A New Perspective on the Organization of Neuronal Activity in Neocortex

and Its Implications for Neuronal Information Processing and Coding

von der Fakultät für Mathematik und Naturwissenschaften der

CARL VON OSSIETZKY UNIVERSITÄT OLDENBURG

zur Erlangung des Grades und Titels eines
Doktors der Naturwissenschaften (Dr. rer. nat.)
angenommene Dissertation von

Kai S. Gansel
geboren am 23.12.1974
in Wangen im Allgäu

2014

Gutachter: Prof. Dr. Josef Ammermüller
Zweitgutachter: Prof. Dr. Dr. h. c. mult. Wolf Singer

Tag der Disputation: 4. Dezember 2014

Acknowledgements

This thesis was funded in part by the Max Planck Society and the German Research Foundation as part of the Collaborative Research Center 269 "Molecular and cellular bases of neuronal organization" and was carried out at the Max Planck Institute for Brain Research in Frankfurt am Main.

The thesis was supervised by Prof. Wolf Singer, whom I would like to thank for giving me the unique opportunity to pursue my own research ideas in a first-rate working environment, for allowing me to be part of a fascinating research community, and for being a prudent teacher. After a long and frustrating day in the lab, he pointed out that if one does not find any meaningful information in the data, one is probably asking the wrong question – simple but true! This comment became the crystallization point from which the ideas presented in this thesis evolved.

I would like to extend my gratitude to the many colleagues and staff at the institute who have helped me in my work in a multitude of ways. Specifically, I thank Dr. Valérie Pothin for introducing me to the techniques of how to prepare acute brain slices, Prof. Sergio Neuenschwander for providing the data recording software, and Prof. Nan-Hui Chen for providing the most elaborate spike sorting software I have ever had the pleasure to work with. I am also thankful for the excellent technical support provided by Dieter Herzberger, Andreas Umminger and Walter Ankenbrand, without whom the construction of the recording setup would have been practically impossible. Above all technical aspects, I would especially like to thank Dr. Matthias Munk who encouraged me to accept the unexpected and reveal its meaning.

Aside from my colleagues in Frankfurt, I am much obliged to Prof. Raoul-Martin Memmesheimer, then at the Max Planck Institute for Dynamics and Self-Organization in Göttingen, who volunteered to carry out the most time- and computer memory-consuming parts of the data analysis.

Finally, I would like to express my sincerest thanks and appreciation to Prof. Josef Ammermüller, who continues to be a kindly mentor and generously agreed to serve as the primary reviewer of my thesis.

I am greatly indebted to Sonja and to my family for their continual support.

Kai Gansel, October 2014

Contents

Abbreviations

ACSF	artificial cerebrospinal fluid
AMPA	α-amino-3-hydroxyl-5-methyl-4-isoxazole-propionate
BIC	bicuculline, (5*S*)-5-[(6*R*)-6,8-dihydro-8-oxofuro[3,4-*e*]-1,3-benzodioxol-6-yl]-5,6,7,8-tetrahydro-6,6-dimethyl-1,3-dioxolo[4,5-*g*]isoquinolinium
CCH	carbachol, (2-hydroxyethyl)trimethylammonium chloride carbamate
CGP	CGP-35348, (3-aminopropyl)(diethoxymethyl)phosphinic acid
GABA	γ-aminobutyric acid
IEI	inter-event interval
ISI	inter-spike interval
LFP	local field potential
MD	mean absolute deviation
NMDA	*N*-methyl-*D*-aspartat
PDF	probability density function
SD	standard deviation
SEM	standard error of the mean

Zusammenfassung

Auf der Grundlage der Auffassung, dass jede kognitive Leistung und alle bewussten wie unbewussten psychischen Prozesse des Menschen auf raumzeitlich organisierter elektrischer Hirnaktivität beruhen, versucht diese Arbeit, grundlegende Prinzipien in der Organisation multineuronaler Aktivitätsmuster aufzudecken und zu kognitiven Vorgängen und Funktionen in eine theoretische Beziehung zu setzen. Gegenstand der Betrachtungen ist die raumzeitliche Koordination neuronaler Aktionspotentiale, genauer die Voraussetzungen und Mechanismen für die Entstehung sich wiederholender Erregungsmuster sowie ihre raumzeitliche Ausprägung als potentielle Bestandteile eines neuronalen Codes zur Darstellung und Verarbeitung von Information.

Kapitel I gibt einen kurzen Abriss über klassische Konzepte neuronaler Codierung und diskutiert Argumente für oder gegen verschiedene Theorien zur neuronalen Repräsentation kognitiver Inhalte. Anhand umfangreicher experimenteller Befunde und einer Reihe theoretischer Überlegungen wird dargelegt, dass einzelne kognitive Inhalte (wie die Wahrnehmung eines Objekts, die Planung einer Bewegung oder ein einzelner Gedanke) kaum durch die Aktivität einer einzelnen Zelle darstellbar sind, sondern durch die koordinierte Aktivität einer Gruppe von Zellen repräsentiert werden, dass die von einem Aktionspotential übermittelte Information von seinem genauen Timing im Verhältnis zu anderen Signalen abhängt, und dass synchrone neuronale Oszillationen in Abhängigkeit kognitiver Prozesse das Auftreten und die Weiterleitung neuronaler Signale im Netzwerk beeinflussen und so konstitutiv zur Gestaltung multineuronaler Erregungsmuster beitragen.

Ausgehend von diesen Erkenntnissen ergeben sich im Hinblick auf die gezielte Koordination räumlich verteilter Aktionspotentiale Fragen zu den relevanten Zeitskalen und den zugrundeliegenden neuronalen Mechanismen. So wird seit langem darüber diskutiert, in welchen Zeitfenstern und mit welcher zeitlichen Präzision neuronale Signale aufeinander abgestimmt sind, und welche Rolle neuronalen Oszillationen bei der selektiven Synchronisation von Gruppen von Zellen zukommt. Die vorliegende Arbeit bietet zu diesen Fragen einen experimentellen Zugang, indem sie die Dynamik multineuronaler Erregungsmuster in Hirnschnitten untersucht, in denen das Auftreten synchroner Netzwerk-Oszillationen ausgeschlossen werden kann, ohne die Erregbarkeit von Zellen oder die Funktion von Transmitter-Rezeptoren beeinflussen zu müssen. Zu diesem Zweck ist ein spezieller Messplatz entwickelt worden, der es ermöglicht, unter pharmakologischer Kontrolle die elektrische Aktivität lokaler Zellpopulationen innerhalb des Hirnschnitts mit Hilfe einer Elektroden-Matrix parallel abzuleiten. Ziel war es, die beobachteten Aktivitätsmuster auf verschiedenen Zeitskalen zu analysieren und so zu einer umfassenden Beschreibung der raumzeitlichen Organisation lokaler Erregungsmuster zu gelangen.

Um beurteilen zu können, ob die erfassten Muster zufällig aufgetreten sind oder statistisch signifikante Strukturen aufweisen, ist eine neue Methode für das Auffinden signifikanter, sich wiederholender Muster in parallelen Zeitreihen entwickelt worden, die keinerlei Annahmen über die statistischen Eigenschaften der untersuchten Daten macht und sowohl umschriebene raumzeitliche Muster als auch die zeitliche Abfolge dieser Muster analysiert. Kapitel II stellt diese Methode vor, evaluiert mittels simulierter Daten ihr Leistungsvermögen und vergleicht sie mit ähnlichen Methoden, deren Limitierungen sie größtenteils aufhebt: So bedient sie sich eines flexiblen und effizienten Suchalgorithmus, der sich wiederholende Muster findet, ohne eine spezielle Struktur der Muster vorauszusetzen, und umgeht das Problem der kombinatorischen Explosion in der Mustersuche durch eine zeitaufgelöste, heuristische Analyse der Korrelationen zwischen den jeweils an einem Muster beteiligten Zellen. Die von zufälligen Mustern ausgehende Nullhypothese wird durch einen neuen Typ von Surrogat-Daten dargestellt, der sich wie bisher beschriebene Typen auch aus den originalen Daten ableitet und selektiv deren Struktur unterhalb einer bestimmten Zeitskala randomisiert, dessen Teststärke sich jedoch als vorteilhaft erwiesen hat. Die Zuhilfenahme von Monte-Carlo-Simulationen und der Verzicht auf ein analytisches Modell der Null-

hypothese vermeidet den Fluch der Dimensionalität, der parametrischen Methoden zur statistischen Validierung hochdimensionaler Muster unweigerlich anhaftet. Signifikante Muster werden kontinuierlich registriert und ermöglichen so eine Korrelation mit experimentellen Faktoren mit einer nur durch die Samplingrate der Signale begrenzten Zeitauflösung.

Kapitel III berichtet von der Durchführung paralleler Messungen multineuronaler Erregungsmuster in Hirnschnitten ohne rhythmische Netzwerk-Aktivität und zeigt die Ergebnisse der Analyse der Muster nach der in Kapitel II beschriebenen Methode. Der Fokus dieser Arbeit lag dabei auf der Großhirnrinde (Neocortex) als dem Teil des Gehirns, dem die meisten höheren kognitiven Funktionen zugeschrieben werden. Als Modell diente der visuelle Cortex der Ratte, dessen anatomische Architektur in weiten Teilen derjenigen anderer corticaler Areale auch in anderen Säugetierhirnen einschließlich dem des Menschen entspricht, was eine Übertragbarkeit der Ergebnisse nahelegt. Diese zeichnen ein differenziertes Bild neocorticaler Dynamik: Obwohl multineuronale Feuermuster in Zeitfenstern von bis zu 50 ms registriert wurden, wiesen signifikante Muster typischerweise eine Dauer von weniger als 5 ms auf. Bemerkenswert hieran ist, dass die Dauer der Muster nicht mit der Entfernung der beteiligten Zellen korrelierte, was bedeutet, dass sich auch weit auseinander liegende, unverbundene Zellen, die keinen gemeinsamen Input erhalten, ohne rhythmische Netzwerk-Aktivität gezielt synchronisieren können. Die relative zeitliche Präzision der an den Mustern beteiligten Signale betrug dabei im Median ~ 0.58 ms. Außerdem bildeten aufeinanderfolgende Gruppen synchron aktiver Zellen sich überzufällig oft wiederholende Sequenzen, die mehrere Sekunden lang sein konnten und sich durch eine zeitliche Präzision im Timing der Gruppen von durchschnittlich ~ 25.13 ms auszeichneten. Diese Sequenzen stimmen also nicht mit dem Konzept von "synfire chains" überein, das von einem wesentlich präziseren Timing ausgeht, sondern eher mit dem von Hebb'schen "phase sequences". Sowohl synchrone Feuermuster als auch längere Sequenzen koordinierter Aktivität reflektierten bekannte synaptische Schaltkreise und zeigten eine ausgeprägte Abhängigkeit vom neuromodulatorischen Zustand des Hirnschnitts.

In Kapitel IV wird versucht, anhand dieser Befunde die in Kapitel I erörterten Konzepte neuronaler Informationsverarbeitung und Codierung zu verfeinern und die sich daraus ergebenden Konsequenzen in einen allgemeinen neurokognitiven Kontext einzubetten. Um die beobachtete Synchronisation räumlich verteilter Zellen zu erklären, wird folgende Hypothese formuliert: Neurone, die durch eine wiederholte Abfolge von synaptischen Eingangssignalen überschwellig erregt werden, können durch Anpassung der Effizienz der beteiligten Synapsen lernen, auf die ersten Signale in der Abfolge zu antworten. Solche Zellen, für die die ersten Signale einer wiederholten Serie von Eingängen auf den gleichen Zeitpunkt fallen, würden dann in der Folge synchron aktiv sein, sobald sie die entsprechenden Eingangssignale empfangen. Die Hypothese trifft mehrere Vorhersagen: Erstens wäre es unerheblich, in welchen Schaltkreisen sich die Zellen befinden und woher ihre Eingangssignale kommen, solange diese Signale zusammen auftreten, zweitens sollten sich wiederholende raumzeitliche Erregungsmuster mit der Zeit komprimiert werden, bis alle oder ein Teil der Signale synchronisiert sind, und drittens sollte diese Kompression mit einer Verringerung der Anzahl der beteiligten Signale einhergehen. Am Ende stünde eine selektive Gruppe von Zellen, die auf ein bekanntes Ereignis innerhalb weniger Millisekunden mit dem synchronen Feuern von Aktionspotentialen antworten würden. Ein somit erlerntes Antwortmuster wäre zusätzlich modulierbar durch synchrone Netzwerk-Oszillationen, die für die synaptische Integration verteilter Signale einen flexiblen, raumzeitlichen Kontext bilden. Es wird erläutert, dass Oszillationen im Membranpotential einer Zelle nur in bestimmten Fällen dazu geeignet sind, das Timing ausgehender Aktionspotentiale zu beeinflussen, und in allen anderen Fällen schlicht darüber entscheiden, ob zu einem gegebenen Zeitpunkt die Zelle überschwellig depolarisiert werden kann oder nicht.

Schließlich wird über die potentiellen kognitiven Funktionen synchron aktiver Zellgruppen sowie längerer sich wiederholender Sequenzen koordinierter multineuronaler Aktivität spekuliert. Das von Chalmers formulierte und aus der Not des "hard problem of consciousness" heraus geborene "principle of structural coherence" wird adoptiert und führt zu der Vermutung, dass die kurzzeitige Synchronisation einer selektiven Gruppe von Zellen einem sinnstiftenden Übergang von einem kognitiven Zustand in einen anderen entspricht, wobei sich zuvor getrennte Aspekte einer sinnvoll erscheinenden Information zu einem einheitlichen Erleben zusammenfinden, so wie bei der Wahrnehmung eines Objekts, dem Treffen einer Entscheidung oder dem Bewußtwerden eines Gedankens. In Verbindung mit längeren Sequenzen neuronaler Aktivität wird das Problem der sinnvollen Verknüpfung zweier aufeinanderfolgender Zustände oder Gedanken bei gleichzeitiger Flexibilität des Gedankenganges diskutiert. Das klassische Modell Hebb'scher "phase sequences" trägt dieser Problematik Rechnung und wird durch die hier gefundene Dynamik neocorticaler Aktivitätsmuster gestützt.

Summary

Based on the assumption that all mental activity and all conscious and unconscious cognitive processes rely on spatially and temporally organized electrical brain activity, this work tries to uncover fundamental principles in the organization of multineuronal activity patterns and to relate them to cognitive processes and functions in a theoretical way. Subject of investigation is the spatiotemporal coordination of neuronal action potentials, or more specifically, the preconditions and mechanisms for the emergence of repeating discharge patterns and their spatiotemporal appearance as potential constituents of a neuronal code for the representation and processing of information.

Chapter I provides a brief account of classical concepts of neuronal coding and discusses arguments in favor or against different theories about the neural representation of cognitive contents. Referring to a large number of experimental findings and several theoretical considerations, it is argued that individual cognitive contents (like the perception of an object, the planning of a coordinated movement sequence or a single idea) can hardly be represented by the activity of a single cell, but require the coordinated activity of a group of cells, that the information that is conveyed by a single action potential depends on its precise timing relative to other signals, and that synchronous neuronal oscillations control the propagation of neuronal signals in the network in accord with cognitive processes and thus contribute in a constitutive way to the formation of multineuronal discharge patterns.

Building on these insights, several questions arise with respect to a directed coordination of distributed discharges regarding relevant timescales and the underlying neuronal mechanisms. In particular, it has been asked in which time windows and with which temporal precision neuronal signals could possibly be coordinated, and which role neuronal oscillations precisely play in the selective synchronization of groups of cells. The work at hand provides an experimental approach to these questions by investigating the dynamics of multineuronal discharge patterns in brain slices, in which synchronous network oscillations can be ruled out without the need to interfere with the excitability of cells or the functioning of transmitter receptors. For this purpose, a special recording setup has been designed to record the electrical activity of local cell populations inside the brain slice in a parallel fashion by means of a multielectrode array while having full pharmacological control. It was intended to analyze the observed activity patterns on various timescales to achieve a comprehensive description of the spatiotemporal organization of locally confined multineuronal discharge patterns.

To assess if the recorded patterns occurred by chance or if they show any significant structure, a new method for the detection of statistically significant repeating patterns in parallel time series has been developed that makes no assumptions about the statistical properties of the data and examines both circumscribed spatiotemporal patterns and sequences of these patterns. Chapter II introduces this method, evaluates its performance by making use of simulated data and compares it to similar approaches, whose limitations are largely resolved: It employs a flexible and efficient search algorithm that finds repeating patterns without implying any particular structure of the patterns, and circumvents the problem of the combinatorial explosion in the pattern search through a time-resolved heuristic analysis of the correlations between the cells that participate in a given pattern. Assuming random patterns, the null hypothesis is represented by a new type of surrogate data that, like other types, derives directly from the original data and randomizes its temporal structure below a certain timescale, but which has superior test power. The use of Monte-Carlo simulations instead of an analytical model of the null hypothesis avoids the curse of dimensionality inherent in any parametric method for the statistical validation of high-dimensional patterns. Significant patterns are registered in a continuous fashion and thus allow for a correlation with experimental factors and conditions with a temporal resolution that is limited only by the sampling rate of the investigated signals.

Chapter III reports on parallel recordings of multineuronal discharge patterns in brain slices exhibiting no rhythmic network activity and shows the results of the analysis of the patterns using the method described in Chapter II. Since most higher cognitive functions are attributed to the neocortex, it is the chosen structure of interest. Its anatomical architecture is largely preserved across brain areas and mammalian species, so that one can hope to gain some general insights into cortical neuronal information processing and coding by investigating a specific area in a certain animal. Here, the visual cortex of the rat served as a model system. The results portray a differentiated picture of neocortical dynamics: Although multineuronal firing patterns had been registered in time windows of up to 50 ms, significant patterns typically had durations of less than 5 ms. Importantly, the durations of the patterns were not correlated with the spatial distance between the participating cells, which means that even widely distributed neurons that are unlikely to be directly connected or to receive common input and are not synchronized by network oscillations may align their firing on the timescale of milliseconds. The median temporal precision of the signals that were contributing to the patterns was ~ 0.58 ms. In addition, groups of cells activated in direct succession were organized to a significant degree into repeating sequences that could have durations of several seconds. The median temporal precision of the relative timing of the successively active groups was ~ 25.13 ms. Thus, these sequences do not conform to the concept of "synfire chains", but rather resemble Hebbian "phase sequences". Synchronous firing patterns and longer sequences of coordinated activity reflected the synaptic circuitry and strongly depended on the neuromodulatory state of the cortical slice.

Based on these findings, Chapter IV tries to refine the concepts of neuronal information processing and coding that have been discussed in Chapter I and to incorporate the resultant consequences into a general neurocognitive framework. To explain the observed synchronization of spatially distributed cells, the following hypothesis is proposed: Neurons that are excited by a repeated sequence of synaptic inputs may learn to selectively respond to the very first signals of the sequence by adaptations of the efficiency of the involved synapses. Those cells that receive the first signals of a repeated sequence of synaptic inputs at the same point in time would then be active together. The hypothesis makes several predictions: First, the position of the cells in the network as well as the source of their input signals would be irrelevant as long as the signals arrive at the same point in time; second, repeating spatiotemporal discharge patterns should get more and more compressed until all or some part of the signals are synchronized; and third, this compression should be accompanied by a sparsening of the involved signals. In this way, selective groups of cells could emerge that would respond to some known event with synchronous action potential firing. Such a learned response pattern could further be modulated by synchronous network oscillations that provide a flexible, spatiotemporal context for the synaptic integration of distributed signals. It is argued that oscillations of a cell's membrane potential may influence the timing of action potentials only under certain conditions; if these conditions are not met, membrane potential oscillations act more like a logic gate determining whether a cell can be depolarized above threshold at a given point in time or not.

Finally, some speculations about potential cognitive functions of synchronously active cell groups and long repeating sequences of coordinated multineuronal activity are presented. The "principle of structural coherence", suggested by Chalmers as a resort from the "hard problem of consciousness", is adopted and leads to the conjecture that the transient synchronization of a selective group of cells corresponds to a meaningful transition from one cognitive state to another, while previously separate aspects of some seemingly meaningful information combine to give way to a holistic experience, like perceiving an object, reaching a decision, or becoming aware of some thought or idea. The problem of how to maintain meaning in the succession of states or thoughts while preserving flexibility of the mind is discussed in connection with the sequential coordination of neuronal activity. The classical model of Hebbian "phase sequences" accommodates this problem and is corroborated by the found dynamics of neocortical activity patterns.

Prologue

"Imagine suddenly entering a dreamlike plane where the usual laws of information derived from perception seem not applicable. Your eyes are the quickest to adapt and yield the first impression: Sweat drips on four screaming cellos as their players, a quartet of concentrated young men, headbang in digital unison. As your senses are ripped open to absorb this primal and somewhat scary atmosphere, you hear the harsh sound of bows striking tortured strings. After penetrating your skull the decibel storm of raucous riffs and blistering glissandos starts rearranging the synapses in your brain. […]*"*

E. Toppinen (2000). ***Cult***, Mercury Records.

The grand challenge

The human mind in all its facets, varying from sense to sensibility, subconscious processes to conscious awarenesses and subjective experience to collective cognition through social interactions, is truly one of the most complicated and complex things we may try to understand. Step into the above scene:

You are watching four men playing music on their cellos. While your senses are attuning to the scenery, various details of different modality enter your mind and are cognitively bound to a holistic percept full of emotional content. In our present understanding, this percept is by no means a simple blueprint of reality, rather reflecting an idea of it emanating from the interplay of the incoming information with the observer's internal state. All sensory input is interpreted in the light of preexisting conceptions and its meaning is attributed accordingly. As a consequence, we cannot perceive what we cannot imagine. We can, however, call into consciousness anything we may think of purely by imagination, be it memories of past events or sheer fantasy – as is directly demonstrated here. Following a few informative sentences, the reader is left with a vivid impression of the whole scene. The words used in the description are functioning as abstract symbols that are associated with certain prelearned meanings. From this perspective, language and the inner structure of cognitive conceptions are intimately related.

Besides perception, action is another aspect of human cognition that is illustrated in this sequence. The four cello players are engaged in a concerted exercise that is highly temporally structured. In order to succeed, they need to continuously align their playing with a multitude of acoustic signals on different timescales and act in a precisely timed manner. Moreover, the task requires full attention. If they lose concentration, they will fail.

Of course, these examples do by far not cover the whole variety of psychological phenomena, but they are well suited to exemplify some of the most fundamental questions in the cognitive sciences: How do sensory stimuli that are distributed in space and time and across modalities converge towards coherent perceptions? How do we create mental images from abstract symbols or pure imagination? How are our cognitive conceptions shaped by learning? How do we coordinate our actions in time? How do attention and expectancy control our conscious awarenesses? And finally, how does consciousness come about, and how do we reach perceptual and other decisions? The work at hand does not address any of these issues directly, but focuses on the underlying neural function. A concise rationale for this approach has been given by Hebb a long time ago (Hebb 1949, pp. XIII–XIV): "Modern psychology takes completely for granted that behavior and neural function are perfectly correlated, that one is completely caused by the other. There is no separate soul or lifeforce to stick a finger into the brain now and then and make neural cells do what they would not otherwise. […] One cannot logically be a determinist in physics and chemistry and biology, and a mystic in psychology. […] If one is to be consistent, there is no room here for a mysterious agent that is defined as not physical and yet has physical effects […] "Mind" can only be regarded, for scientific purposes, as the activity of the brain, and this should be mystery enough for anyone".

Two things need to be clarified in this context. First, the notion of a deterministic brain perfectly agrees with the observation that precisely predicting neural activity is hardly possible, for the same reason as chaos is not based on random processes. While the Heisenberg uncertainty principle applies to particles, uncertainty in an antideterministic sense may not apply to the organizational levels that constitute brain function[23,36] (this statement will later be refined). Second, back in 1949, psychology was still largely dominated by behaviorism, which is supposedly why Hebb was using the term "behavior" in contrast to cognition in a broader sense. Strictly following a stimulus-response paradigm, the behavioristic approach confined itself to the study of observable and quantifiable aspects of behavior and deliberately ignored subjective mental phenomena because of a suspicious "smell of animism" associated

with them. Striving to avoid assumptions about any kind of interaction between mind and behavior, it lies at the very heart of behaviorism to mistrust the many degrees of freedom that internal variables would inevitably add to the stimulus-response regime. Hebb, however, explicitly wanted to include "central processes" in his neuropsychological theory and to overcome the need for assuming a mysterious interdependence between an immaterial mind and neural signal transmission by directly relating volition, motivation, emotion and the like to neural activity[55]. The grand challenge that Hebb was facing and that we are still concerned with today would hence be to develop a comprehensive cognitive theory that is physiologically intelligible, or equivalently, to arrive at an interpretation of neural activity that consistently explains various aspects of cognition. This thesis intends to advance this undertaking by trying to elucidate organizational principles of neuronal activity in the brain and to integrate them into a suitable neurocognitive framework that is open to all sorts of cognitive processes, without concentrating on a single one.

Admittedly, "understanding the brain entails knowing about thousands of brain structures, billions of constituent neurons, exquisitely complex patterns of connectivity, and sophisticated computations mediated by synaptic inputs and spike trains that in turn rely on intricate molecular signaling cascades", as Van Essen framed it[507]. How can one begin to conceptualize an information processing system of this complexity? The general idea of brain functioning as we see it today supposes that discrete cognitive contents correspond to circumscribed spatiotemporal activity patterns in the brain, nested into each other on several spatial and temporal scales. These transient activity patterns are embedded in the functional architecture of the neural network and are shaped by external and internal factors alike. Because neurons are plastic and change the properties of their cellular components in response to activity[124], synaptic strengths and neural connections are continually being modified, revealing a conceptual interchangeability of structure and function. Just as we cannot step into the same river twice, the dynamic process of the brain's functional organization is unidirectional and constantly maps new information onto the existing structure. The resulting neuronal connectivity resembles a distributed meshwork that is characterized by massive feedback and interactivity.

Inasmuch as the brain's architecture is inherently flexible, reflexive and adaptable, it is capable to self-organize and gives rise to highly nonlinear dynamics that cannot be understood from individual cells alone [456]. Their operations are collaborative and creative and readily lead to the emergence of unforeseeable new patterns of activity. By rearranging subpatterns and their elements on any spatiotemporal scale, a virtually infinite number of items and relations can be represented[454,455]. If we now ask for the specific configuration of these representations and the rules that govern their emergence and transformation, we have begun to search for the neuronal code.

Chapter I

Principles of neural coding: classical concepts and current debates

Theories of neural coding

Following Christen[82], a code is generally defined by a relation between two sets of symbols and consists of the so-called code input set, the codeword set and the code relation. The code input set is the whole of all symbols that are accepted as an input to the code relation function, which expresses the rules that are used to convert any input symbol to its corresponding image. The collectivity of all possible symbols that may be generated by the code relation function constitute the codeword set. Applied to the brain, we therefore need to define both the attributes of neural activity that are supposed to carry distinct elements of information (the symbols) and the neuronal mechanisms mediating their formation and conversion (the code relation) to achieve a full description of a neuronal code.

Basic units of neural information processing

In principle, any such definition may involve any spatial or temporal scale, depending on the kind of information that is to be processed. Since the general acceptance of the neuron doctrine in the 1930s, however, the assumption that the neuron is the basic unit of information processing was rarely challenged (for an overview of the historical roots of information theoretic concepts and their implementation into the theory of brain functioning see Christen 2006). Neuronal computation is based on the integration of synaptic inputs[308] to produce sub- or suprathreshold electrical potentials that are transmitted actively or passively to the neuron's postsynaptic targets. In doing so, neuronal cells not simply convey information from one synaptic stage to the next, they rather interpret the input signals by synaptic filtering and nonlinear summation, adding a new value to the incoming information. This process of selecting, integrating, modulating and broadcasting electrical signals is mediated synergistically by dynamic pre- and postsynaptic structures[1], fluctuating intrinsic membrane properties[224,225,312,423,502], and cellular compartmentalization[81,205,206,293,298,472,483]. Although all these elements can be independently regulated through neuronal plasticity[1,298,312,493] and might therefore be seen as separate computational entities, it is only their combined action that produces the output signal of the cell. If the neuron is equipped with a sufficient number of voltage gated ion channels and if the summation of postsynaptic currents depolarizes the cell membrane above spike threshold, this output signal takes the form of an action potential, representing the digital result of an analog computation. While basic arithmetic functions and nonlinear transformations like normalization and gain control can be applied to analog signals directly and economically[444], the all-or-none action potential is more resistant to noise from stochastic ion channels[277,399] and can be transmitted over longer distances at high speed. Unsurprisingly, neuronal communication is mediated extensively by action potentials (although also analog signals can travel along axons and modulate transmitter release at the presynaptic terminal[13,439] or directly enter another cell through electrical synapses[85]). Thus, individual nerve cells are regarded as the brain's basic computational units, and the essential step in neuronal computation is the constantly updated "decision" of a single cell to fire an action potential or to not fire[264] (exceptions exist e.g. in the retina where photoreceptors, horizontal cells and bipolar cells, and to a lesser extent also amacrine cells, transmit information via graded potentials).

Independent and coordinated coding and the level of abstraction

On these grounds, various neural coding theories have been developed. One issue concerns the independence of single cells and can be expressed by asking if all information is available by simply pooling together the activity of a number of basically independent neurons, "as in an election", or if the relevant signals are assembled by a directed coordination of neuronal activity, "as in a symphony"[92]. Independent coding implies that all of the information that can be obtained from one neuron can be obtained from that one neuron alone, without reference to the activities of others. In contrast, the hypothesis of coordinated coding assumes that information processing involves some concerted action among neurons that may only be decoded by relating the activity of single cells to the activity of their peers. Another issue concerns the level of abstraction in

representing irreducible cognitive contents: Single neuron coding refers to the concept of complex feature detectors in sensory areas that emerge as a result of converging pathways[32], and likewise to the concept of specialized command neurons in motor areas that were thought to provide the temporal pattern of impulses needed for a coordinated activation of muscle fibers to form a behavior[270]. Alternatively, population coding refers to a distributed representation of information[385].

Coding by "cell assemblies"

Delage was probably the first who anticipated ensembles of coactive neurons to be the physiological equivalent of what he called "a single idea"[209]. Driven by direct interactions, the members of the ensemble would leave on the physical connections among them a trace, a "relic", that would facilitate their future cooperation. Some thirty years later, and with no reference to the work of Delage which apparently had been forgotten by that time, Hebb elaborately formulated what became known as the "cell assembly hypothesis"[207]. He conjectured that through "some growth process or metabolic change", repeated coactivation of a group of neurons causes the formation of a "cell assembly" – an anatomically dispersed set of neurons among which excitatory connections have been potentiated. As a consequence, repeating activation patterns in a way translate into assembly formation, and are henceforth represented by the activity of the assembly. Given that repeating excitation patterns most likely carry some meaning, each cell assembly is proposed to be a correlate of some discrete, cognitively meaningful item of information. Hebb's concept has been reviewed many times and refined ever since[49,155,169,198,409,457]. In particular, the strict connectivity-based definition has been relaxed in favor of a purely temporal one[155,169,457]: From a downstream point of view, there is no need for the neurons in the assembly to be directly connected – all that matters is their synchronous activity within a critical time window. Common to most cell assembly models is the assumption of an interdependent coordination of cells. An exception is the work by Shaw and colleagues who proposed that the relevant information be distributed across "approximately 30" independently active neurons[432].

The progression of multineuronal activity patterns in time

The firing of a cell assembly may be initiated either by external events through the sensory periphery or by internal processes that are represented by the activities of other assemblies. Hebb supposed that as the excitation of an assembly fades, it triggers the subsequent activation of a new assembly, resulting in a chain of interconnected assemblies termed "phase sequence"[207]. The progression of assemblies in the phase sequence represents successive steps in a serial computation and is the hypothesized substrate of internal cognitive processes; however, the fundamental currency of information processing is the firing of a single assembly, not the sequence[198]. Some decades later, Abeles presented a similar but much more stringent concept: He proposed that groups of synchronously active neurons, each emitting a single spike, follow each other in a precisely timed manner, forming a well-defined structure called "synfire chain" in which synchronous activity in the sending node induces synchronous activity in the receiving node through connections with identical delays[3]. The synfire chain concept was taken up and extended by Bienenstock who referred to diffuse asynchronous firing resulting from differing transmission delays as a "synfire braid"[47], an idea that was worked out further by Izhikevich. He coined the term "polychronization" to indicate temporally dispersed but precisely coordinated firing[232]. In terms of a dynamical system, the progression of neuronal activation patterns reflects the system's movement in an extremely high-dimensional state space[454-456]. The evolving trajectory is thought to transiently visit metastable states without ever being trapped in a fixed point or limit cycle [11,21,104,392,393]. The information that is processed would thus be represented in a self-organized manner by a sequence of transient states that depends on the system's history, rather than by eventually reached attractors or steady states[110]. The concept of "neuronal avalanches" utilizes the notion of self-organized criticality to explain the fractal appearance of propagating synchronous activity[40,382].

Time as a coding dimension

A spirited debate concerns the temporal resolution at which information is represented by individual action potentials. The rate coding hypothesis holds that information is mainly conveyed via the instantaneous firing rate[429,430] – the mere number of spikes in some time window – whereas the temporal coding hypothesis assumes that the precise placement of the spikes in time is also significant[128,129,175,283,449,459,488]. On the single cell level, slowly modulated or constant firing rates, but also the duration of a burst of impulses, the number of spikes within a burst, the rhythm of firing, irregular temporal patterns or simply single spikes may be used for information transmission and selective communication between neurons[110,233]. On the population level, profiles of instantaneous firing rates, the activation order of cells within some time window[162,509], or temporally precise multineuronal spike patterns[8,29,232,447,504] could function as coding entities.

The role of neural oscillations

Synchronously discharging neurons often produce oscillatory rhythms of various frequencies, generated by networks of diverse sizes[65,68,267]. Theoretically, synchronous oscillations might simply be an unavoidable byproduct of neuronal network dynamics without any particular computational role. Alternatively, they could directly contribute to the representation of information, for example by providing the timing for an internal clock[231] or as a reference signal relative to which spike times become meaningful, or they could actively regulate the flow of information in neural circuits by inter-

fering with the action potential generation and temporally link neurons into assemblies[68,142,425]. Encoding by phase and synchrony has highly attractive computational properties[217,218,348]. It has been proposed that phase encoding might effect the temporal segmentation of several working memory items[238,291], and that waves of activity might serve to tag sensory input at different spatial locations with a unique phase[120]. The addition of phase information may be used as a means to segment and categorize parallel inputs. In a similar way, top-down processes could shape spiking activity by coordinating subthreshold membrane potential fluctuations to establish selective functional relationships between neurons during states of anticipation[114]. The idea that the formation of dynamic links mediated by synchrony over multiple frequency bands subserves neuronal communication[25,51,448,511] was dubbed "communication through coherence" by Fries[140]. Rhythmic excitability fluctuations are thought to confine neural signal transmission such that only coherently oscillating neuronal groups can interact effectively, in the sense that their excitability peaks need to coincide to facilitate the propagation of spikes. The resulting effective communication structure may flexibly be rearranged through shifts in attention or other cognitive processes that come along with alterations in the oscillation patterns, which in turn would alter the selective linking of distributed representations[427].

Coherent oscillations could provide a mechanism to solve the so-called "binding problem"[497]: If we assume that some irreducible percept or thought or motor plan is represented by a group of neurons on a dynamical basis, what is the signature that transiently binds their activity into a unified whole? Milner proposed that cells selectively segregate their firing in time to signal their functional relationships[341], and von der Malsburg formulated the "correlation theory of brain function" based on the same rationale[522-524]. Singer and coworkers adopted these concepts[117,180,458] and advanced the "binding by synchrony" hypothesis that suggests that functional relations between neurons are encoded by synchronous firing in the millisecond range, brought about by the phase-locking of distributed oscillations[116,445,447,452]. The idea behind is that elementary relations are represented by the firing of individual neurons mediated through appropriate convergence of input connections, and that more complex relations are represented by the activity of cell assemblies generated by dynamic associations of cells[111,404,457].

Key questions

These concepts lead us back to the initial question if irreducible cognitive contents are represented mainly by single neurons or by neuronal populations. If a number of neurons is involved, do they coordinate their firing, or are they independent? If they act in concert, how is their activity organized? Does the timing of spikes reflect a rate code or a temporal code? What is the temporal precision of neuronal firing? Finally, do neural oscillations contribute to the representation and transformation of information, and if so, how?

Arguments and evidence

Single neuron coding vs. population coding

Throughout the brain, neighboring neurons often share similar information because they share similar inputs. In principle, the resulting redundancy is a useful mechanism to protect against the loss of information. However, given the high metabolic demands of neuronal operations, such redundancy comes at some cost. Moreover, diverging excitatory pathways may recruit large populations of neurons, so that a single message may engage a considerable part of the network. These problems could be mitigated by sparse coding through fast convergence of signals to neuronal detectors of highly specific complex contents[213]. In fact do some neurons in the human medial temporal lobe selectively respond to visually presented persons or objects irrespective of their size or position in the visual field or the viewing angle, and in some cases even to letter strings with their names[390]. Although it is unclear if these cells are driven exclusively by the tested stimuli and if other cells respond to the same stimuli as well, these findings exemplify a remarkable invariant and abstract representation of visual contents. It has to be questioned, however, if a scheme that relies on single neuron coding qualifies as a universal method for representing information in the brain, due to a number of conceptual shortcomings[447]. First of all, such a scheme implies that a selective neuron is available *a priori* for every *possible* percept or mental object, which is simply impractical[63]. Furthermore, it becomes increasingly difficult to encode compositionality and syntactic relations and to establish semantic associations the more the information gets concentrated. Finally, highly abstract representations entail the "bottleneck problem": After convergence, how could the encoded information be decomposed by downstream neurons? Following these arguments, it is obvious that the nervous system needs to maintain some form of population coding.

Signatures of coordinated coding

Through the distribution of information over many neurons, each receiving a redundant but unique combination of inputs, the resolution in representing that information is enhanced. It can be expected that the brain evolved to optimally balance metabolic demands and computational capacity and flexibility. In visual cortical areas, adjacent neurons have been found to carry an average of between 40 % and virtually no redundant information[105,163,164,518,551], depending on the stimulus. How independent are the messages of single neurons in other brain areas? From simultaneous responses of retinal ganglion cells, stimuli can be reconstructed with high accuracy, even if correlations between cells are left unconsidered[127,356]. Nevertheless do retinal ganglion cells synchronize their firing beyond what can be expected from shared visual input[333,334], and significantly more information can be extracted

from their activity if correlations are taken into account[380]. The same is true in other parts of the brain [87,112,199,204,255], demonstrating that neurons are involved in directed interactions that could provide additional information. From this perspective, the signaling of single cells might be imperfect or even completely meaningless unless related to the activity of others.

Further support for a correlational code is given by the widespread observation that multineuronal spiking activity is actively coordinated: Aside from the retina, synchronized firing in excess of what would be predicted from the discharge rates has been recorded from the reticular formation[289] and accompanies responses to fixation onset in primary visual cortex[310] and preparatory processes in primary motor cortex[179,400,401]. Prolonged patterns of coordinated firing have been found in forebrain areas related to cardiac and respiratory control[152] and correlate with auditory[487,515,516] and olfactory[56,139,278] processing, up state onset[301], and behavior [4,389,487,505,516]. In the hippocampus, sequences of place cell activity are replayed forward or in reverse order during brief pauses in waking behavior[90,101,135,251] and during sleep[299,347,349,366], possibly reflecting processes of memory consolidation and retrieval. Overall, these coordinated firing patterns were precise to within three to thirty milliseconds while spanning up to a few hundred milliseconds and have been taken as an indication of functional cell assemblies.

Although simulations have shown that synchronized action potentials can reliably propagate within a cortical-like network[102], neocortical spike recordings have never been thoroughly investigated with respect to higher order temporal structures like synfire chains or phase sequences. Abeles and colleagues tried to infer synfire chain activity from structured firing in the frontal cortex of monkeys, with limited success. They concluded that the patterns were generated by "reverberations in a synfire mode" within self-exciting cell assemblies[4,389]. Ribeiro and co-workers demonstrated the dependence of cell assembly activation sequences recorded simultaneously from the hippocampus and primary visual and somatosensory cortical areas of rats on the behavioral state of the animal[15]. The most direct evidence for functional sequences of discrete firing patterns found so far is the characteristic succession of transiently synchronized neuron ensembles during an odor response in the antennal lobe of the locust[531].

Timescales and accuracy of neuronal signaling

Energy supply critically limits signaling in the brain[250]. For cerebral cortex, the volume of signal traffic that can be supported by the brain's metabolic rate was calculated to be about five action potentials per neuron and second in rat and less than one per neuron and second in human[277,282]. Considering the speed of neural computations, the permissible signaling rate is remarkably low, and this metabolic limit must affect the way in which information is processed. Recordings from sensory cortices suggest that the nervous system has countered this natural constraint by distributing signals sparsely in time and space[100,220,236,258,529,534]. The conclusion that at most a few discharges per neuron are available to convey a message is confirmed by the finding that sensory information is transmitted quickly along feed-forward connections[491], requiring only ten to fifteen milliseconds per processing stage[496]. Therefore, it was argued that information can only be represented by short, fast responses forming a sparse population code. In fact, reliable decoding of stimulus features is possible based on the relative timing of the first spikes elicited in individual neurons in the retina[178,510], the olfactory system[246], the somatosensory system[243,378], and even in cell cultures[431]. But how reliable is the initiation of action potentials in single neurons, and what is their temporal precision? Membrane potential fluctuations induced by stochastic ion channel gating and probabilistic release of synaptic vesicles are potential sources of random variations in spike generation and timing[123,399]. So, the probability that an arriving presynaptic nerve impulse fails to evoke a postsynaptic response is remarkably high, between 0.5 and 0.9[14,277]. However, because of the great number of synapses, failures do not necessarily lose information. Variability introduced by nondeterministic processes acting on the level of single molecules may average out on the cellular level[23] and may even sharpen the signal due to stochastic resonance[328,329]. The amplitude and exact timing of somatic potentials in response to a particular input would be expected to approach a Gaussian distribution, giving rise to precisely timed action potentials in most cases while occasionally failing to cause a spike in time. This is indeed what could be observed by repeatedly injecting irregular depolarizing currents into cortical neurons in vitro[309], and simulations suggest that the same is true for the axonal propagation of action potentials, leading to small, mostly submillisecond variations in spike timing over distances of millimeters[122]. High reliability of spiking has also been demonstrated in the visual[210,249] and in the auditory system[99] in vivo.

The temporal precision of neuronal communication crucially depends on a number of basic cellular properties. Spike-timing-dependent plasticity rules for modifications in synaptic strength indicate that postsynaptic potentials are effectively integrated within only twenty to thirty milliseconds[72,89]. Such short integration times mainly result from rapidly deactivating AMPA receptors that can have deactivation time constants of less than a millisecond[165,212,308] and indirectly control the kinetics of NMDA receptor currents by only allowing for a correspondingly short release of the magnesium block[216,367]. In addition, disynaptic feedforward inhibition may confine the effective integration time window in the soma to a few milliseconds[386]. Backpropagating action potentials coinciding at the synapse with excitatory postsynaptic potentials may trigger dendritic calcium spikes and in this way cause highly nonlinear responses[276,293,418,478,479]. Another nonlinear element is the spike threshold which is inversely related to the rise time of the action potential, endowing neurons with an enhanced sensitivity to synchronous inputs[27,197]. With increasing input rates, both the amplitude and duration of somatic potentials in response to synaptic input is reduced, resulting in a shortening of the temporal inte-

gration window and requiring a yet higher precision of presynaptic signals to drive the neuron to fire[28,281]. Finally, many synapses operate most reliably at certain frequencies of presynaptic firing and display depression or facilitation of postsynaptic responses[159,398,489,538]. Such synapses effectively detect changes in the firing rate, but report frequency of maintained activity poorly [2,26,490,500]. Besides effects of repetitive signaling on the release probability of vesicles at the presynapse, the kinetics of transmitter binding and channel gating of postsynaptic NMDA receptors produces currents with distinct waveforms depending on pulse frequency[384], leading to the long known fact that the postsynaptic response is sensitive to the exact timing of successive input signals[424]. One important consequence of this sensitivity is that modifications in synaptic strength depend not only on the relative spike timing between the neurons, but also on the spiking pattern within each neuron, with the timing of the first spike in each burst being dominant in determining the ensuing synaptic modifications[146]. Taken together, these properties make neurons susceptible for transient signals and precise spike timing codes and the transmission of a continuous rate signal difficult[266].

Nevertheless managed Shadlen and Newsome to devise a model in which the membrane potential undergoes a random walk to the spike threshold so that any temporal structure in the input is lost[430]. Based on the model, they concluded that the summation of postsynaptic potentials in cortical neurons is too imprecise to support precise spike timing codes, thus leaving as the only coding dimension the firing rates of neurons. A reliable estimate of the instantaneous firing rate would then require the simultaneous readout of a population of neurons[162], implying ergodicity and independence of cells. Assuming uncorrelated inputs, the model predicts a Gaussian distribution of membrane potential with only small membrane potential fluctuations. In vivo recordings, however, revealed highly non-Gaussian membrane potential dynamics displaying quiescent periods interrupted by large, brief excursions consistent with coordinated presynaptic firing[100,295]. As has been shown both theoretically and experimentally, organization of presynaptic input into synchronous volleys is also necessary to explain the irregular output firing of neurons[465,475]. These findings are in conflict with basic assumptions of the model and seriously question its validity. To make things even worse, correlations between cells would compromise the ensemble representation of firing rate especially at high frequencies, imposing severe constraints on the temporal accuracy of neural computations[325]. Again, neuronal dynamics and theoretical considerations are at odds with a firing rate code, but there is ample evidence for millisecond-precise spike timing depending on sensory input, behavior, or internal state in a variety of different brain areas like the frontal cortex[505], motor cortex[203,400,436,437], somatosensory cortex[134,372], auditory cortex[91,119], visual cortex[30,181,265,492], thalamus[87,95,503], retina[44,183,234], and the hippocampus[199,402]. We are thus led to a view of neural activity as being basically and essentially characterized by sparse, temporally precise, coordinated firing.

The functional relevance of network rhythms and their role in coordinating multineuronal activity

Given that the brain, like every system which has opposing forces such as excitation and inhibition, almost inevitably will generate oscillations[65,527], it is hard to believe that it did not evolve to make use of them. But how do neural oscillations relate to the processing of information? The idea that oscillations could serve as an internal clock has been dismissed in favor of a model using high-dimensional network states for encoding time[61,252,324]. The difficulty in assigning functional relevance to synchronous oscillations lies in the correlative nature of most of the investigations done so far. There are some exceptions, though. In a series of experiments, Laurent and colleagues used picrotoxin (a GABA antagonist) to disrupt synchronous oscillations in the olfactory systems of insects[307] and so were able to demonstrate that the selective desynchronization of projection neurons in the antennal lobe degrades the selectivity of downstream neurons[306] and impairs the animal's ability to discriminate molecularly similar odorants[477]. In mammals, however, the situation is less clear. Mice lacking the $GABA_A$ receptor β_3 subunit produce enhanced oscillations in the olfactory bulb and after training are better than normal in discriminating closely related odorants, but worse in discriminating odorant mixtures[364]. In the rat olfactory bulb, oscillatory power appears to be actively modulated depending on the molecular similarity of odorants that the rat has to distinguish, suggesting a role of enhanced network oscillations in stimulus disambiguation[45]. On the other hand, newborn rats who have very few GABAergic granule cells do not produce synchronous oscillations and yet are as good in making odor discriminations as older ones who have developed interneurons and do produce oscillatory activity in response to a stimulus[133]. The two most convincing studies that tried to establish a causal link between synchronous oscillations and behavioral performance in vertebrates again relied on pharmacological interference with normal neuronal functioning: In the frog retina, a subclass of oscillating ganglion cells responding to expanding dark objects gets out of sync when exposed to bicuculline ($GABA_A$ antagonist), leading to the failure of an escape behavior as it is normally induced by such stimuli[230]. In the rat hippocampus, cannabinoids cause a decrease in oscillatory power in various frequency bands without affecting average firing rates, which in turn impairs memory formation[402]. These examples show that neural network function is often associated with neuronal oscillations, but in most cases it is unclear how exactly they contribute to the processing of information. Nevertheless, on a mechanistic level, they play a consequential role in coordinating multineuronal activity.

Network oscillations are carried by rhythmic inhibitory input originating from synchronized interneuronal spiking[33,67,138,202,446,466,537]. It has been shown both experimentally[96,275,331,417,521] and in simulations[218] that the ensuing subthreshold membrane potential fluctuations interact with excitatory inputs such that the timing of

action potentials becomes a function of the oscillation phase and is made less variable. As a result, discharges are temporally coordinated and related to the network rhythm. Examples in which the phase of firing carries significant information have been reported from prefrontal[441], auditory[257] and visual cortex[344] and include the "phase precession" in hippocampal place cells [200,226,365,426,461] and entorhinal grid cells[193]. In cortical area V4 of monkeys, the frequency-dependent strength of the phase-locking of spikes is modulated by attention[143,185]. Moreover, the interaction of neuronal groups has been found to depend on the phase relation between rhythmic activities within the groups[541], consistent with the idea that firing phases and times of increased susceptibility to input need to match the intergroup transmission delays to facilitate the propagation of spikes. Similarly, the strength of the interareal phase synchronization of neuronal activity in monkey V4 and prefrontal cortex was shown to correlate with visual short-term memory performance[286], suggesting that this synchronization subserves intercortical communication and contributes to the maintenance of visual short-term memories. Thus, neural oscillations dynamically shape suprathreshold activity and flexibly arrange signaling pathways in concert with cognitive processes[69,527].

A new approach

Although much has been said and done since the days of Hebb and the introduction of the cell assembly concept, it appears that we still lack a complete, comprehensive understanding of the dynamic organization of multineuronal activity. The temporal precision of firing and the timescales on which neuronal activity is coordinated are a matter of ongoing debate[24,92,198,492]. Without a clear characterization of the spatiotemporal structure of concerted neuronal firing on short timescales – that is, the definition of a differentiated signature of neural assemblies – also no superordinate structure possibly representing cognitive processes on longer timescales can be found.

What is hence needed is an approach to assess and precisely characterize higher order correlations among multiple neurons on a moment-by-moment basis. Due to the absence of suitable analysis methods[24,54], investigations have been restricted to pairwise correlations or to some special case of functional organization or did not include at all a test for the statistical significance of the observed activity patterns. The second chapter is therefore devoted to the development of a method for the detection of multineuronal discharge sequences in parallel recordings that provides a precise description of their spatiotemporal organization and allows for a continuous correlation of the activity patterns with the ongoing information processing. It is intended to answer the question if nerve cells fire independently or depending on each other, if repeating spatiotemporal patterns show significant structure, what the relevant timescales are, and if short firing patterns are arranged in coherent sequences.

For understanding the neural code, as important as the signature of neural assemblies – the potential information-carrying symbols – are the mechanisms that mediate their formation and conversion. It has been argued that network oscillations tend to synchronize action potentials in coherently oscillating cells and so create a signature of functional relatedness[447,452]. They naturally arise from the interplay of recurrent excitatory and inhibitory connections and the resonant properties of individual neurons[74,182,225,292]. Synchronization of signals is supported locally by the coupling of cells via gap junctions[42]. Remote populations may engage in zero phase lag oscillations despite long conduction delays if coupled reciprocally to a relay population of cells[79,514,519], and it has been suggested that thalamic nuclei may play an according role in mediating synchrony among distant brain regions[244,435,514]. Another mechanism by which neuronal activity is organized is the shaping of the functional network through synaptic plasticity. Theoretical studies have demonstrated that neurons equipped with spike-timing-dependent plasticity[72] may tune to repeating spatiotemporal input patterns by potentiating synaptic weights on afferents that consistently fire early, thereby steadily decreasing postsynaptic response latency with respect to the onset of the pattern, until it reaches a minimal value[192,321,322]. Given appropriate input firing patterns and plasticity mechanisms, multineuronal spike sequences should therefore progressively be compressed in time and eventually become synchronized[482], if the participating neurons respond to coherent input. How can the influences of these different mechanisms on the generation of precisely timed discharge sequences be disentangled and quantified?

As Buzsáki accurately pointed out, "the acid test for providing a definite proof for the essential role of brain rhythms in computation and brain function would be to selectively eliminate them and examine what is left after the complete lack of oscillatory timing"[65]. However, oscillations are an emergent network property and do not have "receptors" that can be targeted by drugs or other means; only individual neurons do. It is therefore impossible to selectively eliminate a rhythm without fundamentally interfering with the elementary properties of the parts that gave rise to it. Modifying the function of certain receptors for neurotransmitters is likely to radically change the flow of electrical signals in the network and to also affect all other activity patterns[65]. This criticism applies to the aforementioned experiments that used picrotoxin and $GABA_A$ receptor β_3 subunit knock-out mice to disrupt or alter network oscillations, and it also applies to more direct manipulations of the activity of subpopulations of neurons by optogenetic methods[466,552]. An alternative way to study the organization of neural activity in the absence of neural rhythms could be to record from brain slices: By disconnecting some part of the network from the rest of the brain, chances are high that the remaining network is too small to generate synchronous oscillations[480,543]. On the downside, neurons encounter a lack of neuromodulators, but these can in principle be applied externally; the important difference of an in vitro approach

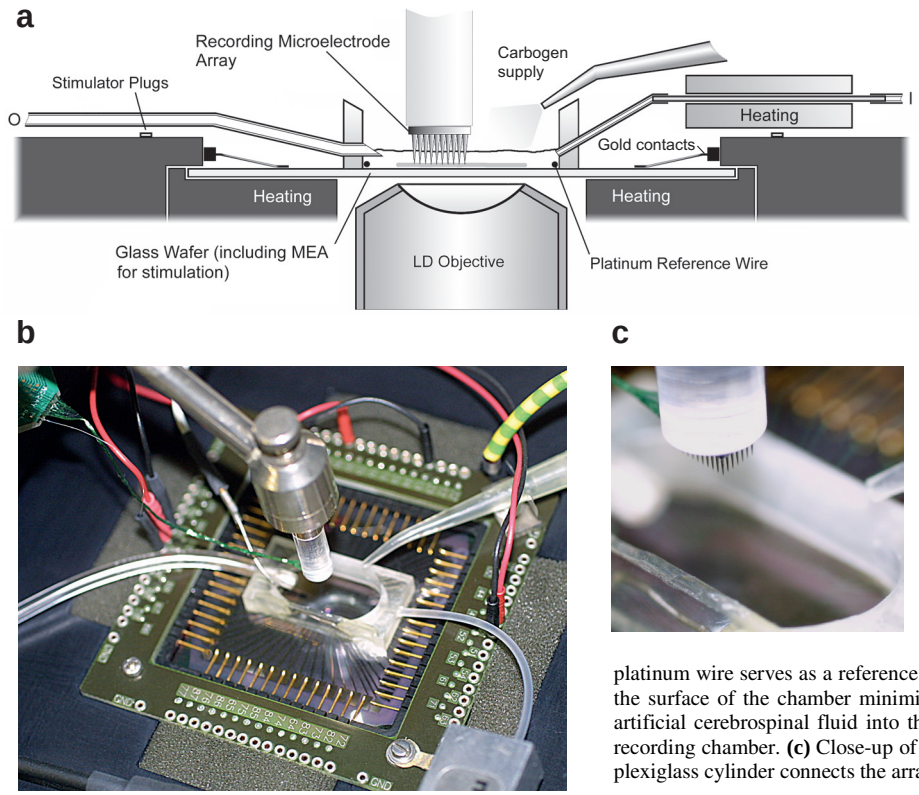

Fig. I-1 Experimental setup (a) Sketch of the recording chamber. A custom-made plexiglass chamber is glued to a glass wafer that holds an integrated multielectrode array with 59 flat electrodes (30 µm diameter) used for electrical stimulation. Gold contacts connect the electrodes to a stimulus generator. The wafer rests on a heating element that precisely controls the temperature of the chamber. Optical stimulation and control can be exerted from below via an inverted microscope. Continuous perfusion with artificial cerebrospinal fluid is provided by a peristaltic pump through the inlet (I) and outlet (O) tubes. A second heating element accurately controls the temperature of the fluid. The recording electrode array consists of 63 sharpened silicon electrodes (spaced at 400 µm) and is inserted into the tissue from the top. In this way, the brain slice is held in place, and tight contact with the stimulus electrodes is established. A circular platinum wire serves as a reference, and a constant supply of carbogen to the surface of the chamber minimizes the diffusion of oxygen from the artificial cerebrospinal fluid into the ambient air. **(b)** Photograph of the recording chamber. **(c)** Close-up of the recording microelectrode array. A plexiglass cylinder connects the array to a micromanipulator.

compared to the elimination of oscillatory activity in vivo is that receptor function and neuronal excitability can be left untouched and unaffected. Following these arguments, the third chapter centers on the dynamics of neural activity in non-oscillating brain slices. Since most higher cognitive functions are attributed to the neocortex, it is the chosen structure of interest. Its anatomical architecture is largely preserved across brain areas and mammalian species[508], so that one can hope to gain some general insights into cortical neuronal information processing and coding by investigating a specific area in a certain animal. For no particular reason except its good accessibility and the extensive knowledge that we have of its circuitry, this study examines the organization of neuronal activity in the primary visual cortex of the rat.

The aim to observe coordinated discharges at sub-millisecond time resolution makes it obligatory to simultaneously record from multiple single neurons with multiple electrodes[64,339,450]. Until now, multi-site recordings of single-unit spike activity in acute brain slices have been reported on only a few occasions and did not follow any standardized approach. Problems arise in particular when using flat electrodes because spikes can be recorded only from the surface of the slice where most cells are damaged as a result of the slicing procedure, and because spike recording requires auxiliary techniques to assure proper contact of the tissue with the electrodes[109]. To resolve these problems, a novel experimental setup has been designed and employed in this study (Fig. I-1a, b) that enables the recording of spikes from cells in the middle of the slice using a matrix of 1.5 millimeters long, sharpened electrodes (Fig. I-1c). The setup allows the observation of a large, random set of neurons of which a subset might participate in a given neuronal assembly[170]. Subsequent analysis, then, allows interference of assembly properties. The basic idea is not to search for any predefined spatiotemporal structure in distributed discharges, but to systematically test for a coordination of spike timing on several timescales and to characterize the organization of local cortical spiking activity in a comprehensive way. The question that is raised is the following: Does the local cortical circuitry give rise to stereotypical activation patterns, and if so, what does their spatiotemporal organization imply with respect to cortical information processing and coding?

Chapter II

Detecting multineuronal temporal patterns in parallel spike trains

Introduction

One of the most fundamental issues in neuroscience is the nature of the neural representation of information. While it is widely appreciated that informational contents are carried by the activities of a large number of neurons, there is dissent about the independence of cells and the relevant timescales of their firing[92]. Do neurons jointly encode information by forming functional cell assemblies[155,169,207,409,522]? Does precise spike timing significantly contribute to neuronal communication[129,175,449,488]? As a consequence, are neuronal assemblies distinguished by a covariation of firing rates[198] or by a time-locked sequence of polychronous[232] or synchronous[457] spiking events? Finally, are the activities of neuronal assemblies, whatever their particular structure may be, arranged sequentially and coherently in time to form superordinate patterns[3,47,207]?

Although studying these questions is a statistical and computational challenge[54,254,374], a variety of methods have successfully been applied to reveal clusters of functionally related cells without characterizing their temporal structure[43,168,171,259,297], to define groups of cells firing in synchrony[172,187-189,253,379,381,413,421,495], to detect spatiotemporal firing patterns[5,151,280,353,415,420,462,463,485,486,546], and to find signatures of synfire chain activity[173,422]. Barring some difficulties in finding appropriate representations of the associated null hypotheses[167,186,190], these different methods and their applications are able to analyze relevant properties of multineuronal activity, but an all-embracing approach is missing.

To achieve a comprehensive and conceptually unrestricted description of multineuronal spiking, I present a new method for analyzing consistent relations between discharges of simultaneously recorded neurons on arbitrary timescales that are referred to as spatiotemporal firing patterns and pattern sequences. Adopting a maximally naïve view on multineuronal suprathreshold activity, repeating spatiotemporal firing patterns are registered with user-defined precision by sliding a temporal window of interest along the parallel spike trains. In addition, series of patterns are scanned for repeating sequences. The significance of repeating firing patterns is estimated individually and globally by comparing the numbers of their occurrences with the numbers that would be expected if the cells' firing were independent on the given timescales. For that purpose, a new type of surrogate data is introduced that allows for variability and sparseness of spiking events and is superior to common resampling methods in terms of statistical test performance. Another difficulty when searching for recurring spatiotemporal patterns in massively parallel recordings arises from the mutual masking of actually unrelated patterns that are arranged in the same window. To avoid the combinatorial explosion that results from testing every single possible subpattern, an algorithm is proposed that separates coincident events based on the preferences with which a neuron joins its various peers in coincident firing[*].

Methods

The following subsections first describe the algorithms for the detection of spatiotemporal firing patterns and pattern sequences, including the proposed procedure of separating subpatterns. Then, a Monte Carlo-based approach to determine the statistical significance of the found patterns and sequences is presented, together with some common and a new resampling technique and the corresponding hypothesis tests. Finally, I briefly comment on the technical implementation of the method.

Detection of spatiotemporal firing patterns

What constitutes a multineuronal spike pattern? As long as we do not explicitly know the relevant timescales of the data under investigation, we should not restrict the analysis to any special scale. The method presented here is therefore designed to provide full flexibility with regard to the temporal organization of the data: In a straightforward approach, I focus on the activation sequence of cells as the essential signature of a pattern[431,509] and define patterns by registering the first spikes of all units within a certain time window W with a certain precision τ (Fig. II-1). Both timescales – the maximal length of the pattern and the spike timing precision – can be arbitrarily chosen and jointly determine which aspects of the data are investigated. By

[*]The terms "spike" and "event" are used here interchangeably, emphasizing either their biological or statistical meaning, respectively.

applying several parameter combinations successively, the data can be scanned for a range of very diverse spatiotemporal patterns.

Given any W and any τ (with τ being an integer fraction of W), patterns are captured by systematically sliding the onset of the time window W from spike to spike along the parallel traces. They are represented by a vector indicating the constituent units ranked by appearance (spikes co-occurring at the same sampling point are ranked by their unit number), optionally followed by the corresponding timing information. Thus, two modes for representing a pattern can be used: a time-resolved mode (Fig. II-1a) and a representation that is simply given by the temporal order of the participating units (Fig. II-1c, d). In the time-resolved version, the scale of the registered spike timing is set by dividing the window into equal bins of length τ, using their respective indices to specify each spike's position in time (Fig. II-1a). Since the central purpose of the analysis is to detect coordinated firing activity among a population of cells, spike patterns comprising only one unit are skipped (Fig. II-1b).

It is important to note that patterns do not necessarily cover the whole space allowed by a given combination of parameters, especially if restrictions are minimized by analyzing the data using wide limits (long W and $\tau = W$). A subsequent analysis of the found patterns may then reveal some characteristic spatiotemporal structure covering only a part of the search space, making it a particularly strong finding if some structure is found that has not been explicitly searched for. Once the data are known to contain repeating patterns on a specific temporal scale, the search space can be adapted to yield a better statistical accuracy.

Peer validation and pattern separation

Because every timestamp marks the start of a new window, allowing each event to participate in multiple patterns on multiple positions, the search is exhaustive in the temporal domain. In order to thoroughly scan the data in the spatial domain, one would need to break down every pattern into all possible subpatterns and to independently assess their individual significance. With an increasing number of units and events, however, this would result in a combinatorial explosion requiring prohibitively large amounts of computer memory (as an example, 20 events can be combined in more than a million ways, forming exactly 2^{20} subsets minus the empty set and the 20 singleton sets). On the other hand, unraveling simultaneously occurring but independent subpatterns is essential to reveal any repetitive structure in larger data sets. As a practical solution, I therefore propose to determine the probability with which a neuron joins its various peers in coincident firing and to split the events that coincide in any given time window W accordingly. To do so, the empirical count of coincidences of any two units during some period t of length T (with $T \gg W$) is compared against a threshold given by the expected count of coincidences and a global support value to classify them as being functionally coupled or uncoupled, thus providing every unit with a set of "validated peers". Because the functional coupling may vary over time[10], it is necessary to choose T appropriately (e.g., one minute) and to currently adjust the correlation values by dividing the data into successive intervals of corresponding length. Formally, raw correlations are expressed as $C_{ij}^{(t)}$ which is the number of coincidences of units i and j as revealed by the pattern search in time interval t. The chance level of spurious coincidences is roughly estimated as

$$P_{ij}^{(t)} = \frac{W}{T} n_i^{(t)} n_j^{(t)} \tag{1}$$

with $P_{ij}^{(t)}$ being the expected number of coincidences of units i and j in time interval t, and $n_i^{(t)}$ and $n_j^{(t)}$ being the numbers of events of units i and j in time interval t (see Appendix 1 for a derivation and necessary conditions). In case of low rates the resulting values may be too low to function as a threshold. To assure that more than one coincidence per unit pair is required to label peers as valid, an additional minimum support value may be applied. Hence, peers are validated according to

$$V_{ij}^{(t)} = \begin{cases} \text{valid, if } C_{ij}^{(t)} > \max(P_{ij}^{(t)}, A) \\ \text{invalid, otherwise} \end{cases} \tag{2}$$

with $V_{ij}^{(t)}$ characterizing units i and j as being functionally coupled or uncoupled during time interval t and A being an arbitrary global threshold referred to as absolute peer criterion that simply denotes the number of coincidences in any time interval t required to validate the functional coupling of any pair of units, irrespective of the event rates. The resulting sets of validated peers indicate which units preferentially take part in concerted firing patterns. To separate coincident events accordingly, all peers that are invalid with respect to a chosen unit are removed from a pattern. The procedure is repeated for every unit that participates in the parent pattern, potentially producing several distinct subpatterns. Finally, non-repeating patterns are dropped. After all repeating patterns thus detected have been registered, they are subjected to a search for some superordinate patterning.

Detection of sequences of patterns

It has repeatedly been hypothesized that neuronal spiking activity be organized into superordinate patterns comprising coherent sequences of circumscribed spatiotemporal firing patterns that signify functional cell assemblies[3,47,207]. As was pointed out by Schrader and colleagues[422], detecting those sequences means collating the previously identified patterns appropriately and variously and searching for new emerging structures – a task that has not been tried yet. Here we present such a method for the detection of repeating pattern sequences that is completely independent of the particular temporal organization of the constituent patterns and makes no a priori assumptions about the spatiotemporal structure of the resulting sequences.

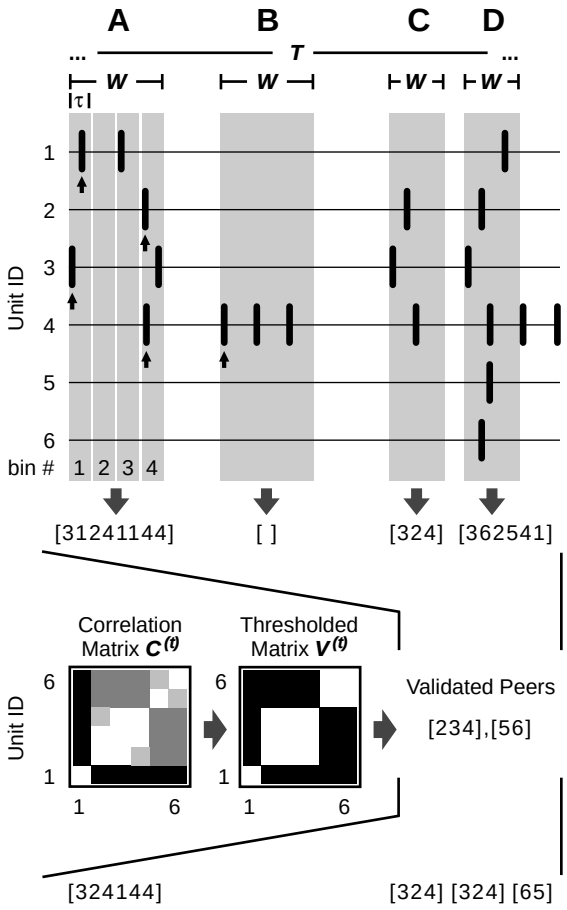

Fig. II-1 Detection of spatiotemporal firing patterns. Illustrated are six simultaneously recorded spike trains and four separately detected patterns (**a**-**d**) as examples. An arbitrary time window *W* (highlighted in gray) is used in each case to define the spatiotemporal activity pattern. The units that coincide in the given time window are further split into subgroups according to previously validated peers: Based on the number of coincident events of any two units during some period *t* of length *T* (raw correlation matrix $C^{(t)}$) and a threshold, units are classified as being functionally coupled or uncoupled (thresholded correlation matrix $V^{(t)}$, see text for details). In this example, units 2, 3 and 4 are correlated, as are units 5 and 6, and unit 1 is not correlated with any other unit. (**a**) A 60 ms window containing six spikes falling into different 15 ms bins given by τ. The resulting pattern is represented by a vector indicating the constituent units ranked by appearance ([3124...]), followed by the corresponding bin numbers of their first spikes ([...1144], see arrows). After comparison with the sets of validated peers, unit 1 is excluded from the pattern. (**b**) Since a pattern consists of at least two spikes and only the first spike of each unit inside the window is considered (see arrow), patterns comprising only one unit are skipped (W = 60 ms). (**c**) and (**d**) If no binning is applied, the vector representing the pattern indicates only the temporal order of the participating units (W = 35 ms). (**c**) After comparison with the sets of validated peers, the pattern is left unmodified. (**d**) After comparison with the sets of validated peers, the pattern is split into two subpatterns and unit 1 is excluded.

In a first step, the vector representation of every repeating pattern is replaced by a hash value indicating the pattern's identity, which helps a lot to alleviate computer memory consumption. Because the significance of a single pattern is statistically distinct from the significance of a sequence of patterns, all repeating patterns are included. Along with the pattern ID, the timestamps of the first and last event are recorded so that sequences can be clearly identified and represented by a vector of successive IDs. However, since patterns are captured with a sliding window and potentially are subdivided as a result of the peer validation procedure, they may overlap in time. To register series of temporally non-overlapping, directly consecutive patterns it is therefore necessary to look for the *very next initiation* of a pattern *after the last event* of the preceding pattern (Fig. II-2a). The resulting sequences may comprise an arbitrary number of patterns and include all corresponding subsequences (Fig. II-2b). As the process is repeatedly started *at every pattern*, the detection of sequences is exhaustive up to the analyzed length. In a last step, shorter sequences that are *always* part of *the same* longer sequence as well as non-repeating sequences are discarded. Importantly, this method does not imply any constraints concerning the exact timing of consecutive patterns (provided that they are temporally separated) or the overall duration of the whole sequence – solely the succession of pattern IDs identifies a sequence.

Statistical significance estimation

Statistical hypothesis

Following the detection of recurring firing patterns and pattern sequences, one may characterize their spatiotemporal properties and relate them to the experimental conditions. However, their mere recurrence does not imply that they occur more often than expected by chance, and both patterns and sequences have to be considered irrelevant unless an appropriate statistical test demonstrates that they recur significantly often. To do so, a non-parametric approach is proposed that can be expressed in the following way: The null hypothesis (H_0) states that the registered patterns and sequences appear by chance, or in other words, that patterns occur independently and coordination of events is random on the timescales that were used to identify a pattern. If this is the case, then varying the timing of events on that scale or rearranging the order of patterns should not affect any statistic extracted from the parallel event trains or the pattern series. To test the probability that this null hypothesis holds, the distribution of pattern and sequence counts is calculated from surrogate data with randomized event timing and randomized pattern sequences, respectively, using a Monte Carlo method. If the value obtained from the original data exceeds the surrogate count with empirical probability x (and falls below that count less often), then the probability that the data are consistent with the null hypothesis is 1 – x. The alternative hypothesis (H_A) states the opposite and assumes that patterns show some systematic interdependence and events some degree of coordination on the corresponding timescales.

Generation of surrogate data

The problem of developing a non-parametric method is to produce surrogate data that differ from the original data in exactly one property, namely the one that is addressed by the alternative hypothesis. In the past,

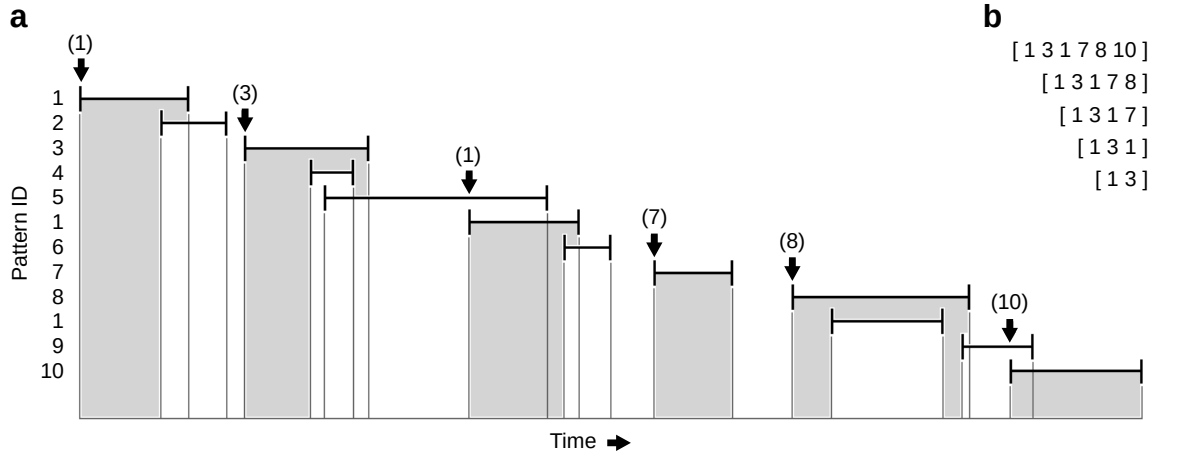

Fig. II-2 Detection of sequences of patterns (sketch). (**a**) Vertical bars indicate first and last spikes of previously identified repeating spatiotemporal firing patterns, horizontal bars indicate their duration. A unique ID is assigned to every pattern, and sequences comprising an arbitrary number of temporally nonoverlapping, directly consecutive patterns (highlighted in gray) are registered by looking for the *very next initiation* of a pattern *after the last spike* of the preceding pattern (see arrows). The process is repeatedly started *at every pattern* with all subsequences being listed. (**b**) Vector representations of the pattern sequences shown in (**a**) when starting at the first pattern.

several procedures have been proposed to create suitable surrogate data for testing the significance of coordinated spike events by repeatedly modifying the original spike trains (for a review see Grün 2009). One possibility is to dither the time of every individual event randomly and independently on a certain scale, thereby destroying the temporal structure contained across as well as within event trains up to that scale[203] (Fig. II-3a). Although it is not necessary to change the interval structure if the intention is to disarrange coordinated events, the approach is intuitively appealing. It has, however, some complications, as was revealed by Gerstein[167]. If the event times are dithered uniformly within some symmetric window (e.g. ± 20 ms), short intervals are added to the interval distribution and its peak is lowered. In terms of gamma distributions, such a surrogate is a move to lower order and hence produces an inappropriately low number of patterns (the order parameter is connected to regularity – the higher the order the more repeating patterns are expected). Gerstein proposed to use a non-uniform dithering instead that is based on the square roots of the adjacent intervals, which he found to produce interval distributions remarkably similar to the original. To circumvent these problems, Pipa and colleagues offered an even simpler method: If all spike times within one train are dithered by the same amount, the spike trains are effectively shifted against each other, and coordinated firing is eliminated up to the corresponding timescale while the full auto-structure is kept intact[381] (Fig. II-3b). A third possibility is to randomly shuffle the inter-spike intervals, which means destroying the temporal structure while exactly preserving the original interval distribution. However, if all intervals are included in the shuffling, the rate profile might be changed to an unacceptable degree. As a solution, a variant of this method is introduced that only shuffles short intervals in between longer intervals that exceed the dither window (compare Hirata et al. 2008). Since in this way all events adjacent to longer intervals keep their position, the spike trains are additionally shifted against each other (Fig. II-3c). Through the combination of shuffled inter-spike intervals and misaligned spike trains, the resulting surrogates become even more dissimilar from the original data, which might make it more likely for individual patterns to be recognized as being significant.

Of great importance is the timescale that is chosen to dither single events or to shift event trains or to dissociate between short and long inter-event intervals: The resulting *average* displacement of an event should closely correspond to the timing precision that is used to define patterns to yield the best compromise between an extensive disarrangement of potentially coordinated events and the preservation of rate modulations on slower timescales (for a brief discussion on this topic see Pipa et al. 2008; see also Pazienti et al. 2008). Three different methods of dithering single events as well as the described methods of shifting event trains with and without additional shuffling of inter-event intervals were applied to simulated data and real recordings and will be evaluated in the results section (see Appendix 2 for a formal description). None of them is however able to assess the significance of pattern sequences. As mentioned before, these are statistically distinct from single spatiotemporal patterns and have to be tested independently. Since sequences are defined solely by a succession of pattern IDs irrespective of their temporal structure, we may generate appropriate surrogate data simply by randomizing the order of IDs in circumscribed stretches of data (for convenience, the same intervals as for the calculation of the functional coupling of units are used), in this way eliminating any potential dependencies between consecutive patterns while approximately preserving each pattern's rate profile. Once the surrogate data are constructed, we can compare the original pattern and sequence counts to those that would be expected given independent events.

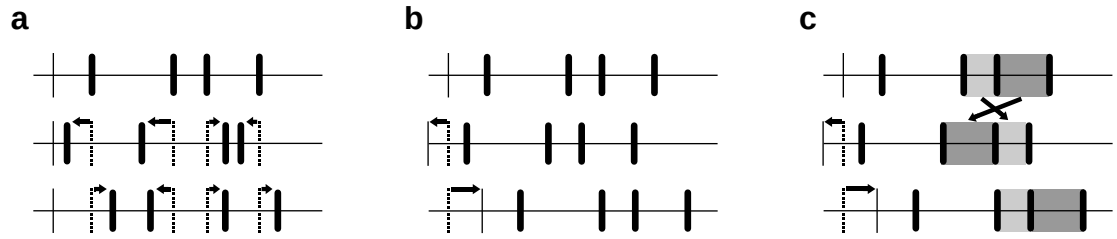

Fig. II-3 Three different ways of creating surrogate data for testing the significance of coordinated events in parallel time series. Top traces depict original event trains, subjacent traces depict surrogate trains that have been derived from the original data (thin vertical bars: onset). (**a**) The time of every individual event is dithered randomly and independently on a certain scale, thereby destroying the temporal structure contained across as well as within event trains up to that scale. (**b**) Whole event trains are shifted randomly against each other, in that way eliminating coordination of events up to the corresponding timescale while preserving the full auto-structure. (**c**) Same as in (**b**), but with additional random shuffling of consecutive inter-event intervals that are equal to or shorter than the maximal allowed shift (highlighted in light and dark gray).

Statistical test

If the null hypothesis was true, the numbers of repeating patterns and sequences extracted from the original data should be approximately the same in the surrogate data. To test the probability that the hypothesis holds, first the occurrences n_o of every repeating individual pattern and pattern sequence in the original data are counted and compared to the frequency of occurrence n_s of the same pattern or sequence in the surrogate data. Testing every pattern and sequence individually is necessary to rule out that its appearance is merely due to the frequencies of its components, the significance of the data as a whole notwithstanding.

Under the null, $n_o < n_s$ and $n_o > n_s$ are equally likely. Given the distribution of surrogate counts, the statistical significance of any pattern or sequence could thus be estimated by testing the relative frequency of $n_o > n_s$ against the expected value of 0.5 with an appropriate binomial test. For example, an exact binomial test gives a probability of at most ~ 6.3×10^{-23} for 95 surrogate counts out of 100 being lower than the original count by chance. A Bonferroni correction for multiple comparisons would in that case allow for ~ 7×10^{20} (!) parallel comparisons while maintaining a significance level of ~ 4.4% (likewise, the correction would allow for ~ 2×10^3 parallel comparisons at a significance level of ~ 4% if 19 surrogate counts from a total of 20 were required to be lower than the original count). It follows that large numbers of individual tests may be performed in parallel at reasonable significance levels if a sufficient number of surrogate data sets are taken into account. As an alternative to the binomial test, a simple heuristic is therefore employed that inherently allows for multiple comparisons and considerably reduces the computational complexity: If at least 95% of the counts from the surrogate data fall below the original count, then H_0 is rejected at a designated significance level of 5% or less, and one can conclude that the original count is unusually high. The minimal necessary number of surrogate data sets is accordingly given by the number of parallel tests and the desired significance level (see examples). After evaluation of their individual statistical significance, insignificant patterns and sequences are discarded.

In a second step, the coordination of events is assessed on a global level. To do so, every pattern and sequence that appears more than once in any surrogate data set is individually tested for significance in the very same way as those occurring in the original data, capitalizing on the assumption that the generation of the surrogate data did not affect the statistic under investigation and that all data sets, including the original one, are essentially indistinguishable with regard to the patterning of events. As a result, every data set is characterized by a certain number X of patterns or sequences that recur unexpectedly often, given their frequencies in the rest of the data sets. Their combined occurrences

$$N_d = \sum_{i=1}^{X_d} n_{id} \qquad (3)$$

(with n_i being the number of occurrences of the ith pattern or sequence that is statistically significant on the individual level and d being the index of the respective data set) are then subjected to a second level analysis to evaluate the overall significance of a coordination of firing events in the original data. The null hypothesis is rejected on the global level if at least 95% of the numbers N from the surrogate data fall below the number from the original data.

Of course, only data that pass the second level test can be considered to contain coordinated firing patterns, namely those patterns and sequences that have been found to be significant on the individual level. It has to be pointed out, though, that the original data might well be statistically distinguishable from the surrogate data despite lacking significant numbers of repeating patterns or sequences, because other statistics may be extracted as well and might prove to be different. One example is the rate of joint-spike events calculated per unit time and complexity (Fig. II-4). Here, however, the intention is to assess the significance of individual activity patterns based on a global evaluation of the whole data set.

Technical implementation

The whole analysis including the detection and statistical evaluation of spatiotemporal firing patterns and

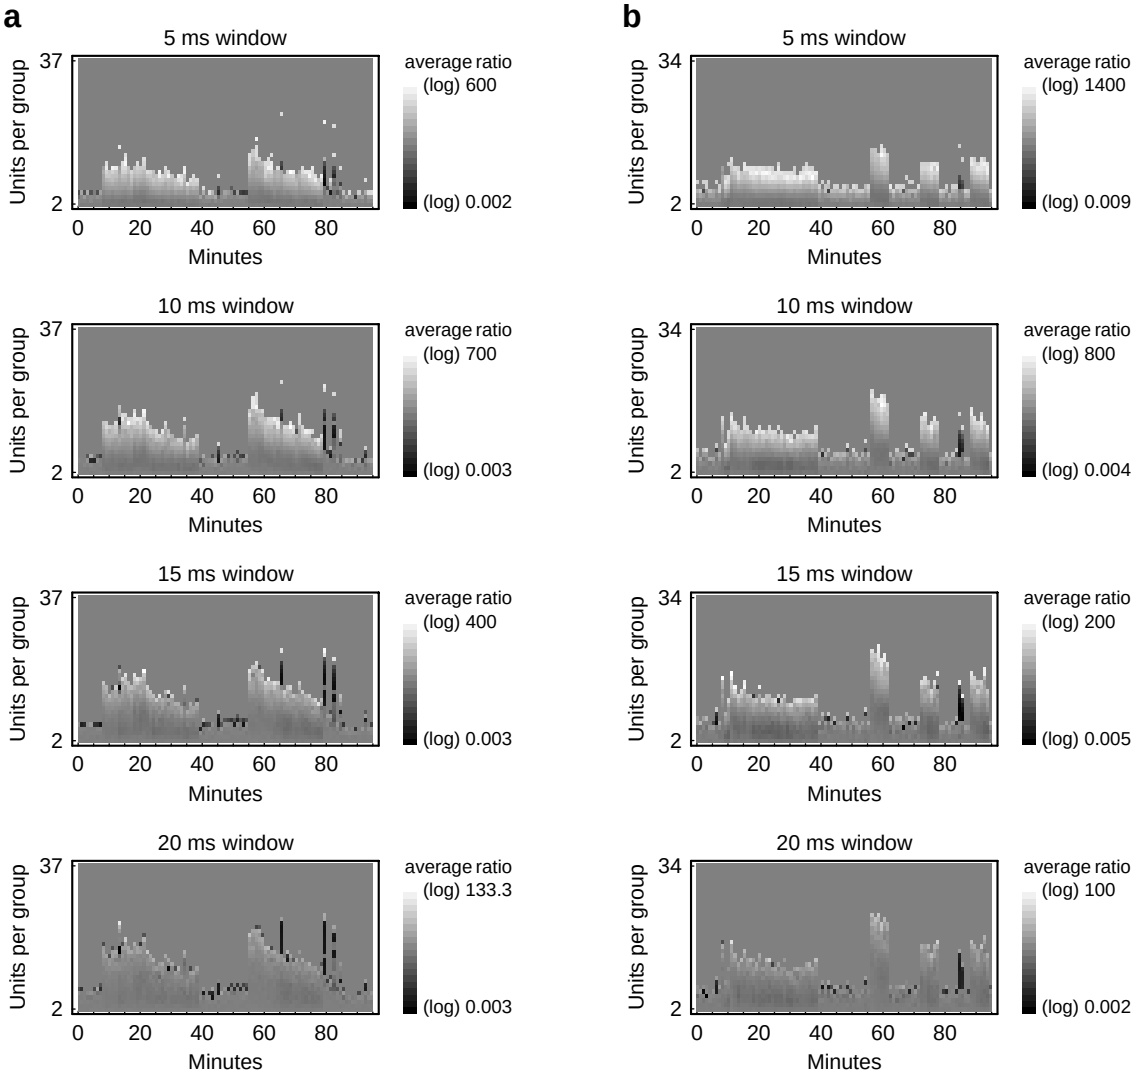

Fig. II-4 Relative frequencies of joint-spike events depending on time and complexity. A sliding window was moved from spike to spike to collect groups of coincident spikes (joint-spike events) in parallel recordings of neuronal activity from rat visual cortex in vitro. Their frequency of occurrence is compared against the corresponding values from 100 surrogate data sets that have been obtained by randomly shifting the original spike trains by an average amount of 5-7 ms. Relative frequencies are expressed on a logarithmic grayscale depending on the number of units per group (complexity) in data segments of one minute duration. The respective numbers of joint-spike events have been validated by an exact binomial test (50% gray indicates no significant difference between the original data and the surrogate data at a significance level of 0.05). (**a**) Data 1 recorded simultaneously from 101 units in 95 minutes. (**b**) Data 2 recorded simultaneously from 102 units in 95 minutes.

pattern sequences was programmed in *MATHEMATICA* (Wolfram Research, Champaign, Illinois) as a single computational process. Adjustable parameters are the duration T of the intervals that form the basis for calculating the correlation matrix and for generating surrogate data, the size W of the temporal window that is used for the pattern detection, the precision τ of the registered spike timing, the criterion A for the peer validation, the width w of the dither window, the number of surrogate data sets, the desired significance level, and the maximal sequence length. All repeating patterns and sequences are saved in text files along with the statistical results for further analysis.

Results

The following subsections give an account of the statistical and computational properties of the method by applying it to simulated data and multielectrode recordings from slices of rat visual cortex. The main findings are that dithering event times with the "square root dither" method is likely to change the interval distribution in a way that produces inappropriate surrogate data, that the new resampling method proposed here yields a slightly lower rate of false positives and is significantly more sensitive than the methods it has been compared to, that the detection of patterns is considerably facilitated by the flexible search algorithm and the controlled separation of concurrent events, and that the associated computational load can easily be handled by a conventional personal computer. See appendices 3 and 4 for details regarding the generation of simulated data and the recording of spiking activity, respectively.

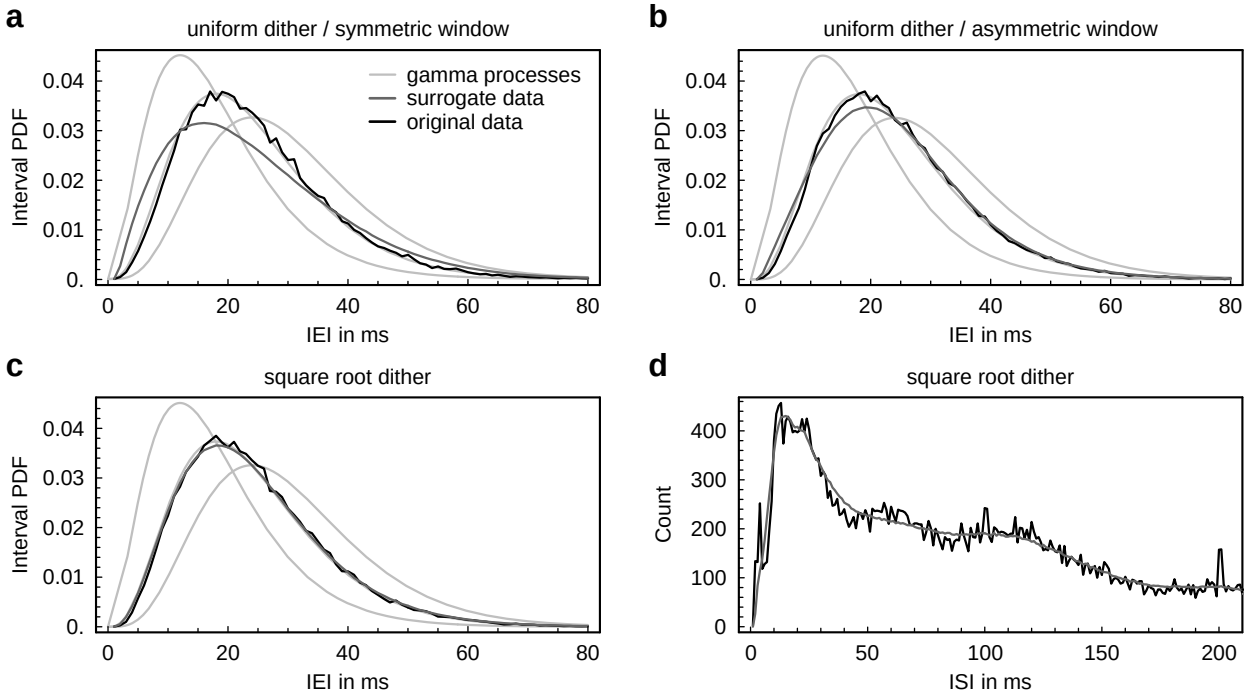

Fig. II-5 Effect of dithering event times on the interval distribution. (**a**, **b** and **c**) Simulated event trains were generated by gamma processes of order 4, and event times were dithered individually and independently using three different methods (number of surrogates: 100). The resulting average probability density functions (PDF) of the inter-event intervals (IEI) are shown in comparison with the original data and with theoretical interval distributions of gamma processes of orders 3, 4 and 5. (**a**) Event times were dithered randomly and uniformly within a symmetric window of maximally ± 20 ms, but assuring a "refractory period" of 1 ms between events. The average absolute displacement of an event resulted in 4.1 ms. (**b**) Same as in (**a**), but allowing for an asymmetric window. The average absolute displacement of an event resulted in 6.1 ms. (**c**) Event times were dithered according to the "square root dither" method (Gerstein 2004) again using a maximal window of ± 20 ms and assuring a "refractory period" of 1 ms between events. The average absolute displacement of an event resulted in 3.9 ms. (**d**) Same as in (**c**), but dithering was applied to real spike trains recorded simultaneously from 67 units. The resulting average distribution of inter-spike intervals (ISI) is shown in comparison with the original data (bin size = 1 ms). The average absolute displacement of an event resulted in 6.1 ms.

Effect of dithering event times on the interval distribution

To elucidate potential complications when creating surrogate data by dithering event times independently, the impact of three different dithering procedures on the interval distribution was investigated (Fig. II-5). In the first instance, simulated event trains were generated by gamma processes of order 4 with a mean event rate of ~ 40 Hz, and a number of surrogates were constructed using each method, always assuring a "refractory period" of 1 ms between events as an upper bound for their displacement (see Appendix 2). If event times are relocated randomly and uniformly within a symmetric window, short intervals are added to the probability distribution and its peak is lowered, in this way producing surrogates with an inappropriately low number of repeating patterns (Fig. II-5a). If we allow for an asymmetric window, the effect gets attenuated and the resulting average probability density function exhibits more regular intervals, but still does not conform to the original data (Fig. II-5b). As previously shown by Gerstein[167], randomizing the event timing within a window given by the square roots of the adjacent intervals and squaring the resulting offset while keeping its sign accurately preserves the original interval distribution (Fig. II-5c). Ensuring a "refractory period" between events alone does not suffice, at any rate, to get appropriate surrogate data.

Since neuronal firing statistics typically defy analytical formulation, we also examined the impact of dithering event times independently on a distribution of real inter-spike intervals to check if it can be as accurately preserved as simulated interval distributions (Fig. II-5d). Surrogate data were again constructed using the "square root dither" method[167] while assuring a "refractory period" of 1 ms between spikes. In contrast to simulated interval distributions, the examined distribution of real inter-spike intervals exhibits local minima and maxima and is conspicuously smoothed as a result of the dithering (the small peaks around 100 ms and 200 ms are due to indirect electrical stimulation). It is unclear if and how this affects the number of repetitive patterns in the surrogate data, but one might suspect that locally decreasing the regularity of the intervals again produces inappropriately low pattern counts.

Error levels in pattern detection and validation

To assess the probability of false positives when patterning of events is actually at chance level, sets of 30 parallel simulated time series generated by inhomogeneous gamma processes were analyzed. Spatiotemporal patterns and sequences were detected and tested both during independent modulations and covariations of event rates using three different techniques for creating surrogate data (dithering single event times with the "square root dither" method and shifting event trains

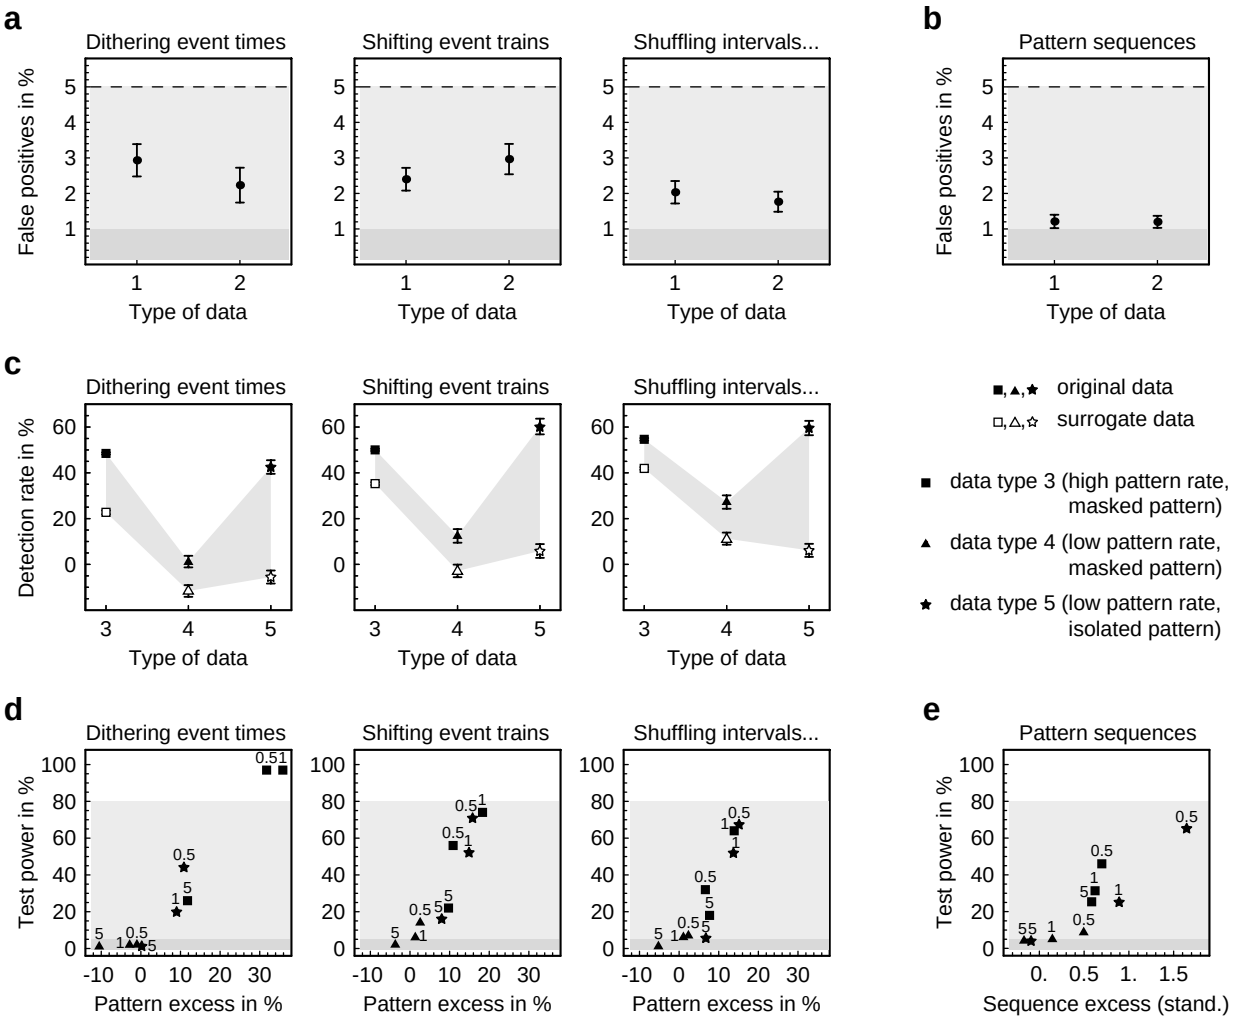

Fig. II-6 Error levels in pattern detection and validation. Sets of 30 parallel simulated time series were generated by gamma processes featuring independent rate modulations (data type 1), rate covariations (data type 2) and precisely repeating spatiotemporal patterns and pattern sequences (data types 3, 4 and 5) (see Appendix 3 for further description). Error levels were estimated for each parameter combination based on 100 independently simulated data sets. (**a**) Rate of false positives (mean ± SEM) when analyzing spatiotemporal patterns using 30 different pattern definitions given by ten different time windows W (5, 10, 15, 20, 25, 30, 35, 40, 45 and 50 ms) and three different relative timing precisions (10 bins, 5 bins and rank order). (**b**) Rate of false positives (mean ± SEM) when analyzing pattern sequences using the same pattern definitions as in (**a**) three times each. (**c**) Detection rates of spatiotemporal patterns (mean ± SEM) calculated as the difference of significant pattern counts from data with and without additional patterns, normalized with respect to the number of inserted patterns, and averaged across 100 data sets using three different pattern definitions (5 ms window, timing precisions of 10 bins, 5 bins and rank order). (**d**) Test power for an excess of spatiotemporal patterns using a window W of 5 ms and timing precisions τ of 0.5 ms, 1 ms and 5 ms (rank order) as indicated (excess averaged across 100 data sets). (**e**) Test power for an excess of pattern sequences using the same parameters as in (**d**) (excess displayed as the standardized mean difference and averaged across 300 data sets).

with and without additional interval shuffling, see Appendix 2) and 30 different pattern definitions corresponding to the combinations of ten different time windows W (5, 10, 15, 20, 25, 30, 35, 40, 45 and 50 ms) and three different relative timing precisions (10 bins, 5 bins and rank order) (cf. Fig. 1). The generation of the surrogate data was balanced such that the average displacement of an event always resulted in 6-8 ms; also the remaining parameters were kept constant (significance level $p \leq 0.05$, number of surrogate data sets: 20, absolute peer criterion: 2, maximal sequence length: 10). The rate of false positives is expressed as the mean percentage of 100 simulated data sets that passed the second level test. It turned out to be considerably below 5% in all cases (Fig. II-6a, b), demonstrating that the presented method constitutes a conservative statistical test for the evaluation of spatiotemporal patterns and pattern sequences even when the instantaneous event rates closely covary. Interestingly, creating surrogate data by combined interval shuffling and event train shifting seems to produce a little less false positive estimates than shifting event trains alone or the "square root dither" technique (Fig. II-6a).

To assess the test power of the method under various conditions, sets of 30 parallel simulated time series were analyzed that were generated by homogeneous gamma processes including different numbers of inserted recurring spatiotemporal patterns and pattern sequences (see Appendix 3 and Fig. 9 for details). The same three procedures for creating surrogate data as before together with three different pattern definitions (5 ms window, timing precisions of 10 bins, 5 bins and rank order) were used to uncover the hidden structure in the parallel event trains. Again, the generation of the

surrogate data was balanced such that the average displacement of an event resulted in 6-8 ms. The remaining analysis parameters were retained as well (significance level $p \leq 0.05$, number of surrogate data sets: 20, maximal sequence length: 10) except for the absolute peer criterion which was matched to the frequencies of the inserted patterns ($A = 5$ when patterns appeared every second, $A = 1$ when patterns appeared every five seconds).

First of all the number of significant spatiotemporal patterns from data containing precisely repeating patterns was compared to the number from data that had been simulated using the same parameters but featuring rate covariations instead of predefined patterns at corresponding positions (Fig. II-6c). As expected, the data containing precise patterns exhibit an excess of significant patterns on the investigated timescales. Remarkably, this holds both for original and for surrogate data, indicating that randomly varying the event timing by an average amount that only slightly exceeds the timescale on which events are coordinated still preserves some patterning of events on that same scale (compare Pazienti et al. 2008). Nevertheless, the surrogate data contain smaller numbers of significant patterns than the original data, particularly if patterns are not masked by collateral events. The *relative* rate of significant patterns decreases with a decreasing *absolute* number of inserted patterns and increases drastically if all unrelated events are removed from the patterns, emphasizing the importance of separating concurrent but independent events. Furthermore, the detection rate is always higher when shifting complete event trains to produce surrogate data as compared to dithering single events with the "square root dither" method ($p < 0.02$, exact binomial test), and it is yet higher when additional interval shuffling is applied in cases where patterns are masked by concurrent events ($p < 0.03$, exact binomial test), especially if pattern rates are low. A controlled combination of shifting whole event trains and shuffling inter-event intervals thus seems to be advantageous over previous resampling methods, particularly under statistically more demanding conditions. (Note that the displayed detection rates do not directly reflect the fraction of inserted patterns that have been found to be significant. In fact, *every* pattern that repeats without collateral events will be detected.)

The test power is expressed as the percentage of 100 simulated data sets that passed the second level test (Fig. II-6d, e). For spatiotemporal patterns, it reaches 80% at an excess of patterns in the original data of about 20% as compared to the surrogate data and strongly depends on the pattern rate, the separation of patterns from unrelated events, and the adaptation of the registered timing precision to the actual precision of the patterns (Fig. II-6d). The same is true for pattern sequences which require an excess of approximately 1-2 standard deviations to be reliably detected (Fig. II-6e). The choice of the method for creating surrogate data does not seem to substantially influence the test power, except for the "square root dither" technique of dithering event times independently which is superior if pattern rates are high and inferior if pattern rates are low. Considering the fact that on average no more than 11.8% (SD 1.2) and 2.6% (SD 0.3) of the events were coordinated in data featuring high and low pattern rates, respectively, it appears that the presented method has impressive power to detect precisely repeating patterns. Nonetheless, it relies heavily on the correct isolation of spatiotemporal patterns and the matching of the timescales on which events are coordinated and registered, most notably when sequences of patterns are concerned. From this perspective, having the possibility to accurately specify the event timing in spatiotemporal patterns with arbitrary precision, as provided by the method presented here, is essential.

Selectivity in pattern detection and validation

To evaluate the performance of the method when analyzing real data and varying the way of detecting or the way of validating patterns, the dependence of the number N of significant patterns and sequences on the peer criterion A, the method for creating surrogate data, and the number of surrogates was investigated from parallel recordings of neuronal spiking activity (Fig. II-7). The data were scanned for repeating patterns and sequences using ten different time windows W (5, 10, 15, 20, 25, 30, 35, 40, 45 and 50 ms) and three different relative timing precisions (10 bins, 5 bins and rank order) for representing patterns, a significance level of $p \leq 0.05$, a maximal sequence length of 10, and an interval T of 1 minute as the basis for calculating the correlation matrix and for generating surrogate data. The generation of the surrogate data was balanced such that the average displacement of a spike always resulted in 5-9 ms.

Applying a global peer criterion in addition to the threshold given by equation (1) and varying it between 1 (no further splitting) and 8 (at least eight coincident spikes within one minute required to label peers as valid) has a dramatic effect on the numbers of significant patterns and sequences (Fig. II-7a), clearly indicating that the proposed method of extracting subpatterns from larger spatiotemporal patterns can improve pattern detection. Given the data and the chosen time interval T, the largest number of significant spatiotemporal firing patterns could be detected using a peer criterion of 2, suggesting that the expected numbers of coincident spikes per unit pair have in some cases been calculated to be smaller. Further incrementing the criterion essentially removes more and more units from the patterns and increasingly impairs their information content, so that their numbers of occurrences approximate chance level. Sequences of patterns, however, apparently become increasingly significant if the threshold for labeling peers as valid is raised, and the distance between original and surrogate counts is concurrently growing.

Creating surrogate data by randomly shifting the spike trains against each other yields marginally less significant spatiotemporal firing patterns if additional interval shuffling is applied, but the difference is practically negligible (Fig. II-7b). In contrast, dithering single spike times with the "square root dither" method

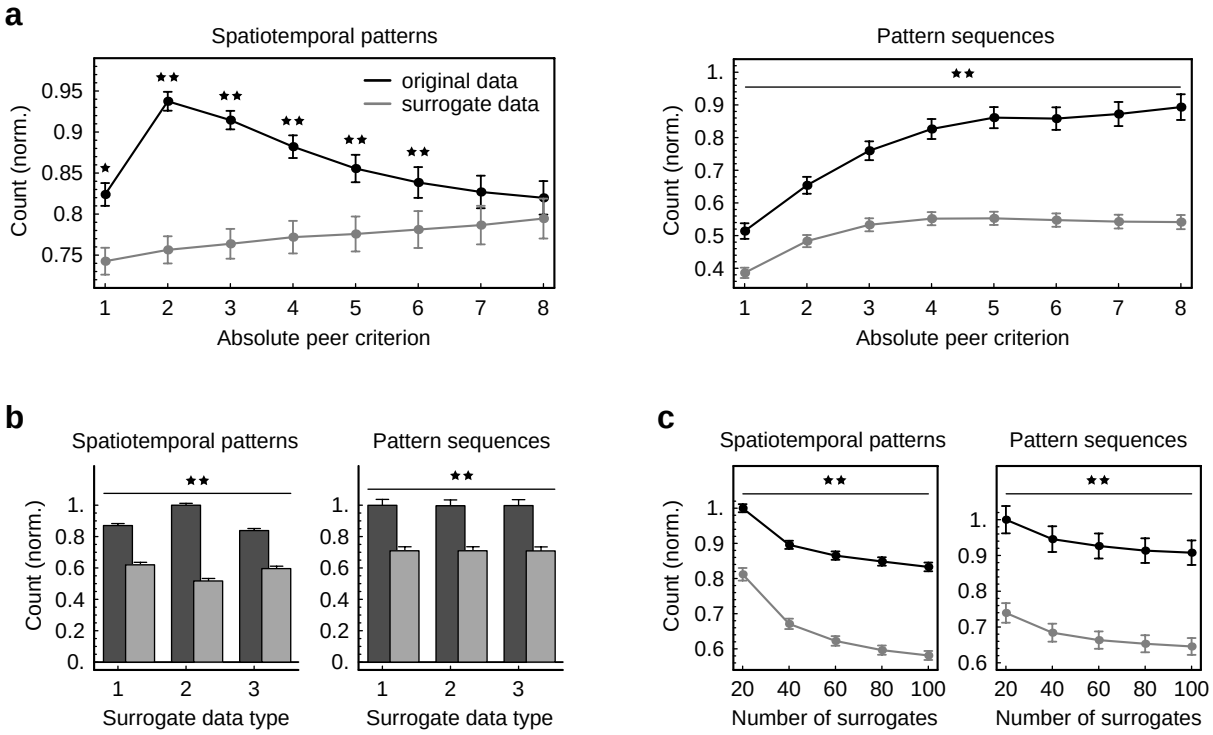

Fig. II-7 Dependence of the number *N* of significant patterns and pattern sequences on analytical parameters. Multineuronal spiking activity was scanned for spatiotemporal patterns using 30 different definitions given by ten different time windows *W* (5, 10, 15, 20, 25, 30, 35, 40, 45 and 50 ms) and three different relative timing precisions (10 bins, 5 bins and rank order). The numbers of significant patterns and sequences were normalized to be comparable across data sets (mean ± SEM across 10 data sets and 30 pattern definitions). Stars denote a significant difference between original and surrogate data (★ $p < 0.004$, ★★ $p < 0.00007$, exact binomial test). (**a**) Counts depending on the absolute peer criterion *A* applied in addition to the threshold given by equation (1) (number of surrogates: 20, surrogate data generated by combined spike train shifting and interval shuffling). (**b**) Counts depending on the surrogate data type (1: shifting spike trains, 2: dithering individual spike times with the "square root dither" method, 3: shifting spike trains and shuffling intervals, number of surrogates: 100, absolute peer criterion: 2). Original counts are displayed in dark gray, surrogate counts in light gray. (**c**) Counts depending on the number of surrogates (absolute peer criterion: 2, surrogate data generated by combined spike train shifting and interval shuffling).

while assuring a refractory period of 1 ms between spikes significantly increases the average number of significant patterns in original data and decreases the average number in surrogate data compared to the other techniques ($p < 8 \times 10^{-28}$, exact binomial test), possibly due to smoothing of the inter-spike interval distribution which likely introduces a bias towards less repeating patterns. Since sequences of patterns are evaluated independently of the significance of individual patterns, the numbers of significant sequences are unaffected by the type of surrogate data used to validate spatiotemporal patterns. Increasing the number of surrogates markedly reduces the average significant pattern counts obtained both from original data and from surrogate data by about the same amount (Fig. II-7c), reflecting the fact that the significance estimation becomes increasingly conservative.

Comparing the counts from original data and from surrogate data shows that randomly displacing spikes by only a few milliseconds leads to a significant drop of repeating spatiotemporal firing patterns within a wide range of parameters (Fig. II-7), demonstrating a remarkably precise coordination of multineuronal suprathreshold activity in the local cortical circuitry. Moreover, directly consecutive firing patterns are organized to a significant degree into repetitive sequences, revealing some superordinate temporal structure beyond the cell assembly concept. Importantly, repeating firing patterns and repeating sequences of patterns occur both spontaneously and in response to electrical stimulation, suggesting that they are an inherent feature of intracortical signaling.

Computational demands

To illustrate the computational requirements of the method, the memory consumption and the processing time for a complete analysis of two example recordings of multineuronal spiking activity was measured depending on the number of surrogate data sets and the window size *W* used for pattern detection (Fig. II-8). The data were scanned for spatiotemporal firing patterns and pattern sequences using the temporal order of spikes to define patterns and a maximal sequence length of 10. Surrogate data were constructed by combined spike train shifting and interval shuffling. The computation was carried out as a single process on a 32-bit machine with a "Pentium 4" CPU running at 2.4 GHz using *MATHEMATICA* version 5.2.

As projected, the processing time grows linearly with the number of surrogates and strongly depends on the average pattern complexity that results from the applied window size and the spike rate (Fig. II-8a). The same approximately holds for the memory consumption (measured using the *MATHEMATICA* command "MaxMemoryUsed"), but here the slope accompanying

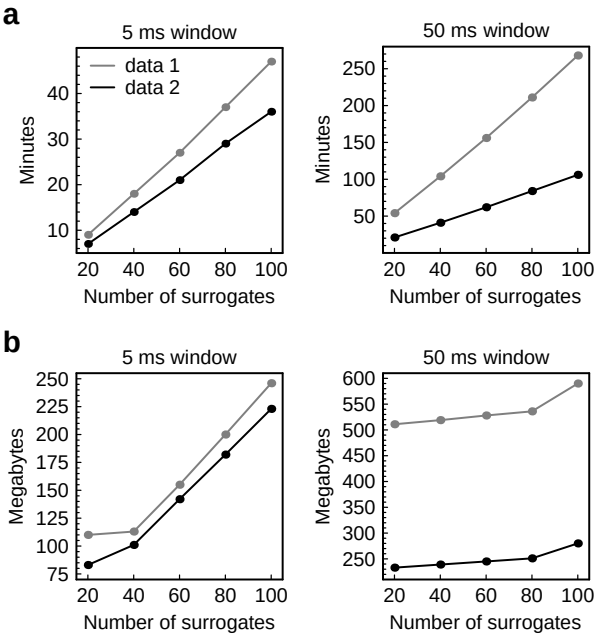

Fig. II-8 Computational demands of the analysis. Parallel recordings of multineuronal spiking activity were scanned for significant spatiotemporal firing patterns and pattern sequences using the temporal order of spikes to define patterns. Data 1 comprises ~ 194000 spikes recorded simultaneously from 66 units in 96 minutes, data 2 comprises ~ 93000 spikes recorded simultaneously from 73 units in 96 minutes. (**a**) Dependence of the processing time on the number of surrogates and the window size used for pattern detection. (**b**) Dependence of the memory consumption on the number of surrogates and the window size used for pattern detection.

an increase in the number of surrogate data sets is less pronounced (Fig. II-8b). The absolute values show that the computational load can be handled with ease by a single personal computer, even in the case of large data sets.

Discussion

The quest for the neuronal code has led to extensive controversies about the relevant timescales and particular organization of neuronal firing and is commonly considered an unresolved issue[25,92,128,129,175,429,449]. To achieve a complete understanding of neuronal information processing, we need to precisely characterize the dynamic dependencies between cells and the temporal relationships between their discharges beyond pairwise correlations[24,54,198].

For this purpose, I presented a straightforward and computationally efficient method for detecting temporally coordinated firing events in parallel spike trains. The method is generally applicable and implies no assumptions about the statistical properties of the data or the spatiotemporal structures contained therein. Focussing on the activation sequence of cells, activity patterns are captured on arbitrary timescales and may or may not show signatures of functional cell assemblies, synchrony, synfire chains or synfire braids. By utilizing carefully modified versions of the original spike trains to assess the significance of any detected activity pattern, the method allows for variability and sparseness of spiking events as well as the analysis of very short data segments. The temporal offset of events in the surrogate data selectively separates the timescales on which coordination of spikes is disturbed and preserved, respectively, and provides a means to directly address the temporal coding and rate coding hypotheses. In effect, the method acts as a filter revealing repetitive spatiotemporal patterns and pattern sequences amongst distributed discharges, yielding a comprehensive description of the neuronal activity on the selected timescales. Through a subsequent analysis of significant patterns, coordinated firing may be characterized in relation to neuronal state changes and information processing with single-spike resolution. First results obtained in the analysis of simultaneous recordings from rat visual cortex demonstrate a millisecond-precise coordination of neuronal spiking and reveal some superordinate patterning beyond the cell assembly concept.

Conceptually, the method is not restricted to the analysis of multiple spike trains. In principle, any parallel time series can be investigated just by adapting the temporal scales. For example, stimulus times or behavioral events may readily be included. Another possible application area is the analysis of so-called multi-voxel patterns in functional magnetic resonance imaging data[358]. In the following, the method is contrasted with existing approaches, and a number of related issues is discussed in more detail.

Comparison with other methods

Common approaches in the analysis of correlational structures in parallel spike trains differ with respect to the particular property in question. While this paper deals with the detection of spatiotemporal firing patterns, other methods investigate the functional coupling between neurons on longer timescales[43,259,296] or concentrate specifically on synchronous firing[187-189,253,381]. It is important to further distinguish the method presented here from methods trying to identify genuine higher order correlations between neurons[191,317,419,433,474]. Whereas the latter aim at discovering multineuronal interactions, the focus of the present work is more directed towards dynamically changing activity patterns arising from these interactions.

To date, only a few publications have specifically addressed the issue of detecting and evaluating spatiotemporal firing patterns. In a pioneer work, Abeles and Gerstein developed a method to capture precisely repeating spike patterns and estimate the significance of classes of patterns including up to six spikes using a parametric approximation[5]. Although partly relying on the same statistical assumptions, Tetko and Villa subsequently resolved the limit on the complexity of the patterns[485] and succeeded in assigning significance to patterns that had been grouped by matching single occurrences with a predefined template[486]. A different approach was pursued by Frostig and colleagues who collected significant patterns of increasing complexity in a cascaded fashion using Fisher's exact test[151]. In the

original formulation, higher order patterns are only detected if all subpatterns formed by at least three spikes are themselves statistically significant (including all possible configurations would lead to an unmanageable combinatorial explosion). Along similar lines, Sastry and Unnikrishnan proposed to identify frequent temporal patterns by incrementally assembling larger patterns from smaller ones that have been found to be significant under some stationarity assumptions[415].

Instead of looking for temporally precise repetitions of multineuronal firing patterns, Lee and Wilson emphasized the significance of the activation order of neurons and presented an elegant way to quantify and statistically rate the degree of matching between a found firing sequence and a preselected reference pattern[280] (see also Smith et al. 2006, 2010). Yet another avenue to uncovering hidden structure in multineuronal spiking is to identify the predominant activity pattern. Yamada et al. took an information theoretic approach to reconstruct the directed functional connectivity among a number of neurons, including exact delays[546]. Nikolić and co-workers exploited pairwise measurements of characteristic temporal relations to detect more complex firing sequences, expressed as precise inter-spike intervals[420] or relative temporal order of spikes[353]. Analyses of this kind are suited to detect differences in the preferred firing sequences between different experimental conditions, but they fail to reproduce the full dynamics of neuronal interactions within short time intervals. So, none of these methods is designed to systematically search for all repeating spatiotemporal patterns on any given timescale and evaluate them both individually and globally without making any statistical assumptions, quite contrary to the method presented here.

The work by Gerstein and colleagues has hitherto been the only attempt to provide a way for directly detecting superordinate activity patterns forming synfire chains[173,422]. The authors utilized specially devised recurrence plots to indicate pairs of temporal bins containing approximately repeating firing patterns across the whole population of recorded neurons. Synfire chains can readily be identified if the repeating patterns are not masked by too many concurrent spikes from cells that do not participate in the same link, if the applied bin width matches the temporal organization of the chain, and if the interlink propagation delays are stable (as postulated for synfire chains). In contrast, I propose to register spatiotemporal firing patterns with a sliding window in the first instance and to subsequently search for recurring sequences of temporally non-overlapping patterns irrespective of their particular spatiotemporal structure, trying to make as few assumptions about the organization of multineuronal activity as possible.

Pending problems in pattern detection

Any way of defining the identity of a pattern, be it by using some template for the pattern search or by specifying the pattern as it is captured, inevitably poses the problem of where to draw the line between similar patterns. For example, it is completely unclear how to divide a continuum of patterns if the classification scheme permits missing or extra spikes. The same principle applies to registering the spike timing: The assignment of exclusive time bins, even the usage of "smooth" templates imposes perfectly arbitrary boundaries on the temporal patterns to be detected. For these reasons, it is most straightforward to require patterns to contain no extra or missing spikes to be unambiguously identified, and it appears that the only way to avoid temporal aliasing effects is to specify spike times solely by their temporal order.

Another problem, as mentioned earlier, is the mutual masking of concurrent but independent patterns. To reliably unravel simultaneously occurring but unrelated patterns, one would have to systematically list all possible subpatterns and to subsequently distill the significant ones. If more than a few units are involved, the resulting combinatorial explosion may however render the computation impractical. Therefore, I propose to group simultaneously active cells according to their joint firing probabilities, which essentially produces a subset of all possible subpatterns. In doing so, some significant patterns might however be missed, and the overall significance of the data might be underestimated, resulting in reduced test power (Fig. II-6). The functional coupling between cells can easily be assessed by comparing the number of synchronous spikes to the chance level of coincident firing as estimated by equation (1), assuming serial independence and stationarity of firing events. As an alternative, one could employ conventional cross-correlation techniques, among which information theoretic approaches seem to be most sensitive[547]. Both methods fail to give meaningful results if events are scarce, which is why an additional global threshold with respect to the joint firing probabilities was implemented to identify each neuron's peers. Besides, other methods to quantify the correlation between two spike trains exist which may as well be applied, like for example fitting pairwise maximum entropy models to the data[408] or determining the degree of synchrony between spike trains (for a comparison of different innovative measures see Kreuz et al. 2007). The particular advantage of the method of estimating the time-dependent functional coupling between cells proposed here is its technical simplicity and ease of computation. Despite its heuristic character, it has clearly demonstrated its potential to improve pattern detection (Fig. II-7a). Nevertheless, the thorough evaluation of *all* possible combinations of events that coincide in the given time window is to be preferred whenever it is technically feasible.

Statistical issues

Since the complexity of multineuronal spiking activity typically prevents analytical approaches, resampling methods have to be applied to test a certain null hypothesis[186,473]. Here, three different surrogate data types for testing the significance of coordinated events in

parallel time series were compared (Fig. II-3). Dithering event times randomly and independently almost inevitably changes the interval distribution and introduces a bias towards inappropriately low levels of coordination if uniform dithering is used, regardless of whether the dither window is always centered on the event or not (Fig. II-5a, b). The only known technique that randomizes event times independently and at the same time does not significantly change the interval distribution of a gamma process is the "square root dither" method proposed by Gerstein[167] (Fig. II-5c). However, this technique conspicuously flattens the more modulated distribution of real inter-spike intervals (Fig. II-5d), rendering them less regular. As a consequence, the surrogate data can again be expected to contain an inappropriately low number of coordinated events, which in turn would result in an increased number of spuriously detected patterns in the original data (Fig. II-7b).

If the intention is to disarrange coordinated events, I therefore suggest to randomly shift whole event trains against each other, preferably in combination with random shuffling of short inter-event intervals (Fig. II-3b, c). While preserving the complete auto-structure is the only way to perfectly account for effects of e.g. spike bursts or oscillatory processes on the apparent coordination of events, shuffling selected intervals in addition to shifting event trains increases the distance between original data and surrogate data and markedly facilitates the detection of patterns (Fig. II-6c). At the same time, the surrogates also become more dissimilar from each other, which leads to an increased number of statistically significant patterns in every surrogate data set and makes the second level test slightly more conservative with respect to the original data (Fig. II-6a, d). Recently, Harrison and Geman presented the idea to randomly and independently shift segments of spike trains so that inter-spike intervals are changed only in between them[201]. How such a procedure affects the validation of patterns and the corresponding error levels remains to be investigated.

Several factors inherent in the presented method should be highlighted that tend to result in conservative estimates of the significance of patterns: First, patterns of a higher complexity may be split into multiple patterns with lower complexity as a consequence of the surrogate data generation (reflected in Fig. II-4) which may produce misleadingly high numbers of repeating patterns in the surrogate data. Second, the surrogate data sets are more dissimilar from each other than from the original data, because the original data constitutes their common source. This consequently implies an increased potential to contain unique, statistically significant patterns, introducing a bias towards higher pattern counts in the surrogate data. The finding that under certain circumstances the original data contain significantly less repeating patterns than the respective surrogate data (Fig. II-6d) is arguably attributable to these first two factors. A third potential source of conservative estimates is the statistical test itself: Directly calculating the percentage of counts from the surrogate data that fall below the original count is the most simple test one might think of. When applying this test, one might confidently do so without explicitly correcting for multiple comparisons. The same does not hold if more refined tests are used, like for example the binomial test, which is perfectly applicable and may serve as an alternative but would clearly require a correction. The *t*-test, like the Wilcoxon signed rank test, is unsuited because of its sensitivity to skewed, discontinous distributions[476] and especially inapplicable if many zero counts generate a floor effect, as is to be expected when looking for complex patterns.

Because the statistical significance of individual patterns and sequences is assessed based on the overall number of their occurrences in the whole data set, they may be significant despite being rare; all that matters for them to be detected as being non-random is that they occur significantly more often in the original data than in the surrogate data, irrespective of other criteria like their particular time of occurrence or their complexity. This distinguishes our method from others in which the statistical evaluation is based on trials or relies on an exact calculation of the probability of occurrence under the null hypothesis. In the presented Monte Carlo approach, the chance level of pattern and sequence occurrences is derived directly from the original data, thus avoiding the "curse of dimensionality" inherent in many other approaches that relate the original data to some exact model and require larger data volumes for assessing the significance of complex patterns. Besides, effects of firing rate on the statistical performance of the method are minimized by disregarding all but the first spike of every unit inside the respective time window and can further be counteracted by adapting the parameters for the pattern search. In conclusion, there are no specific requirements regarding the minimum amount of investigated data, although the probability of a rare pattern to recur and thus to be detected clearly increases with increasing recording time.

Possible technical improvements

Although the analysis has been programmed as a single computational process, it can easily be adapted to parallel computing environments by insertion of only a few extra commands that distribute the separate Monte Carlo steps over the available cores. In the very near future, the use of multi-core systems and the integration of general-purpose GPUs will enable the evaluation of data comprising hundreds of thousands of events in at most a few minutes. At the same time, the availability of 64 gigabytes of working memory or more in a single workstation should make it possible to systematically analyze all spatial subpatterns and hence to directly detect genuine higher order correlations between units also in large data sets. After all, the problem of detecting significant spatiotemporal patterns in massively parallel time series is not a conceptual one, but one of computational resources.

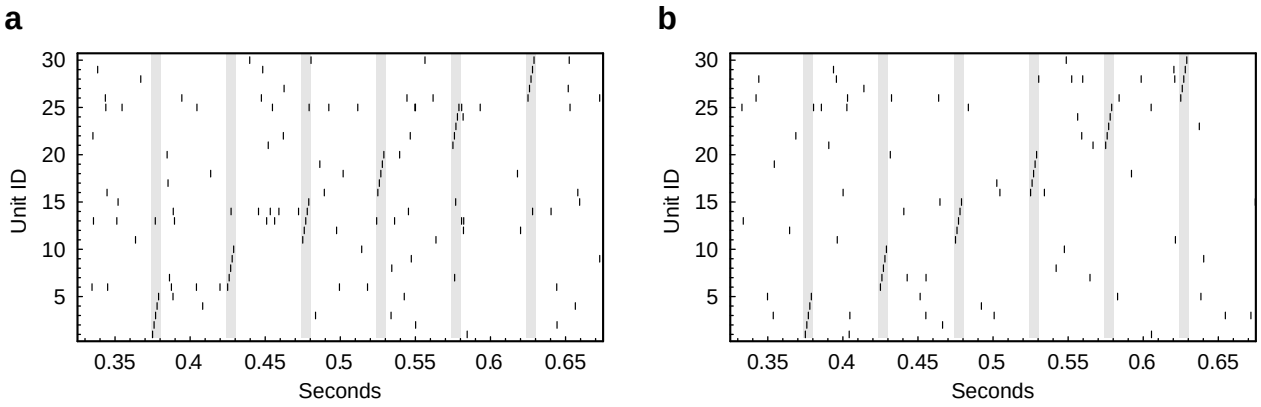

Fig. II-9 Spike raster plots exemplifying simulated data segments including inserted spatiotemporal patterns and pattern sequences. 30 parallel spike trains were generated by gamma processes with rate parameter β = 49 and random shape parameter α = 0.7-7. Patterns consisted of 5 spikes with millisecond intervals and followed each other with an onset delay of 50 ms, giving rise to a synfire chain-like structure made up of 6 distinct patterns with exactly 1 spike per unit (periods of pattern occurrences highlighted in gray). (**a**) Some events co-occur by chance with the inserted patterns, complicating their detection (data types 3 and 4). (**b**) Events concurrent to but not participating in the inserted patterns have been removed to ensure the correct identification of the patterns (data type 5).

Appendix 1: Derivation of the rate-based chance level of spurious coincidences

Let $n_i^{(T)}$ be the (known) number of events of unit i in time interval T, then the number $k_i^{(W)}$ of events of unit i in time window W (with $W < T$) can be approximated as

$$k_i^{(W)} = \frac{W}{T} n_i^{(T)}, \tag{4}$$

assuming serial independence and stationarity of events. Under the condition that $T/W \geq n_i^{(T)}$, the probability P of coinciding events of M different units labeled $1...M$ in time window W is

$$P_{1...M}^{(W)} = \prod_{i=1}^{M} k_i^{(W)} \ . \tag{5}$$

Since T/W is the number of time windows W in time interval T, multiplying equation (5) with this number gives the expected number of coincidences of M different units in time interval T:

$$P_{1...M}^{(T)} = \frac{T}{W} \prod_{i=1}^{M} k_i^{(W)} \ . \tag{6}$$

Insertion of equation (4) and conversion yields

$$P_{1...M}^{(T)} = \left(\frac{W}{T} \right)^{M-1} \prod_{i=1}^{M} n_i^{(T)} \ . \tag{7}$$

In the case of $M = 2$, equation (7) becomes

$$P_{ij}^{(T)} = \frac{W}{T} n_i^{(T)} n_j^{(T)} \tag{8}$$

with $P_{ij}^{(T)}$ being the expected number of coincidences of units i and j in time interval T, and $n_i^{(T)}$ and $n_j^{(T)}$ being the numbers of events of units i and j in time interval T.

Appendix 2: Surrogate data generating procedures

Formally, the methods for generating surrogate data are expressed as follows. Definitions include t_i being the ith timestamp of an event train $t_{1...n}$ with n events labeled $1...n$, the preceding inter-event interval $\delta_p = t_i - t_{i-1}$, the subsequent inter-event interval $\delta_s = t_{i+1} - t_i$, the maximal width w of the dither window, and a random number $r \in \mathbb{R}$.

Single event times were dithered randomly and independently according to

$$t_i^{(S)} = t_i + r \ . \tag{9}$$

The random number r is bounded above by υ_s and below by $-\upsilon_p$ such that $-\upsilon_p \leq r \leq \upsilon_s$. The bounds were defined depending both on the adjacent inter-event intervals and on w. When the resulting dither window was required to be centered on the event ("symmetric dither"), the bounds were specified corresponding to

$$\upsilon_p = \upsilon_s = \min(\delta_p - 1 \text{ ms}, \delta_s - 1 \text{ ms}, w) / 2. \tag{10}$$

In cases where the window center was allowed to deviate from the time of the event ("asymmetric dither"), the bounds were calculated independently from each other:

$$\upsilon_p = \min(\delta_p - 1 \text{ ms}, w) / 2, \tag{11}$$

$$\upsilon_s = \min(\delta_s - 1 \text{ ms}, w) / 2. \tag{12}$$

In both cases, the probability for an event to occur is distributed uniformly over the dither window. By contrast, the "square root dither" method (Gerstein 2004) relocates event times randomly within a window that is composed of the square roots of υ_p and υ_s as given by equations (11) and (12):

$$\upsilon_p = \sqrt{\min(\delta_p - 1 \text{ ms}, w) / 2} \ , \tag{13}$$

$$\upsilon_s = \sqrt{\min(\delta_s - 1 \text{ ms}, w) / 2} \ . \tag{14}$$

The offset r results from drawing a random number $q \in \{-\upsilon_p; \upsilon_s\}$ and squaring its absolute value while keeping its sign, such that $r = q \times |q|$ with $|...|$ denoting the absolute value.

Event trains were shifted randomly and independently according to

$$t_{1...n}^{(S)} = t_{1...n} + r \qquad (15)$$

with $-w/2 \leq r \leq w/2$. When additional interval shuffling was applied, consecutive inter-event intervals $\delta \leq w/2$ were randomly rearranged.

Appendix 3: Data simulation

Sets of 30 parallel simulated spike trains were generated on biologically plausible timescales by gamma processes featuring independent rate modulations (data type 1), rate covariations (data type 2) and precisely repeating spatiotemporal patterns and pattern sequences (data types 3, 4 and 5). By default, the rate parameter was fixed ($\beta = 49$), while the shape parameter (order) varied randomly between 0.7 (bursty) and 7 (regular) for any given spike train, resulting in mean firing rates of ~ 3-30 Hz. In rate modulated data (data type 1), the rate parameter was changed to a random value between 24 and 74 for five consecutive inter-event intervals chosen randomly from every twenty-five inter-event intervals, resulting in transient firing rates of ~ 2-60 Hz. Rate covariations (data type 2) were realized by jointly randomizing the rate parameter between 24 and 74 for one second every five seconds. Exactly repeating spatiotemporal spike patterns, arranged in a synfire chain-like structure (Fig. II-9), were inserted into the data every second (data type 3) or every five seconds (data types 4 and 5), leading to an average fraction of coordinated events of 11.8% (SD 1.2) and 2.6% (SD 0.3), respectively. Importantly, a distinction is made between data containing spikes concurrent to but not participating in the inserted patterns (data types 3 and 4, Fig. II-9a) and data with these events removed (data type 5, Fig. II-9b). The spike trains were truncated at 50 seconds and divided into 5 second intervals as the basis for calculating the correlation matrix and for generating surrogate data.

Appendix 4: Data acquisition

Coronal slices (400 µm) were prepared from visual cortices of juvenile (P17–22) Wistar rats. Recordings were performed at 37° C in a submersion chamber continuously perfused with oxygenated artificial cerebrospinal fluid containing (in mM) 110 NaCl, 3.75 KCl, 1.25 NaH_2PO_4, 25 $NaHCO_3$, 1 $CaCl_2$, 1 $MgCl_2$, and 17.5 glucose. In certain periods, carbachol (20 or 50 µM), bicuculline (30 µM), CGP-35348 (10 µM) and KCl (10 mM) were added. In addition, intermittent electrical stimulation was applied to layer IV or layer V cells using a 59-electrode array (Multichannel Systems, Reutlingen, Germany) with flat electrodes spaced at 200 µm, integrated in the bottom of the chamber. One to three electrode pairs were selected for weak bipolar stimulation (rectangular pulse, ± 100 or ± 200 µA, 200 µs) of the neuronal tissue at frequencies ranging from 0.5 to 40 Hz using a programmable stimulator (STG 1008, Multichannel Systems).

Spontaneous and evoked activity was recorded with a silicon-based multielectrode array (Bionic Technologies, Salt Lake City, Utah) consisting of 1.5 mm long, sharpened electrodes arranged in a regularly spaced (tip distance 400 µm) matrix. The recording electrode tips (2 µm diameter, 0.1–0.8 MΩ impedance at 1 kHz) were placed in the middle of the slice and covered an area of 3.2 mm horizontally and 1.2 mm vertically, including all six cortical layers. Single-unit spiking activity was extracted from 32 sampled channels by offline sorting. The analysis encompasses data obtained from 10 animals comprising between ~ 40000 and ~ 270000 spikes recorded simultaneously from 56 to 125 units in 40 to 120 minutes.

Chapter III

Organization of cell assemblies in the neocortex

Introduction

Information processing in the cerebral cortex involves the coordination of distributed neuronal activity on a fine temporal scale[91,204,236,405,436,505]. The intrinsic spatiotemporal organization of these coordinated activity patterns as well as the neural mechanisms responsible for their coordination are subjects of intense investigation[47,92,155,198,389,457]. Specifically, the temporal coordination of distributed responses may be assisted by network oscillations that tend to adjust the timing of action potentials and potentially link neurons into synchronous assemblies[68,116,142,218,275,521]. However, disentangling the role of network rhythms and of the functional network architecture is complicated by the fact that oscillations are an emergent network property and cannot be selectively suppressed without affecting the excitability of subpopulations of cells, which would fundamentally alter the functional organization of the whole system[65].

One option to study the temporal coordination of distributed neuronal activity in the absence of neural rhythms but without affecting neuronal excitability is to record from brain slices. Whereas in vivo neural networks may display rich oscillatory dynamics, in vitro preparations generally lack network oscillations unless being activated by appropriate pharmacological or electrical stimulation. Until now, approaches to record multineuron activity from cortical slices with single cell resolution relied on calcium imaging at a temporal resolution of typically 0.1 to 1 seconds [86,227,305]. In the only study monitoring the activity of multiple single cells on a millisecond timescale, Mao and colleagues recorded calcium signals corresponding to single action potentials from up to fourteen neighboring neurons in layer 5 of the primary visual cortex of the mouse and found significant correlations at both short (3-11 ms) and long (up to several seconds) delays[311]. Direct simultaneous recordings of the electrical activity of multiple single cells in cortical slices have not yet been systematically explored.

The work presented here fills this gap by using multiple electrodes to simultaneously record the electrical activity of multiple spatially distributed neurons in slices of the visual cortex of the rat. For this purpose, a novel experimental setup has been designed to record spikes from cells at any depth in the slice using a matrix of sharpened silicon electrodes. Spontaneous and electrically evoked activity was recorded in parallel from all six cortical layers of both areas 17 and 18a under different pharmacological conditions.

To give a comprehensive description of the observed spiking activity, a recent method for the detection and evaluation of multineuronal temporal patterns in parallel spike trains is employed that uncovers coordinated activity on arbitrary timescales and investigates if short firing patterns are arranged in coherent sequences (see Chapter II). The precision and spatiotemporal patterning of concerted neuronal firing is characterized by systematically searching for significant coactivations on timescales ranging from 0.5 to 50 milliseconds. It is shown that even in the absence of oscillatory timing and external stimulation, cortical spiking activity tends to be precisely synchronized with the composition of the synchronous assemblies depending on the functional state of the network. Moreover, assemblies activated in direct succession are organized to a significant degree into repeating sequences, revealing some superordinate temporal structure beyond the level of individual cell assemblies. These findings provide new insights into the intrinsic dynamics of the cortical circuitry and might have profound implications for cortical coding in general.

Methods

Slice preparation

Juvenile Wistar rats aged 17 to 22 days were quickly decapitated under isoflurane anesthesia and their brains carefully removed and placed in ice-cold low-calcium artificial cerebrospinal fluid (ACSF) containing (in mM) 100 NaCl, 2.5 KCl, 1.25 NaH_2PO_4, 25 $NaHCO_3$, 0.5 $CaCl_2$, 6.5 $MgCl_2$, and 25 glucose, saturated with carbogen (5% CO_2, 95% O_2). Coronal slices (400 µm) of the visual cortex were cut on an HR2 vibratome (Sigmann Elektronik, Heidelberg, Germany) and incubated at 37° C for 30 to 60 minutes before being transferred to the recording stage. Recordings were performed at 37° C in a submersion chamber where slices were continuously superfused at 5 ml/min with ACSF

of slightly different composition (in mM): 110 NaCl, 3.75 KCl, 1.25 $NaH2PO_4$, 25 $NaHCO_3$, 1 $CaCl_2$, 1 $MgCl_2$, and 17.5 glucose, also saturated with carbogen (5% CO_2, 95% O_2). Osmolarity was confirmed to be 290 mosmol l^{-1} and pH to be 7.4 in both solutions.

Stimulation

The first experimental condition was characterized by the absence of neuromodulators in the solution. To induce a distinct functional state of the network, the cholinergic agonist carbachol (CCH) was added to the ACSF (20 or 50 µM) in certain periods. The resulting activation of muscarinic receptors is known to change the dynamics in cortical circuits by attenuating excitatory as well as inhibitory synaptic transmission while increasing the excitability of some neurons and hyperpolarizing others[260,261,345,544]. To examine the role of inhibition in coordinating distributed discharges, a (presumably) complete blockade of $GABA_A$ and $GABA_B$ receptors was established by concurrent application of bicuculline (BIC, 30 µM) and CGP-35348 (CGP, 10 µM) during the final period of some recordings (n = 8). At the same time, the concentration of KCl was elevated to 10 mM to facilitate depolarization of cells. All drugs were obtained from Sigma-Aldrich (Taufkirchen, Germany).

To increase activity levels, electrical stimulation was applied intermittently during all three pharmacological conditions to layer IV (8 recordings) or layer V cells (2 recordings) using a 59-electrode array with flat, round electrodes spaced at 200 µm (Multichannel Systems, Reutlingen, Germany), integrated in the bottom of the chamber. One to three electrode pairs were selected for single-site or sequential multi-site bipolar stimulation of the neuronal tissue. Rectangular biphasic current pulses (± 100 or ± 200 µA, 200 µs) were delivered at frequencies ranging from 0.5 to 40 Hz using a programmable stimulator (STG 1008, Multichannel Systems). Since activity was recorded from cells located approximately in the middle of the slice, they were unlikely to be directly activated by the current[359], but likely received mono- and polysynaptic input from cells located more towards the surface in response to a stimulus (see the Supplement for additional arguments on this issue). Electrode positions were registered with a C2400-77 CCD camera (Hamamatsu, Hamamatsu City, Japan), mounted on an inverted Axiovert 35 microscope (Zeiss, Oberkochen, Germany).

Recording

Spontaneous and evoked activity was recorded with a silicon-based multielectrode array (Bionic Technologies, Salt Lake City, Utah) consisting of 1.5 mm long, sharpened electrodes arranged in a regularly spaced (tip distance 400 µm) matrix[357]. The recording electrode tips (2 µm diameter, 0.1–0.8 MΩ impedance at 1 kHz) were gently lowered down into the tissue and positioned at a depth of ~ 200 µm, covering a region of 3.2 mm horizontally and 1.2 mm vertically, including all six cortical layers of areas 17 and 18a[376].

Spiking activity of several neurons and the local field potential (LFP) were simultaneously obtained from 32 channels by amplifying (10000 ×) and band-pass filtering (spikes 0.5–3 kHz or 0.5–5 kHz, LFP 0.1–100 Hz) the recorded signals using custom-made preamplifiers and two MCP Plus amplifiers (Alpha Omega Engineering, Nazareth, Israel). The signals were digitized by two acquisition boards (E series, National Instruments, Austin, Texas) at resolutions of 12 bit / 32 kHz (spikes) and 16 bit / 1 kHz (LFP) and stored using a custom-made acquisition software ("SPASS") written in LabVIEW (National Instruments). Spikes were detected by amplitude thresholding and registered as single waveforms (2 ms duration) and corresponding timestamps (peak sample). The threshold was set interactively after online visualization of the waveforms (typically 1–2 standard deviations above noise level). After recording, electrode positions were reproduced by chemically fixing the slices for 30 minutes in ACSF containing 4% paraformaldehyde and subsequent photomicroscopy.

Offline spike sorting was performed using a dynamic template matching method implemented in a custom software package ("Smart Spike Sorter"). Initially, up to twelve different clusters were automatically defined by an artificial neural network based on the adaptive resonance theory[75]. Various cluster properties like autocorrelations of spike times and recording stabilities of spike waveforms were monitored and considered in conjunction with the shape of the waveforms to guide decisions about which clusters to merge or delete. Only clusters visibly separated in 3D principal component space were assigned to single units. Accuracy of spike assignment was validated by objective measurements of cluster separation provided by the J3 and Pseudo-F statistics[536]. Based on these criteria, only well-isolated putative single units were considered for further analysis.

Analyses

All analyses were carried out using the *MATHEMATICA* system (Wolfram Research, Champaign, Illinois).

Local field potentials

The acquired local field potentials are a measure of combined electrical activity within a volume of neural tissue that approximately includes sources at distances of up to 250 µm from the electrode[256]. As such, the LFP signals display local network rhythms and oscillations, if present. To check whether the firing of single units is embedded in some oscillatory population pattern, spike-centered segments comprising 256 samples were extracted from the LFP traces (Fig. III-1a), multiplied with a Kaiser window, and decomposed into their spectral components using the discrete Fourier transform,

$$\tilde{x}_k = \frac{1}{\sqrt{N}} \sum_{n=1}^{N} e^{2\pi i(n-1)k/N} x_n \omega_n ,$$

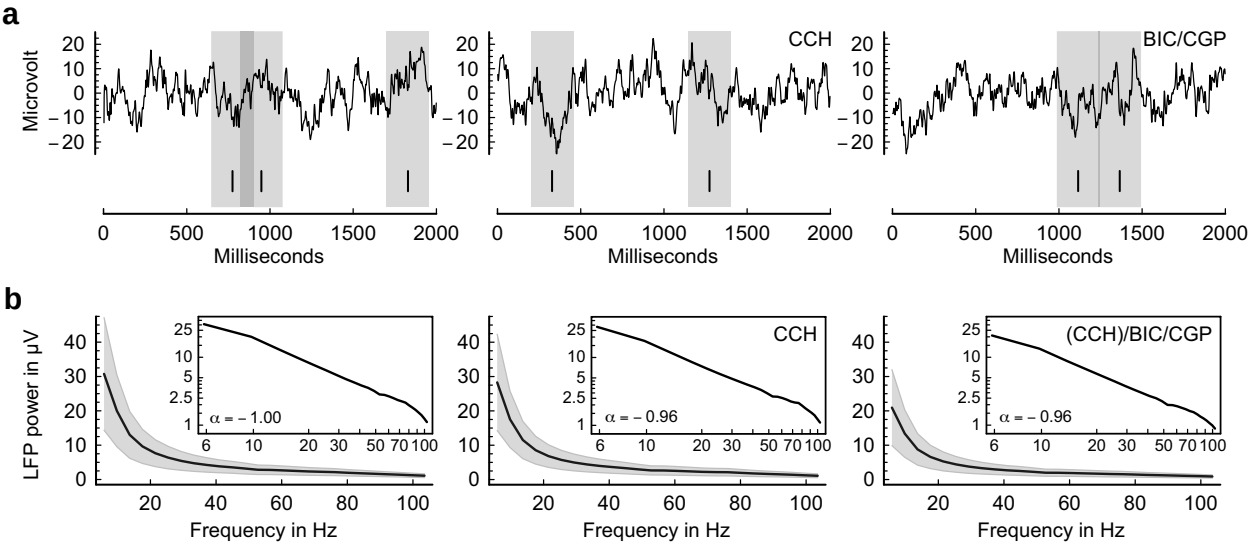

Fig. III-1 Local field potentials during spontaneous activity in the absence of neuromodulators, with carbachol (CCH) added to the ACSF, and with bicuculline (BIC) and CGP-35348 (CGP) added to the ACSF. (**a**) Example traces of local field potentials. Spike-centered windows comprising 256 samples were defined to select segments for analysis as exemplified by the shaded regions (vertical lines indicate spikes). (**b**) Average spike-triggered power spectral densities of local field potentials (mean ± MD across 25940, 121670 and 43878 segments, respectively). Insets show power law scaling with scaling exponent α.

with $k = 0,\ldots, N-1$, N being the number of samples, $x_{1\ldots N}$ being the corresponding voltage values, and $\omega_{1\ldots N}$ being a Kaiser window with shape parameter $\beta = 8$. Accordingly, the complex modulus of $\tilde{x}_k$ denotes the power of the signal per frequency bin k. Power spectral densities were computed for every segment surrounding a spike recorded at the respective electrode that participates in statistically significant sequences of spatiotemporal firing patterns.

Spiking activity

The parallel spike trains obtained through spike sorting were scanned for repeating spatiotemporal firing patterns and sequences of patterns as described in detail elsewhere (see Chapter II). Briefly, patterns were detected by capturing the first spikes of all units within a sliding time window W (with $W = 5$, 10, 15, 20, 25, 30, 35, 40, 45 or 50 ms) and specifying each spike's position in time either by its rank ($\tau = W$) or by the index of the respective sub-window that the spike falls into ($\tau = W/5$ or $W/10$). If applicable, patterns were split into subpatterns based on the joint firing probabilities of the constituent units in time intervals of $T = 60$ s, requiring at least $A = 2$ coincidences per interval to classify them as being functionally coupled or uncoupled. Sequences of temporally non-overlapping, directly consecutive patterns were registered up to a length of 10.

To estimate the significance of the detected patterns and sequences, their individual numbers of occurrences were used as a test statistic and compared to the counts that were to be expected under the null hypothesis of independent firing. For this purpose, 100 surrogate data sets were constructed by randomly and independently shifting the spike trains relative to each other within a time window of ± 10 ms and randomly rearranging consecutive intra-train inter-spike intervals that are shorter than or equal to 10 ms, which resulted in an average displacement of spike events of 5-9 ms. In doing so, the timing of spikes was randomized up to this timescale while rate modulations on slower timescales as well as the interval distributions were left unchanged. Thus, no assumptions about the statistical properties of the data were implied. Likewise, pattern sequences were resampled 100 times by randomizing the order of patterns within intervals of $T = 60$ s, in this way eliminating any potential dependencies between consecutive patterns while approximately preserving each pattern's rate profile. The null hypothesis was rejected if the original count exceeded the counts from the resampled data in at least 95% of the cases. After evaluation of their individual statistical significance, insignificant patterns and sequences were discarded.

In a second step, coordination of spikes was assessed on a global level. To do so, all repeating patterns and sequences from resampled data were tested for significance as well, capitalizing on the assumption that the generation of the surrogate data did not affect the statistic under investigation and that all data sets, including the original one, are essentially indistinguishable regarding the patterning of events. As a result, every data set could be characterized by a certain number X of patterns or sequences that recurred unexpectedly often, given their frequencies in the remaining data sets. Finally, their combined occurrences

$$N_d = \sum_{i=1}^{X_d} n_{id}$$

(with n_i being the number of occurrences of the ith pattern or sequence that is statistically significant on the individual level and d being the index of the respective data set) were subjected to a second level analysis to evaluate the overall significance of a coordination of firing events in the original data.

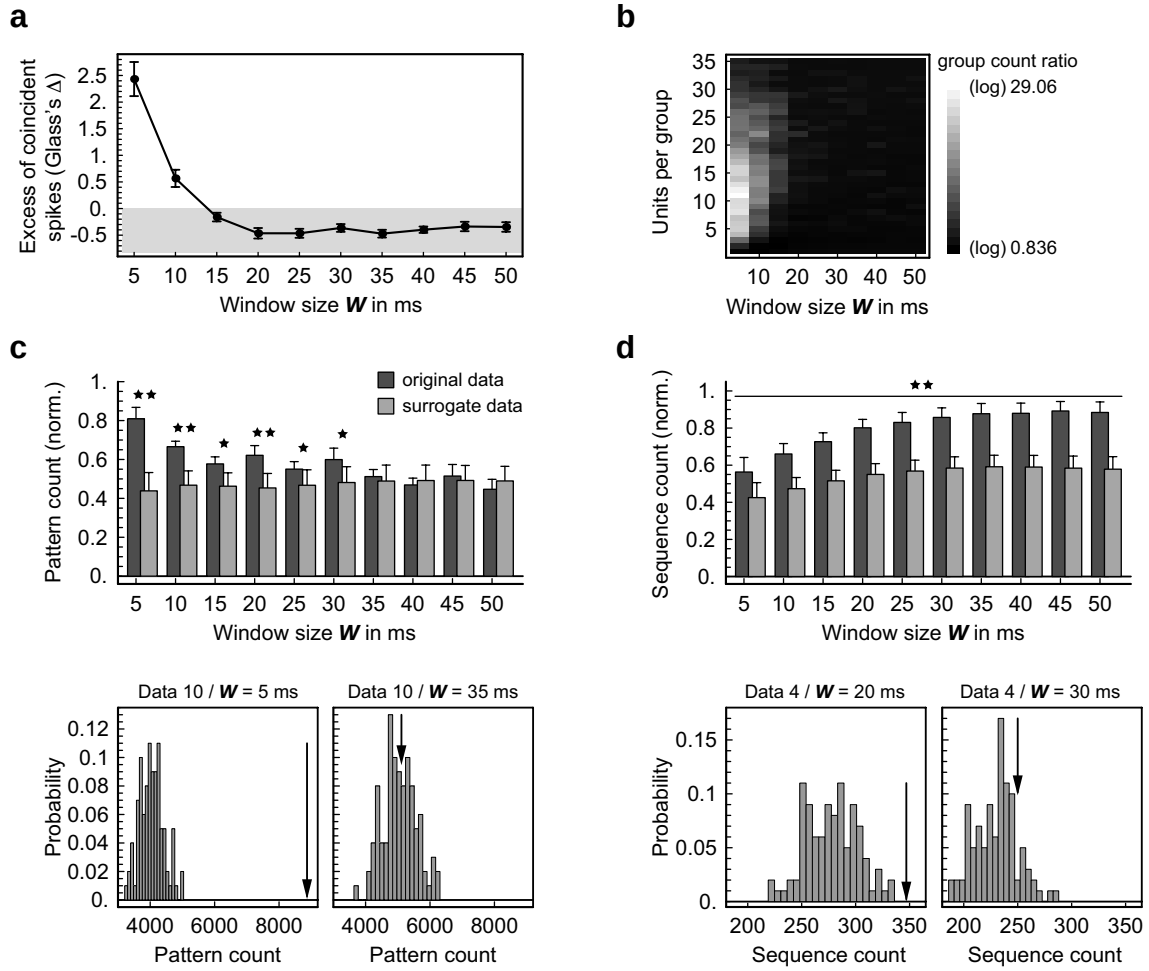

Fig. III-2 Indications of a precise coordination of spike events. (**a**) Excess of the number of coincident spikes in the original data as compared to the surrogate data, expressed as the standardized mean difference (Glass's Δ, mean ± SEM across 10 data sets). (**b**) In the course of detecting spatiotemporal firing patterns, groups of coincident spikes were collected with a sliding window. Their frequency of occurrence, depending on complexity and window size, is shown on a logarithmic grayscale as the ratio of the original count compared to the average number obtained from the surrogate data (mean across 10 data sets). (**c**) and (**d**) Normalized numbers N of significant spatiotemporal firing patterns and pattern sequences (mean + SEM across 10 data sets). Patterns were defined by the temporal order of the participating units within a given time window W ($\tau = W$). Stars denote a significant difference between original and surrogate data (★ $p < 0.0005$, ★★ $p < 0.0000002$, exact binomial test). The lower row shows examples of distributions of surrogate counts along with the corresponding original count (indicated by the arrow).

If $N_{original}$ exceeded at least 95% of the counts from the resampled data, all patterns and sequences that were significant on the individual level were regarded as genuine signatures of concerted neuronal firing and were subsequently analyzed with respect to their spatial and temporal organization and their dependence on electrical stimulation and the neuromodulatory state. In particular, patterns and sequences were classified as belonging to one of four groups according to their spatial structure, namely "clustered" (arranged at a single recording site), "columns" (oriented perpendicular to the pia), "layers" (oriented parallel to the pia) or "dispersed" (no orientation). Chance distributions of group memberships and of the spatial extent of patterns and sequences, measured as the distance between the two most remote units, were calculated by randomly reassigning recording sites or recombining patterns (preserving the number of spikes per electrode and the number of patterns per sequence), repeated ten times. To quantify the temporal precision of firing, a template was computed for every pattern that had been defined by the temporal order of spike events (using $\tau = W$) representing their average timing with respect to the onset of the pattern. The precision of individual spikes was determined as the absolute deviation from their respective average timing. The chance level was estimated by randomly relocating spike times within the time window W that was used for capturing the pattern and recalculating the precision of events, repeated ten times. The accuracy of pattern onsets in sequences of patterns was analyzed analogously. To investigate whether patterns and pattern sequences reflect distinct functional states, they were related to the experimental conditions and said to be selective if their specificity for a particular condition exceeded 50%, i.e., if more than 50% of the occurrences of a single pattern or sequence correlated with a particular pharmacological setting and at the same time with the presence or absence of electrical stimulation. The expected specificity was assessed by repeatedly (10 ×) distributing patterns and sequences randomly over the conditions while maintaining their total number per condition.

Spatiotemporal firing patterns

τ \ W	5 ms	10 ms	15 ms	20 ms	25 ms	30 ms	35 ms	40 ms	45 ms	50 ms
W	8	7	6	7	5	5	5	4	3	3
W / 5	9	8	8	8	7	7	7	7	6	6
W / 10	8	9	8	9	9	7	7	8	7	6

Pattern sequences

τ \ W	5 ms	10 ms	15 ms	20 ms	25 ms	30 ms	35 ms	40 ms	45 ms	50 ms
W	8	9	9	9	8	8	8	7	9	8
W / 5	8	8	9	8	8	8	9	7	8	8
W / 10	8	8	8	6	6	6	8	6	6	8

Tab. III-1 Data sets containing a significant number *N* of firing patterns or pattern sequences depending on pattern parameters *W* and τ.

Results

The analysis encompasses data obtained from 10 slices / animals comprising between ~ 40 000 and ~ 270 000 spikes per set that were simultaneously recorded from 56 to 125 units (Σ = 851) over time intervals of 40 to 120 minutes. The average spontaneous firing rate per neuron was typically below 1 Hz and increased from ~ 0.14 Hz (SD ~ 0.17) to ~ 0.31 Hz (SD ~ 0.28) in the presence of carbachol. Interestingly, application of the $GABA_{A/B}$ receptor antagonists bicuculline and CGP-35348 initially raised activity levels as expected, but did not lead to a sustained increase in average firing rate (~ 0.28 Hz, SD ~ 0.26) above that obtained with carbachol. Electrical stimulation evoked a brief, transient response with latencies of ~ 4 ms that was followed by an increase in average population activity by about 0.2 spikes/ms (with 100 µA stimuli) or 0.4 spikes/ms (with 200 µA stimuli) in the interval of 5-15 ms post stimulus and returned to baseline 20-40 ms after the stimulus (Fig. III-9). Thus, the recorded responses were typically elicited by neuronal signals transmitted via mono- and polysynaptic pathways and reflect a short, single wave of activity traveling through the network (see Supplement – Correlation of spike events with electrical stimuli). Analysis of individual and average power spectra of local field potentials showed no signs of rhythmic population activity in any of the examined conditions (Fig. III-1). The activity under study is thus characterized by sparse spiking in a non-oscillatory regime.

Evidence for precise coordination of neuronal discharges

To test whether the recorded cells fire in a coordinated fashion, the spatiotemporal patterning of firing events was compared to surrogate data that had been derived from the original data by randomly displacing spike times by an average amount of 5-9 ms. The surrogate data represent the null hypothesis that spikes are independent and coordination of spikes is random on timescales below ~ 5 ms. If this was the case, then varying the timing of events on that scale should not affect any statistic extracted from the parallel spike trains.

However, as it turned out, it does. To get a glimpse of the temporal structure inherent in the multineuronal firing activity, the overall number of spikes that coincide with any other spike within a given time window was computed as a rough measure of the temporal clustering of events. The result shows that the original data contain evidently more spikes that are coincident to within a few milliseconds than the surrogate data, giving a first indication that randomizing spike times on the timescale of milliseconds considerably disrupts the temporal coordination of spike events (Fig. III-2a).

In a next step, the particular composition of the groups of coincident spike events was analyzed first in terms of the mere number of the participating units and second with regard to the specific spatiotemporal firing pattern. To do so, activity patterns were registered by capturing coincident events with a sliding window and recording the position of each spike in time with some preset precision (if not stated otherwise, results were obtained using thirty different pattern definitions given by ten different time windows *W* (5, 10, 15, 20, 25, 30, 35, 40, 45 and 50 ms) and three different relative timing precisions τ (*W*/10, *W*/5 and *W*)). Considering all the patterns thus collected, it appears that the surrogate data do not reflect the original distribution of pattern sizes but show a lack of more complex patterns within time windows shorter than 20 ms (Fig. III-2b), which closely agrees with the previous finding. Moreover, not only global pattern properties like the complexity of joint-spike events, but also the frequency of precisely repeating patterns is affected by disturbing the fine temporal structure in the original data. Searching for firing patterns that recur more often than expected by chance yielded numbers of statistically significant patterns that exceed those obtained from the surrogate data in most cases (see Tab. III-1 and Fig. III-2c for cases where $\tau = W$), providing direct evidence for a precise coordination of neuronal discharges that is susceptible to millisecond variations in spike timing.

Finally, the firing activity was scanned for repeating sequences of circumscribed spike patterns to examine its organization on longer timescales. Using the same range of pattern definitions as before, they were almost always found in excess of chance levels estimated from randomly shuffled data (Fig. III-2d), revealing a higher

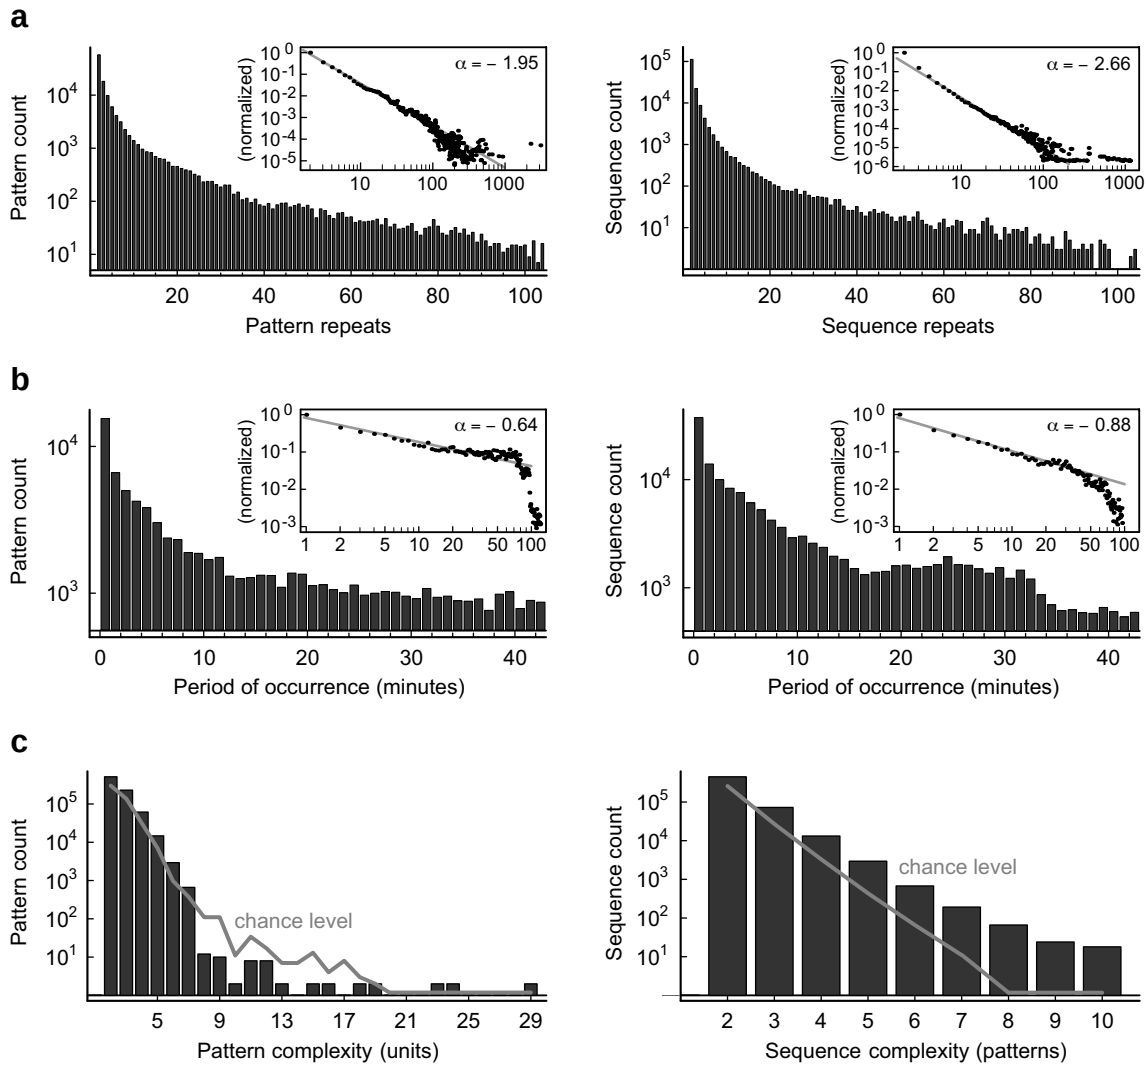

Fig. III-3 Properties of significant spatiotemporal firing patterns and pattern sequences. (**a**) and (**b**) Distributions of repeats and of the periods of occurrence. Insets show normalized distributions averaged across 10 data sets having approximate power law scaling with scaling exponent α (gray lines indicate fitted functions). (**c**) Distributions of complexity in terms of units and patterns, respectively (gray lines indicate chance distributions estimated from resampled data).

temporal structure inherent in the succession of transient activity patterns. The numbers of data sets (out of 10) containing a significant number of repeating firing patterns and pattern sequences are summarized in Tab. III-1 as a function of pattern parameters W and τ.

The distributions of the number of repeats of individual statistically significant patterns and sequences can be approximated by a power law and show mostly low counts while also including cases of more than one hundred repeats (Fig. III-3a). The same is true for the durations over which single patterns or sequences occur, ranging from fractions of a second to more than one hour (Fig. III-3b). This indicates that although most patterns and sequences reoccur only few times and only during short intervals, there is no typical scale associated with the number of repeats and the period of occurrence. Also, linear regression analysis revealed no correlation between repeats and the observed "lifetime" of patterns and sequences ($R^2 \approx 0.07$ and 0.01, respectively), demonstrating that firing patterns may recur after long periods of ongoing activity without requiring frequent replay.

Significant sequences of patterns were found at all lengths tested, but show a marked decrease in overall frequency with increasing complexity. Similarly, the number of significant patterns rapidly drops with increasing number of participating units (Fig. III-3c). Comparing the distributions of pattern complexities between real and resampled data, it appears that – given the number of simultaneously recorded units – patterns including up to seven units are more frequent, whereas larger patterns occur less often than expected. The underrepresentation of significant patterns containing more than seven units suggests that the size of functional cell assemblies in the cortical network is actively limited.

The patterns and sequences identified as significant all violate the null hypothesis of independent firing and therefore have to be regarded as genuine signatures of temporally coordinated activity. In the following these patterns and sequences are analyzed with respect to their spatial and temporal organization and their dependence on electrical stimulation as well as the neuromodulatory state of the slice.

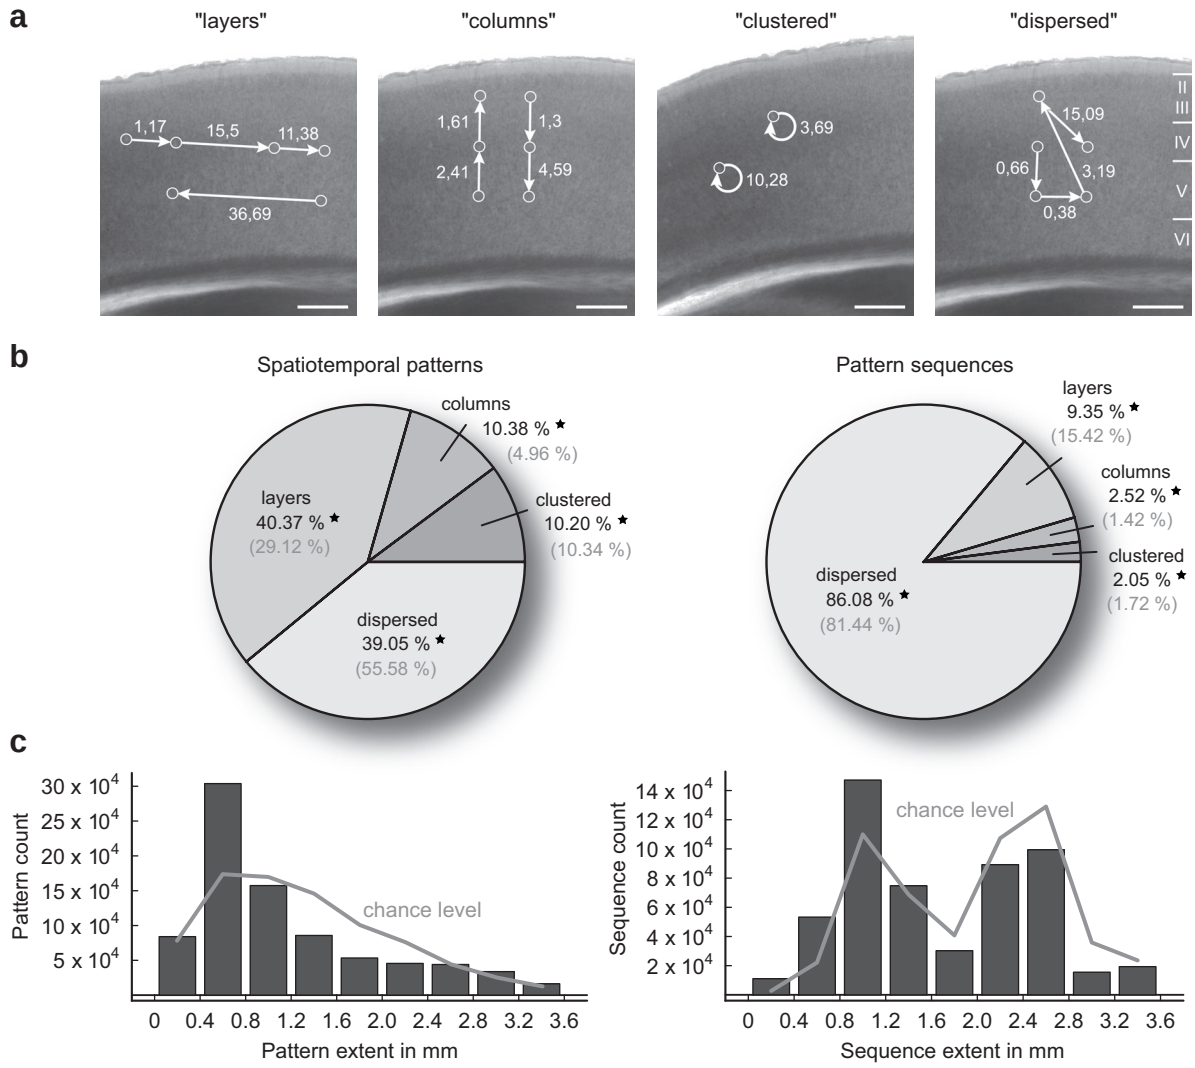

Fig. III-4 Spatial organization of significant spatiotemporal firing patterns and pattern sequences. (**a**) Examples of significant spatiotemporal firing patterns projected onto micrographs of the brain slice representing four different spatial classes. Circles indicate electrode positions with approximate recording horizons, arrows indicate the order of firing, and numbers indicate mean delays in ms (roman numbers on the right denote cortical layers). The slice has been stained with methylene blue to visualize the cortical layering. Scale bar is 400 µm. (**b**) Fractions covered by the four classes exemplified in (**a**) (values in brackets: expected fractions estimated from randomized data) (★ $p < 0.0004$, exact binomial test). (**c**) Distributions of spatial extent (gray lines indicate chance distributions estimated from randomized data).

Spatial organization of multineuronal spiking activity

Coordinated firing patterns and sequences of patterns could be classified according to their spatial structure as being confined to a single recording site, being dispersed both in the vertical and in the horizontal plane, or being strictly aligned perpendicular or parallel to the pia (Fig. III-4a; note that ordered firing does not imply that a signal is transmitted directly between the cells). To investigate whether the topography of statistically significant patterns and sequences reflects the cortical connectivity or is random, they were compared to patterns with randomly reassigned recording sites and sequences with randomly recombined patterns (Fig. III-4b). While the fractions covered by the four classes differed significantly between the original and randomized data (★ $p < 0.0004$, exact binomial test), the spatial configuration of pattern sequences could virtually be reproduced by random combinations of the patterns. Individual firing patterns were confined to vertical columns of cells or to a particular cortical layer much more often than would be expected from a random arrangement of the constituent units, being in line with and possibly reflecting axonal projection patterns[60]. In addition, patterns and sequences primarily extended over shorter distances than calculated from randomized data (Fig. III-4c). Single spatiotemporal patterns were typically confined to directly neighboring recording sites, but could also stretch across the whole array. In contrast, the spatial extent of sequences of patterns shows a distinct bimodal distribution with the two peaks matching the arrangement of functional synaptic connections between areas 17 and 18a[362].

Temporal organization of multineuronal spiking activity

The maximal duration of spatiotemporal firing patterns was limited by the length *W* of the analysis window;

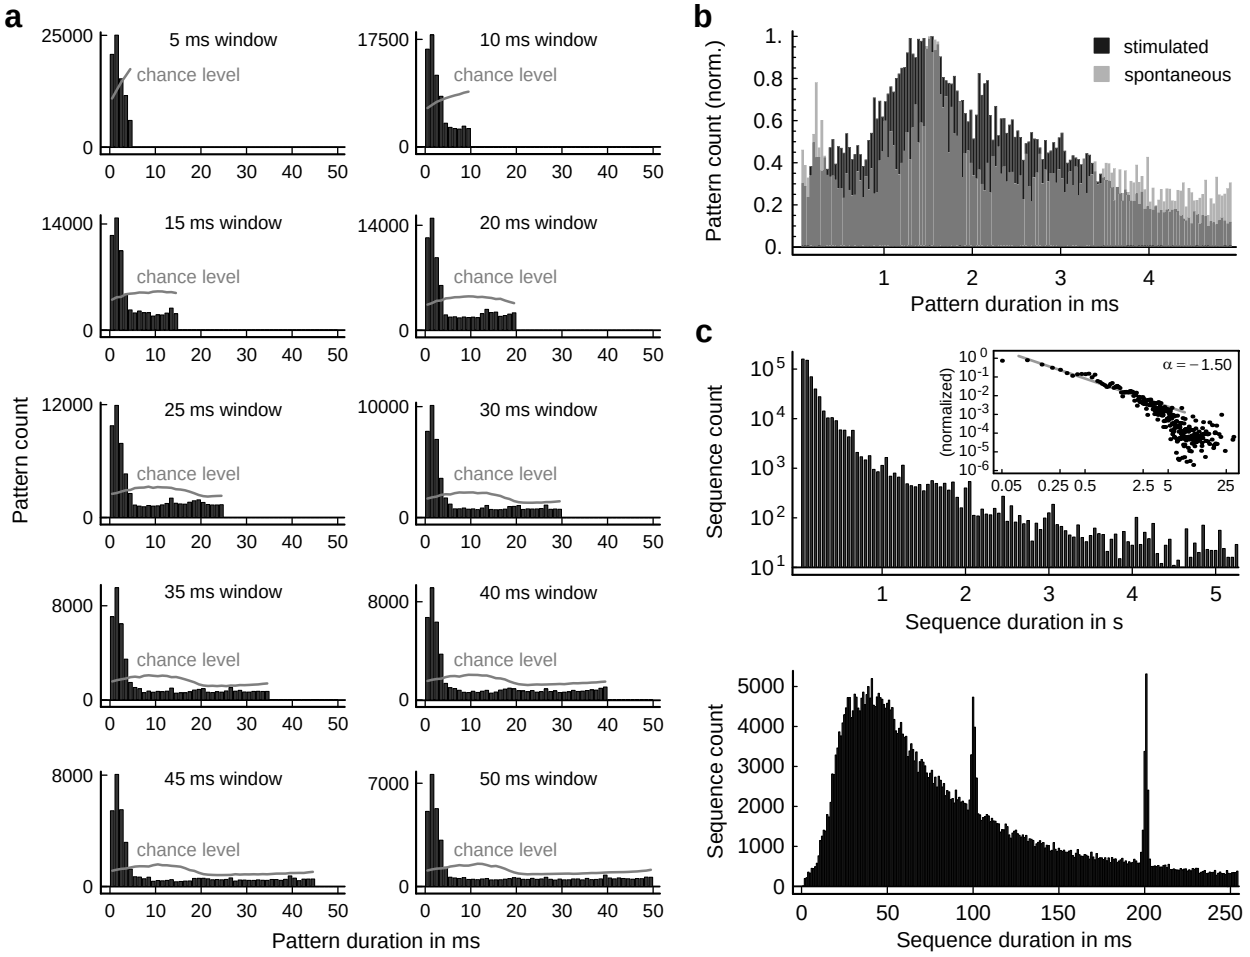

Fig. III-5 Durations of significant spatiotemporal firing patterns and pattern sequences. (**a**) Distributions of pattern durations depending on different time windows (gray lines indicate expected distributions estimated from resampled data). (**b**) Normalized distributions of pattern durations in epochs with and without electrical stimulation. (**c**) Distribution of the durations of sequences of patterns. Inset shows normalized distribution averaged across 10 data sets having approximate power law scaling with scaling exponent α (gray line indicates fitted function).

sequences of patterns, however, were not restricted in time. If neuronal discharges were coordinated on arbitrary timescales, one would expect the distribution of pattern durations to be flat over the entire window that was used to capture the patterns. On the contrary, the temporal extent of significant patterns was typically confined to less than 5 ms, being evenly distributed only on longer timescales (Fig. III-5a). Taking a closer look at the distribution of pattern durations at single sample resolution revealed a prominent peak between 1 and 2 ms (Fig. III-5b). Comparing the pattern durations to those obtained from the resampled data shows that disrupting the coordination of spike events within time windows of ± 10 ms produces patterns that lack any typical temporal scale. Coordinated spiking activity in local cortical circuits thus tends to be synchronized within a few milliseconds even in the absence of input from other areas and of oscillatory timing! Importantly, synchrony does not rely on common drive by some external stimulus, but is maintained also during sparse spontaneous activity (Fig. III-5b).

The integration of these groups of synchronously active neurons into coherent activation sequences gives rise to a coordination of neuronal firing on longer timescales. In the given sample, the distribution of the durations of significant sequences ranged from a few milliseconds to several seconds and peaked at ~ 40 ms (Fig. III-5c). For durations exceeding ~ 50 ms it can be approximated by a power law, suggesting that beyond a particular duration coherent pattern sequences no longer follow a particular scaling rule. The same does not hold, however, if the ongoing activity is modulated by electrical stimulation: Varying sets of cells responding to successive stimuli at short latencies gave rise to repeating sequences of multineuronal firing patterns that were time-locked to the stimulus (reflected in Fig. III-5c, lower part; peaks around 100 ms and 200 ms correspond to stimulus frequencies).

How does the temporal extent of short firing patterns and prolonged pattern sequences relate to their spatial extent? Interestingly, they are not related at all: The duration of both patterns and sequences neither correlates with the distance between the two most remote units nor with their complexity ($R^2 < 0.08$ in all cases, linear regression), meaning that even widely distributed neurons that are unlikely to be directly connected or to receive common input and are not synchronized by network oscillations may align their firing on the timescale of milliseconds. It follows that coordinated firing does not imply any direct communication between the involved cells but may in fact arise from concerted parallel processes.

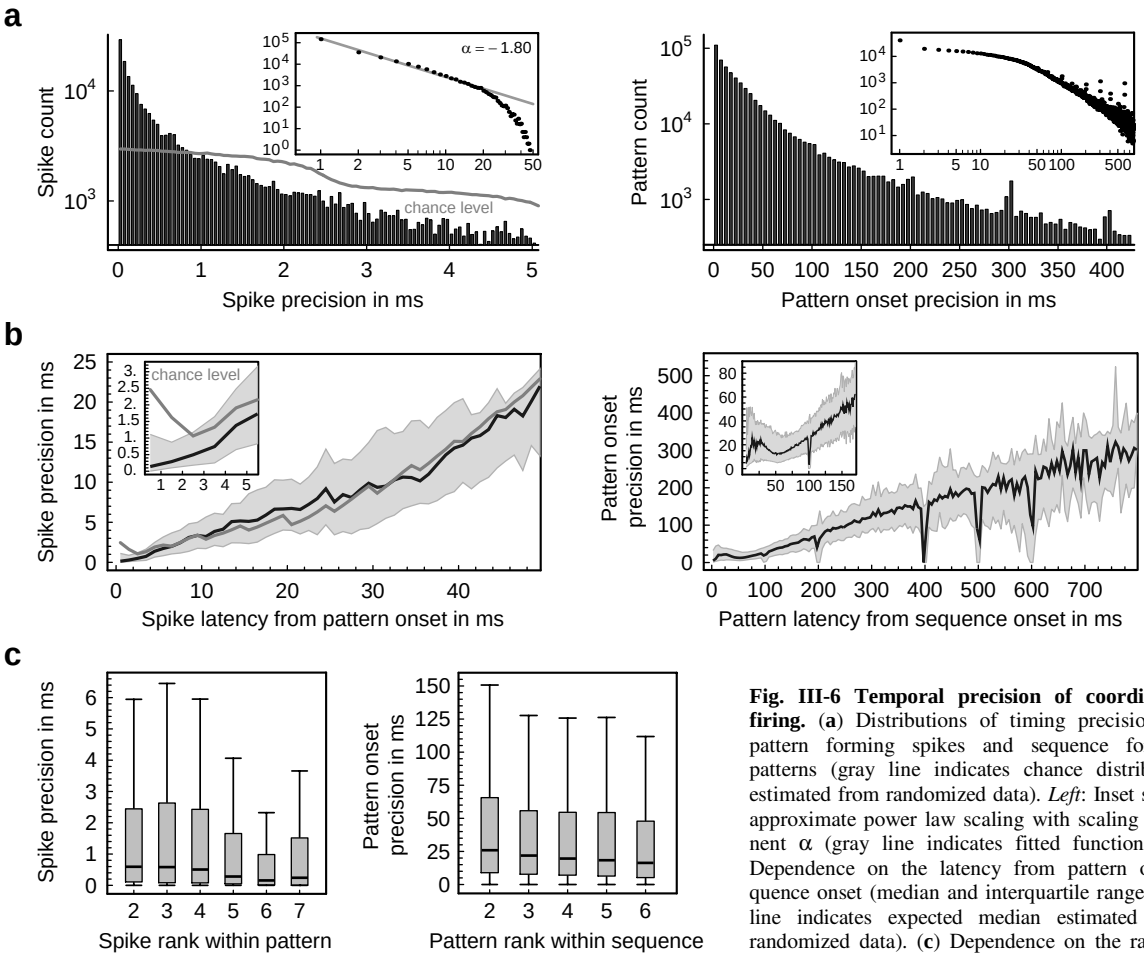

Fig. III-6 Temporal precision of coordinated firing. (**a**) Distributions of timing precisions of pattern forming spikes and sequence forming patterns (gray line indicates chance distribution estimated from randomized data). *Left*: Inset shows approximate power law scaling with scaling exponent α (gray line indicates fitted function). (**b**) Dependence on the latency from pattern or sequence onset (median and interquartile range, gray line indicates expected median estimated from randomized data). (**c**) Dependence on the rank of the event within the pattern or sequence (median and interquartile range, excluding outliers).

In addition to the time span over which neuronal discharges are coordinated, the temporal precision of coordination is likely relevant for cortical processing. To assess this precision, spike times were evaluated relative to a template representing the average timing of spikes with respect to the onset of the pattern (only significant patterns that had been defined using $\tau = W$ were considered). The accuracy of pattern onsets in significant sequences of patterns was analyzed in the same way. It turned out that the distribution of the absolute spike time jitter follows a power law, suggesting that – under the given conditions – the temporal precision of coordinated firing is not constrained to any typical scale. Interestingly, the median jitter was ~ 0.58 ms. By defining spike timing precision with respect to the concurrent spiking activity, this result reveals an accuracy of concerted neuronal signaling that exceeds the duration of a single spike! Moreover, disrupting the original spike timing by randomly relocating events within the time window W used to capture the patterns yields a completely different distribution of timing jitter containing much less events with sub-millisecond precision (Fig. III-6a), suggesting that the original distribution arises from non-randomly placed spikes.

The temporal precision of pattern onsets differed from the precision of spikes within a pattern. While many pattern onsets were precise within a few milliseconds, the distribution of onset jitters includes deviations from the average timing of several hundred milliseconds and shows a median absolute value of ~ 25.13 ms. Remarkably, patterns triggered by electrical stimulation seemed to sometimes skip stimuli and still form repeating sequences, demonstrating that previous activity patterns may promote the emergence of certain subsequent activity patterns without dictating their precise occurrence in time (Fig. III-6a, right part; local maxima around multiples of 100 ms reflect timing of electrical stimuli). Thus, the observed succession of multineuronal firing patterns cannot be accounted for by the concept of synfire chains, in which the propagation delays between nodes are supposed to have a temporal jitter of less than a millisecond[3]. In fact, the precision of pattern onsets depends linearly on the latency from the start of the respective sequence with the median jitter increasing from ~ 12 ms after 50 ms to ~ 300 ms after 700 ms. The linear relationship is violated only at latencies below ~ 50 ms, which could be an edge effect linked to activity rates, and when cells respond to successive electrical stimuli (Fig. III-6b, right part; local minima indicate stimulus intervals). Similarly, the spike timing precision in repeating firing patterns depends approximately linearly on the latency from pattern onset, with the median jitter increasing from ~ 0.15 ms within the first millisecond to ~ 21 ms after nearly 50 ms. This measured precision deviates from the precision expected from random spike times within

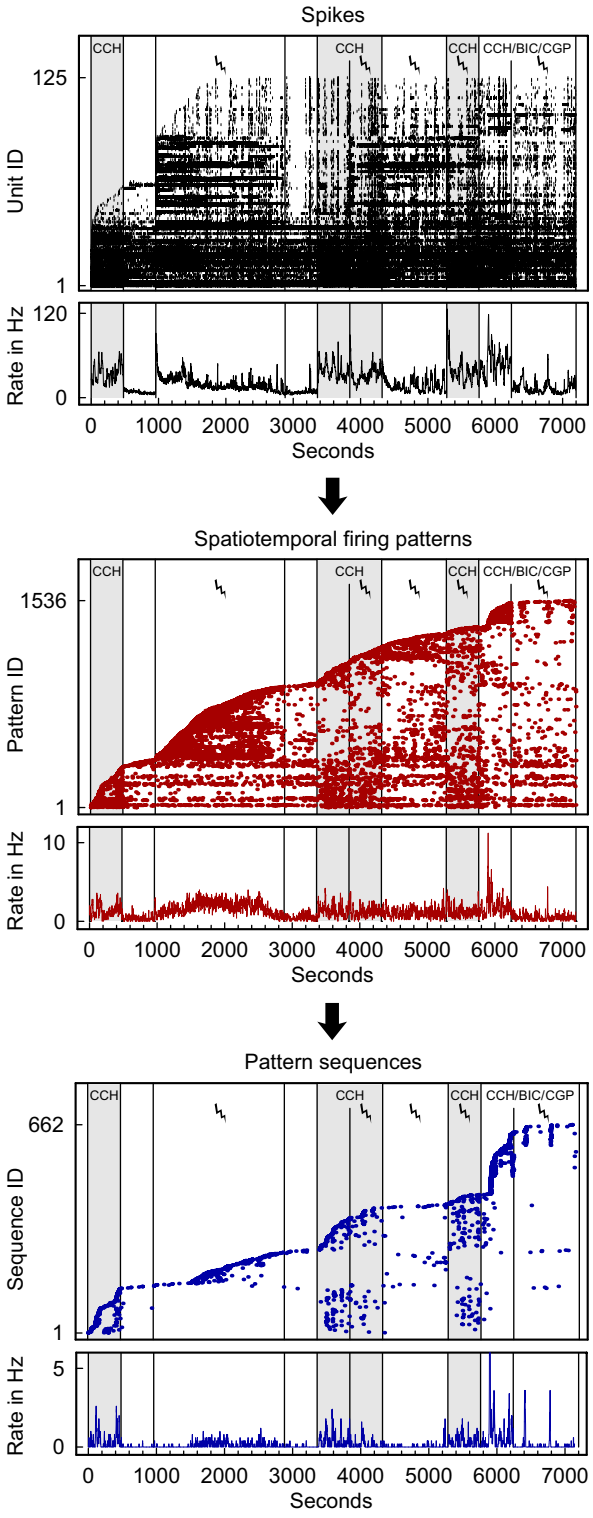

Fig. III-7 State dependence of spatiotemporal firing patterns and pattern sequences. Illustrated are spike times and times of occurrences of significant firing patterns and pattern sequences along with the corresponding population rates from one example data set. Units, patterns and sequences are sorted by appearance. Experimental conditions included electrical stimulation ( ) and application of carbachol (CCH), bicuculline (BIC) and CGP-35348 (CGP) as indicated. Patterns were captured using $W = 30$ ms and $\tau = 3$ ms.

windows W only in about the first three milliseconds after pattern onset (Fig. III-6b, left part), supporting the notion that precise synchronization of neuronal discharges is a significant feature of cortical activity (cf. Fig. III-5a, b), and that exact spike timing is rather irrelevant on longer timescales compared to spike order. At the same time, the average temporal jitter did *not* accumulate over successive spikes within a pattern or consecutive patterns within a sequence, which fits with the finding that the duration and the complexity of firing patterns are not related (Fig. III-6c; note that the first element never displays any timing jitter by definition).

State dependence of recurring multineuronal spike patterns

If precisely repeating spatiotemporal spike patterns as well as repetitive sequences of these patterns reflect the functional architecture of the neuronal network, they are expected to change when the functional state of the network is altered by application of neuromodulatory transmitters. An important modulatory transmitter in the cortex is acetylcholine, which is involved in the control of arousal and attention[211] and known to radically shift intracortical signaling pathways[260,544]. To investigate the dependence of significant firing patterns and pattern sequences on cholinergic modulation, they were related to the application of the cholinergic agonist carbachol. It turned out that individual patterns and sequences had a high tendency to selectively occur either in the absence or in the presence of carbachol (Fig. III-7, shaded periods).

To examine the role of inhibition in the structuring of firing patterns, GABAergic transmission was blocked during the final period of 8 experiments by application of bicuculline and CGP-35348. The effect was that most but not all of the patterns and sequences that had occurred before disappeared. In return, a number of new spatiotemporal patterns and pattern sequences came up (Fig. III-7). Relative spike timing was still precise down to the sub-millisecond scale (median jitter ~ 0.67 ms), but unlike patterns that also occurred during intact inhibition, those firing patterns that were restricted to the disinhibited condition exhibited a loss in synchrony of coordinated spikes. It thus seems that the excitatory cortical network is capable to produce and maintain very precise multineuronal activity patterns during disrupted or diminished GABAergic signaling. These patterns, however, reflect a functional network architecture that differs fundamentally from the natural situation, so they carry no biological meaning; interestingly, this lack of meaning goes along with the lack of a typical timescale over which discharges are coordinated.

Specific spatiotemporal spike patterns and sequences of patterns could also be induced by focal electrical stimulation of the neuronal tissue (Fig. III-7; note the divergence of spike and pattern rates during prolonged stimulation). Considering the neuromodulatory state as well as the presence or absence of electrical stimuli, ~ 73% of all significant patterns and ~ 72% of all significant pattern sequences selectively occurred under a certain experimental condition (chance levels: ~ 51% and ~ 44%, $p < 0.002$ in both cases, exact binomial test), showing an average specificity of ~ 90% and ~

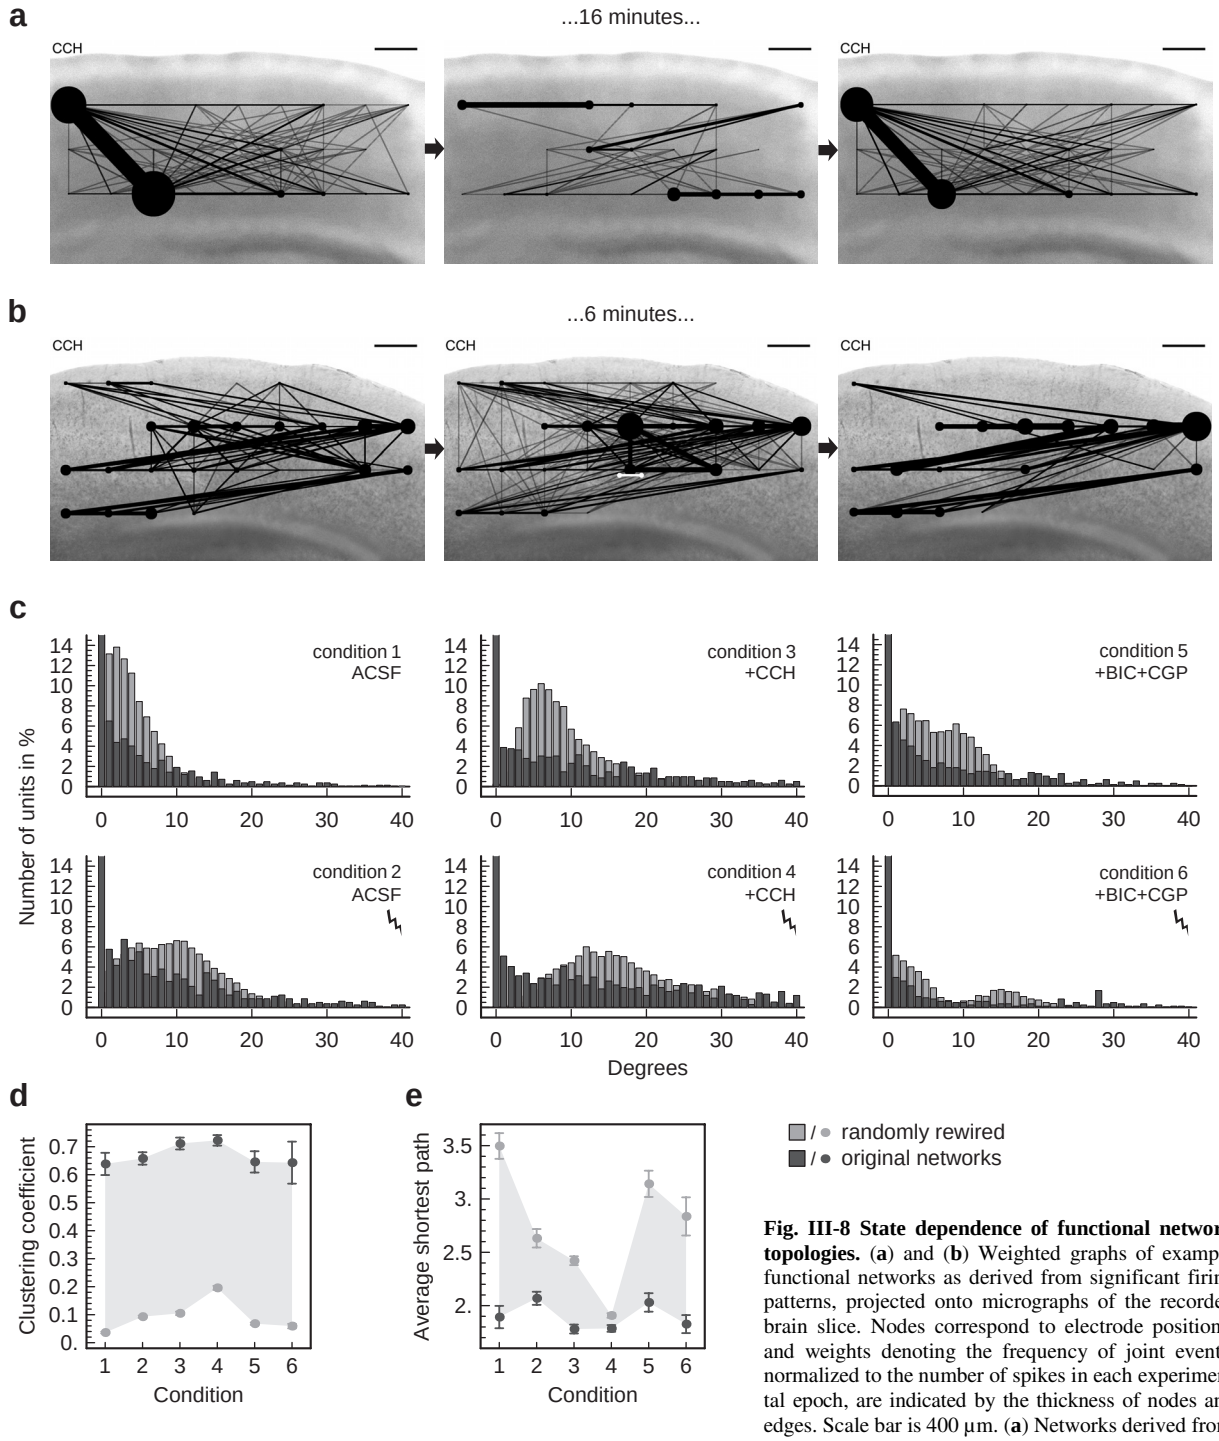

Fig. III-8 State dependence of functional network topologies. (**a**) and (**b**) Weighted graphs of example functional networks as derived from significant firing patterns, projected onto micrographs of the recorded brain slice. Nodes correspond to electrode positions, and weights denoting the frequency of joint events, normalized to the number of spikes in each experimental epoch, are indicated by the thickness of nodes and edges. Scale bar is 400 µm. (**a**) Networks derived from spontaneous activity in three consecutive epochs depending on bath application of carbachol (CCH). (**b**) Networks depending on focal electrical stimulation (white dots indicate positions of stimulus electrodes). (**c**), (**d**) and (**e**) Properties of functional neuronal networks derived from significant firing patterns during different experimental conditions including electrical stimulation (ϟ) and application of carbachol (CCH), bicuculline (BIC) and CGP-35348 (CGP) as indicated. (**c**) Grand degree distributions (zero degrees bin truncated). (**d**) Clustering coefficients (mean ± SEM across data sets). (**e**) Average path lengths (mean ± SEM across data sets).

95%, respectively (chance levels: ~ 82% and ~ 91%, $p < 0.002$ in both cases, exact binomial test). Although plasticity mechanisms acting on functional neuronal connections continually reshape multineuronal activity (Figs. III-7, III-3b), ~ 35% of all selective firing patterns and ~ 24% of all selective pattern sequences recurred after experimental interference and may hence be regarded as the "fingerprints" of a particular condition.

In addition to the sensitivity of individual activity patterns to the state of the network, the influence of electrical stimulation and of the pharmacological setting on the topology of functional neuronal networks was investigated. To do so, edges were drawn between electrodes or units that took part in the same firing pattern during a given experimental epoch or condition (only significant patterns were considered). Like the activity patterns themselves, the resulting functional connectivity maps depended systematically on the action of carbachol or GABA and focal stimulation of the neuronal tissue (Fig. III-8a, b). Interestingly, they

exhibit a small-world architecture: First, the degree distributions reveal many sparsely connected neurons and a long tail of increasingly interconnected units, suggesting the existence of "hubs" in the network. In contrast, randomly rewired networks having the same number of nodes and edges yield much less skewed degree distributions as they are typical for random networks (Fig. III-8c; bimodal random distributions obtained under suppressed inhibition result from the merging of data sets). Furthermore, the clustering coefficients[530] range roughly from 0.43 to 0.82, clearly deviating from the values that would be expected by chance (0 to 0.3) (Fig. III-8d). Finally, the average path lengths fluctuate closely around 1.9, a number notably below those derived from randomized data (Fig. III-8e). However, neither a massive change in cholinergic or GABAergic modulation nor interference of the ongoing activity with electrical stimulation had any substantial effect on these network parameters, although spike rates and firing patterns changed considerably following drug application or washout and in response to electrical stimuli (Fig. III-7).

These findings demonstrate a remarkable stability of functional neuronal network properties with respect to diverse modes of synaptic integration, providing extensive freedom to visit a multitude of functional states while seamlessly maintaining a small-world architecture. As a consequence, the ability to rapidly process and integrate information both locally and globally – a highly desirable property for neuronal computations – arises naturally and appears to be an inherent feature of cortical connectivity that is independent of cholinergic and even GABAergic modulation.

Discussion

Cortical information processing invariably involves the activation of networks of distributed neurons. At present, it is unclear how information is integrated and how coherent representational states are established in cortical networks. To study the spatiotemporal structure of distributed discharges in local cortical circuits and to disentangle the role of network rhythms and of the functional network architecture in shaping it, multineural activity was recorded from visual cortices in vitro and analyzed with respect to the occurrence of temporally coordinated spike events on timescales ranging from 0.5 to 50 ms. It was found that population patterns occurring spontaneously and in response to focal electrical stimulation displayed a statistically significant structure which is characterized by synchronous firing within a few milliseconds, while network rhythms were absent. The formation of these synchronous cell assemblies is likely related to the history of network activation and learning and may constrain or bias the recruitment of local cortical circuits during perception and active behavior. What is more, assemblies activated in direct succession were organized to a significant degree into repeating sequences, revealing a superordinate temporal structure beyond the cell assembly level. These sequences, however, do not conform to the concept of synfire chains[3] because of imprecise delays between the successive groups. Instead, they may be described as Hebbian phase sequences[207] composed of loosely coupled cell assemblies[198] that are distinguished by synchronous firing within typically up to five milliseconds.

Previous work in vivo has suggested that cortical neurons can coordinate their firing with millisecond precision in relation to behavior[4,389,436,516]. In these studies, only few cells were recorded simultaneously, and a comprehensive description of the patterning of multineuronal activity could not be provided. Precise spatiotemporal spike patterns also appear in cortical cell cultures[406], suggesting that they are a universal emergent feature of self-organizing neuronal networks. Calcium imaging of small local cell populations in brain slices further revealed that the firing of neighboring cortical neurons may be significantly synchronized within a few milliseconds during spontaneous, non-oscillatory activity[311]. The present work confirms and extends these results to larger neuronal populations that are distributed across cortical areas and layers. Given the fact that the spatial arrangement of synchronous firing patterns found here agrees with axonal projection patterns[60], it seems likely that distributed coactive cells reflect functional cortical subnetworks[76,553].

Progression of transient activation patterns

The organization of successive spatiotemporal firing patterns into coherent recurring sequences suggests an additional superordinate, "modular" temporal structure in cortical spiking activity that resembles similar dynamics on coarser timescales[227]. The basic neural mechanisms that provide neuronal networks with the potential to produce such long sequential activity patterns are yet unknown. A promising candidate for such a mechanism is the classical spike-timing-dependent plasticity rule[72] that has been demonstrated to produce long, diverse activity sequences in simulation experiments even without sequential training inputs[130,263]. In addition, long-range cytoplasmic signaling within the presynaptic neuron may lead to a retrograde propagation of synaptic potentiation along excitatory pathways [484] and thus to the formation of chains of strong and effective synaptic connections. The only experimental finding of circumscribed cortical cell assemblies firing in functional sequences has been reported by Ribeiro and colleagues[15]. They analyzed synchronously active groups of cells recorded simultaneously from the hippocampus and primary visual and somatosensory cortical areas of rats and demonstrated the dependence of cell assembly activation sequences on experimental periods and the behavioral state of the animal. The present results go beyond these findings by showing that individual sequences repeat significantly often and strongly depend on the neuromodulatory state of the cortical network. The observed scale-free distribution of sequence durations might be a natural consequence of competitive spike-timing-dependent synaptic plasticity[72,467] and resulting network connectivities, as simulations suggest[130].

Previous work has focused on the detection of synfire chain activity by searching for precise repetitions of firing patterns with durations of several hundred milliseconds and reported on "reverberations in a synfire mode" based on these patterns[4,389]. While the present analysis considers spike patterns up to a length of 50 ms, it cannot be ruled out that precisely coordinated spike events do occur on longer timescales. However, the detected decrease in temporal precision with increasing latency from the reference spike and imprecise delays between successively active neuronal groups suggest a different mode of network dynamics in which signals reverberating through recurrent connections would provide the network with a "fading memory" of past activity[355]. The succession of neuronal activation patterns could be described as a heteroclinic sequence[11,392] linking transient metastable network states across time and cortical areas in a highly flexible way[58]. Here, transients are defined by synchronous neuronal groups and represent "saddle points" along the heteroclinic channel. Accordingly, the course of activity would be neither chaotic nor evolve into any fixed point or limit cycle, thus combining flexibility and reliability.

Mechanisms of synchronization and consequences on cortical function

Although it is unclear if the synchronous firing of a set of cells actually *codes* for any relevant content, synchronization of discharges is thought to jointly increase the saliency of the synchronized responses[526] and to thereby serve a number of different functions such as the binding of distributed cells into a transient functional unit[447,457], the selection of signals by attention[144] and the mediation of spike-timing-dependent plasticity [72]. Accordingly, disturbance of synchrony has been identified as one of the correlates of brain disorders[501]. Neuronal synchrony is commonly assumed to arise from correlated sensory signals, common input, or from synchronous network oscillations[504]. Neither of these causes applies in the present case: First, there are no sensory signals in brain slices. Second, the majority of synchronously active cells were located too far apart from each other to be directly connected or to receive common input[208,319,323], suggesting that their activity was coordinated through concerted distributed processes. Third, local field potentials showed no sign of network oscillations, but were rather consistent with white noise and the passive filtering properties of brain tissue [37-39] (under the same experimental conditions, intracellular recordings have also shown a lack of subthreshold membrane potential oscillations[535]). This raises the question which other mechanisms could mediate the synchronous firing of distributed cells in local cortical circuits.

Theoretical studies have demonstrated that neurons endowed with spike-timing-dependent plasticity may "learn" to selectively fire at the onset of a repeating spatiotemporal input pattern by potentiating synaptic weights on afferents that consistently fire early[192,321,322]. As a consequence, distributed neurons repeatedly receiving coherent input would decrease the latency of their firing in a concerted fashion and could finally become synchronized when responding to this input[482]. Since synaptic modifications in rat visual cortex are sensitive to the relative timing of pre- and postsynaptic activity[146,335], the observed synchronization of neuronal responses might indeed be caused by a coordinated adjustment of response latencies and thus reflect the reactivation of specific cortical circuits shaped through experience. Being restricted to local and short-range connections within and between neighboring cortical areas, the synchrony generating circuitry could have been triggered either spontaneously or through focal electric stimulation targeted at layer IV or V (Fig. 5b). This would have major consequences for the respective roles of the functional neuronal network and ongoing activity patterns in coordinating multineuronal firing: While the presented data suggest that synchronous firing of distributed cells is inherently generated by the functional cortical network structure, network rhythms and other activity patterns would control the propagation of signals through the circuitry and effectively select cells and synchronous groups of cells to become activated. Cortical computation, then, unfolds in the interplay between neuronal activity and functional neuronal connectivity[141].

Precision of spike timing coordination

If spike synchrony serves any function in information processing, spike timing needs to be coordinated with appropriate precision. The observed precise synchronization of discharges fits with previous findings showing that spike generation in cortical neurons as well as signal transmission in local cortical circuits can be highly reliable and temporally precise[309,350] and adds to the evidence that neuronal communication may rely on millisecond-precise signaling despite stochastic gating of ion channels and probabilistic release of synaptic vesicles[123]. The result that the accuracy of coordinated spikes decreases as the interval increases has been reported previously and may be explained by synaptic variability on the (sub-)millisecond range[343,350] and by network effects on longer timescales[57]. A crucial factor determining the temporal precision of neuronal signal transmission is the effective time window over which postsynaptic potentials are integrated[403]. Short integration times may arise from rapidly deactivating AMPA receptors that can have deactivation time constants of less than a millisecond[212,308] and indirectly control the kinetics of NMDA receptor currents by only allowing for a correspondingly short release of the magnesium block[216,367]. The resultant spike timing may achieve a resolution shorter than the rise time of the excitatory potential if synaptic inputs arrive synchronously[174]. In addition, disynaptic feedforward inhibition can confine the effective integration time window in the soma to a few milliseconds[386]. Taken together, these results show that the possibility of a code based on precise spike timing is at least not precluded by the biophysical substrate. The presented data demonstrate that spikes can be coordinated with sub-millisecond precision[19]

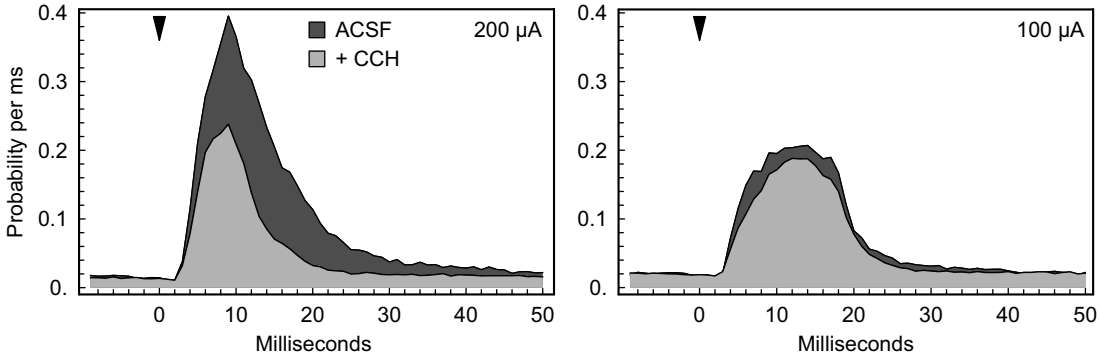

Fig. III-9 Latencies of stimulus responses. Following focal electrical stimulation, the probability of the occurrence of spikes within the range of the recording electrodes transiently increased depending on stimulus intensity (100 or 200 µA) and on muscarinic activation (ACSF alone or with carbachol, CCH) as shown. The arrow indicates the stimulus at time zero. The analysis is based on 25768 and 8788 stimuli (with and without carbachol) at 200 µA and on 38160 and 17040 stimuli (with and without carbachol) at 100 µA.

even in purely excitatory networks lacking effective GABAergic signaling. In this case, activity levels may be regulated by self-inhibition of pyramidal cells[314] and by synaptic depression[512].

Interestingly, the obtained distribution of spike time jitter can be approximated by a power law, which has been regarded as a signature of self-organized criticality[31]. It is tempting to suppose that under the given conditions cortical networks self-organize into a critical state in which the noise affecting the timing of action potentials is tuned to a critical point. Whether this explanation holds, by what mechanisms such a point might be reached, and whether spike time jitter in vivo also follows a power law, remains to be investigated.

Functional implications of the network topology

Albeit conceptually distinct, the functional connections between the recorded cells and anatomical connections between neurons in the rat visual cortex appear to have several features in common: For both types, connected neurons tend to cluster together, and connections are concentrated in strength and number among a subset of cells[468]. The same small-world architecture has also been detected from functional connections in cat visual cortex in vivo[554]. As simulations have shown, small-world connectivity may automatically emerge from the reorganization of synaptic connections through spike-timing-dependent plasticity if the balance between excitation and inhibition is adjusted to a critical level[434], which might explain the intriguing robustness of the small-world property during changes in the functional circuitry through neuromodulation. This property, in turn, is known to support synchronous activity and fast signal propagation[184,271,530] and enhances the tolerance of the network against the failure of single cells[12]. Thus, the arrangement of functional neuronal connections in local cortical circuits not only defines functional states of the network, but also seems optimized for robust and efficient information processing.

Supplement – Correlation of spike events with electrical stimuli

Extracellular electrical stimulation of the cortical gray matter, as applied in this study, is known to effectively excite axonal segments in some volume surrounding the site of stimulation[360,361]. The extent of this volume depends on the strength of the stimulus. In rat visual cortex, the spread of the current has been shown to have a maximal radius of ~ 200 µm and ~ 300 µm, respectively, at currents of 100 µA and 200 µA, with about half of the number of axons being activated at half the maximal radius[359]. Considering the given distances between the recording electrodes and the stimulating electrodes, we thus could not expect many of the recorded cells to be directly activated by the stimulus. Nonetheless, electrical stimulation of fibers en passage may trigger action potentials that travel in both orthodromic and antidromic direction and could be recorded with some delay as far as millimeters away[215]. It is thus necessary to differentiate between signals that could have been evoked by the stimulus and those that could not. Given the latencies of the spike responses (Fig. III-9), the respective distances of the recording electrode tips to the stimulation sites, and the axonal conduction velocities in rat visual cortex of ~ 150-550 µm/ms[346], the average fraction of spikes per data set that could be explained by direct activation through electrical stimuli was calculated to be ~ 1.3 % (SD ~ 1.0). However, if an axon is excited directly by the stimulus current, the resulting spike responses should have a temporal jitter in the submillisecond range. Crosscorrelations of single unit spike events with electrical stimuli revealed no such precise responses, but showed temporal jitter of ~ 8-20 ms. We can thus be sure that the recorded spiking activity was exclusively generated by genuine neuronal signals, with the recorded cells being separated by at least one synaptic stage from the cells that were directly activated by electrical stimulation.

Another line of arguments suggests that the observed spike signals represent real neuronal network activity rather than trivial stimulus responses. To begin with, activation of muscarinic receptors through carbachol is expected to allow a stronger response to afferent input due to a reduction in spike-frequency accommodation

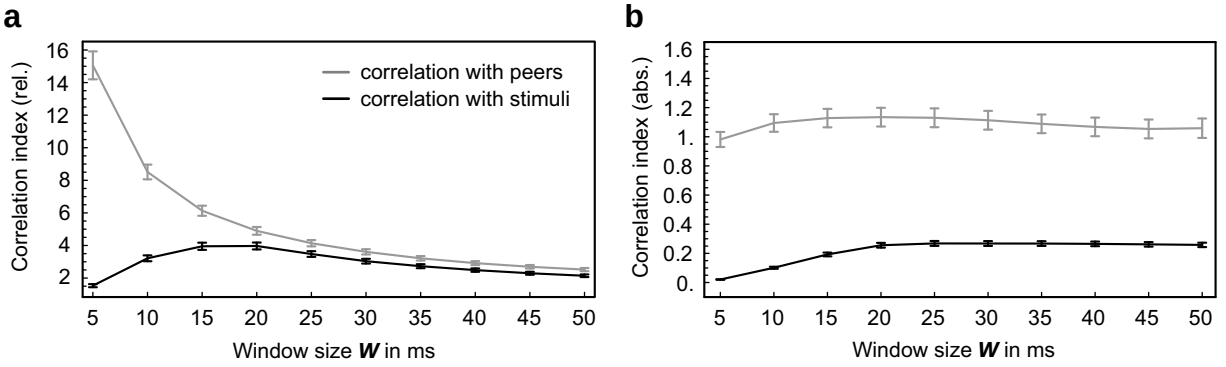

Fig. III-10 Correlation of unit activity with the firing of peers and with electrical stimuli. The number of spikes following electrical stimuli (100 µA) or discharges of other units (peers) within given time windows *W* is compared against the expected number given independence and stationarity of events (mean ± SEM across 358 units). (**a**) Relative correlation of events, expressed as the ratio of the actual number of succeeding spikes compared to the expected number. (**b**) Absolute correlation of events, expressed as the number of preceding events that appear to be related to the occurrence of each spike. (See text for a formal description.)

and inhibitory feedback, but also to reduce recurrent excitatory signaling in rat visual cortex[261,345]. Depending on the involvement of excitatory synaptic transmission in the generation of the recorded spike responses, we should therefore observe a decrease of the response probability when carbachol is applied. This is indeed what was found throughout the recordings (Fig. III-9), indicating that the functional properties of the synaptic connections in the network contribute to the shaping of the recorded spike activity. To quantify the dependence of recorded spike events on electrical stimulation and on the ongoing network activity, the number of spikes following electrical stimuli or discharges of other cells (referred to here as peers) within given time windows was compared against the expected number assuming independence and stationarity of events. The rationale behind is that a deviation of the temporal coherence of events from chance level would be suggestive of some functional relationship between the preceding and the succeeding event. For each recorded unit, a so-called "relative" correlation index was computed according to

$$Corr_j^{(W)} = \frac{1}{M}\sum_{t=1}^{M}\left(\sum_{i=1}^{N} C_{ij}^{(W,t)} \Big/ \sum_{i=1}^{N} P_{ij}^{(W,t)}\right)$$

with $Corr_j^{(W)}$ being the ratio of the real and the expected number of spikes of unit *j* following any spike of any other unit (or any electrical stimulus) with maximal latency *W*, averaged across *M* one-minute time periods, $C_{ij}^{(W,t)}$ being the number of spikes of unit *j* following discharges of unit *i* (or electrical stimuli) with maximal latency *W* in time period *t*, $P_{ij}^{(W,t)}$ being the respective expected number, and *N* being the number of peers (or $N = 1$ in case of a correlation with electrical stimuli). The chance level of the temporal coherence of events, assuming independence and stationarity, was estimated according to

$$P_{ij}^{(W,T)} = \frac{W}{T}\, n_i^{(T)} n_j^{(T)}$$

with $P_{ij}^{(W,T)}$ being the expected number of spike events of unit *j* following events of unit *i* (or electrical stimuli) with maximal latency *W* in time interval *T*, $n_i^{(T)}$ being the number of events of unit *i* (or electrical stimuli) in time interval *T*, and $n_j^{(T)}$ being the number of events of unit *j* in time interval *T* (see Chapter II, Appendix 1 for a derivation). The necessary condition for the equation to hold is that $T/W \geq n_i^{(T)} \wedge T/W \geq n_j^{(T)}$, i.e., the number of the respective events in time interval *T* should not be larger than 1200 if *T* is set to one minute and *W* is set to 50 ms. In the data presented here, this condition was always met. The index is a measure of the degree to which a neuron's firing is coupled to the activity of its peers or to the electrical stimulus. It turned out that spike events strongly depend on the ongoing network activity especially on short timescales (Fig. III-10a), corroborating previous results (Fig. III-2a, b, c). In time windows of 5 ms and in the absence of any externally applied neuromodulators, spike events were preceded by discharges of other cells about 15 times as often as was expected given their rates. With increasing time windows *W*, the relative correlation of spike events decreases, but remains considerably above chance level (which corresponds to $Corr^{(W)} = 1$) even at $W = 50$ ms. In addition, spike events are also clearly correlated with electrical stimuli, but here the relative correlation is strongest in time windows of 15-20 ms (cf. Fig. III-9) and only reaches a maximal mean index of ~ 4 (Fig. III-10a). Most importantly, the average correlation of spikes with preceding stimuli is much lower than with spikes of other units in time windows of 5-15 ms, suggesting that within these windows the ongoing network activity exerted a much stronger influence on the firing of cells than did electrical stimulation.

While the relative correlation index compares the real and the expected number of spikes following electrical stimuli or other events in given time windows, it does not reveal anything about their actual incidence. To specify the actual amount of spike events that are likely to be related to preceding stimuli or discharges of other neurons, an additional "absolute" correlation index was computed according to

$$Cora_j^{(W)} = \frac{1}{M}\sum_{t=1}^{M}\left(\left(\sum_{i=1}^{N} C_{ij}^{(W,t)} - \sum_{i=1}^{N} P_{ij}^{(W,t)}\right) \Big/ n_j^{(t)}\right)$$

with $Cora_j^{(W)}$ denoting the number of preceding spikes (or electrical stimuli) in time windows W that appear to be related to each spike of unit j, averaged across M one-minute time periods, $n_j^{(t)}$ being the overall number of spikes of unit j in time period t, and $C_{ij}^{(W,t)}$, $P_{ij}^{(W,t)}$ and N corresponding to the previous notation. Interestingly, the absolute correlation of spike events turned out to be almost independent of the timescale: Given the recorded units, every spike was preceded by about one spike of another cell that was most likely related to its appearance, regardless of the time window (Fig. III-10b). In contrast, only about every fourth spike was preceded by an electrical stimulus that was probably related to its occurrence in time windows of 20 ms or more, whereas in time windows of 5 ms, this can only be presumed for about 2 out of 100 spikes (cf. Fig. III-9). These results substantiate the idea that the recorded spiking activity is governed to a greater extent by the activity of the neuronal network than by the applied electrical stimulation also in absolute terms.

Chapter IV

Towards a unifying neurocognitive theory of brain function

Stephen Colbert: "How does the brain work? Five words or less!"
Steven Pinker: "Brain cells fire in patterns."

The Colbert Report February 7, 2007.

Theories of neural coding revisited: new light through old windows

Taken on its own, a single action potential represents a single bit of information. If all brain cells were active independently, neural coding could be fully described by defining the mechanisms of spike generation and characterizing the firing patterns of single cells, and no extra information would be contained in multineuronal activity patterns. Yet, as this work and many others before have shown, neurons dynamically coordinate their firing. Neural coding thus involves the formation of multineuronal firing patterns that may function as information carrying "symbols". To detect signatures of such potential symbols in local cortical networks, I recorded multineuronal activity from slices of rat visual cortex and systematically searched for coordinated spike events in time windows of up to 50 ms. Activity patterns revealed a significant spatiotemporal structure characterized by synchronous firing within typically up to 5 ms, a timescale that has been found to be optimal for reliable signal transmission[158]. Hence, neural coding in the cortex might involve the selective synchronization of cells. Synchronization of neuronal discharges on the millisecond scale has long been recognized as a prevalent and functionally important attribute of neural activity[100,283,445,475,501,504]. What is unique about the data presented here is the fact that the observed synchrony cannot be explained by common input or synchronous network oscillations. The present findings thus call for an alternative mechanism causing synchronous firing of cells that is intrinsic to the local cortical circuitry.

The "synchrony through synaptic plasticity" hypothesis

As proposed in Chapter III, the activity of distributed cells may be coordinated through synaptic modifications induced by correlated spiking of pre- and postsynaptic neurons. In various neural circuits and a variety of species ranging from insects to humans the strengthening and weakening of synapses has been shown to depend on the relative timing of pre- and postsynaptic spiking in narrow time windows[46,89] that in turn depend on dendritic location[147,149,284,460] and the amplitude and decay time constant of the postsynaptic potential[154]. This so-called spike-timing-dependent plasticity[72] has profound functional implications. Under conditions in which synaptic potentiation occurs if incoming signals slightly precede postsynaptic depolarization, inputs that consistently fire the postsynaptic neuron with short latency develop strong synapses, while synapses of less effective inputs are weakened[316,467]. As a consequence, response latencies of potentiated synapses become shortened[48,332], causing a backward shift of the critical time window and bringing earlier inputs into effect. In this way, neurons could become responsive to ever earlier signals of a recurring input pattern[192,321], so long as the temporal delay between succeeding input spikes does not exceed the critical time window for synaptic plasticity. Multiple neurons being consistently driven by (parts of) the same repeating spike pattern could thus actively synchronize their firing by learning to selectively respond to this pattern at corresponding temporal positions[482]. This possibly gets to the core of what assembles cell assemblies and presents a simple and effective mechanism for coordinating multineuronal activity in the brain. By excluding other sources of synchronous firing, the recordings described in the preceding chapters provide strong indirect support for this "synchrony through synaptic plasticity" hypothesis in local cortical networks.

The hypothesis makes several predictions. First, neurons that coordinate their firing in response to a recurring activity pattern do not need to be physically connected, and also do not need to receive the same input signals. All that matters for the mechanism to work is that they are driven by a coherent pattern of activity, i.e., that their inputs are correlated. The ensuing spike time coordination may therefore extend over any distance, providing a simple solution for the problem of how information can be integrated on different spatial scales in the brain. In the visual and somatosensory system, for example, information about sensory stimuli is known to be concurrently represented and processed in several cortical areas[326,351], and neuronal firing has been found to be significantly coordinated across these areas in both sequential[498] and synchronous[395] activity.

Such widespread spatial integration of information can be achieved through coordinated adjustments of synaptic strengths in response to correlated inputs, although the tendency of distributed cells to engage in coherent activity patterns is likely to decrease with increasing distance[368].

A second prediction is that, given the right plasticity mechanisms, recurring multineuronal spike sequences should progressively be compressed in time if the involved neurons respond to coherent input. Evidence for a temporal compression of repeating activity patterns has indeed been found both in vitro[227] and in vivo[121] and is thought to reflect functional modifications within the neural circuitry. Assuming that the neurons that participate in the spike sequence receive a succession of correlated input signals, the whole process would stop when the cells learned to respond to the very onset of their respective input pattern. If these input patterns have the same temporal origin (which might be related to a sensory stimulus, a motor command or some other cognitive event), the cells would henceforth respond in unison. In this way, any recurring spatiotemporal activity pattern could eventually be translated into the synchronous spiking of a certain set of cells.

This, in turn, would result in a significant sharpening of neuronal representations and an increasingly concise layout of information, which is another prediction of the hypothesis: Repetitive activation patterns would be transformed to short-latency volleys of synchronous spikes[482], whereas new spatiotemporal arrangements of signals would provoke temporally dispersed responses. Further, the synchronization would most likely involve a sparsening of related spike events, optimizing energy efficiency[282] while enhancing the memory capacity of the network[337]. The shaping of the functional circuitry through synaptic plasticity might thus contribute to the establishment of a sparse coding scheme[277,290,369] as it appears to be implemented in several sensory cortical areas[220,236,383,518,534]. Sparse activation of small neuronal populations and even of single cells in neocortex has also been shown to evoke distinct movements[50] and actually drive behavior[219,222], demonstrating a possible functional role of sparse cortical activity that might be explained by the ability of single discharges to initiate both widespread excitation and inhibition[248,285,343,539].

Coordinating neuronal activity: mechanisms and functions

It has been argued that this sensitivity of the cortical network to single action potentials would cause relatively large random membrane potential fluctuations and so entail a reduction of the signal-to-noise ratio in neuronal communication[294], implicitly assuming that spike generation is inherently noisy. Reversing the argument, however, one might as well conclude that this very sensitivity requires signal propagation to be accurately controlled, and one might suspect that the brain has evolved to make optimal use of its limited resources and has developed adaptive mechanisms to prevent the processing of signals that carry no meaning or, even worse, affect information processing. Cortical computation would then imply sparse representations and a very selective routing of signals. To integrate and segregate distributed information efficiently, neurons would have to coordinate their firing and engage in coherent activity patterns while maximizing the repertoire of functional states[469,494], which amounts to operating in a critical regime[40] between total independence and perfect functional unity. Not surprisingly, the brain is endowed with a variety of features and components controlling neuronal cooperation and the activity flow.

The diversity of neuronal cell types and the intrinsic heterogeneity of their biophysical properties has been found to increase the system's coding capacity through a decorrelation of the firing of cells[370]. Likewise, the dynamics in neuronal populations become more complex through the divergent and convergent actions of various transmitters and neuromodulators[106,313]. Another factor that has an impact on the functional repertoire of the network is the mere number of its neurons and the degree of their connectedness. These and other cellular and network properties also play a role in regulating the activity flow in the brain. First and foremost, the propagation of signals is confined by the functional anatomy of the network, that is, synaptic connections and their effective strengths[470]. Given the dense meshwork of axonal and dendritic processes, neural signals are often presumed to potentially take any direction at any time; however, they do not. Rather, synapses may temporarily be silenced[525], and the overall distribution of synaptic strengths in local cortical circuits comprises relatively few strong connections embedded in a "sea" of weaker ones[468,549], constraining the range of possible signaling pathways[302] and again maximizing the network's capacity to produce and retain stable activity patterns[80]. Besides the effective synaptic connectivity, a neuron's functional state is relevant to the routing of signals. In particular, neuronal responsiveness is controlled through the rate and balance of excitatory and inhibitory inputs[195]. Synaptic bombardment can cause continuing depolarizations and an increased variance of the membrane potential and thereby raise a neuron's sensitivity especially to inputs of small amplitude[327,438], while the concomitant drop in input resistance leads to substantial dendritic attenuation of electrical signals, with distal synapses only having minimal effects at the soma[97,126]. This high conductance state[98] is thought to be generated through local recurrent excitation[414] and has been linked to the encoding of sensory information in primary visual cortex[18]. In addition to these input-driven fluctuations between a hyperpolarized "down" state and a depolarized "up" state, the intrinsic ability of neurons to respond selectively to inputs at preferred frequencies[225] affects their integrative properties: Both synaptic mechanisms[489] and ionic conductances[292] may create a resonance effect[224,423,502] that influences spike timing and information transmission between cells[125]. Finally, the propagation of signals directly depends on the ongoing activity pattern, with neuronal oscillations playing a prominent role in defining temporal windows for effective excitation[73,140-142]. The interplay of concurrent excitatory and inhibitory inputs[229], in conjunction with ephaptic transmission of electrical potentials[17,533],

dynamically determines the possible impact of incoming signals[28,78,394]. Acting together, all these structural and functional elements gate the information flow in the brain[7,150,237,412,446,520] and produce both the network and the activity patterns that then give rise to cognitive functions. The work presented in the preceding chapter emphasizes the importance of the functional synaptic connectivity for the dynamic coordination of multineuronal firing[287,302]. In the following, I will further discuss the role of synchronous oscillatory activity therein.

As already mentioned, network oscillations naturally arise from the interplay of recurrent excitatory and inhibitory connections and the resonant properties of individual neurons[74,182,225,292]. Since they are a built-in feature of virtually every neural system, it can be supposed that controlling them is one of the brain's most basic functions[65]. Poised on the brink of instability[21], cortical networks generate synchronous discharges at various frequencies involving variable numbers of cells and thereby distinguish varying functional states for the processing of information[68]. By rapidly balancing excitation with inhibition, the oscillation frequency can be instantaneously modulated[22]. In participating cells, rhythmic synaptic inputs and oscillations of the local electrical field restrict effective excitation to the depolarizing phase of the oscillation cycle, thus adding a dynamic, temporal gate for the transmission of signals to the spatial gates given by the functional connections in the network[141]. Depending on the interactions of concurrent rhythms[70,272,407], spiking activity may hence be orchestrated on multiple spatiotemporal scales in parallel by modulations of the phase and amplitude of distributed oscillations[71,239,268,506], possibly even across subjects when they engage in coordinated actions[288].

Taking a closer look at how oscillatory network activity interferes with a single cell's firing, the question arises to what extent the timing of an action potential can be controlled by oscillatory input. It is known that through the interplay between the magnitude of dendritic excitation and rhythmic inhibition of the somatic region, the more excited cells tend to fire earlier in the oscillation cycle[200,517,521], such that the phase of firing corresponds to the excitatory drive of the neuron[331]. On these grounds, it has been proposed that the interaction of subthreshold membrane potential oscillations with incoming excitatory signals could serve as a fundamental computational mechanism for the implementation of a temporal coding scheme in which information is encoded by the precise timing of a spike relative to the phase of the ongoing oscillation[142,217,348]. Yet, although the particular phase in which a neuron fires can contain significant information[257,344,441], such a coding scheme would necessarily be limited to the timescale on which rhythmic membrane potential fluctuations can advance or delay the spike timing in a systematic way without completely suppressing spike generation. In vitro recordings have shown that this timescale dynamically depends on the average absolute membrane potential, the time constant of the membrane, the strength of the input signal, and the frequency and amplitude of the membrane potential fluctuations[331,521]. According to these studies, only neurons receiving tonic excitatory drive, combined with slow oscillatory input having a relatively long period compared to the membrane time constant, may produce output signals whose timing is smoothly scaled across the whole depolarizing phase. If, however, there is only transient excitatory input or the period of the rhythmic modulation approaches the time constant of the membrane, neural oscillations act essentially as a logic gate relaying incoming excitatory signals only within accordingly narrow time windows. In so doing, network oscillations provide context to afferent signals by selectively routing information in the brain in a dynamic and state-dependent way[59,84,391]. Playing a complementary role to neuronal connectivity, rhythmic modulations of the membrane potential may also synchronize multineuronal firing when paired with prolonged excitatory input; in this case, spike timing is largely determined and actively controlled by the phase of the modulation and the overall activation level of the cell[275,315,521]. In addition, oscillations of the membrane potential may improve action potential precision by imposing defined temporal windows for the effective integration of excitatory inputs[218,383,417]. Thus, rhythmic excitability fluctuations are able to dynamically control the routing and the timing of neuronal signals[235], and it might not be a coincidence that the phase of ongoing network oscillations in the human brain has been found to correlate with the perception of sensory stimuli[62,303]. Whether the interference of oscillatory activity with afferent inputs and the resulting spike timing constitute a temporal code or a binary code then depends on the timescale of their effective interaction.

Neuronal cooperation and the timescale of cell assembly activation

The above arguments on the coordination of neuronal activity through synaptic plasticity and the dynamic gating of signals suggest that the emergence of synchronous cell assemblies and their selective activation may be central to cortical information processing. The following considerations are concerned with the likely organization of these assemblies in time and space.

The first issue relates to the temporal precision with which neurons could be expected to coordinate their firing. In the past, spike timing was assessed relative to the timing of external stimuli or other spikes emitted by the same cell, a praxis that stems from the classic approach to explain brain function by characterizing the response properties of single cells and simply ignores coordination of spike events across populations of cells. These studies led to the notion that neuronal spiking was generally unreliable and imprecise, which in turn led to the – still widely accepted – conclusion that information cannot be represented by the precise timing of spikes. Yet, this view might just be a misconception of the temporal organization of neuronal firing that follows from not taking into account spike timing across cells, which would require the parallel recording of multiple single cells and appropriate measures of their spike time coordination[320]. We know from combined voltage sensitive dye imaging and intracellular recordings that the firing of a cortical

neuron strongly depends on the present activity pattern in the surrounding area[499] and that the large variability of responses to sensory stimulation arises from a quite deterministic interaction of afferent signals with the ongoing activity[20]. Sensory evoked neural activity thus represents the modulation of ongoing circuit dynamics by sensory afferents, rather than directly reflecting the structure of the input signal[132]. This means that it might be more instructive to relate a neuron's firing to the activity of its peers than to some external event[112,371]. The analysis of multineuronal spiking activity presented in the preceding chapter provides just that: By defining spike timing precision with respect to the surrounding network activity, neuronal discharges were shown to be coordinated with a median accuracy of less than a millisecond. This result is categorically different from the results obtained in single neuron studies and suggests a possible role of precise spike timing in cortical information processing. Furthermore, coordinated discharges typically tended to be synchronized to within one or two milliseconds, involving varying sets of distributed neurons. The idea that these neurons form parts of functional cell assemblies whose activation represents cognitively meaningful units of information and that these assemblies are distinguished by synchronous firing within typically two milliseconds fits with the finding that animals can exploit differences in the timing of cortical signals that are as short as three milliseconds to guide decisions[221,548]. Taken together, these observations suggest that neural assemblies may evolve and take effect on a timescale of a few milliseconds.

The second issue concerns the role of single neurons in information representation. How independent is the message a single cortical neuron conveys by sending an action potential down its axon from the signaling of others? If it was perfectly related to the firing of any other cell, their signaling would be totally redundant, reducing both the network's coding capacity and efficiency. If, on the other hand, it was fully independent, then this would imply that it had a fixed meaning that is unmodifiable by collateral signals, like in a labeled line code. However, several arguments suggest that the information carried by the spiking of a particular neuron may not be invariant but be dependent on the functional state of the network as a whole.

First of all, neural connections and synaptic strengths are plastic and subject to continuing modifications[342], resulting in an ever-changing functional structure of the neuronal network[9,10,366]. Together with other factors that control the integration of synaptic inputs on short timescales like ongoing network oscillations and dynamic excitability changes[83], this leads to a substantial variability in the receptive fields of neurons[148,481,532,542] and thus in the information that could be conveyed by a particular cell. Moreover, neurons in sensory cortical areas may adapt their receptive fields to the properties of the current sensory input[136,550] and were shown to be sensitive also to the larger stimulus pattern[161,194], variations of central states[336,528], stimuli presented simultaneously in other modalities[103,273,513], shifts in spatial and feature selective attention[145,540], and reward values[53,440]. The responsiveness of single neurons thus depends on the concurrent sensory or cognitive context[177]. Again, this means that neurons dynamically coordinate their firing and represent information by their joint activity. As pointed out in the previous section, the synergistic encoding of information by ensembles of neurons, as opposed to single-cell coding, would be favorable also from a theoretical point of view: The variable binding of distributed cells into functionally coherent groups would maximize the repertoire of functional states and in this way dramatically improve the coding capacity of the network. Furthermore, dynamic changes in neuronal firing correlations observed during sensory processing or working memory operations[156,368,410,505] seem to confirm that individual cells flexibly take part in multineuronal representations. Given the large number of diverging and converging connections in the brain, correlated activity indeed appears to be the rule rather than an exception, as it is commonly caused by shared input[274,464]. Yet at the same time, neighboring neurons in sensory cortical areas may actively decorrelate their firing[105,164,396,397,464] in a stimulus-dependent way[163,518,551]. Cortical neurons thus dynamically change the partners with which they share coherent information, implying that the information that is transmitted by a single cell may vary as a function of the activity of its peers and cannot possibly be decoded without taking into account the concurrent population pattern.

The dynamic interdependence of neuronal responses and the ensuing formation of multineuronal representations bring us back to the binding problem[404,458,497,523]: If a neuron's firing, taken on its own, does not unambiguously indicate some specific feature or state of the inner or outer world, what are the mechanisms that are responsible for the emergence of defined, meaningful firing patterns across multiple cells, and what is the resulting spatiotemporal structure of these patterns[457]? In the preceding parts of this thesis, I tried to argue that synchronous firing within a few milliseconds is what defines the members of a functional cell assembly, and that their collective activation is what probably defines irreducible units of information. A possible mechanism that could mediate the selective formation of synchronous cell assemblies as carriers of coherent information is given by the "synchrony through synaptic plasticity" hypothesis which consistently explains the signaling of functional relations between neurons by synchronous spikes *and* the computation of these signals[428]: Whenever multiple cells repeatedly receive correlated input, the associated changes of synaptic efficacies may eventually create a set of cells that synchronously respond to a specific activity pattern and thus share a common, complex receptive field which cannot be reduced to the receptive fields of single neurons[52,242]. Other sources of synchronous spiking and their possible role in neural coding will be discussed in the following section.

Neural correlates of cognitive processes

If the "synchrony through synaptic plasticity" hypothesis is correct, groups of synchronously active cells will

spontaneously emerge in response to repeating excitation patterns in the brain and consequently indicate a "known" event, suggesting that these synchronous cell assemblies may function as symbols of a neuronal code that serve as a correlate of some discrete cognitive content. In what follows, I will speculate about possible roles of these multineuronal signals in cognitive processes.

Plasticity, degeneracy and the hard problem

When relating neural activity to cognitive functions and phenomena, two things need to be considered. The first concern applies to any physical account of conscious experiences and has been addressed as "the hard problem of consciousness" by Chalmers[77]. Although it is evident that experience arises from a physical basis, we have no satisfactory explanation of why and how it so arises. We may well specify the mechanisms that are responsible for the performance of certain neural or cognitive functions, but why brain activity gives rise to subjective experiences or "qualia" is entirely unclear. In principle, any neuronal process could be instantiated in the absence of experience, or as Chalmers[77] put it: "Experience may *arise* from the physical, but it is not *entailed* by the physical". It follows that no account of the physical brain processes will tell us why and how they lead to the emergence of qualia. Even if we succeed in mechanistically explaining the ability of the brain to discriminate, to categorize and to appropriately react to environmental stimuli, to shift attention and to deliberately control behavior, to learn and to adapt, and to selectively combine, memorize and recall pieces of information, we are still limited to a *phenomenological* correlation of observed brain dynamics and subjective conscious experiences[166]. (Nevertheless, we may hope to identify some psychophysical principles connecting the properties of neuronal processes to the properties of related experiences, as I will exemplify below.)

A second problem concerning the relation of neural activity and cognitive functions is caused by the tendency of the brain to display degeneracy[108], that is, the ability of structurally different elements to perform the same function or yield the same output. Degeneracy can be found on virtually any organizational level in the brain and is an inherent feature of intra- and intercellular signaling, synaptic plasticity, motor commands and body movements, and also inter-subject communication (there might be large or sometimes even infinite numbers of ways to transmit the same message, a situation most obvious in language). For instance, different combinations of ionic conductances affecting the integration of dendritic input signals may lead to the generation of identical output signals of a cell[6]. On the network level, different configurations of connection strengths and cellular properties may produce the same population activity patterns[388], and different ensembles of neurons may be dynamically configured to initiate the same behavior[338]. Accordingly, we do not lose the perception of a seen object just because its image is slid across the retina – we may continuously perceive it as the same object although changing populations of neurons receive and carry the associated information (eye movements actually serve to *maintain* a stable visual perception[318]). Further evidence supports the idea that there is no simple one-to-one relationship between particular activity patterns in the brain and certain cognitive processes: In visual cortex, neuronal responses to repeated presentations of the same stimulus are highly variable and are strongly determined by the ongoing activity[20,132]. In addition, synaptic connections continuingly undergo extensive remodeling[124,342], providing the brain with an adaptive yet inherently unstable functional structure. This implies that neural codes may change with time through learning, and that the same activity pattern may be interpreted differently (or evoke a different behavior) later in the day[110].

Synchrony and the principle of structural coherence

Despite these difficulties in establishing well-defined relationships between mind and brain activity, neuronal signaling in the brain is highly organized and far from being random, as are cognitive processes, and some relationship – i.e., some code relating a given activity pattern to a particular cognitive function – must exist between the two. In this thesis, I took a naive approach towards characterizing the spatiotemporal organization of cortical spiking activity and found a prevalence of precisely synchronous spike events among distributed discharges in the visual cortex of the rat. Although a variety of differences exist between different cortical areas in general and between rat and human cortices in particular regarding cell types, cell morphology and cellular connectivity[93,113], we can expect the same synchronization of neuronal signaling to occur also in the human brain and in other neural systems. According to the "synchrony through synaptic plasticity" hypothesis, any two neurons equipped with Hebbian spike-timing-dependent plasticity receiving correlated input may end up responding to this input in synchrony, as explained above, through a mechanism that is largely unaffected by variations in neuronal circuitry[508].

Given the synchronous activity of a selective set of cells, what could be its cognitive correlate? Chalmers argued that information, when being processed by the brain, has two basic aspects, a physical aspect and a phenomenal aspect, and that its physical representation should have a structure that corresponds directly to the differences between phenomenal states[77]. If the argument holds, this principle of structural coherence might provide a fundamental link between the characteristics of cognitive processes on one hand and the organization of neuronal activity on the other. Applied to neuronal synchrony, it means that the directed coactivation of a group of distributed cells could be expected to be paralleled by a meaningful convergence of information on the cognitive level, such as reaching a decision or creating a coherent perception based on distributed signals[117,180,447,458]. It has also been suggested that the same mechanisms that mediate the synchronization of distributed discharges are responsible for giving rise to conscious experiences[118,451,453]. Becoming aware of a percept or an idea necessarily involves a transition of

one functional state to a different state, as has been noted by Dennett[94]: "It seems obvious that there has to be a time before which we are not conscious of some item and after which we are conscious of it. In some sense, then, we *become* conscious of various features of our experience, so there must be some kind of transition, if not arrival at a place or crossing of a boundary, then a change of functional state of one sort or another." Although it is difficult to demonstrate that conscious perception requires neuronal synchrony, it should be clear that the transient synchronization of a non-random assembly of cells is precisely that sort of neuronal activity that indicates non-random changes in functional states on timescales fast enough to comply with all kinds of cognitive processes, and thus satisfies the principle of structural coherence also with respect to the processes underlying conscious experiences.

Any functional interpretation of neuronal synchrony has to include and obviously depends on the specific mechanism that is responsible for the synchronization of neuronal activity. One reason for synchronous firing could be concurrent activation through sensory stimuli. In the mammalian visual system, retinal ganglion cells may fire synchronous action potentials simply because they are driven by a common stimulus. In the lateral geniculate nucleus, corticothalamic efferents come into play, and synchrony of afferent signals becomes a little less trivial[442,443]. After reaching cortical areas, sensory signals are modulated by the ongoing cortical activity and follow a multitude of diverging and converging pathways within both local and widespread cortical circuits. In other words, the more central the signals, the more their timing depends on the functional architecture of the network, and the less trivial is the interpretation of neuronal synchrony relating to the sensory stream[287].

In neocortex, neural connections have several inbuilt features that directly support the synchronization of neuronal signals. The most significant feature is the coupling of inhibitory interneurons by gap junctions, forming large, continuous, cell type-specific syncytia [16,157,176]. Although the amplitude of electrotonic signals quickly falls off with distance, this electrical coupling facilitates the widespread synchronization of rhythmic inhibitory activity, which in turn constrains the firing of entire populations of pyramidal cells to narrow time windows[41,42,67,160,537]. As will be explained below, this mechanism is important for the instantaneous coordination of multineuronal spiking (and hence cognitive processes) on very short timescales. Another feature of many thalamocortical and corticocortical connections is the compensation for differing lengths of a cell's axonal branches by adjusting the degree of myelination and the diameter of the fibers such that all postsynaptic targets receive the signal at the same time[228,262,411]. This makes sense if one assumes that synchrony is a tag for "belonging together" or "being one". Finally, the cortical network might also be endowed with synchronizing mechanisms other than common input or synchronous network oscillations. The "synchrony through synaptic plasticity" hypothesis explains how synchronous cell assemblies could emerge through correlated changes in synaptic efficacies in response to repeating excitation patterns. The gradual build-up of synchronous groups of cells through learning would be vital for creating "sense out of fact"[242] and might help to distinguish between familiar and unfamiliar experiences[482]. Importantly, though, those synchronous groups would be different from a classical Hebbian cell assembly in that the participating cells do not need to be directly connected[169,207], but are fully characterized by the transient (and non-random) synchrony of their discharges[137,198]. From all the cells that happen to be activated by the same repeating excitation pattern, certain subsets could be selected as synchronous groups through coordinated changes in synaptic strengths and latencies[107] based on the relative onset of a cell's input signals associated with that excitation pattern. Their joint activity would then signify the onset of that exact neural (or cognitive) event with shortest possible latency. As suggested by the arrangement of synchronous cell ensembles in the visual cortex of the rat (see Chapter III), neurons being recruited into synchronous groups most likely derive from already existing functional subnetworks[553] within which synaptic connections are relatively frequent and strong. The related reshaping of the cortical network is thought to underlie the consolidation of newly acquired "knowledge" and goes along with the fact that neuronal populations in sensory areas exhibit similar activation patterns both spontaneously and in response to sensory input[240,300,302,305], suggesting that sensory responses are drawn from a limited "vocabulary" of possible activity patterns given by the intracortical functional synaptic connectivity. These findings again support the idea that the functional layout of cortical synaptic connections plays a major role in coordinating and synchronizing cortical activity[287].

There is, however, a problem: Adaptive changes of synaptic connections occur on a much longer timescale than most cognitive processes and are unable to represent changes in sensory information in real time. This inability to instantly reconfigure the functional network in response to afferent signals has been referred to as the "learning-time barrier"[523] and calls for an additional mechanism that can coordinate multineuronal activity on the timescales on which cognitive processes take place, i.e., within milliseconds. It is obvious that such a fast mechanism can only be realized through dynamic activity patterns emerging from and interacting with the functional neural circuitry. Although any complex activation pattern could in principle serve to selectively excite a certain set of cells while inhibiting others, the most prominent and ubiquitous population pattern that is known to flexibly and coherently modulate the excitability of distributed cells on a millisecond timescale is synchronous network oscillations. As explained in previous sections, network oscillations naturally arise from the resonant properties of individual neurons and from the interplay of recurrent excitatory and inhibitory connections. The propensity to produce synchronous oscillations is higher when recurrent feedback is strong [527], which adds to the reason why cortical gamma-band oscillations tightly correlate with hemodynamic signals indicating an increase in energy consumption[352]. This

means that network oscillations are not inherently an energetically "cheap" way to achieve neuronal synchrony, as has been suggested[65]. In fact, the "cheapest" way to synchronize the activity of neuronal ensembles would be to arrange functional synaptic connections such that some selected sets of cells will be synchronously activated by a certain preceding (and possibly sparse) activity pattern, as proposed by the "synchrony through synaptic plasticity" hypothesis. Nevertheless do cortical networks readily engage in oscillatory activity, thus serving the need for a fast and dynamic coordination of multineuronal firing. On a mechanistic level, the principal function of synchronous oscillations comes down to a temporal modulation of the effective neuronal connectivity through rhythmic fluctuations of the excitability of cells. Synchronized rhythmic activity and functional synaptic connections thus combine in a complementary way to allow for a spatially and temporally selective transmission of signals and hence for a selective activation of neuronal ensembles at any point in time.

On a more cognitive level, synchronous oscillations effectively reduce the system's degrees of freedom and restrict the space of possible activity patterns, so as to concentrate on some signals and the information they carry while disregarding others. Indeed do neurons in macaque area V4 that are activated by an attended stimulus engage in enhanced gamma-band synchronization compared with neurons activated by a distracter, pointing to a functional role of synchronous network oscillations in attentional stimulus selection[143,144]. For network rhythms to synchronize the activity of neuron groups, though, it is unimportant if they exhibit a stable phase and frequency – all that matters for an effective coordination of multineuronal signals is the limitation of neuronal discharges to narrow time windows by alternating volleys of synchronous excitation and inhibition[22,229,354]. Could such rhythmic network activity in principle be sufficient for a selective synchronization of neuronal discharges, irrespective of the functional synaptic connectivity? It clearly can not, first because synchronous oscillatory activity is coherent across cell populations[115,241] and thus lacks the spatial selectivity needed for efficient neural coding, and second because meaningful neuronal synchrony can only arise through experience and learning, which involves adjustments of synaptic efficacies and connections. It is thus evident that while dynamic activation patterns are needed to flexibly arrange synchronous cell assemblies on short timescales, functional adaptations of selected synaptic connections are required to allow for a selective synchronization of cells in the first place. According to the "synchrony through synaptic plasticity" hypothesis, the directed assembly of cells into synchronous groups could be based on the detection of repeating activity patterns and hence on recognizing recurrence as a fundamental property of behaviorally relevant events.

The emergence of selective neuronal synchrony as a potential carrier of information bears the question of how this synchrony is interpreted in subsequent processing stages. It has been argued that parallel synaptic inputs arriving synchronously at a postsynaptic neuron summate more effectively and for this reason transmit their signals more reliably than temporally dispersed inputs[266,526]. Although pyramidal neurons are indeed more sensitive to coincident inputs especially at high activity levels[27,28,387], it should be clear that any synchronous discharge pattern in the brain will fan out in both time and space through a multitude of converging and diverging connections, meaning that its impact on downstream cell populations is not determined by the synchronicity of the signals *per se*. The question thus becomes what distinguishes meaningful synchronous discharge patterns from accidental neuronal synchrony, and the only possible answer is their usefulness in interpreting sensory information and generating appropriate behavior[66] by informing downstream populations of neurons about their current functional coherence in an "understandable" and meaningful way. The ability to do so can be expected to depend on prior adaptations of the functional network to enable the immediate recognition and classification of the associated information.

As mentioned before, a typical property of functional brain networks is the dense local clustering of synaptic connections and the linking of distant cell populations through direct long-range projections[247,468] (see Chapter III). A consequence of this so-called small-world architecture[34,471,530] is a small number of intermediary synaptic connections in the transmission route between any pair of neurons and hence an almost immediate global integration of information. Assuming at most six synaptic processing steps for transmitting a signal from any neuron to any other neuron and a maximal average delay of 30 ms between pre- and postsynaptic firing[496], it appears that it takes no more than 180 ms to potentially involve the entire brain. Interestingly, this is a typical timescale for shifts in attention, saccadic eye movements, and reaction times to sensory stimuli[496], so we are led to infer that every minute of our experience is based on a brain-wide evaluation of neural activity. This is exactly where the classical concept of sender and receiver, as applied to neural coding, is bound to fail – if the whole network is involved in distilling the meaning of a given activity pattern, a "message" and its subsequent interpretation would be separated in time, but would not rely on any dedicated communication channel. At any one moment, the brain *as a whole* would integrate all the available information and tend to some coherent state that consistently combines all the different aspects of its current condition on various spatiotemporal scales. Information processing in the brain could thus be conceived as a continuous transformation of transient activity patterns on all of these scales[196], while sending and receiving elements cannot be distinguished.

Emergence of the phase sequence and the "problem of the direction of thought"

The ongoing transformation of spatiotemporal activity patterns and its role in neuronal coding is specifically addressed by the concept of reservoir computing, also known as echo-state or liquid computing[304]. The basic

idea of this computational framework is the mapping of dynamical input states onto a high-dimensional state space of the neuronal network, referred to as reservoir, followed by a classification of the distributed activity patterns through a trained readout that makes the computation meaningful and problem-oriented. The readout can in principle be realized by single neurons that sample the activity from multiple upstream cells in the network and adjust the strength of their incoming synaptic connections so that they are selectively activated by a particular state of the reservoir. Alternatively, the readout stage could comprise cell assemblies that are selectively activated by specific states of the network and in turn directly control effectors without entailing the famous "bottleneck problem". The virtue of a high-dimensional state space in information representation is the fact that stimulus-evoked patterns occupy much lower-dimensional subspaces, i.e., involve only small fractions of the neurons in the network, and can thus easily be discriminated based on their spatiotemporal signatures[456]. The important bit is that stable internal states are not required for giving a stable output, since transient internal states can be transformed by readout neurons into defined output signals[377] due to the high dimensionality and inherent degeneracy of the dynamical system[304]. The response of the reservoir to external perturbations strongly depends on the current functional state of the network[58], which in turn depends on its past activity. Distributed signals reverberating through recurrent connections in sufficiently large and heterogeneous neural circuits provide the reservoir with a fading memory of recent inputs that allows it to integrate input sequences over both time and space[355]. The linear decline of relative spike timing precision with increasing spike delays reported in Chapter III could be an indication of such a fading memory affecting the generation of precisely timed discharge sequences in neocortical circuits.

While cortical networks feature an extremely high-dimensional state space and naturally give rise to reservoir dynamics, repeated stimuli would induce plastic changes of their functional synaptic connectivity such that the resulting response patterns become more and more stable, and different responses to different stimuli become more and more segregated[279]. By transitioning "from vague to crisp"[375], response patterns may evolve into attractor-like network states[35,86,416], thus improving their discriminability and instantiating learning on the network level[377,482]. As a consequence, cortical activity is not chaotic, but is governed by contextual priors that are embedded in the functional network architecture and that bias multineuronal firing towards "expected" or "anticipated" pre-learned response patterns[263]. This inherent bias, in turn, is influenced by attention and the spatial and temporal context of incoming stimuli[131,340] and is instrumental in associating preceding network states (or "ideas") with succeeding states in a coherent and meaningful way.

The conception of a serial dependence of successive, relatively short-lived network states refers to a crucial point in any comprehensive neurocognitive theory: Since many cognitive functions and phenomena are of sequential nature (think of the production and reception of speech and music, episodic memory, or any serial motor task), how could the organization of neuronal network activity be conceptualized to account for coherent sequences of cognitive events? In dynamical systems terms, favored activation patterns that emerged through learning represent saddle points in an attractor landscape that are preferentially visited by the system while it is moving around in state space[21,104]. The ordered succession of visited saddle points or neuronal activation patterns, then, would give rise to a heteroclinic sequence linking transient metastable network states across time and space[11,393] in some meaningful way. Importantly, classical attractor states like fixed points and limit cycles are biologically implausible because they constrain the system's dynamics to low-dimensional subspaces and are therefore unlikely to be realized in biological circuits[392]. The information that is processed would thus be organized into a sequence of transient states that depends on the system's history, rather than by eventually reached attractors or steady states[110,245].

This modern view of sequences of brain activity patterns as dynamically evolving trajectories in state space [330] agrees well with the Hebbian conception of "phase sequences" that are proposed to arise from the sequential activation of cell assemblies, where each assembly represents a semi-stable state and is active only transiently[198,207]. Hebb conjectured that the firing of a single cell assembly corresponds to a single idea as a basic cognitive entity, and that the successive activation of different cell assemblies represents the neural basis of the thought process. An essential property of the phase sequence concept is the relatively loose coupling and flexible timing of successional cell assemblies, which distinguishes it from the concepts of synfire chains and synfire braids which assume temporally precise firing over relatively long stretches of time[3,47,232]. However, from a theoretical point of view, such strictly defined long firing sequences would reduce the flexibility and effective dimensionality of the system in an unsuitable and unnecessary way[102] and are thus unlikely to occur in real neuronal networks. On the other hand, the rather flexible organization of phase sequences poses another problem that Hebb was well aware of and that has been addressed as the "problem of the direction of thought" by Humphrey early on. In his 1939 article[223], he noted that "there is the problem of why one member of the train [of thoughts] succeeds another, why a particular idea or image occurs, more or less relevantly to the point at issue, instead of the million and one irrelevancies that might occur". This problem, of course, gets to the core of human cognition, because it is asking for a mechanism that would explain the organization of the human thought process. Although we are far from any complete explanation of why and how one thought or idea gives rise to another, one might well speculate that there exist some basic neural mechanisms that provide neuronal networks with the potential to produce long sequential activity patterns and to direct their trajectory in state space to some degree. A promising candidate for such a mechanism is the classical spike-timing-

dependent plasticity rule, according to which synapses are strengthened if they reliably take part in driving the postsynaptic cell with short latency[88]. In addition, long-range cytoplasmic signaling within the presynaptic neuron may lead to a retrograde propagation of synaptic potentiation along excitatory pathways[484]. The consequent formation of chains of relatively strong and effective synaptic connections has been demonstrated to produce long, diverse activity sequences in simulation experiments[130,263] and probably constitutes the main neural mechanism of temporal order learning in neuronal networks[153]. By directing the propagation of signals within the network, chains of strong synaptic connections might also be responsible for the cue-triggered recall of learned temporal sequences: After conditioning through repeated sequential stimulation of distributed groups of neurons, neuronal networks have been shown to reproduce the complete sequential response pattern upon being exposed to only the initial part of the stimulus sequence[263,363,545].

If synaptic plasticity indeed underlies the emergence of Hebbian phase sequences, the resulting functional network structure should allow for a flexible propagation of signals along potentiated pathways through the network, so that the succession of visited states or multineuronal activity patterns is probabilistic, rather than fixed. Experimental evidence for non-random yet variable activity sequences in cortical networks is very scarce, though. The only finding of circumscribed cell assemblies firing in functional sequences has been reported by Ribeiro and colleagues[15]. They analyzed synchronously active groups of cells recorded simultaneously from the hippocampus and primary visual and somatosensory cortical areas of rats and demonstrated the dependence of cell assembly activation sequences on experimental periods and the behavioral state of the animal. The results presented in Chapter III go beyond these findings by showing that individual sequences repeat significantly often and strongly depend on the neuromodulatory state of the cortical network. In view of theoretical considerations about cortical dynamics and beginning evidence from multielectrode recordings we apparently have to assume that sequential activity patterns in neocortex are controlled by the functional state of the network and are organized into repeating sequences that correlate with behavior, thus fully conforming to the concept of Hebbian phase sequences.

References

1. Abbott, L.F., and W.G. Regehr (2004). Synaptic computation. *Nature* 431 (7010): 796-803.

2. Abbott, L.F., J.A. Varela, K. Sen, and S.B. Nelson (1997). Synaptic depression and cortical gain control. *Science* 275 (5297): 220-224.

3. Abeles, M. (1991). *Corticonics - Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Melbourne.

4. Abeles, M., H. Bergman, E. Margalit, and E. Vaadia (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology* 70 (4): 1629-1638.

5. Abeles, M., and G.L. Gerstein (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *Journal of Neurophysiology* 60 (3): 909-924.

6. Achard, P., and E. De Schutter (2006). Complex parameter landscape for a complex neuron model. *PLoS Computational Biology* 2 (7): e94.

7. Adesnik, H., and M. Scanziani (2010). Lateral competition for cortical space by layer-specific horizontal circuits. *Nature* 464 (7292): 1155-1160.

8. Aertsen, A.M.H.J., M. Diesmann, M.O. Gewaltig, S. Grün, and S. Rotter (2001). Neural dynamics in cortical networks: precision of joint-spiking events. *Complexity in Biological Information Processing* (*Novartis Foundation Symposia* 239), John Wiley & Sons, Chichester, pp. 193-207.

9. Aertsen, A.M.H.J., M. Erb, and G. Palm (1994). Dynamics of functional coupling in the cerebral cortex - an attempt at a model-based interpretation. *Physica D: Nonlinear Phenomena* 75 (1-3): 103-128.

10. Aertsen, A.M.H.J., G.L. Gerstein, M.K. Habib, and G. Palm (1989). Dynamics of neuronal firing correlation: modulation of "effective connectivity". *Journal of Neurophysiology* 61 (5): 900-917.

11. Afraimovich, V.S., V.P. Zhigulin, and M.I. Rabinovich (2004). On the origin of reproducible sequential activity in neural circuits. *Chaos* 14 (4): 1123-1129.

12. Albert, R., H. Jeong, and A.L. Barabási (2000). Error and attack tolerance of complex networks. *Nature* 406 (6794): 378-382.

13. Alle, H., and J.R.P. Geiger (2006). Combined analog and action potential coding in hippocampal mossy fibers. *Science* 311 (5765): 1290-1293.

14. Allen, C., and C.F. Stevens (1994). An evaluation of causes for unreliability of synaptic transmission. *Proceedings of the National Academy of Sciences of the United States of America* 91 (22): 10380-10383.

15. Almeida-Filho, D.G., V. Lopes-dos-Santos, N.A.P. Vasconcelos, J.G.V. Miranda, A.B.L. Tort, and S.T. Ribeiro (2014). An investigation of Hebbian phase sequences as assembly graphs. *Frontiers in Neural Circuits* 8 (34): 1-13.

16. Amitai, Y., J.R. Gibson, M. Beierlein, S.L. Patrick, A.M. Ho, B.W. Connors, and D. Golomb (2002). The spatial dimensions of electrically coupled networks of interneurons in the neocortex. *Journal of Neuroscience* 22 (10): 4142-4152.

17. Anastassiou, C.A., R. Perin, H. Markram, and C. Koch (2011). Ephaptic coupling of cortical neurons. *Nature Neuroscience* 14 (2): 217-223.

18. Anderson, J., I. Lampl, I. Reichova, M. Carandini, and D. Ferster (2000). Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. *Nature Neuroscience* 3 (6): 617-621.

19. Ariav, G., A. Polsky, and J. Schiller (2003). Submillisecond precision of the input-output transformation function mediated by fast sodium dendritic spikes in basal dendrites of CA1 pyramidal neurons. *Journal of Neuroscience* 23 (21): 7750-7758.

20. Arieli, A., A. Sterkin, A. Grinvald, and A.M.H.J. Aertsen (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273 (5283): 1868-1871.

21. Ashwin, P., and M. Timme (2005). When instability makes sense. *Nature* 436 (7047): 36-37.

22. Atallah, B.V., and M. Scanziani (2009). Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron* 62 (4): 566-577.

23. Atmanspacher, H., and S. Rotter (2008). Interpreting neurodynamics: concepts and facts. *Cognitive Neurodynamics* 2 (4): 297-318.

24. Averbeck, B.B., P.E. Latham, and A. Pouget (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience* 7 (5): 358-366.

25. Averbeck, B.B., and D. Lee (2004). Coding and transmission of information by neural ensembles. *Trends in Neurosciences* 27 (4): 225-230.

26. Avissar, M., A.C. Furman, J.C. Saunders, and T.D. Parsons (2007). Adaptation reduces spike-count reliability, but not spike-timing precision, of auditory nerve responses. *Journal of Neuroscience* 27 (24): 6461-6472.

27. Azouz, R., and C.M. Gray (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proceedings of the National Academy of Sciences of the United States of America* 97 (14): 8110-8115.

28. Azouz, R., and C.M. Gray (2003). Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron* 37 (3): 513-523.

29. Bair, W. (1999). Spike timing in the mammalian visual system. *Current Opinion in Neurobiology* 9 (4): 447-453.

30. Bair, W., and C. Koch (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation* 8 (6): 1185-1202.

31. Bak, P., C. Tang, and K. Wiesenfeld (1988). Self-organized criticality. *Physical Review A* 38 (1): 364-374.

32. Barlow, H.B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1 (4): 371-394.

33. Bartos, M., I. Vida, and P. Jonas (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience* 8 (1): 45-56.

34. Bassett, D.S., and E. Bullmore (2006). Small-world brain networks. *The Neuroscientist* 12 (6): 512-523.

35. Bathellier, B., L. Ushakova, and S. Rumpel (2012). Discrete neocortical dynamics predict behavioral categorization of sounds. *Neuron* 76 (2): 435-449.

36. Beck, J.M., W.J. Ma, X. Pitkow, P.E. Latham, and A. Pouget (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74 (1): 30-39.

37. Bédard, C., H. Kröger, and A. Destexhe (2004). Modeling extracellular field potentials and the frequency-filtering properties of extracellular space. *Biophysical Journal* 86 (3): 1829-1842.

38. Bédard, C., H. Kröger, and A. Destexhe (2006). Does the 1/f frequency scaling of brain signals reflect self-organized critical states? *Physical Review Letters* 97 (11): 1-4.

39. Bédard, C., H. Kröger, and A. Destexhe (2006). Model of low-pass filtering of local field potentials in brain tissue. *Physical Review E* 73 (5): 051911.

40. Beggs, J.M. (2008). The criticality hypothesis: how local cortical networks might optimize information processing. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 366 (1864): 329-343.

41. Beierlein, M., J.R. Gibson, and B.W. Connors (2000). A network of electrically coupled interneurons drives synchronized inhibition in neocortex. *Nature Neuroscience* 3 (9): 904-910.

42. Bennett, M.V., and R.S. Zukin (2004). Electrical coupling and neuronal synchronization in the mammalian brain. *Neuron* 41 (4): 495-511.

43. Berger, D., C. Borgelt, S. Louis, A. Morrison, and S. Grün (2010). Efficient identification of assembly neurons within massively parallel spike trains. *Computational Intelligence and Neuroscience* 2010: 439648.

44. Berry, M.J., D.K. Warland, and M. Meister (1997). The structure and precision of retinal spike trains. *Proceedings of the National Academy of Sciences of the United States of America* 94 (10): 5411-5416.

45. Beshel, J., N.J. Kopell, and L.M. Kay (2007). Olfactory bulb gamma oscillations are enhanced with task demands. *Journal of Neuroscience* 27 (31): 8358-8365.

46. Bi, G.Q., and M.M. Poo (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review of Neuroscience* 24 (1): 139-166.

47. Bienenstock, E. (1995). A model of neocortex. *Network: Computation in Neural Systems* 6 (2): 179-224.

48. Boudkkazi, S., E. Carlier, N. Ankri, O. Caillard, P. Giraud, L. Fronzaroli-Molinieres, and D. Debanne (2007). Release-dependent variations in synaptic latency: a putative code for short- and long-term synaptic dynamics. *Neuron* 56 (6): 1048-1060.

49. Braitenberg, V. (1978). Cell assemblies in the cerebral cortex. *Theoretical Approaches to Complex Systems* (*Lecture Notes in Biomathematics* 21), Springer Verlag, Berlin/Heidelberg/New York, pp. 171-188.

50. Brecht, M., M. Schneider, B. Sakmann, and T.W. Margrie (2004). Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature* 427 (6976): 704-710.

51. Bressler, S.L., R. Coppola, and R. Nakamura (1993). Episodic multiregional cortical coherence at multiple frequencies during visual task performance. *Nature* 366 (6451): 153-156.

52. Brette, R. (2012). Computing with neural synchrony. *PLoS Computational Biology* 8 (6): e1002561.

53. Brosch, M., E. Selezneva, and H. Scheich (2011). Representation of reward feedback in primate auditory cortex. *Frontiers in Systems Neuroscience* 5 (5): 1-12.

54. Brown, E.N., R.E. Kass, and P.P. Mitra (2004). Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neuroscience* 7 (5): 456-461.

55. Brown, R.E., and P.M. Milner (2003). The legacy of Donald O. Hebb: more than the Hebb synapse. *Nature Reviews Neuroscience* 4 (12): 1013-1019.

56. Brown, S.L., J. Joseph, and M. Stopfer (2005). Encoding a temporally structured stimulus with a temporally structured neural representation. *Nature Neuroscience* 8 (11): 1568-1576.

57. Buonomano, D.V. (2003). Timing of neural responses in cortical organotypic slices. *Proceedings of the National Academy of Sciences of the United States of America* 100 (8): 4897-4902.

58. Buonomano, D.V., and W. Maass (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience* 10 (2): 113-125.

59. Burchell, T.R., H.J. Faulkner, and M.A. Whittington (1998). Gamma frequency oscillations gate temporally coded afferent inputs in the rat hippocampal slice. *Neuroscience Letters* 255 (3): 151-154.

60. Burkhalter, A. (1989). Intrinsic connections of rat primary visual cortex - laminar organization of axonal projections. *Journal of Comparative Neurology* 279 (2): 171-186.

61. Burr, D., and M.C. Morrone (2006). Time perception: space-time in the brain. *Current Biology* 16 (5): R171-R173.

62. Busch, N.A., J. Dubois, and R. VanRullen (2009). The phase of ongoing EEG oscillations predicts visual perception. *Journal of Neuroscience* 29 (24): 7869-7876.

63. Bushdid, C., M.O. Magnasco, L.B. Vosshall, and A. Keller (2014). Humans can discriminate more than 1 trillion olfactory stimuli. *Science* 343 (6177): 1370-1372.

64. Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience* 7 (5): 446-451.

65. Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press, New York.

66. Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron* 68 (3): 362-385.

67. Buzsáki, G., and J.J. Chrobak (1995). Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Current Opinion in Neurobiology* 5 (4): 504-510.

68. Buzsáki, G., and A. Draguhn (2004). Neuronal oscillations in cortical networks. *Science* 304 (5679): 1926-1929.

69. Cannon, J., M.M. McCarthy, S. Lee, J. Lee, C. Börgers, M.A. Whittington, and N. Kopell (2014). Neurosystems: brain rhythms and cognitive processing. *European Journal of Neuroscience* 39 (5): 705-719.

70. Canolty, R.T., E. Edwards, S.S. Dalal, M. Soltani, S.S. Nagarajan, H.E. Kirsch, M.S. Berger, N.M. Barbaro, and R.T. Knight (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science* 313 (5793): 1626-1628.

71. Canolty, R.T., K. Ganguly, S.W. Kennerley, C.F. Cadieu, K. Koepsell, J.D. Wallis, and J.M. Carmena (2010). Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies. *Proceedings of the National Academy of Sciences of the United States of America* 107 (40): 17356-17361.

72. Caporale, N., and Y. Dan (2008). Spike timing-dependent plasticity: a Hebbian learning rule. *Annual Review of Neuroscience* 31: 25-46.

73. Cardin, J.A., M. Carlén, K. Meletis, U. Knoblich, F. Zhang, K. Deisseroth, L.H. Tsai, and C.I. Moore (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* 459 (7247): 663-667.

74. Cardin, J.A., L.A. Palmer, and D. Contreras (2005). Stimulus-dependent gamma (30-50 Hz) oscillations in simple and complex fast rhythmic bursting cells in primary visual cortex. *Journal of Neuroscience* 25 (22): 5339-5350.

75. Carpenter, G.A., and S. Grossberg (1987). Discovering order in chaos - stable self-organization of neural recognition codes. *Annals of the New York Academy of Sciences* 504: 33-51.

76. Ch'Ng, Y.H., and R.C. Reid (2010). Cellular imaging of visual cortex reveals the spatial and functional organization of spontaneous activity. *Frontiers in Integrative Neuroscience* 4 (20): 1-9.

77. Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2 (3): 200-219.

78. Chance, F.S., L.F. Abbott, and A.D. Reyes (2002). Gain modulation from background synaptic input. *Neuron* 35 (4): 773-782.

79. Chawla, D., K.J. Friston, and E.D. Lumer (2001). Zero-lag synchronous dynamics in triplets of interconnected cortical areas. *Neural Networks* 14 (6-7): 727-735.

80. Chen, W., J.P. Hobbs, A. Tang, and J.M. Beggs (2010). A few strong connections: optimizing information retention in neuronal avalanches. *BMC Neuroscience* 11 (3): 1-14.

81. Chiu, C.Q., G. Lur, T.M. Morse, N.T. Carnevale, G.C.R. Ellis-Davies, and M.J. Higley (2013). Compartmentalization of GABAergic inhibition by dendritic spines. *Science* 340 (6133): 759-762.

82. Christen M. (2006). The role of spike patterns in neuronal information processing - A historically embedded conceptual clarification. Dissertation, ETH Zürich.

83. Civillico, E.F., and D. Contreras (2012). Spatiotemporal properties of sensory responses in vivo are strongly dependent on network context. *Frontiers in Systems Neuroscience* 6 (25): 1-20.

84. Colgin, L.L., T. Denninger, M. Fyhn, T. Hafting, T. Bonnevie, O. Jensen, M.B. Moser, and E.I. Moser (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* 462 (7271): 353-357.

85. Connors, B.W., and M.A. Long (2004). Electrical synapses in the mammalian brain. *Annual Review of Neuroscience* 27: 393-418.

86. Cossart, R., D. Aronov, and R. Yuste (2003). Attractor dynamics of network UP states in the neocortex. *Nature* 423 (6937): 283-288.

87. Dan, Y., J.M. Alonso, W.M. Usrey, and R.C. Reid (1998). Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neuroscience* 1 (6): 501-507.

88. Dan, Y., and M.M. Poo (2004). Spike timing-dependent plasticity of neural circuits. *Neuron* 44 (1): 23-30.

89. Dan, Y., and M.M. Poo (2006). Spike timing-dependent plasticity: from synapse to perception. *Physiological Reviews* 86 (3): 1033-1048.

90. Davidson, T.J., F. Kloosterman, and M.A. Wilson (2009). Hippocampal replay of extended experience. *Neuron* 63 (4): 497-507.

91. deCharms, R.C., and M.M. Merzenich (1996). Primary cortical representation of sounds by the coordination of action potential timing. *Nature* 381 (6583): 610-613.

92. deCharms, R.C., and A.M. Zador (2000). Neural representation and the cortical code. *Annual Review of Neuroscience* 23: 613-647.

93. DeFelipe, J. (2011). The evolution of the brain, the human nature of cortical circuits, and intellectual creativity. *Frontiers in Neuroanatomy* 5 (29): 1-17.

94. Dennett, D.C. (2009). Multiple Drafts Model. *The Oxford Companion to Consciousness*, Oxford University Press, New York, pp. 452-454.

95. Desbordes, G., J. Jin, C. Weng, N.A. Lesica, G.B. Stanley, and J.M. Alonso (2008). Timing precision in population coding of natural scenes in the early visual system. *PLoS Biology* 6 (12): e324.

96. Desmaisons, D., J.D. Vincent, and P.M. Lledo (1999). Control of action potential timing by intrinsic subthreshold oscillations in olfactory bulb output neurons. *Journal of Neuroscience* 19 (24): 10727-10737.

97. Destexhe, A., and D. Paré (1999). Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *Journal of Neurophysiology* 81 (4): 1531-1547.

98. Destexhe, A., M. Rudolph, and D. Paré (2003). The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience* 4 (9): 739-751.

99. DeWeese, M.R., M. Wehr, and A.M. Zador (2003). Binary spiking in auditory cortex. *Journal of Neuroscience* 23 (21): 7940-7949.

100. DeWeese, M.R., and A.M. Zador (2006). Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *Journal of Neuroscience* 26 (47): 12206-12218.

101. Diba, K., and G. Buzsáki (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience* 10 (10): 1241-1242.

102. Diesmann, M., M.O. Gewaltig, and A.M.H.J. Aertsen (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402 (6761): 529-533.

103. Driver, J., and T. Noesselt (2008). Multisensory interplay reveals crossmodal influences on 'sensory-specific' brain regions, neural responses, and judgments. *Neuron* 57 (1): 11-23.

104. Durstewitz, D., and G. Deco (2008). Computational significance of transient dynamics in cortical networks. *European Journal of Neuroscience* 27 (1): 217-227.

105. Ecker, A.S., P. Berens, G.A. Keliris, M. Bethge, N.K. Logothetis, and A.S. Tolias (2010). Decorrelated neuronal firing in cortical microcircuits. *Science* 327 (5965): 584-587.

106. Edeline, J.-M. (2012). Beyond traditional approaches to understanding the functional role of neuromodulators in sensory cortices. *Frontiers in Behavioral Neuroscience* 6 (45): 1-14.

107. Edelman, G.M. (1978). Group selection and phasic reentrant signalling: a theory of higher brain function. *The Mindful Brain - Cortical Organization and the Group-Selective Theory of Higher Brain Function*, MIT Press, Cambridge/MA, pp. 51-100.

108. Edelman, G.M., and J.A. Gally (2001). Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences of the United States of America* 98 (24): 13763-13768.

109. Egert, U., D.H. Heck, and A.M.H.J. Aertsen (2002). Two-dimensional monitoring of spiking networks in acute brain slices. *Experimental Brain Research* 142 (2): 268-274.

110. Eggermont, J.J. (1998). Is there a neural code? *Neuroscience & Biobehavioral Reviews* 22 (2): 355-370.

111. Eichenbaum, H. (1993). Thinking about brain cell assemblies. *Science* 261 (5124): 993-994.

112. Eldawlatly, S., and K.G. Oweiss (2011). Millisecond-timescale local network coding in the rat primary somatosensory cortex. *PLoS ONE* 6 (6): e21649.

113. Elston, G.N. (2003). Cortical heterogeneity: implications for visual processing and polysensory integration. *Journal of Neurocytology* 31 (3-5): 317-335.

114. Engel, A.K., P. Fries, and W. Singer (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience* 2 (10): 704-716.

115. Engel, A.K., P. König, C.M. Gray, and W. Singer (1990). Stimulus-dependent neuronal oscillations in cat visual cortex - intercolumnar interaction as determined by cross-correlation analysis. *European Journal of Neuroscience* 2 (7): 588-606.

116. Engel, A.K., P. König, A.K. Kreiter, T.B. Schillen, and W. Singer (1992). Temporal coding in the visual cortex - new vistas on integration in the nervous system. *Trends in Neurosciences* 15 (6): 218-226.

117. Engel, A.K., P.R. Roelfsema, P. Fries, M. Brecht, and W. Singer (1997). Role of the temporal domain for response selection and perceptual binding. *Cerebral Cortex* 7 (6): 571-582.

118. Engel, A.K., and W. Singer (2001). Temporal binding and the neural correlates of sensory awareness. *Trends in Cognitive Sciences* 5 (1): 16-25.

119. Engineer, C.T., C.A. Perez, Y.T.H. Chen, R.S. Carraway, A.C. Reed, J.A. Shetake, V. Jakkamsetti, K.Q. Chang, and M.P. Kilgard (2008). Cortical activity patterns predict speech discrimination ability. *Nature Neuroscience* 11 (5): 603-608.

120. Ermentrout, G.B., and D. Kleinfeld (2001). Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. *Neuron* 29 (1): 33-44.

121. Euston, D.R., M. Tatsuno, and B.L. McNaughton (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* 318 (5853): 1147-1150.

122. Faisal, A.A., and S.B. Laughlin (2007). Stochastic simulations on the reliability of action potential propagation in thin axons. *PLoS Computational Biology* 3 (5): 783-795.

123. Faisal, A.A., L.P.J. Selen, and D.M. Wolpert (2008). Noise in the nervous system. *Nature Reviews Neuroscience* 9 (4): 292-303.

124. Feldman, D.E. (2009). Synaptic mechanisms for plasticity in neocortex. *Annual Review of Neuroscience* 32: 33-55.

125. Fellous, J.M., A.R. Houweling, R.H. Modi, R.P.N. Rao, P.H.E. Tiesinga, and T.J. Sejnowski (2001). Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. *Journal of Neurophysiology* 85 (4): 1782-1787.

126. Fellous, J.M., M. Rudolph, A. Destexhe, and T.J. Sejnowski (2003). Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience* 122 (3): 811-829.

127. Fernández, E., J. Ferrandez, J. Ammermüller, and R.A. Normann (2000). Population coding in spike trains of simultaneously recorded retinal ganglion cells. *Brain Research* 887 (1): 222-229.

128. Ferster, D., and N. Spruston (1995). Cracking the neuronal code. *Science* 270 (5237): 756-757.

129. Fetz, E.E. (1997). Temporal coding in neural populations? *Science* 278 (5345): 1901-1902.

130. Fiete, I.R., W. Senn, C.Z.H. Wang, and R.H.R. Hahnloser (2010). Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* 65 (4): 563-576.

131. Fischer, J., and D. Whitney (2014). Serial dependence in visual perception. *Nature Neuroscience* 17 (5): 738-743.

132. Fiser, J., C. Chiu, and M. Weliky (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431 (7008): 573-578.

133. Fletcher, M.L., A.M. Smith, A.R. Best, and D.A. Wilson (2005). High-frequency oscillations are not necessary for simple olfactory discriminations in young rats. *Journal of Neuroscience* 25 (4): 792-798.

134. Foffani, G., B. Tutunculer, and K.A. Moxon (2004). Role of spike timing in the forelimb somatosensory cortex of the rat. *Journal of Neuroscience* 24 (33): 7266-7271.

135. Foster, D.J., and M.A. Wilson (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440 (7084): 680-683.

136. Fournier, J., C. Monier, M. Pananceau, and Y. Frégnac (2011). Adaptation of the simple or complex nature of V1 receptive fields to visual statistics. *Nature Neuroscience* 14 (8): 1053-1060.

137. Freiwald, W.A., A.K. Kreiter, and W. Singer (2001). Synchronization and assembly formation in the visual cortex. *Progress in Brain Research* 130: 111-140.

138. Fricker, D., and R. Miles (2001). Interneurons, spike timing, and perception. *Neuron* 32 (5): 771-774.

139. Friedrich, R.W., and G. Laurent (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science* 291 (5505): 889-894.

140. Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences* 9 (10): 474-480.

141. Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience* 32: 209-224.

142. Fries, P., D. Nikolić, and W. Singer (2007). The gamma cycle. *Trends in Neurosciences* 30 (7): 309-316.

143. Fries, P., J.H. Reynolds, A.E. Rorie, and R. Desimone (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291 (5508): 1560-1563.

144. Fries, P., T. Womelsdorf, R. Oostenveld, and R. Desimone (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *Journal of Neuroscience* 28 (18): 4823-4835.

145. Fritz, J.B., M. Elhilali, S.V. David, and S.A. Shamma (2007). Auditory attention - focusing the searchlight on sound. *Current Opinion in Neurobiology* 17 (4): 437-455.

146. Froemke, R.C., and Y. Dan (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416 (6879): 433-438.

147. Froemke, R.C., J.J. Letzkus, B.M. Kampa, G.B. Hang, and G.J. Stuart (2010). Dendritic synapse location and neocortical spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience* 2 (29): 1-14.

148. Froemke, R.C., M.M. Merzenich, and C.E. Schreiner (2007). A synaptic memory trace for cortical receptive field plasticity. *Nature* 450 (7168): 425-429.

149. Froemke, R.C., M.M. Poo, and Y. Dan (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434 (7030): 221-225.

150. Fröhlich, F., and D.A. McCormick (2010). Endogenous electric fields may guide neocortical network activity. *Neuron* 67 (1): 129-143.

151. Frostig, R.D., Z. Frostig, and R.M. Harper (1990). Recurring discharge patterns in multiple spike trains. I. Detection. *Biological Cybernetics* 62 (6): 487-493.

152. Frostig, R.D., R.C. Frysinger, and R.M. Harper (1990). Recurring discharge patterns in multiple spike trains. II. Application in forebrain areas related to cardiac and respiratory control during different sleep-waking states. *Biological Cybernetics* 62 (6): 495-502.

153. Fu, Y.X., K. Djupsund, H. Gao, B. Hayden, K. Shen, and Y. Dan (2002). Temporal specificity in the cortical plasticity of visual space representation. *Science* 296 (5575): 1999-2003.

154. Fuenzalida, M., D. Fernandez de Sevilla, and W. Buno (2007). Changes of the EPSP waveform regulate the temporal window for spike-timing-dependent plasticity. *Journal of Neuroscience* 27 (44): 11940-11948.

155. Fujii, H., H. Ito, K. Aihara, N. Ichinose, and M. Tsukada (1996). Dynamical cell assembly hypothesis - theoretical possibility of spatio-temporal coding in the cortex. *Neural Networks* 9 (8): 1303-1350.

156. Fujisawa, S., A. Amarasingham, M.T. Harrison, and G. Buzsáki (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience* 11 (7): 823-833.

157. Fukuda, T., T. Kosaka, W. Singer, and R.A. Galuske (2006). Gap junctions among dendrites of cortical GABAergic neurons establish a dense and widespread intercolumnar network. *Journal of Neuroscience* 26 (13): 3434-3443.

158. Galán, R.F., G.B. Ermentrout, and N.N. Urban (2008). Optimal time scale for spike-time reliability: theory, simulations, and experiments. *Journal of Neurophysiology* 99 (1): 277-283.

159. Galarreta, M., and S. Hestrin (1998). Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex. *Nature Neuroscience* 1 (7): 587-594.

160. Galarreta, M., and S. Hestrin (2001). Spike transmission and synchrony detection in networks of GABAergic interneurons. *Science* 292 (5525): 2295-2299.

161. Gallant, J.L., C.E. Connor, and D.C. Van Essen (1998). Neural activity in areas V1, V2 and V4 during free viewing of natural scenes compared to controlled viewing. *Neuroreport* 9 (9): 2153-2158.

162. Gautrais, J., and S.J. Thorpe (1998). Rate coding versus temporal order coding: a theoretical approach. *Biosystems* 48 (1-3): 57-65.

163. Gawne, T.J., T.W. Kjaer, J.A. Hertz, and B.J. Richmond (1996). Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cerebral Cortex* 6 (3): 482-489.

164. Gawne, T.J., and B.J. Richmond (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience* 13 (7): 2758-2771.

165. Geiger, J.R.P., J. Lübke, A. Roth, M. Frotscher, and P. Jonas (1997). Submillisecond AMPA receptor-mediated signaling at a principal neuron-interneuron synapse. *Neuron* 18 (6): 1009-1023.

166. Gelbard-Sagiv, H., R. Mukamel, M. Harel, R. Malach, and I. Fried (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. *Science* 322 (5898): 96-101.

167. Gerstein, G.L. (2004). Searching for significance in spatio-temporal firing patterns. *Acta Neurobiologiae Experimentalis* 64 (2): 203-207.

168. Gerstein, G.L., and A.M.H.J. Aertsen (1985). Representation of cooperative firing activity among simultaneously recorded neurons. *Journal of Neurophysiology* 54 (6): 1513-1528.

169. Gerstein, G.L., P. Bedenbaugh, and A.M.H.J. Aertsen (1989). Neuronal assemblies. *IEEE Transactions on Biomedical Engineering* 36 (1): 4-14.

170. Gerstein, G.L., and K.L. Kirkland (2001). Neural assemblies: technical issues, analysis, and modeling. *Neural Networks* 14 (6-7): 589-598.

171. Gerstein, G.L., D.H. Perkel, and J.E. Dayhoff (1985). Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. *Journal of Neuroscience* 5 (4): 881-889.

172. Gerstein, G.L., D.H. Perkel, and K.N. Subramanian (1978). Identification of functionally related neural assemblies. *Brain Research* 140 (1): 43-62.

173. Gerstein, G.L., E.R. Williams, M. Diesmann, S. Grün, and C. Trengove (2012). Detecting synfire chains in parallel spike data. *Journal of Neuroscience Methods* 206 (1): 54-64.

174. Gerstner, W., R. Kempter, J.L. van Hemmen, and H. Wagner (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383 (6595): 76-81.

175. Gerstner, W., A.K. Kreiter, H. Markram, and A.V. Herz (1997). Neural codes: firing rates and beyond. *Proceedings of the National Academy of Sciences of the United States of America* 94 (24): 12740-12741.

176. Gibson, J.R., M. Beierlein, and B.W. Connors (1999). Two networks of electrically coupled inhibitory neurons in neocortex. *Nature* 402 (6757): 75-79.

177. Gilbert, C.D., and W. Li (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience* 14 (5): 350-363.

178. Gollisch, T., and M. Meister (2008). Rapid neural coding in the retina with relative spike latencies. *Science* 319 (5866): 1108-1111.

179. Grammont, F., and A. Riehle (1999). Precise spike synchronization in monkey motor cortex involved in preparation for movement. *Experimental Brain Research* 128 (1-2): 118-122.

180. Gray, C.M. (1999). The temporal correlation hypothesis of visual feature integration: still alive and well. *Neuron* 24 (1): 31-25.

181. Gray, C.M., A.K. Engel, P. König, and W. Singer (1992). Synchronization of oscillatory neuronal responses in cat striate cortex - temporal properties. *Visual Neuroscience* 8 (4): 337-347.

182. Gray, C.M., and D.A. McCormick (1996). Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* 274 (5284): 109-113.

183. Greschner, M., A. Thiel, J. Kretzberg, and J. Ammermüller (2006). Complex spike-event pattern of transient ON-OFF retinal ganglion cells. *Journal of Neurophysiology* 96 (6): 2845-2856.

184. Grinstein, G., and R. Linsker (2005). Synchronous neural activity in scale-free network models versus random network models. *Proceedings of the National Academy of Sciences of the United States of America* 102 (28): 9948-9953.

185. Grothe, I., S.D. Neitzel, S. Mandon, and A.K. Kreiter (2012). Switching neuronal inputs by differential modulations of gamma-band phase-coherence. *Journal of Neuroscience* 32 (46): 16172-16180.

186. Grün, S. (2009). Data-driven significance estimation for precise spike correlation. *Journal of Neurophysiology* 101 (3): 1126-1140.

187. Grün, S., M. Diesmann, and A.M.H.J. Aertsen (2002). Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Computation* 14 (1): 43-80.

188. Grün, S., M. Diesmann, and A.M.H.J. Aertsen (2002). Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Computation* 14 (1): 81-119.

189. Grün, S., M. Diesmann, F. Grammont, A. Riehle, and A.M.H.J. Aertsen (1999). Detecting unitary events without discretization of time. *Journal of Neuroscience Methods* 94 (1): 67-79.

190. Gütig, R., A.M.H.J. Aertsen, and S. Rotter (2002). Statistical significance of coincident spikes: count-based versus rate-based statistics. *Neural Computation* 14 (1): 121-153.

191. Gütig, R., A.M.H.J. Aertsen, and S. Rotter (2003). Analysis of higher-order neuronal interactions based on conditional inference. *Biological Cybernetics* 88 (5): 352-359.

192. Guyonneau, R., R. VanRullen, and S.J. Thorpe (2005). Neurons tune to the earliest spikes through STDP. *Neural Computation* 17 (4): 859-879.

193. Hafting, T., M. Fyhn, T. Bonnevie, M.B. Moser, and E.I. Moser (2008). Hippocampus-independent phase precession in entorhinal grid cells. *Nature* 453 (7199): 1248-1252.

194. Haider, B., M.R. Krause, A. Duque, Y. Yu, J. Touryan, J.A. Mazer, and D.A. McCormick (2010). Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron* 65 (1): 107-121.

195. Haider, B., and D.A. McCormick (2009). Rapid neocortical dynamics: cellular and network mechanisms. *Neuron* 62 (2): 171-189.

196. Halpern, B.P. (2000). Sensory coding, decoding, and representations: unnecessary and troublesome constructs? *Physiology & Behavior* 69 (1-2): 115-118.

197. Harnett, M.T., J.K. Makara, N. Spruston, W.L. Kath, and J.C. Magee (2012). Synaptic amplification by dendritic spines enhances input cooperativity. *Nature* 491 (7425): 599-602.

198. Harris, K.D. (2005). Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* 6 (5): 399-407.

199. Harris, K.D., J. Csicsvari, H. Hirase, G. Dragoi, and G. Buzsáki (2003). Organization of cell assemblies in the hippocampus. *Nature* 424 (6948): 552-556.

200. Harris, K.D., D.A. Henze, H. Hirase, X. Leinekugel, G. Dragoi, A. Czurkó, and G. Buzsáki (2002). Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature* 417 (6890): 738-741.

201. Harrison, M.T., and S. Geman (2009). A rate and history-preserving resampling algorithm for neural spike trains. *Neural Computation* 21 (5): 1244-1258.

202. Hasenstaub, A., Y. Shu, B. Haider, U. Kraushaar, A. Duque, and D.A. McCormick (2005). Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* 47 (3): 423-435.

203. Hatsopoulos, N.G., S. Geman, A. Amarasingham, and E. Bienenstock (2003). At what time scale does the nervous system operate? *Neurocomputing* 52 (4): 25-29.

204. Hatsopoulos, N.G., C.L. Ojakangas, L. Paninski, and J.P. Donoghue (1998). Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proceedings of the National Academy of Sciences of the United States of America* 95 (26): 15706-15711.

205. Häusser, M., and B. Mel (2003). Dendrites: bug or feature? *Current Opinion in Neurobiology* 13 (3): 372-383.

206. Häusser, M., N. Spruston, and G.J. Stuart (2000). Diversity and dynamics of dendritic signaling. *Science* 290 (5492): 739-744.

207. Hebb, D.O. (1949). *The Organization of Behavior - A Neuropsychological Theory*. John Wiley & Sons, New York.

208. Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological Cybernetics* 82 (2): 111-121.

209. Herculano-Houzel, S. (1999). Yves Delage: Neuronal assemblies, synchronous oscillations, and Hebbian learning in 1919. *The Neuroscientist* 5 (5): 341-345.

210. Herikstad, R., J. Baker, J.P. Lachaux, C.M. Gray, and S.C. Yen (2011). Natural movies evoke spike trains with low spike time variability in cat primary visual cortex. *Journal of Neuroscience* 31 (44): 15844-15860.

211. Herrero, J.L., M.J. Roberts, L.S. Delicato, M.A. Gieselmann, P. Dayan, and A. Thiele (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 454 (7208): 1110-1114.

212. Hestrin, S. (1993). Different glutamate-receptor channels mediate fast excitatory synaptic currents in inhibitory and excitatory cortical neurons. *Neuron* 11 (6): 1083-1091.

213. Hirabayashi, T., D. Takeuchi, K. Tamura, and Y. Miyashita (2013). Microcircuits for hierarchical elaboration of object coding across primate temporal areas. *Science* 341 (6142): 191-195.

214. Hirata, Y., Y. Katori, H. Shimokawa, H. Suzuki, T.A. Blenkinsop, E.J. Lang, and K. Aihara (2008). Testing a neural coding hypothesis using surrogate data. *Journal of Neuroscience Methods* 172 (2): 312-322.

215. Histed, M.H., V. Bonin, and R.C. Reid (2009). Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. *Neuron* 63 (4): 508-522.

216. Holbro, N., Å. Grunditz, J.S. Wiegert, and T.G. Oertner (2010). AMPA receptors gate spine Ca^{2+} transients and spike-timing-dependent potentiation. *Proceedings of the National Academy of Sciences of the United States of America* 107 (36): 15975-15980.

217. Hopfield, J.J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376 (6535): 33-36.

218. Hopfield, J.J. (2004). Encoding for computation: recognizing brief dynamical patterns by exploiting effects of weak rhythms on action-potential timing. *Proceedings of the National Academy of Sciences of the United States of America* 101 (16): 6255-6260.

219. Houweling, A.R., and M. Brecht (2008). Behavioural report of single neuron stimulation in somatosensory cortex. *Nature* 451 (7174): 65-68.

220. Hromádka, T., M.R. DeWeese, and A.M. Zador (2008). Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biology* 6 (1): e16.

221. Hromádka, T., and A.M. Zador (2009). Representations in auditory cortex. *Current Opinion in Neurobiology* 19 (4): 430-433.

222. Huber, D., L. Petreanu, N. Ghitani, S. Ranade, T. Hromádka, Z.F. Mainen, and K. Svoboda (2008). Sparse optical microstimulation in barrel cortex drives learned behaviour in freely moving mice. *Nature* 451 (7174): 61-64.

223. Humphrey, G. (1939). The problem of the direction of thought. *British Journal of Psychology* 30: 183-196.

224. Hutcheon, B., R.M. Miura, and E. Puil (1996). Subthreshold membrane resonance in neocortical neurons. *Journal of Neurophysiology* 76 (2): 683-697.

225. Hutcheon, B., and Y. Yarom (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends in Neurosciences* 23 (5): 216-222.

226. Huxter, J.R., T.J. Senior, K. Allen, and J. Csicsvari (2008). Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. *Nature Neuroscience* 11 (5): 587-594.

227. Ikegaya, Y., G.B. Aaron, R. Cossart, D. Aronov, I. Lampl, D. Ferster, and R. Yuste (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304 (5670): 559-564.

228. Innocenti, G.M., P. Lehmann, and J.C. Houzel (1994). Computational structure of visual callosal axons. *European Journal of Neuroscience* 6 (6): 918-935.

229. Isaacson, J.S., and M. Scanziani (2011). How inhibition shapes cortical activity. *Neuron* 72 (2): 231-243.

230. Ishikane, H., M. Gangi, S. Honda, and M. Tachibana (2005). Synchronized retinal oscillations encode essential information for escape behavior in frogs. *Nature Neuroscience* 8 (8): 1087-1095.

231. Ivry, R.B. (1996). The representation of temporal information in perception and motor control. *Current Opinion in Neurobiology* 6 (6): 851-857.

232. Izhikevich, E.M. (2006). Polychronization: computation with spikes. *Neural Computation* 18 (2): 245-282.

233. Izhikevich, E.M., N.S. Desai, E.C. Walcott, and F.C. Hoppensteadt (2003). Bursts as a unit of neural information: selective communication via resonance. *Trends in Neurosciences* 26 (3): 161-167.

234. Jacobs, A.L., G. Fridman, R.M. Douglas, N.M. Alam, P.E. Latham, G.T. Prusky, and S. Nirenberg (2009). Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences of the United States of America* 106 (14): 5936-5941.

235. Jacobs, J., M.J. Kahana, A.D. Ekstrom, and I. Fried (2007). Brain oscillations control timing of single-neuron activity in humans. *Journal of Neuroscience* 27 (14): 3839-3844.

236. Jadhav, S.P., J. Wolfe, and D.E. Feldman (2009). Sparse temporal coding of elementary tactile features during active whisker sensation. *Nature Neuroscience* 12 (6): 792-800.

237. Jefferys, J.G.R. (1995). Nonsynaptic modulation of neuronal activity in the brain - electric currents and extracellular ions. *Physiological Reviews* 75 (4): 689-723.

238. Jensen, O. (2006). Maintenance of multiple working memory items by temporal segmentation. *Neuroscience* 139 (1): 237-249.

239. Jensen, O., and L.L. Colgin (2007). Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Sciences* 11 (7): 267-269.

240. Jermakowicz, W.J., X. Chen, I. Khaytin, A.B. Bonds, and V.A. Casagrande (2009). Relationship between spontaneous and evoked spike-time correlations in primate visual cortex. *Journal of Neurophysiology* 101 (5): 2279-2289.

241. Jia, X., M.A. Smith, and A. Kohn (2011). Stimulus selectivity and spatial coherence of gamma components of the local field potential. *Journal of Neuroscience* 31 (25): 9390-9403.

242. Johannesma, P., A.M.H.J. Aertsen, H. van den Boogaard, J.J. Eggermont, and W. Epping (1986). From synchrony to harmony: ideas on the function of neural assemblies and on the interpretation of neural synchrony. *Brain Theory*, Springer Verlag, Berlin/Heidelberg/New York, pp. 25-47.

243. Johansson, R.S., and I. Birznieks (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience* 7 (2): 170-177.

244. Jones, E.G. (2002). Thalamic circuitry and thalamocortical synchrony. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357 (1428): 1659-1673.

245. Jones, L.M., A. Fontanini, B.F. Sadacca, P. Miller, and D.B. Katz (2007). Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proceedings of the National Academy of Sciences of the United States of America* 104 (47): 18772-18777.

246. Junek, S., E. Kludt, F. Wolf, and D. Schild (2010). Olfactory coding with patterns of response latencies. *Neuron* 67 (5): 872-884.

247. Kaiser, M., and C.C. Hilgetag (2006). Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. *PLoS Computational Biology* 2 (7): e95.

248. Kapfer, C., L.L. Glickfeld, B.V. Atallah, and M. Scanziani (2007). Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex. *Nature Neuroscience* 10 (6): 743-753.

249. Kara, P., P. Reinagel, and R.C. Reid (2000). Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 27 (3): 635-646.

250. Karbowski, J. (2014). Constancy and trade-offs in the neuroanatomical and metabolic design of the cerebral cortex. *Frontiers in Neural Circuits* 8 (9): 1-16.

251. Karlsson, M.P., and L.M. Frank (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience* 12 (7): 913-918.

252. Karmarkar, U.R., and D.V. Buonomano (2007). Timing in the absence of clocks: encoding time in neural network states. *Neuron* 53 (3): 427-438.

253. Kass, R.E., R.C. Kelly, and W.-L. Loh (2011). Assessment of synchrony in multiple neural spike trains using loglinear point process models. *The Annals of Applied Statistics* 5 (2B): 1262-1292.

254. Kass, R.E., V. Ventura, and E.N. Brown (2005). Statistical issues in the analysis of neuronal data. *Journal of Neurophysiology* 94 (1): 8-25.

255. Katz, D.B., S.A. Simon, and M.A.L. Nicolelis (2002). Taste-specific neuronal ensembles in the gustatory cortex of awake rats. *Journal of Neuroscience* 22 (5): 1850-1857.

256. Katzner, S., I. Nauhaus, A. Benucci, V. Bonin, D.L. Ringach, and M. Carandini (2009). Local origin of field potentials in visual cortex. *Neuron* 61 (1): 35-41.

257. Kayser, C., M.A. Montemurro, N.K. Logothetis, and S. Panzeri (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron* 61 (4): 597-608.

258. Kerr, J.N.D., D.S. Greenberg, and F. Helmchen (2005). Imaging input and output of neocortical networks in vivo. *Proceedings of the National Academy of Sciences of the United States of America* 102 (39): 14063-14068.

259. Kim, S., D. Putrino, S. Ghosh, and E.N. Brown (2011). A Granger causality measure for point process models of ensemble neural spiking activity. *PLoS Computational Biology* 7 (3): e1001110.

260. Kimura, F. (2000). Cholinergic modulation of cortical function: a hypothetical role in shifting the dynamics in cortical network. *Neuroscience Research* 38 (1): 19-26.

261. Kimura, F., and R.W. Baughman (1997). Distinct muscarinic receptor subtypes suppress excitatory and inhibitory synaptic responses in cortical neurons. *Journal of Neurophysiology* 77 (2): 709-716.

262. Kimura, F., and C. Itami (2009). Myelination and isochronicity in neural networks. *Frontiers in Neuroanatomy* 3 (12): 1-5.

263. Klampfl, S., and W. Maass (2013). Emergence of dynamic memory traces in cortical microcircuit models through STDP. *Journal of Neuroscience* 33 (28): 11515-11529.

264. Koch, C., and I. Segev (2000). The role of single neurons in information processing. *Nature Neuroscience* 3 (supplement): 1171-1177.

265. König, P., A.K. Engel, P.R. Roelfsema, and W. Singer (1995). How precise is neuronal synchronization? *Neural Computation* 7 (3): 469-485.

266. König, P., A.K. Engel, and W. Singer (1996). Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends in Neurosciences* 19 (4): 130-137.

267. Kopell, N.J., G.B. Ermentrout, M.A. Whittington, and R.D. Traub (2000). Gamma rhythms and beta rhythms have different synchronization properties. *Proceedings of the National Academy of Sciences of the United States of America* 97 (4): 1867-1872.

268. Kopell, N.J., M.A. Kramer, P. Malerba, and M.A. Whittington (2010). Are different rhythms good for different functions? *Frontiers in Human Neuroscience* 4 (187): 1-9.

269. Kreuz, T., J.S. Haas, A. Morelli, H.D.I. Abarbanel, and A. Politi (2007). Measuring spike train synchrony. *Journal of Neuroscience Methods* 165 (1): 151-161.

270. Kupfermann, I., and K.R. Weiss (1978). The command neuron concept. *Behavioral and Brain Sciences* 1: 3-39.

271. Lago-Fernández, L.F., R. Huerta, F. Corbacho, and J.A. Sigüenza (2000). Fast response and temporal coherent oscillations in small-world networks. *Physical Review Letters* 84 (12): 2758-2761.

272. Lakatos, P., A.S. Shah, K.H. Knuth, I. Ulbert, G. Karmos, and C.E. Schroeder (2005). An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *Journal of Neurophysiology* 94 (3): 1904-1911.

273. Lalanne, C., and J. Lorenceau (2004). Crossmodal integration for perception and action. *Journal of Physiology (Paris)* 98 (1-3): 265-279.

274. Lampl, I., I. Reichova, and D. Ferster (1999). Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* 22 (2): 361-374.

275. Lampl, I., and Y. Yarom (1993). Subthreshold oscillations of the membrane potential - a functional synchronizing and timing device. *Journal of Neurophysiology* 70 (5): 2181-2186.

276. Larkum, M.E., J.J. Zhu, and B. Sakmann (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398 (6725): 338-341.

277. Laughlin, S.B., and T.J. Sejnowski (2003). Communication in neuronal networks. *Science* 301 (5641): 1870-1874.

278. Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nature Reviews Neuroscience* 3 (11): 884-895.

279. Lazar, A., G. Pipa, and J. Triesch (2009). SORN: a self-organizing recurrent neural network. *Frontiers in Computational Neuroscience* 3 (23): 1-9.

280. Lee, A.K., and M.A. Wilson (2004). A combinatorial method for analyzing sequential firing patterns involving an arbitrary number of neurons based on relative time order. *Journal of Neurophysiology* 92 (4): 2555-2573.

281. Léger, J.F., E.A. Stern, A.M.H.J. Aertsen, and D.H. Heck (2005). Synaptic integration in rat frontal cortex shaped by network activity. *Journal of Neurophysiology* 93 (1): 281-293.

282. Lennie, P. (2003). The cost of cortical computation. *Current Biology* 13 (6): 493-497.

283. Lestienne, R. (2001). Spike timing, synchronization and information processing on the sensory side of the central nervous system. *Progress in Neurobiology* 65 (6): 545-591.

284. Letzkus, J.J., B.M. Kampa, and G.J. Stuart (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *Journal of Neuroscience* 26 (41): 10420-10429.

285. Li, C.Y., M.M. Poo, and Y. Dan (2009). Burst spiking of a single cortical neuron modifies global brain state. *Science* 324 (5927): 643-646.

286. Liebe, S., G.M. Hoerzer, N.K. Logothetis, and G. Rainer (2012). Theta coupling between V4 and prefrontal cortex predicts visual short-term memory performance. *Nature Neuroscience* 15 (3): 456-462.

287. Lima, B., W. Singer, N.H. Chen, and S. Neuenschwander (2010). Synchronization dynamics in response to plaid stimuli in monkey V1. *Cerebral Cortex* 20 (7): 1556-1573.

288. Lindenberger, U., S.C. Li, W. Gruber, and V. Müller (2009). Brains swinging in concert: cortical phase synchronization while playing guitar. *BMC Neuroscience* 10 (22): 1-12.

289. Lindsey, B.G., K.F. Morris, R. Shannon, and G.L. Gerstein (1997). Repeated patterns of distributed synchrony in neuronal assemblies. *Journal of Neurophysiology* 78 (3): 1714-1719.

290. Linster, C., and T.A. Cleland (2010). Decorrelation of odor representations via spike timing-dependent plasticity. *Frontiers in Computational Neuroscience* 4 (157): 1-11.

291. Lisman, J.E., and M.A.P. Idiart (1995). Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* 267 (5203): 1512-1515.

292. Llinás, R.R. (1988). The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Science* 242 (4886): 1654-1664.

293. London, M., and M. Häusser (2005). Dendritic computation. *Annual Review of Neuroscience* 28: 503-532.

294. London, M., A. Roth, L. Beeren, M. Häusser, and P.E. Latham (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* 466 (7302): 123-127.

295. Long, M.A., D.Z. Jin, and M.S. Fee (2010). Support for a synaptic chain model of neuronal sequence generation. *Nature* 468 (7322): 394-399.

296. Lopes-dos-Santos, V., S. Conde-Ocazionez, M.A.L. Nicolelis, S.T. Ribeiro, and A.B.L. Tort (2011). Neuronal assembly detection and cell membership specification by principal component analysis. *PLoS ONE* 6 (6): e20996.

297. Lopes-dos-Santos, V., S. Ribeiro, and A.B. Tort (2013). Detecting cell assemblies in large neuronal populations. *Journal of Neuroscience Methods* 220 (2): 149-166.

298. Losonczy, A., J.K. Makara, and J.C. Magee (2008). Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature* 452 (7186): 436-441.

299. Louie, K., and M.A. Wilson (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29 (1): 145-156.

300. Luczak, A., P. Barthó, and K.D. Harris (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62 (3): 413-425.

301. Luczak, A., P. Barthó, S.L. Marguet, G. Buzsáki, and K.D. Harris (2007). Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences of the United States of America* 104 (1): 347-352.

302. Luczak, A., and J.N. MacLean (2012). Default activity patterns at the neocortical microcircuit level. *Frontiers in Integrative Neuroscience* 6 (30): 1-6.

303. Luo, H., and D. Poeppel (2007). Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron* 54 (6): 1001-1010.

304. Maass, W., T. Natschläger, and H. Markram (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Computation* 14 (11): 2531-2560.

305. MacLean, J.N., B.O. Watson, G.B. Aaron, and R. Yuste (2005). Internal dynamics determine the cortical response to thalamic stimulation. *Neuron* 48 (5): 811-823.

306. MacLeod, K., A. Bäcker, and G. Laurent (1998). Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 395 (6703): 693-698.

307. MacLeod, K., and G. Laurent (1996). Distinct mechanisms for synchronization and temporal patterning of odor-encoding neural assemblies. *Science* 274 (5289): 976-979.

308. Magee, J.C. (2000). Dendritic integration of excitatory synaptic input. *Nature Reviews Neuroscience* 1 (3): 181-190.

309. Mainen, Z.F., and T.J. Sejnowski (1995). Reliability of spike timing in neocortical neurons. *Science* 268 (5216): 1503-1506.

310. Maldonado, P.E., C. Babul, W. Singer, E. Rodriguez, D. Berger, and S. Grün (2008). Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. *Journal of Neurophysiology* 100 (3): 1523-1532.

311. Mao, B.Q., F. Hamzei-Sichani, D. Aronov, R.C. Froemke, and R. Yuste (2001). Dynamics of spontaneous activity in neocortical slices. *Neuron* 32 (5): 883-898.

312. Marder, E., L.F. Abbott, G.G. Turrigiano, Z. Liu, and J. Golowasch (1996). Memory from the dynamics of intrinsic membrane currents. *Proceedings of the National Academy of Sciences of the United States of America* 93 (24): 13481-13486.

313. Marder, E., and V. Thirumalai (2002). Cellular, synaptic and network effects of neuromodulation. *Neural Networks* 15 (4-6): 479-493.

314. Marinelli, S., S. Pacioni, A. Cannich, G. Marsicano, and A. Bacci (2009). Self-modulation of neocortical pyramidal neurons by endocannabinoids. *Nature Neuroscience* 12 (12): 1488-1490.

315. Markowitz, D.A., F. Collman, C.D. Brody, J.J. Hopfield, and D.W. Tank (2008). Rate-specific synchrony: using noisy oscillations to detect equally active neurons. *Proceedings of the National Academy of Sciences of the United States of America* 105 (24): 8422-8427.

316. Markram, H., J. Lübke, M. Frotscher, and B. Sakmann (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275 (5297): 213-215.

317. Martignon, L., G. Deco, K. Laskey, M.E. Diamond, W.A. Freiwald, and E. Vaadia (2000). Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Computation* 12 (11): 2621-2653.

318. Martinez-Conde, S., S.L. Macknik, and D.H. Hubel (2004). The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience* 5 (3): 229-240.

319. Mason, A., A. Nicoll, and K. Stratford (1991). Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. *Journal of Neuroscience* 11 (1): 72-84.

320. Masquelier, T. (2013). Neural variability, or lack thereof. *Frontiers in Computational Neuroscience* 7 (7): 1-7.

321. Masquelier, T., R. Guyonneau, and S.J. Thorpe (2008). Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS ONE* 3 (1): e1377.

322. Masquelier, T., R. Guyonneau, and S.J. Thorpe (2009). Competitive STDP-based spike pattern learning. *Neural Computation* 21 (5): 1259-1276.

323. Matsuzaki, M., G.C.R. Ellis-Davies, and H. Kasai (2008). Three-dimensional mapping of unitary synaptic connections by two-photon macro photolysis of caged glutamate. *Journal of Neurophysiology* 99 (3): 1535-1544.

324. Mauk, M.D., and D.V. Buonomano (2004). The neural basis of temporal processing. *Annual Review of Neuroscience* 27 (1): 307-340.

325. Mazurek, M.E., and M.N. Shadlen (2002). Limits to the temporal fidelity of cortical spike rate signals. *Nature Neuroscience* 5 (5): 463-471.

326. McClurkin, J.W., L.M. Optican, B.J. Richmond, and T.J. Gawne (1991). Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 253 (5020): 675-677.

327. McCormick, D.A., Y. Shu, A. Hasenstaub, M.V. Sanchez-Vives, M. Badoual, and T. Bal (2003). Persistent cortical activity: mechanisms of generation and effects on neuronal excitability. *Cerebral Cortex* 13 (11): 1219-1231.

328. McDonnell, M.D., and D. Abbott (2009). What is stochastic resonance? Definitions, misconceptions, debates, and its relevance to biology. *PLoS Computational Biology* 5 (5): e1000348.

329. McDonnell, M.D., and L.M. Ward (2011). The benefits of noise in neural systems: bridging theory and experiment. *Nature Reviews Neuroscience* 12 (7): 415-425.

330. McKenna, T.M., T.A. McMullen, and M.F. Shlesinger (1994). The brain as a dynamic physical system. *Neuroscience* 60 (3): 587-605.

331. McLelland, D., and O. Paulsen (2009). Neuronal oscillations and the rate-to-phase transform: mechanism, model and mutual information. *Journal of Physiology (London)* 587 (4): 769-785.

332. Mehta, M.R., C.A. Barnes, and B.L. McNaughton (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences of the United States of America* 94 (16): 8918-8921.

333. Meister, M. (1996). Multineuronal codes in retinal signaling. *Proceedings of the National Academy of Sciences of the United States of America* 93 (2): 609-614.

334. Meister, M., L. Lagnado, and D.A. Baylor (1995). Concerted signaling by retinal ganglion cells. *Science* 270 (5239): 1207-1210.

335. Meliza, C.D., and Y. Dan (2006). Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron* 49 (2): 183-189.

336. Metherate, R., C.L. Cox, and J.H. Ashe (1992). Cellular bases of neocortical activation: modulation of neural oscillations by the nucleus basalis and endogenous acetylcholine. *Journal of Neuroscience* 12 (12): 4701-4711.

337. Meunier, C., H. Yanai, and S. Amari (1991). Sparsely coded associative memories: capacity and dynamical properties. *Network: Computation in Neural Systems* 2 (4): 469-487.

338. Meyrand, P., J. Simmers, and M. Moulins (1991). Construction of a pattern-generating circuit with neurons of different networks. *Nature* 351 (6321): 60-63.

339. Miller, E.K., and M.A. Wilson (2008). All my circuits: using multiple electrodes to understand functioning neural networks. *Neuron* 60 (3): 483-488.

340. Miller, J.F., M. Neufang, A. Solway, A. Brandt, M. Trippel, I. Mader, S. Hefft, M. Merkow, S.M. Polyn, J. Jacobs, M.J. Kahana, and A. Schulze-Bonhage (2013). Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science* 342 (6162): 1111-1114.

341. Milner, P.M. (1974). A model for visual shape recognition. *Psychological Review* 81 (6): 521-535.

342. Minerbi, A., R. Kahana, L. Goldfeld, M. Kaufman, S. Marom, and N.E. Ziv (2009). Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. *PLoS Biology* 7 (6): e1000136.

343. Molnár, G., S. Oláh, G. Komlósi, M. Füle, J. Szabadics, C. Varga, P. Barzó, and G. Tamás (2008). Complex events initiated by individual spikes in the human cerebral cortex. *PLoS Biology* 6 (9): e222.

344. Montemurro, M.A., M.J. Rasch, Y. Murayama, N.K. Logothetis, and S. Panzeri (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Current Biology* 18 (5): 375-380.

345. Murakoshi, T. (1995). Cholinergic modulation of synaptic transmission in the rat visual cortex in vitro. *Vision Research* 35 (1): 25-35.

346. Murakoshi, T., J.Z. Guo, and T. Ichinose (1993). Electrophysiological identification of horizontal synaptic connections in rat visual cortex in vitro. *Neuroscience Letters* 163 (2): 211-214.

347. Nádasdy, Z. (2000). Spike sequences and their consequences. *Journal of Physiology (Paris)* 94 (5-6): 505-524.

348. Nádasdy, Z. (2009). Information encoding and reconstruction from the phase of action potentials. *Frontiers in Systems Neuroscience* 3 (6): 1-20.

349. Nádasdy, Z., H. Hirase, A. Czurkó, J. Csicsvari, and G. Buzsáki (1999). Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience* 19 (21): 9497-9507.

350. Nawrot, M.P., P. Schnepel, A.M.H.J. Aertsen, and C. Boucsein (2009). Precisely timed signal transmission in neocortical networks with reliable intermediate-range projections. *Frontiers in Neural Circuits* 3 (1): 1-11.

351. Nicolelis, M.A.L., A.A. Ghazanfar, C.R. Stambaugh, L.M.O. Oliveira, M. Laubach, J.K. Chapin, R.J. Nelson, and J.H. Kaas (1998). Simultaneous encoding of tactile information by three primate cortical areas. *Nature Neuroscience* 1 (7): 621-630.

352. Niessing, J., B. Ebisch, K.E. Schmidt, M. Niessing, W. Singer, and R.A. Galuske (2005). Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science* 309 (5736): 948-951.

353. Nikolić, D. (2007). Non-parametric detection of temporal order across pairwise measurements of time delays. *Journal of Computational Neuroscience* 22 (1): 5-19.

354. Nikolić, D., P. Fries, and W. Singer (2013). Gamma oscillations: precise temporal coordination without a metronome. *Trends in Cognitive Sciences* 17 (2): 54-55.

355. Nikolić, D., S. Häusler, W. Singer, and W. Maass (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biology* 7 (12): e1000260.

356. Nirenberg, S., S.M. Carcieri, A.L. Jacobs, and P.E. Latham (2001). Retinal ganglion cells act largely as independent encoders. *Nature* 411 (6838): 698-701.

357. Nordhausen, C.T., E.M. Maynard, and R.A. Normann (1996). Single unit recording capabilities of a 100 microelectrode array. *Brain Research* 726 (1-2): 129-140.

358. Norman, K.A., S.M. Polyn, G.J. Detre, and J.V. Haxby (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences* 10 (9): 424-430.

359. Nowak, L.G., and J. Bullier (1996). Spread of stimulating current in the cortical grey matter of rat visual cortex studied on a new in vitro slice preparation. *Journal of Neuroscience Methods* 67 (2): 237-248.

360. Nowak, L.G., and J. Bullier (1998). Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter. I. Evidence from chronaxie measurements. *Experimental Brain Research* 118 (4): 477-488.

361. Nowak, L.G., and J. Bullier (1998). Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter. II. Evidence from selective inactivation of cell bodies and axon initial segments. *Experimental Brain Research* 118 (4): 489-500.

362. Nowak, L.G., A.C. James, and J. Bullier (1997). Corticocortical connections between visual areas 17 and 18a of the rat studied in vitro: spatial and temporal organisation of functional synaptic responses. *Experimental Brain Research* 117 (2): 219-241.

363. Nowotny, T., M.I. Rabinovich, and H.D.I. Abarbanel (2003). Spatial representation of temporal information through spike-timing-dependent plasticity. *Physical Review E* 68 (1 Pt. 1): 011908.

364. Nusser, Z., L.M. Kay, G. Laurent, G.E. Homanics, and I. Mody (2001). Disruption of $GABA_A$ receptors on GABAergic interneurons leads to increased oscillatory power in the olfactory bulb network. *Journal of Neurophysiology* 86 (6): 2823-2833.

365. O'Keefe, J., and M.L. Recce (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3 (3): 317-330.

366. O'Neill, J., T.J. Senior, K. Allen, J.R. Huxter, and J. Csicsvari (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience* 11 (2): 209-215.

367. Oertner, T.G. (2009). How do synapses measure milliseconds? *Frontiers in Computational Neuroscience* 3 (7): 1-2.

368. Ohiorhenuan, I.E., F. Mechler, K.P. Purpura, A.M. Schmid, Q. Hu, and J.D. Victor (2010). Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* 466 (7306): 617-621.

369. Olshausen, B.A., and D.J. Field (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology* 14 (4): 481-487.

370. Padmanabhan, K., and N.N. Urban (2010). Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. *Nature Neuroscience* 13 (10): 1276-1282.

371. Panzeri, S., and M.E. Diamond (2010). Information carried by population spike times in the whisker sensory cortex can be decoded without knowledge of stimulus time. *Frontiers in Synaptic Neuroscience* 2 (17): 1-14.

372. Panzeri, S., R.S. Petersen, S.R. Schultz, M.A. Lebedev, and M.E. Diamond (2001). The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron* 29 (3): 769-777.

373. Pazienti, A., P.E. Maldonado, M. Diesmann, and S. Grün (2008). Effectiveness of systematic spike dithering depends on the precision of cortical synchronization. *Brain Research* 1225: 39-46.

374. Perkel, D.H., G.L. Gerstein, and G.P. Moore (1967). Neuronal spike trains and stochastic point processes. 2. Simultaneous spike trains. *Biophysical Journal* 7 (4): 419-440.

375. Perlovsky, L. (2013). Learning in brain and machine - complexity, Gödel, Aristotle. *Frontiers in Neurorobotics* 7 (23): 1-3.

376. Peters, A. (1985). The visual cortex of the rat. *Visual Cortex* (*Cerebral Cortex* 3), Plenum Press, New York, pp. 19-80.

377. Peters, A.J., S.X. Chen, and T. Komiyama (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* 510 (7504): 263-267.

378. Petersen, R.S., S. Panzeri, and M.E. Diamond (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron* 32 (3): 503-514.

379. Picado-Muiño, D., C. Borgelt, D. Berger, G. Gerstein, and S. Grün (2013). Finding neural assemblies with frequent item set mining. *Frontiers in Neuroinformatics* 7 (9): 1-15.

380. Pillow, J.W., J. Shlens, L. Paninski, A. Sher, A.M. Litke, E.J. Chichilnisky, and E.P. Simoncelli (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454 (7207): 995-999.

381. Pipa, G., D.W. Wheeler, W. Singer, and D. Nikolić (2008). NeuroXidence: reliable and efficient analysis of an excess or deficiency of joint-spike events. *Journal of Computational Neuroscience* 25 (1): 64-88.

382. Plenz, D., and T.C. Thiagarajan (2007). The organizing principles of neuronal avalanches: cell assemblies in the cortex? *Trends in Neurosciences* 30 (3): 101-110.

383. Poo, C., and J.S. Isaacson (2009). Odor representations in olfactory cortex: "sparse" coding, global inhibition, and oscillations. *Neuron* 62 (6): 850-861.

384. Popescu, G., A. Robert, J.R. Howe, and A. Auerbach (2004). Reaction mechanism determines NMDA receptor response to repetitive stimulation. *Nature* 430 (7001): 790-793.

385. Pouget, A., P. Dayan, and R. Zemel (2000). Information processing with population codes. *Nature Reviews Neuroscience* 1 (2): 125-132.

386. Pouille, F., and M. Scanziani (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293 (5532): 1159-1163.

387. Prescott, S.A., S. Ratté, Y. De Koninck, and T.J. Sejnowski (2006). Nonlinear interaction between shunting and adaptation controls a switch between integration and coincidence detection in pyramidal neurons. *Journal of Neuroscience* 26 (36): 9084-9097.

388. Prinz, A.A., D. Bucher, and E. Marder (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience* 7 (12): 1345-1352.

389. Prut, Y., E. Vaadia, H. Bergman, I. Haalman, H. Slovin, and M. Abeles (1998). Spatiotemporal structure of cortical activity: properties and behavioral relevance. *Journal of Neurophysiology* 79 (6): 2857-2874.

390. Quian Quiroga, R., L. Reddy, G. Kreiman, C. Koch, and I. Fried (2005). Invariant visual representation by single neurons in the human brain. *Nature* 435 (7045): 1102-1107.

391. Quilichini, P.P., and C. Bernard (2012). Brain state-dependent neuronal computation. *Frontiers in Computational Neuroscience* 6 (77): 1-7.

392. Rabinovich, M.I., R. Huerta, and G. Laurent (2008). Transient dynamics for neural processing. *Science* 321 (5885): 48-50.

393. Rabinovich, M.I., A. Volkovskii, P. Lecanda, R. Huerta, H.D.I. Abarbanel, and G. Laurent (2001). Dynamical encoding by networks of competing neuron groups: winnerless competition. *Physical Review Letters* 87 (6): 068102.

394. Radman, T., Y. Su, J.H. An, L.C. Parra, and M. Bikson (2007). Spike timing amplifies the effect of electric fields on neurons: implications for endogenous field effects. *Journal of Neuroscience* 27 (11): 3030-3036.

395. Reed, J.L., P. Pouget, H.X. Qi, Z. Zhou, M.R. Bernard, M.J. Burish, J. Haitas, A.B. Bonds, and J.H. Kaas (2008). Widespread spatial integration in primary somatosensory cortex. *Proceedings of the National Academy of Sciences of the United States of America* 105 (29): 10233-10237.

396. Reich, D.S., F. Mechler, and J.D. Victor (2001). Independent and redundant information in nearby cortical neurons. *Science* 294 (5551): 2566-2568.

397. Renart, A., J. de la Rocha, P. Barthó, L. Hollender, N. Parga, A.D. Reyes, and K.D. Harris (2010). The asynchronous state in cortical circuits. *Science* 327 (5965): 587-590.

398. Reyes, A.D., R. Lujan, A. Rozov, N. Burnashev, P. Somogyi, and B. Sakmann (1998). Target-cell-specific facilitation and depression in neocortical circuits. *Nature Neuroscience* 1 (4): 279-285.

399. Ribrault, C., K. Sekimoto, and A. Triller (2011). From the stochasticity of molecular processes to the variability of synaptic transmission. *Nature Reviews Neuroscience* 12 (7): 375-387.

400. Riehle, A., F. Grammont, M. Diesmann, and S. Grün (2000). Dynamical changes and temporal precision of synchronized spiking activity in monkey motor cortex during movement preparation. *Journal of Physiology (Paris)* 94 (5-6): 569-582.

401. Riehle, A., S. Grün, M. Diesmann, and A.M.H.J. Aertsen (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278 (5345): 1950-1953.

402. Robbe, D., S.M. Montgomery, A. Thome, P.E. Rueda-Orozco, B.L. McNaughton, and G. Buzsáki (2006). Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nature Neuroscience* 9 (12): 1526-1533.

403. Rodriguez-Molina, V.M., A.M.H.J. Aertsen, and D.H. Heck (2007). Spike timing and reliability in cortical pyramidal neurons: effects of EPSC kinetics, input synchronization and background noise on spike timing. *PLoS ONE* 2 (3): e319.

404. Roelfsema, P.R. (2006). Cortical algorithms for perceptual grouping. *Annual Review of Neuroscience* 29: 203-227.

405. Roelfsema, P.R., A.K. Engel, P. König, and W. Singer (1997). Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385 (6612): 157-161.

406. Rolston, J.D., D.A. Wagenaar, and S.M. Potter (2007). Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. *Neuroscience* 148 (1): 294-303.

407. Roopun, A.K., M.A. Kramer, L.M. Carracedo, M. Kaiser, C.H. Davies, R.D. Traub, N.J. Kopell, and M.A. Whittington (2008). Temporal interactions between cortical rhythms. *Frontiers in Neuroscience* 2 (2): 145-154.

408. Roudi, Y., S. Nirenberg, and P.E. Latham (2009). Pairwise maximum entropy models for studying large biological systems: When they can work and when they can't. *PLoS Computational Biology* 5 (5): e1000380.

409. Sakurai, Y. (1999). How do cell assemblies encode information in the brain? *Neuroscience & Biobehavioral Reviews* 23 (6): 785-796.

410. Sakurai, Y., and S. Takahashi (2006). Dynamic synchrony of firing in the monkey prefrontal cortex during working-memory tasks. *Journal of Neuroscience* 26 (40): 10141-10153.

411. Salami, M., C. Itami, T. Tsumoto, and F. Kimura (2003). Change of conduction velocity by regional myelination yields constant latency irrespective of distance between thalamus and cortex. *Proceedings of the National Academy of Sciences of the United States of America* 100 (10): 6174-6179.

412. Salinas, E., and T.J. Sejnowski (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience* 2 (8): 539-550.

413. Samonds, J.M., and A.B. Bonds (2004). Real-time visualization of neural synchrony for identifying coordinated cell assemblies. *Journal of Neuroscience Methods* 139 (1): 51-60.

414. Sanchez-Vives, M.V., and D.A. McCormick (2000). Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nature Neuroscience* 3 (10): 1027-1034.

415. Sastry, P.S., and K.P. Unnikrishnan (2010). Conditional probability-based significance tests for sequential patterns in multineuronal spike trains. *Neural Computation* 22 (4): 1025-1059.

416. Scarpetta, S., M. Yoshioka, and M. Marinaro (2008). Encoding and replay of dynamic attractors with multiple frequencies: analysis of a STDP based learning rule. *Dynamic Brain - From Neural Spikes To Behavior* (*Lecture Notes in Computer Science* 5286), Springer Verlag, Berlin/Heidelberg, pp. 38-60.

417. Schaefer, A.T., K. Angelo, H. Spors, and T.W. Margrie (2006). Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. *PLoS Biology* 4 (6): e163.

418. Schiller, J., Y. Schiller, and D.E. Clapham (1998). NMDA receptors amplify calcium influx into dendritic spines during associative pre- and postsynaptic activation. *Nature Neuroscience* 1 (2): 114-118.

419. Schneider, G., and S. Grün (2003). Analysis of higher-order correlations in multiple parallel processes. *Neurocomputing* 52 (4): 771-777.

420. Schneider, G., M.N. Havenith, and D. Nikolić (2006). Spatiotemporal structure in large neuronal networks detected from cross-correlation. *Neural Computation* 18 (10): 2387-2413.

421. Schnitzer, M.J., and M. Meister (2003). Multineuronal firing patterns in the signal from eye to brain. *Neuron* 37 (3): 499-511.

422. Schrader, S., S. Grün, M. Diesmann, and G.L. Gerstein (2008). Detecting synfire chain activity using massively parallel spike train recording. *Journal of Neurophysiology* 100 (4): 2165-2176.

423. Schreiber, S., I. Erchova, U. Heinemann, and A.V. Herz (2004). Subthreshold resonance explains the frequency-dependent integration of periodic as well as random stimuli in the entorhinal cortex. *Journal of Neurophysiology* 92 (1): 408-415.

424. Segundo, J.P., T.H. Bullock, G.P. Moore, and L.J. Stensaas (1963). Sensitivity of neurones in Aplysia to temporal pattern of arriving impulses. *Journal of Experimental Biology* 40 (4): 643-667.

425. Sejnowski, T.J., and O. Paulsen (2006). Network oscillations: emerging computational principles. *Journal of Neuroscience* 26 (6): 1673-1676.

426. Senior, T.J., J.R. Huxter, K. Allen, J. O'Neill, and J. Csicsvari (2008). Gamma oscillatory firing reveals distinct populations of pyramidal cells in the CA1 region of the hippocampus. *Journal of Neuroscience* 28 (9): 2274-2286.

427. Serences, J.T., and S. Yantis (2006). Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences* 10 (1): 38-45.

428. Shadlen, M.N., and J.A. Movshon (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 24 (1): 67-77.

429. Shadlen, M.N., and W.T. Newsome (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology* 4 (4): 569-579.

430. Shadlen, M.N., and W.T. Newsome (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience* 18 (10): 3870-3896.

431. Shahaf, G., D. Eytan, A. Gal, E. Kermany, V. Lyakhov, C. Zrenner, and S. Marom (2008). Order-based representation in random networks of cortical neurons. *PLoS Computational Biology* 4 (11): e1000228.

432. Shaw, G.L., E. Harth, and A.B. Scheibel (1982). Cooperativity in brain function - assemblies of approximately 30 neurons. *Experimental Neurology* 77 (2): 324-358.

433. Shimazaki, H., S. Amari, E.N. Brown, and S. Grün (2012). State-space analysis of time-varying higher-order spike correlation for multiple neural spike train data. *PLoS Computational Biology* 8 (3): e1002385.

434. Shin, C.W., and S. Kim (2006). Self-organized criticality and scale-free properties in emergent functional neural networks. *Physical Review E* 74 (4 Pt. 2): 045101.

435. Shipp, S. (2003). The functional logic of cortico-pulvinar connections. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 358 (1438): 1605-1624.

436. Shmiel, T., R. Drori, O. Shmiel, Y. Ben Shaul, Z. Nádasdy, M. Shemesh, M. Teicher, and M. Abeles (2005). Neurons of the cerebral cortex exhibit precise interspike timing in correspondence to behavior. *Proceedings of the National Academy of Sciences of the United States of America* 102 (51): 18655-18657.

437. Shmiel, T., R. Drori, O. Shmiel, Y. Ben Shaul, Z. Nádasdy, M. Shemesh, M. Teicher, and M. Abeles (2006). Temporally precise cortical firing patterns are associated with distinct action segments. *Journal of Neurophysiology* 96 (5): 2645-2652.

438. Shu, Y., A. Hasenstaub, M. Badoual, T. Bal, and D.A. McCormick (2003). Barrages of synaptic activity control the gain and sensitivity of cortical neurons. *Journal of Neuroscience* 23 (32): 10388-10401.

439. Shu, Y., A. Hasenstaub, A. Duque, Y. Yu, and D.A. McCormick (2006). Modulation of intracortical synaptic potentials by presynaptic somatic membrane potential. *Nature* 441 (7094): 761-765.

440. Shuler, M.G., and M.F. Bear (2006). Reward timing in the primary visual cortex. *Science* 311 (5767): 1606-1609.

441. Siegel, M., M.R. Warden, and E.K. Miller (2009). Phase-dependent neuronal coding of objects in short-term memory. *Proceedings of the National Academy of Sciences of the United States of America* 106 (50): 21341-21346.

442. Sillito, A.M., and H.E. Jones (2002). Corticothalamic interactions in the transfer of visual information. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357 (1428): 1739-1752.

443. Sillito, A.M., H.E. Jones, G.L. Gerstein, and D.C. West (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature* 369 (6480): 479-482.

444. Silver, R.A. (2010). Neuronal arithmetic. *Nature Reviews Neuroscience* 11 (7): 474-489.

445. Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology* 55: 349-374.

446. Singer, W. (1996). The changing face of inhibition. *Current Biology* 6 (4): 395-397.

447. Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24 (1): 49-65.

448. Singer, W. (1999). Striving for coherence. *Nature* 397 (6718): 391-392.

449. Singer, W. (1999). Time as coding space? *Current Opinion in Neurobiology* 9 (2): 189-194.

450. Singer, W. (2000). Why use more than one electrode at a time? *Trends in Neurosciences* Dec. 2000 (special issue): 12-17.

451. Singer, W. (2001). Consciousness and the binding problem. *Annals of the New York Academy of Sciences* 929: 123-146.

452. Singer, W. (2007). Binding by synchrony. *Scholarpedia* 2 (12): 1657.

453. Singer, W. (2007). Large scale temporal coordination of cortical activity as prerequisite for conscious experience. *The Blackwell Companion to Consciousness*, Blackwell Publishing, Malden/MA, pp. 605-615.

454. Singer, W. (2007). Understanding the brain. *EMBO Reports* 8 (special issue): S16-S19.

455. Singer, W. (2009). The brain, a complex self-organizing system. *European Review* 17 (2): 321-329.

456. Singer, W. (2013). Cortical dynamics revisited. *Trends in Cognitive Sciences* 17 (12): 616-626.

457. Singer, W., A.K. Engel, A.K. Kreiter, M.H.J. Munk, S. Neuenschwander, and P.R. Roelfsema (1997). Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences* 1 (7): 252-260.

458. Singer, W., and C.M. Gray (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* 18: 555-586.

459. Singer, W., A.K. Kreiter, P. Fries, S. Herculano, M. Volgushev, and S. Neuenschwander (1998). The putative role of precise timing in cortical processing. *Basic Mechanisms in Cognition and Language*, Elsevier, Amsterdam, pp. 63-85.

460. Sjöström, P.J., and M. Häusser (2006). A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. *Neuron* 51 (2): 227-238.

461. Skaggs, W.E., B.L. McNaughton, M.A. Wilson, and C.A. Barnes (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6 (2): 149-172.

462. Smith, A.C., V.K. Nguyen, M.P. Karlsson, L.M. Frank, and P. Smith (2010). Probability of repeating patterns in simultaneous neural data. *Neural Computation* 22 (10): 2522-2536.

463. Smith, A.C., and P. Smith (2006). A set probability technique for detecting relative time order across multiple neurons. *Neural Computation* 18 (5): 1197-1214.

464. Smith, M.A., and A. Kohn (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience* 28 (48): 12591-12603.

465. Softky, W.R., and C. Koch (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience* 13 (1): 334-350.

466. Sohal, V.S., F. Zhang, O. Yizhar, and K. Deisseroth (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature* 459 (7247): 698-702.

467. Song, S., K.D. Miller, and L.F. Abbott (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience* 3 (9): 919-926.

468. Song, S., P.J. Sjöström, M. Reigl, S.B. Nelson, and D.B. Chklovskii (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology* 3 (3): e68.

469. Sporns, O., and R. Kötter (2004). Motifs in brain networks. *PLoS Biology* 2 (11): e369.

470. Sporns, O., G. Tononi, and G.M. Edelman (2000). Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral Cortex* 10 (2): 127-141.

471. Sporns, O., and J.D. Zwi (2004). The small world of the cerebral cortex. *Neuroinformatics* 2 (2): 145-162.

472. Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience* 9 (3): 206-221.

473. Stark, E., and M. Abeles (2005). Applying resampling methods to neurophysiological data. *Journal of Neuroscience Methods* 145 (1-2): 133-144.

474. Staude, B., S. Grün, and S. Rotter (2010). Higher-order correlations in non-stationary parallel spike trains: statistical modeling and inference. *Frontiers in Computational Neuroscience* 4 (16): 1-17.

475. Stevens, C.F., and A.M. Zador (1998). Input synchrony and the irregular firing of cortical neurons. *Nature Neuroscience* 1 (3): 210-217.

476. Stonehouse, J.M., and G.J. Forrester (1998). Robustness of the t and U tests under combined assumption violations. *Journal of Applied Statistics* 25 (1): 63-74.

477. Stopfer, M., S. Bhagavan, B.H. Smith, and G. Laurent (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390 (6655): 70-74.

478. Stuart, G.J., and M. Häusser (2001). Dendritic coincidence detection of EPSPs and action potentials. *Nature Neuroscience* 4 (1): 63-71.

479. Stuart, G.J., N. Spruston, B. Sakmann, and M. Häusser (1997). Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends in Neurosciences* 20 (3): 125-131.

480. Sun, J.J., and H.J. Luhmann (2007). Spatio-temporal dynamics of oscillatory network activity in the neonatal mouse cerebral cortex. *European Journal of Neuroscience* 26 (7): 1995-2004.

481. Sun, W., and Y. Dan (2009). Layer-specific network oscillation and spatiotemporal receptive field in the visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* 106 (42): 17986-17991.

482. Suri, R.E., and T.J. Sejnowski (2002). Spike propagation synchronized by temporally asymmetric Hebbian learning. *Biological Cybernetics* 87 (5-6): 440-445.

483. Takahashi, N., K. Kitamura, N. Matsuo, M. Mayford, M. Kano, N. Matsuki, and Y. Ikegaya (2012). Locally synchronized synaptic inputs. *Science* 335 (6066): 353-356.

484. Tao, H.Z.W., L.I. Zhang, G.Q. Bi, and M.M. Poo (2000). Selective presynaptic propagation of long-term potentiation in defined neural networks. *Journal of Neuroscience* 20 (9): 3233-3243.

485. Tetko, I.V., and A.E.P. Villa (1997). Fast combinatorial methods to estimate the probability of complex temporal patterns of spikes. *Biological Cybernetics* 76 (5): 397-407.

486. Tetko, I.V., and A.E.P. Villa (2001). A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 1. Detection of repeated patterns. *Journal of Neuroscience Methods* 105 (1): 1-14.

487. Tetko, I.V., and A.E.P. Villa (2001). A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 2. Application to simultaneous single unit recordings. *Journal of Neuroscience Methods* 105 (1): 15-24.

488. Theunissen, F.E., and J.P. Miller (1995). Temporal encoding in nervous systems - a rigorous definition. *Journal of Computational Neuroscience* 2 (2): 149-162.

489. Thomson, A.M. (2003). Presynaptic frequency- and pattern-dependent filtering. *Journal of Computational Neuroscience* 15 (2): 159-202.

490. Thomson, A.M., A.P. Bannister, A. Mercer, and O.T. Morris (2002). Target and temporal pattern selection at neocortical synapses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357 (1428): 1781-1791.

491. Thorpe, S.J., D. Fize, and C. Marlot (1996). Speed of processing in the human visual system. *Nature* 381 (6582): 520-522.

492. Tiesinga, P.H.E., J.M. Fellous, and T.J. Sejnowski (2008). Regulation of spike timing in visual cortical circuits. *Nature Reviews Neuroscience* 9 (2): 97-109.

493. Tønnesen, J., G. Katona, B. Rózsa, and U.V. Nägerl (2014). Spine neck plasticity regulates compartmentalization of synapses. *Nature Neuroscience* 17 (5): 678-685.

494. Tononi, G., G.M. Edelman, and O. Sporns (1998). Complexity and coherency: integrating information in the brain. *Trends in Cognitive Sciences* 2 (12): 474-484.

495. Torre, E., D. Picado-Muiño, M. Denker, C. Borgelt, and S. Grün (2013). Statistical evaluation of synchronous spike patterns extracted by frequent item set mining. *Frontiers in Computational Neuroscience* 7 (132): 1-13.

496. Tovée, M.J. (1994). How fast is the speed of thought? *Current Biology* 4 (12): 1125-1127.

497. Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology* 6 (2): 171-178.

498. Truccolo, W., L.R. Hochberg, and J.P. Donoghue (2010). Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nature Neuroscience* 13 (1): 105-111.

499. Tsodyks, M.V., T. Kenet, A. Grinvald, and A. Arieli (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286 (5446): 1943-1946.

500. Tsodyks, M.V., and H. Markram (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences of the United States of America* 94 (2): 719-723.

501. Uhlhaas, P.J., G. Pipa, B. Lima, L. Melloni, S. Neuenschwander, D. Nikolić, and W. Singer (2009). Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience* 3 (17): 1-19.

502. Ulrich, D. (2002). Dendritic resonance in rat neocortical pyramidal cells. *Journal of Neurophysiology* 87 (6): 2753-2759.

503. Usrey, W.M. (2002). Spike timing and visual processing in the retinogeniculocortical pathway. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357 (1428): 1729-1737.

504. Usrey, W.M., and R.C. Reid (1999). Synchronous activity in the visual system. *Annual Review of Physiology* 61: 435-456.

505. Vaadia, E., I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A.M.H.J. Aertsen (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373 (6514): 515-518.

506. van Aerde, K.I., E.O. Mann, C.B. Canto, T.S. Heistek, K. Linkenkaer-Hansen, A.B. Mulder, M. van der Roest, O. Paulsen, A.B. Brussaard, and H.D. Mansvelder (2009). Flexible spike timing of layer 5 neurons during dynamic beta oscillation shifts in rat prefrontal cortex. *Journal of Physiology (London)* 587 (21): 5177-5196.

507. Van Essen, D.C. (2007). Neuroinformatics - What's in it for you? *Neuroscience Quarterly* Spring 2007: online.

508. Van Hooser, S.D. (2007). Similarity and diversity in visual cortex: Is there a unifying theory of cortical computation? *The Neuroscientist* 13 (6): 639-656.

509. VanRullen, R., R. Guyonneau, and S.J. Thorpe (2005). Spike times make sense. *Trends in Neurosciences* 28 (1): 1-4.

510. VanRullen, R., and S.J. Thorpe (2001). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Computation* 13 (6): 1255-1283.

511. Varela, F.J., J.P. Lachaux, E. Rodriguez, and J. Martinerie (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience* 2 (4): 229-239.

512. Varela, J.A., K. Sen, J.R. Gibson, J. Fost, L.F. Abbott, and S.B. Nelson (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *Journal of Neuroscience* 17 (20): 7926-7940.

513. Varga, A.G., and D.W. Wesson (2013). Distributed auditory sensory input within the mouse olfactory cortex. *European Journal of Neuroscience* 37 (4): 564-571.

514. Vicente, R., L.L. Gollo, C.R. Mirasso, I. Fischer, and G. Pipa (2008). Dynamical relaying can yield zero time lag neuronal synchrony despite long conduction delays. *Proceedings of the National Academy of Sciences of the United States of America* 105 (44): 17157-17162.

515. Villa, A.E.P., and M. Abeles (1990). Evidence for spatiotemporal firing patterns within the auditory thalamus of the cat. *Brain Research* 509 (2): 325-327.

516. Villa, A.E.P., I.V. Tetko, B. Hyland, and A. Najem (1999). Spatiotemporal activity patterns of rat cortical neurons predict responses in a conditioned task. *Proceedings of the National Academy of Sciences of the United States of America* 96 (3): 1106-1111.

517. Vinck, M., B. Lima, T. Womelsdorf, R. Oostenveld, W. Singer, S. Neuenschwander, and P. Fries (2010). Gamma-phase shifting in awake monkey visual cortex. *Journal of Neuroscience* 30 (4): 1250-1257.

518. Vinje, W.E., and J.L. Gallant (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287 (5456): 1273-1276.

519. Viriyopase, A., I. Bojak, M. Zeitler, and S. Gielen (2012). When long-range zero-lag synchronization is feasible in cortical networks. *Frontiers in Computational Neuroscience* 6 (49): 1-22.

520. Vogels, T.P., and L.F. Abbott (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature Neuroscience* 12 (4): 483-491.

521. Volgushev, M., M. Chistiakova, and W. Singer (1998). Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. *Neuroscience* 83 (1): 15-25.

522. von der Malsburg, C. (1986). Am I thinking assemblies? *Brain Theory*, Springer Verlag, Berlin/Heidelberg/New York, pp. 161-176.

523. von der Malsburg, C. (1995). Binding in models of perception and brain function. *Current Opinion in Neurobiology* 5 (4): 520-526.

524. von der Malsburg, C., and W. Schneider (1986). A neural cocktail-party processor. *Biological Cybernetics* 54 (1): 29-40.

525. Voronin, L.L., M. Volgushev, M. Chistiakova, U. Kuhnt, and W. Singer (1996). Involvement of silent synapses in the induction of long-term potentiation and long-term depression in neocortical and hippocampal neurons. *Neuroscience* 74 (2): 323-330.

526. Wang, H.P., D. Spencer, J.M. Fellous, and T.J. Sejnowski (2010). Synchrony of thalamocortical inputs maximizes cortical reliability. *Science* 328 (5974): 106-109.

527. Wang, X.J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews* 90 (3): 1195-1268.

528. Wang, Z., and D.A. McCormick (1993). Control of firing mode of corticotectal and corticopontine layer V burst-generating neurons by norepinephrine, acetylcholine, and 1S,3R-ACPD. *Journal of Neuroscience* 13 (5): 2199-2216.

529. Waters, J., and F. Helmchen (2006). Background synaptic activity is sparse in neocortex. *Journal of Neuroscience* 26 (32): 8267-8277.

530. Watts, D.J., and S.H. Strogatz (1998). Collective dynamics of 'small-world' networks. *Nature* 393 (6684): 440-442.

531. Wehr, M., and G. Laurent (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384 (6605): 162-166.

532. Weinberger, N.M. (1995). Dynamic regulation of receptive fields and maps in the adult sensory cortex. *Annual Review of Neuroscience* 18: 129-158.

533. Weiss, S.A., and D.S. Faber (2010). Field effects in the CNS play functional roles. *Frontiers in Neural Circuits* 4 (15): 1-10.

534. Weliky, M., J. Fiser, R.H. Hunt, and D.N. Wagner (2003). Coding of natural scenes in primary visual cortex. *Neuron* 37 (4): 703-718.

535. Wespatat, V., F. Tennigkeit, and W. Singer (2004). Phase sensitivity of synaptic modifications in oscillating cells of rat visual cortex. *Journal of Neuroscience* 24 (41): 9067-9075.

536. Wheeler, B.C. (1999). Automatic discrimination of single units. *Methods for Neural Ensemble Recordings* (*Frontiers in Neuroscience* 1), CRC Press, Boca Raton, pp. 61-77.

537. Whittington, M.A., R.D. Traub, N.J. Kopell, G.B. Ermentrout, and E.H. Buhl (2000). Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *International Journal of Psychophysiology* 38 (3): 315-336.

538. Williams, S.R., and S.E. Atkinson (2007). Pathway-specific use-dependent dynamics of excitatory synaptic transmission in rat intracortical circuits. *Journal of Physiology (London)* 585 (3): 759-777.

539. Wolfe, J., A.R. Houweling, and M. Brecht (2010). Sparse and powerful cortical spikes. *Current Opinion in Neurobiology* 20 (3): 306-312.

540. Womelsdorf, T., K. Anton-Erxleben, F. Pieper, and S. Treue (2006). Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nature Neuroscience* 9 (9): 1156-1160.

541. Womelsdorf, T., J.M. Schoffelen, R. Oostenveld, W. Singer, R. Desimone, A.K. Engel, and P. Fries (2007). Modulation of neuronal interactions through neuronal synchronization. *Science* 316 (5831): 1609-1612.

542. Wörgötter, F., K. Suder, Y.Q. Zhao, N. Kerscher, U.T. Eysel, and K. Funke (1998). State-dependent receptive-field restructuring in the visual cortex. *Nature* 396 (6707): 165-168.

543. Wu, C., W.P. Luk, J. Gillis, F. Skinner, and L. Zhang (2005). Size does matter: generation of intrinsic network rhythms in thick mouse hippocampal slices. *Journal of Neurophysiology* 93 (4): 2302-2317.

544. Xiang, Z., J.R. Huguenard, and D.A. Prince (1998). Cholinergic switching within neocortical inhibitory networks. *Science* 281 (5379): 985-988.

545. Xu, S., W. Jiang, M.M. Poo, and Y. Dan (2012). Activity recall in a visual cortical ensemble. *Nature Neuroscience* 15 (3): 449-455.

546. Yamada, S., K. Matsumoto, M. Nakashima, and S. Shiono (1996). Information theoretic analysis of action potential trains. II. Analysis of correlation among n neurons to deduce connection structure. *Journal of Neuroscience Methods* 66 (1): 35-45.

547. Yamada, S., M. Nakashima, K. Matsumoto, and S. Shiono (1993). Information theoretic analysis of action potential trains. I. Analysis of correlation between two neurons. *Biological Cybernetics* 68 (3): 215-220.

548. Yang, Y., M.R. DeWeese, G.H. Otazu, and A.M. Zador (2008). Millisecond-scale differences in neural activity in auditory cortex can drive decisions. *Nature Neuroscience* 11 (11): 1262-1263.

549. Yassin, L., B.L. Benedetti, J.S. Jouhanneau, J.A. Wen, J.F.A. Poulet, and A.L. Barth (2010). An embedded subnetwork of highly active neurons in the neocortex. *Neuron* 68 (6): 1043-1050.

550. Yeh, C., D. Xing, P.E. Williams, and R.M. Shapley (2009). Stimulus ensemble and cortical layer determine V1 spatial receptive fields. *Proceedings of the National Academy of Sciences of the United States of America* 106 (34): 14652-14657.

551. Yen, S.C., J. Baker, and C.M. Gray (2007). Heterogeneity in the responses of adjacent neurons to natural stimuli in cat striate cortex. *Journal of Neurophysiology* 97 (2): 1326-1341.

552. Yizhar, O., L.E. Fenno, T.J. Davidson, M.Z. Mogri, and K. Deisseroth (2011). Optogenetics in neural systems. *Neuron* 71 (1): 9-34.

553. Yoshimura, Y., J.L. Dantzker, and E.M. Callaway (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature* 433 (7028): 868-873.

554. Yu, S., D. Huang, W. Singer, and D. Nikolić (2008). A small world of neuronal synchrony. *Cerebral Cortex* 18: 2891-2901.

Publications

Gansel, K., and W. Singer (2005). Replay of second-order spike patterns with millisecond precision in the visual cortex. *Society for Neuroscience Abstracts* 2005: 276.8, Washington/DC.

Gansel, K., and W. Singer (2006). Detection and validation of first- and second-order multineuronal spike patterns. *FENS Abstracts* 3: A216.7, Wien.

Gansel, K., and W. Singer (2006). Recurring second-order spike patterns concur with increased symmetry in functional neuronal couplings. *Society for Neuroscience Abstracts* 2006: 634.15, Atlanta/GA.

Gansel, K., and W. Singer (2007). Timescale-dependent criticality of repeating spatiotemporal spike patterns. *Proceedings of the 7th Meeting of the German Neuroscience Society – 31st Göttingen Neurobiology Conference*: T16-4B, Göttingen.

Gansel, K., and W. Singer (2007). Repeating spatiotemporal spike patterns reflect functional network states in the visual cortex. *Proceedings of the 1st NCCD Meeting*, Hossegor.

Gansel, K., and W. Singer (2008). Motifs of spiking network activity in visual cortex recur predominantly with sub-millisecond precision. *FENS Abstracts 4*: 153.4, Geneva.

Gansel, K., and W. Singer (2008). Temporal organization and time compression of recurring spike sequences in the visual cortex. *Society for Neuroscience Abstracts* 2008: 568.12, Washington/DC.

Gansel, K., and W. Singer (2009). Independence of functional neuronal network architecture from cholinergic and GABAergic modulation: No escape from the small world. *Proceedings of the 8th Meeting of the German Neuroscience Society – 32nd Göttingen Neurobiology Conference*: T23-1B, Göttingen.

Gansel, K.S., and W. Singer (2012). Detecting multineuronal temporal patterns in parallel spike trains. *Frontiers in Neuroinformatics* 6(18): 1-16.

Nikolić, D., and K. Gansel (2012). Der Kontext entscheidet. *Gehirn und Geist* (5): 20-23.

Gansel, K., and W. Singer (2013). Acquisition of multineuronal spike events from brain slices. *Proceedings of the 10th Meeting of the German Neuroscience Society – 34th Göttingen Neurobiology Conference*: T27-1B, Göttingen.

Sigalas, C., K. Gansel, W. Singer and I. Skaliora (2014). Spatiotemporal propagation patterns of cortical synchronised activity in vitro. *Proceedings of AREADNE*, Santorini.

Gansel, K.S., and W. Singer (2014). Organization of cell assemblies in the neocortex. *Neuron* (submitted).

Curriculum Vitae

Kai S. Gansel
born on December 23, 1974 in Wangen im Allgäu, Germany

Education & academic career

since 2001	Research associate, Max Planck Institute for Brain Research, Frankfurt am Main Department of Neurophysiology
2001-2004	Studies in Medicine, University Frankfurt am Main
1995-2001	Studies in Biology and Psychology, University Oldenburg Diploma in Biology
1994	Abitur, Gymnasium Bad Zwischenahn

Fellowships & prizes

2001-2004	Fellowship, Research Training Group "Neuronal plasticity", funded by the German Research Foundation
2001	Invitation to the young scientists meeting “Junge Elite” (CeBIT, Hannover) and travel grant from the Nord/LB Bank
2000	Peter-Waskönig-Prize
2000	OLB/EWE Prize
1995-2000	Scholarship from the Erich Bruns Foundation

Memberships

Neurowissenschaftliche Gesellschaft
Federation of European Neuroscience Societies

Invited talks

2013	BioMediTech, Tampere University of Technology: "Cortical dynamics revisited: new avenues of unraveling the neural code"
2008	Institute of Physiology, University Mainz: "Dynamic formation of spatiotemporal spike patterns in rat visual cortex in vitro"
2007	Max Planck Institute for Dynamics and Self-Organization, Göttingen: "Timescale-dependent criticality of repeating spatiotemporal spike patterns"

Hiermit erkläre ich, dass ich die Arbeit selbständig verfasst und keine anderen als die angegebenen Hilfsmittel verwendet habe.

Kai Gansel, Oktober 2014