2022

Preliminary investigation of the temporal specificity of sequence learning in the primary visual cortex through predictive coding

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Thesis

PRELIMINARY INVESTIGATION OF THE TEMPORAL SPECIFICITY OF SEQUENCE LEARNING IN THE PRIMARY VISUAL CORTEX THROUGH PREDICTIVE CODING

by

ANTHONY AMIN KHOUDARY

Submitted in partial fulfillment of the

requirements for the degree of

Master of Science

2022

Approved by

First Reader _____

Jeffrey Gavornik, Ph.D.
Assistant Professor of Biology

Second Reader _____

Ian Davison, Ph.D.
Associate Professor of Biology

Acknowledgements

I am extremely grateful for the continuous help and support from all members of the Gavornik Lab throughout this process. Specifically, I would like to thank Dr. Jeffery Gavornik for his invaluable advice and mentorship over the years. The opportunity to be in the lab throughout the constant uncertainty of recent years has been a privilege that many other students have not been fortunate enough to enjoy. This research experience has been one of the most influential parts of my academic career and I am glad I was able to see its fulfillment with this thesis. I would like to extend special thanks to Byron Price for allowing to me to collaborate with him on these experiments. I would not have been able to complete this project without his teaching and guidance. Additionally, I want to thank Cambria Jensen for her constant support and advice with many of the methodologies and logistics involved in this project. Lastly, and arguably most importantly, I would like to thank my family, specifically my parents and my sister Maria. Thank you for always believing in me and motivating me to do my best every day.

PRELIMINARY INVESTIGATION OF THE TEMPORAL SPECIFICITY OF SEQUENCE LEARNING IN THE PRIMARY VISUAL CORTEX THROUGH PREDICTIVE CODING

ANTHONY AMIN KHOUDARY

ABSTRACT

The primary visual cortex (V1) has been classically viewed as an immutable feature detector, with robust responses to low-level characteristics of objects in the visual field. Recent studies have shown the capacity of this cortical area to perform more complex computations. Nominally, the phenomenon of sequence learning relies on the ability of V1 to encode the serial order and temporal frequency of a spatiotemporal visual sequence. Investigating the mechanisms driving this phenomenon through the lens of predictive coding will further the understanding of how V1 operates locally to encode time and learns to predict the future based on minimal sensory information. Through *in vivo* multi-unit recordings from awake mice, this study sought to isolate neural evidence for predictive processing within the paradigm of sequence learning. Seventy unique units were isolated from forty-two mice subjected to experimentation. Preliminary analyses revealed a significant effect that agrees with the initial report on sequence learning but contradicts predictive processing theory. Further investigation is required to draw more robust conclusions about the predictive computations that occur during sequence learning. Increased sample size and refinement of data analysis will likely lead to interesting results.

Table of Contents

List of Figures

I. Introduction

Learning a new task and forming a memory requires some change in neural circuitry. Decades of research has shown that these changes are often the result of synaptic plasticity. Synaptic plasticity can alter the strength of synaptic connections, modulating the information shared among neurons. These changes have been observed across various levels of the neural hierarchy, from individual synapses to functional networks regulating global processes (Abbott & Nelson, 2000). Though many forms of plasticity are well understood, the underlying mechanisms of sequence learning are not. Gavornik and Bear (2014a) have shown that sequence learning contains temporally specific and predictive components. Deducing the mechanistic basis of this observed plasticity in the mouse will contribute to the general body of knowledge relating to how the brain encodes time and predicts the future, with broader implications for the human visual cortex. The ultimate goal is to build machines that learn as effectively as humans and to understand how learning can be revived in certain neuropsychiatric disorders.

Traditional perspectives on primary visual cortex function

The visual cortex has been utilized for decades across animal models to investigate mechanisms of cortical function and plasticity. The hierarchical structure of this cortical area allows for the extraction of increasingly complex characteristics from objects in the visual field, represented in ascending visual areas (DiCarlo et al., 2012; Hubel & Wiesel, 1968; Larkum, 2013). In the traditional framework, conscious

perception is the product of hierarchical combinations of distributed neural representations. This is thought to be achieved by bottom-up relay of sensory information and top-modulation by higher order cortical areas (Glickfeld & Olsen, 2017).

The role of the primary visual cortex (V1) in the traditional theory of visual processing is based on a series of seminal experiments by Hubel and Wiesel (1962). They showed that V1 receives binocular input from the lateral geniculate nucleus (LGN) with similar retinotopy and stimulus preference (Hubek & Wiesel, 1962). In other words, presentation of a bar oriented at 30° would elicit V1 activity when presented to the same retinal location in either eye. From the results of further experiments into the receptive fields for these neurons, Hubel and Wiesel proposed three classes of visually responsive cells within V1: simple cells, complex cells, and hypercomplex cells. Simple cells were observed to have spatially distinct excitatory and inhibitory areas. These cells responded maximally to stimuli with stark contrasts across a range of orientations, giving rise to the classic 'center-surround' receptive field. Complex cells were found to have spatially homogenous receptive fields that maintained orientation tuning but showed sensitivity to direction of movement of the stimulus. Lastly, hypercomplex cells were classified as having inhibitory regions flanking an otherwise homogenous receptive field similar to that of complex cells. For example, an oriented bar with a length that lies well within the receptive field of a hypercomplex cell will drive increased responses as its length is increased up to a certain limit. Once the length passes that limit, the response decreases as a function of stimulus length because portions of the stimulus now lie in the flanking inhibitory regions of the receptive field. This phenomenon is known as end stopping

(Hubel & Wiesel, 1968). Later studies showed the receptive fields of neighboring cortical columns to be tuned to similar orientations, allowing for detection of stimuli that may span multiple receptive fields (Hubel et al., 1976).

Although neurons in V1 were shown to receive binocular input, certain populations responded more readily to visual stimulation in one eye compared to the other. This property was coined ocular dominance (OD) and it has been shown to rely heavily on experience dependent synaptic plasticity; demonstrated though monocular deprivation (MD) experiments in kittens (Wiesel & Hubel, 1965). Following a period of MD, responses in V1 were primarily driven by activity in the opened eye even though the projections from the closed eye remained intact, suggesting a competitive learning process that guides early plasticity. This effective rewiring of V1 in response to MD was found to be confined to an early developmental window, similar to the critical period of embryogenesis (Spemann & Mangold, 2001). If normal vision was restored before the kitten reached an age of roughly three months, functionality of and responsiveness to the deprived eye was restored. However, past this point restoration of normal vision leads to minimal recovery in the deprived eye (Hubel & Wiesel, 1970). These findings have collectively shaped the conceptual framework of the mechanistic basis by which V1 operates and develops. Consequently, V1 has been viewed as static feature detector whose properties are largely driven by experience dependent plasticity during a well-defined period of developmental plasticity (Gavornik & Bear, 2014b).

Nearly all early investigations into visual physiology and development were conducted in monkeys and cats, but within the past two decades the rodent has become

the dominant model. Rodent V1 has been shown to share many of the same features as early models, including OD plasticity (Gordon & Stryker, 1996; Sawtell et al., 2003). The rodent primary visual cortex, similar to monkeys and cats, contains a topographic map of the visual field (Smith & Häusser, 2010; Wang & Burkhalter, 2007), has orientation tuned simple and complex cells (Niell & Stryker, 2008; Ohki et al., 2005), and functional specificity at the synapse level (Ko et al., 2011). The rodent model provides a greater level of convenience with regards to its size, cost and amenability to genetic modification. These factors coupled with newly developed technology for *in vivo* recording, imaging, and manipulating neuronal activity has led to a substantial body of work on the structure, function, and development of the rodent central visual pathways (Seabrook et al., 2017).

Beyond simple feature detection: experience dependent plasticity and sequence learning in primary visual cortex

The mechanisms governing experience dependent plasticity within rodent/mouse V1 and the extent of cortical computation that can be carried out by this presumed low level cortical area remain to be discovered. Previous work has shown that neurons in layer 4 of the binocular region of V1 in adult mice can encode stimulus familiarity. This was evidenced by a plateau in the trough-to-peak magnitude of the visually evoked potentials (VEPs) recorded from awake mice in response to repeated presentations of a high contrast phase reversing grating (Cooke & Bear, 2010). This effect, known as stimulus-selective response potentiation (SRP), is spatially specific and requires *N*-

methyl-*D*-aspartate (NMDA) receptor activation in a similar fashion to long-term potentiation (LTP) (Frenkel et al., 2006).

Reports have also shown the capacity of adult rodent V1 to encode reward timing. In a visually cued reward protocol, naïve mice exhibited V1 activity that lasted the duration of the visual cue. After repeated presentations of the visual cue paired with a water reward over the course of days, V1 activity persisted until the time of reward delivery even when the reward was withheld (Shuler, 2006). These results contradict the traditional view of V1 as a simple feature detector, and instead indicate that V1 activity is modulated by experience. Specifically, they suggest that V1 can learn stimulus familiarity and encode the temporal dynamics of past associations – processes previously assumed to be restricted to higher order visual areas.

To add to the growing body of knowledge regarding the dynamic ability of V1 to learn well beyond the critical period and perform complex computations, Gavornik and Bear (2014) investigated the V1 response to sequences of spatiotemporal visual stimuli. Adult mice were subjected to a four-day training period during which they saw repeated presentations of a four-element spatiotemporal sequence. Each element consisted of a unique full screen oriented, high contrast sinusoidal grating that persisted on screen for 150ms, with the full sequence denoted ABCD. On the fifth day mice were presented with the trained stimulus along with perturbations of the stimulus by reversing the serial order (DCBA) or altering the temporal frequency by presenting each element for 300ms ($ABCD_{300}$). The cortical response to each test stimulus was measured as the VEPs in

binocular V1 layer 4 and compared to baseline measurements taken on the first day of training. Relative to the baseline, the largest increase in response magnitude was observed when the trained stimulus was presented with familiar timing (Figure 1). Conversely, there was no significant difference in the response magnitude when mice were presented with reversed stimulus (DCBA). Similarly, no difference in response magnitude was observed when presented with $ABCD_{300}$ (Figure 1). This increase in the magnitude of the VEP in response to repeated presentations of a spatiotemporal visual stimulus over the course of days has been dubbed sequence learning. The sequence learning phenomenon suggests that V1 is capable of learning the serial order and temporal frequency of spatiotemporal stimuli. Confirmation that this process occurs locally in V1 was provided through monocular occlusion experiments in which mice were trained with sequence presentation restricted to one eye and tested with monocular presentation to both eyes. These experiments showed a significant difference between evoked VEP magnitude in the contralateral and ipsilateral hemispheres corresponding to the trained eye when presented with ABCD. This difference was not observed in the trained hemispheres in response to DCBA, nor was it present in the corresponding hemispheres of the untrained eye. These findings indicate that the plasticity occurs at a site where information from the two eyes can be segregated. As there is no evidence that monocular segregation persists beyond V1, it likely that this process occurs locally within V1 (Gavornik & Bear, 2014b).

To differentiate sequence learning from the SRP phenomenon, mice were systematically treated with the NMDA receptor antagonist 3-(2-carboxypiperazin-

4yl)propyl-1-phosponic acid (CPP) and subjected to both the SRP and sequence learning protocols. CPP treated mice did not show the characteristic increase in VEP magnitude across days when subjected to the SRP procedure; however, the potentiated response was observed in mice subjected to the sequence learning protocol. The mechanistic basis of this process was further isolated by treatment with scopolamine, a muscarinic acetylcholine receptor antagonist, which abolished the increase in VEP magnitude after completion of the sequence learning protocol. Thus, sequence learning is independent of SRP and cholinergic input is a necessary component of its mechanism.

These findings clearly demonstrate that V1 was able to recognize the neural representations of the familiar visual stimulus. However, it was unclear whether this paradigm was sufficient for V1 to regenerate the response in the absence of external stimulation. To investigate this, another cohort was subjected to the sequence learning training protocol but shown test stimuli with the second element omitted and replaced with a grey screen along with the trained stimulus (ABCD). The first element was either predictive of the trained second element (A_CD) or a novel element with no predictive value (E_CD). The cortical response to the omitted second element was significantly larger when preceded by the familiar initial element compared to the novel. This suggests that V1 can actively predict expected visual stimuli with a high degree of temporal precision. Further current source density (CSD) analysis, showing the source/sink pattern of the neural response across cortical layers (Kamarajan et al., 2015), confirmed this conclusion. The source/sink pattern in response A_CD was almost identical compared to observed pattern in response to ABCD (Figure 1). In other words, V1 is active when it

expects to see a stimulus even if a stimulus is not physically present. Collectively the experimental basis of sequence learning provides robust evidence that V1 is dynamic cortical area capable of carrying out presumed "higher" cortical function with mechanisms that operate within the local physiology.

Predictive processing as a mechanism of sequence learning

Predictive processing has been posited as a framework for understanding neural function, with a strong basis in the fields of computational and cognitive neuroscience (Clark, 2013; Koster-Hale & Saxe, 2013). Predictive processing theory is centered on the idea that the brain generates an internal model of the world based on ascending sensory information relayed throughout the cortex. This internal model is then used to predict future sensory input (Gregory et al., 1980). This process is thought to assume a hierarchical organization, similar to the traditional perspective V1 function. Higher level areas communicate the predicted input to lower-level areas via top-down projections. These predictions are compared to the actual, bottom-up, sensory input and an estimate of the difference between the two is computed. Substantive differences between the predicted and actual sensory input causes the cortex to update its model so that it can maintain an accurate internal representation of the world.

The notion of comparisons between predicted and actual feedback has been utilized to model the dopaminergic reward system and the functionality of the cerebellum (Schultz et al., 1997; Wolpert et al., 1998). While it's likely that there are predictive processing components to neocortical function, disambiguating between neuronal activity

responsible for relaying sensory information and communicating predictions has proven difficult with current experimental and technological methods. As a result, the underlying neurophysiology is poorly understood. Regardless of this empirical gap, a large amount of theoretical work has been devoted to characterizing the neural architecture necessary for implementing predictive processing in mammalian cortex.

Computing the difference between predicted and actual sensory input requires two types of neurons: internal representation neurons and prediction error neurons. Internal representation neurons project downward in the hierarchy and send predictions about bottom-up input. Conversely, prediction-error neurons project upward in the hierarchy, relaying the difference between the prediction and actual bottom-up sensory input. It has been proposed the response of prediction error neurons is proportional to the magnitude of the difference between ascending sensory input and descending predictions from internal representation neurons. To clarify, if incoming sensory input matches the prediction from internal representation neurons, prediction-error neurons show a decreased response relative to when the two signals are incongruent. To operate effectively in the sensory cortex, both functional classes of neurons are expected to respond selectively to certain stimulus features (G. B. Keller & Mrsic-Flogel, 2018). There are two scenarios when prediction-error neurons will show an increased response. The first, called positive prediction error, occurs when bottom-up input is stronger than predicted (i.e., presentation of an unexpected stimulus). Conversely, negative prediction error occurs when bottom-up input is weaker than predicted (i.e., omission of an expected stimulus). Neurons with sufficiently high baseline firing rates could signal both positive

and negative prediction errors through bidirectional control of their response, with increased firing rates coding for positive prediction error and decreased firing rates signaling negative prediction error. This bidirectional control has been observed in the dopaminergic system (Schultz et al., 1997); however, the relatively low basal firing rates in the neocortex make bidirectional control much less likely (Niell & Stryker, 2008; Sakata & Harris, 2009). It is hypothesized that positive and negative prediction errors signals are communicated through two independent microcircuits in the neocortex (G. B. Keller & Mrsic-Flogel, 2018).

The internal representations postulated by predictive processing are similar to those in the traditional framework, with certain cells responding maximally to certain features in the visual field such as lines and edges. The difference between the two theories is how these internal representations are updated. The traditional representative framework changes internal representations through feature detection as information is relayed to ascending sensory areas. On the other hand, predictive processing updates internal representations through continual comparisons of bottom-up input and top-down predictions based on previously developed internal representations (G. B. Keller & Mrsic-Flogel, 2018). This theory of cortical processing is advantageous over the traditional view because updating the internal representation is not driven by bottom-up input. This allows for modulation of the internal representation solely by top-down signaling, providing a framework by which the cortex can simulate and predict the environment.

The notion that perception is an active process is sufficient to explain a variety of everyday experiences, particularly within the scope of visual consciousness. This is evident through visual illusions such as color constancy and inattentional blindness (Foster, 2011; Simons & Chabris, 1999). Our distorted perception of true reality confers the advantage of prediction. We can anticipate the sensory consequences of our own movements in addition to the dynamics of other objects in the world. For example, we can approximately deduce where a ball flying through the air will land in relation to ourselves and approximately when that will happen. Such predictions depend on detailed knowledge about the properties of objects in the visual field and their context in particular moments in time (G. B. Keller & Mrsic-Flogel, 2018).

Physiological evidence for predictive processing in the neocortex is evident in both classically observed phenomena and new experiments. End-stopping has been modeled within the context of prediction error (Rao & Ballard, 1999) with suppression of the response being a consequence of top-down inhibition. The idea of top-down inhibition has been applied to a variety of classical visual receptive field properties, demonstrating how they can be explained through predictive processing (Spratling, 2010). Additionally, negative prediction error signals have been observed in layer 2/3 of mouse V1 where neuronal populations selectively respond to the absence of visual flow (G. B. Keller et al., 2012) or the absence of an expected visual stimulus (Fiser et al., 2016). Studies have isolated prediction error signals specific to deviations in spatially confined areas of the visual field (Zmarz & Keller, 2016). These prediction error receptive fields were found to be aligned to the retinotopic map of visual cortex and

similar in size to visual receptive fields, suggesting that prediction errors and visual perception are separate aspects of the same cortical computation. The development of predictive responses in the visual cortex have been shown to rely heavily on experience (Makino & Komiyama, 2015; Poort et al., 2015). Passive sensory experience, visuomotor coupling, and exposure to visual cues within a spatial environment all result in predictive responses (Attinger et al., 2017; Fiser et al., 2016; Gavornik & Bear, 2014a; Leinweber et al., 2017; Xu et al., 2012).

As a result, investigating the sequence learning phenomenon through the lens of predictive processing seems promising. This thesis examines the extent to which sequence learning is predictive with regards to the temporal frequency of the trained stimulus. The goal is to further elucidate the mechanisms which shape experience-dependent synaptic plasticity in mouse V1. Relating these findings to the theory of predictive processing will contribute to the greater body of knowledge surrounding predictive processing and its role in neocortical function with potential implications for understanding the human visual cortex.

II. Materials and Methods

Animals

All animals used in this study were wild-type C57BL/6 mice (Charles River Laboratories). Both male and female mice were included. Animals were housed with same sex littermates with up to five animals per cage. Mice were kept on a twelve-hour light/dark cycle with food and water provided *ad libitum*. Experiments were performed during the light cycle at between the hours of 12pm and 5pm and each animal was used for one experiment. Forty-two mice were run in the experimental design. Average age of animals upon completion of the experiment was post-natal day 74 (P74) ± 1 (standard error of the mean, SEM). All procedures were approved by the Institutional Animal Care and Use Committee (IACUC) of Boston University.

Multi-unit Electrode Construction

In order to record neural data *in vivo* from awake mice, custom multi-unit recording electrodes were constructed from Omnetics connectors (Model A79006-001). To begin a 2.25cm piece of silver wire (AM Systems, Inc. Cat#782500. 0.10") was cut and soldered to two channels of the connector to serve as the electrical ground. Six pieces of tungsten wire (California Fine Wire, 0.0009, MO#351990) were to cut to a length of approximately 10cm. Each piece of wire was individually wrapped around a unique channel of the connector, with three wires on either row. The wires were then heated with a heat gun set to 450° F to ensure secure attachment to the pins. The wire tails were pulled together using water until a straight bundle formed close to the connector. A guide

wire was glued to the bundle and the newly connected wires were looped to the opposite side of the connector and fixed with cyanoacrylate glue (Figure 2). Curation time was accelerated using Zipkicker (Robart Manufacturing). The assembled bundle was then cut to a length of approximately 2mm from the tip of the guide wire (Figure 2). Each channel was electroplated in a non-cyanide gold solution using a Nano-Z software to achieve a final impedance of approximately 200 kilo-Ohms at 1000 Hertz (Hz).

Electrode Implantation Surgeries

Mice were anesthetized with an intraperitoneal injection of 50mg/kg of ketamine and 10mg/xylazine. 0.5% to 3% of isoflurane was used as an additional inhalation anesthetic during surgery. Eye lubricant was applied to both eyes to protect from excessive aridity during the procedure. A heat pad was placed below the surgical stage to prevent fluctuations in the animal's body temperature. After removing fur from the animal's head, the exposed skin was swabbed with 70% ethanol and iodopovidone before being cut to reveal the skull. The skull was cleaned with 70% ethanol to remove membranes and any remaining connective tissue was scraped down with a scalpel. Following another swabbing with 70% ethanol, the skull was dried with compressed air. To facilitate head restraint a steel headpost was affixed to the skull anterior to bregma using cyanoacrylate glue. This headpost was used to secure the mouse in a stereotaxic apparatus to ensure consistent and accurate electrode implantation. Electrodes were placed either in the right or left hemisphere during each procedure. Small burr holes, approximately 0.5 mm, were drilled over binocular primary visual cortex (3.1 mm lateral

from lambda) until the cavity filled with cerebrospinal fluid (CSF). The dura was carefully removed, and the electrode was lowered 450 μm below the cortical surface to acquire neural data from layer 4/5 of V1. The ground wire was placed below the dura mater approximately 1 mm lateral from bregma in the same hemisphere as the electrode in V1. All electrodes were rigidly secured to the skull using cyanoacrylate glue cured with Zipkicker. Dental cement was used to enclose exposed skill and electrodes in a protective headcap. 0.1 mg/kg of Buprenex was administered subcutaneously as a postoperative analgesic. After surgery mice were monitored for signs of infection and provided with at least 48 hours of recovery before habituation to the recording and restraint apparatus.

Experimental Design

The experimental setup was comprised of a tube-like viewing apparatus to which mice were head fixed for passive viewing of visual stimuli on a monitor with no ambient lighting (Figure 1). One day prior to the experiment each animal was habituated in the viewing apparatus for thirty minutes. During habituation mice viewed a grey screen on the monitor under experimental conditions (i.e., without ambient lighting). The experiment consisted of a three-day training period during which mice were randomly assigned a specific training stimulus that was shown on each day of training to induce sequence learning. On the fourth day, various test stimuli were presented to investigate the predictive capability of neurons within V1 with regards to temporal frequency.

Stimulus Presentation

All visual stimuli were generated using MatLab to control drawing and timing. Stimuli were presented on a monitor placed 25cm directly in front of the mouse. Each stimulus consisted of a sequence of four visual elements followed by an inter-sequence grey screen. Each visual element was a full screen oriented sinusoidal grating shown at 75% contrast.

A. Training Stimuli

The two training stimuli were composed of the same four visual elements, referred to as ABCD. The following orientations were used across all mice for training: A – 75°, B – 120°, C – 35°, D – 160°. To investigate the temporal component of predictive processing, the training sequence was presented in one of two alternating temporal patterns, denoted 'Long-Short' (LS) or 'Short-Long' (SL). In the LS presentation pattern, element A was shown for 200ms, the long interval, and element B was shown for 100ms, the short interval. Elements C and D were then both shown with long and short intervals respectively, resulting in a unique temporal pattern. Conversely, in the SL presentation pattern element A was shown for the short interval and element B was shown for the long interval. Element C and element D were shown under the same intervals as element A and element B, respectively. This resulted in two training stimuli with opposite presentation times for each individual element. In each training session mice viewed 200 presentations of their randomly assigned form of ABCD (LS or SL) which were grouped into four blocks of 50 presentations, with each presentation

separated by a grey screen shown for 1000ms. Each presentation block was separated by one minute. Each training day thus contained two full minutes of visual stimulus presentation (200 x 600ms) during the approximately ten minutes spent in the viewing apparatus.

B. Test Stimuli

To create periods of hypothesized positive and negative prediction errors with respect to the temporal pattern of the stimulus, eight test stimuli were presented each under the previously discussed LS or SL temporal pattern. The test stimuli consisted of two sequences of high contrast full screen oriented sinusoidal gratings: ABCD and ABBD. The orientation of elements was identical to those shown during training. On Test Day, animals viewed approximately 75 presentations of each test stimulus in randomized order with each presentation separated by a grey screen shown for 1000ms. All 75 presentations of each test stimulus were shown within one block. There was a one-minute break between each block. The total time of visual stimulation on Test Day was roughly six minutes (75 x 600ms) over approximately thirty minutes spent in the viewing apparatus. Showing both temporal patterns of each test stimulus subjected each mouse, regardless of the assigned training stimulus, to periods where the hypothesized positive and negative prediction error signals could be observed.

Data Processing and Analysis

Each electrode contained five or six channels per mouse, depending on the building process. Data from each channel was amplified and digitized using an Open Ephys recording system. Spiking activity was digitized at 30 kilo-Hz (kHz) and bandpass filtered from 300-6000 Hz. Periods of neural activity were isolated by investigating the variance of the voltage trace in each channel. The threshold for spiking activity was set to four standard deviations below the mean voltage. Spike times were isolated as crossings of this threshold, where the initial measured voltage was above the threshold and the subsequently measured voltage was below the threshold. These timestamps were used to create a raster plot and corresponding peristimulus time histogram (PSTH) for each channel in response to each test stimulus. An algorithm determined which units were visually responsive and eliminated responses from units that were highly correlated to isolate distinct units. Waveforms and number of spikes were analyzed for evidence that was suggestive of predictive computations occurring locally in V1 (Figs. 4,6,8).

A. Unpaired Comparisons

Comparison of distinct units' responses to the same test stimulus created windows where negative and positive prediction error signals could be observed. To probe negative prediction error, responses were collected from all units when presented with ABCD-LS (Figure 4). Responses of units that were trained on the presented sequence, ABCD-LS, were compared to responses from units that were trained on the opposing sequence, ABCD-SL, during the same windows interest (Figure 4, black overlays). Unit responses in each trial were normalized to the average firing

rate of that unit over the entire recording on test day. The average response in the windows of interest was taken across trials and that magnitude was compared between training groups (Figure 5).

To probe positive prediction error, responses were collected from all units when presented with ABCD-SL (Figure 6). Responses of units that were trained on the presented sequence, ABCD-SL, were compared to responses from units that were trained on the opposing sequence, ABCD-LS, during the same windows of interest (Figure 6, black overlays). The average normalized response was obtained as described above and compared between training groups (Figure 7).

In both presentation patterns, element C is shown 300ms after sequence onset. To investigate if there was a differential response to C when shown during a novel or trained presentation pattern, responses 350-400ms after sequence onset were collected. The average normalized response in this window was obtained as described earlier and compared between training groups (Figures 5,7). To investigate whether the orientation of the third element affected the response, the average normalized response to the second B in ABBD was compared between training groups in the appropriate presentation pattern (LS for negative prediction error and SL for positive prediction error, Figures 5,7).

B. Paired Comparisons

Investigating the same unit's response to different test stimuli allows for more controlled analysis of the temporal plasticity of sequence learning. To probe this, responses to the same element presented under novel and trained timing in all units were collected. Responses were gathered 50-100ms after the onset of each element (Figure 8). Responses were normalized to each unit's average firing rate during the entire recording on test day. The average of the normalized response was taken across trials and that magnitude was compared when the element was presented under novel or familiar timing (Figures 9, 10). This analysis was stratified by training type to ensure paired comparison of the same units.

Histology

A. Tissue Preparation

Postmortem perfusions were performed to optimize the clarity of histological images that were analyzed by eye to determine the location of electrode tracks within the cortical layers. This served as a validation to confirm the data was recorded from the correct cortical area and layer. Animals were euthanized with an injection of 5mg per 10g of weight of pentobarbital into the abdominal cavity. Perfusions were performed with 4% paraformaldehyde (PFA) in phosphate buffered saline (PBS). The mice were then decapitated, and the entire head was placed in 4% PFA in PBS for a week to complete the fixation process. Following fixation, the brain was extracted and placed

in a solution of 30% sucrose for two days in preparation for sectioning. Brain tissue was sectioned in 50µm on a cryostat machine and stored in PBS until the immunohistochemistry protocol was performed

B. Immunohistochemistry

The slices were placed in a blocking solution (1% Triton-PBS; Normal Goat Serum) consisting of 5 mL of Normal Goat Serum diluted with 45 mL of 0.1% PBS-Triton (NGST) for one hour at room temperature. For the detection of the location of the electrodes, the slices were stained with NeuN, GFAP, and Hoechst. A primary antibody solution was made with 1:500 polyclonal rabbit anti-GFAP and 1:1000 monoclonal mouse anti-NeuN diluted with NGS-T. The slices were incubated in the primary solution at 4℃ overnight and were then rinsed three times with 0.1% Triton in PBS. A secondary antibody solution was made with 1:500 dilution of Alexa Fluor 488 goat anti-mouse and Alexa Fluor 568 goat anti-rabbit. Because the secondary antibodies are light-sensitive, the slices were covered with aluminum foil and were incubated with the secondary solution for 1 hour. After incubation, they were rinsed three times with PBS, and then stored in PBS until mounted on slides for fluorescence imaging.

C. Fluorescent Imaging

All slices were imaged by a Nikon Ni-E fluorescence microscope. The GFAP stain was imaged using the mCherry filter, the NeuN stain was imaged using the GFP filter, and the Hoechst stain was imaged using the DAPI filter (Figure 9).

III. Results

We were successful in isolating seventy distinct units from thirty-nine mice in our experimental setup. Forty-six units were trained on the SL presentation pattern and twenty-four were trained on the LS presentation pattern. To investigate the response of these units to the test stimuli, the average normalized responses to the elements of the sequence were calculated (see Methods). These values will subsequently be referred to as the response magnitude(s).

To analyze neural signals corresponding to negative prediction errors, response magnitudes from distinct units were compared when presented with ABCD-LS (Figure 4). Specifically, two 50ms windows capturing the visual response to the omission of elements B and D were considered (Figure 4, black overlays). The response magnitudes of units trained on the SL presentation pattern (n=46) was compared to response magnitude of units trained on the LS presentation pattern (n=24). The difference in response magnitudes between the two training types was insignificant in the window corresponding to the omission of B (Figure 5, top left panel, Wilcoxon rank-sum test, $p = 0.82$). Opposingly trained units also had similar response magnitudes in the window corresponding to the unexpected omission of D (Figure 5, top right panel, Wilcoxon rank-sum test, $p = 0.74$). Similarly, there was no difference in response magnitude between opposingly trained units in response to C (Figure 5, bottom left panel, Wilcoxon rank-sum test, $p = 0.97$) or in response to the second B in ABBD-SL (Figure 5, bottom right panel, Wilcoxon rank-sum test, $p = 0.85$).

To probe positive prediction errors, the same comparisons of response magnitudes between training types was conducted, however in response to ABCD-SL. The windows of interest represent the visual response to unexpected presentations of elements B and D (Figure 6). The difference in response magnitudes between training types was insignificant in first window (Figure 7, top left panel, Wilcoxon rank-sum test, $p = 0.89$) and second window (Figure 7, top right panel, Wilcoxon rank-sum test, $p = 0.49$). Additionally, the difference in response magnitudes was insignificant in response to C (Figure 7, bottom left panel, Wilcoxon rank-sum test, $p = 0.21$) and in response to the second B in ABBD-LS (Figure 7, bottom right panel, Wilcoxon rank-sum test, $p = 0.72$).

Analysis of the same unit's response to different stimuli allowed for deeper investigation into the temporal plasticity of sequence learning through paired comparisons of the response to elements presented under novel and trained timing (Figure 8). In units trained on the LS pattern (n=24), there was no significant difference in response magnitude following presentations of elements B (Figure 9, top left panel, Wilcoxon sign-rank test, $p = 0.28$), C (Figure 9, bottom left panel, Wilcoxon sign-rank test, $p = 0.71$), and D (Figure 9, top right panel, $p = 0.84$) under novel and trained timing. In similar fashion, there was an insignificant difference in response magnitude following the presentation of the second B in ABBD when the sequence was presented under novel and familiar timing (Figure 9, bottom right panel, Wilcoxon sign-rank test, $p = 0.39$).

In units trained on SL pattern (n=46) there was an insignificant difference in response magnitude following the presentation of elements B (Figure 10, top left panel, Wilcoxon sign-rank test, $p = 0.42$) and D (Figure 10, top right panel, Wilcoxon sign rank test, $p = 0.22$) under novel and trained timing. Similarly, there was no significant difference in response magnitude following presentation of the second B in ABBD when the sequence was presented under novel and familiar timing (Figure 10, bottom right panel, Wilcoxon sign-rank test, $p = 0.71$). However, there was a significant difference in the response magnitudes following the presentation of C. Greater response magnitudes were elicited by presentation of C under trained timing as compared to novel timing (Figure 10, bottom left panel, Wilcoxon sign-rank test, $p < 0.005$).

IV. Discussion

The goal of this project was to examine sequence learning as a predictive process and investigate the ability of V1 to encode and regenerate the temporal structure of a sequential stimulus. The experimental design was created to allow for periods where both positive and negative prediction errors could be observed and analyzed.

Preliminary data analysis yielded mostly inconclusive results. However, it is important to note that analyses are ongoing. A first pass analysis, done without removing redundant units, showed significant differences in the unpaired response magnitude in opposingly trained units 450-500ms after sequence onset when presented with ABCD-LS. This corresponds to the second window of interest, capturing the V1 response to the omission of element D and potentially represents a negative prediction error signal (see Figure 4). After removing redundant units, this effect was found to be insignificant as reported in this thesis (Figure 5, top right panel). The algorithm employed to remove redundant units uses a low correlation threshold to label units as similar and its automated removal of visually responsive units likely contributed to the lack of the effect observed. Correcting this algorithm and confirming by eye that it has removed the appropriate units should lead to more cohesive results going forward.

A significantly greater response magnitude was observed following presentation of C under trained timing compared to novel timing in units trained on the SL presentation pattern (Figure 10, bottom left panel). The fact that this effect was produced in our first pass analysis and is replicated here inspires confidence in its validity. It is interesting that this differential response to trained C was not seen in LS trained units (Figure 9, bottom

left panel). This could be the result of reduced sample size in the trained LS cohort (n=26) compared to the trained SL cohort (n=46). The training stimulus was randomly assigned on day one of the experiment, leaving the stratification of training to chance. Manually assigning training types in the future would allow for a more balanced analysis.

The factors driving the observed increase in response magnitude following presentation of C under trained timing compared to novel timing remain unknown. It is difficult to explain this finding in the context of predictive processing. If sequence learning is a predictive process, the training period should be sufficient for the posited internal representation neurons within V1 to generate an accurate internal model of the trained sequence (G. B. Keller & Mrsic-Flogel, 2018). Consequently, when presented with the trained stimulus under familiar timing, cortical predictions should accurately match the incoming sensory stimuli resulting in decreased firing from putative prediction error neurons. Contrastingly, Gavornik and Bear (2014a) showed that the trained ssequence robustly drove periods of increased firing rates from multi-units within V1. The increased response magnitude driven by presentation of C under trained timing is puzzling because the onset of C is constant across both presentation patterns. If prediction error signals are present in V1, it follows that they would be observed in the windows of interest investigated in this thesis, not during an invariant period of visual stimulation. Accordingly, this finding is in agreement with the initial report on sequence learning but contradicts predictive processing theories. This discrepancy has potential to be furthered parsed out with additional data from the proposed experiment described in this thesis coupled with a refined analysis pipeline.

While predictive processing has a strong body of experimental evidence, the cortical manifestation of this process is diverse. Studies investigating adaptation as a predictive process have reported that inhibition is a crucial component (Carandini, 2000; A. J. Keller et al., 2017) while others found no evidence of inhibition contributing to the process (Carandini & Ferster, 1997). How predictive processing emerges in sequence learning remains unknown. The lack of significant effects did not allow for a characterization of a distinct response, either in the form of increased or decreased response magnitude.

The goal of this project was to collect data that would show activity of the hypothesized prediction error cells that are central to the predictive processing framework. Data was collected from layer 4 of the binocular region in primary visual cortex which mainly receives thalamic input (Bastos et al., 2012). Although no units in this study showed consistent responses indicative of positive or negative prediction error signaling, it is possible that these cells are present in layer 4 and we simply did not record one or they dropped out of analysis. Other studies have shown predictive processing computations occurring in layer 2/3 of V1 (Adesnik et al., 2012; Angelucci et al., 2017; Saleem et al., 2013; Zmarz & Keller, 2016), thus another possibility is that these cells reside outside of layer 4.

In summary, preliminary analyses of sequence learning as a predictive computation revealed trends that require further experimentation to be fully explained. The data showed evidence that contradicts predictive processing; however, no conclusions can be drawn due to the temporal invariability of the observed effect. Additional data collection

and refinement of data analysis will increase the potential of this experimental design to parse the mechanisms governing sequence learning and their relation to predictive processing.

V. Figures

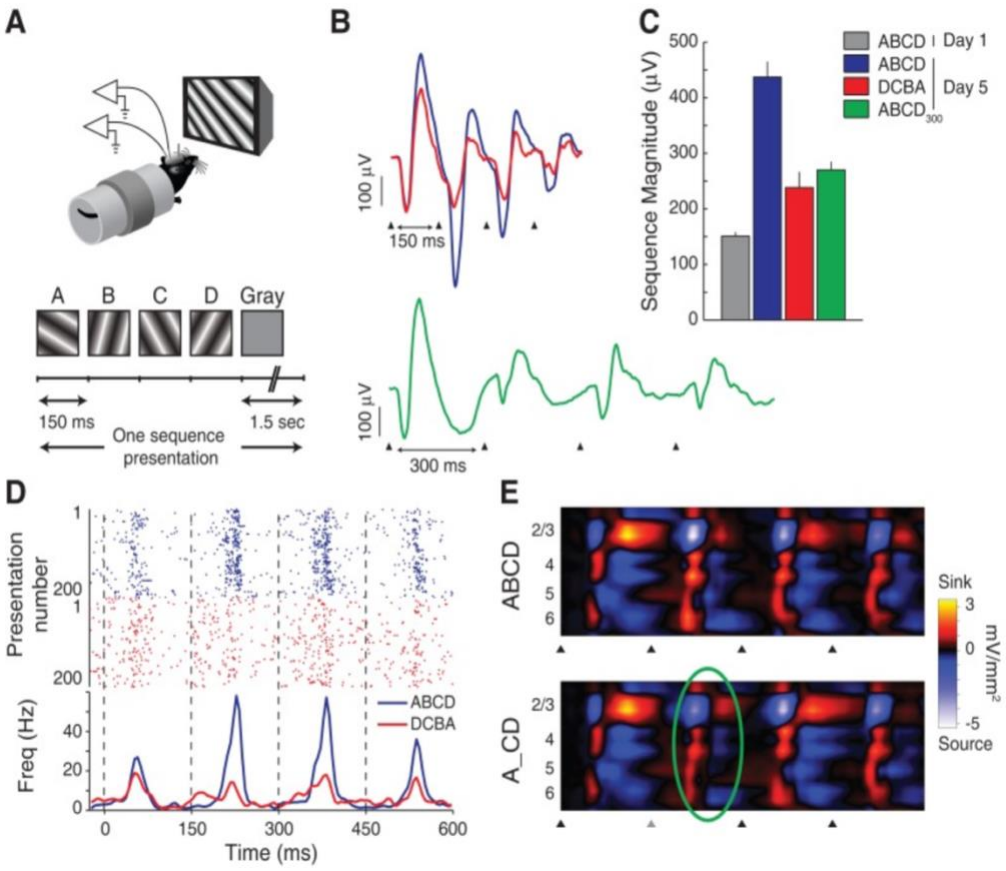

Figure 1. Sequence learning in the primary visual cortex. (a) schematic of experimental set up. (b,c) Visually evoked potentials (VEPs) extracted from extracellular recording electrodes showing waveforms of the response to training and test stimuli (b) and average magnitude of the response on day 1 and day 5 (c). (d) Multi-unit activity evoked by trained and untrained stimuli on day 5 evident through peristimulus time histogram (PSTH) and corresponding raster plot. (e) Current source density analysis (CSD) of the neural activity of V1 in response to the trained stimulus (top) and similar activity in absence of the second element (bottom, green overlay). Reprinted from "Higher brain functions served by the lowly rodent primary visual cortex" J.P. Gavornik and M.F. Bear, 2014, *Learning & Memory*, 21, p. 527–533. Copyright by *Learning & Memory*

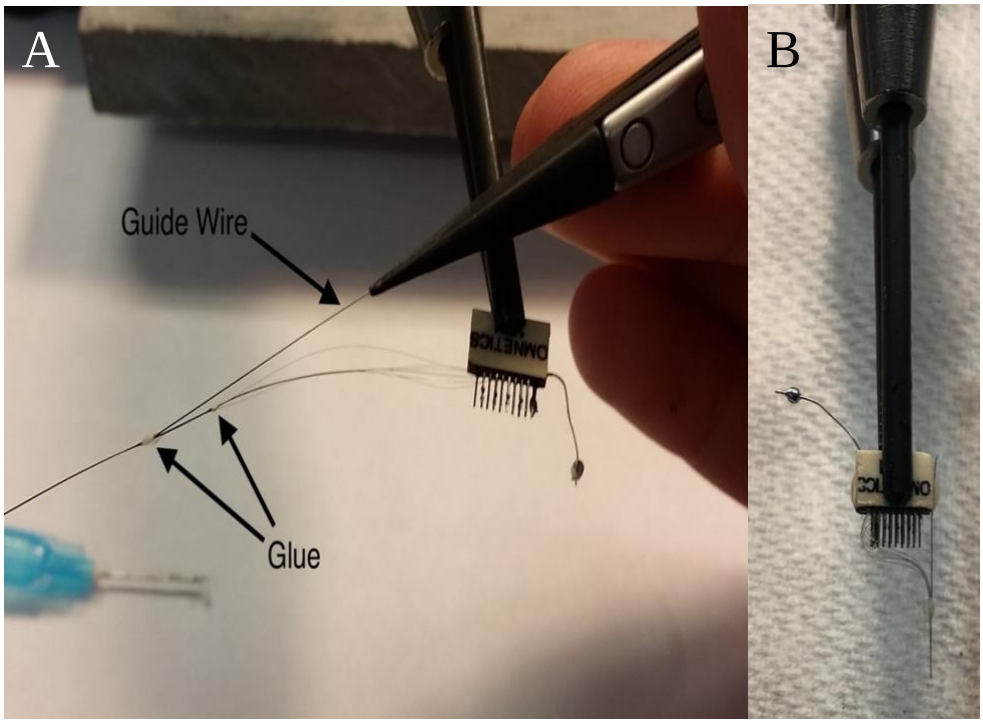

Figure 2. Multi-unit electrode construction. (a) Building process showing attachment of guide wire to assembled bundle and glue placement. (b) Completed electrode ready for surgical implantation

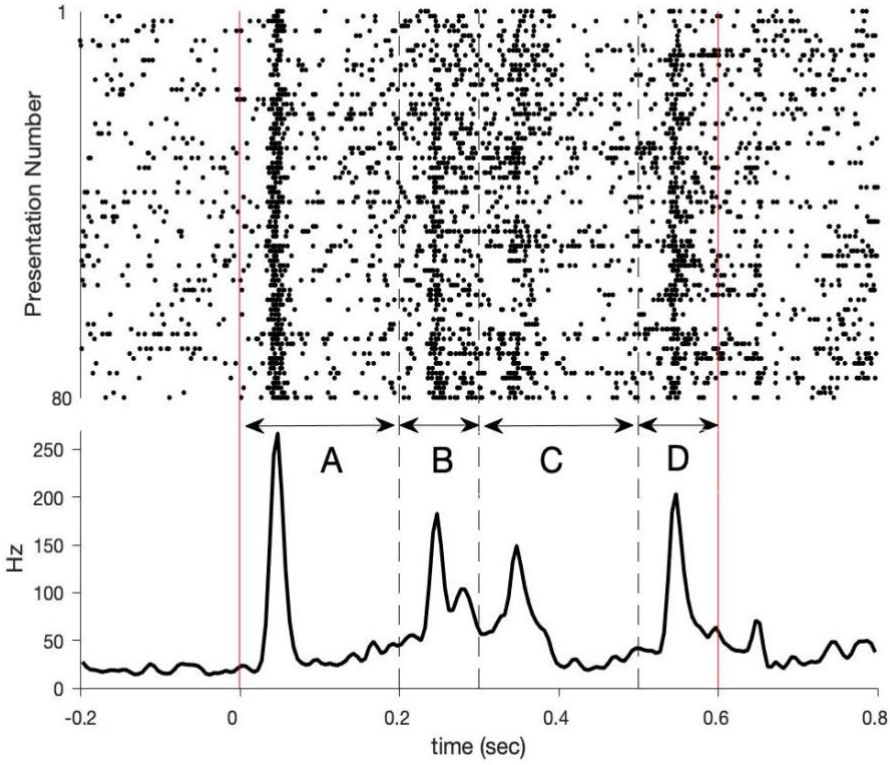

Figure 3. Example of raster plot (top panel) and corresponding peristimulus time histogram (PSTH, bottom panel) in response to ABCD-LS. Each row in the raster is a presentation of the sequence, with each dot representing a spike. Averaging the number of spikes across trials creates the PSTH. Solid red lines at t=0s and t=0.6s indicate sequence on/off. Dashed black lines indicate the onset of each element of the sequence, while horizontal arrows further indicate the duration that each element is held on screen. Clear responses to changes in the visual field are readily observed from multi-unit data 50-100ms after the onset of each element.

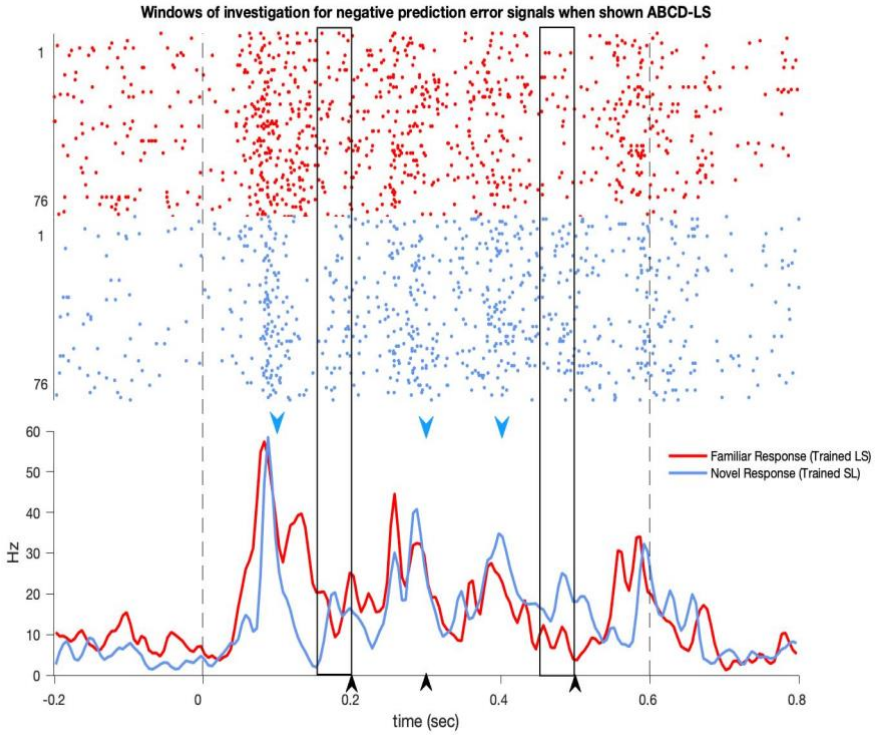

Figure 4. Raster plots and PSTH from distinct units in response to ABCD-LS. Dashed lines indicate sequence on/off. Black arrowheads on x-axis show onset of each element in the presented sequence. The unit shown in red was trained on ABCD-LS whereas the unit shown in blue was trained on ABCD-SL. Unpaired comparisons of multi-unit responses to same stimulus creates two windows (black rectangular overlays) where negative prediction error signals are hypothesized to be observed. As shown, B is presented at 200ms after sequence onset (first black arrowhead, x-axis). Units trained on the SL temporal pattern should expect B to be shown earlier, specifically 100ms after onset of the sequence (first blue arrowhead). Its omission is hypothesized to drive observable neural responses in the window shown by the black overlay. The same conditions apply later in the sequence to element D.

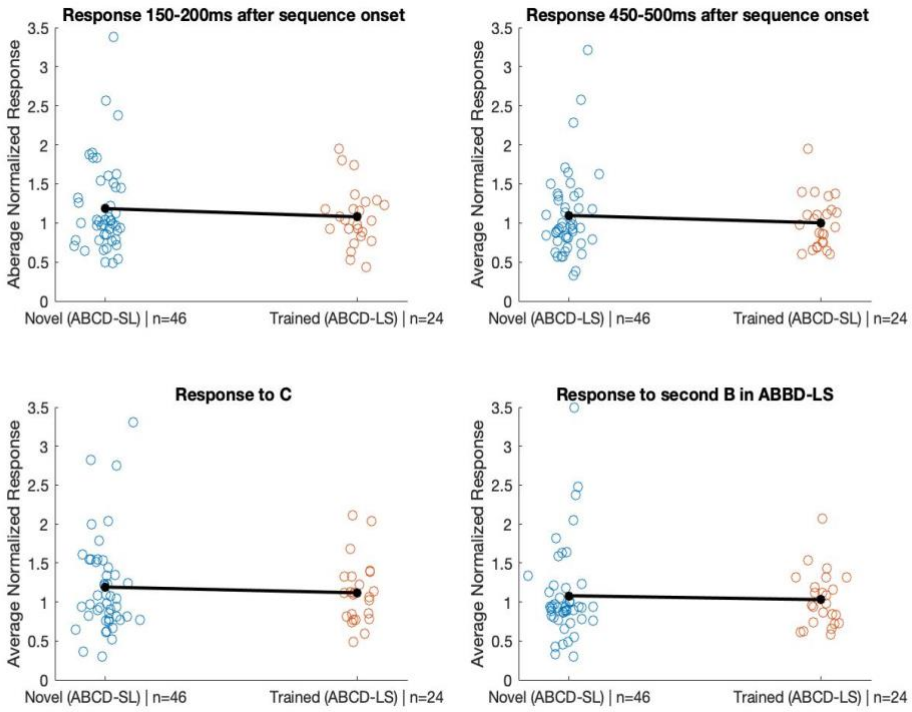

Figure 5. Quantification of unpaired comparisons across units in response to ABCD-LS. Spikes were counted in 50ms windows during which negative prediction error signals should occur (see Figure 4), corresponding to the unexpected omission of elements B (top left, n.s., Wilcoxon rank-sum test, p = 0.82) and D (top right, n.s., Wilcoxon rank-sum test, p= 0.74). Neither the response to C or to the corresponding second B in ABBD-LS were significantly different (Wilcoxon rank-sum tests, p = 0.97 and p =0.85, respectively) between units with opposing training stimuli. Black lines show the mean difference between the two training types.

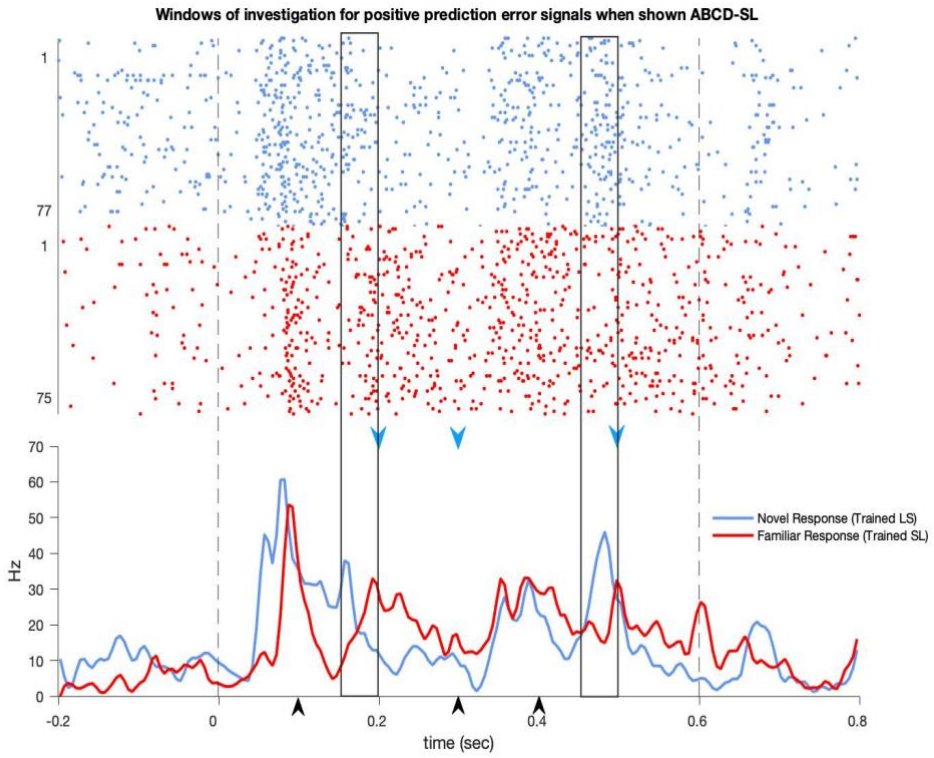

Figure 6. Raster plots and PSTH from distinct units in response to ABCD-SL. Dashed lines indicate sequence on/off. Black arrowheads on x-axis show onset of each element in the presented sequence. The unit shown in red was trained on ABCD-SL whereas the unit shown in blue was trained on ABCD-LS. Unpaired comparisons of multi-unit responses to same stimulus creates two windows (black rectangular overlays) where positive prediction error signals are hypothesized to be observed. As shown, B is presented at 100ms after sequence onset (first black arrowhead, x-axis). Units trained on the LS temporal pattern should expect B to be shown later, specifically 200ms after onset (first blue arrowhead). This unexpected presentation time is hypothesized to drive observable neural responses in window shown by the black overlay. The same conditions apply later in the sequence to element D.

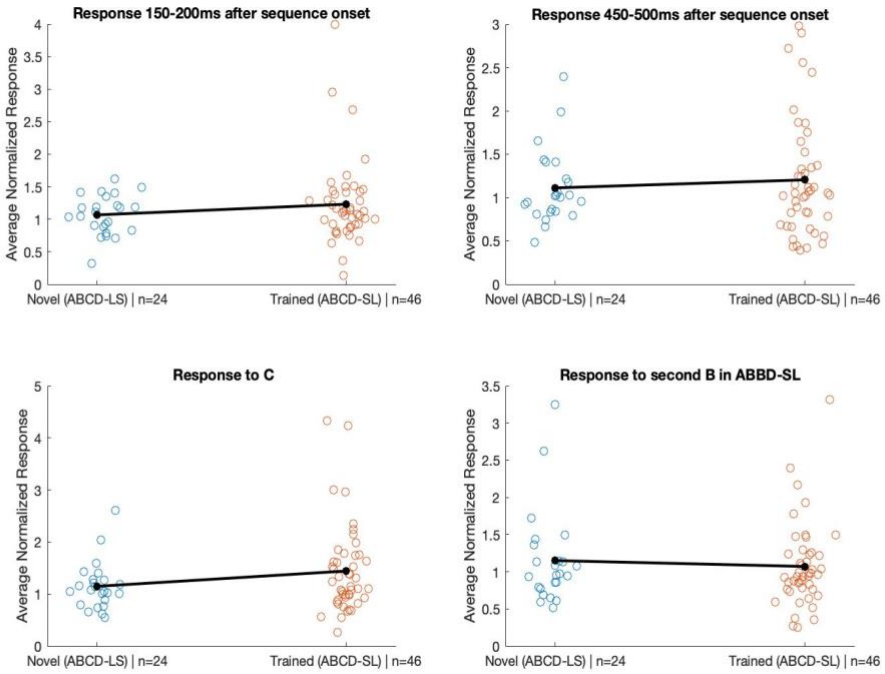

Figure 7. Quantification of unpaired comparisons across all units in response to ABCD-SL. Spikes were counted in 50ms windows during which positive prediction error signals should occur (see Figure 6), corresponding to the unexpected presentation of elements B (top left, n.s., Wilcoxon rank-sum test, $p = 0.49$) and D (top right, n.s., Wilcoxon rank-sum test, $p= 0.89$). Neither the response to C or the corresponding second B in ABBD-SL were significantly different (Wilcoxon rank-sum tests, $p = 0.21$ and $p =0.72$, respectively) between units with opposing training stimuli. Black lines show the mean difference between the two training types.

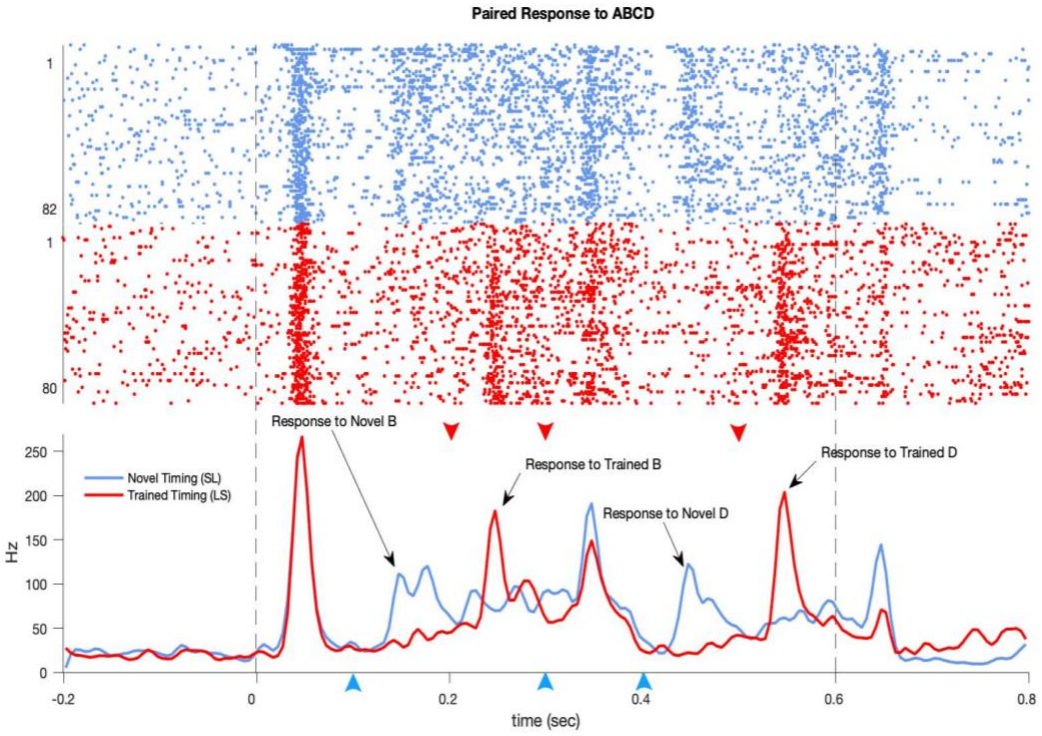

Figure 8. Raster plots and PSTH from the same unit in response to ABCD in both presentation patterns (SL in blue, LS in red). Arrowheads on and above the x axis indicate the onset of each element. This unit was randomly assigned to be trained on the LS pattern making it the trained response. Comparison of the same unit to differing stimuli offers a more controlled analysis of the response to each element, allowing for comparison of the response to each when presented under novel and trained timing, as indicated by the black text arrows.

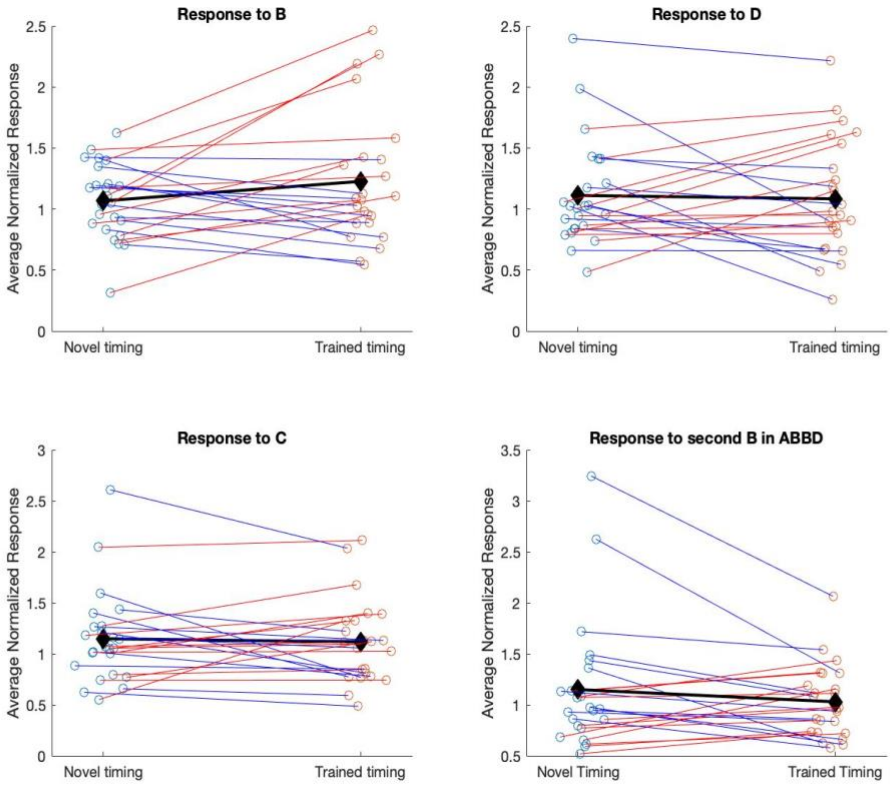

Figure 9. Quantification of each unit's response to elements of sequence when shown under novel and trained timing. These are responses from units trained on LS (n=24) presentation pattern, deeming the SL pattern novel. Paired comparisons are shown through connected data points. If the connecting line is blue, then the response under novel timing was greater than the response under trained timing. Red connections indicate that the trained response was greater than the novel response. Black lines show the mean difference between the two presentation patterns. None of the comparisons were significantly different according to the Wilcoxon sum-rank test ($p > 0.25$ for all).

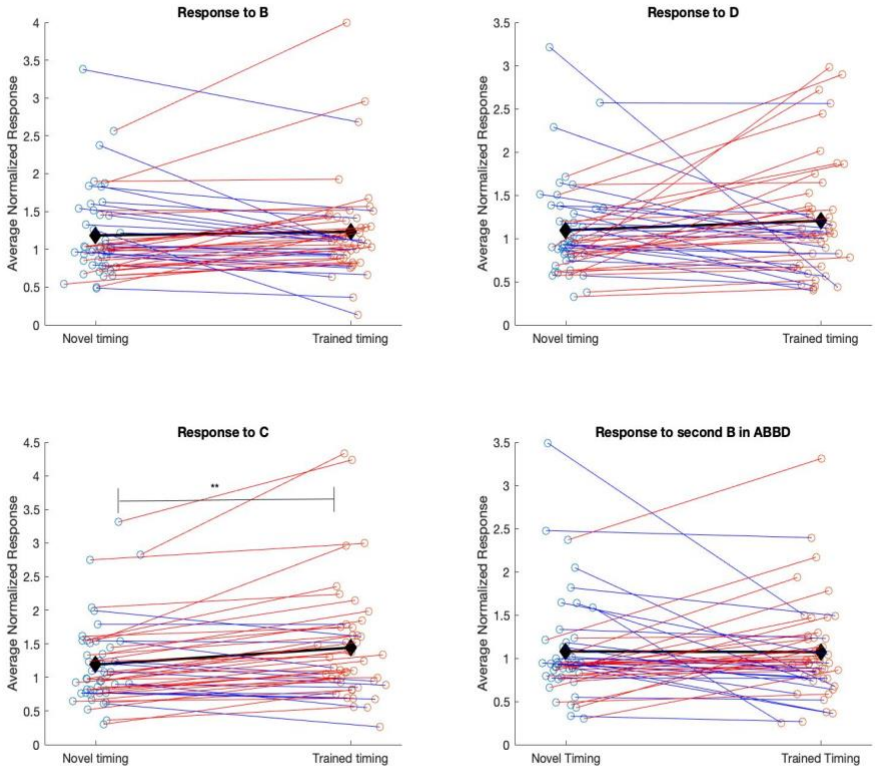

Figure 10. Quantification of each unit's response to elements of sequence when shown under novel and trained timing. These are responses from units trained on the SL (n=46) presentation pattern, deeming the LS pattern novel. Paired comparisons are shown through connected data points. If the connecting line is blue, then the response under novel timing was greater than the response under trained timing. Red connections indicate that the trained response was greater than the novel response. Black lines show the mean difference between the two presentation patterns. Only the response to C was found to be significantly different, with a greater response to trained presentations of the element compared to novel ones (Wilcoxon sum-rank, $p < 0.005$). None of the other comparisons were significantly different according to the Wilcoxon sum-rank test ($p > 0.20$ for all).

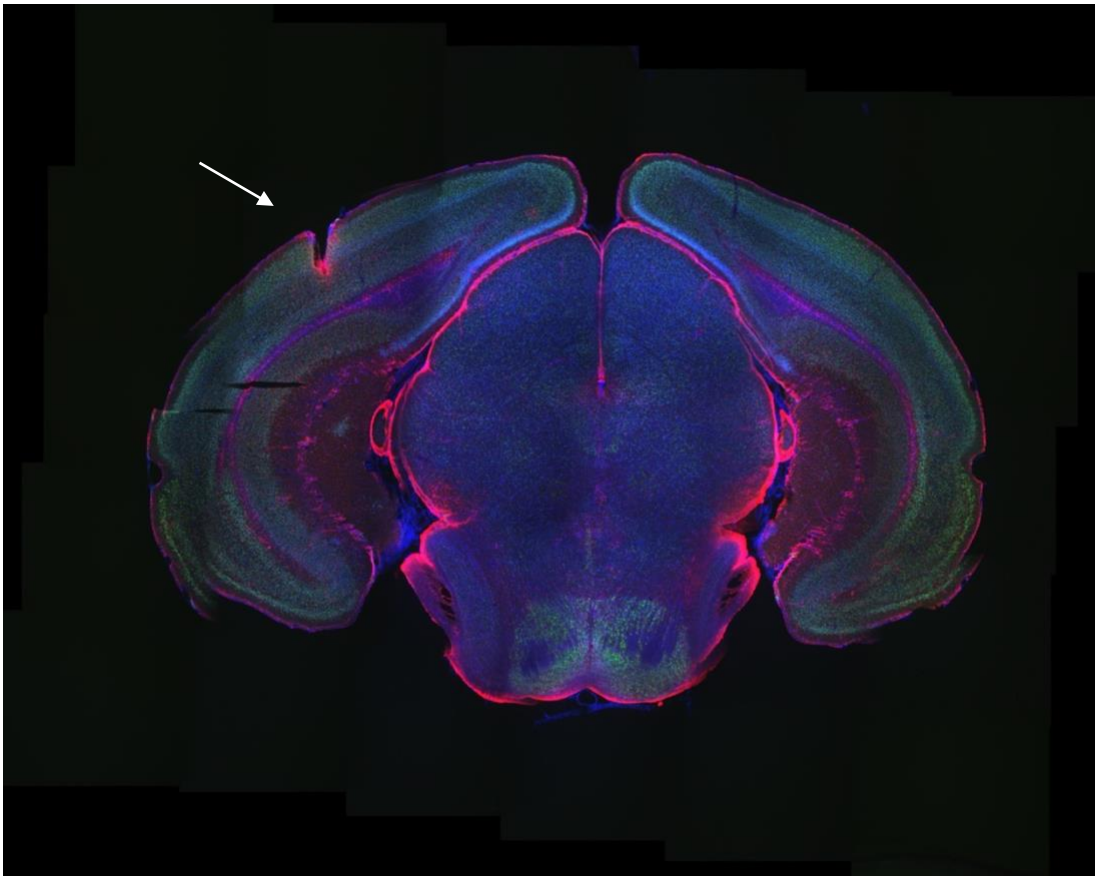

Figure 11. Histology preparation of a 50µm coronal slice of the binocular region of the primary visual cortex. The slice was stained with NeuN (green), GFAP (red), and Hoechst (blue). White arrow indicates location of electrode tract, confirmed by red staining indicating astrocytosis. This image serves as confirmation that data was properly recorded from layer 4/5 of binocular V1.

VI. Bibliography

Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: Taming the beast. *Nature Neuroscience*, *3*(11), 1178–1183. https://doi.org/10.1038/81453

Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J., & Scanziani, M. (2012). A neural circuit for spatial summation in visual cortex. *Nature*, *490*(7419), 226–231. https://doi.org/10.1038/nature11526

Angelucci, A., Bijanzadeh, M., Nurminen, L., Federer, F., Merlin, S., & Bressloff, P. C. (2017). Circuits and Mechanisms for Surround Modulation in Visual Cortex. *Annual Review of Neuroscience*, *40*(1), 425–451. https://doi.org/10.1146/annurev-neuro-072116-031418

Attinger, A., Wang, B., & Keller, G. B. (2017). Visuomotor Coupling Shapes the Functional Development of Mouse Visual Cortex. *Cell*, *169*(7), 1291-1302.e14. https://doi.org/10.1016/j.cell.2017.05.023

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, *76*(4), 695–711. https://doi.org/10.1016/j.neuron.2012.10.038

Carandini, M. (2000). Visual cortex: Fatigue and adaptation. *Current Biology*, *10*(16), R605–R607. https://doi.org/10.1016/S0960-9822(00)00637-0

Carandini, M., & Ferster, D. (1997). A Tonic Hyperpolarization Underlying Contrast Adaptation in Cat Visual Cortex. *Science*, *276*(5314), 949–952. https://doi.org/10.1126/science.276.5314.949

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. https://doi.org/10.1017/S0140525X12000477

Cooke, S. F., & Bear, M. F. (2010). Visual Experience Induces Long-Term Potentiation in the Primary Visual Cortex. *Journal of Neuroscience*, *30*(48), 16304–16313. https://doi.org/10.1523/JNEUROSCI.4333-10.2010

Dean, P. (1981). Grating detection and visual acuity after lesions of striate cortex in hooded rats. *Experimental Brain Research*, *43*(2), 145–153. https://doi.org/10.1007/BF00237758

DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How Does the Brain Solve Visual Object Recognition? *Neuron*, *73*(3), 415–434. https://doi.org/10.1016/j.neuron.2012.01.010

Ellard, C. G., Goodale, M. A., Scorfield, D. M., & Lawrence, C. (1986). Visual cortical lesions abolish the use of motion parallax in the Mongolian gerbil. *Experimental Brain Research*, *64*(3), 599–602. https://doi.org/10.1007/BF00340498

Fiser, A., Mahringer, D., Oyibo, H. K., Petersen, A. V., Leinweber, M., & Keller, G. B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience*, *19*(12), 1658–1664. https://doi.org/10.1038/nn.4385

Foster, D. H. (2011). Color constancy. *Vision Research*, *51*(7), 674–700. https://doi.org/10.1016/j.visres.2010.09.006

Frenkel, M. Y., Sawtell, N. B., Diogo, A. C. M., Yoon, B., Neve, R. L., & Bear, M. F. (2006). Instructive Effect of Visual Experience in Mouse Visual Cortex. *Neuron*, *51*(3), 339–349. https://doi.org/10.1016/j.neuron.2006.06.026

Gavornik, J. P., & Bear, M. F. (2014a). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature Neuroscience*, *17*(5), 732–737. https://doi.org/10.1038/nn.3683

Gavornik, J. P., & Bear, M. F. (2014b). Higher brain functions served by the lowly rodent primary visual cortex. *Learning & Memory*, *21*(10), 527–533. https://doi.org/10.1101/lm.034355.114

Glickfeld, L. L., & Olsen, S. R. (2017). Higher-Order Areas of the Mouse Visual Cortex. *Annual Review of Vision Science*, *3*(1), 251–273. https://doi.org/10.1146/annurev-vision-102016-061331

Gordon, J. A., & Stryker, M. P. (1996). Experience-Dependent Plasticity of Binocular Responses in the Primary Visual Cortex of the Mouse. *The Journal of Neuroscience*, *16*(10), 3274–3286. https://doi.org/10.1523/JNEUROSCI.16-10-03274.1996

Gregory, R. L., Longuet-Higgins, H. C., & Sutherland, N. S. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, *290*(1038), 181–197. https://doi.org/10.1098/rstb.1980.0090

HUBEL, D. H., & WIESEL, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106–154. PubMed. https://doi.org/10.1113/jphysiol.1962.sp006837

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243. https://doi.org/10.1113/jphysiol.1968.sp008455

Hubel, D. H., & Wiesel, T. N. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *The Journal of Physiology*, *206*(2), 419–436. https://doi.org/10.1113/jphysiol.1970.sp009022

Hubel, D. H., Wiesel, T. N., & LeVay, S. (1976). Functional Architecture of Area 17 in Normal and Monocularly Deprived Macaque Monkeys. *Cold Spring Harbor Symposia on Quantitative Biology*, *40*(0), 581–589. https://doi.org/10.1101/SQB.1976.040.01.054

Ingle, D., Cheal, M., & Dizio, P. (1979). Cine analysis of visual orientation and pursuit by the Mongolian gerbil. *Journal of Comparative and Physiological Psychology*, *93*(5), 919–928. https://doi.org/10.1037/h0077618

Kamarajan, C., Pandey, A. K., Chorlian, D. B., & Porjesz, B. (2015). The use of current source density as electrophysiological correlates in neuropsychiatric disorders: A review of human studies. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, *97*(3), 310–322. https://doi.org/10.1016/j.ijpsycho.2014.10.013

Keller, A. J., Houlton, R., Kampa, B. M., Lesica, N. A., Mrsic-Flogel, T. D., Keller, G. B., & Helmchen, F. (2017). Stimulus relevance modulates contrast adaptation in visual cortex. *ELife*, *6*, e21589. https://doi.org/10.7554/eLife.21589

Keller, G. B., Bonhoeffer, T., & Hübener, M. (2012). Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse. *Neuron*, *74*(5), 809–815. https://doi.org/10.1016/j.neuron.2012.03.040

Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive Processing: A Canonical Cortical Computation. *Neuron*, *100*(2), 424–435. https://doi.org/10.1016/j.neuron.2018.10.003

Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjöström, P. J., & Mrsic-Flogel, T. D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature*, *473*(7345), 87–91. https://doi.org/10.1038/nature09880

Koster-Hale, J., & Saxe, R. (2013). Theory of Mind: A Neural Prediction Problem. *Neuron*, *79*(5), 836–848. https://doi.org/10.1016/j.neuron.2013.08.020

Larkum, M. (2013). A cellular mechanism for cortical associations: An organizing principle for the cerebral cortex. *Trends in Neurosciences*, *36*(3), 141–151. https://doi.org/10.1016/j.tins.2012.11.006

Lashley, K. S. (1929). Brain mechanisms and intelligence: A quantitative study of injuries to the brain. *Brain Mechanisms and Intelligence: A Quantitative Study of Injuries to the Brain.*, xi, 186–xi, 186. https://doi.org/10.1037/10017-000

Leinweber, M., Ward, D. R., Sobczak, J. M., Attinger, A., & Keller, G. B. (2017). A Sensorimotor Circuit in Mouse Cortex for Visual Flow Predictions. *Neuron*, *95*(6), 1420-1432.e5. https://doi.org/10.1016/j.neuron.2017.08.036

Makino, H., & Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. *Nature Neuroscience*, *18*(8), 1116–1122. https://doi.org/10.1038/nn.4061

Niell, C. M., & Stryker, M. P. (2008). Highly Selective Receptive Fields in Mouse Visual Cortex. *Journal of Neuroscience*, *28*(30), 7520–7536. https://doi.org/10.1523/JNEUROSCI.0623-08.2008

Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., & Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature*, *433*(7026), 597–603. https://doi.org/10.1038/nature03274

Poort, J., Khan, A. G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G. B., Mrsic-Flogel, T. D., & Hofer, S. B. (2015). Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron*, *86*(6), 1478–1490. https://doi.org/10.1016/j.neuron.2015.05.037

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Sakata, S., & Harris, K. D. (2009). Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex. *Neuron*, *64*(3), 404–418. https://doi.org/10.1016/j.neuron.2009.09.020

Saleem, A. B., Ayaz, A., Jeffery, K. J., Harris, K. D., & Carandini, M. (2013). Integration of visual motion and locomotion in mouse visual cortex. *Nature Neuroscience*, *16*(12), 1864–1869. https://doi.org/10.1038/nn.3567

Sawtell, N. B., Frenkel, M. Y., Philpot, B. D., Nakazawa, K., Tonegawa, S., & Bear, M. F. (2003). NMDA Receptor-Dependent Ocular Dominance Plasticity in Adult Visual Cortex. *Neuron*, *38*(6), 977–985. https://doi.org/10.1016/S0896-6273(03)00323-4

Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, *275*(5306), 1593. https://doi.org/10.1126/science.275.5306.1593

Seabrook, T. A., Burbridge, T. J., Crair, M. C., & Huberman, A. D. (2017). Architecture, Function, and Assembly of the Mouse Visual System. *Annual Review of Neuroscience*, *40*(1), 499–538. https://doi.org/10.1146/annurev-neuro-071714-033842

Shuler, M. G. (2006). Reward Timing in the Primary Visual Cortex. *Science*, *311*(5767), 1606–1609. https://doi.org/10.1126/science.1123513

Simons, D. J., & Chabris, C. F. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception*, *28*(9), 1059–1074. https://doi.org/10.1068/p281059

Smith, S. L., & Häusser, M. (2010). Parallel processing of visual space by neighboring neurons in mouse visual cortex. *Nature Neuroscience*, *13*(9), 1144–1149. https://doi.org/10.1038/nn.2620

Spemann, H., & Mangold, H. (2001). Induction of embryonic primordia by implantation of organizers from a different species. 1923. *The International Journal of Developmental Biology*, *45*(1), 13–38.

Spratling, M. W. (2010). Predictive Coding as a Model of Response Properties in Cortical Area V1. *The Journal of Neuroscience*, *30*(9), 3531. https://doi.org/10.1523/JNEUROSCI.4911-09.2010

Wang, Q., & Burkhalter, A. (2007). Area map of mouse visual cortex. *The Journal of Comparative Neurology*, *502*(3), 339–357. https://doi.org/10.1002/cne.21286

Wiesel, T. N., & Hubel, D. H. (1965). EXTENT OF RECOVERY FROM THE EFFECTS OF VISUAL DEPRIVATION IN KITTENS. *Journal of Neurophysiology*, *28*(6), 1060–1072. https://doi.org/10.1152/jn.1965.28.6.1060

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*(9), 338–347. https://doi.org/10.1016/S1364-6613(98)01221-2

Xu, S., Jiang, W., Poo, M., & Dan, Y. (2012). Activity recall in a visual cortical ensemble. *Nature Neuroscience*, *15*(3), 449–455. https://doi.org/10.1038/nn.3036

Zmarz, P., & Keller, G. B. (2016). Mismatch Receptive Fields in Mouse Visual Cortex. *Neuron*, *92*(4), 766–772. https://doi.org/10.1016/j.neuron.2016.09.057

VII. Vita