}^{\mathrm{th0}} - \delta w_{\mathrm{LTP}}^{\mathrm{th0}}) + \lambda\frac{\sigma(\delta w_{\mathrm{LTP}})}{\sqrt{2\pi}} - \lambda\frac{\sigma(\delta w_{\mathrm{LTD}})}{\sqrt{2\pi}}$$

The two last terms grow as ρ_{post}^2 for Poisson processes such that the overall form of the BCM curve remains the same. The remaining terms are constant and depend on the post and presynaptic firing rate history. These terms will produce a shift in the BCM curve in the downward direction that results in a supra-linear shift of the sliding threshold toward the left thus producing the desired dependence on the past activity. This can be written in a more general form

$$\frac{d\langle w(t)\rangle}{dt} = \Phi(\rho_{\mathrm{post}},\rho_{\mathrm{pre}},\delta w_{\mathrm{LTP}}^{\mathrm{th0}},\delta w_{\mathrm{LTD}}^{\mathrm{th0}})\rho_{\mathrm{pre}} \tag{83}$$

19 Discussion

STDP may be only the visible part of the iceberg, and its complexity should be more deeply considered. Biophysical and tractable models, such as the one presented in the previous section, or those considering triplets of spikes (Pfister and Gerstner, 2006), or post-synaptic membrane potential (Clopath et al., 2010), are pushing in favour of a rule that can explain the diversity of the experimental results gathered on plasticity. They both link STDP and the BCM theory, with the introduction of coefficient taken into account past activity of the post-synaptic neuron. Nevertheless, in both cases, the dependence and the supra-linear aspect of the sliding threshold is imposed as an *ad hoc* constraint. Since the BCM formalism requires a supralinear power dependence on this past activity, both models introduce a quadratic term to ensure this behaviour, without any biological explanation. In our model, the exponential behaviour for the thresholds δw_{LTP} and δw_{LTD} is inherited from molecular biology processes and can be linked with some kinetic models, as already made in more realistic models of STDP (Castellani et al., 2001, Zou and Destexhe, 2007). The proposed mechanism is kept as simple as possible to be tractable in large-scale neuronal networks: the key point is to be able to understand, phenomenologically, the pathways leading to potentiation and depression, and how this plasticity is activity dependent. Metaplasticity is crucial in a coherent learning framework. Similarly to short term plasticity and adaptation (Tsodyks et al., 2000), this may prevent the system from being too sensitive and focus its learning capacity on new features.

Indeed, plasticity rules should be compatible with the Asynchronous Irregular regime, and learning should also be robust to the ongoing activity. The rule we developed here was made for that: during ongoing activity, random pairings elicit internal changes in the synapses that do not necessarily trigger plasticity, due to some non linearities and some thresholds. Then, during a learning phase, the system gathers evidences that significant deviation from its spontaneous activity occurs. It can be an increase of firing rate, or incoming correlations, but they both affect the internal variables δw_{LTP} and δw_{LTD}. If the deviations are large enough, thresholds are crossed, and then plasticity changes are validated at the synapse level. If the stimulation is applied continuously, then by metaplasticity the thresholds will be changed, such that the system will slowly settle into a state where ongoing activity is somehow re-estimated (see Figure 36), and plasticity will be turned off.

Stability Note that such a thresholding mechanism get rids of the debate about additive or weight-dependent rules in STDP. In all the previous results in this Part, results were obtained with an additive implementation, but similar results could be expected without. The key point which guarantees the stability of the system is now the thresholds $\delta w^{th}_{\text{LTP/LTD}}$, and not the detailed schema of the weight modifications. This can also help to solve the delay sensitivity investigated in Part IV, and the problem of autapses. Biological evidences for autapses exists (Lübke et al., 1996) ($\simeq 2.3$ per pyramidal neuron in Layer 5 of the rat), and they are problematic with the classical STDP rule: they should, according to the way the delay is modelled, be always depressed or potentiated, but cannot be stable. With homeostasis, and the proposed rule, they could. BCM theory can provide a substrate whereby the sliding threshold based on the post-synaptic firing rate will balance the autapse strength. The rule is therefore able to stabilize statistical deviations from the ongoing activity, and to store them reliably in the network's structure. Changes will be kept constant as long as others patterns deviate enough to trigger synaptic changes by crossing the thresholds, and will be superimposed, as in a Hopfield network. The advantage is that synaptic changes are here made in a continuum, there is

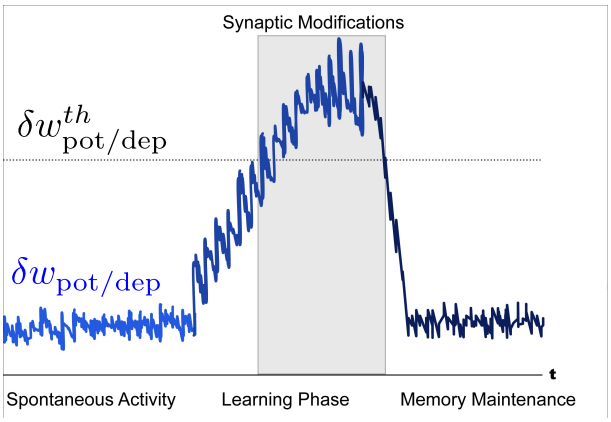

Figure 36: Schematic drawing of the learning based on metaplastic spike timing dependent plasticity. In blue, example of an internal variable δw_{LTP} or δw_{LTD}, fluctuating according to the ongoing activity. The gray line is a schematic threshold $\delta w^{\mathrm{th}}_{\mathrm{LTD}}$ or $\delta w^{\mathrm{th}}_{\mathrm{LTP}}$, which is here constant for the sake of clarity. In the gray shaded region, the threshold is crossed and plasticity is induced. Nevertheless, the threshold adapts, on a slower time scale, to $\delta w_{LTP/LTD}$ (see previous section)

no need to have discrete states to promote and ensure stability. This property emerge from the recurrent inputs and the ongoing activity, keeping the amount of calmodulin and calcineurin in ranges that modulate plasticity changes.

Modulation of the Thresholds Throughout the paper, we stress the fact that the regulation of the thresholds mainly comes from an internal homeostasis, due to the interactions between the LTP and the LTD molecular pathways, both involving calmodulin. But this can be extended to a supervised learning framework. Neuromodulators (such as dopamine, Zhang et al. (2009)) can interact with this internal machinery and affect the values of the thresholds, leading to temporal windows during which plasticity would be possible at the synapse. Gu and Yan (2004) have shown that, in the prefrontal cortex, activation of D4 receptors can regulate the Ca2+/calmodulin-dependent protein kinase II (CaMKII). This modulation is a flexible mechanism allowing dopamine to influence CaMKII and therefore plasticity at the synapse level. Moreover, some authors (Harris, 2008), hypothesis that a raw and slow version of the back-propagation algorithm, used to train artificial neuronal networks, is possible within neurons. A lot of evidence supports the idea that pre-synaptic neurons can know if the post-synaptic one has been potentiated or depressed, by retrograde signalling molecules (such as neurotrophins, or endocannabinoids) in the synaptic cleft. These signals can travel backward along the axon and could also influence the thresholds in our models, by affecting the amount of calmodulin or impacting activation of the transcription factor CREB (cyclic AMP response element binding protein) which is critical for promoting neuronal survival (Riccio et al., 1999) and is known to consolidate LTP (Barco et al., 2002). This impact could also be non homogeneous according to the position of the synapse along the dendritic tree, and therefore lead to a variety of rules according to this position, in line with the results of Letzkus et al. (2006) and Kampa and Stuart (2006). Nevertheless, the dependence of the rule on the position along the dendritic tree is a property which seems to be related to the depolarization level in the post-synaptic neuron. Strong depolarizations can change the shape of the rule (from Hebbian to anti Hebbian in Letzkus et al. (2006)), and the gradient of plasticity along the dendritic tree is in favour of the membrane potential: such a membrane potential gradient along the dendritic

tree exists if we consider the back propagation of the action potential (bAP). Thesee bAPs are one of the best candidates to inform the synapse that a post-synaptic spike has been emitted, but their propagation along the dendrite may be subject to attenuation (Larkum et al., 2001) in certain cell types, such as pyramidal neurons.

Limits of the model The model cannot explain "synaptic tagging" of Frey and Morris (1997), and the fact that some transient LTP can be observed, in hippocampal slices, when stimulating Schaffer collaterals, decaying over a timescale of $\simeq 1$ hour. In the previous paper, we explained those priming results based on the thresholds $\delta w^{\text{th}}_{\text{LTP/LTD}}$. Since they integrate activity over long time scales and slowly evolve, they can be used for keeping a trace of previous stimulation, and impact synaptic changes during the re-activation of the pathway, up to hour later. Pre-activation of the potentation circuitry increases δw_{LTP}, in such a way that another, stronger, stimulation, even a long time after, will trigger a reduced amount of LTP. However, no transient LTP is induced by the model: changes are made or not, but while validated at the synapse level, they are kept constant. Certain *in vitro* data (see Figure 35) suggests that this assumption that once a synaptic change is made, its value should be kept constant is invalid. Depression, for example, seems to be a slowly acting running process, once initiated at the synapse level. Nevertheless, this regulation of the weights, once LTP or LTD changes have been made, could be mediated by other homeostatic mechanisms, such as those described in Turrigiano and Nelson (2004). Processes could act, in parallel, to establish competition between neurons, and in order to get rid of the silent ones (by making them fire, or triggering apoptosis), may scale the weights uniformly, whatever the activity is. This normalization process may take part during sleep, when activity is replayed, and may also help to maintain a dendritic "democracy" within synapses (Roth and London, 2004).

Finally, it has to be noted that some experimental data show that potentiation can be observed even without any post synaptic spikes (Artola and Singer, 1993, Kelso et al., 1986). Such results cannot be explained by our model, which is only based on spike times, an indirect measure of the internal depolarization of the post-synaptic cell. To reproduce, as in Clopath et al. (2010), results based on membrane potentials, we should slightly extend the framework and incorporate the post-synaptic membrane potential as a quantity that will influence the thresholds. Basing the models only on the spike times has some pros and cons. From a computational point of view, this provides an optimized implementation which could be event-based, triggering computations only at pre and post-synaptic spike times. However, the drawback is that weights values are only updated when such events appear.

Experimental predictions Previous work has suggested that a likely candidate to mediate metaplasticity processes is an auto-phosphorylated state of calcium/calmodulin kinase II (Mayford et al., 1995, Bear, 1995, Elgersma et al., 2002, Zhang et al., 2005). In particular, auto-phosphorylation at site Thr305/306 has an inhibitory effect on the calmodulin kinase binding acting on a very slow time scale as described in our model (Elgersma et al., 2002, Zhang et al., 2005). With the previously explained plasticity mechanism, we made an assumption regarding the availability of putative calcineurin for LTD when priming stimulations elicit LTP negative feedback. This provided the facilitation term for LTD, which is an important ingredient of the mSTDP mechanism. This critical prediction can be tested experimentally to directly verify the present molecular model. An *in vitro* experiment should be discussed, and a blockade of the auto-phosphorylation of the calmodulin kinase should test the idea. Calcineurin uncaging should also have an effect, while it is already known that calcium uncaging

can affect the sign of plasticity. More interestingly, not only the amount of uncaged calcium can switch LTP to LTD, but also the kinetics of its release (Yang et al., 1999). These kinetics changes could be explained by the extra time scales added in our model, and integrating past activity of the neuron.

Part VI
The simulation workflow

Summary

20 Introduction **173**
20.1 Neuronal network simulators 174
20.2 The simulations tools 174
20.3 A complex trade-off 176

21 PyNN: A Common Interface for Neuronal Network Simulators. **181**
21.1 Main results 181

22 The analysis workflow in neuroscience **192**
22.1 Heterogeneity of the data 192
22.2 On the analysis tools 192
22.3 The NeuroTools project 192

23 Discussion **194**

20 Introduction

Computational neuroscience is a growing and recent field, compared to better theoretically defined and already explored fields such as mathematics or theoretical physics. A theorem can not be always demonstrated on a sheet of paper since analytical solutions are far from possible when speaking about modelling biology, which is everything except clearly defined and constrained. Approximations, hypotheses, numerical simulations and computational models, like those presented in Part I, such as the integrate-and-fire neuron, are the *sine qua none* conditions for predictions and inferences in large-scale models of neuronal networks. Simulating those networks is crucial for understanding large-scale and generic dynamical properties that may emerge from the complex interactions between all neurons, and since analytical solutions are hardly tractable, one has to deal with numerical approaches. Unlike the first artificial neural networks (ANN) (Rosenblatt, 1958), these networks need to take time into account, each node solving differential equations according to a certain integration scheme, before exchanging, by queuing, information in the form of spike times. The fundamental question that should be raised here, with simulations becoming more and more complex, networks larger and larger, is "are we sure of the simulation tools that are developed"? Complexity in the simulations and sensitivity to the code itself is problematic and cannot be neglected. Indeed, can we be sure that results obtained can be easily reproduced, and if not, what can we do to enhance this? This question is valid not only for the simulation side, but also for the analysis. The more data are produced, the more robust and straightforward the analysis should be to avoid being overwhelmed by inconsistencies that may be hard to spot.

In this Part, I would like to briefly review the large "menagerie" of simulation tools that can be found nowadays in computational neuroscience. As a neuroscientist doing simulations, using a well-established simulator is a good option, instead of reinventing the wheel, because designing a good, efficient, modular and parallel simulator is pure computer science research on its own: since tools exist and are already developed, they should be used. Nevertheless, one should be sure of what they are doing and that some hidden assumptions, of which the developers are unaware or disregard are not introducing bias into the results. There is a growing need for a better explanation of the basic principles in neuroscience modelling. If a clear agreement can be settled when speaking about an integrate-and-fire neuron, this is not always the case for a plastic synapse implementing spike timing dependent plasticity rules, for a particular wiring scheme, for a particular model of a more complex neuron. Definitions and clarifications of concepts would greatly improve the readability of the literature of the field, and enhance reproducibility of the results. This is for example the case of network structure: Nordlie and Plesser (2010) showed how difficult it is from a modelling paper to know exactly what the model is. Materials and Methods sections are not always explicit enough, may contain omissions or typographical errors, and therefore, a clear need is pushing in favour of shared concept and tools. Another example is given by the review of hippocampal neuron models by Lyle Graham (published in the Cerebral Cortex Book series), showing their lack of stability when run with a more systematic parametric space exploration.

In this spirit, I will present in this Part PyNN, a common API for neuronal simulators, that allows enhanced transparency and cross-checking of results by allowing to write the code for a model only once, and to simulate it independently with several simulators. Several advantages of such a strategy will be reviewed. I will then present a collaborative effort to share analysis tools, the NeuroTools project, started with the same goals.

20.1 Neuronal network simulators

Simulations of large scale networks of integrate and fire neurons are not easy computational tasks, and the need for parallel computing starts to become crucial when the sizes of the networks are increased. Although simulations of networks made of fewer than 50000 integrate and fire neurons can be performed on desktop computers, as soon as the number of synapses is increased, scaling non linearly with the number of neurons, simulations face memory and simulation time problems.

It is important to stress that the main bottleneck in those simulations is not the integration of the differential equations, governing the internal dynamics of the neurons. Even for a large number of neurons, as long as integrate and fire neurons are used, the system is not overwhelmed by integration time. The problem is more is the distribution of the spikes, that need to be queued and distributed according to the transmission delay between neurons. Numerous simulation software tools have been designed to allow the creation and the simulation of such large scale networks, and a recent comparative review can be found in Brette et al. (2007). Some of them will be discussed in the following section. Most of these simulation tools are clock-driven. Every time steps, operations are performed to update the state of the neuronal network. The differential equations governing the evolution of the dynamics can be integrated with a fixed or a variable time step, depending on the needed precision. When time is discretized, spike times are often rounded to the time step precision before being queued and delivered. Nevertheless, this is not always the case, and some models provide analytical solutions allowing to interpolated the exact time of the spike, and to avoid this rounding procedure. Finding an efficient way of dealing with this simulation scheme is a ongoing field of research (Morrison et al., 2005, 2007b, Eichner et al., 2009). Other approaches use an so-called event-based framework. If you can provide an analytical solution for the membrane potential dynamics, you can know exactly when a spike is emitted and get rid of the time discretization (Rudolph and Destexhe, 2006, 2007, Mouraud et al., 2006, Mattia and Giudice, 2000, Reutimann et al., 2003). This point is important to stress, because when models are limited to a few neurons, as it is often the case, some differences may emerge and may not be statistically averaged over a large populations. It is important to be sure that errors are not amplified by the resolution scheme and that both approaches always converge to statistically equivalent results, especially when some external phenomena are triggered on the exact spike times (external stimulation, weight updates, ...).

Besides computational simulators implemented on computers, another promising idea is also to design silicon devices that faithfully reproduce the biophysics of their neural counterpart. Building such neuromorphic systems is a hardware engineering problem that was started 20 years ago by Mead and Mahowald (1988), but it can be one option to circumvent the simulation time bottleneck and speed up the exploration of the potentially huge parameter space.

20.2 The simulations tools

As already mentioned in the introduction, a lot of neuronal network simulators have been designed during the last 20 years. An exhaustive review would be tedious, so we will only review briefly the most widely used simulators that can be found nowadays. Some are generic and not especially dedicated to multi-compartmental biophysical spiking neurons (NEURON, GENESIS), while some are made only to simulate point spiking neurons, non-biophysical (NEST, CSIM, Brian, MVASpike, SPLIT, SpikeNet). Dendrites are in this case ignored, and

neurons are considered only from an input/output point of view: if large networks then become computationally tractable, the drawback is that any kind of filtering and/or processing that may be performed in the dendrites is ignored. The choice of one category of simulation tools is closely linked with the problem considered.

NEURON The most popular simulator in neurobiology is NEURON (Hines and Carnevale, 1997), developed in order to model compartmental models of neurons with real morphological structures, ion channel densities, particular kinetics, and so on. NEURON is written in *C* and $C++$, with a few bits of Python, but the definition of a the `hoc` language (and more recently, the Python wrapper) allows model specifications rather easy to implement and to compile. Linked models can then be used, simulated, and analysed since NEURON also comes with a full graphical user interface (GUI) allowing to monitor and record all the simulated variables. Both the power and expressiveness of the `hoc` and the GUI contribute a lot in spreading NEURON in the computational neuroscience community. Nevertheless, the core object of NEURON is a `section` (a compartment where equations can be solved, that can be linked with others to build morphologies), and the simulator was not particularly designed to deal with point-process neurons. Recently, with the BlueBrain project initiated by Markram (2006), a lot of efforts have been devoted to improving the parallel implementation of the software (Hines et al., 2008). The aim of the BlueBrain project is to model, on a BlueGene super computer, a fully detailed cortical column in order to understand the key parameters governing its dynamics. Reconstructed cells are modelled with NEURON, connected, and computations are split over a large number of processors. The compartments of the cells can be distributed, and MPI is used as a communications layer to distribute spiking events.

NEST In contrast to NEURON, NEST (Gewaltig and Diesmann, 2007) is a simulator dedicated to efficient simulation of large networks of point neurons. Entirely written in $C++$, it provides an interpretor layer, the SLI, which uses a stack-oriented language allowing scripts and on-line commands within a shell. The SLI interface is well documented and offers a lot of functionality, long scripts can be hard to read and especially hard to understand from an external point of view. To solve this issue, NEST now has a Python wrapper, PyNEST, allowing to interface with NEST from Python scripts (Eppler et al., 2008). NEST is primarily made to simulate large-scale models of neuronal networks, so particular efforts are put into memory efficiency, parallelism, and spike time buffering. The building and simulation times are fast and scales almost linearly, to a certain extent, with the number of processors. If we take for example a rather classical network, such as the one made by Vogels and Abbott (2005). The network is composed of 100000 integrate and fire neurons with current-based synapses, sparsely and randomly connected, with a probability of $\varepsilon = 2\%$. Table 4 shows the scaling of NEST, either for the building or the simulation time. This was the simulator used in all the simulations made in this manuscript. In NEST, interactions between neurons can only be event-based. Those are the events, or messages, that are exchanged and transmitted through a MPI communication layer. A complex queuing system allows the simulator to take heterogeneous delays into account.

(P)CSIM CSIM is a $C++$ based simulator, with a Python scripting interface that has been rewritten to give a parallel version, PCSIM (Pecevski et al., 2009), with a Python interface. CSIM took advantage of MATLAB's GUI, and was made for rather small and scripted neuronal networks. In contrast, PCSIM is one of the most complete spiking neuron simulators,

Number of Processors	Building Time	Simulation time
10	244s	19s
20	160s	6.2s
40	147s	5s
80	140s	5s

Table 4: Times for building and simulating for 1 s of simulation time a network of 100000 integrate and fire neurons with current-based synapses. Cell parameters are as in Vogels and Abbott (2005), and the connectivity ratio is $\varepsilon = 2\%$. Hardware is a a cluster of 20 quadri-cores processors (Intel(R) Xeon(R) at 2.33GHz) with 8Gb of memory each.

offering a lot of functionality and models. Deeply object-oriented, core structures are optimized but not easy to modify. The software lacks it its documentation and examples, and even if fast and efficient (personal benchmarks show that PCSIM is as fast as NEST for building and simulating time), the interface may be discouraging for the user. Like in NEST, scaling of the simulation time is almost linear for large-networks when launched in a distributed environment, communications being exchanged with MPI. Note that PCSIM, with its Python based approach, is able to encapsulate Brian code (see below) and therefore already establish a link between simulators.

Brian Brian (Goodman and Brette, 2008) is a recent and very promising neuronal network simulator. Entirely written in Python, it uses the power of an interpreted language to allow easy and user-friendly model specification. Most of its implementation relies on the two core scientific packages of Python, `numpy` and `scipy`. To implement a particular neuron model, the user just needs to enter the differential equations governing the evolution of the membrane potential and nothing more. Parsing is done automatically and units are also handled in such a way that the specified equations are turned into vectorized operations performed at each time step. Brian is dedicated to spiking neurons, using sparse matrices to store spike times. Recent optimizations allow to use the power of the GPU to speed up considerably the resolution of the differential equations (with the help of `pycuda`). Rather small (< 10000 neurons) and homogeneous networks are easy to build and simulate. Brian is not parallel, and spikes cannot be routed between different computational nodes. Networks must fit in a desktop machine's memory, and the simulation memory size is mainly imposed by the sparse connectivity matrix of the network.

20.3 A complex trade-off

Every simulator has its own pros and cons, and simulation needs depend on the problem and the targeted model: it could be the need for a GUI, that a particular model exists for simulator and not another, a need for memory efficiency, and so on. If one wants to build very large scale networks, the need for a simulator which is able to deal with distributed simulations may be necessary. Neurons themselves are not the main bottleneck: problems are encountered when networks are too dense, with too many synapses that may be plastic, and so increase the processing and memory requirements. Spikes buffering, queuing and delivering are the major problems that need to be addressed when dense (connection probability $\varepsilon > 5\%$) networks of more than 50000 neurons are built. When a first model is started with a small toy network,

made with an easy to use simulator, one can later encounter the need to simulate the exact same model but, for example, increase its size. The problem is then that encountering the need for an MPI based implementation, and therefore for another neuronal network simulator, should keep the code already made. To address this issue, PyNN is a meta language that allows network model code to be run on various simulators. This gives several advantages.

Code exchange Computational neuroscience, like all computational sciences, relies on code transparency and reproducibility of results. When building models with home made code, one should in theory provide the complete framework allowing readers to reproduce and extend the results. Documentation, equations, and details of implementation should be available in such a way that people can easily spot bugs or start a new research project with already existing models. However, simulation studies suffer from a lack of code exchange and reproducibility. Since every research group develops its own tools, spending more time on the research than on the documentation, it is too often hard for someone outside the group to dive into the code in order to understand what has been done, and why. Nevertheless, it is crucial to be able to compare, to cross-validate results. PyNN (Python Neuronal Networks), that will be presented in the following paragraph, was made to circumvent this problem. This could be seen as an attempt to unify the field of neuronal network simulation by trying to see what are the common points that can be found between different simulators. Even if all of them have unique and special features, they all share a common ground which should be more properly exposed and stressed, in order to get rid of the simulator dependency. For example, most of them have integrate and fire neurons, are able to connect them with current or conductance based synapses, are able to record membrane potentials or spikes. In order to be sure that the integration methods used by each of them is consistent, and that simulation results are coherent, the idea of PyNN was to provide a simple high-level application programming interface (API) to use several simulators, according to the user's choice. The advantages are numerous. The most obvious one is code exchange. If one research group is more used to code everything with NEST, for example, and another one is most familiar with NEURON, they should still be able to compare code, while keeping the particular knowledge gained on their favourite simulator. To achieve such a goal, one needs to establish a common language between the simulators, because as Figure 37 shows, since every simulator uses its own language, its own definitions, it may be hard to easily convert the code. Different interfaces and language syntax make the understanding more complex.

Using a common language for the code will lead to a gain in transparency and definition of the language itself will help to clarify and set the key concepts shared by all those simulators. Then, by launching the exact same code with several simulators, one can verify that results do not depend on the details of the implementations.

Code reliability PyNN also allows to better understand the hidden differences that may exists between simulators. For example, let us take the case of a plastic synapse, with the spike timing dependent plasticity rule (STDP), seen in Part IV. From a modelling point of view, STDP is ill defined at zero time difference: what should the neuronal simulator do if pre- and post-synaptic neurons fire exactly at exact the same time? If the synapse is hosted by the pre-synaptic neuron, then the post-synaptic spike will always arrive at the synapse after a time, δt, the simulation time step, and the synapse will always be potentiated. On the other hand, if the synapse is hosted by the post-synaptic neuron, then the opposite is true: pre will always precede post, and the synapse will be depressed. STDP is very sensitive to

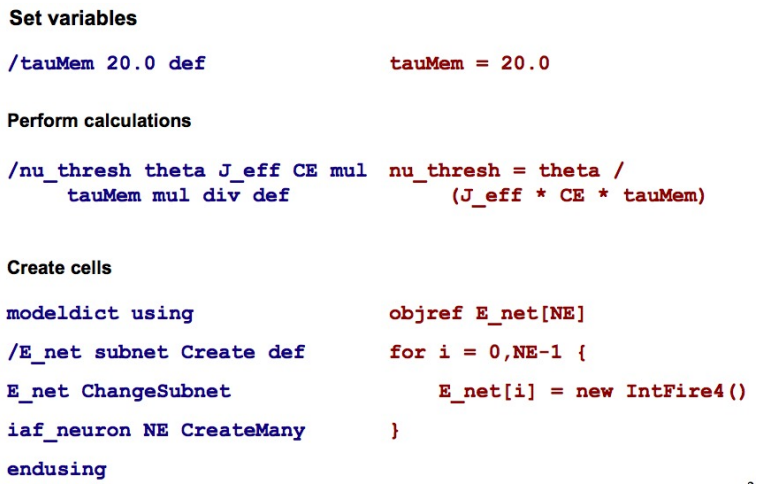

Figure 37: Comparison of SLI code (the native language of NEST) (Left), and NEURON code (Right), for several basic operations, such as creating cells and setting variables.

high frequencies and to small errors of integration made by the simulator. This is where the integration scheme, whether clock driven or event based can make a large difference to the results. The closer in time the spike times are, the higher the synaptic differences. Figure 38 shows the differences for three simulators, with the exact same protocol. Two connected neurons, the first one being just a spike player sending spikes every 2 ms for 600 ms to a target neuron which is a classical integrate and fire neuron, with typical parameters. The synapse between the two neurons is governed by a classical additive STDP rule. As one can see, the exact same short simulation, run with the three simulators, leads to qualitatively similar results, but small differences can already be seen. This is due do rounding errors and differences in integration algorithms, and to how delays are handled by those simulators. One needs to be aware of this, to be sure computational results are not too closely linked with the hidden implementation scheme behind.

To underline these differences, we constructed a toy network of 100 integrate and fire neurons with exponential current based synapses, with an all-to-all connection scheme. The exact detail of the cell parameters is unimportant. The only important point is that synapses within the network are plastic, subject to an additive STDP rule with classical parameters. Every neuron received a particular Poisson input at a fixed frequency, and we run the exact same simulation on three distinct simulators supported by PyNN, taking care that seeds for the Poisson inputs are the same between several runs, so that results should, in theory, be identical between simulators. The simulation was run for 20 s, in order to appreciate eventual divergence between simulators. The neurons were strongly coupled at the beginning, and initial weights were set to 0.5 nS, while w_{max}, for the STDP rule, was set to 1 nS. With NEST, the average frequency of these cells was 20 Hz, 14 Hz with NEURON and 13 Hz with PCSIM. As one can see in Figure 39, some structure emerges in the final connectivity matrix due to the particular patterns imposed in the external spikes trains. It is important to have this in mind while playing with STDP in too small-sized networks. The sensitivity of the rule to small errors in spike times is amplified, and fine structure may be simulator dependent. A more exhaustive analysis of the significativity of these differences should be perfomed.

Benchmarks PyNN also allows benchmarking of simulators, in term of simulation and building times for the exact same network. Of course, these comparisons would require,

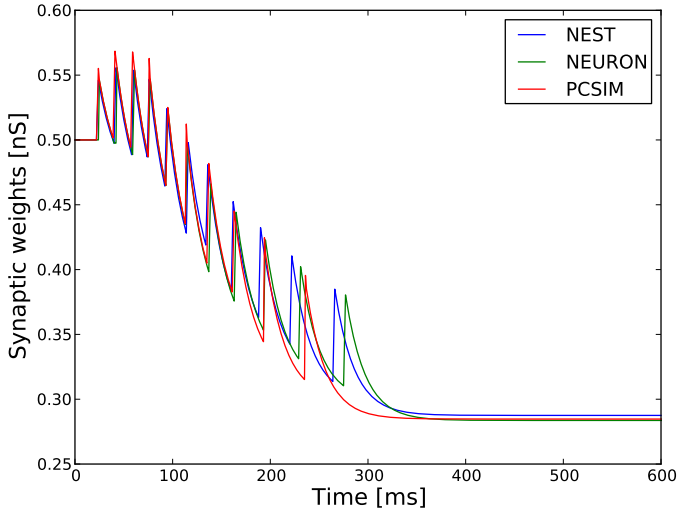

Figure 38: Modification, under an additive STDP rule, of a plastic synapse for an integrate and fire neuron receiving input spike trains at a fixed frequency during 600 ms. The exact same code is launched with several simulators, and the exact spiking patterns of this post-synaptic neuron are different.

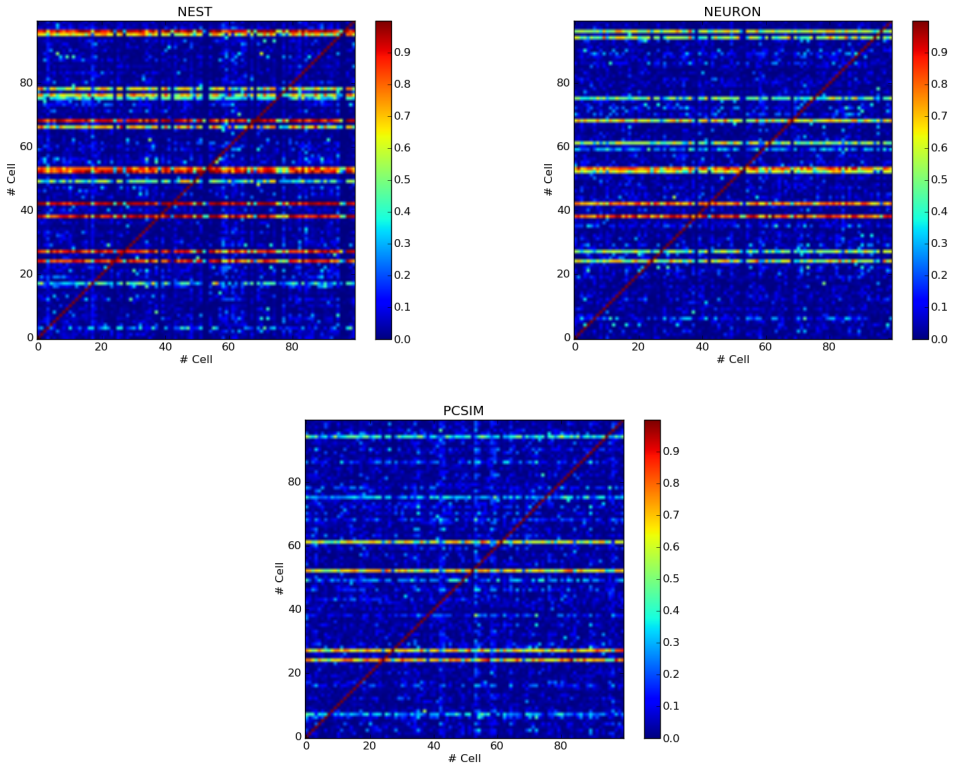

Figure 39: Connectivity matrices after 20 s of simulation, for three neuronal network simulators, of the exact same network with plastic synapses. The network is made of 100 neurons, wired with an all to all connection scheme.

to be fair, that the backends are well-coded and optimized for each simulator. This is currently, in PyNN, only the case for NEST and NEURON. The PCSIM and Brian backends

Simulator	Building Time	Simulation time	Excitatory Rate	Inhibitory rate
NEURON	102 s	624 s	78 Hz	77 Hz
NEST	14 s	63 s	78 Hz	77 Hz
PCSIM	(2000 s)	40 s	73 Hz	73 Hz

Table 5: Comparison of the simulation and building times, with the exact same network, between several simulators. The network is a classical balanced random network of 5000 neurons, as in Brunel (2000). Times in parentheses indicates that they are clearly not optimized (loops are performed in Python). The differences in the firing rates of the populations are due to the integration procedures used by the simulators, since the connectivity is identical.

Simulator	Building Time	Simulation time	Excitatory Rate	Inhibitory rate
NEURON	83 s	489 s	3 Hz	3 Hz
NEST	22 s	22 s	3 Hz	3 Hz
PCSIM	(>1500 s)	9 s	3 Hz	3 Hz
BRIAN	(>1500 s)	12 s	3 Hz	3 Hz

Table 6: Comparison of the simulations and building times, with the exact same network, between several simulators. The network is a classical balanced random network of 10000 neurons, as in Vogels and Abbott (2005) with current based synapses. Times in parentheses indicates that they are clearly not optimized (loops are performed in Python). The differences in the firing rates of the populations are due to the integration procedures used by the simulators.

are working, but still need to be optimized. Simulation times can still be compared because they are only simulator dependent: PyNN can only slow down the building time, if loops or generic functions are used when simulator specific low level functions exist, for example to connect randomly groups of neurons, or set cell parameters. If we take a classical balanced random network, such as the one in Brunel (2000), made with two populations of excitatory (4000 cells) and inhibitory (1000 cells) neurons, connected with a probability $\varepsilon = 10\%$, and classical parameters to be in a asynchronous irregular regime, we have the following results in Table 5 for 1 s of simulation.

Similar benchmarks can be made on a larger and more diluted network, such as the one of Vogels and Abbott (2005), with 10000 neurons and $\varepsilon = 2\%$ (see Table 6 for current based synapses, and Table 7 for conductance based synapses), for 2 s of simulation.

Note that in all cases, the exact same random networks are created, with the same connectivity matrix. The main results here is the simulation times, because as already said, building times are slower in Brian and PCSIM because the backends do not provide optimized $C++$ functions to speed up the wiring procedures (they do, but PyNN does not yet use them). NEST, BRIAN and especially PCSIM are clearly faster than NEURON, because they are dedicated and optimized for such simulations. However, addition of a new model or a new synapse type may be less user friendly and documented than for NEURON. Differences observed in the firing rates are only due to the intrinsic resolution scheme used by the simulators, and this is interesting to note than PCSIM displays small but significant difference with the two others.

Simulator	Building Time	Simulation time	Excitatory Rate	Inhibitory rate
NEURON	83 s	678 s	9 Hz	10 Hz
NEST	22 s	187 s	9 Hz	10 Hz
PCSIM	(>1500 s)	28 s	10 Hz	10 Hz
BRIAN	(>1500 s)	15 s	10 Hz	10 Hz

Table 7: Comparison of the simulations and building times, with the exact same network, between several simulators. The network is a classical balanced random network of 10000 neurons, as in Vogels and Abbott (2005) with conductance based synapses. Times in parenthesis indicates that they are clearly not optimized (loops are performed in Python). The differences in the firing rates of the populations are due to the integration procedures used by the simulators.

21 PyNN: A Common Interface for Neuronal Network Simulators.

21.1 Main results

In the following article, I contribute to the PyNN design of the API, and I have coded part of the implementation of the NEURON and the NEST backends.

frontiers in
NEUROINFORMATICS

ORIGINAL RESEARCH ARTICLE
published: 27 January 2009
doi: 10.3389/neuro.11.011.2008

PyNN: a common interface for neuronal network simulators

Andrew P. Davison[1]*, Daniel Brüderle[2], Jochen Eppler[3,4], Jens Kremkow[5,6], Eilif Muller[7], Dejan Pecevski[8], Laurent Perrinet[6] and Pierre Yger[1]

[1] *Unité de Neurosciences Intégratives et Computationelles, CNRS, Gif sur Yvette, France*
[2] *Kirchhoff Institute for Physics, University of Heidelberg, Heidelberg, Germany*
[3] *Honda Research Institute Europe GmbH, Offenbach, Germany*
[4] *Berstein Center for Computational Neuroscience, Albert-Ludwigs-University, Freiburg, Germany*
[5] *Neurobiology and Biophysics, Institute of Biology III, Albert-Ludwigs-University, Freiburg, Germany*
[6] *Institut de Neurosciences Cognitives de la Méditerranée, CNRS, Marseille, France*
[7] *Laboratory of Computational Neuroscience, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland*
[8] *Institute for Theoretical Computer Science, Graz University of Technology, Graz, Austria*

Edited by:
Rolf Kötter, Radboud University Nijmegen, The Netherlands

Reviewed by:
Graham Cummins, Montana State University, USA
Fred Howell, Textensor Limited, UK

****Correspondence:***
Andrew Davison, UNIC, Bât. 32/33, CNRS, 1 Avenue de la Terrasse, 91198 Gif sur Yvette, France.
e-mail: andrew.davison@unic.cnrs-gif.fr

Computational neuroscience has produced a diversity of software for simulations of networks of spiking neurons, with both negative and positive consequences. On the one hand, each simulator uses its own programming or configuration language, leading to considerable difficulty in porting models from one simulator to another. This impedes communication between investigators and makes it harder to reproduce and build on the work of others. On the other hand, simulation results can be cross-checked between different simulators, giving greater confidence in their correctness, and each simulator has different optimizations, so the most appropriate simulator can be chosen for a given modelling task. A common programming interface to multiple simulators would reduce or eliminate the problems of simulator diversity while retaining the benefits. PyNN is such an interface, making it possible to write a simulation script once, using the Python programming language, and run it without modification on any supported simulator (currently NEURON, NEST, PCSIM, Brian and the Heidelberg VLSI neuromorphic hardware). PyNN increases the productivity of neuronal network modelling by providing high-level abstraction, by promoting code sharing and reuse, and by providing a foundation for simulator-agnostic analysis, visualization and data-management tools. PyNN increases the reliability of modelling studies by making it much easier to check results on multiple simulators. PyNN is open-source software and is available from http://neuralensemble.org/PyNN.

Keywords: Python, interoperability, large-scale models, simulation, parallel computing, reproducibility, computational neuroscience, translation

INTRODUCTION

Science rests upon the three pillars of open communication, reproducibility of results and building upon what has gone before. In these respects, computational neuroscience ought to be in a good position, since computers by design excel at repeating the same task without variation, as many times as desired: reproducibility of computational results ought, then, to be a trivial task. Similarly, the Internet enables almost instantaneous transmission of research materials, i.e. source code, between labs.

However, in practice this theoretical ease of reproducibility and communication is seldom achieved outside of a single lab and a time frame of a few months or years. While a given scientist may easily be able to reproduce a result obtained a few months ago, precisely reproducing a result obtained several years ago is likely to be rather more difficult, and the general experience seems to be that reproducing the results of others is both difficult and time consuming: very many published papers lack sufficient detail to rebuild a model from scratch, and typographic errors are common.

Having available the source code of the model greatly improves the situation, but here still there are numerous barriers to reproducibility and to building upon previously published models. One is that source code can rapidly go out of date as computer architectures, compiler standards and simulators develop. Another is that model source code is often not written with reuse and extension in mind, and so considerable rewriting to modularize the code is necessary. Probably the most important barrier is that code written for one simulator is not compatible with any other simulator.

Although many computational models in neuroscience are written from the ground up in a general purpose programming language such as C++ or Fortran, probably the majority use a special purpose simulator that allows models to be expressed in terms of neuroscience-specific concepts such as neurons, ion channels, synapses; the simulator takes care of translating these concepts into a system of equations and of numerically solving the equations. A large number of such simulators are available (reviewed in Brette et al., 2007), mostly as open-source software, and each has its own programming language, configuration syntax and/or graphical interface, which creates considerable difficulty in translating models from one simulator to another, or even in understanding someone else's code, with obvious negative consequences for communication between investigators, reproducibility of others' models and building on existing models.

However, the diversity of simulators also has a number of positive consequences: (i) it allows cross-checking – the probability of two

different simulators having the same bugs or hidden assumptions is very small; (ii) each simulator has a different balance between efficiency (how fast the simulations run), flexibility (how easy it is to add new functionality; the range of models that can be simulated), scalability (for parallel, distributed computation on clusters or supercomputers), and ease of use, so the most appropriate can be chosen for a given task.

Addressing the problems associated with an ecosystem of multiple simulators while retaining the benefits would greatly increase the ease of reproducibility of computational models in neuroscience and hence make it easier to verify the validity of published models and to build upon previous work.

There are at least two possible (and complementary) approaches to this. One is to enable direct, efficient communication between different simulators at run-time, allowing different components of a model to be simulated on different simulators (Ekeberg and Djurfeldt, 2008). This approach addresses the problem of building a model from diverse components, but still leaves the problem of having to use different programming languages, and does not enable straightforward cross-checking. The other approach is to develop a system for model specification that is simulator-independent. Translation then only has to be done once for each simulator and not once for each model.

Here we can take advantage of the recent, rapid emergence of the Python programming language as an alternative interface to several of the more widely-used simulators. Thus, for example, both NEURON and NEST may be controlled either via their original, native interpreter (Hoc and SLI, respectively) or via Python. More recent simulators (e.g. PCSIM, Brian) have Python as the only available scripting language. This widespread adoption of Python is probably due to a number of factors, including the powerful data structures, clean and expressive syntax, extensive library, maturity of tools for numerical analysis and visualization (allowing use of a single language for the entire modelling workflow from simulation to analysis to graphing), and the ease-of-use of Python as a glue language which allows computation-intensive code written in a low-level language such as C to be transparently accessed within high-level Python code.

Python alone does not address the translation problem (although it does make the translation process easier, since at least simple data structures such as lists and arrays are the same for each simulator), since neuroscience-specific concepts are still expressed differently. However, it is now possible to define a simulator-independent Python interface for neuronal network simulators and to implement automatic translation to any Python-enabled simulator. We have designed and implemented such an interface, PyNN (pronounced "pine"). In this paper we describe its design, concepts, implementation and use. We do not attempt here to provide a complete user guide – this may be found online at http://neuralensemble.org/PyNN.

DESIGN GOALS

When designing and implementing a common simulator interface, the following goals should be taken into account. These are the goals we have kept in mind when designing and implementing the PyNN interface, but they are equally applicable to any other such interface.

Write the code for a model once, run it on any supported simulator or hardware device *without modification*. This is the primary design goal for PyNN.

Support a high-level of abstraction. For example, it is often preferable to deal with a single object representing a population of neurons than to deal with all the individual neurons directly. Each single neuron can be accessed when necessary, but in many cases the population is the more useful abstraction. The advantages of this approach are that (i) it is easier to maintain a conceptual idea of the model, without being distracted by implementation details, and (ii) the internal implementation of an object can be optimized for speed, parallelization or memory requirements without changing the interface presented to the user.

Support any feature provided by at least two supported simulators. The aim is to strike a balance between supporting all features of all simulators (unfeasible) and supporting only the subset of features common to all simulators (overly restrictive).

Allow mixing of PyNN and native simulator code. PyNN should not limit the range of models that can be implemented. Following the two-simulator rule, above, there will be things that are possible in one simulator and not in any other. Although a model implementation consisting of 100% PyNN is the best scenario for running on multiple simulators, an implementation with 50% PyNN code will be easier to convert between simulators than one with no PyNN code.

Facilitate porting of models between simulators. PyNN changes the process of porting a model between simulators from all-or-nothing, in which the validity of the translated model cannot be tested until the entire translation is complete, to an incremental approach, in which the native code is gradually replaced by simulator-independent code. At each stage, the hybrid code remains runnable, and so it is straightforward to verify that the model behaviour has not been changed.

Minimize dependencies, to make installation as simple as possible and maximize flexibility. There are no visualization and few data analysis tools built-in to PyNN, which means the user can use any such tools they wish.

Present a consistent interface on output as well as on input. The formats used for simulation outputs are consistent across simulator back-ends, making it a stable base upon which to build more complex systems of simulation control, data-analysis and visualization.

Prioritize compatibility over optimizations, but allow compatibility-breaking optimizations to be selected by a deliberate choice of the user (e.g. the `compatible_output` flag of the various `print()` methods is `True` by default, but can be set to `False` to get potentially-faster writing of data to file).

API Versioning. The PyNN API will inevitably evolve over time, as more simulators are supported and to take account of the preferences of the community of users. To ensure backwards compatibility, the API should be versioned so that the user can indicate which version was used for a particular implementation. Note that the examples given in this paper use version 0.4 of the API.

Transparent parallelization. Code that runs on a single processor should run on multiple processors (using MPI) without changes.

Some of these goals are somewhat contradictory: for example, having a high level of abstraction and making porting easy.

Reconciling this particular pair of goals has led to the presence in PyNN of both a high-level, object-oriented interface and a low-level, procedural interface that is more similar to the interface of many existing simulators. These will be discussed further below.

USAGE EXAMPLES

Before describing in detail the concepts underlying the PyNN interface, we will work through some examples of how it is used in practice: first a simple example using the low-level, procedural interface and then a more complex example using the high-level, object-oriented interface.

For the simple example, we will build a network consisting of a single integrate-and-fire (IF) cell receiving spiking input from a Poisson process.

First, we choose which simulator to use by importing the relevant module from PyNN:

```
>>> from pyNN.neuron import *
```

If we wanted to use PCSIM, we would just import `pyNN.pcsim`, etc. Whichever simulator back-end we use, none of the code below would change.

Next we set global parameters of the simulator:

```
>>> setup(timestep=0.1, min_delay=2.0)
```

Now we create two cells: an IF neuron with synapses that respond to a spike with a step increase in synaptic conductance, which then decays exponentially, and a "spike source", a simple cell that emits spikes at predetermined times but cannot receive input spikes.

```
>>> ifcell = create(IF_cond_exp,
...                 {'i_offset': 0.11,
...                  'tau_refrac': 3.0,
...                  'v_thresh' : -51.0})
>>> times = map(float, range(5,105,10))
>>> source = create(SpikeSourceArray,
...                 {'spike_times': times})
```

Behind the scenes, the `create()` function translates the standard PyNN model name, `IF_cond_exp` in this case, into the model name used by the simulator, `Standard_IF` for NEURON, `iaf_cond_exp` for NEST, for example and also translates parameter names and units into simulator-specific names and units. To take one example, the `i_offset` parameter represents the amplitude of a constant current injected into the cell, and is given in nanoamps. The equivalent parameter of the NEST `iaf_cond_exp` model has the name `I_e` and units of picoamps, so PyNN both converts the name and multiplies the numerical value by 1000 when running with NEST. Standard cell models and automatic translation are discussed in more detail in the next section.

The `create()` function returns an `ID` object, which provides access to the parameters of the cell models, e.g.:

```
>>> ifcell.tau_refrac
3.0
>>> ifcell.tau_m = 12.5
>>> ifcell.get_parameters()
{'tau_refrac': 3.0, 'tau_m': 12.5,
 'e_rev_E': 0.0, 'i_offset': 0.11,
 'cm': 1.0, 'e_rev_I': -70.0,
 'v_init': -65.0, 'v_thresh': -51.0,
 'tau_syn_E': 5.0, 'v_rest': -65.0,
 'tau_syn_I': 5.0, 'v_reset': -65.0}
```

Having created the cells, we connect them with the `connect()` function:

```
>>> connect(source, ifcell, weight=0.006,
...         synapse_type='excitatory', delay=2.0)
```

Now we tell the system what variable or variables to record, run the simulation and finish.

```
>>> record_v(ifcell, 'ifcell.dat')
>>> run(200.0)
>>> end()
```

The result of running the above model is shown in **Figure 1**, which also shows the degree of reproducibility obtainable between different simulators for such a simple network.

The low-level, procedural interface, using the `create()`, `connect()` and `record()` functions, is useful for simple models or when porting an existing model written in a different language that uses the create/connect idiom. For larger, more complex networks we have found that an object-oriented approach, with a higher-level of abstraction, is more effective, since it both clarifies the conceptual structure of the model, by hiding implementation details, and allows behind-the-scenes optimizations.

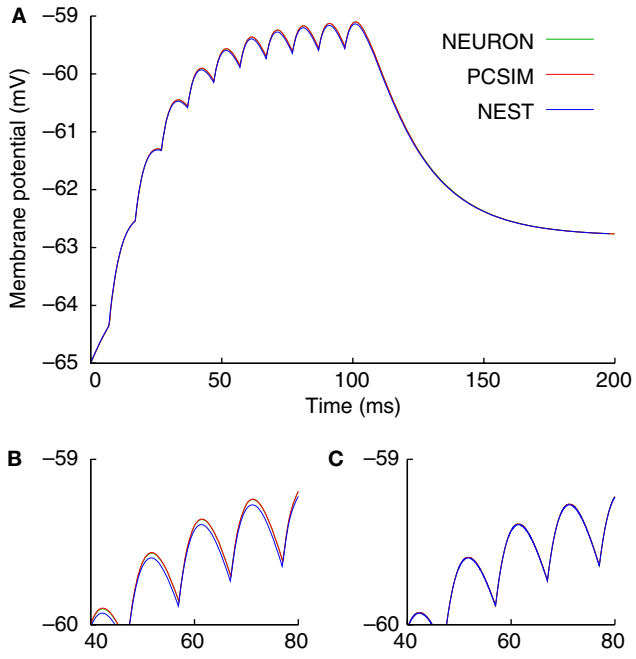

FIGURE 1 | Results of running first example given in the text, with NEURON, NEST and PCSIM as back-end simulators. (A) Entire membrane potential trace with integration time-step 0.1 ms. **(B)** Zoom into a smaller region of the trace, showing small numerical differences between the results of the different simulators. **(C)** Results of a simulation with integration time-step 0.01 ms, showing greatly reduced numerical differences.

To illustrate the high-level, object-oriented interface we turn now from the simple example of a few neurons to a more complex example: a network of several thousand excitatory and inhibitory neurons that displays self-sustained activity (based on the "CUBA" model of Vogels and Abbott (2005), and reproducing the benchmark model used in Brette et al. (2007)). This still is not a particularly complicated network, since it has only two cell types, no spatial structure and no heterogeneity of neuronal or connection properties, but in demonstrating how building such a network becomes trivial using PyNN we hope to convince the reader that building genuinely complex, structured and heterogeneous networks becomes manageable.

Again, we begin by choosing which simulator to use. We also import some classes from PyNN's `random` module.

```
>>> from pyNN.nest2 import *
>>> from pyNN.random import (NumpyRNG,
...                          RandomDistribution)
```

We next specify the parameters of the neuron model (the same model and same parameters are used for both excitatory and inhibitory neurons).

```
>>> cell_params = {
...     'tau_m':     20.0,  'tau_syn_E':   5.0,
...     'cm':         0.2,  'tau_syn_I':  10.0,
...     'v_rest':   -49.0,  'v_reset':   -60.0,
...     'v_thresh': -50.0,  'tau_refrac':  5.0
...     }
```

Parameters with dimensions of voltage are in millivolts, time in milliseconds and capacitance in nanofarads. The units convention is discussed further in the next section.

We now initialize the simulation, this time accepting the default values for the global parameters.

```
>>> setup()
```

Now, rather than creating each cell separately, we just create a `Population` object for each different type of cell:

```
>>> pE = Population(4000, IF_cond_exp,
...                 cell_params,
...                 label="Excitatory")
>>> pI = Population(1000, IF_cond_exp,
...                 cell_params,
...                 label="Inhibitory")
```

By default, all cells of a given Population are created with identical parameters, but these can be changed afterwards. Here we wish to randomize the value of the membrane potential at the start of the simulation to values between −50 and −70 mV.

```
>>> unif_distr = RandomDistribution('uniform',
...                                 [-50,-70])
>>> pE.randomInit(unif_distr)
>>> pI.randomInit(unif_distr)
```

`randomInit()` is a convenience method for randomizing the initial membrane potential. For the more general case of randomizing any cell parameter use `rset()`.

Just as individual neurons are encapsulated within `Populations`, connections between neurons are encapsulated within `Projections`. To create a `Projection` object, we need to specify how the neurons will be connected, either via an algorithm or via an explicit list. Different algorithms are encapsulated in different `Connector` classes, e.g. `FixedProbabilityConnector`, `AllToAllConnector`. An explicit list of connections can be provided via a `FromListConnector` or a `FromFileConnector`.

```
>>> FPC = FixedProbabilityConnector
>>> exc_conn = FPC(0.02, weights=0.004,
...                delays=0.1)
>>> inh_conn = FPC(0.02, weights=0.051,
...                delays=0.1)
```

Note that weights are in microsiemens and delays in milliseconds. Where the delay is not specified, the global minimum delay specified in the `setup()` function is used. Here we set all weights and delays of a `Projection` to the same value, but it is equally possible to pass the constructor a `RandomDistribution` object, as we did above for the initial membrane potential, or an explicit list of values.

To create a `Projection`, we need to specify the pre- and post-synaptic `Populations`, a `Connector` object, and a synapse type. The standard IF cells each have two synapse types, "`excitatory`" and "`inhibitory`". User-defined models can use arbitrary names, e.g. "`AMPA`", "`NMDA`".

```
>>> e2e = Projection(pE, pE, exc_conn,
...                  target='excitatory')
>>> e2i = Projection(pE, pI, exc_conn,
...                  target='excitatory')
>>> i2e = Projection(pI, pE, inh_conn,
...                  target='inhibitory')
>>> i2i = Projection(pI, pI, inh_conn,
...                  target='inhibitory')
```

Having constructed the network, we now need to instrument it, using the `record()` (for recording spikes) and `record_v()` (membrane potential) methods of the `Population` objects. Here we choose to record spikes from 1000 of the excitatory neurons (chosen at random) and all of the inhibitory neurons, and to record the membrane potential of two specific excitatory neurons. We then run the simulation for 1000 ms.

```
>>> pE.record(1000)
>>> pI.record()
>>> pE.record_v([pE[0], pE[1]])
>>> run(1000.0)
```

After running the simulation, we can access the results or write them to file.

```
>>> pI.getSpikes()[:5]
array([[ 715. ,     1.5],
       [ 609. ,     1.6],
       [ 708. ,     1.7],
       [ 796. ,     1.7],
       [  34. ,     1.8]])
```

```
>>> pE.get_v()[:5]
array([[  0.   ,     0.1  ,  -55.073],
       [  1.   ,     0.1  ,  -50.163],
       [  0.   ,     0.2  ,  -55.098],
       [  1.   ,     0.2  ,  -50.212],
       [  0.   ,     0.3  ,  -55.122]])
>>> end()
```

The results of running simulations of the above network with two different simulator back-ends are shown in **Figure 2**.

PRINCIPAL CONCEPTS

To achieve the goal of "*write the code for a model once, run it on any supported simulator without modification*" requires (i) a common interface, (ii) neuron and synapse models that are standardized across simulators, (iii) consistent handling of physical units, (iv) consistent handling of (pseudo-)random numbers. To achieve the twin goals of supporting a high-level of abstraction and facilitating porting of models between simulators requires both an object-oriented and a procedural interface. The implementation of all these requirements is described in more depth in the following. We also illustrate the mixing of PyNN and native simulator code, and how PyNN can support features that are found in only a single simulator back-end, by describing support for multi-compartmental models.

STANDARD CELL MODELS

A fundamental concept in PyNN is the cell type – a given model of a neuron, representable by a set of equations, and comprising sub-threshold behaviour, spiking mechanism and post-synaptic response. The public interface of a cell type is mainly defined by its parameters. Different neurons of the same cell type may have very different behaviour if they have different values of the parameters. For example, the Izhikevich model (Izhikevich, 2003), can reproduce a wide range of spiking patterns, from fast-spiking through regular spiking to multiple types of bursting, depending on the

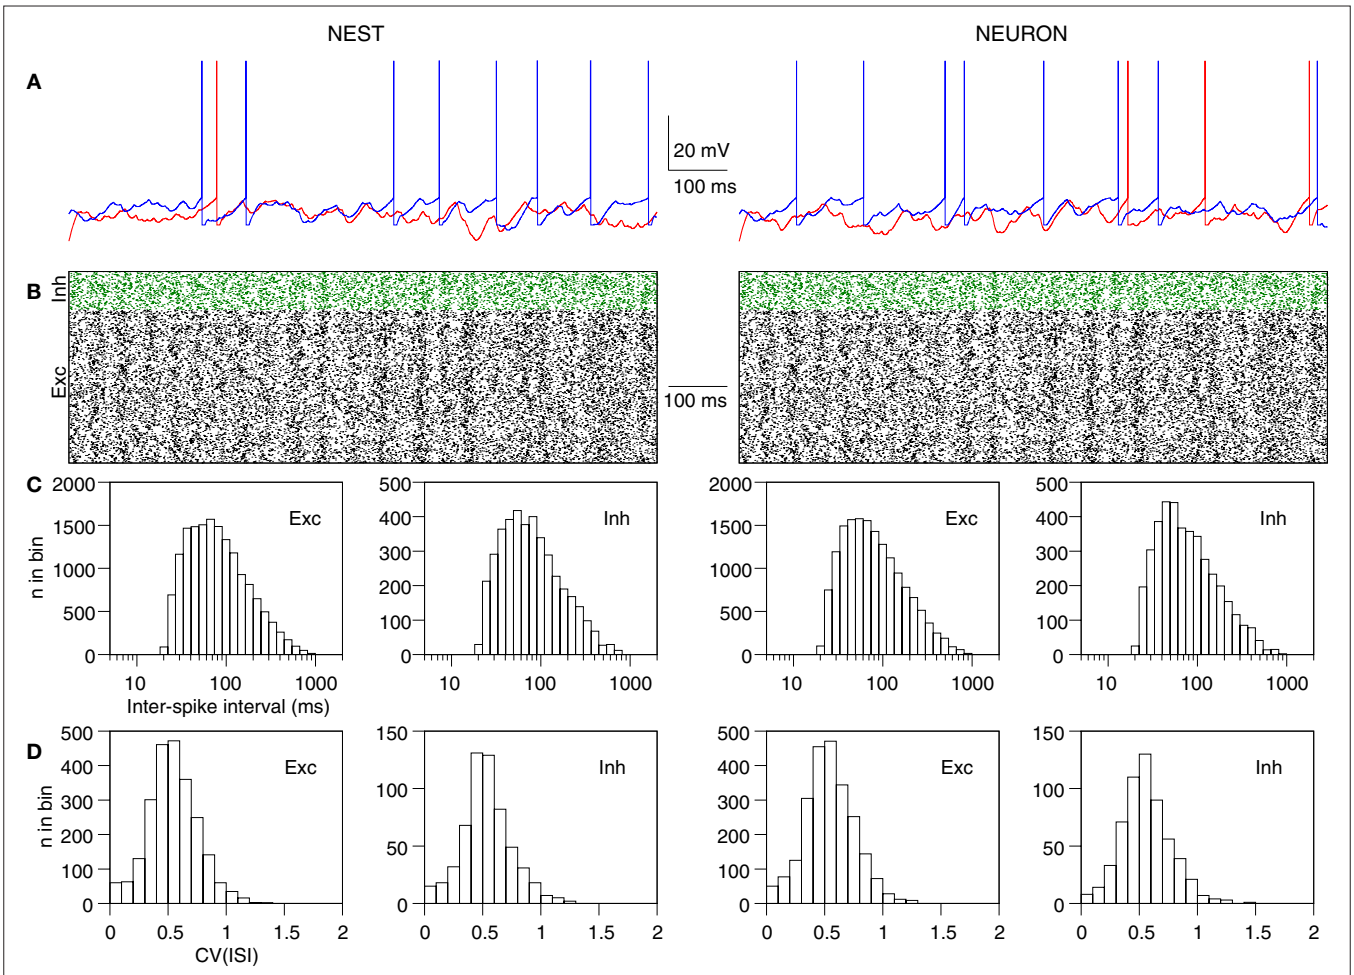

FIGURE 2 | Results of running the second example given in the text, with NEURON and NEST as back-end simulators. Note that the network connectivity and initial conditions were identical in the two cases. **(A)** Membrane potential traces for two excitatory neurons. Note that the NEST and NEURON traces are very similar for the first 50 ms, but after that diverge rapidly due to the effects of network activity, which amplifies the small numerical integration differences. **(B)** Spiking activity of excitatory (black) and inhibitory (green) neurons. Each dot represents a spike and each row of dots a different neuron. All 5000 neurons are shown. **(C)** Distribution of pooled inter-spike intervals (ISIs) for excitatory and inhibitory neurons. **(D)** Distribution over neurons of the coefficient of variation of the ISI [CV(ISI)].

parameter values chosen. A cell type is therefore a model type rather than a biologically defined cell type (such as "Layer V pyramidal neuron", for example).

When using a given simulator back-end, PyNN can work with any cell type that is supported by that simulator. In this case, the cell type is generally represented by a string, holding a model name that is meaningful for that simulator, e.g. "`iaf_neuron`" in NEST.

Of course, such a cell type will only work with one simulator. To create a model that will run on different simulators requires you to use one of PyNN's built-in, standard cell models, each represented by a sub-class of the `StandardCell` class. The models provided by PyNN include various simple IF models, the Izhikevich-like adaptive exponential IF model (Brette and Gerstner, 2005), a single-compartment neuron with Hodgkin–Huxley sodium and potassium channels, and various models that emit spikes (e.g. according to a Poisson process) but cannot receive them.

The `StandardCell` class contains machinery for translating model names, parameter names and parameter units between PyNN standardized values and simulator-specific values. This is particularly useful when the underlying simulators use different unit systems or different parameterizations of the same set of equations, e.g. when one simulator expects the membrane time constant and another the membrane leak conductance. An example of the translations performed by PyNN is given in **Table 1**.

Currently, all the standard cell types are single-compartment or point neuron models, since PyNN currently supports only one simulator for multi-compartmental models (NEURON). Further details on using multi-compartmental models with PyNN's NEURON back-end are given below. We plan in future to allow specifying multi-compartmental cell types using a NeuroML description (Crook et al., 2005).

UNITS

As is clear from the previous section, each simulator back-end has its own convention for which units to use for which physical quantities. The exception to this is Brian, which has a system for explicitly specifying units and for checking that equations are dimensionally consistent. In the future, we plan to adopt Brian's system for PyNN, but for now we have chosen to use a convention, which is similar to that of NEURON and NEST in that the units are those that tend to be used by experimental physiologists. An alternative would have been the convention used by PCSIM (and also by the GENESIS simulator) of using pure SI units with no prefixes. The advantage of the latter convention is that there is no need for checking equations for dimensional consistency. The disadvantage is that numerical values in such a system are often very large or very small, and hence the human intuition for reasonable and unreasonable parameter values is mostly lost.

Irrespective of the relative merits of different conventions, the most important thing is that PyNN now provides a single convention which is valid across simulators. In detail, the convention is as follows: voltage – mV, current – nA, conductance – μS, time – ms, capacitance – nF.

STANDARD SYNAPSE MODELS

In PyNN, the shape and time-course of the elementary post-synaptic current or conductance change in response to a pre-synaptic spike are considered to be a part of the post-synaptic neuron model, while all other properties of a synaptic connection, notably its weight (the peak current or conductance of the synaptic response), delay (for point models, this implicitly includes axonal propagation, chemical transmission and dendritic propagation; more morphologically and/or biophysically detailed models may model explicitly some or all of these sources of delay), and short- and long-term plasticity, are considered to depend on both pre- and post-synaptic neurons, and so are encapsulated in the concept of "synapse type" that mirrors the "cell type" discussed above.

The default type of synaptic connection in PyNN is static, with fixed synaptic weights. To model dynamic synapses, for which the synaptic weight (and possibly other properties, such as rise-time) varies depending on the recent history of post- and/or pre-synaptic activity, we use the same idea as for neurons, of standardized, named models that have the same interface and behaviour across simulators, even if the underlying implementation may be very different.

Where the approach for dynamic synapses differs from that for neurons is that we attempt a greater degree of compositionality, i.e. we decompose models into a number of components, for

Table 1 | Comparison of parameter names and units for different implementations of a leaky integrate-and-fire model with a fixed firing threshold and current-based, alpha-function synapses. This model is called `IF_curr_alpha` in PyNN, `iaf_psc_alpha` in NEST, `LIFCurrAlphaNeuron` in PCSIM and `StandardIF` in NEURON (this is a model template distributed with PyNN and is not in the standard NEURON distribution). Manual conversion of names and units is straightforward but error-prone and time-consuming. PyNN takes care of such conversions transparently.

Parameter	PyNN		NEST		NEURON		PCSIM	
Resting membrane potential	`v_rest`	mV	`E_L`	mV	`v_rest`	mV	`Vresting`	V
Reset membrane potential	`v_reset`	mV	`V_reset`	mV	`v_reset`	mV	`Vreset`	V
Membrane capacitance	`cm`	nF	`C_m`	pF	`CM`	nF	`Cm`	F
Membrane time constant	`tau_m`	ms	`tau_m`	ms	`tau_m`	ms	`taum`	s
Refractory period	`tau_refrac`	ms	`t_ref`	ms	`t_refrac`	ms	`Trefrac`	s
Excitatory synaptic time constant	`tau_syn_E`	ms	`tau_syn_ex`	ms	`tau_e`	ms	`TauSynExc`	s
Inhibitory synaptic time constant	`tau_syn_I`	ms	`tau_syn_in`	ms	`tau_i`	ms	`TauSynInh`	s
Spike threshold	`v_thresh`	mV	`V_th`	mV	`v_thresh`	mV	`Vthresh`	V
Injected current amplitude	`i_offset`	nA	`I_e`	pA	`i_offset`	nA	`Iinject`	A

example for short-term and long-term dynamics, or for the timing-dependence and the weight-dependence of STDP rules, that can then be composed in different ways.

The advantage of this is that if we have n different models for component A and m models for component B, then we require only $n + m$ models rather than $n \times m$, which had advantages in terms of code-simplicity and in shorter model names. The disadvantage is that not all combinations may exist, if the underlying simulator implements composite models rather than using components itself: in this situation, PyNN checks whether a given composite model AB exists for a given simulator and raises an Exception if it does not. The composite approach may be extended to neuron models in future versions of the PyNN interface depending on the experience with composite synapse models.

Currently only a single model exists in PyNN for the short-term plasticity component, the Tsodyks–Markram model (Markram et al., 1998). For long-term plasticity there is a spike-timing-dependent plasticity STDP component, which itself is composed of separate timing-dependence and weight-dependence components.

LOW-LEVEL, PROCEDURAL INTERFACE

We refer to the procedural interface as "low-level" because it deals with a lower level of abstraction – individual neurons and individual synapses – than the object-oriented interface. The procedural interface consists of the functions `create()`, `connect()`, `set()`, `record()` (for recording spikes) and `record_v()` (for recording membrane potential). Each of these functions operates on, or returns, either individual cell `ID` objects or lists of such objects. As was described in the Usage Examples section, as well as being passed around as arguments, the `ID` object may be used for accessing/modifying the parameters of individual neurons, and takes care of parameter translation using the `StandardCell` mechanisms described above.

It is possible to some extent to mix the low-level and high-level interfaces. For example, it is possible to access individual neurons within a `Population` as `ID` objects and then use the `connect()` function to connect them, instead of using a `Projection` object.

Why have both a low-level and high-level interface? Having both is a potential source of confusion for users and is definitely a maintenance burden for developers. The main reason is to support the use of PyNN as a porting tool. The majority of neuronal network models using existing simulators use a procedural approach, and so conversion to PyNN is easier if PyNN supports the same approach. In addition, when developing a PyNN interface for a simulator, or for neuromorphic hardware, that deals primarily with individual cells and synaptic connections, it is easier to implement only the low-level interface, since the high-level interface can be built upon it.

HIGH-LEVEL, OBJECT-ORIENTED INTERFACE

Object-oriented programming has been used for many years in computer science as a method for reducing program complexity. As the ambition and scope of large-scale, biologically detailed neuronal network modelling increases, reducing program complexity will become more and more critical, as the limiting factor in computational neuroscience becomes the productivity of the programmer and not the capacity of the computer (Wilson, 2006). It is for this reason that the preferred interface in PyNN for developing new models is an object-oriented one.

The object-oriented interface is built around three main classes:

Population – a group of cells all with the same cell type (model type). It is generally considered that the cells in a `Population` should all represent the same biological cell type, i.e. although parameter values may vary between cells in the group, all cells should have qualitatively the same firing response. This is not enforced, but is a good guideline to follow for producing understandable code. The `Population` class eliminates tedious iteration over lists of neurons and enables more efficient, array-based management of neuron properties.

Projection – the set of connections of a given synapse type between two `Populations`. Creating a Projection requires specifying the pre- and post-synaptic `Populations`, the synapse type, and the algorithm used to determine which neurons connect to which.

Connector – an encapsulation of the connection algorithm used in creating a Projection. Simple examples of such algorithms are "all-to-all", "one-to-one" and "connect-each-pre-and-post-synaptic-cell-with-a-fixed-probability". It is also possible to provide an explicit list of which cells are to be connected to which others. Each algorithm is defined within a subclass of the `Connector` class. PyNN contains a number of such classes, but it is fairly straightforward for a user to define their own algorithms.

In future development of PyNN, we plan to extend the interface to still higher-level abstractions, such as layers, cortical columns, brain areas and inter-areal projections. We also aim to use the high-level interface as a link between spiking network models and more abstract models that do not represent individual neurons, such as mean-field models.

RANDOM NUMBERS

The central nervous system contains many sources of noise, and activity patterns are often sufficiently complex, and possibly chaotic, to make a stochastic representation a reasonable model.

This can become a problem when comparing the behaviour of a given model run on different simulators, since random differences might obscure real inconsistencies between implementations of the model. Similarly, when performing distributed computations on parallel machines, the model behaviour should not depend on the number of processors used (Morrison et al., 2005), and random differences can conceal real differences between the parallel and serial implementations.

For these reasons, it is important to be able to use identical sequences of random numbers in different simulators, and to have the random number used at a particular point in the program execution be independent of which processor it is running on.

Another consideration is that simulations in most cases use only pseudo-random sequences, and low-quality random number generators (RNGs) may have correlations between different elements of the sequence that can significantly affect the qualitative behaviour of a network. Hence it is necessary to be able to test the simulation with different RNGs.

PyNN supports simulator-independent RNGs and use of different generators – currently any of the generators provided by

the numpy package or by the GNU Scientific Library (GSL) can be used.

This is done by wrapping the numpy and GSL RNGs in classes with a common interface. PyNN's random module contains the classes NumpyRNG and GSLRNG, which both have a single method, next(n, distribution, parameters), which returns n random numbers from a distribution of type distribution with parameters parameters, e.g.

```
>>> from pyNN.random import NumpyRNG, GSLRNG
>>> rngN = NumpyRNG(seed=76847376)
>>> rngG = GSLRNG(seed=87548753)
>>> rngN.next()
0.91457981651574294
>>> rngG.next(5)
array([ 0.02518011, 0.79118205, 0.16679516,
...       0.1902914, 0.66204769])
>>> rngN.next(3, 'gamma', [2.0, 0.5])
array([ 0.48903019, 0.63129009, 0.70428452])
>>> rngG.next(distribution='uniform')
0.93618978746235371
```

Since all PyNN code that uses random numbers accesses the RNG classes only through this next() method, a user can substitute their own RNG simply by defining a wrapper class with such a method.

Since very often one wishes to use the same random distribution repeatedly, rather than changing distribution each time, the random module also provides the RandomDistribution class, which is initialized with the distribution name and parameters, and thereafter the next() method is simplified to take a single argument, the number of values to draw from the distribution, e.g.

```
>>> from pyNN.random import (NumpyRNG,
...                          RandomDistribution)
>>> rng = NumpyRNG(seed=8745753)
>>> gamma_distr = RandomDistribution('gamma',
...                                  [2.0, 0.5],
...                                  rng=rng)
>>> gamma_distr.next(3)
array([ 0.72682412, 0.82490159, 1.03882654])
```

Note that NumpyRNG and GSLRNG distributions may not have the same names, e.g. "normal" for NumpyRNG and "gaussian" for GSLRNG, and the arguments may also differ. One of our future plans is to extend the random module in order to harmonize names across RNGs.

MULTI-COMPARTMENTAL MODELS

PyNN currently supports only a single simulator, NEURON, that is suitable for many-compartment models. Given the principle of supporting simulator-independence only for features that are shared by at least two of the supported simulators, and given PyNN's focus on network modelling, PyNN does not provide an API for specifying simulator-independent multi-compartmental models. This is a possible future development – preliminary work has been done on a PyNN interface to the MOOSE simulator (Ray and Bhalla, 2008) – but a more likely path would be to make use of the NeuroML standards for specifying multi-compartmental models. In this scenario, the filename of a NeuroML level 2 file, specifying a single cell type, would be passed as the cellclass argument to the PyNN create() function or Population constructor.

However, since native and PyNN code can be mixed, the pyNN.neuron module already supports simulations with multi-compartmental models. The pre-synaptic compartment whose voltage is watched to trigger synaptic transmission (e.g. axon terminal) can be specified using the source argument to the Projection constructor, and the post-synaptic mechanism specified with the target argument.

DEBUGGING

Should an error occur in a PyNN simulation, a good first step is to re-run it on another simulator back-end and so narrow down the source of the problem to one back-end in particular. Nevertheless, it has proven to be the case that the additional layers of abstraction provided by PyNN sometimes make it harder to track down sources of errors. To counterbalance this, PyNN traps errors coming from the simulator core and employs Python's introspection capabilities to provide additional information about the error context. For example, if an invalid parameter name is provided to a neuron model, the error message lists all the valid parameter names for that model. Furthermore, logging can be switched on via the init_logging() function in the pyNN.utility module, causing detailed information about what the system is doing to be written to file, a valuable resource for tracking down bugs.

IMPLEMENTATION

PyNN is both a definition of a common simulator interface and an implementation of this interface for each supported simulator. PyNN is implemented as a Python package containing a common module, which defines the API and contains functionality common to all simulator back-ends, a random module (described above), and a module for each simulator back-end, as shown in **Figure 3**. Each simulator module separately implements the API, although it can make use of much shared code in common. In most cases, the simulator modules have been implemented by, or in close collaboration with, the simulator developers.

PyNN currently fully supports the following simulators: NEURON (Carnevale and Hines, 2006; Hines and Carnevale, 1997; Hines et al., 2008), NEST (Eppler et al., 2008; Gewaltig and Diesmann, 2007), PCSIM (http://www.lsm.tugraz.at/pcsim/) and Brian (Goodman and Brette, 2008). Support for MOOSE (Ray and Bhalla, 2008) and for export in NeuroML format (Crook et al., 2005) is under development.

PyNN also supports the Heidelberg neuromorphic hardware system (Schemmel et al., 2007). This illustrates a major benefit of the existence of a common neuronal simulation interface: novel simulation or emulation systems do not need to develop their own programming interface, but can benefit from an existing one that guarantees interoperability with existing tools. Using PyNN as the interface to neuromorphic hardware systems provides the possibility of closing the gap between the two domains of numerical simulation and physical emulation, which have so far coexisted rather separately.

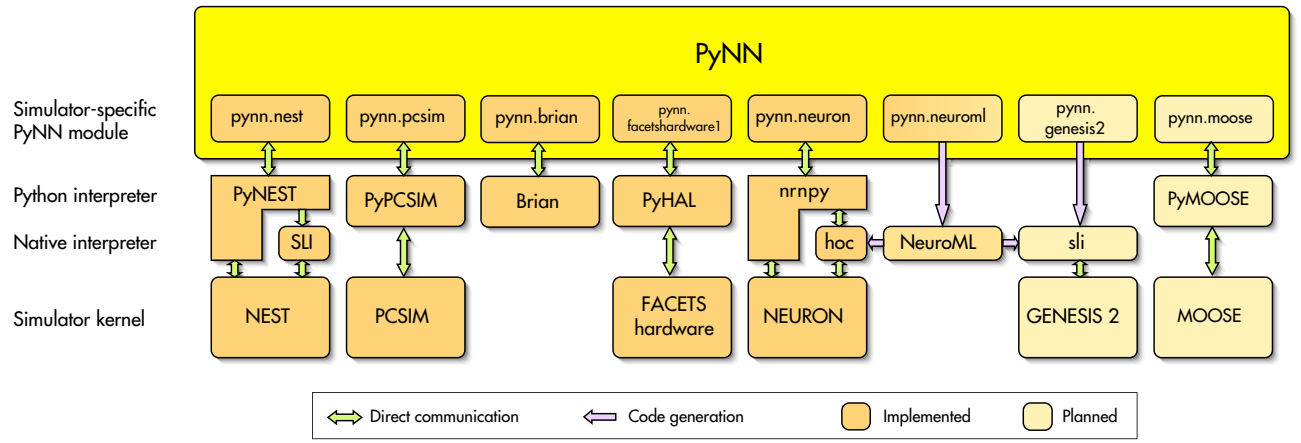

FIGURE 3 | The architecture of PyNN.

LIMITATIONS ON REPRODUCIBILITY

For a given model with a given parameter set run on a given version of a given simulator, it should be possible to exactly reproduce a simulation result, independent of computer architecture (except where this affects the precision of the floating-point representation) or operating system. For parallel systems, results should also be independent of how many threads or processes are used in the computation, although here exact quantitative reproduction is harder to achieve. Reproducibility across different versions of a given simulator is not essential provided the precise version used to generate a given result is specified, but it is of course highly desirable. When running a model on different simulators, exact reproduction is impossible to achieve, except in simple cases, due to round-off errors in floating point calculations. When validating a model implementation by running it on two or more simulators, therefore, what level of reproducibility is achievable, and how can we tell whether any differences are due to round-off error or to implementation errors?

To get a preliminary handle on this problem, we have compared the difference in model activity between two simulators to the difference due to two different initial conditions with the same simulator.

Our test case is the balanced random network, based on Vogels and Abbott (2005), whose implementation was shown above. The activity pattern of this network is very sensitive to initial conditions (chaotic or near-chaotic), and so we cannot use differences in the precise spike pattern to measure reproducibility: we are more interested in the statistical properties of the activity, and so we have chosen to take the distribution of inter-spike intervals (ISIs) of excitatory neurons (see **Figure 2C**) as a measure of network activity.

To measure the difference between the distributions from two different runs we use the Kolmogorov–Smirnov two-sample test. We ran the simulation ten times, each time with a different seed for the RNG used to generate the initial membrane potential distribution, with both NEURON and NEST back-ends. This gave values for the Kolmogorov–Smirnov D-statistic between 0.008 and 0.026 ($n \simeq 19000$) with a mean of 0.015, with associated p-values (probability that the two distributions are the same) between 6.3×10^{-5} and 0.68 with mean 0.15.

We then ran the simulation twenty times just on NEURON, each time with a different RNG seed, to give 10 pairs of distributions. In this case the D-values were in the range 0.007–0.026, mean 0.015, and the p-values in the range 2.8×10^{-5} to 0.77, mean 0.20.

In summary, the differences due to different simulators are in almost exactly the same range as those due to different initial conditions, suggesting that the differences between the simulators are indeed due to round-off errors and that there are not, therefore, any implementation errors in this case.

It is also interesting to note that in most cases the null hypothesis is supported, i.e. the distributions are the same, but that for some initial conditions there are highly significant differences between the ISI distributions. The ISI distribution may not therefore be the best measure for reproducibility in this case.

DISCUSSION

In this article we have presented PyNN, a Python-based common simulator interface, which allows simulator-independent model specification. PyNN is already in use in a number of research groups, and has been a key technology enabling improved communication between labs in a pan-European collaborative project with a major component of modelling and of neuromorphic hardware development (the FACETS project: http://www.facets-project.org).

By providing a standard simulation platform, PyNN also has the potential to act as the foundation for other, simulator agnostic but neuroscience-specific, tools such as analysis, visualization and data-management software.

PyNN is not the only project to address simulator-independent model specification and simulator interoperability (review in Cannon et al., 2007). neuroConstruct (Gleeson et al., 2007) is a tool to develop networks of morphologically-detailed neurons using a graphical user interface (GUI), that can generate code for both the NEURON and GENESIS simulators. A limitation with respect to PyNN is that since it uses code generation rather than a direct interface, neuroConstruct cannot receive information back from the simulator except by reading the data files it

generates. A second limitation is that features that are not available through the GUI cannot be incorporated in a model. The NeuroML standards (Crook et al., 2005, http://www.neuroml.org) are intended to provide an infrastructure for exchanging model specifications between groups in a simulator-independent way. Their scope includes much more detailed levels of modelling, e.g. membrane ion channels and detailed dendritic morphology, than are supported by PyNN. They have the advantage over PyNN of being language-independent, since specifications are written in XML, for which tools exist in all major programming languages. The major disadvantage of purely declarative specifications is lack of flexibility: if a concept or entity is not defined in the standard, it is not possible to specify models that use it, whereas with a procedural/imperative or mixed declarative-procedural specification such as is achievable with PyNN, arbitrary specifications are possible.

Although we emphasize here the differences between the GUI, pure-declarative, and programming-interface approaches to simulator-independent model specification, in fact they are highly complementary. Graphical interfaces are particularly good for beginners, for teaching, for giving high-level overviews of a system, and for integrating analysis and visualization tools. It would be very useful for neuroConstruct to be able to generate PyNN code, for example, in addition to code for NEURON and GENESIS. Declarative specifications reach the highest levels of system-independence, for the range of concepts that are supported. They are also particularly suitable for transformation into human-readable formats and for automated GUI generation. As such, they seem to be best suited for domains in which the modelling approach is fairly stable, e.g. for describing neuron morphologies or non-stochastic ion channel models. In PyNN, we plan to support simulator-independent multi-compartmental models using NeuroML: in this scenario cell models would be specified in NeuroML while PyNN would be used for network specification and for simulation setup and control.

Our main priorities for future development of PyNN are to increase the number of supported simulators (simulator developers who are interested in PyNN support for their simulator are encouraged to contact us), improve the support for multi-compartmental modelling, and extend the interface towards higher-level abstractions, such as cortical columns and more abstract modelling approaches. PyNN is open source software (CeCILL licence, http://www.cecill.info) and has an open development model: anyone who wishes to contribute is welcome and invited to do so.

ACKNOWLEDGEMENTS

This work was supported by the European Union (FACETS project, FP6-2004-IST-FETPI-015879). Jens Kremkow is also supported by the German Federal Ministry of Education and Research (BMBF grant 01GQ0420 to BCCN, Freiburg).

REFERENCES

Brette, R., and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94, 3637–3642.

Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J., Diesmann, M., Morrison, A., Goodman P. H., Harris, F. Jr, Zirpe, M., Natschlager, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., Davison, A., El Boustani, S., and Destexhe, A. (2007). Simulation of networks of spiking neurons: a review of tools and strategies. *J Comput. Neurosci.* 23, 349–398.

Cannon, R., Gewaltig, M., Gleeson, P., Bhalla, U., Cornelis, H., Hines, M., Howell, F., Muller, E., Stiles, J., Wils, S., and De Schutter, E. (2007). Interoperability of neuroscience modeling software: current status and future directions. *Neuroinformatics* 5, 127–138.

Carnevale, N. T., and Hines, M. L. (2006). The NEURON Book. Cambridge, University Press.

Crook, S., Beeman, D., Gleeson, P., and Howell, F. (2005). XML for model specification in neuroscience: an introduction and workshop summary. *Brains Minds Media* 1, bmm228 (urn:nbn:de:0009–3–2282).

Ekeberg, Ö., and Djurfeldt, M. (2008). MUSIC – multisimulation coordinator: request for comments. *Nat. Precedings* http://dx.doi.org/10.1038/npre.2008.1830.1.

Eppler, J., Helias, M., Diesmann, M., and Gewaltig, M.-O. (2008). PyNEST: a convenient interface to the NEST simulator. *Front. Neuroinform.* 2. doi: 10.3389/neuron.11.012.2008.

Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia* 2, 1430.

Gleeson, P., Steuber, V., and Silver, R. A. (2007). neuroConstruct: a tool for modeling networks of neurons in 3D space. *Neuron* 54, 219–235.

Goodman, D., and Brette, R. (2008). Brian: a simulator for spiking neural networks in Python. *Front. Neuroinform.* 2. doi: 10.3389/neuron.11.005.2008.

Hines, M. L., and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Comput.* 9, 1179–1209.

Hines, M., Davison, A., and Muller, E. (2008). NEURON and Python. *Front. Neuroinform.* 2. doi: 10.3389/neuron.11.001.2009.

Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Trans Neural Netw.* 14, 1569–1572.

Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Natl. Acad. Sci. USA* 95, 5323–5328.

Morrison, A., Mehring, C., Geisel, T., Aertsen, A., and Diesmann, M. (2005). Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput.* 17, 1776–1801.

Ray, S., and Bhalla, U. (2008). PyMOOSE: interoperable scripting in Python for MOOSE. *Front. Neuroinform.* 2. doi: 10.3389/neuron.11.006.2008.

Schemmel, J., Brüderle, D., Meier, K., and Ostendorf, B. (2007). Modeling synaptic plasticity within networks of highly accelerated I&F neurons. In Proceedings of the 2007 IEEE International Symposium on Circuits and Systems (ISCAS'07). New Orleans, IEEE Press, pp. 3367–3370. doi: 10.1109/ISCAS.2007.378289.

Vogels, T., and Abbott, L. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *J. Neurosci.* 25, 10786–10795.

Wilson, G. (2006). Where's the real bottleneck in scientific computing? *Am. Sci.* 94, 5–6.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 21 September 2008; paper pending published: 21 October 2008; accepted: 22 December 2008; published online: 27 January 2009.

*Citation: Davison AP, Brüderle D, Eppler J, Kremkow J, Muller E, Pecevski D, Perrinet L and Yger P (2009) PyNN: a common interface for neuronal network simulators. Front. Neuroinform. (2009) **2**:11. doi: 10.3389/neuro.11.011.2008*

22 The analysis workflow in neuroscience

22.1 Heterogeneity of the data

In computational neuroscience, one has to deal with simulated data, resulting from *in computo* simulations. Nevertheless, the main point is that these data need to be compared to experimental data, acquired and analysed with various and complex systems, each of them being tuned and dedicated to a particular setup or experiment. The devices used to record data impose a lot of various and heterogeneous custom file formats, and dedicated software reimplements too often similar functions, because manufacturers do not always give the technical specification of their internal structures, in order to force users to use home made toolboxes or software. To address this issue, from the experimental data side, some projects such as NeuroShare [http://neuroshare.sourceforge.net/] tried to design a common API able to read/write several file formats, in order to provide in a transparent manner some high level objects common to all those file formats to a MATLAB interface. Experimental recordings are saved with objects such as Blocks of recordings, Segments, Trials, Epochs, Spikes, and so on. As in the case of PyNN, canonical structures for these objects should be defined in order to cross-validate numerical results and promote code and data exchange.

22.2 On the analysis tools

Some toolboxes have already been made to take advantage of this NeuroShare implementation. This is the case, for example, of FIND, developed by Meier et al. (2008). But even without the NeuroShare framework, a lot of custom analysis toolboxes took advantage of the MATLAB environment, and especially the graphical user interface offered by the software, to explore and visualize the data. To list more of those analysis tools in computational neuroscience that can be found nowadays, one can mention Chronux, Spike Train Analysis Toolkit, sigTOOL. These tools are general, allowing a user to interact with his data, while some other are more dedicated to clustering and spike sorting algorithm (KlustaKwik, MClust, OSort, Wave_clus). In addition, there are also plenty of commercial tools sold with the hardware devices used to acquire the data. Others approaches, like OpenElectrophy (Garcia and Fourcaud-Trocmé, 2009), try to get rid of the MATLAB dependency and build a coherent and consistent environment for data analysis in Python. Python offers a nice and clean definition of objects, and objects saved after a recording session such as Electrode, Segment, SpikeTrain, and so on can be easily turned into Python objects.

22.3 The NeuroTools project

Background of the initiative NeuralEnsemble [http://neuralensemble.org] is a multilateral effort to coordinate and organise neuroscience software development based around the Python programming language into a larger, meta-simulator software system. To this end, NeuralEnsemble hosts services for source code management and bug tracking (Subversion/Trac) for a number of open-source neuroscience tools and organizes an annual workshop devoted to collaborative software development in neuroscience. In this NeuralEnsemble framework, NeuroTools is a collection of tools to support all tasks associated with a neural simulation project which are not handled by the simulation engine. NeuroTools is written in Python,

and works best with PyNN, or one of the growing list of simulation engines with a Python front-end such as NEURON, NEST, PCSIM, FACETS Neuromorphic VLSI, Brian, MOOSE/-GENESIS3, Neurospaces/GENESIS. NeuroTools provides modules to facilitate simulation setup, parametrization, data management, analysis and visualization. The data-related tools are equally suited to analysis of experimental data, although that is not the primary motivation for their development. As with PyNN, its aim is to allow cross-validation of results and code exchange in order to more quickly detect and isolate bugs or errors.

An implementation in Python NeuroTools is developed in Python because, as an interpreted language, it offers a nice opportunity to dynamically interact with the data. Data can be loaded and explored through scripts or direct commands, as is the case with MATLAB. Nevertheless, the Python implementation is much more powerful because it is based on a growing and very active scientific community. It has a much better support for structured programming and has a wide range of tools outside numerics and data handling. Scientific packages such as `numpy`, `scipy`, `matplotlib` are now turning Python into a very efficient and free MATLAB replacement, much more modular. Python provides the power of a clear object-oriented language with a very modern and powerful syntax, and benefits from a growing scientific community providing a lot of various packages as modules that could be plugged or not. The structure of NeuroTools itself was based on this modularity. As one can see in Figure 40, the project is divided into several sections. The `signals` package gathers all the analysis functions for discrete (spike trains) or analog (membrane potential) signals that may be recorded during a simulation or an experiment. The `parameters` module provides a framework for simulation parameters handling, to simplify exploration and backups when simulations with numerous parameters are launched.

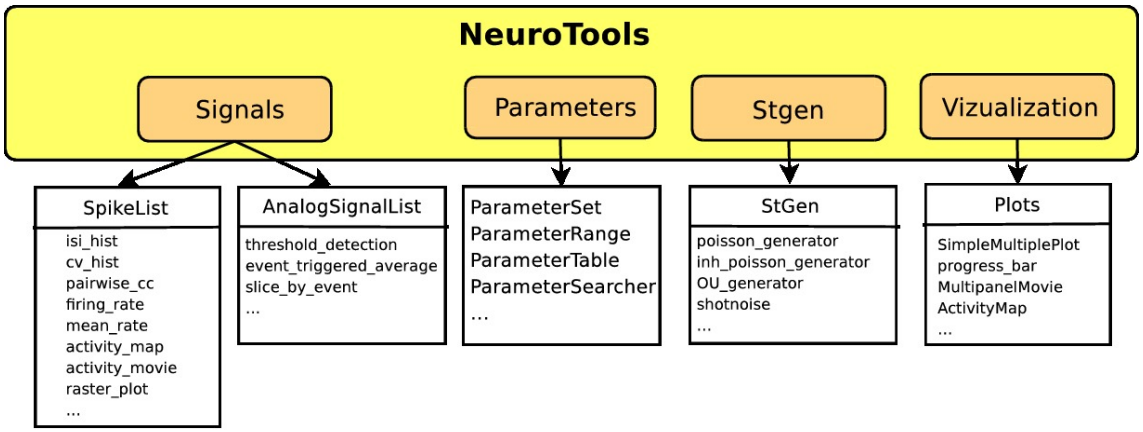

Figure 40: NeuroTools structure. Independent packages can be dynamically loaded to enable tools and/or functions. Among these, the signals package, the visualization package, the parameters package.

In the following, one can see a script example of some very light and easy functions of NeuroTools. The results of the script can be seen in Figure 41. As one can see, the code is compact and the command line provided by Python allows exploration and plots of high-level functions, applied on numerical data. The transparency in the source code and the documentation should increase the confidence in the results.

```
>>> from NeuroTools.signals import *
```

```
>>> from NeuroTools.stgen import *

>>> spikes = load_spikelist('data.gdf',t_start=0,t_stop=500)
>>> spikes.firing_rate(time_bin=5, display=subplot(221))
>>> spikes.raster_plot(1000,display=subplot(222))
>>> pairs = RandomPairs(spikes, spikes, no_silent=True)
>>> spikes.pairwise_cc(1000, pairs_generator=pairs,display=subplot(223))

>>> gen   = StGen()
>>> time  = arange(0,1000.)
>>> rates = sin(time/20)+1
>>> spikes = gen.inh_poisson_generator(100*rates, time, time[-1])
>>> spikes.raster_plot(display=subplot(224), kwargs={'marker':'+'})
>>> plot(time, rates, 'r--')
```

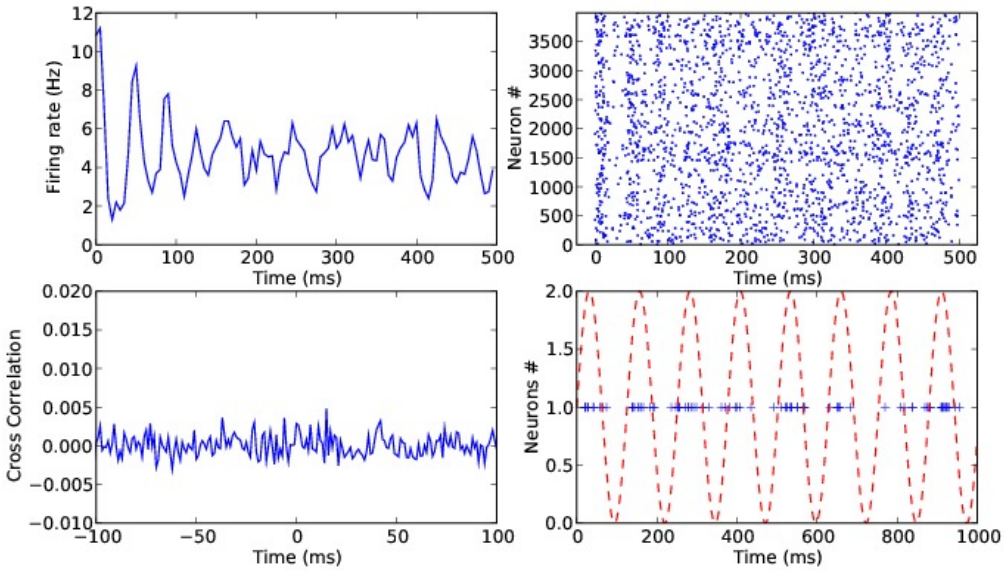

Figure 41: Result of the simple and compact example script written in the text. NeuroTools offers an easy and friendly way to explore and analyze complex data.

23 Discussion

Spreading the tools Combining PyNN and NeuroTools, the idea is to build a coherent and unified framework to load and analyse data, either generated by neuronal network simulators or by real recordings performed during experiments. The current state of this global project is some way from fully achieving its goal. Although PyNN is now more and more popular, offering a clear and well-documented API, the situation is not the same for NeuroTools. Merging it with OpenElectrophy is intended to push it to the next step, i.e. NeuroTools will gain access to a broader spectrum of users, coming from various and heterogeneous backgrounds. Since OpenElectrophy was primarily made to analyse data recorded during *in vivo* and *in vitro* experiments, this is complementary to the core function of NeuroTools, made to analyse data mainly coming from large-scale simulations of networks. Ideally, one could have a single neuroscience software tool that should be able to read/write various file formats, no

matter how they have been acquired, and explore this data in a single consistent environment. The key point, to achieve such a goal, is to reach a certain critical mass where users will be recruited into developers and contribute, by the addition of generic function, to the development of the software. Documentation is essential to make it as simple as possible in order to convince them to contribute to the code without having to dive too much into the complexity of the main kernel. Regarding NeuroTools, this documentation step, although started, is not finished and the key structure of the software is not stabilized yet.

On large-scale neuronal networks In all the simulator literature, the quest is mainly for speed, memory efficiency, and finally the size of the network that can be simulated. Nevertheless, most of those simulations are made with rather generic and homogeneous networks, such as the balanced random ones, and it should be pointed out that, as long as the networks are homogeneous, increasing the size may sometimes be a waste of time. All the results obtained in this manuscript, whether for the Frozen Paradigm (Part II), for the topological study (Part III) of for plasticity (Part IV) were obtained in 12500 cell networks (or network with similar size), but results were always tested in larger networks of 100000 neurons. In all cases, results were qualitatively identical. Since networks are homogeneous, without complex structures embedded in them, increasing the size does not alter nor change the results. Simulating large networks is therefore always something that should be weighted, because it may just make the analysis harder, being overwhelmed by data and recordings. The larger the size is, the more precise the question which is addressed to the system should be. To circumvent this problem, design of non-homogeneous and structured network should be simplified and handling of the data enhanced to speed up the research work flow and parameter exploration.

PyNN evolutions PyNN is still in development, and efforts should be maintained to continue the development of a coherent framework for neuronal network simulations. The API is stable, even if some new structures and functionalities will be added, to promote the design of more and more complex networks. 0.7 release of PyNN will have sub-populations and assemblies (heterogeneous groups of populations) and the connections set algebra. Computational power is not expensive nowadays, so code complexity is more and more the main limitation when one wants to design biologically plausible networks with various cell types, connectivities, and not end up with systems overwhelmed by parameters. This can be tremendously reduced by creating high-level and generic structures in PyNN, such as cortical columns, layered structures, connectivities as functions of neuronal parameter and distances (can be Euclidean distances, or distances in other parameters spaces, such as orientation preference, ...). The goal is to have a simple and robust way to create a realistic neuronal network model, including all the complexity needed to explain and reproduce the huge variety of experimental results obtained in literature.

Part VII
Discussion & Further work

Summary

References 201

In this manuscript, we tried to establish a link between evoked and ongoing spiking activity in cortical-like networks, and to see how and why both could share similar statistics. In a first part, we showed what could be the interest for a chaotic dynamical system near the edge of chaos, in term of reproducibility and information transmission, to be stimulated by inputs sharing the same statistics to those of its ongoing activity. We then explored plausible biological processes that could promote such an unsupervised learning within neuronal networks, in order to establish this match between evoked and ongoing activity. Sensory inputs coming from the periphery have particular spatio-temporal properties that are stored during critical periods of development in the connectivity of the primary sensory areas by unsupervised learning rules, such as STDP. This is the case in V1, where it is known that the receptive fields emergence and development is strongly link to visual experience (Freeman and Pettigrew, 1973). Brain is a labile reservoir which is fed by the external world, and imprinted in its structure by its statistics. Rephrasing the seminal paper of Spinelli and Jensen (1979), plasticity in neocortex is a mirror of sensory experience.

Since plenty of biological evidence demonstrate that connectivity of neuronal network is reflected in the dynamics (see Part II), learned statistics should be spontaneously replayed while no external inputs are present. How this replay is achieved, and what key patterns or songs are useful for the system is still unknown. Are they synfire chains, cortical songs, rate-based measures, the answer may depend on the areas and on the complexity of the task. As we saw, the more the stimuli are complex in V1, the sparser the neuronal responses, and the amount of correlations integrated by the neuron tends to be close to the one present in the ongoing activity. This is why understanding the neuronal code is difficult: in computer science, the coding strategy is fixed, and the basic instructions sent to the processor are always the same, whatever the program. With neuronal networks, all coding schemes can coexist and be alternatively activated as a function of the complexity of the external stimuli. For a "simple" stimulus, with a low informational content, such as drifting grating, the rates of discharge would give all the needed information to decode the input. In contrast, a natural scene with a particular frequency spectrum could trigger finer correlations, and their exact times of occurrences can become crucial for the system.

The Boltzmann machine We propose to see the primary sensory networks as Boltzmann machines, as discussed in Part IV. When particular inputs are presented to a clamped layer (the thalamical neurons), they efficiently drive the global spatio-temporal activity of the whole network composed by the free-units and their own dynamics. Unsupervised learning rules, such as mSTDP explored in Part V, are able to store and to keep a stable trace of external statistics in the ongoing activity, reverberated by the recurrent connections. A more detailed and exhaustive study needs still to be performed in order to have a better understanding on the speed of learning and on the memory capacity of such a learning system. In Part V, we just established a proof of concept that external inputs, with some deviations from the ongoing activity, could be stored and kept by a modified version of STDP, taking homeostasis into account. Mainly firing rates and correlations were explored and stored, but since the Frozen Paradigm worked with precise spikes patterns, this should also be explored. The mSTDP rule should be used to embed, as in (Hosaka et al., 2008), synfire chains or precise patterns within the network. The question of their replay can be accurately addressed with tools already used (Ikegaya et al., 2004, Mokeichev et al., 2007), and it has to be linked with the concept of polychronization(Izhikevich, 2006).

A cascade of such Boltzmann machines would be, conceptually, very close to the "deep belief

networks" (Hinton and Salakhutdinov, 2006). The Frozen Paradigm is a proof that even precisely time coded spatio-temporal patterns could be learned and replayed, even if they may not be used by the system. The brain does not need to implement a proper coding/decoding transformation. The received information should be transformed eventually at a read-out stage into a response, but the exact input values are not relevant any more for the system, as soon as they have imprinted the processing unit dynamics. In that sense, the Frozen Paradigm raises an interesting concern: low dimensional spatio-temporal pattern fragment can reactivate a perseverating complex sequence of precisely time locked action potentials.

Toward realistic V1 models In the previous topographical networks, studied in Part III or IV, cross-correlations between neurons mainly depend on the distance because this was the only relevant information which has been put into the system, imposed by the connectivity profile used to establish the connections. However, as pointed out in Part III, distance is not the only property constraining the connectivity within neurons. In V1, it has been shown that cross-correlations depend also on the orientation preferences (Ts'o et al., 1986), and this would be in favour of even more realistic networks. In this manuscript, we stressed the fact that connectivity is crucial for the network since it constrains the attractor of its ongoing activity. It is therefore important, to design realistic V1 models, to generate realistic connectivities capturing the statistics of the evoked activities in such a way as transfer this knowledge into the wiring. The best way would be to use realistic plasticity rules, such as the mSTDP rule. Such rule, that needs still to be better studied, allows connectivity to keep a persistent trace of the input patterns. By starting from a random network, presentation of spike trains constructed by a retina/LGN model (Wohrer and Kornprobst, 2009) can be used to train the network, and impact the statistics of the emerging receptive fields (Wenisch et al., 2005).

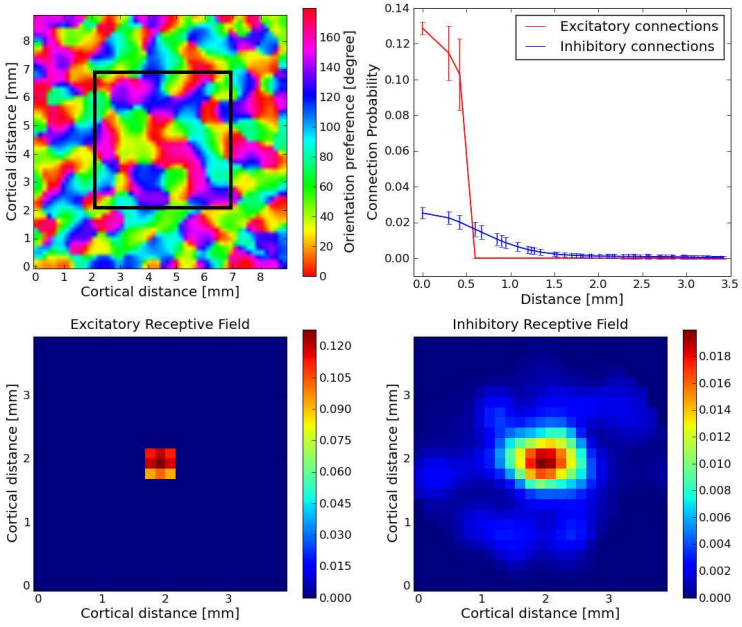

Figure 42: Orientation map obtained after Hebbian learning in Topographica. Upper Right: Profile of the connections from excitatory and inhibitory units, as a function of distance. Lower: example of a receptive field for one excitatory (respectively inhibitory) unit (left, respectively right).

Nevertheless, all these learning algorithms for large scale networks converge slowly, and a practical and easy way to generate more realistic V1 models is to use fast Hebbian learning on firing-rate neurons, which is known to converge efficiently and in a fast manner (Bednar et al., 2004). Software such as Topographica is able to create complex and inhomogeneous features maps as a function of the inputs that have been presented during the development. Such maps can then be used to initialize the connectivity of larger networks of spiking neurons, before turning them into a plastic state, and see how stable they are when STDP (or mSTDP) is active. A conceptual scheme is explained in Figure 42. An orientation preference map obtained with Topographica, with the lateral connections for excitatory and inhibitory cells is shown after learning. Then, each pixel is seen as a cortical column: all the rate-based units are turned into small balanced random networks, and connections strengths between the units are turned into probabilities of connections. We obtained a spiking network with a complex and inhomogeneous structure, embedding several features maps (pinwheels, orientations, ocular selectivity, directions, ...) as a function of its Hebbian learning. The question of the stability of such networks, under spike timing dependent plasticity rules, is essential.

Further works should investigate the dynamics of such networks, with more patchy connectivity. The only problem is that Hebbian learning, in those mainly inhibitory driven network, provides mainly long-range inhibition and short range excitation within V1, while biological evidences suggest a broader and more distributed excitation over long intracortical distances (Stettler et al., 2002).

STDP and backpropagation algorithm Despite all biological evidences for supervised learning in the brain triggered by external neuromodulator release (Romo and Schultz, 1990), one should notice that the neural network algorithm that has been most successful in engineering and real-world applications of supervised learning use mechanisms distinct from the Hebb's rule. This algorithm is known as error backpropagation or *backprop*. Synaptic plasticity in this network requires a training signal to flow backwards along the connections (a retroaxonal signal or a diffusible messenger in the extracellular medium), so that each neuron integrates a training (or error) signal from the same cells to which it sends outputs. With multilayered feed-forward artificial neuronal networks, one can use a gradient descent on a cost function comparing the current output of the neuron to a desired one and adjust the weights by an iterative process backpropagating errors: this approach has been shown to be a universal function approximator. The computing power offered by such a backpropagation method is still a pure engineering concept, and has never been identified in real neurons. It can be therefore interesting to study if biological evidence can support such an algorithmic method. A possibility is to and gain knowledge with biological data recorded in vitro, and study if backpropagation can be used as a generic principle to shape neuronal network structures in order to produce desired pattern of activity, at the spiking level. As reported by some authors (Harris, 2008), biological evidences show how retroaxonal signals are indeed used by neurons, and could support a kind of backprop algorithm in vivo. But even if they are present, retroaxonal signals in real neurons are too slow to implement a real backprop algorithm in vivo. In artificial neuronal networks (ANNs), weight changes are based on the instantaneous difference between the pre-synaptic activity and the retroaxonal error signal, requiring the training signal to arrive while the input pattern is still present, in order to give computational support for the gradient descent. Given that biological synaptic timescales are typically if the order of $\simeq$10ms, while known forms of retroaxonal messages with signalling endosome (Zweifel et al., 2005) are much slower, computation of instantaneous correlations seems impossible.

Nevertheless, such retroaxonal training signals could consolidate causal useful associations in neuronal networks, and be put in the context of metaplasticity (see Abraham (2008) for review) and influence the thresholds of the mSTDP rule: this integrated and distributed feedback coming from target neurons may regulate efficiently the long-term maintenance of local pre-synaptic plasticity.

References

L. F. Abbott. Lapique's introduction of the integrate-and-fire model neuron (1907). *Brain Research Bulletin*, 50 (5/6):303–304, 1999.

L. F. Abbott. Balancing homeostasis and learning in neural circuits. *Zoology (Jena)*, 106(4): 365–371, 2003.

L. F. Abbott and S. B. Nelson. Synaptic plasticity: taming the beast. *Nat Neurosci*, 3 Suppl: 1178–1183, Nov 2000.

M. Abeles. *Corticonics: neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge, 1991.

M. Abeles, Y. Prut, H. Bergman, and E. Vaadia. Synchronization in neuronal transmission and its importance for information processing. *Prog Brain Res*, 102:395–404, 1994.

W. C. Abraham and M. F. Bear. Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci*, 19(4):126–130, Apr 1996.

Wickliffe C Abraham. Metaplasticity: tuning synapses and networks for plasticity. *Nat Rev Neurosci*, 9(5):387, May 2008.

D. H. Ackley, G. E. Hinton, and T. J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9:147–169, 1985.

S. Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol Cybern*, 27(2):77–87, Aug 1977.

Amit, Gutfreund, and Sompolinsky. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys Rev Lett*, 55(14):1530–1533, Sep 1985.

D. J. Amit and N. Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex*, 7(3):237–252, 1997.

F. Amzica and M. Steriade. Short- and long-range neuronal synchronization of the slow ($<$ 1 hz) cortical oscillation. *J Neurophysiol*, 73(1):20–38, Jan 1995.

Rajagopal Ananthanarayanan, Steven K. Esser, Horst D. Simon, and Dharmendra S. Modha. The cat is out of the bag: cortical simulations with 10^9 neurons, 10^{13} synapses. In *SC*, 2009.

R. K. Andjus, K. Suhara, and H. A. Sloviter. An isolated, perfused rat brain preparation, its spontaneous and stimulated activity. *J Appl Physiol*, 22(5):1033–1039, May 1967.

Ian M Andolina, Helen E Jones, Wei Wang, and Adam M Sillito. Corticothalamic feedback enhances stimulus response precision in the visual system. *Proc Natl Acad Sci U S A*, 104 (5):1685–1690, Jan 2007.

A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 273(5283):1868–1871, 1996.

A. Artola and W. Singer. Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci*, 16(11):480–487, Nov 1993.

A. Artola, S. Bröcher, and W. Singer. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature*, 347 (6288):69–72, Sep 1990.

R. Azouz and C. M. Gray. Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J Neurosci*, 19(6):2209–2223, 1999.

Mathilde Badoual, Quan Zou, Andrew P Davison, Michael Rudolph, Thierry Bal, Yves Frégnac, and Alain Destexhe. Biophysical and phenomenological models of multiple spike interactions in spike-timing dependent plasticity. *Int J Neural Syst*, 16(2):79–97, Apr 2006.

Arunava Banerjee, Peggy Seriés, and Alexandre Pouget. Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. *Neural Comput*, 20(4): 974–993, Apr 2008.

Dorit Baras and Ron Meir. Reinforcement learning, spike-time-dependent plasticity, and the BCM rule. *Neural Comput*, 19(8):2245–2279, Aug 2007.

Angel Barco, Juan M Alarcon, and Eric R Kandel. Expression of constitutively active CREB protein facilitates the late phase of long-term potentiation by enhancing synaptic capture. *Cell*, 108(5):689–703, Mar 2002.

H. B. Barlow. Single units and sensation: a neuron doctrine for perceptual psychology? *Perception*, 1(4):371–394, 1972.

M. F. Bear. Mechanism for a sliding synaptic modification threshold. *Neuron*, 15(1):1–4, Jul 1995.

C. Beaulieu and M. Colonnier. The number of neurons in the different laminae of the binocular and monocular regions of area 17 in the cat, canada. *J Comp Neurol*, 217(3):337–344, Jul 1983.

James A Bednar, Amol Kelkar, and Risto Miikkulainen. Scaling self-organizing maps to model large cortical networks. *Neuroinformatics*, 2(3):275–302, 2004.

C. C. Bell, V. Z. Han, Y. Sugawara, and K. Grant. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, 387(6630):278–281, May 1997.

Andrea Benucci, Robert A Frazor, and Matteo Carandini. Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron*, 55(1):103–117, Jul 2007.

N. Bertschinger and T. Natschlager. Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput*, 16(7):1413–1436, 2004.

G. Q. Bi and M. M. Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 18(24): 10464–10472, Dec 1998.

E. Bienenstock. On the dimensionality of cortical graphs. *J Physiol Paris*, 90(3-4):251–256, 1996.

E. L. Bienenstock, L. N. Cooper, and P. W. Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci*, 2 (1):32–48, Jan 1982.

Guy Billings and Mark C W van Rossum. Memory retention and spike-timing-dependent plasticity. *J Neurophysiol*, 101(6):2775–2788, Jun 2009.

Tom Binzegger, Rodney J Douglas, and Kevan A C Martin. A quantitative map of the circuit of cat primary visual cortex. *J Neurosci*, 24(39):8441–8453, Sep 2004.

T. V. Bliss and T. Lomo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol*, 232 (2):331–356, Jul 1973.

Sander M. Bohte and Michael C. Mozer. Reducing spike train variability: A computational theory of spike-timing dependent plasticity. *NIPS*, 2005, 2005.

V. Braitenberg and A. Schüz. *Anatomy of the cortex*. Springer, 1991.

Paul C Bressloff. Bloch waves, periodic feature maps, and cortical pattern formation. *Phys Rev Lett*, 89(8):088101, Aug 2002.

R Brette, M Rudolph, T Carnevale, M Hines, D Beeman, JM Bower, M Diesmann, A Morrison, Jr PH Goodman, FC Harris, M Zirpe, T Natschlager, D Pecevski, B Ermentrout, M Djurfeldt, A Lansner, O Rochel, T Vieville, E Muller, Andrew Davison, Sami El Boustani, and Alain Destexhe. Simulation of networks of spiking neurons: A review of tools and strategies. *J Comput. Neurosci.*, 23:349–98, 2007.

Romain Brette. Generation of correlated spike trains. *Neural Comput*, Aug 2008.

Romain Brette and Wulfram Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J Neurophysiol*, 94(5):3637–3642, Nov 2005.

V. Bringuier, Y. Frégnac, A. Baranyi, D. Debanne, and D. E. Shulz. Synaptic origin and stimulus dependency of neuronal oscillatory activity in the primary visual cortex of the cat. *J Physiol*, 500 (Pt 3):751–774, May 1997.

V. Bringuier, F. Chavane, L. Glaeser, and Y. Frégnac. Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. *Science*, 283(5402):695–699, Jan 1999.

Teresita L Briones, Anna Y Klintsova, and William T Greenough. Stability of synaptic plasticity in the adult rat visual cortex induced by complex environment exposure. *Brain Res*, 1018(1):130–135, Aug 2004.

Carlos D Brody and J. J. Hopfield. Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron*, 37(5):843–852, Mar 2003.

N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*, 8(3):183–208, 2000.

N. Brunel and V. Hakim. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput*, 11(7):1621–1671, Oct 1999.

A. Burkitt, M. Gilson, and J. van Hemmen. Spike-timing-dependent plasticity for neurons with recurrent connections. *Biol Cybern*, Apr 2007.

A. N. Burkitt, H. Meffin, and D. B. Grayden. Study of neuronal gain in a conductance-based leaky integrate-and-fire neuron model with balanced excitatory and inhibitory synaptic input. *Biol Cybern*, 89(2):119–125, Aug 2003.

Daniel A Butts, Chong Weng, Jianzhong Jin, Chun-I. Yeh, Nicholas A Lesica, Jose-Manuel Alonso, and Garrett B Stanley. Temporal precision in the neural code and the timescales of natural vision. *Nature*, 449(7158):92–95, Sep 2007.

Natalia Caporale and Yang Dan. Spike timing-dependent plasticity: a hebbian learning rule. *Annu Rev Neurosci*, 31:25–46, 2008.

G. C. Castellani, E. M. Quinlan, L. N. Cooper, and H. Z. Shouval. A biophysical model of bidirectional synaptic plasticity: dependence on AMPA and NMDA receptors. *Proc Natl Acad Sci U S A*, 98(22):12772–12777, Oct 2001.

Bruno Cessac and Thierry Viéville. On dynamics of integrate-and-fire neural networks with conductance based synapses. *Front Comput Neurosci*, 2:2, 2008.

J. P. Changeux and A. Danchin. Selective stabilisation of developing synapses as a mechanism for the specification of neuronal networks. *Nature*, 264(5588):705–712, 1976.

Gal Chechik. Spike-timing-dependent plasticity and relevant mutual information maximization. *Neural Comput*, 15(7):1481–1510, Jul 2003.

Vivien Chevaleyre, Kanji A Takahashi, and Pablo E Castillo. Endocannabinoid-mediated synaptic plasticity in the cns. *Annu Rev Neurosci*, 29:37–76, 2006.

Claudia Clopath, Lorric Ziegler, Eleni Vasilaki, Lars Büsing, and Wulfram Gerstner. Tag-trigger-consolidation: a model of early and late long-term-potentiation and depression. *PLoS Comput Biol*, 4(12):e1000248, Dec 2008.

Claudia Clopath, Lars Büsing, Eleni Vasilaki, and Wulfram Gerstner. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci*, 13(3):344–352, Mar 2010.

E. E. Clothiaux, M. F. Bear, and L. N. Cooper. Synaptic plasticity in visual cortex: comparison of theory with experiment. *J Neurophysiol*, 66(5):1785–1804, Nov 1991.

D. Contreras and M. Steriade. Cellular basis of EEG slow rhythms: a study of dynamic corticothalamic relationships. *J Neurosci*, 15(1 Pt 2):604–622, Jan 1995.

Diego Contreras. Propagating waves in visual cortex. *Neuron*, 55(1):3–5, Jul 2007.

M. C. Crair, D. C. Gillespie, and M. P. Stryker. The role of visual experience in the development of columns in cat visual cortex. *Science*, 279(5350):566–570, Jan 1998.

Sylvain Crochet, Pablo Fuentealba, Youssouf Cissé, Igor Timofeev, and Mircea Steriade. Synaptic plasticity in local cortical network in vivo and its modulation by the level of neuronal activity. *Cereb Cortex*, 16(5):618–631, May 2006.

Yang Dan and Mu-Ming Poo. Spike timing-dependent plasticity of neural circuits. *Neuron*, 44(1):23–30, Sep 2004.

C. Darwin. *On the Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*. John Murray, 1859.

Andrew P Davison and Yves Frégnac. Learning cross-modal spatial transformations through spike timing-dependent plasticity. *J Neurosci*, 26(21):5604–5615, May 2006.

Andrew P Davison, Daniel Brüderle, Jochen Eppler, Jens Kremkow, Eilif Muller, Dejan Pecevski, Laurent Perrinet, and Pierre Yger. PyNN: A common interface for neuronal network simulators. *Front Neuroinformatics*, 2:11, 2008.

A. F. Dean. The relationship between response amplitude and contrast for cat striate cortical neurones. *J Physiol*, 318:413–427, 1981.

D. Debanne, D. E. Shulz, and Y. Frégnac. Temporal constraints in associative synaptic plasticity in hippocampus and neocortex. *Can J Physiol Pharmacol*, 73(9):1295–1311, Sep 1995.

D. Debanne, D. E. Shulz, and Y. Fregnac. Activity-dependent regulation of 'on' and 'off' responses in cat visual cortical receptive fields. *J Physiol*, 508 (Pt 2):523–548, Apr 1998.

Niraj S Desai, Robert H Cudmore, Sacha B Nelson, and Gina G Turrigiano. Critical periods for experience-dependent synaptic scaling in visual cortex. *Nat Neurosci*, 5(8):783–789, Aug 2002.

R. Desimone, T. D. Albright, C. G. Gross, and C. Bruce. Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci*, 4(8):2051–2062, Aug 1984.

A. Destexhe and D. Paré. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol*, 81(4):1531–1547, Apr 1999.

A. Destexhe, M. Rudolph, J. M. Fellous, and T. J. Sejnowski. Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24, 2001.

Alain Destexhe, Stuart W Hughes, Michelle Rudolph, and Vincenzo Crunelli. Are corticothalamic 'up' states fragments of wakefulness? *Trends Neurosci*, 30(7):334–342, Jul 2007.

Michael R Deweese and Anthony M Zador. Shared and private variability in the auditory cortex. *J Neurophysiol*, 92(3):1840–1855, Sep 2004.

M. Diesmann, M. O. Gewaltig, and A. Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533, 1999.

Hubert Eichner, Tobias Klug, and Alexander Borst. Neural simulations on multi-core architectures. *Front Neuroinformatics*, 3:21, 2009.

S. El Boustani and A. Destexhe. Does brain activity stem from high-dimensional chaotic dynamics? Evidence from the human electroencephalogram, cat cerebral cortex and artificial neuronal networks. *International Journal of Bifurcation and Chaos*, in press, 2009a.

S. El Boustani and Alain Destexhe. A master equation formalism for macroscopic modeling of asynchronous irregular activity states. *Neural Comput*, 21(1):46–100, Jan 2009b.

S. El Boustani, Olivier Marre, Sébastien Béhuret, Pierre Baudot, Pierre Yger, Thierry Bal, Alain Destexhe, and Yves Frégnac. Network-state modulation of power-law frequency-scaling in visual cortical neurons. *PLoS Comput Biol*, 5(9):e1000519, Sep 2009.

Ype Elgersma, Nikolai B Fedorov, Sami Ikonen, Esther S Choi, Minetta Elgersma, Ofelia M Carvalho, Karl Peter Giese, and Alcino J Silva. Inhibitory autophosphorylation of CaMKII controls PSD association, plasticity, and learning. *Neuron*, 36(3):493–505, Oct 2002.

Jochen Martin Eppler, Moritz Helias, Eilif Muller, Markus Diesmann, and Marc-Oliver Gewaltig. PyNEST: A convenient interface to the NEST simulator. *Front Neuroinformatics*, 2:12, 2008.

P. Faure and H. Korn. Is there chaos in the brain? I. concepts of nonlinear dynamics and methods of investigation. *C R Acad Sci III*, 324(9):773–793, Sep 2001.

Isabelle Ferezou, Sonia Bolea, and Carl C H Petersen. Visualizing the cortical representation of whisker touch: voltage-sensitive dye imaging in freely moving mice. *Neuron*, 50(4): 617–629, May 2006.

J. Fiser, C. Chiu, and M. Weliky. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature*, 431(7008):573–578, 2004.

R. Fitzhugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophys J*, 1(6):445–466, Jul 1961.

Michael D Fox and Marcus E Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci*, 8(9):700–711, Sep 2007.

R. D. Freeman and J. D. Pettigrew. Alteration of visual cortex from environmental asymmetries. *Nature*, 246(5432):359–360, Dec 1973.

Yves Frégnac, Marc Pananceau, Alice René, Nazyed Huguet, Oliver Marre, Manuel Levy, and Daniel Schulz. A re-examination of hebbian-covariance rules and spike timing-dependent plasticity in cat visual cortex in vivo. *Frontiers in Synaptic Plasticity*, to appear, 2010.

U. Frey and R. G. Morris. Synaptic tagging and long-term potentiation. *Nature*, 385(6616): 533–536, Feb 1997.

Robert C Froemke and Yang Dan. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, 416(6879):433–438, Mar 2002.

Robert C Froemke, Mu-Ming Poo, and Yang Dan. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature*, 434(7030):221–225, Mar 2005.

Robert C Froemke, Ishan A Tsay, Mohamad Raad, John D Long, and Yang Dan. Contribution of individual spikes in burst-induced long-term synaptic modification. *J Neurophysiol*, 95 (3):1620–1629, Mar 2006.

Y. Frégnac. A tale of two spikes. *Nat Neurosci*, 2(4):299–301, Apr 1999.

Y. Frégnac and M. Imbert. Development of neuronal selectivity in primary visual cortex of cat. *Physiol Rev*, 64(1):325–434, Jan 1984.

Y. Frégnac, D. Shulz, S. Thorpe, and E. Bienenstock. A cellular analogue of visual cortical plasticity. *Nature*, 333(6171):367–370, May 1988.

Y. Frégnac, D. Shulz, S. Thorpe, and E. Bienenstock. Cellular analogs of visual cortical epigenesis. I. plasticity of orientation selectivity. *J Neurosci*, 12(4):1280–1300, Apr 1992.

Y. Frégnac, J. P. Burke, D. Smith, and M. J. Friedlander. Temporal covariance of pre- and postsynaptic activity regulates functional connectivity in the visual cortex. *J Neurophysiol*, 71(4):1403–1421, Apr 1994.

Stefano Fusi and L. F. Abbott. Limits on the memory storage capacity of bounded synapses. *Nat Neurosci*, 10(4):485–493, Apr 2007.

Samuel Garcia and Nicolas Fourcaud-Trocmé. Openelectrophy: An electrophysiological data- and analysis-sharing framework. *Front Neuroinformatics*, 3:14, 2009.

Pablo Garcia-Lopez, Virginia Garcia-Marin, and Miguel Freire. The histological slides and drawings of cajal. *Front Neuroanat*, 4:9, 2010.

A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner. Neuronal population coding of movement direction. *Science*, 233(4771):1416–1419, Sep 1986.

W. Gerstner. Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. *Neural Comput*, 12(1):43–89, Jan 2000.

W. Gerstner, R. Kempter, J. L. van Hemmen, and H. Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–81, Sep 1996. doi: 10.1038/383076a0. URL `http://dx.doi.org/10.1038/383076a0`.

Wulfram Gerstner and Werner M Kistler. Mathematical formulations of hebbian learning. *Biol Cybern*, 87(5-6):404–415, Dec 2002.

M.-O. Gewaltig and M Diesmann. NEST (NEural Simulation Tool). *Scholarpedia*, 2(4):1430, 2007.

Z. Gil, B. W. Connors, and Y. Amitai. Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability. *Neuron*, 23(2):385–397, Jun 1999.

Nace L Golding, Nathan P Staff, and Nelson Spruston. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature*, 418(6895):326–331, Jul 2002.

Dan Goodman and Romain Brette. Brian: a simulator for spiking neural networks in python. *Front Neuroinformatics*, 2:5, 2008.

Michael Graupner and Nicolas Brunel. STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLoS Comput Biol*, 3(11):e221, Nov 2007.

C. M. Gray, P. König, A. K. Engel, and W. Singer. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338(6213):334–337, Mar 1989.

David S Greenberg, Arthur R Houweling, and Jason N D Kerr. Population imaging of ongoing neuronal activity in the visual cortex of awake rats. *Nat Neurosci*, 11(7):749–751, Jul 2008.

G. W. Gross, A. N. Williams, and J. H. Lucas. Recording of spontaneous activity with photoetched microelectrode surfaces from mouse spinal neurons in culture. *J Neurosci Methods*, 5(1-2):13–22, Jan 1982.

Zhenglin Gu and Zhen Yan. Bidirectional regulation of ca2+/calmodulin-dependent protein kinase II activity by dopamine D4 receptors in prefrontal cortex. *Mol Pharmacol*, 66(4): 948–955, Oct 2004.

R. Gütig, R. Aharonov, S. Rotter, and Haim Sompolinsky. Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J Neurosci*, 23(9):3697–3714, May 2003.

Robert Gütig and Haim Sompolinsky. The tempotron: a neuron that learns spike timing-based decisions. *Nat Neurosci*, 9(3):420–428, Mar 2006.

Julie S Haas, Thomas Nowotny, and Henry D I Abarbanel. Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J Neurophysiol*, Aug 2006.

Bilal Haider, Alvaro Duque, Andrea R Hasenstaub, and David A McCormick. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *J Neurosci*, 26(17):4535–4545, Apr 2006.

Bilal Haider, Matthew R Krause, Alvaro Duque, Yuguo Yu, Jonathan Touryan, James A Mazer, and David A McCormick. Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron*, 65(1): 107–121, Jan 2010.

Feng Han, Natalia Caporale, and Yang Dan. Reverberation of recent visual experience in spontaneous cortical waves. *Neuron*, 60(2):321–327, Oct 2008.

Kenneth D Harris. Stability of the fittest: organizing learning through retroaxonal signals. *Trends Neurosci*, 31(3):130–136, Mar 2008.

Andrea Hasenstaub, Yousheng Shu, Bilal Haider, Udo Kraushaar, Alvaro Duque, and David A McCormick. Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron*, 47(3):423–435, Aug 2005.

Yuki Hashimotodani, Takako Ohno-Shosaku, and Masanobu Kano. Ca(2+)-assisted receptor-driven endocannabinoid release: mechanisms that associate presynaptic and postsynaptic activities. *Curr Opin Neurobiol*, 17(3):360–365, Jun 2007.

D.O. Hebb. *The organization of behavior*. Wiley, 1949.

B. Hellwig. A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol Cybern*, 82(2):111–121, Feb 2000.

M. L. Hines and N. T. Carnevale. The NEURON simulation environment. *Neural Comput*, 9 (6):1179–1209, Aug 1997.

Michael L Hines, Hubert Eichner, and Felix Schürmann. Neuron splitting in compute-bound parallel network simulations enables runtime scaling with twice as many processors. *J Comput Neurosci*, 25(1):203–210, Aug 2008.

G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, Jul 2006.

A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*, 117(4):500–544, Aug 1952.

J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*, 79(8):2554–2558, Apr 1982.

Ryosuke Hosaka, Osamu Araki, and Tohru Ikeguchi. STDP provides the substrate for igniting synfire chains by spatiotemporal input patterns. *Neural Comput*, 20(2):415–435, Feb 2008.

Toshihiko Hosoya, Stephen A Baccus, and Markus Meister. Dynamic predictive coding by the retina. *Nature*, 436(7047):71–77, Jul 2005.

Y. Y. Huang, A. Colino, D. K. Selig, and R. C. Malenka. The influence of prior synaptic activity on the induction of long-term potentiation. *Science*, 255(5045):730–733, Feb 1992.

D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat's striate cortex. *J Physiol*, 148:574–591, Oct 1959.

Yuji Ikegaya, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste. Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 304(5670):559–564, Apr 2004.

E.M. Izhikevich. Simple model of spiking neurons. *IEEE Trans Neural Netw.*, 14(6):1569–1572, 2003.

Eugene M Izhikevich. Polychronization: computation with spikes. *Neural Comput*, 18(2): 245–282, Feb 2006.

Eugene M Izhikevich. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb Cortex*, 17(10):2443–2452, Oct 2007.

Eugene M Izhikevich and Niraj S Desai. Relating STDP to BCM. *Neural Comput*, 15(7): 1511–1523, Jul 2003.

Vincent Jacob, Daniel J Brasier, Irina Erchova, Dan Feldman, and Daniel E Shulz. Spike timing-dependent synaptic depression in the in vivo barrel cortex of the rat. *J Neurosci*, 27 (6):1271–1284, Feb 2007.

Adam L Jacobs, Gene Fridman, Robert M Douglas, Nazia M Alam, Peter E Latham, Glen T Prusky, and Sheila Nirenberg. Ruling out and ruling in neural codes. *Proc Natl Acad Sci U S A*, 106(14):5936–5941, Apr 2009.

Herbert Jaeger and Harald Haas. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science*, 304(5667):78–80, Apr 2004.

Janneke F M Jehee and Jaap M J Murre. The scalable mammalian brain: emergent distributions of glia and neurons. *Biol Cybern*, 98(5):439–445, May 2008.

Renaud Jolivet, Timothy J Lewis, and Wulfram Gerstner. Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy. *J Neurophysiol*, 92(2):959–976, Aug 2004.

Mouhamad Jradeh. A new bidimensional neural field model with heterogeneous connection topology. *Neural Comput*, Jul 2010.

Björn M Kampa and Greg J Stuart. Calcium spikes in basal dendrites of layer 5 pyramidal neurons during action potential bursts. *J Neurosci*, 26(28):7424–7432, Jul 2006.

Erich R. Kandel, James H. Schwartz, and Thomas M. Jessell. *Principles of neural science*. McGraw-Hill, New York, 2000.

Kimitaka Kaneki, Osamu Araki, and Minoru Tsukada. Dual synaptic plasticity in the hippocampus: Hebbian and spatiotemporal learning dynamics. *Cogn Neurodyn*, 3(2):153–163, Jun 2009.

T. Kasamatsu and J. D. Pettigrew. Depletion of brain catecholamines: failure of ocular dominance shift after monocular occlusion in kittens. *Science*, 194(4261):206–209, Oct 1976.

L. C. Katz and C. J. Shatz. Synaptic activity and the construction of cortical circuits. *Science*, 274(5290):1133–1138, Nov 1996.

J. P. Kelly and D. C. Van Essen. Cell structure and function in the visual cortex of the cat. *J Physiol*, 238(3):515–547, May 1974.

S. R. Kelso, A. H. Ganong, and T. H. Brown. Hebbian synapses in hippocampus. *Proc Natl Acad Sci U S A*, 83(14):5326–5330, Jul 1986.

Tal Kenet, Dmitri Bibitchkov, Misha Tsodyks, Amiram Grinvald, and Amos Arieli. Spontaneously emerging cortical representations of visual attributes. *Nature*, 425(6961):954–956, Oct 2003.

Frank Kirchhoff. Neuroscience. questionable calcium. *Science*, 327(5970):1212–1213, Mar 2010.

A. Kirkwood, M. C. Rioult, and M. F. Bear. Experience-dependent modification of synaptic plasticity in visual cortex. *Nature*, 381(6582):526–528, Jun 1996.

Katsunori Kitano and Tomoki Fukai. Variability v.s. synchronicity of neuronal activity in local cortical network models with different wiring topologies. *J Comput Neurosci*, 23(2): 237–250, Oct 2007.

Manfred G Kitzbichler, Marie L Smith, Søren R Christensen, and Ed Bullmore. Broadband criticality of human brain network synchronization. *PLoS Comput Biol*, 5(3):e1000314, Mar 2009.

Y. Komatsu and M. Iwakiri. Long-term modification of inhibitory synaptic transmission in developing visual cortex. *Neuroreport*, 4(7):907–910, Jul 1993.

Henri Korn and Philippe Faure. Is there chaos in the brain? II. experimental evidence and related models. *C R Biol*, 326(9):787–840, Sep 2003.

Gregor Kovacic, Louis Tao, Aaditya V Rangan, and David Cai. Fokker-planck description of conductance-based integrate-and-fire neuronal networks. *Phys Rev E Stat Nonlin Soft Matter Phys*, 80(2 Pt 1):021904, Aug 2009.

Birgit Kriener, Tom Tetzlaff, Ad Aertsen, Markus Diesmann, and Stefan Rotter. Correlations and population dynamics in cortical networks. *Neural Comput*, 20(9):2185–2226, Sep 2008.

Birgit Kriener, Moritz Helias, Ad Aertsen, and Stefan Rotter. Correlations in spiking neuronal networks with distance dependent connections. *J Comput Neurosci*, Jul 2009.

Alexandre Kuhn, Ad Aertsen, and Stefan Rotter. Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Comput*, 15(1):67–101, Jan 2003.

A. Kumar, S. Rotter, and A. Aertsen. Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *J Neurosci*, 28(20):5268–5280, 2008a.

Arvind Kumar, Sven Schrader, Ad Aertsen, and Stefan Rotter. The high-conductance state of cortical networks. *Neural Comput*, 20(1):1–43, Jan 2008b.

V. A. Lamme and P. R. Roelfsema. The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci*, 23(11):571–579, Nov 2000.

I. Lampl, I. Reichova, and D. Ferster. Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron*, 22(2):361–374, Feb 1999.

L. Lapicque. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarization. *J Physiol. Pathol. Gen.*, 9:620–635, 1907.

M. E. Larkum, J. J. Zhu, and B. Sakmann. Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. *J Physiol*, 533(Pt 2):447–466, Jun 2001.

Albert K Lee, Ian D Manns, Bert Sakmann, and Michael Brecht. Whole-cell recordings in freely moving rats. *Neuron*, 51(4):399–407, Aug 2006.

Robert Legenstein, Dejan Pecevski, and Wolfgang Maass. A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Comput Biol*, 4(10):e1000180, Oct 2008.

Johannes J Letzkus, Björn M Kampa, and Greg J Stuart. Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *J Neurosci*, 26(41):10420–10429, Oct 2006.

Cheng-Yu T Li, Mu-Ming Poo, and Yang Dan. Burst spiking of a single cortical neuron modifies global brain state. *Science*, 324(5927):643–646, May 2009.

Longnian Lin, Guifen Chen, Kun Xie, Kimberly A Zaia, Shuqing Zhang, and Joe Z Tsien. Large-scale neural ensemble recording in the brains of freely behaving mice. *J Neurosci Methods*, 155(1):28–38, Jul 2006.

J. Lisman. A mechanism for the hebb and the anti-hebb processes underlying learning and memory. *Proc Natl Acad Sci U S A*, 86(23):9574–9578, Dec 1989.

J. Lisman. The CaM kinase II hypothesis for the storage of synaptic memory. *Trends Neurosci*, 17(10):406–412, Oct 1994.

J. Lisman, H. Schulman, and H. Cline. The molecular basis of CaMKII function in synaptic and behavioural memory. *Nat Rev Neurosci*, 3(3):175–190, Mar 2002.

John Lisman and Nelson Spruston. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nat Neurosci*, 8(7):839–841, Jul 2005.

R. R. Llinás, A. A. Grace, and Y. Yarom. In vitro neurons in mammalian cortical layer 4 exhibit intrinsic oscillatory activity in the 10- to 50-hz frequency range. *Proc Natl Acad Sci U S A*, 88(3):897–901, Feb 1991.

Michael London, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E Latham. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466 (7302):123–127, Jul 2010.

Artur Luczak, Peter Bartho, and Kenneth D Harris. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–425, May 2009.

J. Lübke, H. Markram, M. Frotscher, and B. Sakmann. Frequency and dendritic distribution of autapses established by layer 5 pyramidal neurons in the developing rat neocortex: comparison with synaptic innervation of adjacent neurons of the same class. *J Neurosci*, 16 (10):3209–3218, May 1996.

Wolfgang Maass, Thomas Natschlager, and Henry Markram. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput*, 14(11):2531–2560, Nov 2002.

Christian K Machens, Tim Gollisch, Olga Kolesnikova, and Andreas V M Herz. Testing the efficiency of sensory coding with optimal stimulus ensembles. *Neuron*, 47(3):447–456, Aug 2005.

Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.

Roberto Malinow and Robert C Malenka. AMPA receptor trafficking and synaptic plasticity. *Annu Rev Neurosci*, 25:103–126, 2002.

H. Markram and M. Tsodyks. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382(6594):807–810, Aug 1996.

Henry Markram. The blue brain project. *Nat Rev Neurosci*, 7(2):153–160, Feb 2006.

O. Marre, S. El Boustani, Y. Frégnac, and A. Destexhe. Prediction of spatiotemporal patterns of neural activity from pairwise correlations. *Phys Rev Lett*, 102(13):138101, Apr 2009a.

Olivier Marre, Pierre Yger, Andrew Davison, and Yves Frégnac. Reliable recall of spontaneous activity patterns in cortical networks. *Journal of Neuroscience*, 29(46):14596–14606, 2009b.

M. Mattia and P. Del Giudice. Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses. *Neural Comput*, 12(10):2305–2329, Oct 2000.

M. Mayford, J. Wang, E. R. Kandel, and T. J. O'Dell. CaMKII regulates the frequency-response function of hippocampal synapses for the production of both LTD and LTP. *Cell*, 81(6):891–904, Jun 1995.

Alberto Mazzoni, Frédéric D Broccard, Elizabeth Garcia-Perez, Paolo Bonifazi, Maria Elisabetta Ruaro, and Vincent Torre. On the dynamics of the spontaneous activity in neuronal networks. *PLoS One*, 2(5):e439, 2007.

Carver Mead and Misha Mahowald. A silicon model of early visual processing. *Neural Networks*, 1(1):91–97, 1988.

Carsten Mehring, Ulrich Hehl, Masayoshi Kubo, Markus Diesmann, and Ad Aertsen. Activity dynamics and propagation of synchronous spiking in locally connected random networks. *Biol Cybern*, 88(5):395–408, May 2003.

M. R. Mehta, A. K. Lee, and M. A. Wilson. Role of experience and oscillations in transforming a rate code into a temporal code. *Nature*, 417(6890):741–746, Jun 2002.

Ralph Meier, Ulrich Egert, Ad Aertsen, and Martin P Nawrot. FIND–a unified framework for neural data analysis. *Neural Netw*, 21(8):1085–1093, Oct 2008.

Bruce Mockett, Christine Coussens, and Wickliffe C Abraham. NMDA receptor-mediated metaplasticity during the induction of long-term depression by low-frequency stimulation. *Eur J Neurosci*, 15(11):1819–1826, Jun 2002.

Majid H Mohajerani, David A McVea, Matthew Fingas, and Timothy H Murphy. Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *J Neurosci*, 30(10):3745–3751, Mar 2010.

Alik Mokeichev, Michael Okun, Omri Barak, Yonatan Katz, Ohad Ben-Shahar, and Ilan Lampl. Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. *Neuron*, 53(3):413–425, Feb 2007.

C. Monier, J. Fournier, and Y. Frégnac. In vitro and in vivo measures of evoked excitatory and inhibitory conductance dynamics in sensory cortices. *J Neurosci Methods*, 169(2):323–365, Apr 2008.

Johanna M Montgomery and Daniel V Madison. Discrete synaptic states define a major mechanism of synapse plasticity. *Trends Neurosci*, 27(12):744–750, Dec 2004.

G.E. Moore. Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1):82–85, January 1998. ISSN 0018-9219.

A. Morrison, C. Mehring, T. Geisel, A.D. Aertsen, and M. Diesmann. Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput.*, 17:1776–1801, 2005.

Abigail Morrison, Ad Aertsen, and Markus Diesmann. Spike-timing-dependent plasticity in balanced random networks. *Neural Comput*, 19(6):1437–1467, Jun 2007a.

Abigail Morrison, Sirko Straube, Hans Ekkehard Plesser, and Markus Diesmann. Exact subthreshold integration with continuous spike times in discrete-time neural network simulations. *Neural Comput*, 19(1):47–79, Jan 2007b.

Abigail Morrison, Markus Diesmann, and Wulfram Gerstner. Phenomenological models of synaptic plasticity based on spike timing. *Biol Cybern*, 98(6):459–478, Jun 2008.

Anthony Mouraud, Hélène Paugam-moisy, and Didier Puzenat. A distributed and multi-threaded neural event driven simulation framework. In *In Proc. of PDCN'06, Parallel and Distributed Computing and Networks*, 2006.

Scott O Murray, Daniel Kersten, Bruno A Olshausen, Paul Schrater, and David L Woods. Shape perception reduces activity in human primary visual cortex. *Proc Natl Acad Sci U S A*, 99(23):15164–15169, Nov 2002.

Richard Naud, Nicolas Marcille, Claudia Clopath, and Wulfram Gerstner. Firing patterns in the adaptive exponential integrate-and-fire model. *Biol Cybern*, 99(4-5):335–347, Nov 2008.

Ian Nauhaus, Laura Busse, Matteo Carandini, and Dario L Ringach. Stimulus contrast modulates functional connectivity in visual cortex. *Nat Neurosci*, 12(1):70–76, Jan 2009.

Martin P Nawrot, Clemens Boucsein, Victor Rodriguez Molina, Alexa Riehle, Ad Aertsen, and Stefan Rotter. Measurement of variability dynamics in cortical spike trains. *J Neurosci Methods*, 169(2):374–390, Apr 2008.

J. I. Nelson, P. A. Salin, M. H. Munk, M. Arzi, and J. Bullier. Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Vis Neurosci*, 9(1):21–37, Jul 1992.

Thomas Nevian, Matthew E Larkum, Alon Polsky, and Jackie Schiller. Properties of basal dendrites of layer 5 pyramidal neurons: a direct patch-clamp recording study. *Nat Neurosci*, 10(2):206–214, Feb 2007.

A. Ngezahayo, M. Schachner, and A. Artola. Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *J Neurosci*, 20(7):2451–2458, Apr 2000.

Sheila Nirenberg and Peter E Latham. Decoding neuronal spike trains: how important are correlations? *Proc Natl Acad Sci U S A*, 100(12):7348–7353, Jun 2003.

Eilen Nordlie and Hans Ekkehard Plesser. Visualizing neuronal network connectivity with connectivity pattern tables. *Front Neuroinformatics*, 3:39, 2010.

J. O'Keefe and M. L. Recce. Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3(3):317–330, Jul 1993.

Michael Okun and Ilan Lampl. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci*, 11(5):535–537, May 2008.

M. W. Oram, M. C. Wiener, R. Lestienne, and B. J. Richmond. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J Neurophysiol*, 81(6):3021–3033, Jun 1999.

D. Pare, E. Shink, H. Gaudreau, A. Destexhe, and E. J. Lang. Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons in vivo. *J Neurophysiol*, 79(3):1450–1460, Mar 1998.

Dejan Pecevski, Thomas Natschläger, and Klaus Schuch. PCSIM: A parallel simulation environment for neural circuits fully integrated with Python. *Front Neuroinformatics*, 3:11, 2009.

L. M. Pecora and T. L. Carroll. Synchronization in chaotic systems. *Phys Rev Lett*, 64(8): 821–824, 1990.

Laurent Perrinet. *Comment déchiffrer le code impulsionnel de la Vision ? Étude du flux parallèle, asynchrone et épars dans le traitement visuel ultra-rapide.* PhD thesis, Université Paul Sabatier, 2003.

Jean-Pascal Pfister and Wulfram Gerstner. Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci*, 26(38):9673–9682, Sep 2006.

Jonathan W Pillow, Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M Litke, E. J. Chichilnisky, and Eero P Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999, Aug 2008.

M D Plumbley. A hebbian/anti-hebbian network which optimizes information capacity by orthonormalizing the principal subspace, June 02 1993.

Alexandre Pouget, Peter Dayan, and Richard S Zemel. Inference and computation with population codes. *Annu Rev Neurosci*, 26:381–410, 2003.

R. Quian Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107, Jun 2005.

Kanaka Rajan, L. F. Abbott, and Haim Sompolinsky. Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E*, 82(1):011903, Jul 2010.

W. Rall. Membrane time constant of motoneurons. *Science*, 126(3271):454, Sep 1957.

R. P. Rao and D. H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci*, 2(1):79–87, Jan 1999.

R. P. Rao and T. J. Sejnowski. Spike-timing-dependent hebbian plasticity as temporal difference learning. *Neural Comput*, 13(10):2221–2237, Oct 2001.

Alfonso Renart, Jaime de la Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D Harris. The asynchronous state in cortical circuits. *Science*, 327(5965): 587–590, Jan 2010.

Jan Reutimann, Michele Giugliano, and Stefano Fusi. Event-driven simulation of spiking neurons with stochastic dynamics. *Neural Comput*, 15(4):811–830, Apr 2003.

Alex D Reyes. Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nat Neurosci*, 6(6):593–599, Jun 2003.

A. Riccio, S. Ahn, C. M. Davenport, J. A. Blendy, and D. D. Ginty. Mediation by a CREB family transcription factor of NGF-dependent survival of sympathetic neurons. *Science*, 286(5448):2358–2361, Dec 1999.

Dario L Ringach. Spontaneous and driven cortical activity: implications for computation. *Curr Opin Neurobiol*, 19(4):439–444, Aug 2009.

Per E Roland, Akitoshi Hanazawa, Calle Undeman, David Eriksson, Tamas Tompa, Hiroyuki Nakamura, Sonata Valentiniene, and Bashir Ahmed. Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. *Proc Natl Acad Sci U S A*, 103(33):12586–12591, Aug 2006.

R. Romo and W. Schultz. Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. *J Neurophysiol*, 63(3):592–606, Mar 1990.

F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev*, 65(6):386–408, Nov 1958.

Arnd Roth and Michael London. Rebuilding dendritic democracy. focus on "equalization of synaptic efficacy by activity- and timing-dependent synaptic plasticity". *J Neurophysiol*, 91(5):1941–1942, May 2004.

Jonathan E Rubin, Richard C Gerkin, Guo-Qiang Bi, and Carson C Chow. Calcium time course as a signal for spike-timing-dependent plasticity. *J Neurophysiol*, 93(5):2600–2613, May 2005.

Michelle Rudolph and Alain Destexhe. Analytical integrate-and-fire neuron models with conductance-based dynamics for event-driven simulation strategies. *Neural Comput*, 18 (9):2146–2210, Sep 2006.

Michelle Rudolph and Alain Destexhe. How much can we trust neural simulation strategies? *Neurocomputing*, 70(10-12):1966–1969, 2007.

C. D. Salzman, K. H. Britten, and W. T. Newsome. Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346(6280):174–177, Jul 1990.

A. Saudargiene, B. Porr, and F. Wörgötter. Synaptic modifications depend on synapse location and activity: a biophysical model of STDP. *Biosystems*, 79(1-3):3–10, 2005.

Cristina Savin, Prashant Joshi, and Jochen Triesch. Independent component analysis in spiking neurons. *PLoS Comput Biol*, 6(4):e1000757, 2010.

Nathaniel B Sawtell and Curtis C Bell. Adaptive processing in electrosensory systems: links to cerebellar plasticity and learning. *J Physiol Paris*, 102(4-6):223–232, 2008.

Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialekshea. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440 (7087):1007–1012, Apr 2006.

James Schummers, Hongbo Yu, and Mriganka Sur. Tuned responses of astrocytes and their influence on hemodynamic signals in the visual cortex. *Science*, 320(5883):1638–1643, Jun 2008.

M. N. Shadlen and J. A. Movshon. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron*, 24(1):67–77, 111–25, Sep 1999.

M. N. Shadlen and W. T. Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*, 18(10):3870–3896, 1998.

Maoz Shamir and Haim Sompolinsky. Nonlinear population codes. *Neural Comput*, 16(6): 1105–1136, Jun 2004.

C. J. Shatz. Emergence of order in visual system development. *Proc Natl Acad Sci U S A*, 93 (2):602–608, Jan 1996.

G. M. Shepherd, R. K. Brayton, J. P. Miller, I. Segev, J. Rinzel, and W. Rall. Signal enhancement in distal cortical dendrites by means of interactions between active dendritic spines. *Proc Natl Acad Sci U S A*, 82(7):2192–2195, Apr 1985.

Harel Z Shouval, Gastone C Castellani, Brian S Blais, Luk C Yeung, and Leon N Cooper. Converging evidence for a simplified biophysical model of synaptic plasticity. *Biol Cybern*, 87(5-6):383–391, Dec 2002.

Yousheng Shu, Andrea Hasenstaub, and David A McCormick. Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–293, May 2003.

W. Singer and C. M. Gray. Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci*, 18:555–586, 1995.

Benoît Siri, Hugues Berry, Bruno Cessac, Bruno Delord, and Mathias Quoy. A mathematical analysis of the effects of hebbian learning rules on the dynamics and structure of discrete-time random recurrent neural networks. *Neural Comput*, 20(12):2937–2966, Dec 2008.

P. J. Sjöström, G. G. Turrigiano, and S. B. Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164, Dec 2001.

Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. Neocortical ltd via coincident activation of presynaptic nmda and cannabinoid receptors. *Neuron*, 39(4):641–654, Aug 2003.

Matthew A Smith and Adam Kohn. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci*, 28(48):12591–12603, Nov 2008.

H. Sompolinsky, A. Crisanti, and H. J. Sommers. Chaos in random neural networks. *Phys Rev Lett*, 61(3):259–262, 1988.

S. Song and L. F. Abbott. Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, 32(2):339–350, Oct 2001.

R. W. Sperry. Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proc Natl Acad Sci U S A*, 50:703–710, Oct 1963.

D. N. Spinelli and F. E. Jensen. Plasticity: the mirror of experience. *Science*, 203(4375): 75–78, Jan 1979.

Nicholas C Spitzer. Electrical activity in early neuronal development. *Nature*, 444(7120): 707–712, Dec 2006.

Henning Sprekeler, Christian Michaelis, and Laurenz Wiskott. Slowness: An objective for spike-timing dependent plasticity? In *Proc. 2nd Bernstein Symposium for Computational Neuroscience 2006, Berlin, October 1–3*, page 24. Bernstein Center for Computational Neuroscience (BCCN) Berlin, October 2006. Zusammenfassung.

Dominic Standage, Sajiya Jalil, and Thomas Trappenberg. Computational consequences of experimentally derived spike-time and weight dependent plasticity rules. *Biol Cybern*, Apr 2007.

P. K. Stanton and T. J. Sejnowski. Associative long-term depression in the hippocampus induced by hebbian covariance. *Nature*, 339(6221):215–218, May 1989.

G. S. Stent. A physiological mechanism for hebb's postulate of learning. *Proc Natl Acad Sci U S A*, 70(4):997–1001, Apr 1973.

M. Steriade, A. Nunez, and F. Amzica. Intracellular analysis of relations between the slow ($<$ 1 hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *J Neurosci*, 13(8):3266–3283, Aug 1993a.

M. Steriade, A. Nunez, and F. Amzica. A novel slow ($<$ 1 hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J Neurosci*, 13(8):3252–3265, Aug 1993b.

M. Steriade, I. Timofeev, and F. Grenier. Natural waking and sleep states: a view from inside neocortical neurons. *J Neurophysiol*, 85(5):1969–1985, May 2001.

Dan D Stettler, Aniruddha Das, Jean Bennett, and Charles D Gilbert. Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron*, 36(4):739–750, Nov 2002.

D. Sulzer and S. Rayport. Dale's principle and glutamate corelease from ventral midbrain dopamine neurons. *Amino Acids*, 19(1):45–52, 2000.

M. Sur, A. Angelucci, and J. Sharma. Rewiring cortex: the role of patterned activity in development and plasticity of neocortical circuits. *J Neurobiol*, 41(1):33–43, Oct 1999.

David Sussillo and L. F. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557, Aug 2009.

Harvey A Swadlow and Jose-Manuel Alonso. Spikes are making waves in the visual cortex. *Nat Neurosci*, 12(1):10–11, Jan 2009.

J. Teramae and Tomoki Fukai. Temporal precision of spike response to fluctuating input in pulse-coupled networks of oscillating neurons. *Phys Rev Lett*, 101(24):248105, Dec 2008.

Alex M Thomson and A. Peter Bannister. Interlaminar connections in the neocortex. *Cereb Cortex*, 13(1):5–14, Jan 2003.

Alex M Thomson, A. Peter Bannister, Audrey Mercer, and Oliver T Morris. Target and temporal pattern selection at neocortical synapses. *Philos Trans R Soc Lond B Biol Sci*, 357 (1428):1781–1791, Dec 2002.

S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381(6582):520–522, Jun 1996.

P. H. Tiesinga, J. V. José, and T. J. Sejnowski. Comparison of current-driven and conductance-driven neocortical model neurons with hodgkin-huxley voltage-gated channels. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics*, 62(6 Pt B):8413–8419, Dec 2000.

Paul Tiesinga, Jean-Marc Fellous, and Terrence J Sejnowski. Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci*, 9(2):97–107, Feb 2008.

I. Timofeev, F. Grenier, M. Bazhenov, T. J. Sejnowski, and M. Steriade. Origin of slow cortical oscillations in deafferented cortical slabs. *Cereb Cortex*, 10(12):1185–1199, 2000.

Taro Toyoizumi, Jean-Pascal Pfister, Kazuyuki Aihara, and Wulfram Gerstner. Optimality model of unsupervised spike-timing-dependent plasticity: Synaptic memory and weight distribution. *Neural Comput*, 19(3):639–671, Mar 2007.

Nicolas X Tritsch, Eunyoung Yi, Jonathan E Gale, Elisabeth Glowatzki, and Dwight E Bergles. The origin of spontaneous activity in the developing auditory system. *Nature*, 450(7166):50–55, Nov 2007.

D. Y. Ts'o, C. D. Gilbert, and T. N. Wiesel. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci*, 6(4):1160–1170, Apr 1986.

M. Tsodyks, T. Kenet, A. Grinvald, and A. Arieli. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, 286(5446):1943–1946, 1999.

M. Tsodyks, A. Uziel, and H. Markram. Synchrony generation in recurrent networks with frequency-dependent synapses. *J Neurosci*, 20(1):RC50, 2000.

Gina G Turrigiano and Sacha B Nelson. Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci*, 5(2):97–107, Feb 2004.

Thanos Tzounopoulos, Yuil Kim, Donata Oertel, and Laurence O Trussell. Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nat Neurosci*, 7(7):719–725, Jul 2004.

Marius Usher, Martin Stemmler, Christof Koch, and Zeev Olami. Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. *Neural Computation*, 6(5):795–836, 1994.

E. Vaadia, I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A. Aertsen. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373 (6514):515–518, 1995.

M. C. van Rossum, G. Q. Bi, and G. G. Turrigiano. Stable hebbian learning from spike timing-dependent plasticity. *J Neurosci*, 20(23):8812–8821, Dec 2000.

C. van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, Dec 1996.

C. van Vreeswijk and H. Sompolinsky. Chaotic balanced state in a model of cortical circuits. *Neural Comput*, 10(6):1321–1371, 1998.

T. P. Vogels, K. Rajan, and L. F. Abbott. Neural network dynamics. *Annu Rev Neurosci*, 28: 357–376, 2005.

Nicole Voges, Ad Aertsen, and Stefan Rotter. Statistical analysis of spatially embedded networks: From grid to random node positions. *Neurocomputing*, 70(10-12):1833–1837, 2007.

C. von der Malsburg. Binding in models of perception and brain function. *Curr Opin Neurobiol*, 5(4):520–526, Aug 1995.

Huai-Xing Wang, Richard C Gerkin, David W Nauen, and Guo-Qiang Bi. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat Neurosci*, 8(2): 187–193, Feb 2005.

Oliver G Wenisch, Joachim Noll, and J. Leo van Hemmen. Spontaneously emerging direction selectivity maps in visual cortex through STDP. *Biol Cybern*, 93(4):239–247, Oct 2005.

T. N. Wiesel. Postnatal development of the visual cortex and the influence of environment. *Nature*, 299(5884):583–591, Oct 1982.

T. N. Wiesel and D. H. Hubel. Extent of recovery from the effects of visual deprivation in kittens. *J Neurophysiol*, 28(6):1060–1072, Nov 1965.

H. R. Wilson and J. D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J*, 12(1):1–24, Jan 1972.

W. D. Winters. Effects of drugs on the electrical activity of the brain: anesthetics. *Annu Rev Pharmacol Toxicol*, 16:413–426, 1976.

Adrien Wohrer and Pierre Kornprobst. Virtual retina: a biological retina model and simulator, with contrast gain control. *J Comput Neurosci*, 26(2):219–249, Apr 2009.

W. T. Wong, J. R. Sanes, and R. O. Wong. Developmentally regulated spontaneous activity in the embryonic chick retina. *J Neurosci*, 18(21):8839–8852, Nov 1998.

F. Wörgötter, K. Suder, Y. Zhao, N. Kerscher, U. T. Eysel, and K. Funke. State-dependent receptive-field restructuring in the visual cortex. *Nature*, 396(6707):165–168, Nov 1998.

S. N. Yang, Y. G. Tang, and R. S. Zucker. Selective induction of LTP and LTD by postsynaptic [ca2+]i elevation. *J Neurophysiol*, 81(2):781–787, Feb 1999.

M. Yokoi, K. Mori, and S. Nakanishi. Refinement of odor molecule tuning by dendrodendritic synaptic inhibition in the olfactory bulb. *Proc Natl Acad Sci U S A*, 92(8):3371–3375, Apr 1995.

Jennie Z Young, Carolina Isiegas, Ted Abel, and Peter V Nguyen. Metaplasticity of the late-phase of long-term potentiation: a critical role for protein kinase a in synaptic tagging. *Eur J Neurosci*, 23(7):1784–1794, Apr 2006.

Joshua M Young, Wioletta J Waleszczyk, Chun Wang, Michael B Calford, Bogdan Dreher, and Klaus Obermayer. Cortical reorganization consistent with spike timing-but not correlation-dependent plasticity. *Nat Neurosci*, 10(7):887–895, Jul 2007.

Ji-Chuan Zhang, Pak-Ming Lau, and Guo-Qiang Bi. Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. *Proc Natl Acad Sci U S A*, 106(31):13028–13033, Aug 2009.

L. I. Zhang, H. W. Tao, C. E. Holt, W. A. Harris, and M. Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395(6697):37–44, Sep 1998.

Lian Zhang, Timo Kirschstein, Britta Sommersberg, Malte Merkens, Denise Manahan-Vaughan, Ype Elgersma, and Heinz Beck. Hippocampal synaptic metaplasticity requires inhibitory autophosphorylation of ca2+/calmodulin-dependent kinase II. *J Neurosci*, 25 (33):7697–7707, Aug 2005.

E. Zohary, M. N. Shadlen, and W. T. Newsome. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370(6485):140–143, Jul 1994.

Quan Zou and Alain Destexhe. Kinetic models of spike-timing dependent plasticity and their functional consequences in detecting correlations. *Biol Cybern*, 97(1):81–97, Jul 2007.

Larry S Zweifel, Rejji Kuruvilla, and David D Ginty. Functions and mechanisms of retrograde neurotrophin signalling. *Nat Rev Neurosci*, 6(8):615–625, Aug 2005.