

Visual computations in the superior colliculus

Thesis by
Kyu Hyun Lee

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2020
Defended May 26, 2020

© 2020

Kyu Hyun Lee

ORCID: 0000-0001-6483-9444

All rights reserved

ACKNOWLEDGEMENTS

I want to thank Markus Meister, who taught me how to be a scientist. Looking back, I realize how lucky I was to receive his training, and I'm grateful that he took a chance on a student who knew so little about neuroscience. I will always remember the emphasis on rigorous, quantitative thinking, as well as the fun conversations we had over the years.

I also thank the committee members: Doris Tsao, whose work served as an inspiration to the project in Chapter 2; David Anderson, whose insightful questions pushed me to think harder, both in committee meetings and in class; and Ueli Rutishauser, who encouraged me and gave me much helpful advice on the projects.

I thank members of the Meister lab for their friendship, especially Zeynep Turan, Alvita Tran, and Yu-Li Ni with whom I have collaborated. I thank Hiroki Asari, who mentored me when I started in the Meister lab and taught me much about electrophysiology and data analysis. I also thank our collaborators at Janelia, without whom the project in Chapter 3 would not have been possible: Tim Harris, Marius Pachitariu, Jennifer Colonell, and Bill Karsh.

Finally, I thank my family and Hyun-Gi Yun for their love and support.

ABSTRACT

This thesis presents two projects related to large-scale extracellular recordings of neural signals. The first project asks how the brain sifts the onslaught of sensory information to identify the few bits that are relevant for guiding behavior. This question is studied in the context of the looming reaction, an innate defensive behavior against an approaching aerial predator. Interestingly, the mouse responds very selectively to the looming stimulus regardless of changes in orthogonal features, such as its position. The neural basis of this phenomenon is investigated with extracellular recordings in the superior colliculus, a midbrain visual area known to mediate the looming reaction. A detailed analysis of the difference between the superficial and deeper layers of the superior colliculus highlights a core function of visual processing: to discard information intelligently.

The second project presents electrode pooling, a novel method to increase the yield of extracellular recordings with the modern silicon electrode array. The fundamental constraint of wire volume in these devices is identified, and a solution that makes use of the switching circuitry and the sparseness of the neural signal in the time axis is described. Specifically, the method proposes to intelligently choose many recording sites that carry signal and connect them to a single wire via manipulating the switches. This pooled recording is subsequently un-mixed by a spike-sorting algorithm. The method is implemented in a state-of-the-art silicon neural probe, and its effect on signal and noise is analyzed by theory and experiment. Recommendations on the design of silicon devices are made to facilitate the incorporation of this method in the future.

PUBLISHED CONTENT AND CONTRIBUTIONS

- [1] Kyu Hyun Lee, Yu-Li Ni, and Markus Meister. Electrode pooling: How to boost the yield of switchable silicon probes for neuronal recordings. *bioRxiv*, 2019. doi: 10.1101/851691.
K.H.L. participated in the conception of the project, performed the experiments, analyzed the data, and wrote the manuscript.
- [2] Kyu Hyun Lee, Alvita Tran, Zeynep Turan, and Markus Meister. Sifting of visual information in the superior colliculus. *eLife*, 9:e50678, 2020. doi: 10.7554/eLife.50678.
K.H.L. participated in the conception of the project, performed the experiments, analyzed the data, and wrote the manuscript.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	v
List of Illustrations	vii
List of Tables	viii
Chapter I: Introduction	1
1.1 Sifting of visual information in the superior colliculus	1
1.2 Electrode pooling	5
Bibliography	7
Chapter II: Sifting of visual information in the superior colliculus	11
2.1 Abstract	11
2.2 Introduction	11
2.3 Results	13
2.4 Discussion	20
2.5 Figures	24
2.6 Methods	38
Bibliography	49
Chapter III: Electrode pooling	58
3.1 Abstract	58
3.2 Introduction	58
3.3 Results	61
3.4 Discussion	70
3.5 Figures	72
3.6 Methods	85
Bibliography	89
Chapter IV: Concluding remarks	95
4.1 Visual sifting in the SC	95
4.2 Electrode pooling	97
Bibliography	97

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.1 The emergence of selectivity, invariance, and stimulus-specific habituation along the depth of SC	25
S2.1 Histological and electrophysiological identification of SC layers	26
2.2 Selectivity to looming stimulus	27
S2.2 Looming selectivity over other figural stimuli	28
2.3 Invariance to stimulus position	29
2.4 Stimulus-specific habituation	30
S2.3 Suppression of a familiar stimulus is not permanent	31
S2.4 Enhanced response to first stimulus is not a simple consequence of motor activity	32
S2.5 A mutant mouse that lacks the neocortex and the hippocampus	33
2.5 Population decoding of distinct stimulus features	34
2.6 Model of selectivity, invariance, and stimulus-specific habituation	35
S2.6 A putative local looming detector	36
2.7 The logic of selectivity and invariance	37
3.1 Strategies for using a single wire to serve many recording sites in switchable silicon probes	73
3.2 Pooling of signal and noise	74
3.3 Maximal pool size M_{\max} as a function of the parameters α and β that characterize spike signals and noise	75
3.4 Time line of an experiment with electrode pooling	76
3.5 Overview of electrode pooling in the Neuropixels probe	77
3.6 Unmixing pooled signal	78
3.7 Characterization of signal pooling <i>in vivo</i>	79
S3.1 Measurement of pooling coefficient in saline	80
3.8 Characterization of noise pooling	81
3.9 Simulation of pooling	82
S3.2 Distribution of spike amplitudes in three brain regions from the Allen Institute Brain Observatory dataset	83
3.10 Hardware schemes for flexible connection between electrodes and wires	84

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1 Parameter values used for the model in Figure 2.6, as defined by Eqns 2.1-2.9.	48

Chapter 1

INTRODUCTION

This thesis presents two projects that are related by a widely used method in neuroscience: extracellular recording of neural signal. In this chapter, I briefly introduce both projects and provide relevant background information.

1.1 Sifting of visual information in the superior colliculus

One striking observation in sensory neuroscience is the appearance of neurons that select for a specific stimulus feature while ignoring others. They are particularly impressive given the almost infinite number of unique stimuli that the brain might face. A classic example of this phenomenon is the face cell in the inferotemporal cortex of the primate brain. These neurons respond only to images of faces, and do so regardless of their view angle, illumination, or position [10]. Like a miner who sifts through a slew of dirt in search of a few nuggets of gold, these neurons discard all but the few bits of information important for guiding behavior. Neuroscientists have marveled at this phenomenon for a long time, and some have even proposed that our conscious perceptual experience is nothing more than the firing of these neurons [1].

Despite the enduring interest in this topic, we still know very little about how this computation is implemented by biological neurons. The goal of the first project is to open up a new approach to solving this problem by investigating it in an area that has often been neglected by sensory neuroscientists: the superior colliculus. Chapter 2 will discuss the angle of attack we have chosen, the main results of the experiments, and the major lessons learned. The remaining parts of this section will provide further background information about selectivity and invariance in sensory processing.

Losing information intelligently

Before pressing on, it is worth clarifying what makes selectivity and invariance interesting phenomena. After all, one could argue that these apply trivially to any neuron. Just consider the way a neuron generates an output: it integrates dendritic inputs and fires action potentials when they cross the spiking threshold. So at an abstract level, any neuron is selective to stimuli that cause a suprathreshold combi-

nation of inputs and invariant to every other stimuli. Even the lowly photoreceptor is selective to a change in light intensity in its narrow receptive field and invariant to everything else that happens outside of it [17].

Furthermore, the fact that neurons farther removed from the sensory periphery convey less information about the stimulus is completely expected from the data processing inequality theorem [7]: given a Markov chain $X \rightarrow Y \rightarrow Z$, $I(X;Y) \geq I(X;Z)$, i.e. the information about X can only decrease as one goes down the processing chain. Another way to see this is by the following thought experiment: imagine that you give a photo of the Beckman Lawn to Alice to draw it; then Alice gives her drawing to Bob to produce his own drawing, and Bob gives his drawing to Charlie, and so on. We can expect that by the time we compare Charlie's drawing to the original photo, much of the pixel-level information about Beckman Lawn would have been lost, even if some of the participants are expert artists. So what is interesting about the fact that the neurons of the inferotemporal cortex have discarded most of the information present in the photoreceptor layer?

The answer, of course, lies in the nature of the few bits of information that are retained. What's surprising is not that the face cell only responds to a narrow class of stimuli and ignores others; it is that the stimuli it cares about happen to be *faces*, which are obviously important for the monkey's survival. Remarkably, evolution has generated a nervous system that loses information intelligently. So truly appreciating selectivity and invariance requires understanding the ecological problems the animal must solve. This focus on the reduction of information differs sharply from previous work that seeks to decode variables of interest from the sensory system (e.g. [6, 25]). While sensory neurons may carry such information, these studies neglect the fact that the ultimate goal of sensory processing is to inform behavior, which takes place at a much lower information rate than the input.

Selectivity and invariance: What is known?

This section will briefly review previous studies on neural computations of selectivity and invariance. The literature is truly vast and covers many sensory modalities. Here the focus is to highlight a few illuminating examples and the computational strategies they uncover.

As mentioned before, one area in which the issue of selectivity and invariance has been studied is the primate inferotemporal cortex. In addition to face cells, this region harbors neurons that are tuned to the identity of specific objects [8]. How

does such a response property arise? Because the ventral stream of the primate visual pathway consists of many cortical areas, experimental elucidation of how the input is transformed at each successive stage has been difficult. One recent approach that has received much attention is to train deep convolutional neural networks on object recognition tasks and “look under the hood” to see if the intermediate layers of the model can explain the neural activity found in intermediate visual cortical areas [18, 22, 27]. These models indeed are more successful than previous ones [27], but what does it tell us about the mechanism for selectivity and invariance? The convolution operation in these models can be interpreted in two ways. If the previous layer consists of elementary features, then taking their convolution can be viewed as implementing a coincidence detector (an AND gate) that becomes active only when multiple such features are present. This confers selectivity. If the previous layer, on the other hand, consists already of complex feature detectors, then taking their convolution is akin to pooling over them across space (an OR gate). This confers local invariance to position. A neuron can switch between these two functions by controlling its nonlinearity: generally an AND operation would require a higher threshold than an OR operation. By repeatedly stacking these AND and OR operations [8, 17], these models are able to achieve a high degree of selectivity at the output that is tolerant to changes in orthogonal features.

This basic idea is consistent with examples where the circuit is better understood experimentally, such as the jamming avoidance response of the weakly electric fish. These animals use an electric organ to discharge electric field and sense its distortion with electroreceptors on their skin to identify nearby objects. This requires having a good internal model of what is expected from its own electric discharge [11]. When a conspecific is nearby, however, the interference from the jamming signal (the electric discharge of the other fish) disrupts this expectation, and the fish modulates the frequency of its own discharge to make it more different; e.g. when the jamming signal is higher in frequency, the fish lowers the frequency of its own discharge. This requires detecting the difference in the frequency of the electric discharge regardless of the orientation of the jamming signal. The fish seems to solve this problem by first combining the patterns of change in amplitude and phase locally at each sensory location, and then pooling them across space [20]: another instance of an AND-OR cascade. A more recent study has reported that the fish may use a similar strategy for other invariant representations, such as the phase of a communication signal [14].

Is there another way to build selectivity and invariance? One notable proposal

comes from [16], in which an attention-like mechanism is used to dynamically route the input to a central pattern detector. This is done by modulating the synaptic weights of the network at different stages with a set of control neurons, whose linear combination produces the desired pattern of connectivity. By adjusting the position and the size of the attentional window, this model achieves invariance to position and scale. It also completely separates invariance from selectivity: the few neurons with high selectivity are at the top of the network, and the inputs are switched by synaptic modulations. In general, distinguishing this model from the one previously discussed may be difficult. The work presented in Chapter 2, however, will be able to do so, thanks to an additional feature.

Superior colliculus and looming reaction

Here I provide a brief background on the looming reaction and the role of the superior colliculus in controlling it, which motivates the project described in Chapter 2. Although the looming reaction is widely observed across many species, including invertebrates [5], I focus on what is known about it in the mouse.

The looming reaction is a defensive behavior in response to an approaching aerial predator. It is observed even in the laboratory mouse that has never faced threats of predation [28], which suggests that the behavior is mediated by a stereotyped and dedicated circuit. Although the actual hunting patterns of aerial predators (e.g. owls) can vary [24], the looming reaction is reliably triggered by a dark expanding disk presented above the mouse with an LCD monitor. The behavioral output is twofold: the animal either rapidly flees to the nest (when it is present) with a latency of a few tenths of a second, or freezes in place for a prolonged period (> 10 s) [28]. An intriguing feature of this behavior is its specificity to the black expanding disk. Other similar stimuli, such as the white receding disk that shares moving dark edges, are much less effective at driving the behavior [28]. Furthermore, behavioral response is similar regardless of the position of the looming stimulus in the upper visual field. This combination of selectivity and invariance is reminiscent of the face cell discussed earlier.

The neural substrate most closely associated with this behavior is the superior colliculus (SC), a midbrain area that is conserved across vertebrates. It consists of anatomically distinct layers, with the superficial layers receiving input from retinal ganglion cells, and projecting to the deeper layers that integrate other sensory modalities to issue behavioral commands [2]. Most previous studies of the SC

are devoted to its role in controlling eye movement [12], although recent work has expanded the scope to attention, a form of mental orienting [13].

Electrical stimulations of the SC [21], as well as optogenetic stimulations of certain collicular cell types [23, 26], lead to an escape or freezing response, implying that the SC may control this behavior. One recent report has shown that the escape probability varies with the contrast of the looming stimulus [9]. Based on this observation, the authors modeled the behavior as a form of evidence accumulation: the stimulus is scaled by its contrast and drives an increase in an internal variable (“threat”) represented by neurons in the deeper SC. When this rises above a threshold, a motor output is triggered by a downstream brain area (PAG). Though intriguing, this proposal is unconvincing for several reasons. First, it seems inconsistent with the short latency of the behavior. Animals begin flight within 250 ms of stimulus onset, when the looming disk is only about 10° wide. So in the normal mode of operation, there is just not enough time to accumulate threat, and the reduced escape probability to low contrast stimuli could simply reflect a reduced sensory drive that requires a greater number of repetitions to be detected. Furthermore, the differential equation governing the time evolution of threat incorporates the stimulus simply as its diameter multiplied by the contrast. By this logic, a large flashing dark disk would be more effective than the looming stimulus, as it immediately reaches maximal diameter at maximal contrast. But such a stimulus rarely induces flight. In any case, none of these studies have addressed the central mystery of the looming reaction: how does the animal selectively respond to the looming stimulus while ignoring others that share similar features, across a large swath of its visual field? Chapter 2 delves into this question.

1.2 Electrode pooling

The second project develops electrode pooling, a novel method to boost the yield of extracellular neural recordings. A detailed introduction to the project can be found in Chapter 3. Here I provide a few remarks on extracellular electrophysiology that sets the stage for a more detailed description of the work.

What are we measuring?

Recording extracellular signal in the brain is an old technique, but one that has received a lot of attention, as it enables the recording of many neurons at a high spatiotemporal resolution. But what exactly does it measure? The short answer is the electric potential created by transmembrane currents through ion channels [4],

the most important of which are action potentials. Consider a situation where a recording electrode is placed near a neuron's cell body, and a reference electrode is placed farther away, closer to the dendrites. Shortly after the neuron's membrane potential passes a spiking threshold, sodium current rushes into the cell body. This causes the intracellular membrane potential (compared to the reference electrode outside the cell) to become more positive, which generates a capacitive current that moves from the dendrites back to the cell body to balance the sodium influx. This implies that the voltage at the extracellular electrode is more negative (again, compared to the reference electrode by the dendrites), and we see this as a downward deflection. During the hyperpolarization phase, this cycle runs in reverse: potassium current rushes out of the cell and is offset by the capacitive current returning to the dendrites, making the intracellular potential more negative and the extracellular potential more positive (we see this as an upward deflection back to baseline).

What factors influence our measurement?

Extracellular recording is affected by a number of factors, many of which cannot be controlled by the experimenter. The amplitude of the signal is influenced by variables such as the size of the neuron and the distance between the neuron and the electrode. Typically the larger the neuron and closer the electrode, the larger the signal ¹. This is because a larger neuron has a larger surface area and a greater number of ion channels, which together conduct a larger current. Placing the electrode closer to the neuron also helps, because most of the voltage drop occurs right at the cell body where the capacitive currents from the dendrites accumulate. The shape of the signal (i.e. the waveform of the action potential) is affected by the kinds of ion channels expressed by the cell, which determines the dynamics of transmembrane currents [3], as well as the geometry of the dendritic and axonal processes, which determines how the transmembrane currents sum in the extracellular space. The placement of the electrode is also important. Axonal spikes, for example, tend to be triphasic because the electrode detects the outward capacitive current as the axon segment charges up before the downward deflection driven by the sodium current. The shape of the signal is important for spike-sorting, a pre-processing step that assigns spikes detected by the electrode to distinct single units.

What about the noise? One source of noise is the electrode impedance. When a metal

¹This also points to a limitation of extracellular recording: small neurons such as the cerebellar granule cells are largely invisible to the extracellular electrode because they do not generate a sufficiently strong extracellular field.

electrode is placed in a conductive medium, a thin layer of positive charge forms at the interface to balance the negatively charged electrode. This “double-layer” at the metal-liquid junction behaves like a RC circuit [15, 19]. Along with the resistance due to the extracellular solution between the recording and reference electrodes, it generates thermal noise. Another source of noise is the neurons whose signals are too far from the electrode to be detected as action potentials. Their superposition generates “biological noise,” which is often the largest source of noise. Finally, there is noise from the recording electronics that are downstream of the electrode, such as the amplifier, digitizer, and multiplexer. Although this electronic noise is typically negligible compared to the other sources, it can be significant in some recording systems, as we will see in Chapter 3.

References

- [1] Horace B. Barlow. Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, 1(4):371–394, 1972. doi: 10.1068/p010371. URL <https://doi.org/10.1068/p010371>. PMID: 4377168.
- [2] Michele A. Basso and Paul J. May. Circuits for Action and Cognition: A View from the Superior Colliculus. *Annual Review of Vision Science*, 3(1): annurev-vision-102016-061234, October 2017. ISSN 2374-4642. doi: 10.1146/annurev-vision-102016-061234. URL <http://www.annualreviews.org/doi/10.1146/annurev-vision-102016-061234>.
- [3] Bruce P. Bean. The action potential in mammalian central neurons. *Nature Reviews Neuroscience*, 8(6):451–465, Jun 2007. ISSN 1471-0048. doi: 10.1038/nrn2148. URL <https://doi.org/10.1038/nrn2148>.
- [4] György Buzsáki, Costas A Anastassiou, and Christof Koch. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420, 2012. doi: http://www.nature.com/nrn/journal/v13/n6/suppinfo/nrn3241_S1.html. URL <http://dx.doi.org/10.1038/nrn3241>.
- [5] Gwyneth Card and Michael H. Dickinson. Visually mediated motor planning in the escape response of drosophila. *Current Biology*, 18(17):1300 – 1307, 2008. ISSN 0960-9822. doi: <https://doi.org/10.1016/j.cub.2008.07.094>. URL <http://www.sciencedirect.com/science/article/pii/S0960982208010488>.
- [6] Yuzhi Chen, Wilson S. Geisler, and Eyal Seidemann. Optimal temporal decoding of neural population responses in a reaction-time visual detection task. *Journal of Neurophysiology*, 99(3):1366–1379, 2008. doi: 10.1152/jn.

- 00698.2007. URL <https://doi.org/10.1152/jn.00698.2007>. PMID: 18199810.
- [7] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006. ISBN 0471241954.
- [8] James J DiCarlo, Davide Zoccolan, and Nicole C Rust. How does the brain solve visual object recognition? *Neuron*, 73(3):415–434, 2012. ISSN 1097-4199. doi: 10.1016/j.neuron.2012.01.010. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=22325196<http://www.ncbi.nlm.nih.gov/pubmed/22325196><http://www.ncbi.nlm.nih.gov/pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3306444>.
- [9] Dominic A. Evans, A. Vanessa Stempel, Ruben Vale, Sabine Rühle, Yaara Lefler, and Tiago Branco. A synaptic threshold mechanism for computing escape decisions. *Nature*, 558(7711):590–594, June 2018. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-018-0244-6. URL <http://www.nature.com/articles/s41586-018-0244-6>.
- [10] W A Freiwald and D Y Tsao. Functional Compartmentalization and Viewpoint Generalization Within the Macaque Face-Processing System. *Science*, 330(November):845–851, 2010. ISSN 1095-9203. doi: 10.1126/science.1206034. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=21051642.
- [11] Ann Kennedy, Greg Wayne, Patrick Kaifosh, Karina Alvina, L. F. Abbott, and Nathaniel B Sawtell. A temporal basis for predicting the sensory consequences of motor commands in an electric fish. *Nature Neuroscience*, 17(3):416–422, 2014. doi: 10.1038/nn.3650<http://www.nature.com/neuro/journal/v17/n3/abs/nn.3650.html#supplementary-information>. URL <http://dx.doi.org/10.1038/nn.3650>.
- [12] Eileen Kowler. Eye movements: The past 25 years. *Vision Research*, 51(13): 1457–1483, 2011. ISSN 0042-6989. doi: 10.1016/j.visres.2010.12.014.
- [13] Richard J Krauzlis, Lee P Lovejoy, and Alexandre Zénon. Superior Colliculus and Visual Spatial Attention. *Annual Review of Neuroscience*, 36(1):165–182, 2013. doi: doi:10.1146/annurev-neuro-062012-170249. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-neuro-062012-170249>.
- [14] Michael G Metzen, Volker Hofmann, and Maurice J Chacron. Neural correlations enable invariant coding and perception of natural stimuli in weakly electric fish. *eLife*, 5:e12993, 2016. doi: 10.7554/eLife.12993.

- [15] Joana P. Neto, Pedro Baião, Gonçalo Lopes, João Frazão, Joana Nogueira, Elvira Fortunato, Pedro Barquinha, and Adam R. Kampff. Does impedance matter when recording spikes with polytrodes? *Frontiers in Neuroscience*, 12: 715, 2018. ISSN 1662-453X. doi: 10.3389/fnins.2018.00715. URL <https://www.frontiersin.org/article/10.3389/fnins.2018.00715>.
- [16] B A Olshausen, C H Anderson, and D C Van Essen. A neurobiological model of visual-attention and invariant pattern-recognition based on dynamic routing of information. *Journal of Neuroscience*, 13(11):4700–4719, 1993. ISSN 0270-6474.
- [17] Xaq Pitkow and Markus Meister. Neural computation in sensory systems. In Michael S. Gazzaniga and George R. Mangun, editors, *The Cognitive Neurosciences V*, chapter 28. MIT Press, 2014.
- [18] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11):1019–25, 1999. ISSN 1097-6256. doi: 10.1038/14819. URL <http://www.ncbi.nlm.nih.gov/pubmed/10526343>.
- [19] D. A. Robinson. The electrical properties of metal microelectrodes. *Proceedings of the IEEE*, 56(6):1065–1071, 1968.
- [20] G. J. Rose, M. Kawasaki, and W. Heiligenberg. ‘recognition units’ at the top of a neuronal hierarchy? *Journal of Comparative Physiology A*, 162(6): 759–772, Nov 1988. ISSN 1432-1351. doi: 10.1007/BF00610965. URL <https://doi.org/10.1007/BF00610965>.
- [21] N Sahibzada, P Dean, and P Redgrave. Movements resembling orientation or avoidance elicited by electrical stimulation of the superior colliculus in rats. *Journal of Neuroscience*, 6(3):723–733, 1986. ISSN 0270-6474.
- [22] T. Serre, A. Oliva, and T. Poggio. A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci USA*, 104(15): 6424–6429, 2007. ISSN 0027-8424. doi: 10.1073/pnas.0700622104. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=17404214http://www.pnas.org/cgi/doi/10.1073/pnas.0700622104.
- [23] Congping Shang, Zhihui Liu, Zijun Chen, Yingchao Shi, Qian Wang, S. Su Liu, Dapeng Li, and Peng Cao. A parvalbumin-positive excitatory visual pathway to trigger fear responses in mice. *Science*, 348(6242):1472–1477, 2015. ISSN 0036-8075. doi: 10.1126/science.aaa8694. URL <http://www.sciencemag.org/cgi/doi/10.1126/science.aaa8694>.
- [24] Eran Shiffman and David Eilam. Movement and direction of movement of a simulated prey affect the success rate in barn owl tyto alba attack. *Journal of Avian Biology*, 35(2):111–116, 2004. doi: 10.1111/j.0908-8857.2004.03257.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.0908-8857.2004.03257.x>.

- [25] Carsen Stringer, Michalis Michaelos, and Marius Pachitariu. High precision coding in visual cortex. *bioRxiv*, 2019. doi: 10.1101/679324. URL <https://www.biorxiv.org/content/early/2019/11/04/679324>.
- [26] Pengfei Wei, Nan Liu, Zhijian Zhang, Xuemei Liu, Yongqiang Tang, Xiaobin He, Bifeng Wu, Zheng Zhou, Yaohan Liu, Juan Li, Yi Zhang, Xuanyi Zhou, Lin Xu, Lin Chen, Guoqiang Bi, Xintian Hu, Fuqiang Xu, and Liping Wang. Processing of visually evoked innate fear by a non-canonical thalamic pathway. *Nature Communications*, 6:6756, 2015. doi: 10.1038/ncomms7756. URL <http://dx.doi.org/10.1038/ncomms7756>.
- [27] D. L. K. Yamins, H. Hong, C. F. Cadieu, E. A. Solomon, D. Seibert, and J. J. DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014. ISSN 0027-8424. doi: 10.1073/pnas.1403112111. URL <http://www.pnas.org/cgi/doi/10.1073/pnas.1403112111>.
- [28] Melis Yilmaz and Markus Meister. Rapid innate defensive responses of mice to looming visual stimuli. *Current Biology*, 23(20):2011–2015, 2013. ISSN 1879-0445 (Electronic) 0960-9822 (Linking). doi: 10.1016/j.cub.2013.08.015. URL <http://www.ncbi.nlm.nih.gov/pubmed/24120636>.

*Chapter 2***SIFTING OF VISUAL INFORMATION IN THE SUPERIOR COLLICULUS**

Kyu Hyun Lee, Alvita Tran, Zeynep Turan, and Markus Meister. Sifting of visual information in the superior colliculus. *eLife*, 9:e50678, 2020. doi: 10.7554/eLife.50678.

K.H.L. participated in the conception of the project, performed the experiments, analyzed the data, and wrote the manuscript.

2.1 Abstract

Much of the early visual system is devoted to sifting the visual scene for the few bits of behaviorally relevant information. In the visual cortex of mammals, a hierarchical system of brain areas leads eventually to the selective encoding of important features, like faces and objects. Here we report that a similar process occurs in the other major visual pathway, the superior colliculus. We investigate the visual response properties of collicular neurons in the awake mouse with large-scale electrophysiology. Compared to the superficial collicular layers, neuronal responses in the deeper layers become more selective for behaviorally relevant stimuli; more invariant to location of stimuli in the visual field; and more suppressed by repeated occurrence of a stimulus in the same location. The memory of familiar stimuli persists in complete absence of the visual cortex. Models of these neural computations lead to specific predictions for neural circuitry in the superior colliculus.

2.2 Introduction

Even though the human eye takes in about one gigabit of raw visual information every second, we end up using only a few tens of bits to guide our behavior [52]. Of course, those bits are carefully selected from the scene, and which specific bits get used depends entirely on the context and goals. All this happens in a processing time of about a tenth of a second [65, 68]. How the visual brain sifts the onslaught of visual data for the few behaviorally relevant nuggets has been an enduring mystery. Much research in this area has focused on the primate visual system, and specifically the phenomena of invariant object recognition. For example, certain neurons in the inferotemporal cortex respond selectively to a specific individual's face regardless

of its position or view angle [18], or to the concept of a specific celebrity regardless of how that concept arises [53]. An impressive body of theory and computational modeling has been developed to explain how this sifting for important bits from the retinal output may be implemented [12, 63]. However, empirical evidence regarding the actual biological microcircuits has been difficult to obtain.

In rodent vision, a prominent example of visual sifting is the defensive reaction of a mouse to an approaching aerial predator [10, 75]. Freezing or escape can be triggered reliably by an overhead display of an expanding dark disk. Effectively, the alarm circuits in the mouse's visual system extract from the overall visual display just one or two bits of information needed to initiate action. To function properly, such an alarm system must be highly selective for the trigger feature. Indeed, the mouse does not respond to expanding white disks, or to dimming dark disks, or to contracting white disks [75]. All these innocuous stimuli share some low-level features with the expanding dark disk, but not the overall configuration. Furthermore, the behavior is invariant to irrelevant features. For example a mouse will freeze in response to looming stimuli presented anywhere in the upper visual field. It is unknown how this invariance to location arises, and how it can coexist with high selectivity for the local stimulus features.

Recent research on rodents suggests that the visual drive for these defensive behaviors arises not in the thalamo-cortical pathway but in the superior colliculus [16, 64]. The superior colliculus (SC) is an evolutionarily ancient midbrain structure that mammals share with birds, fish, and amphibians [3, 6]. The superficial layers receive inputs from the retina and in mammals also from the visual cortex, organized in a precise retinotopic map [62]. Neurons there project to the deep layers of the SC as well as other brain areas including the lateral geniculate nucleus and pulvinar. The deep layers also receive signals from other sensory modalities including hearing and touch. Neurons in the deep SC represent pre-motor signals and project broadly to many brain areas in both ascending and descending pathways. Generally speaking, neural processing in the SC identifies salient points in the environment and coordinates the orienting of the animal towards or away from such locations. In the primate brain, this has been studied extensively for the special case of eye movements [34], but the primate SC also helps control head, arm, and body movements. Furthermore, the SC contributes to a type of "internal" orienting, namely when we direct our attention to a specific part of the scene without overt eye movements [35].

To better understand how visual sifting proceeds in the SC, we recorded spike trains simultaneously from neurons throughout all layers of this structure in the awake mouse. The set of visual displays included visual threats that reliably elicit defensive reactions, and closely related stimuli that do not. We report on three kinds of neural computations that separate behaviorally relevant from irrelevant stimuli, and we trace their emergence from the superficial to the deep layers of the SC: (1) an increasing selectivity for the threat stimulus; (2) an increasing invariance to location of that stimulus; and (3) the suppression of neural responses to a familiar stimulus. In particular, this memory of familiar stimuli is stimulus-specific, lasts for a behaviorally relevant timescale, and does not require input from the visual cortex. To explain these computations, we consider several circuit models, some of which can be eliminated based on the population recordings. These results suggest how circuits of the SC can effectively distill the ecologically relevant information that guides behavior.

2.3 Results

Emergence of new response properties from superficial to deep layers

To track visual computations in the mouse SC, we recorded from hundreds of neurons simultaneously in all layers of the structure using multi-electrode silicon prongs [14]. The animal was head-fixed, awake, and moving on a running wheel, but not trained to perform any specific task, so we could best observe the autonomous visual functions of the SC. The recording electrodes were aimed at the dorso-medial portion of the SC, which processes stimuli in the upper visual field. Over the course of several hours we presented a battery of visual displays, ranging from abstract stimuli like flickering checkerboards to those with ecological significance, like overhead looming disks.

In analyzing neuronal responses to these stimuli we observed a systematic progression from the superficial layers that receive retinal input to the deep layers of the SC. To illustrate the dramatic change in how stimuli are represented, Figure 2.1 compares recordings from two sample neurons, one in the superficial SC and the other in the deep SC.

The superficial neuron responded well to many different kinds of displays, such as an expanding dark disk (the classic “looming” stimulus), a contracting white disk, a moving disk, or a dimming disk. By contrast, the deep neuron was quite selective for the looming stimulus (Figure 2.1C). Second, the superficial neuron had a small

and precisely circumscribed receptive field roughly 10° in diameter. It responded only when the stimulus invaded that region. By contrast the deep neuron responded strongly to any looming stimulus presented over a wide region that encompasses much of the visual hemifield (Figure 2.1D). Third, the superficial neuron responded reliably to repeated presentation of the identical stimulus. By contrast the deep neuron fired only on the first presentation and failed to respond entirely to the subsequent ones (Figure 2.1C, top row).

The three characteristics found in the deep SC neuron's responses – selectivity for the looming feature, spatial invariance, and habituation to familiar stimuli – are all distinct from the signals transmitted by the retina. For example, an "approach-sensitive" retinal ganglion cell (RGC) has been reported in the mouse retina [46], but later studies have found that it is actually the Off-transient alpha cell [56] that responds to many other Off-type stimuli in addition to the looming stimulus [36]. RGCs also have local receptive fields ranging up to 10° at most [36], which can be readily mapped with white noise stimuli such as flickering checkerboards or bars [76]. Finally, although RGCs show complex adaptation properties, the timescale of adaptation is typically on the order of 0.1–10 s [2, 72], whereas the habituation we find in the deep SC lasts on the order of minutes. In the following sections, we elaborate on these response properties and how they may arise in the circuitry of the SC.

Selectivity for looming stimuli

In an attempt to measure the visual receptive fields of all the recorded neurons, we applied a flickering checkerboard stimulus and then computed the spike-triggered average (STA) stimulus [8]. This is a common procedure that works well for retinal ganglion cells and neurons in the early stages of visual cortex [45, 47]. In the superficial SC, the STA analysis yielded linear receptive fields that resembled those of retinal ganglion cells (Figure 2.2A-B). They were sharply defined in space, with the smallest only $\sim 5^\circ$ across. They frequently showed an antagonistic and delayed surround, and some displayed orientation- and direction-selectivity [17, 28]. The great majority of these neurons ($\sim 90\%$) were Off cells based on the shape of the STA. By contrast, neurons in the deep SC did not produce sustained responses to the flickering checkerboard (Figure 2.2A), and thus contained no structure in the STA (Figure 2.2B). Nevertheless these same deep SC neurons did respond strongly to certain figural stimuli, like the expanding dark disk (Figure 2.2A, C-D).

Among the various figural stimuli we tested, many neurons showed some selective tuning (Figure 2.1C, Figure 2.2D, Supplementary Figure S2.2). We focus here on the comparison of an expanding dark disk with a contracting white disk (Figure 2.2D). These two stimuli are closely related in terms of local features: both contain an advancing dark edge. But the ecological interpretations are quite different: one indicates an approaching dark object and the other a receding white object. Freely moving mice take an evasive action to an expanding dark disk, but are unimpressed by a contracting white disk [75]. Compared to superficial SC, neurons in the deep SC indeed became more selective for the expanding dark disk (Figure 2.2D). This can be seen as sifting what is likely the most behaviorally relevant signal in the upper visual field from other distracting stimuli.

Invariance to stimulus position

Whereas superficial SC neurons often had sharp receptive fields just 5–10° in diameter, deep SC neurons generally responded to stimuli over a large part of the visual field. We probed this tendency with expanding dark disks presented at many different locations, as these were the most effective stimuli in the deep SC. With increasing depth in the SC, neurons showed larger receptive fields, growing by a factor of 6 in area or more (Figure 2.3A-B). Note that the resolution of the receptive field measurement with expanding dark disks is ~15°, and as a result these receptive fields are larger than those measured by the flickering checkerboard (Figure 2.2B).

Despite this wide spatial range, deep SC neurons responded with a remarkably short latency to looming stimuli at any location (Figure 2.3A). By the time such a neuron starts firing, the expanding dark disk has only covered a few retinal ganglion cells. In contrast, for superficial neurons the latency varied depending on the location of the expanding disk stimulus and it often exceeded the latency of deep SC neurons. (Figure 2.3A). Figure 2.3C plots this variation in the latencies across the SC depth. One possible interpretation is that a widefield neuron in the deep SC pools over many local neurons in the superficial SC, such that it becomes sensitive with the same latency at every point in its receptive field. Indeed, such an interlaminar pathway has been demonstrated previously in slice preparations [23, 38]. We consider this possibility more thoroughly below.

In any case, it appears that certain widefield neurons in the deep SC have solved the problem of threat detection to a large degree: they signal the looming stimulus rapidly and sensitively without false alarms from stimuli that share some low-level

features but not the behavioral significance.

Habituation to familiar stimuli

Neurons in the superficial layers generally produced a spike burst of comparable firing rate with every repeat of the stimulus (Figure 2.1C). By contrast, some neurons in the deep layers responded with a sharp burst only to the first presentation; the response to all subsequent repeats was suppressed (Figure 2.4A). The degree of habituation to repeated stimuli was greater in the deeper SC compared to the superficial SC (Figure 2.4B).

The onset of this habituation is immediate and already affects the response ~ 1 s later (Figure 2.1C, Figure 2.4A). The suppression then lasts for minutes: many deep SC neurons showed less than 50% recovery even after ~ 120 s (Figure 2.4D). While we have not measured the exact time course of recovery, we found that the suppression was not permanent. In general, neurons recovered the full sensitivity to the first presentation when probed again about an hour later (Supplementary Figure S2.3). Furthermore, the burst of spikes was not driven simply by a change in locomotor output or pupil size as a secondary consequence of the visual threat (Supplementary Figure S2.4).

Remarkably, this habituation was strictly specific to the stimulus that caused the response. As reported above, widefield neurons in the deep SC can be triggered by looming disks at many different locations (Figure 2.1C, Figure 2.3A). Figure 2.4C shows the response of a single neuron to a looming stimulus whose location was chosen randomly on every trial. By comparing the sequence of responses at one location to that at another one can test whether the habituation transfers across space. As shown in the bottom left panel of Figure 2.4C, a stimulus at one location did not suppress the subsequent response of the same neuron to a stimulus at another location, even separated by as little as 15° . One interpretation is that the habituation takes place in local circuits spanning $\sim 15^\circ$ in width before their output gets pooled by the widefield neuron.

Given that the memory for familiar stimuli can last two minutes or longer, we considered whether the hippocampus or the neocortex play a role in storing this information, perhaps by modulating the gain of collicular signals through the extensive projections from visual cortex [77]. Thus we repeated the experiments in a mutant mouse that lacks all of the dorsal forebrain, including the hippocampus and most of the neocortex [31] (Supplementary Figure S2.5). Intriguingly, the mu-

tant also showed long-lasting suppression of repeated stimuli in deep neurons of the SC (Figure 2.4E), to a degree that matched the suppression seen in the normal mouse (Figure 2.4C bottom right and Figure 2.4D). This is consistent with a local mechanism for habituation within the SC.

The preceding analyses of single-neuron responses suggest that the neural population deep in the SC selectively represents those bits of information that may be of immediate relevance to defensive reactions, while other aspects of the visual display get discarded. To test this directly, we applied a linear decoder to the population vector from neurons in superficial and deep SC. From single stimulus trials, the decoder easily read out the precise location of a visual stimulus from the population in superficial SC, but much less so from neurons in deep SC (Figure 2.5, left). By contrast, the deep SC represented explicitly whether a stimulus appeared at a novel or a familiar location, whereas that information was barely available in the superficial SC (Figure 2.5, right). Of course a decoder with access to the entire history of responses could decode stimulus novelty also from the superficial SC. By contrast, in the deep SC that information is available on individual trials. In the next section we explore how the information about stimulus history may be stored by the collicular circuit.

A working model for circuit mechanisms of visual sifting

The microcircuitry of the SC is still poorly understood, at least compared to that of the retina. One can distinguish about 5 to 10 neuronal types based on morphology and gene expression [5, 19], but their synaptic connectivity is largely unknown. Furthermore the SC interacts through long-range connections with other brain regions, notably the visual cortex [62]. Nevertheless, it is useful to consider what circuit mechanisms may produce the observed visual responses of SC neurons. The functional evidence we have gathered here makes some potential explanations unlikely, and supports others as a guide in future studies of synaptic connectivity. Here we focus on explaining three aspects of visual processing encountered in some deep SC neurons: the selectivity for looming stimuli, the invariance to spatial location, and the long-lasting stimulus-specific habituation. None of these phenomena occur in responses of retinal ganglion cells, and thus they must arise from post-retinal circuitry.

One circuit model that accounts for all the observed effects is shown in Figure 2.6A (“the working model”). It starts with input signals from retinal ganglion cells.

Those are combined to produce neurons selective for a local looming stimulus. The outputs of many such local looming detectors are pooled to produce neurons with widefield sensitivity and position invariance. Finally the input synapses to those widefield neurons undergo a short-term synaptic depression that accounts for the stimulus-selective habituation.

To simulate the function of this circuit we modeled each of the neurons as a Linear-Nonlinear element [8], and the synapses according to a widely used formalism for short-term plasticity [69]. This model correctly recapitulates the preference for looming over other stimuli (Figure 2.6E); the position invariance; and the habituation to familiar stimuli (Figure 2.6F). It even accounts for detailed dynamics of the looming response in deep neurons, such as the short latency (Figure 2.3A) and the rapid quenching of the response caused by synaptic depression (Figure 2.1C, Figure 2.4A).

While a successful circuit model seems promising, one learns something useful only from comparing different explanations. Here we consider several alternative microcircuits to account for the looming selectivity and the stimulus-selective habituation.

The working model (Figure 2.6A) builds on local looming-selective neurons. We encountered multiple cells in the superficial SC that match this profile: a local receptive field, looming selectivity, and little habituation (Supplementary Figure S2.6). In the working model this selectivity is achieved by combining signals from retinal ganglion cells (RGCs) with different dynamics: excitation from a fast and transient Off-cell forms the receptive field center, and inhibition from slow and sustained Off-cells forms the surround. Since RGCs are excitatory, the inhibition requires interneurons in the SC, and the slow dynamics of the surround may well result from filtering by those interneurons. In either case the concentric organization of fast excitation and slow inhibition produces selectivity for looming over contracting white or moving or dimming stimuli (Figure 2.6E).

As an alternative explanation, could the looming selectivity already originate in RGCs? As we noted previously, the “approach-sensitive” Off-cell that has been previously reported [46] is now known to correspond to the Off-transient alpha cell [56] which – while sensitive to looming stimuli – responds equally well to dimming and flashing spots [36]. Therefore these RGCs do not qualify as the local looming detectors.

Another possibility is that looming selectivity results from a radial organization of direction-selective (DS) neurons, each of which reports a segment of the advancing dark edge (Figure 2.6B). Supposing those DS inputs come from the retina, the only candidates are the On-Off DS RGCs [60], which would be equally sensitive to On edges. Thus the looming detectors in the SC should respond to an expanding white disk as well, unlike what we observed (Figure 2.1C). If, on the other hand, the DS signals are generated *de novo* in the SC, one would expect to find such interneurons with all possible preferred directions. Instead, DS neurons in a given region of the superficial SC have a strong bias for just one or two preferred directions [11]. In summary, both of the considered alternative microcircuits for looming selectivity seem unlikely given the available evidence.

In the working model (Figure 2.6A), the stimulus-selective habituation is produced by activity-dependent depression of the synapses that convey the local looming signals to the widefield neuron. A plausible alternative mechanism would involve long-lasting inhibition of the looming detector from a neuron triggered by that same local stimulus (Figure 2.6C). This neuron would need to exhibit a sustained activity following a single stimulus. In our database of collicular recordings, we never encountered a neuron that matches this description. Another possibility is that local looming detectors – in addition to exciting the widefield neuron – also inhibit it via an interneuron (Figure 2.6D). Then the long-lasting habituation could be explained by the potentiation of the inhibitory synapse, rather than depression of the excitatory synapse. In that case, one might expect that repeated looming stimuli should produce a suppression of the ongoing baseline firing during later stimulus periods. We never observed such a suppression (Figure 2.4A). Instead the firing generally increased during stimulus intervals (r_{stim}) compared to inter-stimulus intervals (r_{isi}) (for 15 deep SC neurons with baseline firing > 10 spikes/s, median r_{stim}/r_{isi} : 1.28, 25th-75th percentile range: 1.03-1.85).

In summary, several alternative explanations for the basic phenomena observed in deep SC neurons seem less likely than the working model that we propose, based on our database of extracellular recordings. We suggest that the key components of the working model in Figure 2.6A, namely the microcircuit for looming selectivity and the long-lasting synaptic depression, are fruitful targets for further investigation.

2.4 Discussion

Summary

The superior colliculus (SC) presents an interesting interface between purely sensory representations and pre-motor signals. Our goal here was to follow systematically how the sensory inputs from the retina get digested and filtered in the SC. As a guiding problem we chose a robust visually-triggered behavior: the defensive reaction elicited by an overhead looming stimulus. By following visual responses of neurons from superficial to deep layers we documented three aspects of the sifting process: (1) an increasing selectivity for the behaviorally relevant looming stimulus over other innocuous stimuli with similar low-level features (Figure 2.2); (2) an increasing invariance to other aspects of the visual display, such as the precise location of the threat stimulus (Figure 2.3); and (3) an increasing selectivity for novel over familiar stimuli (Figure 2.4). We considered how this filtering may be achieved by neural circuits and arrived at a plausible model of circuitry in the SC (Figure 2.6) that accounts for all three of the phenomena of visual sifting considered here. Moreover, several alternative circuit-level mechanisms were found to be inconsistent with the neural signals we encountered.

Relation to earlier work

Some of the phenomena reported here have been described before in a wide range of species. A common theme is that neurons in deep SC respond over larger regions of the visual field, while retaining a preference for small stimulus features within that region [9, 13, 21, 27, 29]. Also, the remarkably persistent habituation to repeated stimuli has been noted previously, even in the earliest recordings from optic tectum [9, 13, 25, 39, 54, 67, 74]. Another repeated observation is that the visual cortex appears dispensable for many aspects of visual processing in the SC [24, 27, 43], although it does play a subtle modulatory role [77]. Looming stimuli are particularly effective for many neurons in the superficial SC [77]. Interestingly the early literature missed this, perhaps because of the technical difficulty of generating an expanding dark disk with the commonly used hand-held slide projector [13]. Our present report places these disjoint observations into a common context, namely the animal's need to distill a specific signal of ecological value from the broad range of visual stimuli. We show that SC neurons are not only sensitive to looming stimuli but become increasingly selective in deep layers, an essential requirement for an alarm system. Further we analyze the neural code at the population level throughout this brain region, which reveals the gradual progression of stimulus filtering. Finally

we consider how these aspects of neural representation relate to neural circuitry, and evaluate alternative hypotheses for such circuits. The results allow a broader consideration of how selectivity and invariance come about in brain processing, to be pursued further below.

Ethological significance

The present study focused on stimuli presented in the upper visual field and recordings performed from the corresponding medial region of the SC. Arguably the most behaviorally relevant event in the upper visual field is the impending arrival of a bigger animal, such as an aerial predator. The imminent threat that these events pose may account for the profuse responses to dark looming stimuli among SC neurons in this region (Figure 2.2, [77]). Of course the threats must be distinguished from innocuous events, like the movement of overhead foliage, or the obscuring of the sky when the animal moves under shelter. The increased selectivity to the expanding dark disk in the deeper SC can account for that selectivity (Figure 2.2C, Figure 2.4B).

How should one interpret the profound habituation to repeated stimuli in this context? For one, the habituation does not interfere with the alarm response, since the animal must react to the first occurrence of a clear looming stimulus [75]. If the animal escapes or freezes, and the predator approaches a second time, this is likely in a different part of the visual field, and thus unaffected by the location-specific habituation. On the other hand, if the same stimulus recurs periodically in the same location, it is more likely caused by a leaf waving in the wind. Thus the habituation can be seen as another processing strategy to reject innocuous events from the alarm pathway.

In the lower visual field the animal has different behavioral needs, such as picking out seeds against a cluttered background, following small moving prey [26], perhaps identifying urine marks [30], and tracking optic flow. Furthermore the connectivity between SC and other brain areas seems to differ in the upper and lower visual fields [61]. Thus one expects a corresponding difference in the rules by which visual stimuli are sifted there, a fertile area for future study.

Selectivity, invariance, and habituation

One remarkable phenomenon in sensory processing is the emergence of neuronal responses that are both highly selective and broadly invariant. For example, certain “face cells” in the primate visual cortex respond selectively to one person’s face

regardless of the view angle, scale, or illumination [18]. How do these seemingly conflicting characteristics arise within sensory circuits? In the working model we propose here (Figure 2.6A) the answer is “first selectivity, then invariance.” An AND operation across input neurons with different dynamics generates a local looming-selective neuron. These pattern detectors are distributed across the visual field. Then an OR operation pools across many local pattern detectors to produce the position-invariant response of the widefield neurons (Figure 2.7A).

This seems to be the scheme in other neural systems where the circuitry is understood. For example, in the auditory brain of the barn owl certain high-order neurons are selective for a particular interaural time delay, but invariant to the frequency of the sound [33]. These appear to arise from OR pooling over lower-order neurons that are selective for the same time delay but still tuned to different frequency bands. Those delay detectors in turn arise from an AND combination of signals derived from the two ears [7]. A similar processing scheme applies in the electrolocation circuits of weakly electric fish that exhibit a jamming avoidance response sensitive to frequency but invariant to many other parameters of the electric field [22].

However, this is not the only solution. In the case of face recognition, for example, it seems implausible that the brain should build separate pattern detectors for each face at each retinal location, and then pool over those to achieve invariance. An alternative scheme produces invariance first and then selectivity (Figure 2.7B). Here there exists only a single pattern detector. But the inputs to this neuron are routed to “look at” different spatial locations through a shifting circuit. The sudden appearance of any stimulus could engage these shifter circuits to route the corresponding low-level visual signals into the pattern detector [48, 71].

The observation of habituation and its specificity to location seems to greatly favor one of these schemes. Recall that habituation is seen prominently among neurons in the deep SC that are already highly pattern-selective. The “selectivity first” scheme places the gain modulation somewhere prior to the output of the pattern detector, which is the last spatially localized signal (Figure 2.7A). By contrast, the “invariance first” scheme requires the gain modulation to occur in low-level visual neurons that are not yet pattern-selective (Figure 2.7B). This conflicts with our observations of neurons in the superficial SC that do not show location-specific habituation (Figure 2.4B). In summary, the robust observation of location-selective habituation in neurons of the deep SC favors a circuit model that develops selectivity before invariance.

Of course, one can also envision intermediate solutions. For example, there is speculation that the visual cortex implements an alternation of AND and OR stages through a hierarchy of anatomical areas [12, 55]. Seeing that most vertebrate species do not have a neocortex, yet must solve the same problems of invariant pattern recognition, the SC seems like a promising arena for the study of high-level visual processing.

Circuit mechanisms of sensory sifting

While it is tempting to suppose that the observed reduction of visual data is in fact performed within the SC, such anatomical localization is not a binding conclusion. In the extreme, the neurons of the deep SC, with their selectivity for fine spatio-temporal features and localized habituation, may simply reflect the output of a computation that occurs elsewhere. The SC interacts with many other anatomical structures [3, 61], often in a reciprocal fashion. The most prominent such input, namely from the visual cortex, is likely not required for the essentials of visual sifting, based on our results with mutant mice (Figure 2.4E) and prior work with cortical silencing [77] and ablation [24, 27]. To contribute to sifting, the partner areas should retain a spatial resolution of the stimulus on the order of 10° . This constraint eliminates some small nuclei, but leaves several candidates in place, for example the thalamic area LP [1] and the parabigeminal nucleus. Given the position of the SC as a hub of brain pathways, it is an open question whether one can ultimately assign discrete computational functions to discrete anatomical areas.

On a finer level, one may ask how the circuit models of Figure 2.6 map onto neuron types in the SC. About five cell types have been distinguished in the superficial SC of mammals based on morphology alone [37, 44], and more recent studies have connected these types to visual responses and electrophysiological properties [19]. The most compelling by their visual appearance are the so-called widefield or bottlebrush neurons. These cells have a dendritic fan that extends towards the surface of the SC and spreads out laterally to cover a large area in the retinorecipient layers. Each dendrite terminates in a bottlebrush-shaped ending, and the overall morphology is startlingly similar across birds and mammals [40, 42]. The widefield neurons of mammals project to the pulvinar, and the axon forms multiple collaterals in the SC that could propagate the output to the deep layers [3, 42].

By virtue of their broad dendritic tree, these widefield neurons offer themselves as the substrate for pooling across spatial locations, as in the working model of Figure

2.6A. Two further features recommend such an identification: First, the dendrites of widefield neurons generate spikes that propagate to the soma [15, 41]. In this way, the neuron truly implements an OR operation across its inputs (Figure 2.7A): when any of its inputs fire, the output will fire. Second, experiments on chick tectum showed that each dendritic input undergoes a profound synaptic depression that lasts several seconds, but does not affect the function at another dendrite [41]. This could account for the location-specific habituation as in the model of Figure 2.6A. However, there is some question whether this synaptic depression also happens in the mouse [20]. Also we found a substantial increase of invariance below the anatomical stratum where the widefield neurons reside (Figure 2.3).

In summary, the visual response properties of deep SC cells differ dramatically from any signal that emerges from the retina, and it is tempting to associate this transformation with the bottlebrush neuron that is shaped unlike anything in the retina. Some caution is in order, of course. The diagram of Figure 2.6A should be viewed as a conceptual scheme rather than an explicit circuit with one-to-one corresponding real neurons. Perhaps the selectivity and invariance are accomplished in multiple stages, or with the contribution of other brain areas. Or the local looming detectors may be nonlinear dendrites, and ion channels with long-lasting inactivation [70] may play the role of depressing synapses. The increasing availability of genetic handles for cell types in the SC [5, 19] should help in cracking some of these microcircuits.

2.5 Figures

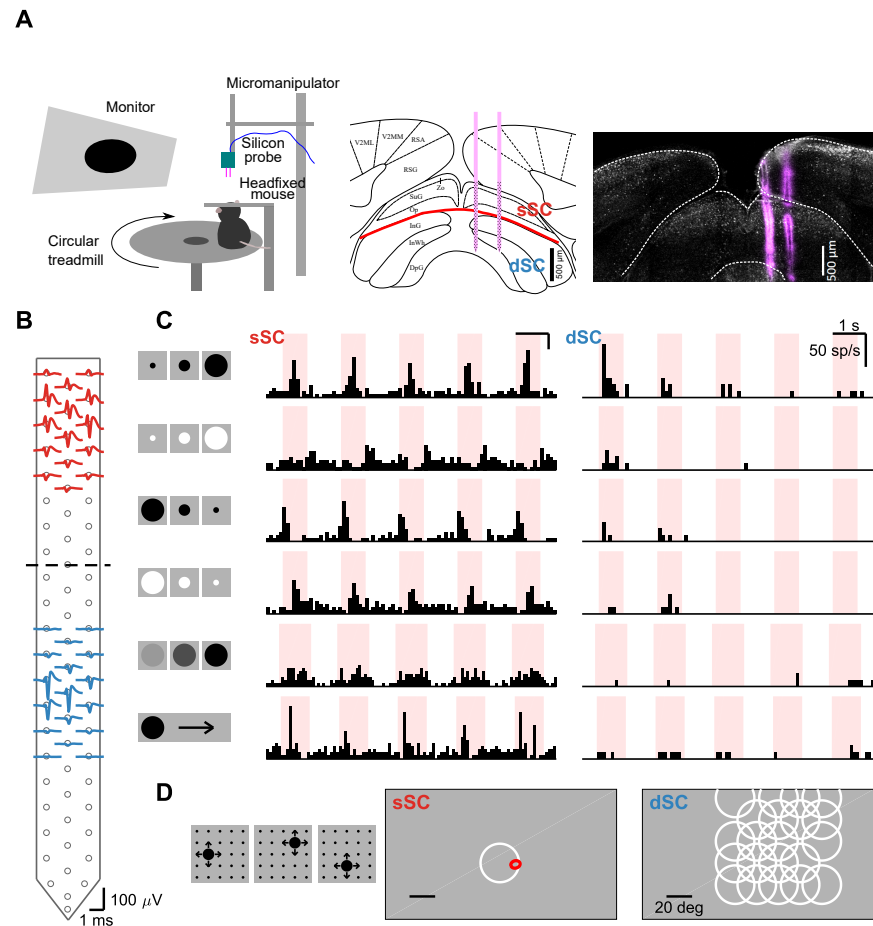



Figure 2.1: The emergence of selectivity, invariance, and stimulus-specific habituation along the depth of SC. (A) Left: Experimental setup. Silicon neural probes with 128 channels were implanted into the SC of a headfixed mouse viewing visual stimuli. The mouse was free to run on a circular treadmill. Middle: Diagram of a coronal section showing the anatomically defined layers of the SC (from [50]). sSC: superficial SC; dSC: deep SC. Right: Corresponding histological section recovered after neural recording, showing tracks of two electrode prongs. Magenta: DiI; white: anti-Calb1. (B) Extracellular spike waveforms of sample sSC (red) and dSC (blue) neurons recorded simultaneously on the silicon probe. Dots indicate the location of recording sites. Dashed line indicates boundary along the electrode array between sSC and dSC (see Methods and Supplementary Figure S2.1). (C) Response of neurons from (B) to visual stimuli. The sSC neuron (middle) responds to many types of figural stimuli (left icons: expanding black, expanding white, contracting black, contracting white, dimming, and moving black disk), whereas the dSC neuron (right) is highly selective to the expanding black disk. The sSC neuron responds robustly to every trial, whereas the dSC neuron responds primarily to the first presentation. (D) In an experiment in which looming stimuli appear from many locations (left), the sSC neuron from (B) (middle) is driven only by stimuli that cross its receptive field, whereas the dSC neuron from (B) (right) responds to stimuli placed at many more locations. White: final size of looming stimuli that elicited significant response from the cell; red: 1 standard deviation outline of spatial receptive field recovered by spike-triggered average method.



Supplementary Figure S2.1: Histological and electrophysiological identification of SC layers. (Top) Brain section (coronal, 100 μm thick) showing the recorded area from an experiment. (Bottom left) Boxed area in the top panel is enlarged. The probe track is marked with DiI (magenta). The anti-Calb1 antibody (green) stains superficial gray layer (SGS) and upper parts of the intermediate gray layer but does not stain the optic layer (SO). White lines mark the top and bottom outline of the SGS and SO. (Bottom right) Current source density analysis of the same recording displayed at the same spatial scale as the section on the left. Looming stimulus was delivered during the time window shaded in gray. Dashed line marks the inflection point that separates current sink (blue) below and current source (red) above. This corresponds to the lowermost point of SGS, as previously reported by [66].

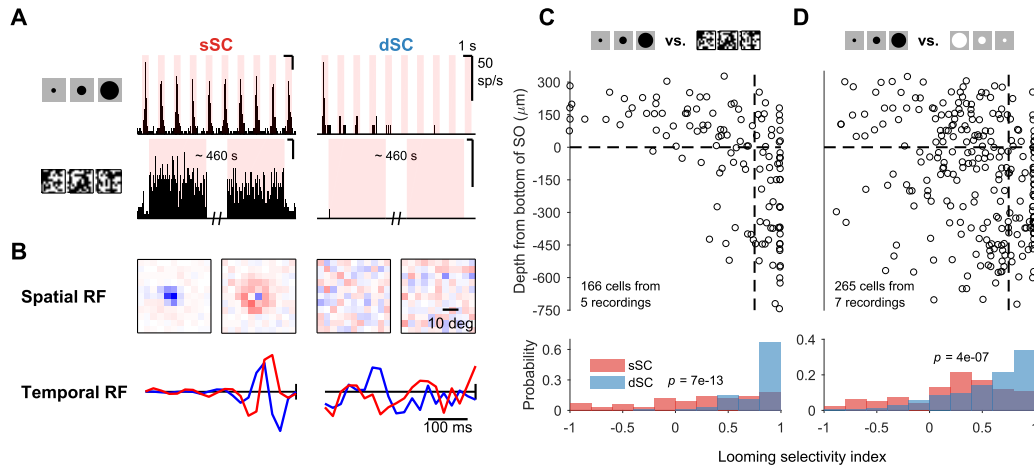
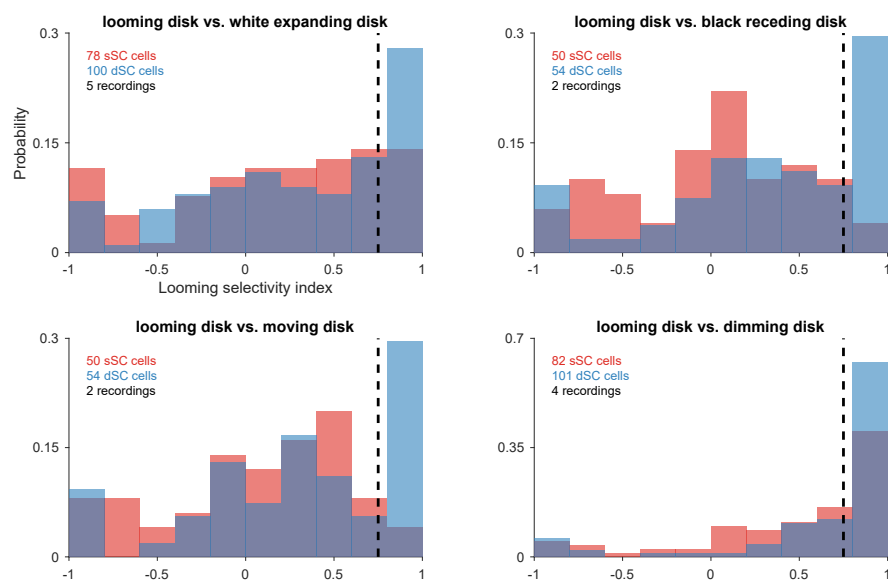


Figure 2.2: Selectivity to looming stimulus. (A) Response of sample sSC (middle) and dSC (right) neurons to looming stimulus (top) and flickering checkerboard (bottom). sSC neuron is driven strongly by both, but dSC neuron is almost completely silent to the checkerboard stimulus. (B) Spatial (top) and temporal (bottom) receptive fields of the sSC (left) and dSC (right) neurons in (A) based on spike-triggered average analysis. In each subpanel, left: spatial center; right: spatial surround; bottom blue: temporal center; bottom red: temporal surround. In the temporal RF panels, the vertical line represents the time of the spike. (C) Population summary of selectivity to looming stimulus over checkerboard stimulus along the depth of SC. Horizontal dashed line indicates the boundary between sSC and dSC. Vertical dashed line separates neurons with high selectivity index (>0.75) from others. The p -value (two-sample Kolmogorov-Smirnov test) indicates that the distributions of sSC and dSC neurons differ significantly. (D) Same as (C), comparing responses to looming stimulus and contracting white disk. Selectivity index is defined as $(r_L - r_O)/(r_L + r_O)$ where r_L refers to response to looming stimulus and r_O refers to response to checkerboard stimulus (C) or contracting white disk (D).



Supplementary Figure S2.2: Distribution of looming selectivity index for neurons in the superficial and deep SC vs. expanding white disk, receding dark disk, moving disk, and dimming disk. Dotted line separates neurons with high looming selectivity index (>0.75) from others.

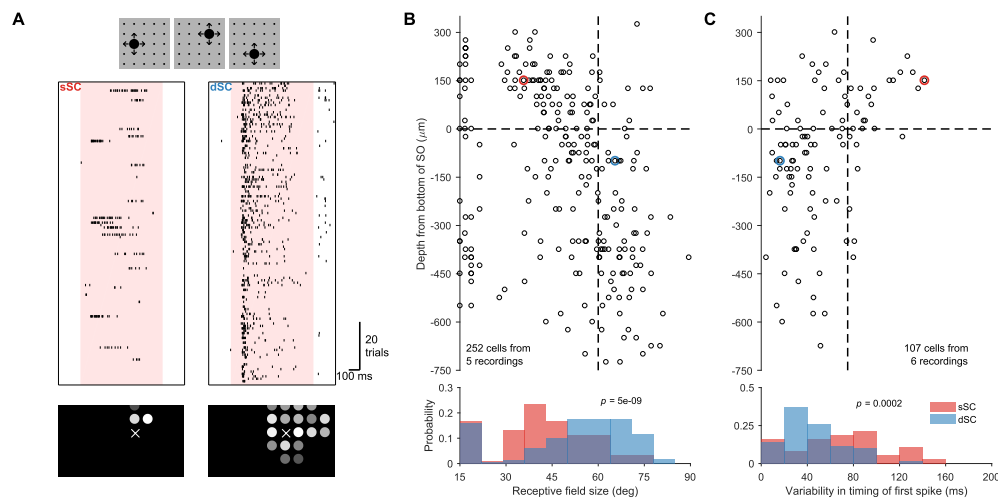


Figure 2.3: Invariance to stimulus position. (A) Raster plot of sample sSC (left) and dSC (right) neurons recorded simultaneously during an experiment in which looming stimuli appear randomly in one of 25 locations (small black dots in cartoon) in each trial. These locations are $\sim 15^\circ$ apart. The dSC neuron responds to many more locations than the sSC neuron and with an invariant latency. Bottom: The response amplitude at each location is reported by the brightness of the circle. X indicates a location that received no stimulus. (B) Population summary of receptive field size estimated from the experiment in (A). Vertical dashed line is at 60° . (C) Population summary of variability in the timing of the first spike from the experiment in (A). Vertical dashed line is at 75 ms. In both (B) and (C), the horizontal dashed line separates sSC and dSC. The red and blue circles denote the sSC and dSC neurons from (A). The p -values (two-sample Kolmogorov-Smirnov test) indicate that the distributions of sSC and dSC neurons differ significantly.

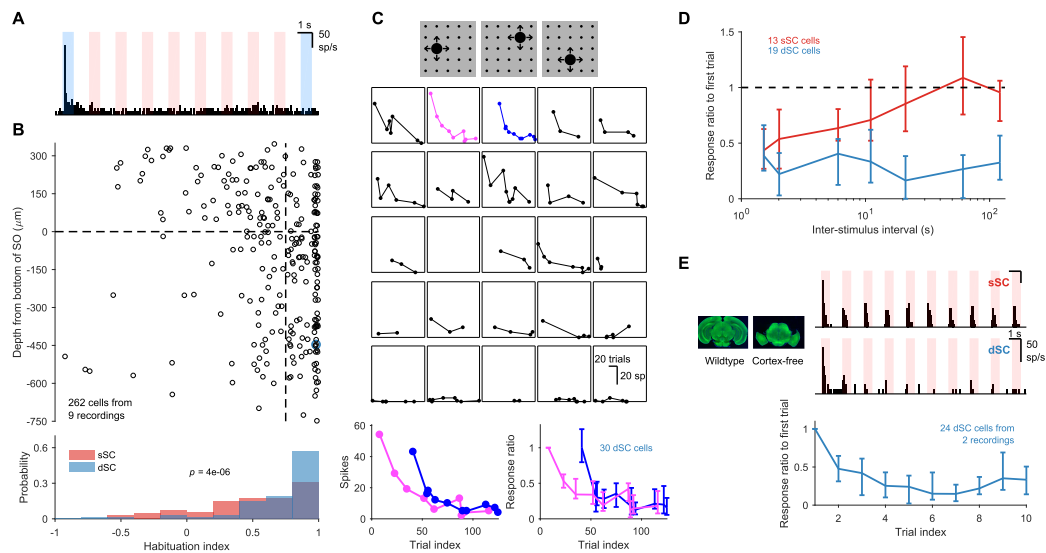

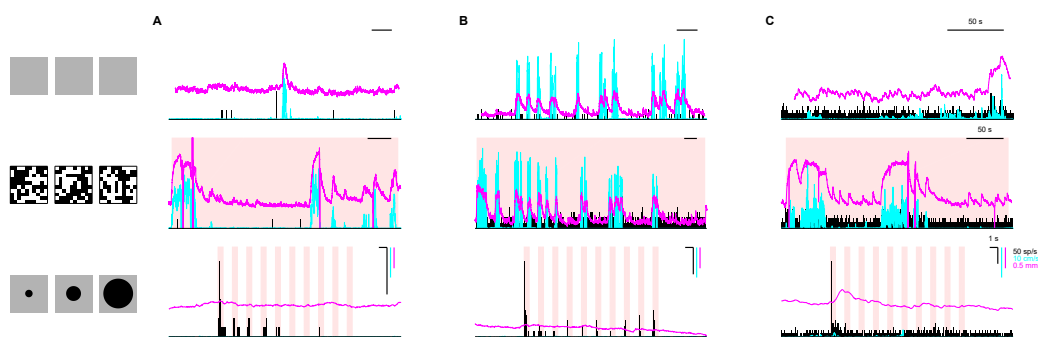



Figure 2.4: Stimulus-specific habituation. (A) Response of a sample dSC neuron to a series of 10 looming stimuli. The first and the 10th trials are shaded in blue. Note that this neuron has a maintained baseline firing rate, which is unchanged by the stimulus on all but the first trial. (B) Population summary of habituation to repeated looming stimulus. The habituation index is defined as $1 - r_1/r_{10}$ where r_i refers to the number of spikes fired in i -th trial after subtracting background activity. The horizontal dashed line separates sSC and dSC. The vertical dashed line is at 0.75. The blue circle is the sample dSC neuron from (A). The p -value (two-sample Kolmogorov-Smirnov test) indicates that the distributions of sSC and dSC differ significantly from each other. (C) Response of a sample dSC neuron to ~ 100 presentation of looming stimuli delivered in random sequence. Each subpanel represents response to stimuli at one of the 25 locations. Bottom left: two of the response traces from above. Even after the neuron has habituated to stimuli at one location (magenta), it responds strongly to the first stimulus at another location (blue). Bottom right: response of all dSC neurons in this recording, normalized by response to first trial of the magenta trace. Data points are medians and error bars range from 25th to 75th percentiles. (D) Summary of time to recover from habituation for a group of simultaneously recorded sSC and dSC neurons. Even after ~ 120 s, dSC neurons do not recover beyond 50% of the initial response. Data points are medians and error bars range from 25th to 75th percentiles. (E) Sample sSC (top right) and dSC (middle right) neurons recorded in a mutant mouse that does not develop the neocortex or the hippocampus (left). The dSC neuron in the mutant mouse also shows habituation. Bottom: population response of dSC neurons to 10 presentations of the looming stimulus, normalized by the response to the first presentation. Data points are medians and error bars range from 25th to 75th percentiles.



Supplementary Figure S2.3: Suppression of a familiar stimulus is not permanent. The response of a sample dSC neuron to looming stimuli that were presented in the beginning (left) and the end (right) of a recording. Other stimuli (e.g. checkerboard) were presented to the animal in the ~50 min that separates these two blocks. Note that this neuron eventually does recover (24 vs. 26 spikes to the first presentation), indicating that the memory of looming stimulus is not permanent.



Supplementary Figure S2.4: Enhanced response to first stimulus is not a simple consequence of motor activity. The response of three sample dSC neurons from Figure 2.1 (A), Figure 2.3A (B), and Figure 2.4A (C) to no stimulus (top), checkerboard stimulus (middle), and looming stimulus (bottom) are shown, along with simultaneously recorded locomotion (cyan) and pupil diameter (magenta). The neurons are silent during no stimulus and checkerboard stimulus even though the locomotion and pupil size are modulated. This suggests that the strong response in the first presentation of the looming stimulus is not simply a result of motor output or change in arousal level in the absence of visual threat. In one case (C, bottom), the burst of firing to the looming stimulus precedes a significant change in the pupil size, suggesting an increased level of arousal.



Supplementary Figure S2.5: A mutant mouse that lacks the neocortex and the hippocampus. A series of *Emx1-Cre:Pals1^{flox/flox}* mouse brain coronal sections (100 μm thick), labeled from most anterior (1) to most posterior (9). The left half of panel 7 shows a corresponding wild-type mouse coronal section with a fully developed neocortex and hippocampus. Note that these structures are lacking in the mutant mouse. Autofluorescence (green, red), DAPI (blue).




Figure 2.5: Population decoding of distinct stimulus features. Linear decoders were trained with simultaneously recorded sSC and dSC neurons to predict location (left) and novelty (i.e. whether the stimulus has appeared at a location for the first time) (right) of stimuli in the experiment described in Figure 2.3. Dashed line: chance performance; error bars: 1 standard deviation across different subsamples of cells.

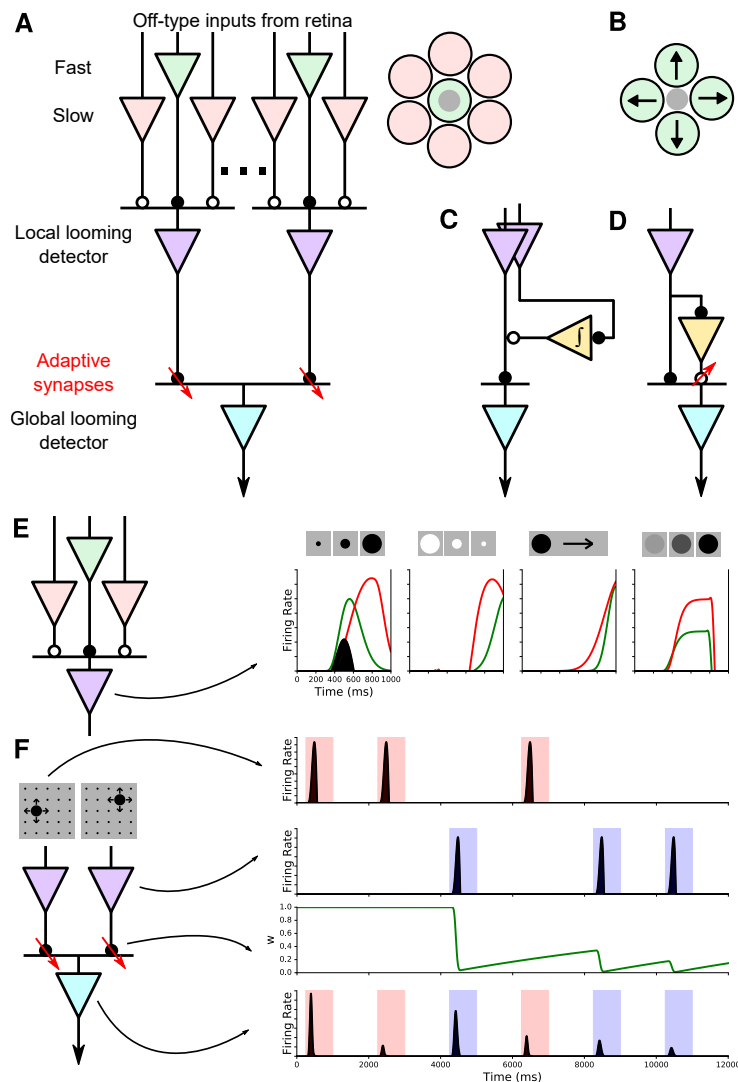
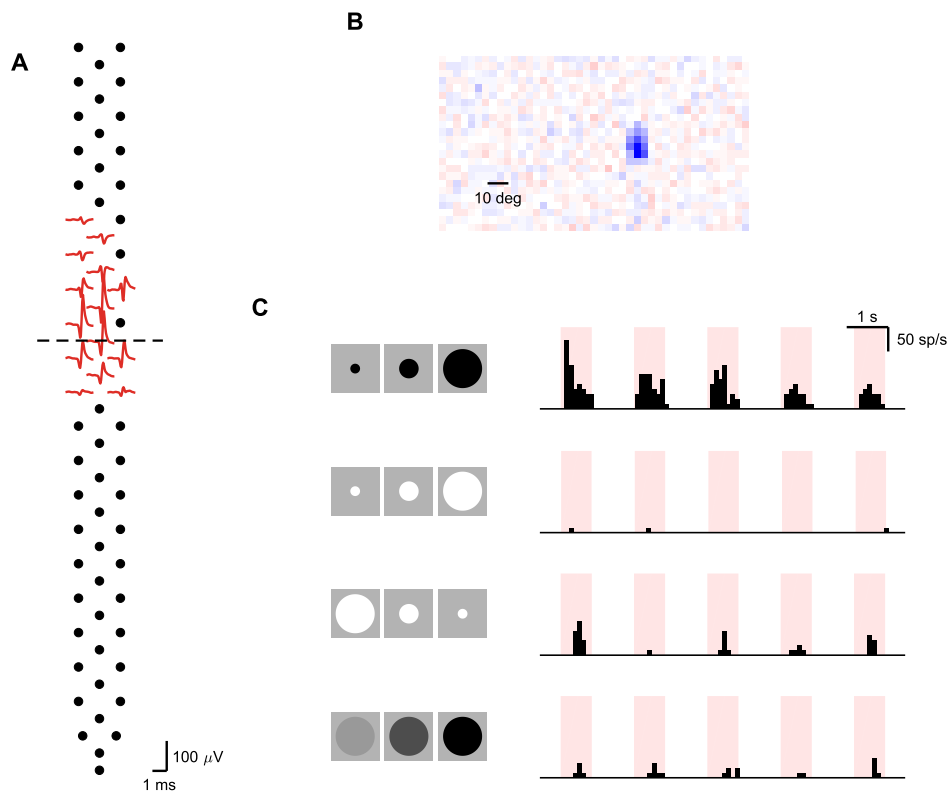


Figure 2.6: Model of selectivity, invariance, and stimulus-specific habituation. (A) The “working model” of how selectivity, invariance, and habituation arise in the dSC. Looming selectivity is generated by combining fast and slow Off-type retinal inputs (green and pink) in the local looming detector (purple) in sSC. Inset on right shows spatial layout of these inputs. Invariance arises from pooling these local looming detectors to a single global looming detector (cyan) in the deep layers. The stimulus-specific habituation is achieved by synapses that undergo activity-dependent short-term depression (red downward arrows). Solid circles: excitation; open circles: inhibition. (B) An alternative model of looming selectivity based on pooling directionally tuned inputs. (C, D) Alternative models of stimulus-specific habituation: the same input as the excitation drives a persistent inhibition (C) or a facilitating inhibitory synapse (D). (E) Simulation of responses to various figural stimuli. Green: excitation from center; red: inhibition from surround; shaded black: net response. (F) Simulation of stimulus-specific habituation. Each local looming detector connects to the global looming detector with a synapse whose strength w decays rapidly and recovers slowly.



Supplementary Figure S2.6: One of several putative local looming detector identified in the superficial SC (A), with a local receptive field (B) and selectivity for looming stimulus without a significant habituation to repeated stimuli (C). Dashed line: boundary between sSC and dSC.

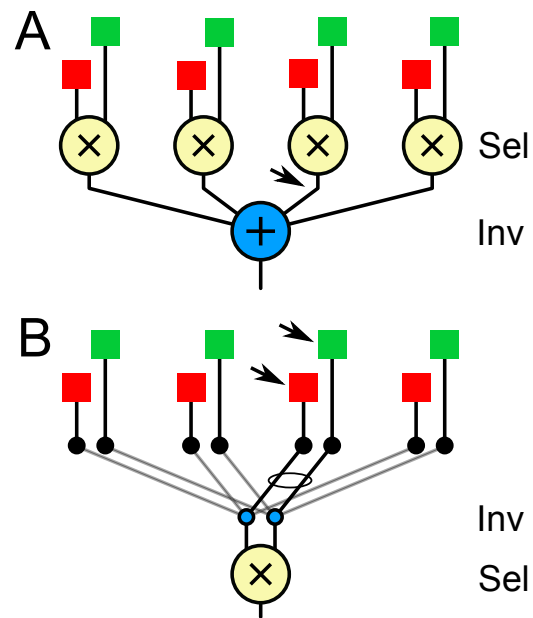


Figure 2.7: The logic of selectivity and invariance. In (A) feature selectivity is accomplished by combining local input signals (red and green) with AND logic (X). Then invariance arises from combining many of those feature signals with OR logic (+). In (B) there is only a single feature computation (X). Invariance is achieved by routing its inputs to local signals in different parts of the visual field. Arrows indicate where the stimulus-specific habituation must take place.

2.6 Methods

Mouse, surgery, neural recording and spike-sorting

We used C57BL/6 mice (RRID:IMSR_JAX:000664) aged 3–10 months (both males and females, Jackson Labs) for electrophysiological recordings. To prepare an animal for an experiment, we first implanted a metal headplate to the skull with a dental adhesive (3M Scotchbond) under anesthesia (2 % isoflurane). After three days of recovery, the animal was habituated to being head-fixed on a circular treadmill for ~30 min/day for three days. On the day of recording, the animal was again anesthetized and a craniotomy (< 1 mm diameter) was made over the SC (0.2–0.4 mm anterior to lambda, ~0.5 mm lateral from midline). A small hole was made over the cerebellum to insert a silver reference wire. The craniotomy was then closed with a silicone elastomer (Kwik-Cast, WPI). After 6–8 hours of recovery, the animal was head-fixed and the craniotomy was exposed. A silicon neural probe was then lowered slowly into the brain (< 5 $\mu\text{m/s}$) and the depth from brain surface was recorded. The craniotomy was then covered with mineral oil to prevent drying of the exposed tissue. We waited a short period (15–30 min) for signals to stabilize before starting the recording. A typical recording session lasted 2–3 hours. All procedures were performed in accordance with institutional guidelines and approved by the Caltech IACUC.

The silicon neural probes were obtained from Sotiris Masmanidis (UCLA) [14]. For the majority of experiments, probe types 128A, 128AN, and 128DN were used. For data acquisition we used the RHD2000 128-channel amplifier board and the RHD2000 USB interface board (Intan). Auxiliary signals including the movement of the running wheel, timing of the stimulus, and timing of pupil video recording were collected concurrently with the neural signal. We used KiloSort [49] for spike-sorting of the data. The output of the automatic template-matching algorithm of KiloSort was visualized and manually curated on Phy [57, 58].

To test if the long-lasting stimulus-specific habituation requires the neocortex and the hippocampus, we also recorded in mutant mice that developmentally lack these brain areas [31]. These animals were bred by conditional knockout of exon 3 of Pals1 gene in cortical progenitor cells during embryonic development, achieved by crossing Pals1^{flox/flox} mice with LoxP sites inserted upstream and downstream of exon 3 with Emx1-Cre animals (Jackson Labs, Strain 005628) expressing Cre recombinase in the cortical progenitor cells. Conditional knockout of both copies of Pals1 due to Cre-mediated recombination during development resulted in Emx1-

Cre:Pals1^{flox/flox} homozygous progeny used in this study (Supplementary Figure S2.5).

Behavioral measures

The animal's pupil diameter and locomotion on the circular treadmill were recorded along with the neural signals. The animals were not trained in any particular task and varied in their tendency to run on the treadmill. When looming stimuli were presented, the animals sometimes reacted by stopping (if the stimulus had arrived during a movement bout) or showing an increase in the pupil size (Supplementary Figure S2.4C), but no characteristic behavioral output was consistently observed. However, we could rule out the possibility that the strong response of deep SC neurons to the first presentation of the looming stimulus is a simple consequence of motor output or change in pupil size, as they were usually not modulated by these factors in the absence of the looming stimulus (Supplementary Figure S2.4). We also tracked the position of the pupil to monitor the eye movements. In many cases the eyes were very stable, as demonstrated by the sharp ($\sim 5^\circ$), circular receptive fields we recovered (Figure 2.2B) in the superficial SC by spike-triggered average analysis.

Post-hoc identification of the recorded brain area

Prior to implanting into the brain, the tip of the silicon probe was covered with fluorescent lipophilic dye (DiD or DiI, Invitrogen). Immediately after recording, the animal was anesthetized and perfused with saline and 4 % PFA. The brain was harvested and fixed with 4 % PFA (Electron Microscopy Sciences) for 24–48 hours at 4 °C, after which it was sectioned coronally at 100 μm thickness with a vibratome (Leica). The sections were then stained with anti-Calb1 antibody (Swant, CB-38a, 1:1000 dilution), which has been previously reported to label the superficial gray layers of the SC [59]. Following secondary antibody staining (AlexaFluor 488, donkey-anti-rabbit, 1:1000 dilution), sections were mounted with Vecta-Shield:DAPI and imaged using a confocal microscope (LSM800, Zeiss). From this, the location of the probe relative to SC layers could be estimated (Supplementary Figure S2.1). This histology-based method of localizing the probe relative the SC layers was complemented with current source density (CSD) analysis. First, the raw, broadband recording was low-pass filtered (150 Hz cutoff) to isolate the LFP band. Then the Laplacian of a column of spatially contiguous electrodes was computed and smoothed with a Gaussian kernel. This revealed a series of current sources and

sinks in response to visual stimulation (Supplementary Figure S2.1). By comparing this CSD analysis to the histological localization, we confirmed the results from [66] that the inflection point between the current source and sink marks the bottom of the superficial gray layer (SGS). We then defined the boundary between the superficial and deep layers as 100 μm below the inflection point (corresponding to 0 depth in Figure 2.2C-D, Figure 2.3B-C, and Figure 2.4B) to account for the thickness of the optic layer.

Stimuli

Visual stimuli were programmed using the Psychtoolbox [4, 32] package in MATLAB (Mathworks) and presented on a gamma-corrected monitor (IPS231, LG) at baseline luminance of $\sim 25 \text{ cd/m}^2$. The position of the monitor was slightly adjusted in each experiment such that the receptive fields of the neurons being recorded were located near the center of the monitor. Usually this was at $\sim 35^\circ$ in elevation and $\sim 45^\circ$ in azimuth (to the left) from the rostro-caudal axis of the animal. The monitor was located 15–20 cm from the animal and covered $\sim 120^\circ$ of the horizontal field of view. The visual stimuli were synchronized to the neural recording by using a photodiode to send timing pulses from the monitor to the data acquisition board.

Before presenting the stimuli, we used a small flickering spot to map the part of the monitor that elicited strong neural responses (“response zone”). Figural stimuli were then presented at these locations. The following is a description of each stimulus type during the stimulus period.

Definition of stimulus period

Throughout this report, the periods during which the stimulus was presented on the screen are called stimulus periods and are marked as pink sections in the PSTHs. Outside the pink sections, the screen was uniformly gray.

Looming stimulus

The looming stimulus expanded from 0° to $\sim 30^\circ$ at a linear expansion rate of $\sim 30\text{--}60^\circ/\text{s}$ and then remained stationary for another 250 ms before disappearing. It was presented at the full contrast achievable by the monitor. This repeated for 5–10 trials at the same location. The inter-stimulus interval was 1–3 s, except when the time to recover from habituation was explicitly tested (Figure 2.4D).

Other figural stimuli

The contracting black, expanding white, and contracting white disks were presented with similar parameters as the looming stimulus. The stationary period of 250 ms was always at the end of the expansion or the contraction. The dimming and the moving dark disks were the same size as the final size of the looming stimulus. The rate of change in contrast of the dimming disk and the trajectory and the speed of the moving dark disk were set such that they had roughly the same duration as the looming stimulus. The moving dark disk traveled at $\sim 40\text{--}70^\circ/\text{s}$, with the response zone in the middle of the trajectory. Several different movement directions were tried.

Flickering checkerboard

During the flickering checkerboard stimulus, the entire screen was divided into square checkers ($\sim 3^\circ$) whose intensity changed randomly between black and white in every frame at a refresh rate of 60 Hz. The duration ranged from 300 to 600 s, but often 300 s was enough for the spike-triggered average analysis.

Random loom

In the “random loom” experiment, 25 locations (in a 5x5 grid) around the response zone were selected, with $\sim 15^\circ$ between adjacent locations (measured from center to center). In each trial, one looming stimulus was presented in one of these locations with the parameters described above. The sequence of stimulus locations was determined with a pseudorandom number generator. The inter-stimulus interval was 3 s and $\sim 60\text{--}120$ trials were presented in total.

Analysis

The progression of visual response properties with depth in the SC was discovered in early exploratory experiments. A subsequent round of recordings was performed to validate the initial observations. The present manuscript analyzes data from only these replication experiments. All analysis scripts were written in MATLAB R2016b (Mathworks) unless otherwise noted.

Definition of neural response and background activity

Throughout our analysis, the neural response is defined as the number of spikes that a neuron fired during the stimulus period (as described above). Some neurons had a

maintained baseline firing rate. The background activity is defined as the expected number of spikes contributed by the baseline firing rate during the stimulus period. To compute this, we estimated the baseline firing rate by counting the spikes fired during the ~5–10 s-long period just preceding the stimulus and dividing by length of this period. We then multiplied this by the stimulus period to get the background activity. The background activity was used to test if the neural response was visually driven (see below).

Identification of visually responsive neurons

Many of the recorded neurons had no clear response to visual stimuli. In a typical neurophysiology experiment, visually responsive neurons can be separated from others by presenting the stimulus many times and choosing only those that respond consistently across repetitions. In our experiments, we did not have the luxury of repeating the stimuli, as many neurons (esp. in the deep SC) showed significant habituation after just a single presentation (Figure 2.4A-B). To identify visually responsive neurons from single trials, we instead used a statistical method. First, we computed the neural response and the background activity (see above). We then computed a p -value for the neural response based on a Poisson noise model whose mean was the background activity. If the p -value was less than the pre-set cutoff of 0.005, we considered the response to be visually driven. In cases where the background activity was very low (< 1 spike), the mean of the Poisson model was set to 1 so that chance firing of 1-2 spikes during the stimulus period would not be considered as a visual response.

This significance criterion was used to select neurons to include in the analysis shown in Figure 2.2, Figure 2.3, and Figure 2.4 (see below). When computing quantities of interest (e.g. selectivity index), we first subtracted the background activity from the neural response. In Figure 2.3, the analysis required identification of significant responses from a series of stimulus presentations. To compensate for this multiple comparison, we applied a Bonferroni correction by dividing the p -value cutoff by the number of stimulus presentations.

Receptive field analysis with flickering checkerboard

To measure the spatio-temporal receptive field (Figure 2.2B), we computed the spike-triggered average stimulus (STA) with the neural response to the flickering checkerboard [45]. In many neurons that had a strong STA, we could separate

the center and the surround of the receptive field by performing singular value decomposition (SVD) on the STA [73]. SVD expresses the spatio-temporal STA as a sum of terms, each of which is a product of a purely spatial and a purely temporal function. The terms are ordered by decreasing contribution to the overall variance in the STA data. We found that often the first term corresponded to the spatial and temporal profile of the center, and the second term to those of the surround.

Stimulus selectivity

To analyze the selectivity to a looming stimulus over other stimuli (Figure 2.2C, D), we computed the looming selectivity index defined as $(r_L - r_O)/(r_L + r_O)$ with $r = r' - \mu$, where r' refers to the number of spikes that a neuron fired during the first presentation of the stimulus, μ refers to the number of spikes expected during the stimulus period from the neuron's baseline firing rate, and the subscripts L and O refer to the looming stimulus and another stimulus (e.g. contracting white disk), respectively. For the comparison to flickering checkerboard (Figure 2.2C), $r_O = \langle r_C \rangle t_L - \mu$, where $\langle r_C \rangle$ is the average firing rate of the neuron during flickering checkerboard and t_L is the duration of the looming stimulus. Only neurons that were significantly responsive to either of the two stimuli being compared based on the Poisson significance criterion outlined above were included in the analysis.

Position invariance

To analyze the invariance to stimulus location (Figure 2.3B), we estimated the receptive field of recorded neurons from the results of the “random loom” experiment in which looming stimuli appeared randomly at one of 25 locations (5x5 grid) in each presentation. The looming stimulus was chosen because unlike the checkerboard stimulus, it reliably drove both sSC and dSC neurons. First, we defined the function $r(\mathbf{x})$ that specifies the maximum response (in spikes) of a neuron to a stimulus at location $\mathbf{x} = (x_1, x_2)$. Then we (1) set to zero the responses that did not deviate significantly from baseline activity; and (2) subtracted the expected number of spikes during stimulus period due to baseline activity from $r(\mathbf{x})$. To capture the width of the receptive field given by the remaining responses, we computed the mean radial distance $\Delta = (\sum \|\mathbf{x} - \mathbf{c}\| r(\mathbf{x})) / \sum r(\mathbf{x})$ where $\mathbf{c} = (\sum \mathbf{x} r(\mathbf{x})) / \sum r(\mathbf{x})$ is the center of mass of the receptive field and $\|\cdot\|$ is the Euclidean distance. Then we defined the receptive field size as 2Δ , i.e. twice the mean radial distance from the center of mass. Based on this method, neurons that respond to stimuli at only a single location

would have a receptive field size of zero, as $\Delta = 0$. We corrected this by adding the inter-center distance between stimuli (often $\sim 15^\circ$) to the estimated receptive field size of all neurons, as this determined the spatial resolution of our experiment.

Variability in response latency

In addition, we analyzed the variability of response latency during this experiment (Figure 2.3C). We defined the latency as the timing of the first spike during the stimulus period. We included only the neurons that met the following conditions: (i) background activity (as defined above) is less than 1; and (ii) shows statistically significant response to at least five trials in the random loom experiment. Condition (i) is required by our definition of latency. Condition (ii) is required because we define the variability of latency as the standard deviation of the timing of first spike, and this requires some number of samples to compute. 41 sSC and 128 dSC neurons that met condition (ii) but not (i) were discarded, and the final plot in Figure 2.3C shows 37 sSC and 70 dSC neurons. Finally, to avoid including spikes not due to visual stimulation, we required the first spike to not occur earlier than 30 ms since stimulus onset.

Stimulus-specific habituation

To analyze the stimulus-specific habituation (Figure 2.4B), we computed the habituation index defined as $1 - r_i/r_1$ where $r_i = r'_i - \mu$ refers to the number of spikes a neuron fired in the i -th repetition of the looming stimulus (r'_i) after subtracting the expected number of spikes due to baseline activity (μ). Analysis with $i = 4, 7$, and 10 did not yield significantly different results (Figure 2.4B uses $i = 10$). Only the neurons whose initial response to the looming stimuli met the significance criterion were included in the analysis.

Statistical test

Furthermore, we tested if the empirical distributions of sSC and dSC neurons differ significantly from each other in Figure 2.2C-D, Figure 2.3B-C, and Figure 2.4B. To do so we applied the two-sample Kolmogorov-Smirnov test using the MATLAB function `kstest2`. In all cases the computed p -values were less than the pre-set cutoff of 0.005 and were reported within the figure panels.

Recovery from habituation

To analyze the time to recover from the habituation (Figure 2.4D), a series of looming stimuli was presented at a single location with inter-stimulus intervals of 1.5, 2, 6, 11, 21, 61, and 121 s, in this order. The extent of recovery was defined as r_i/r_1 where $r_i = r'_i - \mu$ refers to the number of spikes a neuron fired in the i -th repetition of this series (r'_i) after subtracting expected number of spikes due to baseline activity (μ). This was done for simultaneously recorded sSC and dSC neurons that met the significance criterion. The 25th, 50th (median), and 75th percentiles were then computed separately for sSC and dSC neurons and plotted in Figure 2.4D.

Decoding analysis

We analyzed the population of neurons from superficial and deep SC to decode stimulus variables in the “random loom” experiment (Figure 2.5). Specifically, we asked if the population activity contains information about the location (i.e. in which of the 25 possible locations did the stimulus appear?) and novelty (i.e. is this the first stimulus to appear at a location?) of the stimuli.

To do so, we first pooled neurons from three recordings that used similar parameters of the “random loom” experiment. Because of retinotopy in the SC, superficial SC neurons recorded by a single shank of the silicon probe tend to have overlapping receptive fields. As a result, decoding stimulus location from the superficial SC neurons requires sampling them throughout the retinotopic map, which is difficult to do experimentally. Working on the assumption that different parts of the map contain equivalent neural representations, we augmented the data by generating virtual neurons whose response profiles were spatially shifted copies of actual neural responses. Specifically, each copy shifted the response profile to one of the 8 adjacent locations in a 3x3 grid with the original response profile in the center. The neurons were then divided into two groups (sSC and dSC) based on the depth of the channel with maximum waveform. This augmentation process increased the number of neurons used in this analysis from 106 (38 sSC and 68 dSC) to 963 (342 sSC and 621 dSC). Some neurons whose response profile after shifting lay outside the stimulus presentation area were discarded.

After this, the data consisted of neural response of the augmented sSC and dSC populations in each of the ~100 trials. In the case of the location decoder, the labels were multi-class and ranged from 1 to 25 (one for each stimulus location). In the case of the novelty decoder, the labels were binary (stimuli that were novel, i.e. the

first to appear at a location, were 1; others were 0). The performance measure was the mean four-fold cross validation score. The chance performance for the location decoder is the maximum of the number of times the stimulus appeared at each of the 25 locations, divided by the total number of presentations (i.e. $\max_i \{n_i / \sum_j n_j\}$, where n_i refers to the number of times the stimulus appeared at location i). In our data, this was roughly 10%. The chance performance for the novelty decoder is the number of non-novel presentations divided by the total number of presentations. Given that there were 25 possible locations and 100 trials, this was roughly 75%.

We then subsampled sets of 5, 10, 30, 70, 150, and 300 neurons from each of the two groups and used their responses to train the location and novelty decoders. This was done with the `LogisticRegression` class in the scikit-learn package [51] in Python using the following parameters: $penalty = 'l2'$, $C = 1.0$, $max_iter=5000$. This process was repeated with 100 random subsamples, and the mean and standard deviation of this ensemble were computed and plotted in Figure 2.5.

Model

In the circuit of Figure 2.6A, we modeled each input neuron as a linear-nonlinear (LN) element. The neuron's response was calculated as

$$r(t) = N(g(t)) \quad (2.1)$$

where

$$g(t) = s(x, y, t) * k(x, y, t) = \int_x \int_y \int_{t'=-\infty}^t s(x, y, t') k(x, y, t - t') dt' dy dx \quad (2.2)$$

is the convolution of the stimulus s with the spatio-temporal receptive field k . The receptive field $k(x, y, t)$ was parametrized as

$$k(x, y, t) = F(x, y)T(t) \quad (2.3)$$

$$F(x, y) = \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \quad (2.4)$$

$$T(t) = \left(\frac{t}{\tau_1}\right)^{n_1} \exp(-n_1(t/\tau_1 - 1)) - b \left(\frac{t}{\tau_2}\right)^{n_2} \exp(-n_2(t/\tau_2 - 1)) \quad (2.5)$$

The nonlinear transformation was a half-wave rectifier:

$$N(g) = \max(0, mg - \theta) \quad (2.6)$$

where θ is a threshold and m is a scaling factor. The firing rate of the local looming detector neuron (LD) was computed from the difference between the responses of

the center and surround neurons:

$$r_{LD}(t) = N(r_c(t) - r_s(t)) \quad (2.7)$$

and the response of the widefield neuron (WF) was computed from the various local detectors as

$$r_{WF}(t) = \sum_i w_i r_{LD,i}(t) \quad (2.8)$$

where w_i is the synaptic weight from local neuron i onto the widefield neuron. We modeled the habituation in the synapse between local detectors and the widefield neuron with a differential equation of three parameters for short-term synaptic depression and recovery:

$$\frac{d}{dt}w = \frac{1-w}{\tau} - a(w - w_{\min})r(t) \quad (2.9)$$

where τ is the time constant for synaptic recovery, a is a gain factor for depression, and w_{\min} is a floor on synaptic strength. The simulation in Figure 2.6F used $a = 1$ and $w_{\min} = 0$.

The temporal kernels used for the center and surround neurons feeding the local looming detector were taken from the measured receptive fields of mouse alpha retinal ganglion cells [36]. Table 2.1 lists the parameter values chosen. We arranged local looming detectors on a grid with 15° spacing between the centers of adjacent cells.

Table 2.1: Parameter values used for the model in Figure 2.6, as defined by Eqns 2.1-2.9.

<i>Receptive field (Eqns 2.3-2.5)</i>		
Parameter	Center	Surround
σ	4.00°	10.0°
τ_1	104 ms	84.6 ms
n_1	2.77	1.24
τ_2	91.2 ms	79.7 ms
n_2	3.94	1.87
b	1.34	1.33
<i>Nonlinearity (Eqn 2.6)</i>		
Parameter	Value	
m	1	
θ	0	
<i>Synaptic depression (Eqn 2.9)</i>		
Parameter	Value	
a	1	
w_{\min}	0	

References

- [1] Annette E. Allen, Christopher A. Procyk, Michael Howarth, Lauren Walm-
sley, and Timothy M. Brown. Visual input to the mouse lateral posterior
and posterior thalamic nuclei: Photoreceptive origins and retinotopic or-
der. *Journal of Physiology*, 594(7):1911–1929, 2016. ISSN 14697793. doi:
10.1113/JP271707.
- [2] Stephen A. Baccus and Markus Meister. Fast and slow contrast adaptation
in retinal circuitry. *Neuron*, 36(5):909–919, 2002. ISSN 08966273. doi:
10.1016/S0896-6273(02)01050-4.
- [3] Michele A. Basso and Paul J. May. Circuits for Action and Cognition: A
View from the Superior Colliculus. *Annual Review of Vision Science*, 3(1):
annurev-vision-102016-061234, October 2017. ISSN 2374-4642. doi: 10.
1146/annurev-vision-102016-061234. URL [http://www.annualreviews.
org/doi/10.1146/annurev-vision-102016-061234](http://www.annualreviews.org/doi/10.1146/annurev-vision-102016-061234).
- [4] David H. Brainard. The Psychophysics Toolbox. *Spatial Vision*, 1997. ISSN
01691015. doi: 10.1163/156856897X00357.
- [5] Haewon Byun, Soohyun Kwon, Hee-Jeong Ahn, Hong Liu, Douglas Forrest,
Jonathan B. Demb, and In-Jung Kim. Molecular features distinguish ten
neuronal types in the mouse superficial colliculus. *Journal of Comparative
Neurology*, 524(11):2300–2321, 2016. ISSN 10969861. doi: 10.1002/cne.
23952. URL <http://www.ncbi.nlm.nih.gov/pubmed/21858823>.
- [6] Jianhua Cang, Elise Savier, Jad Barchini, and Xiaorong Liu. Visual Function,
Organization, and Development of the Mouse Superior Colliculus. *Annual
Review of Vision Science*, 4(1):239–262, May 2018. ISSN 2374-4642, 2374-
4650. doi: 10.1146/annurev-vision-091517-034142. URL [https://www.
annualreviews.org/doi/10.1146/annurev-vision-091517-034142](https://www.annualreviews.org/doi/10.1146/annurev-vision-091517-034142).
- [7] C E Carr and M Konishi. A circuit for detection of interaural time differences in
the brain stem of the barn owl. *The Journal of Neuroscience*, 10(10):3227–46,
1990. ISSN 0270-6474. URL [http://www.ncbi.nlm.nih.gov/pubmed/
2213141](http://www.ncbi.nlm.nih.gov/pubmed/2213141).
- [8] E J Chichilnisky. A simple white noise analysis of neuronal light responses.
Network, 12(2):199–213, 2001. ISSN 0954-898X. doi: 10.1080/713663221.
URL [http://www.ncbi.nlm.nih.gov/pubmed/11405422%5Cnhttp://
informahealthcare.com/doi/abs/10.1080/net.12.2.199.213](http://www.ncbi.nlm.nih.gov/pubmed/11405422%5Cnhttp://informahealthcare.com/doi/abs/10.1080/net.12.2.199.213).
- [9] M Cynader and N Berman. Receptive-field organization of monkey superior
colliculus. *Journal of Neurophysiology*, 35(2):187–201, 1972. ISSN 0022-
3077.

- [10] Gioia De Franceschi, Tipok Vivattanasam, Aman B. Saleem, and Samuel G. Solomon. Vision Guides Selection of Freeze or Flight Defense Strategies in Mice. *Current Biology*, 26(16):2150–2154, 2016. ISSN 0960-9822. doi: 10.1016/j.cub.2016.06.006. URL <http://linkinghub.elsevier.com/retrieve/pii/S0960982216306182>.
- [11] Daniel de Malmazet, Norma K. Kühn, and Karl Farrow. Retinotopic Separation of Nasal and Temporal Motion Selectivity in the Mouse Superior Colliculus. *Current Biology*, 28(18):2961–2969.e4, 2018. ISSN 09609822. doi: 10.1016/j.cub.2018.07.001.
- [12] James J DiCarlo, Davide Zoccolan, and Nicole C Rust. How does the brain solve visual object recognition? *Neuron*, 73(3):415–434, 2012. ISSN 1097-4199. doi: 10.1016/j.neuron.2012.01.010. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=22325196<http://www.ncbi.nlm.nih.gov/pubmed/22325196><http://www.ncbi.nlm.nih.gov/pubmedcentral/articlerender.fcgi?artid=PMC3306444>.
- [13] U.C. Dräger and D.H. Hubel. Responses to visual stimulation and relationship between visual, auditory, and somatosensory inputs in mouse superior colliculus. *Journal of Neurophysiology*, 38(3):690–713, 1975. ISSN 0022-3077. URL <http://www.ncbi.nlm.nih.gov/pubmed/1127462>http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=1127462.
- [14] J Du, T J Blanche, R R Harrison, H A Lester, and S C Masmanidis. Multiplexed, high density electrophysiology with nanofabricated neural probes. *PLoS ONE*, 6(10):e26204, 2011. doi: 10.1371/journal.pone.0026204. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=22022568<http://www.ncbi.nlm.nih.gov/pubmed/22022568>.
- [15] Toshiaki Endo, Etsuko Tarusawa, Takuya Notomi, Katsuyuki Kaneda, Masumi Hirabayashi, Ryuichi Shigemoto, and Tadashi Isa. Dendritic Ih Ensures High-Fidelity Dendritic Spike Responses of Motion-Sensitive Neurons in Rat Superior Colliculus. *Journal of Neurophysiology*, 99(5):2066–2076, 2008. ISSN 0022-3077. doi: 10.1152/jn.00556.2007. URL <http://www.ncbi.nlm.nih.gov/pubmed/18216232>.
- [16] Dominic A. Evans, A. Vanessa Stempel, Ruben Vale, Sabine Rühle, Yaara Lefler, and Tiago Branco. A synaptic threshold mechanism for computing escape decisions. *Nature*, 558(7711):590–594, June 2018. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-018-0244-6. URL <http://www.nature.com/articles/s41586-018-0244-6>.
- [17] Evan H Feinberg and Markus Meister. Orientation columns in the mouse superior colliculus. *Nature*, 519(7542):229–232, 2015. doi: 10.1038/nature14103.

URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=25517100<http://dx.doi.org/10.1038/nature14103>.

- [18] W A Freiwald and D Y Tsao. Functional Compartmentalization and Viewpoint Generalization Within the Macaque Face-Processing System. *Science*, 330(November):845–851, 2010. ISSN 1095-9203. doi: 10.1126/science.1206034. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=21051642.
- [19] Samuel D Gale and Gabe J Murphy. Distinct Representation and Distribution of Visual Information by Specific Cell Types in Mouse Superficial Superior Colliculus. *Journal of Neuroscience*, 34(40):13458–13471, 2014. doi: 10.1523/jneurosci.2768-14.2014. URL <http://www.jneurosci.org/content/34/40/13458.abstract>.
- [20] Samuel D Gale and Gabe J Murphy. Active Dendritic Properties and Local Inhibitory Input Enable Selectivity for Object Motion in Mouse Superior Colliculus Neurons. *Journal of Neuroscience*, 36(35):9111–9123, 2016. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.0645-16.2016. URL <http://www.ncbi.nlm.nih.gov/pubmed/27581453>.
- [21] B Gordon. Receptive fields in deep layers of cat superior colliculus. *Journal of Neurophysiology*, 36(2):157–178, 1973. ISSN 0022-3077. doi: 10.1152/jn.1973.36.2.157.
- [22] W Heiligenberg. Coding and processing of electrosensory information in gymnotiform fish. *J Exp Biol*, 146:255–275, 1989. ISSN 0022-0949. URL <http://www.ncbi.nlm.nih.gov/pubmed/2689565>.
- [23] Matthew C Helms, Gülden Özen, and William C Hall. Organization of the Intermediate Gray Layer of the Superior Colliculus. I. Intrinsic Vertical Connections. *Journal of Neurophysiology*, 91(4):1706–1715, 2004. doi: 10.1152/jn.00705.2003. URL <http://jn.physiology.org/jn/91/4/1706.full.pdf>.
- [24] Gabriel Horn and R M Hill. Effect of removing the neocortex on the response to repeated sensory stimulation of neurones in the mid-brain. *Nature*, 211(5050):754–755, 1966. ISSN 0028-0836. doi: 10.1038/211754a0. URL <http://www.ncbi.nlm.nih.gov/pubmed/5962123>.
- [25] Gabriel Horn and Richard M. Hill. Responsiveness to sensory stimulation of units in the superior colliculus and subjacent tectotegmental regions of the rabbit. *Experimental Neurology*, 14:199–223, 1966. ISSN 1359-0286. doi: 10.1016/j.cossms.2010.07.001. URL <http://dx.doi.org/10.1016/j.cossms.2010.07.001>.

- [26] Jennifer L. Hoy, Iryna Yavorska, Michael Wehr, and Cristopher M. Niell. Vision Drives Accurate Approach Behavior during Prey Capture in Laboratory Mice. *Current Biology*, 26(22):3046–3052, 2016. ISSN 0960-9822. doi: 10.1016/j.cub.2016.09.009. URL <http://dx.doi.org/10.1016/j.cub.2016.09.009>.
- [27] N.K. Humphrey. Responses to visual stimuli of units in the superior colliculus of rats and monkeys. *Experimental Neurology*, 20(3):312–340, 1968. ISSN 0014-4886. doi: 10.1016/0014-4886(68)90076-9.
- [28] Samsoun Inayat, Jad Barchini, Hui Chen, Liang Feng, Xiaorong Liu, and Jianhua Cang. Neurons in the Most Superficial Lamina of the Mouse Superior Colliculus Are Highly Selective for Stimulus Direction. *Journal of Neuroscience*, 35(20):7992–8003, 2015. doi: 10.1523/jneurosci.0173-15.2015. URL <http://www.jneurosci.org/content/35/20/7992.abstract>.
- [29] Shinya Ito, David A. Feldheim, and Alan M. Litke. Segregation of visual response properties in the mouse superior colliculus and their modulation during locomotion. *Journal of Neuroscience*, 37(35):3689–16, June 2017. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3689-16.2017. URL <http://www.jneurosci.org/content/37/35/8428><http://www.jneurosci.org/content/37/35/8428.full.pdf><http://www.ncbi.nlm.nih.gov/pubmed/28760858><http://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.3689-16.2017>.
- [30] Maximilian Joesch and Markus Meister. A neuronal circuit for colour vision based on rod–cone opponency. *Nature*, pages 1–15, 2016. ISSN 0028-0836. doi: 10.1038/nature17158. URL <http://www.nature.com/doifinder/10.1038/nature17158>.
- [31] S Kim, M K Lehtinen, A Sessa, M W Zappaterra, S H Cho, D Gonzalez, B Boggan, C A Austin, J Wijnholds, M J Gambello, J Malicki, A S LaMantia, V Broccoli, and C A Walsh. The apical complex couples cell fate and cell survival to cerebral cortical development. *Neuron*, 66:69–84, 2010. doi: 10.1016/j.neuron.2010.03.019. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=20399730.
- [32] M Kleiner, D Brainard, Denis Pelli, A Ingling, R Murray, and C Broussard. What’s new in psychtoolbox-3. *Perception*, 36, 2007.
- [33] Masakazu Konishi. Coding of auditory space. *Annu Rev Neurosci*, 26(1):31–55, 2003. ISSN 0147-006X. doi: 10.1146/annurev.neuro.26.041002.131123. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=14527266<http://www.annualreviews.org/doi/10.1146/annurev.neuro.26.041002.131123>.

- [34] Eileen Kowler. Eye movements: The past 25 years. *Vision Research*, 51(13): 1457–1483, 2011. ISSN 0042-6989. doi: 10.1016/j.visres.2010.12.014.
- [35] Richard J Krauzlis, Lee P Lovejoy, and Alexandre Zénon. Superior Colliculus and Visual Spatial Attention. *Annual Review of Neuroscience*, 36(1):165–182, 2013. doi: doi:10.1146/annurev-neuro-062012-170249. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-neuro-062012-170249>.
- [36] Brenna Krieger, Mu Qiao, David L. Rousso, Joshua R. Sanes, and Markus Meister. Four alpha ganglion cell types in mouse retina: Function, structure, and molecular signatures. *PLoS ONE*, 12(7):1–21, 2017. ISSN 19326203. doi: 10.1371/journal.pone.0180091.
- [37] T. P. Langer and R. D. Lund. The upper layers of the superior colliculus of the rat: a Golgi study. *The Journal of Comparative Neurology*, 158(4): 418–435, 1974. ISSN 0021-9967. doi: 10.1002/cne.901580404. URL <http://www.ncbi.nlm.nih.gov/pubmed/4615112>.
- [38] Psyche H Lee, Matthew C Helms, George J Augustine, and William C Hall. Role of intrinsic synaptic circuitry in collicular sensorimotor integration. *Proceedings of the National Academy of Sciences*, 94(24):13299–13304, 1997. URL <http://www.pnas.org/content/94/24/13299.abstract>.
- [39] J.Y. Lettvin, H.R. Maturana, W.H. Pitts, and W.S. McCulloch. Two remarks on the visual system of the frog. In Walter A Rosenblith, editor, *Sensory communication: contributions to the Symposium on Principles of Sensory Communication*, pages 757–776. MIT Press, 1961.
- [40] H Luksch, K Cox, and H J Karten. Bottlebrush dendritic endings and large dendritic fields: motion-detecting neurons in the tectofugal pathway. *The Journal of Comparative Neurology*, 396(3):399–414, 1998. ISSN 0021-9967. URL <http://www.ncbi.nlm.nih.gov/pubmed/9624592>.
- [41] Harald Luksch, Reza Khanbabaie, and Ralf Wessel. Synaptic dynamics mediate sensitivity to motion independent of stimulus details. *Nature Neuroscience*, 7(4):380–388, 2004. ISSN 1097-6256. doi: 10.1038/nn1204. URL <http://www.ncbi.nlm.nih.gov/pubmed/14990932>.
- [42] D E Major, H Luksch, and H J Karten. Bottlebrush dendritic endings and large dendritic fields: motion-detecting neurons in the mammalian tectum. *J Comp Neurol*, 423:243–260, 2000. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=10867657.
- [43] R H Masland, K L Chow, and D L Stewart. Receptive-field characteristics of superior colliculus neurons in the rabbit. *Journal of Neurophysiology*, 34(1):148–156, 1971. ISSN 0022-3077. doi: 10.1152/jn.1971.34.1.148. URL <http://www.ncbi.nlm.nih.gov/pubmed/5540576>.

- [44] Paul J May. The mammalian superior colliculus: laminar structure and connections. *Progress in Brain Research*, 151:321–378, 2006. doi: [http://dx.doi.org/10.1016/S0079-6123\(05\)51011-2](http://dx.doi.org/10.1016/S0079-6123(05)51011-2). URL <http://www.sciencedirect.com/science/article/pii/S0079612305510112>.
- [45] Markus Meister, Jerome Pine, and Denis A. Baylor. Multi-neuronal signals from the retina: acquisition and analysis. *Journal of Neuroscience Methods*, 51(1):95–106, 1994.
- [46] Thomas A Münch, Rava Azeredo da Silveira, Sandra Siegert, Tim James Viney, Gautam B Awatramani, and Botond Roska. Approach sensitivity in the retina processed by a multifunctional neural circuit. *Nature Neuroscience*, 12(10):1308–1316, 2009. doi: http://www.nature.com/neuro/journal/v12/n10/supinfo/nn.2389_S1.html. URL <http://dx.doi.org/10.1038/nn.2389>.
- [47] Cristopher M Niell and Michael P Stryker. Highly selective receptive fields in mouse visual cortex. *J Neurosci*, 28(30):7520–7536, 2008. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.0623-08.2008. URL <http://www.ncbi.nlm.nih.gov/pubmed/18650330><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3040721&tool=pmcentrez&rendertype=abstract>.
- [48] B A Olshausen, C H Anderson, and D C Van Essen. A neurobiological model of visual-attention and invariant pattern-recognition based on dynamic routing of information. *Journal of Neuroscience*, 13(11):4700–4719, 1993. ISSN 0270-6474.
- [49] Marius Pachitariu, Nicholas A. Steinmetz, Shabnam N. Kadir, Matteo Carandini, and Kenneth D. Harris. Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems*, 29(Nips):4448–4456, 2016. ISSN 10495258. URL <https://papers.nips.cc/paper/6326-fast-and-accurate-spike-sorting-of-high-channel-count-probes-with-kilosort.pdf>.
- [50] George Paxinos and Keith B.J. Franklin. *The Mouse Brain in Stereotaxic Coordinates*. Academic Press, 2nd edition, 2001.
- [51] Fabian Pedregosa, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Vincent Dubourg, Fabian Pedregosa, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Fabian Pedregosa, and Ron Weiss. Scikit-learn : Machine Learning in Python To cite this version : Scikit-learn : Machine Learning in Python. *Journal of Machine Learning Research*, 2011. ISSN 0021-9355. doi: <https://dl.acm.org/citation.cfm?id=2078195>.
- [52] Xaq Pitkow and Markus Meister. Neural computation in sensory systems. In M S Gazzaniga and G R Mangun, editors, *The Cognitive Neurosciences*, pages 305–318. MIT Press, Cambridge, MA, 5th edition, 2014.

- [53] R. Quian Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 435 (7045):1102–1107, 2005. ISSN 0028-0836. doi: 10.1038/nature03687. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15973409<http://www.nature.com/doifinder/10.1038/nature03687>.
- [54] Amit Reches and Yoram Gutfreund. Stimulus-specific adaptations in the gaze control system of the barn owl. *Journal of Neuroscience*, 28(6):1523–1533, 2008. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3785-07.2008. URL <https://www.jneurosci.org/content/28/6/1523>.
- [55] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11):1019–25, 1999. ISSN 1097-6256. doi: 10.1038/14819. URL <http://www.ncbi.nlm.nih.gov/pubmed/10526343>.
- [56] Botond Roska and Markus Meister. The retina dissects the visual scene into distinct features. In *The New Visual Neurosciences*, pages 163–182. MIT Press, Cambridge, MA, 2014.
- [57] Cyrille Rossant. Phy: interactive visualization and manual spike sorting of large-scale ephys data, April 2017. URL <http://phy.cortexlab.net>.
- [58] Cyrille Rossant, Shabnam N Kadir, Dan F M Goodman, John Schulman, Maximilian L D Hunter, Aman B Saleem, Andres Grosmark, Mariano Beluscio, George H Denfield, Alexander S Ecker, Andreas S Tolias, Samuel Solomon, György Buzsáki, Matteo Carandini, and Kenneth D Harris. Spike sorting for large, dense electrode arrays. *Nature Neuroscience*, 19(4):634–641, April 2016. ISSN 1097-6256, 1546-1726. doi: 10.1038/nn.4268. URL <http://www.nature.com/articles/nn.4268>.
- [59] David L Rousso, Masahito Yamataga, and Joshua R. Sanes. Organization of neuronal subtypes and retinal inputs in the mouse superior colliculus. In *Society for Neuroscience*, 2014.
- [60] Joshua R. Sanes and Richard H. Masland. The types of retinal ganglion cells: current status and implications for neuronal classification. *Annual Review of Neuroscience*, 38(1):221–246, 2015. ISSN 0147-006X. doi: 10.1146/annurev-neuro-071714-034120. URL <http://www.ncbi.nlm.nih.gov/pubmed/25897874>.
- [61] Michael Anthony Savage, Richard McQuade, and Alexander Thiele. Segregated fronto-cortical and midbrain connections in the mouse and their relation to approach and avoidance orienting behaviors. *Journal of Comparative Neurology*, 525(8):1980–1999, 2017. ISSN 10969861. doi: 10.1002/cne.24186.
- [62] Tania A. Seabrook, Timothy J. Burbridge, Michael C. Crair, and Andrew D. Huberman. Architecture, Function, and Assembly of the Mouse Visual

- System. *Annual Review of Neuroscience*, 40(1):499–538, December 2017. ISSN 0147-006X. doi: 10.1146/annurev-neuro-071714-033842. URL <http://www.annualreviews.org/doi/10.1146/annurev-neuro-071714-033842><http://www.annualreviews.org/doi/pdf/10.1146/annurev-neuro-071714-033842>.
- [63] T. Serre, A. Oliva, and T. Poggio. A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci USA*, 104(15): 6424–6429, 2007. ISSN 0027-8424. doi: 10.1073/pnas.0700622104. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=17404214<http://www.pnas.org/cgi/doi/10.1073/pnas.0700622104>.
- [64] Congping Shang, Zijun Chen, Aixue Liu, Yang Li, Jiajing Zhang, Baole Qu, Fei Yan, Yaning Zhang, Weixiu Liu, Zhihui Liu, Xiaofei Guo, Dapeng Li, Yi Wang, and Peng Cao. Divergent midbrain circuits orchestrate escape and freezing responses to looming stimuli in mice. *Nature Communications*, 9(1):1232, 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-03580-7. URL <http://dx.doi.org/10.1038/s41467-018-03580-7>.
- [65] Terrence R Stanford, Swetha Shankar, Dino P Massoglia, M Gabriela Costello, and Emilio Salinas. Perceptual decision making in less than 30 milliseconds. *Nature Neuroscience*, 13(3):379–385, June 2010. ISSN 1546-1726. doi: 10.1038/nn.2485. URL <https://www.nature.com/articles/nn.2485>.
- [66] Iain Stitt, Edgar Galindo-Leon, Florian Pieper, Gerhard Engler, and Andreas K. Engel. Laminar profile of visual response properties in ferret superior colliculus. *Journal of Neurophysiology*, 110(6):1333–45, 2013. ISSN 1522-1598. doi: 10.1152/jn.00957.2012. URL <http://www.ncbi.nlm.nih.gov/pubmed/23803328>.
- [67] M. Straschill and K.P. Hoffmann. Functional aspects of localization in the cat’s tectum opticum. *Brain Research*, 13:274–283, 1968.
- [68] S Thorpe, D Fize, and C Marlot. Speed of processing in the human visual system. *Nature*, 381:520–522, 1996. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=8632824.
- [69] Misha Tsodyks, Klaus Pawelzik, and Henry Markram. Neural Networks with Dynamic Synapses. *Neural Computation*, 10:821–835, 1998.
- [70] Werner Ulbricht. Sodium Channel Inactivation: Molecular Determinants and Modulation. *Physiological Reviews*, 85(4):1271–1301, 2005. ISSN 0031-9333. doi: 10.1152/physrev.00024.2004.
- [71] S Ullman and S Soloviev. Computation of pattern invariance in brain-like structures. *Neural Networks*, 12(7-8):1021–1036, 1999. ISSN 1879-2782. URL <http://www.ncbi.nlm.nih.gov/pubmed/12662643>.

- [72] Barry Wark, Adrienne Fairhall, and Fred Rieke. Timescales of Inference in Visual Adaptation. *Neuron*, 61(5):750–761, 2009. ISSN 08966273. doi: 10.1016/j.neuron.2009.01.019. URL <http://dx.doi.org/10.1016/j.neuron.2009.01.019>.
- [73] Jonathan Wolfe and Larry A. Palmer. Temporal diversity in the lateral geniculate nucleus of cat. *Visual Neuroscience*, 15(4):653–675, April 1998. ISSN 0952-5238, 1469-8714. doi: 10.1017/S0952523898154068. URL https://www.cambridge.org/core/product/identifier/S0952523898154068/type/journal_article.
- [74] E.J. Woods and B.J. Frost. Adaptation and habituation characteristics of tectal neurons in the pigeon. *Experimental Brain Research*, 27(3-4):347–354, 1977. ISSN 0014-4819. URL <http://www.ncbi.nlm.nih.gov/pubmed/880990>.
- [75] Melis Yilmaz and Markus Meister. Rapid innate defensive responses of mice to looming visual stimuli. *Current Biology*, 23(20):2011–2015, 2013. ISSN 1879-0445 (Electronic) 0960-9822 (Linking). doi: 10.1016/j.cub.2013.08.015. URL <http://www.ncbi.nlm.nih.gov/pubmed/24120636>.
- [76] Yifeng Zhang, In-Jung Kim, Joshua R Sanes, and Markus Meister. The most numerous ganglion cell type of the mouse retina is a selective feature detector. *Proceedings of the National Academy of Sciences*, 109(36):E2391–E2398, 2012. doi: 10.1073/pnas.1211547109. URL <http://www.pnas.org/content/109/36/E2391.abstract>.
- [77] Xinyu Zhao, Mingna Liu, and Jianhua Cang. Visual Cortex Modulates the Magnitude but Not the Selectivity of Looming-Evoked Responses in the Superior Colliculus of Awake Mice. *Neuron*, 84(1):202–213, 2014. doi: 10.1016/j.neuron.2014.08.037. URL <http://www.sciencedirect.com/science/article/pii/S0896627314007387>http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=25220812.

Chapter 3

ELECTRODE POOLING

Kyu Hyun Lee, Yu-Li Ni, and Markus Meister. Electrode pooling: How to boost the yield of switchable silicon probes for neuronal recordings. *bioRxiv*, 2019. doi: 10.1101/851691.

K.H.L. participated in the conception of the project, performed the experiments, analyzed the data, and wrote the manuscript.

3.1 Abstract

State-of-the-art silicon probes for electrical recording from neurons have thousands of recording sites, but only a fraction of them can be used simultaneously due to the forbiddingly large volume of the associated wires. To overcome this fundamental constraint, we propose a novel method called *electrode pooling* that uses a single wire to serve many recording sites. Multiple electrodes are connected to a single wire through a set of controllable switches. Here we present the framework behind this method and an experimental strategy to support it. We then demonstrate its feasibility by implementing electrode pooling on the Neuropixels electrode array and characterizing its effect on signal and noise. Finally we use simulations to explore the conditions under which electrode pooling saves wires without compromising the content of the recordings. We make recommendations on the design of future devices to take advantage of this strategy.

3.2 Introduction

Understanding brain function requires monitoring the complex pattern of activity distributed across many neuronal circuits. To this end, the BRAIN Initiative has called for the development of technologies for recording “dynamic neuronal activity from complete neural networks, over long periods, in all areas of the brain,” ideally “monitoring all neurons in a circuit” [5]. Recent advances in the design and manufacturing of silicon-based neural probes have answered this challenge with new devices that have thousands of recording sites [6, 12, 23, 31]. But in many such devices, only a small fraction of the recording sites can be used at once. The reason is that neural signals must be brought out of the brain via wires, which take up much more volume than the recording sites themselves. For example, in one

state-of-the-art silicon shank, each wire displaces thirty times more volume than a recording site once the shank is fully inserted in the brain [12]. The current silicon arrays invariably displace or destroy more neurons than they record, and thus the goal of "monitoring all neurons" seems unattainable by simply scaling the present approach ¹. Clearly we need a way to increase the number of neurons recorded while avoiding a concomitant increase in the number of wires that enter the brain.

Time-division multiplexing

A common approach by which a single wire can convey multiple analog signals is "time-division multiplexing" [20]. A rapid switch cycles through the N input signals and connects each input to the output line for a brief interval (Figure 3.1A). At the other end of the line, a synchronized switch or sampling system can demultiplex the N signals again. In this way, a single wire carries signals from all its associated electrodes interleaved in time. The cycling rate of the switch is constrained by the sampling theorem [29]: it should be at least twice the highest frequency component present in the signal. The raw voltage signals from extracellular electrodes include thermal noise signals that extend far into the Megahertz regime. Therefore an essential element of any such multiplexing scheme is an analog low-pass filter associated with each electrode. This anti-alias filter removes the high-frequency noise above a certain cut-off frequency. In practice, the cut-off is chosen to match the bandwidth of neuronal action potentials, typically 10 kHz. Then the multiplexer switch can safely cycle at a few times that cut-off frequency.

Given the ubiquity of time-division multiplexing in communication electronics, what prevents its use for neural recording devices? One obstacle is the physical size of the anti-alias filter associated with each electrode. When implemented in CMOS technology, such a low-pass filter occupies an area much larger than the recording site itself [28], which would force the electrodes apart, and thus prevent any high-density recording ². What if one simply omitted the low-pass filter? In that case, aliasing of high-frequency thermal fluctuations will increase the noise power in the recording by a factor equal to the number of electrodes N being multiplexed. One such device with a multiplexing factor of $N = 128$ has indeed proven unsuitable for recording action potentials, as the noise drowns out any signal [8]. A recent design with a more modest $N = 8$ still produces noise power 4-15 times higher than in

¹But see [13] for a wildly optimistic proposal.

²One recent report claims to implement a 4 kHz low-pass filter with an electrode spacing of just 28 μm , but the underlying circuit has never been revealed [1]

comparable systems without multiplexing [22].

Other issues further limit the use of time-division multiplexing: the requirement for amplification, filtering, and rapid switching right next to the recording site means that electric power gets dissipated on location, which leads to local heating of exactly the neurons one wants to monitor. Furthermore, the active electronics in the local amplifier are sensitive to light. This can produce artifacts when combined with bright light flashes for optogenetic stimulation [12, 14].

Static selection

An alternative approach involves static electrode selection (Figure 3.1B). Again, there is an electronic switch that connects the wire to one of many electrodes. However the switch setting remains unchanged during the electrical recording. In this way, the low-pass filtering and amplification can occur at the other end of the wire, outside the brain, where the base of the shank expands to offer virtually unlimited silicon space. The switch itself requires only minimal circuitry that fits comfortably under each recording site, even at a pitch of 20 μm or less. Because there is no local amplification or dynamic switching, the issues of heat dissipation or photosensitivity do not arise. This method has been incorporated recently into flat electrode arrays [19] as well as silicon prongs [12, 17]. It allows the user to choose one of many electrodes intelligently, for example because it carries a strong signal from a neuron of interest. However, it does not serve to increase the number of neurons per wire.

Electrode pooling

On this background, we introduce a third method of mapping electrodes to wires: selecting multiple electrodes with suitable signals and connecting them to the same wire (Figure 3.1C). Instead of rapidly cycling the intervening switches, as in multiplexing, simply leave all those switches closed. This creates a "pool" of electrodes whose signals are summed and transmitted on the same wire. At first, this approach seems counterproductive, as it mixes together recordings that one would like to analyze separately. How can one ever reconstruct which neural signal came from which electrode? Existing multi-electrode systems avoid this signal mixing at all costs, often quoting the low cross-talk between channels as a figure of merit. Instead, we will show that the pooled signal can be unmixable if one chooses the switch settings carefully during the recording session. Under suitable conditions, this method can record many neurons per wire without appreciable loss of information.

3.3 Results

Motivation for electrode pooling: Spike trains are sparse in time

A typical neuron may fire ~ 10 spikes/s on average [2]. Each action potential lasts for ~ 1 ms. Therefore this neuron's signal occupies less than 1% of the time axis in an extracellular recording (e.g., see Figure 3.5). Sometimes a second neuron lies close enough to the same electrode to produce a large spike. That still leaves 98% of the time axis to transmit the action potentials of other neurons. Electrode pooling gives the experimenter the freedom to add more neurons to that signal by choosing other electrodes that carry large spikes. Eventually a limit will be reached when the spikes of different neurons collide and overlap in time so they can no longer be distinguished. Fundamentally electrode pooling works because extracellular recordings are naturally sparse in time. The method can be seen as a variant of compressed sensing, which similarly relies on sparsity and randomly mixes multiple dimensions of sparse data into a common signal [9]. However, one benefits greatly from setting the switches intelligently, rather than leaving them up to random chance.

The effects of pooling on spikes and noise

What signal actually results when one connects two electrodes to the same wire? Figure 3.2 shows the relevant circuit for a hypothetical electrode array that allows electrode pooling. Here the common wire is connected via programmable switches to two recording sites. At each site i , the extracellular signal of nearby neurons reaches the shared wire through a total impedance R_i . This electrode impedance has contributions from the CMOS switch, the metal/bath interface, and the external electrolyte bath [24, 27]. The CMOS switches themselves have low impedance, typically $R_{\text{swi}} \approx 100 \Omega$ [27]. The metal/bath interface is typically modeled as an RC element with a total impedance on the order of $R_{\text{int}} \approx 100 \text{ k}\Omega$ [24, 27]. Finally the external bath impedance is typically a smaller term, on the order of tens of $\text{k}\Omega$ for an electrode of size $\sim 10 \mu\text{m}$. This circuit acts as a simple voltage divider. If a total of M electrodes are connected to the shared wire, the output voltage (V_P) is the average of the signals at the recording sites, weighted inversely by the electrode impedances,

$$V_P = \sum_{i=1}^M c_i V_i \quad (3.1)$$

where

$$c_i = \frac{1/R_i}{\sum_{j=1}^M 1/R_j} \quad (3.2)$$

is defined as the pooling coefficient for electrode i . If all electrodes have the same size and surface coating, they will have similar impedance, and in that limit one expects the simple relationship

$$V_P = \frac{1}{M} \sum_{i=1}^M V_i. \quad (3.3)$$

Thus an action potential that appears on only one of the M electrodes will be attenuated in the pooled signal by a factor $\frac{1}{M}$.

In order to understand the trade-offs of this method, we must similarly account for the pooling of noise (Figure 3.2). There are three relevant sources of noise: (1) thermal ("Johnson") noise from the impedance of the electrode; (2) biological noise ("hash") from many distant neurons whose signals are too small to be resolved; and (3) electronic noise resulting from the downstream acquisition system, including amplifier, multiplexer, and analog-to-digital converter. The thermal noise is private to each electrode, in the sense that it is statistically independent of the noise at another electrode. The biological noise is similar across neighboring electrodes that observe the same distant populations [10]. For widely separated electrodes the hash will be independent and thus private to each electrode, although details depend on the neuronal geometries and the degree of synchronization of distant neurons [25]. Finally the noise introduced by the amplifier and data acquisition is common to all the electrodes that share the same wire.

For any given electrode i the thermal noise and biological noise are additive and statistically independent, so they sum in quadrature to a total private noise with RMS amplitude

$$N_{\text{pri},i} = \sqrt{N_{\text{el},i}^2 + N_{\text{bio},i}^2}. \quad (3.4)$$

In the course of pooling, this signal gets attenuated by the pooling coefficient c_i (Eqn 3.2). Then it gets added to the common noise from data acquisition, which

again is statistically independent of the other noise sources. Thus the total noise at the output has RMS amplitude

$$N_{\text{tot}} = \sqrt{N_{\text{amp}}^2 + \sum_{i=1}^M c_i^2 N_{\text{pri},i}^2}. \quad (3.5)$$

If all electrodes have similar noise properties and impedances, this simplifies to

$$N_{\text{tot}} = \sqrt{N_{\text{amp}}^2 + N_{\text{pri}}^2/M}. \quad (3.6)$$

Theoretical limits of pooling

Now we are in a position to estimate the maximal possible benefits from electrode pooling. Suppose that using conventional split recording (1 electrode per wire, $M = 1$), we find a range of spikes on the electrode array: from the largest, with spike amplitude S_{max} , to the smallest that can still be sorted reliably from the noise, with amplitude S_{min} . To create an effective pool, we find the electrodes with the largest spikes and add them to the pool. Eventually, the spikes in the pooled signal will be so attenuated that they are no longer sortable from the noise. Pooling is beneficial as long as the signal-to-noise ratio of spikes in the pooled signal is larger than that of the smallest sortable spikes in a split-mode recording, namely

$$\frac{S_{\text{max}}/M}{\sqrt{N_{\text{com}}^2 + N_{\text{pri}}^2/M}} > \frac{S_{\text{min}}}{\sqrt{N_{\text{com}}^2 + N_{\text{pri}}^2}}. \quad (3.7)$$

This leads to a limit on the pool size M ,

$$M < M_{\text{max}} = \sqrt{\left(\frac{\beta^2}{2}\right)^2 + (1 + \beta^2)\alpha^2} - \frac{\beta^2}{2} \quad (3.8)$$

where

$$\alpha = S_{\text{max}}/S_{\text{min}}, \quad \beta = N_{\text{pri}}/N_{\text{com}} \quad (3.9)$$

So the maximal beneficial pool size depends on two parameters: the ratio of private to common noise, and the ratio of largest to smallest useful spike amplitudes. These parameters vary across applications, because they depend on the target brain area, the

recording hardware, and the spike-sorting software. In general, users can estimate the parameters α and β from their own experience with conventional recordings, and find M_{\max} from the lookup table in Figure 3.3.

Note these calculations focus on the discrimination of spikes from noise, not on distinguishing different spikes. As more neurons join the pool, two different cells may produce spikes with similar waveforms that are difficult to separate. Furthermore, as the combined firing rate of all neurons increases, action potentials start to overlap in time. These issues are difficult to capture analytically but will be explored in simulation (Figure 3.9). For now, the above expression for M_{\max} (Eqn 3.8) should be considered an upper bound on the optimal pool size.

Acquisition and analysis of pooled recordings

As mentioned above, a key requirement for this method to work is that the experimenter can control the switches that map electrodes to wires. This map should be adjusted to the unpredictable contingencies of any particular neural recording session. In fact the experimenter will benefit from using different switch settings during the same session.

In Figure 3.4 we propose an overall workflow for experiments using electrode pooling. A recording session begins with a short period of acquisition in "split mode" with only one electrode per wire. The purpose is to acquire samples of the spike waveforms from all neurons that might be recorded by the entire array. If the device has E electrodes and W wires, this sampling stage will require E/W segments of recording to cover all electrodes. For each segment the switches are reset to select a different batch of electrodes. Each batch should cover a local group of electrodes, ensuring that the entire "footprint" of each neuron can be captured.

After this sampling stage the experimenter performs a quick analysis to extract the relevant data that will inform the pooling process. In particular this yields a catalog of single neurons that can be extracted by spike-sorting. For each of those neurons one has the spike waveform observed on each electrode. Finally, for every electrode, one gets the private noise N_{pri} . The common noise N_{com} can be assessed ahead of time, because it is a property of the recording system. Thus the experimenter now has all the information needed to form useful electrode pools. Below are some principles one should consider in this process:

1. Pool electrodes that carry large signals.

2. Pool electrodes with different spike waveforms.
3. Don't pool neighboring electrodes that share the same hash noise.
4. Don't pool electrodes that carry dense signals with high firing rates.

After allocating the available wires to effective electrode pools one begins the main recording session in pooled mode. Ideally this phase captures all neurons with spike signals that are within reach of the electrode array.

In analyzing these recordings the goal is to detect spikes in the pooled signals and assign each spike correctly to its electrode of origin. This can be achieved by using the split-mode recordings from the early sampling stage of the experiment. From the spike waveforms obtained in split-mode one can predict how the corresponding spike appears in the pooled signals. Here it helps to know all the electrode impedances R_i so the weighted mix can be computed accurately (Eqn 3.1). This prediction serves as a search template for spike-sorting the pooled recording.

By its very nature electrode pooling produces a dense neural signal with more instances of temporal overlap between spikes than the typical split-mode recording. This places special demands on the methods for spike detection and sorting. The conventional cluster-based algorithm (peak detection - temporal alignment - PCA - clustering) does not handle overlapping spikes well [16]. It assumes that the voltage signal is sparsely populated with rare events drawn from a small number of discrete waveforms. Two spikes that overlap in time produce an unusual waveform that cannot be categorized. Recently some methods have been developed that do not force these assumptions [21, 32]. They explicitly model the recorded signal as an additive superposition of spikes and noise. The algorithm finds an efficient model that explains the signal by estimating both the spike waveform of each neuron and its associated set of spike times. These methods are well suited to the analysis of pooled recordings.

Because the spike templates are obtained from split-mode recordings at the beginning of the experiment, they are less affected by noise than if one had to identify them *de novo* from the pooled recordings. Nonetheless it probably pays to monitor the development of spike shapes during the pooled mode recording. If they drift too much, for example because the electrode array moves in the brain [11], then a recalibration by another split-mode session may be in order (Figure 3.4).

Electrode pooling with the Neuropixels probe

To test the biophysical assumptions underlying electrode pooling *in vivo*, we used the Neuropixels probe version 3B1 [12, 17]. This electrode array has a single silicon shank with 960 recording sites that can be connected to 384 wires via controllable switches (Figure 3.5). The electrodes are divided into three banks (called Bank 0, Bank 1, and Bank 2 from the tip to the base of the shank). Banks 0 and 1 contain 384 recording sites, while Bank 2 contains only 192. Each electrode has a dedicated switch by which it may connect to an adjacent wire. Each wire is assigned to either 2 or 3 electrodes at the same relative location within a bank, for example the bottom left recording sites in each bank share one wire. Under conventional operation of Neuropixels [12] each wire connects to only one electrode at a time. However, with modifications of the firmware on the device and the user interface we engineered independent control of all the switches. This made electrode pooling possible. Specifically the first 192 wires can pool up to 3 electrodes (one in each of the banks) and the remaining wires can pool up to two electrodes.

Given this geometry of the Neuropixels probe, one needs to implant more than one bank into the brain in order to accomplish any electrode pooling. In the pilot experiment analyzed here, the probe was inserted into the brain of a head-fixed, awake mouse to a depth of approximately 6 mm. This corresponds to all of Bank 0 and roughly half of Bank 1, and covered numerous brain areas from the medial preoptic area at the bottom to retrosplenial cortex at the top. Following the work flow proposed in Figure 3.4, we then recorded for ~10 min each from Bank 0 and Bank 1 in split mode (without pooling), followed by ~10 min of recording from both banks simultaneously in pooled mode.

Unmixing the pooled signal

Each of the three recordings (split Bank 0, split Bank 1, and pooled Banks 0 + 1) was then processed by spike-sorting to isolate single units. Then we paired each pooled-mode unit with the split-mode unit that had the most similar waveform, based on the cosine similarity of their waveform vectors (Figure 3.6B-C). In most cases the match was unambiguous, even when multiple units were present in the two banks with similar electrode footprints (Figure 3.6A). Generally the second-best match of a waveform was considerably worse. The matching algorithm proceeded iteratively until the similarity score for the best match dropped below 0.9 (Figure 3.6C). All remaining units in the pooled recording were omitted from further analysis. We corroborated the resulting matches by comparing other statistics of the identified

units, such as the mean firing rate and inter-spike-interval distribution (Figure 3.6A).

When spike-sorting the pooled-mode recording there is of course a strong expectation for what the spike waveforms will be, namely some combination of spikes from the two split-mode recordings. This suggests that one might jump-start the sorting of the pooled signal by building in the prior knowledge from sorting the split-mode recordings. Any such regularization could be beneficial, not only to accelerate the process but to compensate for the lower SNR in the pooled signal. We explored this possibility by running the template-matching function of KiloSort2 on the pooled-mode recording with templates from split-mode recordings ("hot sorting"). Then we compared the outcomes of three procedures (Figure 3.6D): (1) sorting each recording separately, using KiloSort1 with manual curation ("manual"); (2) sorting each recording separately using KiloSort2 with no manual intervention ("cold sorting"); and (3) cold sorting split-mode recordings first and then using their templates to sort the pooled-mode recordings ("hot sorting"). The new "hot sorting" procedure greatly outperformed the other methods, recovering $\sim 60\%$ more neurons from the pooled recordings without manual intervention. The above sequence of operations demonstrates that the pooled-mode recording can be successfully unmixed and the resulting units assigned to their locations along the multi-electrode shank.

Characterization of signal pooling

While the present Neuropixels device is rather limited in its options for electrode pooling, it does allow a detailed characterization of what happens to signals and noise, and a cross-check with the expectations from the simple circuit diagram in Figure 3.2. In that circuit the impedances of the two pooled electrodes act like a voltage divider. The pooled signal is a weighted average of the signals at the two sites, with the pooling coefficients determined by the relative impedances (Eqn 3.2) and summing to 1. For example, if two recording sites have the same impedance, their weights in the pooled signal should each be 0.5.

One can assess those pooling coefficients from *in vivo* recordings by measuring how much an action potential gets attenuated from split mode to pooled mode (Figure 3.7). For all the recording sites observed this way, the pooling coefficients fluctuated around 0.5. Furthermore, the weights of corresponding electrodes that contribute to the same pool summed to 1 to good approximation (Figure 3.7B). This is consistent with the description of pooling as a two-resistor voltage divider (Figure 3.2). For example, if the common wire in Figure 3.2 had introduced a significant

capacitive shunt, because it runs the full length of the silicon prong, then the pooling coefficients would sum to more than 1.

Interestingly, the pooling coefficients were spread over a wide range and they were systematically larger in bank 0 than in bank 1 (Figure 3.7C). The electrode impedance in these devices has been reported as $149 \pm 6 \text{ k}\Omega$ when measured in saline [12]. Based on this 4% standard deviation in impedance across electrodes one expects the pooling coefficients to vary with a standard deviation of only 0.03, whereas the observed variation is much larger (for Bank 0, mean = 0.59, SD = 0.13; Figure 3.7C).

To investigate this issue further, we measured the pooling coefficients in saline (Supplementary Figure S3.1) where the environment is more tightly controlled. Using a pair of flat electrodes we produced several voltage gradients along the electrode array and recorded them in both split and pooled mode. The pooling coefficients still differed systematically between Banks 0 and 1, but their variance was now considerably smaller (for Bank 0, mean = 0.60, SD = 0.06; Supplementary Figure S3.1C). This suggests that the higher variance of pooling coefficients measured *in vivo* is due to the interaction with brain tissue. For example some inactive element, such as a glial cell, may bind tightly to the electrode surface and thus raise the effective impedance. This would lower the pooling coefficient of the affected electrode and raise that of its partners. In any case, the experimenter will benefit from measuring these pooling coefficients rather than simply assuming they are 0.5. This will enable a more precise matching of unit waveforms obtained in pooled-mode and split-mode recordings.

Characterization of noise pooling

In Section 3.3 we described a noise model that takes into account three sources of noise: thermal noise at the electrodes, biological noise from distant neurons, and electronic noise from the acquisition system. By manipulating these sources independently we were able to test the predicted effects of noise pooling (Figure 3.8).

We began by measuring the noise with the electrode array immersed in saline using split-mode recording. One can vary the impedance of all electrodes by simply changing the concentration of saline, as this changes the bath resistance. Indeed, at low concentrations the noise increased dramatically (Figure 3.8A), and this reflects the thermal noise N_{ele} . At high concentrations the noise approached an asymptotic

level of 5–8 μV . We interpreted this as the electronic noise N_{com} that originates in the acquisition system. If this is the case, one predicts that the asymptotic noise level at high saline concentration should be identical across two electrodes that share the same wire, and that was indeed confirmed to high accuracy (Figure 3.8B).

Having established the independent measurements of N_{ele} for each electrode and N_{com} for each wire, one can now predict the noise that should be observed in a pooled recording, based on Eqn 3.6. These predictions agreed remarkably well with the measurements in pooled mode (Figure 3.8C). One surprising outcome from this analysis is that the electronic noise is remarkably high and varies considerably across wires (Figure 3.8B). It greatly exceeds the contribution from thermal noise at each electrode (Figure 3.8B). This differs from the behavior of other electrode recording systems where electronic noise tends to be negligible (e.g. Intan’s RHD2132 amplifier chip has input-referred RMS noise of 2.4 μV).

In addition, we also quantified the thermal noise present at 1X PBS by subtracting from it the noise at 10X PBS, which we assumed to be the electronic noise. The distribution of thermal noise is shown in Figure 3.8D.

We performed a similar analysis for noise measured *in vivo*. After spike-sorting, we identified recording sites that contained no spikes and analyzed the effect of pooling on the measured noise. Again, the additive model of Eqn 3.6 accounted well for the total noise values after pooling (Figure 3.8E). In addition, we used N_{com} and N_{ele} to estimate the biological noise N_{bio} . Figure 3.8F shows a distribution of the RMS biological noise, which had different amplitude on the two different banks (some 3.8 mm apart) and exceeded both the thermal and electronic noise.

With these insights about the behavior of signal and noise in the switched Neuropixels device we can reconsider the limits on beneficial pooling. The above measurements of noise contributions suggest that for this electrode array $\beta < 1$. Regarding the range of usable signal amplitudes we inspected public data from the Allen Institute from neurons recorded with Neuropixels arrays (Supplementary Figure S3.2): Across several brain areas the sortable spikes ranged in amplitude from 50 μV to 400 μV , suggesting that $\alpha = 8$. Under those conditions Eqn 3.8 predicts that the maximal useful pool size should be 7-8.

Simulations of pooling

How many electrodes could an experimenter pool and still sort every neuron with a high accuracy? In Section 3.3 we derived a theoretical limit to electrode pooling

based solely on the amplitude of signal and noise. To explore the influence of other factors, we turned to simulations, where we have perfect control over all aspects of the recording. We began by simulating data from a groups of 4 electrodes (“tetraode”) that together carry spikes from a single neuron. The tetraode was then pooled, channel by channel, with another tetraode, and spike-sorted with KiloSort2 to determine the number of cells that can be recovered with an accuracy greater than 0.8. We repeated this by iteratively pooling more tetraodes up to 12. Results are summarized in Figure 3.9. Importantly, we tested the effect of varying key parameters, such as the spike amplitude, firing rate, and the level of biological noise. This showed that across a diverse set of conditions, we could consistently recover 3-7 units (Figure 3.9B-D). This agrees well with the theoretical prediction based on the estimated signal and noise ratios in Section 3.3.

3.4 Discussion

Summary of results

This work presents the concept of electrode pooling as a way to multiply the yield of large electrode arrays. We show how the signals from many recording sites can be combined onto a small number of wires, and then recovered by a combination of experimental strategy and powerful spike-sorting software. The reduced requirement for wires coursing through the brain will lead to slender array devices that cause less damage to the neurons they are meant to observe. We developed the theory behind electrode pooling, analyzed the trade-offs of the approach, derived a mathematical limit to pooling, and developed a recipe for experiment and analysis that implements the procedure (Section 3.3). We also verified the basic assumptions about signal mixing and unmixing using a real existing device: the Neuropixels probe (Section 3.3). We showed that signals from different neurons can be reliably disambiguated and assigned back to the electrodes of origin. We also used simulations to show that under conditions that resemble real recordings, the spike-sorting algorithm can identify every neuron from a pool of 3-7 electrodes with a high accuracy. We surveyed publicly available data sets from the Allen Brain Observatory and derived spike amplitude distributions in different brain regions. Depending on the experimenters’ region of interest, the corresponding amplitude statistics would affect the ideal pool number. Nonetheless, electrode pooling of 3-7 should be a realistic range in practice, which agrees with our theoretical predictions. For the optimal design of electrode pools and to analyze the resulting data, it is advantageous to gather precise information about impedance and noise properties of the device.

Future developments

Hardware

The ability to service multiple recording sites with a single wire opens the door for much larger electrode arrays that nevertheless maintain a slim form factor and don't require any onboard signal processing. In the current version of the Neuropixels device [12] the ratio of electrodes to wires is only 2.5, and thus there is little practical benefit to be gained from electrode pooling. In most circumstances the user can probably use static selection to pick 40% of the electrodes and still monitor every possible neuron. However, devices are already under development with an electrode:wire ratio of 10 or greater. With pooled recording in mind one should contemplate future devices with much higher ratios still (see Figure 3.3 for guidance).

The design of effective electrode pools requires some flexibility in how recording sites are connected to wires. In the current Neuropixels each electrode has only one associated switch, and thus only one candidate wire. Electrode pools are limited to sites spaced 3.8 mm apart. The CMOS switch itself is small, but the local memory to store the switch state occupies some silicon space [26]. Nonetheless one can implement 3 switches per electrode even on a very tight pitch [7]. When arranged in a hierarchical network [19], these switches could generate a rich diversity of pooling schemes adapted to the specifics of any given experiment (Figure 3.10). For example, one could route any one electrode among a group of four to any one of three wires with two 1:4 switches (Figure 3.10C). This requires just 1 bit of storage per electrode, as in the current Neuropixels probe [12].

Another hardware design feature could greatly increase the capacity for electrode pooling: An optional inverter at each electrode (Figure 3.10D). This is a simple CMOS circuit that changes the sign of the waveform [3] depending on a local switch setting. If half of the electrodes in a pool use the inverter, that helps to differentiate the spike shapes of different neurons. Because extracellular signals from cell bodies generally start with a negative voltage swing, this effectively doubles the space of waveforms that occur in the pooled signal. In turn this would aid the spike-sorting analysis, ultimately allowing even more electrodes to share the same wire.

Software

Electrode pooling will also benefit from further developments in spike-sorting algorithms. For example, as we discussed, a promising strategy is to acquire all the spike shapes present on the electrode array using split-mode recordings, compute

the expected pooled-mode waveforms, and use those as templates in sorting the pooled signals. We have implemented this method in KiloSort2 and have shown that this can greatly increase the number of split-mode cells that are found in the pooled recordings (Figure 3.6). This idea may also be extended to clustering-based sorting algorithms, e.g. by guiding the initialization of the clustering step. Indeed, knowing which waveforms to look for in the recording would help any spike-sorter. We expect this method will also improve resolution of temporally overlapping spike waveforms.

Another interesting software challenge lies in the optimal design of electrode pools. As one envisions experiments with 10,000 or more recording sites, it becomes imperative to automate this process, so that the user wastes no time before launching into pooled recording. In Section 3.3 we lay out some heuristic rules one might follow, but turning these into an effective algorithm, and making use of the full noise and impedance specifications of the device, remains an open problem.

General strategies for using switchable probes

The switching circuitry in today's silicon probes opens up new ways of operating these devices. Static switching, time-division multiplexing, and electrode pooling are just a few of many possible modes of operation (Figure 3.1). Although this chapter focused on the benefits of electrode pooling, experimenters and hardware designers could in principle mix and match these strategies. Indeed they are fully complementary: static switching takes advantage of the redundancy in the neural signal across neighboring electrodes; time-division multiplexing uses the fact that a wire has much higher bandwidth than a neuron; and electrode pooling exploits the sparseness of spiking neural signals on the time axis. Exploring the best combination of these strategies for each use case may dramatically enhance the yield of extracellular recordings in the future.

3.5 Figures




Figure 3.1: Strategies for using a single wire to serve many recording sites in switchable silicon probes. (A) Time-division multiplexing. Rapidly cycling the selector switch allows a single wire to carry signals from many recording sites interleaved in time. Triangles represent anti-aliasing filters. (B) Static switching. A single wire connects to one of many possible recording sites through a selector switch. (C) Electrode pooling. Many recording sites are connected to a single wire through multiple controllable switches.




Figure 3.2: Pooling of signal and noise. An equivalent circuit model for two electrodes connected to a common wire along with downstream components of the signal chain, such as the amplifier, multiplexer, and digitizer. R_1 , R_2 : total impedance for electrodes 1 and 2, with contributions from the switch (R_{swi}), the metal/bath interface (R_{int}), and the external bath (R_{bat}). Noise sources include biological noise from distant neurons (N_{bio}); thermal noise from the total electrode impedance (N_{ele}), and common electronic noise from the amplifier and downstream components (N_{com})

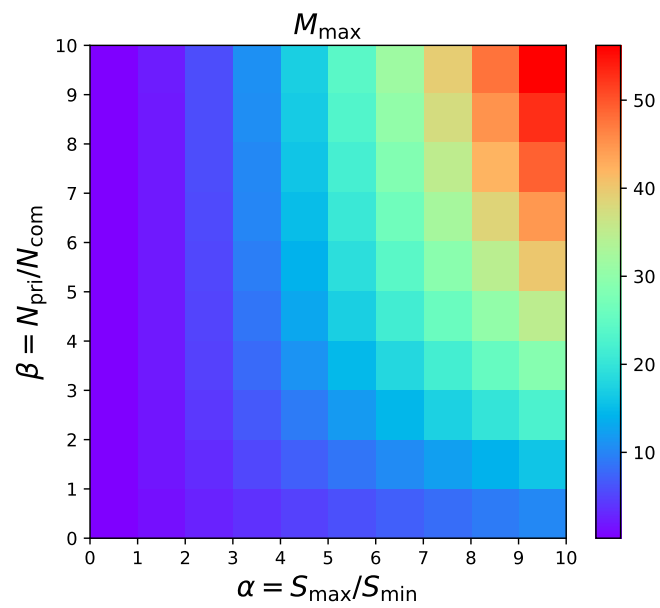


Figure 3.3: Maximal pool size M_{\max} as a function of the parameters α and β that characterize spike signals and noise.

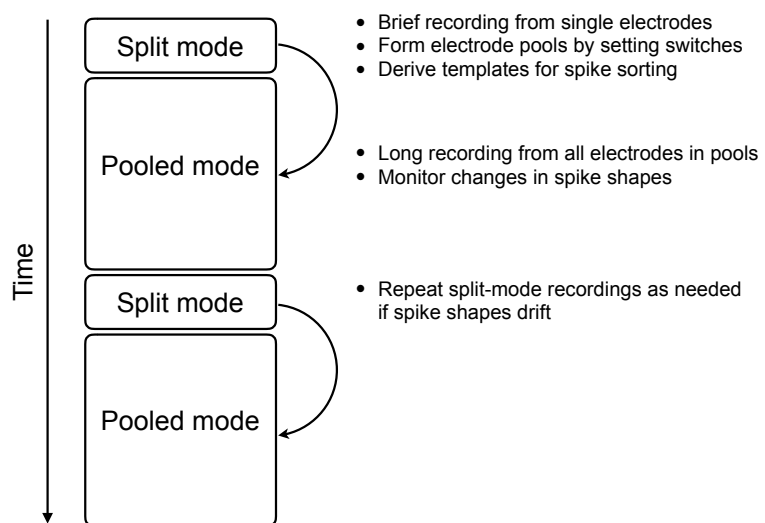


Figure 3.4: Time line of an experiment with electrode pooling.

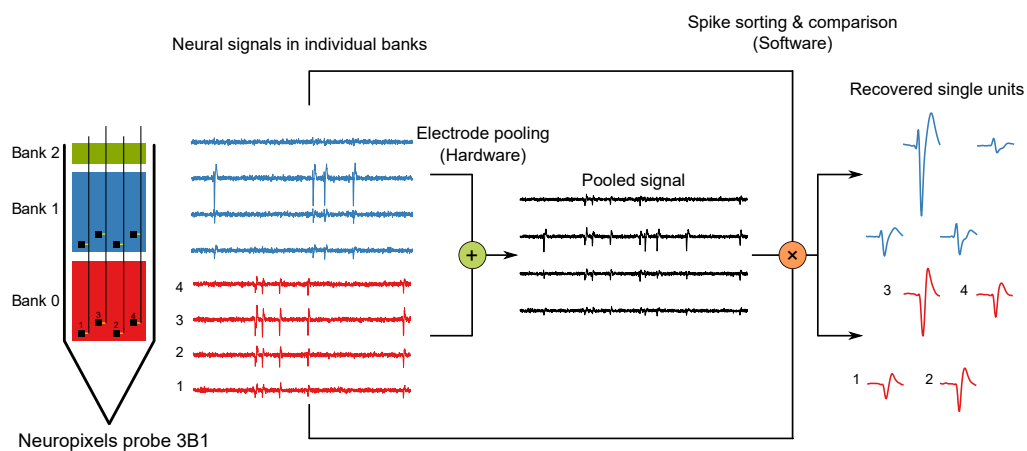


Figure 3.5: Overview of electrode pooling in the Neuropixels probe. Recording sites (black squares, numbered from 1 to 4) in the same relative location of each bank can be pooled to a single wire by closing the switches (yellow). This generates the pooled signal (black), which is a weighted average of the signals detected in each bank (red and blue traces). From the pooled signal, distinct spike shapes are recovered by spike-sorting. A comparison to the spike shapes observed in split-mode recordings allows the correct allocation of each spike to the electrodes of origin.

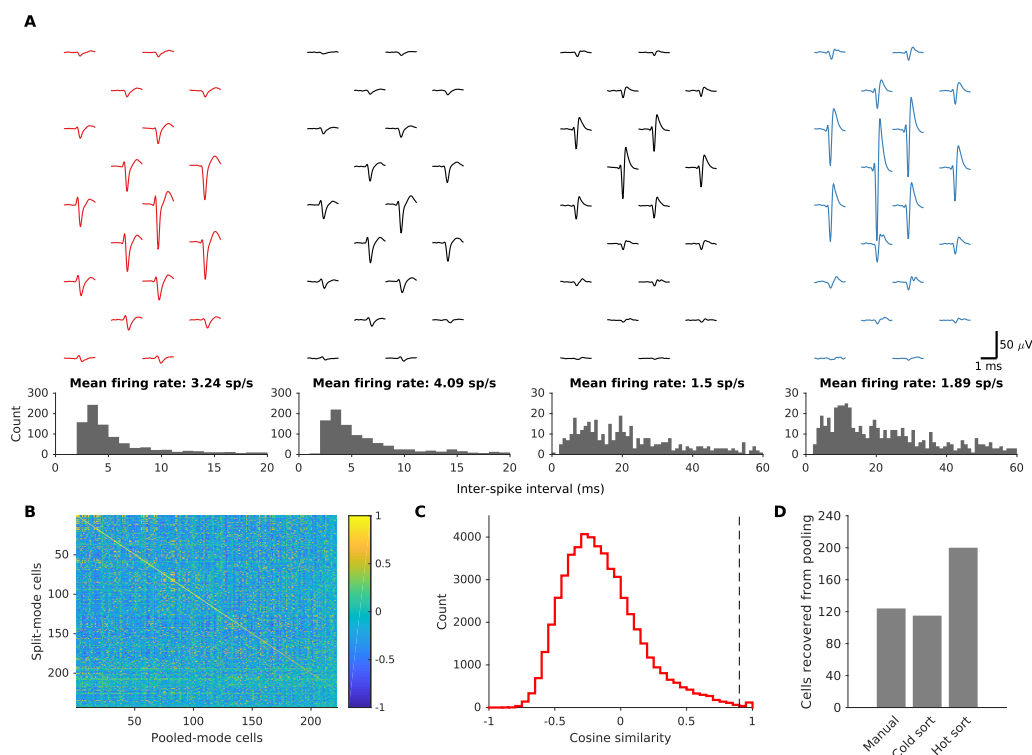


Figure 3.6: Unmixing pooled signal. (A) Examples of matches between cells from split- and pooled-mode recordings. Top: Waveforms of two sample units (middle, black) detected by an overlapping set of electrodes. They have been sorted and matched to a unit from Bank 0 (red) and a unit from Bank 1 (blue) based on the shape of their waveforms. Bottom: the mean firing rates as well as the interspike-interval distributions are also similar between the matched pairs. (B) Matrix of the cosine-similarity between units recorded in pooled mode and in split mode. The rows and columns are ordered from the best match to the worst. (C) Distribution of the cosine similarity. Dotted line is at 0.9. (D) Sorting pooled recording with templates from split-mode recordings. The number of cells from split-mode recordings whose matches were found in the pooled recording are plotted for three different sorting conditions: sorting all recordings by KiloSort1 followed by manual curation (left), sorting all recordings by KiloSort2 (middle), and sorting the pooled recording by KiloSort2 with templates initialized from the split recordings (right).

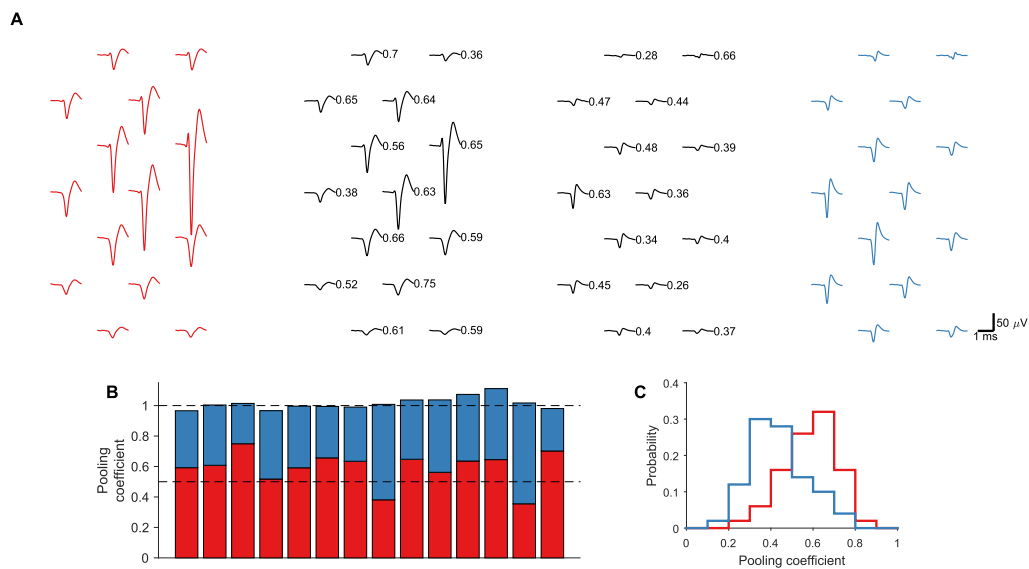
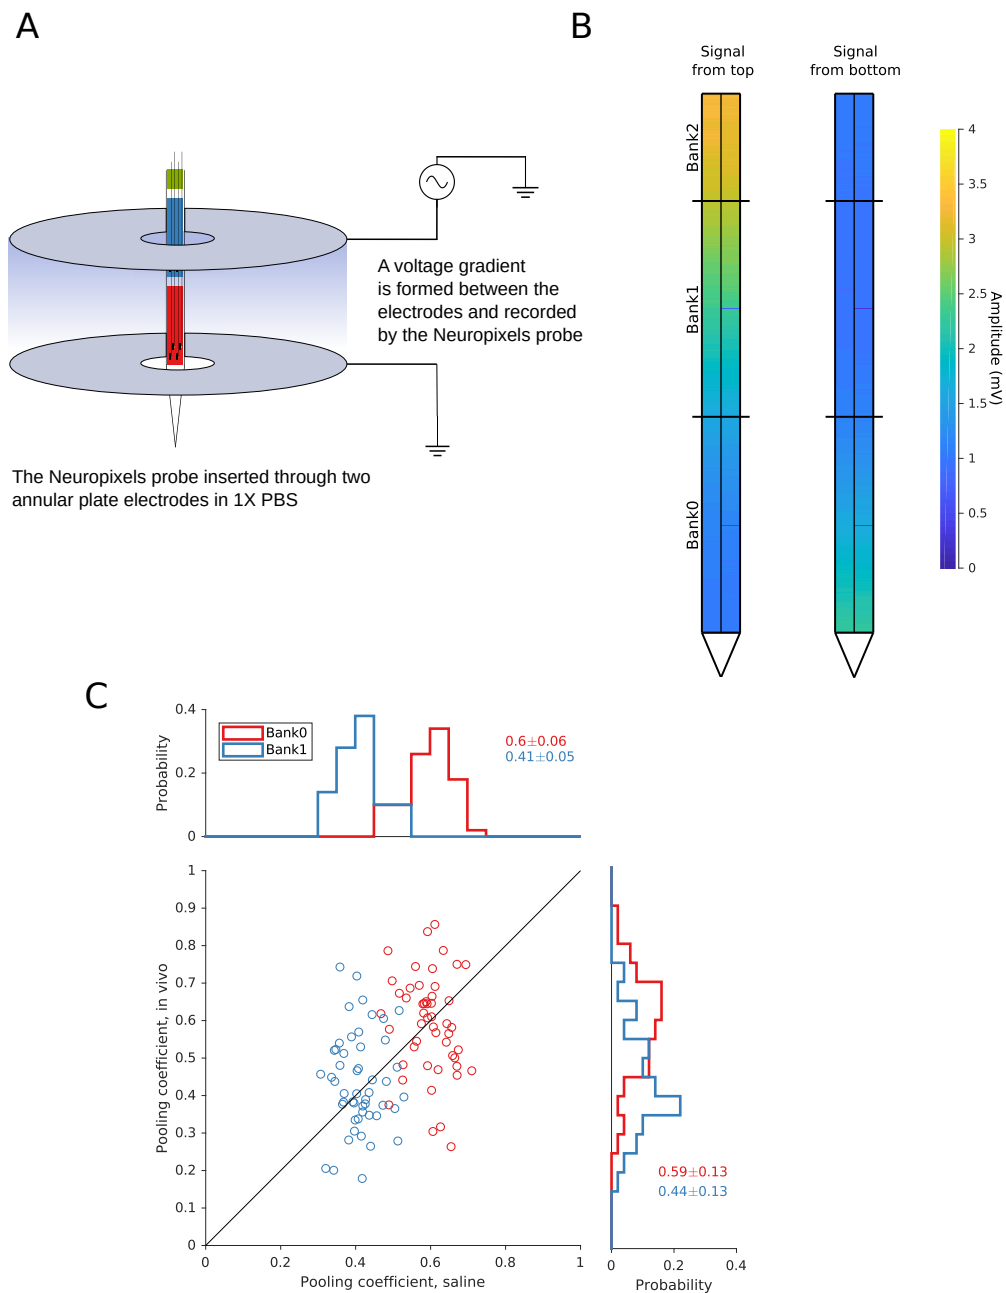


Figure 3.7: Characterization of signal pooling *in vivo*. (A) The spike waveforms from two neurons recorded on corresponding electrodes in Banks 0 and 1. Red, blue: Split-mode recordings from each bank. Black: Pooled-mode recording from both banks, with the two neurons separated by spike-sorting. The numbers next to the waveform at each recording site denote the pooling coefficients, namely the amplitude ratio of the pooled-mode signal to the corresponding split-mode signal. (B) Pooling coefficients for the recording sites shown in panel (A). Dotted lines are at 0.5 and 1. (C) Distribution of pooling coefficients from 50 pairs of sites in Bank 0 (red) and Bank 1 (blue).



Supplementary Figure S3.1: Measurement of pooling coefficient in saline. (A) Setup used to generate a voltage gradient in saline with annular electrodes. (B) Signal amplitude detected at every recording site of the Neuropixels probe when the gradient was applied from the top (left) and the bottom (right). (C) Distribution of the pooling coefficients of recording sites from *in vivo* recording (red and blue, same as in Figure 3.7C) and in saline (black). Text label is mean \pm standard deviation.

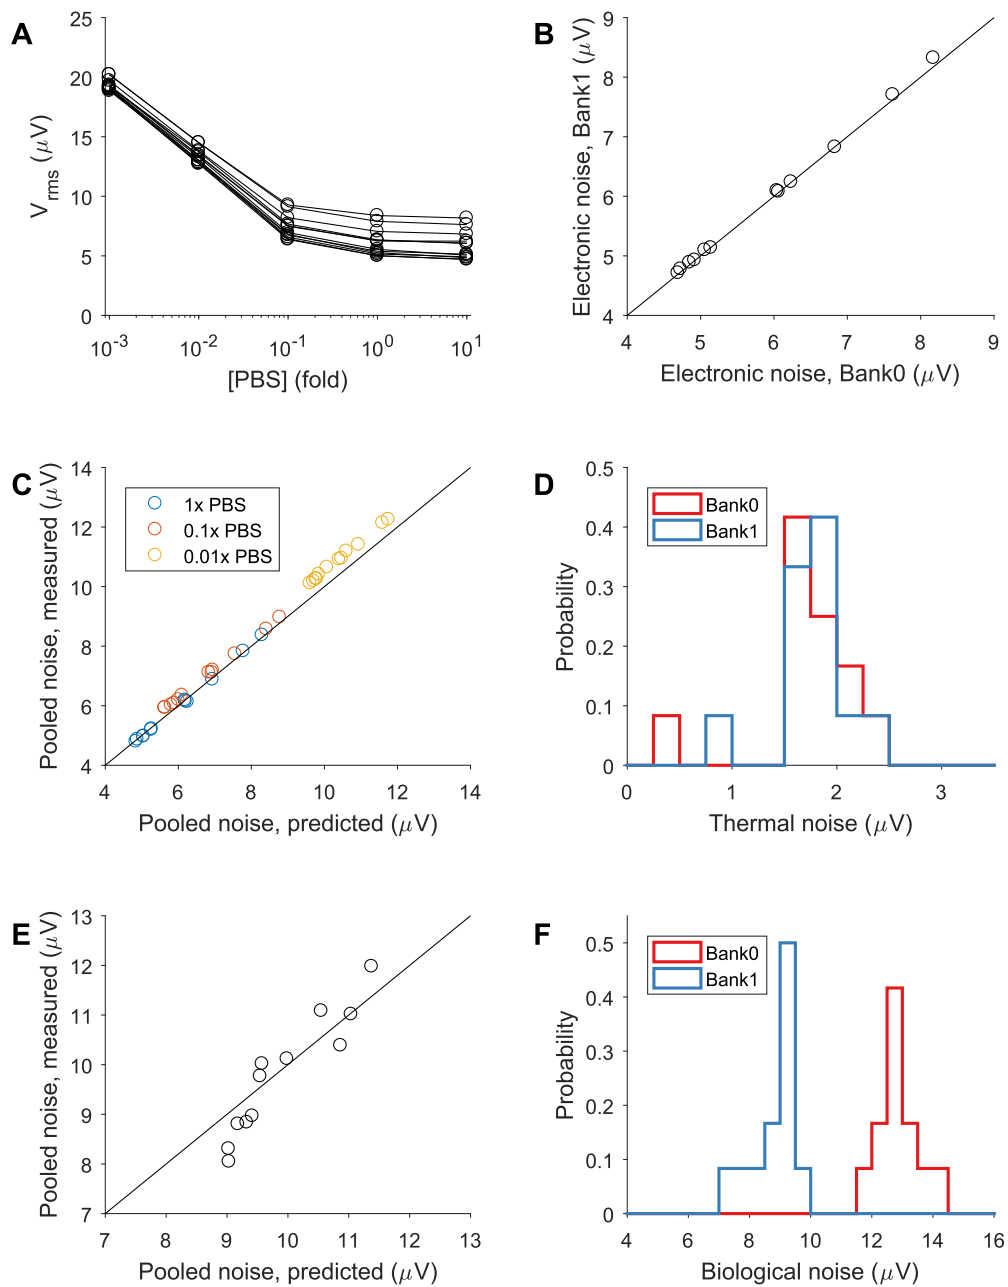


Figure 3.8: Characterization of noise pooling. (A) Noise measured at 12 recording sites in different concentrations of PBS. (B) The electronics noise, defined as the noise level at 10X PBS in (A) at both Bank 0 and Bank 1. (C) Prediction of the pooled noise in PBS from noise measured on each bank prior to pooling. (D) Distribution of thermal noise. (E) Same as (C), *in vivo*. (F) Distribution of biological noise.

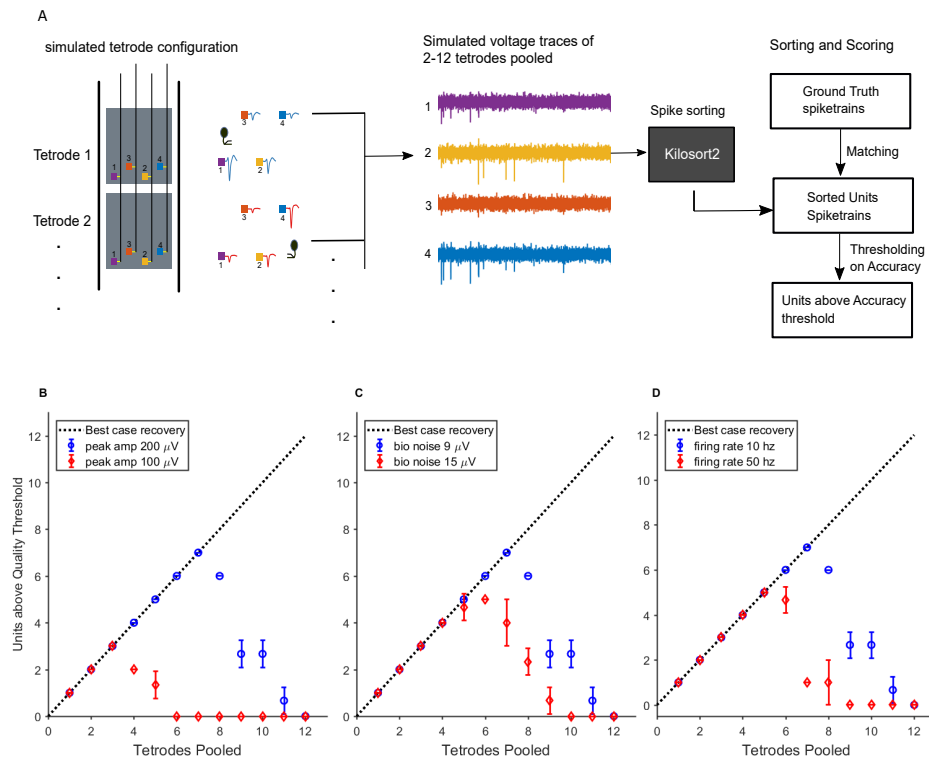
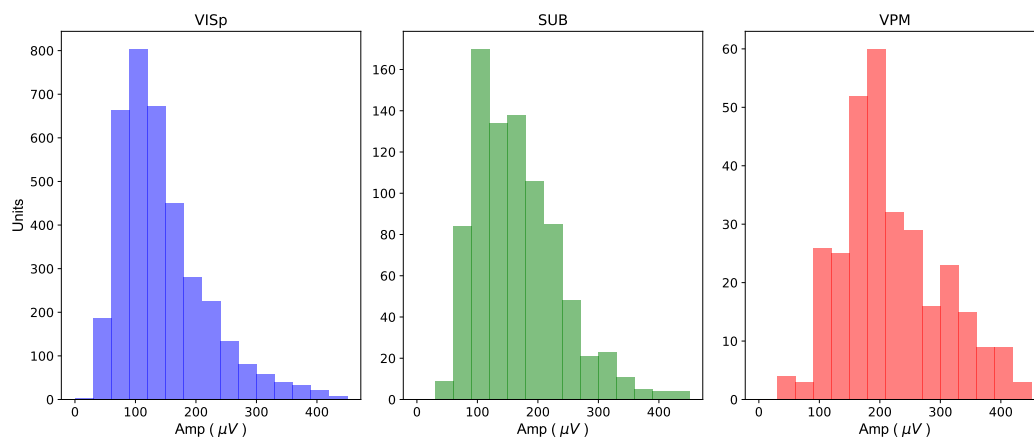


Figure 3.9: Simulation of pooling. (A) Schematic of the simulation. Each set of four electrodes (“tetrodes”) carries signal from one simulated unit. Pooling of up to 12 of these tetrodes is simulated with a realistic level of noise and sorted with KiloSort2. Spike trains of the sorted units are then matched with the ground truth and given an accuracy score. Units with accuracy > 0.8 are defined as being above the quality threshold. (B-D) Number of units recovered after pooling under different conditions. Each parameter combination was simulated 3 times by resampling the noise. Error bars are standard deviations.



Supplementary Figure S3.2: Distribution of spike amplitudes in three brain regions from the Allen Institute Brain Observatory dataset [30]. VISp, primary visual cortex; SUB, Subiculum; VPM, ventral posteromedial nucleus of the thalamus.

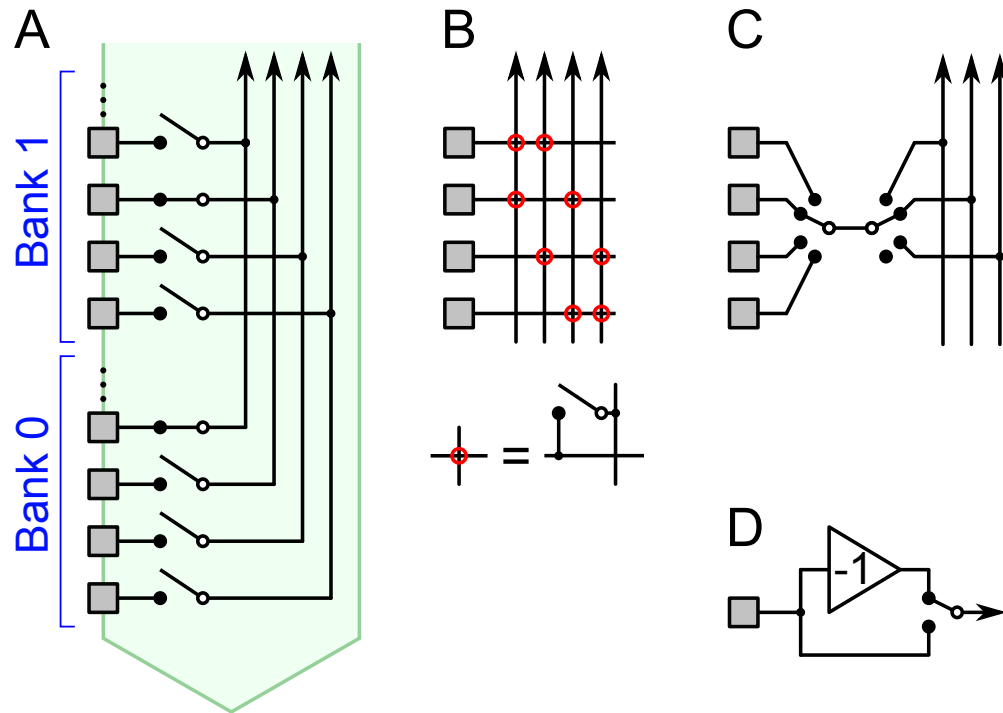


Figure 3.10: Hardware schemes for flexible connection between electrodes and wires. (A) In the current Neuropixels array, each electrode can be connected to just one wire with a controllable switch. (B) Two switches per electrode would allow a choice of 2 wires, enabling many more pooling configurations. (C) Because neighboring electrodes often carry redundant signals, one may want to choose just one from every group of 4. This switch circuit matches that choice with one of 3 (or no) wires. (D) An optional inverter for each electrode, controlled by a local switch.

3.6 Methods

All analysis was performed with Matlab R2016b (Mathworks) unless otherwise stated.

Control of switching circuitry

The Neuropixel probe version 3B1 has 960 recording sites that can be connected to 384 wires via controllable switches. The conventional mode of operation (split-mode) was to connect one electrode to one wire at a time. To implement electrode pooling, both the API of Neuropixels and the user-interfacing software (SpikeGLX) were modified such that more than one electrode can be connected to a single wire. This was done by our collaborators at Janelia.

Characterization of signal pooling (Figures 3.6, 3.7, Supplementary Figure S3.1)

In vivo recording

A Neuropixels probe [12] version 3B1 was used to record neural signals from a head-fixed mouse (C67/BL6, male, 9 months old). The probe entered the brain at 400 μm from midline and 3.7 mm posterior from bregma at $\sim 45^\circ$ and was advanced for ~ 6 mm, which corresponded to all of Bank 0 and roughly half of Bank 1. This covered many brain areas, from retrosplenial cortex at the top to medial preoptic nucleus at the bottom. Detailed description of the mouse surgery, probe implantation, and post hoc histology and imaging of probe track can be found in a previous report [15]. All procedures were done in accordance with institutional guidelines and approved by the Caltech IACUC.

Once the probe was implanted, the recording was done in the following order: (1) split-mode in Bank 0 (i.e. all 384 wires connected to recording sites in Bank 0); (2) split-mode in Bank 1; (3) pooled-mode across Banks 0 and 1. Each recording lasted for ~ 10 min.

Spike-sorting (Figure 3.6D)

KiloSort1 (downloaded from <https://github.com/cortex-lab/KiloSort> on Apr 10, 2018) was used to sort the *in vivo* recordings. The automatic template-matching step was run, followed by merging and designation of cells with high quality by manual inspection. This was done separately for each of the three recordings (split-mode Bank 0, split-mode Bank 1, pooled-mode). The setting for

running the automatic step was identical for all three recordings and is available in the code accompanying this manuscript.

The “hot sorting” feature was implemented in KiloSort2 (downloaded from <https://github.com/MouseLand/Kilosort2> on Mar 19, 2020). This had an additional benefit of not requiring manual curation, as KiloSort2 automatically performs a series of merges and splits after template matching. The two split-mode recordings were first sorted, and their templates were used to initialize the fields W and U of `rez2` before running the main template-matching function on the pooled recording. Cells designated as high quality (`KSLabel` of `Good`) by KiloSort2 were isolated and matched by the algorithm based on cosine similarity described in Section 3.6. This was compared with the result from “cold sorting,” in which the pooled recording was sorted on its own, as well as to the sorting that includes a manual curation step described in Section 3.6 (Figure 3.6D).

Unmixing pooled signal (Figures 3.6, 3.7A)

After sorting the split and pooled recordings, we computed the average waveform of every cell. Specifically, for each cell we averaged over the first n spikes, where n was the lesser of 7500 or all the spikes the cell fired during the recording.

We then sought to identify every cell in the pooled recordings with a cell in the split-mode recordings. This was done by the following procedure.

Let S denote a cell sorted from the split-mode recording ($S \in \mathcal{S}$) and S_i its waveform at channel i . Although i can range from 1 to 384 (the total number of wires available in the Neuropixels probe), we only focus on the channels with significantly large signal, i.e. $K = \{i : A(S_i) > a\}$ where $A(S_i)$ is the peak-to-peak amplitude of the pooled waveform detected at channel i and a is a pre-set threshold (typically $25 \mu\text{V}$). We wish to find the cell P from the pooled-mode recordings ($P \in \mathcal{P}$) that is closest to S . To do so, we compute the cosine similarity score for each pair (S, P) :

$$\Sigma(S, P) = \frac{\mathbf{S} \cdot \mathbf{P}}{\|\mathbf{S}\| \|\mathbf{P}\|} \quad (3.10)$$

where \mathbf{S} and \mathbf{P} are column vectors obtained by concatenating every S_k and P_k ($k \in K$), respectively, and $\|\cdot\|$ is the ℓ^2 norm. Σ is a $|\mathcal{S}|$ -by- $|\mathcal{P}|$ matrix. We identify the largest element of Σ , which corresponds to the most similar pair of S and P . We then update Σ by removing the row and column of this largest element. We repeat

this process until every $P \in \mathcal{P}$ is given a best match. By manual inspection, we found that pairs with similarity score greater than 0.9 were good matches (Figure 3.6B-C), and included only such pairs in the following analysis.

Estimating pooling coefficients *in vivo* and their distribution (Figure 3.7B-C)

Once each $P \in \mathcal{P}$ was assigned a match $S \in \mathcal{S}$, the pooling coefficient (c) was computed by solving the optimization problem below for each i with a least squares method (`mldivide` in Matlab).

$$\arg \min_{c_i} \|S_i - c_i P_i\| \quad (3.11)$$

Sometimes a single recording site detected action potentials from multiple cells. As a result, its pooling coefficient could be estimated from the signal of each of these cells. Typically, these estimates deviated from each other by less than 0.1. In these cases, we assigned the average of these values as the pooling coefficient of the recording site.

When two recording sites that share a wire in pooled-mode both carry significant signal, it enables the estimation of both of their pooling coefficients. Examples of such sites are shown in Figure 3.7A. From this, we could check whether their pooling coefficients sum to 1, as expected from Section 3.3 (Figure 3.7B). The distribution in Figure 3.7C only includes pooling coefficients of these sites (50 pairs in Banks 0 and 1).

Estimating pooling coefficient in saline (Supplementary Figure S3.1)

We estimated pooling coefficients in saline by generating a voltage gradient of a 1100 Hz sinusoidal signal with annular electrodes (1-inch outer diameter) (Supplementary Figure S3.1A). We then measured the signal amplitude at each recording site in split and pooled modes (Supplementary Figure S3.1B) by fitting a sine wave to a 50-ms segment of the band-pass filtered voltage trace with `lsqcurvefit` function in Matlab. The pooling coefficients of the recording sites in Banks 0 and 1 were estimated by measuring the signal amplitude with two different voltage gradients and then solving the following system of linear equations for each electrode with the `mldivide` function in Matlab:

$$\begin{bmatrix} V_{0,\text{top}} & V_{1,\text{top}} \\ V_{0,\text{bottom}} & V_{1,\text{bottom}} \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \end{bmatrix} = \begin{bmatrix} V_{P,\text{top}} \\ V_{P,\text{bottom}} \end{bmatrix}$$

Supplementary Figure S3.1C compares the distribution of pooling coefficients from Figure 3.7C to the distribution obtained by the above procedure at the same 50 recording sites.

Characterization of noise pooling (Figure 3.8)

A series of noise measurements in phosphate-buffered saline (Sigma-Aldrich P4417; 1X PBS contains 0.01 M phosphate buffer, 0.0027 M potassium chloride and 0.137 M sodium chloride, pH 7.4, at 25 °C) were made with the same probe used in the *in vivo* recording, in both split and pooled modes at various concentrations of PBS (10^{-3} , 10^{-2} , 10^{-1} , 1, and 10X). The RMS voltage in each condition was estimated by computing the standard deviation of ~10 s-long recording. Based on Eq. 3.6, the pooled noise was predicted from split-mode measurement (Figure 3.8C). To find the distribution of the thermal noise due to R_{bat} , the noise at 10X PBS was subtracted from noise at 1X PBS in quadrature (Figure 3.8D).

To estimate biological noise, 12 electrodes from the *in vivo* recording that had no action potentials were identified. The estimated electronic and thermal noise from Figure 3.8B and D were subtracted from the noise at these electrodes to obtain the distribution of the biological noise (Figure 3.8F). Again, Eq. 3.6 was used to predict the pooled noise *in vivo* (Figure 3.8E).

Simulation

600-s long voltage traces from groups of 4 electrodes ("tetrodes") (Figure 3.9A). Each tetrode carried signal from one unit. The signal (spikes) and the noise were generated separately and pooled according to Eqs. 3.3 and 3.6.

The spike times of the simulated unit were generated from a Poisson process with a 2 ms absolute refractory period (mean of either 10 Hz or 50 Hz). The spike waveform for each tetrode was randomly chosen from a bank of sample waveforms from a previous Neuropixels recording. One of the electrodes had the maximum spike amplitude (either 100 μV or 200 μV). The spike amplitude at the each of the other three electrodes was scaled down from the maximum by a factor chosen randomly from a uniform distribution with support from 0 to 1.

Each of the three sources of noise - N_{bio} , N_{ele} , N_{com} - were simulated as Gaussian white noise. The standard deviations for N_{ele} and N_{com} were $2 \mu\text{V}$ and $5 \mu\text{V}$ respectively, while two different parameters for N_{bio} were tried: $9 \mu\text{V}$ and $15 \mu\text{V}$, based on the noise measurements in Figure 3.8. N_{ele} was generated independently for each electrode, whereas N_{bio} was the same for all four electrodes in a tetrode.

Each parameter combination was simulated 3 times by resampling the noise.

The simulated data was sorted by KiloSort2. When a large number of simulated tetrodes were pooled, the signal-to-noise-ratio became too low for KiloSort2 to form templates in the preprocessing steps and led to premature terminations. In these cases the number of units recovered was set to zero.

Following methods in the field, [4, 18], the spike timings of the sorted units were matched to those of the ground truth units via a confusion matrix. The accuracy for each unit was quantified as the following:

$$\text{Accuracy} = \frac{\text{\#Matches}}{\text{\#Matches} + \text{\#Misses} + \text{\#False positives}} \quad (3.12)$$

Sorted units with an accuracy score greater than 0.8 were included as units above quality threshold.

References

- [1] Gian Nicola Angotzi, Fabio Boi, Aziliz Lecomte, Ermanno Miele, Mario Malerba, Stefano Zucca, Antonino Casile, and Luca Berdondini. SiNAPS: An implantable active pixel sensor CMOS-probe for simultaneous large-scale neural recordings. *Biosensors and Bioelectronics*, 129:355–364, 2019.
- [2] D. Attwell and S. B Laughlin. An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab*, 21:1133–45, October 2001.
- [3] Woorham Bae. CMOS Inverter as Analog Circuit: An Overview. *Journal of Low Power Electronics and Applications*, 9:26, August 2019. doi: 10.3390/jlpea9030026.
- [4] Alex H. Barnett, Jeremy F. Magland, and Leslie F. Greengard. Validation of neural spike sorting algorithms without ground-truth information. *Journal of Neuroscience Methods*, 264:65 – 77, 2016. ISSN 0165-0270. doi: <https://doi.org/10.1016/j.jneumeth.2016.02.022>. URL <http://www.sciencedirect.com/science/article/pii/S0165027016300036>.

- [5] BRAIN Working Group. BRAIN 2025: A Scientific Vision. Technical report, 2014. URL <https://braininitiative.nih.gov/strategic-planning/brain-2025-report>.
- [6] George Dimitriadis, Joana P. Neto, Arno Aarts, Andrei Alexandru, Marco Ballini, Francesco Battaglia, Lorenza Calcaterra, Francois David, Richárd Fiáth, João Frazão, Jesse P Geerts, Luc J. Gentet, Nick Van Helleputte, Tobias Holzhammer, Chris van Hoof, Domonkos Horváth, Gonçalo Lopes, Carolina M. Lopez, Eric Maris, Andre Marques-Smith, Gergely Márton, Bruce L. McNaughton, Domokos Meszéna, Srinjoy Mitra, Silke Musa, Hercules Neves, Joana Nogueira, Guy A. Orban, Frederick Pothof, Jan Putzeys, Bogdan Raducanu, Patrick Ruther, Tim Schroeder, Wolf Singer, Paul Tiesinga, Istvan Ulbert, Shiwei Wang, Marleen Welkenhuysen, and Adam R. Kampff. Why not record from every channel with a CMOS scanning probe? *bioRxiv*, March 2018. doi: 10.1101/275818. URL <http://biorxiv.org/lookup/doi/10.1101/275818>.
- [7] J. Dragas, V. Viswam, A. Shadmani, Y. Chen, R. Bounik, A. Stettler, M. Radivojevic, S. Geissler, M. E. J. Obien, J. Müller, and A. Hierlemann. In Vitro Multi-Functional Microelectrode Array Featuring 59 760 Electrodes, 2048 Electrophysiology Channels, Stimulation, Impedance Measurement, and Neurotransmitter Detection Channels. *IEEE Journal of Solid-State Circuits*, 52(6): 1576–1590, June 2017. ISSN 0018-9200. doi: 10.1109/JSSC.2017.2686580.
- [8] B. Eversmann, M. Jenkner, F. Hofmann, C. Paulus, R. Brederlow, B. Holzapfl, P. Fromherz, M. Merz, M. Brenner, M. Schreiter, R. Gabl, K. Plehnert, M. Steinhauser, G. Eckstein, D. Schmitt-Landsiedel, and R. Thewes. A 128 x 128 cmos biosensor array for extracellular recording of neural activity. *IEEE Journal of Solid-State Circuits*, 38(12):2306–2317, December 2003. ISSN 0018-9200. doi: 10.1109/JSSC.2003.819174. URL <http://ieeexplore.ieee.org/document/1253878/>.
- [9] Surya Ganguli and Haim Sompolinsky. Compressed Sensing, Sparsity, and Dimensionality in Neuronal Information Processing and Data Analysis. *Annual Review of Neuroscience*, 35(1):485–508, 2012. doi: doi:10.1146/annurev-neuro-062111-150410. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-neuro-062111-150410>.
- [10] K. D. Harris, D. A. Henze, J. Csicsvari, H. Hirase, and G. Buzsaki. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol*, 84:401–14, July 2000.
- [11] James Jun, Catalin Mitelut, Chongxi Lai, Sergey L Gratiy, Costas A Anastasiou, and Timothy D Harris. Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. *bioRxiv*, pages 1–29, 2017. doi: <https://doi.org/10.1101/101030>.

- [12] James J. Jun, Nicholas A. Steinmetz, Joshua H. Siegle, Daniel J. Denman, Marius Bauza, Brian Barbarits, Albert K. Lee, Costas A. Anastassiou, Alexandru Andrei, Çağatay Aydin, Mladen Barbic, Timothy J. Blanche, Vincent Bonin, João Couto, Barundeb Dutta, Sergey L. Gratiy, Diego A. Gutnisky, Michael Häusser, Bill Karsh, Peter Ledochowitsch, Carolina Mora Lopez, Catalin Mitelut, Silke Musa, Michael Okun, Marius Pachitariu, Jan Putzeys, P. Dylan Rich, Cyrille Rossant, Wei Lung Sun, Karel Svoboda, Matteo Carandini, Kenneth D. Harris, Christof Koch, John O’Keefe, and Timothy D. Harris. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551 (7679):232–236, 2017. ISSN 14764687. doi: 10.1038/nature24636. URL <http://dx.doi.org/10.1038/nature24636>.
- [13] David Kleinfeld, Lan Luan, Partha P. Mitra, Jacob T. Robinson, Rahul Sarpeshkar, Kenneth Shepard, Chong Xie, and Timothy D. Harris. Can One Concurrently Record Electrical Spikes from Every Neuron in a Mammalian Brain? *Neuron*, 103(6):1005–1015, September 2019. ISSN 08966273. doi: 10.1016/j.neuron.2019.08.011. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627319306956>.
- [14] Takashi D. Y. Kozai and Alberto L. Vazquez. Photoelectric artefact from optogenetics and imaging on microelectrodes and bioelectronics: New challenges and opportunities. *Journal of Materials Chemistry B*, 3(25):4965–4978, June 2015. ISSN 2050-7518. doi: 10.1039/C5TB00108K.
- [15] Kyu Hyun Lee, Alvita Tran, Zeynep Turan, and Markus Meister. The sifting of visual information in the superior colliculus. *eLife*, 9, April 2020. ISSN 2050-084X. doi: 10.7554/eLife.50678.
- [16] Michael S. Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network*, 9 4:R53–78, 1998.
- [17] Carolina Mora Lopez, Jan Putzeys, Bogdan C. Raducanu, Marco Ballini, Shiwei Wang, Alexandru Andrei, Veronique Rochus, Roeland Vandebriel, Simone Severi, Chris Van Hoof, Silke Musa, Nick Van Helleputte, Refet F. Yazicioglu, and Srinjoy Mitra. A Neural Probe With Up to 966 Electrodes and Up to 384 Configurable Channels in 0.13 μm SOI CMOS. *IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS*, 11(3):510–522, 2017.
- [18] Jeremy F. Magland, James J. Jun, Elizabeth Lovero, Alexander J. Morley, Cole L. Hurwitz, Alessio P. Buccino, Samuel Garcia, and Alex H. Barnett. Spikeforest: reproducible web-facing ground-truth validation of automated neural spike sorters. *bioRxiv*, 2020. doi: 10.1101/2020.01.14.900688. URL <https://www.biorxiv.org/content/early/2020/01/14/2020.01.14.900688>.
- [19] Jan Müller, Marco Ballini, Paolo Livi, Yihui Chen, Milos Radivojevic, Amir Shadmani, Vijay Viswam, Ian L. Jones, Michele Fiscella, Roland Diggelmann,

- Alexander Stettler, Urs Frey, Douglas J. Bakkum, and Andreas Hierlemann. High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. *Lab on a Chip*, 15(13):2767–2780, 2015. ISSN 1473-0197, 1473-0189. doi: 10.1039/C5LC00133A. URL <http://xlink.rsc.org/?DOI=C5LC00133A>.
- [20] Marie Engelen J. Obien, Kosmas Deligkaris, Torsten Bullmann, Douglas J. Bakkum, and Urs Frey. Revealing neuronal function through microelectrode array recordings. *Frontiers in Neuroscience*, 8, January 2015. ISSN 1662-453X. doi: 10.3389/fnins.2014.00423. URL <http://journal.frontiersin.org/article/10.3389/fnins.2014.00423/abstract>.
- [21] Marius Pachitariu, Nicholas A. Steinmetz, Shabnam N. Kadir, Matteo Carandini, and Kenneth D. Harris. Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems*, 29(Nips):4448–4456, 2016. ISSN 10495258. URL <https://papers.nips.cc/paper/6326-fast-and-accurate-spike-sorting-of-high-channel-count-probes-with-kilosort.pdf>.
- [22] Bogdan C. Raducanu, Refet F. Yazicioglu, Carolina M. Lopez, Marco Ballini, Jan Putzeys, Shiwei Wang, Alexandru Andrei, Marleen Welkenhuysen, Nick Van Helleputte, Silke Musa, Robert Puers, Fabian Kloosterman, Chris Van Hoof, and Srinjoy Mitra. Time multiplexed active neural probe with 678 parallel recording sites. *European Solid-State Device Research Conference*, 2016-Octob:385–388, 2016. ISSN 19308876. doi: 10.1109/ESSDERC.2016.7599667.
- [23] Gustavo Rios, Evgueniy V. Lubenov, Derrick Chi, Michael L. Roukes, and Athanassios G. Siapas. Nanofabricated Neural Probes for Dense 3-D Recordings of Brain Activity. *Nano Letters*, 16(11):6857–6862, November 2016. ISSN 1530-6984, 1530-6992. doi: 10.1021/acs.nanolett.6b02673. URL <https://pubs.acs.org/doi/10.1021/acs.nanolett.6b02673>.
- [24] D.A. Robinson. The electrical properties of metal microelectrodes. *Proceedings of the IEEE*, 56(6):1065–1071, 1968. ISSN 0018-9219. doi: 10.1109/PROC.1968.6458. URL <http://ieeexplore.ieee.org/document/1448388/>.
- [25] Erik W. Schomburg, Costas A. Anastassiou, György Buzsáki, and Christof Koch. The Spiking Component of Oscillatory Extracellular Potentials in the Rat Hippocampus. *Journal of Neuroscience*, 32(34):11798–11811, August 2012. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0656-12.2012.
- [26] Karsten Seidl, Stanislav Herwik, Tom Torfs, Herc Neves, Oliver Paul, and Patrick Ruther. CMOS-Based High-Density Silicon Microprobe Arrays for Electronic Depth Control in Intracortical Neural Recording. *Journal of Microelectromechanical Systems*, 20:1439–1448, 2011.

- [27] Karsten Seidl, Michael Schwaerzle, Istvan Ulbert, Herc P. Neves, Oliver Paul, and Patrick Ruther. CMOS-Based High-Density Silicon Microprobe Arrays for Electronic Depth Control in Intracortical Neural Recording—Characterization and Application. *Journal of Microelectromechanical Systems*, 21(6):1426–1435, December 2012. ISSN 1057-7157, 1941-0158. doi: 10.1109/JMEMS.2012.2206564. URL <http://ieeexplore.ieee.org/document/6249712/>.
- [28] Farzaneh Shahrokhi, Karim Abdelhalim, Demitre Serletis, Peter L. Carlen, and Roman Genov. The 128-Channel Fully Differential Digital Integrated Neural Recording and Stimulation Interface. *IEEE Transactions on Biomedical Circuits and Systems*, 4(3):149–161, June 2010. ISSN 1932-4545, 1940-9990. doi: 10.1109/TBCAS.2010.2041350. URL <http://ieeexplore.ieee.org/document/5471738/>.
- [29] C. E. Shannon. Communication in the Presence of Noise. *Proceedings of the IRE*, 37(1):10–21, January 1949. ISSN 0096-8390. doi: 10.1109/JRPROC.1949.232969.
- [30] Joshua H. Siegle, Xiaoxuan Jia, Séverine Durand, Sam Gale, Corbett Bennett, Nile Graddis, Gregory Heller, Tamina K. Ramirez, Hannah Choi, Jennifer A. Luviano, Peter A. Groblewski, Ruweida Ahmed, Anton Arkhipov, Amy Bernard, Yazan N. Billeh, Dillan Brown, Michael A. Buice, Nicolas Cain, Shiella Caldejon, Linzy Casal, Andrew Cho, Maggie Chvilicek, Timothy C. Cox, Kael Dai, Daniel J. Denman, Saskia E. J. de Vries, Roald Dietzman, Luke Esposito, Colin Farrell, David Feng, John Galbraith, Marina Garrett, Emily C. Gelfand, Nicole Hancock, Julie A. Harris, Robert Howard, Brian Hu, Ross Hytten, Ramakrishnan Iyer, Erika Jessett, Katelyn Johnson, India Kato, Justin Kiggins, Sophie Lambert, Jerome Lecoq, Peter Ledochowitsch, Jung Hoon Lee, Arielle Leon, Yang Li, Elizabeth Liang, Fuhui Long, Kyla Mace, Jose Melchior, Daniel Millman, Tyler Mollenkopf, Chelsea Nayan, Lydia Ng, Kiet Ngo, Thuyahn Nguyen, Philip R. Nicovich, Kat North, Gabriel Koch Ocker, Doug Ollerenshaw, Michael Oliver, Marius Pachitariu, Jed Perkins, Melissa Reding, David Reid, Miranda Robertson, Kara Ronellenfitch, Sam Seid, Cliff Slaughterbeck, Michelle Stoecklin, David Sullivan, Ben Sutton, Jackie Swapp, Carol Thompson, Kristen Turner, Wayne Wakeman, Jennifer D. Whitesell, Derric Williams, Ali Williford, Rob Young, Hongkui Zeng, Sarah Naylor, John W. Phillips, R. Clay Reid, Stefan Mihalas, Shawn R. Olsen, and Christof Koch. A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas. *bioRxiv*, 2019. doi: 10.1101/805010. URL <https://www.biorxiv.org/content/early/2019/10/16/805010>.
- [31] Tom Torfs, Arno Aarts, Mehmet Akif Erismis, Junaid Aslam, Refet Firat Yazicioglu, Robert Puers, Chris Van Hoof, Herc Neves, Istvan Ulbert, Balazs Dombovari, Richard Fiath, Balint Peter Kerekes, Karsten Seidl, Stanislav Herwik, and Patrick Ruther. Two-dimensional multi-channel neural probes

with electronic depth control. In *2010 Biomedical Circuits and Systems Conference (BioCAS)*, pages 198–201, Paphos, Cyprus, November 2010. IEEE. ISBN 978-1-4244-7269-7. doi: 10.1109/BIOCAS.2010.5709605. URL <http://ieeexplore.ieee.org/document/5709605/>.

- [32] Pierre Yger, Giulia Lb Spampinato, Elric Esposito, Baptiste Lefebvre, Stéphane Deny, Christophe Gardella, Marcel Stimberg, Florian Jetter, Guenther Zeck, Serge Picaud, Jens Duebel, and Olivier Marre. A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *eLife*, 7, March 2018. ISSN 2050-084X. doi: 10.7554/eLife.34518.

*Chapter 4***CONCLUDING REMARKS**

Here I briefly summarize the main results from Chapters 2 and 3 and describe ways to build on them in the future.

4.1 Visual sifting in the SC

In the first project, we presented the results of large-scale electrophysiology in the SC of the head-fixed mouse, with a focus on the processing of behaviorally relevant stimuli such as the dark expanding disk. By recording from every layer of the SC simultaneously, we showed that compared to the superficial SC, the neurons in the deeper SC become more selective to the looming stimulus and more invariant to its position. Furthermore, these neurons show a rapid and strong habituation to repeated stimuli that is specific to the stimulus position and lasts for a behaviorally relevant timescale, even in the absence of the cerebral cortex. We have explored a possible mechanism of this phenomenon with a simple circuit model and ruled out several others based on our data.

Perhaps the most striking of the three response properties is the stimulus-specific habituation. It is also the most puzzling, as it is not obvious why it should belong with the selectivity and invariance. What function does it serve? One speculation is that the deeper SC neurons we described not only serve as a sensory representation (“here comes a predator”) but also a behavioral command (“here comes a predator, so run!”). In the case of the face cell from the primate inferotemporal cortex, such a habituation is not observed. Presumably this is because there are many ways such a representation can be used, e.g. it can be used to identify your grandmother from a group of people, or to draw her face in an art class. On the other hand, the representation of an approaching predator in the the mouse brain may have no real value other than to issue the appropriate behavioral response. And given the need to execute this behavior particularly rapidly, the brain may have merged the representation with its use in a single neuron, such that when triggered by a threatening stimulus, it drives a whole cascade of motor outputs. It may also be that once such a motor program is activated, the neuron must not issue another command, unless a new predator appears in a different part of the visual field and requires a course correction. This, of course, means that the fidelity of the sensory

representation is compromised; if we equate the activity of these neurons with “seeing” a predator, then the animal only sees it the first time it appears at a given location. But maybe this is a reasonable trade-off, since there is no *need* for the animal to see the predator once the behavioral command to flee the scene has been issued – in fact, investing neural resources to continuously monitor the predator’s location during escape might be wasteful. Interpreted this way, the stimulus-specific habituation reflects a constraint on the sensory representation by the downstream circuit that uses it ¹.

There are many potential future directions to follow this project. One line of investigation is a more in-depth analysis of the circuit described in Chapter 2. Although the model proposed in Figure 2.6 is more likely than the alternative mechanisms we considered, the dearth of information about cell types and their connectivity in the SC means it remains largely a hypothesis. In particular, it would be interesting to further study the mechanism for the long-lasting stimulus-specific habituation. We have pointed to a particular cell type as a likely candidate for the global looming detector, and a Cre line for this type has been previously reported [2]. This could provide a genetic handle to further investigate the circuit mechanism, for example, by intracellularly recording the dynamics of the excitatory and inhibitory synaptic currents, or by imaging the activity of this cell type in a freely moving animal during the looming stimulus.

Another line of investigation is to explore other sifting functions of the SC. Given its access to the entirety of retinal output [1], the SC is well-poised to implement many other visual computations. As mentioned earlier, it may be involved in tracking small moving objects like bugs in the lower visual field to guide hunting behavior [3]. In addition, it may be involved in other high-level operations usually associated with the cortex, such as the modulation of visual attention [5]. The presence of a mutant mouse line that developmentally lacks the cerebral cortex [4] may be useful for parsing the function of the SC from the reciprocally connected cortical areas. It would also be interesting to study how the SC processes more naturalistic visual stimuli rather than the artificial ones used to probe its function so far.

Finally, one could follow the neural processing of the looming reaction further and ask: how is the visual information sifted by the SC converted to behavior? As previously discussed, there are two possible responses to the looming reaction: flight

¹We refrained from this speculation in Chapter 2 because the recordings were done in head-fixed mice and no behavioral output was consistently observed in response to the looming stimulus in that state.

and freezing. How the animal incorporates previous knowledge of its environment (e.g. the presence of a nest) and makes this decision in a few tenths of a second is unknown. Another question has to do with the planning of motor output once the decision is made. In the case of flight, how does brain use the collicular output to plan and execute a trajectory back to the nest? Despite intense investigations into the neural basis of navigation, much remains unexplained and few studies have examined such a rapid homing behavior in the mouse. This question is particularly interesting given that the animal forms a memory of the location of the nest after just a single visit [6]. The interplay between the SC and other brain areas such as the hippocampus in this context would be interesting to study.

4.2 Electrode pooling

In the second project, we have presented a new method for boosting the yield of extracellular recordings in silicon neural probes, called electrode pooling. We have analyzed its effect on the signal and noise theoretically and verified the predictions experimentally by implementing it in the Neuropixels probe. In addition, we have explored the limits of electrode pooling with simulations and made suggestions on the design of future devices to take advantage of it.

Perhaps the most important future direction is the development of an algorithm that generates an intelligently chosen set of electrode pools during an experiment. This would be crucial for making electrode pooling useful to the user, as silicon probes in the future will have orders of magnitude more recording sites and any attempt to manually pool electrodes in this setting would be futile. Although we have identified some heuristics for intelligent pooling (e.g. pool electrodes with large, sparse action potentials), turning them into an algorithm that performs robustly in many physiological conditions would require a more extensive investigation both *in vivo* and *in silico*. Details of an algorithm may also depend on what is physically possible to pool with the hardware, so the development of both the hardware and the software would have to move in lockstep. A new version of the Neuropixels probe will be made available in the near future, with many more recording sites and a greater electrode-to-wire ratio. It would be interesting to test electrode pooling in this device and analyze the performance gains *in vivo* over conventional recordings.

References

- [1] Erika M. Ellis, Gregory Gauvain, Benjamin Sivyer, and Gabe J. Murphy. Shared and distinct retinal input to the mouse superior colliculus and dorsal lateral

- geniculate nucleus. *Journal of Neurophysiology*, 116(2):602–610, 2016. ISSN 0022-3077. doi: 10.1152/jn.00227.2016. URL <http://jn.physiology.org/lookup/doi/10.1152/jn.00227.2016>.
- [2] Samuel D Gale and Gabe J Murphy. Distinct Representation and Distribution of Visual Information by Specific Cell Types in Mouse Superficial Superior Colliculus. *Journal of Neuroscience*, 34(40):13458–13471, 2014. doi: 10.1523/jneurosci.2768-14.2014. URL <http://www.jneurosci.org/content/34/40/13458.abstract>.
- [3] Jennifer L. Hoy, Iryna Yavorska, Michael Wehr, and Cristopher M. Niell. Vision Drives Accurate Approach Behavior during Prey Capture in Laboratory Mice. *Current Biology*, 26(22):3046–3052, 2016. ISSN 0960-9822. doi: 10.1016/j.cub.2016.09.009. URL <http://dx.doi.org/10.1016/j.cub.2016.09.009>.
- [4] S Kim, M K Lehtinen, A Sessa, M W Zappaterra, S H Cho, D Gonzalez, B Boggan, C A Austin, J Wijnholds, M J Gambello, J Malicki, A S LaMantia, V Broccoli, and C A Walsh. The apical complex couples cell fate and cell survival to cerebral cortical development. *Neuron*, 66:69–84, 2010. doi: 10.1016/j.neuron.2010.03.019. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=20399730.
- [5] Richard J Krauzlis, Lee P Lovejoy, and Alexandre Zénon. Superior Colliculus and Visual Spatial Attention. *Annual Review of Neuroscience*, 36(1):165–182, 2013. doi: doi:10.1146/annurev-neuro-062012-170249. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-neuro-062012-170249>.
- [6] Ruben Vale, Dominic A Evans, and Tiago Branco. Rapid Spatial Learning Controls Instinctive Defensive Behavior in Mice. *Current Biology*, 27(9):1342–1349, October 2017. ISSN 0960-9822. doi: 10.1016/j.cub.2017.03.031. URL [http://www.cell.com/current-biology/abstract/S0960-9822\(17\)30329-9](http://www.cell.com/current-biology/abstract/S0960-9822(17)30329-9).