Bachelor's Thesis

Probabilistic Encoding and Decoding of Neural Activity in Rat Primary Visual Cortex

October 2002

Holger Carl Bringmann
Cognitive Science, University of Osnabrück
Contact: hb@lambdabase.de
Supervisor: Dr. Christian Eurich, University of Bremen

Abstract: According to the information processing paradigm in the Cognitive Sciences, one of the nervous system's most important functions is to encode information about the environment. Understanding the *neural code* means understanding the relationship between brain states and real events in the outer world.

In this work, a probabilistic framework for encoding (i.e., mapping events to neural response) and decoding (i.e., reconstruction of events from neural responses) of neural activity is presented and applied to theoretical and empirical data. Reconstruction is based on a Bayesian filter method, allowing the propagation of multi modal densities and the use of non-linear stimulus dynamics.

The model is tested against empirical data recorded from rat primary visual cortex. Cells in this area are reported to be orientation-selective and their neural response characteristics can be learned by presenting oriented gratings to the rat's eye. The learned firing model is then used to reconstruct an unknown *random walk* stimulus from neural activity.

Contents

1 Introduction

According to the information processing paradigm in the Cognitive Sciences, one of the nervous system's most important functions is to encode information about the environment. It is one of the main goals of computational neuro science to explore how brain states come to represent and are altered by information present in the environment [35].

In this work, a probabilistic framework for neural coding and decoding of visual stimuli is applied to and tested against experimental data recorded from rat primary visual cortex.

1.1 The Neural Code

Understanding the *neural code* means understanding the relationship between brain states and real events in the outer world. For example, how does the brain state change when sensory signals reach the central nervous system, motor actions are performed or changes of internal states (e.g. attention) occur? The first problem in quantifying brain states with respect to a certain event is to find those entities of the brain that play a significant role in encoding it.

In the discourse of research on the neural code, different opinions have given raise to an ongoing intense discussion. Existent main stream theories claim [5, 21, 17] that:

- a part of the neural code can be found in *spike trains* of neurons, that is, lists of discrete times where the neuron has produced an action potential. Further information could be coded by complex dendritic processes or metabolic processes in attached glia cells, both generally going along without the production of spikes.

- different spike trains are generated for repeated presentations of the same event. This means the neural code appears to be stochastic. Although a complex neural system could still be deterministic in its nature, only a conditional probability of a neural response, given a certain event, can be measured over multiple trials, rather than a one-to-one mapping between event and response. In section 3.1, potential sources of stochasticity are discussed.

- information is usually encoded by neural populations (*population coding*) rather than single cells.

However, it is still topic of debate [5] whether or not

- correlations of spikes within a spike train carry any relevant information in terms of representing an event. While the existence of refractory and bursting behaviour of neurons shows that spikes are not generated independently of each other, it is not clear whether or not these characteristics carry information necessary for decoding [21, 5].

- correlations of neurons within a population carry any relevant information in terms of representing an event. Due to synchronisation effects, neurons of a population can be highly correlated. However, whether inter cell correlation has a positive or negative impact on coding accuracy is case-dependent and whether these effects are actually used by the brain, is questionable [4, 1].

- firing rates (*rate code*) or exact spike times (*temporal code*) should be considered. Regarding the rate code of a single neuron, precise temporal information about the spike times is lost in favour of a firing rate that is expressed as the number of spikes per time interval. However, the *population rate code* allows a more precise temporal analysis than the rate code of a single cell [21].

In this work, we assume a probabilistic, temporal population code with independent spikes and independent neurons. Although action potentials can vary in duration, amplitude and shape, they are treated as identical stereotyped events. However, the neural model (i.e., a theoretical framework describing the relationship between stimuli and neural response) as derived in chapter 3 is ready to be extended with respect to the named assumptions of independency.

1.2 Neural Recording

Empirical results form the fundament of theoretical modelling. To collect data about brainbrain function, electrophysiology is one of the most successful methods. Electrophysiological recordings measure electric fields induced by synaptic currents and potentials of cell membranes. These recordings are fast enough to cope with the temporal resolution of spiking neurons. Spatial accuracy can be enhanced greatly when using invasive methods. Non-invasive and invasive techniques have been used in numberless studies since the last century [35].

In this work, data have been collected from electrophysiological multi-cell recordings in the primary visual cortex (V1) of the rat. Data have been recorded at the Institute for Brain Research, University of Bremen by W. Freiwald and his research team.

Neurons in the striate cortex are known to be orientation-selective [34, 3, 24, 12]. To trigger a response in V1, oriented gratings are presented using a view space filling monitor picture. From voltage curves recorded by an array of invasive electrodes, spike times are extracted and assigned to cells.

Two kinds of stimulus presentations are applied: First, single stimuli in a row are presented in a controlled way over multiple trials. The neural response characteristics elicited from these presentations are used to instantiate the theoretical neural model with concrete values.

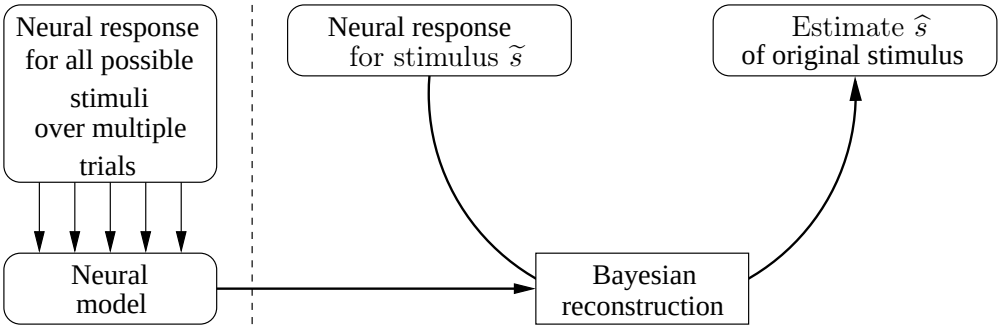

Figure 1.1: Overview of the coding and decoding process. The neural model for coding is generated offline from a controlled set of stimulus presentations. As soon as the neural model is given, decoding can be done online for dynamic stimuli.

Second, a *random walk* stimulus (see section 2.2 for details) is presented once. Recordings from this stimulus trajectory will be used to test the model as described in the next section.

1.3 Reconstruction as Model Test

Given the response characteristics and a neural model[1], the next step is to analyse the results. To check if the recorded cells have encoded a stimulus in a way predicted by our neural model, a straightforward approach is to try to reconstruct an unknown stimulus from arbitrary neural activity.

Interpreting neural population activity by reconstruction has been done in several ways; these can be divided into two classes [37]: basis function methods (i.e. population vector coding, optimal linear estimation and template matching) and Bayesian methods (with and without continuity constraint).

A basis function method uses a linear combination of fixed, but arbitrariliy chosen template functions $\phi_i(\widetilde{s})$ over stimulus space with the corresponding coefficients c_i proportional to the neural response (i.e. the firing rate or the number of spikes). The reconstructed stimulus $\widehat{s}$ is then chosen to be the peak of the linear sum:

$$\widehat{s} = \arg\max_{\widetilde{s}} \sum_i c_i \phi_i(\widetilde{s}).$$

Although these methods are not robust, they have been applied successfully in some cases [25, 33, 5]. Here, we will not further address this issue.

A second approach, taking into account the probabilistic nature of the neural response, is to use Bayesian inference for reconstruction. Bayes' rule directly addresses the relationship between the probability of a response, given a stimulus (which is known from experimental data) and the probability of a stimulus, given

[1] We refer to *neural model* as a mathematical framework to predict neural responses on the basis of simple assumptions about neural spiking behaviour.

a response:

$$p(\text{stimulus}|\text{response}) \propto p(\text{response}|\text{stimulus}) \cdot p(\text{stimulus}).$$

All that can be inferred about the stimulus is present in $p(\text{stimulus}|\text{response})$ and $\widehat{s}$ can be reconstructed from this distribution by estimation. Introducing knowledge about the stimulus dynamics (i.e. the temporal propagation of the stimulus) can increase the performance of this approach further. In this context, the role of the prior distribution $p(\text{stimulus})$[2], will be discussed further in the theoretical chapters (section 4.2).

This so-called *Bayesian Filtering* has been used successfully in recent theoretical and experimental research [37, 16, 2, 11] and will be used (in a generalized version) in this work as well.

1.4 Structure and Scope of this Work

Summarizing, we will use theoretical considerations and empirical data to create a neural model that is capable of encoding stimuli in terms of a neural response as described. To inspect the correctness of this model, we derive an 'inverse' function of the encoding process via Bayesian reconstruction and try to estimate an unknown original stimulus, given an arbitrary neural activity recorded from rat primary visual cortex. Once the response characteristics have been measured and the neural model has been set up, decoding works online for dynamic stimuli (see figure 1.1).

In the subsequent chapter, an overview will be given from what is known about the rat visual cortex. Moreover the experimental setup of data-recording will be explained. Chapters 3 and 4 form the theoretical background of this work, deriving a probabilistic framework for neural coding and decoding. The feasibility of this framework is shown in a simulation of reconstruction, using artificial data. Finally, in chapter 5, the experimental data is analysed. Assumptions made in the theoretical chapters are checked to hold true for the data and an attempt is made to reconstruct the presented stimulus trajectory from recorded data.

[2] the unconditional probability of a certain stimulus to occur.

2 Neural Recording from Visual Cortex

While the visual system of larger animals like cats or monkeys has been thoroughly studied, only few publications considering the rat can be found. Concerning the electrophysiology and response properties of the primary visual cortex in rats, the comprehensive work of Girman et al. [12] has to be named at first instance. The authors used single cell recordings to quantify receptive field properties for various stimulus dimensions (orientation, temporal and spatial frequency, contrast). In another work, multicell recordings have been used to investigate the dynamics of the horizontal functional connectivity [31]. The most important results that are relevant for this work are outlined in the following section.

The sections 2.2 to 2.4 describe how data used in this work have been recorded at the Institute for Brain Research, University of Bremen by W. Freiwald and his research team. Among others, research objectives were to reconstruct and quantify the coding accuracy of stimuli's representations in rat primary visual cortex [26, 27, 10]. In this work we will focus on the reconstruction of the angle of oriented gratings that are presented to the rat.

2.1 Response Properties of Rat Visual Cortex

To quantify the behaviour of single neurons, Girman et al. presented moving oriented gratings on a computer monitor (similar to the monitor pictured in figure 2.1) to the rat's eye (or, more precisely, to the cells' receptives fields). These gratings were altered systematically in different dimensions (i.e. their orientation, while the direction of movement was always orthogonal to the gratings) and the cell's responses were recorded with electrophysiological setup. These recordings resulted in *tuning curves* that relate the average spike count per time interval (i.e. the *firing rate*) to a range of different stimuli[1].

The primary visual cortex significantly showed adaptive responses to the modification of these visual inputs. According to Girman et al., cells found at a depth of about $200\,\mu$m to $800\,\mu$m (encompassing layers II and III of neocortex), had mainly negligable spontaneous activity, with high *reaction index*, *orientation index*, *direction index* and *sharpness of orientation*. The latter four quantities are defined as follows (all frequencies relate to the cell's tuning curves):

[1] A picture of a tuning curve is given in figure 3.1. Further information about tuning curves and their definition can be found in section 3.2.

- The reaction index
$$RI := \frac{|f_{max_{dev}} - f_{sa}|}{f_{max_{dev}} + f_{sa}},$$
where $f_{max_{dev}}$ is the response rate to a specific stimulus with maximum deviation from f_{sa}, the spontaneous rate. The spontaneous firing rate is the response rate elicited by a presentation of an uniform field of averaged luminance (compared to the gratings).

- The orientation index
$$OI := \frac{f_{max_{ori}} - f_{orth_{ori}}}{f_{max_{ori}}},$$
where $f_{max_{ori}}$ and $f_{orth_{ori}}$ are the firing rates given the optimal orientation and the orientation orthogonal to the optimal orientation respectively. Orientation measures the angle between the presented gratings and the monitor horizontal (ranging from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$).

- The direction index
$$DI := \frac{f_{max_{dir}} - f_{opp_{dir}}}{f_{max_{dir}}},$$
where $f_{max_{dir}}$ and $f_{opp_{dir}}$ are the firing rates given the optimal direction and opposite to the optimal orientation respectively. Direction measures angle between the movement of the presented gratings and the monitor horizontal (ranging from $-\pi$ to π).

- The sharpness of orientation SO is defined as the tuning curve's bandwith at half height. In addition to the orientation index, this measure quantifies the quality of orientation or direction tuning as well: a cell with a high orientation index and a low sharpness of orientation carries less information about the presented angle than the same cell with a high sharpness of orientation.

Stimulated with black and white gratings of optimal contrast and spatial frequency, 70% of these cells showed an reaction index that exceeded 0.7. 77% orientation selective cells ($OI \geq 0.7$) and 37% direction selective cells ($DI \geq 0.7$) were found in [12]. Few individual cells showed tetramodal tuning curves with optimal orientations nearly orthogonal to each other. Most of the tuned cells showed a sharpness of orientation of about 60°. Furthermore, a preference of cells tuned to horizontal stimuli could be found in rats, as in other species, too. Orientation columns (implying a regular change of orientation tuning across the cortex) could not be found. These findings imply that a large part of the cells in rat's visual cortex are sensitive to the orientation of simple visual stimuli and that their behaviour might carry enough information to reconstruct a presented orientation from cell responses.

However, the proportion of oriented neurons varies greatly (30 to 80%) among older studies (e.g. [34, 3, 24]) and is generally worse than the presented results in [12]. In addition to different criteria for tuning significance, unsuitable anaesthetics or oversized craniotomies used in these older works might have led to

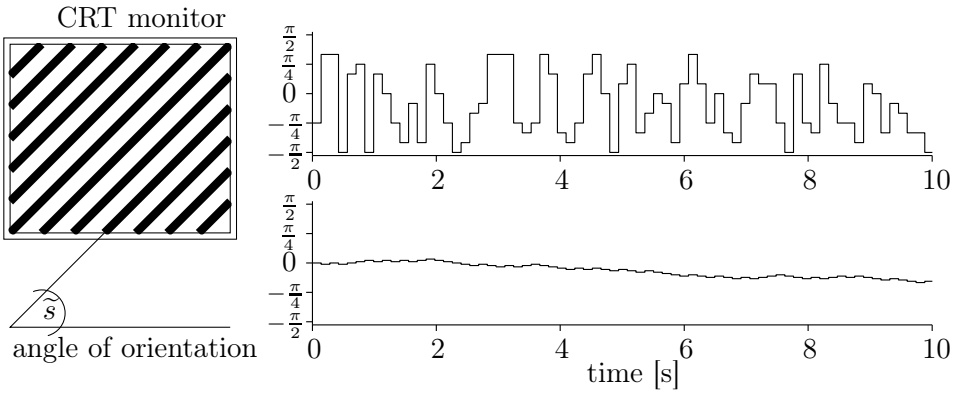

Figure 2.1: Presentation of visual stimulus. On the left side, the monitor picture as presented to the rat is drawn for $\widetilde{s} = \frac{\pi}{4}$. The curves on the right side show a cutout of the temporal stimulus dynamics, i.e. the change of orientation over time. The upper diagram shows the random stimulus used to calculate the response characteristics. The lower diagram shows the random walk that shall be reconstructed.

different or suboptimal results. Furthermore, as rats are less visual animals than cats or monkeys, data quality can generally be expected to be worse than in these animals.

Based on the fact of sparse and partly contradictory data about the tuning properties of the rat visual cortex, the following sections of this chapter describe the experimental setup used by Freiwald et al.. This setup used for presenting visual stimuli, measuring neural activity and postprocessing raw data is aiming at a reconstruction of the original stimuli from neural responses [26, 27, 10].

2.2 Visual Stimulation

Similar to Girman et al., the visual stimuli consisted of fullscreen black and white gratings with a fixed spatial freqency that were presented in different orientations ranging from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$. Stimuli were generated on a CRT monitor with a vertical frequency of 85Hz. Luminance was kept constant over different stimuli and the distance between monitor and the animal's eye was chosen to cover the rat's field of vision completely. The gratings did not move in any direction during the presentation of one orientation.

Two different kinds of stimulus trajectories have been presented (see figure 2.1).

The first consisted of equally distributed random orientations in steps of $\frac{\pi}{12}$, where each orientation was presented for 12 monitor frames (about 141 ms). 120000 frames were shown, resulting in 10000 orientation presentations. The neural activity elicited by these presentations was used to calculate the response

characteristics.

Second, a stimulus trajectory was presented, where the stimulus (i.e. the grating's orientation) followed a random path up and down in steps of $\frac{2\pi}{180}$. In this *random walk*, the difference of orientation of two consecutive presentations is always $\frac{2\pi}{180}$:

$$\widetilde{s}_{\overline{t}} = \left(\left(\widetilde{s}_{\overline{t}-1} \pm \frac{2\pi}{180} \right) \bmod \pi \right) - \frac{\pi}{2}, \tag{2.1}$$

where $\overline{t} = 1, 2, 3, ...$ is the presentation count. Again, 10000 orientations were shown in 120000 frames. This is the stimulus that will be reconstructed online in chapter 4 (from artificial spike trains) and chapter 5 (from recorded spike trains), using the response characteristics drawn from the first stimulus type.

2.3 Animal Preparation

All data have been recorded from mature anaesthetised rats of the Long Evans or Brown Norway strain. To avoid stressing factors, inital anaesthesia was induced using an Isoflurane/Laughing Gas mixture. For surgery, a Ketamine-based anaesthetic was then injected in a body weight-dependent dose. To avoid mucus production in the respiratory tracts and to stabilize the cardiovascular system, intramuscular injections of Atropine were applied. Electrocardiogram and depth of aneaesthesia have been controlled continuously during surgery.

After fixation of the rat in a stereotactic apparatus, its upper cranium surface was excavated using a sagittal transection. The surface was cleaned from any tissue and three small trepanations were drilled to set screws needed for further fixation and electrophysiological setup. The visual cortex was accessed by a left-side craniotomy of size 2x2 mm rostro-lateral from *lambda*, which is the intersection of the fissures between the parietal and occipital bones. The last surgical step was to remove the *dura mater*, allowing the electrode shank to enter smoothly.

Subsequently, the head was fixated with the right eye aligned to the monitor and the left eye closed. To avoid desiccation, the left eye was protected with a non-correcting lens and the visual cortex was moistened with a Ringer/Agar solution. The electrode was positioned at the place of representation of the intersection of the field of view's horizontal and vertical meridian.

As Ketamine has an negative impact on cortical acitivity, the animal was brought back to Isoflurane anaesthesia during stimulus presentation and recording.

2.4 Data Recording and Postprocessing

Recordings were performed with two types of silicon-based multiprobes: a single-shanked electrode with 16 channels in linear alignment and tetrodes with 4 channels on each shank. Signals were amplified (1000-5000 fold) and filtered (0.3 to 5 kHz) using standard electrophysiological instrumentation.

For offline analysis, the analog voltage signal was digitised and transferred to a computer at a sampling rate of 25kHz using a conventional PC analog-digital

transducer. To extract spikes from the voltage curve, a positive and negative bias had been chosen manually. If the signal exceeded the limits, it was treated as a spike event and the signal shape in the temporal neighbourhood has been saved in a matrix for spike sorting.

Spike sorting is necessary, as each channel is likely to record a composite signal of multiple cells. Sorting assigns each spike event to a cell, which is supposed to have a characteristic spike wave form. This has been done computer-aided by clustering two-dimensional scatterplots of different wave-form-describing parameter pairs (i.e., amplitude, bandwith, etc.).

The outcome of this procedure is a list of discrete spike times for each cell that has been recorded.

3 A Model for Neural Activity

In our context, measured neural activity elicited by stimulus presentation can be described by a model that maps values from stimulus space (angles of oriented gratings in our case) to response space (neural activity, i.e. spike trains of recorded cells). For reasons to be discussed in the subsequent section, this model is assumed to be of probabilistic nature. Therefore, we consider the probability of a recorded response given a certain stimulus or, interpreted as a function of the stimulus, the likelihood $p(\text{response}|\text{stimulus})$.

In this chapter, we develop a likelihood term by combining a deterministic data-fitting model for tuning curves and an inhomogeneous Poisson spike generator. Tuning curve and the process of spike generation together provide a complete description of a cell in our model. These findings are put together to model a population likelihood function. Intuitively, information made available by the likelihood will be crucial for the decoding stage described in the following chapter.

3.1 Nomenclature and Model Assumptions

We refer to $\mathbb{S} = \left[-\frac{\pi}{2}, \frac{\pi}{2}\right)$ as a one-dimensional stimulus space representing the angles of orientation of presented gratings, where $\widetilde{s} = -\frac{\pi}{2}$ denotes vertical bars and $\widetilde{s} = 0$ denotes horizontal bars.

As it is convenient to use discrete times on computational architectures, a time intervall $\mathbb{T} = [0, T_{max}]$ is divided into T subintervals of length Δt. $t = 1, ..., T$ denotes the t-th timestep within $\mathbb{T}$, which is the time interval $[t \cdot \Delta t, (t+1) \cdot \Delta t)$.

A discrete-time dynamic stimulus $\mathfrak{s} = \widetilde{s}_1, ..., \widetilde{s}_T$ is defined as time series of stimuli. Without tilde, s is used to denote all possible values of $s \epsilon \mathbb{S}$ to remind that some distributions occuring in the next chapters are functions of s. I.e., $p(s_t)$ denotes the probability *distribution* of each possible stimulus to occur at time step t.

R_i is the response of cell $i = 1, .., N$, which is a spike train (an ordered list of spike times with n_i elements) $R_i := \left\{t_i^1, ..., t_i^{n_i}\right\}$ and $t_i^1 \leq ... \leq t_i^{n_i}$. Although action potentials can vary in duration, amplitude and shape, they are treated as identical stereotyped events. Since stimuli are generally not coded by single cells, we turn the attention to the population response R, which is represented by a list of N recorded cell's spike trains. With respect to the introduction, it should be stressed again that here a further and more general assumption is made. Namely, spike trains of specific cells are considered to represent the *neural code*, i.e. are significant entities carrying information about the dynamic stimulus.

While presenting the same stimulus, the cell and population response will ususally vary from trial to trial. Assuming that a population carries information

about the stimulus, potential sources of stochasticity in our setting seem to fall into four classes: *a*) the population might code for other qualities or modalities that are not part of a controlled change during the experiment; *b*) the physics of a cell could be intrinsically probabilistic; *c*) Somatic changes of the cell's or population's surrounding could affect their internal states. This includes changes of intra- and extravasal substance concentrations as well as changes due to other modulating cognitive processes (like arousal or attention) these cells are connected to via nerve fibres; *d*) probes penetrating the area might interact in an uncontrolled way. Naturally, these classes do not exclude each other [21, 16, 5].

For our purpose, possible reasons of stochasticity do not need to be discussed furthermore: the fact of probabilistic behaviour of the neural system will be reflected in the mathematical framework that forms the model. Within this model, assertions about stimulus/response relations are made in the presence of stochasticity, independently from its reasons.

3.2 Tuning Curves

A neuron's tuning curve f is a function from stimulus space into the positive reals, mapping each stimulus presented to an experimentally measured mean firing rate provoked by this stimulus. The mean firing rate is calculated by averaging the number of spikes caused by a certain stimulus over an infinite number of trials, divided by the duration of the trial. With $r_1(\widetilde{s}), r_2(\widetilde{s}), ...$ being the measured rates of single trials, the tuning curve is defined by

$$f(\widetilde{s}) := \langle r(\widetilde{s}) \rangle \,. \tag{3.1}$$

In practise, the experimenter is confronted with two problems. First, the number of trials is always finite. Therefore, the mean firing rate may have to be estimated from measurements that contain considerable jitter. Second, in continuous stimulus space, only a finite set of all possible stimuli can be used to measure mean firing rates. These are used as sampling points to fit (parametrically or non-parametrically) shapes of continuous tuning curves. Although comparatively simple representations as Gaussians or cosines have been used successfully in recent research to describe tuning curves [2, 28, 5, 37], they are only suitable to depict unimodal and symmetric shapes. For multimodal experimental data (as in our case), it is necessary to use more complex functions, such as a van Mises function [28, 27, 6] or even non-parametric models [11].

Therefore, fitting a tuning curve generally involves three tasks: Choosing a model (i.e. a function class), estimating the sampling points from data and calculating the model's parameters to fit these vectors. The latter task usually involves an optimization process, which is in most cases faced with the problem of avoiding local minima in parameter space.

A straightforward approach to obtain a tuning function would be to estimate the mean rate by averaging over the given trials and filling the gaps in stimulus space using linear interpolation. The resulting tuning curve from K datapoints (a datapoint $\langle \widetilde{s}_k, r_k \rangle$ is a stimulus/response pair) and L trials can be described

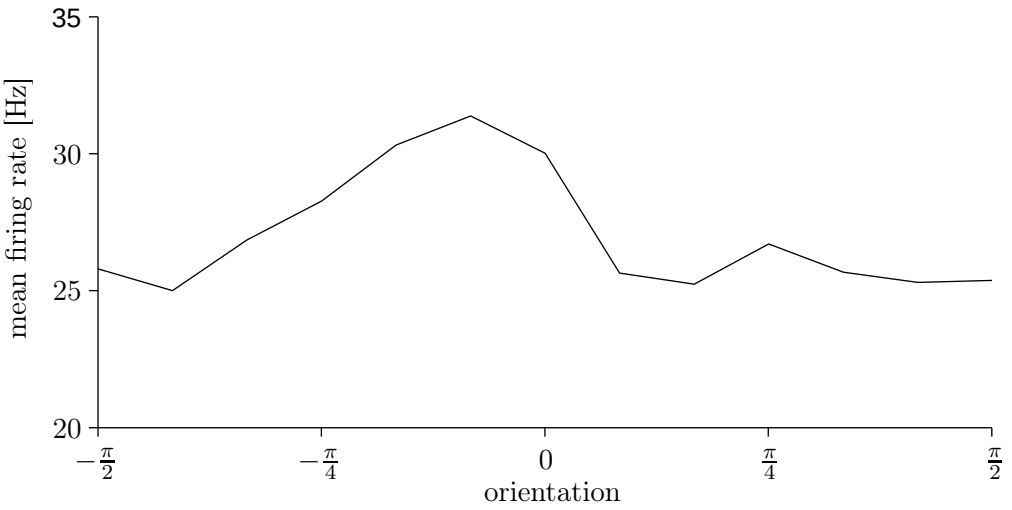

Figure 3.1: A neuron's tuning curve generated by linear interpolation of sampling points. Each sampling point has been calculated by averaging the recorded firing rate of a rat V1 neuron for each orientation over multiple trials.

as a piecewise linear function

$$f(s) = \overline{r}_k + (s - \widetilde{s}_k) \cdot \frac{\overline{r}_{k+1} - \overline{r}_k}{\widetilde{s}_{k+1} - \widetilde{s}_k} \quad \widetilde{s}_k \leq s \leq \widetilde{s}_{k+1}, \tag{3.2}$$

with $k = 1, ..., K-1$, $l = 1, ..., L$ and $\overline{r}_k = \frac{1}{L} \sum_{l=1}^{L} r_k^l$, which is the average rate measured for stimulus $\widetilde{s}_k$ over L trials (figure 3.1 shows a typical tuning curve recorded from rat visual cortex). The shortcoming of this method is that outliers tend to be overestimated[1], as all values have the same weight contributing to $\overline{r}$ and that a biological tuning curve is supposed to be smooth rather than piecewise linear.

Both problems can be solved by choosing an algorithm to fit tuning curves using trigonometric polynomials as described in [6]. The idea is to write each response r as

$$r_{kl} = a_0 + \sum_{k=1}^{K\,div\,2} (a_k \cos(2k\theta_l) + b_k \sin(2k\theta_l)) + \varepsilon_{kl}, \tag{3.3}$$

with

$$K\,div\,2 = \begin{cases} \frac{K}{2} & K \text{ even} \\ \frac{K-1}{2} & K \text{ odd} \end{cases}$$

and to find the parameters a and b that minimise the residuals ε in a ranked-based norm: to reduce the impact of single outliers, this norm applies a ranking to all datapoints in a way that underweighs outliers. The optimization algorithm

[1] This is a problem of commonly used least-square fits as well: minimizing the squared error between each sampling point and corresponding data points, all errors have the same weight; even outliers that might have been sorted out by hand [6].

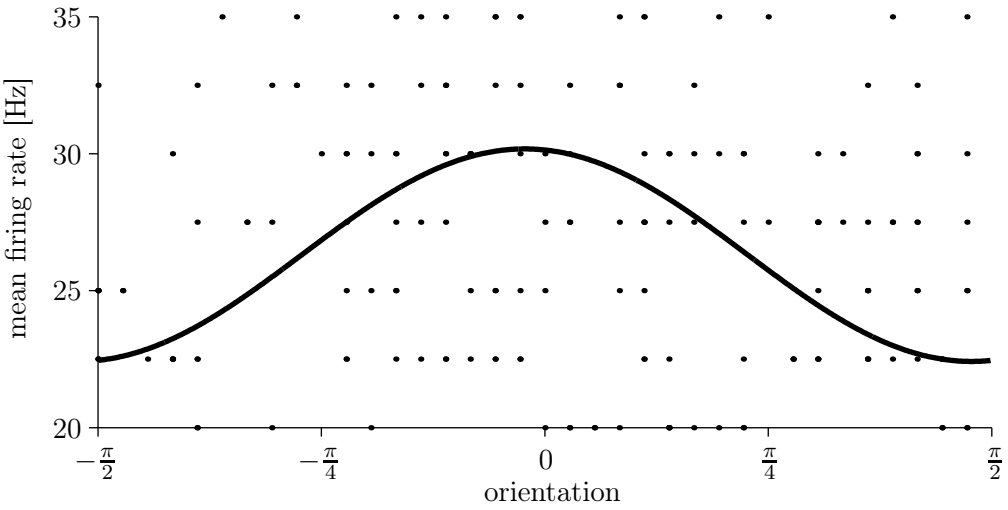

Figure 3.2: Scatterplot and tuning curve of lowest well-fitting degree. The scatterplot shows the distribution of firing rates for different stimuli. The bold line shows a trigonometric polynomial of degree 2 fitting the data. The same empirical data was used to generate the linearly interpolated tuning curve from figure 3.1.

deterministically finds the best approximation of the tuning curve. Finally, the degree of the polynomial is reduced as long as the coefficients of the highest order do not account significantly to fit the data. Hence, although the approximation of maximal degrees allows to have asymmetries and several maxima, a simpler model is chosen if possible (figure 3.2).

3.3 Spike Generation

In addition to a mean firing rate provided by tuning curves, our model for neural activity needs to make a probabilistic prediction of exact times of spike events as well. Point processes are useful for modeling random processes that appear at irregularly spaced times and which may possibly occur anywhere in a continuous time interval. Such a process can be divided into two independent parts: First, the number of spike events in a certain time intervall is chosen on the basis of some distribution and second, a discrete time is assigned to each event.

A simple point process for spike generation that uses a cell's tuning curve as intensity function is the commonly used inhomogeneous Poisson process [21, 5]. Its most important characteristic is that the generation of each spike is independent of all other spikes. The correctness of this *independent spike hypothesis* has been both supported and criticised in experimental and theoretical discourse [14, 5, 21], depending on the data it has been applied to. If the Poisson assumption is a sensible choice for the data used in this work will be subject of discussion in section 5.2.

For cell i, the probability of a single spike occurring in a small interval of length Δt within $\mathbb{T}$ is given by $\frac{\Delta t}{T_{max}}$. Hence, placing k spikes in $\mathbb{T}$, the probability of

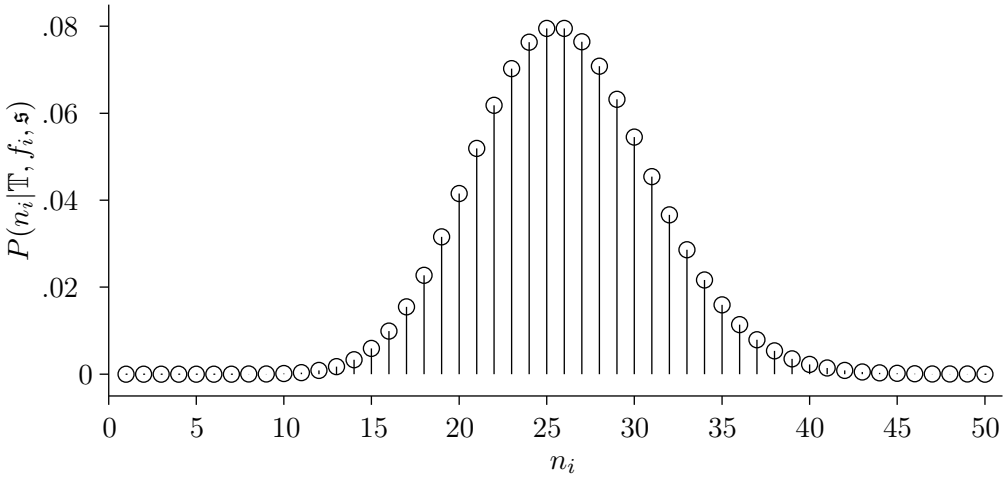

Figure 3.3: Spike count distribution for a Poisson process. The stem plot shows the probability of finding n_i spikes for $\sum_{t=1}^{T} f_i(\widetilde{s}_t)\Delta t = 25$. The spike count distribution describes the probability of a certain number of spikes to occur, given the cell's tuning curve and the stimulus trajectory on a time interval $\mathbb{T}$.

finding n_i of them in a time window of length Δt is described by the binomial formula

$$P(n_i|\Delta t) = \frac{k!}{(k-n_i)!n_i!}\left(\frac{\Delta t}{T_{max}}\right)^{n_i}(1-\frac{\Delta t}{T_{max}})^{k-n_i.} \tag{3.4}$$

Keeping the ratio $\frac{k}{T_{max}}$ constant, it can be shown for $k \to \infty$ that

$$P(n_i|\Delta t) = e^{-\frac{k}{T_{max}}\Delta t}\frac{(\frac{k}{T_{max}}\Delta t)^{n_i}}{n_i!}. \tag{3.5}$$

This is the Poisson probability density function. Interpreting $\frac{k}{T_{max}}$ as the mean firing rate and inserting equation (3.1) yields

$$P(n_i|\Delta t, f_i, \widetilde{s}_t) = e^{-f_i(\widetilde{s}_t)\Delta t}\frac{(f_i(\widetilde{s}_t)\Delta t)^{n_i}}{n_i!}. \tag{3.6}$$

Finally, integrating with respect to $\mathbb{T}$ (as the stimulus is considered to be constant in each time window this amounts to summing up all mean rates for each time window) results in the spike count distribution for an inhomogeneous Poisson process we have been looking for (figure 3.3):

$$P(n_i|\mathbb{T}, f_i, \mathfrak{s}) = e^{-\sum_{t=1}^{T} f_i(\widetilde{s}_t)\Delta t}\frac{(\sum_{t=1}^{T} f_i(\widetilde{s}_t)\Delta t)^{n_i}}{n_i!}. \tag{3.7}$$

The spike count distribution describes the probability of a certain number of spikes to occur, given the cell's tuning curve and the stimulus trajectory on a time interval $\mathbb{T}$.

Next, the exact spike times for a spike train $R_i = \left[t_i^1, ..., t_i^{n_i}\right]$, $t_i^j \epsilon \{1, ..., T\}$ have to be calculated. As for each of n_i spikes, a spike time t_i^j is chosen independently from all other spikes and the spike sequence can be created in any

order, the probability for a certain spike train configuration can be described by the product of the probabilities for the selection of a single spike time which can occur in all permutations [16]:

$$p(R_i|\mathbb{T}, f_i, \mathfrak{s}, n_i) = \prod_{j=1}^{n_i} p(t_i^j|\mathbb{T}, f_i, \mathfrak{s}) \cdot n_i!. \tag{3.8}$$

The probability density for selecting a certain time step t_i^j (i.e. a certain time interval of length Δt, see section 3.1) is proportional to the mean firing rate and described by the following distribution:

$$p(t_i^j|\mathbb{T}, f_i, \mathfrak{s}) = \frac{f_i(\widetilde{s}_{t_i^j})}{\sum_{t=1}^{T} f_i(\widetilde{s}_t)}. \tag{3.9}$$

As already mentioned, the two subprocesses of a poisson point process (spike number generation and assignment to discrete spike times) are independent as well. Thus, combining equations (3.7), (3.8) and (3.9), the conditional probability for a single cell response derived from its tuning function f_i and a Poisson process for spike generation simplifies to

$$\begin{aligned} p(R_i|\mathbb{T}, f_i, \mathfrak{s}) &= p(R_i|\mathbb{T}, f_i, \mathfrak{s}, n_i) \cdot P(n_i|\mathbb{T}, f_i, \mathfrak{s}) & (3.10)\\ &= \prod_{j=1}^{n_i} f_i(\widetilde{s}_{t_i^j}) \cdot e^{-\sum_{t=1}^{T} f_i(\widetilde{s}_t)\Delta t}. & (3.11) \end{aligned}$$

Interpreted as a function of s, this yields the likelihood distribution for a fixed but arbitrary single cell response:

$$p(R_i|\mathbb{T}, f_i, s) = n_i \cdot f_i(s) \cdot e^{-\sum_{t=1}^{T} f_i(s)\Delta t}. \tag{3.12}$$

3.4 Population Response

The next step after deducing the likelihood term of a single cell is to model the response of a whole population of cells to obtain the likelihood term $p(R|\mathbb{T}, f, \mathfrak{s})$ we were initially looking for. In a recurrent neural network like the cerebral cortex, interdependencies between the activities of connected neurons may occur and can, in principle, be measured using multicell recordings. In [1], these intercell correlations are expressed as mutual deviation from their mean firing rate

$$r_{i,t} = f_i(\widetilde{s}_t) + \eta_i, \tag{3.13}$$

where the random terms η_i are generated from a Gaussian probability distribution with zero mean and covariance matrix Q[2], which might depend on s_t or not. According to theoretical results [7, 1, 8], for different shapes of Q, correlations can generally both improve and degrade the accuracy of a population in

[2]For example, $Q_{ij} = \sigma^2 \cdot (\delta_{ij} + c(1 - \delta_{ij}))$ employs a simple model for variance increasing as a function of the degree of correlation c. The base variances σ^2 are identical for both neurons i and j, with δ_{ij} representing the Kronecker matrix. A value of c ranging from 0.1 to 0.2 has been found in area MT of monkey cortex [1].

encoding a certain stimulus. Whether the influence of intercell dependencies is positive or negative depends on the precise configuration of the neural system. This information is not available from the experiments relevant for this work.

Furthermore, it is not clear that even if correlations among the activities of neurons are advantageous from an information-theoretic perspective, the organism actually makes use of this advantage (i.e., if it is part of the neural code [21]). Therefore, and for reasons of simplicity we assume $\eta_i = 0$ and employ a population of conditionally independent neurons. In this case, the probability of a population response, which is the joint probablity of all single cell responses, is simply the product of single cell response probabilities from equation (3.12):

$$\begin{aligned} p(R|\mathbb{T}, f, s) &= \prod_{i=1}^{N} p(R_i|\mathbb{T}, f_i, s) && (3.14) \\ &= \prod_{i=1}^{N} n_i \cdot f_i(s) \cdot e^{-\sum_{t=1}^{T} f_i(s)\Delta t} && (3.15) \\ &= e^{-\sum_{t=1}^{T}\sum_{i=1}^{N} f_i(s)\Delta t} \cdot \prod_{i=1}^{N} n_i \cdot f_i(s) && (3.16) \end{aligned}$$

Summarizing, equation (3.16) characterises all what we assume about neural coding of our dynamic stimulus $\mathfrak{s}$: The response R is a set of spike trains produced by a population of cells with arbitrary tuning curves, a Poisson spike count density and mutual independency. As all of the following distributions are conditioned on f and $\mathbb{T}$, they will not be explicitly mentioned anymore (i.e., we write $p(R|\mathfrak{s})$ instead of $p(R|\mathbb{T}, f, \mathfrak{s})$ from now on).

3.5 Encoding of a Dynamic Stimulus

As an example for the spike generation process, we will encode a dynamic stimulus that changes linearly from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$ with a population of 30 computer-simulated neurons with Gaussian tuning curves. The tuning functions' mean values are equally distributed from $-\frac{3\pi}{4}$ to $\frac{3\pi}{4}$ to cover the stimulus space completely. The maximum firing rate of these Gaussians is choosen to be 30 Hz and their variance is set to $\frac{\pi}{8}$.

As the maximum firing rate is finite, a sufficiently small interval Δt can be choosen to make the probability of more than one spike to occur in this interval ignorable [14]. The *absolute refractory period* (the duration after an action potential where a neuron is not excitable, usually lasting about $1ms$ [17]) that could be used as a lower bound, is not modelled by the Poisson process. Hence, Δt has to be chosen to be smaller than all occuring interspike time intervals generated by the Poisson process. For our simulations, $\Delta t \leq 0.001s$ has been a sensible choice.

For small enough Δt, equation (3.6) simplifies in approximation to

$$P(n_i = 1|\Delta t, f_i, \widetilde{s}_t) = f_i(\widetilde{s}_t)\Delta t.$$

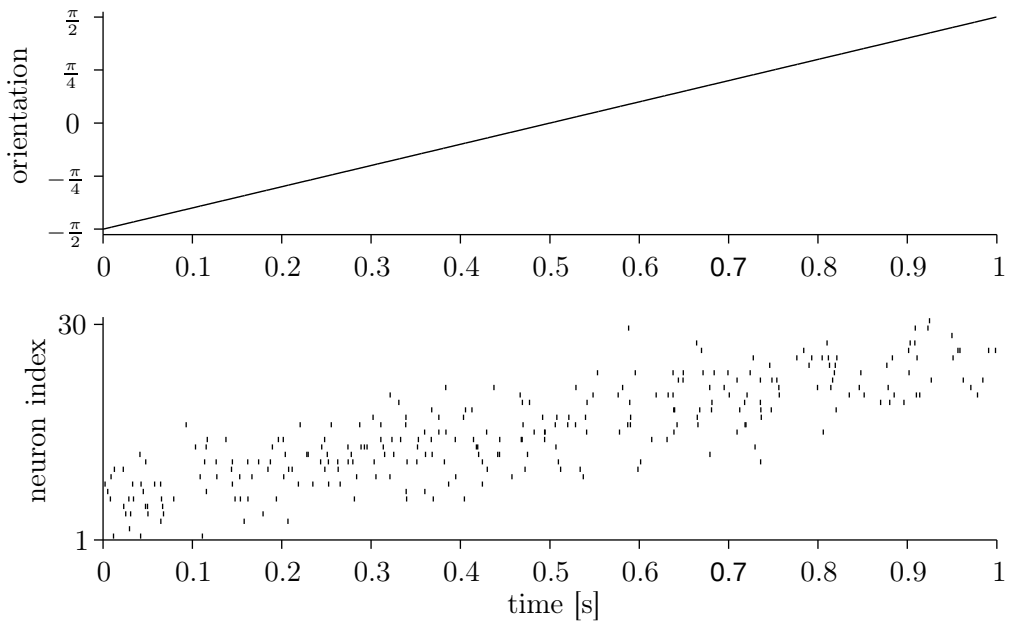

Figure 3.4: Dynamic stimulus and computer-simulated neural response. The upper plot shows the linear trajectory of the stimulus. The lower diagram is a raster plot of 30 neurons with Gaussian tuning functions that encode this stimulus. As the tuning functions' mean values are distributed linearly over stimulus space, the spiketrain pattern resembles the stimulus trajectory.

Hence, to find out wether or not a single spike occurs in each of $t = 1, .., T$ intervals of length Δt, the spiking condition for each interval can be written as

$$\text{rand}[t] \leq f_i(\widetilde{s}_t)\Delta t,$$

where rand$[\cdot]$ is an array of uniformely distributed random numbers in the $[0, 1]$ interval.

According to the previous section, spike trains can be drawn independently for each cell. An example of a computer-simulated neural response to the described stimulus is shown in figure 3.4.

4 Stimulus Reconstruction

The purpose of this chapter is to establish a probabilistic framework for the reconstruction of a dynamic stimulus. Reconstruction in our context means approximating the true stimulus $\widetilde{s}$ by an estimate $\widehat{s}$, given the neural response R and model assumptions made in chapter 3.

In the first section, Baye's rule is used to deduce a posterior probability distribution $p(s|R)$ from which $\widehat{s}$ is extracted. Second, non-linear filtering is used to introduce information available from temporal propagation of the stimulus. Finally, a simulation example will show the theoretical feasibility of this approach.

4.1 Nomenclature and Model Assumptions

As reconstruction can be performed in an online fashion (e.g., only the response of the actual and preceding timesteps are used for estimating the stimulus), we will consider population responses R_t at each timestep t. Responses R_t are assumed to be mutually independent and are not conditioned on the history of temporal stimulus dynamics:

$$p(R_t, ..., R_1|s_t, ..., s_1) = \prod_{t'=1}^{t} p(R_{t'}|s_{t'}). \tag{4.1}$$

Furthermore, stimuli are uncorrelated and no underlying model for the temporal stimulus dynamics itself is given:

$$p(s_t, ..., s_1) = \prod_{t'=1}^{t} p(s_{t'}). \tag{4.2}$$

The second assumption will be subject of change in section 4.3.

4.2 Bayesian Inference

While our model of neural coding from chapter 3 describes the probability of a reponse, given a certain stimulus, it is in the interest of reconstruction to infer the inverse distribution that yields the probability of a stimulus, given a certain (measured) response. The joint probability distribution $p(s_t, R_t)$ can be written as $p(R_t|s_t) \cdot p(s_t)$ as well as $p(s_t|R_t) \cdot p(R_t)$. From this simple equality, Bayes rule states that

$$p(s_t|R_t) = \frac{p(R_t|s_t) \cdot p(s_t)}{p(R_t)}, \tag{4.3}$$

where $p(s_t|R_t)$ is the so-called *posterior distribution*, which, in our context, shapes the probability of each $s_t \epsilon \mathbb{S}$ given an (arbitrary but fixed) response. Hence, it is proportional to the product of the likelihood $p(R_t|s_t)$ and the prior probability distribution $p(s_t)$ (denoting the probability of an stimulus to occur, independent of any data from the neural system):

$$p(s_t|R_t) = \lambda_t \cdot p(R_t|s_t) \cdot p(s_t), \tag{4.4}$$

with $\lambda_t \cdot \int_{\mathbb{S}} p(R_t|\widetilde{s}_t) \cdot p(\widetilde{s}_t)\, d\widetilde{s}_t = 1$.

Since all information about the stimulus is contained in the posterior distribution, the final step is to collapse the posterior distribution $p(s_t|R_t)$ in some way to obtain a discrete value for $\widehat{s}_t$. The Bayesian decision-theoretic approach is to use a loss function $L(\widetilde{s}_t, \widehat{s}_t)$ to measure the loss incurred by estimating the value of the stimulus to be $\widehat{s}_t$ when its true value is $\widetilde{s}_t$. The estimate $\widehat{s}_t$ is chosen to minimise the expected loss, that is, optimizing the loss over all $\widetilde{s}_t$ with respect to the posterior distribution:

$$\begin{aligned} \widehat{s}_t &= \arg\min_s E\left(L(\widetilde{s}_t, s)\right) && (4.5) \\ &= \arg\min_s \int_{\mathbb{S}} L(\widetilde{s}_t, s) \cdot p(\widetilde{s}_t|R_t) d\widetilde{s}_t. && (4.6) \end{aligned}$$

Inserting the commonly used quadratic error loss function $L_{sq}(\widetilde{s}_t, \widehat{s}_t) = (\widetilde{s}_t - \widehat{s}_t)^2$ in equation (4.6) results in

$$\widehat{s}_t = \arg\min_s \int_{\mathbb{S}} (\widetilde{s}_t - s)^2 \cdot p(\widetilde{s}_t|R_t) d\widetilde{s}_t, \tag{4.7}$$

which can be solved by finding the minimum of the first derivative with respect to s:

$$\begin{aligned} \widehat{s}_t &= \arg_s \left\{ 2 \int_{\mathbb{S}} (s - \widetilde{s}_t) \cdot p(\widetilde{s}_t|R_t) d\widetilde{s}_t = 0 \right\} && (4.8) \\ &= \arg_s \left\{ \int_{\mathbb{S}} \widetilde{s}_t \cdot p(\widetilde{s}_t|R_t) d\widetilde{s}_t = s \right\} && (4.9) \\ &= \int_{\mathbb{S}} \widetilde{s}_t \cdot p(\widetilde{s}_t|R_t) d\widetilde{s}_t && (4.10) \\ &= E\left(s_t; p(s_t|R_t)\right). && (4.11) \end{aligned}$$

Hence, quadratic error loss is minimised by taking $\widehat{s}_t$ to be the *posterior mean* [22, 18, 16][1].

So far, we have not discussed the role of the prior distribution $p(s_t)$, the unconditional probability distribution of a certain stimulus to occur. If not given by an experimental setting, this distribution is not known, and therefore usually considered as constant:

$$p(s|R_t) = \lambda_t \cdot p(R_t|s) \cdot p(\widetilde{s}), \tag{4.12}$$

[1] Another, straightforward approach is to use the so-called MAP estimator, where the maximimum of the posterior distribution is chosen.

with $\widetilde{s}\epsilon\mathbb{S} = const.$ Although for a large amount of experimental data choosing an ignorant prior does not have any negative effects for the estimate's value, a sensible choice for small data sets is important to obtain sharper posterior distributions [16]. In the next section this will be done by applying assumptions about the stimulus propagation in time.

4.3 Non-linear Filtering

At this point, equations (4.12) and (4.11) describe a complete way to estimate the originial stimulus at each timestep. However, it is not satisfactory to use an ignorant prior: in the physical world most states of macroscopic entities change gradually over time. Intuitively, this means that there is a connection between the stimuli's values at adjacent timesteps, i.e. it is quite probable that $\widetilde{s}_t$ lies in the neighbourhood of $\widetilde{s}_{t-1}$.

Hence, a somewhat general assumption is that the stimulus dynamics form a temporal Markov chain (with $\widetilde{s}_t$ being the hidden variable), where a direct dependence exists only between two temporally contiguous stimuli. This will replace the assertion made in equation (4.2) with

$$p(s_t|\widetilde{s}_{t-1}, ..., s_1) = p(s_t|\widetilde{s}_{t-1}). \tag{4.13}$$

Assuming a dependence between two subsequent timesteps implies that the current stimulus s_t depends now conditionally not only on R_t but indirectly on the complete response history $R_{t-1}, ..., R_1$ as well. This requires a reformulation of equation (4.4), since now the posterior distribution $p(s_t|R_t, ..., R_t)$ represents all available information that can be deduced at time t. In general, trying to estimate the state of a hidden variable of a discrete-time controlled process like this is called *filtering*.

While the mean of the posterior distribution can still be used as an optimal estimator with respect to a quadratic loss function, the posterior term itself has to be revised: employing the Markov assumption and that of mutual response independency from equation (4.1) we obtain

$$\begin{aligned} p(s_t|R_t, ..., R_1) &= \lambda_t \cdot p(R_t|s_t, R_{t-1}, ..., R_1) \cdot p(s_t|R_{t-1}, ..., R_1) && (4.14) \\ &= \lambda_t \cdot p(R_t|s_t) \cdot p(s_t|R_{t-1}, ..., R_1), && (4.15) \end{aligned}$$

with $\lambda_t \cdot \int_{\mathbb{S}} p(R_t|\widetilde{s}_t) \cdot p(\widetilde{s}_t|R_{t-1}, ..., R_1)\, d\widetilde{s}_t = 1$ and the *temporal prior distribution* $p(s_t|R_{t-1}, ..., R_1)$ that can be transformed into a recursive formulation (again using equation (4.1) for independency between s_t and previous responses):

$$\begin{aligned} p(s_t|R_{t-1}, ..., R_1) &= \int_{\mathbb{S}} (p(s_t|\widetilde{s}_{t-1}, R_{t-1}, ..., R_1) \cdot \\ &\quad\; p(\widetilde{s}_{t-1}|R_{t-1}, ..., R_1))\, d\widetilde{s}_{t-1} && (4.16) \\ &= \int_{\mathbb{S}} p(s_t|\widetilde{s}_{t-1}) \cdot p(\widetilde{s}_{t-1}|R_{t-1}, ..., R_1) d\widetilde{s}_{t-1}. && (4.17) \end{aligned}$$

Combining equations (4.15) and (4.17), the process of inferring the posterior distribution at timestep t can be interpreted as a prediction of the new state that is corrected by the actual measurement:

$$p(s_t|R_t, ..., R_1) =$$

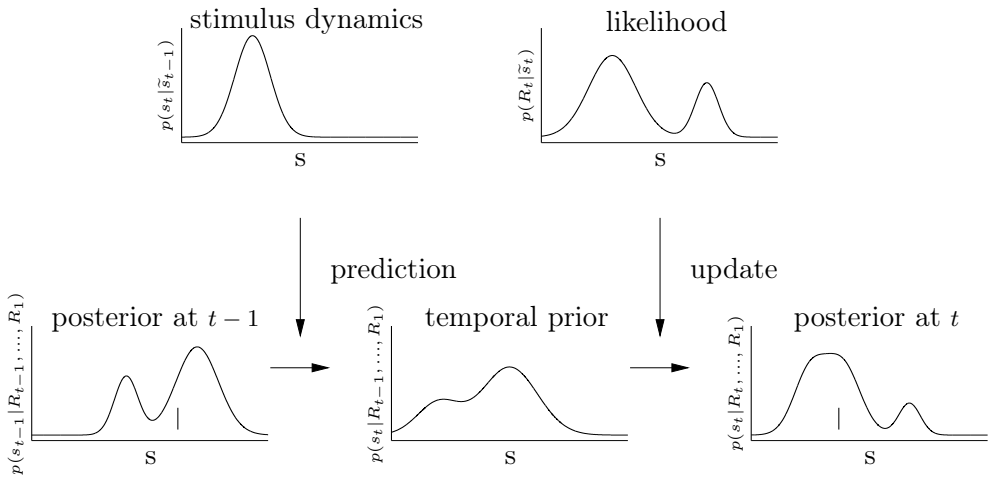

Figure 4.1: Computer-simulated filtering as density propagation. The temporal prior is a convolution of the posterior at time $t-1$ and the density describing the stimulus dynamics. This can be seen as a prediction of the posterior at time t, which is updated by multiplication with the likelihood. The vertical lines in the posterior distributions indicate their mean value, which is our estimate $\widehat{s}$.

$$\lambda_t \cdot \underbrace{p(R_t|s_t)}_{\text{update}} \cdot \underbrace{\int_{\mathbb{S}} d\widetilde{s}_{t-1} \quad \underbrace{p(s_t|\widetilde{s}_{t-1})}_{\text{stim. dynamics}} \cdot \underbrace{p(\widetilde{s}_{t-1}|R_{t-1},...,R_1)}_{\text{posterior at t-1}}}_{\text{prediction}}.$$

The conditional density $p(s_t|\widetilde{s}_{t-1})$ is a complete description of the stimulus dynamics. It shapes the way the stimulus is expected to change without knowledge of the actual measurement. A prediction of s_t, given the reponse history $R_{t-1},...,R_1$ (which is the temporal prior term at timestep t) is made by integrating with respect to all possible preceding stimuli $\widetilde{s}_{t-1}$, that is, convolving the stimulus dynamics and the preceding posterior distribution. From a different point of view this change can be seen as deterministic drift that shifts the posterior distribution towards a new position in stimulus space, and a stochastic diffusion, adding incertainty to the result [15]. To obtain the posterior distribution at timestep t, the prediction is then updated by multiplication with the measurement at timestep t, which is represented by the likelihood distribution. A computer simulation of the whole process of density propagation is visualised in figure 4.1.

So far, no restricting assumptions about the shape of $p(R_t|s_t)$, $p(s_t|\widetilde{s}_{t-1})$ and $p(s_{t-1}|R_{t-1},...,R_1)$ have been made. It is well known that unless normality of these distributions is assumed, an explicit expression for the filtering algorithm cannot be derived [29, 30]. Presupposing Gaussians for all involved densities exploits a thouroughly studied engineering tool: the *Kalman filter* [29, 23]. The standard Kalman filter model propagates only the first two moments (i.e. mean and variance) of the underlying distributions, which is sufficient for Gaussian

distributions, as they can be completely described in terms of these moments. This removes the main computional burden of calculating the convolution of $p(s_t|\widetilde{s}_{t-1})$ and $p(s_{t-1}|R_{t-1},...,R_1)$, as convolving two Gaussians amounts simply to adding their means and variances. Furthermore, multiplication of two Gaussians can be expressed analytically and the propagation of the optimal estimate with respect to a quadratic loss function is already built in.

Although the Kalman filter has been used successfully in neural reconstruction problems [36, 2, 16], it is known to show non-robust behaviour if there is a large discrepancy between the likelihood and the temporal prior [29]. More important, it cannot account for non-normal and possibly multimodal distributions as they occur in our experiment [29, 15].

Different efforts have been made to find more powerful extensions of the Kalman filter (e.g. [32, 29, 30, 15]) and one of them, the so-called *Particle filter*, has already been used for neural decoding of multicell recordings in motor cortex [11]. Our approach is to sample the densities in high resolution as a set of piecewise constant functions and to calculate the mean of the posterior distribution numerically[2].

Summarizing, our final filter equation works with any distribution as it may evolve in time and no restrictions to the stimulus dynamics are made (i.e. non-linear stimulus transitions of non-normal distributions are possible):

$$\begin{aligned}
\widehat{s}_t &= E\left(s_t; p(s_t|R_{t-1},...,R_1)\right) && (4.18)\\
&= E\left(s_t; \lambda_t \cdot p(R_t|s_t) \cdot \int_{\mathbb{S}} p(s_t|\widetilde{s}_{t-1}) \cdot p(\widetilde{s}_{t-1}|R_{t-1},...,R_1) d\widetilde{s}_{t-1}\right) && (4.19)\\
&= E\left(s_t; p(R_t|s_t) \cdot \int_{\mathbb{S}} p(s_t|\widetilde{s}_{t-1}) \cdot p(\widetilde{s}_{t-1}|R_{t-1},...,R_1) d\widetilde{s}_{t-1}\right). && (4.20)
\end{aligned}$$

4.4 Random Walk Simulation

To show the filter's functionality and the feasibility of this approach, we create an artificial population of neurons to encode the random walk stimulus as shown in the experiment. Then, the filter from the previous section is used to reconstruct the stimulus trajectory.

An idealised neuron in simulation could be represented by a bimodal von Mises function, which is suitable for shaping an orientation tuning curve in circular data. However, using this function would be to complex to provide a comprehensible analytic solution. While this is not a problem for the numeric computer simulation, we want to be able to follow the reconstruction process in a symbolic way: for didactic reasons, we again choose the representation to be a Gaussian tuning curve[3]. To obtain a more realistic setting compared to section

[2] Calculating the convolution of arbitrary densities $p(s_t|\widetilde{s}_{t-1})$ and $p(s_{t-1}|R_{t-1},...,R_1)$ can be done efficiently by multiplying the fourier-transformed distributions. A faster fourier transformation (FFT algorithm) can be applied if the number of nodes in numerical integration is a power of two [20].

[3] Naturally, Gaussians are not suitable for circular data. To avoid connecting the boundaries of $\mathbb{S}$, they are virtually extended in numerical simulation. Hence, in simulation the tuning

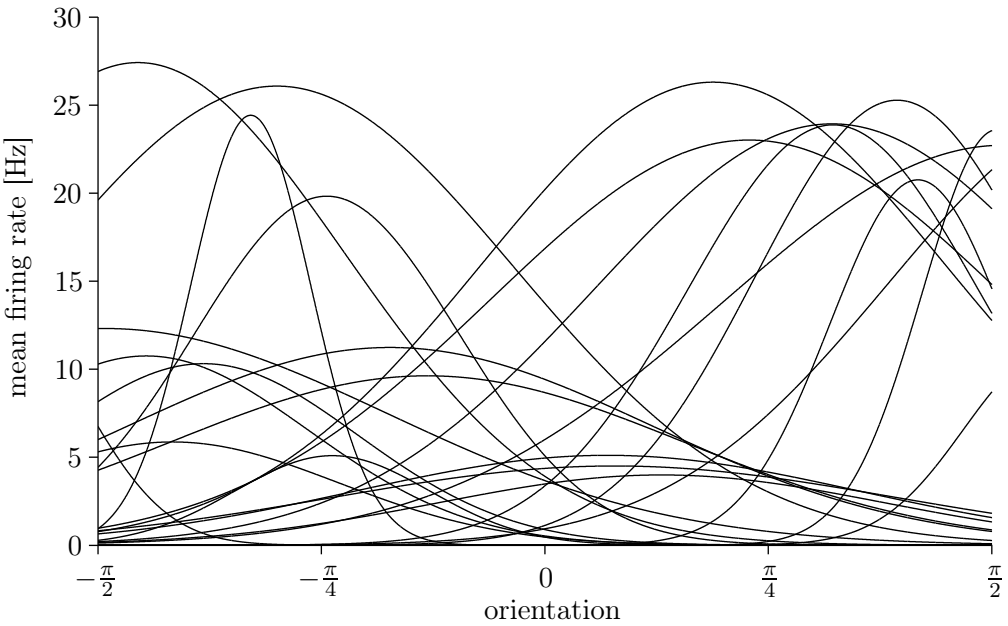

Figure 4.2: The plot shows 25 of 100 randomly computer-generated Gaussian tuning curves. Maximum mean firing rates range uniformly distributed between 0 and 30 Hz. Means in stimulus space range from $-\pi$ to π, variances from 0.2 to 1. All curves are scaled to have the same area as a population of 100 neurons with a peak firing rate of 20 Hz and a variance of 0.6.

3.5, maximum mean firing rates and variances are chosen randomly from some interval and 100 of these neurons are distributed randomly over stimulus space as suggested in figure 4.2. In analogy to the linear stimulus in section 3.5, the population is then used to encode a part of the random walk stimulus from the experiment as shown in figure 4.3.

According to equation (2.1), the stimulus is modeled to change randomly in a positive or negative direction from one timestep to the next. Although we know the absolute value of change to be constant, we discard this knowledge and model the change over time as Gaussian noise. Then, the stimulus dynamics can be expressed as a Gaussian with zero mean and variance v_w:

$$p(s_t|\widetilde{s}_{t-1}) = \kappa \cdot e^{-\frac{1}{2}\frac{(s_t-\widetilde{s}_{t-1})^2}{v_w^2}}, \tag{4.21}$$

with κ as a normalizing constant. As explained later in this section, a good choice for v_w is important for a successful reconstruction: the variance of the Gaussian noise has to fit the step size of the presented stimulus (see equation 2.1), which is unknown to the reconstruction algorithm.

curves are not cut off at $-\frac{\pi}{2}$ and $\frac{\pi}{2}$ as suggested by figure 4.2. The main advantage of using Gaussians is that the product and sum of two Gaussians is still a Gaussian: a simple curve, completely describable with its mean and variance.

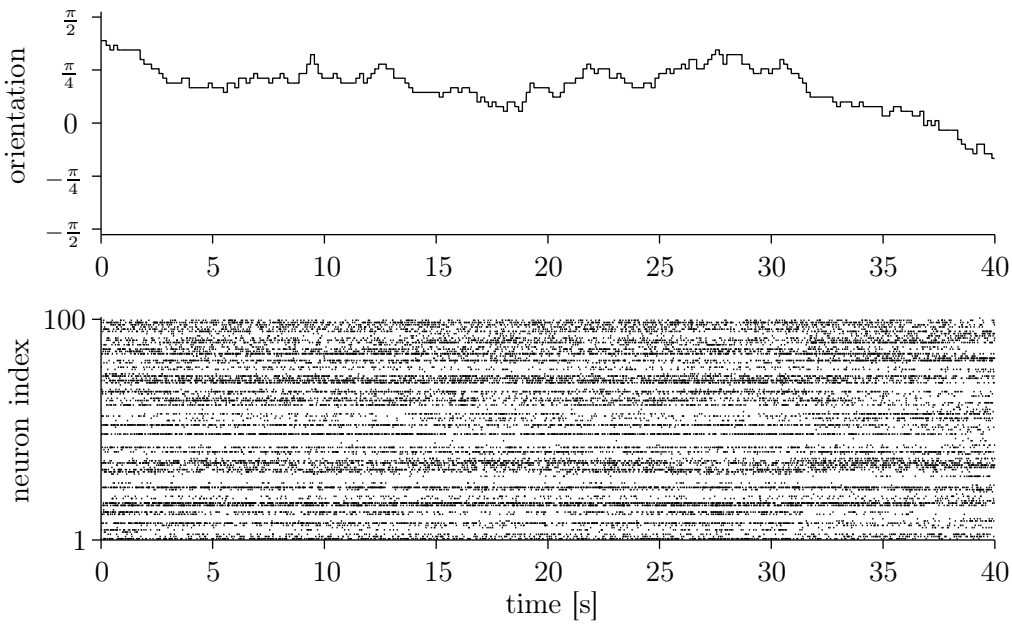

Figure 4.3: Random walk stimulus and computer-simulated neural response. The upper plot shows the trajectory of the random walk stimulus. The lower diagram is a rastergram of 100 neurons with tuning functions as described in figure 4.2. Differently from figure 3.4, no obvious pattern is visible in the spike trains, as the cells are distributed randomly over stimulus space.

Revising the filter model and adopting equations (3.16) and (4.21) we obtain

$$\widehat{s}_t = E(s_t; \text{update} \cdot \text{prediction}) \quad \text{with} \tag{4.22}$$

$$\text{update} = e^{-\sum_{i=1}^{N} f_i(s_t)\Delta t} \cdot \prod_{i=1}^{N} n_i \cdot f_i(s_t) \quad \text{and} \tag{4.23}$$

$$\text{prediction} = \int_{\mathbb{S}} \kappa \cdot e^{-\frac{1}{2}\frac{(s_t - \widetilde{s}_{t-1})^2}{v_w^2}} \cdot p(\widetilde{s}_{t-1}|R_{t-1}, ..., R_1) d\widetilde{s}_{t-1}. \tag{4.24}$$

The process of density propagation in general has already been described in the previous section. To gain a better understanding of the filtering process, we consider the time sampling interval Δt to be small enough to hold exactly one spike or none at all.

If there is no spike, the update term simplifies to $e^{-\sum_{i=1}^{N} f_i(s_t)\Delta t}$ because n_i is zero. If we assume $\sum_{i=1}^{N} f_i(s_t)\Delta t$ to be constant (which is reasonable for larger populations because of the tuning curve's overlap) then the update term shapes a uniform distribution: the probability for each stimulus from stimulus space is the same. Hence, it doesn't have any influence on the estimate $\widehat{s}_t$.

The prediction term, however, is not based on measurement but on knowledge about the old stimulus estimate and its propagation over time. As already stated, tuning curves and stimulus dynamics are Gaussians. Since the sum (convolution) and product of Gaussians yield Gaussians, all densities that occur

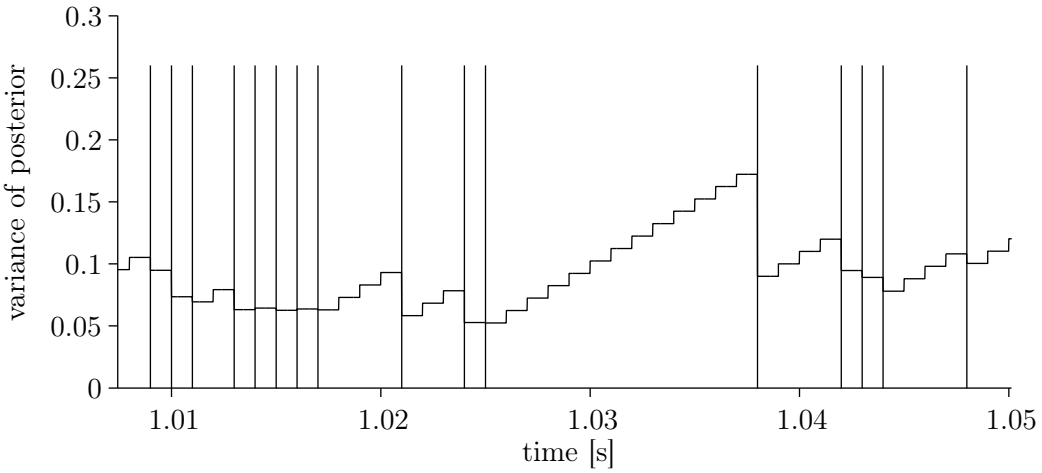

Figure 4.4: Variance of the posterior distribution in temporal close-up. Vertical lines show incoming spikes from any cell. During interspike intervals, the variance increases linearly. The curve is plotted as a stepped function to show the numerical approach with a temporal sampling resolution of 1ms.

in the filter equations are Gaussians and are propagated over time. Particularly, we can say

$$p(s_t|\widetilde{s}_{t-1}) = N(\widetilde{s}_{t-1}, v_w) \quad \text{and} \tag{4.25}$$

$$p(s_{t-1}|R_{t-1}, ..., R_1) = N(\widehat{s}_{t-1}, v_{t-1}), \tag{4.26}$$

with v_{t-1} being the variance of the posterior distribution at timestep $t-1$. Evaluating the prediction term amounts to convolving these distributions, that is, adding their means and variances:

$$\text{prediction} = \int_{\mathbb{S}} N(s_t - \widetilde{s}_{t-1}; 0, v_w) \cdot N(s_t; \widehat{s}_{t-1}, v_{t-1}) d\widetilde{s}_{t-1} \tag{4.27}$$

$$= N(s_t; \widehat{s}_{t-1}, v_{t-1} + v_w). \tag{4.28}$$

While the mean value of the prediction doesn't change with temporal propagation, the variance increases linearly with the time of missing spikes: After k timesteps without any spike, the variance of the posterior distribution will be $v_{t-1} + k \cdot v_w$.

Hence, the Gaussian noise of the random walk models a diffusion process - without a new measurement, the uncertainty about the actual stimulus increases. A new spike provides a meaningful update term that generally decreases the variance of the posterior distribution by contributing the information of its cell's tuning curve (figure 4.4).

Increasing variance of the posterior does not change the estimate $\widehat{s}_t$, since the distribution's mean stays constant[4]. However, by blurring the prior density (which is the posterior density of the previous time step respectively) further

[4]Unless, for $t \to \infty$, the posterior distribution will become completely flat. This means that all knowledge about the stimulus has faded.

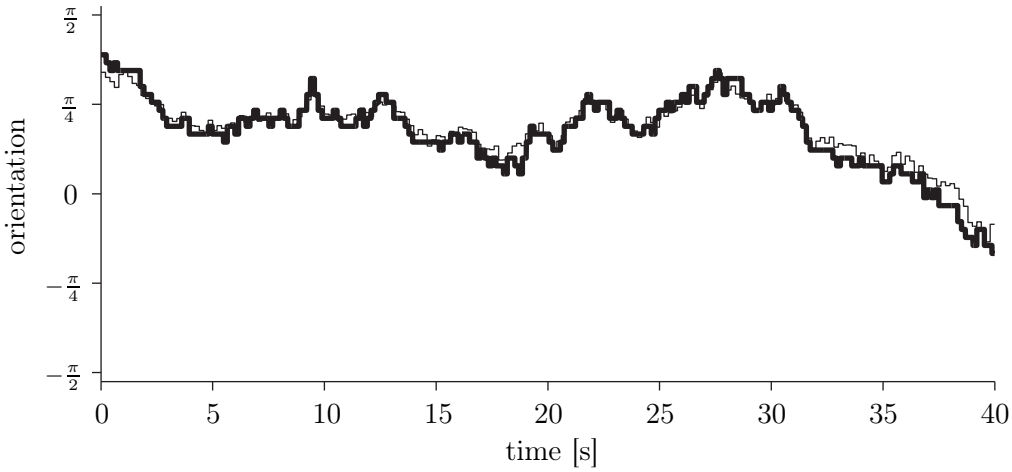

Figure 4.5: Reconstruction results of computer simulation. For comparison, orginal stimulus (thick line) and reconstructed trajectory (thin line) are plotted on top of each other. The random walk variance v_w has been optimised over several trials to 0.1, resulting in a reconstruction mean square error of 0.034 per timestep (1 ms).

and further, the memory of the response history fades. The choice of the random walk variance v_w plays a crucial role in decoding quality, as it controls the speed of fading memory [16]. If v_w is chosen to be small, the stimulus is expected to change slowly over time and the response history fades slowly. Spikes of predecent time steps still have an impact on the actual estimate. This is desirable if the stimulus in fact changes slowly, because a lot of spikes contribute to the stimulus estimation. However, for a fast changing stimulus, a too small v_w acts like a momentum term, delaying the proper response for the actual time step.

Choosing a good v_w is an optimization task that tries to minimise the mean square error of the deviation of the original stimulus trajectory from the estimated one. We define the reconstruction mean square error as

$$\text{RMSE} = \frac{\sum_{e=1}^{E} \sum_{t=1}^{T} (\widetilde{s}_t^e - \widehat{s}_t^e)^2}{E \cdot \Delta t}.$$

The RMSE describes the averaged (over E trials) square error per timestep Δt. A reconstruction result for $v_w = 0.1$ is shown in figure 4.5.

5 Experimental Results

This chapter informs about the empirical results taken from Freiwald et al.'s experiments as described in sections 2.2 to 2.4 and applies them to the reconstruction framework. Assumptions about the neural system of the rat visual cortex made in the theoretical part are tested against the empirical data. As outlined in section 5.3, the reconstruction of the random walk stimulus failed.

5.1 Visual Pathway Latency

To analyse the correlations between a stimulus and its response, the first step is to find out which stimuli belong to which part of the neural response. When a visual stimulus is presented, the signal is processed along the visual pathway (i.e., in the retina and the lateral geniculate of the thalamus) before its representation can be recorded from visual cortex. Hence, it is important to measure the temporal delay between stimulus presentation and cortical response to properly associate stimulus/response pairs.

A technique to do this can be derived from the well-known *spike-triggered average stimulus* $C(\tau)$, which is the average stimulus τ timesteps before a spike is fired. To calculate $C_i(\tau)$ of a cell i, the stimulus at τ timesteps before any spike time t_i^j, $j = 1, ..., n_i$ is averaged over all spikes and trials [5]:

$$C_i(\tau) = \left\langle \frac{1}{n_i} \sum_{j=1}^{n_i} \widetilde{s}_{t_i^j - \tau} \right\rangle . \tag{5.1}$$

Now, we consider the distribution of triggered stimulus values for a fixed τ, i.e. the number of occurances of each stimulus. If none of the stimuli did trigger the response (that is, τ is not the correct temporal delay), the variance of the distribution of frequencies will be low, as all stimuli will occur about the same number of times (as we are averaging over all spikes and trials). In opposite, if τ has been chosen correctly, only some of the stimuli will trigger the response (assumed the cell in question is tuned to a certain subset of stimulus space). The stimulus count variance will increase, because the stimuli the cell is tuned for are much more likely to trigger a response than other stimuli.

A typical result from the set of recorded cells that are used for reconstruction is shown in figure 5.1. A temporal delay of about 85ms (the location of the variance maximum) has been found for other tuned cells as well and has been used for tuning curve calculation and reconstruction.

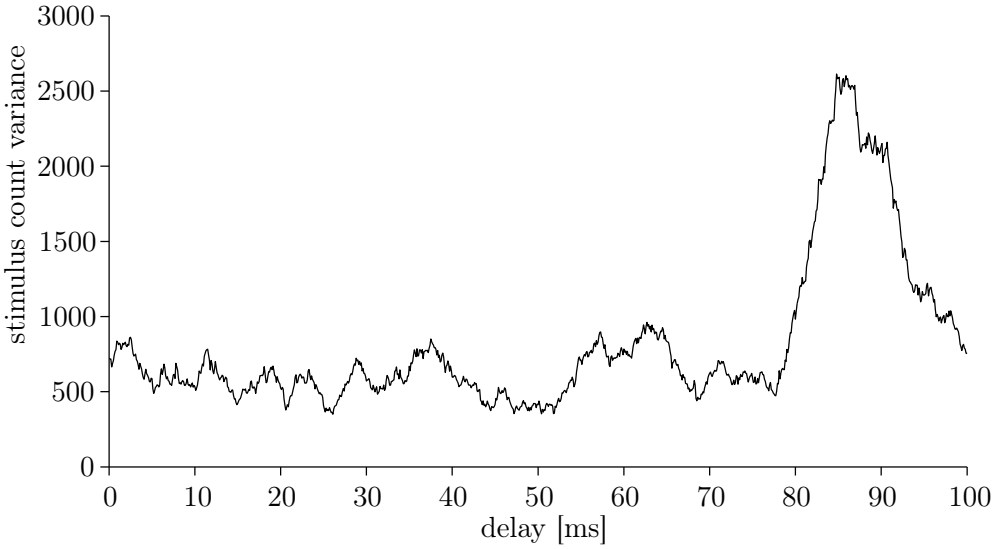

Figure 5.1: Stimulus count variance as a function of temporal delay between stimulus and measured neural response. For a delay around 85ms the variance increases significantly, indicating the correct response latency.

5.2 Review of Assumptions

Two major assumptions about the neural reponse characteristics have been made in our neural model. First, the spike count density is assumed to be a Poisson distribution (see equation (3.7)), and second, simultaneously recorded cells are assumed to react independent of each other (see equation (3.16)). Both can be tested in the light of empirical data.

An important characteristic of a Poisson distribution is that the distribution's mean and variance are equal. This means that the so-called Fano factor, the ratio of mean spike count and spike count variance equals 1 for a given time window. The left subplot of figure 5.2 shows the variance of spike counts for a 141 ms counting period plotted against the mean spike count. Data are from 22 cells recorded under a variety of stimulus conditions. It can be seen that the data points are scattered around the coordinate system's bisection line, which is the prediction of the Poisson model: the average Fano factor of the recorded cells is about 1.

Another test of a presumably underlying Poisson process is to examine the interspike intervals. In a Poisson process, the probability of a certain delay between two events decays exponentially. In the right subplot of figure 5.2, the theoretical curve is plotted on top of the empirical results. The curve calculated from data of a typical cell differs in two major aspects from the presumptive one: on the one hand, there is a frequent occurance of intervals smaller than 25ms. This means that there is an increased tendency for bursting response (fast consecutiveness of action potentials) behaviour which is not predicted by the model. On the other hand, a further deviation from the Poisson model is

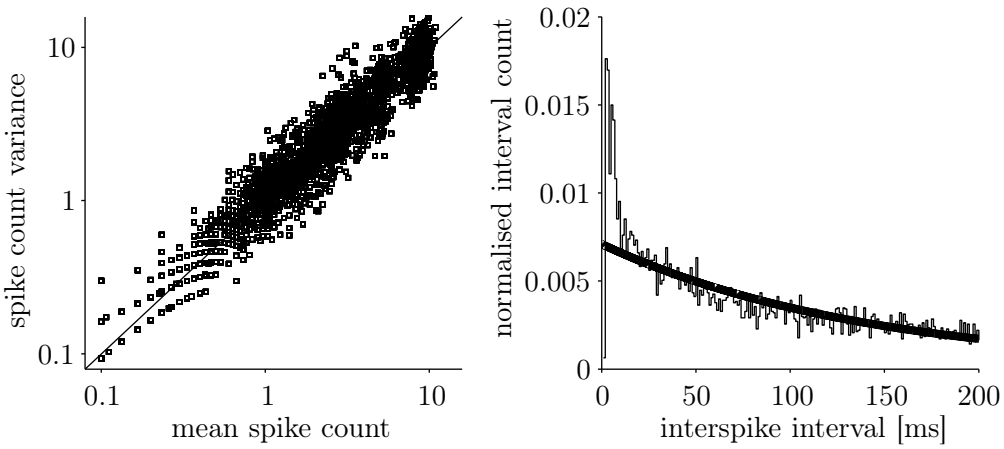

Figure 5.2: Examination of assumptions made by the Poisson model. The left diagram shows the variance of spike counts for a 141ms counting period plotted against the mean spike count. The straight line is the prediction of the Poisson model. Data are from 22 cells recorded under a variety of stimulus condiftions. The right plot shows the interspike interval histogram of a cell firing at 7 Hz (thin line), overlaying the predicted exponential curve (thick line). Bursting and refractory behaviour of the cell are not reproduced by the model.

the decreased number of occurances of intervals smaller than 5ms. This reflects the existence of an absolute and relative refractory period, which the Poisson model cannot account for.

Summarizing, assumptions made by the Poisson spike generator are insufficient for modelling the response characteristics of neurons in the primary visual cortex of rat adequately. However, it is not clear whether the physiological behaviour not mapped by the model is important for coding.

The second assumption to test is the mutual independency of cell responses within a population. The response correlation coefficient $r_{i,j}$ of cell i and j can be calculated using Pearson's product-moment correlation:

$$r_{i,j} = \frac{c_{i,j}}{\sqrt{c_{i,i} \cdot c_{j,j}}}, \tag{5.2}$$

where $c_{i,j}$ is element of the covariance matrix $C = \mathrm{Cov}(X)$ and X is the matrix containing all observations for different stimuli, for one cell per column. Correlation coefficients have been pictured for $N = 22$ simultaneously recorded cells in figure 5.3.

To test the significance of the Pearson product-moment correlation coefficients with respect to the hypothesis of an uncorrelated response behaviour, the coefficients are transformed to $\mathbf{t}$ distribution [19]. If the true correlation within the population described by X is $\rho = 0$, and if the size of the sample, N, on which an observed value of r is based is equal to or greater than 3, then the

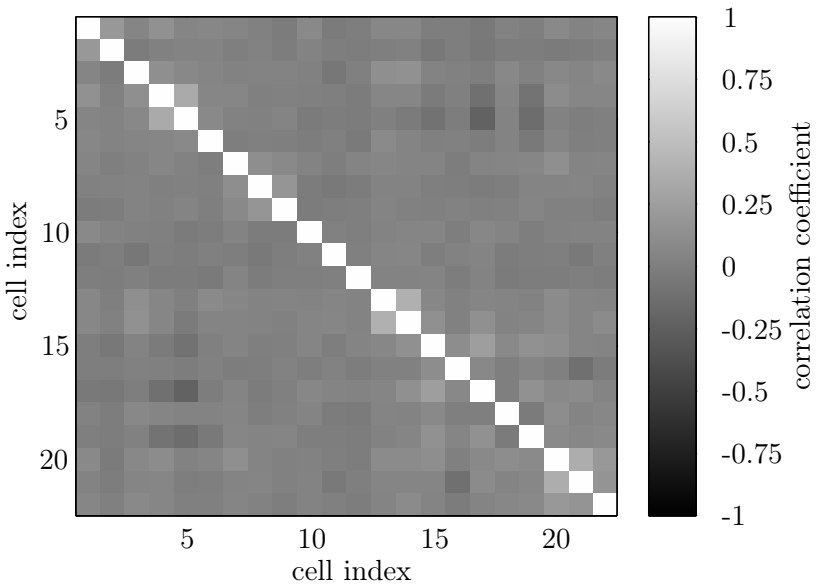

Figure 5.3: Pearson product-moment correlation coefficients of 22 simultaneously recorded cells. No pair of two different cells shows a significantly correlated response. Critical values from [13].

quantity

$$t_{i,j} = \frac{r_{i,j}}{\sqrt{\frac{1-r_{i,j}^2}{N-2}}} \tag{5.3}$$

is distributed approximately as $\mathbf{t}$ with $N-2$ degrees of freedom. Application of this formula to any particular observed sample value of r will accordingly test the null hypothesis that the observed value comes from a population in which $\rho = 0$.

At a 5% level, for none of the cells the null hypothesis could be rejected, i.e. no cell is significantly correlated to any other cell of the subpopulation. However, as can be seen in figure 5.3, both positive and negative correlations do occur. Although not significant with respect to the applied test, a physiological meaning of these correlations cannot be ruled out.

5.3 A Reconstruction of Random Walk

The tuning properties of the 22 cells already mentioned in the previous section turned out to be insufficient for reconstruction of the random walk stimulus. Tuning curves have been calculated using different approaches proposed by Dayan and Abbott [5], including different time window sizes, both overlapping and non-overlapping. The best results emerged from the use of an overlapping time window of 100 ms (see figure 5.4). All tuning curves appear to be too flat (compared to their base frequency) to carry significant orientation tuning. Trying to fit any of these curves with a trigonometric polynomial as described in section 3.2 results in polynomials of degree zero, that is, a flat line located near

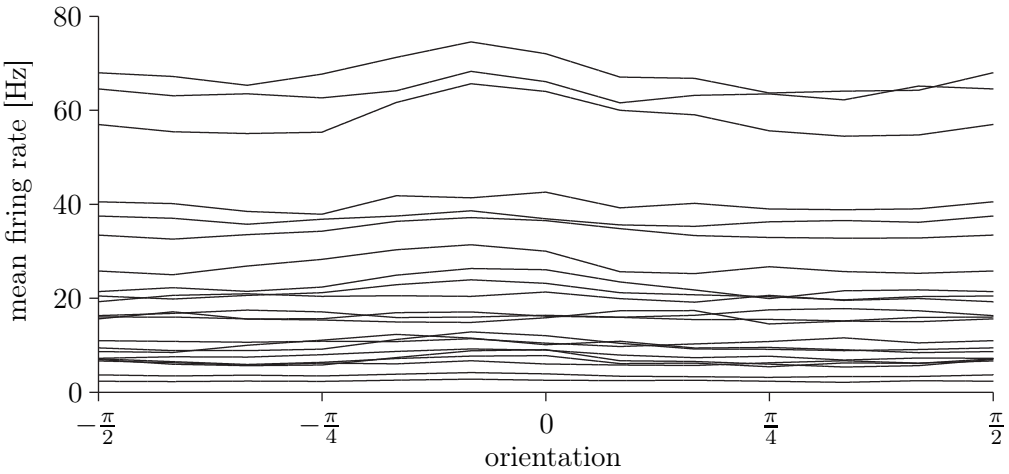

Figure 5.4: Linearly interpolated tuning curves based on empirical data recordeded from 22 cells. All tuning curves show a flat shape, being rather unspecific in their orientation tuning properties.

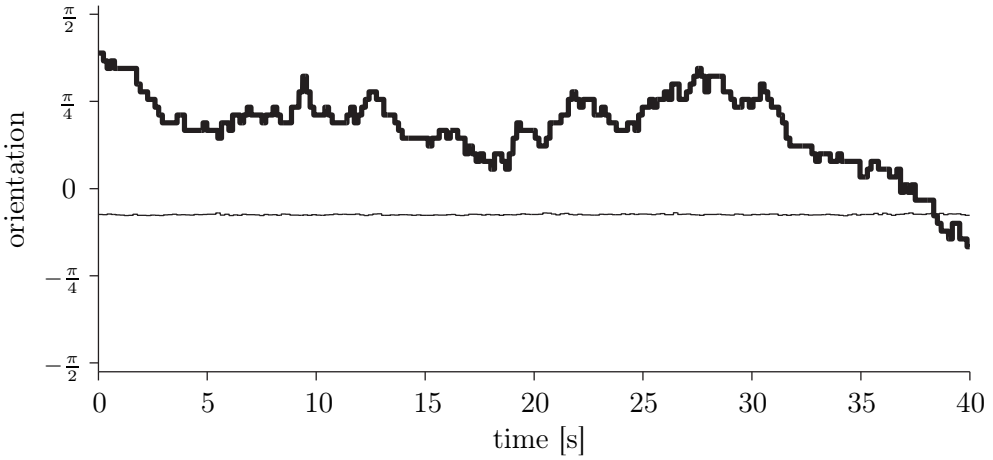

Figure 5.5: Reconstruction results using empirical data recorded from 22 cells. Orginal stimulus (thick line) and reconstructed trajectory (thin line) are plotted on top of each other. The optimal random walk variance $v_w = 0.1$ found in simulation has been used here as well. The reconstruction was unsuccessful and resulted in an amlost flat line at $\widetilde{s} = -0.24$.

the average base frequency.

Therefore, linearly interpolated tuning curves as shown in figure 5.4 were used for reconstruction, avoiding to level out the sparse information available. But still, reconstruction using the measured spiketrains and tuning curves was unsuccessful resulted in an almost flat line at an angle of $\widetilde{s} \cong -0.24$ (figure 5.5).

To test wether the information content of the tuning curves could potentially be sufficient for a reconstruction, we generated an artificial answer an the basis of the same measured tuning curves. Similar to the simulation in section 4.4, this will show if a reconstruction is possible, given the 'optimal' response based on

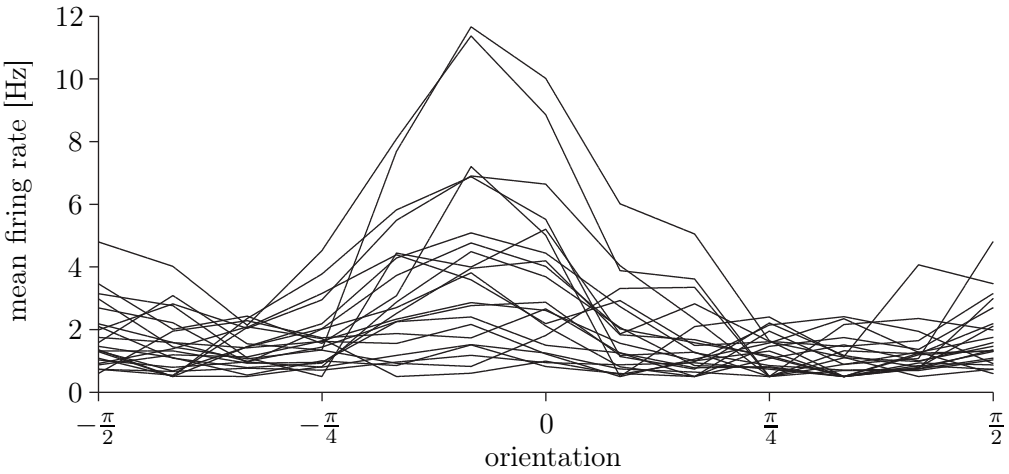

Figure 5.6: Linearly interpolated tuning curves with substracted minimum rate. Most tuning curves have their peak at $\widetilde{s} = -0.24$. Consequently, even a reconstruction with these modified tuning curves leads to the same negative result seen in figure 5.5. As all cells are not or only weakly tuned to $\widetilde{s} = -0.24$, this is the reconstruction result, whatever the stimulus is.

measured tuning curves. The result was negative as well, with the reconstructed trajectory being very similar to that already shown in figure 5.5. This implies that the problem does not lie in the evaluation of wrong data (i.e. the use of wrong parts of the spiketrain for reconstruction), but in the measured tuning curves.

Analysing the tuning curves' shapes, a low reaction index (see section 2.1) is a first characteristic that draws trough all cells. Subtracting the minimum rate from each tuning curve leads to a clear picture of what is the problem: in figure 5.6, it becomes evident that most of the tuning curves have their peak at $\widetilde{s} \cong -0.24$. Consequently, even a reconstruction with these modified tuning curves leads to the same negative result. As all cells are not or weakly tuned to $\widetilde{s} \cong -0.24$, this is the reconstruction result whatever the stimulus is.

In other publications facing similar problems [26, 11], an artificial translation of tuning curves and responses is considered to increase the coverage of stimulus space, which facilitates positive reconstruction results. Here, it is not necessary to do the same, as it could only show the feasablity of the decoding approach. This has already be done in chapter 4, especially section 4.4. The reasons for the negative tuning results have to be discussed in the next chapter.

6 Discussion

In this work, theoretical and empirical results have been presented. On the theoretical side, we have developed a probabilistic framework for coding and decoding of neural activity. The model extends approaches presented by Brown et al. (1998) and Jaekel (2001), since the filtering process as a part of the reconstruction algorithm is no longer based on Gaussian distributions. By contrast, arbitrarily shaped tuning curves (e.g. multimodal) and stimulus dynamics (e.g. nonlinear) can be propagated with the presented algorithm. A reconstruction filter with similar generality has been used so far in one recent work by Gao et al. (2002).

Two significant simplifications have been made: spike events are generated in an inhomogeneous Poisson process and intercell firing correlations have been assumed to be irrelevant for coding and decoding. Both simplifications can be easily replaced by more complex features within the framework. The feasabilty of the framework to reconstruct a random walk stimulus (as an example for an arbitrary dynamic stimulus) has been shown in simulation.

On the empirical side, we have tried to reconstruct the trajectory of a dynamic stimulus from recordings made by Freiwald et al. in the rat primary visual cortex. As already stated in section 5.3, the reconstruction failed because of insufficient measured tuning curves. The reason(s) for failure may lie in at least one of the following three domains - the rat visual cortex itself, the visual stimulation and the measuring chain:

1. Concerning the tuning properties of the rat visual cortex, it is not unlikely that the rat is not the optimal animal for exploring a complex question as the reconstruction of dynamic visual stimuli. It is well known from biology that compared to cats and monkeys, visual processing in rats is less important for their typical behaviour. Missing orientation columns and significantly differing results from independent research teams seem to underline this assumption. Furthermore, the use of comparatively (e.g. to Girmal et al.) light anaesthetics might have led to generally higher spontaneous firing rates [9].

2. It is questionable if the presented stimuli were adequate to test the orientation tuning properties. While both stimuli have similar spatial characteristics compared to previous research it is not clear how long a single orientation should be presented for optimal results (which might be a question of applied anaesthetics as well). In opposite of Girman et al., non-shifting gratings at a fixed spatial phase and frequency have been used for these sets of experiments and further systematic tests might be necessary for the given experimental setup. Moreover, hypermetropia or myopia of the rat's eye has been ignored.

3. It cannot be ruled out that the measuring chain is deficient. This includes every possible error from electrophysiological setup to data conversion on the computer. Besides stimulus-response synchronisation of trigger and measurement, which is susceptible even for small errors, wrong spike sorting might cause large errors: if the rat visual cortex does not have any orientation columns, directly neighboured cells might be tuned for completely different orientations. If the spike sorting process cannot distinguish these cells, the tuning information for single cells will be lost. All recorded 'cells' would then show the average tuning of the area - which has a preference for horizontal stimuli. This could explain the common peak at a nearly horizontal orientation of most tuning curves shown in section 5.3.

Summarizing, the theoretical part is left open for various extensions: the Poisson spike generator could be replaced by a more realistic model to better fit measured spike count and interspike interval distributions; intercell correlations could be considered. On the filter side, the Markov structure could be substituted by a more complex model.

For the experimental setup, next steps should be to verify the functionality of the measurement chain with definitive tests and to systematically test different stimulus parameters to optimise neural responses.

Bibliography

[1] L.F. Abbott and P. Dayan. The Effect of Correlated Variablity on the Accuracy of a Population Code. *Neural Computation*, 11:91–101, 1999.

[2] E. Brown, L. Frank, D. Tang, M. Quirk, and M. Wilson. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *The Journal of Neuroscience*, 18:7411–7425, 15 September 1998.

[3] R. A. Burne, J. G. Parnavelas, and C. S. Lin. Response properties of neurons in the visual cortex of the rat. *J. Neurophysiol.*, 28:1073–1090, 1965.

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

[5] P. Dayan and L.F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press Cambridge, Massachusetts, 2001.

[6] A. Etzold, C.W. Eurich, and H. Schwegler. Tuning Properties of Noisy Cells with Application to Orientation Selectivity in Rat Visual Cortex. *in press*, 2002.

[7] C. W. Eurich and S. D. Wilke. Multidimensional Encoding Strategy of Spiking Neurons. *Neural Computation*, 14:155–189, 2001.

[8] C. W. Eurich and S. D. Wilke. On the functional role of noise correlations in the nervous system. *Neurocomputing (in press)*, 2002.

[9] W. Freiwald. Personal communication, 2002.

[10] W. Freiwald, H. Stemmann, A. Wannig, A. K. Kreiter, U. G. Hofmann, M. D. Hills, G.T.A. Kovacs, D. T. Kewley, J. M. Bower, C. W. Eurich, and S. D. Wilke. Stimulus representation in rat primary visual cortex: multielectrode recordings with micro-machined silicon probes and estimation theory. *Neurocomputing*, pages 44–46, 407–416, 2002.

[11] Y. Gao, M. J. Black, E. Bienenstock, S. Shoham, and J. Donoghue. Probabilistic inference of arm motion from neural activity in motor cortex. *Advances in Neural Information Processing Systems*, 14, 2002.

[12] S.V. Girman, Y. Sauve, and R. D. Lund. Receptive Field Properties of Single Neurons in Rat Primary Visual Cortex. *Journal of Neurophysiology*, 82:301–311, 1999.

[13] G.V. Glass and J.C. stanley. *Statistical methods in education and psychology.* New Jersey: Prentice-Hall, Englewood Cliffs, 1970.

[14] D. Heeger. Poisson Model of Spike Generation. Handout, University of Standford, 5 September 2000.

[15] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. Journal of Computer Vision, 1998.

[16] F. Jaekel. Decoding Neural Activity of Dynamic Stimuli. Bachelor Thesis, University of Osnabrück, 27 September 2001.

[17] E.R. Kandel, J. H. Schwartz, and T. M. Jessel. *Principles of Neural Science.* McGraw Hill, 4th edition, 2000.

[18] K.R. Koch. *Einführung in die Bayes-Statistik.* Springer Verlag, 2000.

[19] R. Lowry. *Concepts and Applications of Inferential Statistics.* Vassar College, New York, 2002.

[20] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C.* Cambridge University Press Massachusetts, 2nd edition, 1993.

[21] F. Rieke, D. Warland, R. de Ruyter van Stevenick, and W. Bialek. *Spikes. Exploring the Neural Code.* MIT Press, Cambridge, Massachusetts, 1998.

[22] S. Russel and P. Norvig. *Artificial Intelligence. A Modern Approach.* Prentice Hall, Upper Saddle River, New Jersey, 1995.

[23] W. Schneider. *Der Kalmanfilter als Instrument zur Diagnose und Schaetzung variabler Parameter in ökonometrischen Modellen.* Physica Verlag, Heidelberg, 1986.

[24] C. Shaw, U. Yinon, and E. Auerbach. Receptive field and response properties of neurons in the rat visual cortex. *Vision Res.*, 15:203–208, 1975.

[25] G. Stanley, F. Li, and Y. Dan. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *J. Neurosci.*, 19(18):8036–8042, 1999.

[26] H. Stemmann. Multizellableitungen im primären visuellen Cortex (V1) der Ratte: Untersuchungen zur Populationskodierung von visuellen Reizen unter Anwendung der Bayesschen Rekonstruktion. Diplomarbeit, Fachbereich Biologie/Chemie der Universität Bremen, July 2002.

[27] H. Stemmann and A. Wannig. Multizellableitungen im Areal V1 von Ratten. Technical report, University of Bremen, 8 October 2001.

[28] N.V. Swindale. Orientation tuning curves: empirical description and estimation of parameters. *Biological Cybernetics*, 78:45–56, 1998.

[29] H. Tanizaki. *Nonlinear Filters. Estimation and Applications*. Springer Verlag, 1993.

[30] H. Tanizaki. Nonlinear and Non-Gaussian Space Modeling using Sampling Techniques. *Ann. Inst. Statist. Math.*, 53:63–81, 2001.

[31] C. van der Togt, V. A. F. Lamme, and H. Spekreijse. Functional connectivity within the visual cortex of the rat shows state changes. *European Journal of Neuroscience*, 10:1490–1507, 1998.

[32] E. Wan and R. van der Merwe. The unscented Kalman filter for nonlinear estimation. in Proc. of IEEE Symposium 2000 (AS-SPCC), Lake Louise, Alberta, Canada, October 2000.

[33] J. et al Wessberg. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, (408):361–365, 2000.

[34] Z. Wiesenfeld and E. E. Kornel. Receptive fields of single cells in the visual cortex of the hooded rat. *Brain Res.*, 16:823–827, 1976.

[35] R. A. Wilson and F. Keil, editors. *The MIT Encyclopedia of Cognitive Sciences*. Massachusetts Institute of Technology, 1999.

[36] W. Wu, M. J. Black, Y. Gao, Bienenstock, M. E., Serruya, and J. P. Donoghue. Inferring Hand Motion from Multi-Cell Recordings in Motor Cortex using a Kalman Filter. *SAB'02-Workshop on Motor Control in Humans and Robots: On the Interplay of Real Brains and Artificial Devices*, October 2002.

[37] K. Zhang, I. Ginzburg, B. McNaughton, and T. Sejnowski. Interpreting neuronal population activity by reconstruction: A unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79:1017–1044, 1998.

Erklärung

Hiermit erkläre ich, dass ich diese Arbeit selbständig angefertigt und nur die angegebenen Hilfsmittel benutzt habe.

Holger Bringmann, 24.10.2002, Osnabrück