

Eberhard Karls Universität Tübingen
Mathematisch-Naturwissenschaftliche Fakultät
Wilhelm-Schickard-Institut für Informatik

Bachelor Thesis Cognitive Science

Neural modulation by temporal predictability in the mouse primary visual cortex

Sindy Löwe

August 27, 2015

Reviewer

Prof. Dr. Hanspeter A. Mallot
Department of Biology
University of Tübingen

Supervisor

Dr. Masataka Watanabe
University of Tokyo

Sindy Löwe

Neural modulation by temporal predictability in the mouse primary visual cortex

Bachelor Thesis Cognitive Science
Eberhard Karls Universität Tübingen

Period: April 30 - August 30

Abstract

The influential theoretical framework of predictive coding states that low-level sensory areas subtract top-down predictions from bottom-up signals to calculate prediction errors, which are used for perceptual inference. So far, only few single-unit studies have directly tested predictions of this model. Here, we measured responses of single neurons in mouse primary visual cortex (V1) and asked how temporal predictability affects the processing of a visual stimulus.

Head-fixed mice were free to run or sit on a disc in front of a grey screen, on which we flashed a visual stimulus (black square, 100ms duration) in the absence of any task. To manipulate stimulus predictability, we presented a visual cue (black square, 100ms duration, outside of the receptive field) shortly before the target stimulus appeared. In the standard condition (80% of the trials), the stimulus onset asynchrony (SOA) between cue and target was fixed at 500ms. The remaining 20% of the trials comprised of a non-cued condition, where only the target was presented which was therefore non-predictable, shifted SOA conditions, where the SOAs deviated from the standard value (± 100 ms, ± 200 ms), and a target-absent condition, where only the cue was given.

We recorded from V1 neurons simultaneously across cortical layers and found that stimulus predictability significantly decreased V1 responses, mainly in the initial transient component of neuronal activity.

Although no robust statistical significance could be reached, we observed a tendency that the standard SOA condition leads to the strongest decrease, resulting in a "U-shaped" curve of evoked activity versus SOAs. This suggests the existence of a neural mechanism that attenuates evoked activity dependent on previously exposed timing. Since we did not find neurons in V1 showing interval coding behavior, our results are in favor of a higher-level feedback mechanism that acquires the timing offline during the exposure epoch and provides temporal prediction online.

To investigate the local neuronal circuitry, we expressed Archaeorhodopsin (Arch) in V1 parvalbumin-positive inhibitory interneurons. We suppressed interneuron activity by shining yellow light (590nm, 500ms duration) at the time of stimulus onset. As a result, the attenuation effect vanished as predictable and non-predictable stimuli evoked similar neuronal responses. This suggests that local inhibitory neurons play a crucial role in mediating the effect of predictability.

Furthermore, we replaced the visual cue by an auditory cue (pure tone, 5kHz, 20ms duration). Here, overall effects of predictability and target omission remained the same.

One interesting observation is that we found no modulation of activity in

the target-absent condition. This finding indicates that prediction and sensory input are multiplied, rather than subtracted, which is inconsistent with the predictive coding model.

In conclusion, our results provide further evidence for prediction based attenuation of evoked neural activity at the single neuron level, and indicate the involvement of feedback mechanisms. However, the lack of an omission response contradicts the state of the art framework of predictive coding, which assumes subtractive neural processing for obtaining prediction errors. It raises an interesting question whether previously observed attenuation effects of prediction are actually due to predictive coding, and if so, how the framework can be modified to fit our observations, which suggest multiplicative modulation.

Acknowledgements

First of all, I would like to express my sincere gratitude to my supervisor Dr. Masataka Watanabe for the continuous support, for his patience, motivation, and immense knowledge.

Besides my supervisor, I would like to thank my reviewer Prof. Dr. Hanspeter A. Mallot for his support and for giving me the opportunity to write my thesis within the context of this interesting and challenging project.

My sincere thanks also goes to Dr. Laura Busse and Dr. Steffen Katzner at the Werner Reichardt Centre for Integrative Neuroscience (CIN), who provided me an opportunity to join their team as a research assistant, and who gave access to the laboratory and research facilities.

Last but not the least, I would like to thank my family and friends for their support and understanding.

Contents

List of Figures	vii
List of Tables	ix
1 Introduction	1
2 Methods	7
2.1 Brief explanation of the applied techniques	7
2.1.1 Extracellular Recordings	7
2.1.2 Optogenetics	8
2.2 Surgical Preparation	9
2.3 Experiment	10
2.3.1 Experimental setup	10
2.3.2 Procedure and Stimuli	11
2.4 Analysis	14
2.4.1 Unit extraction and spike sorting	14
2.4.2 Preprocessing	14
3 Results	15
3.1 Visual cue experiment	15
3.1.1 Standard SOA cued condition against non-cued condition	15
3.1.2 Temporally shifted conditions	19
3.1.3 Target-absent condition	23

3.1.4	Standard cued against non-cued condition with optogenetic stimulation	27
3.2	Auditory cue experiment	29
3.2.1	Standard SOA cued condition against non-cued condition	29
3.2.2	Target-absent condition	31
4	Discussion	35
4.1	General effect	36
4.2	Origin and neuronal circuitry of modulation	36
4.3	Error calculation	38
4.4	Conclusion	39
	Bibliography	41

List of Figures

2.1	Experimental Setup	10
2.2	Identified receptive field	12
2.3	Time course for visual cue experiments	13
2.4	Time course for auditory cue experiments	13
3.1	Single neuron activity in cued and non-cued condition (visual cue)	16
3.2	Evoked activity in cued and non-cued condition (visual cue)	17
3.3	Normalized population activity in cued and non-cued condition (visual cue)	18
3.4	Single neuron activity in shifted SOA conditions (visual cue)	21
3.5	Normalized population activity for shifted SOA conditions, cued and non-cued condition (visual cue)	22
3.6	Single neuron activity in the target-present and target-absent condition (visual cue)	24
3.7	Evoked activity in target-absent and target-present condition (visual cue)	25
3.8	Normalized population activity in target-absent and target-present condition (visual cue)	26
3.9	Single neuron activity in the cued and non-cued condition using optogenetic stimulation (visual cue)	27
3.10	Normalized population activity in cued and non-cued condition using optogenetic stimulation (visual cue)	28

3.11	Single neuron activity in the cued and non-cued condition (auditory cue)	29
3.12	Evoked activity in cued and non-cued condition (auditory cue) .	30
3.13	Normalized population activity in cued and non-cued condition (auditory cue)	31
3.14	Single neuron activity in the target-present and target-absent condition (auditory cue)	32
3.15	Evoked activity in target-absent and target-present condition (auditory cue)	33
3.16	Normalized population activity in target-absent and target-present condition (auditory cue)	34
4.1	Responses of interval coding neurons	37

List of Tables

3.1	Results of single neuron testing on shifted SOA conditions (visual cue)	19
3.2	Results for testing population activity on shifted SOA conditions (visual cue)	20

Chapter 1

Introduction

We operate in a world of sensory uncertainty, due to various factors that limit the reliability of sensory information. 3D objects of the world are mapped onto a 2D image in our retina, neuronal noise is introduced in early stages of sensory coding and there are other structural constraints on neuronal representations and computations, such as the blind spot of the retina, where the optic nerve and blood vessels leave the eye (Gregory & Cavanagh, 2011). As a result, our brains have to effectively deal with the resulting uncertainty to generate a congruent representation of the world. This has led to the idea that perception is a process of inference, where sensory inputs are combined with prior knowledge (Von Helmholtz, 1867).

A modern approach on how the brain implements probabilistic inference is the model of predictive coding (Kawato *et al.*, 1993; Mumford, 1992; Rao & Ballard, 1999). It postulates that inferences of higher-level areas, coined "predictions" within the model, are compared with incoming sensory information in lower areas through cortical feedback.

In the outline of Rao & Ballard (1999) the model is based on a hierarchical network, where each level attempts to predict the activity at the next lower level via feedback connections. This prediction is subtracted from the bottom-up sensory representation, which results in an error that is sent back to the higher level via feedforward connections. There, the error is used to correct the estimate of the bottom up sensory signal. Thus, the neuronal activity in lower areas should decrease when predictions are accurate and as a result redundancy is reduced by removing the predictable. These prediction and error-correction cycles exist concurrently throughout the hierarchical network and therefore top-down information consisting of predictions influence lower-level estimates and bottom-up information based on the sensory input influence higher-level estimates. Higher-level modules predict and estimate the responses of several lower-level modules and therefore work on a bigger spatial scale than lower levels. As a consequence, effective receptive field (RF)

size increases progressively in the ascending hierarchy. Ultimately the goal of a system becomes optimally estimating the hidden causes for each input on each level and on a longer time scale, designing an efficient internal model of natural inputs.

Now, if we want to apply the model of predictive coding to the visual cortex, we may assume that predictions of the expected neuronal activity in V1 are carried from a higher to a lower area (i.e. V2 to V1) via feedback connections, while the residual activity in V1 that was not predicted is forwarded back to V2 (Mumford, 1992). There the model, which yields the prediction, is updated and the loop starts over again. Theoretically, the visual cortex would provide the structure needed for the predictive coding framework, being hierarchically organized and having reciprocal corticocortical connections (Felleman & Van Essen, 1991). Additionally, rendering predictions is heavily facilitated by the properties of natural images, where neighboring pixel intensities tend to be correlated, as well as the values of a given pixel over time.

On top of this, Rao & Ballard (1999) showed that a three-level hierarchical network, constructed according to the hierarchical model described above and trained on several thousand natural image patches, could elicit some extra-classical RF effects with similar properties to those found in the visual cortex. For example, 28 out of their 32 model error-detecting neurons showed endstopping properties, where firing is enhanced when a dark bar lies within the RF of the neuron, but suppressed when it extends beyond it. Most interestingly, disabling feedback connections eliminated this effect in most of these neurons in a manner qualitatively similar to that observed in cortical neurons after inactivation of layer 6 (Bolz & Gilbert, 1986). Thus, their simulation not only suggests that some extra-classical RF effects are an emergent property of the visual cortex using a predictive coding framework, but also provides strong evidence for the applicability of the predictive coding model to the neocortex.

To further test whether the model of predictive coding is applicable to the biological brain, several studies have been conducted. A number of functional magnetic resonance imaging (fMRI) studies have shown, that predictable stimuli evoke smaller responses in early visual areas (Murray *et al.*, 2006; Harrison *et al.*, 2007; Summerfield *et al.*, 2008). Murray *et al.* (2006), for example, used individual elements to show that when they formed coherent shapes, activity in higher object-processing areas increases while activity in the primary visual cortex (V1) is reduced concurrently. Since the predictable attributes of the stimuli are not represented in early sensory cortex, local adaptation cannot explain the suppression effect. Therefore this exemplary study provides strong evidence for the notion, that feedback connections exhibit inhibitory effects on lower level areas as predictability improves.

Meanwhile, there are attempts to relate classical effects, such as mis-

match negativity, to predictive coding. Mismatch negativity can be found both in vision and audition in electro-encephalography (EEG) and magneto-encephalography (MEG) (Tales *et al.*, 1999; Garrido *et al.*, 2009) and consists of enhanced neuronal responses to deviant stimuli, which violate sensory predictions established by a regular stimulus sequence. These enhanced responses are thought to reflect an inability of higher cortical areas to predict, and thereby suppress, the activity of populations encoding prediction error (Wacongne *et al.*, 2011).

However, in this case, it is possible that the suppression of predicted stimuli is due to adaptation effects or repetition suppression. To test for this possibility there is an interesting variant of the mismatch paradigm, where the expected stimulus is omitted, rather than replaced with another stimulus. Again extensive responses locked to the absence of a predictable stimulus can be found both in the auditory and visual modality when measured with EEG or MEG (Czigler *et al.*, 2006; Yabe *et al.*, 1997; Raji *et al.*, 1997). These omission responses are difficult to explain by adaptation effects, but fit quite well in the predictive coding framework assuming that stimulus-evoked brain activity represents the difference between a sensory signal and its top-down prediction. Then, if the predicted stimulus is omitted, the evoked activity should reflect the pure prediction signal that is left. Omission responses therefore might constitute a critical test of the predictive coding framework. Even so, there are no single unit studies addressing this topic yet.

Since most of the described studies so far made use of EEG, MEG or fMRI, they only provide circumstantial evidence that the cortex works according to the predictive coding framework. Thus, dedicated electrophysiological studies are awaited for further compelling evidence at the scale of single neurons, neuronal circuits and interareal interactions.

To our knowledge there is only one study that directly validated the existence of predictive coding using electrophysiology. Meyer & Olson (2011) exposed monkeys to pairs of images presented in fixed sequences, so that each leading image became a strong predictor for its corresponding subsequent picture. After prolonged training, they measured neuronal responses in the inferotemporal cortex (ITC) to image sequences that either obeyed or violated the learned transitional rules. They found that inferotemporal neurons respond much stronger to unpredicted transitions than to predicted ones. The study demonstrates that single units are affected by experience-based predictions and react accordingly when the prediction is invalid.

We may summarize that the end result of predictive coding, an attenuation of neural activity for predicted sensory stimuli, has been demonstrated at various levels, but whether predictive coding is the actual mechanism behind the effect and if so, how it is actually implemented in the brain remains an open

question. Mouse vision would provide an optimal platform with its modern measurement and manipulation techniques to pinpoint the neural mechanism of such effects.

In this study we examined the effect of predictability at the single neuron level in the mice primary visual cortex (V1) using extracellular recordings. To manipulate stimulus predictability, we exposed the animals to a temporal sequence of two stimuli, such that the first stimulus ("cue") became a predictor of the second ("target"). In the standard condition, where the animal was exposed to the sequence repeatedly in the order of thousands of trials, the stimulus onset asynchrony (SOA) was set to 500ms. By omitting the cue we obtained a control condition, which revealed the activity for a non-predictable stimulus. According to the framework of predictive coding and previous experimental studies examining the model, the predictable stimulus should evoke an attenuated response compared to the non-predictable one.

Furthermore, we tried to identify the origin of such prediction effects. For this purpose, we examined the temporal nature of modulation, namely, whether it is a temporally broad adaptation effect that takes effect anytime post-cue for a certain period, or an effect that incorporates precise timing acquired during exposure to the fixed temporal sequence. To this end, we introduced temporally shifted targets, which were preceded by a cue, but had different SOAs in relation to the standard condition. We hypothesized that the primary visual cortex alone would not be sufficient to store offline and calculate online the temporal predictions at the scale of few hundreds of milliseconds, and if we do observe such effects, we may conclude that the modulation originates from higher levels (but see Hangya & Kepecs, 2015, and discussion).

Additionally, to further investigate the origin of modulation, we adopted temporal cues of different sensory modalities, namely, visual and auditory.

Meanwhile, as an attempt to pinpoint the local neuronal circuitry of modulation, we conducted experiments with optogenetic manipulation. In the mammalian brain, extrinsic corticocortical connections are exclusively excitatory but may mediate selective attenuation via their termination in inhibitory neurons (Bastos *et al.*, 2012). Therefore, manipulation of local inhibitory mechanisms should alter the attenuation effect, if inhibitory feedback is involved as assumed in the predictive coding model. We suppressed activity of parvalbumin-positive interneurons (PV+), a subtype of neurons that constitutes the majority of GABAergic inhibitory neurons in mouse V1, using optogenetic stimulation. Combined with future work manipulating other types of neurons, it would not only provide interesting insights into the network structure employed for attenuation effects by prediction, but would also answer the question, whether and how predictive coding is implemented in the neural architecture.

In relation, we tested whether the omission of a predicted stimulus elicits a neuronal response. According to the standard predictive coding framework, attenuation of sensory evoked responses occurs due to the subtraction of top-down prediction from bottom-up input. Meanwhile, various neuronal modulatory effects are known to be multiplicative (e.g., Reichardt & Rosenblith, 1961; Gabbiani *et al.*, 2002; Salinas & Thier, 2000; Mehaffey *et al.*, 2005). Thus, omission responses to predicted stimuli observed at the scale of single neurons would offer clear evidence whether previously observed neural attenuation effects are subtractive as in the predictive coding framework, or multiplicative as observed in many other neural modulatory effects, and provide insights into the question whether predictive coding exists in the biological brain in the form of previously proposed models.

This bachelor thesis is divided into four parts: In the first part, the introduction, we have outlined the current state of research on the topic of predictive coding, which questions still remain open and how we will try to answer them. In the next section, the methods, we will provide general information about the techniques applied in this study before describing the surgical preparations, experimental setup and data analysis. The third section contains the results of the data analysis for our experiment, which are then discussed and embedded into current research in the fourth and final part, the discussion.

Chapter 2

Methods

Experiments were performed on adult mice. Procedures were in accordance with the standards of the Society for Neuroscience and the German Law for Protection of Animals and were approved by the local authorities following appropriate ethics review.

2.1 Brief explanation of the applied techniques

Before describing our surgical preparations, experimental setup and stimuli, we want to outline some general information about the techniques used in this study.

2.1.1 Extracellular Recordings

There is both large- and small-scale organization in the brain and different functions take place on higher and lower levels. In order to understand computational rules of the brain studying neural activity at the network (sub-millimeter) level is considered most suitable (Petsche *et al.*, 1984). Emerging properties of neuronal interactions can be measured within this domain and models such as predictive coding can be developed. A straightforward approach to simultaneously record from a large number of neurons are extracellular recordings (Nadasdy *et al.*, 1998).

In contrast to intracellular recordings, such as voltage clamp, the tip of the microelectrode used for extracellular recordings is immediately adjacent to, but outside of a neuron. In this way, current fields generated by action potentials in a cell in close proximity can be detected as small voltage deflections (typically $0.1 - 1mV$).

There are several advantages to this technique. First of all, it provides a more intact preparation *in vivo* and thus brain regions or neurons can be studied in their normal complement of inputs and targets within their natural milieu. Moreover the cells themselves are usually not severed or damaged and have developed normally in an intact organism. These factors can add credibility and reduce caveats to results obtained *in vivo*. In addition several experimental questions require an intact organism and cannot be pursued *in vitro*. For example, effects of sensory inputs typically must be examined in the intact brain.

Still, there are some disadvantages to extracellular recordings. For a start there may be confounds, such as immobilization stress that could alter the normal electrophysiological responses. Additionally when using a multi-unit recording, where the activity seen on an individual electrode can be generated from several nearby neurons, the results are more difficult to interpret. The obtained spikes have to be matched with multiple single units, which can be a fault-prone process (cf. Bloom *et al.*, 1995).

2.1.2 Optogenetics

At a basic level, the nervous system can be thought of as a highly complex electrical circuit. A variety of pump and channel proteins maintain a negative membrane potential in the resting neuron by controlling the flow of ions across the membrane. If activation signals cause positively-charged ions to flow into the cell via these channel proteins, depolarizing the membrane potential and if sufficient crossing the threshold potential, an action potential is triggered. This causes sodium ions to rapidly flow into the neuron, which reverses the voltage inside the cell. Consequently a chain reaction of sodium-ion influx is initiated that propagates down the length of the axon, eventually stimulating or inhibiting the production of electrical impulses in neighboring neurons (cf. Bamberg, 2010).

Historically, there have been some attempts to artificially manipulate neurons and thereby trigger or suppress action potentials. Microelectrodes, for example, have been used for direct stimulation in neurophysiological studies, but exhibit a poor resolution limit. "Caged" neurotransmitters that remain inactive unless triggered by laser illumination and chemically modified photo-switchable ion channels notably improved the precision for functional studies, but possessed limited possibilities for application.

On account of this, the discovery of channelrhodopsin, an algae protein that can act as an "on" switch by allowing the influx of positive neurons in response to illumination with blue light, revolutionized the field. Later more proteins were found, for example halorhodopsin which could also provide an

”off” switch, as it triggers the influx of negatively-charged chlorine ions in response to yellow light. There are various techniques to introduce these proteins into the target cells. In this study, for example, we injected a viral vector associated with the Archaeorhodopsin (Arch) protein into V1 of transgenic mice expressing Cre recombinase under the control of the parvalbumin promoter (PVCre). As a consequence the respective light sensitive ion channels express within the parvalbumin-positive interneurons (PV+ interneurons) and when shining yellow light on V1, activity of the PV+ interneurons will be suppressed. GABAergic interneurons (GABA, γ -aminobutyric acid), such as the PV+ interneurons, may be a minority cell type since they only represent about 10–20% of all neurons, but nevertheless they are vital for normal brain function as they regulate the activity of principal neurons. PV+ interneurons themselves are fast-spiking cells that contribute to feedback and feedforward inhibition and can convert an excitatory input signal into an inhibitory output signal within a millisecond (Hu *et al.*, 2014). Thus, suppressing PV+ interneurons will lead to an overall enhancement of activity in V1.

2.2 Surgical Preparation

In two PVCre mice, anesthesia was induced by 3% isoflurane and maintained during surgery using the same chemical (1–2%). Animal temperature was kept at 37°C via a feedback controlled heating pad (WPI). A custom designed headpost was mounted to the skull using dental cement (Tetric EvoFlow, Ivoclar Vivadent). A reference wire was placed into the cerebellum and a ground wire under the skin. During surgery, buprenorphine (0.1mg/kg sc) was used for analgesia, atropine (0.3mg/g sc) to reduce bronchial secretions, baytril (5mg/kg sc) was given as antibiotic and the eyes were protected with an ointment (Bepanthen). The following 3 days the same antibiotics and longer lasting analgesics (carprofen, 5mg/kg sc) were injected.

Additionally, they were injected with the adeno-associated viral vector AAV9 (Vector Core, University of Pennsylvania) through a small craniotomy 3mm lateral to the midline and 1.5mm in front of the anterior margin of the transverse sinus. A glass pipette connected to a Picospritzer III (Parker) was slowly lowered to approximately 750 μ m below the brain surface, and a total of 100–150nl of virus was injected every 100 μ m while gradually retracting the pipette. The pipette was left in place for an additional five minutes to allow viral diffusion, and the craniotomy was covered with Kwik-Cast (WPI). Neurophysiological recordings with optogenetic stimulation were performed at least three to four weeks after virus injection.

Thereafter, the animals were habituated to being head-fixed and placed on a disc, before a craniotomy ($< 1mm^2$) located 3mm lateral to the mid-

line and 1.1mm in front of the anterior margin of the transverse sinus (Wang *et al.*, 2011) was performed, which was sealed with Kwik-Cast (WPI) after each recording session. Recordings typically were continued on consecutive days as long as unit isolation remained of high quality.

2.3 Experiment

2.3.1 Experimental setup

Head-fixed mice were free to sit or run on a disc, while recording from V1 were obtained. Extracellular recordings were performed with 32-channel silicon probes (Neuronexus, $A1 \times 32 - 5\text{mm} - 25 - 177 - A32$) that were inserted perpendicular to the brain surface and lowered to approximately $1000\mu\text{m}$. The setup was enclosed with a dark curtain and there was no ambient light other than that of the stimulus monitor. In order to monitor the exact stimulus onset, a photodiode was placed in the upper left corner of the screen and a black square was displayed below this diode synchronous to each presented stimulus (Fig. 2.1).



Figure 2.1: Experimental Setup: the head-fixed mouse was free to sit and run on a disc with a distance of 10cm to the screen while extracellular recordings were performed in V1

Eye position was monitored under infrared illumination using a camera (Guppy AVT; frame rate 50Hz) coupled to a zoom lens (Navitar Zoom 6000). Although eye movements were occasionally observed, their amplitude was small and less of a concern, because we were able to consistently evoke visual responses to the target stimulus, while the visual cue stimulus was placed far enough away not to be affected by eye movements.

For conducting the optogenetic experiments, photostimulation was carried out using a fiber-coupled light-emitting diode (LEDs; Doric lenses) with a

center wavelength of $590nm$, driven by an LED driver (LEDD1B, Thorlabs). The fiber was mounted on a manual manipulator and positioned less than $1mm$ from the craniotomy. LED light was delivered with a light intensity of $224mW/mm^2$ measured at the tip of a $1000\text{-}\mu m$ -diameter core.

2.3.2 Procedure and Stimuli

Stimuli were presented on a calibrated liquid crystal display (LCD) monitor (mean luminance $46cd/m^2$) $10cm$ in front of the animal’s eyes. To estimate RF position before conducting the experiment, we used custom software (EXPO; <https://sites.google.com/a/nyu.edu/expo/home>) for stimulus presentation and mapped ON and OFF subfields of RFs using a sparse noise stimulus (Liu *et al.*, 2010). This stimulus consisted of white or black squares flashed for $180ms$. Following Liu *et al.* (2010), we fitted ON and OFF subfields separately with a two-dimensional Gaussian:

$$f(x, y) = \frac{A}{2\pi ab} \exp\left(-\frac{x'^2}{2a^2} - \frac{y'^2}{2b^2}\right) \quad (2.1)$$

where A is the maximum amplitude, a and b are half-axes of the ellipse, and x' and y' are transformations of the stimulus coordinates x and y , taking into account the angle θ and the coordinates of the center (xc, yc) of the ellipse. For each electrode contact, we computed a single RF center by averaging coordinates of the best-fit ON and OFF subfield (explained variance $> 70\%$). Subsequent target stimuli were presented at the average RF center across recording sites (Fig 2.2).

For stimulus presentation in the main experiment we used MATLAB (MathWorks, Inc.) with the Psychophysics Toolbox Version 3 (PTB-3). Throughout the experiment a grey background (luminance $46cd/m^2$) was maintained. The target stimuli consisted of a black square that was displayed for $100ms$ at the position of the average RF center identified earlier. To manipulate stimulus predictability, we presented a cue shortly before the stimulus appeared. Depending on the experiment the cue either consisted of a visual or an auditory stimulus.

The visual cue comprised a single black square stimulus similar to the target, but positioned outside the RF. We made sure the visual cue did not elicit any neuronal response in any of the recording sites in order to prevent effects caused by refractory periods of the neurons. The visual cue was displayed for $100ms$ with a stimulus onset asynchrony (SOA) of $500ms$ to the target (Fig. 2.3).

In the auditory experiment two cues were presented in each trial each consisting of a $5kHz$ pure tone played for $20ms$. Between the first and the second

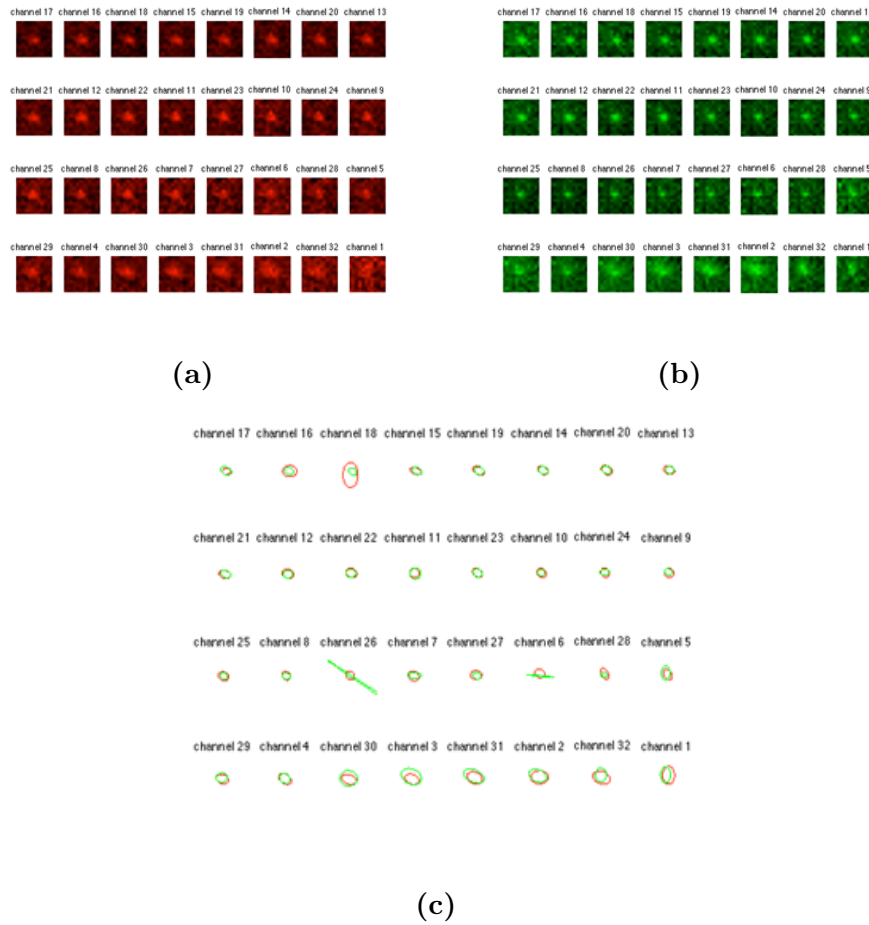


Figure 2.2: Identified receptive fields (RFs) over all 32 channels of the electrode: **(a)**, **(b)** color maps of ON **(a)** and OFF **(b)** subfields **(c)** RF contours depicting the fitted outline of the spike subfields

cue, as well as between the second cue and the target the SOA was set to 250ms (Fig. 2.4).

Irrespective of whether the auditory or visual cue was used, the procedure was designed as follows: First, we repeated the cue-target temporal sequence about 500 times, to make sure the animal learns the transition between cue and target. After this time, the standard condition was reduced to comprise 80% of the trials and other conditions for testing purposes were intermingled. First, we included a non-cued condition, where only the target was displayed, to test whether the neural response to the predictable stimuli differs to that to the non-predictable one. Then, in the target-absent condition solely the cue was presented and the target itself was omitted, to see whether there is any change in neuronal activity associated with the timing of the missing target. In the visual cue experiment we further introduced four temporally shifted




Figure 2.3: Time course for visual cue experiments: the visual cue stimulus was presented for 100ms and with a SOA of 500ms the target stimulus was displayed for the same duration. Gray dots represent the possible timings of the target in the temporally shifted conditions




Figure 2.4: Time course for auditory cue experiments: the auditory cue stimuli were played for 20ms, the first at 0ms, the second at 250ms and at 500ms the visual target stimulus was displayed for 100ms

conditions, where the SOA between the cue and target was deviated from the standard value ($\pm 100ms$ and $\pm 200ms$), to test the temporal precision of the prediction. In all trials the intertrial interval (ITI) was randomized between one and two seconds.

Furthermore, we tried to manipulate the effects of predictability by optogenetic stimulation. In this experiment, only the standard SOA cued and non-cued conditions with the visual cue were used, to test whether the overall effect remains the same. Associated with the visual stimuli, yellow light was linearly ramped up to full intensity and maintained for 500ms to avoid optogenetic offset contamination of the data.

2.4 Analysis

2.4.1 Unit extraction and spike sorting

Wideband extracellular signals were digitized at $30kHz$ (Blackrock microsystems) and analyzed using the NDManager software suite (Hazan *et al.*, 2006). In order to isolate single neurons from linear arrays, we grouped neighboring channels into five equally sized "virtual octrodes" (eight channels per group with two channel overlap for the 32 channel probes). Using an automatic spike detection threshold (Quiroga *et al.*, 2004) multiplied by a factor of 1.5, spikes were extracted from the high-pass filtered continuous signal for each group separately. The first three principal components of each channel were used for semi-automatic isolation of single neurons with KlustaKwik (Henze *et al.*, 2000). Clusters were manually refined with Klusters (Hazan *et al.*, 2006). We assigned each unit to the contact with the largest waveform, and gave a single score based on the subjective rating of the manual sorter, the firing rate, the cleanness of the refractory period, and the stability over time.

2.4.2 Preprocessing

In order to match the neural activity with single trials, we determined their onset timings using the information provided by the photodiode and their respective condition by a trial schedule that had been created and used for each experiment. Then we sorted single spikes in $10ms$ bins for each trial and condition. To convert the spike count to Hertz, we divided it by the bin width (i.e. $10ms$) and if several trials were combined by the number of trials. If a single bin within a trial exceeded a spike rate of $200Hz$, the trial was discarded to eliminate potential noise. In the next step, we determined the visually responsive neurons by conducting a two-sample t-test ($\alpha = 0.05$) on single units with firing rates higher than $4Hz$ comparing the baseline before stimulus onset with the evoked response for the target (time window: 0.0 to 0.25 seconds from stimulus onset) in the standard condition in each trial and subdivided these neurons into two groups: group one ("positively responsive neurons") evoked a stronger response during the observed time window compared to the baseline, whereas the second group ("negatively responsive neurons") had a significant suppression effect elicited by the stimulus. This way we obtained an overall number of 109 positive and 12 negative responsive neurons out of 398 single units. Furthermore we normalized the activity of each channel separately by subtracting the baseline measured before stimulus onset and dividing the result by the mean of the five highest values within the $10ms$ bins in the standard condition.

Chapter 3

Results

We wanted to examine, whether predictability modulates neuronal activity in mice V1 by presenting a fixed temporal sequence of two stimuli ("cue" and "target") such that the cue became a predictor for the target. Furthermore we tried to identify the locus of such effects and the involved neural circuitry. Lastly, we were interested in the neuronal response elicited by the omission of the predicted target.

3.1 Visual cue experiment

In the visual cue experiment, cue and target comprised of the same stimulus, each presented for 100ms. The target was displayed inside the RF, while the cue was positioned outside, in order to avoid mere adaptation effects.

3.1.1 Standard SOA cued condition against non-cued condition

First we investigated, whether predictability attenuates neuronal responses to visual stimuli. As a control to the standard cued condition described above, we introduced a non-cued condition in which only the target was presented.

In the majority of neurons observed during the visual cue experiment, the target stimulus clearly evoked a smaller response in the cued condition compared to the non-cued condition. In figure 3.1 we can see the activity of two example neurons in the cued and non-cued condition, where the peristimulus time histograms (PSTHs) reveal a clear attenuation effect in the cued condition.

To determine the timing at which this effect is maximized, we plotted the evoked activity of each neuron in the cued condition against the non-cued

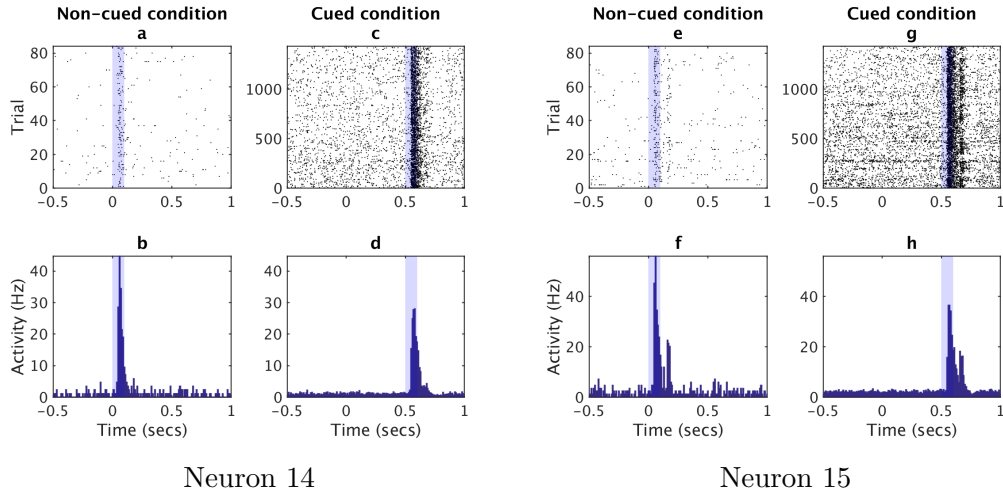


Figure 3.1: Single neuron activity in the cued and non-cued condition: (a),(c) show rasterplots for the non-cued and cued condition for neuron 14 (b),(d) show the respective peristimulus time histograms (PSTHs) (e)-(h) respective plots for neuron 15. blue shades indicate stimulus presentation

condition for four different time windows ($0.00-0.05s$, $0.05-0.10s$, $0.10-0.15s$ and $0.15-0.20s$ from stimulus onset; Fig. 3.2). One can see that attenuation is strongest for the second time window ($0.05-0.10s$), where most of the dots are shifted downwards. Therefore we chose this time window for the following statistical tests.

Using a single-tailed two-sample t-test ($\alpha = 0.05$) 20 out of 45 visually responsive neurons showed a significant reduction of neuronal firing rate for the cued compared to the non-cued condition. This effect was also significant at the population level (single-tailed, paired-sample t-test, $n = 45$, $p < 0.05$) and can be seen in figure 3.3 where we plotted the normalized population activity as a function of time in trials in which the cue was either present or absent (see chapter 2.4.2 for details on the normalization process). The figure reveals a clear difference between population activity in the cued and non-cued condition in the initial transient component of neuronal activity.

Our results therefore suggest, that predictability attenuates the magnitude of neuronal responses.




Figure 3.2: Evoked activity (Hz) in cued and non-cued condition for different time windows, where red dots represent positively responsive neurons, blue dots negatively responsive neurons: **(a)** time window: 0.00 - 0.05s **(b)** time window: 0.05 - 0.10s **(c)** time window: 0.10 - 0.15s **(d)** time window: 0.15 - 0.20s




Figure 3.3: Normalized population activity in cued (blue curve) and non-cued (red curve) condition. Target stimulus was presented from 0.00 to 0.10s. Shaded regions indicate 95% confidence interval

3.1.2 Temporally shifted conditions

As a second point, we tried to identify the temporal nature of modulation, namely, whether it is a simple adaptation effect that generally exists post-cue, or an effect that is dependent on the exposed timing. For this purpose, we introduced four temporally shifted targets, which were preceded by a cue, but had different SOAs (300ms, 400ms, 600ms and 700ms) in relation to the standard condition (500ms).

In some neurons, the target stimulus clearly evoked a smaller response in the standard SOA compared to the shifted SOA conditions. In figure 3.4 we can see the activity of two example neurons in the temporally shifted and standard SOA conditions, as well as in the non-cued condition. The PSTHs reveal a stronger attenuation effect in the standard cued condition compared to the shifted SOA conditions.

To examine, how many neurons showed an attenuation effect based on precise timing, we tested 45 visually responsive neurons on the same time window relative to stimulus onset ($0.05s - 0.10s$) using a single-tailed two-sample t-test ($\alpha = 0.05$) for each shifted SOA condition separately. The results can be seen in table 3.1 and show that only few neurons show a significant difference in activity for the temporally shifted compared to the standard SOA condition.

SOA	number of neurons with significant modification
300ms	4
400ms	4
600ms	3
700ms	0

Table 3.1: Results of single neuron testing on shifted SOA conditions, time window: 0.05s to 0.10s from stimulus onset, a total amount of 45 neurons was tested


In figure 3.5 we plotted the normalized population activity as a function of time in trials of the shifted SOA conditions, the standard SOA condition and the non-cued condition (see chapter 2.4.2 for details on the normalization process). One can see the tendency, that the standard SOA condition evokes the smallest response among the cued conditions.

Next, we tested whether the observed tendencies at the neural population level is statistically significant (single-tailed, paired-sample t-test, $n = 45$, $\alpha = 0.05$). The results, shown in table 3.2, depict that only the 300ms SOA condition resulted in a significantly larger response compared to the standard SOA condition.

SOA	test result	p-value
300ms	1	< 0.05
400ms	0	0.06
600ms	0	0.45
700ms	0	0.75

Table 3.2: Results for testing population activity on shifted SOA conditions compared to standard SOA condition: time window: 0.05s to 0.10s from stimulus onset, n=45

In conclusion, we did not find a significantly robust effect for precise timing. The lack of statistical significance may be due to an insufficient number of trials in the temporally shifted conditions, and thus future work is required to reach a firm conclusion.



Neuron 13

Neuron 23

Figure 3.4: Single neuron activity in shifted SOA conditions compared to standard SOA condition and non-cued condition: (a),(c),(e),(g),(i),(k) show rasterplots for the cued conditions with SOAs 300ms, 400ms, 500ms (standard), 600ms and 700ms and the non-cued condition for neuron 13 (b),(d),(f),(h),(j),(l) show the respective PSTHs (m)-(x) respective plots for neuron 23. blue shades indicate stimulus presentation




Figure 3.5: Normalized population activity for shifted and standard SOA conditions and non-cued condition: standard SOA cued-condition - blue curve, non-cued condition - red curve, SOA 300ms - violet curve, SOA 400ms - cyan curve, SOA 600ms - turquoise curve, SOA 700ms - yellow curve. Shaded regions indicate 95% confidence interval

3.1.3 Target-absent condition

Next, we examined whether the absence of the predicted target elicits any neuronal response. For this purpose, we have included a target-absent condition, where only the predicting cue was presented.

In the activity of single neurons no change could be seen. In figure 3.6 the activity of four example neurons for the target-absent as well as the target-present (standard) condition is shown. One can see, that while the target itself evokes a neuronal response, there is no change in baseline in the target-absent condition.

Using a two-tailed two-sample t-test ($\alpha = 0.05$, $n = 107$) no neuron could be found with a statistically significant change in activity locked to the timing of the missing target (0.05-0.10s from virtual stimulus onset) when compared to baseline. In figure 3.7 the evoked activity of each neuron in the target-absent condition is plotted against the evoked activity in the cued target-present condition. This scatterplot reveals that while there is a considerable variety in response strength to the target, little difference can be seen in the responses of the neurons for the target-absent condition.

There was also no significant change in activity on the population level when tested with a two-tailed, paired-sample t-test ($p = 0.32$). In figure 3.8 we plotted the normalized population activity as a function of time in trials in which the target was either present or omitted.

In summary, it can be said that the target-absent condition did not evoke significant neuronal responses at the timing of the missing target.

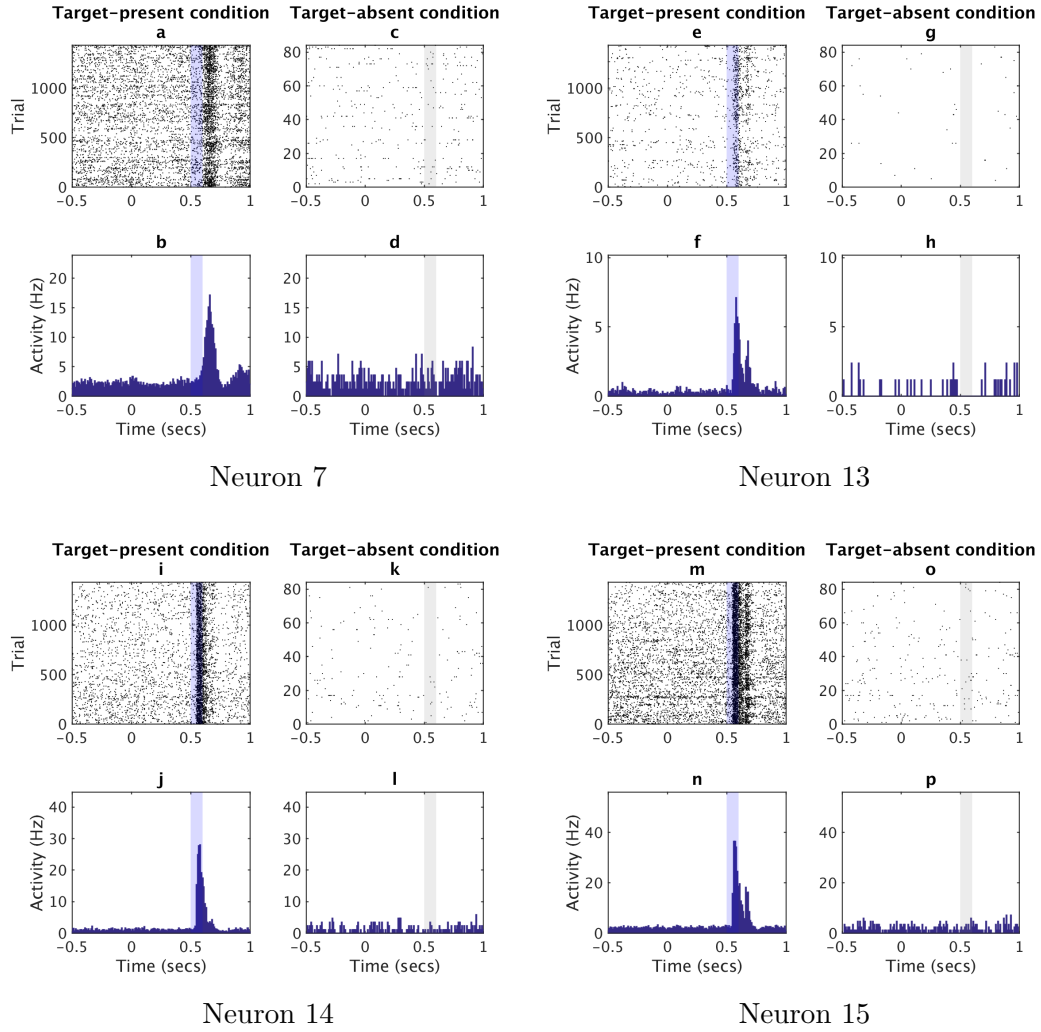


Figure 3.6: Single neuron activity in the target-present and target-absent condition: (a),(c) show rasterplots for target-present and target-absent condition for neuron 7, (b),(d) show the respective peristimulus time histograms (PSTHs) (e)-(h) respective plots for neuron 13 (i)-(l) respective plots for neuron 14 (m)-(p) respective plots for neuron 15. blue shades indicate stimulus presentation, gray shades the respective timing of stimulus omission




Figure 3.7: Evoked activity (Hz) in target-absent and target-present condition: red dots represent positively responsive neurons, blue dots negatively responsive neurons, black dots visually unresponsive neurons




Figure 3.8: Normalized population activity in target-absent (black curve) and target-present (blue curve) condition. Shaded regions indicate 95% confidence interval

3.1.4 Standard cued against non-cued condition with optogenetic stimulation

In order to further investigate the involved circuitry of the attenuation effect, we used optogenetic manipulation to suppress the activity of parvalbumin-positive GABAergic interneurons while comparing the evoked activities between cued and non-cued conditions.

No difference could be seen between the magnitude of evoked activity in the cued and non-cued condition of single neurons. Figure 3.9 shows the activity of two example neurons in the cued and non-cued condition under the influence of optogenetics. Apparently there is no difference in response strength in these two conditions.

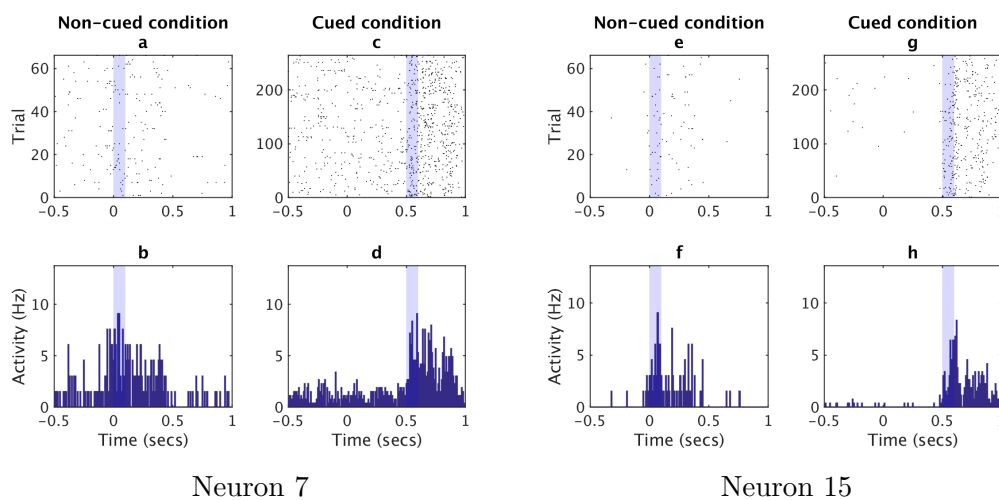


Figure 3.9: Single neuron activity in the cued and non-cued condition using optogenetic stimulation: (a),(c) show rasterplots for non-cued and cued condition for neuron 7 (b),(d) show the respective PSTHs (e)-(h) respective plots for neuron 15. blue shades indicate stimulus presentation

Because the evoked activity seems to be shifted in time between cued and non-cued conditions, we only tested the peak values for significant differences. None out of 18 visually responsive neurons had a significant difference in peak activity between cued and non-cued condition when tested with a two-tailed two-sample t-test ($\alpha = 0.05$).

There was also no significant change in peak activity of the population response when tested with a two-tailed, paired-sample t-test ($n = 18$, $p = 0.85$). In figure 3.10 we plotted the normalized population activity as a function of time in trials in which the cue was either present or omitted (see chapter 2.4.2 for details on the normalization process). Evidently the activities seem

to be shifted in time, but nevertheless show not much difference in their overall activity pattern.




Figure 3.10: Normalized population activity in cued (blue curve) and non-cued (red curve) condition using optogenetic stimulation. Target stimulus was presented from 0.00 to 0.10s. Shaded regions indicate 95% confidence interval

In conclusion, it can be stated that the attenuation effect vanished when the activity of PV+ interneurons is suppressed by optogenetic stimulation, indicating that the modulation of neural activity relies on a local inhibitory neural mechanism.

3.2 Auditory cue experiment

In the auditory cue experiment, the cue comprised of a pure tone (5 kHz, 20ms duration) which was played twice before the visual target of 100ms duration. In the standard condition the SOAs (first cue to second cue, second cue to target) were fixed to 250ms.

3.2.1 Standard SOA cued condition against non-cued condition

First we investigated, whether predictability attenuates neuronal responses to visual stimuli. As a control to the standard cued condition described above, we introduced a non-cued condition in which only the target was presented.

In the majority of neurons observed during the auditory cue experiment, the target stimulus clearly evoked a smaller response in the cued condition compared to the non-cued condition. In figure 3.11 we can see the activity of two example neurons in the cued and non-cued condition, where the peristimulus time histograms (PSTHs) reveal a clear attenuation effect in the cued condition.

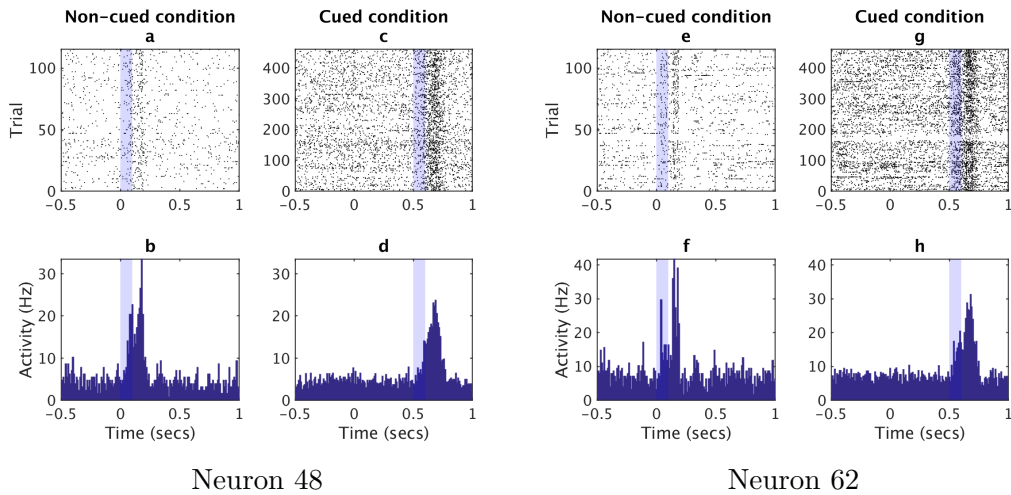


Figure 3.11: Single neuron activity in the cued and non-cued condition: (a),(c) show rasterplots for non-cued and cued condition for neuron 48 (b),(d) show the respective peristimulus time histograms (PSTHs) (e)-(h) respective plots for neuron 62. blue shades indicate stimulus presentation

To determine the timing at which this effect is maximized, we plotted the evoked activity of each neuron in the cued condition against the non-cued condition for four different time windows ($0.00-0.05s$, $0.05-0.10s$, $0.10-0.15s$ and $0.15-0.20s$ from stimulus onset; Fig. 3.12). One can see that attenuation

is strongest for the first time window (0.00 – 0.05), where most of the dots are shifted downwards. Therefore we chose this time window for the following statistical tests.




Figure 3.12: Evoked activity (Hz) in cued and non-cued condition for different time windows, where red dots represent positively responsive neurons, blue dots negatively responsive neurons: **(a)** time window: 0.00 - 0.05s **(b)** time window: 0.05 - 0.10s **(c)** time window: 0.10 - 0.15s **(d)** time window: 0.15 - 0.20s

Using a single-tailed two-sample t-test ($\alpha = 0.05$) 15 out of 34 visually responsive neurons showed a significant reduction of neuronal firing rate for the cued compared to the non-cued condition. This effect was also significant at the population level (single-tailed, paired-sample t-test, $n = 34$, $p < 0.05$) and can be seen in figure 3.13 where we plotted the normalized population activity as a function of time in trials in which the cue was either present or absent (see 2.4.2 for details on the normalization process). The figure reveals a clear difference between population activity in the cued and non-cued condition in the initial transient component of neuronal activity.

Our results therefore suggest, that predictability attenuates neuronal response strength.




Figure 3.13: Normalized population activity in cued (blue curve) and non-cued (red curve) condition. Target stimulus was presented from 0.00 to 0.10s. Shaded regions indicate 95% confidence interval

3.2.2 Target-absent condition

Next, we examined whether the absence of the predicted target elicits any neuronal response. For this purpose, we have included a target-absent condition, where only the auditory cue was presented.

In the activity of single neurons no change could be seen. In figure 3.14 the activity of four example neurons for the target-absent as well as the target-present (standard) condition is shown. One can see, that while the target itself evokes a neuronal response, there is no change in baseline in the target-absent condition.

Using a two-tailed two-sample t-test ($\alpha = 0.05$, $n = 93$) no neuron could be found with a significant change in activity locked to the timing of the missing target (0.00-0.05s from virtual stimulus onset) compared to baseline. In figure 3.15 the evoked activity of each neuron in the target-absent condition is plotted against the evoked activity in the cued target-present condition. This scatterplot reveals that while there is a considerable variety in response strength to the target, little difference can be seen in the responses of the neurons for the target-absent condition.

There was also no significant change in activity at the population level

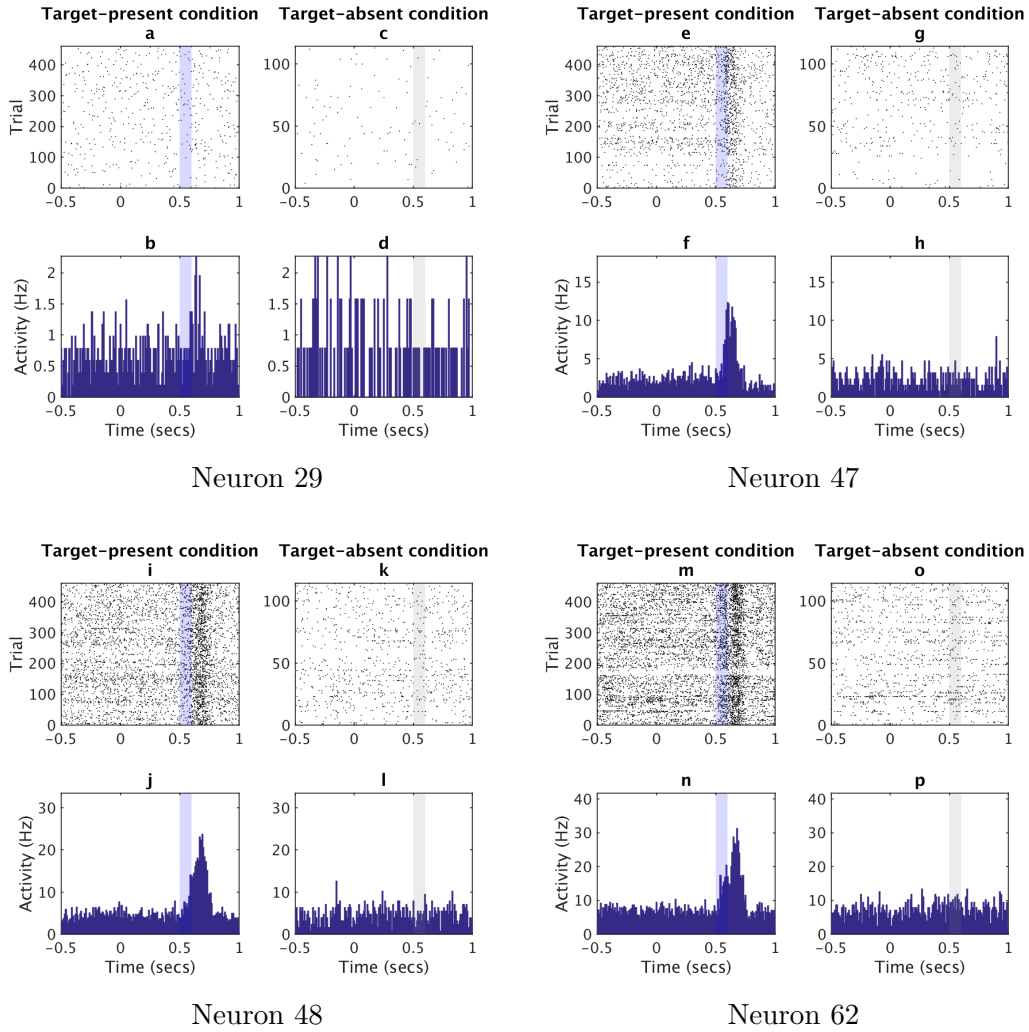


Figure 3.14: Single neuron activity in the target-present and target-absent condition: (a),(c) show rasterplots for target-present and target-absent condition for neuron 29, (b),(d) show the respective peristimulus time histograms (PSTHs) (e)-(h) respective plots for neuron 47 (i)-(l) respective plots for neuron 48 (m)-(p) respective plots for neuron 62. blue shades indicate stimulus presentation, gray shades the respective timing of stimulus omission

when tested with a two-tailed, paired-sample t-test ($p = 0.056$). In figure 3.16 we plotted the normalized population activity as a function of time in trials in which the target was either present or omitted (see 2.4.2 for details on the normalization process).

In summary, it can be said that the target-absent condition did not evoke neuronal responses.




Figure 3.15: Evoked activity (Hz) in target-absent and target-present condition: red dots represent positively responsive neurons, blue dots negatively responsive neurons, black dots visually unresponsive neurons




Figure 3.16: Normalized population activity in target-absent (black curve) and target-present (blue curve) condition. Shaded regions indicate 95% confidence interval

Chapter 4

Discussion

In this study we investigated the effect of predictability on stimulus evoked neuronal activity in mice V1 by exposing the animals to a temporal sequence of stimuli, such that the first ("cue") became a temporal predictor of the second ("target"). Together with a non-cued condition, we compared neuronal activity between predictable and non-predictable stimuli. Following the framework of predictive coding, we assumed that the predictable stimulus would elicit a weaker response than the non-predictable one, and indeed found an attenuation of neuronal response, when the stimulus was predictable.

Furthermore we tried to specify the locus and neuronal circuitry of this predictability effect. For this purpose, we introduced shifted SOA conditions to test the temporal precision of prediction. Although a tendency was noticeable that the standard SOA condition elicited the smallest response, maximization of effect at this timing was not statistically significant.

Additionally we used both visual and auditory cues, to test whether predictability effects can be generated across different sensory modalities. Our results show qualitatively similar effects in both experiments.

In order to investigate the involved circuitry, we tested the influence of optogenetic stimulation on the effect. Suppression of PV+ interneurons led to the disappearance of the attenuation effect, indicating their crucial involvement in mediating predictions.

Lastly we examined the neuronal activity linked to the omission of the target stimulus. In the predictive coding framework, calculation of prediction error is subtractive and thus, negative modulation should occur at the timing of the predicted sensory stimulus. But there is also a number of evidence suggesting that many calculations in the brain are actually multiplicative. In contrast to the predictive coding model, if the attenuation of evoked activity is based on multiplicative computations, the absent target should not evoke any activity, since the prediction would be multiplied by zero, the non-existing

sensory input. The omission response at the single neuron level could therefore provide interesting insights into the computational mechanisms of the brain in regard to stimulus prediction. Interestingly, we did not observe any neuronal response associated with the timing of the missing target.

In the following section, we will discuss our results in regard to these topics.

4.1 General effect

Comparing the activity in the cued and non-cued condition, we observed that temporal predictability manipulates neuronal activity in V1. In both the visual and auditory cue experiment, the initial transient component of neuronal activity to the target stimulus was significantly suppressed when prediction was possible. This finding is consistent with the model of predictive coding, which postulates that an accurate prediction results in a smaller error which in turn corresponds to the lower neuronal activity we measured in V1.

4.2 Origin and neuronal circuitry of modulation

Following this, we tried to specify the origin and neuronal circuitry of the suppression effect of predictability.

First, we could show that the visual and auditory cue induced qualitatively similar effects of predictability. This might be seen as an indicator towards the involvement of higher-level computations. However, there is the possibility of a fast low-level crossmodal connection to V1 that provokes prediction processes within V1.

Then, in order to investigate the temporal nature of the predictability effect, namely, whether it is a simple adaptation effect that exists post-cue, or an effect that is dependent on the precise timing, we introduced shifted SOA conditions. Unfortunately, we did not find robust significant differences, although a general trend is observable that a greater shift of the SOA from the standard value resulted in a reduction of the attenuation effect. In case this specific property of the effect reaches statistical significance after optimizing the experimental protocol (e.g. increasing trials for the shifted conditions), we may claim that the attenuation effect is not based on a mere low-level neural adaptation effect, but requires a precise mechanism that keeps track of the exposed timing.

Although interestingly, according to Hangya & Kepecs (2015), it has been shown that V1 could provide such timing mechanisms. Based on the work of

Shuler & Bear (2006) and Liu *et al.* (2015), they state that V1 is capable of predicting reward timing using a population code that provides an estimate of both the mean and temporal uncertainty about anticipated reward time. They further describe three types of interval coding neurons that were found in V1 and their particular reward-timing activity. Some neurons seem to show sustained activation from cue presentation to the expected time of reward, while others show inhibition during the same time period. A third group exhibits a firing rate peak at the expected reward time (Fig. 4.1). Based on this, we take the position that if the predictability effect is based on a within V1 mechanism, we should find similar activity patterns in the neurons we observe. However, this was not the case. We did not find any neurons that showed the described interval coding behavior. From this, we infer that the attenuation effect of predictability is less likely to be generated by a V1 local mechanism, but rather mediated by higher-level feedback mechanisms.




Figure 4.1: Responses of V1 neurons that anticipate reward time from the visual cue (green dashed line) to the reward (blue dashed line). Three types of interval coding neurons were identified: sustained activated, sustained inhibited and peak firing neurons (figure taken from Hangya & Kepecs (2015))

The results of the experiment using optogenetic stimulation do not contradict the above position. Since extrinsic corticocortical connections are exclusively excitatory, inhibitory feedback effects can only be mediated by employing local inhibitory mechanisms (Bastos *et al.*, 2012). Thus, if the predictability effect is caused by higher-level feedback, it should be eliminated when inhibitory mechanisms in the lower-level area are silenced. We tested this possibility by suppressing the activity of PV+ interneurons, a type of inhibitory neurons that constitutes the majority of GABAergic neurons in mouse V1, using optogenetic stimulation. Indeed the difference between neuronal response strengths to predictable and non-predictable stimuli disappeared under this condition,

revealing that PV+ interneurons play a substantial role for predictability effects in V1. However, this does not provide inevitable evidence for top-down influences, since an intra-areal suppression effect would have to rely on local inhibitory mechanisms as well.

In summary, it can be stated that our results indicate the involvement of higher-level feedback in the prediction effect. In order to produce the temporal precision we observed, an accurate mechanism would be required. This might be provided by V1 itself, but we did not find any neurons showing interval coding as in the work of Hangya & Kepecs (2015) and therefore conclude, that the mechanism is likely to be provided by higher-level sources. Furthermore, the results of the auditory cue experiment and the optogenetic manipulation experiment do not contradict with this view.

It would be interesting to test whether shutting down higher-level feedback influences the attenuation effect. This might be achieved by cooling down the respective brain areas or by means of optogenetics and could probably provide the most compelling evidence for the necessity of these connections.

4.3 Error calculation

In the model of predictive coding, prediction error is calculated by subtracting the prediction from bottom-up sensory input. Omitting a predicted stimulus would therefore result in an error equivalent to the size of prediction, due to its subtraction from zero (i.e. the absence of sensory input). This error should then be seen as a neuronal response temporally locked to the missing stimulus.

In our results, however, we did not find any such activity in the target absent condition. Not one single unit had a significant change in activity at the timing of the missing target. This suggests that the attenuation effect is multiplicative, rather than subtractive. Following predictive coding, the incorrect prediction would be multiplied by zero (i.e. the missing sensory input) and therefore the resulting error equals zero as well and no change in activity arises.

Multiplicative modulation, or gain control, is consistent with a number of studies on computational neuroscience (for review, see Salinas & Thier, 2000). McAdams & Maunsell (1999), for instance, manipulated the response amplitudes of neurons independently of their selectivity or RF characteristics (so called "gain modulation") by guiding the attention of the examined rhesus monkeys. Then they tested, whether the resulting changes in V4 orientation-tuning curves could be described as a multiplication and indeed this was the case. In another study, Mel (1993) also revealed a possible mechanism for

neurons to calculate approximate multiplications based on dendrites containing NMDA-type synaptic conductances and voltage-dependent sodium (Na), potassium (K) and calcium (Ca) channels.

Aside from that, Buzsáki & Mizuseki (2014) argue that most distributions in the brain, such as synaptic weights, firing rates of individual neurons and the number of synaptic contacts between neurons, are skewed and typically lognormal, rather than bell-shaped as it is often assumed. Based on the mathematical rule that the product of many independent variables has an approximately lognormal distribution, it is therefore concluded, that the majority of interactions in the brain are multiplicative and synergistic. Another argument for multiplicative computations in the brain is that they would significantly increase the capabilities of single neuron computations (Mehaffey *et al.*, 2005).

Considering these works on the topic, the most suitable explanation for our results seems to be an attenuation effect based on multiplicative calculations. Nevertheless, we have to consider other possibilities. One possibility would be that the arrival and thus the subtraction of the prediction is triggered by the stimulus itself, such that the error is only produced when a stimulus is present. However, this is unlikely because our results show that predictability mainly affects the very initial neuronal response, suggesting that the prediction is already held available when the stimulus appears.

Assuming that the attenuation effect is multiplicative, it raises a fundamental question whether the mouse visual system utilizes predictive coding in the first place, and if so, how. Further work is required on both the theoretical and experimental side to answer this critical question.

4.4 Conclusion

The present study showed that temporal predictability manipulates neuronal firing in mouse V1. Since we did not find any neurons showing interval coding in V1 and obtained similar effects with both auditory and visual cues, our results are in favor of a higher-level feedback mechanism that mediates attenuation of temporally predicted stimuli. Interestingly, we found no significant modulation when the target was omitted, which contradicts with the current predictive coding framework assuming subtractive error calculation. It raises the fundamental question whether the brain utilizes such predictive coding schemes, and if so, how the framework can be modified to fit our observations which suggest multiplicative modulation.

Bibliography

- Bamberg, Ernst. 2010. *Optogenetics*. <http://www.mpg.de/18011/Optogenetics?page=1>. [Accessed: 16. August 2015].
- Bastos, Andre M, Usrey, W Martin, Adams, Rick A, Mangun, George R, Fries, Pascal, & Friston, Karl J. 2012. Canonical microcircuits for predictive coding. *Neuron*, **76**(4), 695–711.
- Bloom, F.E., Kupfer, D.J., & of Neuropsychopharmacology, American College. 1995. *Psychopharmacology: The Fourth Generation of Progress : An Official Publication of the American College of Neuropsychopharmacology*. Raven Press.
- Bolz, Jürgen, & Gilbert, Charles D. 1986. Generation of end-inhibition in the visual cortex via interlaminar connections.
- Buzsáki, György, & Mizuseki, Kenji. 2014. The log-dynamic brain: how skewed distributions affect network operations. *Nature Reviews Neuroscience*, **15**(4), 264–278.
- Czigler, István, Winkler, István, Pató, Livia, Várnagy, Anna, Weisz, Júlia, & Balázs, László. 2006. Visual temporal window of integration as revealed by the visual mismatch negativity event-related potential to stimulus omissions. *Brain research*, **1104**(1), 129–140.
- Erisken, Sinem, Vaiceliunaite, Agne, Jurjut, Ovidiu, Fiorini, Matilde, Katzner, Steffen, & Busse, Laura. 2014. Effects of locomotion extend throughout the mouse early visual system. *Current Biology*, **24**(24), 2899–2907.
- Felleman, Daniel J, & Van Essen, David C. 1991. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex*, **1**(1), 1–47.
- Gabbiani, Fabrizio, Krapp, Holger G, Koch, Christof, & Laurent, Gilles. 2002. Multiplicative computation in a visual neuron sensitive to looming. *Nature*, **420**(6913), 320–324.
- Garrido, Marta I, Kilner, James M, Stephan, Klaas E, & Friston, Karl J. 2009. The mismatch negativity: a review of underlying mechanisms. *Clinical*

neurophysiology : official journal of the International Federation of Clinical Neurophysiology, **120**(3), 453–63.

- Gregory, R., & Cavanagh, P. 2011. The Blind Spot. *Scholarpedia*, **6**(10), 9618.
- Hangya, Balázs, & Kepecs, Adam. 2015. Vision: How to Train Visual Cortex to Predict Reward Time. *Current Biology*, **25**(12), R490–R492.
- Harrison, LM, Stephan, KE, Rees, G, & Friston, KJ. 2007. Extra-classical receptive field effects measured in striate cortex with fMRI. *Neuroimage*, **34**(3), 1199–1208.
- Hazan, Lynn, Zugaro, Michaël, & Buzsáki, György. 2006. Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. *Journal of neuroscience methods*, **155**(2), 207–216.
- Henze, Darrell A, Borhegyi, Zsolt, Csicsvari, Jozsef, Mamiya, Akira, Harris, Kenneth D, & Buzsáki, György. 2000. Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *Journal of neurophysiology*, **84**(1), 390–400.
- Hu, Hua, Gan, Jian, & Jonas, Peter. 2014. Fast-spiking, parvalbumin+ GABAergic interneurons: From cellular design to microcircuit function. *Science*, **345**(6196), 1255–1263.
- Kawato, Mitsuho, Hayakawa, Hideki, & Inui, Toshio. 1993. A forward-inverse optics model of reciprocal connections between visual cortical areas. *Network: Computation in Neural Systems*, **4**(4), 415–422.
- Liu, Bao-hua, Li, Pingyang, Sun, Yujiao J, Li, Ya-tang, Zhang, Li I, & Tao, Huizhong Whit. 2010. Intervening inhibition underlies simple-cell receptive field structure in visual cortex. *Nature neuroscience*, **13**(1), 89–96.
- Liu, Cheng-Hang, Coleman, Jason E, Davoudi, Heydar, Zhang, Kechen, & Shuler, Marshall G Hussain. 2015. Selective Activation of a Putative Reinforcement Signal Conditions Cued Interval Timing in Primary Visual Cortex. *Current Biology*.
- McAdams, Carrie J, & Maunsell, John HR. 1999. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *The Journal of Neuroscience*, **19**(1), 431–441.
- Mehaffey, W Hamish, Doiron, Brent, Maler, Leonard, & Turner, Ray W. 2005. Deterministic multiplicative gain control with active dendrites. *The Journal of neuroscience*, **25**(43), 9968–9977.

- Mel, Bartlett W. 1993. Synaptic integration in an excitable dendritic tree. *Journal of neurophysiology*, **70**(3), 1086–1101.
- Meyer, Travis, & Olson, Carl R. 2011. Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences*, **108**(48), 19401–19406.
- Mumford, David. 1992. On the computational architecture of the neocortex. *Biological cybernetics*, **66**(3), 241–251.
- Murray, Scott O, Olman, Cheryl A, & Kersten, Daniel. 2006. Spatially specific fMRI repetition effects in human visual cortex. *Journal of neurophysiology*, **95**(4), 2439–2445.
- Nadasdy, Zoltán, Csicsvari, Jozsef, Penttonen, Markku, Hetke, Jamille, Wise, Kensall, & Buzsaki, György. 1998. Extracellular recording and analysis of neuronal activity: from single cells to ensembles. *Neuronal Ensembles: Strategies for Recording and Decoding*, 17–55.
- Petsche, H, Pockberger, H, & Rappelsberger, P. 1984. On the search for the sources of the electroencephalogram. *Neuroscience*, **11**(1), 1–27.
- Quiroga, R Quian, Nadasdy, Zoltan, & Ben-Shaul, Yoram. 2004. Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural computation*, **16**(8), 1661–1687.
- Raij, T, McEvoy, L, Mäkelä, Jyrki P, & Hari, Riitta. 1997. Human auditory cortex is activated by omissions of auditory stimuli. *Brain research*, **745**(1), 134–143.
- Rao, Rajesh PN, & Ballard, Dana H. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, **2**(1), 79–87.
- Reichardt, Werner, & Rosenblith, WA. 1961. Autocorrelation, a principle for evaluation of sensory information by the central nervous system. *Pages 303–317 of: Symposium on Principles of Sensory Communication 1959*. MIT Press.
- Salinas, Emilio, & Thier, Peter. 2000. Gain modulation: a major computational principle of the central nervous system. *Neuron*, **27**(1), 15–21.
- Shuler, Marshall G, & Bear, Mark F. 2006. Reward timing in the primary visual cortex. *Science*, **311**(5767), 1606–1609.
- Summerfield, Christopher, Trittschuh, Emily H, Monti, Jim M, Mesulam, M-Marsel, & Egner, Tobias. 2008. Neural repetition suppression reflects fulfilled perceptual expectations. *Nature neuroscience*, **11**(9), 1004–1006.

- Tales, Andrea, Newton, Philip, Troscianko, Tom, & Butler, Stuart. 1999. Mismatch negativity in the visual modality. *Neuroreport*, **10**(16), 3363–3367.
- Vaiceliunaite, Agne, Erisken, Sinem, Franzen, Florian, Katzner, Steffen, & Busse, Laura. 2013. Spatial integration in mouse primary visual cortex. *Journal of neurophysiology*, **110**(4), 964–972.
- Von Helmholtz, Hermann. 1867. *Handbuch der physiologischen Optik*. Vol. 9. Voss.
- Wacongne, Catherine, Labyt, Etienne, van Wassenhove, Virginie, Bekinschtein, Tristan, Naccache, Lionel, & Dehaene, Stanislas. 2011. Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences*, **108**(51), 20754–20759.
- Wang, Quanxin, Gao, Enquan, & Burkhalter, Andreas. 2011. Gateways of ventral and dorsal streams in mouse visual cortex. *The Journal of Neuroscience*, **31**(5), 1905–1918.
- Yabe, Hirooki, Tervaniemi, Mari, Reinikainen, Kalevi, *et al.* 1997. Temporal window of integration revealed by MMN to sound omission. *Neuroreport*, **8**(8), 1971–1974.

Selbständigkeitserklärung

Hiermit versichere ich, dass ich die vorliegende Bachelorarbeit selbständig und nur mit den angegebenen Hilfsmitteln angefertigt habe und dass alle Stellen, die dem Wortlaut oder dem Sinne nach anderen Werken entnommen sind, durch Angaben von Quellen als Entlehnung kenntlich gemacht worden sind. Diese Bachelorarbeit wurde in gleicher oder ähnlicher Form in keinem anderen Studiengang als Prüfungsleistung vorgelegt.

Ort, Datum

Unterschrift