Developing Tools for Neurobiology: The Retina as a Neuropharmacology Testbed & Electrode Pooling to Boost Extracellular Array Recording

Thesis by
Yu-Li Ni

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

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2023
Defended September 29, 2022

© 2023

Yu-Li Ni
ORCID: 0000-0003-1600-9854

All rights reserved except where otherwise noted

ACKNOWLEDGEMENTS

When I was a medical student, I was intrigued by the elusive phenomenon of general anesthesia as well as how people use it pervasively without any meaningful understanding of the anesthetics. Later in life, I realized the same statement holds true for most aspects of human society, and during the Ph.D. learned why the mechanism stayed elusive for so long. Despite of this, our approach of reverse engineering the neuropharmacology with known neural circuits are revealing the mechanisms in general anesthesia as well as the sub-anesthetic dosage effect of some of the drugs such as ketamine.

For this, I am grateful for Markus being my advisor. Markus is one of the most intelligent person I know who holds professional integrity and great taste, both scientific and secular. I will greatly miss all the discussions with him. In addition, I thank my Committee Henry, Carlos, and Rob. Drilling down this century-old question with you was rewarding.

My family who supported me wholeheartedly. Close friends Hsiao-Yi and Chien-I understand the good and bad of academia and have gone through this expedition together.

My lab. Kyu taught me most of the techniques required for the projects. Brenna helped me navigate grad school and career choices. Yang, Matt, Zeynep, and Tony were great lab mates that shared ideas and resources generously. Younger colleagues Jieuyu, Jiang, and Dan rejuvenated the dynamics of our lab.

My career mentors and figures to look up to. T-Z, my first scientific mentor taught me to appreciate science through ion channels. Jun-Han and Hao showed me passages and the tactics to approach the alternative M.D. career paths. President Kung-Yi Liang of Yang Ming always looked after my career progress. Dr. Chen-Huan, my dean from Yang Ming created the MD-PhD program and elevated his trainees. President Yin who supported Chen-Huan's vision and funded us graciously.

Med school friends Emily, Isaac, Perry, Stanley, Shang-chin, Yu-Jung, and Yen-Chung. High school friends Travis, Sanhe, Pinchi, Hunguo, Tsunghsien, Vince, and Mengda. Grad school friends Ying, GuaGua, YT, Allen, Sandy, Piraner, Hurt, Jon, Koichiro, Kenny, Howard, Sophia, Tinghao, and Lucas. Having you around is such joy and luck.

ABSTRACT

This thesis presents two tool development projects for neurobiology and one explorative project to find organizing principles for autism.

The first project (Chap 2, *Retina as Probe*) was conceived to tackle the problem that there hasn't been a reliable model system for system-level neuropharmacology. We introduce a testbed for this: the mammalian retina. The retina involves many of the known neurotransmitters and modulators. Yet its synaptic wiring is well understood, and quantitative models exist to explain its input-output functions. One can connect the systems-level effects to the underlying cellular and molecular causes. To demonstrate the retina's use, we explored the effects of a range of general anesthetics on the light responses of the mouse retina. At sub-anesthetic doses, we found that certain anesthetics exert a paradoxical effect: they increase the light response of some retinal neurons and suppress the response of others. Notably, this occurred for alcohols and ketamine but not for isoflurane. We traced these effects to transmitter release at a specific synapse and, in one case, to a specific presynaptic ion channel. All the anesthetics silenced the output of the retina completely at concentrations similar to their effective dose for anesthesia in humans. Sedatives reduced retinal sensitivity but did not silence it. Finally, we used specific drugs that target hypothesized molecular mechanisms to probe how much they each contribute to anesthesia of the retina.

The second project which attempted to probe the principles of autism (Chap 3) was conceptually a direct extension of the retina as a testbed. Similar to the situation in seeking for what the mechanism of general anesthesia is, the field of autism research also lacks a good testbed but for systemically comparing gene mutation - circuit defect - behavior outcomes. Similarly, we utilized the retina as a platform to identify circuit defects in four different autism model mice and followed through the different mouse line's behavior readouts using our lab's maze navigation paradigm. We discovered that the different autism mouse lines varied in the retinal circuits and varied in their navigation preferences. Nevertheless, unlike the anesthetic project, there wasn't a simple mechanism to explain why or how these differences are coupled together.

The last project, *Electrode Pooling*, (Chap 4) aimed to boost the yield of extracellular recording electrode arrays with a novel method we named electrode pooling. The

per-implant yield of extracellular recording leaped significantly from the order of tens to the order of hundreds when engineers built multiple electrode arrays based on silicon technology to replace tetrode wires. Unfortunately, this yield-per-site is already maxed out with modern silicon technology. The constraint of the yield is mainly biological, as explained in the chapter, and thus could not be further advanced by improving the manufacturing processes of semiconductors. Our solution utilized an approach that multiplexed the array recording sites (not the bottleneck) onto the readout wires with accompanying filters (the actual bottleneck). Specifically, the method proposes intelligently choosing many recording sites that carry signals and connecting them to a single wire via manipulating the switches and later un-mixed with a spike-sorting algorithm. We demonstrated the first proof-of-principle study that shows that one could get more single-neuron recordings per implant site with electrode pooling, and made recommendations on the hardware design that could facilitate the advancement of probes that use pooling algorithms.

PUBLISHED CONTENT AND CONTRIBUTIONS

Ni, Yu-Li, Norianne Ingram, Alapakkam Sampath, and Markus Meister (2022). "The Retina as a Neuropharmacology Testbed: Reverse Engineering the Circuit Level Mechanisms of General Anesthesia". In: *Manuscript in Prep*.
Y-L.N. participated in the conception of the project, performed the experiments, analyzed the data, and wrote the manuscript.

Bagherian, Dawna, James Gornet, Jeremy Bernstein, Yu-Li Ni, Yisong Yue, and Markus Meister (2021). "Fine-Grained System Identification of Nonlinear Neural Circuits". In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. KDD '21. Association for Computing Machinery, pp. 14–24. DOI: 10.1145/3447548.3467402.
Y-L.N. performed the experiments.

Lee, Kyu Hyun*, Yu-Li Ni*, Jennifer Colonell, Bill Karsh, Jan Putzeys, Marius Pachitariu, Timothy D. Harris, and Markus Meister (2021). "Electrode pooling can boost the yield of extracellular recordings with switchable silicon probes". In: *Nature Communications* 12.1, p. 5245. DOI: 10.1038/s41467-021-25443-4.
Y-L.N. participated in the conception of the project, performed the experiments, analyzed the data, and wrote the manuscript.

Maresca, David, Anupama Lakshmanan, Audrey Lee-Gosselin, Johan M. Melis, Yu-Li Ni, Raymond W. Bourdeau, Dennis M. Kochmann, and Mikhail G. Shapiro (Feb. 2017). "Nonlinear ultrasound imaging of nanoscale acoustic biomolecules". In: *Applied Physics Letters* 110.7, p. 073704. DOI: 10.1063/1.4976105.
Y-L.N. performed the experiments.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

INTRODUCTION

1.1 Retina as Testbed

1.1.1 Neuropharmacology is built on empirical findings and has not developed a systematic approach like other fields in medicine

Imagine you are a medical history book author who is working on a new publication that compares the progress of each field of treatment. You are most interested in two subfields: *oncology* and *neuro-psychiatry* - after all, cancer is the king of all maladies, and the brain is the most complex part of the human body.

Upon reading ancient works of literature, you found that neuro-related treatments had a transient lead 4 millennia ago. Imhotep and his contemporaries were likely prescribing opioid-like substances for analgesia, but only offered condolences to cancer patients as there was nothing they could do — not even empirical methods were available.

Then fast forward 4000 years to 1970. You noticed the situation has flipped. Cancer was still quite a mystery and was treated based on empirical clinical experience — Sidney Farber and his pupils were testing various combinations of empirically found chemicals — but the great minds of the time were close to coming up with the foundations of the pathological principle of cancer. The central dogma of modern biology linking DNA to protein was established. Temin and Baltimore had their hands on the first solid causal-relationship protein substrate of reverse transcriptase. In comparison, Domino et al. were testing ketamine fully empirically by dosing up inmates and crushing their skin with hemostats, which is conceptually no different from the Imhotep age's clinical development except that now, there are more chemicals to test with.

Fast forward again to the 2020s. You noticed that many new therapies via novel modalities beyond the imagination of Farber were available to various cancers that were still untreatable at the turn of the Millenium. In contrast, almost all new pipelines for depression, schizophrenia, and Alzheimer's have failed. Pharma and biotech are laying off neuro programs one after another.

You interviewed a few experts and realized fundamentally, neuropharmacology

hasn't developed core principles and testbeds that have predictive values like oncology's. It is not because developers don't want it – but that neuro-psychiatry functions in circuits – thus the leap in molecular biology which focuses on cellular level modulations largely does not apply to neuro's uses. Instead, clinical research in this field relied on pathological animal models that tracked some features of the disease state - and empirically tried to fix the animal by the best-guess treatments which so far seem to have no predictive power.

1.1.2 Retina for reverse engineering neuropharmacology: From the rationale to setup

What makes a good experimental platform to study these links between molecular activity and systems function? One would like to work on a neural circuit complex enough to exhibit all the hypothesized targets of GA drugs. The system should have a clear functional readout that serves as a proxy for anesthesia. And it should be experimentally accessible, allowing the rapid introduction and withdrawal of drugs. Most importantly, this circuit's cell types and connectivity should be well understood, so one can trace any changes in system function back to cells and synapses. Here we argue that the mammalian retina can serve this role.

The retina is a complex circuit of neurons in the back of the eye. Its task is to convert light into a neural signal, filter and process this neural image, and convey the result to the brain through the optic nerve. The circuit includes about 100 different cell types that fall into 5 major classes defined by synaptic connectivity (Figure 2.1 upper). Much progress has been made to identify the ion channels, receptors, transmitters, and modulators used by retinal neurons. The basic synaptic circuit leading from photoreceptors to retinal ganglion cells is well understood. Indeed, there exist computational models of this circuitry that successfully predict the firing of the output neurons based on the visual image presented at the input.

Experimental access to the retina is superb: the circuit can be explanted intact from the eye and continues to function in a dish for many hours or even days. This allows complete optical access, for example, to present visual stimuli, to visualize fluorescent neurons of a specific cell type, or record their activity by calcium imaging. By placing the retina on a flat multi-electrode array one can record simultaneously the spike trains of many of its output neurons. By delivering visual stimuli and monitoring their transformation into spike trains one obtains a sensitive measure of the system's function. Drugs can be delivered within seconds

by superfusing the retina in the dish, and the resulting effects on circuit function can be recorded instantaneously.

We begin by defining a proxy for anesthesia: the visual response of retinal ganglion cells. We show that this activity gets silenced by diverse GA drugs at the same concentrations which leads to loss of consciousness. Then we explore some commonly-used drugs across a wide series of concentrations. We report a paradoxical effect induced by the alcohols and ketamine but not isoflurane. Taking advantage of the retina's known blueprint, we tie these effects to a specific presynaptic site. Finally, we use some non-GA drugs with known molecular targets to evaluate alternative hypotheses for the action mechanism of ethanol and ketamine. Based on this experience we suggest that the retina will prove a valuable testbed for a broader research program on the systems-level function of general anesthesia.

1.1.3 Probing mechanism of general anesthesia

General Anesthesia (GA) is an integral part of medicine. Every day, hundreds of thousands of people undergo a pharmacologically controlled loss of consciousness during a surgical procedure. Despite the enormous practical utility of anesthesia and the long history of its use dating back to the 19th Century, some deep questions remain about its biological mechanism.

One mystery lies in the extreme diversity of the class of chemicals that act as general anesthetics, ranging from single-atom noble gases to complex molecules with many ring structures. Early workers noted that despite this structural diversity, the efficacy of a general anesthetic is directly proportional to its lipid solubility. The relationship is not perfect, but it accounts for a huge fraction of the explainable variance over many orders of magnitude of solubility. This simple empirical relationship between a complex brain-level effect – loss of consciousness – and a simple molecular property – the lipid partition coefficient – hinted that there might be a single biophysical explanation for general anesthesia. Unfortunately, such a unifying mechanism has so far failed to materialize.

Instead, recent research has focused on how diverse the anesthetic drugs are in their chemical activities, and especially their interaction with various candidate target molecules in the nervous system. Many drugs in the GA class have been associated with neurotransmitter receptors or ion channels that control the polarization state of nerve cells. For instance, ketamine, one of the best-studied GAs, is thought to block NMDA-receptors, a type of channel that activates neurons (Anis et al., 1983).

But it is also thought to block HCN-channels, leading to silence neurons (X. Chen, Shu, and Bayliss, 2009). One would like to understand now how the systems-level phenomenon of anesthesia is related to these molecular-level interactions. The challenge, as illustrated with ketamine, is that the same drug may have multiple activities and that the presumed targets of the drug are distributed through many stages of the nervous system.

1.2 Autism

1.2.1 Current understanding of the autism spectrum disorder

Autism spectrum disorder (ASD) is a developmental disorder thought to derive from the dysfunction of complex neuronal circuits. It has a strong genetic component, and there is great interest in understanding the links between autism-related genes and circuit-level dysfunctions. A large body of research has tackled this domain, leading to a wide range of mechanistic proposals for circuit malfunctions: from developmental abnormalities in the long-range connection between brain areas to the imbalance between excitation and inhibition in local circuits (J. A. Chen et al., 2015)

1.2.2 Rationale of our approach

Mouse models of ASD have played an important role in uncovering such mechanistic relationships. Of particular interest are mouse lines created to carry genetic mutations that are associated with ASD in humans. It is remarkable that many of these mutant mice also show behavioral abnormalities reminiscent of human autism, which spurs some confidence that common mechanisms are at work in the mouse and human brains. The scientific challenge in finding those mechanisms results from the fact that genes and brain function are separated by so many explanatory levels: protein expression, cell biology of neurons and glia, the development of synaptic connectivity, the signaling dynamics of individual neurons, and the collective function of synaptic networks. A typical research study of mutant mice will probe in a directed fashion for abnormalities at one or more of these levels, with the hope that any such observed defects can explain the change in system function. In this enterprise, much of the focus has been on cortical areas, suspected to be the seat of the higher functions that appear perturbed in ASD. Here we propose an alternative and complementary approach.

The retina is one neural system in which the links between genes, proteins, circuits, and neural function are particularly well understood already. In complexity, this

network rivals the rest of the central nervous system. The retina contains ~ 100 different types of neurons, many of which are defined quantitatively and genetically accessible (Zeng and Sanes, 2017). It includes >30 different microcircuits (Sanes and Masland, 2015) — one for each type of output neuron -– and these comprise every imaginable circuit motif (Gollisch and Meister, 2010), using most every neurotransmitter and modulator known to neuroscience (Hoon et al., 2014). Nevertheless, there are now neural circuits leading through the retina, from photoreceptors to ganglion cells, where all the important interneurons and their synapses are known (Gollisch and Meister, 2010; Helmstaedter et al., 2013; Krishnaswamy et al., 2015). Mathematical models of retinal circuits can predict the visual responses of the output neurons in quantitative detail (Baccus et al., 2008). The historical reasons for this depth of understanding compared to other parts of the brain lie in the unusual experimental control one has over the retina, with complete access to its inputs and outputs.

On this background, we propose to screen ASD mouse models for abnormalities in the function of the retina. Our guiding hypothesis is that any mutation or environmental perturbation that causes defects in neuronal circuits will leave a trace in the functions of the retina. As detailed below, such measurements can be extremely sensitive: there are many retinal microcircuits available to test, each with a different configuration of synapses. The toolkit of visual stimuli is virtually unlimited and can be tailored for specific hypotheses. Finally, the neural response measurements are very reproducible across animals, allowing even small effects to be resolved. Once one finds an abnormality in system function – in this case the processing of visual inputs – one can quickly dissect the dysfunction and identify the underlying synaptic and cellular defects. In homing into home in on the likely sources one gets invaluable help from the extensive existing knowledge of retinal microcircuitry. These mechanistic tests will include the classical tools of immunocytochemistry and pharmacology. Effectively the retina can serve as a test bed in which one develops circuit-level explanations that can then be tested in brain areas more likely to be causally involved in ASD. Building on the results of the retina findings, we took advantage of newly emerging opportunities to broaden the scope of the comparative analysis. Specifically, we studied the behavior of mice from these same lines in several natural tasks that are part of the essential rodent repertoire: escape from predators, exploration of an environment, learning, and navigation in space. The aim was to identify differences in these behaviors among the mutants and with respect to the wild-type mouse and correlate any such differences with the class of

phenomena found in the circuit-level analysis of the retina.

1.3 Electrode Pooling

1.3.1 The need for extracellular recording methods

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" (BRAIN Working Group, 2014). 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 (Jun et al., 2017; Dimitriadis et al., 2018; Rios et al., 2016; Torfs et al., 2010; Steinmetz et al., 2021). Still, the best methods sample neural circuits very sparsely, for example recording fewer than 10^4 cells in a mouse brain that has 10^8 (Stevenson, 2013).

1.3.2 Limitations and previous attempts of improving extracellular recording

In many of these electrode array 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 (Jun et al., 2017). The current silicon arrays invariably displace more neurons than they record, and thus the goal of "monitoring all neurons" seems unattainable by simply scaling the present approach (but see Kleinfeld et al., 2019). 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.

A common approach by which a single wire can convey multiple analog signals is time-division multiplexing (Obien et al., 2015). A rapid switch cycles through the N input signals and connects each input to the output line for a brief interval (Figure 4.1a). At the other end of the line, a synchronized switch demultiplexes 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 (Shannon, 1949): it should be at least twice the highest frequency component present in the signal. The raw voltage signals from extracellular electrodes include thermal noise that extends 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 (Shahrokhi et al., 2010), which would force the electrodes apart and prevent any high-density recording. 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 (Eversmann et al., 2003). A recent design with a more modest $N = 8$ still produces noise power 4-15 times higher than in comparable systems without multiplexing (Raducanu et al., 2017).

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, heating up exactly the neurons one wants to monitor. Furthermore, the active electronics in the local amplifier are sensitive to light, which can produce artifacts during bright light flashes for optogenetic stimulation (Jun et al., 2017; Kozai and Vazquez, 2015).

An alternative approach involves static electrode selection (Figure 4.1b). 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 space is less constrained. 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 (Müller et al., 2015) as well as silicon prongs (Jun et al., 2017; Lopez et al., 2017; Steinmetz et al., 2021). It allows the user to choose one of many electrodes intelligently, for example because it carries

a strong signal from a neuron of interest. This strategy can increase the yield of neural recordings, but it does not increase the number of neurons per wire.

1.3.3 Innovation: Electrode pooling

On this background, we introduce a third method of mapping electrodes to wires: select multiple electrodes with suitable signals and connect them to the same wire (Figure 4.1c). 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 averaged and transmitted on the same wire. At first, that 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 unmixed if one controls the switch settings carefully during the recording session. Under suitable conditions, this method can record many neurons per wire without appreciable loss of information.

We emphasize that the ideal electrode array device to implement this recording method does not yet exist. It would be entirely within reach of current fabrication capabilities, but every new silicon probe design requires a substantial investment and consideration of various trade-offs. With this article we hope to make the community of electrode users aware of the opportunities in this domain and start a discussion about future array designs that use intelligent electrode switching, adapted to various applications in basic and translational neuroscience.

References

Anis, N.a. et al. (1983). "The dissociative anaesthetics, ketamine and phencyclidine, selectively reduce excitation of central mammalian neurones by N-methyl-aspartate". en. In: *British Journal of Pharmacology* 79.2, pp. 565–575. ISSN: 1476-5381. DOI: 10.1111/j.1476-5381.1983.tb11031.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1476-5381.1983.tb11031.x (visited on 02/17/2022).

Baccus, Stephen A. et al. (2008). "A Retinal Circuit That Computes Object Motion". In: *Journal of Neuroscience* 28.27, pp. 6807–6817. DOI: 10.1523/jneurosci.4206-07.2008. URL: http://www.jneurosci.org/content/28/27/6807.abstract.

BRAIN Working Group (2014). *BRAIN 2025: A Scientific Vision*. en. Tech. rep., p. 146. URL: https://braininitiative.nih.gov/strategic-planning/brain-2025-report (visited on 11/04/2019).

Chen, Jason A. et al. (2015). "The Emerging Picture of Autism Spectrum Disorder: Genetics and Pathology". In: *Annual Review of Pathology: Mechanisms of Disease* 10.1, pp. 111–144. DOI: 10.1146/annurev-pathol-012414-040405.

Chen, Xiangdong, Shaofang Shu, and Douglas A. Bayliss (Jan. 2009). "HCN1 Channel Subunits Are a Molecular Substrate for Hypnotic Actions of Ketamine". en. In: *Journal of Neuroscience* 29.3. Publisher: Society for Neuroscience Section: Articles, pp. 600–609. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3481-08.2009. URL: https://www.jneurosci.org/content/29/3/600 (visited on 05/20/2021).

Dimitriadis, George et al. (Mar. 2018). "Why not record from every channel with a CMOS scanning probe?" en. In: *bioRxiv*. DOI: 10.1101/275818. URL: http://biorxiv.org/lookup/doi/10.1101/275818 (visited on 08/28/2019).

Eversmann, B. et al. (Dec. 2003). "A 128 x 128 cmos biosensor array for extracellular recording of neural activity". en. In: *IEEE Journal of Solid-State Circuits* 38.12, pp. 2306–2317. ISSN: 0018-9200. DOI: 10.1109/JSSC.2003.819174. URL: http://ieeexplore.ieee.org/document/1253878/ (visited on 08/27/2019).

Gollisch, Tim and Markus Meister (2010). "Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina". In: *Neuron* 65.2, pp. 150–164. ISSN: 08966273. DOI: 10.1016/j.neuron.2009.12.009. URL: http://www.cell.com/neuron/abstract/S0896-6273(09)00999-4.

Helmstaedter, Moritz et al. (2013). "Connectomic reconstruction of the inner plexiform layer in the mouse retina". In: *Nature* 500.7461, pp. 168–174. DOI: 10.1038/nature12346http://www.nature.com/nature/journal/v500/n7461/abs/nature12346.html#supplementary-information. URL: http://dx.doi.org/10.1038/nature12346.

Hoon, Mrinalini et al. (Sept. 2014). "Functional Architecture of the Retina: Development and Disease". In: *Progress in retinal and eye research* 42, pp. 44–84. ISSN: 1350-9462. DOI: 10.1016/j.preteyeres.2014.06.003. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4134977/ (visited on 09/15/2022).

Jun, James J. et al. (2017). "Fully integrated silicon probes for high-density recording of neural activity". In: *Nature* 551.7679, pp. 232–236. ISSN: 14764687. DOI: 10.1038/nature24636. URL: http://dx.doi.org/10.1038/nature24636.

Kleinfeld, David et al. (Sept. 2019). "Can One Concurrently Record Electrical Spikes from Every Neuron in a Mammalian Brain?" en. In: *Neuron* 103.6, pp. 1005–1015. ISSN: 08966273. DOI: 10.1016/j.neuron.2019.08.011. URL: https://linkinghub.elsevier.com/retrieve/pii/S0896627319306956 (visited on 11/05/2019).

Kozai, Takashi D. Y. and Alberto L. Vazquez (June 2015). "Photoelectric Artefact from Optogenetics and Imaging on Microelectrodes and Bioelectronics: New Challenges and Opportunities". en. In: *Journal of Materials Chemistry B* 3.25, pp. 4965–4978. ISSN: 2050-7518. DOI: 10.1039/C5TB00108K.

Krishnaswamy, Arjun et al. (Aug. 2015). "Sidekick 2 directs formation of a retinal circuit that detects differential motion". eng. In: *Nature* 524.7566, pp. 466–470. ISSN: 1476-4687. DOI: 10.1038/nature14682.

Lopez, Carolina Mora et al. (2017). "A Neural Probe With Up to 966 Electrodes and Up to 384 Configurable Channels in 0.13 µm SOI CMOS". In: *IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS* 11.3, pp. 510–522.

Müller, Jan et al. (2015). "High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels". en. In: *Lab on a Chip* 15.13, pp. 2767–2780. ISSN: 1473-0197, 1473-0189. DOI: 10.1039/C5LC00133A. URL: http://xlink.rsc.org/?DOI=C5LC00133A (visited on 11/05/2019).

Obien, Marie Engelene J. et al. (Jan. 2015). "Revealing neuronal function through microelectrode array recordings". en. In: *Frontiers in Neuroscience* 8. ISSN: 1662-453X. DOI: 10.3389/fnins.2014.00423. URL: http://journal.frontiersin.org/article/10.3389/fnins.2014.00423/abstract (visited on 11/05/2019).

Raducanu, Bogdan C. et al. (Oct. 2017). "Time Multiplexed Active Neural Probe with 1356 Parallel Recording Sites". en. In: *Sensors* 17.10, p. 2388. DOI: 10.3390/s17102388.

Rios, Gustavo et al. (Nov. 2016). "Nanofabricated Neural Probes for Dense 3-D Recordings of Brain Activity". en. In: *Nano Letters* 16.11, pp. 6857–6862. ISSN: 1530-6984, 1530-6992. DOI: 10.1021/acs.nanolett.6b02673. URL: https://pubs.acs.org/doi/10.1021/acs.nanolett.6b02673 (visited on 08/28/2019).

Sanes, Joshua R. and Richard H. Masland (July 2015). "The types of retinal ganglion cells: current status and implications for neuronal classification". eng. In: *Annual Review of Neuroscience* 38, pp. 221–246. ISSN: 1545-4126. DOI: 10.1146/annurev-neuro-071714-034120.

Shahrokhi, Farzaneh et al. (June 2010). "The 128-Channel Fully Differential Digital Integrated Neural Recording and Stimulation Interface". en. In: *IEEE Transactions on Biomedical Circuits and Systems* 4.3, pp. 149–161. ISSN: 1932-4545, 1940-9990. DOI: 10.1109/TBCAS.2010.2041350. URL: http://ieeexplore.ieee.org/document/5471738/ (visited on 11/05/2019).

Shannon, C. E. (Jan. 1949). "Communication in the Presence of Noise". In: *Proceedings of the IRE* 37.1, pp. 10–21. ISSN: 0096-8390. DOI: 10.1109/JRPROC.1949.232969.

Steinmetz, Nicholas A. et al. (Apr. 2021). "Neuropixels 2.0: A Miniaturized High-Density Probe for Stable, Long-Term Brain Recordings". en. In: *Science* 372.6539. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.abf4588.

Stevenson, Ian (Oct. 2013). *Tracking Advances in Neural Recording | Statistical Neuroscience Lab*. en-US. https://stevenson.lab.uconn.edu/scaling/.

Torfs, Tom et al. (Nov. 2010). "Two-dimensional multi-channel neural probes with electronic depth control". en. In: *2010 Biomedical Circuits and Systems Conference (BioCAS)*. Paphos, Cyprus: IEEE, pp. 198–201. ISBN: 978-1-4244-7269-7. DOI: 10.1109/BIOCAS.2010.5709605. URL: http://ieeexplore.ieee.org/document/5709605/ (visited on 08/28/2019).

Zeng, Hongkui and Joshua R. Sanes (Sept. 2017). "Neuronal cell-type classification: challenges, opportunities and the path forward". eng. In: *Nature Reviews. Neuroscience* 18.9, pp. 530–546. ISSN: 1471-0048. DOI: 10.1038/nrn.2017.85.

Chapter 2

THE RETINA AS A NEUROPHARMACOLOGY TESTBED: REVERSE ENGINEERING THE CIRCUIT LEVEL MECHANISMS OF GENERAL ANESTHESIA

Ni, Yu-Li, Norianne Ingram, Alapakkam Sampath, and Markus Meister (2022). "The Retina as a Neuropharmacology Testbed: Reverse Engineering the Circuit Level Mechanisms of General Anesthesia". In: *Manuscript in Prep*.

2.1 Abstract

General anesthetics have been widely used since the 19th century and are an essential part of medicine. This is a remarkably heterogeneous class of neuroactive drugs, ranging from noble gases to complex biomolecules. Much has been learned about the molecular targets of these drugs, yet there remain open questions about the systems-level mechanisms of general anesthesia. Here we introduce a systems-level testbed for these biological mechanisms: the mammalian retina. This complex neural circuit involves many of the known neurotransmitters and modulators. Yet its synaptic wiring is well understood and quantitative models exist to explain its input-output functions, so one can connect the systems-level effects to the underlying cellular and molecular causes. We explored the effects of a range of general anesthetics on the light responses of the mouse retina. At sub-anesthetic doses we found that certain anesthetics exert a paradoxical effect: they increase the light response of some retinal neurons and suppress the response of others. Notably, this occurred for alcohols and ketamine, but not for isoflurane. We traced these effects to transmitter release at a specific synapse, and in one case to a specific presynaptic ion channel. All the anesthetics silenced the output of the retina completely at concentrations similar to their effective dose for anesthesia in humans. Sedatives reduced retinal sensitivity but did not silence it. Finally, we used specific drugs that target hypothesized molecular mechanisms to probe how much they each contribute to anesthesia of the retina.

2.2 Introduction

General Anesthesia (GA) is an integral part of medicine. Every day, hundreds of thousands of people undergo a pharmacologically controlled loss of consciousness

during a surgical procedure. Despite the enormous practical utility of anesthesia, and the long history of its use dating back to the 19th Century, some deep questions remain about its biological mechanism.

One mystery lies in the extreme diversity of the class of chemicals that act as general anesthetics, ranging from the single-atom noble gases to complex molecules with many ring structures. Early workers noted that despite this structural diversity, the efficacy of a general anesthetic is directly proportional to its lipid solubility. The relationship is not perfect, but it accounts for a huge fraction of the explainable variance over many orders of magnitude of solubility. This simple empirical relationship between a complex brain-level effect - loss of consciousness - and a simple molecular property - the lipid partition coefficient - hinted that there might be a single biophysical explanation for general anesthesia. Unfortunately, such a unifying mechanism has so far failed to materialize.

Instead, recent research has focused on how diverse the anesthetic drugs are in their chemical activities, and especially their interaction with various candidate target molecules in the nervous system. Many drugs in the GA class have been associated with neurotransmitter receptors or ion channels that control the polarization state of nerve cells. For instance, ketamine, one of the best studied GAs, is thought to block NMDA-receptors, a type of channel that activates neurons (Anis et al., 1983). But it is also thought to block HCN-channels, leading to silence neurons (Chen, Shu, and Bayliss, 2009). One would like to understand now how the systems-level phenomenon of anesthesia is related to these molecular-level interactions. The challenge, as illustrated with ketamine, is that the same drug may have multiple activities, and that the presumed targets of the drug are distributed through many stages of the nervous system.

What makes a good experimental platform to study these links between molecular activity and systems function? One would like to work on a neural circuit complex enough to exhibit all the hypothesized targets of GA drugs. The system should have a clear functional readout that serves as a proxy for anesthesia. And it should be experimentally accessible, allowing the rapid introduction and withdrawal of drugs. Most importantly, this circuit's cell types and connectivity should be well understood, so one can trace any changes in system function back to cells and synapses. Here we argue that the mammalian retina can serve this role.

The retina is a complex circuit of neurons in the back of the eye. Its task is to convert light into a neural signal, to filter and process this neural image, and to convey the

result to the brain through the optic nerve. The circuit includes about 100 different cell types that fall into 5 major classes defined by synaptic connectivity (Figure 2.1 Upper). Much progress has been made to identify the ion channels, receptors, transmitters, and modulators used by retinal neurons. The basic synaptic circuit leading from photoreceptors to retinal ganglion cells is well understood. Indeed, there exist computational models of this circuitry that successfully predict the firing of the output neurons based on the visual image presented at the input.

Experimental access to the retina is superb: The circuit can be explanted intact from the eye and continues to function in a dish for many hours or even days. This allows complete optical access, for example to present visual stimuli, or to visualize fluorescent neurons of a specific cell type, or to record their activity by calcium imaging. By placing the retina on a flat multi-electrode array one can record simultaneously the spike trains of many of its output neurons. By delivering visual stimuli and monitoring their transformation into spike trains one obtains a sensitive measure of the system's function. Drugs can be delivered within seconds by superfusing the retina in the dish, and the resulting effects on circuit function can be recorded instantaneously.

We begin by defining a proxy for anesthesia: The visual response of retinal ganglion cells. We show that this activity gets silenced by diverse GA drugs at the same concentrations that leads to loss of consciousness. Then we explore some commonly used drugs across a wide series of concentrations. We report a paradoxical effect induced by the alcohols and ketamine but not isoflurane. Taking advantage of the retina's known blueprint, we tie these effects to a specific presynaptic site. Finally we use some non-GA drugs with known molecular targets to evaluate alternative hypotheses for the action mechanism of ethanol and ketamine. Based on this experience we suggest that the retina will prove a valuable testbed for a broader research program on the systems-level function of general anesthesia.

2.3 Results

2.3.1 The retina's light response as a proxy for anesthesia

The mouse retina was isolated from the eye, then flat-mounted on a multi-electrode array with the ganglion cell layer facing the electrodes (Fig 2.1). A 256-channel acquisition system was used to record spike trains from many retinal ganglion cells. Meanwhile a digital light projector delivered visual stimuli through an optical system focused on the photoreceptor layer. A gravity superfusion system delivered

oxygenated Ames's solution to the retina. Most of the tested drugs were dissolved in this Ames's solution at the specified concentrations. Isoflurane, which evaporates rapidly, was bubbled through the Ames's solution continuously using an anesthesia machine.

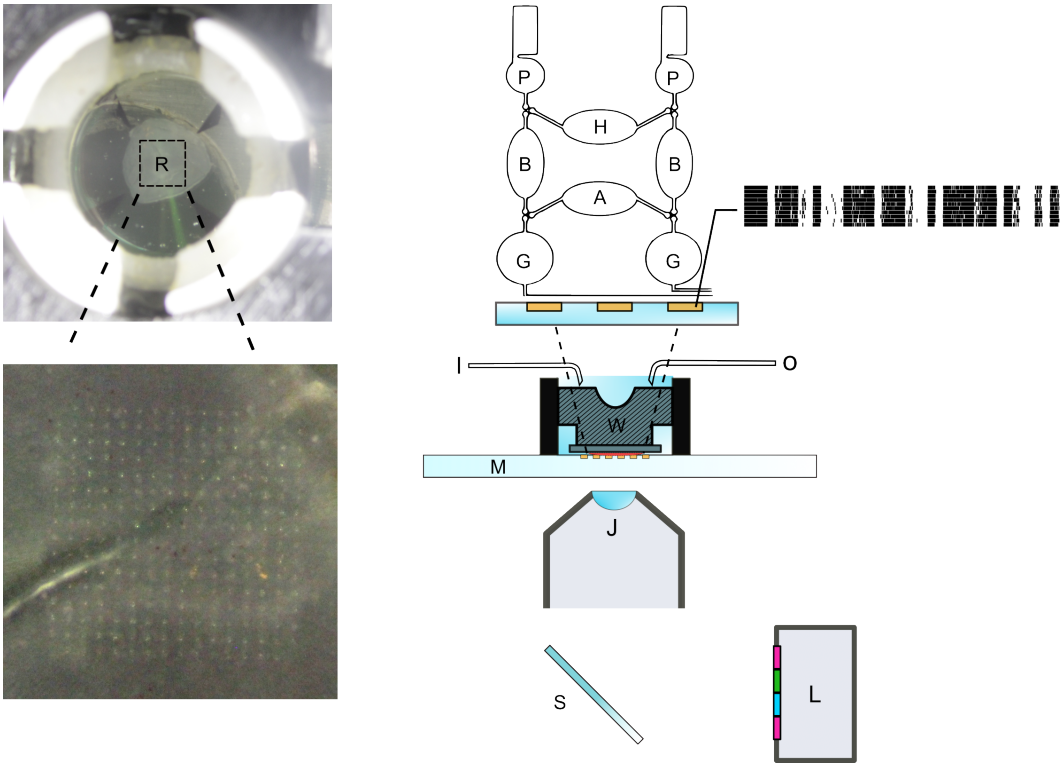

Figure 2.1: **Retinal recording with a multi-electrode array**. **Right:** Schematic of retinal circuitry with photoreceptors (P), bipolar (B), horizontal (H), amacrine (A), and ganglion (G) cells. The retina is placed in a superfusion chamber on top of a multi-electrode array (M). A customized widget (W) ensures contact with the electrodes, and Ames's solution flows through the chamber via glass tubes at the inlet (I) and outlet (O). A digital light projector (L) produces visual stimuli through a mirror (S) and objective lens (J) focused on the photoreceptor (P) layer. **Left:** Top view of the retina in the superfusion chamber (**top**), and enlargement showing the array of electrodes under the tissue (**bottom**). The electrode spacing is 60 μm. **Bottom:** An example of spike trains recorded from one ganglion cell. Each row represents a repeat of the same visual stimulus (see bottom trace). Every tick mark is an action potential. The firing rate curve shows the mean spike count across all the trials.

We stimulated the retina with full-field flashes, alternating light and dark periodically. Each neuron responds to the visual input with a large modulation in firing rate, precisely timed to the light transitions. The firing rate typically peaks within 100 - 200 ms of the excitatory phase of the flash, then decays to a lower value. During the other half of the stimulus cycle, firing is strongly inhibited. The response is very

reproducible across repeats of the same stimulus, ensuring that even small changes over time will be detected reliably.

Then we exposed the retina to increasing concentrations of general anesthetics. As seen in Figure 2.2, the light response of retinal ganglion cells declined and eventually all firing ceased. At intermediate concentrations, the firing rate profile did not simply scale proportionally: Instead the response became more transient, suggesting that only the initial peak in firing was able to break through the suppression. Cells with a more sustained response to the flash required a somewhat higher concentration of the drug for complete silencing.

As a summary measure, we plotted the cummulative spikes within the flash response segments of ganglion cells against the concentration of the drug (Fig 2.2). For all the general anesthetics tested, the response dropped off very steeply over a narrow range of concentrations. Typically the 20-80% change occurred over a factor of 2-4 in concentration. The effective concentration in the retina, namely the dose for half-maximal response suppression, was closely related to the effective dose for anesthesia. For illustration we plot our results from retinal responses along with a classic Meyer-Overton curve based on righting reflexes in animals. Note, for example, that the effective dose of the aliphatic alcohols varies with chain length in the same way as observed for their anesthetic action.

We also tested midazolam, a benzodiazepine commonly used as a sedative. Unlike many of the general anesthetics, midazolam's effect on patients has a gentle concentration dependence with lower risk of overdose. It does not block pain even at deep sedation nor does it shut down breathing centers. Under midazolam the ganglion cell flash response was gently suppressed, with lowered peaks and plateaus of the firing rate. However, these effects saturated at a drug concentration of $\sim 1\ \mu$M (Fig 2.2), similar to the concentration that produces sedation in humans ($\sim 0.5\mu$M). The ganglion cells still responded robustly to light even at 10 μM.

In summary, all the general anesthetics we tested fully inhibited the firing of retinal ganglion cells within a narrow window of concentration. In contrast, the sedative midazolam failed to silence the retina even at concentrations far beyond the clinically relevant dose for sedation. The close correspondence between the concentration-dependence we observed on retinal responses and the observations on human anesthesia in the clinic suggests that light responses of the retina may serve as a useful testbed for anesthesia.

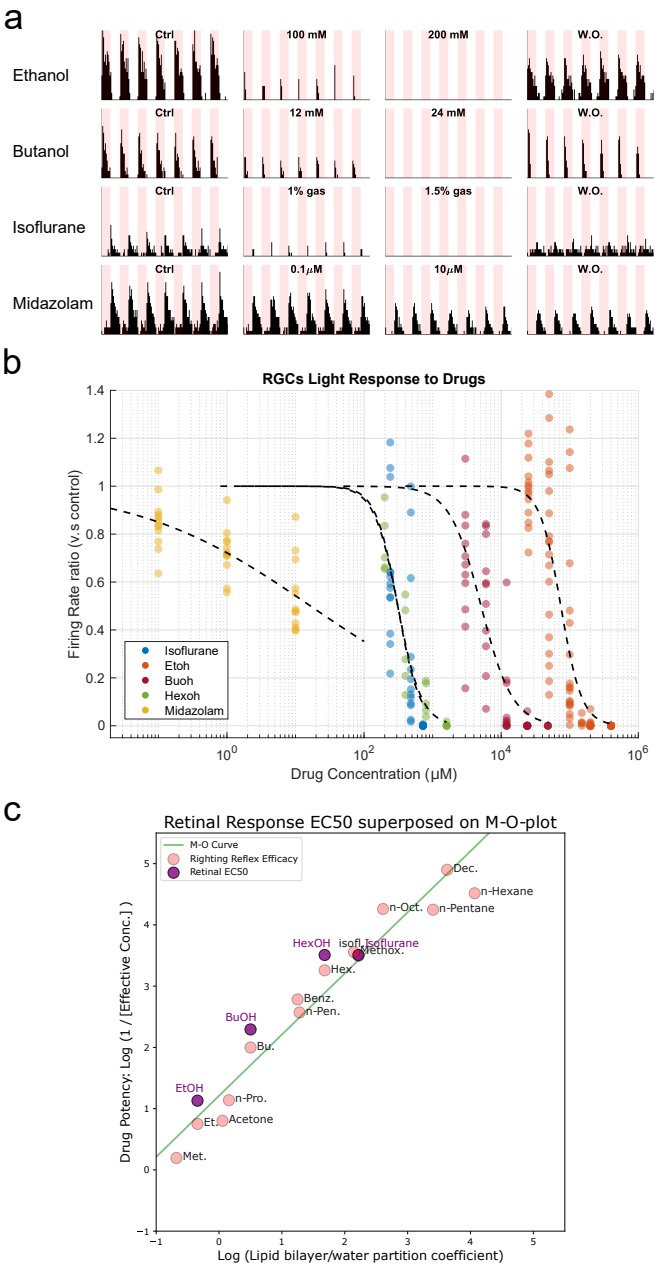

Figure 2.2: **The Retinal Test-bed reflects known macroscopic and microscopic observations**. **(a)** Examples of Retinal ganglion cells (RGCs) firing changes to flashes stimulation in different drugs and concentrations. Repeats of 1-sec whole field bright (pink stripes) then 1-sec dark flashes (white) was projected to stimulated the retina. Several alcohols and Isoflurane fully inhibit the retina whereas midazolam do not even at extremely high concentration. **(b)** Summary of the flash-firing rate dose responses. Firing rate ratios (Y-axis) were derived by normalizing firing rates in the drugs to the corresponding controls. Hill Equation fits were superposed as dotted lines. **(c)** Meyer-Overton (M-O) correlation superposed with the retinal derived "M-O like" correlations. The Y-axis is the potency measured as $\log_{10}(1/EC50)$, where the EC50 is the halved-firing rate concentration derived from Hill Equation fits (dotted lines in panel **(b)**). X-axis is the lipophilicity, measured as partition coefficient $\log_{10}(lipid/water)$. The M-O correlation was adapted from (R. Miller, 2015).

2.3.2 RGCs' firing rate increased paradoxically under ethanol and ketamine, and this phenomenon is coupled with ON-OFF pathway

We've learned that GAs fully inhibited retinal circuits at concentrations that track the dose-response of *loss of righting reflex*. Nevertheless, when the circuits were fully inhibited, the circuit also cease to provide meaningful Input-Output responses for probing mechanisms. Therefore, we then focused on clinically relevant - high clinical concentrations where the retina is still light-responsive.

To our surprise, not all anesthetics behave like simple neural depressants. Amongst the tested GAs, ketamine and ethanol excited the RGCs with increased firing in a window of low-intermediate concentrations (See example cells in Fig 2.3 a, b, leftmost panel). Nevertheless, when we further increased the dosage, these two drugs still suppressed the firing of the ganglion cells and eventually fully inhibited RGC activities at anesthetic concentrations.

After spike-sorting, we learned that actually only around half of the ganglion cells in the intermediate concentrations have increased firing rate under ketamine and ethanol (Fig 2.3 a, b, rightmost panel). The other half of retinal ganglion cells are monotonically inhibited with decreased firing rate across dosage. Interestingly, this dichotomy of opposite firing rate change was strongly coupled to the RGC's ON-OFF selectivity. To be specific, most cells that paradoxically increased activity under low-intermediate dose drugs were ON cells and vice versa (Fig 2.3 d, e).

In contrast, isoflurane and midazolam inhibited both ON and OFF pathways in all the concentrations tested (Fig 2.3 c, f).

2.3.3 The ON-OFF Dichotomy could be used to probe presynaptic v.s postsynaptic actions of drugs

The retina duplicated its signal pathway that encodes the same visual information into two opposite copies, namely the ON and OFF pathways. As the name implies, increased light intensity excites the ON ganglion cell, and inhibits the OFF ganglion cells. The visual signals split into ON and OFF pathways respectively at the photoreceptor-bipolar cell synapse (see Fig 2.4 a).

The photoreceptor synapses release glutamate vesicles steadily in the dark. When the photoreceptor "sees" photons, a series of signal cascades hyperpolarize the photoreceptor and reduce its vesicle release, leading to the reduced activity of the OFF bipolar cell and the OFF Ganglion cells downstream. In contrast, the ON pathway gets excited through a special sign-inverting synapse via metabotropic

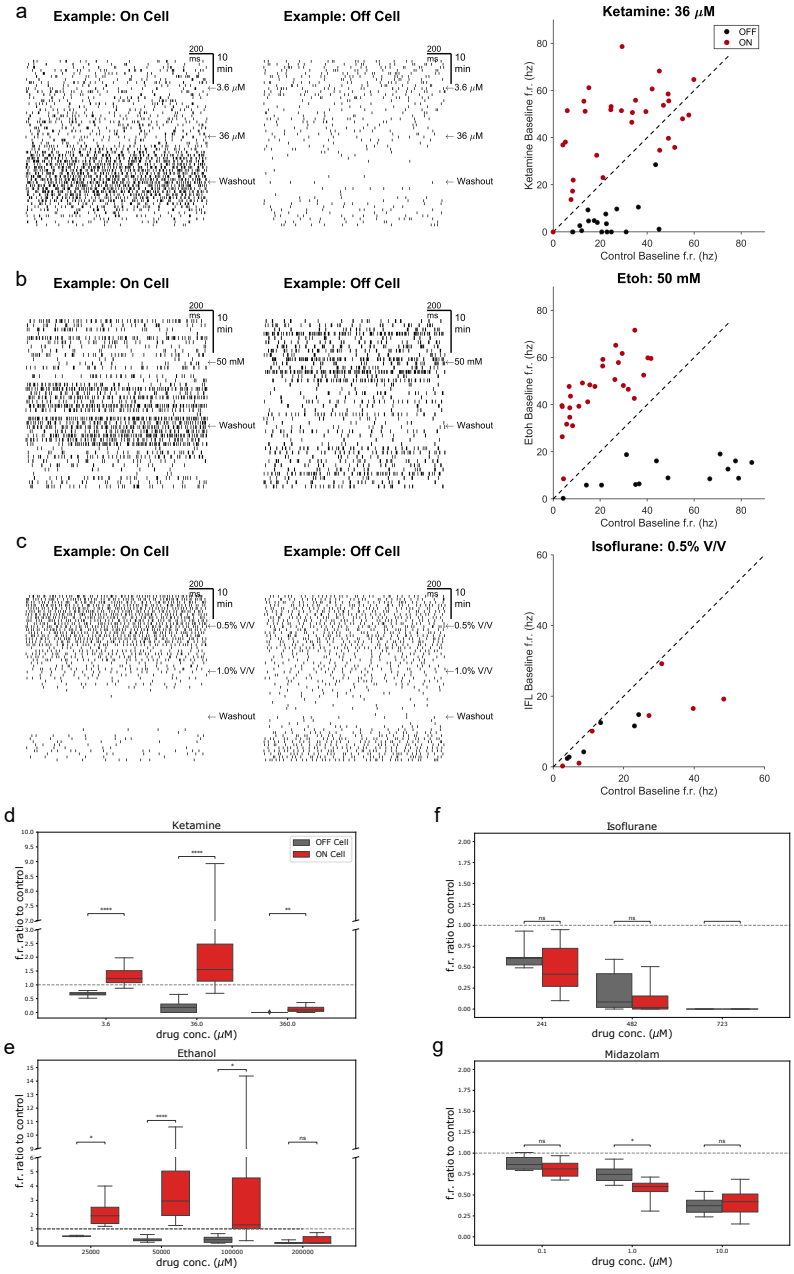

Figure 2.3: **Paradoxical firing rate changes at clinically relevant drug concentrations.** **(a)** Baseline firing change of ganglion cells in ketamine. On cells increased firing (left raster) whereas Off cells decreased firing (right raster) mildly at 3.6 μM then strongly at 36 μM. **(b)** Similar trend of firing rate changes for On vs Off cells in 50 mM Ethanol. **(c)** Isoflurane inhibited both On and Off pathways. **(a, b, c), right panels**. Firing rate of *ctrl* vs *drug* of all recorded RGCs. **(d)** Ketamine trials. At sub-anesthetic concentrations we observed the paradoxical firing rate change. At high concentrations, both On and Off cells are inhibited. **(e)** Ethanol Trials. We observed a similar paradoxical-then-inhibition trend. **(f,g)** Isoflurane and Midazolam inhibited both pathways with a monotonic dose response. Most cells are still firing at a rate around half of the control at an extreme dose of midazolam (10 μM) in **(g)**. **(d)** Ketamine: (OFF, ON cells counts): 3.6 μM, n= (11, 16), 36 μM, n= (18, 31), 360 μM, n= (7, 15). **(e)** Ethanol: 2.5e04 μM n= (4, 6), 5.0e04 μM n= (14, 28), 1.0e05 μM n= (8, 11), 2.0e05 μM n= (8, 11). **(f)** Isoflurane: 241 μM n= (6, 7), 482 μM n= (6, 7), 723 μM n= (7, 5). **(g)** Midazolam: 0.1 μM n= (5, 7), 1.0 μM n= (5, 7), 10 μM n= (5, 7). Box plots show the quartiles of the data. Whiskers extend to the full distribution. Statistics: *two-sided Welch's t-test* calculated with *scipy*. p-val: * : p ≤ 5.0e-02, **: p ≤ 1.0e-02, ***: p ≤ 1.0e-03, ****: p ≤ 1.0e-04, ns: not significant, No annotation: var= 0.

glutamate receptors. (Note that the ON and OFF pathways use inhibitory and excitatory synapses in the photoreceptor-bipolar connection respectively, fig 2.4 a.) The circuit downstream to the photoreceptor-bipolar synapse then stayed relatively symmetric. This implied that the dichotomy of the ON vs OFF cell firing rate changes originated from mechanisms at or before the signal-splitting synapse.

Thus, drugs that reduce synaptic strength (Simplified as "Presynaptic" actions in the remaining text) would mimic seeing photons and enhance the ON-OFF dichotomy. In contrast, mechanisms that inhibit post-synaptic activities such as increased inhibition e.g. GABAa channel activated *or* decreased excitation e.g. AMPA/NMDA receptor blocked) would result in the suppressed firing of both pathways ("Postsynaptic" actions).

Electrophysiology study of the Photoreceptors elucidated Ethanol's paradoxical effect

We first revisited the mechanisms of one of the most commonly used GA throughout history - ethanol. In addition to the frequently reported GABAa agonist mechanism, ethanol also acts presynaptically although less studied. From our recording, we see that ethanol's presynaptic effect played a significant role in changing the circuit output readily at 25mM.

Ethanol's reported presynaptic molecular actions included activating certain potassium channels and/or blocking calcium channels. Either action could result in the paradoxical firing we observed. Therefore to examine which (or both) mechanisms are relevant in the functioning retina, we did patch-clamp recordings on the photoreceptors. We would observe hyperpolarization of the membrane potential if the potassium channel action is dominant; likewise, we would record reduced Ca current should the Ca channel block is pertinent to the paradoxical changes.

As a result, *current clamp* of photoreceptors with ethanol perfusion showed no obvious resting membrane voltage changes, which fluctuated near -40mV (Fig 2.4 b) in control (left most inset) and various time point in ethanol (remaining insets). The hyperpolarizing membrane potential with respect to each flash also remained similar to the control. Increased intensity of hyper-polarizing cascades would have created extended and deepened troughs like (Barrow and Wu, 2009).

Next, we examine the calcium current which gates the synaptic vesicle release. To isolate calcium current, we followed the protocol in Methods 2.5.5, which utilizes

blockers to block the remaining non-Ca currents. Photoreceptor cells were held at -70 mV and then depolarized with a series of voltage steps that covered the working range (-30mV to -50mV) of the photoreceptor (Fig 2.4 c, d). As a result, the peak calcium influx was reduced by two-folds in 25 mM ethanol in the working range. This is equivalent to the retina being stimulated with more light, reducing the presynaptic vesicle release and explaining the paradoxical firing changes.

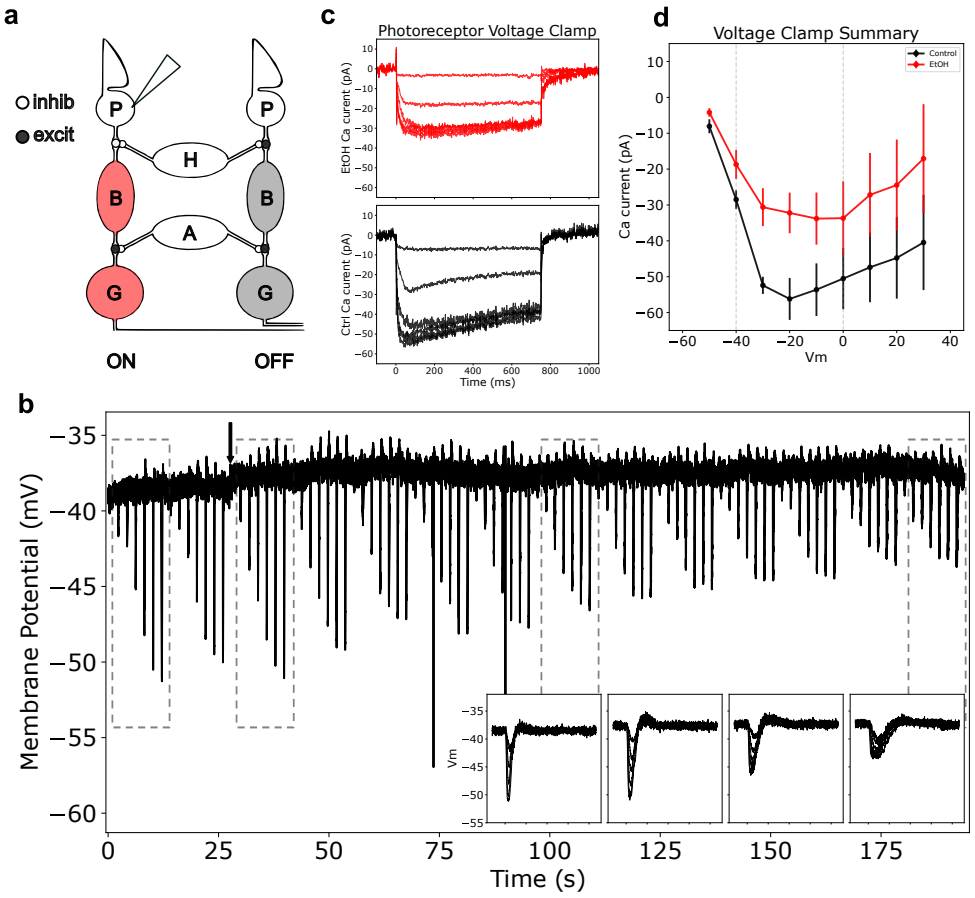

Figure 2.4: **Electrophysiology study at the Photoreceptors verified Ethanol's synaptic block effect**. **(a)** Schematic of the patch clamp studies on the photoreceptors. **(b)** Current clamp of Cone Photoreceptor in control (Ames's solution). Pulses of light hyperpolarize the membrane with transient dips. Perfusion of ethanol does not further hyperpolarize the resting membrane voltage. (Insets). **(c)** Voltage Clamp measurement of the isolated Ca current of control (Lower) and EtOH (Upper). **(d)** I-V measurements derived from of the peak current in **(c)**. EtOH reduced the inward current of photoreceptors in working range (-30mV to -50mV) by roughly 2-folds.

2.3.4 Revisiting our understanding of neuropharmacology by pairing drugs with their thought-to-be targets

One of the main problems that plagued the progress of neuropharmacology was the ambiguity in identifying the actual target(s). Now that we have a testbed with interpretable readouts, we could drive the hypothesized drug targets with a specific agonist/ antagonist and compare the responses of that drug. Retinal response to the drug and the specific target driver would align should the proposed target is the actual drug action site and vice versa.

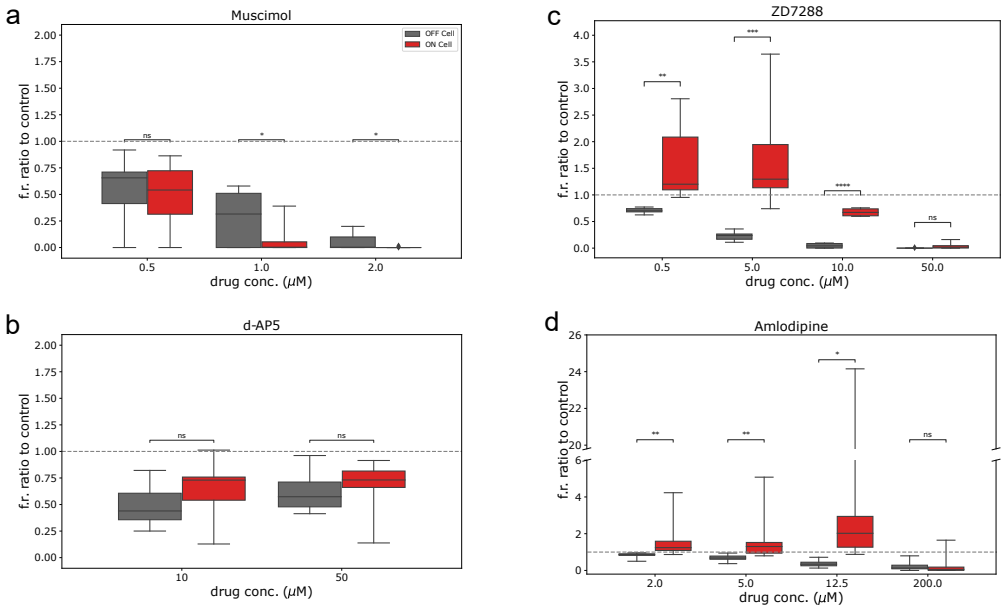

Figure 2.5: **Emulating the anesthetics features with target specific agonist/ antagonists. (a)** Muscimol,(GABAa receptor agonist) inhibits both ON and OFF pathways. At higher concentration, muscimol fully inhibits the RGCs. (OFF, ON cells counts): n= (11, 17). **(b)** d-AP5, (NMDA receptor antagonist) inhibits both ON and OFF pathways. However, the inhibition plateaued and does not result in full inhibition like muscimol, n= (11, 8). **(c)** ZD7288 (HCN channel blocker) Emulates the paradoxical firing of ON cell as well as fully inhibits both pathways in higher concentrations. (OFF, ON cells counts): 0.5 μM, n= (5, 9), 5.0 μM, n= (9, 13), 10.0 μM, n= (4, 4), 50.0 μM, n= (5, 9). **(d)** Amlodipine, (L-type Ca channel blocker). Tracked the paradoxical firing changes, strongly inhibited both pathways in higher concentrations., n= (11, 21).

GABAa receptor agonist matched Isoflurane's circuit perturbation profile

We began with activating the classic post-synaptic inhibitory target GABAa receptor with muscimol. Muscimol reenacted the "orthodoxical" inhibition of both ON and OFF cells similar to what we observed in isoflurane, as well as shuts the retinal circuit at higher concentrations (Fig 2.5 a).

NMDA receptor blocker could not emulate Ketamine's action on the retina

Next, we revisited ketamine. Given that ketamine is thought to exert its effect mainly by blocking NMDA receptors, we paired it with the NMDA receptor blocker d-AP5.

Our results showed that d-AP5 also acted as an orthodoxical depressant and inhibited both pathways (Fig 2.5 b). It did not reproduce the paradoxical firing changes even when the median firing rate of OFF cells were as high as 57% of control (50 µM d-AP5). In comparison, we still see paradoxical firing when the median firing rate of OFF cells was as low as 18% of control at 36 µM ketamine). In addition, at saturating NMDA receptor block concentration (50 µM d-AP5), the firing rate was suppressed but far from fully inhibited like ketamine at anesthetic concentration.

HCN channel blocker emulated Ketamine

HCN channel block is the other advocated anesthetics target of ketamine. We thus examined if circuit manifestations of HCN channel blocker matched ketamine.

We first tested Ivabradine, a commonly used HCN blocker for heart rate control. Ivabradine reproduced the paradoxical changes across all concentrations tested (Supp Fig 2.1). Amongst the tested concentrations, the paradoxical effect dichotomized maximally at 20 µM. The majority of ON cells were still more active than Control at the extreme concentration of 100 µM.

Ivabradine was documented as a relatively safe drug with a ceiling effect in its use of heart rate control. Given that the RGCs were far from fully inhibited like ketamine in anesthetic concentration, we tested another common HCN blocker ZD7288 (Fig 2.5 d). ZD7288 reproduced the paradoxical firing changes at low doses, then fully inhibited both ON and OFF cells by 50 µM.

Ca Channel blocker reenacted Ethanol's action

Last, we were interested in repeating ethanol's Ca Channel blocking effect that supposedly led to its paradoxical firing at low concentrations. Thus we delivered a common hypertension prescription, the L-type Ca channel blocker *Amlodipine* to pair the ethanol recordings.

Interestingly, Amlodipine also drove the retina to fire paradoxically in our tested range of 2 µM - 12.5 µM. At the extreme concentration of 200 µM, Amlodipine also nearly fully inhibits the action potentials.

2.3.5 Mechanism for the full-inhibition feature of General Anesthetics is nuanced

With the standard drug versus anesthetics pairing experiments, we successfully emulated the paradoxical effect of the sub-anesthetics dose ketamine and ethanol. In contrast, the full-inhibition effect that initially appeared as an intuitive result of the anesthetics was more diverse than we expected when reproduced with the target-specific agonist/ antagonist. For example, the inhibition from saturated NMDA channel block with d-AP5 was relatively light. In contrast, muscimol which targets GABAa was extremely effective and fully quenched the retina below its saturating concentration. On the other hand, HCN blocker ZD7288 and the common L-type Ca Channel blocker Amlodipine were peculiarly potent at high dose.

Synaptic block is insufficient to fully inhibit the RGCs from firing

Given this heterogeneity, we are curious about what it takes to fully inhibit the retinal circuit. One logical inference from the observation of the d-AP5 results was that NMDAr are the auxiliary excitatory targets of the synapses, and it would naturally require both the primary targets (AMPAr) and auxiliary targets to be blocked in order to achieve full inhibition via synaptic block.

Therefore, we experimented with the combination of saturating dose of AMPAr blocker CNQX (50 μM) with NMDAr blocker (50 μM). Surprisingly, both this double-combination and the triple-combination (plus pre-synaptic block via the addition of 200 μM Cadmium Chloride) do not completely stop the RGCs from firing (Fig 2.6 a). To check if the receptors were fully saturated by the blocker cocktails, we analyzed the correlation of the retinal ganglion cell light response to the same stimuli. The triple-combination fully abolished spike train structures of the in the RGC responses to stimuli that started with high fidelity in the control solution. (See example RGC in Fig 2.6 b).

Expanded Retinal-Meyer-Overton plot showed duality in how the retinal circuit was fully inhibited

Last, as the spontaneous RGC firing appeared to be robust against synaptic block alone, we decided to analyze light response in the Meyer-Overton range concentrations of drugs, which were much higher than the intended nominal concentration of their targets.

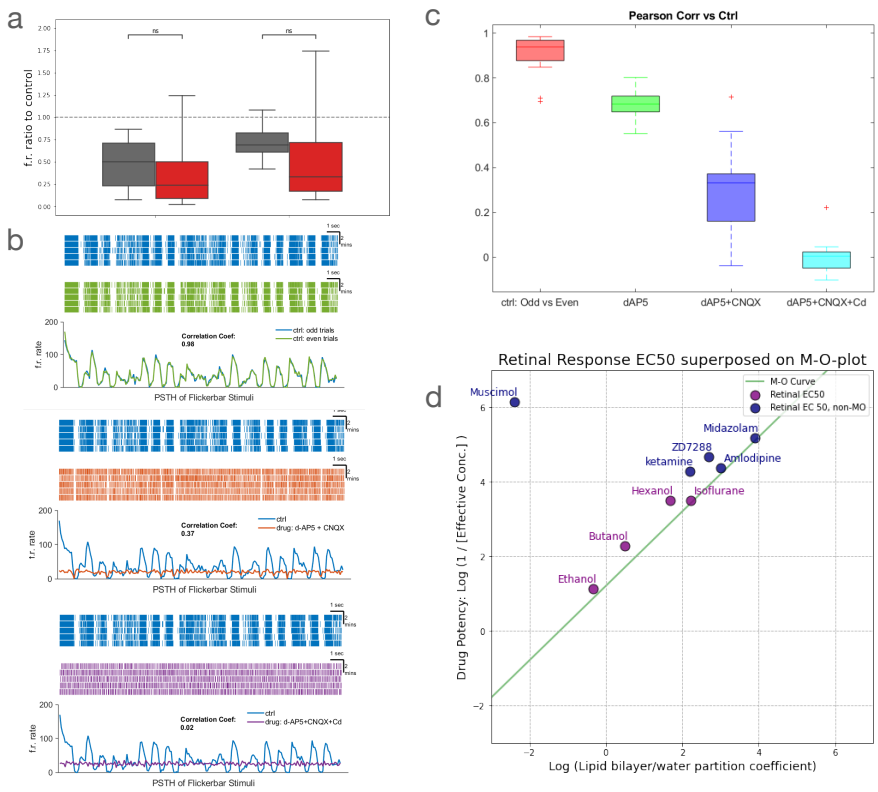

Figure 2.6: **Duality of full RGC inhibition**. **(a)** The combination of d-AP5 (50 µM) + CNQX (50 µM) versus the combination of d-AP5 (50 µM) + CNQX (50 µM) + $CdCl_2$ (200 µM). **(b)** Comparison of Spike trains of the same RGC to the repeated visual stimuli under different drug combinations. The top panel is the odd number runs vs even number runs in Ames' solution, showing high fidelity of between the spike trains to the same stimuli. Pearson's correlation coefficient was calculated by comparing the averaged PSTH of 5 runs binned at 100 *m*s. **(c)** Box plot of the Pearson's correlation coefficient of all RGCs in different conditions versus control. **(d)** The potency of quenching of Retinal response to light flashes under different compounds, plotted against the compound's lipid partition coefficient, superposed on the Meyer-Overton correlation.

When we plotted the potency of RGC inhibition to flash responses like in Fig 2.2 against their lipophilicity, we found a duality amongst the drugs. The GABAa-targeting muscimol fell off the M-O curve approx 1000-folds stronger than what M-O would have predicted. Conversely, the remainders followed the correlation.

2.4 Discussion

2.4.1 Summary of results

In this work, we exploited drugs' interaction with retinal circuits to probe the on-going activities in neuronal circuits under anesthetic and sub-anesthetic concentrations.

At sub-anesthetic doses, we found that certain anesthetics exert a paradoxical effect. They increase the light response of some retinal neurons and suppress the response of others. Notably, this occurred for alcohols and ketamine, but not for isoflurane. We traced these effects to transmitter release at a specific synapse, and in one case to a specific presynaptic ion channel. All the anesthetics silenced the output of the retina completely at concentrations similar to their effective dose for anesthesia in humans. Sedatives reduced retinal sensitivity but did not silence it even at 10-folds above clinical doses. We used specific drugs that target hypothesized molecular mechanisms to probe how much they each contribute to anesthesia of the retina. The specific drug-target pairs verified that HCN channels were a significant, if not the main part of ketamine's neuromodulation, and similarly, Voltage Gated Ca channels in ethanol. Finally, we showed that the inhibition of retinal response under all drugs tested in this project follows the Meyer-Overton correlation when pushed to the relevant dosage, with the exception of muscimol which is several orders more potent than the prediction.

2.4.2 Our observations complement current understanding of the mechanism of general anesthesia

Specific, non-Specific, or both?

The field of anesthetics research is fastly moving away from the *non-specific*[1], membrane perturbation hypothesis that was the mainstream explanation for around a century since Meyer and Overton (Overton, 1901) toward the *specific*, molecular-as-target centered explanation, since the membrane-free assay on luciferase suggested competitive action sites on a protein target could also follow the Meyer-Overton correlations (Ueda and Kamaya, 1973; N. P. Franks and Lieb, 1984). More recently, researchers seemed to have completely accepted that there are a set of molecular targets general anesthetics bind to, and one should move on and explain the functional results of general anesthesia- loss of consciousness, analgesia, immobility - via neuronal networks and brain circuits that involve the thalamus and sleep centers (Nicholas P. Franks, 2008).

Part of our data supports the more modern specific molecular target-brain circuits explanation. Without inputs from other inhibitory nuclei, it is difficult to explain how weak certain drugs are to result in obviously observable behavior manifestations such as sedation. For instance, midazolam affected the retinal circuit only weakly

[1] Non-specific as *not specifically bound to a protein target(s)* as opposed to the specific hypothesis that follows.

even at 10-X the clinical dose. In comparison, anesthetics at their corresponding clinical concentrations perturbed the retina immensely. We know at the 10X clinical dose applied, it would have knocked a person into deep sedation. This suggests that other brain circuits must have involved amplifying the thought-to-be circuit/system level effects from midazolam, which are known to have a direct binding site on the GABAa receptors. And in the other extreme, muscimol the specific GABAa agonist fully quenched RGC spikes below its saturating dose and was much more potent than what the MO correlation would have predicted. These two examples showed that the *non-specific* hypothesis alone is insufficient to capture important features of neuropharmacology.

However, other than muscimol which fell off from the M-O correlation, the rest of the compounds tested largely followed the trend. Given that quite a few of them were specific agonists/antagonists toward certain neuro-transduction targets, this seemed odd.

A proposed explanation

One simple explanation is that Meyer and Overton measured a to-be-explained nonspecific mechanism that perturbed *arbitrary* neuro circuits sufficiently at the corresponding concentration with the form of full circuit inhibition. The mice retina we tested on is merely an excerpt of all possible circuits rather than a special case. The M-O correlation thus should be regarded as the lower bound of concentration for each compound needed to reach the disruption threshold via the non-specific mechanism[2], which is comparatively weak as opposed to compounds that directly interact with essential CNS targets. The direct interference with targets that functions in *circuit inhibition* or *action potential generation* would have disrupted circuits with a much lower concentration. For instance, muscimol in Figure 2.6 D. Likewise, Na blocker blockers such as tetrodotoxin that directly interferes with action potential would also have fallen off the curve. For example, Mao et al., 2001 used 50 nM tetrodotoxin to fully inhibit all spontaneous activity in brain slices. Being a lipophobic agent, it will sit at the left-upper corner of the correlation plot, 4 orders left and at least 1 order above muscimol (logP = -6.2, NIH PubChem, 2022b). In contrast, agents with no significant interactions with the important molecular targets

[2]For simplicity, I will refer to the *to-be-explained nonspecific mechanism* as *non-specific mechanism* that should not be mixed with the original M-O based non-specific hypothesis that focuses on changes of the fluidity of cell membranes.

would sit close to the linear fit, such as the compounds tested in the classical M-O series which are basically simple alkanols.

Unfamiliar as it might appear at first glance, this explanation derived from the *non-specific* lineage has its strengths to complement findings that could not be reconciled by the specific molecular target - brain circuits theory.

First, the concentration applied to human for compounds to *anesthesize* are orders of magnitude higher than the majority of other modern medications. Biomedical scientists who are already used to reagents with high affinity and specificity toward their intended targets would have been appalled by the doses applied by us. The mole fraction of the compounds partitioned into the lipid bilayer is 0.02 to 0.05 (Cantor, 2001). That is 1 drug compound for every 20 surrounding molecules! Take the more familiar compounds that mostly have logP between 2 to 4 for example: It requires 100 mM for ethanol and 200 μM for the more potent isoflurane (Krasowski and Harrison, 1999) - which is still 100 X higher than muscimol - for anesthesia. Admittedly, compounds with high logP would have a predicted lower aqueous concentration needed to anesthetize, but the poor water solubility or high vapor pressure basically prevented meaningful measurements yet is still cited as evidence against the *non-specific hypothesis* which I will further discuss in section 2.4.3.

Second, anesthesia to the same compounds is conserved from worms to humans that are evolutionary split for millions of years across many folds of difference in nervous system complexity. While evolution might have endowed all life forms with a nervous system with a convenient molecular switch(es) for anesthesia that could be activated with hundreds of daily-life chemicals in the first chapters of organic chemistry, it contradicts the trend to explain general anesthesia via specific molecular targets embedded in dedicated brain circuits and nuclei. That is, organisms with or without the proposed dedicated circuits both could be anesthetized. From this perspective, our findings support the conventional explanation that GAs are disrupting circuit processing to the extent that no meaningful information could be relayed.

Third, molecules that could anesthetize take all kinds of forms, ranging from atoms such as Xenon (Lynch et al., 2000; Neice and Zornow, 2016) to complicated compounds of different chemical classes. This diversity of the compounds originally guided people to ponder the mechanism as a non-specific one. I will continue this discussion this in Section 2.4.3 on discovering how one family of exotic proteins followed MO correlation doesn't justify that it is actually the underlying mechanism,

nor preclude that other non-specfic targets in the neuronal circuits. In addition to the noble gases, one exotic yet more commonly encountered example that showcases the diversity of general anesthetics is the *nitrogen narcosis*. Nitrogen has a log P of 0.1, thus should be a harmless compound at normal conditions according to the MO prediction as being two orders weaker than isoflurane (PubChem, 2022a). In fact, we breathe in roughly 70 % v/v nitrogen at 1 a.t.m breath to breath and appear conscious. However, during deep dives where the high pressure forced nitrogen to dissolve into the bloodstream several times higher than the normal atmospheric condition, divers experience symptoms of CNS perturbations as nitrogen accumulates[3]. The symptoms exacerbate with depth, which in extreme depths leads to unconsciousness and death.

The third point indicated almost all compounds in the known world could be general anesthetics when they reach a sufficiently high concentration. This obviously contradicts reality as we rarely see people inebriated by drinking a quart of milk. There are obviously at least two additional criteria that had to be met, namely *the capability to cross the blood-brain barrier* and *sufficiently nontoxic* for the compounds to exert their effect in the CNS even if the non-specific hypothesis holds true. These observations are intuitive for our platform but not necessary for the conventional whole-organism-based behavioral tests and are further explained in the following passage.

Retina as a testbed revealed observations that could have been masked or actively avoided due to biological constraints

1. The need to cross the blood-brain barrier (BBB): Compounds would have to first reach the CNS to exert their effects. With this layer of insulation, around 98 % of all small molecules are not transported across the blood-brain barrier and could not reach the neurons. For diffusion-based transportation, rules of thumb suggested a compound has to be < 400 Da and forms < 8 hydrogen bonds to diffuse through the lipid layers (Pardridge, 2012). This protection from the blood brain barrier immediately confined the observable effect from infinite types of molecules to the ones that remained from iterations of empirical trial-errors in the development of anesthetics: simple oily compounds. The retina *in vivo* is also protected by a similar insulation called

[3]Perhaps oxygen also played a minor role with a logP at -1.1 but compensated with 2.5X more soluble in the blood according to Henry's law.

blood-eye barrier. Being stripped off from the eye cup, our readouts unmasked the bias of the permeability of drugs. Within our tested chemicals, ZD7288 and Amlodipine would not have crossed the BBB, thus would not have been considered as general anesthetics nor as evidence of the non-specific theory. Similarly, people rarely include IV-form compounds when comparing the biophysical properties or effects with the volatile ones due to realistic concerns of pharmacokinetics in addition to the BBB. Our preparation showed that the non-specific explanation is indeed quite invariant to the classes of chemicals and extends to conventionally oral/IV-form drugs that were never considered to be anesthetics.

2. Confined by safety: Another observational bias in addition to the need to cross the BBB is that the neuropharmacology was weeding away lethal compounds that kill the test animals without a safety margin that segregates the treatment dose and lethal doses. For instance, drugs with potency orders stronger than the ones on the MO correlation would in theory be potent GAs but are likely to perturb other systems with excitable membranes such as the heart. These compounds do not make it further down the drug development programs due to their toxicology profile, instead, the inert and boring ones were selected. Interestingly, GABAa agonist, despite being a strong CNS agent, is relatively safe interacting with the heart which was regulated by a different part list (GABAb) and is spared. In contrast, Na channel blockers such as lidocaine which are classified as "local" anesthetics by jamming the peripheral nervous system could likely be a GA but would have simultaneously paralyzed the heart and killed the organism. Interestingly, fish researchers use Na channel blockers routinely as a general anesthetics such as *tricaine* for fish studies (Attili and Hughes, 2014) where the fish cardiology system seems to be using variants of cation channel different from their CNS. Our platform is decoupled from vital organs and would have seen through the non-neurological biases such as application-oriented classifications based on safety rather than the fundamental mechanisms.

Fully inhibiting the retina turned out to be nuanced

At first glance, giving high doses of arbitrary compounds to result in fully inhibiting neural circuits is not surprising. It fits with thinking people's expectation of *dosis sola facit venenum*, the dose makes the poison, after all, we could also kill

cultured cancer cells with a high enough concentration of table salt. Nevertheless, we find it more nuanced to emulate the same level of inhibition with just standard agonist/antagonists applied at their saturating doses with the exception of muscimol. The triple-mix cocktail that fully blocks synaptic transmission with AMPA/NMDA/pre-synpatic blockers resulted in a different circuit state where RGCs were stimuli irresponsive but still robustly firing. With this data, we think the end result from the non-specific concentration effect aligned better with a voltage membrane depression/ action potential block-based mechanism rather than a synaptic transduction-block one.

Given the mounting evidence that there are various observations that could not be explained intuitively by the specific hypothesis better, we decided to review the evidence that the specific protein binding hypothesis supporters raised to decide the non-specific theory outdated. These several prongs were so pervasively cited that they are taken as fact without being further scrutinized or reexamined with modern standards, as there is no one left on the non-specific side.

2.4.3 MO is flawed, but the several prongs of critiques compiled by advocates of direct target theory were often not stringently conducted and were rarely independently verified

1. The first argument people raise against the non-specific hypothesis is the exceptions in M-O correlations. Commonly cited examples include that 1-alkanols have an abrupt cut-off of losing potency to abolish tadpole's righting reflex after > 13 C (Pringle, Brown, and K. W. Miller, 1981), and fluorocarbons larger than carbon tetrafluoride (CF_4) do not anesthetize mice in hyperbaric chambers filled with fluorocarbon gases. Upon revisiting, I found that 1-alkanols > 13 C do not permeate into lipid bilayers. The fluorocarbon experiments were done in extreme conditions with mice in a pressure chamber in order to compensate for the extremely low Henry's law constant that prevents the gases from partitioning into the bloodstream. So extreme that the team predicted the partial pressure required for (CF_4) in the lung alveolar needed to reach 35 atm and never delivered to the nominal pressure, as the mice are killed before 30 atm. For these molecules, not limited to fluorocarbons, with extremely low water solubility yet with vapor pressure as high as several atm, documenting dose response based on the animal's behavioral outcome that required chemicals to diffuse through several layers of interfaces (air-> liquid film on the aveolar-> across lung epithelial cells -> across capillary endothelial cell membrane) is difficult. Changing the subject to tadpole would not ameliorate the difficulty of

having these extreme compounds permeate all the way into neurons. Not to say, the final brain concentration was never reported as quality control for the negative results.

2. "If membrane biophysical properties changes are the cause of anesthesia, then physical perturbations that disrupt the membrane at the same order would also have caused conscious change". Intriguingly, despite this notion being commonly cited as evidence against the membrane fluidity hypothesis, there are various temperature-related CNS examples such as hypothermia or hyperthermia-related loss of consciousness in humans. Exothermic lifeforms such as insects could be readily knocked out by cold. Though there are ample of examples, the underlying nervous system mechanism is too unclear to serve as evidence for or against either hypothesis.

3. Evidence of direct protein-anesthetics binding. People often regard the finding of the firefly luciferase-luciferin catalytic reaction being competitively inhibited by various forms of general anesthetics as the nail in the non-specific hypothesis' coffin. The luciferase assay demonstrated that direct binding sites in a protein target could also follow the Meyer-Overton rule (N. P. Franks and Lieb, 1984). And although we now know that the long chain alkanol experiment mentioned in the first point could be a fluke, the paper intelligently used a lipophilic pocket in the protein to explain why there was a cutoff of potency when the drug is too big to fit into the pocket. However, since the 80s, people have yet to find essential protein targets such as Na, GABA, and other CNS receptor that could be inhibited like the luciferase reaction at clinical relevant concentration (Urban, Bleckwenn, and Barann, 2006). In fact, almost all protein targets could form some kind of log -log correlation according to Urban, Bleckwenn, and Barann, 2006, albeit with different slope and potency. However, the measured effective dose usually missed the clinical dose by several folds. Measurements like this were sporadically reported from time to time for different anesthetics vs arbitrary proteins such as the myoglobin binds with xenon (Schoenborn, Watson, and Kendrew, 1965) [4] that the luciferase seemed like more of a one-of-a-kind hit that just happened to come out within the relevant clinical concentration. In summary, finding an exotic protein that followed the MO correlation does not justify that direct protein-anesthetic interaction is the underlying mechanism, nor precluded other non-specific binding sites within the CNS to be the actual mechanism.

[4] Perhaps one will also find a MO like binding affinity correlation with myoglobins as well after throwing the arsenal of drug onto it

4. Stereo Isomers with different potency. If a molecule's stereoisomer exerted different potency, it almost guarantees that there is an underlying specific mechanism, as the stereoisomers have the same biophysical property that would be invariant in potency should the mechanism is non-specific. The isomer of etomidate R(+) has a higher potency than S(-) and was commonly cited as evidence of this prong (Tomlin et al., 1998; Krasowski and Harrison, 1999). Nevertheless, as the authors should have known, etomidate is technically not a general anesthetics as it is not an analgesic. Etomidate (Not considered as a BZD drug) targets GABAa receptor and thus is not surprising to have different *sedative* potency. It is no more surprising to claim isomers of benzodiazepines also have different strengths in the sleeping drive.

In summary, the dominant evidence cited against the non-specific hypothesis is based on special cases where the compound has extreme biophysical properties, the substrate is too irrelevant to actual protein targets in the brain, or confusing the subject matter with sedatives that are not GAs. Despite that these findings might be novel evidence when first discovered, with neurobiology as a field improved, people should reexamine the stringency of these works, which have formed the current textbook understanding of anesthesiology. Also, one should not ignore the other 90 percent of the oily compounds that followed the MO rule, with or without the 4 prongs listed.

Next, we transition into the sub-anesthetic dose findings. In contrast to the anesthetic range effects that still required further explanations, drugs seem to exert their effect through specific synaptic targets or ion channels. The retina highlighted interesting targets to be important players in the sub-anesthetics dose circuit responses that were previously sporadically reported but not taken seriously due to the first-comer effect of mechanisms proposed earlier (see 2.4.4). Given that we still barely know about how alcohol is acting in the brain, as well as how low-dose ketamine act as a rapid anti-depressant, this part of the study again showed that using well-known circuits as testbed revealed finding that should have been obvious if one simply changes the scale of observations.

2.4.4 Sub-anesthetic dose/ Direct Target: circuit approach highlighted targets that may so far been overlooked

Voltage gate Ca Channels (CaV) and Hyperpolarization-activated cyclic nucleotide-gated (HCN) channels are relevant targets

Observations that led to the explanations for neuropharmacology came with a strong first-comer effect. Often times the first available model system or the earliest tested substrates with a positive result would be hailed as the actual mechanism. Ketamine was coupled with NMDA receptor block when an experiment showed that it specifically blocked NMDA receptors amongst several types of glutamate receptors (Anis et al., 1983). No doubt this was a novel discovery as Ketamine barely interacts with the popular GABAa at relevant concentrations, and a molecular target was in need to explain its effects. Again, logically analogous to the luciferase experiment, ketamine being a specific NMDAr blocker does not automatically justify its dissociative anesthesia property from the NMDAr block. Chen, Shu, and Bayliss, 2009 alternatively showed that ketamine's HCN antagonist effect was more likely the underlying mechanism. It is known to experimental neurobiologists that specific NMDAr blockers have only a mild circuit perturbation effect and do not lead to an anesthetic state. Nevertheless, ketamine as an NMDAr blocker has become the canonical textbook mechanism for effects exerted by ketamine (R. Miller, 2015). We showed that the paradoxical firing caused by low dose ketamine could be reliably repeated by two different standard HCN antagonists, and do not align with specific NMDA blockers. This finding supports the argument that HCN rather than NMDA, is more relevant to ketamine's CNS effect.

On the other hand, the sub-anesthetic dose of ethanol that people encounter frequently for recreation remained largely elusive and does not have a dominant explanation. The recordings we have from extracellular recording and the following patch clamp experiment showed that CaV block is the main cause of the low-dose circuit changes and that CaV blocker Amlodipine could recreate the same effect.

Isoflurane's effect starting at sub-anesthetic dose

In contrast to the ethanol and ketamine, we observed monotonically dose-correlated depression across the concentration applied. As a result isoflurane's interaction with non-GABA targets is unlikely core to its pharmacological property. Isoflurane's data aligned with muscimol and suggest the net effect is inhibiting postsynaptic-ally starting from the subanesthetic dose.

Why haven't people connected GAs with HCN (Or CaV) earlier if they are potent?

Specific HCN blocker ZD7288 is a potent compound that could hyperpolarize membrane potential by 5 mV in hippocampus patch recordings (Gasparini and DiFrancesco, 1997) and abolish spikes. Mao et al. (2001) worked with cortical slices imaging and also observed that HCN channel blockage by ZD7288 decimated the spontaneous firings by several folds. Like us, they also reported surprise at how little effect glutamatergic block had on firing rate.

Despite these prior findings, most HCN (and CaV) blockers do not pass the blood-brain barrier and faced an observational bias like the aforementioned. For the neuropharmacology field that largely relied on heuristic observations of animal behaviour up till the 70s, such as crushing the skin of inmates with hemostats when testing ketamine and its phencyclidine analogs (Domino, Chodoff, and Corssen, 1965), these two classes of drugs would not have exerted CNS manifestations of interest and made it into the next stage.

2.4.5 Connection with the paradoxical firings of the brain in other studies

To the retina paradoxical firing (coined as different names in other studies) is an intuitively understood phenomenon for drugs with presynaptic action. Here are a few other studies that reported unexpected paradoxical firing. The underlying circuit could be different thus the cause of paradoxical firing in these projects might not simply be explained as pre-synaptic.

1. The ketamine and thought-to be NMDA blockers (Widman and McMahon, 2018). Attempt of interpreting how the circuit disinhibition worked was convoluted. An HCN based explanation could result in the same phenomenon with less assumption.

2. Found focal firing core in the hypothalamus that fires paradoxically under many drugs: Iso and sevo-flurane, propofol, ketamine (Jiang-Xie et al., 2019)

3. Isoflurane also elicits paradoxical firing for some cells in the hippocampus (Ou et al., 2020)

2.4.6 Is the paradoxical firing related to the euphoric properties or anti-depressant effect at lower doses?

Intriguingly, the paradoxical firing is hypothesized as the fast anti-depressant basis of low-dose ketamine (Duman and Aghajanian, 2012; Duman, Aghajanian, et al., 2016). There seems to be a gap of a leap of faith in how the connection was made. In addition, ketamine is inheriting its NMDA blocker label, thus researchers directly cite and interpret their increased firing also from NMDA block. Again, this showed that there need to be a neuropharmacology testbed to standardize the interpretation of mechanisms.

2.4.7 Future directions: Testbed beyond the general anesthetics

We demonstrated that by measuring interactions of drugs with a well-studied neural circuit, one could immediately gain insights and interpret the actual targets.

Following this direction, the neuropharmacology field should aim for an overarching goal of seeking more circuit-based platforms to complement molecular and behavioral studies. New circuit based testbeds should be used to reexamine other empirically derived use of chemicals for neuro disorders. For instance, movement disorders.

2.5 Methods and Materials

2.5.1 Extracellular Retina Recording

We harvested retina from C57BL/6 mice (JAX:000664, Jackson Laboratory) aged 6 weeks-6 months. Mice were sacrificed by cervical dislocation after 1-hour dark adaptation. We first remove the eyes from the mouse using curved tweezers. Then, with the eyeball immersed in Ames' solution (Sigma-Aldrich) filled petri dish, we dissected away the cornea and lens to expose the retina. The Ames' filled petri dish was bubbled with carbogen (95 $\%O_2$, 5$\%CO_2$) continuously during the dissection to keep the retina viable.

The retina was then gently peeled off from the remaining eye cup with fine tweezers. We cut the retina into 4 equal-sized slices, and transferred one slice onto a customized widget as in (Fig 2.1 W). The widget held the flattened retina on a piece of transparent hemodialysis membrane (Sigma-Aldrich) like a tightened drum head against the multiple electrode array (MEA: 256MEA60/10iR-ITO, multichannel systems). In this configuration, the retinal ganglion cells faced the electrodes, and the photoreceptors faced the hemodialysis membrane on the widget. The MEA were then mounted onto its data acquisition board (USB-MEA 256, multichannel sys-

tems). The retina were perfused with continuously bubbled Ames' solution through a gravity-driven setup to remain light-responsive.

We waited 45 minutes for the retina to recover in dark. Then we light-adapted the retina with background level light (See Method 2.5.2) for another 20 minutes and started the recordings. The retinal slices that were not reactive to light after adaption were discarded. A typical recording lasted 2-3 hours. All procedures were performed in accordance with institutional guidelines and approved by the Caltech IACUC.

2.5.2 Stimuli

Visual stimuli were programmed using the Psychtoolbox (Brainard, 1997; Kleiner et al., 2007) package in MATLAB (Mathworks) and presented on a gamma-corrected projector (LightCrafter: DLP3000, Texas Instrument).

2.5.3 Pharmacology

All drugs were dissolved in Ames solution for their delivery with the exception of Isoflurane, which was bubbled into the Ames solution with an anesthesia machine. We took the data that was collected 10 min after each transition of drug type or drug concentration as the new steady state measurements.

Compounds	Concentration	Conventionally Declared Target	Brand
Ethanol	25 - 150 mM	GABAa	Sigma Aldrich
Butanol	4 - 40 mM	GABAa	Sigma Aldrich
Hexanol	0.4 – 4 mM	GABAa	Sigma Aldrich
d-AP5	10 - 50 μM	NMDAr	abcam
CNQX	50 μM	AMPAr	abcam
Isoflurane	0.5- 1.5 % V/V	–	Patterson
Muscimol	0.5 - 2 μM	GABAa	abcam
Ivabradine	2 - 100 μM	HCN Channel	Sigma Aldrich
ZD7288	0.5 - 50 μM	HCN Channel	Sigma Aldrich
Midazolam	0.1 - 100 μM	GABAa	Sigma Aldrich
Amlodipine	2 - 200 μM	L-type Ca Channel	Sigma Aldrich
Cadmium Chloride	200 μM	Ca Channel	Sigma Aldrich

Table 2.1: Table of pharmacology experiment conditions

2.5.4 Spikesorting

We sorted spikes with *KiloSort1* (Pachitariu et al., 2016). The sorted data were then manually curated with *phy* (Rossant et al., 2016).

2.5.5 Whole cell Patch clamp

Using technique in Ingram, Sampath, and Fain, 2020, we performed Whole-cell patch clamp in retinal slices (Cx36−/− mouse cones). In brief, mice were sacrificed by cervical dislocation after overnight dark adaptation. The anterior portion of the eye including the lens was removed, and the remaining eyecup was stored at 32°C in a custom, light-tight storage container that allowed for the gassing of solutions. For each slice preparation, half of the eyecup was dissected with a No.10 scalpel, and the retina was gently peeled off from the retinal pigmented epithelium with fine tweezers. The isolated retinal piece was embedded in 3 percent of low-temperature gelling agar in Ames'-HEPES. In cold Ames'-HEPES, 200-μm thick slices were cut with a vibratome (Leica VT-1000S); the retina was cut vertically in an attempt to maintain neural circuitry. Cut slices were either transferred to dishes for immediate recording or stored in the light-tight container with the remaining pieces of the eyecups. During recordings, slices were stabilized with handmade anchors. Bath solution was maintained at 35 ± 1°C. Cones were identified by the position and appearance of their somata, as well as from measurements of membrane capacitance and sensitivity to a moderate-intensity flash. All light stimuli were brief (3–5 ms), monochromatic flashes of 405-nm light, a value near the isosbestic point of the S-cone and M-cone pigments. Monochromatic light was provided by ultra-bright light emitting diodes driven with a linear feedback driver (Opto-LED; Carin Research). Recordings were made with a Cs+ internal (pipette) solution, which contained both Cs+ and TEA in sufficient concentrations to block BK channels.

2.5.6 Analysis Software

All analysis was performed with Matlab R2021a (Mathworks) and Python3.

Supplementary Figures

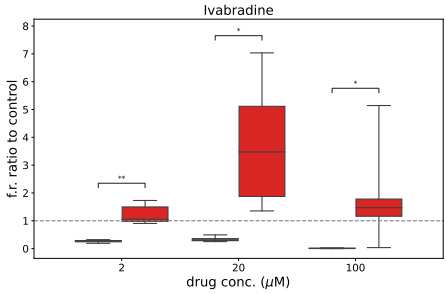

Figure 2.1: **Supp Fig** Ivabradine, (HCN channel blocker). Emulates the paradoxical firing of ON cell while inhibiting the OFF cells in all concentrations tested. It does not fully inhibit the retina in extremely high concentrations, n= (4, 6).

References

Anis, N.a. et al. (1983). "The dissociative anaesthetics, ketamine and phencyclidine, selectively reduce excitation of central mammalian neurones by N-methyl-aspartate". en. In: *British Journal of Pharmacology* 79.2. _eprint: https://onlinelibrary.wiley.com/doi/pdf/ 5381.1983.tb11031.x, pp. 565–575. ISSN: 1476-5381. DOI: 10.1111/j.1476-5381.1983.tb11031.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1476-5381.1983.tb11031.x (visited on 02/17/2022).

Attili, Seetharamaiah and Simon M. Hughes (Aug. 2014). "Anaesthetic Tricaine Acts Preferentially on Neural Voltage-Gated Sodium Channels and Fails to Block Directly Evoked Muscle Contraction". In: *PLoS ONE* 9.8, e103751. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0103751. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4121177/ (visited on 09/08/2022).

Barrow, A. J. and S. M. Wu (May 2009). "Low-Conductance HCN1 Ion Channels Augment the Frequency Response of Rod and Cone Photoreceptors". en. In: *Journal of Neuroscience* 29.18, pp. 5841–5853. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.5746-08.2009. URL: https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.5746-08.2009 (visited on 02/08/2022).

Brainard, D. H. (1997). "The Psychophysics Toolbox". In: *Spatial Vision* 10, pp. 433–436.

Cantor, R S (May 2001). "Breaking the Meyer-Overton rule: predicted effects of varying stiffness and interfacial activity on the intrinsic potency of anesthetics." In: *Biophysical Journal* 80.5, pp. 2284–2297. ISSN: 0006-3495. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1301419/ (visited on 08/30/2022).

Chen, Xiangdong, Shaofang Shu, and Douglas A. Bayliss (Jan. 2009). "HCN1 Channel Subunits Are a Molecular Substrate for Hypnotic Actions of Ketamine".

en. In: *Journal of Neuroscience* 29.3. Publisher: Society for Neuroscience Section: Articles, pp. 600–609. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3481-08.2009. URL: https://www.jneurosci.org/content/29/3/600 (visited on 05/20/2021).

Domino, Edward F., Peter Chodoff, and Guenter Corssen (May 1965). "Pharmacologic effects of CI-581, a new dissociative anesthetic, in man". en. In: *Clinical Pharmacology & Therapeutics* 6.3, pp. 279–291. ISSN: 00099236. DOI: 10.1002/cpt196563279. URL: https://onlinelibrary.wiley.com/doi/10.1002/cpt196563279 (visited on 05/17/2022).

Duman, Ronald S. and George K. Aghajanian (Oct. 2012). "Synaptic Dysfunction in Depression: Potential Therapeutic Targets". In: *Science* 338.6103. Publisher: American Association for the Advancement of Science, pp. 68–72. DOI: 10.1126/science.1222939. URL: https://www.science.org/doi/10.1126/science.1222939 (visited on 09/09/2022).

Duman, Ronald S., George K. Aghajanian, et al. (Mar. 2016). "Synaptic plasticity and depression: new insights from stress and rapid-acting antidepressants". en. In: *Nature Medicine* 22.3. Number: 3 Publisher: Nature Publishing Group, pp. 238–249. ISSN: 1546-170X. DOI: 10.1038/nm.4050. URL: https://www.nature.com/articles/nm.4050 (visited on 09/09/2022).

Franks, N. P. and W. R. Lieb (Aug. 1984). "Do general anaesthetics act by competitive binding to specific receptors?" en. In: *Nature* 310.5978. Number: 5978 Publisher: Nature Publishing Group, pp. 599–601. ISSN: 1476-4687. DOI: 10.1038/310599a0. URL: https://www.nature.com/articles/310599a0 (visited on 09/09/2022).

Franks, Nicholas P. (May 2008). "General anaesthesia: from molecular targets to neuronal pathways of sleep and arousal". en. In: *Nature Reviews Neuroscience* 9.5, pp. 370–386. ISSN: 1471-0048. DOI: 10.1038/nrn2372. URL: https://www.nature.com/articles/nrn2372 (visited on 05/10/2019).

Gasparini, S. and D. DiFrancesco (Nov. 1997). "Action of the hyperpolarization-activated current (Ih) blocker ZD 7288 in hippocampal CA1 neurons". en. In: *Pflügers Archiv* 435.1, pp. 99–106. ISSN: 1432-2013. DOI: 10.1007/s004240050488. URL: https://doi.org/10.1007/s004240050488 (visited on 05/14/2022).

Ingram, Norianne T., Alapakkam P. Sampath, and Gordon L. Fain (Jan. 2020). "Membrane conductances of mouse cone photoreceptors". In: *Journal of General Physiology* 152.3, e201912520. ISSN: 0022-1295. DOI: 10.1085/jgp.201912520. URL: https://doi.org/10.1085/jgp.201912520 (visited on 05/11/2022).

Jiang-Xie, Li-Feng et al. (Apr. 2019). "A Common Neuroendocrine Substrate for Diverse General Anesthetics and Sleep". en. In: *Neuron*. ISSN: 08966273. DOI: 10.1016/j.neuron.2019.03.033. URL: https://linkinghub.elsevier.com/retrieve/pii/S089662731930296X (visited on 05/01/2019).

Kleiner, M. et al. (2007). "What's new in psychtoolbox-3". English (US). In: *Perception* 36.14. Publisher: Pion Ltd., pp. 1–16. ISSN: 0301-0066.

Krasowski, M. D. and N. L. Harrison (Aug. 1999). "General anaesthetic actions on ligand-gated ion channels". en. In: *Cellular and Molecular Life Sciences CMLS* 55.10, pp. 1278–1303. ISSN: 1420-9071. DOI: 10.1007/s000180050371. URL: https://doi.org/10.1007/s000180050371 (visited on 08/31/2022).

Lynch, Carl et al. (Mar. 2000). "Xenon Anesthesia". In: *Anesthesiology* 92.3, pp. 865–870. ISSN: 0003-3022. DOI: 10.1097/00000542-200003000-00031. URL: https://doi.org/10.1097/00000542-200003000-00031 (visited on 09/12/2022).

Mao, Bu-Qing et al. (Dec. 2001). "Dynamics of Spontaneous Activity in Neocortical Slices". English. In: *Neuron* 32.5. Publisher: Elsevier, pp. 883–898. ISSN: 0896-6273. DOI: 10.1016/S0896-6273(01)00518-9. URL: https://www.cell.com/neuron/abstract/S0896-6273(01)00518-9 (visited on 06/21/2022).

Miller, Ronald (2015). *Miller's Anesthesia*. 8th. Saunders. ISBN: 978-0-7020-5283-5.

Neice, A. E. and M. H. Zornow (2016). "Xenon anaesthesia for all, or only a select few?" en. In: *Anaesthesia* 71.11. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/anae.13569, pp. 1267–1272. ISSN: 1365-2044. DOI: 10.1111/anae.13569. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/anae.13569 (visited on 07/02/2020).

Ou, Mengchan et al. (Jan. 2020). "The General Anesthetic Isoflurane Bilaterally Modulates Neuronal Excitability". en. In: *iScience* 23.1, p. 100760. ISSN: 2589-0042. DOI: 10.1016/j.isci.2019.100760. URL: https://www.sciencedirect.com/science/article/pii/S258900421930505X (visited on 06/15/2022).

Overton, Charles Ernest (1901). *Studien über die Narkose: zugleich ein Beitrag zur allgemeinen Pharmakologie*. de. Google-Books-ID: a809AAAAYAAJ. G. Fischer.

Pachitariu, Marius et al. (2016). "Fast and accurate spike sorting of high-channel count probes with KiloSort". In: *Advances in Neural Information Processing Systems* 29.Nips, pp. 4448–4456. ISSN: 10495258. URL: https://papers.nips.cc/paper/6326-fast-and-accurate-spike-sorting-of-high-channel-count-probes-with-kilosort.pdf.

Pardridge, William M (Nov. 2012). "Drug transport across the blood–brain barrier". In: *Journal of Cerebral Blood Flow & Metabolism* 32.11, pp. 1959–1972. ISSN: 0271-678X. DOI: 10.1038/jcbfm.2012.126. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3494002/ (visited on 06/03/2021).

Pringle, Michael J., Kenneth B. Brown, and Keith W. Miller (Jan. 1981). "Can the Lipid Theories of Anesthesia Account for the Cutoff in Anesthetic Potency in Homologous Series of Alcohols?" en. In: *Molecular Pharmacology* 19.1. Publisher:

American Society for Pharmacology and Experimental Therapeutics Section: Article, pp. 49–55. ISSN: 0026-895X, 1521-0111. URL: https://molpharm.aspetjournals.org/content/19/1/49 (visited on 08/30/2022).

PubChem (2022a). *Nitrogen*. en. URL: https://pubchem.ncbi.nlm.nih.gov/compound/947 (visited on 09/13/2022).

– (2022b). *Tetrodotoxin*. en. URL: https://pubchem.ncbi.nlm.nih.gov/compound/11174599 (visited on 09/09/2022).

Rossant, Cyrille et al. (Apr. 2016). "Spike sorting for large, dense electrode arrays". en. In: *Nature Neuroscience* 19.4, pp. 634–641. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.4268. URL: http://www.nature.com/articles/nn.4268 (visited on 07/03/2019).

Schoenborn, Benno P., Herman C. Watson, and John C. Kendrew (July 1965). "Binding of Xenon to Sperm Whale Myoglobin". en. In: *Nature* 207.4992. Number: 4992 Publisher: Nature Publishing Group, pp. 28–30. ISSN: 1476-4687. DOI: 10.1038/207028a0. URL: https://www.nature.com/articles/207028a0 (visited on 09/12/2022).

Tomlin, Sarah L. et al. (Mar. 1998). "Stereoselective Effects of Etomidate Optical Isomers on Gamma-aminobutyric Acid Type A Receptors and Animals". In: *Anesthesiology* 88.3, pp. 708–717. ISSN: 0003-3022. DOI: 10.1097/00000542-199803000-00022. URL: https://doi.org/10.1097/00000542-199803000-00022 (visited on 09/14/2022).

Ueda, Issaku and Hiroshi Kamaya (May 1973). "Kinetic and Thermodynamic Aspects of the Mechanism of General Anesthesia in a Model System of Firefly Luminescence in Vitro". In: *Anesthesiology* 38.5, pp. 425–435. ISSN: 0003-3022. DOI: 10.1097/00000542-197305000-00002. URL: https://doi.org/10.1097/00000542-197305000-00002 (visited on 06/29/2022).

Urban, Bernd W., Markus Bleckwenn, and Martin Barann (Sept. 2006). "Interactions of anesthetics with their targets: Non-specific, specific or both?" en. In: *Pharmacology & Therapeutics* 111.3, pp. 729–770. ISSN: 0163-7258. DOI: 10.1016/j.pharmthera.2005.12.005. URL: https://www.sciencedirect.com/science/article/pii/S0163725806000027 (visited on 06/04/2021).

Widman, Allie J. and Lori L. McMahon (Mar. 2018). "Disinhibition of CA1 pyramidal cells by low-dose ketamine and other antagonists with rapid antidepressant efficacy". In: *Proceedings of the National Academy of Sciences* 115.13. Publisher: Proceedings of the National Academy of Sciences, E3007–E3016. DOI: 10.1073/pnas.1718883115. URL: https://www.pnas.org/doi/10.1073/pnas.1718883115 (visited on 05/16/2022).

Chapter 3

AUTISM STUDY

3.1 Abstract

Similar to the situation in seeking what the mechanism of general anesthesia is, the field of autism research also lacks a good testbed but for systemically comparing gene mutation - circuit defect - behavior outcomes. Likewise, we utilized the retina as a platform to identify circuit defects in 4 different autism model mice and followed through the different mouse line's behavior readouts using our lab's maze navigation paradigm. We discovered that for the different autism mouse lines differed in the retinal circuits and varied in their navigation preferences. Nevertheless, unlike the anesthetic project, there wasn't a simple mechanism to explain why or how these differences are coupled together

3.2 Introduction

Autism spectrum disorder (ASD) is a developmental disorder thought to derive from the dysfunction of complex neuronal circuits. It has a strong genetic component, and there is great interest in understanding the links between autism-related genes and circuit-level dysfunctions. A large body of research has tackled this domain, leading to a wide range of mechanistic proposals for circuit malfunctions: from developmental abnormalities in the long-range connection between brain areas to an imbalance between excitation and inhibition in local circuits (Chen et al., 2015)

Mouse models of ASD have played an important role in uncovering such mechanistic relationships. Of particular interest are mouse lines created to carry genetic mutations that are associated with ASD in humans. It is remarkable that many of these mutant mice also show behavioral abnormalities reminiscent of human autism, which spurs some confidence that common mechanisms are at work in the mouse and human brains. The scientific challenge in finding those mechanisms results from the fact that genes and brain function are separated by so many explanatory levels: protein expression, cell biology of neurons and glia, the development of synaptic connectivity, the signaling dynamics of individual neurons, and the collective function of synaptic networks. A typical research study of mutant mice will probe in a directed fashion for abnormalities at one or more of these levels, with the hope that any such observed defects can explain the change in system function. In this

enterprise, much of the focus has been on cortical areas, suspected to be the seat of the higher functions that appear perturbed in ASD. Here we propose an alternative and complementary approach.

The retina is one neural system in which the links between genes, proteins, circuits, and neural function are particularly well understood already. In complexity, this network rivals the rest of the central nervous system. The retina contains 100 different types of neurons, many of which are defined quantitatively and genetically accessible (Zeng and Sanes, 2017). It includes >30 different microcircuits (Sanes and Masland, 2015) – one for each type of output neuron – and these comprise every imaginable circuit motif (Gollisch and Markus Meister, 2010), using most every neurotransmitter and modulator known to neuroscience (Hoon, Okawa, et al., 2014). Nevertheless, there are now neural circuits leading through the retina, from photoreceptors to ganglion cells, where all the important interneurons and their synapses are known (Gollisch and Markus Meister, 2010; Helmstaedter et al., 2013; Krishnaswamy et al., 2015). Mathematical models of retinal circuits can predict the visual responses of the output neurons in quantitative detail (Baccus et al., 2008). The historical reasons for this depth of understanding compared to other parts of the brain lie in the unusual experimental control one has over the retina, with complete access to its inputs and outputs.

On this background, we propose to screen ASD mouse models for abnormalities in the function of the retina. Our guiding hypothesis is that any mutation or environmental perturbation that causes defects in neuronal circuits will leave a trace in the functions of the retina. As detailed below, such measurements can be extremely sensitive: There are many retinal microcircuits available to test, each with a different configuration of synapses. The toolkit of visual stimuli is virtually unlimited and can be tailored for specific hypotheses. Finally, the neural response measurements are very reproducible across animals, allowing even small effects to be resolved. Once one finds an abnormality in system function – in this case the processing of visual inputs – one can quickly dissect the dysfunction and identify the underlying synaptic and cellular defects. In homing into home on the likely sources one gets invaluable help from the extensive existing knowledge of retinal microcircuitry. These mechanistic tests will include the classical tools of immunocytochemistry and pharmacology. Effectively the retina can serve as a test bed in which one develops circuit-level explanations that can then be tested in brain areas more likely to be causally involved in ASD. Building on the results of the retina findings, we took

advantage of newly emerging opportunities to broaden the scope of the comparative analysis. Specifically, we studied the behavior of mice from these same lines in several natural tasks that are part of the essential rodent repertoire: escape from predators, exploration of an environment, learning, and navigation in space. The aim was to item Identify differences in these behaviors among the mutants and with respect to the wild-type mouse. and correlate any such differences with the class of phenomena found in circuit-level analysis of the retina.

3.3 Results

3.3.1 Retinal Probing Results

A sample of recording and analysis

Figure 3.1 illustrates the range of measurements obtained from a single retinal ganglion cell. When the retina was exposed to periodic flashes of light, this neuron fired a sharp burst of 100 spikes per second following offset of the light (Fig 3.1 A): a classic Off-type cell. A white-noise analysis revealed the spatial profile of its receptive field and the time course of its response kernel (Fig 3.1 B). This neuron had a center diameter of 250 µm and a kernel time-to-peak of 50 ms. Finally, stimulation with moving bars revealed that the neuron responds consistently to the dark edge of a bar, with no overt preference for any movement direction (Fig 3.1 C).

These measurements were performed for several 1000 retinal ganglion cells from the various mouse lines. About 830 neurons survived an initial step of data curation based on the quality of the electrical recording and stability of the light response. Then we compared the key parameters of the light response between mutant animals and their respective controls, as described in the following sections.

Categorical cell classes

Retinal ganglion cells are coarsely classified as On, Off, or On-Off based on whether their firing increases at the onset of a light flash, at the offset (Fig 3.1 A), or both (Carcieri, Jacobs, and Sheila Nirenberg, 2003). This elementary aspect of the neuron's function is directly related to its morphology, specifically, whether the dendrites arborize in the upper or the lower stratum of the inner plexiform layer (Wässle, 2004). These strata contain the axonal terminals of On-type (lower) and Off-type (upper) bipolar cells respectively. We found that some of the mutant lines have an abnormal distribution of these response categories (Fig 3.2). Both the Nlgn3 and the Mecp2 mutant have an abnormally high fraction of On-Off cells, at

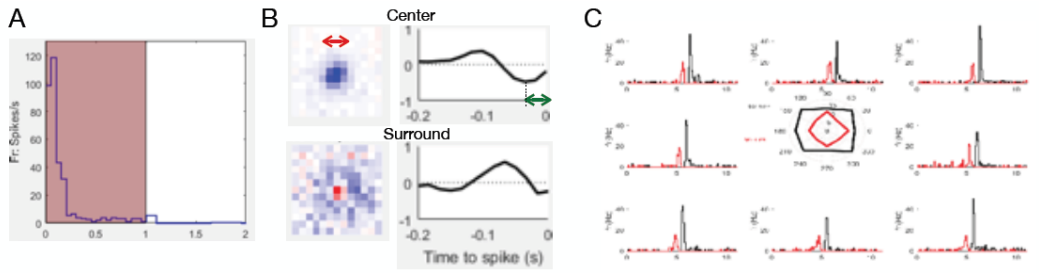

Figure 3.1: A: response to a uniform light flashing off (shaded) and on (white). The neuron's firing rate (measured in spikes per s, or Hz) is plotted as a function of time throughout one period of this stimulus. B: The spatio-temporal receptive field (RF) of the cell, measured with a white-noise flickering checkerboard stimulus. Top left: Profile of the RF center, with the diameter indicated in red. Top right: Time course of the RF center, with time-to-peak indicated in green, Bottom: Same for the RF surround. C: Here the retina was stimulated with a black or white bar moving over a gray background in 8 possible directions. The corresponding 8 graphs around the periphery each plot the cell's firing rate as a function of time. One sees a brief burst of firing when the bar crosses over the neuron's receptive field. A dark bar (black) produces a burst at the leading edge (the bright-to-dark transition) and a white bar (red) at the trailing edge. The polar plot at the center summarizes the response amplitude as a function of the direction of motion; for this cell, the response varied little with direction.

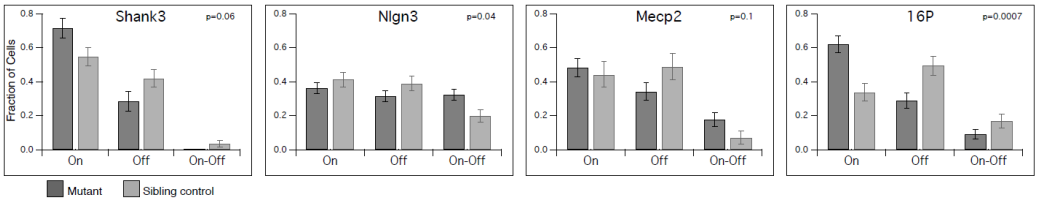

Figure 3.2: The distribution of basic response types among retinal ganglion cells in the four mutant lines investigated. On, Off, and On-Off types were identified by their response to periodic light flashes (Fig 1A). The fraction of each response type in the observed ganglion cell population is plotted along with its standard error. P-values are derived from a chi-square test comparing the distributions in mutant and sibling control animals.

the expense of pure On and Off cells. By contrast, both Shank3 and 16P mutants contain an abnormally high fraction of On cells and fewer Off cells (Fig 3.2 C-D).

Quantitative visual processing

To test for quantitative changes in retinal processing we analyzed the spatio-temporal receptive field (RF) of each retinal ganglion cell (Fig 3.1 B). This is a first-order description of how that neuron encodes visual stimuli (see Fig 3.1). Typically the spatial receptive field includes an excitatory central region and an inhibitory surround. This shape reflects neural pathways for convergence and lateral connectivity

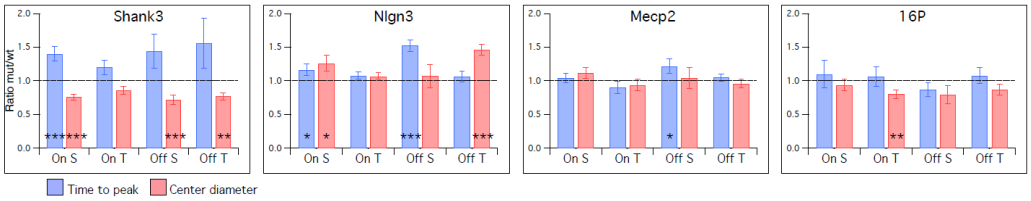

Figure 3.3: The spatio-temporal response properties of Alpha retinal ganglion cells. The diameter and the time-to-peak of the receptive field center were measured for Alpha cells of 4 different types in the four mutant lines and their sibling controls. The ratio between the mutant and control value is plotted, along with its standard error. P-values are derived from a t-test comparing the ratio to 1.0; * p<0.05, ** p<0.01, *** p < 0.001.

from photoreceptors to retinal ganglion cells. By contrast the time course of the temporal receptive field kernel reports on the dynamics of integration within the retina, which is controlled primarily by cellular parameters of the component neurons, such as the membrane time constant and speed of synaptic transmission. The mouse retina comprises at least 30 types of retinal ganglion cell (Sanes and Masland, 2015) with a wide range of receptive field properties. For a reliable comparison across animals one must therefore focus on cells of the same type. Here we report on the so-called alpha retinal ganglion cells. The alpha cells produce large action potentials and have some of the largest receptive fields; thus they can be identified reliably in extracellular recordings. Alpha cells come in four types that are distinguished easily by the categorical differences in their response to light flashes (Krieger et al., 2017) : There are two Off-types and two On-types. Each pair has one type with a transient response to a light step and the other with a sustained response.

For each RGC we summarized the receptive field by just 2 parameters that reflect spatial and temporal pooling respectively: the diameter of the receptive field center and the time-to-peak of the time course (Fig 3.1 B). These were measured for neurons from each of the 4 alpha RGC types in each of the 4 mutant lines and their sibling controls. Then we computed how much that response parameter differs between mutants and their sibling controls. If there were no difference, one expects a ratio of 1.0 in all comparisons. The observed ratios are summarized in Fig 3.3.

The most dramatic result appeared in the Shank3 line: In the mutant, the time-to-peak was considerably longer, and the center diameters were smaller. This held for all 4 alpha RGC types, although more data would help in some cases to confirm statistical significance. Effectively visual processing by these circuits was slower than in the normal retina but also more restricted in space. The effect size is substantial, up to 40 % or 50 % for both the spatial and temporal parameters.

A similar trend was seen in the 16P mutant mice, although the effect size there is smaller. In the Nlgn3 mutant mice, visual processing was again slower, especially in the On-S and Off-S types. Unlike in Shank3, however, the receptive field diameters were larger than in non-mutant animals, especially for the On-S and Off-T types. We conclude that there are substantial effects for some of these mutations that can be resolved as quantitative differences in spatial and temporal processing. Of the 32 comparisons tested in this effort, 10 yielded departures from the normal retina that were significant at alpha = 0.05. At the same time, more data would help to pin these effects down further and perhaps resolve additional ones.

Specialty circuits for motion processing

Motion detection is a vital function for animal behavior, and it begins at the level of the retina. So-called direction-selective ganglion cells (DSGCs) fire when an object moves in the "preferred" direction but not in the opposite "null" direction. Underlying this computation is a rather intricate circuitry, involving several types of bipolar cells, nonlinear dendritic processing in at least two types of amacrine cells, a highly selective synaptic connectivity onto the retinal ganglion cell dendritic tree, and a balanced interplay of excitation and inhibition (Borst and Helmstaedter, 2015). There exist DSGCs of all three major response types: On, Off, and On-Off (Sanes and Masland, 2015). Given this complexity, one expects that DSGCs would be sensitive reporters of any circuit abnormality.

To our surprise, the Nlgn3 mutant displayed all 3 major types of DSGCs (see also Hoon, Krishnamoorthy, et al. (2017), reporting on a different Nlgn3 knockout mutant). In contrast, the Mecp2 mutant had only On-Off DS cells. Furthermore, these DSGCs in this mutant showed distinctly abnormal light responses (Fig 3.4): When probed with a flashing spot they fired at both ON and OFF transitions, as expected. But when probed with moving bars they responded strongly only to the black bar, not the white bar. Clearly, the wiring leading to the ON-OFF DS cells is perturbed in the Mecp2 line, and this will offer a useful handle for follow-up mechanistic studies (Aim 2 of our original proposal). In the Shank3 and 16p.11.2 mutants we have not found any DSGCs so far, but given the low prevalence of DSGCs even in the normal retina, establishing the significance of this absence will require further data collection.

Next We studied four mouse lines with autism-related mutations in the context of natural behaviors. One study tested the so-called looming reaction: a sequence

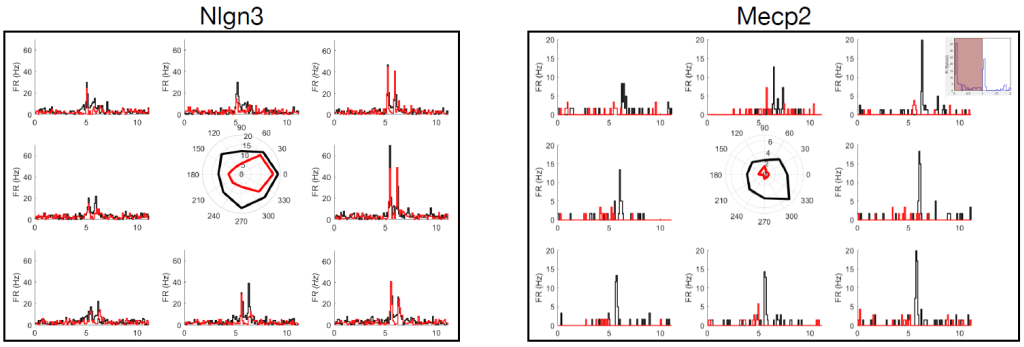

Figure 3.4: Examples of On-Off direction-elective cells in retinas of Nlgn3 and Mecp2 mutants. Responses to moving black or white bars displayed as in Fig 1C. Note the Nlgn3 cell produces two bursts of firing under both black and white bars: one to the leading edge the other to the trailing edge. The Mecp2 cell produces only one burst. Nevertheless it has an On-Off response to flashing lights, see inset top right.

of behaviors by which rodents escape from aerial predators (Yilmaz and Markus Meister, 2013). In the normal mouse, this behavior involves visual perception, discrimination of closely-related stimuli, directed escape, and freezing. The other study observed the mice in a complex labyrinth (Rosenberg et al., 2021). The animals are water-deprived and a source of water is hidden in the labyrinth. The normal mouse engages in a sequence of behaviors including exploration of the environment, the discovery of the water port, and rapid learning of the shortest possible route for navigation to the port.

3.3.2 Looming reaction behavior

MECP2 mutant mice (n=3) showed clear anomalies in their reaction to looming stimuli (Figure 3.5). In particular, they lacked any notable response to black expanding disks, which would typically induce a robust freeze and/or escape, as seen in both the non-mutant litter-mates (n=4) and C57Bl6/J WT mice (Yilmaz and Markus Meister, 2013). In addition, these mutant animals froze and/or escaped in response to white disk stimuli, which indicates that their lack of response is not attributable to gross motor deficits.

We also found that the non-mutant litter-mates differed considerably from the C57Bl6/J strain (Yilmaz and Markus Meister, 2013), producing escape responses also to the ecologically innocuous shrinking disk stimuli (4 of 4 animals, Figure 3.5). Even though Jackson Labs uses C57Bl6/J males for breeding this line, there may be remnants of genetic background other than the mutation of interest that affects the behavior of the mutant animals. We had reported substantial background

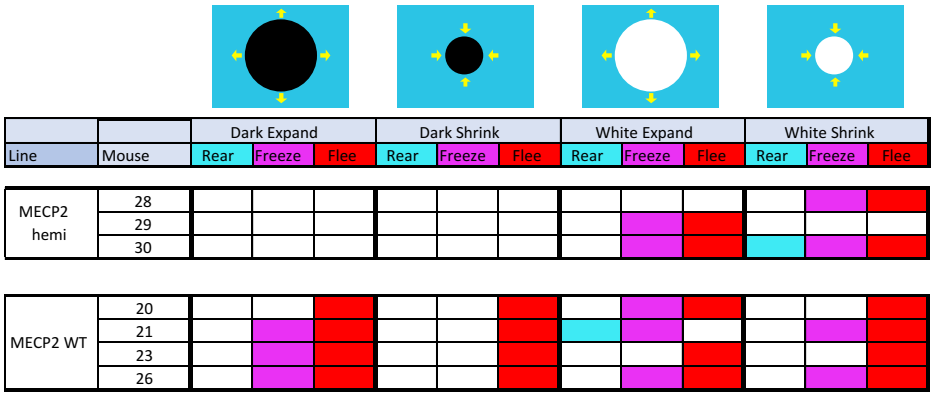

Figure 3.5: Reactions of MECP2 mutant mice and non-mutant siblings to the classic looming stimulus and comparison displays. Each row is an individual mouse.

effects also in our retinal physiology study of Year 1.

By contrast, our analysis of retinal function in the MECP2 animals revealed no overt deficit in visual signaling, as presented in our Year 1 report. In particular, the proportion of ON-type and OFF-type ganglion cells is statistically within the normal range. Also, a detailed analysis of Spatio-temporal receptive fields of the alpha-RGCs revealed no peculiarities. A new study published today claims that one of these alpha cell types (Off-transient) carries the visual signal essential for the looming reaction Wang et al., 2021. We suggest therefore that the behavioral deficit is not caused by retinal malfunction, but arises later in the visual pathway or in the brain centers that coordinate the defensive reaction.

The other 3 mutant lines – NLGN3 (n=4), SHANK3 (n=4), and 16P (n=4) – all showed statistically normal reactions to the four stimuli. This included appropriate escape and freezing behaviors. The execution of these behaviors, for example, the reaction time and speed of escape, appeared within normal limits as well. It is conceivable that a study with larger numbers of animals would reveal some small anomalies in these mutant lines, but at this stage, we find no substantial deficits in the looming reaction. This indicates that visual processing, the discrimination of different stimuli, and the behavioral control of escape and freezing are largely intact.

3.3.3 Maze learning Results

Initial response to the maze

Mice of lines NLGN3 and MECP2 entered the maze within a few minutes of the tunnel opening, similar to WT (C57Bl6/J) mice (Rosenberg et al., 2021). By

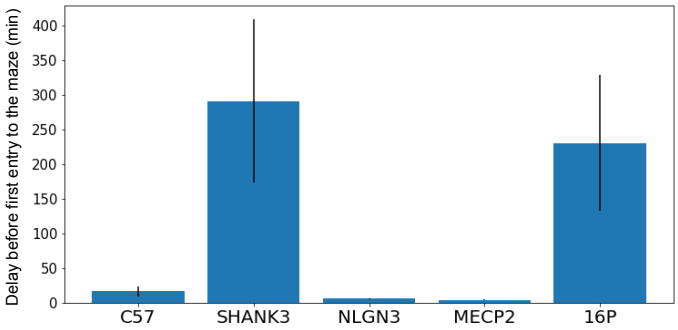

Figure 3.6: Delay between tunnel opening and the animal's first entry into the maze. Mean ± SE across 10 animals for the C57Bl6/J wild type, and 3 animals each for the mutant lines.

contrast mice of lines 16P and SHANK3 were extremely reluctant to enter the maze, requiring several hours before crossing the tunnel for the first time (Figure 3.6).

Some of the mice of lines 16P and SHANK3 never entered the maze over the 7 hours of the study. For these animals, we extended the study to a second night in which we placed the animal directly into the tunnel leading to the labyrinth and temporarily closed off the opening to the home cage. From that point on these "reluctant" mice entered the maze, explored it thoroughly, and then alternated freely between the home cage and maze on a schedule similar to that of WT mice.

Early bouts of exploration

During the first entries into the maze, wild-type mice are somewhat tentative: They spend only a few seconds at a time there, and over several visits, they gradually make their way from the entrance to the first intersection (Rosenberg et al., 2021). Mutant mice from the lines NLGN3, SHANK3, and 16P began their explorations in a similarly hesitant manner (Figure 3.7A). However, the MECP2 mice acted very differently: On the first bout, they typically explored the entire maze to the deepest level, often discovering the water port in the process (Figure 3.7B). Over the first 5 visits to the maze the total time spent in the maze was 11.9 ± 2.5 min for WT mice, but 42.9 ± 10.7 for MECP2 mice. The other 3 mutant lines all acted similarly to the WT animals (Figure 3.7C).

The MECP2 trajectories involved a great amount of repetition. For example, the mouse would focus on one-quarter or one-half of the maze and repeatedly cycle through all the endpoints of the maze there, then move to another location for a similar set of repeat sequences (see Figure 3.7B). Eventually, after several visits to the maze, this repetitive behavior ceased, and the mouse adopted a schedule

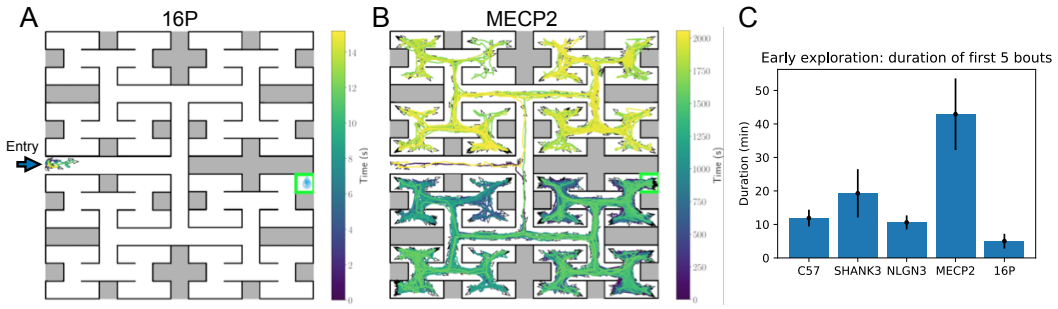

Figure 3.7: Sample trajectories during early visits to the maze for a 16P mouse (A) and a MECP2 mouse (B). Arrowhead marks entry into the maze, drop symbol marks the water port. The trajectory of the animal's nose is color coded with time in maze from purple (early) to yellow (late). (C) Time spent on the first 5 bouts for the WT (C57) and mutant mice, mean ± SE.

of brief visits much like the other mutant lines and the wild-type mice. In this repetitive behavior, MECP2 mice are uncannily similar to mice that lack most of the neocortex and hippocampus, either through acute lesions or developmental mutations (our unpublished data).

Learning of the water location

At the start of the experiment the animal is water-deprived but has no knowledge of the water port in the maze. Eventually, it discovers the port and receives a water reward for poking its nose at it. Then the port remains inactive for a time-out delay of 90 s, and the animal typically leaves the site to explore the maze further or return to the cage (Rosenberg et al., 2021). After a few water rewards, the animal learns the location of the port, and returns to it on a regular schedule. To reach the water port from the maze entrance the animal must make 6 correct turns in a row (each among 3 options, see Figure 3.7A).

Figure 3.8 shows the timing of every water reward received for all the mice in this study, along with a group of 10 wild-type animals (Rosenberg et al., 2021). Note that the time course of learning is very similar across all these mice. They typically discovered the water port after a few minutes and then started collecting water at a steady rate after about 10 rewards. One of the MECP2 animals (ME32) collected rewards at a rather low rate. And the 3 NLGN3 animals did so at a relatively high rate. But overall, the animals were efficient at finding the port and exploiting it, and there is no dramatic difference between the mutants and the wild-type mice.

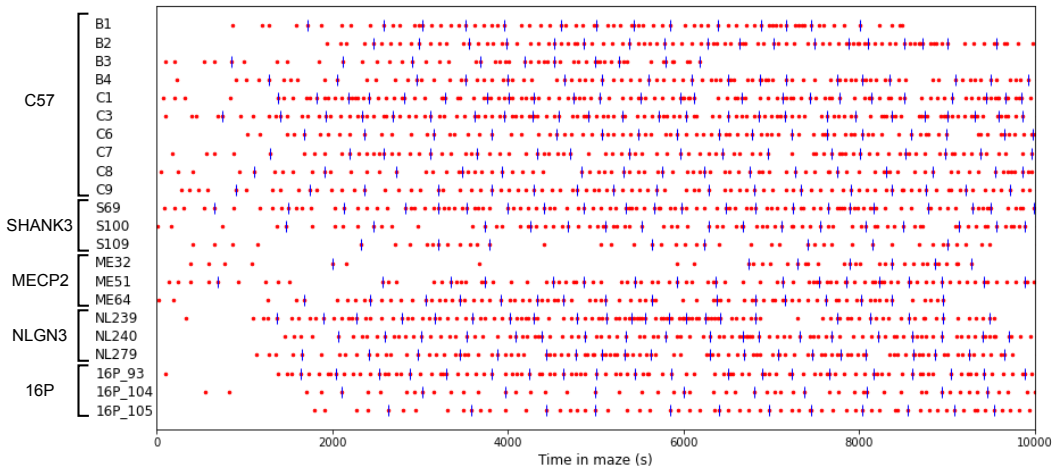

Figure 3.8: Timeline of all water rewards collected by the mice. Each row is one animal. Red dots indicate the collection of a water reward, blue ticks mark every fifth reward. Data from the C57 animals were previously reported in Rosenberg et al., 2021.

Exploration vs exploitation

To inspect the time course of learning and exploitation more closely we analyzed the mouse's 'water runs', from the maze entrance to the first visit to the water port. Normal WT mice follow a characteristic pattern shown in Figure 3.9A: Early on the mouse travels a long distance through the maze, but after a few bouts the animal learns to access the water port via the shortest possible path (6 steps between nodes of the maze). Eventually, these 'perfect runs' become routine. Nevertheless, even at this late time, the WT animal will frequently begin its bout with a long excursion (see the many runs with length>6 in Fig 3.9A). We assume this is for the purpose of patrolling the maze for newly appearing resources or threats Rosenberg et al., 2021.

Among the mutants we inspected, the NLGN3 mice departed significantly from this schedule. They also learned the perfect path to the water port, in fact somewhat faster than the WT mice ($1/e$ time of 7 rewards vs 10). However, once they mastered this skill, they only used the perfect path, almost 100% of the time (Figure 3.9B). Effectively they valued exploitation over exploration much more than the WT mice. This was remarkably consistent for all 3 animals from that line.

3.4 Discussion

The behavioral phenotyping of autism mouse models has largely employed a battery of standardized tasks for which commercial equipment and consensus analysis recipes exist. Many of these tasks, such as swimming across a pool of milk, have only limited ecological relevance for a mouse. As a complementary approach here

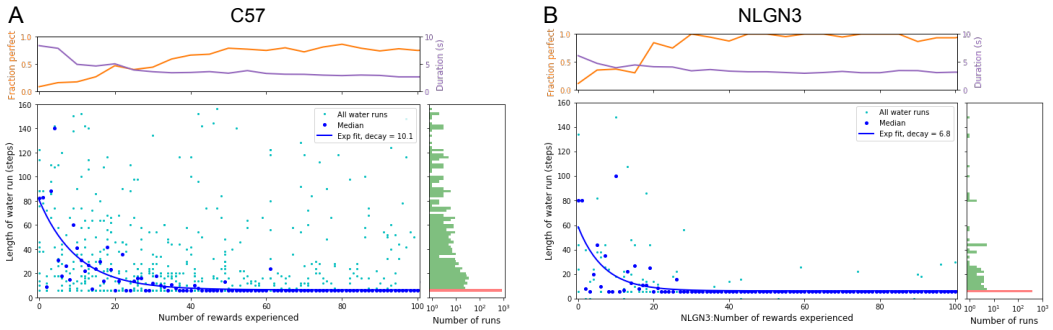

Figure 3.9: The length of runs from the entrance to the water port, measured in steps between nodes, and plotted against the number of rewards experienced. Mice from lines C57 (A) and NLGN3 (B). Main panel: All individual runs (cyan dots) and median overall mice (blue circles). Exponential fit to the median (blue line) decays by $1/e$ over the number of rewards indicated in the legend. Right panel: Histogram of the run length, note log axis; red: perfect runs with the minimum length 6; green: longer runs. Top panel: The fraction of perfect runs (length 6) plotted against the number of rewards experienced, along with the median duration in seconds of those perfect runs.

we focused on behaviors that are closer to the core of the mouse operating system: escaping from predators, finding resources in a complex environment, remembering where they are, and navigating back to those locations. All the behaviors we studied are innate, they require no conditioning, and we report only results from the animal's first exposure to the task. Do the mutations used to model autism disrupt those core mouse functions?

Somewhat surprisingly we found that the four mutant mouse lines examined here were remarkably competent at these innate behaviors, with a few exceptions. In Figure 3.10 we list the most significant abnormalities encountered, including our earlier results on retinal processing. Here we discuss some of these findings.

Looming response of the MECP2 mutants

The MECP2 mutant mouse stands out because it failed to respond to the classic looming stimulus (a black expanding disk presented overhead) while initiating escape in response to other stimuli. At the same time, retinal processing was largely normal in this mutant, with no notable defect in the OFF channel.

The MECP2 null mutation is a mouse model for Rett's syndrome (Chahrour and Zoghbi, 2007). Visual abnormalities have been reported in both Rett's patients and MECP2 mice. Both exhibit lower amplitude visual evoked potentials (VEPs) in response to drifting grating stimuli, as well as deficits in discriminating smaller

	SHANK3	NLGN3	MECP2	16P
RGCs ON/Off Ratio	excess ON			excess ON
α-RGC temporal processing	slow	some slow		
α-RGC receptive field size	small	some large		some small
Looming response			no response to black disks	
Initial reponse to maze	no entry to tunnel			no entry to tunnel
Early exploration			long repetitive bouts	
Exploitation of perfect path		faster learning pure exploitation		

Figure 3.10: A summary of abnormalities encountered in aspects of retinal visual processing, the looming reaction, and maze learning and navigation. Only the most significant departures are listed, obtained by comparison to wild-type or non-mutant sibling mice.

patterns at higher spatial frequencies (LeBlanc et al., 2015). In addition, MECP2-deficient mice have a weak optomotor response to high-contrast moving gratings, suggesting lower visual acuity than the wild-type (Durand et al., 2012). *Vice versa*, male mice with a duplication of the MECP2 gene have been reported with superior visual detection performance, as well as a preference for higher spatial frequency stimuli in V1 neurons, suggesting higher visual acuity and contrast sensitivity (D. Zhang et al., 2017). Taken together with our detailed study of retinal ganglion cell signaling, the literature suggests that a defect in the looming response of MECP2 null mutants likely arises in the visual system downstream of the retina.

The other three mutant lines (NLGN3, SHANK3, 16P) appeared similar to non-mutant sibling controls in their looming reaction. A recent report claims an impaired looming response in the valproic acid mouse model of autism (Hu et al., 2017) and attempts to relate that to the visual reactions of autistic children. Our results emphasize that this behavioral abnormality cannot be used as a common phenotype of autism model mice.

Transient maze-phobia of SHANK3 and 16P mutants

Mice of lines SHANK3 and 16P have a remarkably similar profile in our survey, extending across retinal and behavioral phenotypes (Figure 3.10). Both were highly reluctant to enter the maze at first, an unusual behavioral trait among all the mice we have observed so far. It may be tempting to diagnose this reaction as “anxiety”, which is commonly listed as a phenotype in autism model mice. However, researchers generally measure anxiety in mice based on whether they explore the center of an open field. If you are a mouse then traversing an open field is dangerous, so avoiding it is the rational choice. The mouse’s aversion towards open spaces is innate and cannot be overcome by crossing the center of an open field once.

By contrast, whatever internal state kept these mutant animals from entering the tunnel to the maze, it was overcome entirely by a single forced "exposure therapy”. After being placed in the maze environment involuntarily once, the animals quickly adapted and then explored the maze in a normal fashion on their own initiative. If anything, this adjustment can be seen as a mark of cognitive flexibility and an instance of one-shot learning. It would be interesting to see whether rodents can improve on other cognitive tests as well after a brief behavioral intervention.

Efficient spatial learning in NLGN3 mutants

The NLGN3 mutant line we investigated carries the R451C substitution in the neuroligin-3 protein that is found in a small fraction of human ASD patients (Tabuchi et al., 2007). The original report on this mouse had already noted that it seems to perform better than wild-type in the Morris milk pool test [1]. Specifically, the mutant reaches the hidden platform after swimming a shorter distance and for a shorter time. However, this may reflect better spatial memory in the mutant or greater exploitation of the same spatial memory. While the milk pool leaves that unresolved, the more complex environment of the labyrinth allows a distinction. Late in the experiment, it is clear that both normal and mutant animals have completely mastered the perfect run to the water port; yet the wild-type mouse gets "distracted" by preferring to explore the environment on $\sim 25\%$ of the bouts (Figure 3.9A). The single-mindedness of the NLGN3 mutant mouse may show some parallels to the enhanced cognitive abilities in high-functioning individuals with ASD.

[1] I avoid the conventional term "water maze" here because the test wouldn’t work with water and it isn’t a maze and it may be confused with the real maze in our experiments

3.5 Methods

Animals

We acquired four lines of mutant mice from Jackson Labs: NLGN3 (Stock # 008475), SHANK3 (Stock # 017688), MECP2 (Stock # 003890), and 16P (16p11.2, Stock # 013128). The choice of these lines was partly informed by their use in other SFARI studies, which promises a future integration of our results with other insights. Heterozygous animals were mated and the offspring were genotyped. For experiments, we chose male mice, hemizygous for the X-linked mutations (NLGN3, MECP2), heterozygous (16P), or homozygous (SHANK3), as well as non-mutant litter-mates of each line, and wild-type mice of strain C57Bl6/J.

Retina as a Probe Method

Retinal physiology: The retina was isolated from the eye and placed ganglion-cells-down in a dish with a 256-electrode array (Markus Meister, Pine, and Baylor, 1994; S. Nirenberg and M. Meister, 1997). Movies were projected onto the photoreceptor layer while the electrodes recorded spikes from many retinal ganglion cells. The stimuli were drawn from a broad palette: flashing spots, white-noise flicker, moving bars, shifting gratings (Baden et al., 2016; Y. Zhang et al., 2012). We used the resulting responses both to identify different categories of ganglion cells and to measure their visual function with quantitative parametric models, following established methods (Chichilnisky, 2001).

Looming reaction Methods

The experiments were conducted according to the protocol of Yilmaz (2013) (Yilmaz and Markus Meister, 2013). In brief:

Arena: The mouse was placed in an empty square arena (ca 50 x 50 cm) that included a shelter made of opaque material in one corner. The roof of the arena was a display monitor. The animal was acclimated to the arena for ~10 min. Then the operator triggered a visual stimulus on the display monitor that lasted ~10 s. The animal's reaction was tracked by multiple video cameras.

Stimuli: The classic looming stimulus is an expanding black disk, intended to approximate the appearance of an approaching bird. A stimulus episode consisted of the expanding disk repeated 10 times at intervals of 1 s. For comparison, we also used stimuli that share some visual features but not the ecological significance:

expanding white disk; contracting white disk; contracting black disk. An animal was exposed to only one of these stimuli per day.

Analysis: Based on the video recordings we scored the behaviors of the animal around the time of the visual stimulus. We focused on 3 motifs: 'rearing', a posture by which the animal stretches towards the ceiling; 'escape', a sudden acceleration and rapid locomotion that typically ends in the shelter; and 'freezing', a period of prolonged immobility.

Maze learning Methods

Here we followed the protocol of (Rosenberg et al., 2021). In brief:

Behavior and recording: The mouse's home cage was connected by a short tunnel to a labyrinth (~ 60 x 60 x 5 cm), a system of short corridors about the width of a mouse, connected by T-junctions. The labyrinth's design followed a binary tree, with 6 levels of branches and 64 endpoints (see Figure 3.7). One of these endpoints was fitted with a water port that dispensed a single drop of water in response to a nose poke. The construction materials were visually opaque and the experiments were conducted in darkness. A video camera below the labyrinth monitored the animal's movements via infrared illumination. The mouse was mildly water-deprived and then placed in the home cage or into the labyrinth at the start of the subjective night. The animal was free to enter and leave the labyrinth as desired. We recorded its behavior continuously over the subsequent 7 hours, with no interference by the investigator. All results are from the animal's first encounter with the labyrinth.

Analysis: From the video recording we tracked 7 key points on the mouse body using DeepLabCut (Nath et al., 2019). Subsequent analysis focused on the trajectory of the animal's nose through the labyrinth. For the present report, we analyzed four components of the behavior: (1) The readiness with which the animal first entered the labyrinth; (2) how efficiently it explored the new environment; (3) how fast it learned the location of the water port; (4) how consistently it exploited the shortest route to the port. The behavior of mutant lines was compared to that of C57Bl6/J ("WT") reported (Rosenberg et al., 2021). Details of the analysis including all the code used are available there (Rosenberg et al., 2021).

3.6 Significance

This study introduces a comparative analysis across 4 mutant mouse lines related to ASD. It includes single-neuron measurements of visual processing in the retina,

a visually-guided defensive behavior, and a suite of behaviors that support spatial learning and navigation in a complex labyrinth. This was a pilot study, and its tentative conclusions need to be probed more deeply. However, some potentially impactful results include:

- These natural behaviors involving predator defense, spatial learning, and navigation are largely intact in all these ASD-related mutants. Conceivably the brain implements such vital behaviors more robustly and redundantly than some of the abstract tasks that animals must complete in conventional behavioral phenotyping. Is this a general principle?
- We found a similar physiological and behavioral profile in two lines with entirely different mutations: SHANK3 and 16P. Is this merely a coincidence, or do these two lines perhaps express a common intermediate-level defect?
- We replicated an intriguing case of enhanced cognitive function in the NLGN3 line and propose an alternative interpretation of its nature.

References

Baccus, Stephen A. et al. (2008). "A Retinal Circuit That Computes Object Motion". In: *Journal of Neuroscience* 28.27, pp. 6807–6817. DOI: 10.1523/jneurosci.4206-07.2008. URL: http://www.jneurosci.org/content/28/27/6807.abstract.

Baden, T et al. (2016). "The functional diversity of retinal ganglion cells in the mouse." In: *Nature* 529, pp. 345–350. DOI: 10.1038/nature16468. URL: http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=26735013.

Borst, Alexander and Moritz Helmstaedter (Aug. 2015). "Common circuit design in fly and mammalian motion vision". eng. In: *Nature Neuroscience* 18.8, pp. 1067–1076. ISSN: 1546-1726. DOI: 10.1038/nn.4050.

Carcieri, Stephen M., Adam L. Jacobs, and Sheila Nirenberg (Sept. 2003). "Classification of Retinal Ganglion Cells: A Statistical Approach". In: *Journal of Neurophysiology* 90.3. Publisher: American Physiological Society, pp. 1704–1713. ISSN: 0022-3077. DOI: 10.1152/jn.00127.2003. URL: https://journals.physiology.org/doi/full/10.1152/jn.00127.2003 (visited on 09/15/2022).

Chahrour, Maria and Huda Y. Zoghbi (Nov. 2007). "The Story of Rett Syndrome: From Clinic to Neurobiology". eng. In: *Neuron* 56.3, pp. 422–437. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2007.10.001.

Chen, Jason A. et al. (2015). "The Emerging Picture of Autism Spectrum Disorder: Genetics and Pathology". In: *Annual Review of Pathology: Mechanisms of Disease* 10.1, pp. 111–144. DOI: 10.1146/annurev-pathol-012414-040405.

Chichilnisky, E J (2001). "A simple white noise analysis of neuronal light responses". In: *Network* 12.2, pp. 199–213. 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.

Durand, Severine et al. (Dec. 2012). "NMDA Receptor Regulation Prevents Regression of Visual Cortical Function in the Absence of Mecp2". eng. In: *Neuron* 76.6, pp. 1078–1090. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2012.12.004.

Gollisch, Tim and Markus Meister (2010). "Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina". In: *Neuron* 65.2, pp. 150–164. ISSN: 08966273. DOI: 10.1016/j.neuron.2009.12.009. URL: http://www.cell.com/neuron/abstract/S0896-6273(09)00999-4.

Helmstaedter, Moritz et al. (2013). "Connectomic reconstruction of the inner plexiform layer in the mouse retina". In: *Nature* 500.7461, pp. 168–174. DOI: 10.1038/nature12346http://www.nature.com/nature/journal/v500/n7461/abs/nature12346.html#supplementary-information. URL: http://dx.doi.org/10.1038/nature12346.

Hoon, Mrinalini, Vidhyasankar Krishnamoorthy, et al. (July 2017). "Loss of Neuroligin3 specifically downregulates retinal GABAA2 receptors without abolishing direction selectivity". In: *PLoS ONE* 12.7, e0181011. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0181011. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5510863/ (visited on 09/15/2022).

Hoon, Mrinalini, Haruhisa Okawa, et al. (Sept. 2014). "Functional Architecture of the Retina: Development and Disease". In: *Progress in retinal and eye research* 42, pp. 44–84. ISSN: 1350-9462. DOI: 10.1016/j.preteyeres.2014.06.003. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4134977/ (visited on 09/15/2022).

Hu, Yu et al. (Nov. 2017). "A Translational Study on Looming-Evoked Defensive Response and the Underlying Subcortical Pathway in Autism". In: *Scientific Reports* 7, p. 14755. ISSN: 2045-2322. DOI: 10.1038/s41598-017-15349-x.

Krieger, Brenna et al. (2017). "Four alpha ganglion cell types in mouse retina: Function, structure, and molecular signatures". In: *PLoS ONE* 12.7, pp. 1–21. ISSN: 19326203. DOI: 10.1371/journal.pone.0180091.

Krishnaswamy, Arjun et al. (Aug. 2015). "Sidekick 2 directs formation of a retinal circuit that detects differential motion". eng. In: *Nature* 524.7566, pp. 466–470. ISSN: 1476-4687. DOI: 10.1038/nature14682.

LeBlanc, Jocelyn J. et al. (2015). "Visual Evoked Potentials Detect Cortical Processing Deficits in Rett Syndrome". en. In: *Annals of Neurology* 78.5, pp. 775–786. ISSN: 1531-8249. DOI: 10.1002/ana.24513.

Meister, Markus, Jerome Pine, and Denis A. Baylor (1994). "Multi-neuronal signals from the retina: acquisition and analysis". In: *Journal of Neuroscience Methods* 51.1, pp. 95–106.

Nath, Tanmay et al. (July 2019). "Using DeepLabCut for 3D Markerless Pose Estimation across Species and Behaviors". en. In: *Nature Protocols* 14.7, pp. 2152–2176. ISSN: 1750-2799. DOI: 10.1038/s41596-019-0176-0.

Nirenberg, S. and M. Meister (Apr. 1997). "The light response of retinal ganglion cells is truncated by a displaced amacrine circuit". eng. In: *Neuron* 18.4, pp. 637–650. ISSN: 0896-6273. DOI: 10.1016/s0896-6273(00)80304-9.

Rosenberg, Matthew et al. (Jan. 2021). "Mice in a Labyrinth: Rapid Learning, Sudden Insight, and Efficient Exploration". en. In: *bioRxiv*, p. 2021.01.14.426746. DOI: 10.1101/2021.01.14.426746.

Sanes, Joshua R. and Richard H. Masland (July 2015). "The types of retinal ganglion cells: current status and implications for neuronal classification". eng. In: *Annual Review of Neuroscience* 38, pp. 221–246. ISSN: 1545-4126. DOI: 10.1146/annurev-neuro-071714-034120.

Tabuchi, Katsuhiko et al. (Oct. 2007). "A Neuroligin-3 Mutation Implicated in Autism Increases Inhibitory Synaptic Transmission in Mice". In: *Science (New York, N.Y.)* 318.5847, pp. 71–76. ISSN: 0036-8075. DOI: 10.1126/science.1146221.

Wang, Fei et al. (Apr. 2021). "OFF-Transient Alpha RGCs Mediate Looming Triggered Innate Defensive Response". English. In: *Current Biology* 0.0. ISSN: 0960-9822. DOI: 10.1016/j.cub.2021.03.025.

Wässle, Heinz (Oct. 2004). "Parallel processing in the mammalian retina". en. In: *Nature Reviews Neuroscience* 5.10. Number: 10 Publisher: Nature Publishing Group, pp. 747–757. ISSN: 1471-0048. DOI: 10.1038/nrn1497. URL: https://www.nature.com/articles/nrn1497 (visited on 09/15/2022).

Yilmaz, Melis and Markus Meister (2013). "Rapid innate defensive responses of mice to looming visual stimuli". In: *Current Biology* 23.20, pp. 2011–2015. ISSN: 1879-0445 (Electronic) 0960-9822 (Linking). DOI: 10.1016/j.cub.2013.08.015. URL: http://www.ncbi.nlm.nih.gov/pubmed/24120636.

Zeng, Hongkui and Joshua R. Sanes (Sept. 2017). "Neuronal cell-type classification: challenges, opportunities and the path forward". eng. In: *Nature Reviews. Neuroscience* 18.9, pp. 530–546. ISSN: 1471-0048. DOI: 10.1038/nrn.2017.85.

Zhang, Dinghong et al. (July 2017). "Altered Visual Cortical Processing in a Mouse Model of MECP2 Duplication Syndrome". En. In: *Scientific Reports* 7.1, p. 6468. ISSN: 2045-2322. DOI: 10.1038/s41598-017-06916-3.

Zhang, Yifeng et al. (2012). “The most numerous ganglion cell type of the mouse retina is a selective feature detector”. In: *Proceedings of the National Academy of Sciences* 109.36, E2391–E2398. DOI: 10.1073/pnas.1211547109. URL: http://www.pnas.org/content/109/36/E2391.abstract.

Chapter 4

ELECTRODE POOLING CAN BOOST THE YIELD OF EXTRACELLULAR RECORDINGS WITH SWITCHABLE SILICON PROBES

Lee, Kyu Hyun*, Yu-Li Ni*, Jennifer Colonell, Bill Karsh, Jan Putzeys, Marius Pachitariu, Timothy D. Harris, and Markus Meister (2021). "Electrode pooling can boost the yield of extracellular recordings with switchable silicon probes". In: *Nature Communications* 12.1, p. 5245. DOI: 10.1038/s41467-021-25443-4.

4.1 Abstract

State-of-the-art silicon probes for electrical recording from neurons have thousands of recording sites. However, due to volume limitations there are typically many fewer wires carrying signals off the probe, which restricts the number of channels that can be recorded simultaneously. To overcome this fundamental constraint, we propose a method called electrode pooling that uses a single wire to serve many recording sites 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 1.0 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.

4.2 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" (BRAIN Working Group, 2014). 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 (J. J. Jun et al., 2017; Dimitriadis et al., 2018; Rios et al., 2016; Torfs et al., 2010; Steinmetz et al., 2021). Still, the best methods sample neural circuits very sparsely,

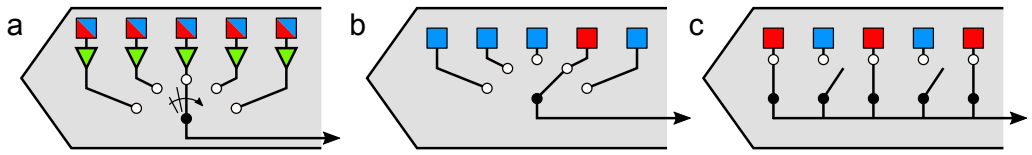

Figure 4.1: Strategies for using a single wire to serve many recording sites in switchable silicon probes. (a) Time-division multiplexing. Rapidly cycling the selector switch allows a single wire to carry signals from many recording sites interleaved in time. Triangles represent anti-aliasing filters. (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.

for example recording fewer than 10^4 cells in a mouse brain that has 10^8 (Stevenson, 2013).

In many of these electrode array 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 (J. J. Jun et al., 2017). The current silicon arrays invariably displace more neurons than they record, and thus the goal of "monitoring all neurons" seems unattainable by simply scaling the present approach (but see Kleinfeld et al., 2019). 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.

A common approach by which a single wire can convey multiple analog signals is time-division multiplexing (Obien et al., 2015). A rapid switch cycles through the N input signals and connects each input to the output line for a brief interval (Figure 4.1a). At the other end of the line, a synchronized switch demultiplexes 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 (Shannon, 1949): It should be at least twice the highest frequency component present in the signal. The raw voltage signals from extracellular electrodes include thermal noise that extends 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 (Shahrokhi et al., 2010), which would force the electrodes apart and prevent any high-density recording. 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 (Eversmann et al., 2003). A recent design with a more modest $N = 8$ still produces noise power 4-15 times higher than in comparable systems without multiplexing (Raducanu et al., 2017).

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, heating up exactly the neurons one wants to monitor. Furthermore, the active electronics in the local amplifier are sensitive to light, which can produce artifacts during bright light flashes for optogenetic stimulation (J. J. Jun et al., 2017; Kozai and Vazquez, 2015).

An alternative approach involves static electrode selection (Figure 4.1b). 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 space is less constrained. 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 (Müller et al., 2015) as well as silicon prongs (J. J. Jun et al., 2017; Lopez et al., 2017; Steinmetz et al., 2021). It allows the user to choose one of many electrodes intelligently, for example because it carries a strong signal from a neuron of interest. This strategy can increase the yield of neural recordings, but it does not increase the number of neurons per wire.

On this background we introduce a third method of mapping electrodes to wires: select multiple electrodes with suitable signals and connect them to the same wire (Figure 4.1c). 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 averaged and transmitted on the same wire. At first that 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 cost, often quoting the low cross-talk between channels as a figure of merit. Instead, we will show that the pooled signal can be unmixed if one controls the switch settings carefully during the recording session. Under suitable conditions this method can record many neurons per wire without appreciable loss of information.

We emphasize that the ideal electrode array device to implement this recording method does not yet exist. It would be entirely within reach of current fabrication capabilities, but every new silicon probe design requires a substantial investment and consideration of various trade-offs. With this article we hope to make the community of electrode users aware of the opportunities in this domain and start a discussion about future array designs that use intelligent electrode switching, adapted to various applications in basic and translational neuroscience.

4.3 Theory

Motivation for electrode pooling: spike trains are sparse in time

A typical neuron may fire ~10 spikes/s on average (Attwell and Laughlin, 2001). 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 4.3b). Sometimes additional neurons lie close enough to the same electrode to produce large spikes. That still leaves most of the time axis unused for signal transmission. 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. These overlaps may be more common under conditions where neurons are synchronized to each other or to external events.

The effects of pooling on spikes and noise

What signal actually results when one connects two electrodes to the same wire? Figure 4.2a shows an idealized 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 electrode impedance Z_i. This impedance

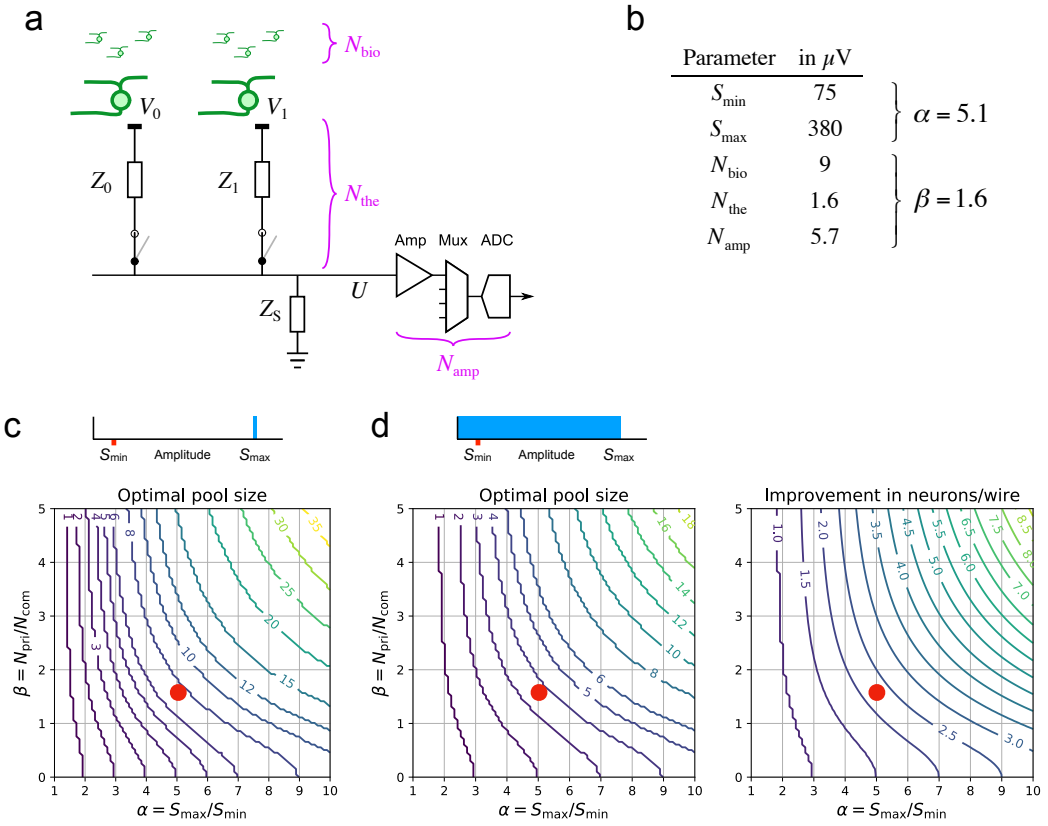

Figure 4.2: Pooling of signal and noise. (a) An idealized circuit for two electrodes connected to a common wire along with downstream components of the signal chain, such as the amplifier, multiplexer, and digitizer. Z_0, Z_1: total impedance for electrodes 0 and 1, with contributions from the metal/bath interface and the external bath. Z_S: shunt impedance at the amplifier input. Noise sources include biological noise from distant neurons (N_{bio}); thermal noise from the electrode impedance (N_{the}), and common electronic noise from the amplifier and downstream components (N_{amp}). (b) Numerical values of the relevant parameters, derived from experiments or the literature (sections 4.4 and 4.5). (c, d) The optimal electrode pool under different assumptions about the spike amplitude distribution (top insets). The contour plots show the optimal pool size and the enhancement of the neuron/wire ratio as a function of the parameters α – the ratio of largest to smallest sortable spike signals – and β – the ratio of private to common noise. (c) Most favorable condition: Each electrode carries a single large spike of amplitude S_{max}, and spikes are sortable down to amplitude S_{min}. In this case the neurons/wire ratio is equal to the pool size. (d) Generic condition: Each electrode carries a uniform distribution of spike amplitudes between 0 and S_{max}. Red dots: Conditions of α and β encountered experimentally, based on the values in panel b.

has contributions from the metal/saline interface and the external electrolyte bath (Seidl, Schwaerzle, et al., 2012; Robinson, 1968), typically amounting to 100 kΩ - 1 MΩ. By comparison the CMOS switches have low impedance, typically ~ 100 Ω (Seidl, Schwaerzle, et al., 2012), which we will ignore.

In general one must also consider the shunt impedance Z_S in parallel to the amplifier input. This can result from current leaks along the long wires as well as the internal input impedance of the amplifier. For well-designed systems, this shunt impedance should be much larger than the electrode impedances; for the Neuropixels device we will show that the ratio is at least 100. Thus one can safely ignore it for the purpose of the present approximations. In that case the circuit acts as a simple voltage divider between the impedances Z_i. If a total of M electrodes are connected to the shared wire, the output voltage U is the average of the signals at the recording sites V_i, weighted inversely by the electrode impedances,

$$U = \sum_{i=1}^{M} c_i V_i \tag{4.1}$$

where

$$c_i = \frac{1/Z_i}{\sum_{j=1}^{M} 1/Z_j} \tag{4.2}$$

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

$$U = \frac{1}{M} \sum_{i=1}^{M} V_i. \tag{4.3}$$

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 4.2a). 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; (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 (Harris et al., 2000). 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 (Schomburg et al., 2012). In that case the private noise

$$N_{\mathrm{pri},i} = \sqrt{N_{\mathrm{the},i}^2 + N_{\mathrm{bio},i}^2}. \tag{4.4}$$

because thermal noise and biological noise are additive and statistically independent.

Finally the noise introduced by the amplifier and data acquisition is common to all the electrodes that share the same wire,

$$N_{\mathrm{com}} = N_{\mathrm{amp}}. \tag{4.5}$$

In the course of pooling, the private noise gets attenuated by the pooling coefficient c_i (Eq 4.2) and summed with contributions from other electrodes. Then the pooled private noise gets added to the common noise from data acquisition, which again is statistically independent of the other noise sources. With these assumptions the total noise at the output has RMS amplitude

$$N_{\mathrm{tot}} = \sqrt{N_{\mathrm{com}}^2 + \sum_{i=1}^{M} c_i^2 N_{\mathrm{pri},i}^2}. \tag{4.6}$$

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

$$N_{\mathrm{tot}} = \sqrt{N_{\mathrm{com}}^2 + N_{\mathrm{pri}}^2 / M}. \tag{4.7}$$

Theoretical benefits of pooling

Now we are in a position to estimate the benefits from electrode pooling. Suppose that the electrode array records neurons with a range of spike amplitudes: 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 the most favorable conditions for pooling one would select electrodes that each carry a single neuron, with spike

amplitude ~ S_{max} (Figure 4.2c). As one adds more of these electrodes to the pool, there comes a point when the pooled signal is so attenuated that the spikes 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 single-site recording, namely

$$\frac{S_{max}/M}{\sqrt{N_{com}^2 + N_{pri}^2/M}} > \frac{S_{min}}{\sqrt{N_{com}^2 + N_{pri}^2}}. \tag{4.8}$$

This leads to a limit on the pool size M,

$$M < M_{max} = \sqrt{\left(\frac{\beta^2}{2}\right)^2 + (1+\beta^2)\alpha^2} - \frac{\beta^2}{2} \tag{4.9}$$

where

$$\alpha = S_{max}/S_{min},\ \beta = N_{pri}/N_{com} \tag{4.10}$$

If one pools more than M_{max} electrodes all the neurons drop below the threshold for sorting. So the optimal pool size M_{max} is also the largest achievable number of neurons per wire. This number depends on two parameters: the ratio of private to common noise, and the ratio of largest to smallest useful spike amplitudes (Eq 4.10). 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 graph in Figure 4.2c.

Next we consider a more generic situation, in which each electrode carries a range of spikes from different neurons (Figure 4.2d). For simplicity we assume a uniform distribution of spike amplitudes between 0 and S_{max}. As more electrodes are added to the pool, all the spikes are attenuated, so the smallest action potentials drop below the sortable threshold S_{min}. Beyond a certain optimal pool size, more spikes are lost in the noise than are added at the top of the distribution and the total number of neurons decreases. By the same arguments used above one finds that the improvement in the number of sortable neurons, n_M, relative to conventional split recording, n_1, is

$$\frac{n_M}{n_1} = \frac{M\left(\alpha - M\sqrt{\frac{1+\beta^2/M}{1+\beta^2}}\right)}{\alpha - 1} \tag{4.11}$$

The optimal pool size $M_{\max}$ is the M which maximizes that factor. The results are plotted in Figure 4.2d.

The benefits of pooling are quite substantial if the user can select electrodes that carry large spikes. For example under conditions of α and β that we have encountered in practice, Figure 4.2c predicts that one can pool 8 electrodes and still resolve all the signals, thus increasing the neuron/wire ratio by a factor of 8. On the other extreme – with a uniform distribution of spike amplitudes – the optimal pool of 4 electrodes increases the neuron/wire ratio by a more modest but still respectable factor of 2.3 compared to conventional recording. The following section explains how one can maximize that yield.

Acquisition and analysis of pooled recordings

With these insights about the constraints posed by signal and noise one can propose an overall workflow for experiments using electrode pooling (Figure 4.3a). A key requirement 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 experiment. In fact the experimenter will benefit from using different switch settings during the same session.

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.

During 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 measures the total noise. The amplifier noise N_{amp} and thermal noise N_{the} can be assessed ahead of time, because they are a property of the recording system, and from them one obtains the biological noise N_{bio}. Now the experimenter has all

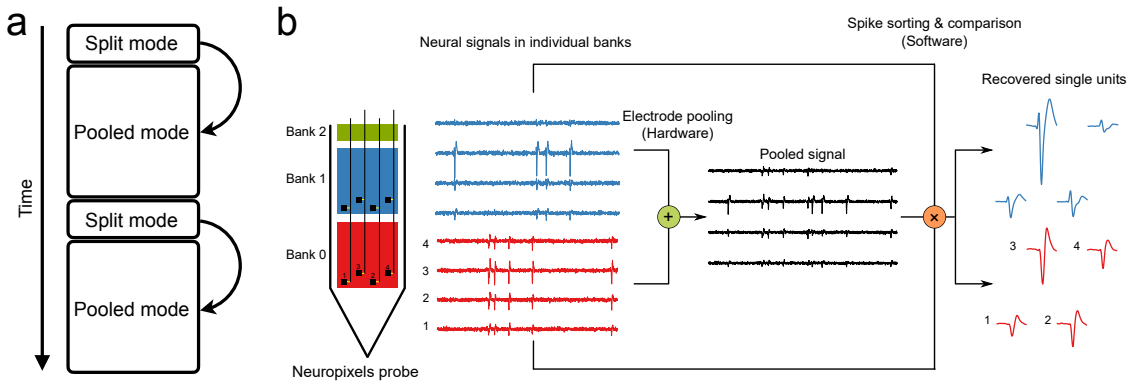

Figure 4.3: Workflow proposed for electrode pooling. (a) Time line of an experiment, alternating short split and long pooled recording sessions. (b) Electrode pooling using the Neuropixels probe. Recording sites (black squares, numbered from 1 to 4) in the same relative location of each bank can be pooled 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 one recovers distinct spike shapes 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.

the information needed to form useful electrode pools. Some principles one should consider in this process:

1. Pool electrodes that carry large signals. Electrodes with smaller signals can contribute to smaller pools.
2. Pool electrodes with distinct spike waveforms.
3. Pool distant electrodes that don't 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 Z_i so the weighted mix can be computed accurately (Eq 4.1). 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 (Lewicki, 1998). 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 (Yger et al., 2018; Pachitariu et al., 2016). 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 recording. If they drift too much, for example because the electrode array moves in the brain (J. Jun et al., 2017), then a recalibration by another split-mode session may be in order (Figure 4.3a). Alternatively electrode drift may be corrected in real time if signals from neighboring electrodes are available (Steinmetz et al., 2021), a criterion that may flow into the selection of switches for pooling. Chronically implanted electrode arrays can record for months on end (Steinmetz et al., 2021), and the library of spike shapes can be updated continuously and scanned for new pooling opportunities.

It should be clear that the proposed workflow relies heavily on automation by dedicated software. Of course automation is already the rule with the large electrode arrays that include thousands of recording sites, and electrode pooling will require little more effort than conventional recording. Taking the newly announced Neuropixels 2.0 as a reference (5120 electrodes and 384 wires): Sampling for 5 minutes from each of the 13 groups of electrodes will take a bit over an hour. Spike-sorting of those signals will proceed in parallel with the sampling and require no additional time. Then the algorithm decides on the electrode pools, and the main recording session starts. Note that these same steps also apply in conventional recording: The user still has to choose 384 electrodes among the 5120 options, and will want to scan the whole array to see where the best signals are. The algorithms we advocate

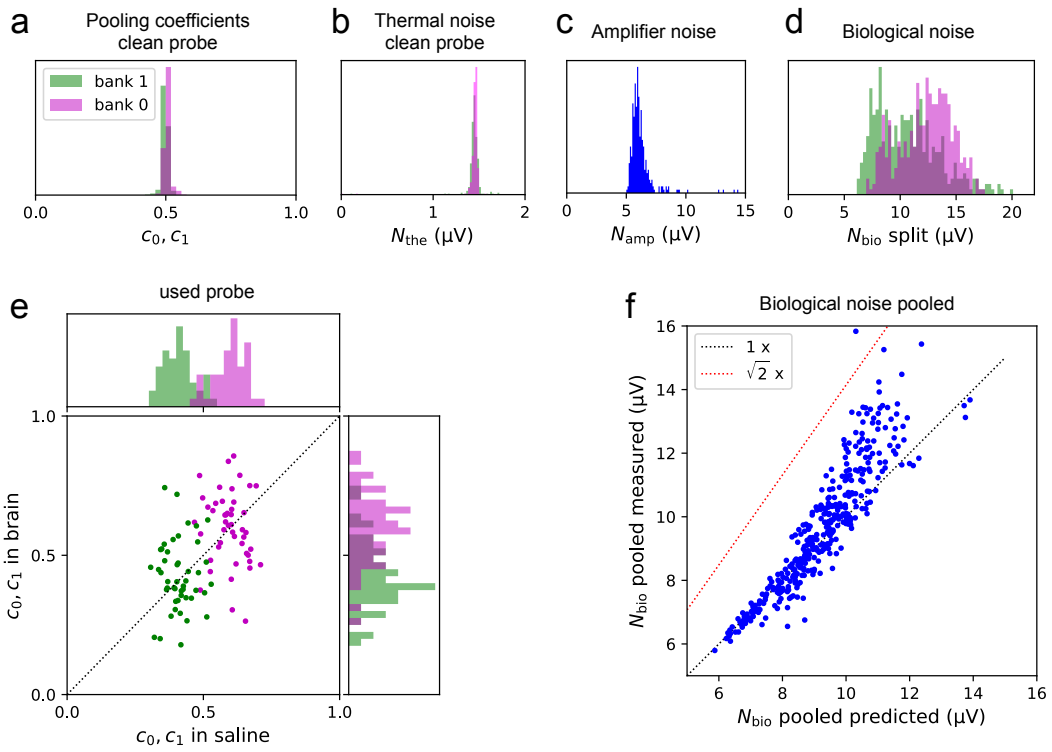

Figure 4.4: Pooling of signal and noise with the Neuropixels 1.0 device. (a) Pooling coefficients on a pristine probe measured in saline, histogram across all sites in banks 0 (red) and 1 (green). (b) Thermal noise (RMS) during split recording in standard saline, histogram across all sites in banks 0 and 1. (c) Amplifier noise (RMS), histogram across all 383 wires. (d) Biological noise (RMS) during brain recordings, histogram across all sites in banks 0 and 1. (e) Pooling coefficients on a used probe, measured in saline (horizontal) vs in brain (vertical). 47 pairs of sites in banks 0 and 1 with suitable action potentials. (f) Biological noise in a pooled recording measured in brain (vertical) vs the prediction derived from assuming uncorrelated noise at the two sites. '1 x': identity. '$\sqrt{2}$ x': expectation for perfectly correlated noise.

to steer electrode pooling will simply become part of the software suite that runs data acquisition.

4.4 Experiments

Pooling characteristics of the Neuropixels 1.0 array

To test the biophysical assumptions underlying electrode pooling, we used the Neuropixels probe version 1.0 (J. J. Jun et al., 2017; Lopez et al., 2017). This electrode array has a single silicon shank with 960 recording sites that can be connected to 384 wires via controllable switches (Figure 4.3b). The electrodes are divided into three banks (called Bank 0, Bank 1, and Bank 2 from the tip to the base of the shank). In the present study only Banks 0 and 1 were used. Banks 0 and 1

each contain 383 recording sites (one channel is used for a reference signal). Each site has a dedicated switch by which it may connect to an adjacent wire. Sites at the same relative location in a bank share the same wire. These two electrodes are separated by 3.84 mm along the shank. Under conventional operation of Neuropixels (J. J. Jun et al., 2017) each wire connects to only one site 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 enabled a limited version of electrode pooling across Banks 0 and 1.

We set out to measure those electronic properties of the device that affect the efficacy of pooling, specifically the split of the noise signal into common amplifier noise N_{amp} (Eq 4.7) and private thermal noise N_{the} (Eq 4.4), as well as the pooling coefficients c_i (Eq 4.2). These parameters are not important for conventional recording, and thus are not quoted in the Neuropixels user manual, but they can be derived from measurements performed in a saline bath (see Methods).

On a pristine unused probe, the pooling coefficients c_0 and c_1 for almost all sites were close to 0.5 (Figure 4.4a), as expected from the idealized circuit (Figure 4.2a) if the electrode impedances are all equal (Eq 4.2). Correspondingly the thermal noise was almost identical on all electrodes, with an RMS value of 1.45 ± 0.10 µV (Fig 4.4b). The amplifier noise N_{amp} exceeded the thermal noise substantially, amounting to 5.7 µV RMS on average, and more than 12 µV for a few of the wires (Figure 4.4c). Because this noise source is shared across electrodes on the same wire, it lowers β in Equation 4.9 and can significantly affect electrode pooling.

Neural recording

Based on this electronic characterization of the Neuropixels probe we proceeded to test electrode pooling in vivo. Recall that each bank of electrodes extends over 3.84 mm of the shank, and one needs to implant more than one bank into the brain to accomplish any electrode pooling. Clearly the opportunities for pooling on this device are limited; nonetheless it serves as a useful testing ground for the method.

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 involved 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 4.3, we then recorded for ~10 min each from Bank 0 and Bank 1 in split mode, followed by ~10 min of recording from both banks

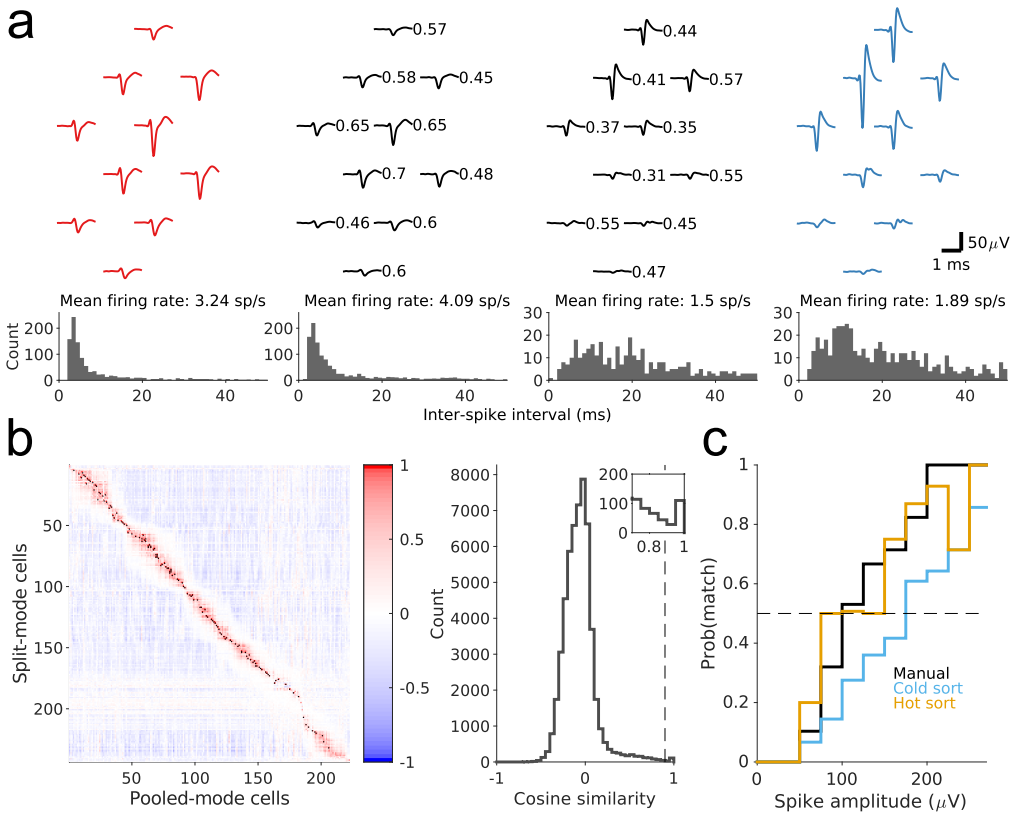

Figure 4.5: Recordings from mouse brain. (a) Matching spike shapes from split- and pooled-mode recordings. Top: Waveforms of two sample units (middle, black) detected in pooled mode on the same set of wires. The left unit was matched to a unit recorded in split mode from Bank 0 (red) and the right unit to one from Bank 1 (blue). Numbers indicate the scaling of the signal of the pooled-mode unit relative to its split-mode signal. Bottom: the mean firing rates and the interspike-interval distributions are similar for the matched pairs. (b) Left: matrix of the cosine similarity between units recorded in pooled- and split-mode, arranged by depth. Black dot indicates greater than the threshold at 0.9. Right: distribution of the cosine similarity. Dashed line indicates threshold at 0.9. Inset zooms into the 0.7-1 range of the distribution. (c) Fraction of units from the two split recordings that are matched to a unit in the pooled recording as a function of spike amplitude. Three different sorting conditions are shown: sorting all recordings by KiloSort1 followed by manual curation (Manual), sorting all recordings by KiloSort2 (Cold sort), and sorting the pooled recording by KiloSort2 with templates initialized from the split recordings (Hot sort). Dashed line indicates 50%, or the 'break-even' point where the pooled-mode yields as many simultaneous recordings as the average split-mode.

simultaneously in pooled mode.

Unmixing a pooled recording

As proposed above, one can unmix the pooled recording by matching its action potentials to the spike waveforms sampled separately on each of the two banks (Figure 4.3b). Each of the three recordings (split Bank 0, split Bank 1, and pooled Banks 0 + 1) was spike-sorted to isolate single units. Then we paired each split-mode unit with the pooled-mode unit that had the most similar waveform, based on the cosine similarity of their waveform vectors (Eq 4.16, Figure 4.5b). In most cases the match was unambiguous even when multiple units were present in the two banks with similar electrode footprints (Figure 4.5a). The matching algorithm proceeded iteratively until the similarity score for the best match dropped below 0.9 (Figure 4.5b). 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 4.5a).

When spike-sorting the pooled-mode recording there is of course a strong expectation for what the spike waveforms will be, namely a scaled version 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 this method to two other procedures (Figure 4.5c): (1) sorting each recording separately, using KiloSort1 with manual curation ("manual") and (2) sorting each recording separately using KiloSort2 with no manual intervention ("cold sorting").

Figure 4.5c illustrates what fraction of the units identified in both split mode recordings combined were recovered from the pooled recording, and how that fraction depends on the spike amplitude. First, this shows that hot sorting significantly outperforms cold sorting, and in fact rivals the performance of manually curated spike sorting. This is important, because manual sorting by a human operator will be unrealistic for the high-count electrode arrays in which electrode pooling may be applied. Second, one sees that the fraction of spikes recovered from the combined split recordings exceeds 0.5 even at moderate spike amplitudes of 100 µV. For spikes

of that amplitude and above the pooled recording will contain more neurons than the average split recording. Clearly, electrode pooling is not restricted to the largest spikes in the distribution, but can be considered for moderate spike amplitudes as well.

Recall that the Neuropixels 1.0 probe is not optimized for electrode pooling, in that it has a fixed switching matrix, and only 2 banks of electrodes fit in the mouse brain. Thus our pilot experiments were limited to brute-force pooling the two banks site-for-site without regard to the design principles for electrode pools. Nonetheless the "hot sorting" method recovered more neurons from the pooled recording (184) than on average over the two split recordings (166). We conservatively focused this assessment only on units identified in the split recordings, ignoring any unmatched units that appeared in the pooled recording. This validates the basic premise of electrode pooling even under the highly constrained conditions. Overall, the above sequence of operations demonstrates that a pooled-mode recording can be productively unmixed into the constituent signals, and the resulting units assigned to their locations along the multi-electrode shank.

Pooling of signal and noise in vivo

Closer analysis of the spike waveforms from split and pooled recordings allowed an assessment of the pooling coefficients in vivo. When spikes are present on the corresponding electrodes in both banks (as in Fig 4.5a) one can measure the pooling coefficients c_0 and c_1 of Eq 4.2. Unexpectedly, instead of clustering near 0.5, these pooling coefficients varied over a wide range (Figure 4.4e), at least by a factor of 3. The two banks had systematically different pooling coefficients, suggesting that the impedance was lower for electrodes near the tip of the array.

Following this in vivo recording we cleaned the electrode array by the recommended protocol (tergazyme / water) and then measured the pooling coefficients in saline. Again the pooling coefficients varied considerably across electrodes, although somewhat less than observed in vivo (Figure 4.4e). Also the bath resistance of the electrodes was larger on average than on an unused probe (30 kΩ as opposed to 13 kΩ). This change may result from the interactions within brain tissue. For example some material may bind to the electrode surface and thus raise its bath resistance. This would lower the pooling coefficient of the affected electrode and raise that of its partners. Because the thermal noise is never limiting (Figure 4.4b-d), such a change would easily go unnoticed in conventional single-site recording. The

precise reason for the use-dependent impedance remains to be understood.

To measure the contributions of biological noise in vivo we removed from the recorded traces all the detected spikes and analyzed the remaining waveforms. After subtracting (in quadrature) the known thermal and electrical noise at each site (Figure 4.4b,c) one obtains the biological noise N_{bio}. This noise source substantially exceeded both the thermal and amplifier noise (Figure 4.4f). It also showed different amplitude on the two banks, presumably owing to differences between brain areas 3.84 mm apart.

Given this large distance between electrodes in the two banks, one expects the biological noise to be statistically independent between the two sites, because neurons near one electrode will be out of reach of the other. To verify this in the present recordings we measured the biological noise in the pooled condition and compared the result to the prediction from the two split recordings, assuming that the noise was private to each site. Indeed the noise in the pooled signal was largely consistent with the assumption of independent noise (Figure 4.4f). It seems likely that the 1-cm shank length on these and similar array devices suffices for finding electrodes that carry independent biological noise.

4.5 Simulations

How many electrodes could experimenters pool and still sort every neuron with high accuracy? Earlier we had derived a theoretical limit to electrode pooling based solely on the signal and noise amplitudes (Figure 4.2). To explore what additional limitations might arise based on the density of spikes in time and the needs of spike sorting we performed a limited simulation of the process (Figure 4.6a). We simulated units with an extracellular footprint extending over 4 neighboring electrodes, and then pooled various such tetrodes into a single 4-channel recording. These pooled signals were then spike-sorted and the resulting spike trains compared to the known ground-truth spike times, applying a popular metric of accuracy (J. Magland et al., 2020). This revealed how many neurons can be reliably recovered depending on the degree of electrode pooling (Figure 4.6b). Then we evaluated the effects of various parameters of the simulation, such as the amplitude of the largest spikes, the biological noise, and the average firing rate.

For simplicity we focused on the favorable scenario of Figure 4.2c: It presumes that the experimenter can choose for pooling a set of tetrodes that each carry a single unit plus noise. The curves of recovered units vs pool size have an inverted-U

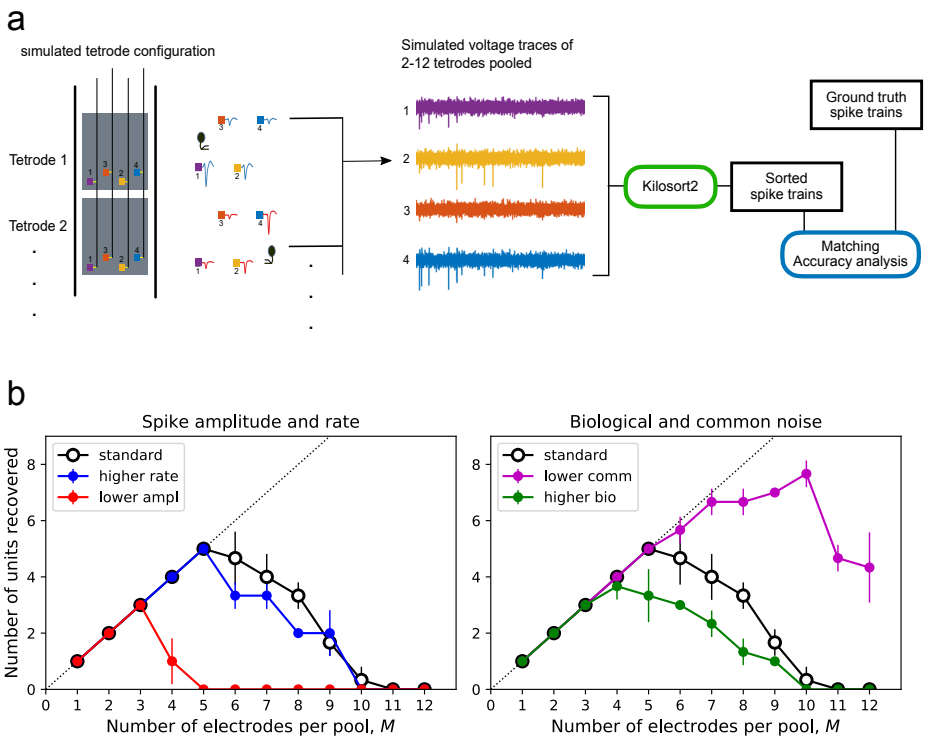

Figure 4.6: Simulations of electrode pooling. (a) Workflow: Groups of 4 recording sites ("tetrodes") each carry a spike train from one simulated unit, superposed with electrode noise and biological noise. Between $M = 1$ and 12 of these tetrodes are then pooled into a single 4-wire recording followed by addition of common noise. The pooled signal is sorted and the resulting single-unit spike trains are matched with the ground truth spike trains from the M tetrodes. Units with an accuracy metric > 0.8 are counted as recovered successfully. (b) Number of units recovered as a function of the pool size, M, under various conditions of simulation. Effects of varying different parameters. The "standard" condition serves as a reference: Spike amplitude $V = 380\ \mu$V, spike rate $r = 10$ Hz, electrode noise $N_{\text{ele}} = 1.6\ \mu$V, common noise $N_{\text{com}} = 5.7\ \mu$V, biological noise $N_{\text{bio}} = 9\ \mu$V. "lower ampl": $V = 205\ \mu$V. "higher rate": $r = 20$ Hz. "higher bio": $N_{\text{bio}} = 15\ \mu$V. "lower com": $N_{\text{com}} = 2.85\ \mu$V. Each parameter combination was simulated 3 times with noise and spike times resampled, error bars are mean ± SD.

shape (Figure 4.6b). For small electrode pools one can reliably recover all the units. Eventually, however, some of the units drop out, and for a large pool size all the recovered units fall below the desired quality threshold. We will call the pool that produces the largest number of recovered units the "optimal pool".

For the "standard" condition of simulations we chose a reasonably large spike amplitude of 380 µV peak-to-peak (the 90-th percentile in a database of recordings by the Allen Institute (Siegle et al., 2019)), a firing rate of 10 Hz, and all the noise values as determined experimentally from the Neuropixels 1.0 device (Figure 4.4). Under these conditions one can pool up to 5 electrodes per wire and still recover all 5

of the units reliably (Figure 4.6b). This optimal pool size is sensitive to the amplitude of the spikes: If the spike amplitude is reduced by a factor of 2, the optimal pool drops from 5 to 3 electrodes. Similarly, if the biological noise increases to 15 µV, the optimal pool is reduced to 4 electrodes. This indicates that the recovery of the units from the pooled signal is strongly determined by the available signal-to-noise ratio at each electrode. By contrast increasing the firing rate two-fold to 20 Hz did not change the optimal pool from 5. Thus the temporal overlap of spikes is not yet a serious constraint. Looking to the future, if the amplifier noise on each wire could be reduced by a factor of 2 the optimal pool would expand significantly from 5 to 7 electrodes or more (Figure 4.6b).

How do these practical results relate to theoretical bounds of Figure 4.2? Recall that this bound depends on the noise properties, but also on the ratio of largest to smallest sortable spikes. In our "standard" simulation with a pool size of 1 (split mode) we found that the smallest sortable spikes had an amplitude of 75 µV. This also corresponds to the low end of sorted spikes reported by the Allen Institute (10-th percentile (Siegle et al., 2019)). With these bounds on large and small spikes, and the measured values of private and common noise, one obtains $\alpha = 5.1$ and $\beta = 1.6$ in Equation 4.9, which predicts an optimal pool of $M_{\max} = 8$ (Fig 4.2c), compared to the actually observed value of 5. The simple theory based purely on signal and noise amplitude gives a useful estimate, but additional practical constraints that arise from temporal processing and spike-sorting lower the yield somewhat from there.

In summary, under favorable conditions where the experimenter can select electrodes, the pooling method may increase the number of units recorded per wire by a factor of 5. Even for significantly smaller spikes or higher biological noise one can expect a factor of 3. And with future technical improvements a factor of 7 or more is plausible.

4.6 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 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 (Figures 4.2, 4.3). We also verified the basic assumptions about signal mixing and unmixing using a real existing device: the Neuropixels 1.0 probe (Figures 4.4, 4.5). We showed that signals from different neurons can be reliably disambiguated and assigned back to the electrodes of origin. 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. In simulations we showed that with a proper selection of electrodes based on the signals they carry, the method could improve the yield of neurons per wire by a factor of 3 to 7 (Figure 4.6).

Electrode pooling is categorically different from most data compression schemes that have been proposed for neural recording systems (Linderman et al., 2008; Olsson and Wise, 2005; Suo et al., 2014). In many of those applications the goal is to reduce the bit rate for data transmission out of the brain, for example using a wireless link. By contrast electrode pooling seeks to minimize the number of electrode wires one needs to stick into the brain to sample the neural signals, thus minimizing biological damage to the system under study. By itself, that doesn't reduce the bit rate, although it produces denser time series. For the optimal wireless recording system, both objectives – lower wire volume at the input and lower data volume at the output – should be combined, and their implementations are fully independent.

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. On the commercially available Neuropixels 1.0 device (J. J. Jun et al., 2017) 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. By contrast, the recently announced Neuropixels 2.0 (Steinmetz et al., 2021) has an electrode:wire ratio of 13.3. Another device, currently in engineering test, will have 4416 sites on a single 45 mm shank, with electrode:wire ratio of 11.5. For the Neuropixels technology,

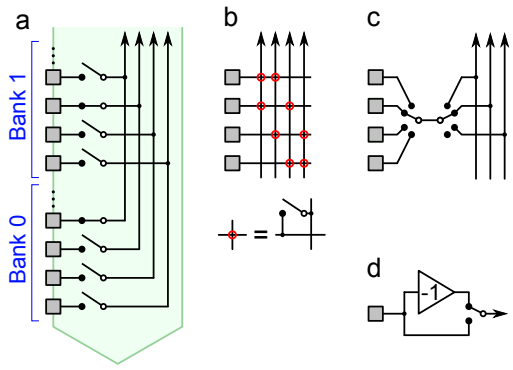

Figure 4.7: Hardware schemes for flexible connection between electrodes and wires. (a) In the current Neuropixels array each electrode can be connected to just one wire with a controllable switch. (b) Two switches per electrode would allow a choice of 2 wires, enabling many more pooling configurations. (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.

the number of sites can grow with shank count and shank length while channel count is limited by base area and trace crowding on the shank. These new probes already offer substantial opportunities to pool electrodes. Indeed, Steinmetz et al (Steinmetz et al., 2021) report an example of pooling two electrode banks, although their approach to unmixing the signals differs from that advocated here.

The design of effective electrode pools requires some flexibility in how recording sites are connected to wires. In the current Neuropixels technology each electrode has only one associated wire, which constrains the choice of electrode pools. The CMOS switch itself is small, but the local memory to store the switch state occupies some silicon space (Seidl, Herwik, et al., 2011). Nonetheless one can implement 3 switches per electrode even on a very tight pitch (Dragas et al., 2017). When arranged in a hierarchical network (Müller et al., 2015) these switches could effect a rich diversity of pooling schemes adapted to the specifics of any given experiment (Figure 4.7). 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 4.7c). This requires just 1 bit of storage per electrode, as in the current Neuropixels probe (J. J. Jun et al., 2017).

Another hardware design feature could greatly increase the capacity for electrode pooling: An optional analog inverter at each electrode (Figure 4.7d). This is a simple CMOS circuit that changes the sign of the waveform (Bae, 2019) 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.

Of course each of these proposals comes with some cost, such as greater power use or added space required for digital logic. The overall design of a probe must take all these trade-offs into account. The several-fold gain in recording efficiency promised by electrode pooling should act as a driver in favor of fully programmable switches, but deciding on the optimal design will benefit from close interaction between users and manufacturers.

Software

Electrode pooling will also benefit from further developments in spike-sorting algorithms. For example, 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 so-called "hot sorting" method in KiloSort2 and have shown that it can greatly increase the number of split-mode cells recovered in the pooled recordings (Figure 4.5c). This idea may also be extended to cluster-based sorting algorithms, by guiding the initialization of the clustering step. Indeed, knowing ahead of time 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.

As one envisions experiments with 10,000 or more recording sites, it becomes imperative to automate the optimal design of electrode pools, so that the user wastes no time before launching into pooled recording (Figure 4.3). The pooling strategy can be adapted flexibly to the statistics of the available neural signals, even varying along the silicon shank if it passes through different brain areas. The user always has the option of recording select sites in conventional mode; for example this might serve to sample local field potentials at a sparse set of locations. Designing an effective algorithm that recommends and implements the electrode switching based on user goals will be an interesting challenge.

High-impact applications

Finally, we believe that the flexible pooling strategy will be particularly attractive in chronic studies, where an electrode array remains implanted for months or years. In these situations, maintaining an updated library of signal waveforms is an intrinsic part of any recording strategy. Round-the-clock recording serves to populate and refine the library, enabling the design of precise spike templates, and effective separation of pooled signals. The library keeps updating in response to any slow changes in recording geometry that may take place.

A second important application for pooling arises in the context of sub-dural implants in humans. Here the sub-dural space forces a low-profile chip with minimal volume for electronic circuitry, whereas one can envision several slender penetrating electrode shafts with thousands of recording sites. We estimate that some devices that are now plausible (no published examples yet) will have an electrode-to-channel ratio near 25. Clearly one will want to record from more than 1/25 of the available sites, and electrode pooling achieves it without increased demand on electronic circuitry.

In summary, while the devices to maximize pooling benefits are not yet available, they soon may be. Consideration of pooling options would benefit the designers and users of these devices. The advantage of pooling grows naturally as the same tissue is recorded across sessions or time. The calculations and demonstrations reported here are intended to inspire professional simulations and design of future devices for a variety of applications, including human implants.

4.7 Methods

All analysis was performed with Matlab R2016b (Mathworks) and Python 3. All the quoted uncertainties are standard deviations.

Control of Neuropixels switching circuitry

The Neuropixels 1.0 probe 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. Electrode pooling was implemented by modifying the Neuropixels API and the GUI software SpikeGLX to allow connecting up to three electrodes to each readout wire.

Neuropixels device measurements

To characterize signal and noise pooling on the Neuropixels 1.0 array, we immersed the probe in a saline bath containing two annular electrodes to produce an electric field gradient (Fig 4.8a). The electrolyte was 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). We recorded from all 383 wires (recall that one wire is a reference electrode), first closing the switches in Bank 0 then in Bank 1, then in both banks (Figure 4.3b).

One set of measurements simply recorded the noise with no external field applied. Then we varied the concentrations of PBS (by factors 10^{-3}, 10^{-2}, 10^{-1}, 1, and 10), which modulated the conductance of the bath electrolyte in the same proportions. For each of the 15 recording conditions (5 concentrations x 3 switch settings) we measured the root-mean-square noise on each of the 383 wires. Then we set to explain these 5 x 3 x 383 noise values based on the input circuitry of the Neuropixels device. After some trial-and-error we settled on the equivalent circuit in Figure 4.8b. It embodies the following assumptions:

- Each electrode is a resistor R_i in series with a capacitor C_i. The resistor is entirely the bath resistance, so it scales inversely with the saline concentration.
- The shunt impedance Z_S across the amplifier input is a resistor R_S in parallel with a capacitor C_S.
- The thermal noise from this R-C network and the voltage noise N_{amp} from the amplifer and acquisition system sum in quadrature.

With these assumptions one can compute the total noise spectrum under each condition. In brief, each resistor in Figure 4.8b is modeled as a white-spectrum Johnson noise source in series with a noiseless resistor (Thevenin circuit). The various Johnson noise spectra are propagated through the impedance network to the output voltage U. That power spectrum is integrated over the AP band (300-10000 Hz) to obtain the total thermal noise. After adding the amplifier noise N_{amp} in quadrature one obtains the RMS noise at the output U. This quantity is plotted in the fits of Figure 4.8c.

The result is rather insensitive to the electrode capacitance C_i because that impedance is much lower than the shunt impedance Z_S. By contrast the bath resistance (R_0, R_1)

has a large effect because one can raise it arbitrarily by lowering the saline concentration. To set the capacitor values, we therefore used the information from the Neuropixels spec sheet that the total electrode impedance at 1 kHz is 150 kΩ,

$$C_i = \frac{1}{2\pi \cdot 1000\,\mathrm{Hz} \cdot \sqrt{(150\,\mathrm{k\Omega})^2 - {R_i}^2}} \tag{4.12}$$

We also found empirically that the shunt impedance is primarily capacitive: R_S is too large to be measured properly and we set it to infinity. Thus the circuit model has only 4 scalar parameters: R_0, R_1, C_S, N_{amp}. Their values were optimized numerically to fit all 15 measurements. This process was repeated for each of the 383 wires. The fits are uniformly good, see Figure 4.8c for examples.

As expected the thermal noise increases at low electrolyte concentration because the bath impedance increases (Figure 4.8c). However, the noise eventually saturates far below the level expected for the lowest saline concentration. This reveals the presence of another impedance in the circuit that acts as a shunt across the amplifier input (Fig 4.2a). We found that $Z_S \approx 20\,\mathrm{M\Omega}$ and is mostly capacitive. Because the shunt impedance far exceeds the electrode impedances ($\sim 150\,\mathrm{k\Omega}$ (J. J. Jun et al., 2017)), it has only a minor effect on signal pooling, which justifies the approximations made in Equation 4.3.

The measured noise voltage also saturates at high saline concentration (Fig 4.8c), and remains far above the level of Johnson noise expected from the bath impedance. That minimum noise level is virtually identical for the two electrodes that connect to the same wire, whether or not they are pooled, but it varies considerably across wires (Figure 4.8d). We conclude that this is the amplifier noise N_{amp} introduced by each wire's acquisition system (Figure 4.2a).

Figure 4.8e shows the best-fit values of the 4 circuit parameters, histogrammed across all the wires on an unused probe. Note they fall in a fairly narrow distribution. The bath impedance of the electrodes (in normal saline) is ~13 kΩ, the shunt capacitance is ~10 pF, and the common noise N_{amp} has a root-mean-square amplitude of ~6 μV integrated over the AP band (300-10000 Hz).

These measurements were performed on both fresh and used Neuropixels devices, with similar results. On a device previously used in brain recordings the bath impedance of the electrodes was somewhat higher: 30 kΩ instead of 13 kΩ).

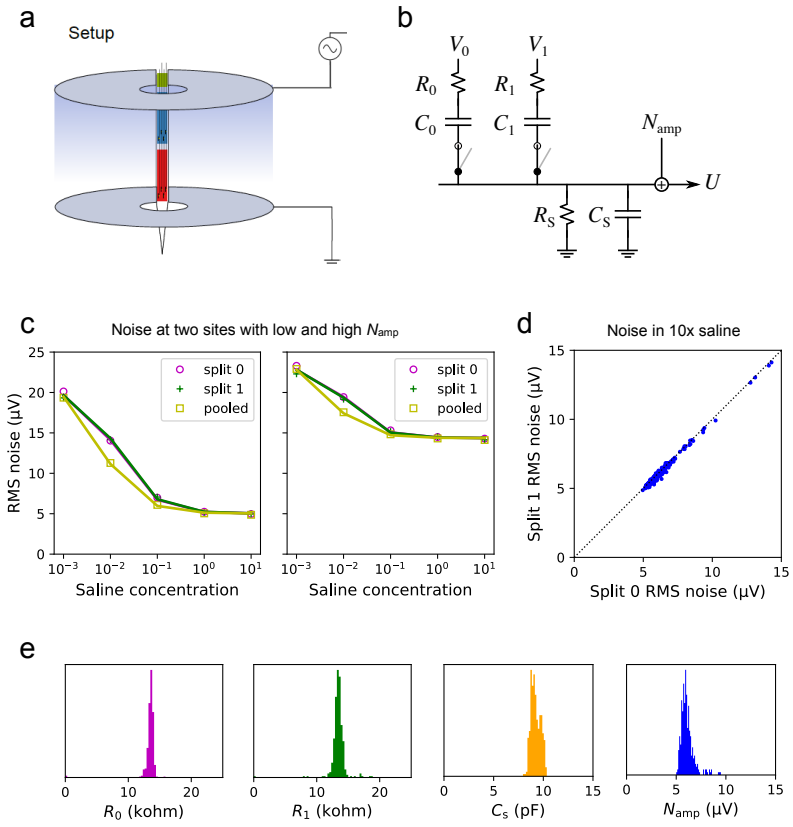

Figure 4.8: Methods for in vitro measurements of Neuropixels function. (a) The probe is immersed in saline, with two annular electrodes producing an electric gradient along the shank. (b) Equivalent circuit model to understand signal and noise pooling for one wire of the array. (c) Measurements of noise only without an external field. RMS noise as a function of the saline concentration under 3 conditions of the switches: split recording from Bank 0, split recording from Bank 1, and pooled recording from both. Examples of two different wires, one with high, the other with low amplifier noise N_{amp}. (d) The noise at the highest saline concentration, recording from electrode 1 vs electrode 0. Each dot is for one of the 383 wires. This limiting noise is identical across the two electrodes on the same wire. (e) Histograms of the best-fit circuit parameters derived for each of the 383 wires on a pristine Neuropixels probe. R_{S} is too large to be measured properly.

To measure the pooling coefficients we applied an oscillating electric field (1000 Hz) along the electrode array with a pair of annular electrodes (Figure 4.8a). From the recorded waveform we estimated the signal amplitude by the Fourier coefficient at the stimulus frequency. Two different field gradients (called A and B) yielded two sets of measurements, each in the two split modes ($U_{0,\mathrm{A}}, U_{1,\mathrm{A}}, U_{0,\mathrm{B}}, U_{1,\mathrm{B}}$) and the pooled mode ($U_{\mathrm{P,A}}, U_{\mathrm{P,B}}$). For each of the 383 wires we estimated the pooling coefficients of its two electrodes by solving

$$\begin{bmatrix} U_{0,\mathrm{A}} & U_{1,\mathrm{A}} \\ U_{0,\mathrm{B}} & U_{1,\mathrm{B}} \end{bmatrix} \begin{bmatrix} k_0 \\ k_1 \end{bmatrix} = \begin{bmatrix} U_{\mathrm{P,A}} \\ U_{\mathrm{P,B}} \end{bmatrix} \tag{4.13}$$

These mixing coefficients k_0 and k_1 express the recorded amplitude U_{P} in terms of the recorded amplitudes U_0 and U_1,

$$U_{\mathrm{P}} = k_0 U_0 + k_1 U_1 \tag{4.14}$$

whereas the pooling coefficients c_0 and c_1 (Eq 4.2) are defined relative to the input voltages V_0 and V_1, namely

$$U_{\mathrm{P}} = c_0 V_0 + c_1 V_1 \tag{4.15}$$

The U_i differ from the V_i only by the ratio of electrode impedance to shunt impedance. Given the measured value of $Z_{\mathrm{S}} \approx 20\,\mathrm{M\Omega}$ that ratio is less than 1%, a negligible discrepancy. So the measured k_0 and k_1 are excellent approximations to the pooling coefficients c_0 and c_1, which in turn reflect the ratio of the two electrode impedances (Eq 4.2).

In vivo recording

We used a Neuropixels 1.0 probe (J. J. Jun et al., 2017) to record neural signals from a head-fixed mouse (C57BL/6J, male, 9 months old). The probe entered the brain at 400 µm from midline and 3.7 mm posterior from bregma at ~45° 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 (Lee et al., 2020). All procedures were in accordance with institutional guidelines and approved by the Caltech IACUC, protocol 1656.

Once the probe was implanted, data were recorded 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.

Following brain recordings the array was cleaned according to recommended protocol by immersion in tergazyme solution and rinsing with water.

Spike-sorting

For "manual" spike-sorting of the in vivo recordings we used KiloSort1 (downloaded from `https://github.com/cortex-lab/KiloSort` on Apr 10, 2018). We ran the automatic template-matching step; the detailed settings are available in the code accompanying this manuscript. This was followed by manual curation, merging units and identifying those of high quality. These manual judgements were based on requiring a plausible spike waveform with a footprint over neighboring electrodes, a stable spike amplitude, and a clean refractory period. This was done separately for each of the three recordings (split-mode Bank 0, split-mode Bank 1, pooled-mode).

We implemented the “hot sorting” feature in KiloSort2 (downloaded from `https://github.com/MouseLand/Kilosort2` on Mar 19, 2020). No manual curation was used in this mode, because (1) we wanted to generate a reproducible outcome, and (2) manual inspection is out of the question for the high-volume recordings where electrode pooling will be applied. We first sorted the two split-mode recordings and used their templates to initialize the fields `W` and `U` of `rez2` before running the main template-matching function on the pooled recording (see the accompanying code for more details). Finally, the splits, merges, and amplitude cutoffs in Kilosort2 ensured that the final output contained as many well-isolated units as possible. We then selected cells designated as high quality (`KSLabel` of `Good`) by KiloSort2, indicating putative, well-isolated single neurons (Stringer et al., 2019).

To elaborate on the internal operations of Kilosort2: Spike-sorted units were first checked for potential merges with all other units that had similar multi-channel waveforms (waveform correlation >0.5). If the cross-correlograms had a large dip (<0.5 of the stationary value of the cross-correlogram) in the range [-1 ms, +1 ms], then the units were merged. At the end of this process, units with at least 300 spikes were checked for refractory periods in their auto-correlograms, which is a measure

of contamination with spikes from other neurons. The contamination index was defined as the fraction of refractory period violations relative to the stationary value of the auto-correlogram. The default threshold in Kilosort2 of 10 percent maximum contamination was used to determine good, well-isolated units.

Following spike sorting, we applied the matching algorithm based on cosine similarity (Figure 4.5b) to determine how many cells identified in split recordings could be recovered from the pooled recording. This was compared with the results from "cold sorting," in which the pooled recording was sorted on its own, as well as to the conventional sorting that includes manual curation (Figure 4.5c).

Unmixing pooled signals

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 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 20 channels above and 20 channels below the channel with the largest amplitude (i'), i.e. $J = [i' - 20, i' + 20]$. 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):

$$\Sigma(S, P) = \frac{\mathbf{S} \cdot \mathbf{P}}{\|\mathbf{S}\|\|\mathbf{P}\|} \tag{4.16}$$

where $\mathbf{S}$ and $\mathbf{P}$ are column vectors obtained by concatenating every S_j and P_j ($j \in J$), 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. This process gets iterated 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.

Estimating pooling coefficients in vivo

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

method (`mldivide` in Matlab).

$$\arg\min_{k_i} \; \|P_i - k_i S_i\| \tag{4.17}$$

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 each carry significant signal, it enables the estimation of both of their pooling coefficients. Examples of such sites are shown in Figures 4.5d-e (up to 50 pairs in Banks 0 and 1).

Simulation

Generating simulated data

We simulated extracellular voltage signals on 12 groups of 4 local electrodes ("tetrodes"). Each time series was sampled at 30,000 samples/s and extended over 600 s. After combining signal and noise as described below, the time series was filtered with a passband of 300-5000 Hz.

Each tetrode carried spikes from a single unit. The spike waveform of the unit was chosen from an actual mouse brain Neuropixel recording, with a different waveform on each tetrode. Within a tetrode, one electrode chosen at random carried this spike at the nominal peak-to-peak amplitude, V (Figure 4.6b). On the other 3 electrodes the spike was scaled down by random factors drawn from a uniform distribution over [0,1]. The spike train was simulated as a Poisson process with a forced 2-ms refractory period, having average firing rate r (Figure 4.6b).

Three sources of noise - biological noise N_{bio}, thermal electrode noise N_{the}, and common amplifier noise N_{com} - were generated as gaussian processes. The quoted noise values (Figure 4.6b) refer to root-mean-square amplitude over the 300-5000 Hz passband. Thermal noise was sampled independently for each electrode, but the biological noise was identical for electrodes within a tetrode, given that they likely observe the same background activity.

Electrode pooling across M tetrodes was implemented by combining the voltage signals of the corresponding electrode on each tetrode, resulting in signals on 4

wires. In the process each electrode signal was weighted by $1/M$, then the amplifier noise was added to the resulting average. Amplifier noise was sampled separately for each wire.

Tetrodes were added to the pool in a sequence determined by the spike shape of their units. We started with the two most dissimilar units as determined by the cosine similarity of their spike waveforms. Then we progressively added the unit that had the lowest similarity with those already in the pool.

Sorting simulated data

The simulated 4-wire time series was sorted using Kilosort2; detailed configuration settings are available in the code accompanying this paper. We found it necessary to turn off the "median voltage subtraction" during preprocessing, because that feature somehow introduced artifacts in the 4 voltage traces. This did not occur when processing electrode array data with many channels, for which the algorithm is intended. We note that an effective means of subtracting the common signal across wires may help suppress the biological noise and lead to better sorting results.

When large numbers of tetrodes were pooled the signal-to-noise ratio dropped to the point where KiloSort2 could not form templates in the preprocessing step. Under those conditions we report zero units recovered (Figure 4.6b).

Scoring simulated data

Following previous reports (Barnett, J. F. Magland, and Greengard, 2016; J. Magland et al., 2020), the spike times of the sorted units and the ground truth units were matched and compared using the confusion matrix algorithm from Barnett, J. F. Magland, and Greengard, 2016. We set the acceptable time error between sorted spikes and ground-truth spikes at 0.1 ms. Then we counted the number of spike pairs with matching spike times, n_{match}, the number of unmatched spikes in the ground-truth unit, n_{miss}, and the number of unmatched false-positive spikes in the sorted unit, n_{fp}.

To assess the quality of the match between ground-truth and sorted units we adopted the *Accuracy* definition in (J. Magland et al., 2020):

$$\text{Accuracy} = \frac{n_{\text{match}}}{n_{\text{match}} + n_{\text{miss}} + n_{\text{fp}}} \tag{4.18}$$

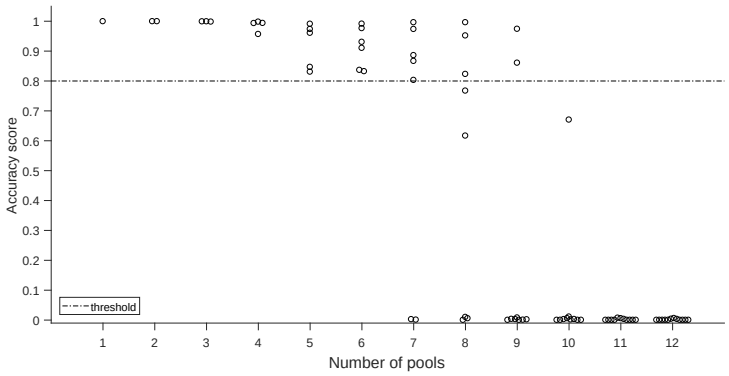

Figure 4.9: Accuracy scores of one "standard" condition simulation, Related to Fig 4.6. Units with accuracy score > 0.8 were counted as recovered.

Figure 4.9 shows the accuracy distribution obtained for various degrees of pooling. Sorted units with accuracy > 0.8 were counted as "recovered" from the pooled signal. For each parameter set we ran the simulation 3 times, randomizing the noise and the spike times. Results from the 3 runs are reported by mean ± SD (Figure 4.6b).

Data availability

All data relevant to the reported results are available in a public repository: https://github.com/markusmeister/Electrode-Pooling-Data-and-Code. An archived version is available from CaltechDATA: http://dx.doi.org/10.22002/D1.2032.

Code availability

All code used to obtain the reported results are available in a public repository: https://github.com/markusmeister/Electrode-Pooling-Data-and-Code. An archived version is available from CaltechDATA: http://dx.doi.org/10.22002/D1.2032.

4.8 Acknowledgements

This work was supported by a grant to MM from NINDS (5R01NS111477) and an award to MM from the Tianqiao and Chrissy Chen Institute for Neuroscience. YLN was supported by the Taipei Veterans General Hospital – National Yang-Ming University Physician Scientists Cultivation Program, No.103-Y-A-003.

References

Attwell, D. and S. B Laughlin (Oct. 2001). "An Energy Budget for Signaling in the Grey Matter of the Brain." In: *J Cereb Blood Flow Metab* 21, pp. 1133–45.

Bae, Woorham (Aug. 2019). "CMOS Inverter as Analog Circuit: An Overview". In: *Journal of Low Power Electronics and Applications* 9, p. 26. DOI: 10.3390/jlpea9030026.

Barnett, Alex H., Jeremy F. Magland, and Leslie F. Greengard (2016). "Validation of neural spike sorting algorithms without ground-truth information". In: *Journal of Neuroscience Methods* 264, pp. 65–77. ISSN: 0165-0270. DOI: https://doi.org/10.1016/j.jneumeth.2016.02.022. URL: http://www.sciencedirect.com/science/article/pii/S0165027016300036.

BRAIN Working Group (2014). *BRAIN 2025: A Scientific Vision*. en. Tech. rep., p. 146. URL: https://braininitiative.nih.gov/strategic-planning/brain-2025-report (visited on 11/04/2019).

Dimitriadis, George et al. (Mar. 2018). "Why not record from every channel with a CMOS scanning probe?" en. In: *bioRxiv*. DOI: 10.1101/275818. URL: http://biorxiv.org/lookup/doi/10.1101/275818 (visited on 08/28/2019).

Dragas, J. et al. (June 2017). "In Vitro Multi-Functional Microelectrode Array Featuring 59 760 Electrodes, 2048 Electrophysiology Channels, Stimulation, Impedance Measurement, and Neurotransmitter Detection Channels". In: *IEEE Journal of Solid-State Circuits* 52.6, pp. 1576–1590. ISSN: 0018-9200. DOI: 10.1109/JSSC.2017.2686580.

Eversmann, B. et al. (Dec. 2003). "A 128 x 128 cmos biosensor array for extracellular recording of neural activity". en. In: *IEEE Journal of Solid-State Circuits* 38.12, pp. 2306–2317. ISSN: 0018-9200. DOI: 10.1109/JSSC.2003.819174. URL: http://ieeexplore.ieee.org/document/1253878/ (visited on 08/27/2019).

Harris, K. D. et al. (July 2000). "Accuracy of Tetrode Spike Separation as Determined by Simultaneous Intracellular and Extracellular Measurements." In: *J Neurophysiol* 84, pp. 401–14.

Jun, James et al. (2017). "Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction". In: *bioRxiv*, pp. 1–29. DOI: https://doi.org/10.1101/101030.

Jun, James J. et al. (2017). "Fully integrated silicon probes for high-density recording of neural activity". In: *Nature* 551.7679, pp. 232–236. ISSN: 14764687. DOI: 10.1038/nature24636. URL: http://dx.doi.org/10.1038/nature24636.

Kleinfeld, David et al. (Sept. 2019). "Can One Concurrently Record Electrical Spikes from Every Neuron in a Mammalian Brain?" en. In: *Neuron* 103.6, pp. 1005–1015. ISSN: 08966273. DOI: 10.1016/j.neuron.2019.08.011. URL: https://

linkinghub.elsevier.com/retrieve/pii/S0896627319306956 (visited on 11/05/2019).

Kozai, Takashi D. Y. and Alberto L. Vazquez (June 2015). "Photoelectric Artefact from Optogenetics and Imaging on Microelectrodes and Bioelectronics: New Challenges and Opportunities". en. In: *Journal of Materials Chemistry B* 3.25, pp. 4965–4978. ISSN: 2050-7518. DOI: 10.1039/C5TB00108K.

Lee, Kyu Hyun et al. (Apr. 2020). "The sifting of visual information in the superior colliculus". eng. In: *eLife* 9. ISSN: 2050-084X. DOI: 10.7554/eLife.50678.

Lewicki, Michael S. (1998). "A review of methods for spike sorting: the detection and classification of neural action potentials." In: *Network* 9 4, R53–78.

Linderman, Michael D. et al. (Jan. 2008). "Signal Processing Challenges for Neural Prostheses". English. In: *Ieee Signal Processing Magazine* 25.1, pp. 18–28. ISSN: 1053-5888. DOI: 10.1109/MSP.2008.4408439.

Lopez, Carolina Mora et al. (2017). "A Neural Probe With Up to 966 Electrodes and Up to 384 Configurable Channels in 0.13 µm SOI CMOS". In: *IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS* 11.3, pp. 510–522.

Magland, Jeremy et al. (May 2020). "SpikeForest, Reproducible Web-Facing Ground-Truth Validation of Automated Neural Spike Sorters". In: *eLife* 9. Ed. by Markus Meister, Ronald L Calabrese, and Markus Meister, e55167. ISSN: 2050-084X. DOI: 10.7554/eLife.55167.

Müller, Jan et al. (2015). "High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels". en. In: *Lab on a Chip* 15.13, pp. 2767–2780. ISSN: 1473-0197, 1473-0189. DOI: 10.1039/C5LC00133A. URL: http://xlink.rsc.org/?DOI=C5LC00133A (visited on 11/05/2019).

Obien, Marie Engelene J. et al. (Jan. 2015). "Revealing neuronal function through microelectrode array recordings". en. In: *Frontiers in Neuroscience* 8. ISSN: 1662-453X. DOI: 10.3389/fnins.2014.00423. URL: http://journal.frontiersin.org/article/10.3389/fnins.2014.00423/abstract (visited on 11/05/2019).

Olsson, R. H. and K. D. Wise (Dec. 2005). "A Three-Dimensional Neural Recording Microsystem with Implantable Data Compression Circuitry". In: *IEEE Journal of Solid-State Circuits* 40.12, pp. 2796–2804. ISSN: 1558-173X. DOI: 10.1109/JSSC.2005.858479.

Pachitariu, Marius et al. (2016). "Fast and accurate spike sorting of high-channel count probes with KiloSort". In: *Advances in Neural Information Processing Systems* 29.Nips, pp. 4448–4456. ISSN: 10495258. URL: https://papers.nips.cc/paper/6326-fast-and-accurate-spike-sorting-of-high-channel-count-probes-with-kilosort.pdf.

Raducanu, Bogdan C. et al. (Oct. 2017). "Time Multiplexed Active Neural Probe with 1356 Parallel Recording Sites". en. In: *Sensors* 17.10, p. 2388. DOI: 10.3390/s17102388.

Rios, Gustavo et al. (Nov. 2016). "Nanofabricated Neural Probes for Dense 3-D Recordings of Brain Activity". en. In: *Nano Letters* 16.11, pp. 6857–6862. ISSN: 1530-6984, 1530-6992. DOI: 10.1021/acs.nanolett.6b02673. URL: https://pubs.acs.org/doi/10.1021/acs.nanolett.6b02673 (visited on 08/28/2019).

Robinson, D.A. (1968). "The electrical properties of metal microelectrodes". en. In: *Proceedings of the IEEE* 56.6, pp. 1065–1071. ISSN: 0018-9219. DOI: 10.1109/PROC.1968.6458. URL: http://ieeexplore.ieee.org/document/1448388/ (visited on 11/04/2019).

Schomburg, Erik W. et al. (Aug. 2012). "The Spiking Component of Oscillatory Extracellular Potentials in the Rat Hippocampus". en. In: *Journal of Neuroscience* 32.34, pp. 11798–11811. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.0656-12.2012.

Seidl, Karsten, Stanislav Herwik, et al. (2011). "CMOS-Based High-Density Silicon Microprobe Arrays for Electronic Depth Control in Intracortical Neural Recording". In: *Journal of Microelectromechanical Systems* 20, pp. 1439–1448.

Seidl, Karsten, Michael Schwaerzle, et al. (Dec. 2012). "CMOS-Based High-Density Silicon Microprobe Arrays for Electronic Depth Control in Intracortical Neural Recording–Characterization and Application". en. In: *Journal of Microelectromechanical Systems* 21.6, pp. 1426–1435. ISSN: 1057-7157, 1941-0158. DOI: 10.1109/JMEMS.2012.2206564. URL: http://ieeexplore.ieee.org/document/6249712/ (visited on 11/05/2019).

Shahrokhi, Farzaneh et al. (June 2010). "The 128-Channel Fully Differential Digital Integrated Neural Recording and Stimulation Interface". en. In: *IEEE Transactions on Biomedical Circuits and Systems* 4.3, pp. 149–161. ISSN: 1932-4545, 1940-9990. DOI: 10.1109/TBCAS.2010.2041350. URL: http://ieeexplore.ieee.org/document/5471738/ (visited on 11/05/2019).

Shannon, C. E. (Jan. 1949). "Communication in the Presence of Noise". In: *Proceedings of the IRE* 37.1, pp. 10–21. ISSN: 0096-8390. DOI: 10.1109/JRPROC.1949.232969.

Siegle, Joshua H. et al. (2019). "A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas". In: *bioRxiv*. DOI: 10.1101/805010. eprint: https://www.biorxiv.org/content/early/2019/10/16/805010.full.pdf. URL: https://www.biorxiv.org/content/early/2019/10/16/805010.

Steinmetz, Nicholas A. et al. (Apr. 2021). "Neuropixels 2.0: A Miniaturized High-Density Probe for Stable, Long-Term Brain Recordings". en. In: *Science* 372.6539. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.abf4588.

Stevenson, Ian (Oct. 2013). *Tracking Advances in Neural Recording | Statistical Neuroscience Lab*. en-US. https://stevenson.lab.uconn.edu/scaling/.

Stringer, Carsen et al. (2019). "Spontaneous behaviors drive multidimensional, brainwide activity". In: *Science* 364.6437. ISSN: 0036-8075. DOI: 10.1126/science.aav7893. eprint: https://science.sciencemag.org/content/364/6437/eaav7893.full.pdf. URL: https://science.sciencemag.org/content/364/6437/eaav7893.

Suo, Yuanming et al. (Oct. 2014). "Energy-Efficient Multi-Mode Compressed Sensing System for Implantable Neural Recordings". English. In: *Ieee Transactions on Biomedical Circuits and Systems* 8.5, pp. 648–659. ISSN: 1932-4545. DOI: 10.1109/TBCAS.2014.2359180.

Torfs, Tom et al. (Nov. 2010). "Two-dimensional multi-channel neural probes with electronic depth control". en. In: *2010 Biomedical Circuits and Systems Conference (BioCAS)*. Paphos, Cyprus: IEEE, pp. 198–201. ISBN: 978-1-4244-7269-7. DOI: 10.1109/BIOCAS.2010.5709605. URL: http://ieeexplore.ieee.org/document/5709605/ (visited on 08/28/2019).

Yger, Pierre et al. (Mar. 2018). "A Spike Sorting Toolbox for up to Thousands of Electrodes Validated with Ground Truth Recordings in Vitro and in Vivo". eng. In: *eLife* 7. ISSN: 2050-084X. DOI: 10.7554/eLife.34518.

Chapter 5

CONCLUSION & FUTURE REMARKS

Here I will briefly discuss about the implications of the projects and future directions of them.

Well-studied circuits as a testbed

We demonstrated that by measuring interactions of drugs with a well-studied neural circuit, one could immediately gain insights and interpret the actual targets.

Following this direction, the neuropharmacology field should aim for an overarching goal of seeking more circuit-based platforms to complement molecular and behavioral studies. Given our own finding that ketamine seems to deliver its action through its HCN effect rather than the deep-rooted NMDAr block, this wave of new testbeds should first also aim to reexamine another neuropharmacology use cases that are currently still in an empirical phase that does not have a good mechanism but somehow works. For instance, movement disorders medications.

In addition to exploring new testbeds and new applications, I think it is a riped time to study the rapid anti-depression effect of ketamine, with the premise that people have to know what they are optimizing for. As mentioned in Chap 2, the depression research field took the default status of NMDAr *IS* as its mechanism of anesthesia and is diving down this rabbit hole without further verifying of the actual mechanism. One should take a step back to the drawing board and verify the actual molecular switch of paradoxical firing in their platforms.

On the other branch of application, we were showing that the non-specific, concentration-based explanation of anesthetics still made sense. Nevertheless, it is also unrealistic to dis-prove each of the specific hypothesis arguments in a whack-a-mole fashion despite it being the current mainstream theory. Therefore, I think it is fair to verify the extreme cases one at a time. Such as the fluorocarbons experiments which are borderline keeping the tested animals alive while using behavioral output as end-points. Instead, a retina on an MEA should function fine in a hyperbaric chamber without concern for animal vitality.

Electrode Pooling

We believe that the flexible pooling strategy will be particularly attractive in chronic studies, where an electrode array remains implanted for months or years. In these situations, maintaining an updated library of signal waveforms is an intrinsic part of any recording strategy. The round-the-clock recording serves to populate and refine the library, enabling the design of precise spike templates, and effective separation of pooled signals. The library keeps updating in response to any slow changes in recording geometry that may take place.

A second important application for pooling arises in the context of sub-dural implants in humans. Here the sub-dural space forces a low-profile chip with minimal volume for electronic circuitry, whereas one can envision several slender penetrating electrode shafts with thousands of recording sites. We estimate that some devices that are now plausible (no published examples yet) will have an electrode-to-channel ratio near 25. Clearly one will want to record from more than 1/25 of the available sites, and electrode pooling achieves it without increased demand on electronic circuitry.

In summary, while the devices to maximize pooling benefits are not yet available, they soon may be. Consideration of pooling options would benefit the designers and users of these devices. The advantage of pooling grows naturally as the same tissue is recorded across sessions or time. The calculations and demonstrations reported here are intended to inspire professional simulations and design of future devices for a variety of applications, including human implants.