

Experience-dependent spatial expectations in mouse visual cortex

Inauguraldissertation

zur

Erlangung der Würde eines Doktors der Philosophie

vorgelegt der

Philosophisch-Naturwissenschaftlichen Fakultät

der Universität Basel

von

Aris Nikolaos Fiser

von Griechenland und USA

Basel, 2019

Originaldokument gespeichert auf dem Dokumentenserver der Universität Basel edoc.unibas.ch

Genehmigt von der Philosophisch-Naturwissenschaftlichen Fakultät auf Antrag von

Dr Thomas Mrsic-Flogel

Dr Georg Keller

Dr Walter Senn

Basel, 21/03/2017

Prof. Dr. M. Spiess, Dekan

Table of Contents

Acknowledgements.....	4
Abstract.....	5
Introduction	6
The visual system	10
Anatomy.....	10
Function	11
Predictive coding.....	12
Theoretical Work	12
Experimental evidence.....	13
Experience-dependent spatial expectations in mouse visual cortex.....	14
V1 activity becomes descriptive of spatial location.....	15
V1 develops predictive responses to upcoming visual stimuli	21
ACC conveys stimulus predictive signals to V1	24
Omitting an expected stimulus drives strong responses in V1.....	27
Discussion.....	29
ACC as part of the brain's generative model	30
Why do visually driven responses persist in the presence of predictions?	31
Why does a familiar stimulus in an unexpected location not evoke a mismatch signal?	33
Why are stimulus predictions also found in V1?	34
Why is spatial decoding in CA1 poor in the virtual environment?	35
Conclusion.....	36
Methods.....	37
Animals and imaging.....	37
Virtual reality and behavior.	38
Experimental design.....	38
Data analysis.	38
V1 activity becomes descriptive of spatial location.....	39
V1 develops predictive responses to upcoming visual stimuli	40
ACC conveys stimulus predictive signals to V1.	40
Omitting an expected stimulus drives strong responses in V1.....	41
Bibliography	42

Acknowledgements

The circumstances which led me to study neuroscience were entirely accidental. Almost two years into my undergraduate studies, when I had changed majors about three times, I stumbled onto a news article about how the brain could possibly be “tricked” out of drug addiction. I suddenly became fascinated with the idea that this great wall of certainty that is our experience of ourselves and our motivations and states was in fact pliable, and could easily be intervened with. This led me to finally make the decisive switch to neuroscience. The topic that attracted me the most was vision, as to us it seems effortless yet a very large fraction of our brain is dedicated to it; and, as optical illusions demonstrate, it relies heavily on the fact that the brain can paint reality with things that aren’t there in order to give us the impression of a seamless whole.

To study how the brain does exactly that, I joined the lab of Georg Keller as a PhD student. I’m very grateful for the opportunities and guidance Georg offered. I also couldn’t have done all the experiments outlined in this thesis without the help of my lab mates, most notably David Mahringer, Hassana Oyibo, Marcus Leinweber and Anders Petersen. I’d also like to thank the members of my thesis committee, Tom Mrsic-Flogel and Walter Senn, for providing me with excellent feedback and support on this project, including ideas for new experiments and analyses.

Finally, I’d like to thank my family for unconditionally supporting me all these years, and my friends for making doing a PhD lots of fun.

Abstract

The presence of optical illusions tells us that what we see - at least a large fraction of it – is influenced by our expectations, which are built on our experience of the environment. This means that, at a perceptual level, our brain fills in parts of the visual world with things that aren't there.

In this work, we have discovered the presence of neurons in the primary visual cortex (V1), the earliest cortical stage of visual processing, that predict the identity of upcoming visual stimuli that mice observed in a virtual environment. We also show that these predictions are dependent on the animals' spatial location, suggesting that an internal representation of space can serve as a scaffold for predictions. Consistent with an influence representations of space on visual processing, we find that the activity of neurons in V1 was modulated by location in the environment, and it is thus possible to decode not only the stimulus the mouse is observing at any given time, but its location as well, in the case where an identical stimulus is presented in multiple locations.

Furthermore, we identified the anterior cingulate cortex (ACC) as a potential source of stimulus predictions to V1, as V1-projecting ACC axons carried stimulus-predictive activity. Finally, omitting an expected visual stimulus drove strong responses in V1. These results are consistent with a predictive coding framework, wherein predictions of future sensory stimuli are compared to ongoing sensory input, and mismatches between the two lead to error signals.

Introduction

In the early days of artificial intelligence, it was thought that the brain functions easiest to reproduce in machines would be the senses, like vision, and the most difficult would be more esoteric concepts like reasoning. The exact opposite turned out to be true: computers are remarkably good at proving (simple) theorems, solving geometry problems and playing complex games, but are incredibly inept at vision¹ (Moravec 1988). However, vision, and sensing the world in general, seem remarkably easy to us. Why, then, does this discrepancy exist? How are humans - and many other animals - so good at vision?

An elegant model of how sensory systems could function came at least as early as the work of Horace Barlow in the 1950's, with the theorization of the “feature detector”, a module in the sensory brain that would be sensitive to specific features in the environment. It was brought about by the discovery of neurons in the frog retina that responded to small moving objects in the visual field (Barlow 1953). Further experimental evidence for such a detector came with the seminal work of David Hubel & Thorsten Wiesel (Hubel and Wiesel 1962), when they found neurons in the primary visual cortex (V1) of the cat that were selectively responsive to bars of specific orientations in their preferred regions of the visual field. They hypothesized that these types of properties could arise by summing over the responses of neurons in the lateral geniculate nucleus of the thalamus (LGN; the primary source of retinal input to V1 (Lund 1988)) with appropriately aligned receptive fields, such that they would best drive the downstream V1 neuron when a dark oriented bar was in that area of visual space (see **Fig. 1** below).

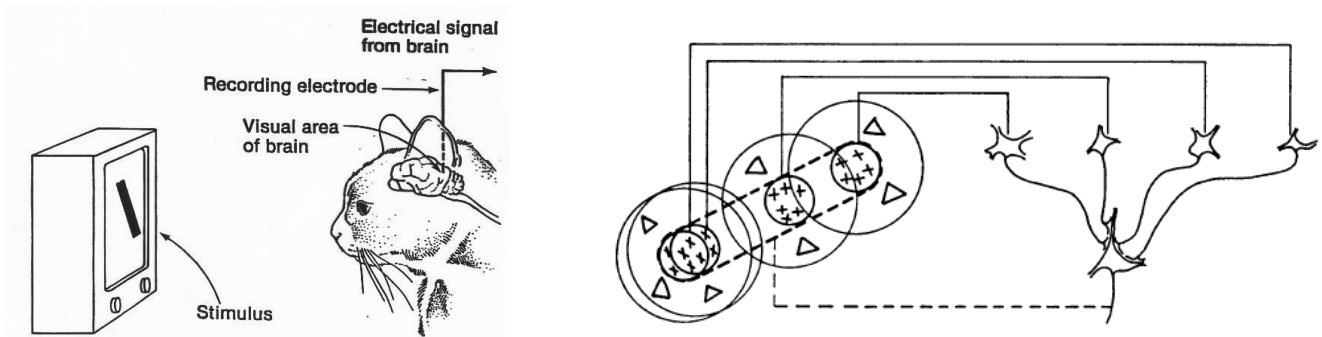


Figure 1: Left: schematic of the experiment. Cats viewed stimuli presented on a screen while the activity of neurons in visual cortex was recorded. Right: Schematic of a possible connectivity pattern that could explain the formation of orientation-selective neurons in V1. Adapted from (Hubel and Wiesel 1962)

¹ This is known as “Moravec’s paradox”. Moravec writes: “it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility”.

One could propagate this idea further and hypothesize that the activity of neurons with diverse orientation selectivity in V1 could be integrated in higher visual areas to give rise to neurons selective to curves, complex shapes, and even faces (indeed, primates seem to have a visual area dedicated to faces (Kanwisher, McDermott, and Chun 1997)). This feed-forward hypothesis has become the dominant view of how sensory processing works. While it carries a lot of credibility based on known response properties in higher visual areas, it is insufficient in explaining how we perceive the world. The most intuitive way to demonstrate this is via optical illusions, such as the following:

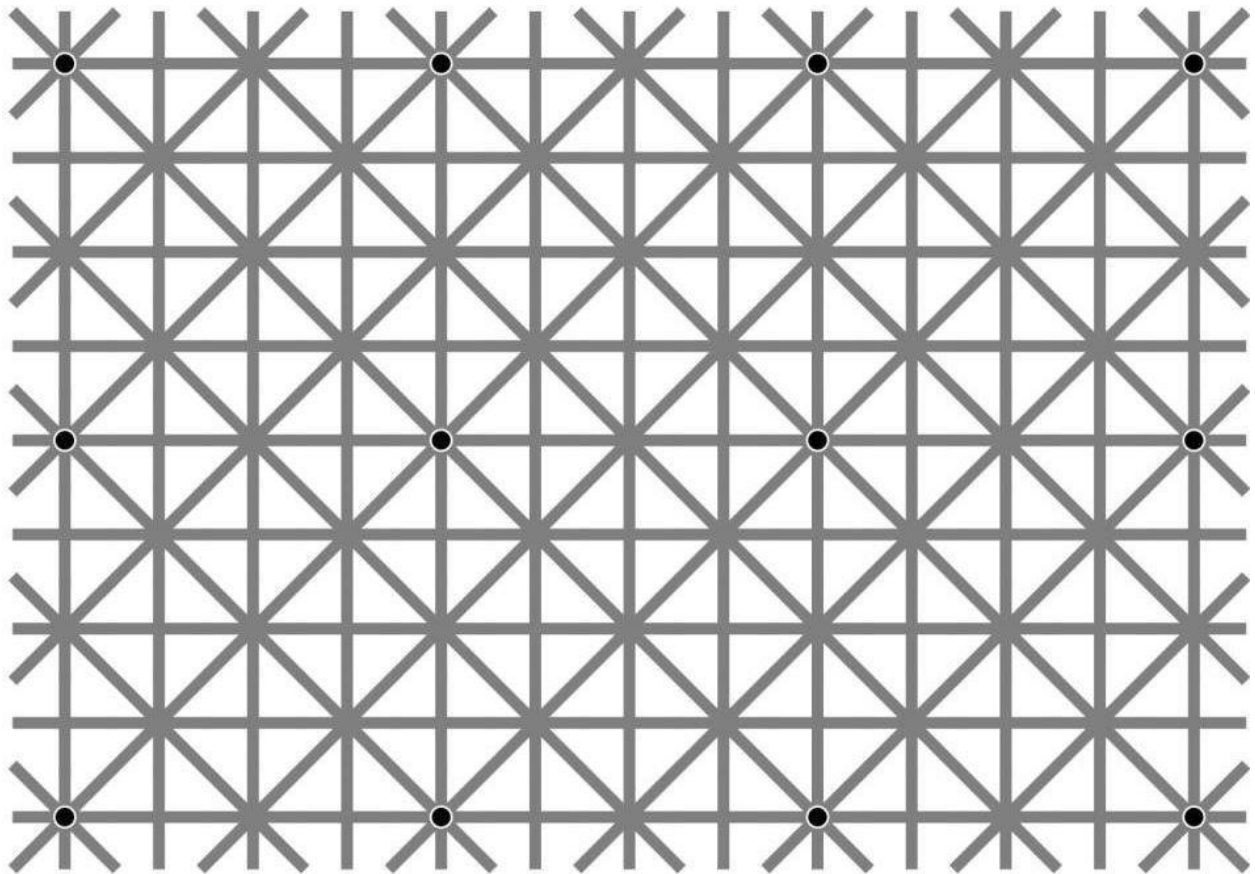


Figure 2: Optical illusion (Will Kerslake²)

In this illusion, you can only see the black dots that you are foveating (focusing your eyes) on. The rest become invisible. This illusion works even when you distance yourself from the image and are able to contain a large part of it in your fovea. A likely explanation is that our brain is aware of the highly regular pattern in the image and imposes this pattern onto our perception on regions we are not focusing on, therefore overriding what our eyes tell it is there. This is an underlying principle for many illusions: we see what we expect to see. Feed-forward summation of objects in the visual scene is insufficient to

² This illusion is traced to twitter post: <https://twitter.com/wkerslake/status/77510533333204992>

explain illusions like this: even though the feed-forward representation of these dots exist, we do not perceive them. Some other processing component, based on knowledge of the visual environment, must then occur. This could be achieved via feeding such knowledge back into earlier visual areas. Indeed, feedback connections to V1 from other cortical regions as well as connections within V1 compose the majority of its input (Kennedy, H., P. Barone, A. Falchier 2000). Their functions, however, are largely unknown.

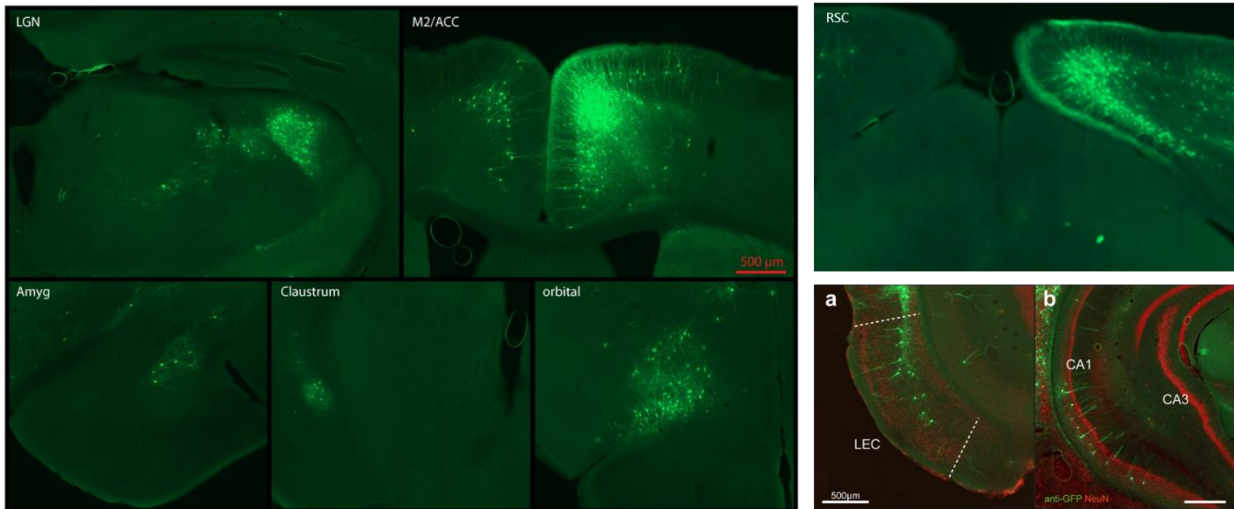


Figure 3: Example images from mouse brains with monosynaptic rabies virus expressing GFP injected into V1. Note the diversity of areas projecting to V1, as well as the relative density of intracortical projections compared to that of the projection from LGN. (Work done by Marcus Leinweber)

A theoretical framework called predictive coding postulates that a role of these feedback projections is to relay predictions of sensory input to earlier sensory areas (Bastos et al. 2012; Rao and Ballard 1999). In these earlier areas, the predictions are then subtracted from the actual input arriving from the environment, and the difference between the two – the error – is fed forward to the same areas that provided the predictions. This error can be used to signal changes or new rules in the environment and update our predictions. For example, consider learning to play tennis. A difficulty that many inexperienced players face is that the racquet frequently does not connect with the ball. Every time this happens, we get an error signal, e.g. in the form of a lack of visual or somatosensory cues triggered by the ball having bounced off the racquet. We expected the racquet to be in a given position that would intercept the ball, based on our knowledge of the position of our arm. However, this was not the case.

³ The virus is complemented with rabies virus glycoprotein in V1 and then “jumps” retrogradely to infect neurons that send axons to V1 (Wickersham et al. 2007).

We therefore use the error signal to update our knowledge of how our arm is positioned and the way we grip the racket translate to how we expect the ball to react.

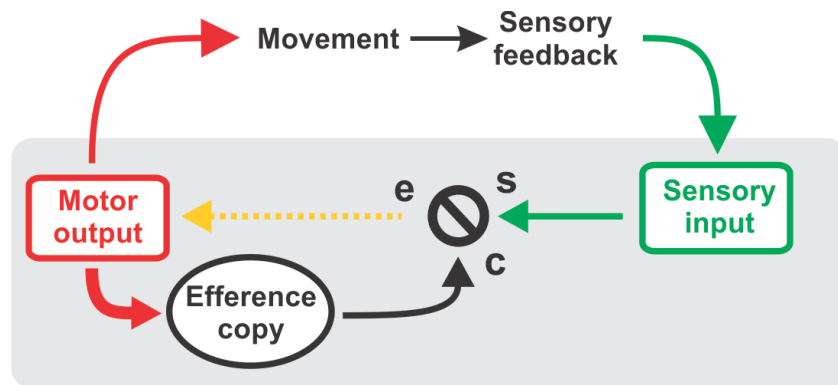


Figure 4: Schematic of sensorimotor learning: incoming sensory input is compared to an efference copy of motor output, and errors are used to guide adjustments. Adapted from (Keller and Hahnloser 2009)

In order to examine the above in functional detail, let's take a closer look at the structure and known functions of primary visual cortex.

The visual system

Anatomy

Retinal ganglion cells send axons from the retina to the lateral geniculate nucleus of the thalamus (LGN) via the optic nerve (Lund 1988). The LGN has classically been thought to be a “relay” center of retinal information to V1, but retinal signals do undergo processing in that area, though the nature of this processing is a subject of debate (Cudeiro and Sillito 2006).

Afferent projections from the LGN terminate in Layer 4 (L4) of primary visual cortex (Felleman and Van Essen 1991). From there, information is largely sent to L2/3, which in turn projects to L5, to L4 of downstream cortical areas, as well as horizontally in L2/3 itself (Felleman and Van Essen 1991). Feedback connections to V1 typically originate in deep layers (L5) of the downstream cortical area, and terminate outside of L4 (Felleman and Van Essen 1991).

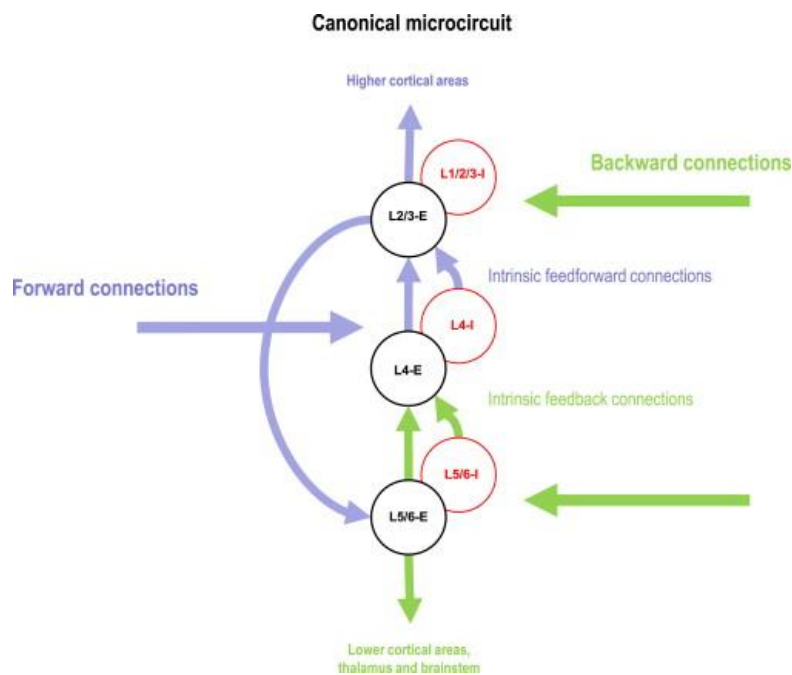


Figure 5: Schematic of a proposed cortical microcircuit (Bastos et al. 2012).

As mentioned previously, although retinal input is communicated to V1 via the LGN, the majority of V1's inputs come from other cortical areas (Salin and Bullier 1995). The role of this input has classically been regarded as modulatory (Crick and Koch 1998; Sherman and Guillery 1998).

However, feedback inputs in the absence of LGN activity and L4 activity are still sufficient to drive activity in L2/3 of V1. Indeed, silencing downstream areas can abolish L2/3 activity in V1 (Mignard and Malpeli 1991).

Function

The prevailing models of visual processing are consistent with the feedforward connectivity of the visual system. In brief, inputs from the LGN are pooled in V1 to produce orientation selectivity for stimuli in each neuron's receptive field. In the "ventral stream", a hypothesized hierarchy of visual areas concerned with the perception and recognition of visual objects (Goodale and Milner 1992), which begins in V1 and terminates in the hippocampal formation, this pattern is repeated as one follows a putative processing hierarchy, such that neurons higher up in this hierarchy are selective to curves, complex shapes (such as faces) and achieve view invariance – the ability to "recognize" an object despite its location and orientation in the environment (Poggio and Riesenhuber 1999; Watanabe et al. 2007). In the "dorsal stream", a similarly envisioned hierarchy specializing in motion perception and the visual guidance of movement (Goodale and Milner 1992), starting from V1 and terminating in parietal multisensory and motor areas in primates, the activity of motion-selective neurons in V1 is fed forward to areas like V3 and MT (Simoncelli and Heeger 1998) and LIP (Blatt, Andersen, and Stoner 1990), which contain complex representations of visual motion.

Traditionally, the role of feedback in the visual system has been confined to a modulatory one, hypothesizing that feedback connections can sharpen or change the gain of neurons representing features in the environment that are important to the animal. This can be done via attentional or other learning-related processes (Zhang et al. 2014). More recently, locomotion has been shown to strongly affect, and indeed drive, responses in L2/3 of V1 (Keller, Bonhoeffer, and Hübener 2012; Niell and Stryker 2010; Polack, Friedman, and Golshani 2013; Saleem et al. 2013), even in the absence of visual input (Keller et al. 2012; Saleem et al. 2013). The role of motor-related inputs to V1 is unclear. It has so far been hypothesized to alter the gain of visual representations, to aid in speed estimation (Saleem et al. 2013), or to confer predictions of visual flow (Keller et al. 2012).

The latter hypothesis is reinforced by the finding of neurons in L2/3 of V1 that signal the mismatch between predicted visual flow (provided by the animal's running speed) and actual visual flow (controlled via the virtual environment) (Keller et al. 2012; Saleem et al. 2013). Indeed, the amplitude of this mismatch signal correlates with the speed of the animal during the visual flow perturbation (Keller et al. 2012). These results are consistent with a predictive coding framework, wherein primary sensory

areas signal the residual errors between actual sensory input and predictions of sensory input are relayed via feedback.

Predictive coding

Imagine sitting on your desk reading this document on your laptop. An inviting cup of coffee lies a bit further away on the tabletop, and you reach for it while still engrossed in the fascinating details of what you're reading. You extend your arm, grasp the cup and bring it to your mouth without ever looking at what you're doing. This is all made possible by the fact that the brain contains an internal representation of the world (Hinton 2000). One important advantage of such an internal model is that it can be used to predict the sensory outcomes of our motor actions, therefore allowing us to adjust movements in real time and to reach for invisible coffee cups.

As in the tennis example earlier, the errors between what we expect to sense based on our previous experience and what actually occurs can be used to update our internal model. As an idea this is uncontroversial. The point of contention lies rather in how this is implemented in the brain. One possibility is that this comparison is done in higher parts of the brain involved in action planning and memory. Predictive coding, on the other hand, hypothesizes that an internal model of the world generates predictions of what we are about to sense based on prior experience, which are fed back to sensory cortices. These predictions are then subtracted from input from the sensory organs, and sensory cortex signals the difference, or error, between the two (Rao and Ballard 1999).

Theoretical Work

This theory is appealing for a number of reasons. Firstly, predictive coding models can account for the emergence of a large fraction of the types of receptive fields found in V1 and other visual areas (Rao and Ballard 1999; Spratling 2010, 2012). The top-down communication of stimulus predictions could possibly also explain the observation that V1 is active during mental imagery (Goebel et al. 1998; Kosslyn et al. 1995) and during illusory perception of contours (Mendola et al. 1999).

Secondly, despite the additional step of forming predictions, it can address issues of redundancy in the sensory world (Huang and Rao 2011). The properties of nearby areas of visual space are typically highly correlated, as objects possess continuity in space. This is true of time as well: objects persist in time. Theoretical work has shown that a predictive coding framework can reduce redundancy in representations by inhibiting the representation of predictable components of visual scenes, both in time and in space (Huang and Rao 2011). A consequence of this is that it maximizes the “surprise”, or information content, of an unexpected stimulus (Friston 2010; Huang and Rao 2011). Furthermore, as

the predicted component inhibits the neurons carrying its representation, the limited dynamic range of individual neurons can be better exploited (Huang and Rao 2011).

Thirdly, predictive coding can be applied to other brain functions, and seen as a fundamental computation of cortex. Predictive coding can account for observations in perceptual decision making (Summerfield et al. 2006; Summerfield and de Lange 2014), and can be extended to motor cortex (Shipp, Adams, and Friston 2013) and account for the properties of the mirror neuron system (Kilner, Friston, and Frith 2007).

Experimental evidence

Despite the abundance of theoretical work on the field, systems neuroscience studies that probe for hypothesized elements of a predictive coding scheme are relatively scarce (Eliades and Wang 2008; Keller et al. 2012; Keller and Hahnloser 2009; Saleem et al. 2013). At least two groups have found a population of neurons in mouse V1 that encode the error between expected and actual visual flow (Keller et al. 2012; Saleem et al. 2013). This work has so far been confined to the exploitation of an internal model of optic flow given motor output as a basis for forming predictions. However, in theory all possible internal representations can be used to generate sensory predictions. Perhaps the most explored such representation is that of space.

Spatial maps have been discovered in the hippocampus (O'Keefe and Dostrovsky 1971) and entorhinal cortex (Hafting et al. 2005), most characteristically in the form of "place cells" and "grid cells", respectively. Place cells in fields CA1 and CA3 of the hippocampus are preferentially active in a specific location of the environment, whereas grid cells in the lateral entorhinal cortex (LEC) tile the environment in a hexagonal pattern.

Internal representations of space are very well poised as experimentally testable sources of predictions in a predictive coding framework, as they have been shown to be updated by sensory cues (Fyhn et al. 2007; O'Keefe and Conway 1978) and persist in virtual environments (Domnisoru, Kinkhabwala, and Tank 2013; Harvey et al. 2009), which allow for experimentally controlled violations of predictions.

The goal of this PhD project is to investigate influences on visual processing from internal representations of space and, in particular, to uncover elements of predictive coding in the form of a) predictions of visual stimuli based on spatial location and b) error signals when these predictions are violated.

Experience-dependent spatial expectations in mouse visual cortex⁴

Indirect evidence for an influence of internal representations on visual processing comes from the findings that hippocampal replay during sleep is accompanied by replay in visual cortex (Ji and Wilson 2007), and from the appearance of theta oscillations in the LFP in visual cortices during locomotion in mice (Niell and Stryker 2010) and during short-term memory tasks in monkeys (Lee et al. 2005). We speculated that if a direct influence of spatial maps on visual processing develops with experience, it could manifest as a prediction of visual stimulus based on spatial location. The underlying conceptual model is that spatial representations of the environment activate the corresponding visual representations of stimuli encountered in specific locations. This would likely be mediated by top-down projections to V1 from areas involved in spatial memory, like the anterior cingulate cortex (ACC) (Frankland 2004; Maviel et al. 2004; Teixeira et al. 2004; Weible et al. 2012). This leads to a number of testable predictions. First, visual representations of the environment should change systematically with increasing experience in a given environment. Second, we should find non-sensory stimulus-predictive responses that are tied to a conjunction of spatial location and the visual stimulus previously encountered at this location. Third, if the stimulus encountered at a given location is different from the one previously encountered at the same location this should lead to detectable mismatch signals.

To probe for the existence of experience-dependent spatial expectations in mouse primary visual cortex, we repeatedly let mice explore a virtual tunnel over the course of several days. Throughout exploration, we chronically recorded the activity of the same 1630 neurons in V1 layer 2/3 of 9 adult C57BL/6 mice, using two-photon imaging of the genetically encoded calcium indicator GCaMP6f (Chen et al. 2013) (AAV2/1-Ef1a-GCaMP6f-WPRE). The experiments in V1 were done by myself, David Mahringer and Hassana Oyibo. For all imaging experiments, mice were head-fixed and free to run on a spherical treadmill (Dombeck et al. 2007; Holscher et al. 2005). Rotation of the spherical treadmill was restricted to forward and backward directions and controlled movement in a virtual tunnel that was projected onto a toroidal screen surrounding the mouse (**Fig. 6a**). Upon reaching the end of the tunnel, mice received a water reward and their position was reset to the beginning of the tunnel. The walls of the virtual tunnel were lined with four different landmark stimuli and five uniform gray areas, marking locations at which one of two orthogonal sinusoidal gratings (henceforth referred to as **A** and **B**) were presented when the mouse reached the corresponding gray area (**Fig. 6b**). This was done to ensure precise control of when the mouse would first see the grating. During the first two sessions the

⁴ All of the data in this chapter have appeared in the publication (Fiser et al. 2016).

sequence of the five grating stimuli was identical (**A-B-A-B-A**) on every traversal (condition 1). In subsequent sessions the identity of the last grating stimulus changed to a **B** on randomly selected traversals (90% **A** and 10% **B** in condition 2; 100% **B** in condition 3; and 10% **A** and 90% **B** in condition 4). In the fifth condition we omitted the grating in position 5 altogether on 10% of randomly selected traversals. Each condition comprised two recording sessions that lasted between 1 and 2 hours and occurred daily (spaced by 24 ± 4 hours, with the exception of condition 5 which immediately followed condition 4). We imaged from the same neurons chronically throughout the duration of the experiment. Animals traversed the tunnel an average of 109 times per session. Each traversal lasted between 10 and 120 seconds. In addition, we measured responses of the same neurons during anesthesia to passive presentations of the tunnel presented at a constant visual flow speed both before the first condition (pre-experience anesthesia) and after the last condition (post-experience anesthesia). In total we recorded the activity of 1147 L2/3 neurons in V1 of 6 animals exposed to conditions 1 through 5 (**Fig. 6b**), of which 899 neurons were responsive to at least one visual element of the tunnel (tunnel responsive; 78.4%, see Methods). We also recorded from 483 neurons in conditions 1 and 2 in an additional 3 animals, of which 436 neurons were classified as tunnel responsive (90.2%; in total 1335 of 1630 or 81.9% of neurons were tunnel responsive).

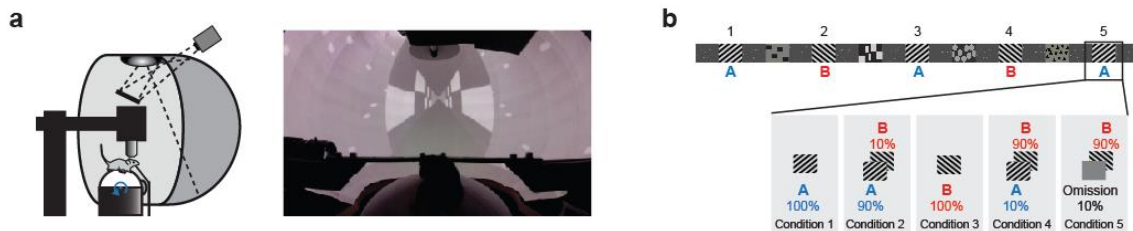


Figure 6: The experimental setup. (a) Left: Schematic of the experimental setup. Right: Photograph of a mouse approaching a landmark stimulus in the virtual tunnel. (b), Schematic representation of the texture lining both walls of the tunnel. Gratings *A* and *B* in positions 1-5 were only shown once the animal reached the corresponding position in the tunnel. In between the grating positions were four permanent landmark stimuli. The probability of encountering an *A* or *B* in position 5 changed with conditions as shown.

V1 activity becomes descriptive of spatial location

To probe for a spatial component in V1 activity, we investigated whether location in the environment modulates neuronal responses to identical visual stimuli. We found that peak calcium fluorescence amplitudes of grating-responsive neurons were different for the presentation of the same grating in different positions in the tunnel (**Fig. 7a,b**). In order to quantify the spatial heterogeneity of neuronal responses in the population, we trained a classifier (Matlab Treebagger, see Methods) to predict which grating location the mouse was traversing in each trial for each behavioral condition using the average population activity within a 667 ms (10 frames) window following each grating onset. Based on V1

activity, the classifier was able to predict not only the identity of the grating the mouse was seeing but also where in the tunnel the mouse was seeing the grating (**Fig. 7c**). Classification performance, measured as the mean of the diagonal of the confusion matrix for each condition (see Methods), significantly increased between conditions 1 and 4 (condition 1: $53.3\% \pm 7.7\%$; condition 4: $81.7\% \pm 4.6\%$, mean \pm s.e.m.; $p = 0.029$, Wilcoxon Rank Sum test). The classifier also performed considerably better in post-experience, compared to pre-experience anesthesia (**Fig. 7c; Fig. 8a**; Ane. Pre: $31.3 \pm 6.2\%$; Ane. Post: $67.2 \pm 7.2\%$; $p = 0.031$, Wilcoxon Rank Sum test). To ensure that the difference in responses to the same stimulus in different locations was not due to running speed tuning (Keller et al. 2012; Niell and Stryker 2010), we trained a classifier to predict the animal's location based on running speed. The classifier did not perform better than chance (**Fig. 8b**). Training a classifier on slow traversals and testing it on fast traversals, and vice versa, yielded classification accuracy that remained well above chance in both cases (**Fig. 8c**), suggesting that speed tuning is not a major contributor of predictive power in the classification. To test if calcium dynamics influence the change in classification performance, we deconvolved raw calcium traces using an exponential kernel with a time constant of 0.5 s (Chen et al. 2013; Yaksi and Friedrich 2006) (**Fig. 8d**; see Methods) and trained the classifier on the deconvolved traces. Average accuracy was slightly decreased when the classifier was trained on deconvolved traces, but the increase between condition 1 and 4 was unchanged (**Fig. 8e**).

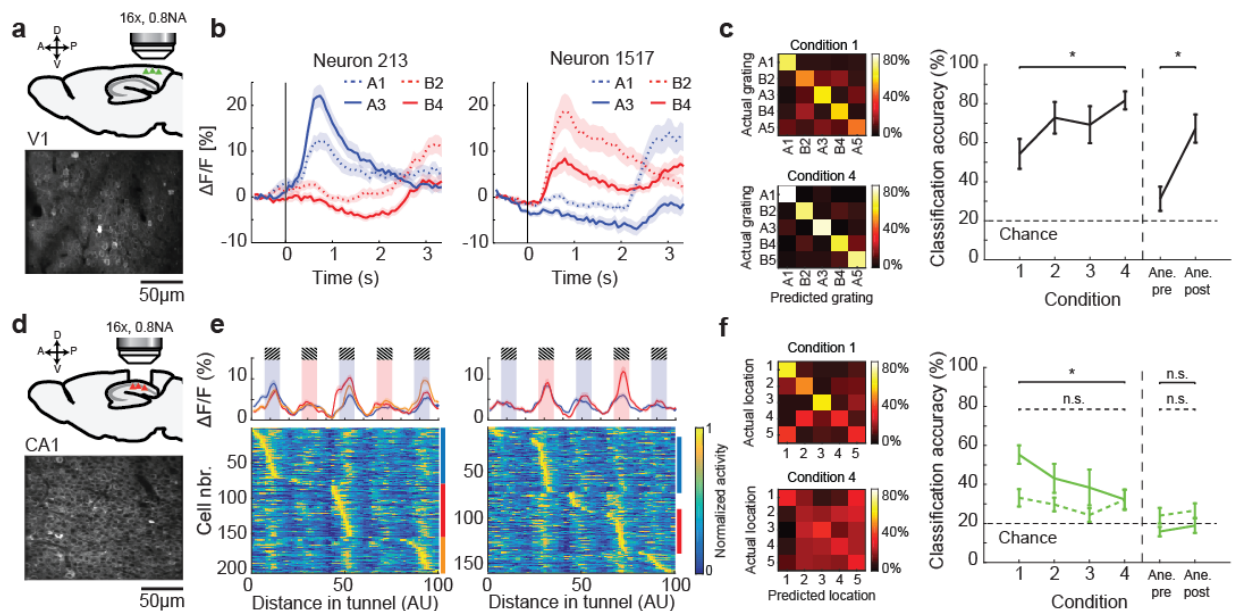


Figure 7: Identical visual stimuli in different spatial locations can elicit different responses in V1 and similar responses in CA1. (a), Top: schematic of V1 imaging strategy. Bottom: example two-photon image of V1 L2/3 neurons. (b), Average responses of an example A-selective neuron (left) and an example B-selective neuron (right) to A1, B2, A3 and B4. Shading indicates s.e.m. across grating presentations (left: 178 presentations; right: 218 presentations). (c), Left: Confusion matrices of the distributions of

classified grating location (x-axis) based on grating onset responses, as a function of actual grating location (y-axis). Right: Mean classification accuracy for all conditions, measured as the mean of the diagonal of the confusion matrix for each condition. Note, for these plots V1 data recorded in conditions 1-2 (from 9 animals) and data recorded in conditions 1-4 (from 6 animals) were combined. Mean \pm s.e.m. across animals. Ane-pre: pre-experience anesthesia; Ane-post: post-experience anesthesia. *: $p = 0.029$ (conditions 1 and 4); $p = 0.031$ (pre- and post-experience anesthesia), Wilcoxon Rank Sum test. (d), Top: Schematic of CA1 imaging strategy. Bottom: Example two-photon image of CA1 pyramidal neurons. (e), Heatmaps showing normalized fluorescence traces of CA1 neurons in condition 1, selective to *A* (left) and *B* (right), sorted by peak position. Traces on top are the mean activity of neurons shown below, highlighted by the blue, red and green vertical bars respectively. (f), Left: As in c, but based on mean response, not grating onset response. Right: as in e, for mean response (solid line) and for grating onset response (dashed line). Ane-pre: pre-experience anesthesia; Ane-post: post-experience anesthesia. Mean \pm s.e.m. across animals ($n = 5$). *: $p = 0.043$, n.s.: $p = 0.931$ (conditions 1 and 4); n.s.: $p = 0.524$, $p = 0.463$ (pre- and post-experience anesthesia), Wilcoxon Rank Sum test.

In addition to a spatial component in V1 activity, an increase in stimulus selectivity could also influence the discriminability of stimuli in the environment. We quantified the selectivity of all neurons to the two grating stimuli *A* and *B* using a selectivity index (SI) as $(R_A - R_B)/(R_A + R_B)$, where R_A is the average response to *A* in positions 1 and 3, and R_B is the average response to *B* in positions 2 and 4; SI was set to 0 for neurons without a significant response to either *A* or *B* (Fig. 9a; see Methods). We found that, with experience, neurons in V1 that are grating-selective become more selective with time (Fig. 9b,c), an effect that cannot be explained by their mean activity (Fig. 9e). Furthermore, the stability of these selective neurons increased with experience, an effect not explained by stability in motor behavior (Fig. 9f).

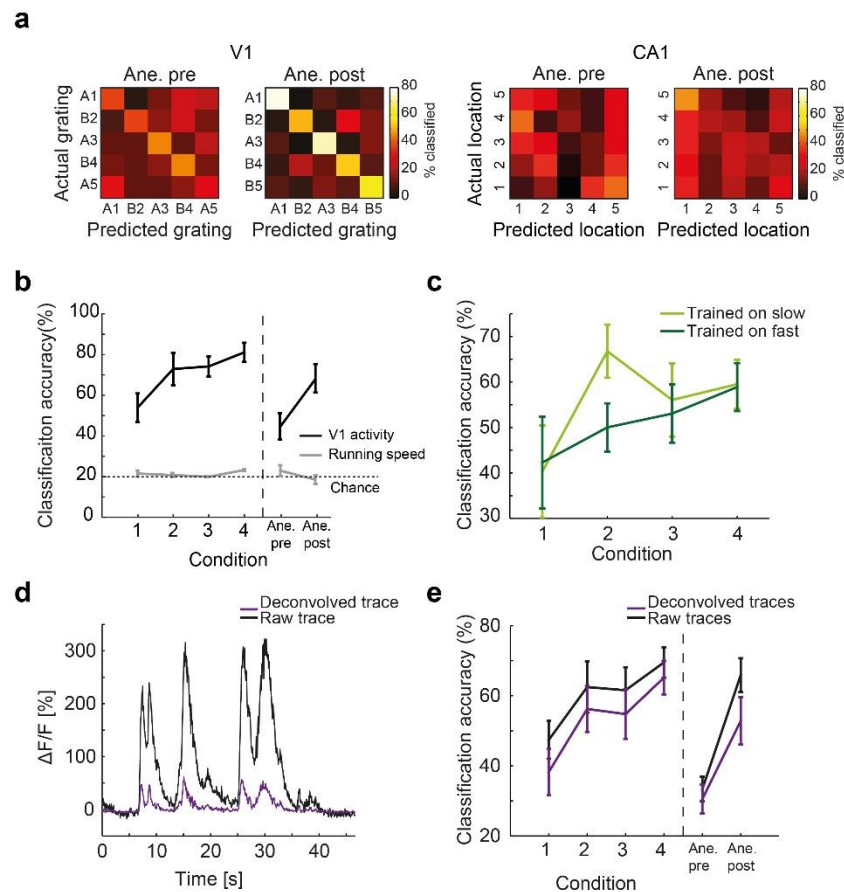


Figure 8: (a), Confusion matrices as in **Fig. 7c,f** for pre- and post-anesthesia in V1 (left) and CA1 (right). In this and subsequent panels, Ane-pre: pre-experience anesthesia; Ane-post: post-experience anesthesia. (b), Accuracy of the classifier predicting the animal's location in the tunnel based on either neural activity (black line) or running speed (gray line) for conditions 1 through 4. Chance level of 20% (dashed line) is given by the 5 possible grating locations in the tunnel. Error bars: s.e.m. (c), Classifier accuracy when trained on neural activity during fast (slow) traversals and tested on slow (fast) traversals. Error bars: s.e.m. (d), Example of a raw activity trace in time (black), and the same trace deconvolved using an exponential deconvolution kernel with a time constant of 0.5 s (magenta). (e), Classifier performance using raw and deconvolved traces. The similar performance suggests that down-sweeps in calcium signals are not the main predictors of spatial location in population activity. Error bars: s.e.m.

To compare dynamics of spatial signals in V1 to potential changes in the spatial map in hippocampus, we chronically recorded the activity of the same 1736 neurons in hippocampal region CA1 in 5 animals exposed to conditions 1 through 5. The experiments in CA1 were performed by myself, David Mahringer and Anders Petersen. Changes in spatial signals in V1 could be the result of changes in the spatial representation in hippocampus, or changes in the way V1 is activated by the spatial representation. In either case, these changes should be reflected in top-down inputs to V1. Activity in CA1 exhibited place-like responses that reflected the pattern of visual stimuli along the tunnel. Neurons showed visually locked responses, and locations with similar visual stimuli elicited similar neural responses (**Fig. 7d,e**). Consistent with previous reports (Ziv et al. 2013), we found that activity patterns were only partially stable over different conditions or days. For 14.8 % of neurons, the location of peak activity in the tunnel was stable over the five behavioral conditions (**10a,c**; within 5% of tunnel length; see Methods). By comparison, in V1, 32% of neurons exhibited a stable location of peak activity (**Fig. 10b,c**). The instability of CA1 activity may have been augmented by the unilateral removal of cortical tissue necessary to image CA1 pyramidal neurons. Previous work, however, has argued that place field responses measured by imaging using similar methods are not different from those measured with electrophysiological techniques (Dombeck et al. 2010). Classification of grating identity based on grating onset responses using CA1 data was only slightly above chance (**Fig. 7f**). This was likely due to the absence of clear grating onset responses (**Fig. 7e**); using mean activity instead of grating onset responses, classification performance in condition 1 was not different from that based on V1 data (**Fig. 7f**; condition 1: $53.9 \pm 4.4\%$; condition 4: $32.3 \pm 4.8\%$, mean \pm s.e.m.; $p = 0.043$, Wilcoxon Rank Sum test). Interestingly however, classification performance decreased with experience, indicating that CA1 activity becomes less informative of spatial location. Furthermore, neurons that were stimulus-selective on average showed decreasing selectivity with experience and maintained high trial-to-trial response variability (**Fig. 9d,f**). This is opposite to the trend we observed in V1 activity where decoding performance increased with experience. The experience-dependent effects found in V1 therefore cannot be explained by a concurrent change of a spatial map in CA1.

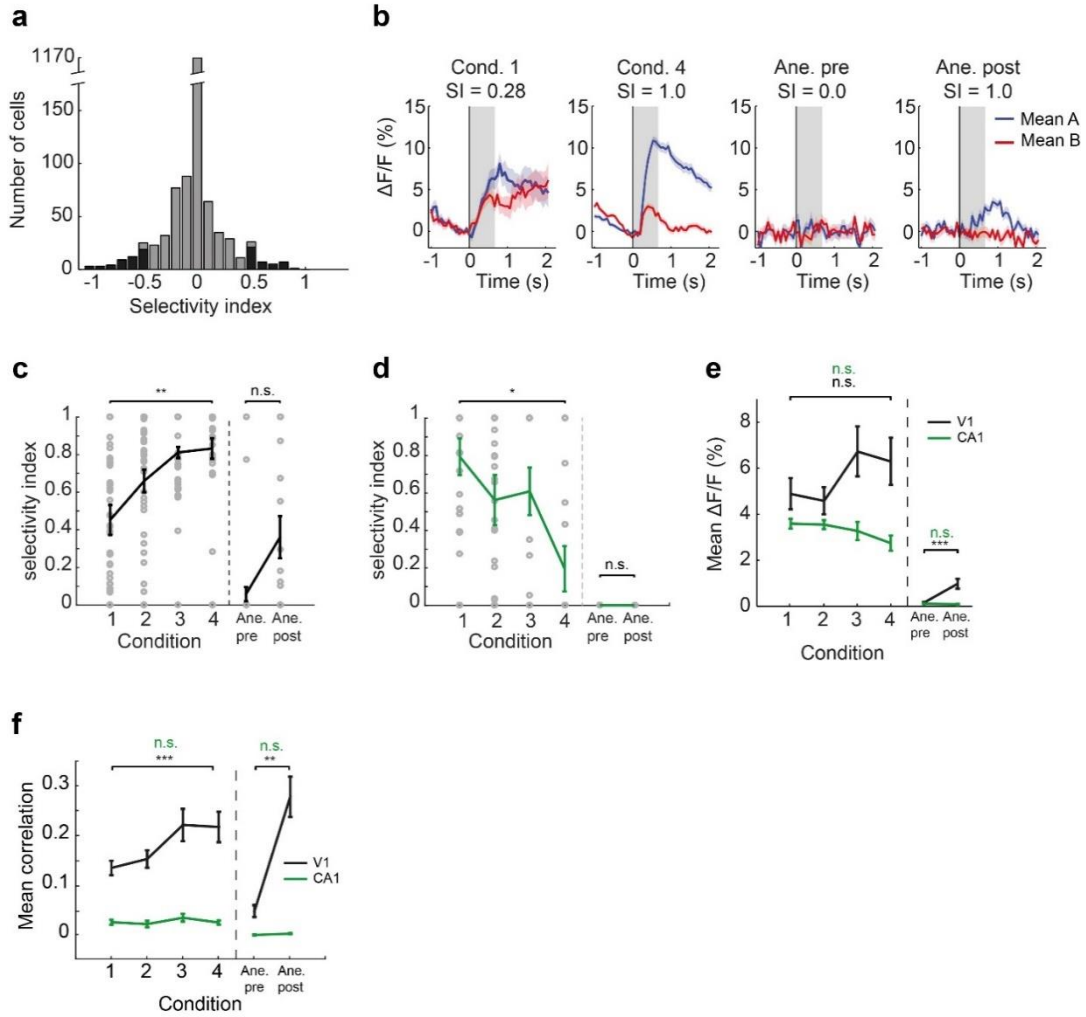


Figure 9: Orientation-selective neurons in V1 (CA1) become more (less) selective with experience. **(a)**, Histogram of the selectivity index (SI) of all neurons averaged across conditions. Black bars: neurons with a $SI \geq 0.5$ (≤ -0.5) were considered to be selective for **A** (**B**). **(b)**, Average responses of an example neuron to **A** (blue lines) and **B** (red lines) across presentations. Gray shading indicates time window used to calculate SI. Blue and red shading indicate s.e.m. **(c)**, Average SI per condition for V1 as in **a**, showing the selectivity indices for each cell included (gray circles). Mean \pm s.e.m. across animals. **: $p = 0.0047$, n.s.: $p = 0.125$, Wilcoxon Rank Sum test. **(d)**, As in **c**, but for CA1. *: $p = 0.009$; n.s.: $p = 1$, Wilcoxon Rank Sum test. **(e)**, Average activity of the same neurons (V1: black; CA1: green) shown in **c** per condition does not increase with time. Mean \pm s.e.m. Conditions 1-4: V1: $p = 0.093$, CA1: $p = 0.12$ (conditions 1-4). Pre- and post-anesthesia: V1: $p = 1.5e-5$; CA1: $p = 1$, Wilcoxon Rank Sum test. **(f)**, Average trial-to-trial stereotypy of activity increased with experience. Shown is the mean correlation coefficient r of activity traces of neurons shown in **c** and **d** in different traversals with the same and across different conditions. Mean \pm s.e.m. (Note, for this analysis traversals were subsampled to match stereotypy of running speed across conditions). Conditions 1-4: V1: $p = 0.00085$; CA1: $p = 0.44$. Pre- and post-anesthesia: V1: $p = 2.5e-8$; CA1: $p = 1$, Wilcoxon Rank Sum test.

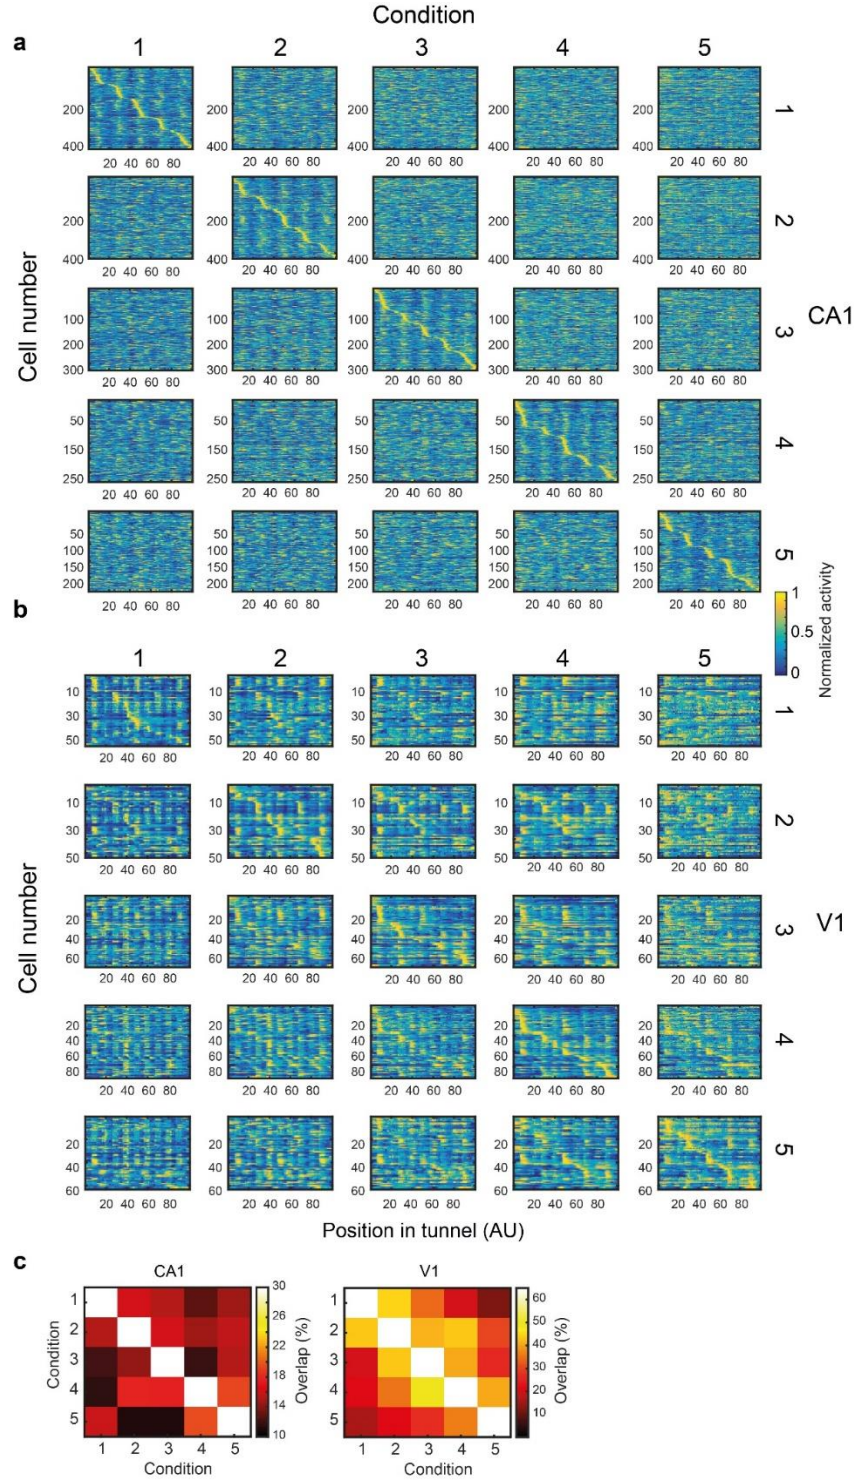


Figure 10: Stability of grating responses between conditions is higher in V1 than in CA1. **(a)**, Normalized activity of grating selective CA1 neurons (SI ≥ 0.1) sorted by position of peak response in the tunnel. Selection of grating responsive neurons and sorting was done on condition 1 for the first row of plots. Data are shown for the same neurons using the same sorting for condition 2 through 5 in the remaining plots of the first row. Similarly for the remaining rows of plots. **(b)**, As in **a**, but for V1 data. **(c)**, Quantification of stability of responses in V1 and CA1. Shown is the fraction of neurons plotted in **a** and **b** that peak within 5% of tunnel length (or one texture length) between the conditions indicated.

V1 develops predictive responses to upcoming visual stimuli

A potential role for the spatial modulation of V1 activity is to enhance the discriminability of similar stimuli in different contexts. In this scenario, a spatial input would trigger predictions of expected visual input at a given spatial position. Indeed, we found a group of neurons (5.6% or 50 of 899 tunnel-responsive neurons) that, with increasing experience in the tunnel, started firing prior to the appearance of the upcoming grating in an **A**- or **B**-selective manner (**Fig. 11a,b**). As responses both preceded the stimulus and signaled the identity of the upcoming stimulus, we will refer to these signals as stimulus-predictive. The response in predictive neurons developed with experience and was absent in the first condition (**Fig. 11b**). In contrast, responses in neurons classified as visual and selective to either **A** or **B** (4.9% or 44 of 899 tunnel-responsive neurons) were present already in the first condition and exhibited a much smaller increase with experience (**Fig. 11c**). Predictive and visual neurons were equally selective for **A** or **B** (predictive neurons: mean SI = 0.79 ± 0.04 ; visual neurons: 0.82 ± 0.03). Using a classifier trained on the activity of predictive neurons preceding the appearance of the stimulus we could predict the identity of the upcoming visual stimulus (**Fig. 11d**; accuracy = $81.4 \pm 5.1\%$, mean \pm s.e.m.). Once present, predictive responses are stable over conditions. The correlation of the mean responses of predictive neurons between conditions 3 and 4 was almost as high as for visual neurons (predictive neurons: $r = 0.81$, $p = 1.1 \cdot 10^{-27}$; visual neurons: $r = 0.95$, $p = 5.4 \cdot 10^{-47}$). Moreover, only one neuron classified as predictive in condition 3 was classified as visual in condition 4.

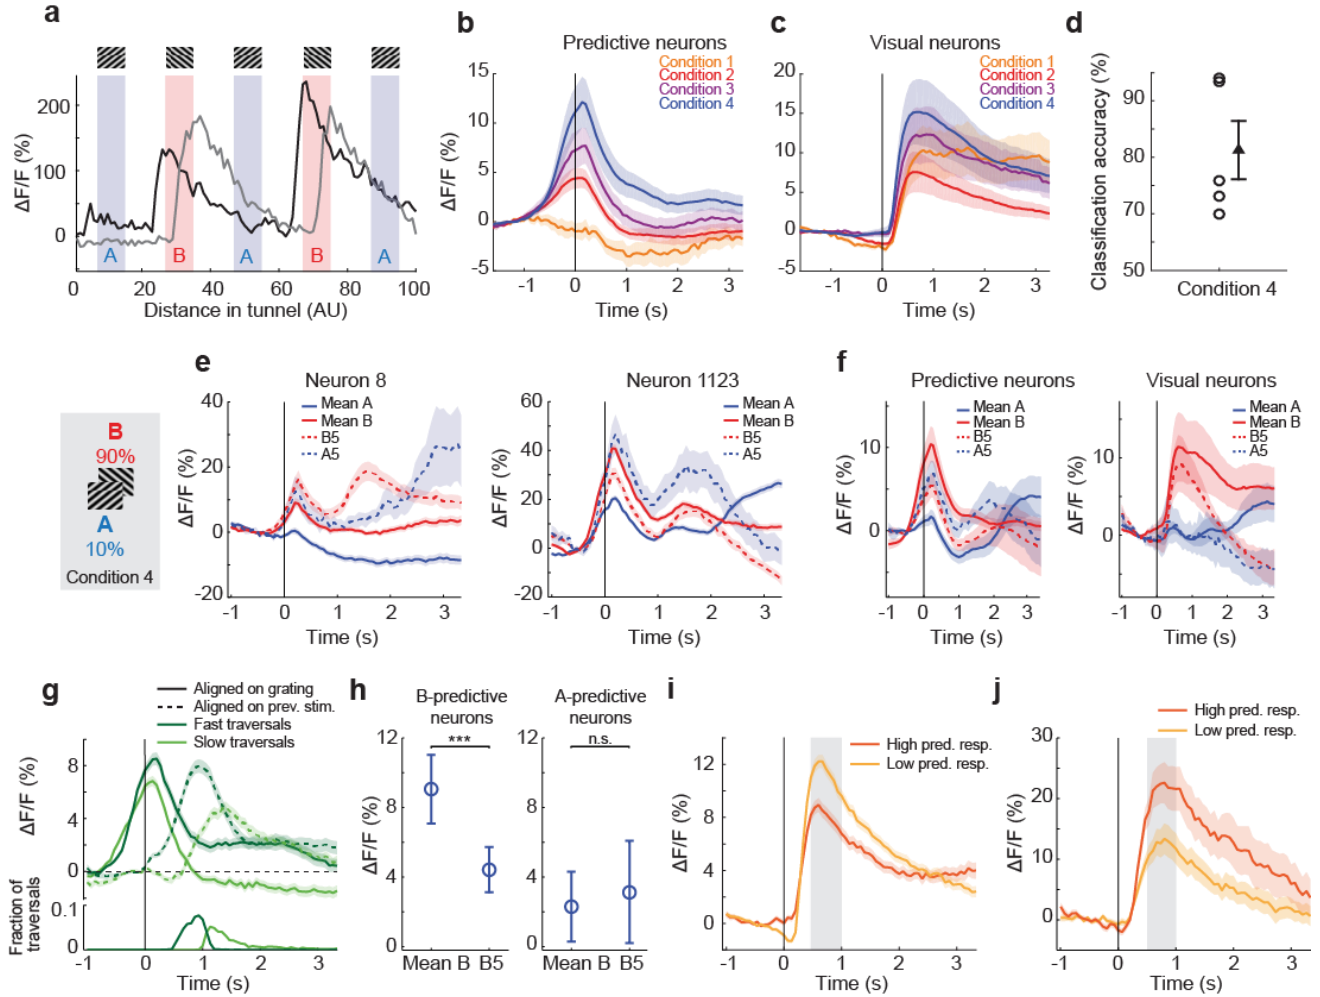


Figure 11: V1 neurons develop predictive responses to approaching visual stimuli with experience. (a), The activity of two **B**-selective neurons during a single traversal of the tunnel. Note that one neuron (black line) fires in anticipation of each **B** presentation, whereas the other fires causally with a delay after the presentation. (b,c), The average response of predictive (b, 50 neurons) and visual (c, 44 neurons) stimulus-selective neurons to their preferred grating orientation in conditions 1-4. In these and all following panels shading along curves indicates s.e.m. (d), Classification accuracy of a classifier trained on the activity of predictive neurons ($n = 50$) to decode grating identity (**A3** vs **B4**) prior to the stimulus (-333ms to 0ms) in condition 4. Circles: individual sites (5 sites); Triangle: Mean; Error bars: s.e.m. Data in d-i are from condition 4. (e), Average responses of two example **B**-selective neurons to mean **A1** & **A3** (blue line), mean **B2** & **B4** (red line), unexpected **A5** (blue dashed), and the expected **B5** (red dashed). (f), Average responses of predictive ($n = 50$) and visual ($n = 44$) neurons to mean **A1** & **A3**, mean **B2** & **B4**, expected **B5** and unexpected **A5**. (g), Upper panel: Responses of predictive neurons ($n = 50$) aligned to either previous landmark stimulus (dashed lines) or upcoming grating stimulus (solid lines) for fast (dark green) and slow traversals (light green). Note that responses for slow and fast traversals align best with upcoming grating onset. Lower panel: histogram of time between previous landmark stimulus and upcoming grating onset for fast (dark green) and slow (light green) traversals. Fast and slow traversals were classified by mean running speed in a window of 467 ms (7 frames) preceding onset of the grating stimulus. (h), Strength of predictive responses of **B**-predictive neurons ($n = 39$, left) and **A**-predictive neurons ($n = 11$, right) to **B2** & **B4** (mean **B**) and **B5** in condition 4 where the animal always encountered **B** in position 2 and 4, but only with 90% probability in position 5. Mean \pm s.e.m. across neurons. (***: $p = 0.00015$, n.s.: $p = 0.7$, Wilcoxon Signed-Rank test). (i), Responses of visual neurons on traversals of high (orange) and low (yellow) activity in predictive neurons. Strong predictive activity before a grating (in the top 20%, 787 grating presentations) correlated with weak visually driven responses, and vice versa (in the bottom 20%, 787 grating presentations). Mean responses are calculated in the window indicated with gray shading. (j), Average responses of visually selective neurons in response to the unexpected **A5** in traversals with weak (yellow; 51 **A5** presentations) and strong (orange; 52 **A5** presentations) activity in predictive **B**-selective neurons. The higher the activity in predictive **B**-selective neurons, the higher the mean visual responses to the unexpected **A** (mean responses are calculated in the window indicated with gray shading).

In conditions 2 and 4 we presented a different grating on 10% of randomly selected traversals in the final location (**Fig. 6b**). On these traversals, with an unexpected grating in the final location, stimulus-predictive neurons fired as if the predicted grating would appear, but visual neurons fired in response to the actual grating shown (**Fig. 11e,f**). Given that stimulus-predictive neurons were as selective for the upcoming stimulus (**A** or **B**) on average as visual neurons, it is unlikely that predictive responses are responses to the preceding stimulus. To confirm this, we aligned responses of stimulus-predictive neurons to either the preceding landmark stimulus or the upcoming grating for fast and slow traversals separately. Alignment of the responses for fast and slow traversals should be best for the stimulus (previous or upcoming) that actually drives the responses. We found that responses were best aligned with the upcoming stimulus (**Fig. 11g**), and thus are best explained by distance, and not time, from the last stimulus. This implies that predictive activity relies on spatial location to signal the upcoming visual stimulus.

Stimulus-predictive signals could reflect the frequency of having encountered a certain stimulus in a specific location. Thus predictive signals should be higher when the same stimulus is always encountered as opposed to when in a specific location different stimuli were encountered during experience. Therefore, the predictive response to **B5** should be lower than the predictive response to mean **B**, as the stimulus presented in position 5 varied with session (**Fig. 6b**), whereas in position 2 and 4 the animal always encountered a **B**. This was indeed the case: in condition 4, **B**-selective predictive neurons were significantly less active prior to grating **B5** (90% **B**) than on average to gratings **B2** and **B4** (100% **B**) (**Fig. 11h**, mean **B**: $\Delta F/F = 9.1\% \pm 2\%$; **B5**: $\Delta F/F = 4.4\% \pm 1.3\%$; $p = 0.00015$, Wilcoxon Signed-Rank test). Conversely, one could argue that **A**-predictive neurons should be more active prior to **B5** as the animal encounters an **A** in this location on 10% of the traversals. A rare encounter, however, did not lead to a measureable increase in predictive activity in V1 (**Fig. 11h**, mean **B**: $\Delta F/F = 2.3\%$; $\pm 2\%$ **B5**: $\Delta F/F = 3.1\% \pm 2.9\%$; $p = 0.7$, Wilcoxon Signed-Rank test).

In predictive coding models, the primary visual cortex communicates the error between predicted and actual visual stimuli to downstream visual areas (Bastos et al. 2012; Rao and Ballard 1999). If predictive activity scales with the frequency of having encountered a visual stimulus in a particular location, then the strength of the visual response to the stimulus may signal the surprise of seeing it. This would be reflected in lower visually-driven activity on trials when prediction of a grating was high. We observed that on traversals with high predictive activity preceding each grating, visually evoked activity to the grating was lower ($\Delta F/F = 8.0\% \pm 0.5\%$, mean \pm s.e.m.) than on traversals with low predictive activity ($\Delta F/F = 11.2\% \pm 0.5\%$, mean \pm s.e.m.; **Fig. 11i**). In sum, this suggests that stronger visual responses report

the discrepancy between predicted and actual visual input, and that activity in stimulus-predictive neurons may lead to a reduction of visual responses.

If activity in stimulus-predictive neurons indeed signals the identity of the upcoming stimulus, one would expect a difference in the visual response when an unpredicted stimulus is encountered. We found that the responses of **A**-selective visual neurons to the unexpected **A** at position 5 were stronger when **B**-predictive neurons fired strongly in anticipation to the grating presentation (**Fig. 11j**). Traversals were split into two groups by median amplitude of the response of predictive neurons (average visual responses on high traversals with high predictive activity: $21.4\% \pm 3.1\% \Delta F/F$; and on traversals with low predictive activity $12.5\% \pm 2.2\% \Delta F/F$; mean \pm s.e.m.; $p = 0.02$, Wilcoxon Rank Sum test). Altogether, these findings indicate that the strength of predictive responses preceding a stimulus strongly affect the visual responses to it, suggesting a dynamic interplay between stimulus prediction and stimulus response.

ACC conveys stimulus predictive signals to V1

As the source of predictive signals must be extra-retinal, one would expect that at least some of the top-down inputs to primary visual cortex exhibit signals that are stimulus-predictive, and that this predictive input develops with experience. One of the candidate structures for such top-down inputs to V1 is the anterior cingulate cortex (ACC). ACC is known to project to V1 (Miller and Vogt 1984; Vogt and Miller 1983; Zhang et al. 2014), and has been shown to be involved in long-term memory storage (Frankland 2004; Maviel et al. 2004; Teixeira et al. 2004; Weible et al. 2012). To test if spatial information could be relayed to V1 via ACC, we recorded the activity of ACC axons in layer 1 of V1 in condition 1 (3513 axons, 5 sites) and in condition 4 (8599 axons, 10 sites) in 5 animals (**Fig. 12a**; see Methods). The ACC recordings were performed by Marcus Leinweber. Note that, unlike for the V1 and CA1 experiments, we were unable to chronically record from the same ACC axons on different days. The combination of the high density of ACC axons in layer 1 of V1 and the low baseline fluorescence made it impossible for us to ensure that we were recording from the same axons on different days. However, there likely was a large overlap between the axons recorded on different days as imaging regions (5 of 10) were realigned based on blood vessel patterns. We imaged activity on the first and the sixth (2 sites, 1 animal) or seventh (8 sites, 4 animals) day in the tunnel. As the total experience in the tunnel between the two imaging time points was comparable to the difference between condition 1 and condition 4 in the V1 and CA1 data we will use the same nomenclature for the ACC data.

When classifying grating position based on the activity of ACC axons we found an increase in classification accuracy with experience, similar to the increase in V1 (**Fig. 12b**; condition 1: $31.7 \pm 18.3\%$; condition 4: $88.6 \pm 2.4\%$, mean \pm s.e.m.; $p = 0.03$, Wilcoxon Rank Sum test). We then compared the activity of ACC axons that exhibit selective responses for either **A** or **B** in early and late conditions. Using the same criteria as for responses in V1 (**Fig. 11b,c**) we were able to classify axons as either predictive or visual in both early and late conditions (**Fig. 12c, d**). We found that in early conditions there were visual responses but no predictive responses and that in late conditions stimulus-predictive responses emerged (**Fig. 12d**). Note, even axons classified as visual exhibited activity that preceded the presentation of the stimulus in condition 4. The contribution of these predictive responses to the total population response to gratings was larger in ACC than in V1 (**Fig. 13**). As in V1 (**Fig. 11h**), the strength of predictive responses depended on the reliability of having encountered a certain stimulus in a specific location. Predictive responses to **B2** and **B4** (100% **B**) were larger than the predictive responses to **B5** (90% **B** – 10% **A**, in condition 4; Mean **B**: $\Delta F/F = 6.5 \pm 0.6\%$; **B5**: $\Delta F/F = 4.2\% \pm 0.8\%$; mean \pm s.e.m.; $p = 0.00028$, Wilcoxon Signed-Rank test), and different from activity in V1, **A**-predictive activity was higher prior to **B5** as compared to **B2** and **B4** (**Fig. 12e**; Mean **B**: $\Delta F/F = 2.2\% \pm 0.3\%$; **B5**: $\Delta F/F = 4.4\% \pm 0.5\%$; mean \pm s.e.m.; $p = 0.00015$, Wilcoxon Signed-Rank test). Predictive inputs from ACC could signal spatial location (in spatial coordinates) or signal the predicted visual stimulus (in visual coordinates). To test if axons classified as visual in condition 4 were actually visually driven, we compared responses of expected and unexpected presentations of **A** or **B** (condition 4). Responses of axons that were **B**-selective had visual responses to an expected **B5** comparable to the mean response to **B**. Responses to the unexpected **A** still showed predictive activity, but diverged from **B** responses following stimulus onset (**Fig. 12f**). Conversely, axons selective for **A** exhibited predictive responses to an expected **A1** or **A3** and only small responses to the expected **B5**, but showed clear visual responses to the unexpected **A5** (**Fig. 12g**). Thus, predictive signals in V1 are likely conveyed by top-down signals carrying an expectation of the visual input based on spatial location.

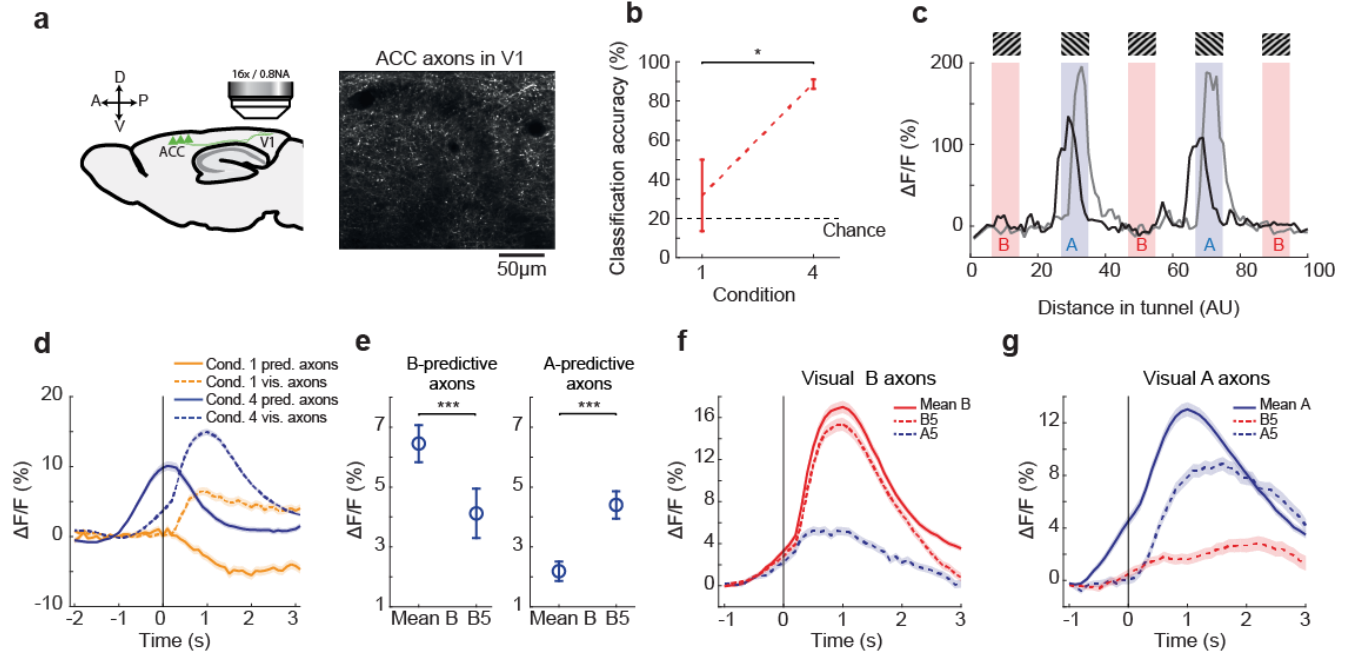


Figure 12: ACC projections to V1 carry visual stimulus predictions. (a), Left: schematic of ACC axon imaging strategy. Right: Example two-photon image of ACC axons in V1. (b), Classification accuracy for a classifier trained to decode grating location based on grating onset responses (as in Fig. 1e,h). The accuracy of the classifier increases with time. Mean \pm s.e.m. across sites (condition 1: 3 sites; condition 4: 10 sites). *: $p = 0.03$, Wilcoxon Rank Sum test. (c), Activity of two A-selective axons during a single traversal of the tunnel. One axon (black line) fires in anticipation of each A, whereas the other (gray line) peaks after each stimulus. (d), As in V1, stimulus-predictive responses emerge with experience. Orange lines indicate mean activity of predictive (solid; $n = 654$) and visual (dashed; $n = 1377$) axons in condition 1, whereas blue lines indicate activity of the corresponding axons in condition 4 (736 predictive and 2559 visual axons). Shading indicates s.e.m. across axons. (e), As in Fig. 2h, for V1-projecting ACC axons. Strength of predictive responses of B-predictive axons (left) and A-predictive axons (right) to B2 & B4 (mean B) and B5 in condition 4 where the animal always encountered B in position 2 and 4, but only with 90% probability in position 5. Mean \pm s.e.m. across axons (312 A-predictive and 500 B-predictive axons). ***: $p = 0.00015$ (A-predictive), $p = 0.00028$ (B-predictive), Wilcoxon Signed-Rank test. (f), Activity of visual B axons ($n = 1175$) to mean B (B2 & B4) (solid red), expected B5 (dashed red) and unexpected A5 (dashed blue). (g), As in f, but for visual A axons ($n = 1384$). Note the visually evoked response to the unexpected A5. Shading in f and g indicates s.e.m.

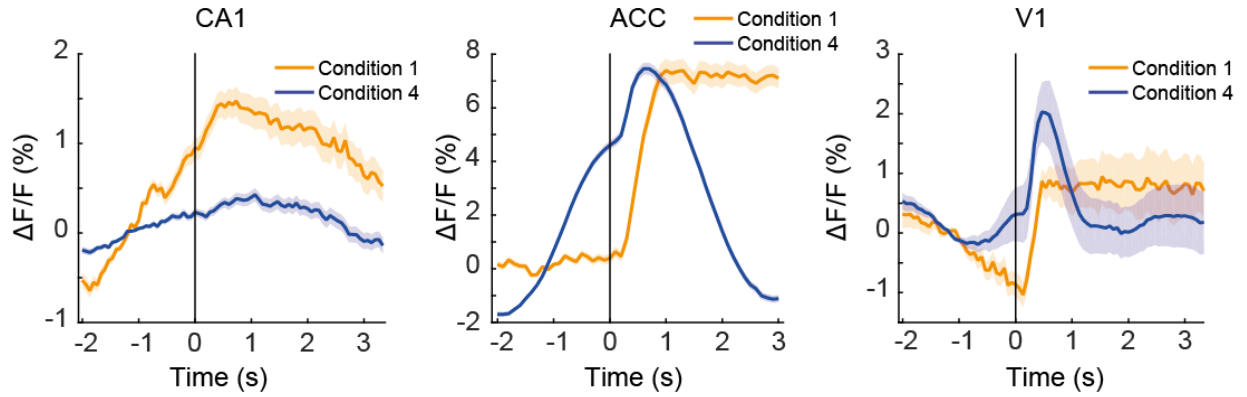


Figure 13: Grating responses in condition 1 (orange) and condition 4 (blue) for grating-selective (SI > 0.1) CA1 and V1 neurons and ACC axons.

Omitting an expected stimulus drives strong responses in V1

To probe if expectations can drive responses in V1 in the absence of a visual stimulus, we omitted the final stimulus altogether in 10% of randomly selected traversals (condition 5). In these omission traversals, no grating would appear as the mouse reached the gray area marking the location of the final grating presentation. The omission of the stimulus elicited a strong response in the V1 population (**Fig. 14a, Fig. 15**). Moreover, a subset of neurons selectively responded to the omission (**Fig. 14b**; 2.3% or 21 of 899 tunnel-responsive neurons). If this omission response indeed signals a deviation between predicted and actual visual input, one would expect the strength of the predictive response to correlate with the omission response. To test this, we split all traversals into two categories depending on how strong the average response of predictive **A**- and **B**-selective neurons was prior to the omission. The average omission response in trials with low predictive activity was significantly smaller than in trials with high predictive activity (**Fig. 4c**; omission-evoked activity on trials with high predictive activity: $\Delta F/F = 35.2\% \pm 8.5\%$; on trials with low predictive activity: $\Delta F/F = 16.6\% \pm 2.2\%$, mean \pm s.e.m.; $p = 0.00018$, Wilcoxon Rank Sum test with bootstrapping, see Methods).

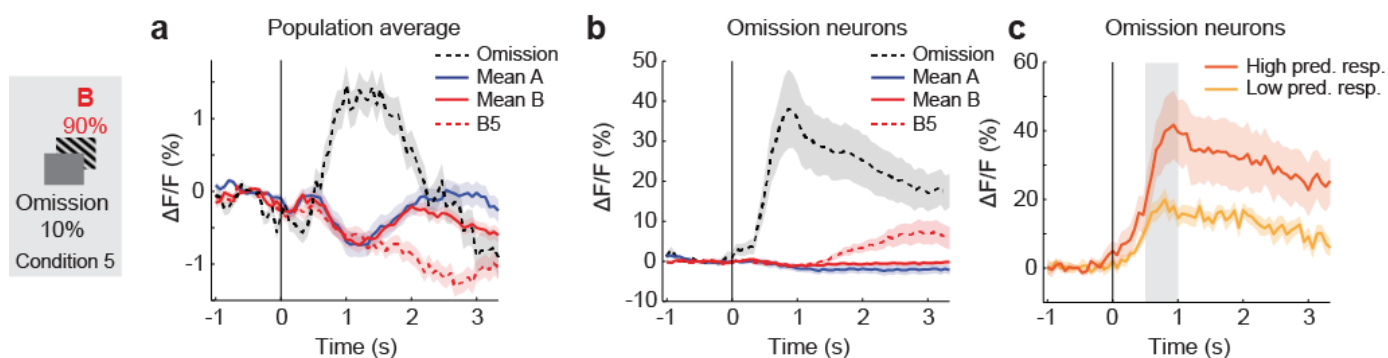


Figure 14: The omission of an expected grating strongly drives activity in V1. **(a)**, Average population response (1147 neurons) to the omission of grating B5 (black dashed line) in comparison to the average response to **A** (blue line), **B** (red line) and **B5** (red dashed line). Shading indicates s.e.m. across neurons. **(b)**, As in **a**, but for omission selective neurons (21 of 899 tunnel-responsive). Shading indicates s.e.m. across neurons. **(c)**, Average omission responses in omission selective neurons on traversals of high activity in predictive neurons (orange line; 13 traversals) and on traversals with low activity in predictive neurons (yellow line; 12 traversals). Gray shading indicates window over which mean activity was calculated. Shading indicates s.e.m. over traversals.

The lack of local visual flow during the grating omissions raises the possibility that these responses are instances of sensorimotor mismatch (Keller et al. 2012). To test this, we compared the omission responses of the omission-selective neurons to their responses to the expected uniform gray stimulus that the animals encountered while running at the beginning and end of the tunnel. We found no response to the expected gray stimulus in omission selective neurons (**Fig. 14b**). Thus, omission responses can best be explained by a deviation between expected and actual visual stimulus based on what the mouse had seen in this position on previous traversals. Furthermore, omission responses were absent in ACC axons (**Fig. 14c**), suggesting that visual cortex compares visual stimulus predictions, relayed by top-down cortical input, to actual visual input.

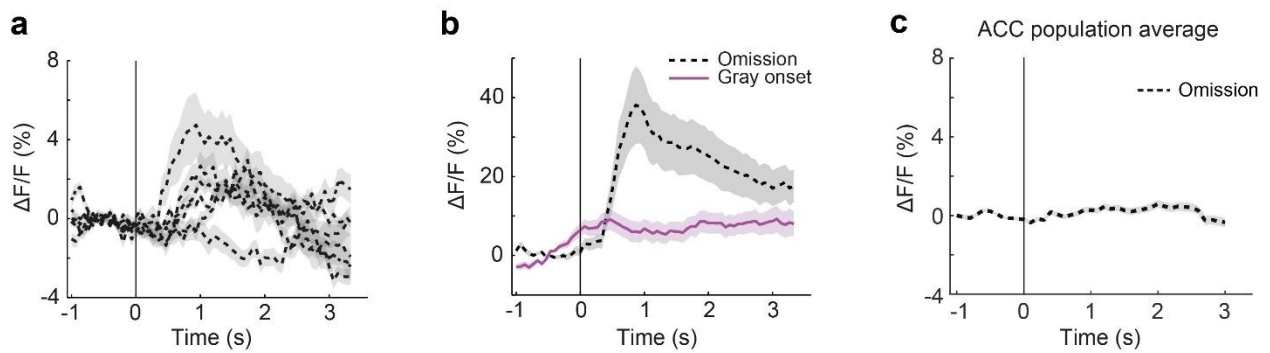


Figure 15: (a), Population mean response to grating omission for each animal (dashed black lines, $n=6$). Shading indicates s.e.m. (b), Responses of omission selective neurons to omission events and onsets of gray areas in the tunnel locations that are always encountered as gray. (c), Average ACC axon population response (8599 axons) to the grating omission.

An additional possibility is that omission-related activity in V1 can be explained via a change in motor activity specific to the stimulus omission. However, running speed during the omission was identical to that during the presentation of the expected B5 grating (**Fig. 16**)

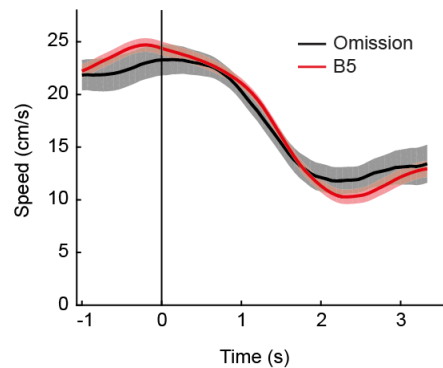


Figure. 16: Average running speed following the expected B5 grating (red) and the stimulus omission (black). Shading indicates s.e.m.

Discussion

In this work, we have shown that experience shapes the activity of neurons in L2/3 of primary visual cortex. Specifically, V1 develops stimulus-predictive responses that are tied to spatial location; the omission of an expected stimulus drives responses in V1 that are stronger than those to visual stimuli; and the activity of V1 neurons becomes more descriptive of both spatial location and stimulus identity. Furthermore, the activity of visually driven responses scales with the activity of stimulus-predictive neurons preceding the stimulus in a manner consistent with the hypothesis that feedforward activity in sensory cortices signals the errors between predicted and actual input (Rao and Ballard 1999).

Additionally, we have identified the anterior cingulate cortex (ACC) as a potential source of these stimulus predictions. Our results are therefore consistent with a predictive coding framework, where an internal model, in this case a model of space, generates predictions of upcoming sensory input. These findings build upon previous work showing that motor-related signals are integrated with sensory signals in primary sensory cortices to generate sensorimotor error signals (Eliades and Wang 2008; Keller et al. 2012; Keller and Hahnloser 2009; Saleem et al. 2013).

Our results cannot be explained by timing dependent recall of activity or reward anticipation. Cue-triggered recall of activity, for example, has been shown to occur in visual cortex after repeated experience of rapid sequences of stimuli (Gavornik and Bear 2014) and fast moving spots (Xu et al. 2012). Sequence learning, however, is specific to the timing used for training; a change in timing of the stimuli of as little as 150 ms abolishes the effect (Gavornik and Bear 2014). Predictive responses in our experiments persisted even though trial to trial differences in traversal times were on the order of tens of seconds, and predictive activity traces aligned to the predicted stimulus (and not to a previous stimulus) in a manner that is invariant to locomotion speed. Moreover, the cue-triggered recall was only observed when the animal was anesthetized, or the animal was awake but cortex was in a synchronized state (characteristic of quiet wakefulness), and the effect was absent when cortex was in a desynchronized state (characteristic of motor behavior) (Xu et al. 2012). Another effect that has been shown to drive activity in visual cortex is reward anticipation (Shuler and Bear 2006). Neurons that code for a rewarded stimulus are selectively activated in anticipation of reward-predicting stimuli (Poort et al. 2015). However, as in our data predictive activity is stimulus selective and omission responses are independent of the reward delivered at the end of the tunnel, reward anticipation cannot explain the spatial modulation we describe here.

ACC as part of the brain's generative model

We have identified ACC as a potential source of predictions of visual input to V1, suggesting that it contains internal representations of the environment and can thus serve as part of the brain's generative model. Moreover, unpublished work from our lab (Leinweber et al., in preparation) suggests that the ACC projection to V1 carries predictions of visual flow. Furthermore, the previously published report that this projection modulates surround suppression (Zhang et al. 2014) is not inconsistent with the idea of low-level stimulus predictions being fed back to V1, as predictive coding models can account for the formation of these types of receptive fields (Rao and Ballard 1999).

An additional argument in favor of this idea is that, as in V1, the activity of V1-projecting ACC axons also scales with the probability of a stimulus being at a certain location (**Fig.12e**), suggesting that its representation is updated based on the history of sensory input. However, this study provided no insight into how an internal representation of the virtual environment in ACC could be updated using feedforward input. A handle on this could be obtained by imaging from cell bodies in ACC while perturbing the activity of ACC-projecting V1 neurons during the animal's experience in the environment.

However, the study described in this thesis does not provide evidence that this projection is necessary for the formation of stimulus-predictive responses in V1. An ideal experiment to perform would be to selectively silence V1-projecting ACC axons optogenetically preceding the grating. This could lead to a reduction of stimulus-predictive signals in V1, and possibly a potentiation of visually driven responses. However, at the time of writing this document, this technique has not been possible in our hands, likely for reasons mentioned in (Mahn et al. 2016). It would still be interesting, however, to locally activate ACC axons in V1 preceding a grating presentation, as it could lead to an error signal in V1.

Furthermore, ACC is only one of the many non-sensory cortical areas providing input to V1. Since there is evidence of more than one type of prediction originating in ACC this begs the question of what types of predictions other downstream areas might be feeding back. Anatomically, based on its connectivity with the hippocampal formation (Sugar et al. 2011), the retrosplenial cortex (RSC) is equally well poised to relay stimulus predictions based on spatial location to V1. The presence of head-direction and other self-motion responses in RSC (Cho and Sharp 2001) suggests that a potential role for this input could be to relay predictions of stimulus based on head orientation and motion cues. The projection from orbitofrontal cortex could relay predictions of stimuli associated with value-based decision making (Feierstein et al. 2006; Schoenbaum, Chiba, and Gallagher 1998).

Another major source of input to V1 is of course secondary visual cortex. It would be of major interest to determine whether predictions follow a progression from high level “abstractions” to low level retinotopic visual representations as they are passed downwards through the visual hierarchy.

Why do visually driven responses persist in the presence of predictions?

In the predictive coding framework, predicted elements of the sensory scene are suppressed. The free-energy formulation of predictive coding (Friston 2010) postulates that the brain seeks to minimize “surprise” and spurious activity. However, in this work and others (Keller et al. 2012; Saleem et al. 2013) that report error signals, visually driven responses persist. The most intuitive explanation for this is that the internal model is imperfect and therefore so are the predictions it generates, thus allowing residual error following every prediction. This is likely a realistic argument to make, given that creating a detailed representation from a general, high-level concept is a poorly posed problem. Since residual errors are used to update the internal model, this would imply that the model is constantly updated even in what should be the absence of errors. Though this seems like an inefficient state of affairs, arguments can be made as to why this could be rational.

1. The world is never static: we never see something exactly the same way twice. Whether due to environmental factors (viewing angle, lighting, changes in the object itself) or simply due to the fact that we constantly perform microsaccades (Rolfs 2009), visual (and other sensory) stimuli are never completely invariant spatially or temporally. This suggests that predictions can only operate within a certain boundary of accuracy by design. The need to constantly update our internal model by having imperfect prediction machinery is therefore useful: it allows us to learn more about an object each time we view it, and potentially allows us to make better predictions about similar objects or contexts further down the road (e.g. viewing things when they are wet).

Additionally, other sensory and non-sensory modalities influence our internal model. We often associate visual stimuli with sounds, textures and contexts. These are also not invariant each time we view a stimulus, and as such may affect our predictions. Furthermore, if generating predictions happens via a process similar to memory retrieval, which is possible, since predictions are by design based in memory, each prediction generated will be different (Bridge and Paller 2012).

This explanation implicitly states that the presence of visually driven responses in the studies mentioned above is also due to the stimuli not being completely identical with each

presentation. A lack of stimulus invariance could therefore likely affect both the internal model and feedforward components.

2. Counterintuitively, having an inaccurate internal model can be not only necessary, due to a lack of invariance, but also beneficial from a computational perspective. A very “sharp” prediction would potentiate feedforward neurons to strongly signal small deviations from prediction, and this could therefore saturate responses in the event that a stimulus strongly deviates.

Predictions could therefore be “fuzzy”. This also allows the brain to generate predictions using finite connectivity.

However, this is only beneficial in the case where predictions are subtracted from feedforward sensory input. In the case where a prediction is stronger than sensory input, for example in the case of running during visual flow halts (Keller et al. 2012; Saleem et al. 2013), the symmetrical computation needs to take place (since neurons most likely cannot compute absolutes): sensory input should also be subtracted from predictions ($E = P - I$; E being the error, P the prediction and I the sensory input). Generating fuzzy predictions in this case would seem counterproductive, as subtracting the sensory signal would always result in a residual error. A number of possibilities exist for how the brain could deal with this conundrum. One is that it uses recurrent activity to “sharpen” the prediction. This could also apply to the first computation described ($E = I - P$), and would introduce a potential role for the stimulus-predictive neurons in V1 introduced in this study (see section “Why are stimulus predictions also found in V1”, below).

Another possibility is that residual errors from the $P - I$ computation are “ignored” until they reach a certain threshold. This would be consistent with the fact that, as mentioned in the introduction, what we see is dominated by our expectations. It is unknown whether a psychophysical asymmetry exists: do we better perceive changes in the sensory world that are of larger magnitude than what we predict? For example, are we better at detecting accelerations of backward visual flow when we move than decelerations?

3. In the event that all feedforward activity is completely silenced, it is conceivable that this would negatively affect perception. This is consistent with studies showing that images that move with the eye, i.e. become spatially and temporally invariant, appear to fade with time (Martinez-

Conde, Macknik, and Hubel 2004). In this scenario, a lot of what we see would be prediction errors themselves.

It could thus be argued that the persistence of visually driven responses in the presence of predictions is not at odds with predictive coding. Instead, they could be a result of ways the brain could beneficially constrain prediction accuracy given limited resources and the variance of the environment.

Why does a familiar stimulus in an unexpected location not evoke a mismatch signal?

We found that omitting the presentation of an expected visual stimulus drives very strong responses in V1. However, presenting a familiar stimulus in an unexpected location does not appear to do the same (**B**-grating at position 5: **Fig. 11e,f**). The response of visually driven neurons to the unexpected stimulus scales with the strength of predictive neuron activity preceding it (**Fig. 11j**), yet effectively this stimulus is an absence of the expected stimulus – just like the stimulus omission.

A possible explanation is that the prediction for the grating at location 5 is uncertain due to its variability, and that the unexpected stimulus is predicted to some extent. This is consistent with the activity of B-predictive V1 neurons and ACC axons being lower to the B grating at position 5 than to the other B gratings, and vice versa for A-predictive neurons and axons (**Fig. 11h**). A possible contributing factor to this uncertainty is the change of sequence between conditions 1-2 and 3. Therefore, presenting the unexpected stimulus in condition 2, the first time an unexpected stimulus was presented in position 5, might be more likely to drive a mismatch response. However, this is not the case (**Fig. 17** below), possibly since stimulus predictions in condition 2 are weaker than in condition 4 (**Fig. 12b**).

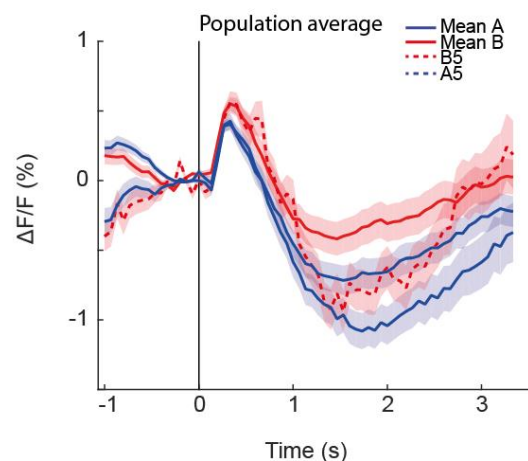


Figure 17: Mean population response (1630 neurons) to the unexpected **B5** (red dashed line), the average response to **A** (blue line), **B** (red line) and **A5** (blue dashed line). Shading indicates s.e.m. across neurons.

Another possibility is that the response to the orthogonal grating is within the uncertainty of the prediction for the expected grating. This seems unlikely, however, since stimulus predictive neurons themselves are stimulus selective.

It is possible, consistent with **Fig. 2i,j**, that the error is communicated via the visually driven response to the unexpected grating. This leads to a bit of a kerfuffle. The type of mismatch signal we observe with stimulus omission is most likely a prediction error of the form $E = P - I$: the prediction drives responses in V1 which are not silenced by expected visual input. However, if the visually driven response to the unexpected grating is a type of error signal, it would be of the form $E = I - P$. This would imply that the result of the $P - I$ subtraction is zero when the unexpected grating is presented, suggesting that both grating orientations are within the bounds of the prediction. Since this would contradict the selectivity seen in the top-down predictions (**Fig. 12**), it is possible that different levels of prediction are relayed for each error computation. Alternatively, the presence of a positive $I - P$ signal could inhibit the $P - I$ signal.

Why are stimulus predictions also found in V1?

The presence of top-down stimulus predictions and bottom-up errors we found in this study is consistent with predictive coding. However, the presence of stimulus predictions within V1 L2/3 does not lie classically within this framework. This raises the question of the significance of this signal. A similar stimulus-specific anticipatory response in V1 has been detected in human subjects using fMRI (Kok, Failing, and de Lange 2014). The hypothesis the authors posit is that the role of top-down input is to drive these neurons in V1, and their activity is used locally in order to calculate deviations from prediction. Previous work, including unpublished work from our lab, has shown that long-range projections from ACC target inhibitory interneurons in L2/3 of V1 (Zhang et al. 2014). It is possible that stimulus-predictive neurons in L2/3 of V1 are inhibitory and exert a direct influence on visually driven responses. Our estimates place nearly 98% of *Ef1α-GCamp6f* expressing neurons in L2/3 as being excitatory (unpublished data from our lab). Approximately 5% of V1 L2/3 neurons were stimulus predictive, roughly in range of this estimate. This allows for the possibility that stimulus predictive neurons directly inhibit visually driven responses. However, this makes the assumption that the activity of the two populations overlaps in time. Though this is hard to estimate with Calcium imaging, stimulus-predictive neurons peak roughly 100 ms following stimulus onset (**Fig. 11b**), which is approximately when visually driven neurons start firing. It is unclear how this type of inhibition would robustly reduce visually driven responses, unless this prediction is maintained in the cortical column through recurrent activity. A more obvious candidate for the stimulus prediction that would need to be subtracted is

therefore the visual response fed to V1 from ACC as, at least at the temporal resolution that Calcium imaging can provide, it strongly coincides with the visually driven responses in V1.

Another possibility is that stimulus predictive activity reflects an error signal of the form $E = P - I$, which is suppressed when the stimulus appears. This would suggest that the activity of stimulus predictive neurons be different during the omission than following grating presentation (**Fig. 18** below).

Unfortunately, the low number of stimulus omissions does not allow for any conclusions to be drawn.

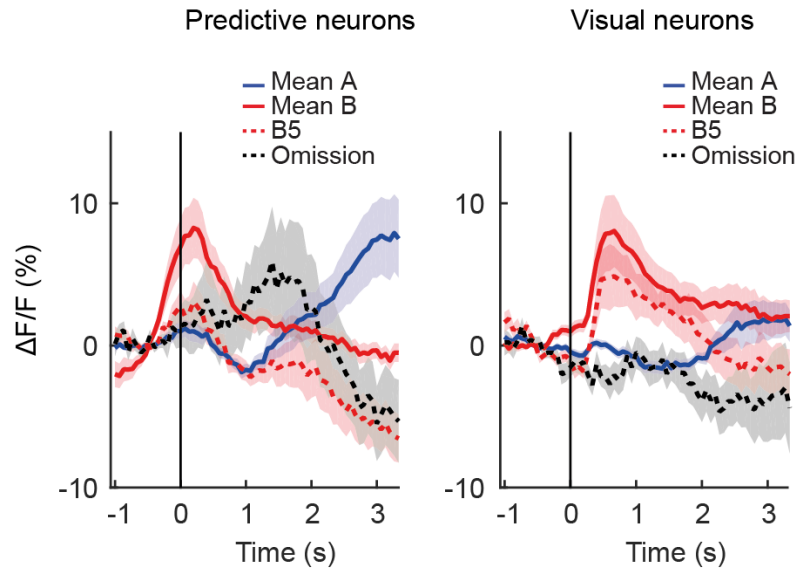


Figure 18: Activity of B-predictive (left) and B-visual (right) neurons to mean **A1** & **A3**, mean **B2** & **B4**, expected **B5** and stimulus omission.

Finally, it is interesting to speculate that stimulus-predictive neurons could be involved in mental imagery, an effect that is consistent with predictive coding (see Introduction).

Why is spatial decoding in CA1 poor in the virtual environment?

A particularly notable finding in this study is that the accuracy in decoding spatial location from neuronal activity is considerably lower in CA1 than in V1 (**Fig. 7f**), which appears to contradict previous research on spatial representations in the hippocampus (J O'Keefe and Dostrovsky 1971). A potential factor contributing to this effect is that, in order to image from CA1, we removed cortical tissue above it, primarily from secondary visual cortex (Dombeck et al. 2010). A deficit in visual input, which is important for the formation of spatial maps in the hippocampus (O'Keefe and Conway 1978), could potentially impair this process. However, activity in CA1 in our experiment was strongly locked to visual input (**Fig. 7e**), suggesting that a lack of visual input may not have been a major contributor to poor decoding

accuracy. However, it could be that the presence of ipsilateral visual cortex is somehow necessary for the stability of spatial representations in CA1.

Another culprit may be head fixation in virtual reality (Aghajan et al. 2014). However, this is not sufficient to explain the extent of the effect, since decoding accuracy was considerably higher in V1 and ACC under the same conditions (**Fig. 7c**, **Fig. 12b**).

An alternative explanation might lie in the nature of the virtual environment itself. Unlike environments typically used in hippocampal studies (Harvey et al. 2009), the corridor used in this study is highly visually repetitive: all textures repeat except for the unique “landmark” stimuli. This might selectively impair the formation of spatial maps in CA1 but not in the cortical regions examined in this study. Alternatively, it could be that spatially tied CA1 activity does not reflect spatial location *per se*, but a conjunction of spatial location and sensory input at that location, e.g. via neurons that are responsive to gratings (**Fig. 7e**). The between-condition stability of CA1 neurons responsive to gratings was similar to that reported for place cells in previous studies (approximately 15%; **Fig. 10a,c**) (Ziv et al. 2013).

Another open question is why decoding accuracy in CA1 decreased with experience (**Fig. 7f**). In addition to decoding, stimulus selectivity (**Fig. 9d**) and stimulus responses (**Fig. 13**) also decreased. A potential explanation is that, over time, the representation of the virtual environment was transferred to cortex via systems consolidation (Dudai 2004). This hypothesis is consistent with the observation that decoding accuracy increases both in V1 and ACC with experience. It can also posit a mechanism for the emergence of stimulus-predictive responses in ACC and V1. Place-like activity in CA1, that is tied to visual elements of the environment, could act as a scaffold for spatially modulated and stimulus-predictive activity in ACC, which could then be relayed to V1 via top-down input.

Conclusion

In this study, we have found evidence that primary visual cortex carries predictions of upcoming visual stimuli that are based on the animal’s location in the environment, and that omitting an expected stimulus evokes a mismatch signal. These results are consistent with a predictive coding framework, and add to recent systems neuroscience literature showing signals in sensory areas expected within this framework (Eliades and Wang 2008; Keller et al. 2012; Keller and Hahnloser 2009; Saleem et al. 2013).

The novelty of this study is twofold: first, it has described a top-down projection that may feed stimulus predictions to an early sensory area. Second, it provides evidence that internal models of space can act as a framework for forming predictions. Though this is, to our knowledge, the first time these findings have been reported in the mouse visual system, work on neural correlates of predictive coding has also

been done using fMRI imaging on humans (Alink et al. 2010; Kok et al. 2014), and is influential as a theoretical framework outside of systems neuroscience (Clark 2013). Further work on the systems neuroscience level is crucial for determining the validity of predictive coding as a fundamental functional principle of cortex, and the many subtle ways it could operate differently than theory suggests.

Methods⁵

Animals and imaging. All experiments were carried out in accordance with protocols approved by the Veterinary Office of the Canton of Basel, Switzerland. For the V1 experiments, we used imaging data from a total of 9 female C57BL/6 mice, aged 75 to 90 days at the start of the imaging series. For CA1 experiments, we used imaging data from a total of 6 female C57BL/6 mice, aged 60 to 80 days at the start of the imaging series, and for ACC experiments, a total of 5 female C57BL/6 mice, aged 71 days at the start of the imaging series. All animals were group-housed in a vivarium (light/dark cycle: 12/12 hours). The data of two V1 animals with visible z-axis motion were discarded. No statistical methods were used to pre-determine sample sizes, but our sample sizes are similar to those generally employed in the field. Mice were water restricted for the duration of the experiment, but allowed to drink water ad libitum for one hour per day in addition to receiving water rewards during experiments. Weight of all mice remained above 80% of starting weight. Viral injections and window implantation were performed as previously described (Leinweber et al. 2014). Briefly, mice were anesthetized using a mix of fentanyl (0.05 mg/kg), medetomidine (0.5 mg/kg) and midazolam (5 mg/kg) for all surgical procedures. For primary visual cortex imaging experiments, a craniotomy was made over visual cortex and AAV2/1-Ef1a-GCaMP6f-WPRE (titer 7.0×10^{10} or 5.0×10^{11} TU/ml) was injected, and the craniotomy was sealed with a 4 mm or 5 mm glass coverslip. For hippocampus imaging experiments, a 3 mm craniotomy was made above either left or right dorsal hippocampus and posterior parts of cortex were aspirated, and AAV2/1-Ef1a-NES-jRGECO1a-WPRE (Dana et al. 2016) (titer 1.2×10^{11} TU/ml) was injected into region CA1. The craniotomy was sealed with a 3 mm cover slip. For ACC axon imaging experiments, a craniotomy was made over visual cortex and sealed with a 4 mm glass coverslip. Additionally, a small craniotomy over ACC (0.3 mm lateral of bregma) was made and AAV2/1-Ef1a-GCaMP6f-WPRE (titer 1.0×10^{11} TU/ml) was injected before the region was sealed with cyanoacrylate. All viral vectors were injected at a volume of ~150 nl per site. A titanium head bar was attached to the skull and stabilized with dental cement. Imaging commenced 10 to 14 days (primary visual cortex), 23 days (hippocampus) or 30 days (ACC) following injection, and was done using a custom built two-photon microscope (Leinweber et al. 2014).

⁵ The methods described in this paper have appeared in (Fiser et al. 2016)

Illumination source was an Insight DS laser (Spectra Physics) tuned to a wavelength of 910 nm. We used a 12 kHz resonance scanner (Cambridge Technology) for line scanning, which enabled frame rates of 60 Hz at 400 x 750 pixels resolution. Imaging of ACC axons was performed using an 8 kHz resonance scanner (Cambridge Technology) resulting in frame rates of 40 Hz at 400 x 750 pixels resolution. In addition, we used a piezo actuator (Physik Instrumente) to move the objective (Nikon 16x, 0.8 NA) in steps of 15 μ m between frames to acquire images at 4 different depths, thus reducing the effective frame rate to 15 Hz for the V1 and CA1 experiments, and 10 Hz for ACC experiments. Note, the use of different setups for axonal and cell body imaging were the result of lab logistics.

Virtual reality and behavior. The behavioral imaging setup was as previously described (Leinweber et al. 2014). Briefly, head-fixed mice were free to run on an air-supported polystyrene ball, the motion of which was restricted to the forward and backward directions by a pin (**Fig. 6a**). The ball's rotation was coupled to linear displacement in the virtual environment that was projected onto a toroidal screen surrounding the mouse. The screen covered a visual field of approximately 240 degrees horizontally and 100 degrees vertically. All elements of the tunnel including the gratings were calibrated to be isoluminant. Gratings were presented on a uniform gray area once the mouse entered the gray area. A reward zone was located at the end of the tunnel. Reaching the reward zone triggered a 5 s timeout during which the mouse could lick for a water reward provided by a spout located in front of the mouse. After the timeout, the mouse was teleported to the beginning of the tunnel to start the next trial. In early traversals animals were encouraged to run by applying occasional mild air puffs to the neck.

Experimental design. Experimental sessions were one to two hours long, and each behavioral condition consisted of two such sessions, one per day, spaced on average 24 hours apart. Behavioral conditions occurred on subsequent days, with the exception of condition 5 (grating omission), which took place immediately following condition 4. For anesthetized recordings, mice were lightly anesthetized using a mix of fentanyl (0.025 mg/kg), medetomidine (0.25 mg/kg) and midazolam (2.5 mg/kg) at half surgical dose, and head fixed on the setup and passively viewed the tunnel which was presented at a constant visual flow speed. The pre-experience anesthesia condition took place one day before condition 1, and post-experience anesthesia took place either immediately after condition 5 or the following day.

Data analysis. Imaging data were full-frame registered using a custom-written software (Leinweber et al. 2014). Neurons were selected manually based on their mean fluorescence or maximum projection. This biased our selection towards active neurons. Regions of interest in the axonal data were

automatically selected by a combination of independent component analysis and image segmentation as previously described (Mukamel, Nimmerjahn, and Schnitzer 2009). Fluorescence traces were calculated as the mean pixel value in each region of interest per frame, and were then median-normalized to calculate $\Delta F/F$. $\Delta F/F$ traces were filtered as previously described (Dombeck et al. 2007). No blinding of experimental condition was performed in any of the analyses.

V1 activity becomes descriptive of spatial location: For all plots of stimulus triggered fluorescence changes, fluorescence traces were mean-subtracted in a window 10 to 5 frames (-667 ms to -333 ms for V1 and CA1; -700ms to -300ms for ACC) preceding the stimulus onset unless noted otherwise. Neurons were classified as tunnel-responsive if their mean $\Delta F/F$ in a 10-frame (667 ms) window post-stimulus onset was at least 1.5% on average for any condition, for any grating or landmark stimulus in the tunnel.

Classification of grating position using neuronal activity was done using Matlab's TreeBagger algorithm. A separate classifier was trained on each condition for each animal. 500 classification trees were used for each classifier. The input to the classifier was the mean $\Delta F/F$ of each neuron for each grating presentation in each traversal within a 10-frame window post-grating onset. Traversals were randomly interleaved, and then activity was averaged over three of the randomly chosen traversals at a time. This was done to diminish the influence of trial-to-trial variability on classification. Traversals including an unexpected grating in position 5 in conditions 2 and 4 were omitted from the classification. Classification of spatial position was done as above except that activity was binned into each fifth of the virtual tunnel. For both paradigms, the training set consisted of one half of the randomly interleaved traversals in the condition. The classifier was tested on the other half of traversals. Performance of the classifier was evaluated as the mean accuracy, i.e. the mean of the diagonal of the confusion matrix for each condition per animal. When determining the effect of traversal speed on classification, we classified traversals with a mean speed of 15 cm/s or below as slow, and 19 cm/s and above as fast. For trace deconvolution, traces were first smoothed (mean in a sliding window of 5 frames (333 ms)). The smoothed traces were then deconvolved using an exponential decay kernel with a time constant $\tau = 0.5$ s.

We considered neurons as having stable activity peaks if the peak of each neuron's average activity in each condition was within 5% of tunnel length (5 bins), i.e. one texture length, between conditions.

Selectivity indices for each neuron were calculated using the mean $\Delta F/F$ in the same window post-grating onset. Neurons were assigned a selectivity index for a condition if at least one of the responses to a grating (mean **A1** + **A3** or mean **B2** + **B4**) throughout the condition was at or above 1.5% $\Delta F/F$.

Otherwise, selectivity for the condition was set to 0. In **Fig. 9**, neurons were considered selective if their

absolute selectivity index (SI) was greater than 0.5 on average across the animals' experience. In order to measure stability of selective neurons we correlated fluorescence traces for each neuron, mapped to spatial coordinates, between traversals. We only included activity in the first 85% of the tunnel, to avoid differences due to the variable grating in position 5.

V1 develops predictive responses to upcoming visual stimuli. Predictive and visual neurons were sorted based on the position of the peak of their average $\Delta F/F$ post grating-onset over each condition. Predictive neurons were selected as exhibiting a peak within 333ms (5 frames) of grating onset, and visual neurons were selected as having a peak after 533 ms (8 frames). The window for calculating selectivity index was 133 ms pre- to 333 ms post grating onset (-2 to 5 frames) for predictive neurons, and 533 ms to 1000 ms (8 to 15 frames) post grating onset for visual neurons. In **Fig. 2b** and **c**, the baseline subtraction window for both plots was -1.33 s to -1 s (-15 to -10 frames). To compute the mean of predictive or visual neurons per traversal, the mean of each predictive neuron was taken in a 7-frame window post-grating onset (0 ms to 467 ms), and the mean of each visual neuron within a 8-frame to 15-frame (533 ms to 1s) window post-grating onset. These windows were always used to compute the mean responses for the respective class of neurons in other figures. To compare visually evoked stimulus responses in trials with high predictive activity versus those with high visual activity, we chose presentations with a predictive activity in the top 20th percentile in the 10-frame window following grating onset, and compared them to presentations with predictive activity in the bottom 20th percentile. We removed outlying traversals (with average activity values above 3σ from the mean, for predictive and visual neuron activity) before making this comparison. In Figure 2j, traversals with average predictive activity below median were classified as low-predictive traversals, and those above as high-predictive traversals.

ACC conveys stimulus predictive signals to V1. In the ACC experiments animals experienced the same tunnel for conditions 1 through 3. Imaging was only done in condition 1 and 4. To simplify comparison of **Fig. 2h** and **Fig. 3e** we use **B** to denote the stimulus the animal saw in position 5 in condition 4. To select for stimulus-predictive ACC axons in **Fig. 3d**, we selected predictive axons as in V1 but on the second half of traversals in each condition, and then plotted the responses of the selected axons in the first half of traversals. For this analysis we excluded 2 (of 5) sites with a low number of traversals in the respective conditions. In all other aspects axon selection was identical to neuron selection in the V1 data. To calculate the mean of **B**-predictive and **A**-predictive axons in **Fig. 12e**, we first mean-subtracted as described above to select neurons, and then measured mean activity after a mean-subtraction of 1 s to 500 ms preceding onset.

Omitting an expected stimulus drives strong responses in V1. Omission-selective neurons were selected as having a mean $\Delta F/F$ response larger than 10% to the omission and as having an omission selectivity index larger than 0.33. The omission selectivity index is defined as $SI_{OM} = (R_{OM} - R_{max}) / (R_{OM} + R_{max})$, where R_{OM} is the mean response to the omission for each neuron and R_{max} is the largest mean response to any other grating in the tunnel. For statistical comparison of omission-selective neuron activity in traversals with high and low predictive activity (**Fig. 13c**), data were bootstrapped 10 times with random replacement then a Wilcoxon rank-sum test was performed on the bootstrapped data.

Bibliography

- Aghajan, Zahra M. et al. 2014. "Impaired Spatial Selectivity and Intact Phase Precession in Two-Dimensional Virtual Reality." *Nature Neuroscience* 18(1):121–28. Retrieved July 10, 2016 (<http://www.nature.com/doifinder/10.1038/nn.3884>).
- Alink, Arjen, Caspar M. Schwiedrzik, Axel Kohler, Wolf Singer, and Lars Muckli. 2010. "Stimulus Predictability Reduces Responses in Primary Visual Cortex" Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus Predictability Reduces Responses in Primary Visual Cortex, 30(8), 2960–2966. doi:10.1523/JNEUROSCI.3730-. 30(8):2960–66.
- BARLOW, H. B. 1953. "Summation and Inhibition in the Frog's Retina." *The Journal of physiology* 119(1):69–88. Retrieved September 28, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/13035718>).
- Bastos, Andre M. et al. 2012. "Canonical Microcircuits for Predictive Coding." *Neuron* 76(4):695–711.
- Blatt, Gene J., Richard A. Andersen, and Gene R. Stoner. 1990. "Visual Receptive Field Organization and Cortico-Cortical Connections of the Lateral Intraparietal Area (Area LIP) in the Macaque." *The Journal of Comparative Neurology* 299(4):421–45. Retrieved September 28, 2016 (<http://doi.wiley.com/10.1002/cne.902990404>).
- Bridge, Donna J. and Ken A. Paller. 2012. "Neural Correlates of Reactivation and Retrieval-Induced Distortion." *The Journal of neuroscience : the official journal of the Society for Neuroscience* 32(35):12144–51. Retrieved September 29, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/22933797>).
- Chen, Tsai-Wen et al. 2013. "Ultrasensitive Fluorescent Proteins for Imaging Neuronal Activity." *Nature* 499(7458):295–300. Retrieved January 20, 2014 (<http://dx.doi.org/10.1038/nature12354>).
- Cho, J. and P. E. Sharp. 2001. "Head Direction, Place, and Movement Correlates for Cells in the Rat Retrosplenial Cortex." *Behavioral neuroscience* 115(1):3–25.
- Clark, Andy. 2013. "Whatever next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *The Behavioral and brain sciences* 36(3):181–204. Retrieved September 29, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/23663408>).
- Crick, Francis and Christof Koch. 1998. "Constraints on Cortical and Thalamic Projections: The No-Strong-LoopsHypothesis." *Nature* 391(6664):245–50. Retrieved September 28, 2016 (<http://www.nature.com/doifinder/10.1038/34584>).

- Cudeiro, Javier and Adam M. Sillito. 2006. "Looking Back: Corticothalamic Feedback and Early Visual Processing." *Trends in Neurosciences* 29(6):298–306.
- Dana, Hod et al. 2016. "Sensitive Red Protein Calcium Indicators for Imaging Neural Activity." *eLife* 5.
- Dombeck, Daniel A., Christopher D. Harvey, Lin Tian, Loren L. Looger, and David W. Tank. 2010. "Functional Imaging of Hippocampal Place Cells at Cellular Resolution during Virtual Navigation." *Nature Neuroscience* 13:1433–1440.
- Dombeck, Daniel a, Anton N. Khabbaz, Forrest Collman, Thomas L. Adelman, and David W. Tank. 2007. "Imaging Large-Scale Neural Activity with Cellular Resolution in Awake, Mobile Mice." *Neuron* 56(1):43–57.
- Domnisoru, Cristina, Amina a Kinkhabwala, and David W. Tank. 2013. "Membrane Potential Dynamics of Grid Cells." *Nature* 495(7440):199–204. Retrieved November 3, 2013 (<http://www.ncbi.nlm.nih.gov/pubmed/23395984>).
- Dudai, Yadin. 2004. "The Neurobiology of Consolidations, Or, How Stable Is the Engram?" *Annual Review of Psychology* 55(1):51–86. Retrieved September 27, 2016 (<http://www.annualreviews.org/doi/10.1146/annurev.psych.55.090902.142050>).
- Eliades, Steven J. and Xiaoqin Wang. 2008. "Neural Substrates of Vocalization Feedback Monitoring in Primate Auditory Cortex." *Nature* 453(7198):1102–6. Retrieved January 26, 2014 (<http://dx.doi.org/10.1038/nature06910>).
- Feierstein, Claudia E., Michael C. Quirk, Naoshige Uchida, Dara L. Sosulski, and Zachary F. Mainen. 2006. "Representation of Spatial Goals in Rat Orbitofrontal Cortex." *Neuron* 51(4):495–507.
- Felleman, D. J. and D. C. Van Essen. 1991. "Distributed Hierarchical Processing in the Primate Cerebral Cortex." *Cerebral Cortex* 1(1):1–47. Retrieved January 30, 2014 (<http://cercor.oxfordjournals.org/content/1/1/1.1.short>).
- Fiser, Aris et al. 2016. "Experience-Dependent Spatial Expectations in Mouse Visual Cortex." *Nature neuroscience*. Retrieved September 23, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/27618309>).
- Frankland, P. W. 2004. "The Involvement of the Anterior Cingulate Cortex in Remote Contextual Fear Memory." *Science* 304(5672):881–83.
- Friston, Karl. 2010. "The Free-Energy Principle: A Unified Brain Theory?" *Nature reviews. Neuroscience* 11(2):127–38. Retrieved November 1, 2013 (<http://www.ncbi.nlm.nih.gov/pubmed/20068583>).
- Fyhn, Marianne, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I. Moser. 2007.

- "Hippocampal Remapping and Grid Realignment in Entorhinal Cortex." *Nature* 446(7132):190–94. Retrieved January 21, 2014 (<http://dx.doi.org/10.1038/nature05601>).
- Gavornik, Jeffrey P. and Mark F. Bear. 2014. "Learned Spatiotemporal Sequence Recognition and Prediction in Primary Visual Cortex." *Nature neuroscience* 17(5):732–37.
- Goebel, Rainer, Darius Khorram-Sefat, Lars Muckli, Hans Hacker, and Wolf Singer. 1998. "The Constructive Nature of Vision: Direct Evidence from Functional Magnetic Resonance Imaging Studies of Apparent Motion and Motion Imagery." *European Journal of Neuroscience* 10(5):1563–73. Retrieved September 28, 2016 (<http://doi.wiley.com/10.1046/j.1460-9568.1998.00181.x>).
- Goodale, Melvyn A. and A.Davi. Milner. 1992. "Separate Visual Pathways for Perception and Action." *Trends in Neurosciences* 15(1):20–25.
- Hafting, Torkel, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. 2005. "Microstructure of a Spatial Map in the Entorhinal Cortex." *Nature* 436(7052):801–6. Retrieved July 11, 2014 (<http://dx.doi.org/10.1038/nature03721>).
- Harvey, Christopher D., Forrest Collman, Daniel A. Dombeck, and David W. Tank. 2009. "Intracellular Dynamics of Hippocampal Place Cells during Virtual Navigation." *Nature* 461(7266):941–46. Retrieved January 24, 2014 (<http://dx.doi.org/10.1038/nature08499>).
- Hinton, Geoffrey E. 2000. "Computation by Neural Networks." *Nature Neuroscience* 3(Supp):1170–1170. Retrieved September 28, 2016 (<http://www.nature.com/doifinder/10.1038/81442>).
- Holscher, C., A. Schnee, H. Dahmen, L. Setia, and HA Mallot. 2005. "Rats Are Able to Navigate in Virtual Environments." *Journal of Experimental Biology* 208(3):561.
- Huang, Yanping and Rajesh P. N. Rao. 2011. "Predictive Coding." *Wiley Interdisciplinary Reviews: Cognitive Science* 2(5):580–93. Retrieved September 28, 2016 (<http://doi.wiley.com/10.1002/wcs.142>).
- HUBEL, D. H. and T. N. WIESEL. 1962. "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex." *The Journal of physiology* 160:106–54. Retrieved January 21, 2014 (<http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1359523&tool=pmcentrez&rendertype=abstract>).
- Ji, Daoyun and Matthew a Wilson. 2007. "Coordinated Memory Replay in the Visual Cortex and Hippocampus during Sleep." *Nature neuroscience* 10(1):100–107.
- Kanwisher, Nancy, Josh McDermott, and Marvin M. Chun. 1997. "The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception." *J. Neurosci.* 17(11):4302–11. Retrieved

- January 27, 2014 (<http://www.jneurosci.org/content/17/11/4302.short>).
- Keller, Georg B., Tobias Bonhoeffer, and Mark Hübener. 2012. "Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse." *Neuron* 74(5):809–15.
- Keller, Georg B. and Richard H. R. Hahnloser. 2009. "Neural Processing of Auditory Feedback during Vocal Practice in a Songbird." *Nature* 457(7226):187–90.
- Kennedy, H., P. Barone, A. Falchier. 2000. "Relative Contributions of Feedforward and Feedback Inputs to Individual Areas." *European Journal of Neuroscience* 12.
- Kilner, James M., Æ. Karl J. Friston, and Æ. Chris D. Frith. 2007. "Predictive Coding : An Account of the Mirror Neuron System." 159–66.
- Kok, Peter, Michel F. Failing, and Floris P. de Lange. 2014. "Prior Expectations Evoke Stimulus Templates in the Primary Visual Cortex." *Journal of cognitive neuroscience* 26(7):1546–54.
- Kosslyn, Stephen M., William L. Thompson, Irene J. Klm, and Nathaniel M. Alpert. 1995. "Topographical Representations of Mental Images in Primary Visual Cortex." *Nature* 378(6556):496–98. Retrieved September 28, 2016 (<http://www.nature.com/doi/10.1038/378496a0>).
- Lee, Han, Gregory V Simpson, Nikos K. Logothetis, and Gregor Rainer. 2005. "Phase Locking of Single Neuron Activity to Theta Oscillations during Working Memory in Monkey Extrastriate Visual Cortex." *Neuron* 45(1):147–56. Retrieved July 8, 2015 (<http://www.sciencedirect.com/science/article/pii/S0896627304008414>).
- Leinweber, Marcus et al. 2014. "Two-Photon Calcium Imaging in Mice Navigating a Virtual Reality Environment." *Journal of Visualized Experiments* (84):e50885.
- Lund, J. S. 1988. "Anatomical Organization of Macaque Monkey Striate Visual Cortex." *Annual Review of Neuroscience* 11(1):253–88. Retrieved September 28, 2016 (<http://www.annualreviews.org/doi/10.1146/annurev.ne.11.030188.001345>).
- Mahn, Mathias, Matthias Prigge, Shiri Ron, Rivka Levy, and Ofer Yizhar. 2016. "Biophysical Constraints of Optogenetic Inhibition at Presynaptic Terminals." *Nature neuroscience* 19(4):554–56. Retrieved September 29, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/26950004>).
- Martinez-Conde, Susana, Stephen L. Macknik, and David H. Hubel. 2004. "The Role of Fixational Eye Movements in Visual Perception." *Nature Reviews Neuroscience* 5(3):229–40. Retrieved September 29, 2016 (<http://www.nature.com/doi/10.1038/nrn1348>).
- Maviel, Thibault, Thomas P. Durkin, Frédérique Menzaghi, and Bruno Bontempi. 2004. "Sites of

- Neocortical Reorganization Critical for Remote Spatial Memory.” *Science (New York, N.Y.)* 305(5680):96–99.
- Mendola, J. D., A. M. Dale, B. Fischl, A. K. Liu, and R. B. Tootell. 1999. “The Representation of Illusory and Real Contours in Human Cortical Visual Areas Revealed by Functional Magnetic Resonance Imaging.” *The Journal of neuroscience : the official journal of the Society for Neuroscience* 19(19):8560–72. Retrieved September 28, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/10493756>).
- Mignard, M. and JG Malpeli. 1991. “Paths of Information Flow through Visual Cortex.” *Science* 251(4998).
- Miller, M. W. and B. A. Vogt. 1984. “Direct Connections of Rat Visual Cortex with Sensory, Motor, and Association Cortices.” *The Journal of comparative neurology* 226(2):184–202. Retrieved July 10, 2015 (<http://www.ncbi.nlm.nih.gov/pubmed/6736299>).
- Moravec, Hans P. 1988. *Mind Children : The Future of Robot and Human Intelligence*. Harvard University Press.
- Mukamel, Eran A., Axel Nimmerjahn, and Mark J. Schnitzer. 2009. “Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data.” *Neuron* 63(6):747–60.
- Niell, Cristopher M. and Michael P. Stryker. 2010. “Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex.” *Neuron* 65(4):472–79.
- O’Keefe, J. and D. H. Conway. 1978. “Hippocampal Place Units in the Freely Moving Rat: Why They Fire Where They Fire.” *Experimental brain research* 31(4):573–90.
- O’Keefe, J. and J. Dostrovsky. 1971. “The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely-Moving Rat.” *Brain research* 34(1):171–75.
- O’Keefe, J. and J. Dostrovsky. 1971. “The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely-Moving Rat.” *Brain Research* 34(1):171–75.
- Poggio, Tomaso and Maximilian Riesenhuber. 1999. “Hierarchical Models of Object Recognition in Cortex.” *Nature Neuroscience* 2(11):1019–25. Retrieved September 28, 2016 (<http://www.nature.com/doifinder/10.1038/14819>).
- Polack, Pierre-Olivier, Jonathan Friedman, and Peyman Golshani. 2013. “Cellular Mechanisms of Brain State-Dependent Gain Modulation in Visual Cortex.” *Nature neuroscience* 16(9):1331–39.
- Poort, Jasper et al. 2015. “Learning Enhances Sensory and Multiple Non-Sensory Representations in Primary Visual Cortex.” *Neuron* 86(6):1478–90. Retrieved June 8, 2015 (<http://www.sciencedirect.com/science/article/pii/S0896627315004766>).

- Rao, Rajesh P. N. and Dana H. Ballard. 1999. "Predictive Coding in the Visual Cortex : A Functional Interpretation of Some Extra-Classical Receptive-Field Effects." *Nature* 79–87.
- Rolfs, Martin. 2009. "Microsaccades: Small Steps on a Long Way." *Vision Research* 49(20):2415–41.
- Saleem, Aman B., Asli Ayaz, Kathryn J. Jeffery, Kenneth D. Harris, and Matteo Carandini. 2013. "Integration of Visual Motion and Locomotion in Mouse Visual Cortex." *Nature neuroscience* 16(12):1864–69.
- Salin, P. A. and J. Bullier. 1995. "Corticocortical Connections in the Visual System: Structure and Function." *Physiological Reviews* 75(1).
- Schoenbaum, Geoffrey, Andrea A. Chiba, and Michela Gallagher. 1998. "Orbitofrontal Cortex and Basolateral Amygdala Encode Expected Outcomes during Learning." *Nature Neuroscience* 1(2):155–59. Retrieved September 30, 2016 (<http://www.nature.com/doi/10.1038/407>).
- Sherman, S. M. and R. W. Guillery. 1998. "On the Actions That One Nerve Cell Can Have on Another: Distinguishing 'drivers' from 'modulators'." *Proceedings of the National Academy of Sciences of the United States of America* 95(12):7121–26. Retrieved September 28, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/9618549>).
- Shipp, Stewart, Rick A. Adams, and Karl J. Friston. 2013. "Reflections on Agranular Architecture: Predictive Coding in the Motor Cortex." *Trends in Neurosciences* 36(12):706–16.
- Shuler, Marshall G. and Mark F. Bear. 2006. "Reward Timing in the Primary Visual Cortex." *Science (New York, N.Y.)* 311(5767):1606–9.
- Simoncelli, Eero P. and David J. Heeger. 1998. "A Model of Neuronal Responses in Visual Area MT." *Vision Research* 38(5):743–61.
- Spratling, M. W. 2012. "Predictive Coding Accounts for V1 Response Properties Recorded Using Reverse Correlation." *Biological Cybernetics* 106(1):37–49. Retrieved September 28, 2016 (<http://link.springer.com/10.1007/s00422-012-0477-7>).
- Spratling, Michael W. 2010. "Predictive Coding as a Model of Response Properties in Cortical Area V1." *The Journal of neuroscience : the official journal of the Society for Neuroscience* 30(9):3531–43. Retrieved September 28, 2016 (<http://www.ncbi.nlm.nih.gov/pubmed/20203213>).
- Sugar, Jørgen, Menno P. Witter, Niels M. van Strien, and Natalie L. M. Cappaert. 2011. "The Retrosplenial Cortex: Intrinsic Connectivity and Connections with the (Para)hippocampal Region in the Rat. An Interactive Connectome." *Frontiers in neuroinformatics* 5:7. Retrieved September 29, 2016

(<http://www.ncbi.nlm.nih.gov/pubmed/21847380>).

- Summerfield, Christopher et al. 2006. "Predictive Codes for Forthcoming Perception in the Frontal Cortex." *Science* 314(5803).
- Summerfield, Christopher and Floris P. de Lange. 2014. "Expectation in Perceptual Decision Making: Neural and Computational Mechanisms." *Nature Reviews Neuroscience* (October). Retrieved October 15, 2014 (<http://www.nature.com/doifinder/10.1038/nrn3838>).
- Teixeira, Cátia M., Stephen R. Pomedli, Hamid R. Maei, Nohjin Kee, and Paul W. Frankland. 2004. "Involvement of the Anterior Cingulate Cortex in the Expression of Remote Spatial Memory." *Science* 26(29):7555–64.
- Vogt, B. A. and M. W. Miller. 1983. "Cortical Connections between Rat Cingulate Cortex and Visual, Motor, and Postsubicular Cortices." *The Journal of comparative neurology* 216(2):192–210. Retrieved June 22, 2015 (<http://www.ncbi.nlm.nih.gov/pubmed/6863602>).
- Watanabe, Keiji Tanaka and Takeo, Charles E. Connor, Scott L. Brincat, and Anitha Pasupathy. 2007. "Transformation of Shape Information in the Ventral Pathway." *Current Opinion in Neurobiology* 17(2):140–47. Retrieved January 27, 2014 (<http://www.sciencedirect.com/science/article/pii/S0959438807000347>).
- Weible, A. P., D. C. Rowland, C. K. Monaghan, N. T. Wolfgang, and C. G. Kentros. 2012. "Neural Correlates of Long-Term Object Memory in the Mouse Anterior Cingulate Cortex." *Journal of Neuroscience* 32(16):5598–5608.
- Wickersham, Ian R. et al. 2007. "Monosynaptic Restriction of Transsynaptic Tracing from Single, Genetically Targeted Neurons." *Neuron* 53(5):639–47.
- Xu, Shengjin, Wanchen Jiang, Mu-Ming Poo, and Yang Dan. 2012. "Activity Recall in a Visual Cortical Ensemble." *Nature neuroscience* 15(3):449–55, S1-2.
- Yaksi, Emre and Rainer W. Friedrich. 2006. "Reconstruction of Firing Rate Changes across Neuronal Populations by Temporally Deconvolved Ca²⁺ Imaging." *Nature Methods* 3(5):377–83. Retrieved August 19, 2015 (<http://dx.doi.org/10.1038/nmeth874>).
- Zhang, Siyu et al. 2014. "Selective Attention. Long-Range and Local Circuits for Top-down Modulation of Visual Cortex Processing." *Science (New York, N.Y.)* 345(6197):660–65. Retrieved April 18, 2015 (<http://www.sciencemag.org/content/345/6197/660.short>).
- Ziv, Yaniv et al. 2013. "Long-Term Dynamics of CA1 Hippocampal Place Codes." *Nature neuroscience*

16(3):264–66.

Aris Nikolaos Fiser

Nationality: Greek

Email: aris.fiser@fmi.ch | **Phone:** +41 78 851 3543

Address: Schorenweg 40, 4058, Basel

Research Experience

Postdoctoral Researcher, Keller Lab

May 2018 – current

Friedrich Miescher Institute for Biomedical Research (FMI)

Project: Functional correlates of immediate early gene expression in hippocampal CA1 || Modeling schizophrenia phenotypes in terms of imbalances between sensory predictions and sensory prediction errors.

Postdoctoral Researcher, Grewe Lab

May 2017 – May 2018

Institute of Neuroinformatics, ETH Zurich

Project: Building neural networks that can learn internal models of the environment and use them to create context-dependent predictions.

PhD student, Keller Lab

Sept. 2013 – Apr. 2017

Friedrich Miescher Institute for Biomedical Research (FMI)

Project: Investigating how visual cortex generates predictions of visual input based on internal representations of space.

Master's Student, Moore Lab

Sept. 2012 – Jun. 2013

Swiss Federal Institute of Technology Lausanne (EPFL)

Project: Explored potential interactions between ATP13A2 and α -Synuclein in the pathogenesis of Parkinson's Disease in *in vivo* models.

Research Assistant, Yaksi Lab

Jun. – Aug. 2012

Neuroelectronics Research Flanders (NERF)

Project: Developed behavioral assays for and investigated neural correlates of olfactory fear conditioning in juvenile zebrafish.

Research Assistant, Visual Thinking Lab

Sept. 2010 – May 2011

Johns Hopkins University

Project: Conducted experiments on visual attention and short-term memory. Focused on Multiple Object Tracking as a study on the constraints of visual working memory.

Education

PhD, Neuroscience, University of Basel | FMI

Sept. 2013 – Apr. 2017

M.Sc., Life Sciences and Technologies, EPFL

Sept. 2011 – Jun. 2013

B.A., Neuroscience, Johns Hopkins University

Sept. 2007 – May 2011

IB, United World College of the American West

Sept. 2005 – May 2007

Publications

A. Fiser*, D. Mahringer*, H. Oyibo*, A. Petersen, M. Leinweber, G.B. Keller. “Experience-dependent spatial expectations in mouse visual cortex”. *Nature Neuroscience* (2016)

G. Daniel, A. Musso, E. Tsika, **A. Fiser**, L. Glauser, O. Pletnikova, B. Schneider & D.J. Moore. “ α -Synuclein-Induced Dopaminergic Neurodegeneration in a Rat Model of Parkinson’s Disease Occurs Independent of ATP13A2 (PARK9)”. *Neurobiology of Disease* (2015)

E. Tsika, L. Glauser, R. Moser, **A. Fiser**, G. Daniel, U-M. Sheerin, A. Lees, J.C. Troncoso, P.A. Lewis, R. Bandopadhyay, B.L. Schneider, D.J. Moore. “Parkinson’s Disease-Linked Mutations in VPS35 Induce Dopaminergic Neurodegeneration”. *Human Molecular Genetics* (2014)

Awards and Honors

PhD Fellowship, Boehringer Ingelheim Fonds

Sept. 2014 – Aug. 2016

M.Sc. Scholarship, Greek National Scholarship Foundation

Sept. 2012 – Jun. 2013

Best Poster Award, Swiss Society for Neuroscience Annual Meeting

Jan. 2017

Languages

Bilingual in English and Greek. Working competence in German and French.

