Intrinsic and Functional Aspects of Neuronal Synchrony in Primary Visual Cortex

Dissertation
zur Erlangung des akademischen Grades des Doktors der Naturwissenschaften (Dr. rer. nat.)

Eingereicht im Fachbereich Biologie, Chemie, Pharmazie der Freien Universität Berlin

vorgelegt von

Denise Berger

aus Berlin

Juni 2009

The research presented in this thesis was carried out from December 2005 until May 2009 at the Institute for Biology, Neurobiology of the Freie University Berlin.

1. Gutachter: Prof. Dr. Martin Nawrot

2. Gutachter: PD Dr. Sonja Grün

Disputation am 03.07.2009

Declaration

Chapters 1, 2, and 3 have been published in peer-reviewed international journals. Chapter 4 is accepted in a peer-reviewed international journal. Chapter 5 will be submitted to peer-reviewed international journals.

The contributions of individual authors were as follows:

Chapters 1 is published as:

> Berger, D.; Warren, D.; Normann, R.; Arieli, A. & Grün, S.
> Spatially organized spike correlation in cat visual cortex.
> Neurocomput, 2007, 70, 2112-2116.

The work has been carried out by D. Berger, except for the experiments which were performed in the laboratory of R. Normann. A. Arieli contributed in the discussions. S. Grün supervised the project.

Chapters 2 is published as:

> Maldonado, P.; Babul, C.; Singer, W.; Rodriguez, E.; Berger, D. & Grün, S.
> Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images.
> J. Neurophysiol., 2008, 100, 1523-1532.

The work was mainly carried out by P. Maldonado and S. Grün. D. Berger designed the method for the detection of eye events. W. Singer and E. Rodriguez contributed in the discussions.

Chapters 3 is published as:

> Grün, S.; Berger, D. & Borgelt, C.
> Identification of neurons participating in cell assemblies
> Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2009, 3493-3496.

The work originated from discussions between D. Berger and the other authors. C. Borgelt proposed the estimation procedure. S. Grün performed the simulations.

Chapters 4 will appear as:

> Berger, D.; Borgelt, C.; Morrison, A. & Grün, S.
> Efficient Identification of Assembly Neurons within Massively Parallel Spike Trains.
> in Computational Intelligence in Neuroscience

Data analyses were carried out by D. Berger. C. Borgelt proposed the estimation procedure. S. Grün supervised the project. Developments of surrogates were developed by D. Berger, C. Borgelt, and S. Grün. A. Morrison performed the network simulations.

Chapters 5 is to be submitted as:

> Berger, D.; Nawrot, M.; Maldonado, P. & Grün, S.
> Statistics of eye movements of monkeys freely viewing natural scenes.

The work has been carried out by D. Berger, except for the experiments which were performed in the laboratory of P. Maldonado. S. Grün and M. Nawrot supervised the project.

Summary

Most neurons in V1 are topologically organized in cortical columns, that coherently respond to a particular, preferred stimulus by increasing their discharge rate. From a statistical point of view, an increase of the frequency of action potentials is not sufficient to explain how neurons in V1 are able to bring together and compute complex stimuli. Modern theories of the primary visual pathway account for spatial and temporal aspects. In particular, the cell assembly hypothesis postulates dynamically interacting groups of neurons as the key mechanism for cortical information processing. As massively parallel extracellular recordings are becoming available, the limited experimental evidence in favor of the temporal coding hypothesis has to be assigned to a large extent to a lack of suitable analysis tools. In this thesis I summarize a work to study the activity of the visual cortex during natural viewing conditions. I, moreover, present new analysis strategies that allow to study the activity of many neurons simultaneously and observe coordinated network activity.

In the first part of this doctoral work results of pairwise correlation analysis from simultaneously recorded single units in V1 of free viewing monkeys are related to their eye events, and simultaneously recorded multi-unit activities from visual cortex area 17 of anesthetized cats are matched with the corresponding orientation maps. The questions I answer in that context are: how can we relate the concept of orientation maps from optical imaging and the temporal coding hypothesis? What is the relationship between correlated neuronal activity and 1. spatial arrangement of the recording sites, and 2. viewing behavior?

Mere pairwise tests are not suitable to conclusively identify assemblies. Existing analysis techniques for higher-order synchrony fail in estimating the necessary parameters, whose number grows exponentially with the number of neurons. Therefore, I designed two novel methods for identifying neurons participating in assembly activity. Both methods enable fast and reliable detection of neurons which participate in correlations.

This work provides new insights into the coding strategy employed by neurons of the primary visual cortex. The demonstrated locking of neuronal synchronization to the time of eye fixation, and the evidence that the spatial arrangement of synchronized neurons is in agreement with orientation tuning maps strongly support the hypothesis that primary visual cortex neurons encode visual information via precisely timed spikes. This work sets also the bases for further tests of the temporal coding hypothesis, by introducing two extremely efficient methods for identification of higher-order synchrony.

Finally, the analysis of eye movement and natural viewing strategies provides novel ap-

proaches to data from free viewing animals and thus open new routes for experiments that aim to relate neuronal activity to natural behavior.

Contents

Introduction

Most of our cognitive functions and perceptual processes are carried out by the neocortex, which is the largest structure of the human brain. Among all somatosensorial areas of the central nervous system, the visual cortex has the most complex neural circuitry. About half of the cortex is involved in visual processing as indicated from monkey studies (Tong, 2003). The primary visual cortex (V1 or area 17) is the most extensively studied and perhaps the best understood area of the whole cerebral cortex. It is the first cortical area to receive visual information from the environment, and is responsible for the initial decomposition and integration of this input. About 90% of projections from the retina are channeled through the lateral geniculate nucleus (LGN), a part of the thalamus, and evetually reach V1. From V1 information is disseminated to various extrastriate visual areas, including areas V2, V3, V4, responsible for further processing.

Receptive Fields Properties in V1 and their retinotopical organization

A single neuron in V1 receives information from a restricted area of the visual environment and become active only when specific features, like orientation, contrast, or motion of an object are presented in that location. This area is referred to as the receptive field of the cortical unit (Hartline 1938). By definition, stimuli presented outside of the receptive field will neither increase nor decrease the ongoing activity of that individual cell. Two classes of neurons in V1 have been defined according to the complexity of their response properties, namely simple and complex cells. Simple cells are the most specific. For optimal activation, they require a bar of light of a certain width, which entirely fills the central area of the field, and which is oriented at a certain angle. Complex cells respond also to such stimuli, however in any position within their receptive field (Hubel and Wiesel, 1962, 1968; Albus, 1975). In addition, some complex cells respond particularly to lines or edges moving in a specific direction across the receptive field.
As a recording electrode moves parallel to the surface of the striate cortex, the position in the receptive field from which the responses are evoked changes in an orderly, predictable manner, thereby preserving the spatial organization of the ganglion cells within the retina. The receptive field location of many cortical cells have been used to map the visual field onto the cortical surface of the monkey (Talbot and Marshall, 1941; Daniel and Whitteridge, 1961; Bilge et al., 1967; Tusa et al., 1978), referred to as the *retinotopic map*.

Orientation and ocular dominance columns

A well known property of V1 is its organization in columns, which can be regarded as functional modules. Hubel and Wiesel showed that cells sharing the preferred stimulus orientation are arranged in vertical columns in cats (Hubel and Wiesel, 1962, 1963) and monkeys (Hubel and Wiesel, 1968). Moreover, detailed mapping of adjacent columns revealed a precise organization from one column to the next (Daniel and Whitteridge, 1961; Hubel and Wiesel, 1962; Albus, 1975). Orientation preference changes smoothly along the cortical surface, with systematic shifts in the axis of orientation from one column to the next, occasionally interrupted by non-responsive foci.

The organization in orientation columns has also been studied by optical imaging techniques in vivo (Blasdel and Salama, 1986; Bonhoeffer and Grinvald, 1991). Using either a voltage-sensitive dye or inherent differences in the light scattering due to blood deoxygenation close to active cells, a highly sensitive camera can detect the pattern of active and inactive orientation columns during the presentation of a bar of light with a specific orientation. Careful inspection revealed the presence of pinwheel centers, focal points at which all orientations come together. From there the system is organizes radially, in an extraordinarily regular manner (Hubel and Wiesel, 1962, 1968; Albus, 1975). Some pinwheels are systematically organized with clockwise progression, others counterclockwise, with each orientation appearing only once in a cylce. One or two such centers, evenly spaced, occur in each square millimeter of cortex (Hubel and Wiesel, 1962, 1968; Albus, 1975) (see Fig. 1).

Another system of alternating columns processes inputs from each eye (Hubel and Wiesel, 1959). Ocular dominance columns represent again an orderly arrangement, preferentially driven by either the left or right eye. A complete sequence of ocular dominance columns and orientation columns is repeated regularly and precisely over the surface of the primary visual cortex, each occupying a region of about 1 mm^2 (Hubel and Wiesel, 1968) (see Fig. 2). This two-fold repeating organization is a striking illustration of the modular organization characteristic of the cerebral cortex. Each module acts as a window on the visual field and each window represents only a tiny part of the visual field, but the whole field is covered by many such windows. Within the processing module all information about that part of the visual world is processed. It has been shown that orientation and ocular dominance columns communicate with one another by means of horizontal connections that link cells within a layer (Gilbert and Wiesel, 1983, 1989; Malach et al., 1993; Schmidt et al., 1997; Kisvarday et al., 1997).

In conclusion, the visual cortex is organized functionally into two sets of intersecting connections, one vertical, consisting of columns spanning the different cortical layers, and

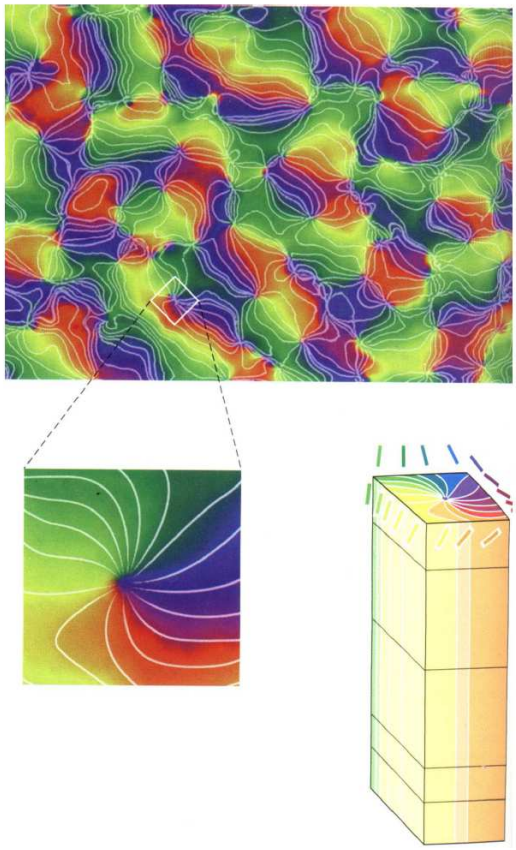

Figure 1: Orientation columns in the visual cortex of the monkey (modified from Kandel et al., 2000). Optical imaging of the cortical surface taken while the monkey viewed contours of different orientations (upper graphic). Bottom, left: Enlargement of a pinwheel-like area. Bottom, right: Three-dimensional organization of orientation columns in a 1 mm $\times$ 1 mm $\times$ 2 mm slab of the primary visual cortex under the square surface region depicted in the left graphic.

the other horizontal, connecting columns with the same response properties. Thus, substantial knowledge has been gained about the anatomical structures and the physiological mechanisms that generate orientation selectivity, but only recently studies focussed on its functional role. In some sense the experimental description of the visual cortical function is well ahead of our theoretical understanding of it. It has now become critical to fully understand the principles of cortical representation and to provide a quantitative description of the neural computation involved.

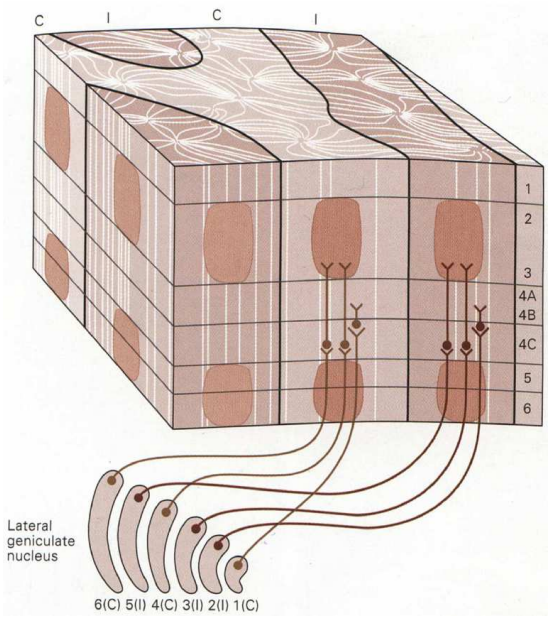

Figure 2: Organization of orientation columns and ocular dominance columns in primary visual cortex (modified from Kandel et al., 2000). An array of functional columns of cells in the visual cortex contains the machinery necessary to analyze a discrete region of the visual field and can be thought of as a functional module. Each module contains one complete set of orientation columns and one set of ocular dominance columns. The entire visual field can be represented in the visual cortex by a regular array of such modules.

Neuronal Coding

Over a century ago the seminal work of Ramon y Cajal (1893) identified single nerve cell as the elementary units of processing in the nervous system. The use of action potentials (spikes) as a mean of communication is the striking feature of neurons in the central nervous system. Since the discovery by Adrian (1926) that action potentials are generated by sensory neurons with a frequency that is substantially determined by the stimulus, the idea of rate coding has become a prevalent paradigm in neuroscience (Perkel, 1968). However, there is emerging evidence that populations of neurons are also able to synchronize their activity in a millisecond time scale and that synchrony plays a role in information processing. Substantial knowledge has been gathered concerning the biophysical (e.g., Gerstner and Kistler, 2002) and functional properties (Rieke et al., 1997). However, one of the key questions in neuroscience remains, how information is encoded in the spiking activity of a single cell and a whole population of cells (Golomb et al., 1994; McClurkin et al., 1991; Rieke et al., 1997; Shadlen and Newsome, 1995, 1998). In this debate, two different views can be distinguished. The rate-coding hypothesis claims that through the mean firing rate - i.e., the average number of spikes in some time bin

- information is carried, whereas the temporal-coding hypothesis claims that the precise placement (*time-locking*) of the spikes is a significant additional mechanism.

Rate code and population code

The frequency code theory was first advocated by Sherrington (1906) who proposed that neurons are *integrate-and-fire* devices (Knight, 1972). According to the integration principle each neuron sums the input activity and processes the response firing rate as a function of the input. This was verified by Adrian (1926), who demonstrated that changes in the firing of a neuron encoded information about the stimuli he presented to the animal. The powerful idea that neurons use a so called *rate code* to transmit information was soon generalized to other systems, and eventually to the entire nervous system. The estimation of firing rate by counting spikes during a response period has been enormously successful in advancing our understanding of sensory information processing and provided the foundation for our current understanding of the cortical code. From the point of view of signal analysis, this approach offers a dramatic simplifications because it implies that an entire spike train - a complex time varying signal comprising a long list of times at which a neuron fired - can be replaced by a single number, the mean rate.

An alternative hypothesis about the role of the frequency of neuronal firing is the *population code* model (Georgopoulos et al., 1982). According to this hypothesis, specific features of a stimulus are encoded by a distributed combination of firing rates over a group of neurons. The stimulus is represented as a distributed configuration of different firing rates. The population coding hypothesis has been shown to be statistically more reliable than observations using single neuron responses.
Evidence for integration over pools of neurons has come from the work of Georgopoulos et al. (1982), who has shown that vector averaging over pools of sequentially recorded motor cortical neurons can accurately reconstruct movement trajectories in monkeys. These response averages are taken over the entire duration of the movement ($\sim$ 500 ms). A criticism to this idea is that population coding is not able to be disentangled from the issue of temporal integration.

Cell assemblies

The first challenge to the rate cording hypothesis in favor of the importance of temporal coding was given by the Canadian psychologist Donald O. Hebb. His theory was truly groundbreaking when it was published in 1949 and it defined the guideline for theoretical neuroscience still prevailing to the current day. Hebb built his theory on the definition of a new concept that he called the “cell assembly” (Hebb, 1949).

A cell assembly is defined as a group of neurons with strong mutual excitatory connections. Hebb postulated the principle "*what fires together, wires together*", a simple synaptic mechanism how new cell assemblies can be formed by groups of cells that are often co-activated. According to this principle, repeated activation of the first cell will generate a strongly connected line of cells that is reliably activated whenever its first cell is active. Thus, once a subset of its cells are stimulated, the cell assembly tends to be activated as a whole. Therefore, it can be considered as an operational unit in the brain. In the framework of Hebb's theory, a number of phenomena described by psychology can be explained by simple mechanisms. An "association" is nothing but the activation of an assembly by either a stimulus or another assembly.

Despite being originally formulated as a postulate with little biological justification, many decades later the hebbian synapse and Hebb's rule have received extensive experimental support (e.g., Bliss & Lomo, 1973; Gustafsson et al., 1987). The concept of the cell assembly as a general principle of neural computation, however, remains hotly debated. Cell assemblies could be the manifestation, at least in principle, of mere covariances between firing rates. However, as it turned out that also the precise relative firing times and temporal correlations in neural ensembles is important for coding and synaptic plasticity (Abeles, 1991; Aertsen et al., 1989; Gerstein et al., 1989; Singer and Gray, 1995; Eckhorn, 1999; van Hemmen and Senn, 2002).

These observations led to a shift in paradigm from the study of rate coding mechanisms to spike based coding in the neurosciences during the last 20 years.
Hebb's ideas have influenced the research of many experimental and theoretical neuroscientists. Various "brain theories" realized aspects of Hebb's work or incorporated his concepts, e.g., Marr (1971); Shaw et al. (1985); Damasio (1989); Abeles (1991); Miller (1996); Mesulam (1998); Wickelgren (1999). One particularly intriguing suggestion for assembly organization is the *synfire chain* first proposed by Abeles (1991), based on an original ideal by Griffith (1963).

Synfire chains

The synfire chain is a network design that embeds the postulation of cell assembly, and features signal propagation and information processing by means of synchronous spike activity. The synfire chain concept is based on the prediction that, given typical cortical connectivity, firing rates, synaptic strengths, and synchronous inputs are likely to drive postsynaptic cells more effectively than asynchronous inputs. Consequently, in each step of the synfire chain the simultaneous firing of a significant pool of neurons raises the potential of the (functionally) next pool of neurons, and eventually bring them to fire.

From a topological point of view, a synfire chain is a feed-forward network composed of groups of neurons, where each group converges and diverges strongly onto the next group. Figure 3 illustrates the connectivity map in a synfire chain. When one pool is ignited by an external stimulus, the whole chain becomes active and packets of synchronous activity propagates between successive groups along the chain. Abeles (1982) argued that due to this architecture synchronous firing of neurons in one layer will reliably cause synchronous responses in the next group. He therefore suggested that synchronous activity can be propagated in a stable manner. Indeed, stable propagation of pulse packets in synfire chains has been shown theoretically (Hertz 1997), numerically (Diesmann et al., 1999), and with simulations (Abeles, 1982; Hayon et al., 2005).

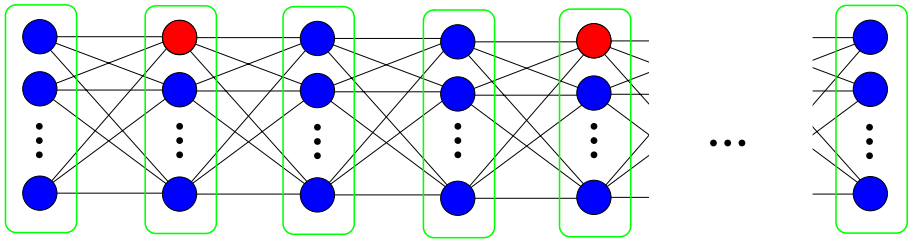

Figure 3: Schematic of a synfire chain. Each neuron is depicted as a blue circle. Black lines symbolize synaptic connections. A pool of neurons is denoted by a green rectangle. Each neuron converges and diverges to all neurons in the previous and next step, respectively. The layout of the chain does not represent the anatomy, but rather the order of activation in time. The same neuron can be part of more than one pool. For example the two red neurons shown in the figure may be physically one and the same neuron. Repeated participation of the same neurons can occur up to the limit of the synfire chains' memory capacity.

The problem of assembly detection

Several issues limits the ability to verify this hypothesis. However, no clear experimental evidence was reported that synfire chains occur in real brains. Experimental support has so far been limited to inference from detection of a few repeating spatio-temporal neuronal firing patterns in multiple single-unit recordings. However, theoretical and modeling studies suggest the number of synchronized neurons to be in the order of 100, with 80% of those neurons participating per activation (Abeles, 1991; Diesmann et al., 1999). Clear evidence of the existence of synfire chains would require simultaneous recording of several neurons from one pool and several from the next pool under conditions in which the synfire chain is repeatedely activated.

For a long time the required techniques of multiple single-unit recording and multivariate data analysis were just unavailable. However, advanced recording techniques became meanwhile available, that allow simultaneous recordings of up to several hundred neurons (electrode methods: Warren et al. 2001, Donoghue 2002, Csicsvari et al. 2003, Rolston et al. 2007, Euston et al. 2007; optical methods: Ohki et al. 2005; Ohki et al. 2006, Sasaki et al. 2007). A drawback of these methods is that recorded neurons are dispersed over a few cortical hyper-columns. With such a small sample, the likelihood to detect neurons either from the same or successive links of a particular chain is vanishingly small. However, multi-electrode recordings are steadily improving, such that recordings of even larger neuronal populations can be expected in future.

Another problem for the verification of synfire chains is the limited number of available analysis techniques. As massive parallel extracellular recordings are becoming more and more available, the limited experimental evidence in favor of the cell assembly hypothesis has be assigned to a large extent to a lack of suitable analysis tools (Brown et al., 2004). Synchronized spiking across large groups is potential signature for active assemblies, predicting higher-order correlations among the spike trains of assemblies members. However, the estimation of the necessary parameters of present analysis technique for higher-order correlations poses serious problems, mainly because the number grows exponentially with the number of recorded neurons. As a consequence, most attempts to detect active cell assemblies resort to pairwise correlations. Indeed, the existence and functional relevance of pairwise interactions has been shown in various cortical systems and behavioral paradigms, including visual (Singer et al., 1997; Kohn and Smith, 2005), motor (Vaadia et al., 1995; Riehle et al., 1997), and prefrontal (Sakurai and Takahashi, 2006) areas.

Pairwise analysis, however, do not differentiate situations in which only pairs fire synchronously from synchronized spikes in large populations. In such a case, the analysis is insensitive to higher-order correlations and sparse synchronous events. Compelling evidence for active assemblies thus require analysis methods that transcend pairwise correlations.
Approaches to analyze correlations between more than two neurons do exist but typically work only for a small number of neurons (e.g., Abeles and Gerstein, 1988; Dayhoff and Gerstein, 1983b; Grün et al., 2002a,b) or consider only pair correlations while analyzing the ensemble (e.g., Gerstein et al., 1978; Shlens et al., 2006; Schneidman et al., 2006). Therefore, limitations of available analysis tools has rendered experimental evidence for their fundamental prediction.

Linking the neuronal code to vision

Much of the information regarding the response properties of cells in the primary visual cortex has been obtained by displaying the optimal visual stimuli within a cell's receptive field while a monkey fixates on a screen (Hubel and Wiesel, 1968; Maunsell and Newsome, 1987) The initial assumption was that the responses thus recorded are not substantially different from those generated in a more natural setting, but many studies have now demonstrated that the responses of neurons in the visual cortex of cats and monkeys significantly modulated their activity when stimuli are presented beyond their classical receptive field, and that these responses are not readily apparent from their receptive field properties (Nelson and Frost, 1978; Zeki, 1983; Allman et al., 1985; Desimone and Schein, 1987; Gulyas et al., 1987). However, natural vision includes active eye movements, and the intent and initiation of eye movements is expected to have an influence on the neuronal activity in the visual system. Nevertheless, few neurophysiological studies have been conducted to investigate the discharge properties of neurons in the visual cortex of animals allowed to freely view natural images.

Every time we inspect a visual scene, we are confronted with a complex pattern of features At each location of gaze, the visual system must parse the scene into component objects. In spite of this complexity, our visual system is able to rapidely identify the features that belong together. A psychophysical experiment on monkeys and humans, where subjects were asked to decide whether a briefly flashed photograph of a natural scene contains a target category such as an animal, food, or a means of transport, has shown that monkeys can respond as early as 160 ms after stimulus presentation, and human subjects around 220 ms (Thorpe et al., 1996; Fabre-Thorpe et al., 1998; Van Rullen et al., 1998a)
To cope with the enormous amount of visual information in our everyday environment, the human visual system uses a mechanism of visual attention and saccadic eye movements to filter and process only the relevant information. The mechanisms that allow attention to search all the possible feature combinations in a reasonable amount of time are, however, still to be unveiled (Millner, 1974). Neisser (1967) suggested that some form of image segmentation should occur that not only happens prior to attention but also serves to attract it. One possibility is that salient features in an image, such as regions of high contrast lead to elevated activity, and the population of cells with the greatest activity captures attention (Koch and Ullmann, 1985). Indeed, experiments in humans have shown that features such as orientation, color, and image intensity significantly draw human attention. On the other side, fixations are less strongly tied to visual saliency when meaningful scenes are viewed during active tasks (Oliva et al., 2003; Land and Hayhoe, 2001; Henderson and Hollingworth, 1999; Turano et al., 2003) It is now

commonly accepted that human eye movements are controlled by both bottom-up and top-down mechanisms. The role of these two mechanisms in the control of eye movements is thought to depend on the amount of semantic content in the image, as suggested by (Parkhurst and Niebur, 2003). Images with high semantic content are likely to be more top-down controlled, since an internal model of the content exists (Parkhurst et al., 2002). The inspection of only a small part of the image in cooperation with the internal model provides the observer with information about what to expect in the rest of the image. In contrast, unfamiliar images, which have little meaning to the observer, stimulus-driven or bottom-up mechanisms are supposed to play a more prominent role.

However, little is known about eye movements in relation to activity of population of neurons. To this purpose, experiments with primates may offer the chance to gain new insights into the mechanisms underlying visual processing in the brain, by combining psychophysical and neurophysiological experiments. Moreover, it remains to be investigated in how far viewing strategies employed by monkeys resemble similarities to those used by humans. It has been shown in a speed categorization task that monkeys perform with only slightly decreased accuracy but greater speed, than human observers (Fabre-Thorpe et al., 1998). These results reveal striking similarity between monkeys and humans regarding rapid processing of natural scenes. Nevertheless, it remains to be investigated whether or not this similarity generalizes to viewing under natural conditions, when attentional processes and eye movements take effects.

Aim and outline of the thesis

The first part of thesis addresses the issue of temporal coordination between neurons in the primary visual cortex. In the first chapter the question of how the concept of cortical maps is related to the concept of temporal coding is approached. To this end we analyzed massively parallel spike recordings from the visual cortex (area 17) of an anesthetized cat under different stimulus conditions, and during spontaneous activity. I addressed on the one hand the issue of the intrinsic mechanisms employed by V1, and on the other hand of the time dependencies between spike times of neurons belonging to the same column and to different columns.
In the second chapter the concept of temporal coding is related to the eye events of free viewing monkeys. Correlated activity of single neurons in V1 is analyzed in respect to eye fixations and saccades. The analysis allowed to address the question whether synchrony is temporally related to viewing behavior, thereby providing a functional mechanism for initial image segmentation.

In the second part of this thesis new approaches towards the detection of higher-order correlation are discussed. Two novel methods for the selection of assembly members in large data sets are presented (Chapter 3 and 4). The methods are tested on a network of 5000 neurons with embedded synfire chain activity, and then applied to the same data set already analyzed in Chapter 1.

In the last part of this thesis I investigate the eye movements of monkeys during free exploration of natural images, in order to develop possible approaches for linking correlated neuronal activity to active vision (Chapter 5). The questions of what guides the monkeys' eye fixations and what strategies they use for image exploration are addressed and related to the case of humans.

Finally, I will review the findings presented in this thesis in detail and bring the different results in an unified framework, before closing the discussion with an outlook on future work.

Note about section numbering

In the following chapters, for maximal conformity with both the published version of the manuscripts and the structure and readability of this dissertation, I use a slightly different notation for numbering subsections, formulae and figures.

The section and subsection numbers assigned in this dissertation are shown in square brackets and are followed by the original sectioning numbers, as found in the published version of the manuscripts. Similarly, for formulae and figures the original number from the published version of the manuscripts is preceded by the dissertation sectioning number. Furthermore, figures are shown in color also when they were originally published in black and white.

Chapter 1

Spatially Organized Spike Correlation in Cat Visual Cortex

Spatially Organized Spike Correlation in Cat Visual Cortex

Denise Berger[1], David Warren[2], Richard Normann[2], Amos Arieli[3] and Sonja Grün[4]

[1]Neuroinformatics, Inst. Biology - Neurobiology, Free University, Berlin, Germany
[2]Dept. Bioengineering, University of Utah, Salt Lake City, USA
[3]Dept. Neurobiology, Weizmann Inst. of Science, Rehovot, Israel
[4]Computational Neuroscience Group, RIKEN Brain Science Institute, Wako City, Japan & Bernstein Center for Computational Neuroscience, Berlin, Germany

This section is published as

Berger, D.; Warren, D.; Normann, R.; Arieli, A. & Grün, S.
Spatially organized spike correlation in cat visual cortex
Neurocomput, 2007, 70, 2112-2116.

Abstract

Inspired by optical recordings from visual cortex which show maps of orientation selectivity, and the finding that very similar patterns of population activity occur when the neurons fire spontaneously (Kenet et al., 2003), we approach the question of how the concept of cortical maps may be related to the concept of temporal coding. To this end we analyzed parallel spike recordings performed using a 10 × 10 electrode grid covering an area of 3.6 mm × 3.6 mm of cat visual cortex for occurrence of spike correlation.
We calculated all possible pairwise correlations between multi-unit activities (MUA) by cross-correlation and extracted significantly correlated pairs using a boot-strap procedure. The MUAs involved in correlated pairs were typically involved in more than a single correlated pair. Using methods of graph theory we found that the whole set of correlated MUAs decomposes into a small number of groups of MUAs that have a high degree of overlap of the mutually correlated pairs.
Mapping these groups back onto the spatial arrangement of the recording electrodes revealed that these also correspond to spatially segregated clusters. The spatial scale of this correlation map is in agreement with the scale of orientation tuning maps found by optical imaging.

Keywords: spike synchronization; maps ; visual cortex

[1].1 Introduction

Most cells in the primary visual cortex are selective for orientation and respond best, i.e., by increasing their rates, to their optimal stimulus. Neurons located in different columns but with same stimulus preference coherently show high activity during the same, optimal stimulus, as indirectly shown by optical imaging. Moreover, it has been shown that the firing rate of a spontaneously active single neuron strongly depends on the instantaneous pattern of the ongoing population activity in a larger cortical area (Tsodyks et al., 1999). Very similar patterns of population activity were observed, both when the neuron fired spontaneously and when driven by its optimal stimulus (Kenet et al., 2003). On the other hand, the temporal coding hypothesis suggests that coordinated spiking activity on a fine temporal scale should occur in the nervous system (Nelson et al., 1992) and is suggested to be used for information coding. Accordingly, modern theories of the primary visual pathway contain spatial and temporal aspects accounting for the observed cell behavior in order to explain visual information processing.
We therefore approach the question of how the concept of temporal coding of single neurons are related to the overall functional architecture. More specifically, we ask if correlated neurons are arranged in a specific spatial organization that may be related to maps of orientation tuning. Parallel spike recordings based on using a 10×10 electrode grid (Utah electrode array, Bionic Technologies, Inc., Salt Lake City, UT, USA) covering an area of 3.6 mm $\times$ 3.6 mm of cat visual cortex (Warren et al., 2001) allow us to address the question of the relation of correlated neuronal activity to the distance and the spatial arrangement of the recording sites. Data were recorded from area 17 of anesthetized cat during spontaneous activity, and under fullflash treatment with two different stimulus intensities. We first analyze the simultaneous multi-unit recordings for pairwise correlations using cross-correlation analysis and evaluate their significance using bootstrap techniques. In a next step, we identify groups of correlated pairs that are highly mutually intra-correlated. Finally, mapping these groups back onto the electrode positions and thereby to cortical space, allows us to relate their spatial arrangement to the spatial scales found for orientation tuning maps.

[1].2 Methods and Results

[1].2.1 Detection of Correlated Spiking Activity

We analyzed the spiking activities that were recorded from the grid of electrodes (Fig. [1].1) for pairwise correlations, during spontaneous activity (SP; no stimulus) and during

full flash stimulation with two different intensities. We segmented epochs of high intensity (HI) and low intensity (LI) into two separate data sets and analyzed them separately. For simplicity we restricted ourselves to the evaluation of the multi-unit activities (MUA), and requested a minimal firing rate of 1 Hz for further consideration. This left us with 80 parallel MUA spike trains during SP and 83 during LI and HI for further analysis. From these we computed cross-correlations (CCH, Perkel et al., 1967) of all possible pairs in the different stimulus conditions. To evaluate the significance of the correlation, we used a boot-strap method that accounts for the firing rate changes of the neurons. Methods that rely on trial repetitions (e.g., Aertsen et al., 1989) cannot be used here since under the spontaneous condition we have only one trial. Thus we generated surrogate data sets (100) that account for the firing rate changes in the data by dithering the individual, original spike times within a given time interval (Hatsopoulos et al., 2003). Dithering is performed for each individual spike of each of the parallel MUAs. Each spike was dithered by random placement (equal probability) within a time interval of pre-defined width positioned symmetrically around the original spike time. Since the widths of the peaks in the original cross-correlograms were typically in the range of 30-50 ms, the dither time interval was chosen such (70 ms) to destroy the temporal correlation while keeping slow rate variations.

For each surrogate data set we computed the cross-correlation in the same way as for the experimental, original data, which finally resulted in a mean cross-correlogram of all surrogates including the standard deviations for each delay τ (Fig. [1].1C). Spike correlation between two MUAs is considered significant, if the center peak of the smoothed CCH (box car kernel of 10 ms width) exceeds two standard deviations (i.e., a significance level of 5%) of the surrogate data. Then a pair is called correlated. A total of 148/3160 pairs were significant during SP, 78/3402 during HI and 203/3402 during LI.

[1].2.2 Groups of Inter-Correlated Pairs

Representing the pair correlations in form of an undirected graph, in which each MUA is represented as a node, and existence of correlation by an edge between the two involved MUAs, we found that the connection degree of edges per node was typically larger than two, i.e., a single MUA is typically correlated with more than two other MUAs simultaneously. This led us to identify cliques of correlated pairs, i.e., groups of all-to-all mutually correlated MUAs. Table [1].1 shows for the case of LI how often a clique of k nodes, i.e., MUAs, were found. Surprisingly, we found cliques with up to $k = 9$ nodes. In addition, we noticed that cliques often have mutually overlapping members.

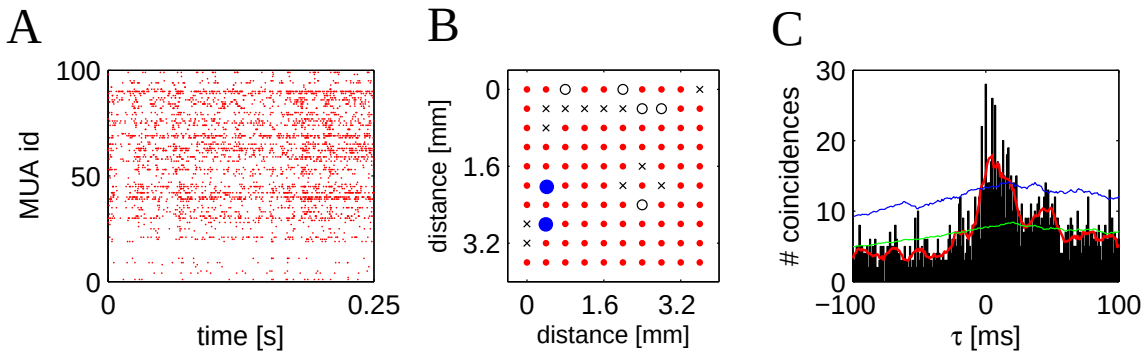

Figure [1].1: A. Dot display of MUA spiking activity under low intensity stimulation measured simultaneously by a grid of electrodes (see B). MUA id corresponds to the electrode position numbered from upper left to lower right, each line from left to right. B. Electrode recording grid. 10 × 10 Utah electrode array covering 3.6 mm × 3.6 mm of the cortex with an inter-electrode distance of 400 μm. Electrodes marked by a black circle were broken and did not deliver any signal. MUA signals with rates below 1 Hz were not considered for further analysis (corresponding electrode positions marked by crosses). C. Cross-correlogram of 2 MUAs (recorded from electrode 52 and 72, marked by blue circles in B.). Raw cross-correlogram (black) computed on the time resolution of 1 ms. The red line shows the smoothed CCH (kernel width: 10 bins of h). The green line indicates the mean of the predictor, the blue line the significance level (mean plus twice the standard deviation). If the center peak of the smoothed cross-correlogram exceeds the significance level, the pair of MUA is considered to be significantly correlated.

k	2	3	4	5	6	7	8	9
#	23	18	9	8	2	2	3	1

Table [1].1:

This led us to define groups of highly intra-correlated MUAs (GIC) according to the following conditions: (1) only cliques with a minimum of $k = 3$ members were considered and, (2) cliques are requested to have a minimal overlap of one member. Interestingly, as a result the graph decomposed into only a small number of completely disjoint subgraphs (3 during SP, 2 during HI and 4 during LI). The GICs contained 4, 13, and 19 MUAs during SP, 7 and 15 MUAs during HI and 14, 3, 11, and 21 during LI (see Fig. [1].2). Although cliques were requested to have a minimal overlap of one member, we actually find that each N-clique within a GIC has an overlap of N-1 nodes of other cliques, if at least one other clique with M≥ N nodes exists. A single MUA may be involved in up to 15 pairwise correlations.

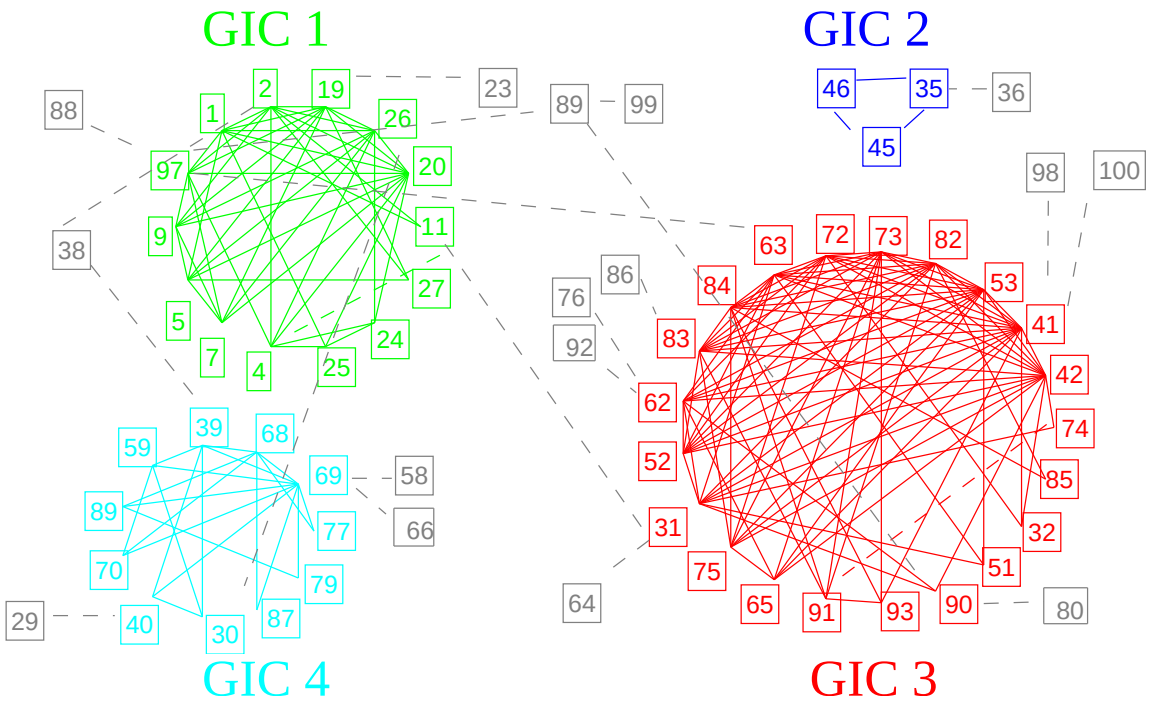

Figure [1].2: Graph of all significant pairwise MUA correlations under LI stimulation. Each MUA is indicated by a square that includes the id of the MUA. These represent nodes, and edges indicate an existing correlation between the two connected MUAs. The four circles of nodes show groups that are composed of different cliques. Conditions to form groups are: a minimal clique size of $k = 3$ and a minimal overlap between cliques of 1. The resulting groups (GIC 1 - GIC 4) contain 14, 3, 11 and 21 MUAs, and are composed of 41, 3, 21 and 91 pairwise correlations, respectively. Additional squares (gray) mark nodes that do not fullfill the above criteria, i.e., do belong to cliques of only $k = 2$ members, and thus do not belong to one of the groups. MUAs 52 and 72, for which we have shown the cross-correlation in Fig. [1].1C, are both members of GIC 3.

[1].2.3 Spatial Organization of Correlated Activity

Now we address the question how the GICs are *spatially arranged*. Therefore we map the members of the identified groups back onto the recording grid and mark members of a group by filled circles of a specific color (Fig. [1].3A). Interestingly, the different GICs form separated, compact clusters in space. Thus, the groups of MUAs that are highly intra-correlated also form localized groups in space. Two of the clusters (green corresponding to GIC 1 and red corresponding to GIC 3, cf. Fig. [1].2) also included one remotely located MUA. Note, that also the remote MUAs are included in multiple correlations with the rest of the MUAs of the corresponding GIC. Circle radii indicate number of MUAs the indicated MUA is correlated to. Our findings clearly show that

the spatial organization of correlations are not randomly or homogeneously distributed in space, but rather cluster in certain spatial domains. The spatial scale of the GICs is directly comparable to the spatial scale of orientation maps found by optical imaging (Fig. [1].3B). Across the different stimulus conditions the number of correlated MUAs differ (see Section [1].2.1): during HI or SP less MUAs appear to be correlated as compared to LI. As a consequence, in these conditions a lower number of GICs are found (see Section [1].2.2). However group memberships stay mainly the same across different conditions. Variation of the dither width had a minor impact on the results. For a smaller dither width as used for the former analysis (30 ms), a smaller number of correlations appeared to be significant. The involved MUAs formed the same spatial clusters as for larger dither widths, but with less degree of connectivity and less members. For increasing dither widths, the clusters became more and more pronounced. Beyond a dither width of 70 ms, cluster members and degree connectivity remained stable.

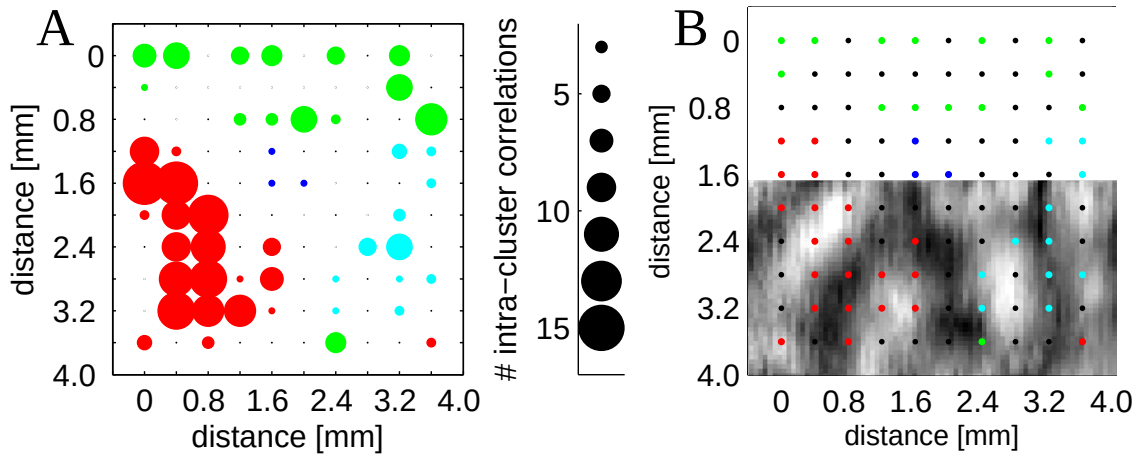

Figure [1].3: A. Spatial arrangement of groups of highly intra-correlated MUAs for the same data as shown in Fig. [1].2 during LI. The array marks the electrode grid used for the recording. Filled circles of a specific color indicate electrode positions of MUAs that are identified as members of a GIC. The radii of the filled circles indicate the number of correlations the MUAs were involved in (see marker bar on the right). GICs result in spatial clusters: GIC 1 corresponds to the green cluster (top, horizontal), GIC 2 to the blue cluster (center), GIC 3 to the red cluster (left, vertical), and GIC 4 to the cyan cluster (right, vertical). B. Comparison of spatial scales. Optical imaging orientation map for one stimulus orientation (full field grating; modified from Kenet et al., 2003) including the electrode grid used in our study are both shown on the same scale.

[1].3 Discussion

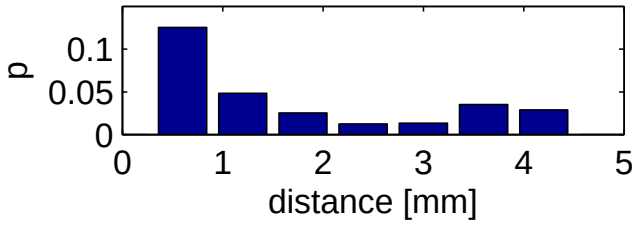

Figure [1].4: Distance dependence of correlations. The histogram shows as a function of electrode distance the probability of finding correlated MUA pairs (during LI, same data as shown in Figs. [1].2 and [1].3).

This study investigated the spatial organization of correlated activity in cat visual cortex. Multi-unit activities were recorded by a grid spanning an area of 3.6 × 3.6 mm^2 allowing us to investigate the relation of correlated MUA pairs to spatial structure. On average 5% of all possible pairs were significantly correlated. To capture the correlation structure, we constructed a graph with each MUA being a node and assigning edges to significantly correlated MUAs. This led to the finding that the large number of pairwise correlations (about 185) separate into a low number of groups (2-4) of highly intra-correlated MUAs with 3-21 members per group (GICs). By mapping the detected GICs back onto cortical space, we found that the members of the individual GICs also clustered in space revealing a relationship between correlation pattern and spatial pattern. The occurrence of correlation between pairs of MUAs appeared to be distance dependent. When adjacent, they occurred most frequently, increasingly less up to about 3 mm, and then again increasingly more (Fig. [1].4). Moreover, the spatial extent of the GICs corresponds to the spatial scale of the orientation tuning map obtained by optical imaging (see Fig. [1].3B, Kenet et al., 2003), as well as by electrophysiological imaging maps (Warren et al., 2004). Thus we conclude that our results in respect to the extent and spatial arrangement of the clusters are in agreement with the spatial organization reported in visual cortex contained by optical imaging (Huebener et al., 1997).

Ts'o et al. (1986) investigated cross-correlations between cortical neurons in respect to long-range horizontal connections, i.e., distances of up to several mm. Their study suggests a high degree of inter-connectivity between cells having the same orientation preference within as well as between columns, thereby strongly supporting our interpretation of our results. Das and Gilbert (1999) investigated short-range interactions and observed that neighboring cells in a range of up to 500 - 800 μm show spike correlations independent of the orientation preference. However, we found correlations on a scale of distance. The GICs have spatial extents that exceed this short-range interactions in all but one GIC.

Furthermore, an additional important difference between their experiments and ours are the type of stimulus. Here we investigated spontaneous activity as well activity under full-flash treatment, whereas Das et al. used isolated optimal bars in the receptive fields of the neurons. More comparable to our results is the study by Tsodyks et al. (1999) who found orientation maps under spontaneous activity that is related to the tuning of the individual cortical neurons and concluded that the underlying functional architecture is associated to tuning properties. Another important observation, adding credence to the possibility that our observed GICs reflect the underlying connection network in area 17 is that these correlation maps are highly similar for the 3 recording conditions. We conclude that the observed correlation patterns can not be considered to be random in view of the overall number of correlations. Indeed, the correlations between MUA pairs appear to be well organized, particularly when distance dependent correlations are compared to orientation tuning in distinct columns. Since an orientation stimulus was not used here, MUA correlation cannot definitively be ascribed to orientation tuning. Nevertheless, we propose that these correlations of neuron pairs are linked according to orientation tuning. For a categorical affirmation these results would have to be verified using an adequate experimental paradigm using oriented stimuli. We also aim to address in future the temporal dynamics in the correlation patterns of the clusters. Since Kenet et al. (2003) found different orientation maps to be activated dynamically and systematically during spontaneous activity, we expect the GICs to be dynamically correlated and alternatingly active. Another open question is related to the order of spike correlations. Due to the limitation of the cross-correlation method to be restricted to pairwise correlations only, we could not determine if the MUAs involved in more than one pair actually reflects the existence of higher-order correlations, i.e., spike correlation patterns that involve more than two neurons at a time (e.g., Schneider and Grün, 2003).

Acknowledgments: Partial funding by the BMBF (BCCN Berlin, grant 01GQ0413), the Stifterverband für die Deutsche Wissenschaft, and the Volkswagen Foundation. This work was carried out while Sonja Grün was based at the Freie Universität in Berlin, Germany.

References

A. M. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J. Neurophysiol.*, 61:900–917, 1989.

A. Das and C. D. Gilbert. Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature*, 399:655–661, 1999.

N. Hatsopoulos, S. Geman, A. Amarasingham, and E. Bienenstock. At what time scale does the nervous system operate? *Neurocomp.*, 52–54:25–29, 2003.

M. Huebener, D. Shoham, A. Grinvald, and T. Bonhoeffer. Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.*, 17:9270–9284, 1997.

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

J. I. Nelson, P. A. Salin, M. H. Munk, M. Arzi, and J. Bullier. Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Vis. Neurosci.*, 9:21–37, 1992.

D. H. Perkel, G. L. Gerstein, and G. P. Moore. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys. J.*, 7:419–440, 1967.

G. Schneider and S. Grün. Analysis of higher-order correlations in multiple parallel processes. *Neurocomp.*, 52–54:771–777, 2003.

D. Y. Ts'o, C. D. Gilbert, and T. N. Wiesel. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.*, 6:1160–1170, 1986.

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

D. J. Warren, E. Fernandez, and R. A. Normann. High-resolution two-dimensional spatial mapping of cat striate cortex using a 100-microelectrode array. *Neuroscience*, 105:19–31, 2001.

D. J. Warren, A. Koulakov, and R. A. Normann. Spatiotemporal encoding of a bar's direction of motion by neural ensembles in cat primary visual cortex. *Ann Biomed Eng*, 32:1265–1275, 2004.

Chapter 2

Synchronization of Neuronal Responses in Primary Visual Cortex of Monkeys Viewing Natural Images

Synchronization of Neuronal Responses in Primary Visual Cortex of Monkeys Viewing Natural Images

Pedro Maldonado[1], Cecilia Babul[1], Wolf Singer[2], Eugenio Rodriguez[2], Denise Berger[3,4], and Sonja Grün[4,5]

[1]Physiology and Biophysics, Faculty of Medicine, University of Chile, Santiago, Chile
[2]Neurophysiology, MPI Brain Research, Frankfurt, Germany
[3]Neuroinformatics, Inst. Biology - Neurobiology, Free University, Berlin, Germany
[4]Bernstein Center for Computational Neuroscience, Berlin, Germany
[5]Computational Neuroscience Group, RIKEN Brain Science Institute, Wako City, Japan

This section is published as

Maldonado, P.; Babul, C.; Singer, W.; Rodriguez, E.; Berger, D. & Grün, S. Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. J. Neurophysiol., 2008, 100, 1523-1532.

Abstract

When inspecting visual scenes, primates perform on average four saccadic eye movements per second, which implies that scene segmentation, feature binding, and identification of image components is accomplished in <200 ms. Thus individual neurons can contribute only a small number of discharges for these complex computations, suggesting that information is encoded not only in the discharge rate but also in the timing of action potentials. While monkeys inspected natural scenes we registered, with multielectrodes from primary visual cortex, the discharges of simultaneously recorded neurons. Relating these signals to eye movements revealed that discharge rates peaked around 90 ms after fixation onset and then decreased to near baseline levels within 200 ms. Unitary event analysis revealed that preceding this increase in firing there was an episode of enhanced response synchronization during which discharges of spatially distributed cells coincided within 5-ms windows significantly more often than predicted by the discharge rates. This episode started 30 ms after fixation onset and ended by the time discharge rates had reached their maximum. When the animals scanned a blank screen a small change in firing rate, but no excess synchronization, was observed. The short latency of the stimulation-related synchronization phenomena suggests a fast-acting mechanism for the coordination of spike timing that may contribute to the basic operations of scene segmentation.

[2].1 Introduction

When inspecting visual scenes it can be anticipated that a novel input constellation needs to be processed each time a saccadic eye movement is performed to shift the gaze to a new position. Primates perform up to four saccades per second, suggesting that feature extraction, perceptual grouping of features belonging to individual objects, and identification of objects can be accomplished within about 200 ms. Similar conclusions on processing speed have been derived from psychophysical experiments with human subjects who had to recognize individual objects in cluttered scenes (Fabre-Thorpe et al., 2001). Based on the evidence that processing can be surprisingly fast, it had further been proposed that the neuronal networks at early stages of visual processing should accomplish functions such as feature extraction, scene segmentation, and perceptual grouping by exchanging only a few spikes per neuron (Guyonneau et al., 2004). In this case, only a limited amount of information can be encoded by varying the discharge rate and it has thus been suggested that, in addition, the cortex might exploit a temporal code. In this case, information would also be encoded in the precise timing of action potentials (Fries et al., 2001, 2007; Grün et al., 2002a; Hopfield, 2004). However, little is known about the responses of cells in the visual cortex when animals are free to explore natural images. In cats and monkeys, neuronal responses to stimuli placed in the cells' receptive field (RF) are modified by stimuli presented outside the "classical" RF (Jones et al., 2001; McAdams and Maunsell, 1999; Sugita, 1999; Trotter and Celebrini, 1999). Recent studies in monkeys exposed to complex visual stimuli that simulate natural viewing conditions have shown, in addition, that neurons in the primary visual cortex (V1) significantly reduce their firing when a larger portion of the visual field is covered with the sample image (Vinje and Gallant, 2000). These context-dependent modulatory processes have been hypothesized to participate in stimulus selection, scene segmentation, and perceptual grouping. However, little is known about how these multiple modulatory influences combine when animals are allowed to freely scan visual scenes and no data are available on precise timing relations among the discharges of simultaneously recorded neurons. We hypothesized that information is encoded not only in discharge rates but also in timing relations among the spikes of individual neurons; thus we should find indications of precise timing, e.g., synchronization of discharges in these data.

Using an array of eight individually adjustable tetrodes we simultaneously recorded the discharges of multiple neurons from the V1 of two adult, male capuchin monkeys (Cebus apella). The animals were presented with a collection of pictures of different natural scenes. Using Unitary Event (UE) analysis, we demonstrate an increased and highly precise spike synchronization that is dissociated from the later increase in neuronal firing

rate. These results suggest that coordination of the timing of discharges may play an important role in the initial processing of signals in V1.

[2].2 Methods

Subjects and recording setup

All experiments followed institutional and National Institutes of Health guidelines for the care and use of laboratory animals. Two adult, male capuchin monkeys (Cebus apella) weighing 3-4 kg served as subjects for this study. All surgical and recording procedures were similar to those in Friedman-Hill et al. (2000). A hard plastic recording chamber with eight independently movable nichrome tetrodes was mounted over the surface of the V1. Guiding tubes were positioned on top of the dura mater. Tetrodes could be lowered $\leq$ 3-4 mm; thus we could record from the top cortex as well as a second section of V1 located underneath. After the end of recordings, the head post, eye coils, and manipulator were removed and after full recovery the animals were donated for adoption.

Visual task and eye movements

The animals were seated in a chamber dimly lit at the low scotopic level and presented with a collection of 15 pictures of different natural scenes that were displayed on a 21-in. computer monitor (frame rate 60 Hz) located 57 cm in front of the animals, subtending 30×40^o of visual angle. Each image was presented at least three times in a single recording session. The animals were allowed to freely explore natural images ("images") or black frames ("blank"). The experimental protocol required the animals to maintain their gaze for $\leq$5 s within the limits of the frames, to be rewarded with a drop of juice. To maintain alertness, we also presented black frames that contained a single central fixation spot that the animals had to fixate for 1 s to be rewarded (1^o window). Vertical and horizontal eye positions were monitored with a search coil driver (DNI Instruments; resolution: 1.2 min of arc) by magnetic induction of an implanted coil (Judge et al., 1980) and then digitized at 2 kHz (see Fig. [2].1). Eye coils were made of three loops of Teflon-coated wire and implanted over the sclera, one in each eye. Eye coils were calibrated using fixation points located at nine different positions on the screen. The calibration procedure was performed at the beginning of each recording session.

Recordings

Neuronal activity of neighboring neurons was recorded with an array of eight individually adjustable custom-fabricated nichrome tetrodes (1- to 2-MΩ impedance). The tetrode wires were 12 μm thick made of nickel chromium (Kanthal, Palm Coast, FL), insulated

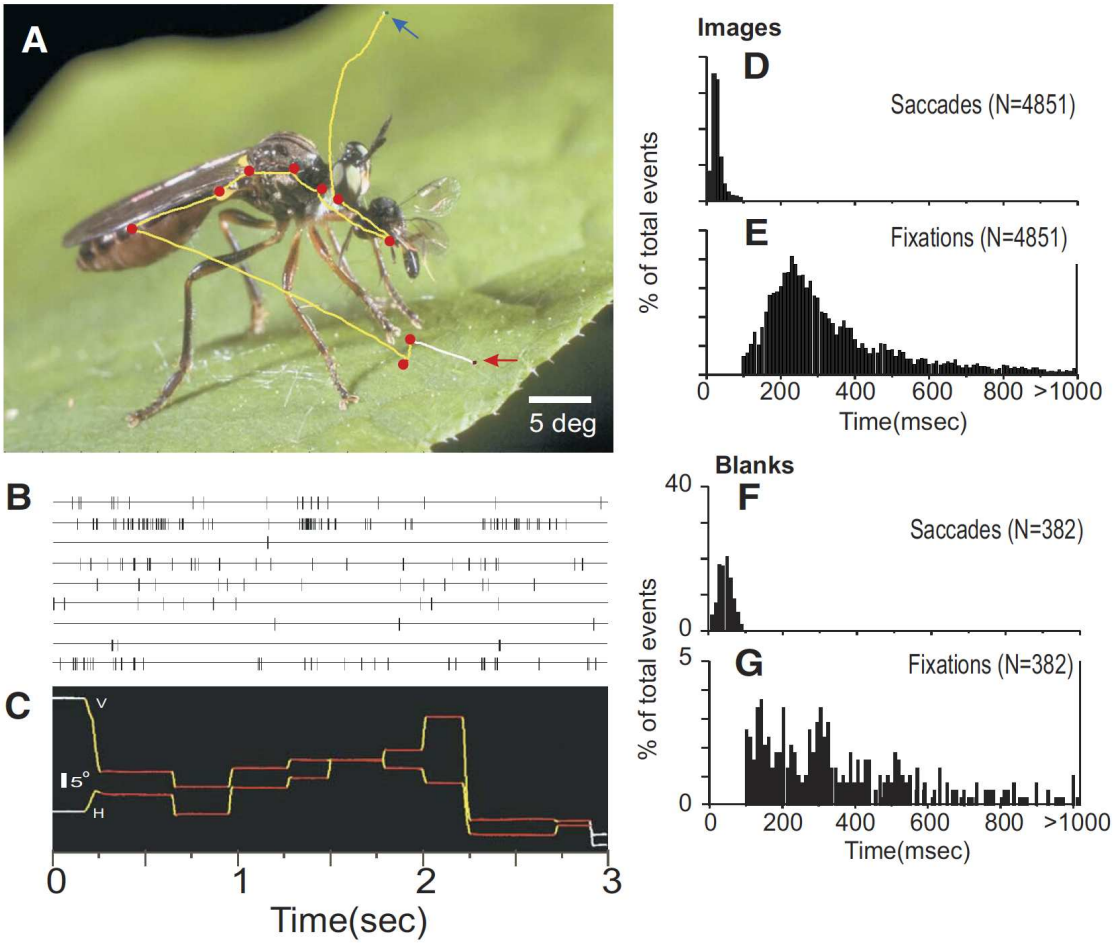

Figure [2].1: Responses of cells in the primary visual cortex of a monkey (Cebus apella) visually exploring a natural scene. A: example of the presented images. Monkeys scanned the picture for $\leq$5 s. In this image presentation of 3 s duration, the colored line represents the eye trajectory. The yellow traces represent the saccades, the red blobs represent the fixations, and the white traces the unclassified eye movements. The blue and red arrows show the initial and final positions, respectively. B: spike trains from 10 simultaneously recorded cells for the time period shown in A. Each spike is represented as a small vertical line. Note the diversity of activity for each visual fixation shown in C. C: vertical and horizontal eye traces recorded during the same image presentation, showing the automated eye classification assignments: yellow and red for saccades and fixation, respectively. The white traces correspond to uncompleted or unclassified eye movements. D-G: frequency distributions of the duration of different eye movements in image (D-E) and blank (F-G) conditions, respectively. Data are from 1219 images and 547 blank frames. Total numbers of the respective eye movements are shown in parentheses. The distributions for both saccades and fixations events differ significantly between the image and blank conditions (Kolmogorov-Smirnov, P $<$ 0.001).

with polyamide, and wound together in a four-wire array. The tetrodes were mounted in a 33G stainless-still tubing, which was beveled in a grinding wheel and gold plated (for construction details, see Gray et al., 1995). The tetrodes were positioned in a circular array, with a center-to-center distance of about 400 μm. The signals were amplified ($\times$10K), band-pass filtered for single-unit activity (0.5-5 kHz), sampled at 25 kHz, and then stored for off-line spike sorting and analysis. Signals were fed through an off-line sorting program to reconstruct the spike trains of units recorded simultaneously by a single tetrode. Since the tetrode array remained in the animals for several days, each successive recording through the same penetration was resumed always $\geq$200 μm deeper than that during the previous session. This sampling procedure was continued until activity could no longer be measured after which the guide tubes were repositioned. Some penetrations crossed V1 twice in the anterior part of the calcarine sulcus, which led to systematic changes in RF position.

Data analysis

EYE MOVEMENTS. An automatic algorithm was developed to extract from the eye traces the different types of eye movements. Saccades were defined as eye movements with an angular velocity $>100^o$/s lasting for $\geq$5 ms. In addition, saccades were required to exhibit a minimum acceleration of $170^o/s^2$. The animals performed saccades with peak velocities that ranged between 100 and $>1{,}000^o$/s; the average peak velocity was around 400^o/s. Fixation periods were classified as such when they lasted $\geq$100 ms with the eye position maintained within 1^o of the gaze location reached at the end of a saccade. To minimize errors, only pairs of unambiguous saccade-fixation sequences were considered for further analysis. The significance of firing rate changes during saccades was assessed by comparing the rates in the first 20 ms and the consecutive 20 ms following the onset of saccades. Significance tests for rate changes during fixations were performed comparing the first and consecutive 50 ms after the onset of the events (t-test).

In this study animals freely directed their visual behavior. Thus after each saccade, cells were stimulated by different portions of the image. The data thus reflect the average responses of large numbers of cells to a large number of different stimuli. We address each individual eye-movement-related event (saccade and fixation) as a "trial" and the measurement period during which electrode positions were kept constant and the constellation of recorded cells was stable as "recording session." Trials of the same event (saccades or fixations) were analyzed individually. Subsequently those occurring during the same viewing condition (image or blank) and within the same recording session were grouped. In Figure [2].2A we illustrate the segmentation of the data from a single image presentation into trials. Since statistical analysis requires that trials have the same

duration, we segmented the data using the onset of an eye-movement-related event as reference (see histograms in Figure [2].1). For fixations, data were included from -25 ms before fixation onset to 325 ms after fixation onset. This epoch was chosen so that the UE analysis performed in sliding windows of 50 ms renders its first value at fixation onset. If fixations were shorter, data from subsequent eye events were included to complete the time window. An alternative would have been to include only the data from within the trial that, however, would lead to a successive decrease of data for increasing fixation duration. This in turn would lead to an increase in the noise. Therefore we decided to include data from subsequent eye events. This probably introduced distortions toward the end of the analysis window beyond 200 ms, but as the data show, the important events on which our conclusions are based all occurred within the first 100 ms after fixation onset. For the analysis of signals related to saccades, data were included from -25 to 75 ms.

ANALYSIS OF SPIKE CORRELATION. For the evaluation of temporal correlations between spikes from different neurons we used the UE analysis, a statistical method that detects the presence of spike coincidences that exceed the level predicted by the respective firing rates and evaluates their statistical significance (see METHODS; Grün et al., 1999, 2002a,b, 2003; Riehle et al., 1997) . Coincident firing that exceed the chance level by a significant amount are called unitary events (Grün et al., 2002a). Although in principle the method is able to handle correlations between more than two neurons, we restricted the analysis here to pairwise correlations. The reason is that the analysis of triplet coincidences or even higher complex synchrony patterns requires a higher number of samples to obtain reliable statistics. Since we are interested in the dynamics of synchrony we use a sliding-window analysis. The width of the window has to be chosen relatively small here (50 ms) to account for the changes in the firing rate. Given these constraints we decided to restrict the analysis to pairs of neurons only.

Spike trains of a pair of neurons recorded simultaneously during a particular trial were aligned to the onset of the respective event (Fig. [2].2B, "EV", here chosen as fixation onset). To account for the nonstationarity of the neurons' firing rates, the data were analyzed with a sliding window (Grün et al., 2002b). An analysis window of fixed duration T_w (here: 50 ms; gray box in Figure [2].2B) was slid along the data such that all data are covered from the beginning to the end of the trial. The window was advanced in steps corresponding to the time resolution h of the data (0.1 ms). The first window position was centered at $t = t_0 = T_w/2$ from trial onset and the last window around time $T_w/2$ before the end of the trial. At each position of the window all the data within the window were analyzed for unitary events and the analysis was repeated sequentially for the next positions $t_k = t_0 + kh$, where k is the index of the window positions. The results obtained

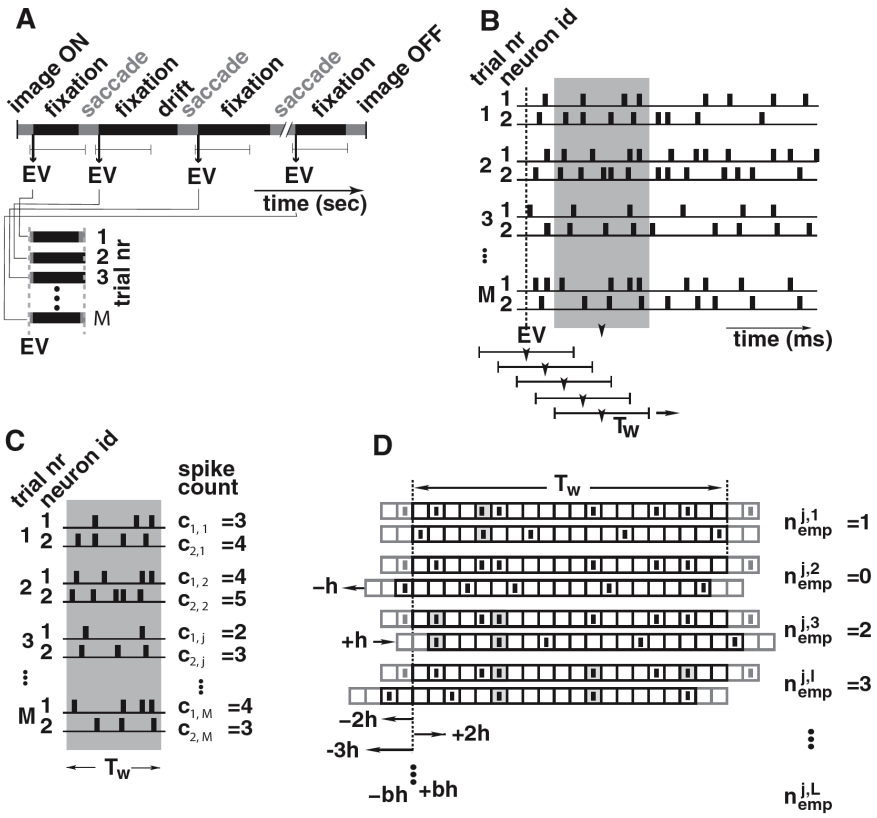

Figure [2].2: Unitary event (UE) analysis. A: trial alignment and sliding window approach. The simultaneous spike trains of 2 neurons are aligned according to the relevant behavioral event (EV, exemplified for fixations). The data extracted in relation to such an event were from -25 ms before until 325 ms after the event. For analysis we applied a sliding window T_w (width 50 ms) that was slid along the data in the time resolution of the data of 0.1 ms. At each position of the window the UE analysis was newly performed. B: analysis within a window. Per trial j and per neuron i the number of spikes $c_{i,j}$ within the window is extracted. From these the data firing probabilities per trial per neuron are calculated by dividing the spike counts by the number of bins N in the window (here: 500 bins of h = 0.1 ms). These probabilities then serve to calculate the expected number of coincidences. C: extracting the empirical number of coincidences in a single trial. Within a single trial the number of coincidences is extracted as the joint spike events of the 2 neurons occurring within the same bin (of 0.1 ms) (marked in light gray). D: to also detect coincidences with small temporal jitter, the spike trains are shifted against each other in units of the bins, with increasing shift. One neuron is kept fixed as a reference, whereas the spike train of the other is shifted forward and backward up to the maximal allowed jitter b (here 5 ms). The total number of coincidences within the trial is then given by the sum of all L = 2(b/h) + 1 individual shifts. For simplicity we omitted the time indices of the window positions in the figures.

for a given window position were plotted at the time point corresponding to the center of the windows (Fig. [2].2B, indicated by arrows) to obtain time-dependent measures of UEs.

The basic operation in the UE method is to compare the empirical number of coincidences to the expected number that would occur by chance given the firing rates of the neurons. Additionally, to account for nonstationary rates across trials (Grün et al., 2003), as are typically present in data of awake behaving animals and in particular in the data treated here, the relevant measures are obtained from single-trial analysis and only subsequently summed across trials. Thus within the analysis window the expected number of coincidences is calculated on the basis of the trial-by-trial firing probabilities $p_{i,j}$, which are estimated by the spike count $c_{i,j}$ of neuron i in trial j divided by the number of bins N within a window: $p_{i,j} = c_{i,j}/\mathrm{N}$, where $\mathrm{N} = T_w/\mathrm{h}$ (Fig. [2].2C). The joint probability for finding a coincidence by chance per trial is calculated by the product of the single-neuron firing probabilities $p_{1,j} \cdot p_{2,j}$. The expected number of coincidences per trial j results from multiplying this probability with the number of bins N that are included in the analysis window: $n^j_{exp} = p_{1,j} \cdot p_{2,j} \cdot N$. The total number of expected coincidences within the window is derived from the sum of the expected numbers per trial: $n_{exp} = \sum_{j=1}^{M} n^j_{exp}$.

The total empirical number of coincidences per window is derived simply by counting the coincident events (spikes of both neurons within the same bin of width h) within a trial and taking the sum across all the trials. However, being interested in detecting coincidences that may occur with a small temporal jitter (e.g., up to b = 5 ms), we applied the "multiple-shift" approach (Grün et al., 1999). In this method the coincidences within a single trial are detected by shifting the second spike train against the first (reference) spike train (see Fig. [2].2D). The shift is stepwise increased from 0 in steps of h = 0.1 ms up to b = 5 ms (for positive and negative shifts). For each shift l of the L = 2(b/h) + 1 shifts, the exact coincidences $n^{j,l}_{emp}$ are counted. The total number of coincidences considering all shifts then yields the total number of empirical coincidences in that trial

$$n^j_{emp} = \sum_{l=1}^{L} n^{j,l}_{emp}$$

(Fig. [2].2D). The total sum of the empirical coincidences in a window is then given by the sum of all coincidences across the trials

$$n_{emp} = \sum_{j=1}^{M} n^j_{emp}$$

Finally, we compared the empirical n_{emp} to the expected number n_{exp} of coincidences to detect significant deviations. To this end, we calculated the joint-p-value jp [i.e., the probability of getting the given number of empirical coincidences (or an even larger number) under the null hypothesis of independent firing]. The distribution that represents the probability to find a given number of coincidences under this null hypothesis is given analytically for the case of Poisson processes. Then the coincidence counts follow a Poisson distribution, with its mean corresponding to the expected number of coincidences (for the analytical derivation see Grün et al., 2002a). (See also the discussion on deviation from the Poisson assumption in Controls in RESULTS.) Then the significance of n_{emp}, given n_{exp}, measured as the joint-p-value, yields (Grün et al., 2002a)

$$jp(n_{emp}, n_{exp}) = \sum_{r=n_{emp}}^{\infty} \frac{n_{exp}^r}{r!} exp(-n_{exp})$$

If its value is below an a priori threshold (here chosen as 5 %) coincident firing (synchrony) is classified as significant and the unitary events identified in this way are expressed as unitary event rate in a particular window. For better visualization of the significance values (see Fig. [2].3D) we use a log-transformed measure known as the joint-surprise (Grün et al., 2002a). In summary, for each window position we calculate the expected number of coincidences, the empirical number of coincidences, and the unitary events, all expressed as rates, thus obtaining time-dependent functions of the respective measures: $n_{emp}(t)$, $n_{exp}(t)$, and $UE(t)$. The described analysis was performed for each pair of simultaneously recorded neurons and separately for each session. To derive population statistics, the time-dependent functions were averaged across all neuron pairs for a particular type of trials (fixations or saccades). This yielded the average coincidence rates (expected and empirical) and the average UE rate as functions of time. An example of a UE analysis for a single pair is shown in Figure [2].3.

[2].3 Results

Eye movements during free viewing

During free viewing of natural images, the animals performed eye movements that were categorized in saccades or fixation periods. Figure [2].1, A-C shows one example of the presented visual scenes together with single-unit responses and the corresponding eye-movement traces. We first examined distributions of the various eye movements for all "image" and "blank" conditions. These distributions differed markedly in the two conditions (compare Fig. [2].1D and E with Fig. [2].1F and G). With images, saccades

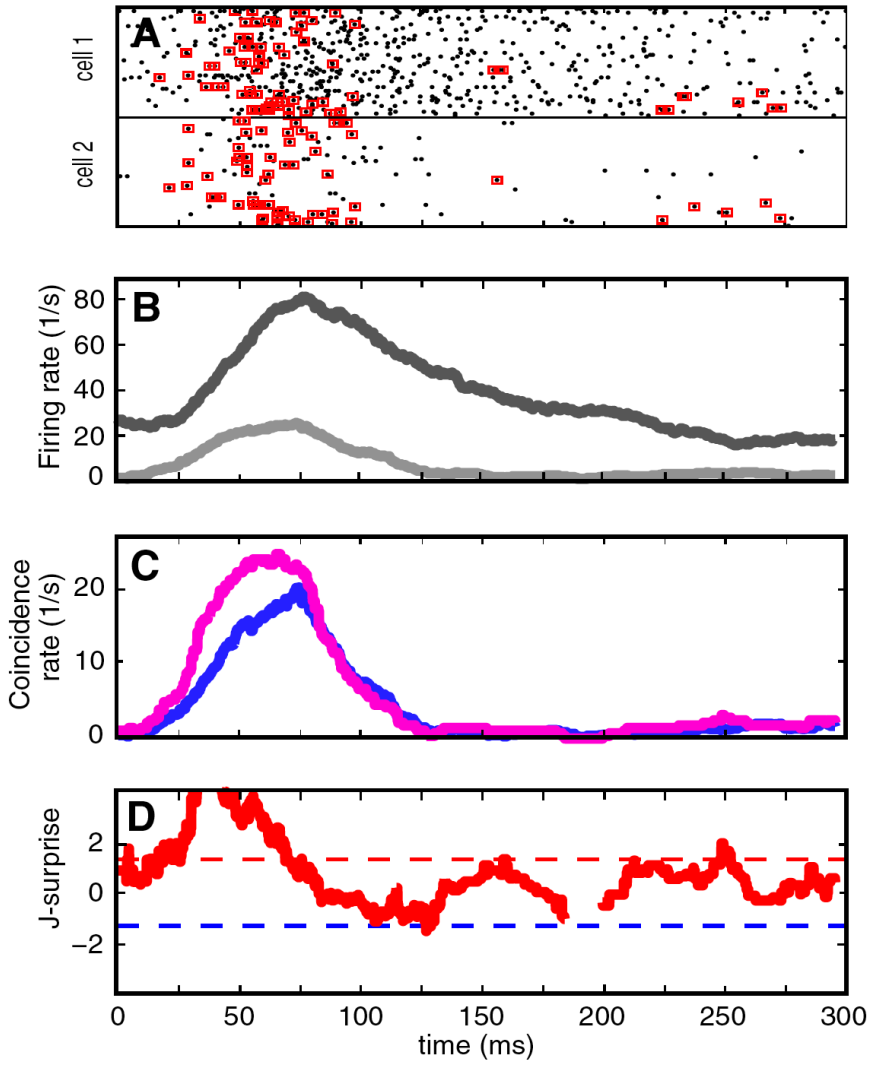

Figure [2].3: Example of pairwise UE analysis. A: raster plots of spike responses to 57 fixations on images of cell 1 and cell 2 recorded simultaneously. Time 0 represents the onset of fixation. Red squares indicate unitary events. UE analysis was performed within sliding windows of $T_w = 50$ ms. B: smoothed poststimulus time histograms (boxcar with a window of the same width as used for UE analysis) of cell 1 (dark gray) and cell 2 (light gray). C: expected number of coincidences (blue) and empirical number of coincidences (pink) expressed as rates. There are many more coincidences than expected at the rising phase of the firing rate response. D: significance of UEs expressed by the surprise measure. The surprise is not shown for values of $-\infty$ for windows in which one neuron does not have any spikes (i.e., 180-200 ms). Dashed lines correspond to a significance level of 5% (used here for UE detection): red for excess coincidences, blue for lacking coincidences.

and fixation periods occurred typically every 250-300 ms, about sixfold more frequently than during the scanning of blank frames. Saccades were significantly longer during the blank condition than during the image condition and fixation times were scattered more broadly (Kolmogorov-Smirnov, $P < 0.001$). The distribution of the fixation durations in the image condition shows a clear peak at 230 ms, with 57 % of fixations lasting between 150 and 350 ms. The corresponding distribution in the blank condition is more dispersed and fixation times are longer.

Firing rates during free viewing

We simultaneously recorded one to three cells per tetrode with at least eight tetrodes and analyzed 418 single units, individually identified by a manual clustering method (Gray et al., 1995). The receptive fields (RFs) of most units were located within 5-10^o from the center of gaze and were <2^o. About 15 % of the recordings were made below the opercular layer. The RFs of these cells were located within 15^o and their sizes were <5^o. We pooled the cells from these two regions of V1 because we assumed that there are no basic differences in firing statistics. No further RF properties were determined to save time for data acquisition during free viewing.

The significance of firing rate changes during saccades was assessed by comparing the rates in the first 20 ms after the onset of saccades with those in the subsequent 20-ms interval using a t-test ($P < 0.05$). For fixation we compared the first 50-ms interval with the consecutive 50-ms interval after the onset of the event. In the image condition, 30 % of the cells significantly increased their firing rate (Table [2].1) during the fixation periods and 23 % of the cells significantly decreased their firing during saccades. In the blank condition significant rate changes occurred in <2 % of the cells (Table [2].1). We examined the time course of the firing rate changes for all units by computing averaged histograms (1-ms bin width) using the onset of saccades or fixations as trigger. Given that each saccade or fixation had a different duration, each rate point (bin) is expressed as the average over the entire sample of recorded neurons, normalized to the number of neurons and trials contributing to that bin. The plots are truncated at time points beyond which <90 % of the neurons contributed to the averages. Figure [2].4A summarizes the results for the image condition. Whenever the animals fixate, average firing rates increase, with a peak at about 90 ms after fixation onset, and then steadily decrease despite ongoing visual fixation. In the grand average, saccadic eye movements were associated with a decrease in firing rate consistent with saccadic suppression (Burr et al., 1994). In the blank condition (Fig. [2].4B) all changes in firing rate started from a lower baseline level than that in the image condition. Fixation periods were associated with an increment in firing rate that was smaller and developed more slowly and was more sustained than that in the image

	Changes in Firing Rate (Percentage of Cells, n=418)			Incidence of Synchronous Firing (Percentage of Total Pairs, n=1369)
	A	B	C	D
	Increase	Decrease	No Change	Unitary Events
Cue fixation	1.4	1.2	97.4	21.3
Image saccades	0.4	23.2	76.4	10.8
Image fixations	30.4	0.2	69.4	34.0
Blank saccades	0.2	1.4	98.4	0.4
Blank fixations	0.7	0.2	99.1	5.0

Summary statistics of the full data set.

Table [2].1: A-C: percentage of cells showing significant changes in firing rates during the the different phases of scene inspection. The significance of changes during saccades was assessed by comparing the rates in the first 20 ms and the consecutive 20 ms following the onset of saccades and for fixations by comparing the first and consecutive 50 ms after the onset of the respective events (t-test). Positive or negative differences were classified as increments and decrements, respectively. D: percentage of cell pairs showing UEs in the intervals from 20 to 50 ms (saccades) and 50 to 100 ms (fixations) after the onset of the respective events.

condition. This modulation of firing rate is most likely due to corollary attention-related activity associated with eye movements (Buisseret and Maffei, 1977; Purpura et al., 2003). However, these changes were significant in only 4 of 418 single units. Saccadic eye movements were not associated with modulations in rate. During fixation of a single fixation spot displayed on an otherwise blank screen, firing rates increased rapidly toward a level corresponding to baseline activity in the image condition and then decayed slowly. This increase was significant in only 2 % of the neurons. In this case measurements were triggered when the axis of gaze entered the fixation box (1^o). Thus the first window contains activity associated with the end of the saccade preceding the fixation.

Temporal relations among the discharges of simultaneously recorded cells (UE analysis)

We hypothesized that rate codes are complemented by a temporal coding strategy and analyzed the data for coordination of spike timing among the discharges of simultaneously recorded neurons. Therefore we examined correlations between the firing patterns of 1,369 cell pairs that were obtained in 80 recording sessions with $\leq$12 simultaneously registered neurons. Altogether, 4851 trials were analyzed with images and 383 in the blank condition. For correlation analysis we examined responses during the first 50 and

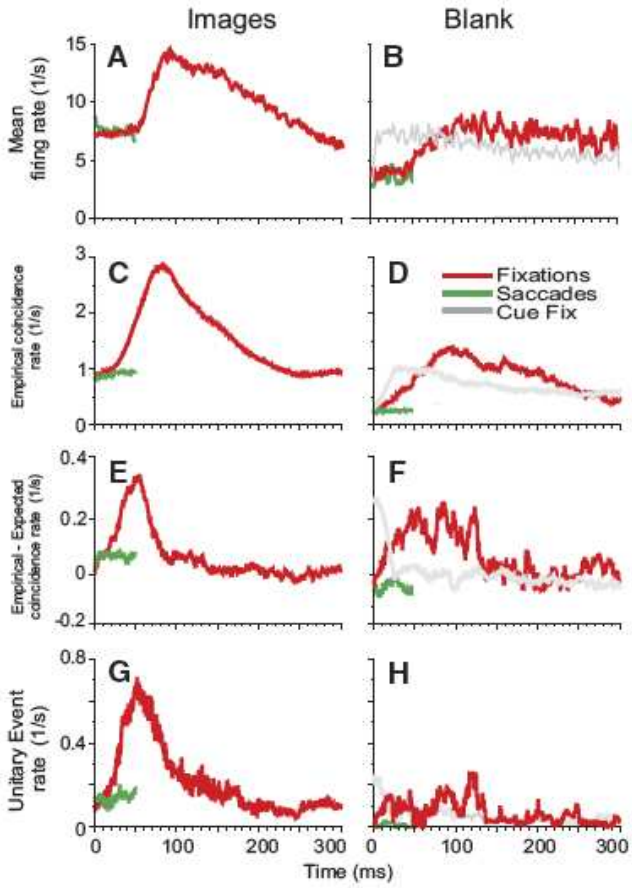

Figure [2].4: Time course of changes in firing rates and synchrony associated with different eye movements. Plots are color coded (see D) for images (left column) and blank condition (right column). The gray lines in the right column refer to the condition where the animals saccade toward and fixate a central spot on an otherwise blank screen. In this condition the cells' receptive fields (RFs) fall on the dark screen. Traces for saccades are shorter than those for fixations because of the inherently shorter duration of these events. A and B: firing rates computed for and averaged over 418 single units. For each 1-ms bin, the averaged rate is computed by normalizing to the number of contributing neurons and to the corresponding number of trials per neuron. Abscissa: time from onset of saccades and fixations. C and D: frequency of empirical coincidences, defined as pairwise joint firing within 5-ms intervals computed from 1,369 simultaneously recorded cell pairs. E and F: difference of the empirically observed and the expected coincidence rates, the latter being derived from the corresponding firing rates. G and H: occurrence of unitary events (UEs), i.e., events of joint firing that significantly exceed (5 %) the level of coincident firing expected from the firing rates (Grün et al., 1999, 2002a,b, 2003).

300 ms following the onset of saccades and fixations, respectively. Because of the variable duration of these eye events, these interval durations were chosen to get a reasonable number of trials for reliable statistics (see METHODS). The estimation of temporal relations among the spikes of simultaneously recorded neurons becomes difficult when discharge rates are low and highly nonstationary, as is the case during the inspection of complex scenes. Therefore we used the UE analysis. Here, we address correlated firing as unitary events if the discharges of two neurons coincide within a window of 5 ms and if these coincidences significantly exceed what is predicted by the actual firing rates. To account for nonstationary firing rates, comparisons are computed between the observed (empirical) number of coincidences and the expected number of coincidences within a sliding window (here: 50 ms wide) that is moved in 0.1-ms steps along the time series. In addition, possible non-stationarities of firing rates across the trials are controlled for by estimating the firing rates and the resulting expected joint-probabilities on a trial-by-trial basis within the sliding windows.

Table [2].1 (under heading D) shows the incidence of UEs, i.e., of coincident firing that significantly exceeds expected levels for all cell pairs. During fixation on images, >34 % of all cell pairs (n = 1,369) exhibited UEs. During cue fixations and saccades on images, 21 and 10 % of all pairs showed UEs, respectively, but the mean rate of UEs was much lower than that during fixations on images (see Fig. [2].4H). In all other viewing conditions the percentage of cell pairs showing UEs was <6 % and thus does not exceed the chosen level of significance (5 %). The eye-movement-related time courses of the observed coincidences and UEs, averaged over all cell pairs, are shown for the image condition in Figure [2].4G and H. As expected, coincidences increase with firing rate but the precise time courses of the respective increases differ (compare Fig. [2].4A and G). The reason is that there is significantly more coincident firing than expected from the respective rates during the early response phase. As the time course of the UEs indicates, this excess synchronization of firing is transient, lasts for only about 40 ms, reaches a maximum already 50 ms after onset of fixation (Fig. [2].4G), and decreases before the firing rates reach their maximum. This excess synchronization is so marked that it is already clearly visible in the difference between the observed coincidences and the coincidences predicted by the actual firing rates (Fig. [2].4E and F). In the condition "fixation on blank" (Fig. [2].4B, D, F, and H), the observed coincidences are slightly higher than expected from the rates (Fig. [2].4F), although this increase is not significant (Fig. [2].4H).

Controls

It can be argued that synchronous events are particularly likely if firing rates increase jointly and abruptly, as occurs in response to rapidly changing stimuli. We examined this

issue by intentionally destroying the exact timing of spikes while keeping the firing rate profiles approximately constant. This is achieved by creating surrogate data by "dithering" each individual original spike of each of the neurons within a given window around the spike and reanalyzing the newly generated parallel time series in exactly the same way as the original data. Similar approaches are described in Date et al. (1998); Gerstein (2004), and Shmiel et al. (2006). As dithering width was increased from 0 to ±50 ms (centered at spike times), the maxima of the difference curves between observed and expected coincidences decreased considerably (Fig. [2].5D). Accordingly, the UE rate dropped to about 75 % at ±5 ms, to 50 % at ±10 ms, and approximated the level expected in the case of independent firing at about ±40 ms (Fig. [2].5E). This decay rate corresponds exactly to theoretical predictions derived for empirical coincidences that are successively dithered under consideration of the size of the coincidence window and the temporal jitter of the coincidences, indicating that the UEs were not due to transient rate co-variations but are indications of excess synchrony (Pazienti et al., 2007, 2008). Pazienti et al. (2008) showed that without the presence of excess synchrony there would be no decay at all. Unavoidably, the firing rate profiles were also affected by the dithering and got smoothed with increasing dither (Fig. [2].5A). As a consequence, the time course of the expected coincidences was also affected (Fig. [2].5C). However, the effect of dithering was much more pronounced for the observed than for the predicted coincidences. Dithering completely abolished the peak of observed coincidences around 50 ms after fixation onset. Thus dithering is effective only on coincidences, which the UE analysis had identified as nonspurious, thereby supporting the interpretation that the UEs arise from active synchronizing processes.

To further corroborate this interpretation, we simulated spike trains that were either completely independent or contained a known number of deliberately inserted coincident spikes and then analyzed the data in exactly the same way as the experimental data. To test whether the nonstationary rate profile of the experimental data (Fig. [2].6B, gray) was causing a UE rate profile similar to that of the experimental data (Fig. [2].6A) we simulated pairs of inhomogeneous Poisson spike trains with a firing rate profile that was identical to that of the original data (Fig. [2].6B, gray). Thus these data included a possible source for false-positive correlations, i.e., covariation of firing rate (Grün et al., 2003). We generated a data set for 100 pairs of neurons and 100 trials per pair. For the independent spike trains, the resulting UE rate (averaged over all pairs) (Fig. [2].6D) shows no distinct peak and fluctuates around the low level expected from the 5 % significance threshold. In contrast, when we injected additional coincident spikes (with a coincidence rate profile shown in Figure [2].6B, black) and reduced the background rate

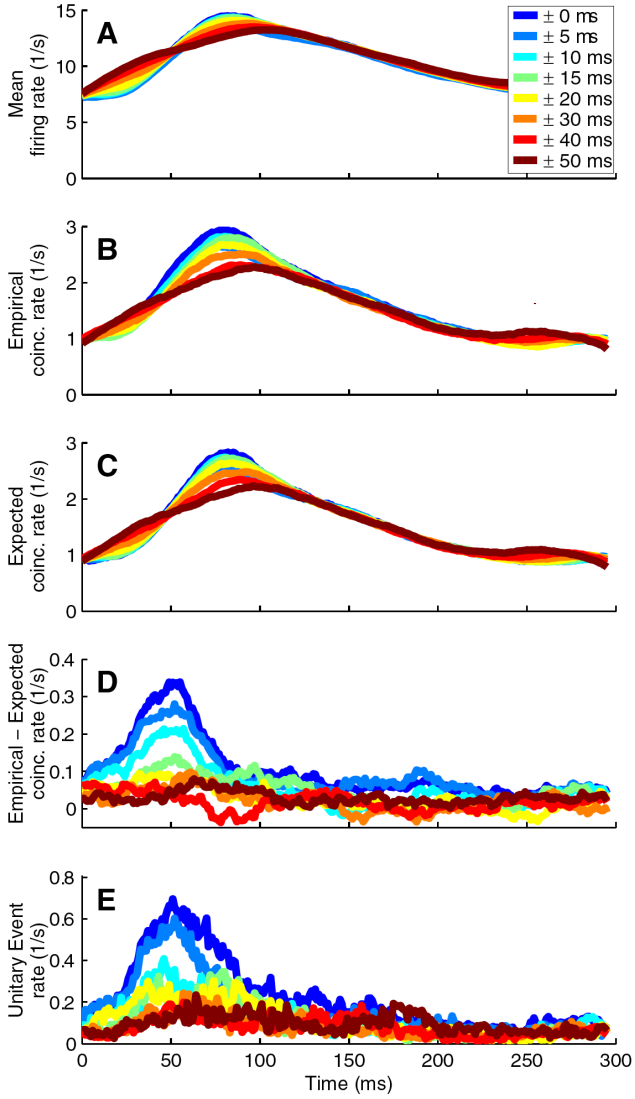

Figure [2].5: Decreased synchronization with increasing dithering of spike times. Dithering was applied to each individual spike of all data and the dithered time series were analyzed for UEs as the original data (same parameters as in Fig. [2].3, i.e., coincidence width, 5 ms; sliding window width, 50 ms; here: shown for alignment relative to fixation onset only). Dither values varied from ±5 to ±20 ms in 10-ms steps, and from ±30 to ±50 ms in 20-ms steps (color code in C). A: average firing rates of all analyzed neurons. Abscissa (for all plots): time from onset of fixations. B: frequency of empirical and (C) expected coincidences. D: difference of the empirical and the expected coincidences. E: significant coincidences/Unitary Events. Note that dithering does not destroy all of the original coincidences. These remaining coincidences depend on the ratio of the bin width (here 5 ms) over the dither width.

such that the total rate and the rate profile of the neurons remained the same as those in the previous simulation, a prominent peak of UEs was detected. The choice for the time course of the coincidence rate was based on one of our theoretical studies in which we tested the detectability of intervals that included coincidences (Grün et al., 2002b). We assumed Poisson statistics in the UE analysis, although being aware of the fact that experimental spike trains typically deviate from that assumption. In the study of Pipa et al. (2008) we showed that the spike train autostructure indeed has an impact on the significance estimation of coincident spike events. We specifically tested for spike trains with interspike interval (ISI) statistics given by a gamma-distribution and systematically varied the shape factor of the ISI distribution, thereby modifying spike trains from highly irregular to increasingly more regular firing. We found the false-positive rate to be at or below the expected level, given the a priori significance level for realistically regular spike trains (coefficient of variation values ranging from 0.2 to 0.8; Nawrot et al., 2008). Only for unrealistically regular or highly irregular spike trains is the false-positive rate slightly enhanced. Thus if spike trains are more regular than Poisson, the significance test assuming Poisson statistics rather underestimates the significance. To test whether our result of excess synchrony was not a result of false positives in the case of spike trains more irregular than Poisson (e.g., as given by bursty processes), we performed an additional control. We debursted the data, by excluding spikes that occurred within a time interval of <5 ms after a former spike and reanalyzed the data again. Although in some data sets the firing rate was considerably reduced, the occurrence and the time course of the UEs, as shown for the original data, were preserved (not shown here). An alternative of the significance test-based on an analytical description of the coincidence distribution for such processes-is not known and hard to derive (see Pipa et al., 2008). Generating the coincidence distribution by surrogate data is prone to errors: trial shuffling is not appropriate due to cross-trial nonstationarity, ISI shuffling would destroy potential second-order structures within the spike trains, and other methods would affect the firing rates.

Grün et al. (2002b) showed that by varying the width of the analysis window it is possible to detect the existence and the width of the interval over which coincident spikes were present. For analysis window sizes smaller that the interval of coincidences, not all coincidences are detected and the UEs are underestimated. At the width of the window that matches the interval of excess coincidences the significance curve exhibits a peak and saturates for even larger windows. The center of the maximum always stays at the same time. This is also evident in the data presented here (Fig. [2].7). By varying the width of the analysis window from 20 to 60 ms we always found a maximum at 50 ms after fixation, with a tendency of saturation after a width of 50 ms. This implies that regardless of the

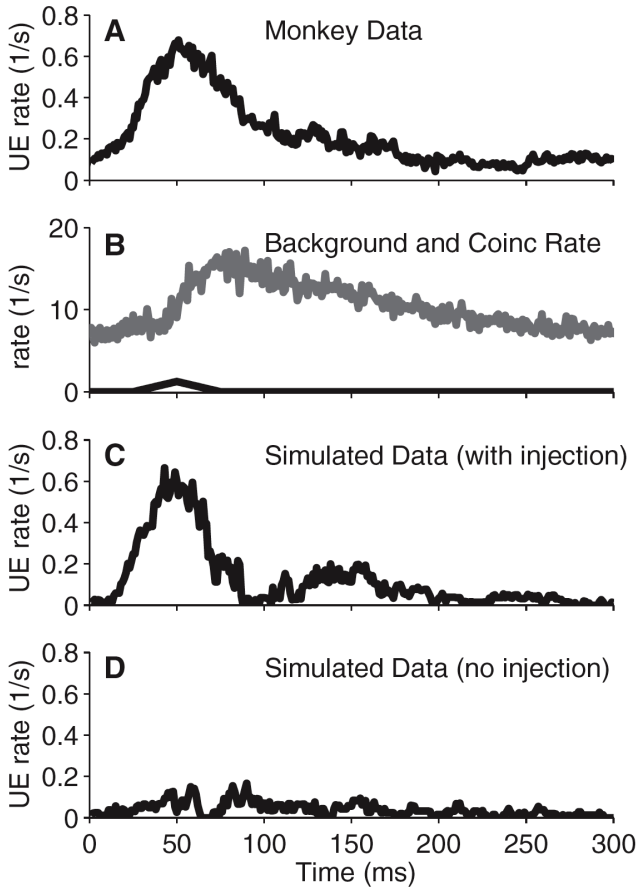

Figure [2].6: Comparison of real and simulated data. A: UE rate of the experimental data during fixation aligned to fixation onset (as also shown in Fig. [2].3). B: firing and coincidence rate profile used for the simulation. The gray curve shows the average firing rate of all neurons (1-ms resolution). This firing rate profile is then used in the simulations. For the case of independent neurons (D), this time course is directly used to simulate inhomogeneous Poisson spike trains. For dependent neurons (C), the firing of the simulated neurons is composed of independent firing (background firing rate) and injected coincidences at a rate corresponding to the profile shown in B (black). In the case of injection of coincident spikes, the background firing rate was reduced by the coincidence rate such that the neurons again had the same total rate as that in the independent case. Each neuron was drawn with the same rate profile and one neuron pair was simulated for 100 trials. Coincidences were inserted with a temporal jitter of 5 ms. In all, 100 neuron pairs were simulated and then analyzed in the same way as the experimental data (sliding window of 50 ms, allowed coincidence width of 5 ms). Time course of the UE rates of the simulated data with (C) and without (D) coincidence injection.

width of the analysis window, the method reliably detects the time course of nonspurious coincidences that seem to be spread over an interval of about 50 ms. Taken together,

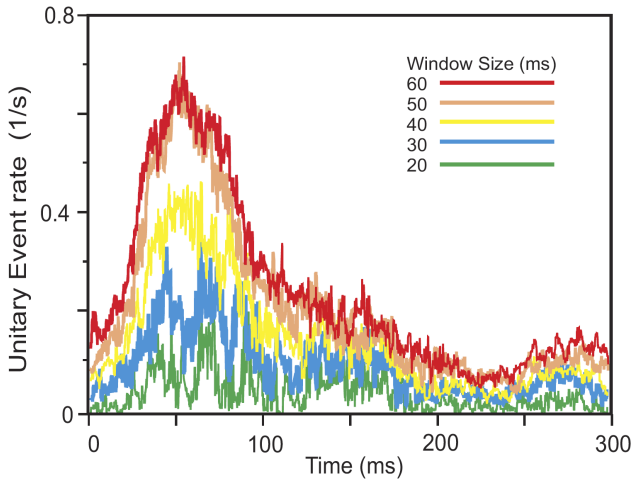

Figure [2].7: UE profile as a function of analysis window width. Data were analyzed as for Figure [2].4G for various analysis window widths (20 to 60 ms in 10-ms steps, color coded). The UE profiles exhibit an increasing maximum for increasing analysis window width, which is most strongly expressed for window size 40-60 ms. The peak height saturates for window sizes >50 ms. The peak of these curves is at the same point in time (50 ms after fixation onset).

these controls indicate that the unitary event analysis applied here reliably detects the incidence and the time course of excess synchrony in nonstationary time series and ignores spurious, rate-dependent coincidences.

[2].4 Discussion

In contrast to previous studies (e.g., Vinje and Gallant, 2000), we have not attempted to establish relations between eye movements and neuronal responses resulting from the passage of individual RFs over particular contours of the visual scene. Rather, we were interested in the relations between response variables such as discharge rate and spike synchronization during the various phases of visual scanning movements. Because our animals were free to direct their gaze to any location on the images, the actual stimulus for the cells changed with each saccade. Thus we measured the average behavior of cells under natural viewing conditions. We reasoned that this should allow us to detect characteristic, eye-movement-related response patterns and the results suggest that this was indeed the case. The application of multielectrode recordings revealed for the first time

a fixation-related, transient increase of response synchronization. This increase in coincident firing is dissociated from the later increase in neuronal firing rate, suggesting that temporal coordination of discharges may play a role in the initial processing of signals in V1.

Methodological considerations

Because of the spike-sorting procedure that we applied to extract single-cell responses from the tetrode signals it is likely that we underestimated the incidence of synchronous firing. If cells recorded from the same tetrode discharge in precise synchrony, the respective spikes superimpose in a complex way and are excluded from analysis (Gray et al., 1995). This limited our ability to detect synchronous firing of pairs of neurons recorded by the same tetrode and it also precluded detection of synchrony among cells recorded from different tetrodes if two cells recorded from one tetrode fired in precise synchrony with a third cell, recorded by another tetrode. Of a total of 1,369 pairs 262 were cells recorded from the same tetrode and 1,107 from distant electrodes. Thus we may have underestimated the incidence of synchronous firing by $\gtrsim 20$ %.

Functional implications

Several arguments suggest that the unexpectedly high incidence of synchronous firing was due to synchronizing neuronal interactions rather than a by-product of stimulus-induced covariation in firing rate. First, the statistical methods developed for the analysis of synchronous firing, the unitary event analysis, is specifically designed to eliminate spurious coincidences caused by joint rate increases. Second, the controls with dithering spike times provided independent evidence that the UEs were not due to rate covariations. Third, our simulations with artificially injected spikes showed that pure rate covariation could not explain the high incidence of UEs.

Interestingly, the phase of enhanced synchrony set in right at the beginning of the rate increases and reached its maximum 40 ms before the peak of the rate increase. Because UEs are virtually absent during saccades and because retinal response latencies cover a range of 30 to 60 ms, this implies that the synchronizing mechanism operates already and preferentially on the very first spikes evoked by the newly fixated parts of the image.

A mechanism capable of achieving such rapid synchronization has recently been identified (Fries et al., 2007). When the visual cortex is engaged in oscillatory activity in the gamma-frequency range, neurons in columns with related feature preferences oscillate in synchrony probably because of strong reciprocal connections between these columns. This has the effect that the onset latencies of responses to stimuli driving coherently oscillating columns become synchronized because they can fire only during the depolarizing phase

of the oscillation cycles (Engel et al., 2001; Fries et al., 2001, 2007, see also Volgushev et al., 1998, for a related in vitro study). As a result this may lead to an increased synchronization of the very first spikes of responses to stimuli that match the preferences of synchronously oscillating columns.

A further mechanism suggesting that visually induced first spikes have a supportive role in synchronization, based on an internal, saccade-locked signal, was recently hypothesized. It was shown that there is a pronounced saccade-related, evoked local field potential (LFP) signal (Bartlett et al., 1976). This in turn could serve as a reference signal to realize the proposed amplitude-to-temporal code conversion suggested by Fries et al. (2007). Thus the saccade-evoked LFP would provide the reference frame for the timing of induced spikes and, in particular, would favor the first incoming spikes. There are first results supporting this hypothesis (Ito et al., 2008) that will be reported elsewhere.

For scene segmentation and object perception, stimulus features that share similarities tend to be bound together. Since neurons encoding related features engage in synchronous gamma oscillations, it has been proposed that these oscillations serve the translation of connectivity patterns into specific constellations of response synchronization that in turn support rapid feature-specific synchronization of discharges (Fries et al., 2001). This proposal is based on the evidence that precise synchronization of spikes selectively raises the saliency of the synchronized discharges and thereby favors their further joint processing (for review see Singer, 1999). The rapid synchronization of spikes in the early postsaccadic fixation could thus be a reflection of such a temporal coding mechanism.

It has been concluded from measurements of recognition times that in V1 computations required for elementary scene analysis have to be accomplished within a few tens of milliseconds (Fabre-Thorpe et al., 2001). Because this allows for only a few spikes per neuron, it has been proposed that information should be encoded not only in the discharge rate but also in the precise temporal relations between the individual discharges of neurons. The present results support this scenario and suggest that the visual cortex possesses a mechanism to complement rate codes with a fast temporal coding strategy adjusting the timing of individual spikes. As suggested by the saccade-related occurrence of oscillatory activity, this precise synchronization could be due to saccade-related oscillatory activity of corollary origin. Several arguments make it appear likely that this corollary signal is related to mechanisms of attention. There is abundant evidence for a close relation between the mechanisms that control attention and saccadic eye movements, respectively (Corbetta et al., 1998; Everling, 2007). Moreover, attentional mechanisms have been shown to modulate the oscillatory patterning of neuronal responses, which in turn modulates synchronous discharges. Focused attention and expectations are associated with an

anticipatory entrainment of cortical networks in both beta and gamma oscillations (Fries et al., 2001; Roelfsema et al., 1997) and responses to attended stimuli exhibit stronger oscillatory patterning in the gamma-frequency range and stronger synchronization than responses to nonattended stimuli (Schoffelen et al., 2005; Steinmetz et al., 2000; Tiitinen et al., 1993). To further examine this possibility, field potential oscillation must be analyzed and related to the timing of spike discharges.

However, at this stage the possibility should not be dismissed that spike synchronization is caused in addition by retinal or thalamic mechanisms. There is evidence that the discharges of retinal ganglion cells are not independent and can become synchronized with millisecond precision in a stimulus-specific way (Meister et al., 1995; Neuenschwander and Singer, 1996). These synchronized volleys are transmitted with high temporal fidelity by the lateral geniculate nucleus (Castelo-Branco et al., 1998; Usrey and Reid, 1999). Moreover, corticofugal projections to the lateral geniculate have been shown to have a synchronizing effect (Sillito et al., 1994). Both the retinal and the corticofugal synchronizing mechanisms induce coincident firing in neurons responding to continuous contours, suggesting the possibility that this early synchronization contributes to scene segmentation (Stephens et al., 2006). So far these synchronizations have been studied with stationary flashed or continuously moving contours. Therefore it is unknown whether the stimulation caused by the sudden arrest of rapidly moving contours, such as those occurring at the end of saccades, can induce such synchronization.

In conclusion, our data revealed that scanning of natural scenes is associated with a rapid succession of distinct fixation-related activation patterns that consist of a transient increase in excess coincident firing that is followed by a transient increase in discharge rate. This suggests the action of a fast-acting synchronizing mechanism that is in good agreement with the evidence that basic operations of scene analysis must be accomplished within a few tens of milliseconds in V1, most likely requiring temporal coding strategies.

[2].5 Grants

This study was supported in part by Volkswagen Stiftung grants to P. Maldonado, W. Singer, and S. Grün; Iniciativa Cientifica Milenio Grant P04-068F, and a Stifterverband für die deutsche Wissenschaft grant to S. Grün.

[2].6 Acknowledgments

We thank the Chilean Primate Center-OMS Catholic University of Chile; C. M. Gray for useful comments on the manuscript; and M. Diesmann for clarifying discussions. Part of the work was performed while S. Grün was located at the Free University, Neuroinformatics, Institute for Biology (Neurobiology), Berlin, Germany.

References

J. R. Bartlett, R. W. Doty, B. B. Lee, and H. Sakakura. Influence of saccadic eye movements on geniculostriate excitability in normal monkeys. *Exp Brain Res*, 25:487–509, 1976.

P. Buisseret and L. Maffei. Extraocular proprioceptive projections to the visual cortex. *Exp Brain Res*, 28:421–425, 1977.

D. C. Burr, M. C. Morrone, and J. Ross. Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature*, 371:511–513, 1994.

M. Castelo-Branco, S. Neuenschwander, and W. Singer. Synchronization of visual responses between the cortex, lateral geniculate nucleus, and retina in the anesthetized cat. *J. Neurosci.*, 18:6395–6410, 1998.

M. Corbetta, E. Akbudak, T. E. Conturo, A. Z. Snyder, J. M. Ollinger, H. A. Drury, M. R. Linenweber, S. E. Petersen, M. E. Raichle, D. C. Van Essen, G. L. Shulman, and D. C. Van Essen. A common network of functional areas for attention and eye movements. *Neuron*, 21:761–773, 1998.

A. Date, E. Bienenstock, and S. Geman. On the Temporal Resolution of Neural Activity (Technical report). *Providence, RI: Brown Univ., Division of Applied Mathematics,*, 1998.

A. K. Engel, P. Fries, and W. Singer. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.*, 2:704–716, 2001.

S. Everling. Where do i look? From attention to action in the frontal eye field. *Neuron*, 56:417–419, Nov 2007.

M. Fabre-Thorpe, A. Delorme, C. Marlot, and S. Thorpe. A limit to the speed of processing in ultra-rapid visual categorization of novel natural scenes. *J Cogn Neurosci*, 13:171–180, 2001.

S. Friedman-Hill, P. E. Maldonado, and C. M. Gray. Dynamics of striate cortical activity in the alert macaque: I. Incidence and stimulus-dependence of gamma-band neuronal oscillations. *Cereb. Cortex*, 10:1105–1116, 2000.

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

P. Fries, D. Nikolic, and W. Singer. The gamma cycle. *Trends Neurosci.*, 30:309–316, 2007.

G. L. Gerstein. Searching for significance in spatio-temporal firing patterns. *Acta Neurobiol Exp (Wars)*, 64:203–207, 2004.

C. M. Gray, P. E. Maldonado, M. Wilson, and B. McNaughton. Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J. Neurosci. Methods*, 63:43–54, Dec 1995.

S. Grün, M. Diesmann, F. Grammont, A. Riehle, and A. Aertsen. Detecting unitary events without discretization of time. *J. Neurosci. Methods*, 94:67–79, 1999.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Comput*, 14:43–80, 2002a.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput*, 14:81–119, 2002b.

S. Grün, A. Riehle, and M. Diesmann. Effect of cross-trial nonstationarity on joint-spike events. *Biol Cybern*, 88:335–351, 2003.

R. Guyonneau, R. Vanrullen, and S. J. Thorpe. Temporal codes and sparse representations: a key to understanding rapid processing in the visual system. *J. Physiol. Paris*, 98:487–497, 2004.

J. J. Hopfield. Encoding for computation: recognizing brief dynamical patterns by exploiting effects of weak rhythms on action-potential timing. *Proc. Natl. Acad. Sci. U.S.A.*, 101:6255–6260, 2004.

J. Ito, P. Maldonado, and S. Grün. Saccade-related lfp oscillations set the stage for processing visually evoked spikes. *Computational and Systems Neuroscience (COSYNE)*, Salt Lake City, UT, 2008.

H. E. Jones, K. L. Grieve, W. Wang, and A. M. Sillito. Surround suppression in primate V1. *J. Neurophysiol.*, 86:2011–2028, 2001.

S. J. Judge, B. J. Richmond, and F. C. Chu. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res.*, 20:535–538, 1980.

C. J. McAdams and J. H. Maunsell. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.*, 19:431–441, 1999.

M. Meister, L. Lagnado, and D. A. Baylor. Concerted signaling by retinal ganglion cells. *Science*, 270:1207–1210, 1995.

M. P. Nawrot, C. Boucsein, V. Rodriguez Molina, A. Riehle, A. Aertsen, and S. Rotter. Measurement of variability dynamics in cortical spike trains. *J. Neurosci. Methods*, 169: 374–390, 2008.

S. Neuenschwander and W. Singer. Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. *Nature*, 379:728–732, 1996.

A. Pazienti, M. Diesmann, and S. Grün. Bounds of the ability to destroy precise coincidences by spike dithering. *Advances in Brain, Vision, and Artificial Intelligence (Lecture Notes in Computer Science). Berlin: Springer Berlin/Heidelberg*, 4729:428–437, 2007.

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

G. Pipa, C. van Vreeswijk, and S. Grün. Auto-structure of spike-trains matters for testing on synchronous activity. *(Abstract II-38). Computational and Systems Neuroscience (COSYNE)*, Salt Lake City, UT, 2008.

K. P. Purpura, S. F. Kalik, and N. D. Schiff. Analysis of perisaccadic field potentials in the occipitotemporal pathway during active vision. *J. Neurophysiol.*, 90:3455–3478, 2003.

A. Riehle, S. Grün, M. Diesmann, and A. Aertsen. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278:1950–1953, 1997.

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

J. M. Schoffelen, R. Oostenveld, and P. Fries. Neuronal coherence as a mechanism of effective corticospinal interaction. *Science*, 308:111–113, 2005.

T. Shmiel, R. Drori, O. Shmiel, Y. Ben-Shaul, Z. Nadasdy, M. Shemesh, M. Teicher, and M. Abeles. Temporally precise cortical firing patterns are associated with distinct action segments. *J. Neurophysiol.*, 96:2645–2652, 2006.

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

W. Singer. Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24:49–65, 1999.

P. N. Steinmetz, A. Roy, P. J. Fitzgerald, S. S. Hsiao, K. O. Johnson, and E. Niebur. Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*, 404:187–190, 2000.

G. J. Stephens, S. Neuenschwander, J. S. George, W. Singer, and G. T. Kenyon. See globally, spike locally: oscillations in a retinal model encode large visual features. *Biol Cybern*, 95:327–348, 2006.

Y. Sugita. Grouping of image fragments in primary visual cortex. *Nature*, 401:269–272, 1999.

H. Tiitinen, J. Sinkkonen, K. Reinikainen, K. Alho, J. Lavikainen, and R. Naatanen. Selective attention enhances the auditory 40-Hz transient response in humans. *Nature*, 364:59–60, 1993.

Y. Trotter and S. Celebrini. Gaze direction controls response gain in primary visual-cortex neurons. *Nature*, 398:239–242, 1999.

W. M. Usrey and R. C. Reid. Synchronous activity in the visual system. *Annu. Rev. Physiol.*, 61:435–456, 1999.

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

M. Volgushev, M. Chistiakova, and W. Singer. Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. *Neuroscience*, 83:15–25, Mar 1998.

Chapter 3

Identification of Neurons Participating in Cell Assemblies

Identification of Neurons Participating in Cell Assemblies

Sonja Grün[1,2], Denise Berger[2,3], and Christian Borgelt[4]

[1]Theoretical Neuroscience Group, RIKEN Brain Science Institute, Wako-shi, Japan
[2]Bernstein Center for Computational Neuroscience, Berlin, Germany
[3]Neuroinformatics, Inst. Biology, Freie Universität, Berlin, Germany
[4]European Center for Soft Computing, Mieres (Asturias), Spain

This section is published as

Grün, S.; Berger, D. & Borgelt, C.
Identification of neurons participating in cell assemblies
Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2009, 3493-3496.

Abstract

Chances to detect assembly activity are expected to increase if the spiking activities of large numbers of neurons are recorded simultaneously. Although such massively parallel recordings are now becoming available, methods able to analyze such data for spike correlation are still rare, because it is often infeasible to extend methods developed for smaller data sets due to a combinatorial explosion. By evaluating pattern complexity distributions the existence of correlated groups can be detected, but their member neurons cannot be identified. In this contribution, we present approaches to actually identify the individual neurons involved in assemblies. Our results may complement other methods and also provide the opportunity for a reduction of data sets to the "relevant" neurons, thus allowing us to carry out a refined analysis of the detailed correlation structure due to reduced computation time.

Keywords: massively parallel spike trains; spike synchrony; higher-order correlation; data mining

[3].1 Introduction

Synchronized presynaptic spiking activity is known to have a higher efficacy in generating output spikes than non-coordinated spike timing (Abeles, 1982). Therefore temporal coordination of spike timing is a commonly accepted signature of neuronal assembly activity (Gerstein et al., 1989; Abeles, 1991; Singer et al., 1997; Harris, 2005). Consequently, approaches to detect assembly activity have focused on the detection of correlated spiking activity on a millisecond time resolution.

With massively parallel recordings becoming available at an accelerating rate (Csicsvari et al., 2003), chances to observe the signature of assembly activity are improving. However, currently we still lack the corresponding analysis tools (Brown et al., 2004). Most of the existing methods are based on pairwise analysis, e.g. (Aertsen et al., 1989; Kohn and Smith, 2005; Shmiel et al., 2006). Approaches to analyze correlations between more than two neurons exist, but typically work only for a small number of neurons (Abeles and Gerstein, 1988; Tetko and Villa, 2001; Dayhoff and Gerstein, 1983a; Grün et al., 2002a,b) or consider only pair correlations when analyzing the ensemble (Schneidman et al., 2006; Shlens et al., 2006; Berger et al., 2007; Eldawlatly et al., 2009) (a set of neurons is seen as an assembly if most of them are pairwise correlated).

It is usually infeasible to simply extend existing methods that identify individual spike patterns to massively parallel data due to a combinatorial explosion. Therefore we tried new approaches that evaluate the complexity distribution (Staude et al., 2007; Grün et al., 2008) or the intersection matrix (Schrader et al., 2008), which can handle massively parallel data and analyze it for higher-order spike patterns. These methods are able to detect the presence of higher-order correlation, but do not identify neurons that participate in the correlation. The goal of the present study is to resolve this issue: we want to directly identify neurons that take part in an assembly as expressed by coincident firing. Our aim is not, however, to determine the order of the correlation in which they are involved, but to provide an efficient tool to reduce the data set to the relevant neurons, which will then be examined in detail in further analysis. We present three different methods, all of which rely on the idea to detect whether an individual neuron is involved more often in any kind of coincident event than can be expected by chance.

[3].2 Generation of Test Data

In order to demonstrate our approaches we make use of four data sets of $N = 100$ parallel spike trains, which contain different types of correlations or independent processes. The data sets are generated by a stochastic model, which has its origins in Grün et al. (1999);

Kuhn et al. (2003), and Grün et al. (2008). The basic assumption is that the activation of an assembly is expressed by synchronous spiking of its member neurons. Due to the typically blind sampling from the cortical tissue, the chances to observe a number of neurons from one assembly are rather small. This enters our modeling by assuming that only a small percentage of neurons are correlated—the rest fires independently.

[3].2.1 Stochastic Model

Simultaneous spike trains are modeled as parallel, binary processes realized as stationary Bernoulli processes. The simplest form realizes fully independent processes with predefined firing rates λ_i per neuron i, thus defining the occupation probability $p_i = \lambda_i \cdot h$ per time bin of length h for each process. Such realizations model the basic activity of the N neurons. Without further insertion of correlated spiking they serve as control data sets ("rate model", Set1 in Fig. [3].1).

Assembly activity is modeled by coincident spiking activity in a subset of m out of the N neurons: a hidden "mother" process of rate α is realized, from which spikes are copied into m selected child processes with probability ϵ. ϵ may be 1, so that all m processes receive a copy of a spike of the hidden process. In this case all m neurons exhibit coincidences of order m (Set2 in Fig. [3].1). Alternatively, and presumably more realistic for experimental data, ϵ can be chosen < 1. In this case the resulting coincidences within the m neurons are on average composed of $\epsilon \cdot m < m$ spikes with a random composition of spiking neurons per event (Set3 in Fig. [3].1).

Finally, the correlated and uncorrelated spike trains are merged. The spike train of a child process is then composed of "background" firing and of spikes involved in coincidences. The total firing rate is $\lambda_i = \lambda_{i,b} + \alpha \cdot \epsilon$, where $\lambda_{i,b}$ is the background rate and $\lambda_{i,c} = \alpha \cdot \epsilon$ the coincidence rate. Trivially, the firing rates of each process can be predefined, and the background firing rate can be adjusted accordingly.

Multiple assemblies can be generated analogously by using one hidden process per assembly. The sets of neurons, to which the spikes are copied from each of these processes, may overlap or not. The total rate of neurons taking part in more than one assembly is composed of the sum of the assembly coincidence rates and the background rate: $\lambda_i = \lambda_{i,b} + \sum_j \lambda_{i,c_j}$ with assembly index j (Set4 in Fig. [3].1).

[3].3 Analysis of Assembly Membership

We explored the performance of three approaches to identify whether a neuron is part of a correlated group of neurons.

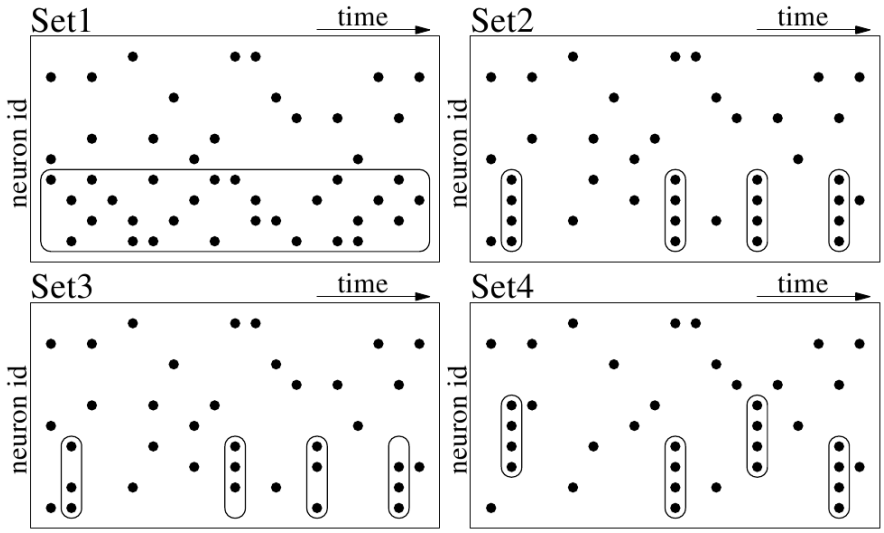

Figure [3].1: Sketches of the data sets used to test our methods. In the realizations, all of the $N = 100$ neurons in Set2, Set3, and Set4 have a total firing rate of $\lambda_i = 20$ Hz, $i \in [1 \dots N]$; coincidence firing rates are $\lambda_c = 5$ Hz. Set1: rate model of independent firing, composed of two groups of neurons of different firing rates ($\lambda_i = 50$ Hz, $i \in [1 \dots m]$, $m = 10$ (box), and $\lambda_i = 20$ Hz, $i \in [m+1 \dots N]$). Set2 and Set3 contain one assembly of $m = 10$ neurons generated with parameters $[\epsilon = 1, \alpha = 5$ Hz$]$ and $[\epsilon = 0.4, \alpha = 12.5$ Hz$]$, respectively. Set4 contains two assemblies ($\epsilon = 1$) with $m_1 = 7$ and $m_2 = 8$ neurons within $m = 10$ neurons, 5 participate in both assemblies. Background firing rates have been chosen in such a way that the total rate of $\lambda_i = 20$ Hz is preserved.

[3].3.1 Spike Shuffling

In all three approaches explored here, we compute a different test statistic from the data, each of which is based on a different basic idea. Since there are certain obstacles to finding the distributions of these test statistics under the null hypothesis that the considered neuron is not part of an assembly, we rely on a spike shuffling procedure instead. First we compute, for a neuron i under consideration, the test statistic from the data. Then we shuffle the spike times of neuron i in order to make it independent of all other (observed) neurons, recompute the test statistic, and compare the results. Spike shuffling and recomputing the test statistic is repeated sufficiently often. Therefore counting the number of times the result of a shuffling trial meets or exceeds the result obtained on the actual data and dividing this number by the total number of trials yields a p-value. This p-value measures the probability that a result as obtained on the actual data would be obtained by chance if neuron i was independent of all other neurons.
Note that we shuffle *only* the spikes of neuron i; all other spikes are left unchanged. As a consequence, the correlation structure among all other neurons is preserved—only

correlations of neuron i with other neurons are destroyed.

[3].3.2 Background Rate Estimation (BRE)

With the model outlined above, the probability that a neuron i fires in a given time bin may be decomposed into two constituents: the background firing probability, denoted by $\theta_i = \lambda_{i,b} \cdot h$ (where h is the length of a time bin), and the coincidence probability, denoted by $\xi_i = \lambda_{i,c} \cdot h$, which captures the joint influence of *all* assemblies neuron i participates in. Since we assume that the hidden processes that generate the background and the coincident spikes are independent, the effective firing probability is $\eta_i = \theta_i + \xi_i - \theta_i\xi_i$.
We want to determine whether neuron i participates in an assembly or not, that is, whether $\xi_i > 0$ (alternative hypothesis) or $\xi_i = 0$ (null hypothesis). Due to the above equation we have $\xi_i = \frac{\eta_i - \theta_i}{1-\theta_i}$. Note that η_i can easily be estimated from the data, namely as $\hat{\eta}_i = k_i/k$, where k_i is the number of time bins in which neuron i fires and k is the total number of time bins. As a consequence we can easily derive a statistical test if we can estimate the background firing probability θ_i.
To derive the estimator for θ_i, we consider a set of N independent neurons without assemblies (i.e. $\forall j; 1 \leq j \leq N : \xi_j = 0$). The probability p_0 that *no* neuron fires in a time bin is $p_0 = \prod_{j=1}^{N}(1-\theta_j)$ and the probability p_{i0} that *only* neuron i fires is $p_{i0} = \theta_i \prod_{j=1, j\neq i}^{N}(1-\theta_j)$. It follows $\frac{p_{i0}}{p_0} = \frac{\theta_i}{1-\theta_i}$ and therefore $\theta_i = \frac{p_{i0}}{p_{i0}+p_0}$. Obviously, p_0 and p_{i0} can easily be estimated, namely as $\hat{p}_0 = k_0/k$ and $\hat{p}_i = k_{i0}/k$, respectively, where k_0 is the number of time bins in which *no* neuron fires and k_{i0} the number of bins in which *only* neuron i fires.
The crucial insight is now that the probabilities p_0 and p_{i0} remain unaffected if neurons participate in assemblies, because both refer to time bins with *at most one* spike. However, coincident firing, by its very definition, means that *more than one* neuron fires in the same time bin. As a consequence, θ_i can be estimated, even in the presence of coincident firing, as $\hat{\theta}_i = k_{i0}/(k_{i0}+k_0)$. The actual test whether neuron i participates in an assembly checks whether $\hat{\xi}_i$ is sufficiently large so that the null hypothesis $\xi_i = 0$ can be rejected. A natural test statistic, estimating the fraction of coincidence spikes, is

$$t_i^{\mathrm{BRE}} = \frac{\hat{\xi}_i}{\hat{\eta}_i} = \frac{\hat{\eta}_i - \hat{\theta}_i}{\hat{\eta}_i(1-\hat{\theta}_i)}.$$

A p-value can then easily be derived by using the spike shuffling procedure described in Section [3].3.1.
As a generalization of this approach one may consider to estimate the background firing rate of neuron i not only from the time bins in which at most neuron i fires, but also

from those time bins, in which a maximum of r, $r \geq 0$, other neurons fire ($r = 0$ yields the case discussed above). This provides an indication whether neuron i participates in assemblies with more than $r + 1$ neurons. However, one should be aware that for $r > 0$ a possible participation in assemblies of smaller size (at most $r+1$ neurons) can obscure the participation in larger assemblies (more than $r + 1$ neurons), because in this case we are not estimating the true background firing probability, but the firing probability resulting from background firing and participation in small assemblies.

[3].3.3 Conditional Pattern Complexities (CPC)

The approach presented in Section [3].3.2 only considers *whether* some other neuron j, $j \neq i$, fires or not, but neglects *how many* and *which* other neurons fire and *how often* they fire together with neuron i. The approaches studied in this and the next section try to exploit such additional information.

The most natural approach is to take into account *how many* other neurons fire, on average, together with neuron i. If neuron i participates in one or more (large) assemblies, there should be several time bins in which it fires together with several other neurons. Hence the average complexity of patterns involving neuron i should be larger than expected by chance. Formally, we use $x_i = \frac{1}{k_i} \sum_{l=1}^{k} \tau(i \in I_l)(|I_l| - 1)$ and $\bar{x}_i = \frac{1}{k} \sum_{l=1}^{k} (|I_l| - \tau(i \in I_l))$, where I_l is the index set of the neurons that fire in the l-th time bin and $\tau(\varphi)$ is 1 if φ is true and 0 otherwise. Obviously, x_i is the average pattern complexity in those time bins in which neuron i fires (with spikes of neuron i removed). On the other hand, $\bar{x}_i$ is simply the overall average pattern complexity (again with spikes of neuron i removed). With this, a natural test statistic is

$$t_i^{\mathrm{CPC}} = \frac{x_i - \bar{x}_i}{\bar{x}_i}.$$

A p-value can again be derived by using the spike shuffling procedure described in Section [3].3.1.

[3].3.4 Conditional Spike Frequencies (CSF)

In a second approach we take into account *how often* other neurons fire together with neuron i. The idea is that if neuron i participates in one or more assemblies, it should fire more often together with certain other neurons (those also in the assemblies) than can be expected by chance. In order to take care of different firing rates, we use the number of excess spikes to form a test statistic: we compute for each neuron j, $j \neq i$, the difference between the number of spikes observed together with a spike of neuron i and the expected

number of such spikes, estimated as $k_i\hat{\eta}_j$. Since only excess spikes tell us about possible correlations, negative differences are ignored. Formally, the test statistic is

$$t_i^{\mathrm{CSF}} = \frac{1}{n} \sum_{j=1, j\neq i}^{n} \tau(k_{ij} > k_i\hat{\eta}_j)(k_{ij} - k_i\hat{\eta}_j),$$

where k_{ij} is the number of time bins in which both neuron i and neuron j fire and (as above) $\tau(\varphi)$ is 1 if φ is true and 0 otherwise. A p-value is derived by the same shuffling procedure described above (see Section [3].3.1).

[3].3.5 Results

Due to limitations of space, we can present detailed results only for one of the three approaches. We chose CPC (Section [3].3.3) because of two reasons: in the first place, BRE (Section [3].3.2) proved to be a fairly weak method in our experiments, at least for $r = 0$. The reason is mainly the small number of time bins in which *only* the considered neuron i fires (small k_{i0}), which leads to an unreliable estimate and thus a high variance. Performance improves for $r = 1$ and becomes reasonably good for $r = 2$, but still falls short of the performance of the other two methods. Therefore we chose to discard BRE, even though we believe that the estimator of the background firing rate used in it is a relevant insight. Secondly, even though the performance of CSF (Section [3].3.4) was actually the best in our experiments, the approach requires an estimate of the expected number of spikes for the other neurons. This can turn out to be a severe hindrance for analyzing real spike trains, for which firing rates may be non-stationarity, and made us favor the CPC approach.

CPC results for the four data sets are depicted in Fig. [3].2. Each diagram shows the value of the test statistic and a box plot indicating the distribution of the spike shuffling results for each of the 100 neurons. For Set1 the test did not produce a significant result for any of the neurons, since they are independent. For the other three data sets (Set2–Set4) it is detected that the first 10 neurons have excess coincidences, thus exactly identifying the neurons involved in assemblies. For Set2, significance is highest, which corresponds to the fact that all correlated neurons participated with probability $\epsilon = 1$ in the coincident events. Significance is slightly reduced for Set3, which is due to the decreased participation probability ($\epsilon = 0.4$) in the coincidences. For Set4 also all neurons taking part in assemblies are reliably detected. Significance is higher for the neurons participating in both assemblies as compared to the ones participating in one of the assemblies only.

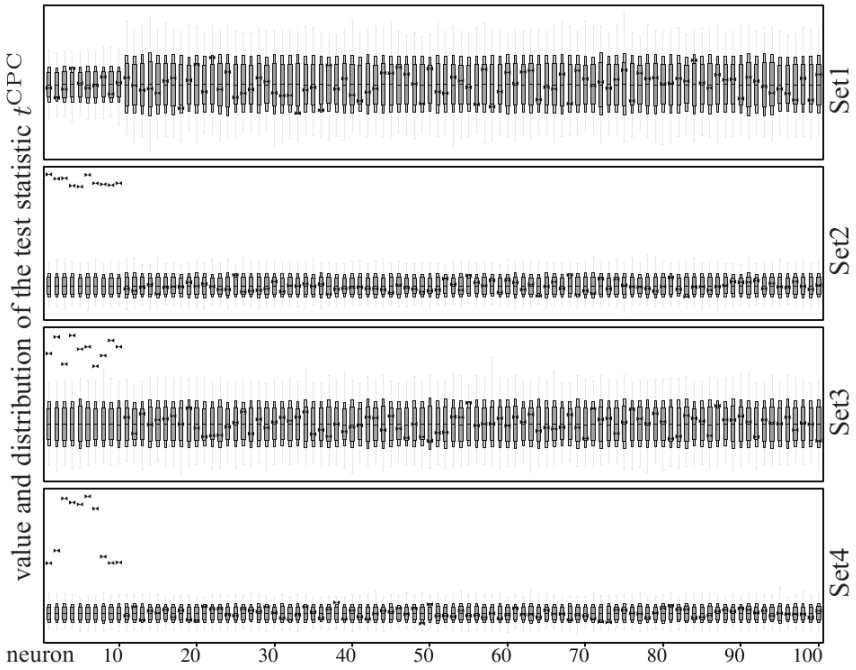

Figure [3].2: Analysis results for Set1 to Set4 using CPC. In each panel, box plots show the distribution (thick box: 5% to 95%, thin box: 1% to 99%, gray whiskers: minimum to maximum) of the shuffling results ($s = 10^5$ trials) for each neuron i (ids ordered, along the x-axis). Outside the whiskers, results are significant on a $1/s$ level. The test statistic value t_i^{CPC} obtained on the actual data is shown as a bowtie.

[3].4 Discussion

We presented three simple test statistics to identify neurons that are involved in assemblies. All of them test for a given neuron whether it is more often involved in a coincidence spike event than can be expected by chance. To do so, BRE relies on a background rate estimator (which we believe to be relevant in its own right), CPC analyses the coincidence complexities of the parallel spike trains, and CSF aggregates pairwise frequency comparisons. To assess their performance, we applied these statistics, using a spike shuffling approach, to massively parallel spike trains generated by stochastic models that allow for defining different spike correlation structures.

Our approaches are a perfect addition to methods that provide information of *presence* of higher-order correlation in data, but do not identify the individual neurons involved. Reducing data sets to the *relevant* neurons only helps to considerably reduce computation time in further analysis steps that aim at identifying the detailed higher-order correlation structures e.g. considering information geometrical measures (e.g., Shimazaki et al., 2009;

Amari, 2001) or approaches based on accretion (Gerstein et al., 1978).

Acknowledgments: Partially funded by the BMBF (grant 01GQ01413) and RIKEN Strategic Programs for R&D.

References

M. Abeles. Role of the cortical neuron: integrator or coincidence detector? *Isr. J. Med. Sci.*, 18:83–92, Jan 1982.

M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, 1st edition, 1991.

M. Abeles and G. L. Gerstein. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.*, 60:909–924, Sep 1988.

A. M. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J. Neurophysiol.*, 61:900–917, 1989.

S. Amari. Information geometry on hierarchy of probability distributions. *IEEE Trans. Info. Theory*, 47:1701–1711, 2001.

D. Berger, D. Warren, R. Normann, A. Arieli, and S. Grün. Spatially organized spike correlation in cat visual cortex. *Neurocomput*, 70:2112–2116, 2007.

E. N. Brown, R. E. Kass, and P. P. Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci*, 7(5):456–461, May 2004. doi: 10.1038/nn1228. URL `http://dx.doi.org/10.1038/nn1228`.

J. Csicsvari, D. A. Henze, B. Jamieson, K. D. Harris, A. Sirota, P. Bartho, K. D. Wise, and G. Buzsaki. Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J. Neurophysiol.*, 90:1314–1323, Aug 2003.

J. E. Dayhoff and G. L. Gerstein. Favored patterns in spike trains. II. Application. *J. Neurophysiol.*, 49:1349–1363, Jun 1983.

S. Eldawlatly, R. Jin, and K. G. Oweiss. Identifying functional connectivity in large-scale neural ensemble recordings: a multiscale data mining approach. *Neural Comput*, 21: 450–477, Feb 2009.

G. L. Gerstein, D. H. Perkel, and K. N. Subramanian. Identification of functionally related neural assemblies. *Brain Res.*, 140:43–62, Jan 1978.

G. L. Gerstein, P. Bedenbaugh, and M. H. Aertsen. Neuronal assemblies. *IEEE Trans Biomed Eng*, 36(1):4–14, Jan 1989.

S. Grün, M. Diesmann, F. Grammont, A. Riehle, and A. Aertsen. Detecting unitary events without discretization of time. *J. Neurosci. Methods*, 94:67–79, 1999.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Comput*, 14:43–80, 2002a.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput*, 14:81–119, 2002b.

S. Grün, M. Abeles, and M. Diesmann. Impact of higher-order correlations on coincidence distributions of massively parallel data. *Lecture Notes in Computer Science*, 5286:96–114, 2008.

K. D. Harris. Neural signatures of cell assembly organization. *Nat. Rev. Neurosci.*, 6: 399–407, May 2005.

A. Kohn and M. A. Smith. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.*, 25:3661–3673, Apr 2005.

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

E. Schneidman, M. J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007–1012, Apr 2006.

S. Schrader, S. Grün, M. Diesmann, and G. L. Gerstein. Detecting synfire chain activity using massively parallel spike train recording. *J. Neurophysiol.*, 100:2165–2176, Oct 2008.

H. Shimazaki, S. Amari, E. Brown, and S. Grün. State-space analysis on time-varying correlations in parallel spike sequences. *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2009.

J. Shlens, G. D. Field, J. L. Gauthier, M. I. Grivich, D. Petrusca, A. Sher, A. M. Litke, and E. J. Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.*, 26:8254–8266, Aug 2006.

T. Shmiel, R. Drori, O. Shmiel, Y. Ben-Shaul, Z. Nadasdy, M. Shemesh, M. Teicher, and M. Abeles. Temporally precise cortical firing patterns are associated with distinct action segments. *J. Neurophysiol.*, 96:2645–2652, 2006.

W. Singer, A. K. Engel, A. K. Kreiter, M. H. J. Munk, S. Neuenschwander, and P. R. Roelfsema. Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences*, 1(7):252–261, 1997.

B. Staude, S. Rotter, and S. Grün. Detecting the existence of higher-order correlations in multiple single-unit spike trains. *Soc. Neurosci. Abstr. 103.9/AAA18, Washington, DC,*, 2007.

I. V. Tetko and A. E. Villa. A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 2. Application to simultaneous single unit recordings. *J. Neurosci. Methods*, 105:15–24, Jan 2001.

Chapter 4

Efficient Identification of Assembly Neurons within Massively Parallel Spike Trains

Efficient Identification of Assembly Neurons within Massively Parallel Spike Trains

Denise Berger[1,2], Christian Borgelt[3], Abigail Morrison[4], and Sonja Grün[1,4]

[1]Bernstein Center for Computational Neuroscience, Berlin, Germany
[2]Neuroinformatics, Inst. Biology, Freie Universität, Berlin, Germany
[3]European Center for Soft Computing, Mieres (Asturias), Spain
[4]RIKEN Brain Science Institute, Wako-shi, Japan

This section will appear in Computational Intelligence in Neuroscience as

Berger, D.; Borgelt, C.; Morrison, A. & Grün, S.
Efficient Identification of Assembly Neurons within Massively Parallel Spike Trains.

Abstract

The chance of detecting assembly activity is expected to increase if the spiking activities of large numbers of neurons are recorded simultaneously. Although such massively parallel recordings are now becoming available, methods able to analyze such data for spike correlation are still rare, as a combinatorial explosion often makes it infeasible to extend methods developed for smaller data sets. By evaluating pattern complexity distributions the existence of correlated groups can be detected, but their member neurons cannot be identified. In this contribution, we present approaches to actually identify the individual neurons involved in assemblies. Our results may complement other methods and also provide a way to reduce data sets to the "relevant" neurons, thus allowing us to carry out a refined analysis of the detailed correlation structure due to reduced computation time.

Keywords: massively parallel spike trains; spike synchrony; higher-order correlation; data mining

[4].1 Introduction

Synchronized presynaptic spiking activity is known to have a higher efficacy in generating output spikes than non-coordinated spike timing (Abeles, 1982). Therefore temporal coordination of spike timing is a commonly accepted signature of neuronal assembly activity (Gerstein et al., 1989; Abeles, 1991; Singer et al., 1997; Harris, 2005). Consequently, approaches to detect assembly activity have focused on the detection of correlated spiking activity on a millisecond time resolution.

With massively parallel recordings becoming available at an accelerating rate (Csicsvari et al., 2003), the likelihood of observing the signature of assembly activity is improving. However, we currently still lack the corresponding analysis tools (Brown et al., 2004). Most of the existing methods are based on pairwise analysis, e.g., Aertsen et al. (1989); Kohn and Smith (2005), and Shmiel et al. (2006). Approaches to analyze correlations between more than two neurons exist, but typically only work for a small number of neurons (Abeles and Gerstein, 1988; Dayhoff and Gerstein, 1983a; Grün et al., 2002a,b; Tetko and Villa, 2001) or only consider pair correlations when analyzing the assembly (Schneidman et al., 2006; Shlens et al., 2006; Berger et al., 2007; Eldawlatly et al., 2009) (in these approaches a set of neurons is seen as an assembly if most of them are pairwise correlated). It is usually infeasible to simply extend existing methods that identify individual spike patterns to massively parallel data due to a combinatorial explosion. Therefore in previous studies we tried new approaches that evaluate the complexity distribution (Staude et al., 2007; Grün et al., 2008) or the intersection matrix (Schrader et al., 2008), which can handle massively parallel data in reasonable computational time and analyze it for higher-order spike patterns. These methods are able to detect the presence of higher-order correlation, but do not identify neurons that participate in the correlation. The goal of the present study is to resolve this issue: we want to directly identify neurons that take part in an assembly as expressed by coincident firing. Our aim is not, however, to determine the order of the correlation in which they are involved, but to provide an efficient tool to reduce the data set to the relevant neurons, which can then be examined in detail in further analysis. We present two different methods, both of which rely on the idea of detecting whether an individual neuron is involved in any kind of coincident event more often than can be expected by chance.

The paper is organized as follows: in Section [4].2 we discuss methods of generating surrogate data from given spike trains, which we need in order to obtain reference distributions for the test statistics that are introduced in Section [4].3. In Section [4].4 we apply our test statistics to several artificial and one real-world data set and assess their performance. Finally, in Section [4].5 we evaluate our findings and draw conclusions about the useful-

ness of our approach. This study is based on a former contribution (Grün et al., 2009), and is extended here by a systemtic study of parameter dependencies, and the analysis of simulated network data and neuronal data.

[4].2 Generation of Surrogate Data

Our methods of detecting neurons that are participating in an assembly consist of two ingredients: a test statistic (following section) and a procedure to generate surrogate data (this section), which is needed to estimate their distribution.

[4].2.1 General Procedure

In all approaches explored in this paper we compute a different test statistic from the data, each of which is based on a different basic idea (see Section [4].3). Unfortunately, there are certain obstacles that prevent us from easily finding the distributions of these test statistics under the null hypothesis that the considered neuron i is not part of an assembly. Therefore we rely on the generation of *surrogate data* from the original data set in order to estimate this distribution. The surrogate data set is created in such a way that a neuron i under consideration, if it is part of an assembly, becomes independent of all other neurons, or at least is considerably less dependent on the other neurons than in the original data set.
The general test procedure is as follows: first we compute, for the neuron i under consideration, the test statistic on the original data set. Then we generate a surrogate data set in one of the ways described below, recompute the test statistic, and compare the results. Generating surrogate data sets and recomputing the test statistic is repeated sufficiently often (unless otherwise stated, 5000 times). Finally, counting the number of times the result of a surrogate run meets or exceeds the result obtained on the actual data and dividing this number by the total number of runs yields a p-value. This p-value measures the probability that a result as obtained on the actual data would be obtained by chance if neuron i was independent of all other neurons (where what exactly is meant by "chance" is made precise by the surrogate generation method).

[4].2.2 Problems of Generating Surrogate Data

In general, the procedure with which surrogate data is generated is fairly critical for the validity of a test for assembly activity, because it implicitly represents the null hypothesis. As a consequence, a common argument against a chosen surrogate generation method is

that it represents a (much) too restricted null hypothesis. In such a case obtaining a significant test result could be explained not only by the original data actually having the property tested for, but also by some other, unrelated property, which should have been covered by the null hypothesis, but was not. In principle, the null hypothesis must cover the full complement of the property tested for (which, of course, represents the alternative hypothesis). The "worst" data model in this complement—that is, the data model that is most likely to produce, by chance, data that looks like as if it was generated by a model having the property tested for—determines the significance of a test result.

A closer analysis reveals that this problem cannot in general be fully overcome for a test for assembly activity, because the complement of assembly activity is too large. In this complement there are data models that are so close to models exhibiting assembly activity (that is, models that can, with a high probability, generate the same data), that it is almost impossible to obtain significant p-values. In order to cope with this problem, one has to restrict the null hypothesis. However, at the same time one must take care to restrict it only to such an extent that the restrictions can reasonably be considered harmless (where "harmless" means that only such models are excluded of which we are fairly certain that they are impossible).

In practice, it is this idea that underlies the desire to preserve, as much as possible, certain properties of the original data, which are not directly connected to the property tested for, but could explain a significant test result. Unfortunately, there is no consensus about which properties should or must be preserved. Examples include the overall and local firing rates of the neurons, the distribution of inter-spike intervals, the number of spikes per time bin etc. For an extensive discussion of this problem see Grün (2009).

In addition, the more properties one tries to preserve, the more complex and more expensive (in terms of computing time) generating surrogate data sets becomes. If one recalls that a large number of surrogate data sets have to be generated in order to be able to compute a sufficiently small and reliable p-value, the required computation time is a serious concern.

As a consequence, we confine ourselves here to fairly simple methods of generating surrogate data with the following argument: we see our method mainly as a preprocessing method and always assume that the findings of our methods are later substantiated with more detailed and statistically more powerful tests. Our predominant goal is to reduce the number of neurons that have to be subjected to this more detailed and thus also (again in terms of computing time) more costly analysis. Therefore we can afford to have a few false positive results, because it can reasonably be assumed that they will be detected as such in subsequent steps. However, we cannot tolerate false negatives, since then neurons that are actually relevant would be excluded from further analysis. Relying on a too restricted

null hypothesis, implemented as an over-simplified surrogate generation method that does not preserve all relevant properties of the original data, may produce additional false positives but will not produce false negatives. For the purposes of preprocessing, we therefore prioritize speed over accuracy and use simple surrogate data generation methods.

[4].2.3 Spike Shuffling in Time

One of the simplest and fastest ways of generating a surrogate data set in which a considered neuron i is independent of all other neurons, is to shuffle the spike times of neuron i. That is, the spikes of neuron i are reassigned from their original time bins to time bins that are randomly chosen from the available ones. Note that we shuffle *only* the spikes of neuron i; all other spikes are left unchanged. As a consequence, the correlation structure between all other neurons is preserved—only correlations of neuron i with other neurons are destroyed.

Obvious advantages of such a shuffling scheme are that it preserves the number of spikes of the considered neuron, changes the number of spikes per time bin by at most 1 (only a spike of the considered neuron can be removed or added), is easy to program, and fast to execute. However, it has the serious disadvantage that in real world data firing rates are not constant over time (for an extensive discussion of this issue see Grün, 2009).

If we assume that the neurons under consideration change their firing rates coherently (the firing rate rises or falls for all neurons in parallel), there is a fairly simple, but nevertheless surprisingly effective approach to cope with varying spike rates: we simply see the number of spikes in a bin as a (very coarse) indicator of the general neuronal activity level. We may then derive an estimate for the probability of a spike occurring in a time bin j as follows: let I_l be the index set of the neurons that fire in the l-th time bin, and let k be the total number of time bins. Then we may define the probability of randomly assigning a spike of some neuron i to a given time bin j as

$$\pi_j = \frac{|I_j| + c}{\sum_{l=1}^{k}(|I_l| + c)},$$

where c is a correction term that in similar contexts is known as the *Laplace correction*. Formally, this correction term can be justified with arguments from Bayesian statistics (Bolstad, 2007), where it is derived from an uninformative prior distribution. This prior distribution is a uniform distribution here, because formally we consider a polynomial distribution over the time bins. This prior distribution is then modified with the data distribution using Bayes' rule. As a consequence, the Laplace correction c introduces a tendency towards a uniform distribution and may be chosen in $[0, +\infty)$. For $c = 0$

the number of spikes in a bin relative to the total number of spikes directly determines the probability of assigning a spike to this time bin. Formally, this yields a maximum likelihood estimation of the time bin probability, because no weight is assigned to the uniform prior distribution; one relies entirely on the available data. For $c \to \infty$ we obtain, as a limiting case, the uniform spike shuffling scheme described above: each time bin has the same probability of having a spike assigned to it, regardless of the number of spikes it contains. Values $0 < c < \infty$ introduce a limited tendency towards a uniform distribution, which becomes stronger for greater values of c.

[4].2.4 Trial Shuffling

In real recordings of parallel spike trains it is often the case that some kind of stimulus is presented to the subject and then the neuronal response to this stimulus is recorded for a certain period. In order to reduce the effects of random influences that may be present in a single instance, several *trials* of presenting the stimulus and recording the response are carried out.

If multiple trials are available, one may use *trial shuffling* to generate surrogate data. In trial shuffling the spike trains of a neuron i under consideration are randomly assigned to other trials, thereby replacing the original spike trains in the respective trials. Even though all trials must, of course, be aligned w.r.t. the onset of the stimulus, the individual spike trains vary sufficiently due to their stochastic nature, such that neuron i becomes independent of the other neurons in terms of the precise timing of spikes across the neurons. At least it is plausible to assume that existing dependences will be considerably reduced due to the independence of the trials.

However, one should be aware that because of similar responses due to the same stimulus, trial shuffling can hide assembly activity that is actually present, despite the independence of the trials. In particular this is the case if the assembly response is sufficiently exact in time (relative to the stimulus onset), so that trial shuffling does not destroy the dependence between the neurons. As a consequence, trial shuffling can lead to false negatives, which is a critical disadvantage for our approach (see Section [4].2.2).

In addition, plain trial shuffling has the disadvantage that the number of available trials limits the achievable p-values: suppose there are s trials. Then one is the original data and the spike train of the neuron i under consideration can be moved to at most $s - 1$ other trials in order to obtain surrogate data. Even if in all of these $s - 1$ cases the obtained statistic value is less than that obtained on the original data, all we can say about the p-value is that it is less than $1/(s-1)$. Since s usually ranges only in the order of a few dozen, it may not even be possible to obtain a significant test result. The fact

that we can use any trial as the original data and all other trials to generate surrogates does not amend this drawback, because this only provides us with s independent p-values (cf. Pipa and Grün, 2003).
In order to improve this situation, we concatenate the individual trials into one long data set and compute the test statistic once on the entire set. Then the trial segments of the spike train of the neuron i under consideration are shuffled to obtain a surrogate data set. As a consequence, the number of surrogate data sets that can be generated is considerably increased, because any permutation of the trial segments of the spike train (except the identity, that is, the case in which all trial segments are at their original positions) yields a surrogate data set.
However, one should be aware that this scheme is also open to criticism: suppose that two different permutations move a given trial segment into the same location. Then the terms of the test statistic that result from the time bins in the corresponding trial are necessarily the same for both permutations. This may have the effect that the obtained distribution of the test statistic exhibits less variation than the actual distribution. If avoiding this effect is made a strict requirement, only $t - 1$ permutations (other than the identity) can be defined. In this case there is no advantage over separate trials.
Nevertheless we believe that such a trial shuffling scheme is reasonable. In the first place, the fraction of terms of the test statistic that are identical for two permutations will be relatively small unless the total number t of trials is very small. Hence the effect on the variability of the test statistic is likely negligible. Secondly, reducing the variability of the test statistic for surrogate data has at most the effect of creating false positive results. However, as already argued in Section [4].2.2, we can afford a few false positive results, because later processing steps should remove them.

[4].3 Test Statistics

We explored the performance of two approaches to identify whether a neuron is part of a correlated group of neurons.

[4].3.1 Conditional Pattern Complexities (CPC)

The first approach to identify neurons participating in assemblies is based on the idea that such neurons should have, on average, more neurons firing together with them in the original data than in the surrogates. In other words, if some neuron i participates in one or more (large) assemblies, there should be several time bins in which it fires together with a larger number of other neurons. Hence the average complexity of patterns involving

neuron i should be larger than can be expected by chance. Formally, we use

$$\mu(i) = \frac{1}{T_i} \sum_{l=1}^{T} \tau(i \in I_l)(|I_l| - 1) \tag{4.1}$$

and

$$\bar{\mu}(i) = \frac{1}{T} \sum_{l=1}^{T} (|I_l| - \tau(i \in I_l)),$$

where I_l is the index set of the neurons that fire in the l-th time bin, T is the total number of time bins, T_i the number of time bins in which neuron i fires, and $\tau(\varphi)$ is 1 if φ is true and 0 otherwise. Thus, μ_i is the average pattern complexity in those time bins in which neuron i fires (with spikes of neuron i removed), while $\bar{\mu}_i$ is simply the overall average pattern complexity (again with spikes of neuron i removed). A natural test statistic is

$$t^{\mathrm{CPC}}(i) = \frac{\mu(i) - \bar{\mu}(i)}{\bar{\mu}(i)}. \tag{4.2}$$

A p-value can be derived by using the surrogate data generation procedures described in Section [4].2.

An obvious way to improve this measure is to weight large complexities more strongly than smaller ones, because large complexities are, intuitively, more indicative of assembly activity. A simple technical means to achieve such weighting is to raise the complexities to a user-specified power α:

$$\mu_\alpha(i) = \frac{1}{T_i} \sum_{l=1}^{T} \tau(i \in I_l)(|I_l| - 1)^\alpha$$

and

$$\bar{\mu}_\alpha(i) = \frac{1}{T} \sum_{l=1}^{T} (|I_l| - \tau(i \in I_l))^\alpha.$$

In other words, instead of a simple mean of the pattern complexities, we employ higher moments. The resulting test statistic is

$$t_\alpha^{\mathrm{CPC}}(i) = \frac{\mu_\alpha(i) - \bar{\mu}_\alpha(i)}{\bar{\mu}_\alpha(i)}. \tag{4.3}$$

[4].3.2 Conditional Spike Frequencies (CSF)

In a second approach we take into account how often other individual neurons fire together with neuron i. The idea is that if neuron i participates in one or more assemblies, it should

fire more often together with certain other neurons (those also in the assemblies) than can be expected by chance. In order to be insensitive to differing firing rates, we use the number of excess spikes to form a test statistic: we compute for each neuron j, $j \neq i$, the difference between the number of spikes observed together with a spike of neuron i and the expected number of such spikes, estimated as $T_i\hat{\eta}_j$ with $\hat{\eta}_j = T_j/T$, where T_i and T_j are the number of time bins in which neurons i and j, respectively, fire and T is the total number of time bins. Since only excess spikes tell us about possible correlations, negative differences are ignored. Formally, the test statistic is

$$t^{\mathrm{CSF}}(i) = \frac{1}{N-1} \sum_{j=1, j\neq i}^{N} \tau(T_{ij} > T_i\hat{\eta}_j)(T_{ij} - T_i\hat{\eta}_j), \tag{4.4}$$

where T_{ij} is the number of time bins in which both neuron i and neuron j fire and (as above) $\tau(\varphi)$ is 1 if φ is true and 0 otherwise. As for the conditional pattern complexity (see above), we may consider weighting a large number of excess spikes more strongly than a small number, as a large number is certainly more indicative of assembly activity. In order to achieve this, we once again introduce a user-specified power α to which the number of excess spikes is raised:

$$t_\alpha^{\mathrm{CSF}}(i) = \frac{1}{N-1} \sum_{j=1, j\neq i}^{N} \tau(T_{ij} > T_i\hat{\eta}_j)(T_{ij} - T_i\hat{\eta}_j)^\alpha. \tag{4.5}$$

As another straightforward variant, one may relate the excess spikes to the expected number of spikes, because a higher firing rate also leads to a higher variation in the actual number of observed spikes and thus makes a large number of excess spikes more likely. In this case, the measure reads (with the optional power α already added):

$$t_\alpha^{\mathrm{CSF}}(i) = \frac{1}{N-1} \sum_{j=1, j\neq i}^{N} \tau(T_{ij} > T_i\hat{\eta}_j)\left(\frac{(T_{ij} - T_i\hat{\eta}_j)}{T_i\hat{\eta}_j}\right)^\alpha.$$

However, it is not immediately clear whether this modification improves or deteriorates the sensitivity of this measure. In this paper we confine our study to the unmodified version.

A p-value is derived by the same surrogate generation procedures described above (see Section [4].2).

[4].4 Results for Different Test Data Sets

We tested our statistics in three scenarios: data that was generated using the stochastic model (Section [4].4.1) that specifies the coordinated activity we desire to detect (see Section [4].4.2), data that was generated by simulating a large network of neurons into which a synfire chain was embedded (see Section [4].4.3) and real spike train recordings from cat visual cortex (see Section [4].4.5).

[4].4.1 Stochastic Model of Assembly Activity

The model of coordinated neural activity we adopt here has its origins in Grün et al. (1999); Kuhn et al. (2003), and Grün et al. (2008). The basic assumption is that the activation of an assembly is expressed by synchronous spiking of its member neurons. Due to the typically blind sampling from the cortical tissue, the chance of observing a number of neurons from one assembly is rather small. This enters our modeling by assuming that only a small percentage of neurons are correlated—the rest fire independently.
In our correlation model simultaneous spike trains are modeled as parallel, binary processes realized as stationary Bernoulli processes. The simplest form realizes fully independent processes with predefined firing rates λ_i per neuron i, thus defining the occupation probability $p_i = \lambda_i \cdot h$ per time bin of length h for each process. Such realizations model the basic activity of the N neurons. Without further insertion of correlated spiking such independent spike trains serve as the control setting ("rate model", labeled "Indep" in Fig. [4].1).

Assembly activity is modeled by coincident spiking activity in a subset of M out of the N neurons: a hidden "mother" process of rate λ_c is realized, from which spikes are copied into the M selected child processes with probability ϵ. If the probability is 1 all M processes receive a copy of each spike of the hidden process. In this case all M neurons exhibit coincidences of order M (labeled "SIP" in Fig. [4].1). A process of this type is also called a *single interaction process* (SIP) (Kuhn et al., 2003).
Alternatively, and presumably more realistically for experimental data, the copy probability can be $\epsilon < 1$. In this case the resulting coincidences within the M neurons are composed of $\epsilon \cdot M < M$ spikes on average, with a random composition of spiking neurons per event (labeled "MIP" in Fig. [4].1). A process of this type is also called a *multiple interaction process* (MIP) (Kuhn et al., 2003).
Finally, the correlated and uncorrelated spike trains are merged. The spike train of a child process is then composed of "background" firing and of spikes involved in coincidences.

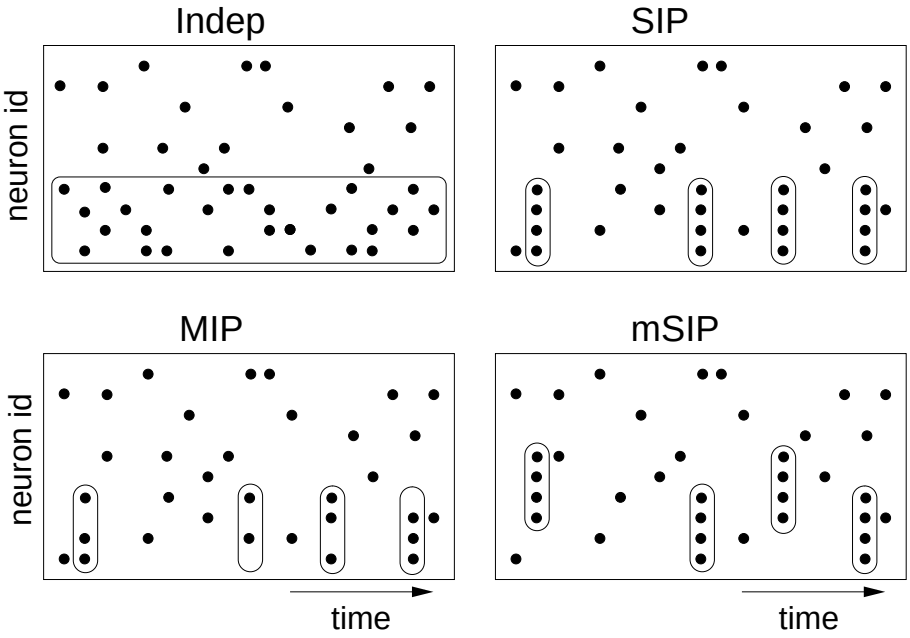

Figure [4].1: Sketches of different settings of our stochastic model. Indep: independent spike trains, but differing firing rates; SIP: single assembly, single interaction process; MIP: single assembly, multiple interaction process; mSIP: multiple assemblies, single interaction processes.

The total firing rate is $\lambda_i = \lambda_{b,i} + \lambda_{c,i}$, where $\lambda_{b,i}$ is the background rate and

$$\lambda_{c,i} = \lambda_c \cdot \epsilon \tag{4.6}$$

the coincidence rate. Trivially, the firing rates of each process can be predefined, and the background firing rate can be adjusted accordingly.

Multiple assemblies of neurons can be generated analogously by using one hidden process per assembly. The sets of neurons to which the spikes are copied from each of these processes may overlap or not. The total rate of neurons that take part in more than one assembly is composed of the sum of the assembly coincidence rates and the background rate: $\lambda_i = \lambda_{b,i} + \sum_{j=1}^{m} \lambda_{c,ij}$ with assembly index j, $1 \leq j \leq m$, where m is the number of assemblies (labeled "mSIP" in Fig. [4].1). Note that multiple assemblies can be modeled—just like single assemblies—with a copy probability $\epsilon = 1$ or with $\epsilon < 1$. In the latter case, each assembly may even use a different copy probability ϵ_j, so that we have $\lambda_{c,ij} = \lambda_{c,j} \cdot \epsilon_j$, where $\lambda_{c,j}$ is the firing rate of the j-th assembly and ϵ_j the corresponding copy probability.

These stochastic models (with no groups, one group, and several groups of correlated neurons) specify what we understand by the assembly activity we desire to detect. Of course, these models are not exhaustive in capturing what may be understood by "assembly

activity", but merely define the focus of our study.

[4].4.2 Results on Stochastic Model Data

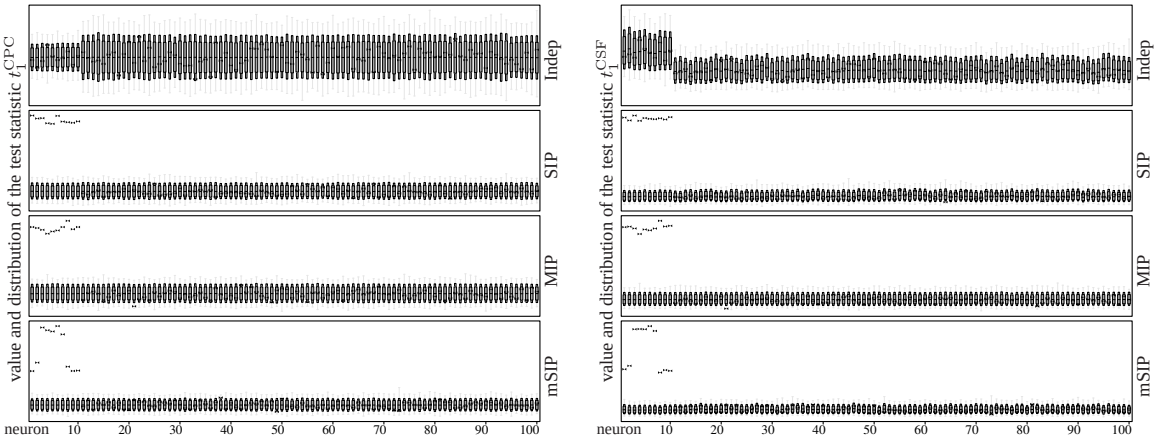

Figure [4].2: Analysis results for data sets representing the settings Indep, SIP, MIP, and mSIP using t_1^{CPC} (left) and t_1^{CSF} (right). All sets are composed of $N = 100$ neurons. In Indep, neurons 1–10 ($M = 10$) have a higher firing rate ($\lambda = 50$ Hz) than the rest of the neurons ($\lambda = 20$ Hz). In SIP, MIP and mSIP the first $M = 10$ neurons are involved in synchronous events, the rest are independent. All neurons have the same firing rate, $\lambda = 20$ Hz. The coincidence firing rate of the assemblies is $\lambda_c = 5$ Hz. The synchronous events in MIP are generated with a copy probability of $\epsilon = 0.8$. In each panel, box plots show the distribution (thick box: 5% to 95%, thin box: 1% to 99%, gray whiskers: minimum to maximum) of the shuffling results ($s = 10^5$ runs) for each neuron i (ids ordered along the x-axis). Beyond the extent of the whiskers, results are significant on a $1/s$ level. The test statistic values t_1^{CPC} and t_1^{CSF} obtained on the actual data are shown as bowties.

[4].4.2.1 Test Statistic Per Neuron

In a first experiment to demonstrate our approaches we make use of data sets generated by the stochastic model (Section [4].4.1) In these data sets we have full control over the statistics of the spike trains (stationary Poisson) and the correlation structure between the neurons. Thus, the analyses of these data serve as a proof of principle.

Figure [4].2 shows results from the four data sets analyzed using the test statistics t^{CPC} (left) and t^{CSF} (right). The data sets are composed of $N = 100$ parallel spike trains of which the first $M = 10$ neurons either have a different firing rate then the rest of the neurons ("Indep") or are correlated via synchronous spike events ("SIP", "MIP", "mSIP"). Each diagram shows the value of the test statistic (bowtie) and a box plot

indicating the distribution of the surrogate data results (spike shuffling in time) for each of the 100 neurons. The spike shuffling procedure was chosen as uniform, because the data are generated as stationary Poisson processes.
For "Indep" the tests did not produce a significant result for any of the neurons, since they are independent. For the other three data sets ("SIP", "MIP", "mSIP") the tests detected that the first 10 neurons have excess coincidences, thus exactly identifying the neurons involved in assemblies. Significance is highest for "SIP", due to the fact that all correlated neurons participated with probability $\epsilon = 1$ in the coincident events. Significance is slightly reduced for "MIP", but all neurons participating in the assembly are still reliably detected.
For the mSIP setting all neurons taking part in assemblies are also reliably detected. Significance is higher for the neurons participating in both assemblies as compared to the ones participating in only one of the assemblies.

[4].4.2.2 False Positive and False Negative Rate as a Function of Assembly Size

In the following we investigate how assembly size affects the quality of detection in two of the four settings identified in Section [4].4.1, namely SIP and MIP. In addition, we make use of the additional parameter α introduced in Section [4].3 and explore $\alpha = 1$ (default, standard form of the measures) and $\alpha = 3$ for both test statistics t^{CPC} and t^{CSF}, and thus consider a total of four measures.
Figure [4].3 depicts the false positive (FP) and false negative (FN) rate for the four test statistics (indicated by different colors) as a function of assembly size M and total number of recorded neurons N. The FN rate is derived as the percentage of neurons that participate in assembly activity but are not detected (with a significance level of 1%; positive y-axis). For a low coincidence rate of $\lambda_c = 1$ Hz and a small number of neurons involved in the assembly ($M = 5$ for $N = 100$ for SIP and MIP, see Figure [4].3a,b) the percentage of FN is very high, in particular for MIP (close to 100%). Interestingly, for larger numbers of observed neurons ($N = 1000$, Fig. [4].3c,d) even for a smaller percentage of neurons involved in assemblies ($M = 10$, i.e., $M/N = 10/1000 = 1/100$ as compared to 5/100 in Fig. [4].3a,b) the percentage of FN is lower. If we consider the case that 5% of neurons are involved in assemblies (i.e., $M = 5$ for $N = 100$ and $M = 50$ for $N = 1000$), significant FNs can be observed for the smaller neuron set, whereas no FNs occur for the larger set. For the SIP cases, increasing the assembly size leads to a fast drop of the FNs, for MIP the decrease is more gradual. This is true for all four test statistics. The percentage of false positives (Fig. [4].3, negative y-axis), i.e., neurons wrongly detected as

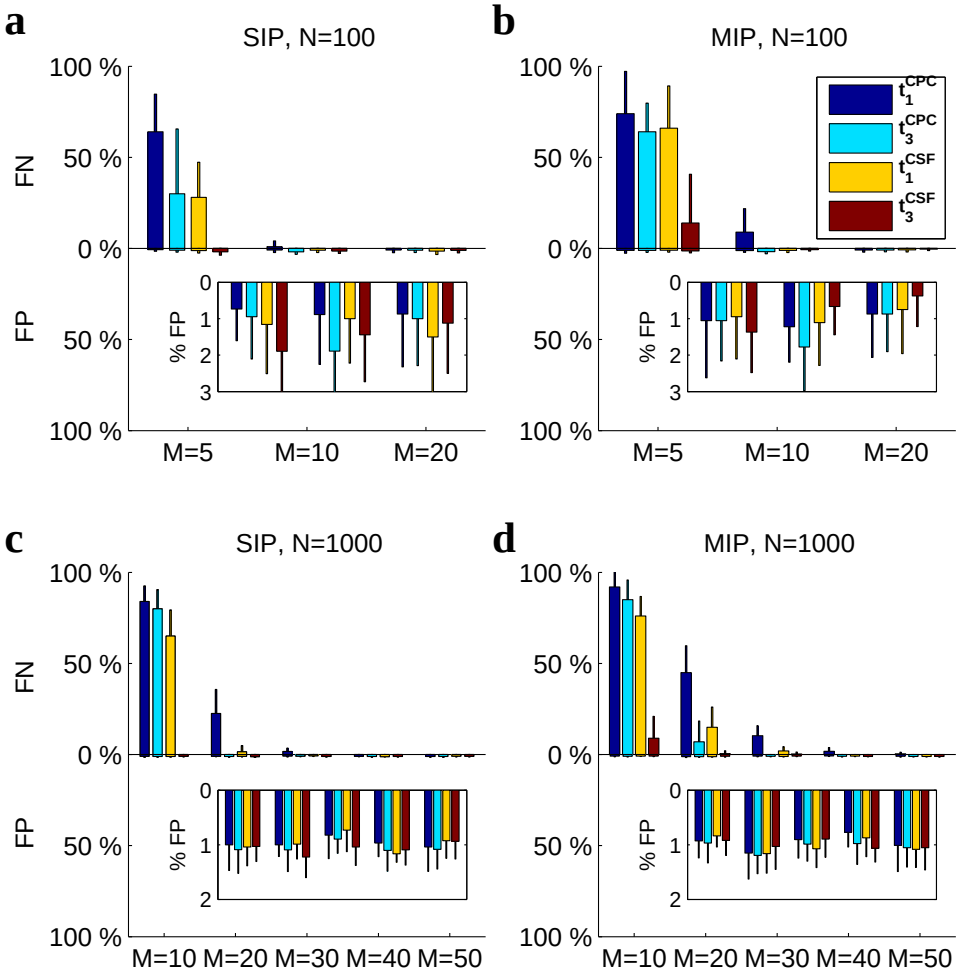

Figure [4].3: False positive and false negative rate as a function of assembly size M. Each panel shows the FN rate as bar plots (positive y-axis) and FP rate (negative y-axis, and inset) for each type of data set (a): SIP, $N = 100$; (b) MIP, $N = 100$, $\epsilon = 0.8$; (c) SIP, $N = 1000$; (d) MIP, $N = 1000$, $\epsilon = 0.8$. Each panel contains a sets of grouped bars (colors indicate the results for the different test statistics t_1^{CPC}, t_3^{CPC}, t_1^{CPC}, t_3^{CSF}; see legend) for each M. Each bar represents the mean of 10 realizations, and the error bar indicates one standard deviation. All data sets have stationary firing rates (all neurons $\lambda = 20$ Hz), a stationary coincidence rate of $\lambda_c = 1$ Hz, and a duration of $k = 10,000$ time steps of $h = 1$ ms. Significance level was 1%. For generation of the surrogates (5000 repetitions) spikes are shuffled with uniform distribution in time.

assembly members, are very low in all cases, i.e., at about the level expected by the chosen significance level of 1%. There is no obvious difference in the number of FPs generated between the different data sets or test statistics.

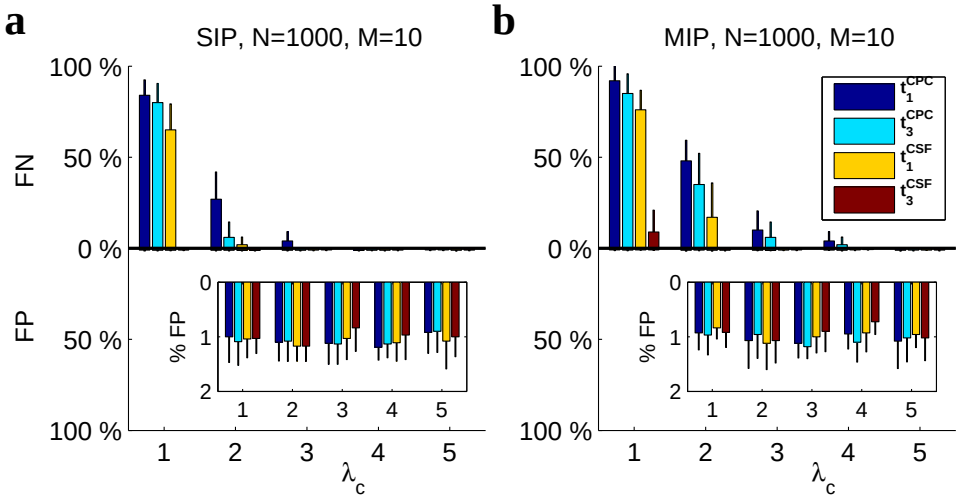

Figure [4].4: FP rate and FN rate as a function of the coincidence rate $\lambda_c = 1, 2, 3, 4, 5$ Hz shown for a) SIP and b) MIP with $N = 1000$ neurons and $M = 10$ assembly members. Display style and other parameters as in Figure [4].3.

[4].4.2.3 Impact of Coincidence Rate on False Positives and False Negatives

In the previous paragraph we showed that the smaller the assembly size, the harder it is to detect all of its members. Here we demonstrate the dependence of the FP and FN rate as a function of the coincidence rate λ_c for a small assembly size ($M = 10$) to discuss a worst case scenario (see Figure [4].4). Here different test statistics lead to quite different results. t_3^{CSF} achieves very good detection (low FN rate, FP $\approx 1\%$) for all coincidence rate levels for both SIP and MIP data. In contrast, the other test statistics result in a high FN rate (over 50%) for the SIP data with a small coincidence rate, but this drops to 0% as the coincidence rate increases from 1 to 2 Hz. Thus the presence of at least 2 Hz $\cdot 10,000$ ms $= 20$ SIP events leads to a perfect detection (Figure [4].4a). For $N = 100$ (not shown here) we found comparable results for increasing coincidence rate. In the case of MIP, the detection of assembly members becomes gradually better with increasing coincidence rate, and all assembly neurons are detected for $\lambda_c \geq 4$ Hz (Fig. [4].4b). t^{CSF} has the fastest decay of FNs, whereas t^{CPC} performs worst.

The excellent performance of t_3^{CSF} for MIP can be explained by the following example. For a coincidence rate of $\lambda_c = 1$ Hz, approximately 10 MIP events are in the data. Due to the copy probability of $\epsilon = 0.8$, on average only $M \cdot \epsilon = 8$ neurons of the assembly participate in a synchronized event. Within these 10 occurring synchronous events a specific pair of neurons has on average $0.8^2 \cdot 10 = 6.4$ coincident spike events. A pair will have in total approximately 10.1 coincidences composed of background $((\lambda - \lambda_c) \cdot h)^2 \cdot T = 3.61$ and injected coincidences (6.4). This exceeds by approx. 6.1 the expected number of

$(\lambda \cdot h)^2 \cdot T = 4$. Since t_α^{CSF} sums over all pairs neuron i has with other neurons, all $M-1$ pairs containing the neurons of the assembly will contribute with this difference.
This can be expressed in more formal way by considering expected values. t^{CSF} takes the sum of the pairwise correlations between neuron i and all other neurons. For coincidences of neuron $i \in \{1,\ldots,M\}$ with neurons also $\in \{1,\ldots,M\}$, the contributions exceed the expected value and $M-1$ terms add to the sum. Inserted coincidences occur at $T_c = T\lambda_c$ times with probability ϵ^2. The rest of the spikes that occur outside of injected events contribute to chance coincidences with their background rate $\lambda - \lambda_c$. The summands of neuron i with the other $N-M$ neurons contribute according to the variance of the independent pairings compared to the expected value p^2T, each with a term n_σ. This is approximated by

$$\begin{aligned} &\frac{1}{N-1}(M-1)(\epsilon^2 T_c + (T-T_c)(h(\lambda-\lambda_c))^2) \\ +&\frac{1}{N-1}(N-M)\cdot n_\sigma. \end{aligned} \tag{4.7}$$

In contrast, if we consider a neuron that is $\notin \{1,\ldots,M\}$, all $N-1$ terms in Eq. 4.7 contain only terms according to the variance of the realizations, i.e., approx.$\frac{1}{N-1}(N-1)\cdot n_\sigma$.
Clearly, for a fixed number of neurons N, the larger the assembly size M the more terms contribute with large coincidence counts, resulting in a larger significance and fewer FNs. For a fixed relation of M/N, the contribution of the excess coincidences (first term) increases as N increases, while the noise term (second term) decreases. Thus, with more neurons observed but the same relative amount of neurons in the assembly (here for $M=5$ in $N=100$ cmp. to $M=50$ in $N=1000$), detection improves.
The coincidence rate λ_c enters as a factor in the term of the excess coincidences, and contributes only to the product of the firing probabilities and thus only little to reducing the chance coincidences originating from background spikes. The noise term is not affected by the injected coincidences. Thus increasing the coincidence rate has a strong effect on the resulting significance, and thereby indirectly reduces the FN rate.
The difference of SIP and MIP is due to the value of ϵ^2. MIP is defined as $\epsilon < 1$, and therefore has a lower contribution to the excess coincidences than SIP, which explains the reduced sensitivity for MIP as compared to SIP processes.
In contrast, t^{CPC} evaluates the number of neurons that fire coincidently with neuron i. We now present an analytical approximation of t^{CPC} considering expected values in the context of the stochastic model by following the line of arguments introduced in Grün et al. (2008). Let us first calculate the average complexity of a neuron that is part of the M neurons receiving injected events. The spikes of that neuron can be divided into spikes

happening at times of occurrences of injected events and at times between the injected events. At injection times the total complexity is composed of the injected events of complexity $\epsilon M - 1$ (the subtraction of 1 is not to count the spikes of neuron i itself) plus the complexity due to independent firing of the $N - M$ independent neurons. Their complexity follows a Binomial distribution (Grün et al., 2008) with probability $p = \lambda h$ and complexities up to $N - M$ and yields on average a complexity of $(\lambda \cdot h) \cdot (N - M)$. These combined events (first term in Eq. 4.1) occur $T_c = \lambda_c T$ times

$$((\epsilon M - 1) + \lambda h(N - M)) \cdot \lambda_c h T. \tag{4.8}$$

At spike times when there was no injected event, coincidence patterns occur by chance. Their complexity is composed of chance coincidences of the $N-M$ neurons with occurrence probability λh as above, and chance coincidences of the M neurons with probability $(\lambda - \lambda_c)h$. The latter component has a complexity $(\lambda - \lambda_c)hM$ (mean of Bernoulli distribution with probability $(\lambda - \lambda_c)h$ in M processes).
In Eq. 4.2 for t^{CPC} we compare the actual values to the values derived from shuffling the spikes of neuron i. This situation is almost the same as considering an independent neuron of the $N - M$ neurons that did not receive injected events. In the following, we assume the case of $i \in \{1, \ldots, M\}$ with shuffled spike times. The shuffled spikes meet the chance coincidences of the $N - M$ neurons and the by chance injected events of order $M\epsilon - 1$ with a probability $\lambda h \frac{\lambda_c}{\epsilon} h$. Thus we get as the sum of these complexities

$$((M\epsilon - 1) + \lambda h(N - M)) \cdot \lambda h \frac{\lambda_c}{\epsilon} h T. \tag{4.9}$$

The rest of the spikes of neuron i that did not fall on an injected event occur with probability $(\lambda - \frac{\lambda_c}{\epsilon})$. At these times chance coincidences may occur with the background spikes of neurons $\in \{1, \ldots, M\}$ and with spikes of the other $N-M-1$. Their complexity and occurrence times are very similar to the case with $i \in \{1, \ldots, M\}$.
Thus the difference between the complexity conditional on spikes of neuron $i \in \{1, \ldots, M\}$ (Eq. 4.8) and the complexity conditional on the shuffled spike times of neuron i (Eq. 4.9) is the number of time points at which they occur. This clearly demonstrates the dependence of the coincidence rate. For increasing N (cmp. Fig. [4].3b and d) the complexity due to the $N - M$ background neurons increases with N for fixed M, but due to the summing across their number of occurrences it still improves the test statistic.
All these arguments also hold for the test statistics for $\alpha > 1$, however, due to the exponential weighting the discussed contributions of the excess coincidences are even more emphasized.

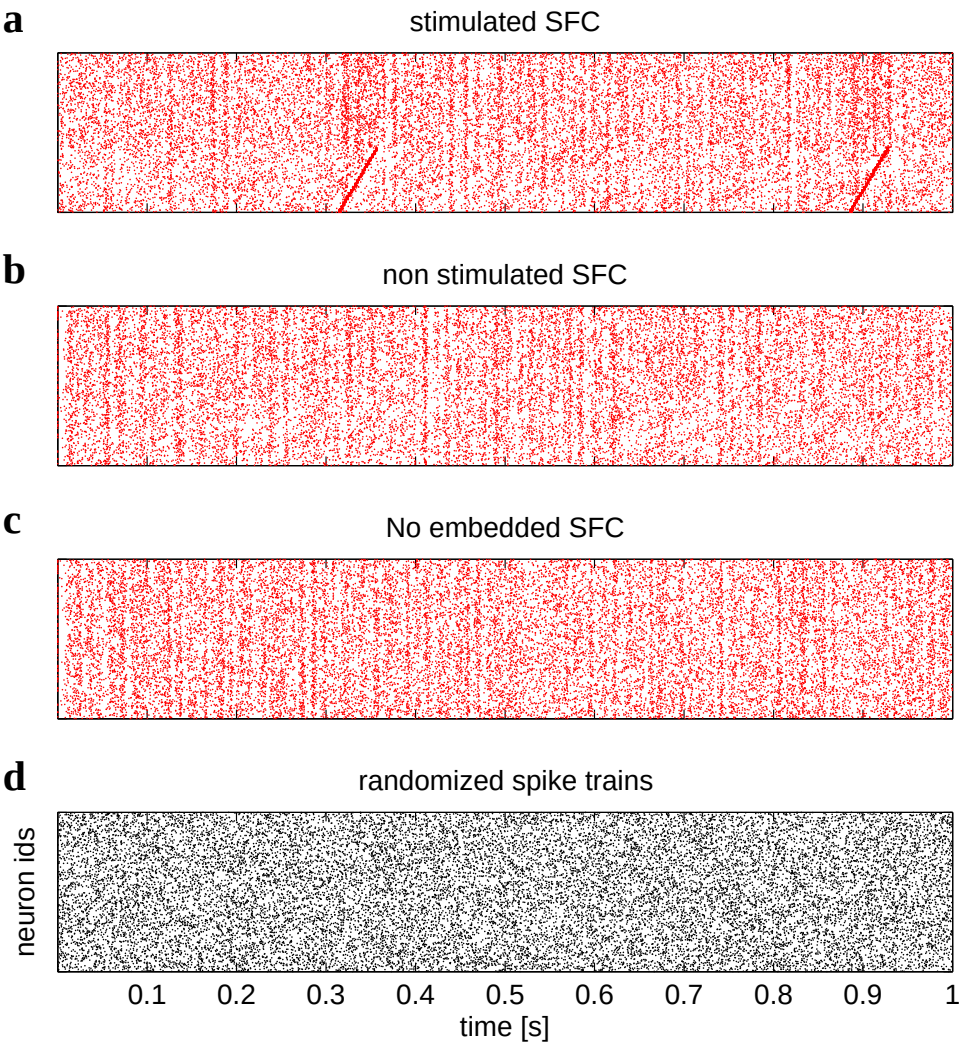

Figure [4].5: Dot displays of data from network simulations. a) Network with embedded synfire chain (5000 neurons are displayed) the neurons in the lower part of the display (here 2000) are part of the synfire chain. Synfire runs can be observed after stimulation at times 0.31s and 0.89 s. b) Same connectivity in the network as in a) but without stimulation of the synfire chain. c) Simulation of a network with purely random connectivity (no embedded synfire chain). d) Visualization of spike time randomization per neuron for the data shown in a). The displays show time segments of 1 sec duration.

[4].4.3 Data Generated from a Synfire Chain Model

The second experiment to illustrate our technique is based on spike data from a simulated neural network containing an embedded synfire chain (Abeles, 1991). The network is based on the balanced random network proposed in Brunel (2000) and contains $64,000$ excitatory and $16,000$ inhibitory neurons firing at 4.2 Hz. The neuron model implements current-based integrate-and-fire dynamics with postsynaptic currents (PSCs) represented

Parameter	**Value**
membrane time constant	**10 ms**
membrane capacitance	**250 pF**
threshold	**20 mV**
refractory period	**0.5 ms**
rise time of PSC	**0.3258 ms**

Table [4].1: Neuron model parameters

as α-functions, for a detailed description see Morrison et al. (2007). The model parameters are given in Table [4].1. Each neuron receives 6000 synaptic inputs. For the inhibitory neurons and excitatory neurons that are not part of the synfire chain, 3840 inputs are drawn randomly from the local excitatory population, 960 from the local inhibitory population, and the rest are considered to be connections from remote excitatory neurons, represented as independent Poisson processes firing at 14 Hz. The peak value of recurrent excitatory and inhibitory PSCs are 38.5 pA and -231 pA respectively.
The embedded chain consists of 20 pools of 100 excitatory neurons. Each neuron in a pool receives synaptic input from each neuron in the previous pool. To ensure robust propagation of synfire activity, the peak strength and synaptic delay of a feed-forward connection are set to 61 pA and 1.5 ms respectively. Each chain neuron also receives 3740 inputs drawn randomly from the local excitatory population with a reduced peak strength of 37.56 pA, and inhibitory and remote connections as described above. Except for the feed-forward connections, synaptic delays are given by $k \cdot h$, where h is the computational step size (0.1 ms) and k is drawn from a uniform distribution between 1 and 29.
To activate synfire chains the first pool in the chain is stimulated with a large synchronous pulse at irregular intervals with a Poisson rate of 1 Hz (SFC). To provide a control case, the embedded synfire chain is not specifically stimulated (SFCu). A further control case is provided by a network that does not contain a synfire chain, i.e., all neurons have the same input statistics (NoSFC). The respective synfire data sets are recordings of 100 s duration of N=5000 neurons from these networks. k of these are randomly sampled from the 2000 synfire chain neurons, the other N-k excitatory neurons are randomly selected from the rest of the network. The sampling degree k was chosen as 10%, 25% and 100%. All simulations were performed in NEST (Gewaltig and Diesmann, 2007); the simulation scripts are available from the authors on request.

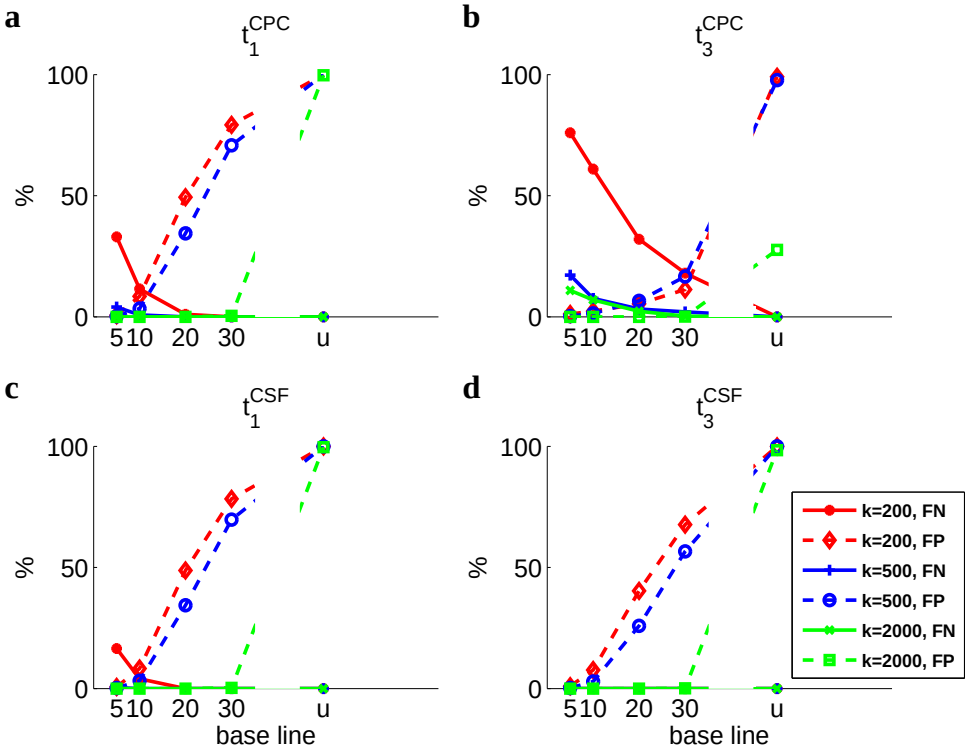

Figure [4].6: FN and FP for data of a SFC network with stimulated SFC. All panels show FP (dashed lines) and FN rates (solid lines) as a function of base line level for the spike time randomization (x-axis). The graphs of different colors in a panel represent different sample sizes of the SFC (k=200,500,2000) within a total of N=5000 observed and analyzed neurons. Each of the different panels displays the results for a specific test statistic (a: t_1^{CPC}, b: t_3^{CPC}, c: t_1^{CSF}, d: t_3^{CSF}).

[4].4.4 Results on Synfire Chain Data

[4].4.4.1 Embedded and Stimulated Synfire Chain

One obvious feature of the simulated network data is the oscillatory nature of the activity (Fig. [4].5a-c) which is a well known feature of balanced recurrent random networks. The key factors influencing the strength and frequency of the oscillations are the comparative strengths of excitatory and inhibitory synapses, the strength of the external stimulation (Brunel, 2000) and the distribution of the synaptic delays (Tetzlaff et al., 2005). In extreme these factors lead to fully synchronous activity. Here we chose parameters that result in asynchronous irregular activity, i.e., there is little synchrony between the neurons and each spike train is irregular. Nonetheless, some oscillatory structure can be observed. Irrespective of the origin of these oscillations, the generation of the surrogates have to be carefully chosen to avoid false outcomes. Surrogates with uniform spike shuffling clearly

destroy the oscillatory feature (Fig. [4].5d), which leads to false positive detections in the data with stimulated SFC (Fig. [4].6, 'u') as well as for non-stimulated SFC or even when no SFC is embedded (see Fig. [4].7, 'u'). Thus a spike shuffling approach that accounts for the coherent rate oscillations of the neurons as introduced in Section [4].2.3 seems to be more appropriate. To this end we consider the population histogram as an estimate for the inhomogeneous weighting of shuffle times including an additional base line. If neurons were not coherently oscillating, the estimator of the population histogram would not be appropriate.

The level of the base line is a free parameter and has to be chosen according to the degree of modulation of the population firing rate. Since we know which neurons are members of the synfire chain we can systematically vary the base line level and evaluate the resulting FP rate. Fig. [4].6 illustrates that the larger the base line, the higher the FP rate (dashed lines). This effect is to a large extent independent of the chosen test statistic (panels a-d), but differences are more prominent dependent on the sample size k of the neurons from the total of 2000 members of the synfire chain. When all synfire chain neurons are sampled (k=2000, green dashed), the FP rate abruptly changes from 0% FP to a high level. For sample sizes smaller than 2000, the increase of FP with base line is more gradual. In summary, a base line level of 5 leads to a 0% FP rate, independently of the test statistic or sample size.

In contrast, the FN rate (Fig. [4].6, solid lines) depends strongly on the test statistic used. The FN rate typically decreases with base line level, and the smaller the sample size k the higher the FN rate. However, for test statistic t_3^{CSF} no FNs occur (Fig. [4].6d), even for small sample sizes. Thus a choice of t_3^{CSF} and surrogates with a base line level of 5 lead to perfect detection of the member neurons of the synfire chain. Different base line levels may be tolerated at the price of a higher FP rate if further analysis emphasizing higher-order analysis of the data is performed.

[4].4.4.2 Controls

Since data from the network containing a non-stimulated synfire chain and the network with no embedded synfire chain also exhibit the oscillatory firing rate coherent across neurons, we tested the FP rate as a function of base line level (Fig. [4].7) for these control data sets as well. The FP rates behave indistinguishably from the data with stimulated synfire chains (Fig. [4].6) and do not show any dependence on the test statistic chosen. This holds true for FPs for the subset of k neurons selected from the synfire neurons and for neurons from the rest of the network (compare FP_1 and FP_2 in Fig. [4].7). Data with no embedded synfire chain show the same results. This confirms that the oscillatory

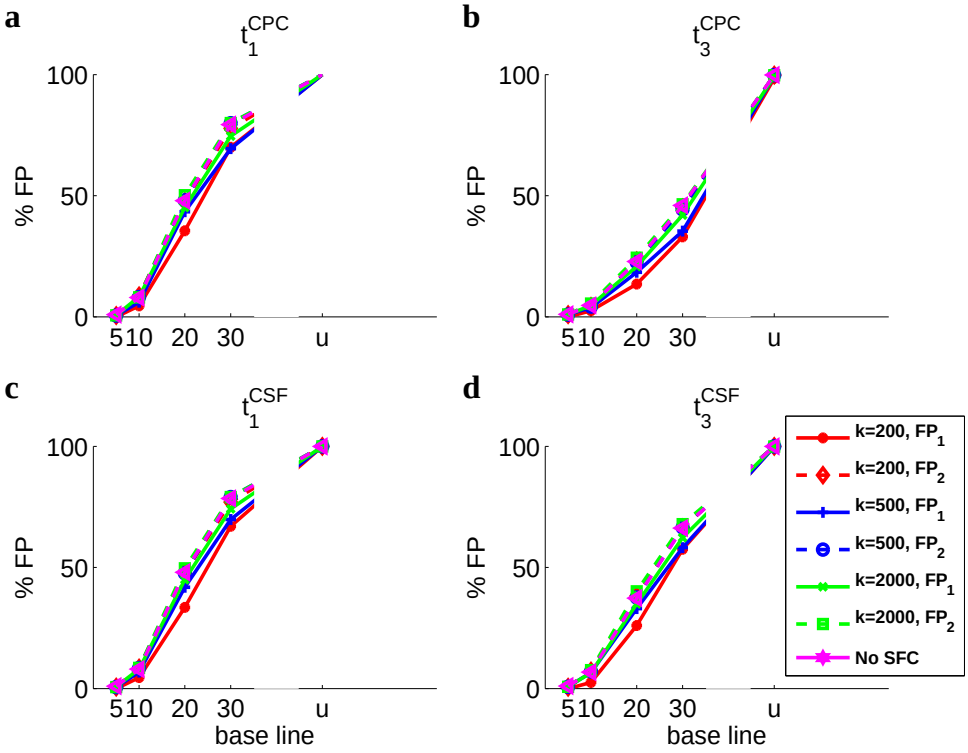

Figure [4].7: FP rates of SFCu and NoSFC data as a function of the base line level of the surrogates (spike time randomization; x-axis). Each panel displays the results for a different test statistic (a: t_1^{CPC}, b: t_3^{CPC}, c: t_1^{CSF}, d: t_3^{CSF}). The curves within each panel show the percentage of FPs for different sampling sizes from SFCu ($r = 200$, 500, 2000 (all), of a total of $N = 5000$ analyzed neurons). False positives are distinguished between FPs from neurons that are members of the embedded chain (FP_1, solid lines) and neurons that are part of the rest of the network (FP_2, dashed lines). The cyan curves show the FP rate for $N = 5000$ randomly selected neurons from a network without an embedded chain (NoSFC).

structure in the data, which is not related to the synfire activation, is responsible for increasing FP rates for increasing homogeneity of the surrogates. The best generator of FPs is uniform shuffling.

[4].4.5 Real World Spike Data

After having tested and calibrated our analysis approach, we finally demonstrate its applicability to neuronal data. We chose to apply it to a data set recorded from cat visual cortex that we have extensively studied with other analysis methods, so that we can compare the results.

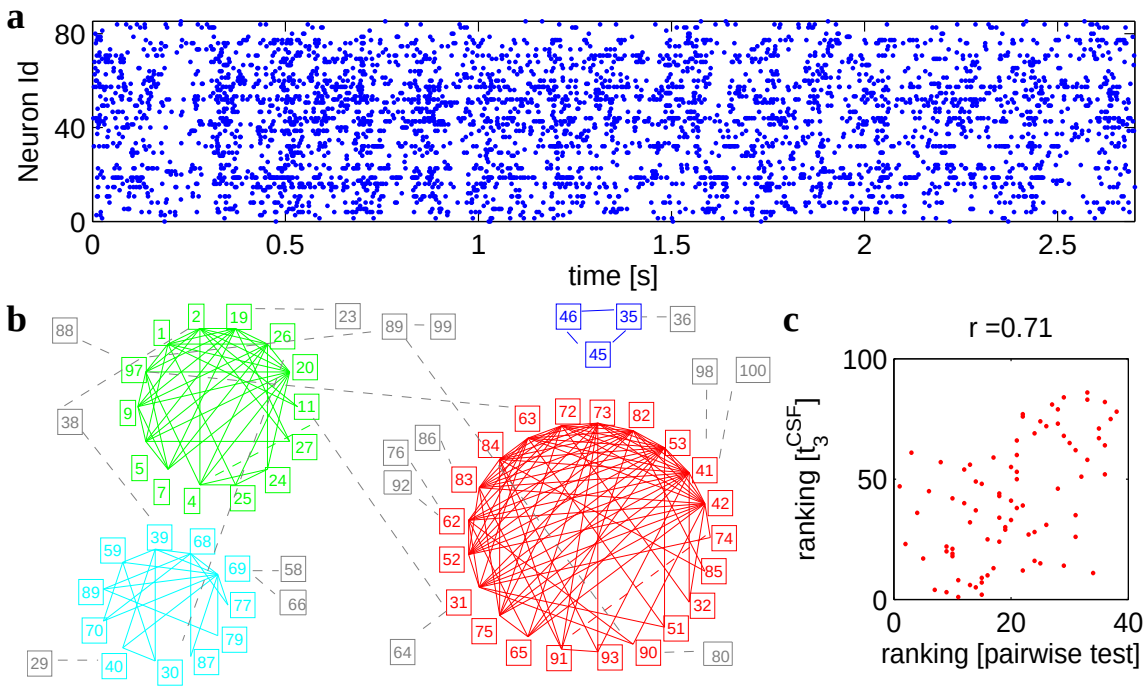

Figure [4].8: Analysis of parallel spike trains from cat visual cortex. a) Dot display of 84 multi-unit spiking activities recorded simultaneously during low intensity full flash stimulation. The dot display shows the 3^{rd} trial of low intensity stimulation lasting for 2.7 sec. b) Results of pairwise correlation analysis and subsequent clustering of groups of inter-correlated MUAs of the same data, but all trials 16 are considered. (Figure modified from Berger et al., 2007) c) Comparison of the results of the pairwise correlation analysis and our identification of potential assembly member neurons using t_3^{CSF}. Surrogates are generated by trial shuffling. Each dot in the scatter display shows the rank of each MUA according to the the number of involvements in significant pairwise correlations (x-axis; rank of 1 corresponds to the highest number) and the ranking according to the value of the test statistic (t_3^{CSF}) using trial shuffling for surrogates (y-axis; rank 1 corresponds to highest value).

[4].4.5.1 Experimental Procedures

Parallel spike recordings were obtained by a 10 × 10 electrode grid (Utah electrode array, Bionic Technologies, Inc., Salt Lake City, UT, USA) covering an area of 3.6 mm × 3.6 mm of cat visual cortex (Warren et al., 2001). Data were recorded from area 17 of anesthetized cat under full flash treatment. Data were recorded either under spontaneous condition, i.e., without visual stimulation, or with a full flash stimulus. The intensity of the latter was changed alternatingly from high intensity to low intensity, each presented for 0.51 sec and 2.7 sec, respectively. Experiments were performed under animal care and experimental guidelines that conformed to those set by the NIH (National Institute of Health). Animals

were obtained from the University of Utah Animal Resource Center. Description of the animal preparation, maintenance, and surgery procedures is fully described in Warren et al. (2001) and Nordhausen et al. (1996).

[4].4.5.2 Previous Results

In Berger et al. (2007) we analyzed the spiking activities recorded from the grid of electrodes for pairwise correlations. High intensity (HI) and low intensity (LI) epochs were separated into different data sets and analyzed separately. For simplicity we restricted ourselves to the evaluation of the multiple-unit activities (MUA), which here were composed of a mixture of typically 3 neurons. The cut-off for considering a MUA for the correlation analysis was a minimal firing rate of 1 Hz. Some electrodes were broken. This left us with 83 parallel MUA spike trains for LI and HI. From these we computed the cross-correlations (CCH) (Perkel et al., 1967) for all possible MUA pairs. To evaluate the significance of the correlation, we used a boot-strap method (spike dithering, Pazienti et al., 2008) that accounts for the changes in firing rate of the neurons. Spike correlation between two MUAs was considered significant if the original CCH exceeded the mean of the surrogate CCHs (100 surrogates, smoothed with a box car kernel of width $10 \cdot h$) by $2\ \sigma$ (i.e., a significance level of 5%). As a result 78/3402 pairs for HI and 203/3402 for LI were significantly correlated.
We noticed that individual MUAs were often not only correlated with a single other MUA, but typically with more than 2 (up to 9). This led us to a graph theoretical analysis to discover whether groups of inter-correlated MUAs can be identified. In a first step we identified cliques with at least more than 2 members, i.e., groups of MUAs that are all-to-all mutually correlated. Since these turned out to exist abundantly and with overlapping membership, we further grouped them using the criteria of a minimal overlap of one member. Interestingly, as a result the graph decomposed into a small number of completely disjoint subgraphs (see Fig. [4].8b). These clusters of highly inter-correlated neurons also clustered in cortical space with a similar space constant to orientation tuning maps found by optical imaging (Kenet et al., 2003), strongly suggesting a relationship between them.

[4].4.5.3 Comparison of Analysis Results

We re-analyzed the neuronal data with the approach presented here to identify potential assembly neurons. The dot display of the simultaneous spike trains of the cat visual cortex data (Fig. [4].8a) indicate non-stationarities in the firing rates. Additionally the firing rates of the different MUAs seem to change their rates in a coherent manner. These could be seen as arguments to use the weighted spike time shuffling procedure for the

test statistic	t_1^{CPC}	t_3^{CPC}	t_1^{CSF}	t_3^{CSF}
p_o (%)	**95**	**89**	**95**	**95**
p_{noo} (%)	**22**	**20**	**23**	**21**
r	**0.71**	**0.70**	**0.74**	**0.71**

Table [4].2: Comparison of the results. The pairwise analysis revealed 67/84 MUAs to be significantly involved in pairwise correlations. The respective analysis with our various test statistics revealed the percentage of neurons that were detected by both methods p_o relative to the total number of neurons detected by the pairwise analysis. Thus the complement indicates what percentage of neurons was not detected by using the test statistics as compared to the pairwise analysis. p_{noo} indicates the percentage of MUAs detected by the test statistics but not with the pairwise analysis, relative to the total number of MUAs detected by the respective test statistic. r indicates the correlation coefficient of the comparison of the rankings of the two approaches (cf. Fig. [4].8c for t_3^{CSF}).

generation of the surrogate data. However, it seems that not all neurons participate in this rate modulation, but either have another modulatory rhythm or do not modulate their rate. Thus, estimating the weighting distribution via the population histogram would not be appropriate for all neurons. Therefore we choose trial shuffling for the generation of the surrogates, and selected the onset of the low intensity stimulus as trigger times.
The correlation structures in the real world data are unknown to us. However, we have already studied the chosen data set extensively using other analysis methods, and it is of interest to determine whether different analysis strategies lead to consistent results. Here we aim to identify whether individual MUAs are members of groups of neurons that exhibit synchronous spiking activity, whereas in the former study we identified groups of MUAs by significant pairwise correlation and the high degree of multiple involvements in correlation and their strong overlap. Table [4].2 demonstrates that the results of the two approaches agree well. The overlap in detection of identical neurons involved in synchronous events was 95% for all but for t_3^{CPC}. Besides these the test statistics also detected MUAs as potential assembly members which were not detected by the former analysis. Their number relative to the total number detected by the respective test statistic is on average 21%. So, if we would perform the pairwise analysis again with the set of MUAs selected by our test statistic, the additionally detected MUAs would be filtered out. However, some of the previously detected MUAs involved in significant correlation are not be detected and thus also the correlation between respective pairs. Candidates for the non-detected ones are MUAs involved in only a few correlated pair correlations since our test statistic exhibits larger values for multiple involvements in

correlated groups (cf. Fig. [4].2, bottom graphs).
We further explored this by comparing the number of involvements of the individual MUAs in pairwise significant correlations to the strength of significance in the test statistic. For the latter we use the actual values of the test statistic normalized to the width of the surrogate distribution represented by the difference of its 0.99 and 0.01 quantile to avoid fluctuations due to differences in the firing rates. These values are used to rank the MUAs. On the other hand, we use the number of pairwise involvements in pairwise correlations for ranking. As a result each MUA has two ranks, which is displayed in form of a scatter plot for all MUAs (Fig. [4].8c). Clearly, the rankings by the two analyses are strongly correlated (correlation coefficient $r = 0.71$) as shown for t_3^{CSF}. Table [4].2 shows similarly high correlation coefficients for the other test statistics as well, which indicates that this result is rather robust. Multiple involvements in correlated activity of a single MUA is reflected in a larger value of the test statistic.
The pairwise correlation analysis revealed multiple involvements in correlated pairs, and grouping of many correlated pairs in clusters of highly intra-correlated members. This may hint at the existence of higher-order correlation between the neurons in one cluster, or at involvement in many, strongly overlapping assemblies. This could not been decided on the basis of the previous analysis, but requires additional higher-order analysis which are currently being developed (Shimazaki et al., 2009; Staude et al., 2009), which make use of similar measures used here for the test statistics.

[4].5 Discussion

We presented two simple test statistics, each parameterized with two different values ($\alpha = 1$ and $\alpha = 3$), to identify neurons that are involved in assemblies. Both of these statistics test for a given neuron whether it is involved in a coincidence spike event more often than can be expected by chance. To do so, t^{CPC} analyses the coincidence complexities of the parallel spike trains, and t^{CSF} aggregates pairwise frequency comparisons. In order to assess their performance, we applied these statistics, using a spike shuffling approach, to massively parallel spike trains (either 100, 1000, or 5000) generated by stochastic models and network simulations, and to real-world neural spike data. The data generated by the stochastic model served to test and calibrate the test statistics, because these models enable us to define different spike correlation structures, rate of correlation occurrence and number of processes involved in the correlation. Our various test statistics are very sensitive with a high detection rate and a low false positive rate that corresponds to the applied significance level. The false negative rate was high for very low coincidence rates

or very small percentage of neurons involved in the correlation, for $N = 100$ and $N = 1000$ respectively, but drops drastically for larger values. The statistic t_3^{CSF} (that is, t^{CSF} with $\alpha = 3$) generally behaves best. We conclude that our approach proves to be a reliable tool for the detection of potential assembly members.

The data generated by network simulations revealed additional properties that were not present in the stochastic model data. The firing rates of the neurons appear to be non-stationary in time, with a high oscillation frequency (about 100 Hz). As a consequence the way to generate the surrogates used for evaluating the significance of our measures had to be adjusted. Our results clearly demonstrate that a uniform spike shuffling is a strong generator of false positives. If instead the spike shuffling is performed according to an estimator that is based on the population firing rate, the false positive rate can be reduced close to zero if the modulation of the shuffling probabilities as a function of time is high (i.e., low base line). The control data, which do not contain any activated synfire chains, also exhibit large FP rates for increasing base line. Obviously, the present coherent rate oscillations generate a high amount of synchronous events, which, if ignored by e.g., uniform shuffling, would be interpreted as the existence of precise spike synchrony. Interestingly, the false negative rate is close to zero for t_3^{CSF}, even for smaller numbers of neurons sampled from the embedded synfire chain. Although a sampling of 200 out of 2000 neurons sounds large as a percentage, the number of neurons sampled from an individual group of the synfire chain, in which synchronous spike events occur, is rather small ($200/20 = 10$), a number comparable to the assembly size M studied for the stochastic model data.

After having tested our approach in simulated data we applied it to neuronal parallel spike data from cat visual cortex. We chose this particular data because we have already studied the chosen data set extensively using other analysis methods, in particular pairwise with subsequent clustering. Thus, we know that the data contain correlated spiking, in particular groups of neurons that are highly intra-correlated. Our analysis using our simple test statistics revealed an almost perfect detection of all the neurons that were involved in significant pairwise analysis. The MUAs that were not detected are likely to be the ones that were involved only in very few pairwise correlations. The number of involvements in pairwise analysis was clearly shown to correlate with the significance of the test statistics. On the other hand, we learned from our test with the stochastic model that the larger the assembly size, i.e., the order of correlation, the better the detection by our test statistic. Preliminary results suggest that the identified clusters on the basis of the former pairwise reflect higher-order correlations between MUAs in a cluster. Grün et al. (2008) and Staude et al. (2007) showed that the complexity distribution indeed indicates the presence of higher-order correlation in data. A preliminary analysis of the

same data based on accretion (Gerstein et al., 1978; Berger et al., 2009a) also revealed strong evidence that higher-order correlated groups of neurons are present. In conclusion, we are confident that our test statistics, which either considers the complexity of patterns or the pairwise correlations, picks up the presence of larger groups of correlated neurons. Existent methods for correlation analysis (Grün et al., 2002a,b; Shimazaki et al., 2009; Amari, 2001) and in particular higher-order correlation analysis of such high-dimensional data is very time consuming and often not applicable due to the immense memory consumption or computation time. In contrast, the test statistics we presented here prove to be very efficient and fast and is able to digest easily recordings of an order of 1000 s of parallel processes. Thus the approach proves to be an efficient method of reducing the data set to the relevant neurons, which could be followed by detailed higher-order analysis spike dithering, (e.g., Shimazaki et al., 2009; Amari, 2001). Due to the efficiency of the method one may also think of a time dependent application that may recover dynamical changes of the correlation structure. Thus, our approaches are a valuable addition to methods that provide information on the *presence* of higher-order correlation in data, but do not identify the individual neurons involved. Reducing data sets to the *relevant* neurons helps to reduce computation time considerably in further analysis steps.

Acknowledgments: Partially funded by the BMBF (grant 01GQ01413) and RIKEN Strategic Programs for R&D. We thank Markus Diesmann for valuable discussions.

References

M. Abeles. Role of the cortical neuron: integrator or coincidence detector? *Isr. J. Med. Sci.*, 18:83–92, Jan 1982.

M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, 1st edition, 1991.

M. Abeles and G. L. Gerstein. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.*, 60:909–924, Sep 1988.

A. M. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J. Neurophysiol.*, 61:900–917, 1989.

S. Amari. Information geometry on hierarchy of probability distributions. *IEEE Trans. Info. Theory*, 47:1701–1711, 2001.

D. Berger, D. Warren, R. Normann, A. Arieli, and S. Grün. Spatially organized spike correlation in cat visual cortex. *Neurocomput*, 70:2112–2116, 2007.

D. Berger, C. Borgelt, M. Diesmann, G. Gerstein, and S. Grün. Spatially organized higher-order spike synchrony in cat area 17. *8th Göttingen Meeting of the German Neuroscience Society*, pages T26–8B, 2009.

B. Bolstad. *Introduction to Bayesian Statistics, 2nd ed.* Wiley Interscience, Chichester, UK, 2007.

E. N. Brown, R. E. Kass, and P. P. Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci*, 7(5):456–461, May 2004. doi: 10.1038/nn1228. URL `http://dx.doi.org/10.1038/nn1228`.

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

J. Csicsvari, D. A. Henze, B. Jamieson, K. D. Harris, A. Sirota, P. Bartho, K. D. Wise, and G. Buzsaki. Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J. Neurophysiol.*, 90:1314–1323, Aug 2003.

J. E. Dayhoff and G. L. Gerstein. Favored patterns in spike trains. II. Application. *J. Neurophysiol.*, 49:1349–1363, Jun 1983.

S. Eldawlatly, R. Jin, and K. G. Oweiss. Identifying functional connectivity in large-scale neural ensemble recordings: a multiscale data mining approach. *Neural Comput*, 21: 450–477, Feb 2009.

G. L. Gerstein, D. H. Perkel, and K. N. Subramanian. Identification of functionally related neural assemblies. *Brain Res.*, 140:43–62, Jan 1978.

G. L. Gerstein, P. Bedenbaugh, and M. H. Aertsen. Neuronal assemblies. *IEEE Trans Biomed Eng*, 36(1):4–14, Jan 1989.

M. Gewaltig and M. Diesmann. Nest (neural simulation tool). *Scholarpedia*, 2:1430, 2007.

S. Grün. Data-driven significance estimation for precise spike correlation. *J. Neurophysiol.*, 101:1126–1140, Mar 2009.

S. Grün, M. Diesmann, F. Grammont, A. Riehle, and A. Aertsen. Detecting unitary events without discretization of time. *J. Neurosci. Methods*, 94:67–79, 1999.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Comput*, 14:43–80, 2002a.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput*, 14:81–119, 2002b.

S. Grün, M. Abeles, and M. Diesmann. Impact of higher-order correlations on coincidence distributions of massively parallel data. *Lecture Notes in Computer Science*, 5286:96–114, 2008.

S. Grün, D. Berger, and C. Borgelt. Identification of neurons participating in cell assemblies. *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3493–3496, 2009.

K. D. Harris. Neural signatures of cell assembly organization. *Nat. Rev. Neurosci.*, 6: 399–407, May 2005.

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

A. Kohn and M. A. Smith. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.*, 25:3661–3673, Apr 2005.

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

A. Morrison, S. Straube, H. E. Plesser, and M. Diesmann. Exact subthreshold integration with continuous spike times in discrete-time neural network simulations. *Neural Comput*, 19:47–79, Jan 2007.

C. T. Nordhausen, E. M. Maynard, and R. A. Normann. Single unit recording capabilities of a 100 microelectrode array. *Brain Res.*, 726:129–140, Jul 1996.

A. Pazienti, P. E. Maldonado, M. Diesmann, and S. GrÃ$\frac{1}{4}$n. Effectiveness of systematic spike dithering depends on the precision of cortical synchronization. *Brain Res.*, 1225: 39–46, 2008.

D. H. Perkel, G. L. Gerstein, and G. P. Moore. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys. J.*, 7:419–440, 1967.

G. Pipa and S. Grün. Non-parametric significance estimation of joint-spike events by shuffling and resampling. *Neurocomputing*, 52-54:31–37, June 2003.

E. Schneidman, M. J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007–1012, Apr 2006.

S. Schrader, S. Grün, M. Diesmann, and G. L. Gerstein. Detecting synfire chain activity using massively parallel spike train recording. *J. Neurophysiol.*, 100:2165–2176, Oct 2008.

H. Shimazaki, S. Amari, E. Brown, and S. Grün. State-space analysis on time-varying correlations in parallel spike sequences. *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2009.

J. Shlens, G. D. Field, J. L. Gauthier, M. I. Grivich, D. Petrusca, A. Sher, A. M. Litke, and E. J. Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.*, 26:8254–8266, Aug 2006.

T. Shmiel, R. Drori, O. Shmiel, Y. Ben-Shaul, Z. Nadasdy, M. Shemesh, M. Teicher, and M. Abeles. Temporally precise cortical firing patterns are associated with distinct action segments. *J. Neurophysiol.*, 96:2645–2652, 2006.

W. Singer, A. K. Engel, A. K. Kreiter, M. H. J. Munk, S. Neuenschwander, and P. R. Roelfsema. Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences*, 1(7):252–261, 1997.

B. Staude, S. Rotter, and S. Grün. Detecting the existence of higher-order correlations in multiple single-unit spike trains. *Soc. Neurosci. Abstr. 103.9/AAA18, Washington, DC,*, 2007.

B. Staude, S. Rotter, and S. Grün. Cubic: cumulant based inference of higher-order correlations in massively parallel spike trains. *submitted*, 2009.

I. V. Tetko and A. E. Villa. A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 2. Application to simultaneous single unit recordings. *J. Neurosci. Methods*, 105:15–24, Jan 2001.

T. Tetzlaff, A. Morrison, M. Timme, and M. Diesmann. Heterogeneity breaks global synchrony in large networks. *Proc 30th Göttingen Neurobiology Conference*, page 206B, 2005.

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

Chapter 5

Statistics of eye movements of monkeys freely viewing natural scenes

Statistics of Eye Movements of Monkeys Freely Viewing Natural Scenes

Denise Berger[1,2], Martin Nawrot[1,2], Pedro Maldonado[3], and Sonja Grün[2,4]

[1]Neuroinformatics, Inst. Biology - Neurobiology, Free University, Berlin, Germany
[2]Bernstein Center for Computational Neuroscience, Berlin, Germany
[3]Physiology and Biophysics, Faculty of Medicine, University of Chile, Santiago, Chile
[4]RIKEN Brain Science Institute, Wako-shi, Japan

This section is to be submitted as

Berger, D.; Nawrot, M.; Maldonado, P. & Grün, S.
Statistics of eye movements of monkeys freely viewing natural scenes.

Abstract

In order to uncover possible mechanisms guiding the eye movements during free viewing, we investigated the statistics of the eye movements of three monkeys freely exploring natural images. Therefore, in a first step we analyze the spatial distributions of the fixation points and correlate them with the saliency features of the images explored. As a result we find that guidance of eye movements is for most of the presented images determined by low level features as represented by the saliency of the images. However, for images that contained primate faces the fixation positions are not explained by the exogenous, stimulus features but indicate involvement of endogenous activity. In a next step, we investigate features of the eye movement trajectories to capture their sequence dependence. Fixations on images cluster in space which are interpreted as regions-of-interest. A Markov chain analysis reveals that these regions are not randomly visited during eye movements, but eye movement trajectories are organized in sequences of local explorations of the regions-of-interest. These result resemble close similarities to visual behavior of humans and thus encourage further work on free viewing monkeys to relate neuronal activity to behavioral paradigms that involve attentional aspects.

[5].1 Introduction

Eye movements are known to indicate attention and internal state of humans (Noton and Stark, 1971a; Biedermann, 1987). Human eye movements in free viewing conditions have been extensively studied in the last decades starting from Buswell (1935) and Yarbus (1967). It has been shown, that humans are strongly attracted by low-level features in the images expressed in the saliency of the images (Itti and Koch, 1999a, 2001). On the other hand, human eye movements are also strongly affected by objects which are meaningful to their requested behavior (Yarbus, 1967; Just and Carpenter, 1967). Although several models have been proposed to functionally account for such visual attention of humans (see, e.g., Wolfe, 1994; Milanse et al., 1995; Tsotsos et al., 1995; Itti et al., 1998), it remains unclear how the brain processes such information and guides the eye movements. Obviously it would be of high interest to understand the brain mechanisms underlying active vision. Electro-physiological studies in animals can access the neuronal level but are typically performed with optimal stimuli in the receptive fields of the recorded neurons under fixation requirement (e.g., Hubel and Wiesel, 1968; Maunsell and Newsome, 1987), or where the task required precise eye movements to pre-defined target locations(Vinje and Gallant, 2000). Thus studies that involve free viewing of natural scenes, are required to allow studying viewing under natural conditions. However, such paradigms are typically not used in animal experiments, since it makes the analysis of neuronal data hard due to the lack of 'trials'. Thus, retrieving a better understanding of eye movement strategies of monkeys can provide new experimental paradigms and insights. First indications of similar visual behavior of monkeys and humans was shown in a speed categorization task. Monkeys perform only with slightly decreased accuracy but greater speed than human observers (Fabre-Thorpe et al., 1998; Thorpe et al., 1996). These results reveal striking similarity between monkeys and humans regarding rapid processing of natural scenes. Nevertheless, it remains to be investigated whether or not this similarity generalizes to viewing under natural conditions, when attentional processes and eye movements take effects.

In order to uncover possible mechanisms guiding the monkeys' fixations during free viewing, we here investigate the eye movements of monkeys, while they were allowed to freely explore natural images. In particular we are analyzing the spatial positions the monkeys fixate on and the strategy they use to explore the images. Thus, in the first part of this paper, we investigate in how far image features defined by their saliency determine the guidance of eye movements and fixation positions in contrast to potentially meaningful items. We find that the eye movements are mainly saliency driven with the exception if

primate faces are contained in the images.
In the second part we investigate how the images are explored, i.e., whether the monkeys employ a certain strategy in the spatial sequence of fixation positions while scanning the images. Using a mean shift algorithm we determine clusters of a high density fixation positions to investigate then the scan path by a Markov chain analysis. We find that image exploration is non-random but involve a strategy where region-of-interests (ROIs) are locally explored before directing attention to a new ROI.

[5].2 Methods

[5].2.1 Experiment

Three adult, male capuchin monkeys (Cebus apella) weighing 3-4 kg served as subjects for this study. All experiments followed institutional and NIH guidelines for the care and use of laboratory animals. All surgical and recording procedures are described in Maldonado et al. (2008). After the animals participated in the recordings, their head post, eye coils, and manipulator were removed. After full recovery the animals were donated for adoption (Maldonado et al., 2008). The animals were seated in a chamber dimly lit at the low scotopic level. A selection from a total of 11 pictures, with a size of 30 x 40 cm, of different natural scenes were presented in random order on a 21-inch computer monitor, with a frame rate of 60 Hz comprising a total visual angle of 30×40 degrees, located 57 cm in front of the animal. Each image was presented up to 3 times in a single recording session. The animals were allowed to freely explore natural images (in the following referred to as images). The experimental protocol required the animals to maintain their gaze for up to 5 s within the limits of the visual image in order to be rewarded. The images were presented interleaved with blank frames containing a single central fixation spot (referred to as fixation cue), in order to maintain alertness, which the animals had to fixate (within a 1^o window) for 1 s in order to be rewarded. Fig. [5].1 shows a sketch of the time course of the experiment. Vertical and horizontal eye position were monitored throughout the experiment by the use of a search coil driver at 2 kHz (DNI Instruments, Resolution: 1.2 min of arc and then digitized (see Judge et al., 1980, for details).

[5].2.2 Detection and Statistics of Eye Events

We developed an automatic algorithm to extract different types of eye movements from the eye traces (Maldonado et al., 2008). Eye movements are categorized in two different

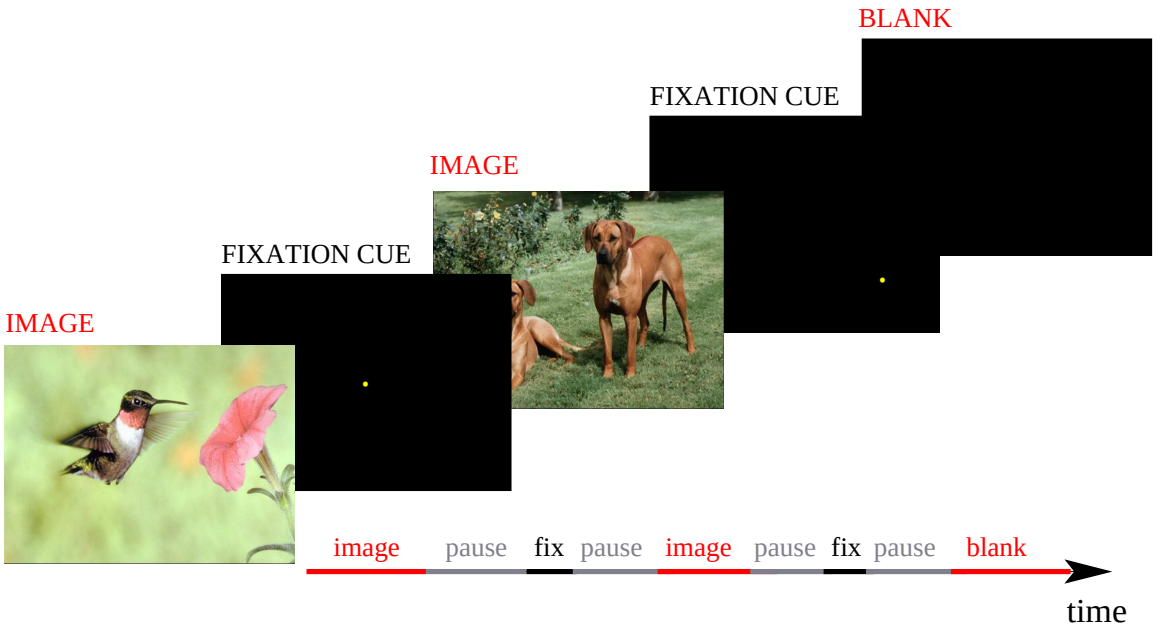

Figure [5].1: Time course of the experiment. Images, blanks and blanks with fixation spot (white dot) were presented in interleaved manner, with a pause of 2-3 s between image and fixation cue. Monkeys had to keep their gaze within the limits of the images or blanks for 3-5 s, and had to fixate the fixation cue for 1 s. Successful behavior was rewarded with a drop of juice. Per session a set of 4-7 different images (out of 11) were presented in random manner. The size of the images and blanks was 30 × 40 cm, which corresponded to a visual angle of 30 × 40 degree.

groups, i.e., saccades and fixations (cf. Fig. [5].2a,b), according to the following criteria: Saccades are defined as eye movements with an angular velocity higher than 150^o/sec, lasting for at least 5 ms. In addition, saccades are required to exhibit a minimum acceleration of $170^o/\text{sec}^2$. Fixation periods are classified as eye positions lasting at least 100 ms within 1^o of the gaze location reached at the end of a saccade. Data which can not been assigned into one of the two categories, are not been taken into account for further analysis. Only pairs of unambiguous saccade-fixation sequences are considered for further analysis.

Basic statistics of fixation and saccade durations pooled over all experiments are shown in Fig. [5].2c,d. The distributions of fixation durations of two monkeys (d and m) reveal high similarities, with a mean fixation duration of 0.31 s similar to the duration of fixations in humans (between 0.26 and 0.33 s) during scene perception (Castelhano and Henderson, 2007). The distribution of fixation durations of one of the monkeys (s) however differs from the others. The fixation durations of that monkey (s) shows a broader distribution, with its mean value of 0.42 s shifted to the right in comparison to that of the other two

monkeys (d and m).

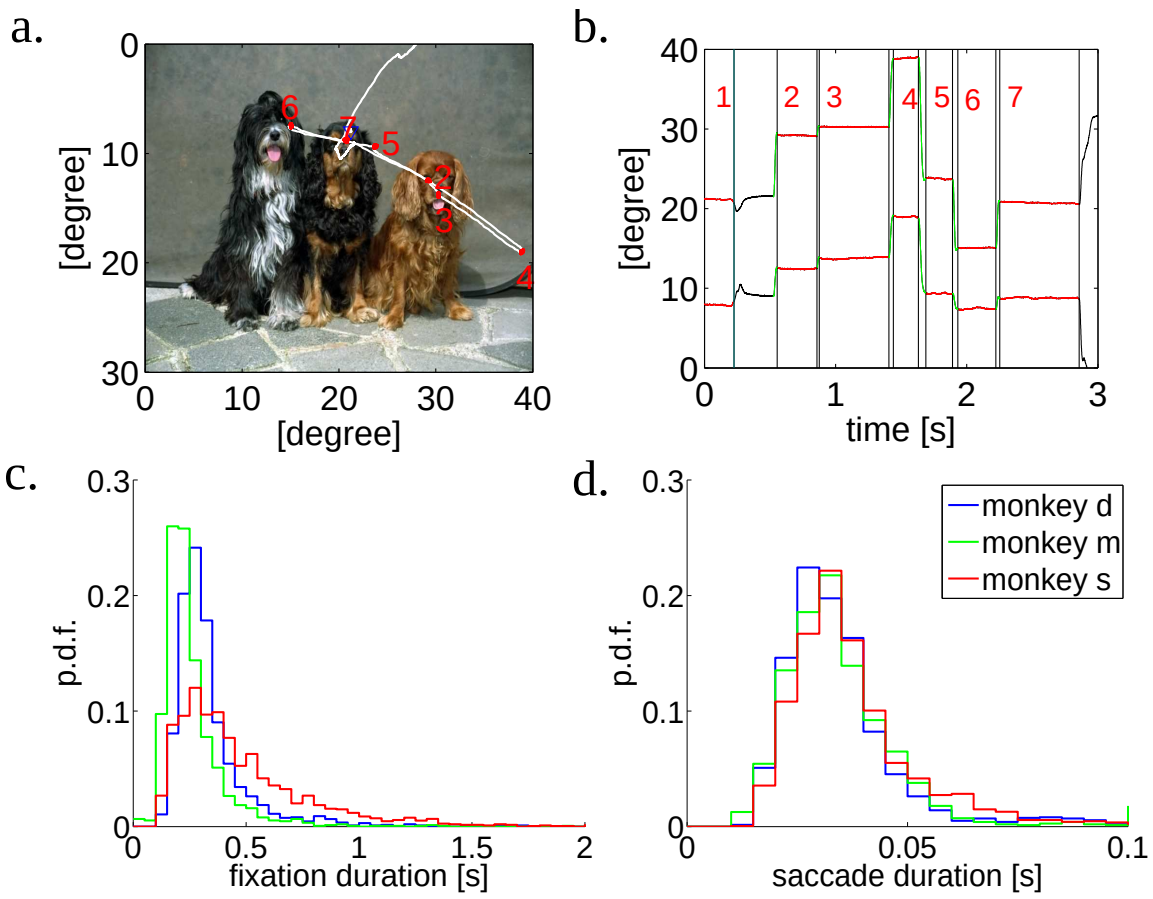

Figure [5].2: Statistics of duration of eye movements. a. Example of a presented image (image no. 3) with the scan path (white lines) during one single image presentation. The red dots indicate the positions of fixations. b. Time course of eye positions during the same trial (upper trace: horizontal, lower trace: vertical). Fixations are marked in red, saccades in green, and unclassifiable events are colored in black. c. Distributions of saccade durations (bin width: 5 ms) and d. distributions of fixation durations (bin width: 50 ms), pooled over all sessions and all images presentations, separately for each monkey (d: blue, m: green, s: red).

[5].2.3 Saliency Maps

In order to relate the visual foci of the monkeys as expressed by the fixation positions on the images, we will finally correlate fixation maps (cf. Section [5].2.4) to the saliency maps. The saliency maps of the individual images is calculated according to the model by Itti et al. (1998). This model is inspired by the neural architecture of the early primate visual system and consists of multi-scale feature maps that allow the detection of local

discontinuities in intensity, color, and orientation.
The images presented to the monkey are segregated into three separate feature maps, one for intensity, color, and orientation, as described in detail by Itti et al. (1998). In a second step, each feature is re-organized into a center-surround arrangement that is characteristic of receptive field organization (Hubel and Wiesel, 1962), and highlights the parts of the scene that strongly differ from their surroundings. This is achieved by the difference between fine and coarse scale applied to the feature maps to extract local enhanced intensities for each feature type. In the last step these conspicuity maps are normalized to the total number of maps and added to yield the final saliency map $s(x, y)$ (see Fig. [5].4a).

[5].2.4 Fixation Maps

As a measure for the regions in the images that preferably attract the interest of the monkeys we compute a fixation map for each image and monkey. All fixations on a particular image are pooled across different experiments and trials (see Fig. [5].5a for examples) to calculate the 2D-probability distribution of the fixations $f(x, y)$. Therefore we discretized the images (original resolution: 600 $\times$ 800 pixel) into 2D-bins (x,y) of 30 $\times$ 40 pixels. Within each bin the number of fixations are counted and the resulting 2D-histogram is normalized by the total number of fixations that entered the histogram to yield a probability distribution.

[5].2.5 Similarity of Maps

To quantify the similarity between the saliency map of an image and the respective fixation map we calculate the Kullback-Leibler distance (KLD) (Kullback and Leibler, 1951) between the two. The KLD is an information theoretical measurement for the distance between two probability density functions (pdfs), here $s(x, y)$ and $f(x, y)$ defined as:

$$KLD(s, f) = \sum_x \sum_y s(x, y) log \frac{s(x, y)}{f(x, y)}$$

The KLD measures the 'closeness' between two pdfs. It is always nonnegative. The smaller the KLD, the higher the similarity between the two pdfs, with its lower bound at zero, if the two pdfs are identical.
To evaluate the significance of the measured KLD value (KLD_{act}) we calculate the distribution of the KLD assuming independence of the two maps (KLD_{ind}). This is achieved by a 'random viewer' where the same number of fixation points as in the original viewing

behavior are distributed randomly (homogeneously) onto the image. The KLD_{ind} between the random viewing fixation maps and the original saliency map are computed for 1000 times to generate the distribution of KLD values under independence, thus allowing to perform the significance test for the original KLD at a level of 1 % (cf. Fig. [5].4b).
For visualization purposes (cf. Fig. [5].4c), we show for each image the difference of the actual KLD value and the mean of the independent values: $\Delta\text{KLD} = KLD_{ind} - KLD_{act}$. Positive values of ΔKLD denote a higher similarity between the actual fixation and saliency map than expected, indicating that the saliency map is a good predictor for eye movements. On the contrary, negative values signify that the distance between the actual fixation map and the saliency map is larger than expected when assuming random viewing. Significant deviations from independence are marked by an asterisk.

[5].2.6 Spatial clustering of fixation positions

In order to investigate the existence of particular regions-of-interest (ROI) defined as areas with high density of fixation positions, we identify spatial clusters of fixations by use of the mean shift algorithm (Funkunaga and Hostetler, 1975; Comaniciu and Meer, 2002) adapted for eye movement data (Santella and DeCarlo, 2004). This is an automatic, entirely data-driven method, that derives the number and arrangement of clusters deterministically.
It starts from a set of N fixation positions with x_i and y_i, $i \in 1, .., N$, being the horizontal and vertical positions of fixation position v_i on the 2D screen. The clustering algorithm proceeds iteratively, while moving with each iteration each of the points to its new position $s(v_i)$ depending on the weighted mean of proximity and density of points v_j with $i \neq j$ around the reference point: $s(v_i) = \sum_j k(v_i - v_j)v_j / \sum_j (v_i - v_j)$. The kernel is defined as a Gaussian with a mean and variance of 0: $k(v) = \exp(x^2 + y^2)/\sigma_s^2$ with σ_s assumed here as 2.5 cm, which correspond to 2.5 degrees of visual angle. Thereby the positions are iteratively moved into denser configurations by estimating the gradient of the density of points at v_j as $s(v_j) - v_j$. The procedure is stopped after convergence. All fixations are assigned to a cluster with those fixations, whose reference points lay within a diameter of 1^o apart. Robustness to extreme outliers is achieved by limiting the support of points at large distances as defined by the kernel $k(v)$.
In order to discard outliers, we additionally apply a significance test to disregard clusters containing only a very small fraction of the data which deviate from expectation of independence. For a significance test on the estimated cluster, we simulate fixation positions (same number as in the original) of a 'random viewer' and apply also on these data the mean shift algorithm. This is repeated 100 times for each image. For each we calculate

the expected number of fixation positions per cluster and their density. An original cluster of fixation positions is identified as significant, when the number of fixations and their density exceeds the mean ± 2 sd of the cluster derived from the random viewers. Only if a cluster is identified as significant, it will enter the analysis described in the following section.

[5].2.7 Transition probabilities between fixation clusters

The statistical properties of the scan paths the monkey chose to explore an image is analyzed by a Markov chain (MC) analysis (Markov, 1913). A MC is a sequence of random variables with propagation through a chain of states in accordance with given transition probabilities. These are estimated from the data by the normalized frequency how often a specific state x_j transits into any particular other state x_l or stays within the same one: $P_{j,l} = P(X_{t+1} = x_j | X_t = x_l)$. Per image, we assume the identified clusters of fixation points as its states. The transition probabilities from any one state of these states to any other are represented in matrix form. The state of the system at step t with $t = 1, ...T - 1$, with T being the total number of fixations on an image, is derived via $Pr(X_{t+1} = x | X_t = x_i, .., X_1 = x_1) = Pr(X_{t+1} = x | X_t = x_i)$ for all states $x_i \in x_1, ..x_n$. Thus we assume that the scan paths of the monkeys satisfy the Markov property, i.e., that, given the present state, the past states are independent.
For better intuition, the MC is visualized by a transition graph, in which the vertices are identified with the states. The graph has oriented edges from each vertex to any other, which are weighted with the corresponding transition probability between the respective states. Whenever the transition probability equals zero ($P_{j,l} = 0$), no connection is made between the two vertices. In addition, each vertex has an edge to itself weighted by the probability of staying within the same state in the next step. For visualization purposes we represent the transition probabilities by the thickness of the edges (cf. Fig. [5].5a), and thereby deviate in the graphical display from conventional transition graphs.
In order to interpret the transition probabilities derived by the MC we compare them to the transition probabilities assuming homogeneous chance probabilities of the transitions between any two fixations. Based on the numbers of fixations f_m of each state x_m, we calculate the expected transition probabilities for each state x_i to every other state x_j by $P(X_{t+1} + x_j | X_t + x_i) = \frac{f_j}{T}$, with T the total number of transition steps.

[5].3 Results

[5].3.1 Topology of Fixation Locations

In a first step we investigate the topology of fixation positions of the different monkeys (s, d, and m) on the individual images (see Fig. [5].3 for examples). For two of the monkeys (d and m), the spatial distribution of fixation locations clearly shows dense clusters of fixations which differ significantly depending on the image. These clusters are assumed to reflect regions-of-interest of the monkey. Overall, monkeys d and m show a similar characteristics in the topology of eye positions for the individual images. However, the fixations of the third monkey (s) show different distributions compared to the other monkeys. In addition, his fixation distributions do not differ for the different images, but form a large cluster within the lower left part of the images, irrespective of the content of the images. Additional clusters are not observed.

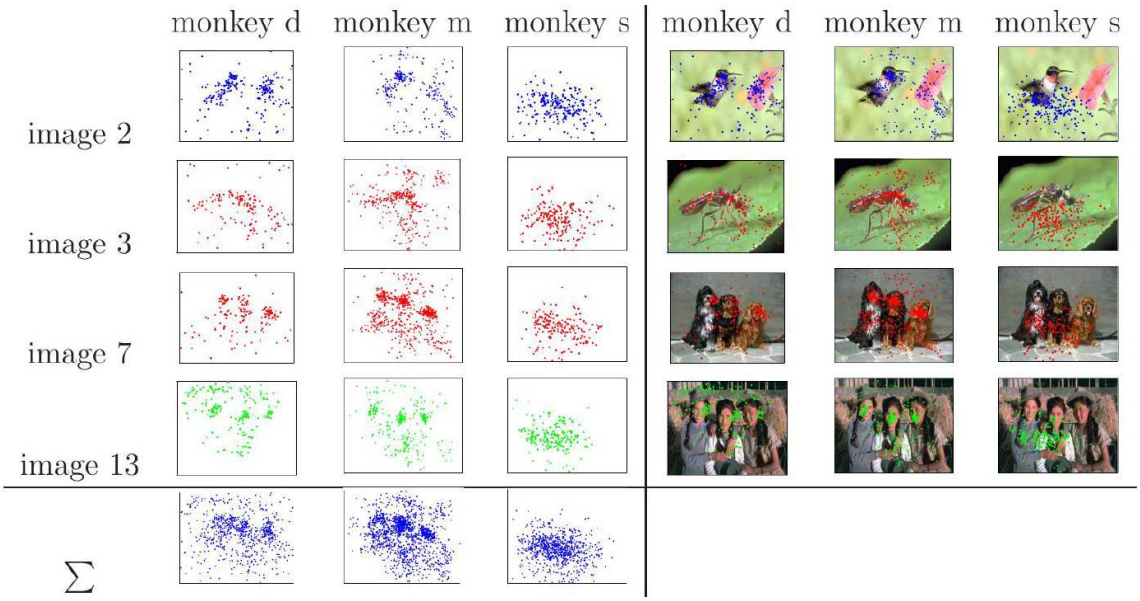

Figure [5].3: Spatial distribution of fixations on example images. Each column corresponds to one of four example images (shown right). Distributions of fixation positions (dots) are collected from all presentations of the respective image, over all sessions. Left three columns: fixation positions per monkey and images; right three columns: same data, with presented images underlaid. Merged fixation positions from the four images above are shown in the bottom row.

[5].3.2 Saliency as driving force for non-primate images

In the next step of our analysis we take into account the features of the images to uncover if the fixation positions are related to the images. In particular we are interested to what extent stimulus-driven or attention related factors influence the gaze position. To capture stimulus-driven components of the fixations we estimate the saliency of the features of the images by computing their saliency map based on the model by Itti et al. (1998) (see Methods, Section [5].2.3). This model combines three different conspicuity maps that consider orientation, intensity, and color into one saliency map (see Fig. [5].4a for examples). To quantify the correlation between fixation positions (see green dots on the saliency maps in Fig. [5].4a) with the saliency maps, we derive for the same images also its fixation maps (see Methods, Section [5].2.4). We capture the similarity between the saliency and the fixation map by their Kullback-Leibler distance (see Methods, Section [5].2.5). Significance of the actual KLD is derived by comparison to the KLDs resulting from random viewing. For better illustration of similarity we express the difference to the KLD of the random viewer by the difference (Δ KLD) of the actual KLD values and the mean of random fixation positions. Figure [5].4c shows the results for each of the images and for each of the monkeys. Positive entries mark images for which the fixation distributions are fully explained by the salient features, whereas negative entries mark images where the fixation positions are not explained by the saliency map. Entries marked by an asterisk signify images for which the monkey's fixation pattern differed significantly from a random viewer (see Methods, Section [5].2.5).

For monkeys d and m the KLD was significantly ($P<0.01$) smaller than a random viewer for about 75% of the images (in 8 of 11 images for monkey d and 4 of 5 for monkey m), i.e., the fixation maps were very similar to the saliency maps. Thus, for these the saliency maps are good predictors for the fixation positions on these images. However, for the remaining images the KLD is significantly larger than expected by a random viewer, i.e., the fixation map differs significantly from the saliency map and thus may be interpreted as avoidance of salient features. The contents of these images differs systematically from the former ones. They show faces of primates, while the former did only contain non-primate animals. Images containing primates were primarily fixated close to the primates eyes, areas that do not exhibit salient features.

For monkey s the results are totally different. The KLDs do not significantly deviate from expectation in all of the cases. This is consistent with the observation that this monkey did not 'explore' the images as the other two monkeys did. Instead he kept his gaze within the lower left part of the image, irrespective of the presented image. Thus, it is highly likely that for this monkey the behavioral training failed. Exploration was not

controlled in the experiments but the monkeys were rewarded when they kept their gaze on the image, also when the monkey fixated only on a restricted area.
In summary, we find for the fixation positions of the monkeys that actively explore the natural scenes are driven by low-level image features, if the images did not show primates. In case of images showing primates the eye movements are driven by other mechanisms, most likely by attentional aspects since primate faces provide to other primates very important cues for their social live.

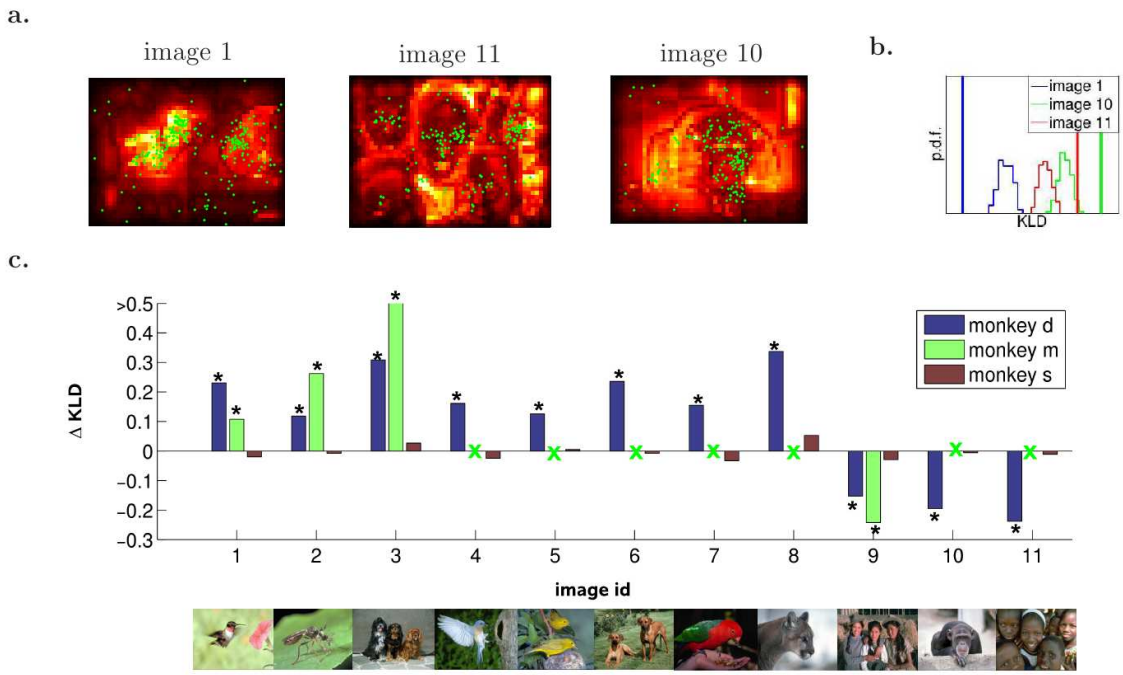

Figure [5].4: Correlation of saliency and fixation maps. a. Saliency maps for 3 example images combined with fixation positions of monkey d (green dots) pooled from all sessions and all trials. b. KLDs resulting from correlating saliency maps and fixations maps (same data as in a.). Actual KLDs are marked by vertical lines, KLD distributions of random fixation positions (1000 drawings) are shown as histograms. Different colors mark the three example images shown in a. c. Differences (Δ KLD) of actual KLD value and mean of random fixation positions for all images (x-axis, images below). Colored bars per image correspond to the Δ KLDs from the different monkeys (blue: monkey d; green: monkey m; brown: monkey s). Stars mark Δ KLDs that deviate significantly from expectation (P<0.01). Green X mark images which had not been presented to monkey m.

[5].3.3 Sequence of ROI exploration

The cluster analysis of monkey d and m reveals that the fixation positions on a particular image cluster in space and do not appear homogeneously in space (cmp. Fig. [5].3). Application of an automatic clustering approach (see Methods, Section [5].2.6) revealed typically 3-5 significant fixation clusters on the images presented to monkey d and m. For monkey s no significant clusters could be extracted. Thus for this reasons and the fact that monkey s did obviously not actively explore the images, we exclude this monkey from the following analysis which concentrates on mechanisms underlying the scan path. In order to investigate the chronological order in which these ROIs are visited, we explore the scan paths of image explorations by applying a Markov chain (MC) analysis on the eye movement trajectories (see Methods, Section [5].2.7). We assume the extracted ROIs as Markov states and derive by the MC analysis the transition probabilities for staying within a cluster and for transition from one cluster to the next. Data underlying one transition graph are pooled from all eye movement sequences to all presentations of the same image to one monkey. Figure [5].5a shows examples of eye movement sequences for a few presentations of the same image. The identified fixation clusters are illustrated in Figure [5].5b and the transition graphs between these three identified clusters are shown in Figure [5].5c. To derive significance of the actual transition probabilities, we compare them to chance transition probabilities assuming homogeneous probabilities between any two fixations. As a result we find for the majority of the identified clusters across all images (87 % (40 out of 46) for monkey d, and 95% (19 out of 20) for monkey m). that the probabilities to stay within the clusters are significantly higher than expected by chance. On the contrary, the transition probabilities between clusters are for the majority of the cases (60/97 transitions for monkey d, and 83/134 transitions for monkey m significantly lower than chance. These results are summarized as distributions in Figure [5].5c and d. The comparison of the distributions of the actual and the expected transition probabilities reveals that they are significantly different for both conditions, and holds true for both monkeys (Wilcoxon test $P< 10^3$).

To summarize, the Markov chain analysis reveal that the monkeys primarily saccade between fixation positions within the same, local cluster, and are saccading out of the cluster into others or to the background with lower probability. These results show no dependence on the contents of the images, in particular with respect to primates. Thus the monkeys seem to perform local explorations, interrupted by jumps to new regions-of-interest for further exploration.

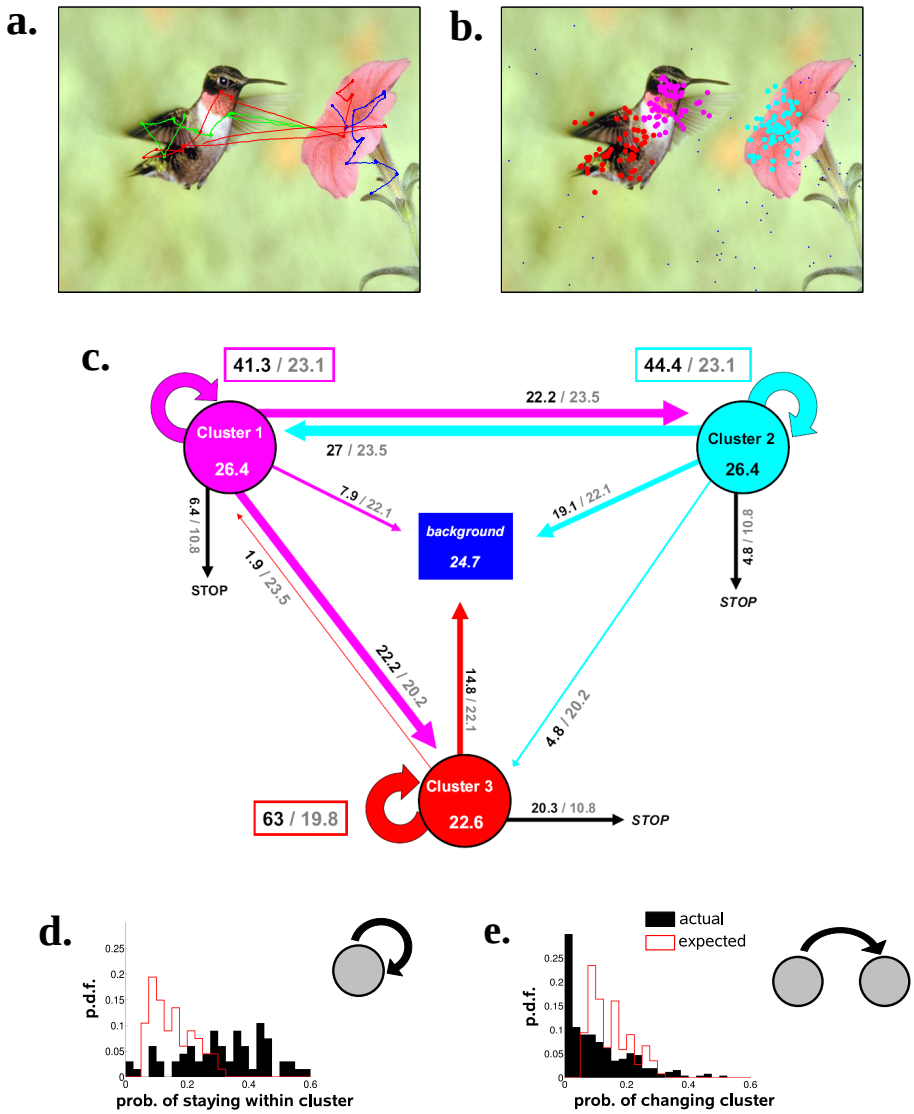

Figure [5].5: Transition probabilities between clusters of fixations for monkey d viewing image 2. a. Example of eye movement trajectories during exploration in 22 different trials (different colors). b. Spatial clusters (n=3) of fixation positions extracted by the mean shift algorithm. Fixation points in the respective clusters are colored in red, cyan, and magenta. c. Transition graphs between identified clusters (b.) interpreted as the states of the model (same colors as in b.). Fixations that do not fall into any of the identified clusters are collected into the background state (dark blue box). Experimental transition probabilities are indicated in black, expected probabilities assuming random transitions in gray. STOP refers to trial ends after exceeding the maximal image presentation time. d. Experimental (black) and expected (red) distributions of the probability to stay within the same cluster and of e. transition probabilities to change between clusters

[5].4 Discussion

Investigation of the topology of fixation locations showed in two out of the three investigated monkeys (monkey m and d) specific spatial patterns of fixation positions dependent on the presented images. One of the three monkeys (s), however, showed a completely different viewing behavior. His fixations did not differ for the different images, but were rather restricted to a limited area of the image frame. Given the consistent behavior of the other two monkeys, we have to assume that this monkey did not learn to explore the images but rather to keep his gaze within the lower left quadrant of the frame, irrespective of the contents of the image. This is highly likely since the experiment did not ask for a certain task besides staying with the gaze somewhere within the image, thus the monkey was nevertheless awarded. This view is also supported by the correlation analysis of fixation positions and saliency maps, which shows no indication that the fixations of the monkey were driven by features of the images. Therefore we restricted our analysis of the sequence analysis of the scan paths to the monkeys that explored the images, and therefore will in the following discuss the results of monkeys m and d.

Their fixations cluster in space, i.e., on locations of the images that appear to be relevant to the monkeys and we interpret these clusters as region-of-interests. Similar viewing behavior was found in humans which observe complex scenes by restricting their fixations to locations rich of details instead of distributing their attention evenly across the entire scene. The scene, therefore, can be divided into informative features and redundant regions as for humans (Mackworth and Morandi, 1967; Yarbus, 1967; Krieger et al., 2000). To get insight into the mechanisms that guide the visual exploratory behavior of the monkeys, we analyzed the images for salient features. Theses are are likely candidates for the attraction of the monkeys' gaze since no particular task was requested in contrast to studies in humans where the subjects were involved in search tasks or had to judge the content of the images (Yarbus, 1967). We find that for approximately 75% of the images their saliency maps are good predictors for the fixation positions suggesting that the eye movements are mainly driven by low-level features.
However, for the remaining images the eye movements seem not directly related to salient features of the images. Common to the images is that they contain faces of primates (monkeys or humans). This viewing behavior may be rather attributed to top-down mechanisms, where attention to meaningful details of the objects determines the location of gaze. These results resemble similarity to those found in humans where the choice of fixation positions was shown to be either driven by bottom-up (exogenous) or by top-down (endogenous), internal or external factors (McKelvie, 1976). Thereby, faces play a

particular role and are probably the most important visual stimuli in primate social communications (Bruce and Young, 1998) as these can provide significant cues to intention and mental state of other individuals (Bruce and Young, 1998; Emery, 2000; Andrew, 1963; Anderson, 1998). Similar observations exist in non-human primates. At an early age, monkeys respond appropriately to the expressions of other individuals (Mendelson et al., 1982), and have been shown to be able to recognize the faces of other individuals (Rosenfeld and Van Hoesen, 1979).

In psychological studies it was shown that the sequences of saccades and fixations are relevant for perception (Noton and Stark, 1971b). We find that the monkeys fixate preferably at certain locations on the images, and the eye movements between these regions-of-interests are not random. The Markov chain analysis revealed that the monkeys primarily saccade within positions of a region-of-interest, i.e., locally explore, before directing their eye movement by a larger saccade to a new region-of-interest (global exploration). This finding is consistent with the hypothesis that large saccades to new areas are followed by local, short saccades to near-by positions for refinement of the percept (Körner et al., 1999; Ullman, 1995). Thus our results indicate that the monkeys actively explore the visual scenes.

Experiments including active vision, i.e., without the request that eyes fixate on a predefined position, are hard to find in experiments that involve electro-physiological recordings since they do not contain repetitive, identical trials and thus are hard to analyze. This study provides new approaches to data from free viewing animals and thus open new routes for experiments that aim to relate neuronal activities to natural behavior. Electrophysiological studies that involve the presentation of natural stimuli such as movies to fixating animals already show that the perspective of a simple stimulus-response relation is becoming questionable (Vinje and Gallant, 2000; Yen et al., 2008). Data appear much more variable which cannot be simply related to stimulus features, but higher-order brain areas and attentional effects obviously play a role. Active vision includes self-initiated eye-movements and thus naturally involves a combination of internal and external driving forces. Active vision is fast: within the duration of a fixation (about 200 ms) visual input enters the system, visual information is processed and the next new eye movement is initiated. This requires fast processing and leaves an individual stage of the nervous system only very limited time for computation (about 10 ms, Thorpe et al., 1996). Thus, electro-physiological studies of active vision need to include the dynamics of processing as suggested by models of the visual system (Körner et al., 1999; Van Rullen et al., 1998b) which predict temporal coordination of neuronal activities. Recently, Maldonado et al. (2008) presented first evidence that spike synchrony is involved early in the processing in

the visual system. Ito et al. (2009) found evidence that a signal activating large populations of the visual system (local field potential) occur locked to saccade onset thereby providing an internal reference signal for the coordination of neuronal activity induced by visual input. It is highly likely that this signal is modulated along the scan-path or as a function of attention thus providing the ground for context-dependent neuronal processing.

References

J. R. Anderson. Social stimuli and social rewards in primate learning and cognition. *Behavioural Processes*, 42:159–175, 1998.

R. J. Andrew. Evolution of facial expression. *Science*, 142:1034–1041, Nov 1963.

I. Biedermann. Recognition-by-components: a theory of human image understanding. *Psychol Rev*, 94:115–147, 1987.

V. Bruce and A. Young. *In the eye of the beholder: the science of face perception.* Oxford, England; New York: Oxford University Press, 1998.

G. Buswell. *How people look at pictures: A study of the psychology of perception in art.* Univ. Chicago Press, Chicago, 1935.

M. S. Castelhano and J. M. Henderson. Initial scene representations facilitate eye movement guidance in visual search. *J Exp Psychol Hum Percept Perform*, 33:753–763, Aug 2007.

D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence archive*, 24(5):603 – 619, May 2002.

N. J. Emery. The eyes have it: the neuroethology, function and evolution of social gaze. *Neurosci Biobehav Rev*, 24:581–604, Aug 2000.

M. Fabre-Thorpe, G. Richard, and S. J. Thorpe. Rapid categorization of natural images by rhesus monkeys. *Neuroreport*, 9:303–308, Jan 1998.

K. Funkunaga and L. D. Hostetler. The estimation of the gradient of a density function, with applications to pattern recognition. *IEEE Trans. Information Theory*, pages 32–40, 1975.

D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.)*, 160:106–154, Jan 1962.

J. Ito, P. Maldonado, W. Singer, and S. Grün. First spikes in v1 are phase locked to saccade-related lfp oscillations. *submitted*, 2009.

L. Itti and C. Koch. A comparison of feature combination strategies for saliency-based visual attention systems. *Proc. SPIE Human Vision and Electronic Imaging IV (HVEI'99)*, 3644:473–482, 1999a.

L. Itti and C. Koch. A model of saliency-based visual attention for rapid scene analysis. *Nature*, 2:194–203, 2001.

L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20 (10):1254–1259, November 1998.

S. J. Judge, B. J. Richmond, and F. C. Chu. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res.*, 20:535–538, 1980.

M. A. Just and P. A. Carpenter. Eye fixations and cognitive processes. *Cognitive Psychology*, 8:441–480, 1967.

E. Körner, M.-O. Gewaltig, U. Körner, A. Richter, and T. Rodemann. A model of computation in neocortical architecture. *Neural Networks*, 12:989–1005, 1999.

G. Krieger, I. Rentschler, G. Hauske, K. Schill, and C. Zetzsche. Object and scene analysis by saccadic eye-movements: an investigation with higher-order statistics. *Spat Vis*, 13: 201–214, 2000.

S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.

N. H. Mackworth and A. Morandi. The gaze selects informative details within pictures. *Perception and Psychophysics*, 2:547–552, 1967.

P. Maldonado, C. Babul, W. Singer, E. Rodriguez, D. Berger, and S. Grün. Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. *J. Neurophysiol.*, 100:1523–1532, 2008.

A. Markov. An example of statistical study on text of eugeny onegin illustrating the linking of events to a chain. *illyustriruyuschij svyaz' ispytanij v cep', Izvestiya Akademii Nauk*, 6:153–162, 1913.

S. J. McKelvie. The role of eyes and mouth in the memory of a face. *The American Journal of Psychology*, 89:311–323, 1976.

M. Mendelson, M. M. Haith, and P. Goldman-Rakic. Face scanning and responsiveness to social cues in infant rhesus monkeys. . *Developmental Psychology*, 18:222–228, 1982.

R. Milanse, S. Gil, and T. Pun. Attentive mechanisms for dynamic and static scene analysis. *Optical Eng.*, 34(4):2428–2434, 1995.

D. Noton and L. Stark. Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vision Res.*, 11:929–942, Sep 1971a.

D. Noton and L. Stark. Scanpaths in eye movements during pattern perception. *Science*, 171:308–311, Jan 1971b.

S. A. Rosenfeld and G. W. Van Hoesen. Face recognition in the rhesus monkey. *Neuropsychologia*, 17:503–509, 1979.

A. Santella and D. DeCarlo. Robust clustering of eye movement recordings for quantification of visual interest. *Eye Tracking Research and Application (ETRA)*, pages 27–34, 2004.

S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381:520–522, 1996.

J. K. Tsotsos, S. M. Culhane, W. Y. K. Wai, Y. H. Lai, N. Davis, and F. Nuflo. Modeling visual attention via selective tuning. *Artifical Intelligence*, 78:507–545, 1995.

S. Ullman. Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb. Cortex*, 5:1–11, 1995.

R. Van Rullen, J. Gautrais, A. Delorme, and S. Thorpe. Face processing using one spike per neurone. *BioSystem*, 48:229–239, 1998.

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

J. M. Wolfe. Guided search 2.0: A revised model of visual search. *Psychonomic Bull. Rev.*, 1:202–238, 1994.

A. L. Yarbus. Eye movements and vision. *Plenum Press*, 1967.

S. C. Yen, J. Baker, J. P. Lachaux, and C. M. Gray. Responses of primary visual cortical neurons to natural movies in anesthetized cat. *BMC Neuroscience*, page 125, 2008.

Discussion

Correlated spiking activity in the primary visual cortex

The first issue analyzed in this thesis concerns the question of the temporal coordination of neuronal activity in the primary visual cortex. I presented correlation analyses from (1) massively parallel recorded multi-unit activities (MUAs) from anesthetized cats during spontaneous activity and visual "full-flash" stimulation (Chapter 1), and (2) simultaneously recorded single-unit activities from three monkeys during free exploration of natural scenes (Chapter 2). Both analyses showed that neurons of the primary visual cortex significantly correlate their spiking activity, with a high precision (<10 ms), beyond the effect attributable to co-variations of the firing rates. On the one hand, synchrony occurs spontaneously, between MUAs, which form distinct highly intra-correlated groups, and is also present during full-flash treatment. The occurrence of synchronization during spontaneous activity unveils an intrinsic property of the network embedded in the primary visual cortex. On the other hand, the locking of the synchronization to behavior gives insights into the functional mechanisms of visual information processing.

In the first study (Chapter 1, Berger et al., 2007), we could find several, distinct highly intra-correlated MUA groups (referred to as GICs) in the absence of stimuli and during full-flash treatment. The correlation effect is stronger during low intensity full-flash stimulation than during spontaneous activity. This can be explained as during this protocol all neurons within the visual field received the same external stimulus. This should therefore lead to an equal elevation of the intrinsic network. However, during high-intensity full-flash stimulation these synchronization effects are reduced in terms of numbers of correlated MUA pairs, as well as in number and size of the GICs. One possible explanation could be that if the stimulation is too high all neurons are strongly activated, and the high discharge rate may become indistinguishable from synchrony.
The occurrence of correlations between adjacent MUAs appears to originate in the columnar architecture of the primary visual cortex, where neurons with the same orientation preference are organized in a column (Hubel and Wiesel, 1962, 1963, 1968). As orientation preferences change smoothly moving tangentially in the primary visual cortex, orientation columns are repeated regularly and precisely over the cortex. Thus, neurons recorded from adjacent electrodes have in general the same or at least similar orientation preferences. If neurons with the same orientation tuning are synchronized, then there must be correlation between both, MUAs located in the same column and MUAs which lie in another column with the same orientation preference. This is reflected in the measurement of distance dependences of correlated MUA pairs during spontaneous activity (cf. Fig. [1].4), and during stimulation. The occurrence of correlation between MUA pairs appeared to be distance dependent. Nearby MUAs are correlated, but the synchronization decreases as

the distance increases. When the distance becomes comparable to the typical distance between identically oriented columns, the frequency of synchronization between neurons increases again. Indeed, the spatial scale of the GICs corresponds to that of orientation columns found by optical imaging, using voltage-sensitive dye, in anesthetized cat area 17, as observed e.g., by Kenet et al. (2003) (cf. Section [1].2.3, Fig. [1].3). The authors of this study showed that different orientation maps activate dynamically and systematically during spontaneous activity. We expect therefore the GICs to be dynamically correlated and alternatingly active. Preliminary analysis has shown that this is indeed the case for higher-order synchronized MUA groups, which match the GICs emerged in the pairwise analysis (Berger et al., 2009a, see Fig. 4a). Another important observation, adding credence to the possibility that GICs are a reflection of the underlying network organization in area 17, is that correlation maps were highly similar for all three different recording conditions (see Fig. 4b).

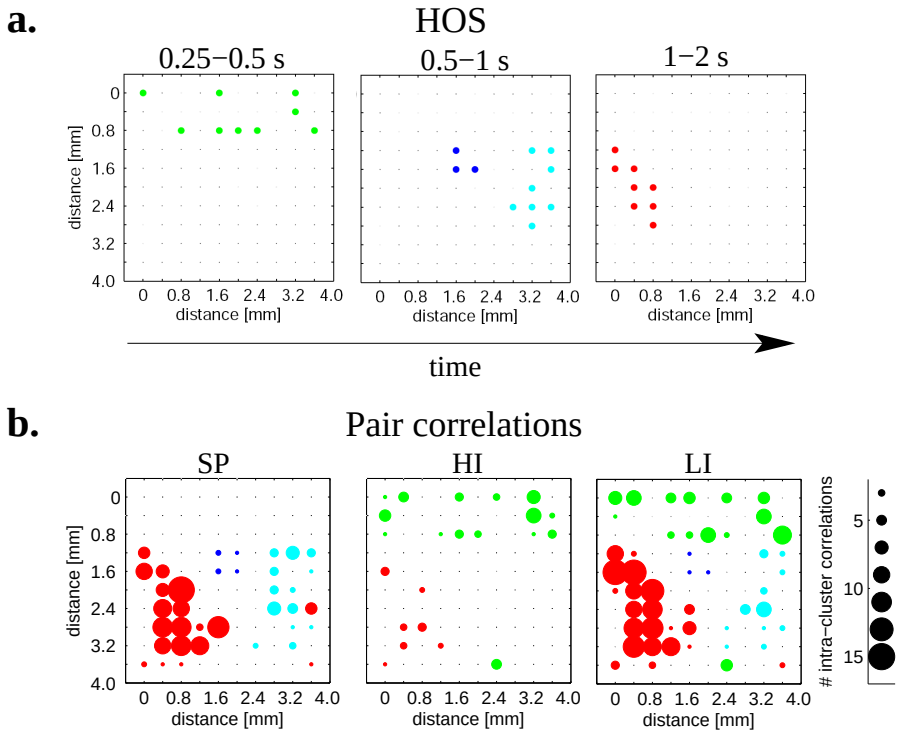

Figure 4: a. Higher-order synchrony (HOS) in cat visual cortex during low-intensity full-flash stimulation, after segmentation of data into quasi-stationary time segments. At different times, different sets of MUAs exhibit HOS. Verification by trial shuffling. b. Pairwise analysis (cf. Chapter 1 for details) for low intensity (LI), spontaneous (SP) and high intensity (HI) activity.

Our results are in agreement with a series of other studies. Ts'o et al. (1986) investigated

cross-correlations between cortical neurons in anesthetized cat in respect to long-range horizontal connections, i.e., distances of up to several mm. Their study suggests a high degree of inter-connectivity between cells that have the same orientation preference within - as well as between - columns, therefore strongly supporting the interpretation of our results. Das and Gilbert (1999) investigated short-range interactions in anesthetized cat area 17 and observed that neighboring cells in a range of up to 500 - 800 μm show spike correlations independently of the orientation preferences. However, we found correlations on a broader scale of distance. All but one GICs have spatial extents that exceed such a short-range interaction. An additional and important difference between their experiments and ours is the type of stimulus used. Here, we investigated spontaneous activity as well as activity under full-flash treatment, whereas Das and Gilbert (1999) used isolated optimally tuned bars in the receptive fields of the neurons. More comparable to our results is the study by Tsodyks et al. (1999) in anesthetized cat area 17/18 who found orientation maps under spontaneous activity that is related to the tuning of the individual cortical neurons using voltage-sensitive dye imaging. He concluded that the underlying functional architecture is associated to tuning properties. In Chapter 4 we applied a test statistics for identifying potential assembly members on the same data sets. The analysis resulted in an almost perfect detection of the MUAs detected by the pairwise analysis, confirming the results found by pairwise analysis (Berger et al., 2009b). The only undetected MUAs were not members of the GICs (see discussion below).
Spontaneous activity is believed to play an important role for generation of synchronous responses (Fries et al., 2001; Kenet et al., 2003; Blumenfeld et al., 2006; Singh et al., 2008), providing the basis for rapid synchronization during stimulus input. However, whether or not our results point towards a functional relevant role of synchronization, or solely reflecting the underlying network architecture would need to be systematically tested.

The second study of correlation analysis between neurons of the primary visual cortex presents an analysis of data from three awake and behaving monkeys (Chapter 2, Maldonado et al., 2008). The monkeys were allowed to freely explore natural scenes (presented as still images). In order to relate the correlations to the eye events of fixations and saccades, unitary events have been analyzed in a time resolved manner, aligned to the eye events. Moreover, since the actual stimulus changed with each saccade and the recorded neurons were different for each session, the average neuronal behavior was evaluated, allowing us to detect characteristic eye movement-related response patterns.
We could show that precise correlations were strongly locked to the onset of fixations during image presentation and occurred between 30 to 90 ms after fixation onset. This effect cannot be explained by the rate change, as the rate and the synchrony show dif-

ferent temporal dynamics, the discharge rate peak being at 90 ms after fixation onset and then decreasing to near baseline levels within 200 ms, whereas the synchrony episode started 30 ms after fixation onset and ended by the time the discharge rate had reached its maximum (cf. Section [2].3, Fig. [2].4). On the contrary, during the control period –when the monkeys scanned a blank screen– only small changes in firing rates occurred, and no excess synchrony was detected. These results suggest a functional relevance of synchronization between neurons.

With each onset, a new visual stimulus is conveyed to the brain, with a fixation period typically lasting 100 to 200 ms. Thus, feature extraction and object identification must be accomplished within approximately 200 ms. Support of this hypothesis is given by Fabre-Thorpe et al. (2001). In their study, psychophysical experiments involving with humans performing object identification led to similar conclusions about the processing speed of primary visual cortex. The authors claimed that in early stages of visual processing computations required for elementary scene analysis have to be achieved within a few tens of milliseconds. Based on the hypothesis that initial image processing occurs within the first tens of milliseconds after image presentation, Guyonneau et al. (2004) proposed that feature extraction, scene segmentation, and object recognition should be accomplished by exchanging only a few spikes per neurons. In this scenario, integration of firing rates across a neuronal population would not be feasible. Therefore, it has been proposed that information is encoded in the precise timing of action potentials.

Our results support this hypothesis. Comparison between the dynamics of rate modulation and synchrony showed that the rate change is much slower than the occurrence of excess synchrony, as it reached its maximum approximately 40 ms before the peak of the rate increase. Response latencies cover a range of 30 to 60 ms, which implies that potential synchronizing mechanisms operate already and preferentially on the very first spikes evoked by the newly fixated parts of the image.

A mechanism capable of achieving such rapid synchronization has recently been identified by Fries et al. (2001, 2007). They observed that neurons with related feature preferences are engaged in synchronous oscillations. This has the effect that the onset latencies of responses to stimuli driving coherently oscillating columns become synchronized, where neurons fire only during the depolarizing phase of the oscillation cycles (Engel et al., 2001; Fries et al., 2001, 2007). A recent analysis by Ito et al. (2009) on the same data set we have analyzed, showed that there is a strong beta-band oscillation in the local field potential. The oscillation is strongly locked to saccade onset and continues into the successive fixation period, and visually induced spikes are phase locked to this oscillation. The authors suggest that the oscillations serve as a corollary signal enabling precise timing of spikes in V1 and thereby providing a mechanism for synchronization.

Fries et al. (2001, 2007), and Engel et al. (2001) suggested that fast synchronization occurs between neurons with related features. In our study, the analysis was performed by averaging across all neurons; thus the link between orientation tuning and synchrony cannot be directly drawn. A challenging next step would be to test whether neurons with the same orientation preference are primarily involved in the synchronization we have found.
Another challenging question would be to test in how far the synchronous activity found during spontaneous activity resembles the activity during free viewing behaviour. Ito et al. (2009) suggested that the initiation of the saccade generates oscillations, which in turn synchronizes the neurons during vision. Synchrony during spontaneous activity may be an intrinsic phenomenom, solely reflecting the connectivity of V1, and not functionally involved in viewing. However, as argued above many studies suggest that spontaneous activity is important for the initiation of rapid synchronization. Moreover, it has been suggested that spontaneous activity reflects the internal state of the brain and might be influenced by context, attention or perceptual memories (Blumenfeld et al., 2006). When it comes to our results, it is worth noting that there are some important differences between the first and the second study presented. Indeed, the first analysis was performed on anesthetized cats, the second other on on awake and behaving on monkeys. Further experiments would be needed to complete the framework, e.g. by recording multi-electrode arrays in the absence of behavior and during free viewing.

In conclusion, our work suggests that synchrony is mechanism for information processing employed by the visual cortex, that may operate by fine tuning the timing of neurons. Correlations between MUA pairs appear to be well organized even in the absence of stimuli, particularly when distance dependent correlations are compared to orientation tuning in distinct columns. Since oriented stimuli were not used here, MUA correlation is not to be attributed to orientation tuning. Nevertheless, it can be surmised from these results that these correlations of neuron pairs are linked according to orientation tuning. Moreover, the visual cortex possesses a mechanism to complement rate code with a fast temporal coding strategy, adjusting the timing of individual spikes for fast initial image processing.

On the approach to identifying higher-order correlations

The second topic addressed in this thesis is the identification of higher-order synchrony in neuronal data (Chapters 3 and 4). The cell assembly hypothesis (Hebb, 1949) predicts higher-order correlations, i.e., involving large sets of neurons. However, there are two problems in verifying this hypothesis, causing lack of convincing evidence. Firstly, the

combinatorial explosion of the parameter space dimension is the major limitation of all available methods for analyzing more than two neurons. Even for pairwise approaches the task is hard, as the number of parameters grows quadratically with the population size. Secondly, restrictions of electrophysiological recording tools allowed only recently neuronal recordings of more than a few neurons. Abeles (1991) and Diesmann et al. (1999) suggest that the number of synchronized neurons are in the order of 100, with 80% of those neurons participating per activation. However, even with big arrays of electrodes, the chance to record all the neurons participating in an assembly within a set of only a few hundreds of recorded neurons is rather low, since the recorded neurons will be dispersed over a few cortical hyper-columns. Conversely, it is likely that most of the processes within the set are statistically independent. That is, the chance is high that most of the recorded neurons do not participate in correlations. To address these issues, we developed the following two test statistics, enabling selection and identification of correlated neurons: (1) conditional pattern complexity (t^{CPC}), which evaluates the number of neurons firing in coordination with a tested cell, and (2) conditional spike frequency (t^{CSF}), which evaluates the average pairwise correlation the tested neuron is involved in (Chapters 3 and 4, Grün et al., 2009; Berger et al., 2009b). Each one of the methods tests for a given neuron whether it whether a neuron is involved in coincidence spike events more often than expected by mere chance. The methods are able to rapidly and reliably detect neurons participating in assembly activity. A limitation is that the methods do not provide information about which neurons a given test neuron is correlated to, but only whether the neuron participates in correlations at all. Grouping and clustering of correlated neurons should be addressed by further steps of analyses case by case.
t^{CPC} and t^{CSF} were tested on networks of different types and size: 100 and 1000 neurons network generated by stochastic models, and 5000 neurons network generated by computer simulations. In the case of the stochastic model network, we tested the performance of our methods by varying the size and the number of the correlated assemblies, the rate of correlations, and the number of the simulated neurons. Our results showed that both methods are extremely powerful, showing a high detection rate and a low false positive (FP) rate. The latter is of the order of the applied significance level. The false negative (FN) rate was high only in the case where the coincidence rates or the percentage of neurons involved in the correlation were very low (for $N = 100$ and $N = 1000$, respectively). However, FN rates dropped drastically and reached 0% for larger values of at least one out of the two parameters (cf. Section [4].4.2, Fig. [4].3 and [4].4). Interestingly, t^{CSF} achieves a lower FN rate even for the worst case scenarios. As the method is thought to be a pre-processing tool, it is crucial to achieve a FN rate of 0%, i.e., not missing any correlated neurons. The FP rate, on the other hand, is not so critical, as FPs can

be excluded in subsequent steps of analysis. Over all, for reasonable assembly sizes and correlation rates the methods achieve perfect detection of all neurons potential assembly members and FP rates of 1-2 %, proving to be both effective and reliable tools, with t^{CSF} generally slightly outperforming t^{CPC}.

The data set generated by network simulations contained a synfire chain embedded in a network of 64000 excitatory and 16000 inhibitory neurons. The embedded chain consisted of 20 pools of 100 excitatory neurons, thus involving 2000 neurons. A set of 5000 neurons containing k randomly sampled synfire chain neurons was sampled from the network and analyzed, with k=200, 500, or 2000. As a control, the same analysis was performed on (1) the same neuronal network with the synfire chain not being activated, and (2) a network which did not contain any synfire chain. These data sets contained additional features not present in the stochastic model networks. For example, the firing rates were non-stationary in time, with a high oscillation frequency (cf. Section [4].4.3, Fig. [4].5). Our results clearly demonstrated that uniform spike shuffling is a strong generator of FPs, which was the case for the data with embedded assembly activity as well as for the control data. Obviously, the coherent rate oscillations in the data generate a high amount of synchronous events by chance, which, if ignored (as by uniform shuffling) would be interpreted as the existence of precise spike synchrony. Shuffling the spikes according to an estimator based on the population firing rate, the FP rate can be reduced to almost zero, if in addition the modulation of the shuffling probabilities as a function of time is kept high (i.e., low base line). Interestingly, the FN rate is close to zero for t^{CSF}, even for smaller numbers of neurons sampled from the embedded synfire chain. Although a sampling of 200 out of 2000 neurons sounds large as a percentage, the number of neurons sampled from an individual synfire chain group, i.e., synchronous at a given time, is, with an average of 10 neurons, rather small (200/20).

Having tested our approach in simulated data we applied both methods to the same neuronal parallel spike data from cat visual cortex, which was already analyzed for pair correlations in Chapter 1 (Berger et al., 2007) and discussed above. This resulted in an almost perfect detection of all the MUAs that were involved in significant pairwise analysis. The MUAs that were not detected were those involved in only a few pairwise correlations, and were not members of the pre-detected highly intra-correlated MUA groups (GICs). The number of involvements in pairwise correlations was clearly shown to correlate with the significance of the test statistics (t^{CPC} and t^{CSF}) (cf. Fig. [4].8). On the other hand, tests with the stochastic model showed that the larger the assembly size, i.e., the order of correlation, the better the detection by our test statistic. The results presented in Chapter 1 suggest that the identified clusters reflect higher-order correlations between MUAs

(see discussion above). Indeed, preliminary analysis of the same data based on accretion (Gerstein et al., 1978; Berger et al., 2009a) likewise revealed strong evidence of higher-order correlated groups of neurons. Furthermore, identical results were obtained on the reduced data set, i.e., after application of the test statistics, and on the complete one, without any loss of significant correlated neurons. Thus, data reduction to the relevant neurons leads to a substantial improvement in terms of computation time.
In conclusion, our methods successfully pick up the presence of large groups of correlated neurons. Although both test statistics showed an overall good performance, t^{CSF} generally performs better, achieving the lowest FN rates even in the worst case scenarios. This method allowed the detection of all synfire chain neurons, even if only a small number of neurons are sampled. Existing methods of correlation analysis (Grün et al., 2002a,b; Shimazaki et al., 2009; Amari, 2001), particularly higher-order correlation analysis of such high-dimensional data, are very demanding in terms of computation time and memory consumption, and often may not be applicable at all. On the contrary, the test statistics we presented proved to be highly efficient and able to easily digest recordings of several thousand of parallel processes. Our approaches are a valuable addition to previous methods that provide information on the *presence* of higher-order correlation in data, but do not identify the individual neurons involved (Staude et al., 2009). Moreover, they are designed for reducing the data set to the relevant neurons, which can be followed by a more detailed higher-order analysis (e.g., Shimazaki et al., 2009; Amari, 2001). Future work will address the next processing stages. These involve grouping the detected significantly correlated neurons into the different pools of the synfire chain. Due to the efficiency of the method one may also think of a time dependent application that may recover dynamical changes of the correlation structure.

From neuronal mechanisms to behaviour

In the last part of my thesis, I presented a detailed behavioral study of eye movements of monkeys freely exploring natural scenes (Chapter 5, Berger et al., 2009c). The eye movements were recorded at the same time as the activity of single neurons. The correlation analysis of the single-units, aligned to fixation onset, was already presented in Chapter 2 and discussed above. There, correlations were averaged across fixation onsets, irrespective of the image content or attended image object, as well as across neurons, and animals. Two reasons motivated us to investigate the eye movement behavior in details, in particular to relate the synchronous events to the different classes of fixations. First, it has been shown that attentional mechanisms modulate the oscillatory patterns of neuronal responses in V1, which in turn modulates synchronous discharges (see, e.g.,

Luck et al., 1997; Chawla et al., 1999; Brefczynski and DeYoe, 1999; Gandhi et al., 1999). Second, each fixation provides a novel stimulus to the brain, but averaging across fixations this distinctions are lost. However, during active vision the eye positions are voluntarily chosen by the subjects, which results in the absence of identical, repetitive trials. This in turn makes statistical analyses not feasible. To approach the question how the stimulus affects correlated activity during image exploration, we investigated the eye movements of the monkeys. By classifying fixations into different categories we created "trials", thus providing an approach that allows statistical analysis of electrophysiological data.

The investigation of the topology of fixation locations resulted in specific spatial patterns of fixation positions, which depend on the presented images, in two out of the three monkeys (cf. Fig. [5].3). The eye fixations of the third monkey, however, did not alter for different images, but were rather restricted to the lower left quadrant of the image frame. As the experiment did not assign any specific task (besides keeping the gaze within the image), we concluded that the third monkey did not learn to explore the images. Therefore, further analyses were restricted to the first two monkeys.

The fixation positions form clusters in space, which we interpreted as region-of-interests (ROIs). Similar viewing behavior was found in humans. These observe complex scenes by limiting their fixations to details-rich locations instead of distributing their attention evenly across the entire scene. An image, therefore, can be divided into informative and redundant regions (Mackworth and Morandi, 1967; Yarbus, 1967; Krieger et al., 2000). To gain insight into the mechanisms guiding the visual exploratory behavior of the monkeys, we extracted from the images all salient features, considered good candidates for attracting the eye focus of the monkeys. We found that in approximately 75% of the images the resulting saliency maps are good predictors for the fixation points, suggesting that the eye movements are mainly driven by low-level features. In the remaining images the salient regions were visited far less frequently than expected. However, we noticed that the latter images all contained faces of primates. When presented with those images, both monkeys focused primarily on the eyes, thereby avoiding the 'high' salient regions (cf. Fig. [5].4). This second strategy of directed vision might be attributable to top-down mechanism, where attention to meaningful object details determine the gaze distribution. Our results closely resemble those found in humans where the choice of fixation points were shown to be either driven by bottom-up (exogenous) or by top-down (endogenous) - internal or external factors (McKelvie, 1976).

A further analysis concentrated saccades performed between successive fixations by considering the sequence of eye fixations (cf. Fig. [5].5). The use of a Markov chain model revealed that the monkeys are primarily making saccades within a region of interest, i.e., *locally* exploring restricted areas in the scene before directing their attention to a new

region of interest (*global* exploration).

Summarizing, these results indicate that monkeys explore natural scenes by making directed eye movements to regions that provide them visual information, either in terms of saliency or in terms of meaningful items, such as eyes, in cases of images containing primates. Overall, characteristics in the fixation patterns in two out of three monkeys showed similarities to those in humans. Previous results have shown striking similarities between primates and humans on the rapid processing of natural scenes. However, here we analyzed for the first time the eye movements of monkeys during prolonged viewing of natural scenes. Our results suggest that the similarities in eye movement behavior between humans and monkeys applies even more generally.
Our work was made possible by combining detailed eye movements measurements and the absence of a pre-defined task (free viewing). Due to difficulties in analyzing electrophysiological data not containing repetitive, identical trials, experiments including active vision typically require subjects to fixate on a pre-defined position. The studies presented here provide new approaches of analyses and open new routes for experiments that aim to relate neuronal activities to natural behavior. It is now possible to relate the neuronal activity to the different features which attract the individual fixations. A challenging question that remains to be addressed involves the analysis of images of primates and images containing non-primate objects in order to contrast the corresponding neuronal activities, and relate these results to correlation analysis (Chapter 2). Another subject of future investigation is the comparison between the first fixation in a ROI to the subsequent ones, as well as to the last fixation in the ROI and to fixations on the background. Further insights into attentional mechanisms would be provided by the comparison of neuronal responses between the two monkeys attentively exploring the scenes and the third monkey, who did not explore the different scenes.

A final word

This thesis presented evidence of synchronous firing of primary visual cortex neurons suggesting synchrony as a mechanism employed by the cortex for information encoding and decoding. On the one hand, the presence of correlations during spontaneous activity reveals that synchrony is an intrinsic property of the network, supporting the hypothesis that correlation is linked to the orientation tuning of the neurons. On the other hand, neuronal correlations in awake monkeys during free viewing is temporally strongly locked to behavior, suggesting that synchronization provides a mechanism for fast image processing. Whether or not synchronized neurons share the same orientation preference remains to be tested systematically. This requires additional experiments with adequate stimuli.

Specifically testing in how far the correlation maps found during spontaneous activity are altered during vision compared to spontaneous activity is similarly required. To approach the question how the stimulus affects correlated activity in absence of repetitive trials, fixations were classified in different categories. The results showed that two out of three monkeys were actively exploring the presented images. As attention has been shown to affect neuronal behavior, comparison between the attentive monkeys and the non-attentive monkey would give further insights into the attentional mechanisms employed by the visual cortex. Moreover, the two monkeys actively exploring the scenes seem to distinguish and employ different viewing strategies for images containing primate faces and images containing non-primate objects. In the first image class, the attention of both monkeys was significantly drawn to meaningful objects - such as the eyes and mouth of the primates - thereby avoiding salient regions of the images. In the second class of images, the monkeys eye position have been shown to be significantly drawn to the image saliency. Comparing neuronal behavior in relation to these classes of images would be highly interesting, and may answer the question whether attention is a top-down driven (case of primate images) or bottom-up driven (for non-primate object cues) mechanism. The sequence of fixation reveals in both monkeys and for all images, that exploration is not random, but a high level strategy of first local exploration within a region-of-interest (ROI) is followed by a shift to a new ROI (global exploration). Therefore, fixations can be further divided according to the novelty of input, i.e., the first fixation within a ROI versus the subsequent ones. Analysis of neuronal behavior which takes into account the different fixation and image classes would be important for further understanding of the strategy employed by the visual cortex.

Another topic we addressed in this thesis concerned the higher-order correlations in neuronal data. Due to the high number of correlations found in the highly correlated MUA groups (cf. Chapter 1), we believe that MUAs are not only involved in pairwise correlations but also exhibit higher-order correlations. In this framework we provided two test statistics for detection and identification of assembly activity in massively parallel data sets (Chapters 3 and 4, respectively). Application on simulated data showed that these methods are able to select and detect the correlated neurons within large data sets. The methods do not provide information about assembly membership, which needs to be analyzed in a further study. However, they provide a fast way to select relevant neurons, and thereby significantly reduce a data set. This solves the issues of all existent methods for higher-order correlations, often not applicable due to the immense memory consumption or computation time. The test statistics we presented proved to be very efficient, fast, and reliable, and are able to analyze several thousands parallel processes.

Application of these test statistics to the huge data sets of cat visual cortex yield an

almost perfect detection of the MUAs previously identified by pairwise analyses. Only the MUAs that were not members of any GIC were not detected. As the methods performed better with increasing assembly size we can conclude that they are tuned to better identify higher-order synchronized neurons rather than neurons which participate only in pair correlations. Indeed, preliminary analysis of higher-order correlations corroborate this hypothesis, and suggests the presence of higher-order synchrony in the data (Berger et al., 2009a). Moreover, higher-order correlated groups are in agreement with the GICs found by pairwise analysis. The results on the reduced data set after application of the test statistics, and on the whole data set are identical. Therefore, the reduction of the data set does not lead to a loss of neurons, but rather helps to reduce computation time considerably in further analysis.

References

M. Abeles. Role of the cortical neuron: integrator or coincidence detector? *Isr. J. Med. Sci.*, 18:83–92, Jan 1982.

M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, 1st edition, 1991.

M. Abeles and G. L. Gerstein. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol.*, 60:909–924, Sep 1988.

E. D. Adrian. The impulses produced by sensory nerve endings: Part I. *J. Physiol. (Lond.)*, 61:49–72, Mar 1926.

A. M. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J. Neurophysiol.*, 61:900–917, 1989.

K. Albus. A quantitative study of the projection area of the central and the paracentral visual field in area 17 of the cat. ii. the spatial organization of the orientation domain. *Exp Brain Res*, 24(2):181–202, Dec 1975.

J. Allman, F. Miezin, and E. McGuinness. Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annu Rev Neurosci*, 8:407–430, 1985. doi: 10.1146/annurev.ne.08.030185.002203. URL `http://dx.doi.org/10.1146/annurev.ne.08.030185.002203`.

S. Amari. Information geometry on hierarchy of probability distributions. *IEEE Trans. Info. Theory*, 47:1701–1711, 2001.

J. R. Anderson. Social stimuli and social rewards in primate learning and cognition. *Behavioural Processes*, 42:159–175, 1998.

R. J. Andrew. Evolution of facial expression. *Science*, 142:1034–1041, Nov 1963.

J. R. Bartlett, R. W. Doty, B. B. Lee, and H. Sakakura. Influence of saccadic eye movements on geniculostriate excitability in normal monkeys. *Exp Brain Res*, 25:487–509, 1976.

D. Berger, D. Warren, R. Normann, A. Arieli, and S. Grün. Spatially organized spike correlation in cat visual cortex. *Neurocomput*, 70:2112–2116, 2007.

D. Berger, C. Borgelt, M. Diesmann, G. Gerstein, and S. Grün. Spatially organized higher-order spike synchrony in cat area 17. *8th Göttingen Meeting of the German Neuroscience Society*, pages T26–8B, 2009a.

D. Berger, C. Borgelt, A. Morrison, and S. Grün. Efficient identification of assembly neurons within massively parallel spike trains. *Computational Intelligence in Neuroscience*, page accepted, 2009b.

D. Berger, M. Nawrot, P. Maldonado, and S. Grün. Statistics of eye movements of monkeys freely viewing natural scenes. *to be submitted*, 2009c.

I. Biedermann. Recognition-by-components: a theory of human image understanding. *Psychol Rev*, 94:115–147, 1987.

M. Bilge, A. Bingle, K. N. Seneviratne, and D. Whitteridge. A map of the visual cortex in the cat. *J Physiol*, 191(2):116P–118P, Jul 1967.

G. G. Blasdel and G. Salama. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321(6070):579–585, 1986. doi: 10.1038/321579a0. URL `http://dx.doi.org/10.1038/321579a0`.

B. Blumenfeld, D. Bibitchkov, and T. M. Neural network model of the primary visual cortex: From functional architecture to lateral connectivity and back. *J. Comput Neurosci*, 20:219–241, 2006.

B. Bolstad. *Introduction to Bayesian Statistics, 2nd ed.* Wiley Interscience, Chichester, UK, 2007.

T. Bonhoeffer and A. Grinvald. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature*, 353(6343):429–431, Oct 1991. doi: 10.1038/353429a0. URL `http://dx.doi.org/10.1038/353429a0`.

J. A. Brefczynski and E. A. A. DeYoe. A physiological correlate of the 'spotlight' of visual attention. *Nat. Neurosci*, 2:370–374, 1999.

E. N. Brown, R. E. Kass, and P. P. Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci*, 7(5):456–461, May 2004. doi: 10.1038/nn1228. URL `http://dx.doi.org/10.1038/nn1228`.

V. Bruce and A. Young. *In the eye of the beholder: the science of face perception.* Oxford, England; New York: Ocford University Press, 1998.

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

P. Buisseret and L. Maffei. Extraocular proprioceptive projections to the visual cortex. *Exp Brain Res*, 28:421–425, 1977.

D. C. Burr, M. C. Morrone, and J. Ross. Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature*, 371:511–513, 1994.

G. Buswell. *How people look at pictures: A study of the psychology of perception in art.* Univ. Chicago Press, Chicago, 1935.

M. S. Castelhano and J. M. Henderson. Initial scene representations facilitate eye movement guidance in visual search. *J Exp Psychol Hum Percept Perform*, 33:753–763, Aug 2007.

M. Castelo-Branco, S. Neuenschwander, and W. Singer. Synchronization of visual responses between the cortex, lateral geniculate nucleus, and retina in the anesthetized cat. *J. Neurosci.*, 18:6395–6410, 1998.

D. Chawla, G. Rees, and K. J. Friston. The physiological basis of attentional modulation in extrastriate visual areas. *Nat. Neurosci.*, 2:671–676, 1999.

D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence archive*, 24(5):603 – 619, May 2002.

M. Corbetta, E. Akbudak, T. E. Conturo, A. Z. Snyder, J. M. Ollinger, H. A. Drury, M. R. Linenweber, S. E. Petersen, M. E. Raichle, D. C. Van Essen, G. L. Shulman, and D. C. Van Essen. A common network of functional areas for attention and eye movements. *Neuron*, 21:761–773, 1998.

J. Csicsvari, D. A. Henze, B. Jamieson, K. D. Harris, A. Sirota, P. Bartho, K. D. Wise, and G. Buzsaki. Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J. Neurophysiol.*, 90:1314–1323, Aug 2003.

A. R. Damasio. Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. *Cognition*, 33:25–62, Nov 1989.

P. M. Daniel and D. Whitteridge. The representation of the visual field on the cerebral cortex in monkeys. *J Physiol*, 159:203–221, Dec 1961.

A. Das and C. D. Gilbert. Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature*, 399:655–661, 1999.

A. Date, E. Bienenstock, and S. Geman. On the Temporal Resolution of Neural Activity (Technical report). *Providence, RI: Brown Univ., Division of Applied Mathematics,*, 1998.

J. E. Dayhoff and G. L. Gerstein. Favored patterns in spike trains. II. Application. *J. Neurophysiol.*, 49:1349–1363, Jun 1983a.

J. E. Dayhoff and G. L. Gerstein. Favored patterns in spike trains. I. Detection. *J. Neurophysiol.*, 49:1334–1348, Jun 1983b.

R. Desimone and S. J. Schein. Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *J. Neurophysiol.*, 57:835–868, Mar 1987.

M. Diesmann, M. O. Gewaltig, and A. Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402:529–533, Dec 1999.

R. Eckhorn. Neural mechanisms of scene segmentation: recordings from the visual cortex suggest basic circuits for linking field models. *IEEE Trans Neural Netw*, 10(3):464–479, 1999. doi: 10.1109/72.761705. URL `http://dx.doi.org/10.1109/72.761705`.

S. Eldawlatly, R. Jin, and K. G. Oweiss. Identifying functional connectivity in large-scale neural ensemble recordings: a multiscale data mining approach. *Neural Comput*, 21: 450–477, Feb 2009.

N. J. Emery. The eyes have it: the neuroethology, function and evolution of social gaze. *Neurosci Biobehav Rev*, 24:581–604, Aug 2000.

A. K. Engel, P. Fries, and W. Singer. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.*, 2:704–716, 2001.

S. Everling. Where do i look? From attention to action in the frontal eye field. *Neuron*, 56:417–419, Nov 2007.

M. Fabre-Thorpe, G. Richard, and S. J. Thorpe. Rapid categorization of natural images by rhesus monkeys. *Neuroreport*, 9:303–308, Jan 1998.

M. Fabre-Thorpe, A. Delorme, C. Marlot, and S. Thorpe. A limit to the speed of processing in ultra-rapid visual categorization of novel natural scenes. *J Cogn Neurosci*, 13:171–180, 2001.

S. Friedman-Hill, P. E. Maldonado, and C. M. Gray. Dynamics of striate cortical activity in the alert macaque: I. Incidence and stimulus-dependence of gamma-band neuronal oscillations. *Cereb. Cortex*, 10:1105–1116, 2000.

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

P. Fries, D. Nikolic, and W. Singer. The gamma cycle. *Trends Neurosci.*, 30:309–316, 2007.

K. Funkunaga and L. D. Hostetler. The estimation of the gradient of a density function, with applications to pattern recognition. *IEEE Trans. Information Theory*, pages 32–40, 1975.

S. P. Gandhi, D. J. Heeger, and G. M. Boynton. Spatial attention affects brain activity in human primary visual cortex. *Proc. Natl. Acad. Sci. USA*, 96:3314–3319, 1999.

A. P. Georgopoulos, J. F. Kalaska, R. Caminiti, and J. T. Massey. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.*, 2:1527–1537, Nov 1982.

G. L. Gerstein. Searching for significance in spatio-temporal firing patterns. *Acta Neurobiol Exp (Wars)*, 64:203–207, 2004.

G. L. Gerstein, D. H. Perkel, and K. N. Subramanian. Identification of functionally related neural assemblies. *Brain Res.*, 140:43–62, Jan 1978.

G. L. Gerstein, P. Bedenbaugh, and M. H. Aertsen. Neuronal assemblies. *IEEE Trans Biomed Eng*, 36(1):4–14, Jan 1989.

W. Gerstner and W. M. Kistler. Mathematical formulations of hebbian learning. *Biol Cybern*, 87(5-6):404–415, Dec 2002. doi: 10.1007/s00422-002-0353-y. URL `http://dx.doi.org/10.1007/s00422-002-0353-y`.

M. Gewaltig and M. Diesmann. Nest (neural simulation tool). *Scholarpedia*, 2:1430, 2007.

C. D. Gilbert and T. N. Wiesel. Clustered intrinsic connections in cat visual cortex. *J. Neurosci.*, 3:1116–1133, May 1983.

C. D. Gilbert and T. N. Wiesel. Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J. Neurosci.*, 9:2432–2442, Jul 1989.

D. Golomb, D. Kleinfeld, R. C. Reid, R. M. Shapley, and B. I. Shraiman. On temporal codes and the spatiotemporal response of neurons in the lateral geniculate nucleus. *J Neurophysiol*, 72(6):2990–3003, Dec 1994.

C. M. Gray, P. E. Maldonado, M. Wilson, and B. McNaughton. Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J. Neurosci. Methods*, 63:43–54, Dec 1995.

J. S. Griffith. On the stability of brain-like structures. *Biophys. J.*, 3:299–308, Jul 1963.

S. Grün. Data-driven significance estimation for precise spike correlation. *J. Neurophysiol.*, 101:1126–1140, Mar 2009.

S. Grün, M. Diesmann, F. Grammont, A. Riehle, and A. Aertsen. Detecting unitary events without discretization of time. *J. Neurosci. Methods*, 94:67–79, 1999.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Comput*, 14:43–80, 2002a.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput*, 14:81–119, 2002b.

S. Grün, A. Riehle, and M. Diesmann. Effect of cross-trial nonstationarity on joint-spike events. *Biol Cybern*, 88:335–351, 2003.

S. Grün, M. Abeles, and M. Diesmann. Impact of higher-order correlations on coincidence distributions of massively parallel data. *Lecture Notes in Computer Science*, 5286:96–114, 2008.

S. Grün, D. Berger, and C. Borgelt. Identification of neurons participating in cell assemblies. *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3493–3496, 2009.

B. Gulyas, G. A. Orban, J. Duysens, and H. Maes. The suppressive influence of moving textured backgrounds on responses of cat striate neurons to moving bars. *J Neurophysiol*, 57(6):1767–1791, Jun 1987.

R. Guyonneau, R. Vanrullen, and S. J. Thorpe. Temporal codes and sparse representations: a key to understanding rapid processing in the visual system. *J. Physiol. Paris*, 98:487–497, 2004.

K. D. Harris. Neural signatures of cell assembly organization. *Nat. Rev. Neurosci.*, 6: 399–407, May 2005.

N. Hatsopoulos, S. Geman, A. Amarasingham, and E. Bienenstock. At what time scale does the nervous system operate? *Neurocomp.*, 52–54:25–29, 2003.

G. Hayon, M. Abeles, and D. Lehmann. A model for representing the dynamics of a system of synfire chains. *J Comput Neurosci*, 18(1):41–53, 2005. doi: 10.1007/s10827-005-5479-1. URL `http://dx.doi.org/10.1007/s10827-005-5479-1`.

D. O. Hebb. *Organization of behavior. A neurophysiological theory.* John Wiley & Sons, New York, 1949.

J. M. Henderson and A. Hollingworth. High-level scene perception. *Annu. Rev. Psychol.*, 50:243–71, 1999.

J. J. Hopfield. Encoding for computation: recognizing brief dynamical patterns by exploiting effects of weak rhythms on action-potential timing. *Proc. Natl. Acad. Sci. U.S.A.*, 101:6255–6260, 2004.

D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat's striate cortex. *J Physiol*, 148:574–591, Oct 1959.

D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.)*, 160:106–154, Jan 1962.

D. H. Hubel and T. N. Wiesel. Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *J Neurophysiol*, 26:994–1002, Nov 1963.

D. H. Hubel and T. N. Wiesel. Receptive fields and functional architecture of monkey striate cortex. *J Physiol*, 195(1):215–243, Mar 1968.

M. Huebener, D. Shoham, A. Grinvald, and T. Bonhoeffer. Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.*, 17:9270–9284, 1997.

J. Ito, P. Maldonado, and S. Grün. Saccade-related lfp oscillations set the stage for processing visually evoked spikes. *Computational and Systems Neuroscience (COSYNE)*, Salt Lake City, UT, 2008.

J. Ito, P. Maldonado, W. Singer, and S. Grün. First spikes in v1 are phase locked to saccade-related lfp oscillations. *submitted*, 2009.

L. Itti and C. Koch. A comparison of feature combination strategies for saliency-based visual attention systems. *Proc. SPIE Human Vision and Electronic Imaging IV (HVEI'99)*, 3644:473–482, 1999a.

L. Itti and C. Koch. A model of saliency-based visual attention for rapid scene analysis. *Nature*, 2:194–203, 2001.

L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20 (10):1254–1259, November 1998.

H. E. Jones, K. L. Grieve, W. Wang, and A. M. Sillito. Surround suppression in primate V1. *J. Neurophysiol.*, 86:2011–2028, 2001.

S. J. Judge, B. J. Richmond, and F. C. Chu. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res.*, 20:535–538, 1980.

M. A. Just and P. A. Carpenter. Eye fixations and cognitive processes. *Cognitive Psychology*, 8:441–480, 1967.

E. R. Kandel, J. H. Schwartz, and T. M. Jessell. *Principles of Neural Science*. Fourth edition, 2000.

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

Z. F. Kisvarday, E. Toth, M. Rausch, and U. T. Eysel. Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cereb Cortex*, 7(7):605–618, 1997.

B. W. Knight. Dynamics of encoding in a population of neurons. *J. Gen. Physiol.*, 59: 734–766, Jun 1972.

C. Koch and S. Ullmann. Shifts in visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, 4:219–227, 1985.

A. Kohn and M. A. Smith. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.*, 25:3661–3673, Apr 2005.

E. Körner, M.-O. Gewaltig, U. Körner, A. Richter, and T. Rodemann. A model of computation in neocortical architecture. *Neural Networks*, 12:989–1005, 1999.

G. Krieger, I. Rentschler, G. Hauske, K. Schill, and C. Zetzsche. Object and scene analysis by saccadic eye-movements: an investigation with higher-order statistics. *Spat Vis*, 13: 201–214, 2000.

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

S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.

M. F. Land and M. Hayhoe. In what ways do eye movements contribute to everyday activities? *Vision Res*, 41(25-26):3559–3565, 2001.

S. J. Luck, L. Chelazzi, S. A. Hillyard, and R. Desimone. Neural mechanisms of spatial selective attention in areas v1, v2, and v4 of macaque visual cortex. *J. Neurophysiol.*, pages 24–42, 1997.

N. H. Mackworth and A. Morandi. The gaze selects informative details within pictures. *Perception and Psychophysics*, 2:547–552, 1967.

R. Malach, Y. Amir, M. Harel, and A. Grinvald. Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc Natl Acad Sci U S A*, 90(22):10469–10473, Nov 1993.

P. Maldonado, C. Babul, W. Singer, E. Rodriguez, D. Berger, and S. Grün. Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. *J. Neurophysiol.*, 100:1523–1532, 2008.

A. Markov. An example of statistical study on text of eugeny onegin illustrating the linking of events to a chain. *illyustriruyuschij svyaz' ispytanij v cep', Izvestiya Akademii Nauk*, 6:153–162, 1913.

D. Marr. Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci*, 262(841):23–81, Jul 1971.

J. H. Maunsell and W. T. Newsome. Visual processing in monkey extrastriate cortex. *Annu Rev Neurosci*, 10:363–401, 1987. doi: 10.1146/annurev.ne.10.030187.002051. URL `http://dx.doi.org/10.1146/annurev.ne.10.030187.002051`.

C. J. McAdams and J. H. Maunsell. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.*, 19:431–441, 1999.

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

S. J. McKelvie. The role of eyes and mouth in the memory of a face. *The American Journal of Psychology*, 89:311–323, 1976.

M. Meister, L. Lagnado, and D. A. Baylor. Concerted signaling by retinal ganglion cells. *Science*, 270:1207–1210, 1995.

M. Mendelson, M. M. Haith, and P. Goldman-Rakic. Face scanning and responsiveness to social cues in infant rhesus monkeys. . *Developmental Psychology*, 18:222–228, 1982.

M. M. Mesulam. From sensation to cognition. *Brain*, 121 (Pt 6):1013–1052, Jun 1998.

R. Miller. Neural assemblies and laminar interactions in the cerebral cortex. *Biol Cybern*, 75(3):253–261, Sep 1996.

A. Morrison, S. Straube, H. E. Plesser, and M. Diesmann. Exact subthreshold integration with continuous spike times in discrete-time neural network simulations. *Neural Comput*, 19:47–79, Jan 2007.

M. P. Nawrot, C. Boucsein, V. Rodriguez Molina, A. Riehle, A. Aertsen, and S. Rotter. Measurement of variability dynamics in cortical spike trains. *J. Neurosci. Methods*, 169: 374–390, 2008.

U. Neisser. Cognitive psychology. *Appleton, New York*, 1967.

J. I. Nelson and B. J. Frost. Orientation-selective inhibition from beyond the classic visual receptive field. *Brain Res.*, 139:359–365, Jan 1978.

J. I. Nelson, P. A. Salin, M. H. Munk, M. Arzi, and J. Bullier. Spatial and temporal coherence in cortico-cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Vis. Neurosci.*, 9:21–37, 1992.

S. Neuenschwander and W. Singer. Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. *Nature*, 379:728–732, 1996.

C. T. Nordhausen, E. M. Maynard, and R. A. Normann. Single unit recording capabilities of a 100 microelectrode array. *Brain Res.*, 726:129–140, Jul 1996.

D. Noton and L. Stark. Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vision Res.*, 11:929–942, Sep 1971a.

D. Noton and L. Stark. Scanpaths in eye movements during pattern perception. *Science*, 171:308–311, Jan 1971b.

A. Oliva, A. Torralba, M. Castelhano, and J. Henderson. Top-down control of visual attention in object detection. In *Proc. International Conference on Image Processing ICIP 2003*, volume 1, pages I–253–6 vol.1, 2003. doi: 10.1109/ICIP.2003.1246946.

D. Parkhurst, K. Law, and E. Niebur. Modeling the role of salience in the allocation of overt visual attention. *Vision Res.*, 42:107–123, Jan 2002.

D. J. Parkhurst and E. Niebur. Scene content selected by active vision. *Spat Vis*, 16: 125–154, 2003.

A. Pazienti, M. Diesmann, and S. Grün. Bounds of the ability to destroy precise coincidences by spike dithering. *Advances in Brain, Vision, and Artificial Intelligence (Lecture Notes in Computer Science). Berlin: Springer Berlin/Heidelberg*, 4729:428–437, 2007.

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

D. H. Perkel, G. L. Gerstein, and G. P. Moore. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys. J.*, 7:419–440, 1967.

L. L. Perkel. Medical treatment of peptic ulcer. *J Med Soc N J*, 65(6):242–248, Jun 1968.

G. Pipa and S. Grün. Non-parametric significance estimation of joint-spike events by shuffling and resampling. *Neurocomputing*, 52-54:31–37, June 2003.

G. Pipa, C. van Vreeswijk, and S. Grün. Auto-structure of spike-trains matters for testing on synchronous activity. *(Abstract II-38). Computational and Systems Neuroscience (COSYNE)*, Salt Lake City, UT, 2008.

K. P. Purpura, S. F. Kalik, and N. D. Schiff. Analysis of perisaccadic field potentials in the occipitotemporal pathway during active vision. *J. Neurophysiol.*, 90:3455–3478, 2003.

S. Ramon y Cajal. *Estructura del asta de ammon y fascia dentata*, volume 22. 1893.

A. Riehle, S. Grün, M. Diesmann, and A. Aertsen. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278:1950–1953, 1997.

F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code.* Cambridge, MA: The MIT Press, 1997.

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

S. A. Rosenfeld and G. W. Van Hoesen. Face recognition in the rhesus monkey. *Neuropsychologia*, 17:503–509, 1979.

Y. Sakurai and S. Takahashi. Dynamic synchrony of firing in the monkey prefrontal cortex during working-memory tasks. *J. Neurosci.*, 26:10141–10153, Oct 2006.

A. Santella and D. DeCarlo. Robust clustering of eye movement recordings for quantification of visual interest. *Eye Tracking Research and Application (ETRA)*, pages 27–34, 2004.

K. E. Schmidt, D. S. Kim, W. Singer, T. Bonhoeffer, and S. Löwel. Functional specificity of long-range intrinsic and interhemispheric connections in the visual cortex of strabismic cats. *J Neurosci*, 17(14):5480–5492, Jul 1997.

G. Schneider and S. Grün. Analysis of higher-order correlations in multiple parallel processes. *Neurocomp.*, 52–54:771–777, 2003.

E. Schneidman, M. J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007–1012, Apr 2006.

J. M. Schoffelen, R. Oostenveld, and P. Fries. Neuronal coherence as a mechanism of effective corticospinal interaction. *Science*, 308:111–113, 2005.

S. Schrader, S. Grün, M. Diesmann, and G. L. Gerstein. Detecting synfire chain activity using massively parallel spike train recording. *J. Neurophysiol.*, 100:2165–2176, Oct 2008.

M. N. Shadlen and W. T. Newsome. Is there a signal in the noise? *Curr. Opin. Neurobiol.*, 5:248–250, Apr 1995.

M. N. Shadlen and W. T. Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.*, 18: 3870–3896, May 1998.

G. L. Shaw, D. J. Silverman, and J. C. Pearson. Model of cortical organization embodying a basis for a theory of information processing and memory recall. *Proc Natl Acad Sci U S A*, 82(8):2364–2368, Apr 1985.

C. S. Sherrington. Observations on the scratch-reflex in the spinal dog. *J. Physiol. (Lond.)*, 34:1–50, Mar 1906.

H. Shimazaki, S. Amari, E. Brown, and S. Grün. State-space analysis on time-varying correlations in parallel spike sequences. *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2009.

J. Shlens, G. D. Field, J. L. Gauthier, M. I. Grivich, D. Petrusca, A. Sher, A. M. Litke, and E. J. Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *J. Neurosci.*, 26:8254–8266, Aug 2006.

T. Shmiel, R. Drori, O. Shmiel, Y. Ben-Shaul, Z. Nadasdy, M. Shemesh, M. Teicher, and M. Abeles. Temporally precise cortical firing patterns are associated with distinct action segments. *J. Neurophysiol.*, 96:2645–2652, 2006.

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

W. Singer. Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24:49–65, 1999.

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

W. Singer, A. K. Engel, A. K. Kreiter, M. H. J. Munk, S. Neuenschwander, and P. R. Roelfsema. Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences*, 1(7):252–261, 1997.

G. Singh, F. Memoli, T. Ishkhanov, G. Sapiro, G. Carlsson, and D. L. Ringach. Topological analysis of population activity in visual cortex. *J Vis*, 8:1–18, 2008.

B. Staude, S. Rotter, and S. Grün. Detecting the existence of higher-order correlations in multiple single-unit spike trains. *Soc. Neurosci. Abstr. 103.9/AAA18, Washington, DC,*, 2007.

B. Staude, S. Rotter, and S. Grün. Cubic: cumulant based inference of higher-order correlations in massively parallel spike trains. *submitted*, 2009.

P. N. Steinmetz, A. Roy, P. J. Fitzgerald, S. S. Hsiao, K. O. Johnson, and E. Niebur. Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*, 404:187–190, 2000.

G. J. Stephens, S. Neuenschwander, J. S. George, W. Singer, and G. T. Kenyon. See globally, spike locally: oscillations in a retinal model encode large visual features. *Biol Cybern*, 95:327–348, 2006.

Y. Sugita. Grouping of image fragments in primary visual cortex. *Nature*, 401:269–272, 1999.

S. A. Talbot and W. A. Marshall. P hysiological studies on neural mechanisms of visual localisation and discrimination. *Am. J. Opthalmol.*, 24:1255–1264, 1941.

I. V. Tetko and A. E. Villa. A pattern grouping algorithm for analysis of spatiotemporal patterns in neuronal spike trains. 2. Application to simultaneous single unit recordings. *J. Neurosci. Methods*, 105:15–24, Jan 2001.

T. Tetzlaff, A. Morrison, M. Timme, and M. Diesmann. Heterogeneity breaks global synchrony in large networks. *Proc 30th Göttingen Neurobiology Conference*, page 206B, 2005.

S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381:520–522, 1996.

H. Tiitinen, J. Sinkkonen, K. Reinikainen, K. Alho, J. Lavikainen, and R. Naatanen. Selective attention enhances the auditory 40-Hz transient response in humans. *Nature*, 364:59–60, 1993.

F. Tong. Primary visual cortex and visual awareness. *Nature Neuroscience*, 4, 2003.

Y. Trotter and S. Celebrini. Gaze direction controls response gain in primary visual-cortex neurons. *Nature*, 398:239–242, 1999.

D. Y. Ts'o, C. D. Gilbert, and T. N. Wiesel. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.*, 6:1160–1170, 1986.

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

K. A. Turano, D. R. Geruschat, and F. H. Baker. Oculomotor strategies for the direction of gaze tested with a real-world activity. *Vision Res*, 43(3):333–346, Feb 2003.

R. J. Tusa, L. A. Palmer, and A. C. Rosenquist. The retinotopic organization of area 17 (striate cortex) in the cat. *J Comp Neurol*, 177(2):213–235, Jan 1978. doi: 10.1002/cne.901770204. URL `http://dx.doi.org/10.1002/cne.901770204`.

S. Ullman. Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb. Cortex*, 5:1–11, 1995.

W. M. Usrey and R. C. Reid. Synchronous activity in the visual system. *Annu. Rev. Physiol.*, 61:435–456, 1999.

E. Vaadia, I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A. Aertsen. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373(6514):515–518, Feb 1995. doi: 10.1038/373515a0. URL `http://dx.doi.org/10.1038/373515a0`.

J. L. van Hemmen and W. Senn. Hebb in perspective. *Biol Cybern*, 87:317–318, Dec 2002.

R. Van Rullen, J. Gautrais, A. Delorme, and S. Thorpe. Face processing using one spike per neurone. *BioSystem*, 48:229–239, 1998a.

R. Van Rullen, J. Gautrais, A. Delorme, and S. Thorpe. Face processing using one spike per neurone. *BioSystem*, 48:229–239, 1998b.

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

M. Volgushev, M. Chistiakova, and W. Singer. Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. *Neuroscience*, 83:15–25, Mar 1998.

D. J. Warren, E. Fernandez, and R. A. Normann. High-resolution two-dimensional spatial mapping of cat striate cortex using a 100-microelectrode array. *Neuroscience*, 105:19–31, 2001.

D. J. Warren, A. Koulakov, and R. A. Normann. Spatiotemporal encoding of a bar's direction of motion by neural ensembles in cat primary visual cortex. *Ann Biomed Eng*, 32:1265–1275, 2004.

W. A. Wickelgren. Webs, cell assemblies, and chunking in neural nets: introduction. *Can J Exp Psychol*, 53(1):118–131, Mar 1999.

J. M. Wolfe. Guided search 2.0: A revised model of visual search. *Psychonomic Bull. Rev.*, 1:202–238, 1994.

A. L. Yarbus. Eye movements and vision. *Plenum Press*, 1967.

S. C. Yen, J. Baker, J. P. Lachaux, and C. M. Gray. Responses of primary visual cortical neurons to natural movies in anesthetized cat. *BMC Neuroscience*, page 125, 2008.

S. Zeki. Colour coding in the cerebral cortex: the reaction of cells in monkey visual cortex to wavelengths and colours. *Neuroscience*, 9(4):741–765, Aug 1983.

Zusammenfassung

Die meisten Neurone des V1 weisen eine topologische Anordnung in Kolumnen auf, die bei optimaler Reizung ihre neuronale Aktivität erhöhen. Enkodierungsprozesse komplexer Stimuli lassen sich, vom statistischen Standpunkt aus gesehen, nicht abschließend durch Frequenzerhöhungen der Aktionspotentiale erklären. Moderne Theorien der primären Sehbahn beinhalten räumliche sowie zeitliche Aspekte zur Informationsverarbeitung. Insbesondere die Hypothese der Zellassemblies postuliert dynamisch interagierende Gruppen von Neuronen als Grundbaustein kortikaler Informationsverarbeitung. Während simultane Ableitungen von vielen Neuronen immer häufiger Anwendung finden, ist der fehlende experimentelle Beweis für die zeitliche Kodierungshypothese den fehlenden adäquaten Analysemethoden zuzuschreiben.
Die vorliegenden Dissertationsschrift untersucht die grundlegenden Mechanismen neuronaler Aktivität im visuellen Kortex während des Sehvorgangs. In diesem Kontext werden neue Strategien vorgestellt, die die Analyse einer großen Anzahl simultan abgeleiteter Neuronen ermöglicht.

Im ersten Teil dieser Doktorarbeit präsentiere ich Ergebnisse einer paarweisen Korrelationsanalyse von simultan abgeleiteten V1 Neuronen von Affen während der Exploration von natürlichen Szenen. Desweiteren wird die Korrelationsstruktur von multi-unit Aktivitäten des primären visuellen Kortex anästhesierter Katzen analysiert. In diesem Zusammenhang beantworte ich die folgenden Fragen: In welcher Beziehung steht die kortikale Karte von Orientierungsselektivitäten zu der zeitlichen Kodierungshypothese? Wie ist die Beziehung zwischen 1. korrelierter neuronaler Aktivität, Distanz und räumlicher Anordnung auf dem Ableitgitter und 2. korrelierter neuronaler Aktivität und Sehverhalten ?

Eine paarweise Korrelationsanalyse ist nicht ausreichend, um Assemblies herauszufiltern. Existierende Methoden für die Analyse von Korrelationen höherer Ordnung sind jedoch aufgrund ihres exponentiell wachsenden Aufwandes mit steigender Neuronenanzahl nicht anwendbar. Aus diesem Grund entwickelte ich zwei neue Methoden zur Identifikation von Assemblyneuronen. Beide Methoden detektieren schnell und zuverlässig Neurone, die an Korrelationen beteiligt sind.

Die vorliegende Arbeit bietet neue Einblicke in die Kodierungsstrategien von Neuronen des primären visuellen Kortex. Die gezeigte Kopplung zwischen Synchronisation und Fixationsbeginn, sowie die räumliche Anordnung synchronisierter Neurone, welche mit der von Orientierungskolumnen übereinstimmt, untermauern die Hypothese, dass Neurone visuelle Informationen über zeitpräzise Aktionspotentiale kodieren. Desweiteren schafft diese Arbeit eine Grundlage für weitere Korrelationsanalysen mit Hilfe von zwei hoch

effizienten Methoden für die Identifizierung von Synchronisationen höherer Ordnungen. Die Augenbewegungsanalyse stellt neue Herangehensweisen an Daten von explorierenden Tieren vor, und eröffnet somit neue Wege um neuronale Aktivitäten mit dem Verhalten in Verbindung zu bringen.

Acknowledgments

I would like to thank Prof. Nawrot for being the supervisor of my PhD and for having given me the opportunity to conduct this thesis at this Institute. I am grateful for the help and support he has given me throughout this work, and especially during the final preparation of the thesis.

I am deeply indebted to Sonja Grün. I thank her for all the encouragements and support she has given me. She introduced me into science, especially into the field of neuroscience, and supervised and guided my work. I am very grateful to have had her as my supervisor.

I would like to thank the collaborators I have had the pleasure to work with: Pedro Maldonado, Christian Borgelt, George Gerstein, and Amos Arieli. I thank them for the fruitful discussions, for data provisions, and for making science so much fun.

I also want to thank the members of the research groups I had the pleasure to work with in Berlin, and in Tokyo. In particular I thank Benjamin, Michael, Ina, and Hideaki for all the discussions and help throughout these years which were very important for my work.

My deep acknowledgment goes to Antonio. He has not only supported me troughout these years, but has also accompanied me during the final preparation. I want to thank him for his constant encouragement, and believe in me.

Finally I want to thank my family and friends for all their support and encouragement throughout these years. Especially, I thank my mom and Leila for their incredible help during the final preparation. My thanks also go to Rob.