5-14-2024

Stimulation of Peripheral Nerves in M. Mulatta for Somatosensory Psychophysics and VNS Applications

Tyler Schlichenmeyer
Washington University – McKelvey School of Engineering

WASHINGTON UNIVERSITY IN ST. LOUIS

McKelvey School of Engineering
Department of Biomedical Engineering

Dissertation Examination Committee:
Daniel W Moran, Chair
Dennis Barbour
Harold Burton
Ilya Monosov
Wilson Z Ray

Stimulation of Peripheral Nerves in M. Mulatta for Somatosensory Psychophysics and VNS Applications

by

Tyler Schlichenmeyer

A dissertation presented to
the McKelvey School of Engineering
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

May 2024
St. Louis, Missouri

Table of Contents

List of Figures . . . . . **v**

List of Tables . . . . . **vii**

Acknowledgments . . . . . **viii**

Abstract . . . . . **x**

Chapter 1: Introduction **1**
1.1 An Outline of This Thesis 4

Chapter 2: Background **5**
2.1 Neural Interfaces for Rehabilitation 5
2.1.1 Neural Plasticity After Injury 6
2.1.2 Long-Term Variation in Neural Interface Characteristics 7
2.1.3 Neural Adaptation During Tool Usage 8
2.1.4 The Macaque Model 8
2.1.5 The Sensorimotor Feedback Loop 9
2.1.6 Vagus Nerve Stimulation (VNS) 9
2.2 Peripheral Nerve Interface (PNI) Design 10
2.2.1 Advanced Stimulation Strategies 10
2.2.2 Limits of PNI Control 12
2.2.3 Assessing Interface Designs and Stimulation Strategies 12
2.3 Psychophysics 14
2.3.1 Foundations - Fechner and Weber 14
2.3.2 The Curse of Dimensionality and Non-Stationarity 16
2.3.3 Bayesian Inferencing 17

Chapter 3: Detection and Discrimination of Electrical Stimuli from an Upper Limb Cuff Electrode in M. Mulatta **19**
3.1 Introduction 19
3.2 Materials and Methods 20
3.2.1 Ethical Treatment of Animals 20
3.2.2 Surgical Procedure 20

3.2.3 Implant Characteristics . . . 22
3.2.4 Stimulation . . . 22
3.2.5 Psychophysical Procedures . . . 24
3.2.6 Analysis . . . 27
3.3 Results . . . 28
3.3.1 Detection Thresholds . . . 29
3.3.2 Discrimination . . . 34
3.4 Discussion . . . 34
3.4.1 Stability Assessment . . . 37
3.4.2 Strength-Duration Findings . . . 38
3.4.3 Weber Analysis . . . 38
3.4.4 Limitations . . . 39
3.4.5 Future Directions . . . 41
3.4.6 Conclusion . . . 42

Chapter 4: PsychoAnalyze: An Open-Source Toolset for Psychophysics Analysis . . . 43
4.1 Introduction . . . 43
4.1.1 Comparable Work . . . 44
4.2 Core Features . . . 45
4.2.1 Dashboard . . . 46
4.2.2 Python Package . . . 49
4.2.3 Notebooks . . . 49
4.3 Development Environment . . . 50
4.4 Future Direction and Roadmap . . . 53

Chapter 5: Vagus Nerve Stimulation, Fear Extinction, and Post-Traumatic Stress Disorder . . . 54
5.1 Introduction . . . 54
5.1.1 Vagus Nerve Stimulation (VNS) . . . 54
5.1.2 Parameterization of the Electrical Stimulus Delivered to the Vagus Nerve . . . 56
5.1.3 The Fear Extinction Model . . . 56
5.2 Methods . . . 58
5.2.1 Ethical Treatment of Animals . . . 58
5.2.2 Stimulus Conditioning . . . 59
5.2.3 Surgical Procedure . . . 60
5.2.4 Delivering VNS During Learning and Extinction . . . 61
5.2.5 Analysis . . . 62
5.3 Results . . . 62
5.4 Discussion . . . 64
5.4.1 Retrospective . . . 64
5.4.2 Status of VNS as a Clinical Intervention for PTSD . . . 64

Chapter 6: Discussion: Towards a Performant, Integrated Model of Peripheral Nerve Encoding . . . **69**
6.1 Introduction . . . 69
6.1.1 Modern Data Practices and the Bayesian Paradigm . . . 70
6.1.2 The Hybrid Model of the Peripheral Nerve and its Limitations . . . 71
6.1.3 Finite Element Modeling . . . 71
6.1.4 Neuron Compartment Model . . . 73
6.2 Bayesian Methods and their Applications in Neural Engineering . . . 73
6.2.1 Psychometric Curve Fitting . . . 73
6.2.2 "Smart" Bayesian Sampling of the Parameter Domain . . . 75
6.2.3 The Open Road Ahead: Novel Applications of Bayesian Methods in Neuroscience . . . 76
6.3 Software Tooling . . . 78
6.3.1 All-in-One/Integrative Software . . . 79
6.3.2 Probabilistic Computing . . . 80
6.3.3 Data Visualization . . . 81
6.3.4 Psychophysical Modeling . . . 83
6.3.5 Data Manipulation . . . 83
6.3.6 Neuron Simulation . . . 85
6.3.7 Finite Element Modeling Software (and Related Tooling) . . . 86
6.3.8 Honorable Mentions to other Programming Languages . . . 87
6.4 A Bayesian Workflow Proposal to Guide Next Steps . . . 88
6.4.1 Choosing an Initial Model . . . 88
6.4.2 Scaling and Transforming the Parameters . . . 88
6.4.3 Prior Predictive Checking . . . 89
6.4.4 Running the Model . . . 89
6.4.5 Trying Other Stimulus Dimensions . . . 90
6.4.6 Tuning the Model for a Single Session . . . 90
6.4.7 Introducing Hierarchy . . . 91
6.5 Conclusion . . . 91

References . . . **93**

List of Figures

Figure 1.1: PNIs in closed-loop feedback control 3

Figure 2.1: Spatial resolution vs invasivity 11

Figure 3.1: Electrode geometry and impedance data. 21

Figure 3.2: A typical pulse train delivered by the cuff nerve interface. 23

Figure 3.3: Yes/No joystick task 25

Figure 3.4: Summary of fitted parameters - Detection experiments 30

Figure 3.5: Amplitude detection thresholds 31

Figure 3.6: Strength-duration characteristics of behavioral thresholds 33

Figure 3.7: Discrimination and Weber coefficients 36

Figure 4.1: A screenshot of the *PsychoAnalyze* dashboard. 47

Figure 4.2: Screenshot of PsychoAnalyze documentation. 52

Figure 5.1: Schematic views of cranial nerves and the vagus nerve 55

Figure 5.2: Conditioning Task 59

Figure 5.3: Matching task 60

Figure 5.4: Photographs of the vagus nerve before and after implant placement. . . 61

Figure 5.5: Eye Tracker Recordings 63

Figure 5.6: Distributions of Conditioning Metrics across trials 65

Figure 5.7: Pupillometry Data 66
Figure 5.8: Behavioral Measurements from an Appetitive/Aversive Task 67

List of Tables

Table 3.1: Experimental sessions performed by each monkey 29

Table 3.2: Strength-Duration Data 32

Table 3.3: Weber's Law estimates 35

Acknowledgments

Portions of Chapters 2 and 3 are recreations of previously published work in the *Journal of Neural Engineering* [87]. A special thank you to Dr. Harold Burton, who assisted with many revisions of the manuscript and provided valuable feedback and guidance throughout the project.

Portions of Chapter 4 are recreations of work submitted for publication in the *Journal of Open Source Software*.

Portions of this work were sponsored by the Defense Advanced Research Projects Agency (DARPA) Biological Technologies Office (BTO) Hand Proprioception and Touch Interfaces (HAPTIX) program under the auspices of Dr. Doug Weber through the DARPA Contracts Management Office Cooperative Agreement No. HR0011-15-2-0007.

This work was also supported by funding from the *Cognitive, Computational, and Systems Neuroscience* (CCSN) training grant (NIH-T32) from the CCSN Pathway at Washington University in St. Louis.

My colleague Kara Donovan contributed significantly to the VNS data analyses presented in Chapter 5.

Myelinated nerve SVG art in Figure 1.1 was sourced from *Reactome*[32] and brain cross section in Figure 2.1 is from DBCLS[1]. Fig 5.1 of thoracic features from *Jkwchu* [2]. These images and other clip art images are available for use under the Creative Commons License[3].

Tyler Schlichenmeyer

Washington University in St. Louis
May 2024

[1] https://togotv.dbcls.jp/en/pics.html - Frontal_plane_of_the_brain_amygdala
[2] https://creativecommons.org/licenses/by-sa/3.0, via Wikimedia Commons
[3] https://creativecommons.org/licenses/by/4.0/

A very special thanks to:

My parents for their unconditional love.

The best vet tech in the world, Donna Reedy, for her mentorship and companionship.

Maintainers of open source software everywhere.

Cody "Codes" Murray and the "Couple More Crew" for the background soundtrack to this disseration and a community that made long nights in the COVID era just a bit less isolating.

My advisor, Dr. Dan Moran, and my collaborators and mentors, Dr. Harold Burton and Dr. Erik Zellmer, whose collective kindness, brilliance, and unwavering support kept me going through this journey.

ABSTRACT OF THE DISSERTATION

Stimulation of Peripheral Nerves in M. Mulatta for Somatosensory Psychophysics and VNS Applications

by

Tyler Schlichenmeyer

Doctor of Philosophy in Biomedical Engineering

Washington University in St. Louis, 2024

Professor Daniel W Moran, Chair

Cuff-style peripheral nerve interfaces (PNIs) communicate with the central nervous system and offer interesting advantages when compared with other neural interfaces, including brain-computer interfaces (BCIs) and non-cuff PNI designs. PNI development must consider design factors such as volume of neural activity being driven or recorded, specificity of stimulus activation, and longevity of the device. This dissertation provides an analytical framework for PNI experiments, showcasing two experiments where PNIs were implanted in highly trained macaque monkeys. We applied foundational psychophysical procedures and other behavioral paradigms to examine the relationship between the input electrical stimulus delivered by a cuff electrode and the resulting behavioral response during a classical conditioning task. We studied somatosensory encoding in three monkeys in a longitudinal study over several months. Additionally, we studied vagus nerve stimulation (VNS) in a single monkey in a study to examine the effects of VNS stimulation during trace extinction, since the facilitation of fear extinction might improve outcomes for VNS treatment given to patients with post-traumatic stress disorder (PTSD). To support data analysis for these experiments, I developed the open source Python package *PsychoAnalyze*, which powers an interactive dashboard for navigating psychophysical datasets. As researchers integrate modern data engineering approaches, tools, and workflows with our existing protocols and

models, there is significant opportunity for researchers to build increasingly sophisticated computational models that improve VNS outcomes and our understanding of the peripheral nervous system.

Chapter 1

Introduction

The peripheral nerve interface (PNI) lies at a curious intersection of neuroscience and engineering. For context, the PNI may be considered alongside its well-recognized functional counterpart, the brain-computer interface (BCI) (Figure 1.1). Both technologies have been developed in an effort to relieve individuals experiencing various degrees of neurological damage and affliction. The applications of neural interfaces are not restricted to sensorimotor pathologies such as amputation and spinal cord injury: neuromodulatory devices have shown to be effective for afflictions of higher-order neural processes such as Parkinson's, epilepsy, depression, PTSD, and more. Furthermore, the promise of a closed-loop sensorimotor prosthetic limb device capable of volitional control and haptic feedback has attracted public curiosity[4] and significant funding efforts from veterans' organizations and federal programs [84]. Still, challenges remain as engineers seek to develop neural interfaces that are performant, accessible, and cost-effective enough to sustain widespread clinical use and obtain acceptable rates of adoption among amputees and other patients [13]. What should our expectations be for the performance ceiling of these devices? How can we leverage engineering principles to guide us through growing complexity as we seek to understand the neural systems involved and "close the feedback loop?" This dissertation aims to address these questions and more by blending simple and traditional experimental frameworks with cutting-edge data architecture and tooling.

In many practical applications, including clinical contexts, PNIs can supplement or replace functionality provided by BCIs. While both device categories are generally responsible for obtaining input signals (via recording) and delivering output signals (via stimulation) to the neural system, with control processing usually mediated via a computational algorithm, their respective positions in the neural pathway make comparing them somewhat "apples-to-oranges." Understanding the ways that they complement each other is important when

[4] https://en.wikipedia.org/wiki/Prosthetics_in_fiction

determining experimental design or clinical aims. In comparison to BCIs, PNIs demonstrate improved device longevity and signal stability while maintaining clinically-relevant performance. BCIs often sacrifice this durability for more direct access to the signals from the brain, where, depending on device architecture, they might either record stochastic ensemble signals from a subset of the billions of neurons and quadrillions of synapses in the brain, or obtain high-resolution single unit recordings from intracortical microelectrode arrays. In contrast, PNIs record signals from a subset of only hundreds to thousands of electrically aligned and topographically organized neurons.

A performant neural interface would feature a closed-loop, bidirectional mechanism which executes both the decoding of afferent signals and encoding of efferent signals. The reduced complexity of PNI systems suggests they may provide a tractable model for some clinical applications and research questions in this regard. Additionally, the exceptional adaptive capabilities of the nervous system often dovetail nicely with modern machine-learning approaches which, among other adaptive procedures, are able to "intelligently" sample and stimulate sensorimotor fields. When these developments are integrated in tandem, it becomes feasible to achieve device adaptability and generalizability that promotes successful outcomes across a wide range of users. While peripheral nerve anatomy and neural pathways vary considerably from subject to subject, a certain degree of functional topographic organization of the nerve relaxes the selectivity requirements of the interface. If researchers seek to replicate natural-feeling sensorimotor function via computationally-managed control signals, the properties of PNI systems make them a worthy area of our engineering focus, providing an opportunity to develop sophisticated models with a growing degree of accuracy.

In addition to clinical benefits, the reduced complexity of the PNI system makes it a useful tool for neuroscientific inquiry in general. To assess quality and progress in our device designs, we must undertand the limits of our biology, our materials, and more generally, the physical laws that govern our universe[5]. How can we utilize engineering principles and the scentific method to contextualize and quantify the generally *subjective* experience of sensation, to assess the performance of neural interfaces and the design trade-offs that must be made during their development? How can our models of neural systems inform our treatment and control strategies? In so many words, we must rigorously examine and understand the fundamental models of the sensory experience. To begin to answer these questions, we might turn to the rich history in the field of *psychophysics*–a subfield of psychology and neuroscience

[5] https://xkcd.com/435/

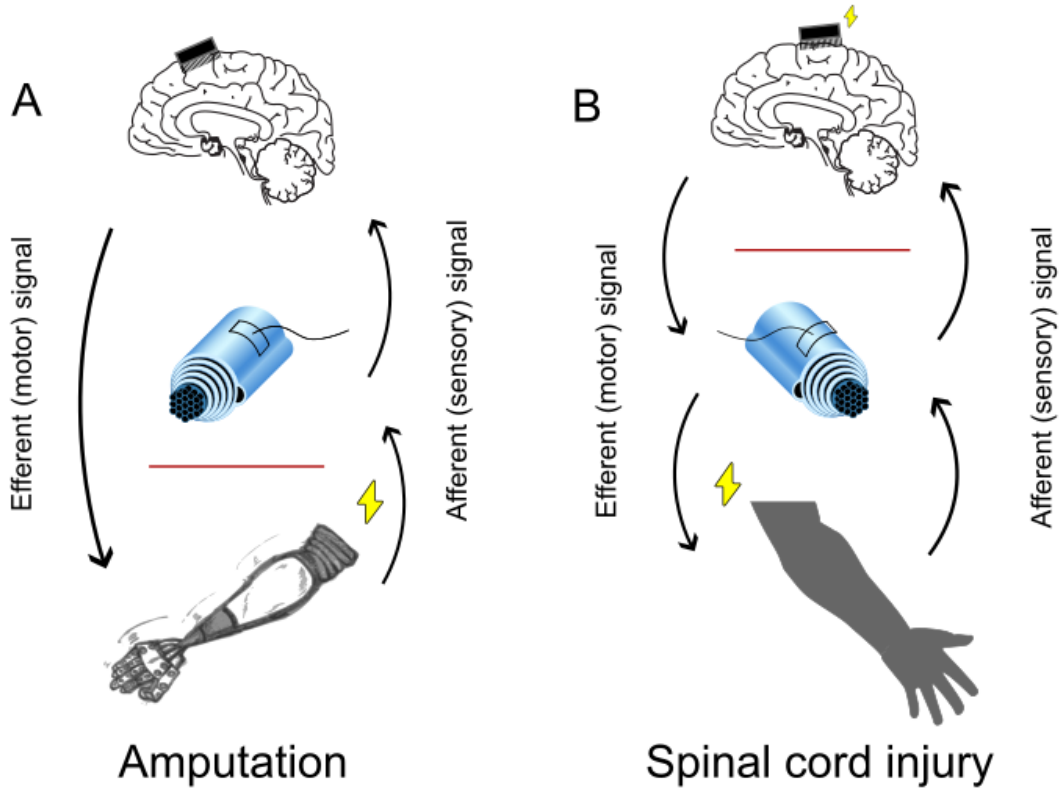

Figure 1.1: *PNIs in closed-loop feedback control.* Simplified examples of closed-loop feedback control using peripheral nerve interfaces (PNIs) and brain-computer interfaces (BCIs) in instances of amputation (A) and spinal cord injury (B). Red lines convey site of injury. **A.** In a simple example of a prosthetic limb device for amputees, a BCI implant sends motor signals to an actuator on the prosthetic limb device, while pressure sensors send sensory signals to intact peripheral nerve proximal to injury. **B.** In cases of spinal cord injury, The brain sends volitional motor signals to an PNI implant, which stimulates the peripheral nerve to drive muscle fiber activation. Sensory signals are recorded from the PNI device and are sent to a BCI device which stimulates the sensory cortex.

devoted to the art and science of measuring sensation–where surprisingly simple principles have led to outsizingly impactful discoveries.

1.1 An Outline of This Thesis

In *Chapter 2*, I review foundational research related to the design of peripheral nerve interfaces, including an introduction to *psychophysics*, historical challenges in PNI design, and the current state of the art of PNI development.

In *Chapter 3*, I describe a long-scale study of somatosensory encoding via a PNI in three monkeys over several months.

In *Chapter 4*, I describe the development of the open source Python package *PsychoAnalyze*, a software suite capable of automating data analyses for psychophysics experiments, designed with customizability and extensibility in mind.

In *Chapter 5*, I describe a brief pilot study assessing the effects of vagus nerve stimulation (VNS) on trace conditioning and extinction in a single monkey.

In Chapter 6, I conclude my arguments with an exposition on modern data engineering approaches, toolsets, and workflows, notably Bayesian workflows, to demonstrate how we may leverage modern computational advancements to facilitate exploration of new frontiers in neuroengineering.

In summary, this dissertation describes preliminary results and methodological findings from exploratory behavioral experiments involving non-human primates fitted with peripheral nerve interfaces. Particularly, we provide insight on the challenges researchers may face during peripheral nerve interface development. In addition to reporting raw measures from model parameter estimation for the NHP models studied, I argue that combining traditional psychophysical methods and models with modern analysis techniques would facilitate improvements in device design and experimental protocols.

Chapter 2

Background

2.1 Neural Interfaces for Rehabilitation

Neural interfaces, by the nature of the electrochemical processes that enable their function as a neuroscientific tool, are capable of both stimulating the nervous system and recording from it. Many high-profile breakthroughs in neural interface development were developments in the area of motor BCIs; the convenient location of the M1 atop the central sulcus and the functionally topographic arrangement of the cells in response to motor signals made it an appealing area for development. However, neural interfaces are functionally limited without robust sensory feedback mechanisms in place. The development of sensory interfaces has introduced unique challenges related to the decoding/encoding of neural signals related to sensation. Regardless, progress in the field has brought us closer to "closing the feedback loop" – an important, albeit nebulous, milestone wherein somatosensory feedback is strong enough, localized enough, and fast enough to provide the entire nervous system with the adequate level of control it needs to perform everyday tasks.

PNIs and BCIs have been most notably applied to the restoration of sensorimotor function after injury such as amputation or spinal cord injury [98, 102, 99, 34, 18, 109], but applications for these interfaces extend to other sensory domains such as with cochlear implants for hearing loss [114] and retinal implants for vision loss [63]. Additionally, we have seen remarkable treatments in the form of dynamic neuromodulation devices such as in *deep brain stimulation* for the treatment of Parkinson's; meanwhile, PNIs have shown a similar ability to treat "higher-order" ailments (e.g. mood disorders) via *vagus nerve stimulation* (VNS) [86, 69].

2.1.1 Neural Plasticity After Injury

To understand how engineered devices can replace lost neural function due to injury or disease, we should first understand how the nervous system responds to such disruptions in cases of natural healing and existing treatments. The nervous system adapts to injury through so-called *plasticity*, the ability of the nervous system to change its structure and function in response to changes in the environment [27]. Structural changes occur in the morphology of synapses and axons, which manifest new neural pathways in response to patterned input from the subject's environment. These changes are how the nervous system adapts to change, including injury. These morphological changes are largely responsible for adaptive processes in the nervous system, including classical conditioning and other forms of learning, skill acquisition such as with the usage of tools, and their and subsequent bodily "incorporation" [56]. These new pathways are designed to encode patterns in the sensory experience to help us navigate the world. As the adage goes, neurons that fire together, wire together (i.e. Hebbian learning).

Previous studies have shown clear evidence of neural plasticity in somatosensory systems in response to peripheral nerve damage [52, 78]. *Merzenich et al.* transected the median nerve in several non-human primates and measured the effects of nerve transection via cortical recordings. They demonstrated that cortical representations of the hand rearrange themselves in response to the lack of neural input after nerve transection [60] and amputation [61]. When nerves were blocked from regenerating, neighboring cortical representations of the hand expanded into the cortical area that was previously occupied by the deafferented area. When the nerves were allowed to regenerate, the cortical representations of the hand returned to their original locations, albeit in a slightly abnormal manner [107]. For example, 65% of responsive neurons were responsive to multiple receptive fields, and 25% of the remaining responsive neurons had receptive fields scattered across the cortical map. Neurons in both of these groups featured receptive fields smaller than pre-injury receptive fields. An additional 4% of responsive neurons in Area 3b had sensitive, "pacinian-like" behavior with large receptive fields, behavior not observed in normal macaques in this area. Even still, some receptive field areas did indeed end up mirroring their pre-injury counterparts. Areas 3b and 1 in the same animals showed different reorganization patterns indicating cortical mechanisms independent of the state of the periphery. Interestingly, nerves recovering from crush injury do not seem to recover as abnormally as what is seen in cases of complete nerve transection [106].

In summary, these results indicate that that there is a wide degree of reorganization happening on relatively short time scales in response to perturbations of the peripheral nervous system, and while neural recovery mechanisms are sometimes capable of returning to pre-injury state, the degree to which recovery is possible may be limited by neurobiological circumstances of the healing process. When modeling these processes, changes happening in the cortex need to be reconciled with changes simultaneously happening in the periphery. For example, we might ask to what degree these abnormal cortical patterns are a result of shuffling the connection mappings between mending peripheral fibers, as opposed to adaptation happening solely in the central nervous system. These differences in recovery mechanisms may be further examined and differentiated by inspecting the course of changes in observed qualities over time, not only in the observation of receptive fields via neural recordings, but through observation of *behavior* as well.

2.1.2 Long-Term Variation in Neural Interface Characteristics

A 1990 study measured long-term fluctuations in the ability of monkeys to detect a sound stimulus during usage of a cochlear implant after induced deafness, as measured by a *detection threshold* [77]. These fluctuations came in a variety of patterns that likely convey similar adaptation mechanisms as seen in peripheral nerve transection. In the discussion of these results, Pfingst hypothesizes that the observed changes were not due to changes in tissue morphology near the electrode, supported by a lack of frequency-dependent fluctuations and correlations between threshold and impedance measurements. Additionally, sharp decreases in threshold observed in the initial stages of observations are likely not due simply to learning of the electric stimulus, since we know that changes related to plasticity usually happen on much longer time scales. In many ways, the experiment described in Chapter 3 of this dissertation provides a corrolary to these experiments, applying related lines of questioning to the somatosensory system.

The existence of neural reorganization patterns months and years after injury and device implantation demonstrate a clear clinical need for calibration mechanisms available to patients who expect to use neuroassistive devices on a long-term basis. A deeper understanding of neural mechanisms of recovery and plasticity will help researchers developing interfaces that do not require frequent manual calibration from clinicians or technicians.

2.1.3 Neural Adaptation During Tool Usage

Maravita & Iriki [56] demonstrate how our neural system adapts to tools outside of its preconceived notion of physical self, examining receptive fields during tool use in non-human primates, normally-functioning humans, and humans with brain damage. We can refer to this study and similar studies to conceptualize how the nervous system adapts to the use of a prosthetic device and how interfacing with peripheral nerves can be used to facilitate this adaptation. Remapping in response to targeted peripheral nerve simulation results in downstream remapping in the cortex. This remapping is consistent with the remapping observed in response to natural sensory input.

The ability of our nervous system to adapt to injury, and its highly adaptible nature in general, adds an interesting and encouraging characteristic of open-endedness when it comes PNI design and engineering. By understanding the ways that the neural system adapts to change, we can better emulate or replace these mechanisms using the signals recorded and emitted from our devices to induce the creation of pathways that close the sensimotor feedback loop. However, research indicates we do not have to perfectly model or mimic naturalistic mechanisms in order to create functional devices, especially when adaptation from the nervous system can "meet us halfway" (or whichever percentage of the way that biology dictates). Rouse et al. [85] demonstrated that a monkey could gain 2D control of a cursor by learning to modulate neural signals from a relatively "naive" choice of electrode placement above the cortex. Additionally we might experiment with "virtual" anatomies and non-naturalistic prosthetic designs which might adapt more fittingly to the PNI system.

2.1.4 The Macaque Model

Non-human primates (NHPs) are a common model organism for neural interface research. NHPs are used in research for a variety of reasons, including their close evolutionary relationship to humans, their similar nervous system and upper limb/hand anatomy, and their ability to be trained to perform complex and repetitive tasks. In the sensory domain, the macaque model offers a convenient anatomical advantage; in humans, many important areas of primary somatosensory cortex (Areas 2 and 3) are obscured by their location within the sulci of the brain. However, macaque brains do not feature similar sulci and their primary somatosensory cortex is more easily accessed for recording and stimulation. One characteristic

of macaque studies, especially in recent decades, is that of "low n": the number of subjects in a study is often by necessity lower than what is expected for studies with adequate statistical power to generalize results within that species. This must be taken into consideration during experimental design, where principles that can be deduced "within-subject," such as temporal developments, are of particular importance. Despite "low n" when measured by number of subjects, macaques can complete thousands of trials per experimental session over many consecutive days, providing high-frequency feedback that is often prohibitively difficult and expensive to obtain with human studies. In summary, while macaque experiments are generally costly, they provide key advantages over both human and non-primate studies [73].

2.1.5 The Sensorimotor Feedback Loop

The sensorimotor feedback loop is the mechanism by which the nervous system continuously receives sensory input from the environment and uses it to update its internal model of the world, which is then used to generate motor commands and behavior in general. Sensory feedback is an important component of both simple and complex behavior, voluntary and involuntary. The circuits involving this control loop are extremely complex both anatomically and in the context of engineering control systems [101]. Importantly, adequate sensory feedback has been demonstrated to be a key indicator of successful prosthetic limb use [13].

2.1.6 Vagus Nerve Stimulation (VNS)

The vast majority of peripheral nerves in the human body mediate sensorimotor processes and are connected to the brain via the spinal cord. However, there are a specialized subset of twelve nerves that interface directly with the brain, which are designated as *cranial nerves*. The nerve designated as the tenth cranial nerve (cranial nerve X) is the *vagus* nerve, which mediates several connections along the throat, heart and stomach. The nerve consists of about 80% afferent fibers and 20% efferent fibers that are said to convey various sensory information about visceral organs [6].

Within the past few decades, researchers have discovered that electrical stimulation of the vagus nerve, delivered in a similar manner as with spinal nerve stimulation, can alleviate

neurological disorders such as epilepsy, major depression, and other anxiety disorders [25, 76]. Further background relevant to our studies are discussed in Chapter 5.

2.2 Peripheral Nerve Interface (PNI) Design

There are clear engineering tradeoffs when it comes to peripheral nerve interface design. By examining popular interface designs in the literature, we can better understand the tradeoffs and design decisions that must be made when designing a peripheral nerve interface and optimizing for performance.

PNIs encounter many of the same high-level performance tradeoffs as seen in BCIs (Figure 2.1). BCIs can measural neuronal activity in the cortex at a high spatial resolution, but they require invasive surgical procedures. The most invasive BCI electrodes, intracortical arrays, penetrate the cortex and induce inflammatory responses that limit the lifetime of the device [115]. Other BCI designs such as electrocorticography (ECoG) implants and electroencephalography (EEG) implants are less invasive, but they record with lower spatial resolution as the signal degrades across the dura and the skull. PNIs face similar considerations regarding penetration of the epineurium: implants that penetrate the outer fascicular layer of the nerve obtain higher spatial resolution at the cost of durability (e.g. TIME implants), while implants that circumvent the epineurium tend to achieve lower spatial resolution but are more durable (e.g. cuff electrodes).

Helpfully, axon bundles in human peripheral nerve, especially in the upper limbs, are *topographically* organized; bundles of nerve fibers are spatially organized into functionally similar and anatomically distinct regions called *fascicles* (See outlines of subregions in peripheral nerve schematic drawing in Figure 2.1). Thus, selectivity at the individual neuron level is likely not necessary for functional feedback from the device. It is unclear from our current models what degree of selectivity is sufficient to achieve naturalistic sensations.

2.2.1 Advanced Stimulation Strategies

One approach employed to achieve higher spatiotemporal selectivity with cuff electrodes is via the use of multi-electrode implants to shape the electric field across the nerve. The timing and

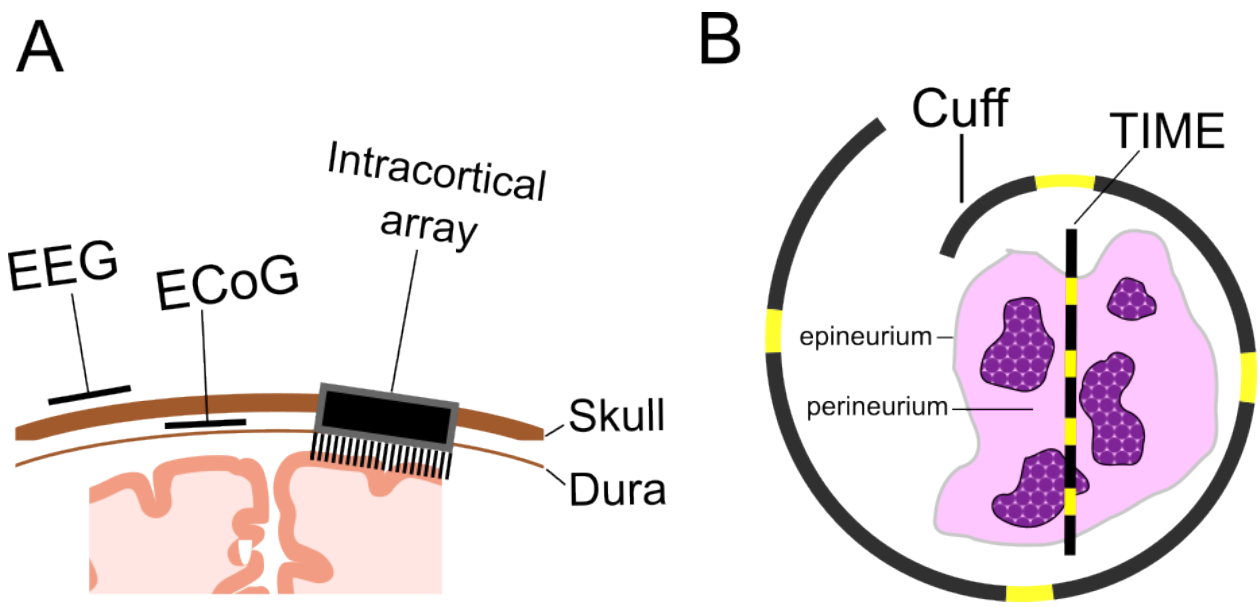

Figure 2.1: **Relative invasivity of neural interface designs.A.** Three examples of brain-computer interface (BCI) designs that demonstrate the tradeoffs between device invasivity and signal quality and/or stability. Electroencephalographic (EEG) signals may be acquired noninvasively, but signal quality is diminished through the skull. Electrocoricographic (ECoG) signals may sit above the dura or be placed below the dura at a higher risk of inflammation response. Intracortical arrays penetrate neural tissue but can record single-unit activity at hundreds of electrode sites. **B.** Similar design trade-offs occur in peripheral nerve devices. Cross-sections of two representative electrode designs are shown here. In a spiral cuff design (outer trace, yellow areas indicate electrode sites), one or multiple electrodes are placed in silicone tubing that circumvents the outer layer of the peripheral nerve (epineurium). Conversely, an *intrafascicular* design such as the Transverse Intrafascicular Multichannel Electrode (TIME), penetrate the perineurium but provide a higher degree of spatial selectivity.

amplitude of each electrode contact can be adjusted to achieve a desired spatial distribution of the current field, inducing pockets of higher and lower current densities delivered to different functional areas [104, 14]. This technique is referred to as *multipolar current steering* or *current focusing*. Multipolar stimulation has been shown to be effective in achieving a wide range of spatial distributions of current fields [5].

While neural interface stimulation is stereotypically parameterized as a charge-balanced biphasic square pulse, other unique stimulation strategies have been deployed with interesting results, including periodic envelopes [99].

2.2.2 Limits of PNI Control

Some limitations of PNI performance are consequences of electrode engineering and design, as discussed in the previous section. The accuracy and precision of the functionality delivered by the devices are limited by the spatiotemporal selectivity provided by the electrode or implant, whose design is constrained by the physics and biology of the electrode-tissue interface. However, there are further limitations imposed by our lack of understanding of the neural code. Even presupposing a magic device that offers high spatiotemporal selectivity *and* longitudinal stability, the question remains if we have a detailed enough understanding of the sensory code such that we could develop a "training program" that could *efficiently* and *sufficiently* train the central nervous system to interpret and embody the appropriate signals delivered by the interface. With this in mind, how should we try to model the input-output relationship of information transfer happening between the components of this system? How can we characterize downstream perceptual effects of modulating input stimulation patterns, particularly in a high-dimensional input space where practical concerns and limitations of experimentation quickly obfuscate the path forward?

2.2.3 Assessing Interface Designs and Stimulation Strategies

Comprehensive reviews and comparisons of various PNI and BCI designs are plentiful in the literature. A couple example designs are provided here to illustrate particular design trade-offs.

Consider the *current focusing* technique introduced earlier. The physical mechanisms of this technique can be easily demonstrated using simple electrostatic modeling; the gradient of the electric field is made larger while the cross-sectional area of activation is made smaller. While we can accurately characterize the manner in which parameterization of the device modulates this phenomenon at the electric output of the interface, it is more difficult to characterize how the output of the device affects the downstream perceptual experience. In the case of current focusing, we can imagine that sensation is some multiplicative combination of the intensity and spatial spread of activation, and a "current-focused" stimulus would create higher volage gradients over smaller spatial area. However, *in vivo* experimentation is necessary to model a relationship between these variables and the resulting sensation, since sensation is a subjective experience.

Consider another set of common design choices: the size, number, and spacing of the electrode contacts in the device. The size of the electrode contact is a trade-off between the spatial resolution of the device and the current density at the interface. Smaller electrodes provide higher spatial resolution, but they are not able to activate as many fibers, and also increase the current density at the interface, which can lead to tissue damage and electrode corrosion. Additional electrode contacts may allow for higher spread of neural activation, but come with additional engineering constraints such as algorithmic performance and power draw. Spacing electrodes too close together may lead to cross-talk between electrodes, but spacing them too far apart may result in less functional control of topographically organized fibers. In experimental settings, these choices are often hastily made, governed by intuition and precedent.

However, a principled approach to experimental design can help us better understand the downstream effects of these design choices. Firstly, we should be interrogating computational models to inform aspects of experimental design such that we can narrow down our input parameter space and thus obtain more relevant feedback earlier in the experiment. By supplementing *in silico* models with real, induced behavioral measurements, we can hone in on mechanistic hypotheses instead of wandering aimlessly in the 6th-dimensional plane. To address these challenges we should apply appropriate statistical methods to our experimental design; in Chapter 6 I examine how Bayesian models and methods are powerful tools in this regard.

Ultimately, we have the opportunity to build more powerful computational models that can build upon the elegant simplicity of a centuries-old, tried-and-true, first-principles approach to the measurement of sensation, a domain of psychology and neuroscience referred to as *psychophysics*.

2.3 Psychophysics

Psychophysics generally refers to a foundational set of experimental methods and theoretical models developed as part of an ongoing effort to accurately and meaningfully quantify the subjective sensory experience. In essence, in order to quantify the entire realm of sensory experience on an interpretable scale, we must meticulously manipulate the conditions under which we observe a subject's response to modulations in sensory input. For example, a cornerstone of psychophysical research is the "two-alternative" design, wherein experimenters do not concern themselves with the complexities of measuring sensations in how they relate to 2D or 3D space, but instead reduce the subjective response of the experimental subject to a binary outcome (e.g. "yes" or "no", "left" or "right").

Helpfully, the psychophysical framework allows us to attempt to measure and quantify the sensory experience both within and between sensory modalities, e.g. vision, audition, and somatosensation. Mathematically, it attempts to probabilistically model the mechanism of response from our sensory systems by formulating a relationship between the evoked response and the form and intensity of the sensation-evoking stimulus.

While simple experiments and models have provided foundational explanations of our sensory systems, challenges inherent in the data structure of more complex experiments have provided roadblocks in the quest for a comprehensive model, such as one that might be part of a sensorimotor prosthetic feedback system that can realistically recreate the sensorimotor experience of a lost limb.

2.3.1 Foundations - Fechner and Weber

In the late 19th and early 20th century, a paradigm shift around some of these questions evolved out of the seminal work of Gustav Fechner and Ernst Weber. Weber's Law, first

proposed in 1834[38], is a simple model that describes the relationship between the intensity of a stimulus and the intensity of the sensation it evokes. In essense, the "law" states that the just-noticeable difference (JND) in stimulus intensity is proportional to the original stimulus intensity.

$$\frac{\Delta I}{I} = k \tag{2.1}$$

where delta I is the just-noticeable difference (JND) in stimulus intensity, I is the original stimulus intensity, and k is a constant.

The Weber-Fechner Law has stood the test of time as an illustrative, if incomplete, example of psychophysical principles. It comfortably balances accuracy with utility, as many simple models do. For example, it is easy to conceptualize the classic example demonstrated by Weber that a person might feel a 2% difference in weight; if the 2% ratio is accurate, they would notice if 0.2 pounds were added to the weight of carrying 10 pounds, would notice if 2 pounds were added to the weight of 100 pounds, etc. It's also fairly straightforward to surmise the limitations of this model at its extremes.

In the history of psychophysics, the experimental methods that developed over time have been as impactful as the models that were derived from them. Weber began to articulate a formal relationship between stimulus and sensation in parallel to thoughtfully and methodically developing the experimental protocols necessary to properly evaluate such proposals. At a high level, experiments quickly revealed the predictive value of the model, while paricularities of the model were refined over time by his student Gustav Fechner and again later by SS Stevens[96]. For example, while the system clearly modeled some exponential relationship, the formulation of this relationship had limitations at edge cases; SS Stevens improved the precision of the model by identifying the relationship governing the "Weber Fraction" constant as "power law" relationship.

Despite, or perhaps because of, the simplicity of the framework, the application of psychophysics directly led researchers to groudbreaking validations of many different hypotheses in sensory neuropsychology. While still subject to all of the higher-order and potentially non-linear neural processing that occurs between sensation and behavioral response, the paradigm is adept at unearthing patterns in neural systems that require no invasive procedure or neural recordings of any kind. For example, in the visual system, psychophysics

has predicted properties of the physiology of the system well before they were verified with histology or cortical recordings. By testing subjects' ability to detect a visual stimulus at varying durations of time spent in the dark, Hecht et al. provided convincing evidence of two separate types of receptor cells in the eye: rods and cones [105]. The trichromatic theory of vision was similarly discovered in 1802 [117]! In the somatosensory domain, psychophysical experiments allowed us to deduce, very early on, that there are multiple types of mechanoreceptors in the skin that respond differently to different types of stimuli [46]. From Weber-Fechner to the present-day, countless applications of these basic psychophysical principles have amassed a wide body of evidence applicable to various neural models [82].

Psychophysics conveniently also provides us with a framework to translate certain behavioral lines of questioning to animal models. For example, while it's naturally difficult to infer qualitative dimensions of sensation such as *wetness* or *dryness* from animals, adhering to the psychophysical framework enables researchers to make reasonable inductions about subjective experiences by compiling responses to systematically administered yes/no or same/different tasks.

2.3.2 The Curse of Dimensionality and Non-Stationarity

The high-impact psychophysical studies described in the previous sections were executed in tightly controlled lab conditions. However, as researchers seek to model more complex processes, we might need to assess if the psychophysical paradigm is sufficient to model sensory processes at higher levels of complexity. Although the aim sould be to model a variety of stimulation strategies that might represent real-world processes, the complexity of the data collected from more complex experiments can quickly become overwhelming, for both the computational algorithms involved and the humans trying organize and process the data in order to run such analyses. The number of experimental conditions and stimulus parameterizations which can be manipulated in a psychophysical experiment (herefore referred to collectively as *dimensions*) can quickly grow to the point where the number of samples needed to provide valuable information is much greater than the number of trial samples that can be obtained in a reasonable amount of time. This mathematical conundrum is known as the *curse of dimensionality*, and cannot be overcome with naive experimental design. In this way, a robust computational model built on foundational hypotheses may help us navigate the multidimensional space.

The essence of the *bias-variance tradeoff*, colloquially referenced in terms of overfitting and underfitting, is the tendency of more complex models to find an "optimum" solution in terms of describing the training data at the cost of its generalizable predictive power. Simpler (*think: linear*) models tend to generalize fairly well in machine learning for a wide variety of problems. We clearly observe nonlinear relationships at various stages of neural processing, but perhaps by interfacing with a manageable number of organized neurons on one end, and by restricting our behavioral measurements to binary outcomes on the other, we can reduce the system to one that is linear in performance, if not in reality.

One other experimental condition that has limited the application of the psychophysical approach to more complicated experiments, alluded to previously but here made explicit, is that of *non-stationarity* in longitudinal experiments. *Non-stationarity* is a term in statistics and probability that simply conveys change in the nature of a probabilistic distribution over time. In the context of psychophysical experiments, we might observe (or infer) changes in the distribution of an estimated measure such as *threshold* over time, and the corresponding violation of assumptions of *stationarity* make it that much more difficult to isolate neural mechanisms and biological properties during our deduction process [21]. Additionally relevant to longitudinal psychophysical studies is the idea of a time-balanced study; in order to calculate simple correlations (e.g. Pearson's) in observations between study subjects, for example, it is very helpful mathematically to have made the observations on the same time scale. For example, if we take psychophysical measurements on days 1, 5, and 10 for one subject, our correlations will be more accurate if we take them on days 1, 5, and 10 in another subject. For many obvious reasons, this is difficult to coordinate as the number of observations and the length of the study increases, although there are various, more complex approaches to try to circumvent this limitation [83]. Developing an analytical framework that takes these time-related considerations into account is crucial to the future application of these methods to longitudinal experiments, where we may want to ask more complex questions about mechanisms of learning and adaptation.[31].

2.3.3 Bayesian Inferencing

The results of psychophysical studies are surprisingly deterministic when there is a high level of precise control over laboratory conditions, resulting in curve fits with extremely small residual errors. However, in a majority of cases where these careful conditions are not applied

but we still would like to obtain some knowledge about the system, there are often many potential sources of variance that may contribute to the overall observed variance. In this context, Bayesian statistics is an appropriate framework to not only provide deeper insight into the various sources of uncertainty but to utilize this information to make more robust predictors. Bayesian methods have become the predominant framework for curve-fitting for the most predominant software libraries designed to assist with this task, and more advanced application of Bayesian methods such as the hierarchical variety are beginning to be applied to analyses as well [79, 62, 40, 54, 39]. In Chapter 6, I will clearly demonstrate the utility of Bayesian methods for psychophysics and neural engineering and its direct application to state-of-the-art models. First, however, the following chapter will outline a comprehensive, multi-year study monitoring PNI usage in non-human primates.

Chapter 3

Detection and Discrimination of Electrical Stimuli from an Upper Limb Cuff Electrode in M. Mulatta

3.1 Introduction

Peripheral nerve interfaces seek to restore nervous system function through electrical stimulation of peripheral nerves. In clinical use, these devices should function reliably for years or decades. In this study, we assessed evoked sensations from multi-channel cuff electrode stimulation in macaque monkeys up to 711 days post-implantation.

Three trained macaque monkeys received multi-channel cuff electrode implants at the median or ulnar nerves in the upper arm. Electrical stimuli from the cuff interfaces evoked sensations, which we measured via standard psychophysical tasks. We adjusted pulse amplitude or pulse width for each block with various electrode channel configurations to examine the effects of stimulus parameterization on sensation. We measured detection thresholds and just-noticeable differences (JNDs) at irregular, near-daily intervals for several months using Bayesian inferencing from trial data. We examined data trends using classical models such as Weber's Law and the strength-duration relationship using linear regression.

Detection thresholds were similar between blocks with pulse width modulation and blocks with pulse amplitude modulation when represented as charge per pulse, the product of the amplitude and the pulse width. Conversely, Weber fractions—calculated as the slope of the regression between JND charge values and reference stimulus charge—were significantly different between pulse width and pulse amplitude modulation blocks for the discrimination task.

Weber fractions were lower in blocks with amplitude modulation than in blocks with pulse width modulation, suggesting that pulse amplitude modulation allows finer resolution of sensory encoding above threshold. Consequently, amplitude modulation may enable a greater dynamic range for sensory perception with neuroprosthetic devices.

3.2 Materials and Methods

We directly stimulated nerves using multi-channel cuff-style electrodes in previously trained macaques. Cuff electrodes durably surround fascicles located proximal to an injury or transection and tend to cause minimal nerve damage compared to intrafascicular electrodes [30]. Standard psychophysical measures assessed how stimulus parameterization affected behavioral outcomes. We adjusted independent variables of pulse amplitude, pulse width, and electrode selection in pursuit of estimates of detection thresholds and just-noticeable differences (JNDs) at near-daily intervals across extended study on a clinically relevant time scale. We evaluated longitudinal patterns in the data and examined effects of stimulus parameterization based on established sensory models such as Weber's Law [28] and the strength-duration relationship [8].

3.2.1 Ethical Treatment of Animals

The Institutional Animal Care and Use Committee approved the experimental paradigm design, surgical procedures, neurophysiological stimulations, and daily animal care, following all guidelines set by the Association for Assessment and Accreditation of Laboratory Animal Care and the Society for Neuroscience. We sedated animals with ketamine, intubated, and anesthetized with isoflurane before and during surgery under the supervision of Department of Comparative Medicine veterinary staff at Washington University in St Louis.

3.2.2 Surgical Procedure

Implant cuffs occupied a different upper arm location for each of three rhesus macaques. Monkey Y had a 2 mm electrode implanted at the ulnar nerve (Figure 3.1C). Monkeys

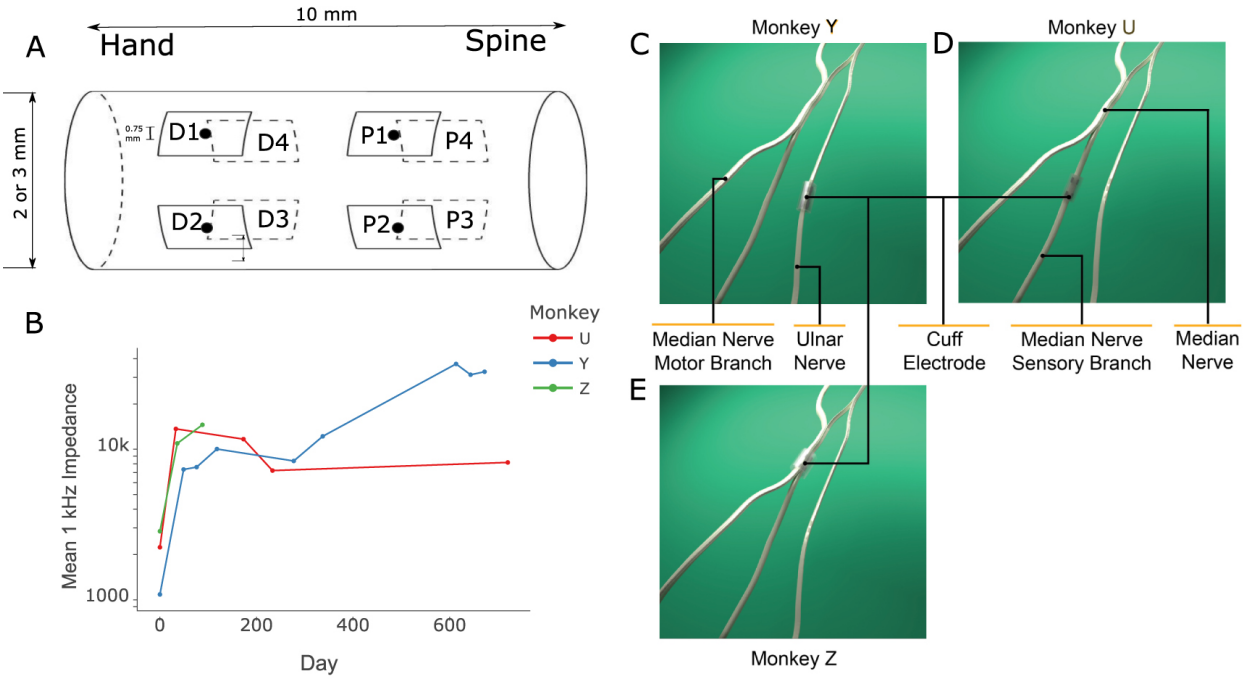

Figure 3.1: *Electrode geometry and impedance data.* **A.** Spiral cuff implants had four proximal (P1-4) electrodes and four distal (D1-D4) electrodes evenly spaced on a silicon conduit. Dimensions not to scale. **B.** Impedance measurements at 1 kHz before and during experiments. Measurements were from each electrode channel and summarized by the mean of impedances from all channels from a given session. Impedances remained generally stable through the course of experiments with some increase in Monkey Y's implant towards the end of experiments. **C.** The implant in Monkey Y surrounded the ulnar nerve. **D.** The implant in Monkey U surrounded a surgically isolated sensory branch of the median nerve at a proximal location. **E.** The implant in Monkey Z surrounded the median nerve.

U and Z had 3 mm electrodes implanted at the median nerve (Figures 3.1D and 3.1E). During the surgical procedure in Monkey U, we blunt dissected the sensory branch of the median nerve away from the nerve trunk, in the proximal portion of the upper arm where topographical organization of the nerve enabled such separation. The electrode encircled the entire nerve bundle in Monkeys Y and Z. We subcutaneously routed wire leads connecting cuff electrodes in the arm to a connector port embedded in an acrylic head cap at the scalp. Experiments began several weeks following cuff implantation surgeries.

3.2.3 Implant Characteristics

We utilized multi-contact spiral cuff electrodes (Ardiem Medical, Indiana, PA) with inner diameters of 2- and 3-mm and lengths of 10 mm (Figure 3.1A). The electrodes had eight circular channel sites placed equidistant in four quadrants of the electrode's inner surface. These channel contacts were 0.75 mm in diameter and each set of four channels occupied a circular plane on either the proximal or distal side of the conduit. We monitored electrode impedances before and periodically after surgical implantation to assess predicted bioelectrical properties of the interface. An Autolab Potentiostat/Galvanostat PGSTAT12 (Eco Chemie, Utrecht, Netherlands) with a built-in frequency response analyzer (FRA2, Brinkmann Instruments, Westbury, NY) measured impedances at 1 kHz.

3.2.4 Stimulation

Figure 3.2 displays a typical electrochemical pulse delivered at one or more electrode channels in the interface. Platinum electrode sites delivered pulse trains to the nerves generally of 500 ms duration. Pulses were biphasic, consisting of a leading cathodic phase and a zero-delay charge-balancing anodic phase [33]. The pulse waveform of the charge-balancing phase was of 10x decreased amplitude and 10x increased pulse width to mitigate formation of virtual cathodes. One or more proximal electrodes (P1-P4 in Figure 3.1A) served as the current sink (delivering cathodic stimulation) for all four distal electrodes (D1-D4), which served as current sources and provided an anode block. We designed and administered stimuli using hardware and software from Tucker-Davis Technologies (Alachua, FL).

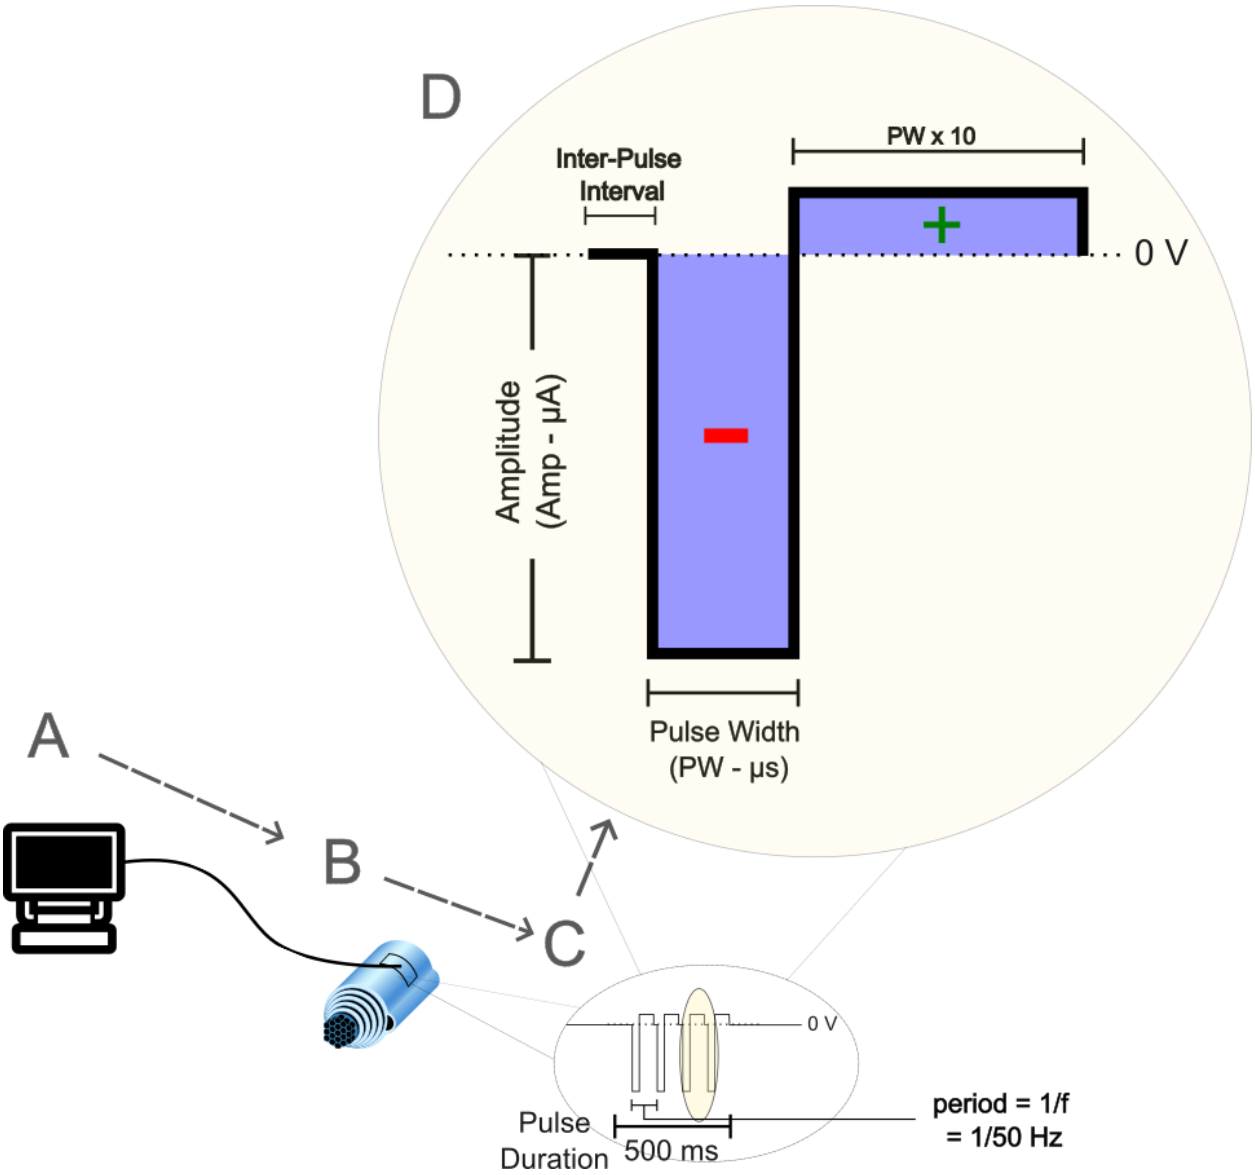

Figure 3.2: *A typical pulse train delivered by an electrode channel.* **A**. Pulse train parameters were governed by a programmable pulse train generator (Tucker-Davis Technologies, Alachua, FL). **B**. The parameterized waveform was delivered at one or multiple electrode sites at the nerve/device interface. **C** A single pulse train, usually lasting 500 milliseconds, was delivered as a test stimulus. The frequency of pulse repetition was generally set to 50 Hz. **D**. We primarily parameterized the waveform with *pulse width* and *amplitude* parameters. Inter-pulse interval was determined by the *frequency* parameter as shown in C. Each pulse consisted of a cathodic activation phase (-) followed by an anodic recovery phase (+). We arbitrarily chose a ratio of 10 to scale the pulse width and amplitude appropriately during the anodic phase. Figure emphasizes waveform features and is not drawn according to scale.

3.2.5 Psychophysical Procedures

Training on a Yes/No Task

Before implantation of the cuff device, each monkey learned to use a joystick in a standard center-out task through training with positive reinforcement operant conditioning. Following cuff implantation, we ascertained a safe baseline current amplitude prior to the next phase of training by gradually increasing the current amplitude from zero in increments of 10 to 50 µA until we observed slight muscle twitching in the fingers and hand. We established this baseline maximum at about 1 mA, using a 500 ms pulse train of 50 pulses per second and a 200 µs pulse width. Training on the subsequent task began slightly below this level.

Next, each monkey learned a yes-no detection task (Figure 3.3). The monkeys initiated a sequence by using a joystick to move the cursor to a circular target centered on the screen. A ring surrounded the cursor upon arrival at the center target. The monkey was then able to initiate a trial by maintaining the cursor in the center position for 750 to 1500 ms. We then delivered a *test trial* stimulus or a *catch trial* non-stimulus with equal probability. There were one of four trial outcomes. For *test trials*, the monkey received 1) a reward for moving the cursor to the outer ring (*Hit*), or 2) a time-out penalty of 1 to 3 seconds for remaining in the center of the ring (*Miss*). For *catch trials*, the monkey received 3) a reward for retaining the cursor in the center of the ring (*Correct Rejection*), or 4) a penalty for moving the cursor to the outer ring before trial conclusion (*False Alarm*). We applied a post-trial delay of 2 seconds before allowing initiation of the next sequence.

Each monkey trained for several weeks on the yes/no task, beginning at stimulation levels just below the motor twitch threshold and with a visual cue. We noted successful task learning when each monkey rapidly improved performance to at least 90% for all trials, including *catch* trials. Once a monkey reliably performed the prior task, we removed the visual cue and gradually decreased the current amplitude to levels near threshold, as observed by declining performance. During all training and experiments, each monkey executed the task for approximately 1000 trials per day.

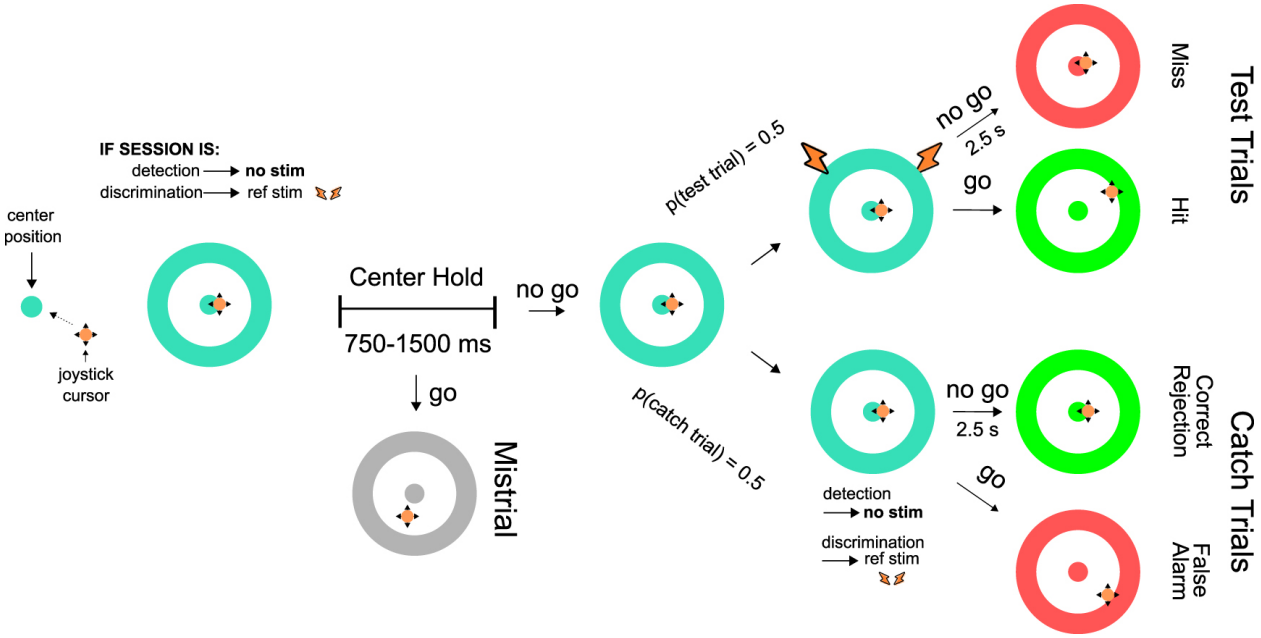

Figure 3.3: *Yes/No joystick task.* Monkey initiated a trial by moving a cursor to a circular target displayed in the center of the screen, at which point a *reference* stimulus was delivered (I=0 for detection experiments, I≠0 for discrimination experiments). After a brief delay period, a *comparison* stimulus was presented. In 50% of trials, designated *catch* trials, we delivered a comparison stimulus equivalent to the test stimulus. In the other 50% of trials, designated *test* trials, the comparison stimulus was one of eight stimulus levels higher than the reference stimulus. Proportions of four outcomes (Hit, Miss, False Alarm, and Correct Rejection) conveyed task performance.

Method of Constant Stimuli (Detection Task).

Fully trained monkeys performed detection threshold experiments administered using the *method of constant stimuli* as follows. We chose conditions of each block of trials such as the modulated dimension and electrode channel configuration in an exploratory fashion. The stimulus level for a given trial was chosen randomly from a set of 8 or 9 evenly spaced stimulus levels over a designated range in the modulated dimension while the other stimulus dimension was held constant. We held pulse width constant during amplitude modulation blocks, usually at 200 µs. For pulse width modulation blocks, we held amplitude constant at whichever intensity allowed us to sweep the range of pulse width values bound by the frequency of the pulse train and hardware limitations. At the beginning of each session, we monitored responses to the task in real time and determined an appropriate stimulus range, usually beginning with the previous session's range and making operational adjustments to the stimulus range when necessary. Adjusted stimulus ranges resulted in hit rates near 100% at maximum stimulus levels and near 10% at minimum levels.

Discrimination Task.

After months of detection experiments, we adjusted the *method of constant stimuli* to administer discrimination experiments by concurrently delivering a reference stimulus upon trial initiation and before the delay period. Thus, the behavioral outcomes conveyed whether the monkey could reliably *discriminate* between the reference and test stimuli, in contrast with the previous set of sessions where responses conveyed *detection* of the test stimulus or lack thereof. This study refers to the difference between the absolute stimulus level at the estimated threshold I_{thresh} and the corresponding reference stimulus level I_{ref} as a *just-noticeable difference (JND)*, although other terms are in the literature such as *difference threshold* or *difference lumen*.

$$\text{JND} = I_{thresh} - I_{ref} \tag{3.1}$$

We administered multiple trial blocks at different reference stimulus levels or electrode channel configurations for most sessions. We chose reference stimulus levels for discrimination blocks from a range beginning near the detection threshold level and working up to previously

elicited motor activation levels, avoiding higher stimulus levels where the monkey became hesitant to initiate trials, suggesting discomfort. We inspected the shape and consistency of each resulting psychometric curve before proceeding to the next parameter set.

3.2.6 Analysis

We used hit rate (proportion of *hits* from *test trials*) as the performance metric for each block of trials. We estimated the parameterization of the psychometric function (Figure 3.4A) using Bayesian inferencing [50]. These estimations, provided by the *psignifit4* Python package [88], fit the results to the equation:

$$\Psi(x; \alpha, \beta, \gamma, \lambda) = \gamma + (1 - \gamma - \lambda)F(x; \alpha, \beta) \tag{3.2}$$

where Ψ is the chosen performance metric, x is the magnitude of the stimulus along the modulated dimension, α is the threshold, β is the width of the function between 5% and 95% performance, γ is the guess rate, λ is the lapse rate and F is a sigmoid function (we used the cumulative Gaussian function). We used *psignifit4*'s default prior distributions.

We analyzed aggregate data with lab-developed Python data visualization software to detect anomalies, identify trends, and compare experimental conditions[6]. Each data point that was fed to regressions represented a maximum-a-posteriori (MAP) point estimate of the threshold from each experimental block, and we explain limitations of this choice in the *Discussion* section of this chapter. We performed regressions using the *statsmodels* Python package and performed Analysis of Covariance (ANCOVA) tests with the *aoctool* function in MATLAB's Statistics and Machine Learning Toolbox. We then used the MATLAB tool *multcompare* to perform a Tukey-Kramer multi-comparison on the results.

We analyzed strength-duration patterns of nerve activation by estimating amplitude thresholds at several fixed pulse width levels and, inversely, estimating pulse width thresholds at several fixed amplitude levels. We used these estimates to perform ordinary least squares (OLS) linear regression, adapting Weiss's linear formulation for the strength-duration curve [8]:

[6]Repository at https://github.com/psychoanalyze/psychoanalyze

$$Q_p = I_0(t_{PW} + \tau_{SD}) \tag{3.3}$$

where t_{PW} is the pulse width and Q_p is charge per pulse ($t_{PW} \times I$). The regression provides estimates of strength-duration time constant τ_{SD} and an asymptotic current level I_0. When thresholds represent stimulation levels that trigger action potentials from single nerve fibers or compound action potentials from fiber populations, τ_{SD} is historically referred to as the *chronaxie* and I_0 is referred to as the *rheobase*. Our study extended this line of analysis through the central nervous system by measuring thresholds via learned behavior [91].

Finally, we examined the results of discrimination experiments using Weber's Law. Weber's Law is regarded as a generalized principle, which states that the perceived difference between stimuli is proportional to the size of the stimuli. Often, modeling this principle uses the equation:

$$k \propto \Delta I/I \approx \text{JND}/I_{ref} \tag{3.4}$$

where k is the Weber fraction/coefficient, and $\Delta I/I$ is interpreted as JND/I_{ref} for our experimental context. We can rearrange this equation into the familiar $y = mx + b$ form as:

$$\text{JND} = I_0 + kI_{ref} \tag{3.5}$$

While potential mechanisms behind Weber's Law and its general applicability for sensory models are debated in neuroscience and psychology [1, 47, 97], our study utilizes the model as a simple, linear approach to contextualize data in a variety of circumstances which likely conveys some true subjective experience by the monkey.

3.3 Results

We conducted more than 400 experimental sessions ranging from 87 to 711 days post-implantation. Table 3.1 lists the number of experimental sessions performed by each monkey, the number of days since implantation on the first day of experiments, and the number of

Table 3.1: Experimental sessions performed by each monkey

Monkey	Number of Sessions	Start Day*	End Day*
U	236	94	711
Y	170	174	630
Z	25	87	126

* *Start Day* and *End Day* relative to day of implant surgery

days since implantation on the last day of experiments. Implants in Monkeys U and Y delivered detectable signals through 1-2 years of chronic usage. The implant in Monkey Z exhibited stable signals up to 4 months post-implantation until lead failure near the site of the head cap terminated further experiments in this case.

Figure 3.1B shows average electrode site impedance data at 1 kHz for the three implants across the experiments. Average electrode site impedance increased by approximately an order of magnitude from the days before and after implantation. Trends showed stable implant impedances for more than hundreds of days, but the implant in Monkey Y exhibited higher impedances near the end of the experiments. All electrode sites remained functional throughout experiments for obtaining psychophysical data.

Figures 3.4B-D show the empirical cumulative distribution functions (eCDFs) that summarize the parameter fit estimates from all blocks. The parameters of guess rate (γ) and lapse rate (λ) were generally low and in line with previously reported rates in animal experiments [95]. Low values for these "nuisance" parameters served as indicators that the monkey was performing the task honestly and reliably. For example, Monkey Y exhibited higher guess rates on average than Monkey U, but still maintained a guess rate below 10% for approximately 80% of sessions (Figure 3.1B).

3.3.1 Detection Thresholds

Figure 3.5 shows trends in detection thresholds from each monkey over the course of amplitude modulation experiments. Monkey U (split median nerve) completed the greatest number of experimental sessions over the longest period. This monkey's thresholds decreased from initial levels >1 mA down to about 200 µA, interrupted by sharp threshold increases near days 140, 210, and 260. Monkey Y (ulnar nerve) needed less current to detect

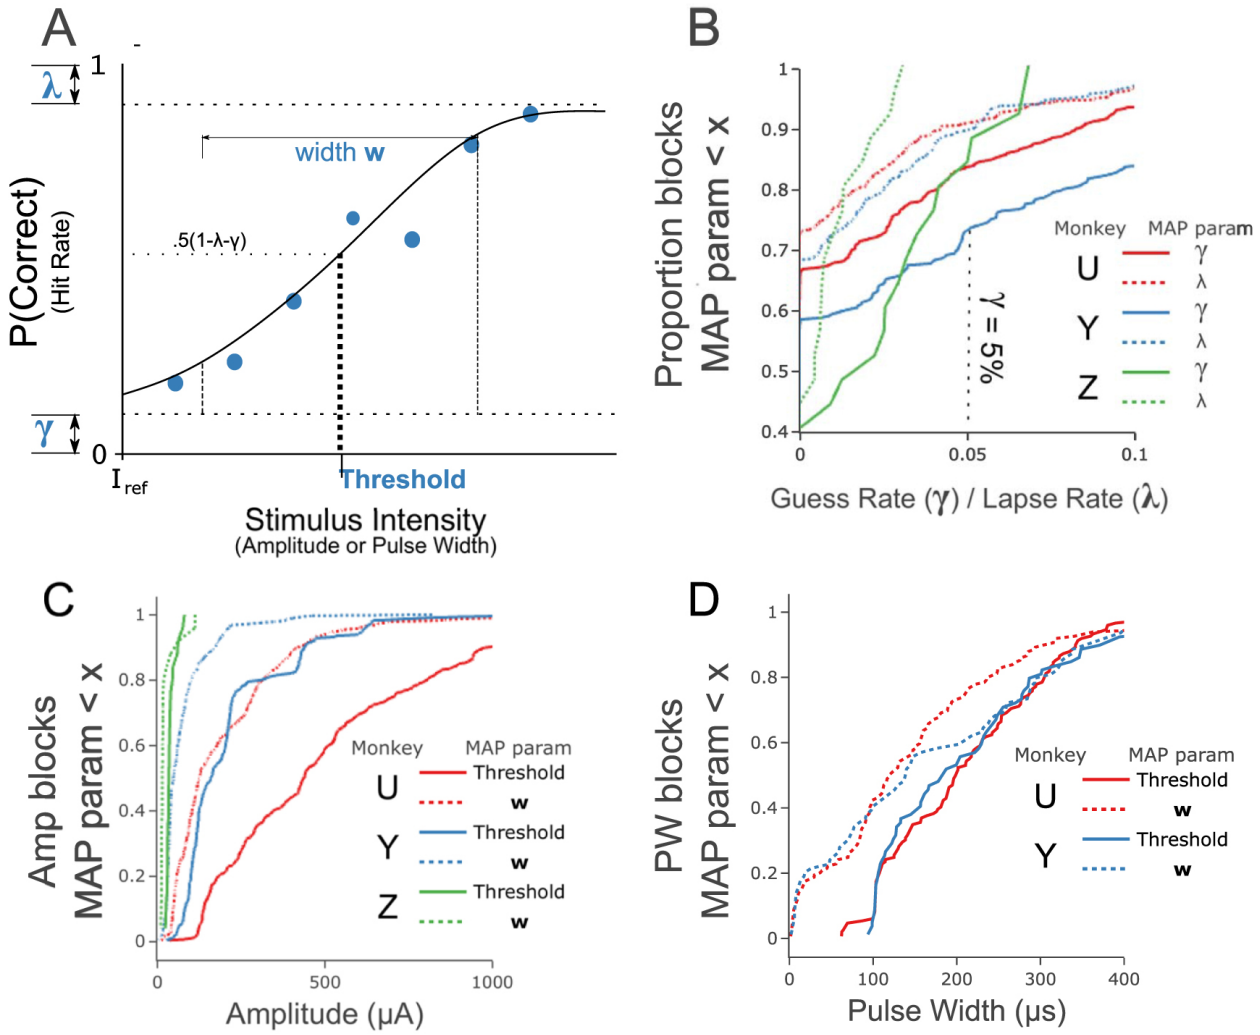

Figure 3.4: *Summary of fitted parameters - Detection Experiments* **A.** Stereotyped psychometric function. Four parameters—threshold (**Threshold**), pulse width **w**, guess rate (γ), and lapse rate (λ)—determined the shape of the function. We chose *hit rate*—(n *hits* / n *test trials*)—as our performance metric (P(Correct)). **B.** *Empirical cumulative distribution functions (eCDFs) for γ and λ estimates.* eCDFs convey the raw distribution of observed measurements. The eCDF value at x communicates the proportion of blocks where the fitted value was less than or equal to x. For example, the highlighted example in **B** shows that Monkey Y had a guess rate less than 5% for 70% of blocks. All monkeys showed guess rates and lapse rates near zero for >40% of sessions (lower left of figure) and showed rates below 10% for >80% of sessions (upper right of figure). **C, D.** eCDFs of *threshold* (solid lines) and *width* parameters (dotted lines) for amplitude modulation blocks (*C*) and pulse width modulation blocks (*D*) for all detection blocks.

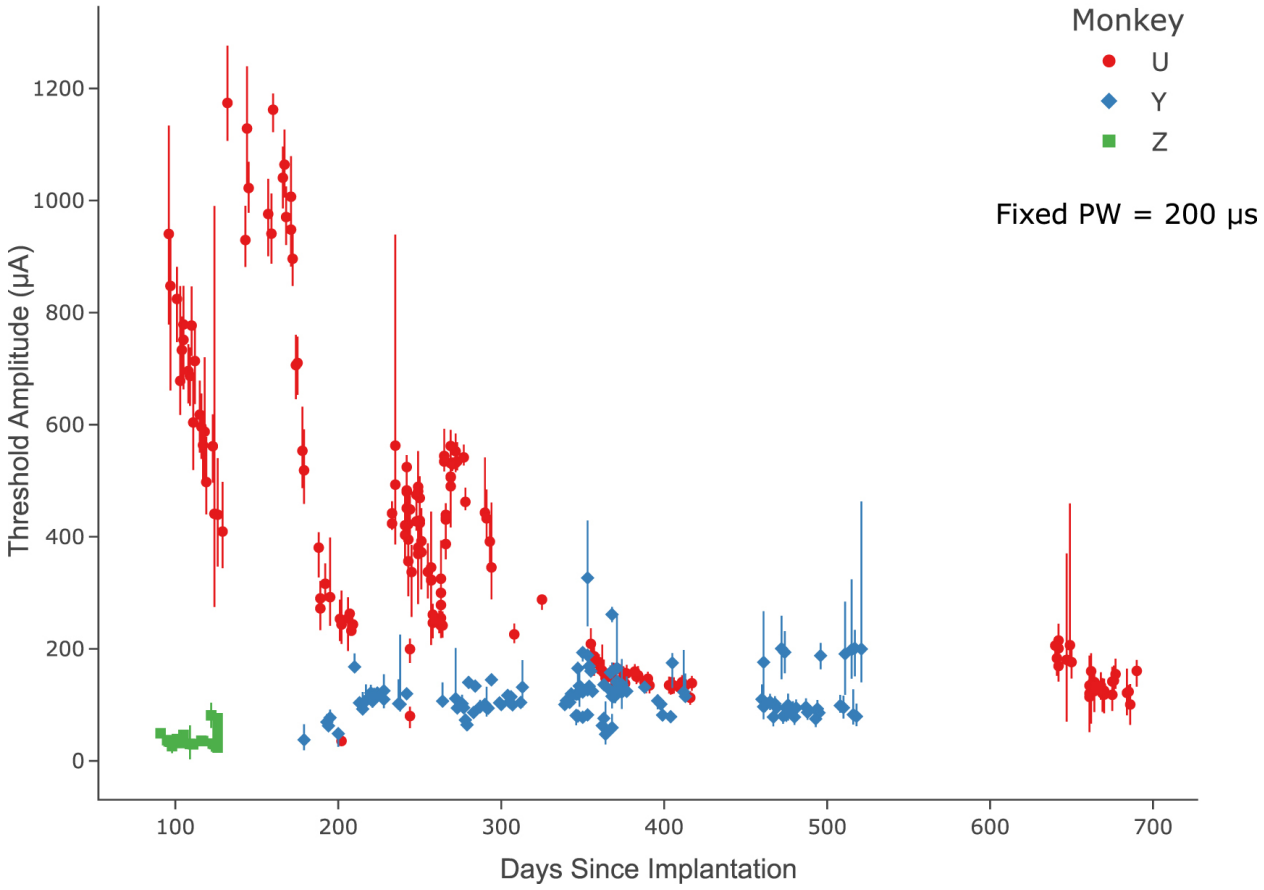

Figure 3.5: *Amplitude detection thresholds.* Longitudinal plot of detection thresholds for amplitude modulation blocks. Point estimates represent *maximum a posteriori* estimates of the parameters from curve fitting, and error bars represent 95% confidence intervals. For the data shown, pulse width was held constant at 200 μs. Monkey U (red circles) performed for nearly two years and showed highly fluctuating thresholds that settled to a steady state at about 350 days. Monkey Y (blue diamonds) maintained much lower thresholds for nearly a year of experiments. Monkey Z (green squares) exhibited the lowest thresholds, but experiments ended prematurely due to a lead failure near the site of the head cap.

Table 3.2: Strength-duration estimates, visualized in Fig3.6

Monkey	Date Range	Rheobase (μA)	Chronaxie (μs)	R^2
U	62-75	48.5	18.5	0.983
U	110-114	53.1	4.46	0.959
U	131, 195	20.2	8.41	0.993
U	201-210	38.0	9.11	0.990
Y	56-71	5.91	10.4	0.939

threshold, averaging pulse amplitudes of 115 ± 31 µA at threshold from the beginning of experiments. Monkey Z (intact median nerve) had the lowest detection thresholds (36 ± 13 µA) across all experiments, however, these experiments ended prematurely due to an electric lead failure at the head-cap.

Figure S1 in the supplemental materials of the publication presenting this work[87] shows the same data as in Figure 3.5 delineated by the channel configuration chosen for each block, with data shown on a log scale to emphasize within-subject differences across time. Thresholds in Monkey U displayed correlation between channel thresholds despite large fluctuations in threshold. We could not perform typical correlation tests because the time series data was unbalanced, meaning that they were sampled at different relative time points for each condition. However, the appearance of correlation in Monkey U paired with the stability and spread of data seen in Monkey Y suggests that the choice to pool data across channel configurations did not introduce significant levels of bias in the measurements.

Figure 3.6 conveys thresholds estimated similarly for various choices of a) modulated stimulus dimension (*amplitude* or *pulse width*) and b) the fixed stimulus level in the non-modulated dimension. Figure 3.6A (amplitude blocks) and Figure 3.6B (pulse width blocks) convey the inverse relationship between amplitude and pulse duration traditionally observed in strength-duration data. The same data from Figures 3.6A and 3.6B are shown transformed and regrouped on the shared *charge* axis in Figures 3.6C (Monkey U) and Figure 3.6D (Monkey Y) according to Equation 3. Strength-duration coefficients (slopes and intercepts, Figures 3.6C, 3.6D) were nearly identical between amplitude blocks and pulse width blocks within subjects across time within the observed time bins. Table 3.2 provides the regression coefficients for strength-duration time constant τ_{SD} and an asymptotic current level I_0 for these time bins.

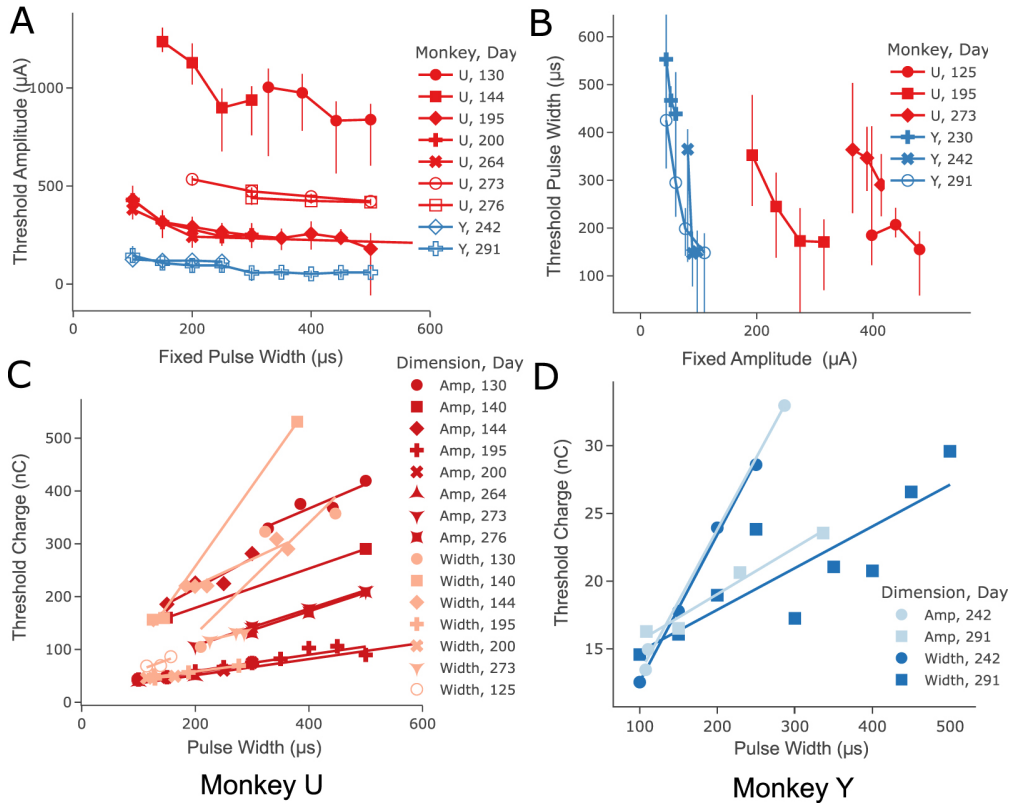

Figure 3.6: *Strength-duration characteristics of behavioral thresholds.* **A-B.** We employed the *method of constant stimuli* at a range of fixed pulse widths for **amplitude modulation blocks** (**A**) and at a range of fixed amplitudes for **pulse width modulation blocks** (**B**) in both monkeys. Error bars indicate standard deviations around the mean threshold when pooled at each stimulus level. Each trace indicates an iso-contour at detection threshold between waveform parameters amplitude and pulse width. These contours displayed the hallmark inverse relationship observed in strength-duration data. **C-D.** *Weiss (linear) formulation of the strength-duration relationship.* We combined thresholds for each monkey from the previous two subplots and displayed data from both amplitude modulation and pulse modulation sessions on a shared *charge* axis according to Weiss's linear formation of the strength-duration relationship (Equation 3.3). Linear regression coefficients (slope and intercept) for blocks with amplitude and pulse width modulations were nearly identical and correlated across time within subjects.

3.3.2 Discrimination

Figure 3.7A shows average JND values against reference charge levels over all discrimination blocks, pooled at each reference stimulus level for plotting only. The slope of the fitted lines represents the Weber fraction over all pooled data for each monkey and each stimulus dimension (solid lines/markers = amplitude modulation blocks, dotted lines/open markers = pulse width modulation blocks). Figure 3.7B shows the results of an ANCOVA test on the data from Figure 3.7A. Under the assumptions of the test, the aggregate Weber fraction was significantly lower in blocks with amplitude modulations than with pulse width modulations for both monkeys, as demonstrated by non-overlapping comparison intervals in Figure 3.7B. Additionally, the Weber fractions were not significantly different for the same dimension across different monkeys, as demonstrated by overlapping intervals. Table 3.3 shows the precise values for Weber fractions estimated from each experimental configuration and monkey. Supplemental Figure S2 from the publication presenting this work [87] shows the distribution of Weber fractions from all individual blocks.

We pooled data across all channel configurations, as with data from detection experiments. Supplemental Figure S3 from the publication presenting this work [87] shows data from Figure 3.7 grouped by channel configuration. We observed simultaneously that a) many configurations had similar outcomes, demonstrated by the proximity of the fits to each other in relation to the error bars, while b) some configurations still showed deviations from the rest—*e.g.* Monkey Y, Channel 2—that indicate some variance is introduced by choice of channel configuration.

3.4 Discussion

Electrical stimulation, administered through cuff electrodes around upper arm peripheral nerves, evoked durable sensations in three different macaques. The findings demonstrated long-term durability of a sensory cuff electrode in the upper limbs of macaques during psychophysical experiments for months to years. Three monkeys showed graded sensory activation from stimulus patterns scaled by pulse width and current amplitude, despite differences between cases including nerve anatomy, electrode diameter, and location of the implant. We took these measurements with a frequency on the scale of days, i.e., with high enough

Table 3.3: Weber's Law regression coefficients for data shown in figures

	Monkey	Channel	Intercept	Slope
Amplitude	U	average	73.7	0.1
	Y	1	-6.78	0.17
		2	-1.007	0.129
		3	4.97	0.0886
		average	3.46	0.109
Pulse duration	U	1	-28.2	0.500
		3	-18.9	0.364
		4	-49.0	0.631
		average	-29.4	0.453
	Y	1	-61.7	0.638
		2	-100.0	1.00
		3	-14.0	0.427
		4	-35.3	0.548
		average	-35.3	0.548
Charge	U	1	14.4	0.129
		2	-26.6	0.719
		3	-7.54	0.364
		4	-13.9	0.450
	Y	1	-4.20	0.311
		2	-17.3	0.874
		3	2.39	0.0743
		4	3.76	0.0873
Charge (Amp)	U	average	22.6	0.0752
	Y	average	0.913	0.102
	average		11.8	0.0888
Charge (PW)	U	average	-14.8	0.507
	Y	average	-13.4	0.776
	average		-14.1	0.642

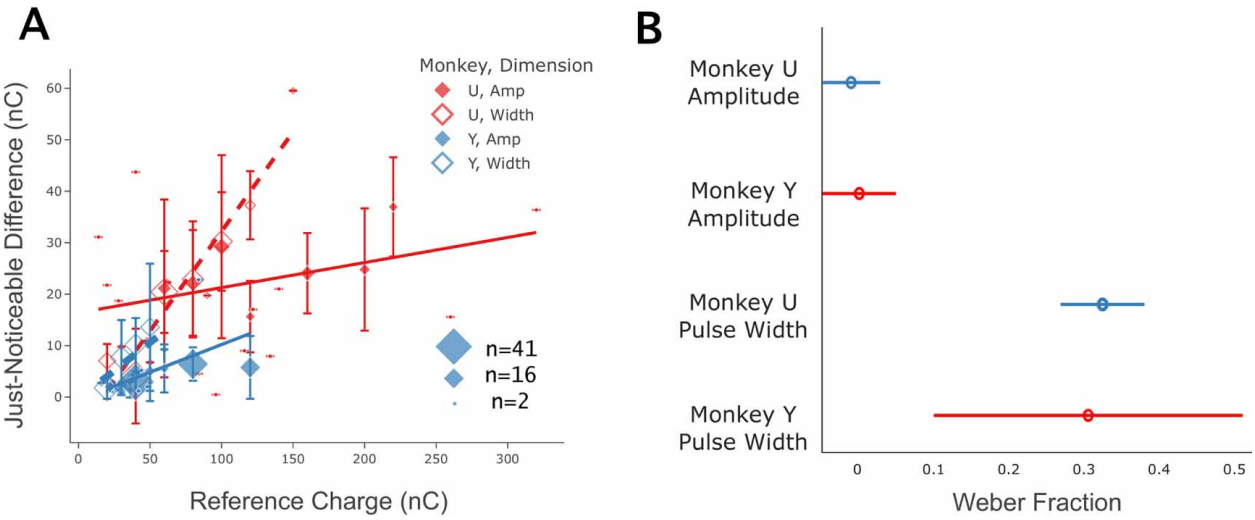

Figure 3.7: **A.** *Discrimination and Weber coefficients.* We assessed the monkeys' abilities to discriminate between stimulus levels by observing responses and estimating JNDs at various reference stimulus levels. Plots show JNDs relative to reference stimulus levels according to Weber's Law. Marker size indicates the number of sessions at each reference stimulus and error bars show standard deviations of each grouped set of sessions. The slope of each regression line represents the group-level Weber coefficient. **B.** Results of an analysis-of-covariance (ANCOVA) test on the Weber fractions from *A*. The width of the lines represents comparison intervals such that non-overlapping bars are statistically significant. Weber fractions were significantly different between amplitude blocks and pulse width blocks but did not show significant differences for each dimension across monkeys.

frequency to inform speculation on potential mechanisms of variation in the data, including between- and within-subject differences, across months.

Almost all electrode channels remained functional throughout the experiments. Only irreparable damage to a percutaneous lead compromised enduring results from Monkey Z. However, collected data indicated excellent sensory detections even in this truncated case. Future experiments and clinical applications might utilize wireless implant technology to mitigate the risk of lead damage and increase PNI durability.

Our study suggests that the ratio of pulse width to pulse amplitude in the stimulus waveform may not affect detection thresholds, but the choice of modulating amplitude vs pulse width at stimulation above threshold has significant effect (see *Limitations* for caveats). If this conclusion holds, then amplitude modulation may have more resolution and dynamic range than pulse width modulation, e.g. for simple 1-D prosthetic control.

3.4.1 Stability Assessment

Thresholds in Monkey U, compared to Monkeys Y and Z, showed greater day-to-day variability and larger, slower fluctuations. Only Monkey U underwent a different surgical procedure to isolate sensory nerve bundles (see *Methods*). The major discrepancies in threshold characteristics between Monkey U and the other monkeys, including sudden perturbations in threshold, might have reflected ongoing inflammation, nerve damage, and healing in the nerve and surrounding tissue in Monkey U stemming from the surgical procedure or abnormal stimulation [12]. By isolating the sensory branch, we reduced the fiber count contained within the cuff electrode. The potentially increased distance from the electrode contacts to the nerve fibers in closest proximity might cause electrical shunting and high, unstable thresholds. In addition to or separate from this effect, axonal damage possibly occurred during the implant procedure. Alterations in fiber properties resulting from axonal damage could increase activation thresholds to stimuli applied from a cuff electrode transiently or permanently [22, 65]. This would assume the cuff was around a section of the nerve where such damage had occurred. Over time, inflammation and subsequent encapsulation of loose extraneural space [37], in addition to the recovery of axonal properties, might contribute to the gradually stabilizing stimulus amplitudes we see in Monkey U (Figure 3.5). However,

we did not observe a significant increase in impedances (Figure 3.1B). We also did not observe significantly higher thresholds in Monkey U compared to the other monkeys once they stabilized (right side, Figure 3.5), as expected for encapsulated tissue.

3.4.2 Strength-Duration Findings

We applied the Weiss strength-duration formula (Figures 3.6C-D, Equation 3.3), which previously modeled thresholds for single axon activation [8] and compound action potentials [64]. Roughly, we observed fits from these regressions that were similar or nearly identical on short time scales of fewer days within a monkey, but coefficients drifted in a correlated manner over the span of weeks or months.

As is the case with strength-duration parameters from single fibers and compound action potentials, these estimates may convey some present state of the system, potentially rapidly measurable for calibration of peripheral nerve devices [100]. The apparent correlation between the strength-duration coefficients and detection threshold (Figures 3.5, 3.6C) warrants further examination.

The estimated strength-duration coefficients gleaned from detection thresholds were remarkably similar for each monkey on a given day or short span of days regardless of the modulated stimulus dimension. Thus, a model for simple detection threshold phenomena may be able to collapse the dimensions of amplitude and pulse width into a single dimension, charge-per-pulse (Q_p).

3.4.3 Weber Analysis

We compared the Weber coefficients obtained from blocks of amplitude modulation vs pulse width modulations. This analysis examined the hypothesis that stimulus intensity, encoded by charge per pulse Q_p, was independent of the ratio of amplitude to pulse width. While the structure of this study did not lend itself to straightforward hypothesis testing (see *Limitations*), we generally observed lower Weber fractions when modulating amplitude vs pulse width in both monkeys.

In contrast with the effects of modulated dimension on strength-duration estimates, signficant differences in the Weber fraction between amplitude-modulation blocks and pulse-width-modulation blocks in discrimination experiments suggested that the relative contributions of amplitude and pulse width may affect the quality of sensations and/or the nerve recruitment function at levels above the detection threshold. This evidence suggests the two modulation approaches would result in different fiber recruitment profiles in the nerve [33, 36, 59, 15, 89]. If amplitude modulation shows lower Weber fractions than pulse width modulation across comparable conditions and various experimental and clinical contexts, practical application would suggest amplitude modulation would be more energy efficient and allow for more dynamic range than pulse width modulation for simple (i.e. 1-D) sensory feedback engineering.

3.4.4 Limitations

We applied linear regressions and ANCOVA tests to marginal MAP estimates for each block of data. Grouping choices only weakly attempted to account for violations of test assumptions such as temporal non-stationarity, within-subject pooling, and cross-channel pooling. The data thus reflected this series of grouping decisions, which likely biased the statistical estimates. For example, grouping data across channel configurations violated the test assumption that each measurement came from the same natural probability distribution; our data showed some systematic variation across channels. Similarly, performing regressions against MAP point estimates for threshold did not relay the uncertainty present in these estimates due to sampling size, etc. We attempted to group measures visually if the spread of the data appeared to revolve around a consistent mean and standard deviation to simplify preliminary analysis without introducing extreme biases. However, caveats of overconfidence and Type 1 errors apply to all statistical test measures reported in this study. Notably, we did not exclude groups when performing such pooling to hedge against "data snooping."

While the statistical outcomes of this study come with these caveats, they will be useful for guiding the design of future PNI experiments, which may be used to make stronger assertions about sensory models, including the classical models tested here as well as more advanced models that attempt to explain the component mechanisms of the interface. Follow-up analyses of our data set might apply a *hierarchical Bayesian model* (or its frequentist counterpart, a *Linear Mixed Model*) that would more fully account for violations of our test assumptions [40, 71]. Additionally, analyses that incorporate *signal detection theory* may leverage

the estimated values for our "nuisance" parameters to more fully account for variation introduced by the attentive or motivational state of the monkey [53].

We did not adjust pulse rate (50 pulses per second) or the overall length of the pulse train (0.5 seconds). However, follow-up experiments that examine these parameters may discover further stimulation strategies for PNIs. Particularly, examining the effects of modulating the overall stimulus duration may more fully account for potential adaptation effects not considered in our experimental design.

Monkeys may have partially reacted to efferent-to-afferent volleys resulting from muscle twitching when making judgements. Such effects may only have been applicable for Monkeys Y and Z, whose cuff implants circumscribed the whole nerve, including motor fascicles. Conversely, segregating the sensory branch of the median nerve in Monkey U likely minimized effects from potential muscle twitching. We provide three considerations that suggest minimal contamination from motor activation, even in Monkeys Y and Z:

Context from literature. Early studies have presented data suggesting lower stimulation thresholds for sensory fibers compared to motor fibers [17]. Recent studies have attributed such discrepancies to phenotype-specific differences in axonal membrane conductivities [41]. Computational models of axons with membrane properties that represent each phenotype suggest substantially lower recruitment thresholds for sensory fibers of a given caliber compared to motor fibers [23].

Anode block. We used the distal set of electrodes as anodes, which hyperpolarized axon membranes distal to activation, theoretically further increasing the threshold for motor activation.

Visual inspection. We detected no muscle twitches during experiments or during stimulation sessions where we attentively observed the forearm movements of each monkey. However, it was difficult to observe the monkeys' arms consistently during experimentation.

3.4.5 Future Directions

We provide our data analysis software as an open-sourced Python package (see next chapter) and provide a browser-based data visualization dashboard for simple and responsive psychophysics data analysis, particularly suited for longitudinal and multidimensional data sets. The software may analyze novel data from similarly designed studies to execute common psychophysics analyses performed here and elsewhere. More advanced psychophysical procedures—such as adaptive sampling of multidimensional space to efficiently gain information via Bayesian methods—would effectively supplement the Bayesian components of this study [108]. Additionally, future studies may combine aspects of our approach with brain, nerve, or muscle recordings to supplement behavioral results, possibly enabling usage of more detailed peripheral nerve integration models [66].

Macaque upper limb nerve experiments enabled acquisition of large data sets resulting from stimulation applied to a system similar to human peripheral nerves in size, dexterity, cortical mappings, and re-mappings [10]. Animal data sets may help validate or tune sensory encoding models that would be directly applicable to a human interface. However, the empirical measurements derived from model parameters would likely vary across species. In particular, the increased number of fascicles in the human compared to monkeys would likely result in different thresholds and recruitment profiles [10]. However, some consequences of the anatomical differences between species may be advantageous to clinical translation in this instance. Inducing graded sensations might be less challenging in the larger, more fascicularized human nerve in comparison with the macacque nerve. Specifically, the low conductivity of the perineurium and the spatial separation between fascicles facilitates fascicle-specific activation and may easily reflect distinct sensations due to topographical organization within the nerve.

The larger human nerve trunk may furthermore increase activation thresholds due to the larger size of an appropriately sized cuff and the potentially larger distance between electrode sites and axons. In contrast, the larger fiber calibers present within the human nerves [23, 11] relative to those of found in the macaque [3] would decrease the activation thresholds. The effect of these discrepancies would counteract in sum. It is therefore difficult to speculate about the overall effects of these differences absent supporting data from further studies.

Neural interfaces in use today commonly require calibration routines before daily use. The experimental and statistical groundwork of our study may provide a clearer picture of the calibration requirements of assistive devices and open the door for sophisticated, targeted calibration routines that account for the dynamics of the system.

3.4.6 Conclusion

Successful demonstration of long-term usage of a peripheral nerve interface in a macaque model encourages future experiments with cuff-style peripheral nerve interfaces. The current study had monkeys calibrate a nuanced sensorimotor behavior involving a varied parameter space and identifying matching stimuli. We used simple linear models and common assumptions to examine the effects of many stimulation configurations in a longitudinal study, which more closely resembled clinical usage and time scale than many shorter studies. Sensory feedback from a prosthetic device might similarly utilize electrical nerve stimulation.

Chapter 4

PsychoAnalyze: An Open-Source Toolset for Psychophysics Analysis

4.1 Introduction

Several software packages exist that provide tools for fitting the psychometric function to experimental data, but few provide tools for data manipulation and visualization that accomodate more complex experimental setups that require intricate data flows.

PsychoAnalyze aims to make more advanced psychophysical analysis accessible to researchers without extensive software or data engineering background [35], and to provide an accessible development platform to encourage contributors from all backgrounds to build custom features and plugins for the package.

Because Python has rapidly gained popularity in both data science and the natural sciences, a Python package that properly utilizes and integrates the wealth of developer tools and data tools available in the Python ecosystem[7] can provide researchers and developers with stronger mental models and a tighter feedback loop.

In an effort to address data management challenges encountered in psychophysics experiments, I developed an open-source tool, PsychoAnalyze, which is published at https://github.com/psychoanalyze/psychoanalyze. I developed PsychoAnalyze with the following goals in mind:

- Provide a Python package for querying, manipulating, and analyzing psychophysical data. The package should be *pip*-installable and otherwise adhere to modern Python

[7] https://numfocus.org/sponsored-projects

packaging and distributing conventions. The package should be well-documented; documentation should be built with a static site generator documentation framework and be available on the internet. Additionally, relevant tutorials and demonstrations should be provided in common "notebook" formats such as Jupyter and Quarto. The package should adhere to semantic versioning or a similarly justified versioning scheme when possible.

- Provide a web-hosted dashboard for interactive exploration of psychophysical data, including the ability to simulate or upload custom data sets, and to provide simple plots describing the data with common psychophysical metrics. This tool should be accessible to researchers and students with minimal coding expertise or background in psychophysics theory. The dashboard should be provided both as a no-code tool so that anybody may utilize its features, but also as a documented example of how to effectively use the package and perform the relevant data manipulations at each point of the process.

- Project structure, tooling, and documentation should prioritize beginner-friendliness to facilitate code contributions from the community and the ability to deploy self-hosted versions of the application. *Agile* software development practices, including test-driven development, should be prioritized to facilitate rapid prototyping and iteration of new features upon publication of the software. Students and researchers at all levels of programming and data engineering practices should still be able to contribute to the project via proposal and testing of new features. The quality of the code base should be evaluated by the time to production of new features, bug fixes, and UI improvements.

The following sections provide detailed descriptions of the software in the context of the above goals at the time of the publication of this dissertation. Readers interested in the current state of the project should visit the project's home page at `https://psychoanalyze.io`.

4.1.1 Comparable Work

There are several existing software packages that provide tools for psychophysical analysis. PsychoAnalyze seeks to bridge the gap in the data pipeline between experimental design software such as *PsychoPy/PsychToolbox* and model-fitting software such as *Palamedes/psignifit*.

- PsychoPy [75] is a Python package that provides a complete suite of tools for designing and running psychophysical experiments on a personal computer device.
- PsychToolbox [9] a MATLAB package that similarly provides routines for stimulus presentation and data collection.
- Palamedes [80] is a MATLAB toolbox that provides an advanced set of curve-fitting procedures, including procedures that use subject-level data to use hierarchical Bayesian methods for more accurate estimates.
- psignifit [88] is primarily developed in MATLAB, but there is a Python port of the implementation. It mainly provides methods for fitting the psychometric function using Bayesian methods.
- BayesFit [90] is a Python-first model-fitting library, but is inconsistently maintained by an individual contributor.

4.2 Core Features

PsychoAnalyze provides data manipulation and visualization tools built on modularity and extensibility, aiming to empower researchers to more fully explore and contextualize their data, while minimizing time spent wrestling with custom analysis scripts.

In addition to providing a Python package and a command-line tool, *PsychoAnalyze* offers a web-hosted dashboard demonstrating the capabilities of the package. On its own, the dashboard provides researchers with a no-code interface to fit their data to a psychometric function and visualize the results in an interactive setting. Further, developers may examine the dashboard code to contextualize the API of the package/library functions.

For example, PsychoAnalyze provides convenient methods to:

- Aggregate trial-level data to an appropriate format (*e.g.*, grouped by intensity level of the stimulus) for model-fitting procedures.
- Generate simulation data according to common psychophysical experimental procedures.

- Export results to a variety of formats such as CSV, Parquet, and DuckDB (tables) or PNG, SVG, and PDF (figures).
- Transform model parameters between parameterizations, *e.g.* location μ / scale σ form to intercept β_0 / slope β_1 form:

$$\psi(x) = \frac{1}{1 + e^{-\frac{(x-\mu)}{\sigma}}} \iff \frac{1}{1 + e^{-(\beta_0 + \beta_1 x)}} \tag{4.1}$$

4.2.1 Dashboard

One of the challenges of multi-dimensional and longitudinal data analysis is the ability to quickly and interactively explore the data. As the number of "dimensions" of the input stimulus parameterization and/or the number of experimental conditions being tested grows, the complexity of the data manipulation code needed to properly contexualize and visualize experimental effects becomes ever greater, as complex patterns of abstractions and inheritance become necessary, and the challenges of analyzing and visualizing higher-dimensional data manifest themselves. Interpretability and complexity are at odds, both in the code base and in the resulting analysis.

Development of the software was driven via the dashboard. New features were designed in the dashboard first, and constant operation of the dashboard alongside developments allowed it to serve as a proxy evaluation of "end-to-end" behavior. The dashboard's skeleton can easily be constructed using mock or hard-coded components, and features can be developed in a way that provides quick and integrated visual feedback to developers (helpful for data-centric *test-driven development*!) and prioritizes user friendliness.

The current iteration of the dashboard is built in Python using the Dash framework, which runs in the browser and is powered by the Plotly data visualization package. The dashboard is currently hosted using the Heroku platform, and a link is provided in the README for interested parties to deploy their own self-hosted version of the application.

The dashboard is composed of three main panels, one for user input of experimental and model parameters, one for the visualization of data, and one for table-formatted output and export options (Figure 4.1).

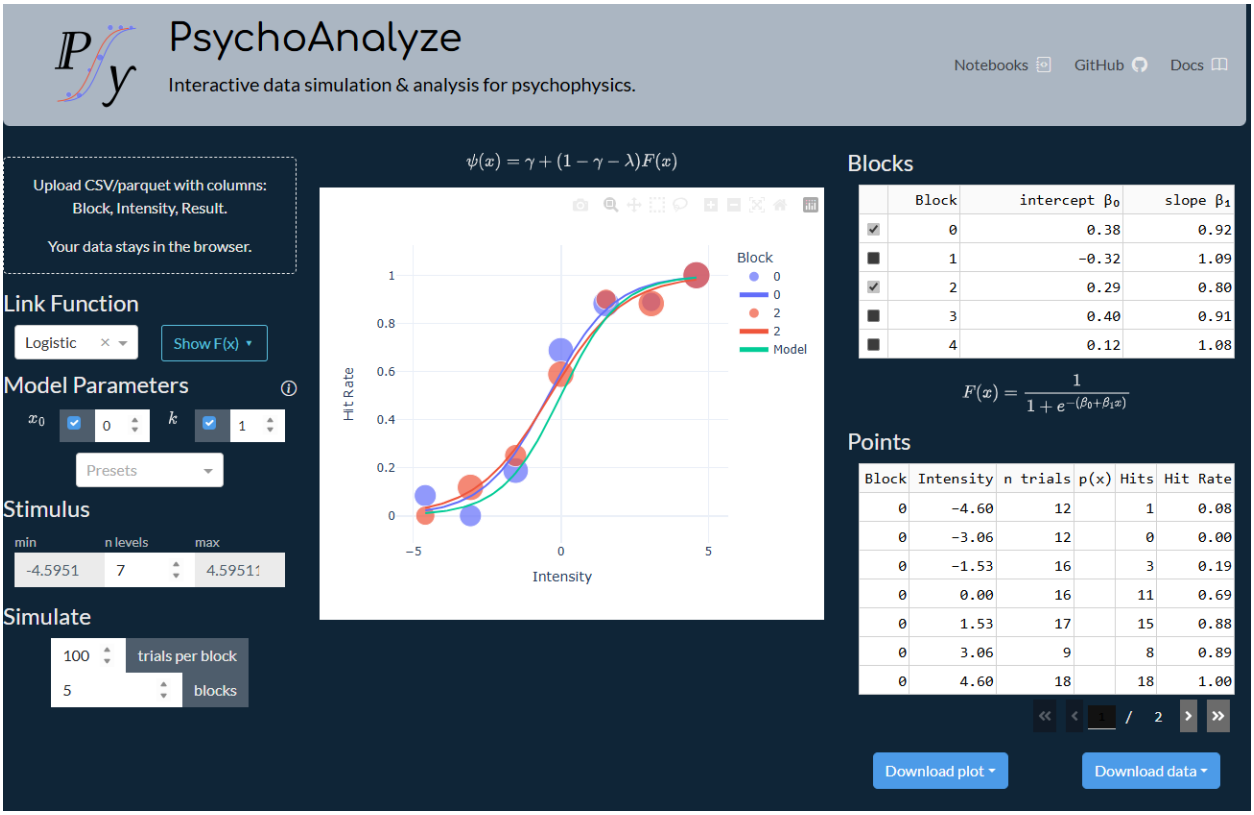

Figure 4.1: A screenshot of the *PsychoAnalyze* dashboard for v1.0.0 of the *PsychoAnalyze* package.

The *Input Panel* provides the user to determine the data processed by the dashboard, either via upload or selection of model parameters and experimental conditions. The first component in the input panel column allows anyone to upload a dataset to be processed by the dashboard. The dataset must be a table either in .csv or .parquet format, and for now is restricted to be in the form of a very simple schema similar to other model-fitting libraries in the ecosystem.

By default, the dashboard displays the results from a simulation of a hypothetical yes-no experiment. A sequence of randomly-sampled trial outcomes are generated, processed, fit to the psychometric function, and visualized, with assistance from PsychoAnalyze's package functions. The simulation estimates the psychometric function's location and scale parameters according to the chosen link function $F(x)$, which is the logistic equation by default:

$$\psi(x) = F(x) = \frac{1}{1 + e^{-\frac{x-\mu}{\sigma}}} \tag{4.2}$$

μ is the location parameter, which represents the psychophysical "threshold". σ is the scale parameter, which represents the slope (i.e., sensitivity) of the psychometric curve near the threshold.

In an upcoming release, PsychoAnalyze will support a variety of link functions. For now, the only supported link function is the logit function. You can toggle the visibility of the equation for the logit function by clicking the Show/Hide button.

You may adjust the parameters μ and σ for model simulations using the sliders and input boxes in the panel. The simulation is completely regenerated in the browser's memory each time any parameter in the Input Panel is adjusted. However, a new feature of Dash has been released where incremental updates to the figure are possible; this will be implemented in a future release.

The Visualization Panel is the central column of the dashboard. It contains the plot of the psychometric function fitted to the simulated or uploaded data. The psychometric function plot is generated using the Plotly Python library. It is an interactive plot that allows the user to zoom in and out, pan, and hover over data points to see their values. Click a legend item to toggle the visibility of the corresponding data series, or double-click a legend item to isolate the trace.

4.2.2 Python Package

The code emerging from the development of the dashboard was translated into a Python package, which is published on PyPI[8]. The package provides a Python API for querying, manipulating, and analyzing psychophysical data. The package is built on top of the Pandas data analysis library, which provides a powerful and expressive API for manipulating and analyzing tabular data. While notebooks provide a popular format for data exploration and visualization, users of PsychoAnalyze should be encouraged to share their notebooks and scripts with the community and to integrate their custom use cases into the package. The *nbdev* package provides an interesting framework for adhering to software engineering principles in a notebook-first development workflow. Early experience with *nbdev* has been positive and it will likely be adopted in future releases.

4.2.3 Notebooks

Notebooks are a popular format for sharing code and data analysis workflows, providing a browser-centric IDE with the ability to organize "cells" of code, Markdown-rendered text, and data visualizations.

PsychoAnalyze has been released with an example notebook that demonstrates the use of the package to analyze data from an example experiment. It is contained in its own GitHub repository at `https://github.com/psychoanalzye/notebooks`. At the time of this dissertation, notebooks in the repository are hosted on a deployed instance of JupyterHub at `https://nb.psychoanalyze.io` on an Azure Kubernetes Service, enabling anyone with Internet acccess to have direct access to an interactive development environment with psychoanalyze and its dependencies installed. Hosted notebooks may be distributed via a different service such as Binder as the cost of hosting is considered, but the notebook files will always be available in the GitHub repository.

[8] `https://pypi.org/project/psychoanalyze/`

4.3 Development Environment

The setup of a local environment, including the installation of library dependencies, can be a major barrier to entry for new developers unfamiliar with the mechanics and installation processes of the specific tooling necessary to run the software, whether they are aiming to run the software locally or need to set up a development environment to make contributions to the code base. Thus, the project takes an opinionated approach to developer tooling with an emphasis on automation and reproducibility.

Python, MATLAB, and R are three high-level languages capable of implementing the robust data workflows necessary for psychophysical data analysis. MATLAB is particularly common in academic engineering departments. The included IDE and ease of development setup also makes it a popular choice. Accordingly, the robust MATLAB library `Palamedes` provides sophisticated analysis procedures to fit psychophysical models to data. However, licensing costs to use MATLAB essentially limit its use to corporations or academic institutions. While the open-source Octave project provides a free alternative, it often lacks the robust third-party library support of MATLAB, as well as core features such as the MATLAB IDE.

The R programming language provides many of the same benefits as MATLAB, including a robust, data-centric IDE (RStudio) and a large library of packages for data analysis (tidyverse). Recently, a robust package, MixedPsy, was published on the R package repository CRAN and provides many similar tools as I seek to provide here, which may be helpful for those used to programming in R.

In summary, however, Python's general-purpose utility for tasks such as web deployment, machine learning, and everything in between give it a leg up on the competition, to say nothing of its very nice collection of general-purpose scientific computing libraries. MATLAB and R's limited capabilities are often a barrier to more robust software development practices. Although they serve a purpose as scripting languages for one-off analyses by non-programmers, in many instances it is beneficial to have common analyses performed by software packages tailored for the job and more thoroughly tested. In these cases, general-purpose programming languages provide libraries and tools that are needed to develop software that can be shared, managed, and developed across the whole scientific community. PsychoAnalyze contains many such more features to facilitate this developer friendliness, including:

- **Dev Containers.** PsychoAnalyze is configured to run in a *Dev Container*, which provides "one-click" development environment setups in Visual Studio Code and other IDEs, including a cloud environment in GitHub Codespaces. Dev Containers allow developers to define a Docker container that contains all of the necessary dependencies to run the software. They allows developers to run the software in a consistent environment regardless of the host operating system, various versions of Python that may be installed, they facilitate collbaration by making development environments more portable between users.

- **pytest.** The *pytest* testing framework provides a structured and flexible approach to automated testing and test-driven development, which is essential to bug-free code, iterable design, and confident collaboration. Having a robust test suite allows developers to refactor code with confidence, knowing that the tests will catch any breaking changes. Test-driven development is a specific method of software design that has proven effective, and goes hand in hand with the ideas of *continuous integration* and *continuous delivery*, which enable developers to rapidly iterate on new features and bug fixes. PsychoAnalyze utilized GitHub Actions and pre-commit hooks to perform many of these sorts of tasks in an automated fashion.

- **Poetry.** Modern package managers such as *Poetry* allow the software to integrate with more sophisticated and customizable third-party packages while performing dependency resolution and virtual environment management. They also make it easier to publish and distribute the package on PyPI.

- **Documentation.** The documentation for PsychoAnalyze is written in simple Markdown, but is translated into beautiful HTML via the *MkDocs* static-site generator and the mkdocs-material theme. The documentation is hosted on GitHub Pages and routed to the custom domain `https://psychoanalyze.io` via a CNAME file in the repository. Finally, API documentation is automatically configured using the *mkdocstrings* extension, which parses the *docstrings* in the Python code and injects them into the documentation.

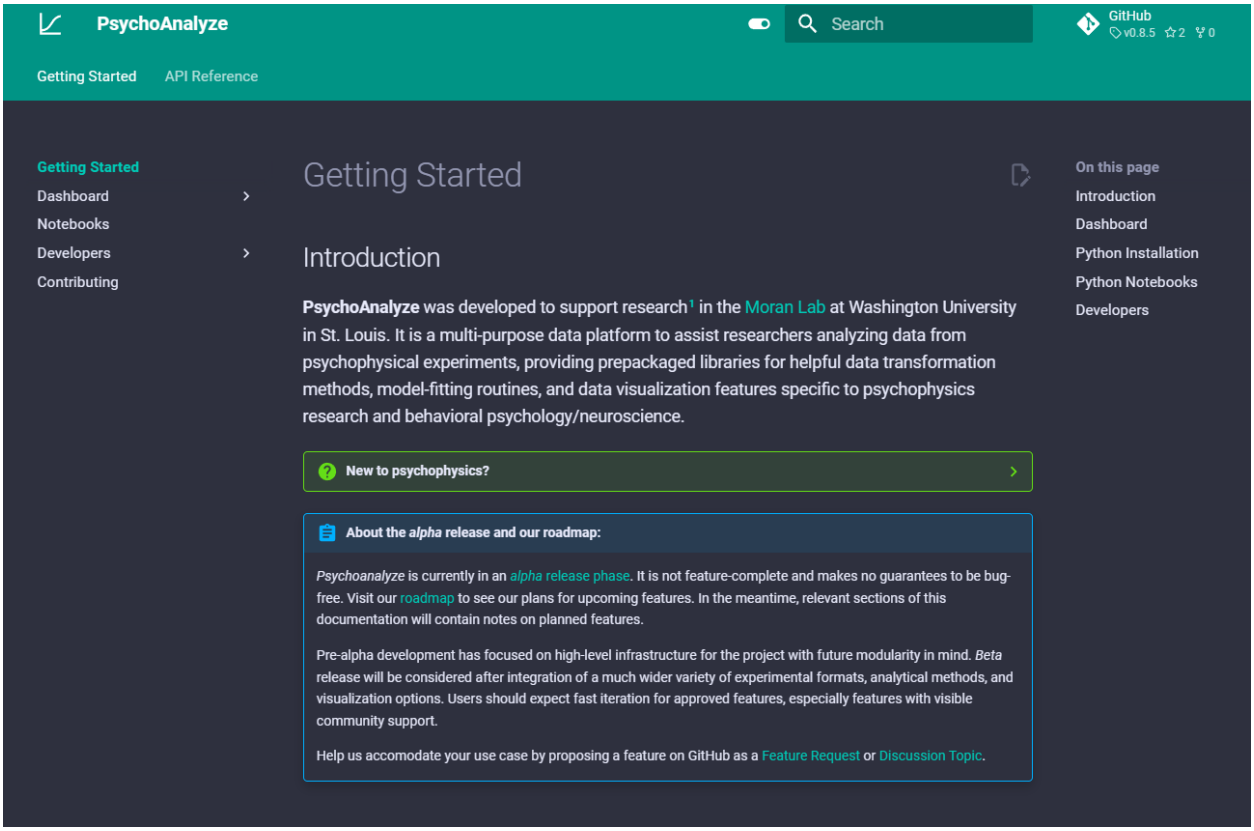

Figure 4.2: Screenshot of PsychoAnalyze documentation.

4.4 Future Direction and Roadmap

The release of PsychoAnalyze corresponding to this submission is labeled as an *alpha* release (tagged v1.0.0), primarily conveying that the software is not feature complete. Software development in the *pre-alpha* phase focused on architecture and extensibility over supporting a wide range of customizable options from the start, although the ability to incorporate such options was factored into project architecture. For example, care has been taken to incorporate abstractions for the *link function* such that other link functions may easily be substituted for the logistic function. Integration of developer tooling and careful project architecture were priorities, with the intention of enabling rapid iteration of features in the next release phase. Development will likely prioritize "plug-in" modules for the software packages in the previous *Comparable Work* section as well as more advanced features outlined in the project roadmap on GitHub.

PsychoAnalyze was developed in support of the research performed in Chapter 3 [87]. Experience from these experiments, in addition to feedback and requests from the community, will inform the next iteration of development.

Chapter 5

Vagus Nerve Stimulation, Fear Extinction, and Post-Traumatic Stress Disorder

5.1 Introduction

5.1.1 Vagus Nerve Stimulation (VNS)

The majority of discussion in the previous chapters was focused on peripheral nerve stimulation in the context of the somatosensory system. Most sensorimotor functions in the human nervous system are mediated by *spinal nerves*, which are the majority subset of peripheral nerves and interface with the spinal cord. Additionally, the human body contains twelve nerves generally designated as *cranial nerves*, which connect peripheral organs directly to structures in the brain. Cranial nerves are responsible for a variety of functions including sensorimotor and autonomic functions. The "tenth" cranial nerve (Cranial Nerve X), is more commonly referred to as the *vagus nerve*. It is the longest cranial nerve and transmits feedback signals for a variety of autonomic functions including heart rate, digestion, and respiration. It provides one of the primary information channels from visceral organs to the brain and is composed mostly of afferent fibers[20].

Many of the principles and techniques of peripheral nerve stimulation in the sensorimotor system are directly translatable to *vagus nerve stimulation* (VNS). Originally, VNS was administered as a treatment for intractable epilepsy, but clinicians noticed beneficial side effects for depression in some cohorts. After some years of further research, VNS was approved by the FDA for the treatment of depression in 2005. The last couple of decades in

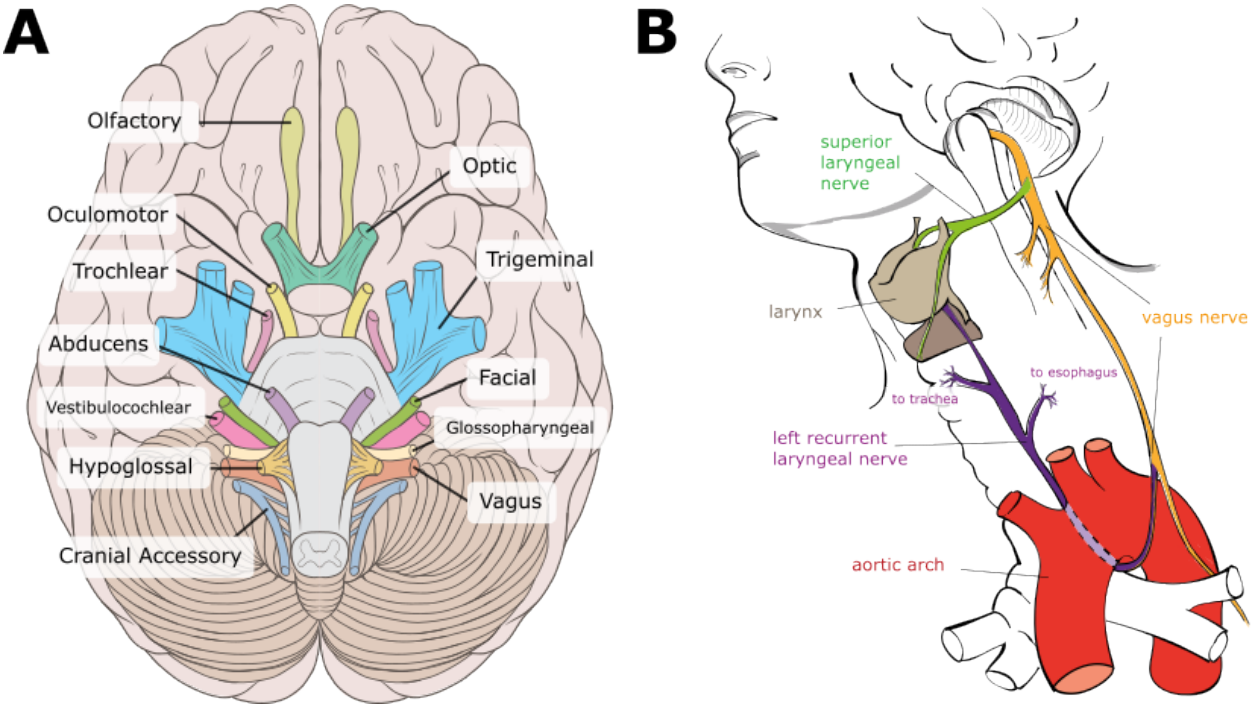

Figure 5.1: **A.** A view of the cranial nerves from the ventral side of the brain. **B.** Diagram of the vagus nerve against other thoracic features. Images are from the public domain and are reprinted under the *Creative Common License* (see Acknowledgements).

VNS research have also proposed some interesting applications for other anxiety and mood regulation disorders such as PTSD [26][94].

5.1.2 Parameterization of the Electrical Stimulus Delivered to the Vagus Nerve

From a clinical standpoint, the course of action for VNS therapy is quite broad, in large part due to a lack of understanding regarding how the parameterization and timing of the delivered electrical stimulus affects downstream neural processes during VNS. Commonly, a cuff electrode is placed around the vagus nerve, and the electrode is connected to a pulse generator that is implanted in the chest. The pulse generator is programmed to deliver a fixed stimulation pattern, typically for long durations of seconds to minutes. The device is typically implanted in the left chest and the leads are wrapped around the left vagus nerve to minimize cardiac side effects. Generally, a limited set of stimulation parameters are programmed for treatment. Often, higher stimulus intensities are administered when more aggressive treament is desired, at the cost of inducing side effects such tingling or pain in the neck, hoarseness, coughing, or cardiac arrhythmias. While some general conclusions can be drawn towards the extremes of the clinically-relevant stimulus parameter space, there is a broad lack of understanding how the stimulus parametrization might have differential effects on various potential neural pathways according to the functional topography of the vagus nerve [45].

5.1.3 The Fear Extinction Model

The primary hypothesis behind PTSD-motivated VNS research is that VNS boosts *fear extinction* processes, perhaps by facilitating the consolidation of replacement memories via changes to global state such as attention or criticality. *Extinction* is the process of disassociating a conditioned stimulus (CS) with an unconditioned stimulus (US). Extinction displays distinct mechanistic differences from other forms of classical conditioning, particularly in ways that CS-US disassociation is context-dependent or time-bound, weakening the effectiveness of targeted extinction treatments when the goal is complete or near-complete

disassociation[68]. Thus, strengthening the fear extinction process during targeted behavioral therapy, such as exposure therapy, is a key aim of VNS research in the context of PTSD treatment.

Exposure to a particularly traumatic event causes the formation of emotional memories associated with stimuli surrounding the event. In some circumstances, these memories are easily inhibited by extinction processes, and the traumatic event no longer causes severe emotional distress. However, in PTSD, the traumatic event causes the formation of emotional memories, but these emotional memories are not extinguished over time. VNS affects many brain regions that are involved in the formation and extinction of emotional memories, and thus may be able to help regulate emotional memories in PTSD patients.

A wealth of research is available on Pavlovian fear extinction models and the corresponding neural dynamics in various animal models including rodents, non-human primates, and humans[111, 16]. These experiments have implicated the basolateral complex of the amygdala (BLA), the ventromedial prefrontal cortex (vmPFC), and the hippocampus, among other neural structures, as key contributors to the neural circuits involved in fear extinction[57, 19, 2]. These conditioning experiments have examined how patterns of conditioning develop in these structure during association or disassociation and have begun delineating differential roles for the various brain structures in the conditioning process. Importantly, the rate of conditioning might be modulated by various models of nervous system state such as excitability, criticality, attention, or arousal.

While there are many experiments that elucidate the role of various neural structures and behavioral processes involved in fear extinction, we can focus our discussion in the context of our research towards a series of experiments from University of Texas-Dallas researchers where this model of VNS action has been extensively examined in the rodent model[70, 92, 93], and for the most part has validated conventional hypotheses about VNS and fear extinction. These experiments applied straightforward conditioning protocols and assessed corresponding processes at various levels of the nerobiological hierarchy, establishing important consequences of input parameterization in the context of VNS and fear extinction. While many of these studies were performed after our research took place, they provide helpful context regarding the aims and potential outcomes of our study.

Work in Dr. Monosov's lab at Washington University in St Louis has established mature experimental protocols and laboratory infrastructure for examining neural mechanisms of

conditioning. Generally, macaque monkeys in the Monosov lab are well-trained in visuotemporal learning tasks administered and monitored via eye-tracking software, and the lab incorporates single-unit studies to asses various mechanisms of neural models, including those involving deep brain structures. We thus initiated a pilot study for investigating mechanisms of associative learning during VNS in the macaque model. While we only analyzed behavioral outcomes of the experiment, omitting neural recordings, we hoped that this exploratory study would provide greater understanding, context, and practical experience for future studies that might incorporate such data.

We executed a brief pilot study wherin we administered basic trace association conditioning protocols and monitored eye movements and pupil behavior during a oculomotor task with a macaque monkey. In contrast with our upper limb experiments which utilized *operant* conditioning, *trace* conditioning does not rely on "correct" responses from the subject in order to deliver reward and thus develop the conditioning association, it simply does so regardless of any action from the subject. We applied vagus nerve stimulation and compared outcomes between treatment and sham blocks to investigate effects of VNS on associative learning. Our protocol closely resembles experiments performed in Peck, Peck and Salzman (2014) [74]. They presented basic visual cues, structuring stimulus presentation in a such a way that enables the disentanglement of the roles of various brain structures during associative learning (such as whether or not the location of the stimulus on the screen, or whether trace or operant conditioning was performed, has an effect on learning rates and outcomes). They examined the affects of these experimental conditions on neural firing in the amygdala and gaze patterns such as the amount of time that the subject spent fixating on the unconditioned stimulus. Because the amygdala is both heavily involved in fear extinction and is a major structure activated by VNS, we examined a similar model of visual association in the macaque during VNS.

5.2 Methods

5.2.1 Ethical Treatment of Animals

The Institutional Animal Care and Use Committee approved the experimental paradigm design, surgical procedures, neurophysiological stimulations, and daily animal care, following

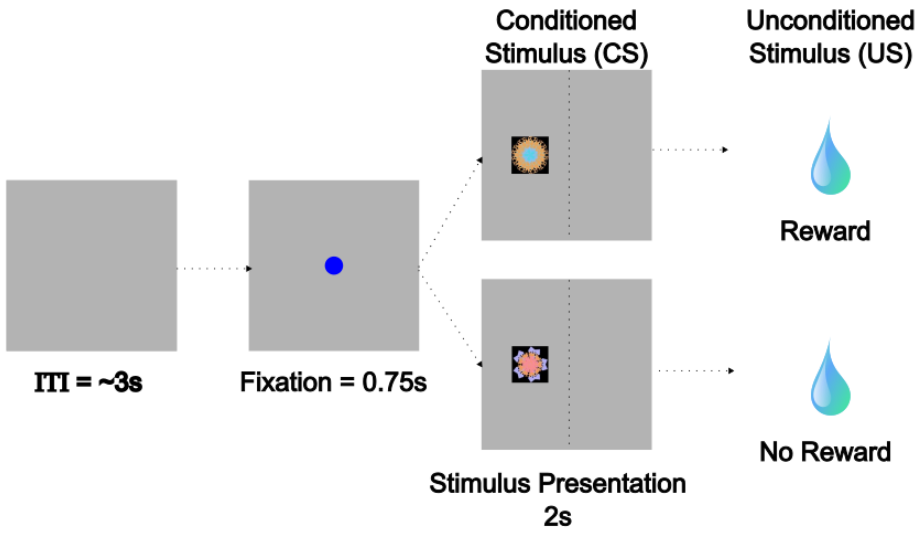

Figure 5.2: After a brief inter-trial interval (ITI), a fixation point was presented to the monkey. After the monkey fixated on the point for 0.75s, a conditioned stimulus (CS) was presented for 2s. Then, the stimulus was removed and a liquid reward was presented as the unconditioned stimulus (US), or was not presented depending on the fractal pattern for the trial. The figure shows two selected examples out of several possible stimulus configurations that include the location of the stimulus, the familiarity of the fractal pattern, or whether or not fixation on the fractal was required for a reward (i.e. in operant conditioning). However, our analysis focused on the simple case of novel vs familiar stimuli during trace conditioning.

all guidelines set by the Association for Assessment and Accreditation of Laboratory Animal Care and the Society for Neuroscience. We sedated animals with ketamine, intubated, and anesthetized with isoflurane before and during surgery under the supervision of Department of Comparative Medicine veterinary staff at Washington University in St Louis.

5.2.2 Stimulus Conditioning

We continuously measured the gaze of a macaque monkey using eye-tracking software (SR Research, Ottowa, Canada) while administering trace conditioning on visual cues. Figure 5.2 conveys two simple example trial configurations out of several combinations of stimulus configurations that were administered. The monkey was trained to fixate on a central point on the screen for 0.75s, at which point a visual stimulus in the form of a fractal pattern was

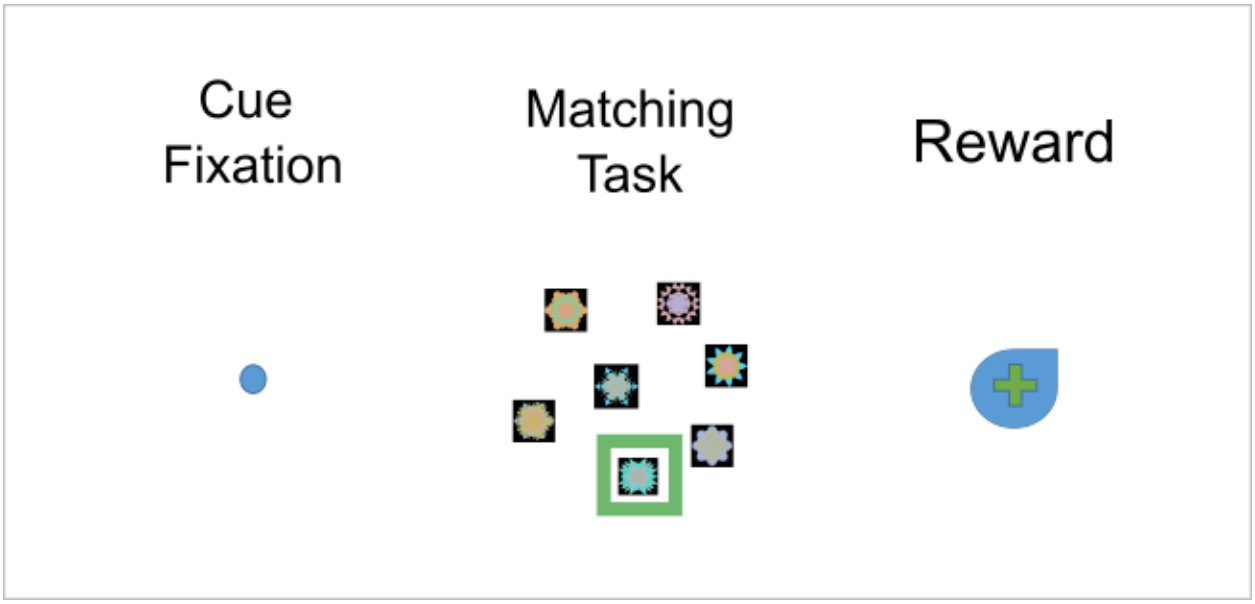

Figure 5.3:
An example of an alternative stimulus presentation in the conditioning experiment. One familiar stimulus was presented alongside several unfamiliar stimuli. This configuration might require increased attentiveness to learn the CS and anticipate the US; this may provide evidence for or against theories of VNS mechanisms that facilitate attention during conditioning and/or extinction.

presented on the screen. The stimulus was chosen from a small set of randomly-generated isoluminant fractal patterns, and trials were administered in blocks such that the monkeys were presented with controlled portions of novel and learned stimuli. After 2 seconds of presentation of the fractal, an unconditioned stimulus was delievered in the form of liquid reward. This process was repeated over several experimental sessions with various experimental protocols, including brief forays into alternate protocols such as those involving aversive stimuli (which were delivered via air puff to the eye area), operant conditioning protocols (which required target fixation for reward), and "matching" tasks where several stimuli were presented simultaneously (e.g. Figure 5.3). Data for these alternate protocols are not included in this analysis but may be available upon request.

5.2.3 Surgical Procedure

Once the monkey had been trained in the conditioning task, an experienced neurosurgeon implanted a VNS device (LivaNova, London, UK) in the monkey's neck over the left vagus

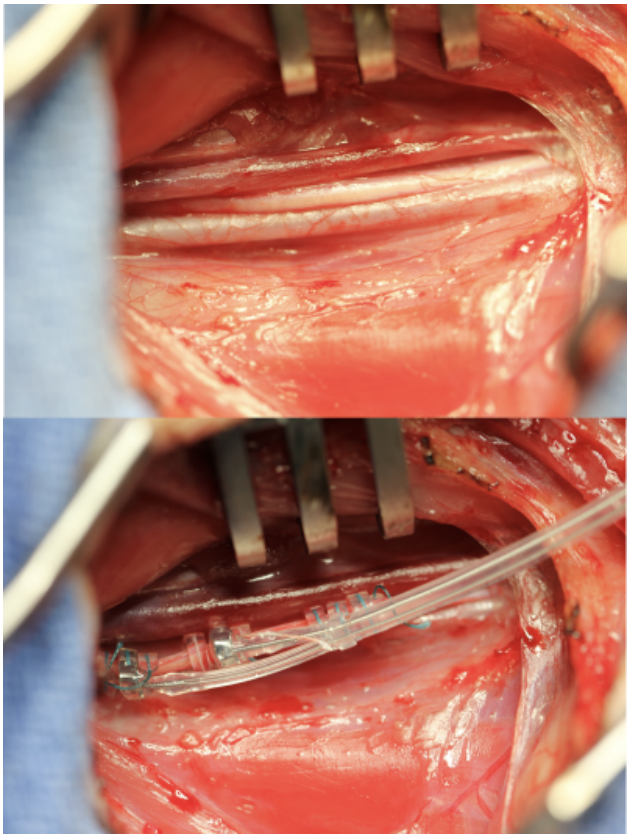

Figure 5.4: Photographs of the vagus nerve before and after implant placement.

nerve. The electrodes were cuff-style electrodes 2-mm in diameter and with 5 mm interelectrode spacing between the proximal and distal electrode sites. Images in Figure 5.4 show the site of the implant before and after the device's placement over the vagus nerve.

5.2.4 Delivering VNS During Learning and Extinction

After the monkey had recovered from surgery, we conditioned the monkey on a new fractal set in week 1 and subsequently performed extinction on half of the conditioned stimuli (Figure 5.3). In week 2 we repeated the process but administered VNS for half of the fractal set .

Experiments were ended after two weeks when impedance readings on the interface suggested implant failure, likely due to poor/severed connections from the wired leads that connected the implant to the stimulation device.

5.2.5 Analysis

Data was analyzed using eye tracking measurements that recorded the x and y positions of the monkey's focal point, sampled at 1000 Hz. We derived a "time fixating" measure that represents the amount of time that the monkey was fixated on the fractal during CS presentation, measured by counting the samples where the focal point occupied the 2D space on the screen where the fractal was located during the corresponding time bin. We also calculated the "response time" as the time between CS onset and first fixation. Other measures such as those from pupillometry data (Figure 5.7) and reaction times were performed but were not a focus of the study; the cursory analysis performed on data from these particular sessions was inconclusive.

Similarly to previous chapters, the analysis code was written in Python and is available on GitHub; the code can be run as an interactive dashboard where one may interactively browse the results[9]. My colleague Kara Donovan significantly contributed to the data processing performed and the analyses presented here.

5.3 Results

An example of data acquired from eye-tracking measurements in a single trial be seen in Figure 5.5. We calculated a "proportion fixating" measure similar to Peck et al. [74] and assessed the measure for each grouping of stimulus type: *novel* fractals that represented a compltely new fractal pattern, *novel learning* fractals that were new for an experimental session but were repeatedly presented over the course of the session, *familiar* fractals that were learned in a previous session, and *familiar learning* fractals that represent the extinguishing scenario. The box-and-whisker plots in Figure 5.6 convey the distributions of the *time fixating* and *reaction time* measures. T-tests performed between equivalent groups from before and after VNS administration yielded insignificant results.

Analysis of looking and blinking behavior during sessions that contained aversive stimuli demonstrated clear conditioned response patterns, however, these results are of limited utility without further investigation. One interesting finding was the nonmonotonicity in the

[9] https://github.com/schlich/vns

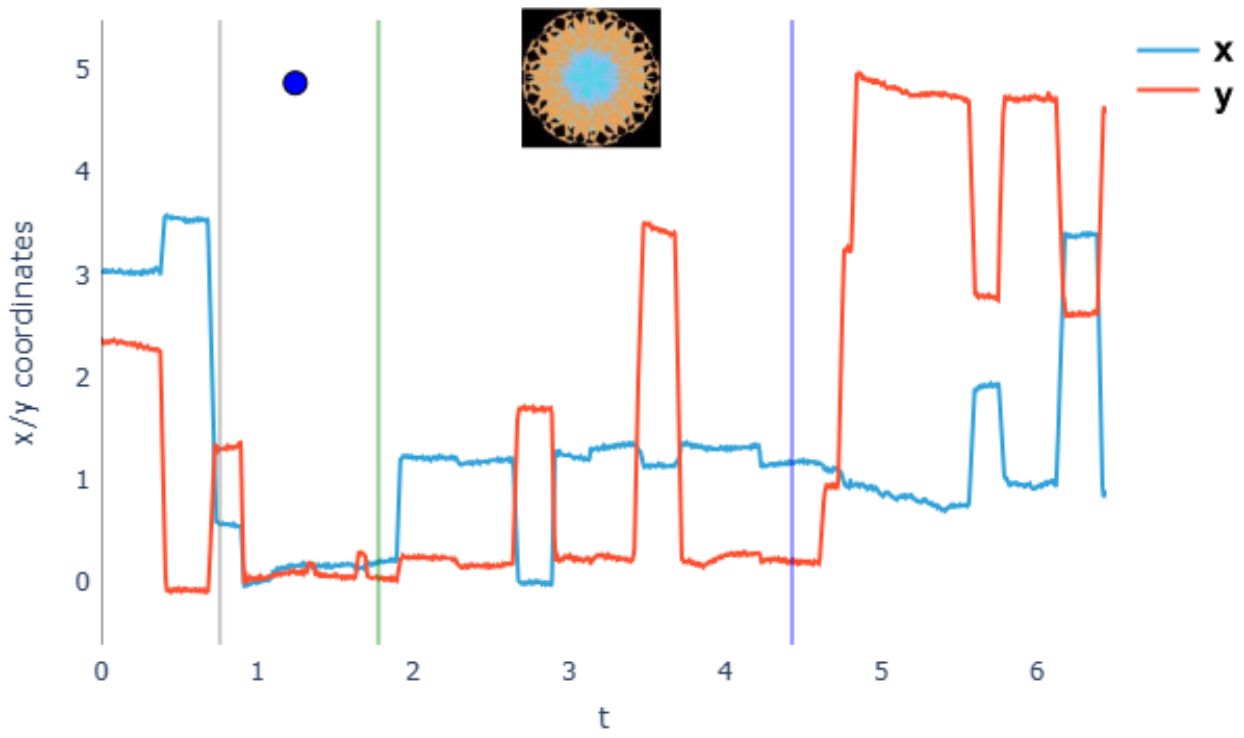

Figure 5.5:
A selected example of eye tracking measurements during a single trial; several of these traces were aggregated to serve as conditioning measures for our analysis. Fixation can clearly be seen at the point (0,0) during the fixation period (blue dot) and at the point (1,0) during CS presentation, coinciding with the center of the fractal image. The "time fixating" measure was derived from measurement samples contained in the 1x1 square (dimensionless units) that the fractal occupied.

conditioned response with respect to the stimulus "value;" both appetitive and aversive stimuli elicited stronger conditioned responses than neutrally-coded stimuli 5.8.

5.4 Discussion

5.4.1 Retrospective

While our pilot study did not yield actionable results, subsequent studies from other researchers administered in the time since this experiment took place have continued to justify the inquiry of fear extinction processes as it relates to vagus nerve stimulation and potential avenues of PTSD treatment. I am personally currently unaware of literature that establishes the nonmonotonic pattern of the conditioned responses with respect to aversive and appetitive stimuli (Fig 5.8), but this finding may warrant further investigation if it is a novel or controversial finding. Similarly to the experiments in the Moran lab with (spinal) peripheral nerve stimulation, this was a completely new line of study for all collaborators involved. Many practical obstacles were learned and encountered in the course of this pilot study. Future researchers are encouraged to formulate clear causal hypotheses from the outset; when employed in combination with a Bayesian approach to experimental design and more sophisticated data management practices, this line of research still has tremendous potential to provide breakthroughs in our mechanistic understanding of well-established neural processes such as classical conditioning and inform approaches that may improve the outcomes of VNS administered in the clinic. Speculation about the nature of such methodological improvements is the focus of the remainder of this dissertation.

5.4.2 Status of VNS as a Clinical Intervention for PTSD

While initial applications of VNS in a fear-extinction paradigm have shown positive results in laboratory settings for a range of species from rodents to humans, researchers have ultimamtely had difficulty translating positive results in animal models into clinical methodologies that reliably produce positive and maintainable results. Notably, experiments from non-invasive *transcutaneous* vagus nerve stimulation have returned limited results. Also notably, there has been progress with other methods of treating therapy-resistant anxiety

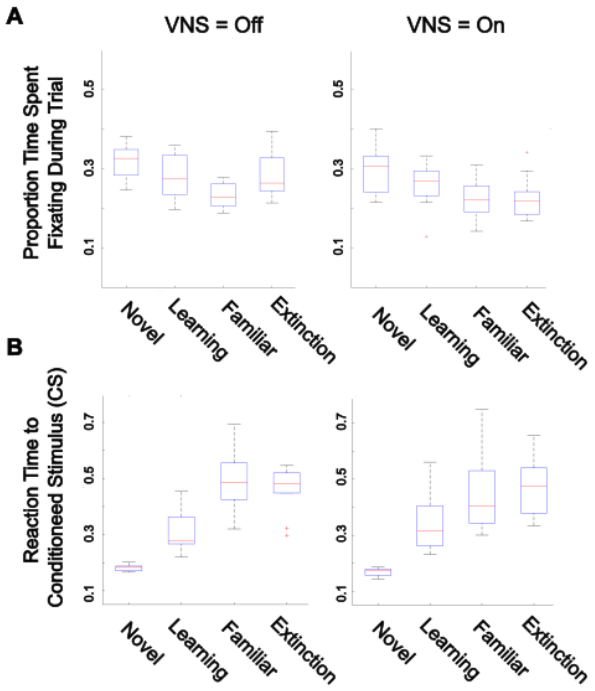

Figure 5.6: *Distributions of Conditioning Metrics across trials.* **A.** We calculated a "proportion fixating" measure from the fractiion of time the monkey gazed at the conditioned stimulus (CS) during its presentation. We assessed the measure for each grouping of stimulus type: *novel* fractals that represented a compltely new fractal pattern, *learning* fractals that were new for an experimental session but were repeatedly presented over the course of the session, *familiar* fractals that were learned in a previous session, and *extinction* fractals that represent the fractals which replaced the Conditioned Stimulus, thereby inhibiting it and causing extinction. Figure 5.6 conveys the distributions of the *time fixating* and *reaction time* measures. T-tests performed between equivalent groups from before and after VNS administration yielded insignificant results.

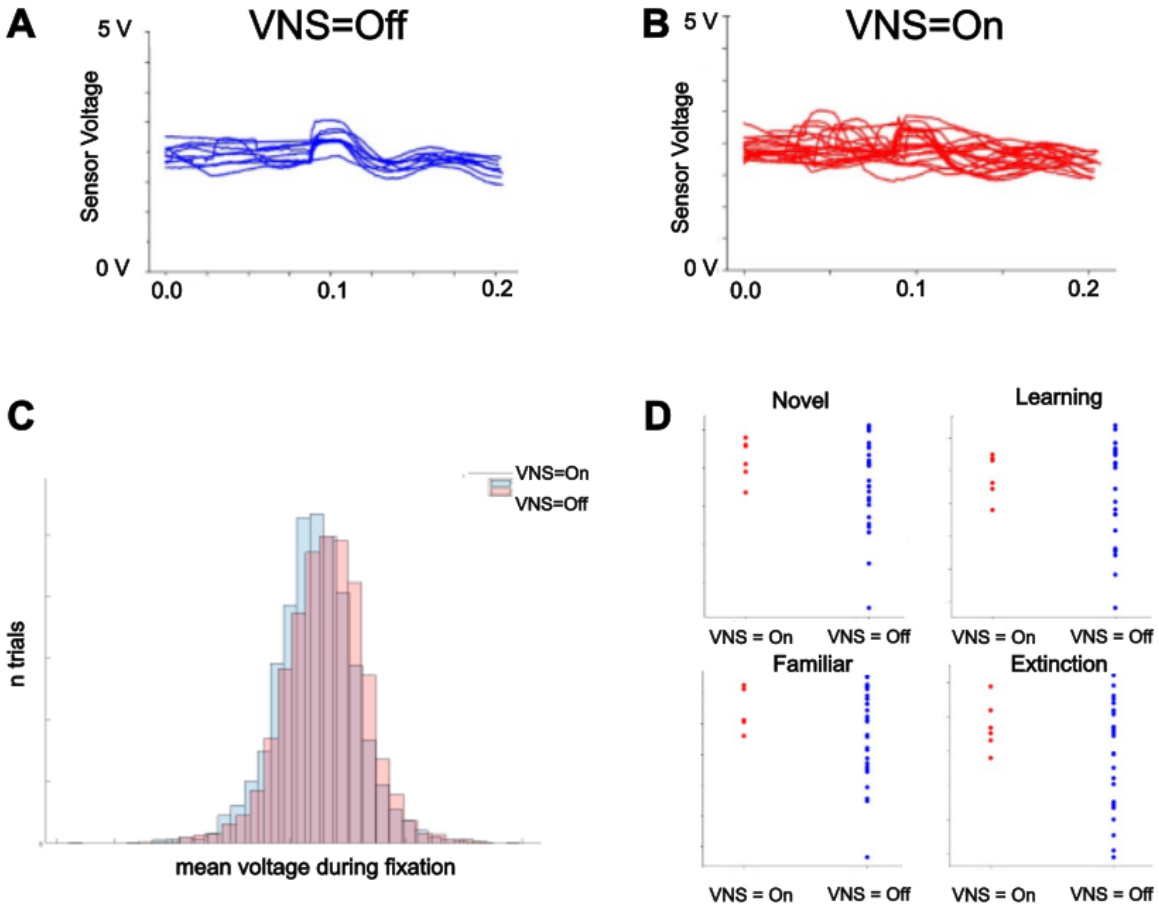

Figure 5.7: Pupillometry data from conditioning experiment. Panels **A** and **B** show pupillometry measurements during conditioned stimulus (CS) presentation in VNS and no-VNS conditions. Measurements from the eye tracking device were recorded as the raw voltage of the input read by the sensor, and thus were not easily interpretable. While the VNS=Off condition seems more patterned, we did not statistically test directly on the traces in the panel. Instead, we calculated the time spent fixating during CS presentation for each trial and took the mean across all trials; a histogram of these means are presented in **C**. In **D**, we examine the spread of these means across sessions, comparing each of the four trial types. No statistical tests showed significant p values from t-tests performed between VNS and no VNS conditions.

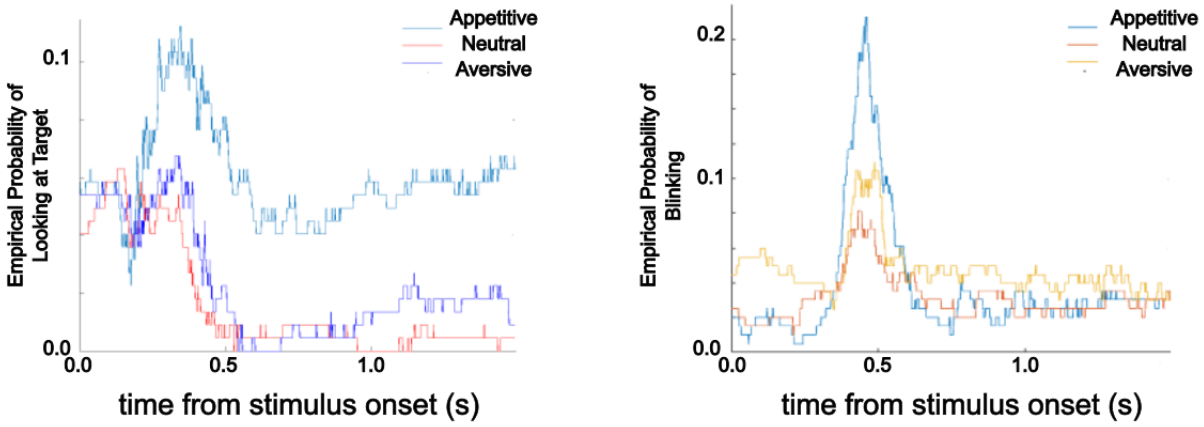

Figure 5.8: *Behavioral Measurements from an Appetitive/Aversive Task.* For some sessions, monkeys were presented with an air puff to the eye to obtain responses to an *aversive* unconditioned stimulus. Data shows clear differences in behavioral outcomes; the monkey clearly spent more time looking at appetitively-conditioned stimuli (**A**) and spent more time blinking for aversely-conditioned stimuli (**B**). Interestingly, blinking behavior shows a non-monotonic relationship to stimulus "value."

disorders such as neurochemical approaches and transcranial magnetic stimulation. Researchers should seek a more comprehensive mechanistic understanding of VNS action and seek to apply more targeted treatment with a fuller understanding of how stimulus parameterization modulates the effects of VNS. This level of scientific understanding seems like a crucial requirement to justify the continued exploration of VNS for PTSD treatment, but these goals are still eminently obtainable. The evidence is clear that vagus nerve stimulation modulates activity in key areas of interest in a graded and optimizable manner.

More specifically, clinical application of VNS for anxiety disorders should understand how the timing, intensity, and characteristics of the electrical stimulus from the VNS implant maps to the desired clinical effects as well as undesired side effects of the treatment (including but not limited to discomfort, hoarseness in the voice/throat, or cardiac arrhythmias). Our pilot experiment performed a very coarse and rudimentary parameterization of VNS delivery that has been utilized in clinical trials for depression; this parameterization, which involve long, continuous pulse trains, could have been adequate enough to yield results, but alternatively might have precluded any significant results from the study, even if all other methods were adequately executed. While calibration of device stimulation parameters for a majority of VNS studies thus far have been coarse, progress continues to be made – one recent study

implements a Bayesian approach to estimate optimal parameterization of VNS in a porcine model of heart failure, and the results and methodology provide a convincing path forward in this area[110]. It stands to reason that more targeted stimulation of the vagus nerve can improve clinical results and provide more precise data on downstream activated pathways during treatment. Studies to this aim are already well under way [14][116]. One particular challenge that may be faced when it comes to paramater optimization is that downstream effects of interest during VNS likely develop at a much longer temporal scale than what is commonly modeled with the peripheral nerve in sensorimotor circuits, which generally operate on the order of milliseconds.

In summary, we can apply many similar PNI techniques developed in the somatosensory system to study the effects of VNS [44]. Modeling studies have assessed stimulation methods designed to target clinically-relevant fibers in the vagus nerve [4]. In the Monosov Lab, we used highly trained monkeys to examine macaque monkeys' performance during an extinction task to see if we could observe any effects of VNS on the monkeys' performance. Results from this study might inform future studies that include single-unit nerve recordings to draw further conclusions from the experimental paradigm[76].

Chapter 6

Discussion: Towards a Performant, Integrated Model of Peripheral Nerve Encoding

6.1 Introduction

The previous chapters outlined two experimental paradigms for peripheral nerve stimulation in macaques. The first focused on sensory stimulation in a longitudinal study, while the second focused on vagus nerve stimulation in a shorter study. Although these exploratory experiments were moderately successful in achieving their objectives, they began without clearly defined hypotheses or expected outcomes related to specific, significant questions about the underlying physiology, which might have enabled more targeted experimental design. Thus, I will discuss the importance of computational modeling when it comes to experimental ideation and design, especially in situations where success is critical and resources are expensive. Modern computational models, and particularly, as I will argue, Bayesian computational models, have the ability to build and tune neuroscientific models, particularly peripheral nerve models, to additional levels of sophistication and complexity.

In addition to the above, I provide a basic overview of the most cutting-edge and effective data processing tools available to data practitioners and neuroengineers in 2024. Ideally, this overview provides readers with an adequate understanding of the landscape and the ability to make effective tooling choices when sophisticated data management and processing techniques are warranted.

To conclude this chapter and this dissertation I will outline a very specific proposal for a Bayesian model workflow. When implemented, this workflow may provide additional

insight into the data sets we have collected and provide context for the design of effective experimental protocols for future iterations of this research.

Interested readers may want to monitor the PsychoAnalyze GitHub organization page for future attempts at developing the ideas and proposals in this chapter into actual code. In the context of the neuroengineering community, several researchers have already made significant progress in executing many different components of Bayesian analysis in a wide range of neuroscientific and psychophysical studies, including VNS studies, peripheral nerve studies, studies for neuroprosthetic control, etc. Taken together, these developments mark an exciting period for neuroengineering research.

6.1.1 Modern Data Practices and the Bayesian Paradigm

Modern data processing software and techniques, undoubtedly accelerated by developments in AI technology and the widespread availability of large datasets, can make complex scientific and statistical analyses both accessible and enjoyable. The significant increases in raw computational power available to researchers in recent decade have unlocked a whole new set of tools and research paths.

One set of these tools are advanced simulation methods that have been developed for Bayesian computation. Bayesian methods may provide guidance for inexperienced and experienced researchers alike who may be asking fundamental question about experimental design and statistical power such as:

- "How many trials are 'good enough' for a session?"
- "How many sessions or subjects are enough for my experiment?"
- "How can I minimize the costs and duration of my experiments to obtain reasonable estimates of the effects of my experimental manipulation?"
- "Which experimental interventions are most likely to uncover impactful results?"
- "How can I evaluate and compare my favorite models of sensory coding?"

Bayesian methods are particularly apt at answering these types of questions. Before going into how this is the case, let's examine how computational peripheral nerve models are currently built.

6.1.2 The Hybrid Model of the Peripheral Nerve and its Limitations

The *hybrid model* is a generally popular and effective method of modeling the peripheral nerve [81]. In essence, the hybrid model combines an electrostatic *finite element model* of the nerve, governed by fundamental equations from calculus and physics, with a *neuron compartment model* that simulates the nerve's response to stimulation. The finite element method is used to determine the distribution of electric potential in the nerve, while the neuron cable/compartment model is used to determine the recruitment of fibers in the nerve based on the simulation of models that govern neuronal firing such as the famous Hodgkin-Huxley model. Historically, these models are very expensive to compute and have very loose applicability to biological systems that are many orders of magnitude more complex. This presents a problem when, in order to create tighter sensory feedback loops delivered via a PNI, we need to perform predictive computations in near-real time. A detailed and performant model must seek alternatives or improvements to the hybrid model.

Implementation considerations of these two components of the hybrid model are discussed in the next two sections.

6.1.3 Finite Element Modeling

The finite element method (FEM) is a numerical method for solving partial differential equations when analytical solutions are not readily available It is commonly used in engineering to solve problems in heat transfer, fluid dynamics, and electromagnetism. In the context of peripheral nerve stimulation, the FEM is commonly used to solve the Laplace equation for the cross-sectional field electric potential in the nerve resulting from applied electrical stimuli. Here are some important components/parameters of FEM that researchers are tasked with specifying, in the context of some of the limitations and trade-offs encountered during model building:

- **Geometry**: Although non-invasive, *in-vivo* imaging of the peripheral nerve is an active area of research [55], usually the anatomy of the peripheral nerve is not known *a priori*, despite anatomy's sizeable impact on the function of the interface. Performant peripheral nerve interfaces should be able to rapidly adapt to individual differences in anatomy. Researchers must make decisions about representations of both the interface and the nerve itself that provide "good-enough" insights given the limited computational resources available during research and/or the computational resources available to the engineered devices. Aspects of peripheral nerve anatomy that are most fundamental to the finite element model include the location, size, and shape of the implant, the nerve, and the fascicles within; usually these are simplified to simple gemotric shapes that make both modeling and computation easier. Histological samples may be used to assess peripheral nerve anatomy with more accurate geometry, however these samples are often obtained posthumously, thus information obtained from these samples would only be applicable in a probabilistic sense.

- **Boundary conditions**: One of the necessary tradeoffs of using a discrete solution to the Laplace equation instead of an analytical one, as modeling software allows us to do, is that the *boundary condition*, i.e., the electric potential at the boundary of the model, must be specified (in analytical solutions it is modeled as V=0 (ground) at infinity). The boundary for peripheral nerves is usually represented with some volume of area around the implant with a saline medium. Generally, the larger the boundary, the more accurate the model, but the more computationally expensive it is to solve.

- **Material properties**: Both the nerve and the implant have inherent electrical properties that partially govern the course of charge through the medium. These properties are usually represented with constants from the literature derived from a combination of model and experimentation.

- **Initial conditions**: Generally, the *initial condition* of the model is needed to perform the computation. In peripheral nerve models, this generally means the current level delivered by the implant electrode channel, including which channels are active in multi-channel implants.

  Almost all FEM models in peripheral nerve research represent *static* snapshots of the system, but temporal dynamics are an important component of peripheral nerve signaling that cannot be completely ignored.

6.1.4 Neuron Compartment Model

The *neuron compartment model* is a numerical method for simulating the response of a neuron to an electric field based on sodium-conductance equations found in the literature. It is commonly used in neuroscience to simulate the response of a neuron to an electric field. In the context of peripheral nerve stimulation, the compartment model is used to determine the recruitment of fibers in the nerve based on the simulation of the model. Here are two common specifications of the model that must be chosen by the researcher:

- **Governing equations**: Generally, a set of differential equations describes the electrochemistry of the neuronal membrane,derived from properties of ion channels in the membrane and physical equations governing electrical circuits and diffusion. The McIntyre-Richardson-Grill (MRG) model [58] is commonly cited as a standard.
- **Distribution of fiber types**: There are some fundamental properties of nerve fibers that make them more or less prone to excitation by an electric field. These properties include the fiber's diameter, myelination, and the distance of Nodes of Ranvier from the electrode site. As with in the finite element model, the distribution of these properties within the nerve is not known *a priori*, but can be estimated from the literature.

6.2 Bayesian Methods and their Applications in Neural Engineering

The idea of applying Bayesian methods and principles to neuroscience and neural engineering is certainly nothing new. As I explain further, their usage is actually quite widespread, and yet still there are entire domains where there is tremendous headroom for further development.

6.2.1 Psychometric Curve Fitting

Many researchers, especially those familiar with psychophysical protocols, are likely already familiar with an extremely common Bayesian method: utilization of a generalized linear

model (GLM) with a sigmoidal link function to perform linear regression on datasets from two-alternative experiment designs for estimation of the *psychometric curve* (specifically, its parameters, which include *threshold* and *slope*). We used this technique to provide estimates of sensory threshold in Chapter 3, and there is quite sophisticated tooling available for this purpose, including the modest implementation described in Chapter 4 which currently borrows its curve-fitting procedure from the *psignifit* library.

However, I suspect that the majority of researchers that use Bayesian tools for curve-fitting treat these tools as so-called *black boxes* that spit out parameter estimates and don't have significant consequences that differentiate it from traditional regression models and similarly simple curve-fitting approaches. This is certainly not a shortcoming of the researchers who use these tools (I, myself, did not understand the mechanics of Bayesian models until well after I was using them to fit curves, and arguably still don't). However, a lack of appropriate emphasis on the Bayesian paradigm within the field has shielded us from understanding the underutilized features of these methods.

The authors of the *Palamedes Toolbox* have compiled excellent resources to supplement their implementation of Bayesian inferencing for psychometric curve-fitting. I recommend these resources to readers who are interested in a practical introduction to the theory behind the methods, although some familiarity with Bayesian statistics in general might be advisable[10]. Importantly, *Palamedes* is the only tool to my knowledge that comes with a turn-key implementation for *hierarchical* Bayesian modeling, which can be insufficiently described as a statistically "conscious" method for pooling psychometric curve fits in experiments with multiple subjects or cohorts. Notably, the publication accompanying the *Palamedes* implementation of hierarchical modeling was only recently published in 2024, further demonstrating the broad-but-shallow application of Bayesian methods in neuroscience as well as the green pastures ahead[11]. For better or for worse, *Palamedes* is MATLAB-only software and thus is limited in all of the usual ways that MATLAB programs are limited in comparison to implementations in general-purpose programming languages such as Python.

At this point, we can begin to formulate a response to the class of questions posed at the beginning of this chapter which relate to experimental design, particularly the questions

[10]For complete Bayesian novices, I recommend *Statistical Rethinking* by Richard McElreath and its accompanying YouTube series.

[11]Regarding the "broad" aspect of this statement: a canonical work on Bayesian curve fitting in psychophysics by Wichmann and Hill (2001)[112], has over 2700 citations as indexed by *Google Scholar*

about sample size (whether in terms of number of trials or number of participants). Researchers may be familiar with the frequentist "power analysis", but the Bayesian approach to these questions operates on fundamentally different assumptions (see *Kruschke and Liddell* (2017)[49]). Because a hierarchical model can precisely model how probabilistic uncertainty should be incorporated into the model both *within* and *across* subjects, it should improve the model's ability to represent the probabilistic landscape of the input domain. For example, a sampling algorithm might choose a point in the parameter space that maximizes the *"expected"*[12] gain in information, such as in how much each additional trial or participant reduces the *uncertaintly* in the estimation of our parameters of interest, and we may then make an informed decision on the costs and trade-offs of additional sampling.

6.2.2 "Smart" Bayesian Sampling of the Parameter Domain

The next most likely Bayesian method a neuroengineer is likely to encounter in their research might be from the family of *adaptive sampling algorithms*. These algorithms rely on online updates to the model as additional data comes in; after each iteration, the sampling algorithm (referred to in the literature as the *acquisition function*) is able to calculate the point in the parameter domain that is most likely to provide meaningful information that improves our estimates and reduces uncertaintly in estimations. Here at Washington University in St Louis, research out of the Barbour lab has utilized these algorithms to effectively sample parameters online in a non-parametric, audiometric setting. Acquisition function algorithms, perhaps the most notable being the "Expected Improvement" (EI) algorithm, are fairly common in research settings as an alternative to nave grid search and are particularly useful in domains where data acquisition is resource-intensive.

Wernisch et al. [110] demonstrate a modern, sophisticated implementation of Bayesian Optimization (a general term for methods that use acquisition functions and other techniques to "optimally" estimate the parameters given the data) applied to vagus nerve stimulation specifically – in a porcine model of heart failure (yet another application of VNS treatment under active research), they demonstrate not only how Bayesian methods can vastly outperform grid search when attempting to estimate model parameters during active data collection, but they also demonstrate how they used Bayesian simulation models to guide

[12] *Expected* is used here in the probabilistic sense, such as in *"expected value"*.

their experimental work *in-vivo*. This leads us to the following discussion on Bayesian methods that are more rare in neuroscience and psychology but can help researchers design and plan their experiments to be efficient before acquiring a single data point.

Pre-existing algorithmic implementations such as the **Quest+** method[108] might be helpful in detemining how adaptive sampling can be performed online during experiments.

6.2.3 The Open Road Ahead: Novel Applications of Bayesian Methods in Neuroscience

To conclude this overview of Bayesian methods in neuroscience, I will briefly describe potential applications that I have not seen implemented in the literature, primarily:

- **Simulation-based Experimental Design.** The design of experiments based on Bayesian simulations, sometimes referred to as "optimal experimental design" (OED), is a line of inquiry that heavily utilizes simulation-centric Bayesian methods in concert with well-established but computationally expensive models, precisely like those that utilize differential equations such as those discussed in the *hybrid model*. The result is a clear picture of the range of possible outcomes of an experiment in response to modulation of parameters of interest, assisting researchers in avoiding acquisition of uninformative data points and arriving to premature conclusions from early stages of data acquisition that have the potential to hijack the experiment. A good overview of OED, including an application of the model to a biochemical signaling pathway, can be found in Vanlier et al. 2012 [103]. Huan and Marzouk describe how this process can be utilized for nonlinear models in particular[42].

- **Hierarchical, multi-modal/multi-model integration.** Kass and Steffey[48] precisely outline an approach to hierarchical Bayesian modeling designed specifically to address the type of shortcut we were forced to take in our reporting of threshold values in Chapter 3. Namely, at the "top" level(s) of the hierarchy, it is common for researchers, in search of aggregate estimates of a parameter e.g. threshold for a cohort, do so from *point* estimates of experiments from each individual, instead of incorporating estimates of uncertainty at the individual level that ultimately improve the performance of the model. While a remedy to this statistical sleight-of-hand, in

the form of *hierarchical models*, has been established for decades, they are tricky to implement and require detailed knowledge of both statistics and the problem domain. However, as we refine the hierarchical Bayesian model for the peripheral nerve, we can conceptualize a model that actually incoporates several levels of natural hierarchies in neural processing. In no particular order, potential hierarchical "levels" can include: cell receptor type (e.g. Type I vs Type II mechanoreceptors), spinal reflex, anatomical brain regions, single unit responses, LFPs, ECoG/ EEG signals, sensory modality (e.g. sight vs touch), genetics, emotional state, demographics, anatomy, and so on. Obviously, incorporating additional level of hierarchy introduces complexity to the model. If interactions between layers of the hierarchy are poorly specified, introducing layers to the model can increase the chance of non-convergence from the model (meaning the size and topography of the parameter space becomes too large and complex to navigate efficiently). However, as we refine our models and interactions between components of neural processing, we can begin formulating increasingly complex Bayesian models, and instead of introducing costs of additional complexity, the additional layers actually contrain our problem set and our parameter space to make optimization *even more efficient*. One can imagine a model that probabilistically incorporates a large body of scientific knowledge from multiple experimental sources, if scientists can find the motivation to collaborate in such a manner. Similarly to the way that multimodal sensory encoding (e.g. devices that incorporate visual signals in addition to touch/pressure signals) improve encoding models, careful integration of disparate neuroscientific models may push the boundaries of the sophistication of neural models in ways that were previously inconceivable.

- **Time-series/longitudinal models.** Another family of Bayesian models yet to be fully explored in modern neuroscience are those that incorporate non-stationary fluctuations across time. While valiant efforts have been made to identify and quantify such stability and their effects on Bayesian estimates, I have not seen a satisfactory implementation of a Bayesian model which investigates any of several time-dependent influences on neural function, such as learning, plasticity, and other forms of adapation. Such models would clearly be applicable to situations such as what we saw in Chapter 3 with Monkey U's unstable threshold estimates. By incorporating time-gnostic Bayesian models, which are quite developed in the form of models such as Gaussian Processes and Hidden Markov Models, we can investigate causes of variation in model

parameters across time and develop more robust devices that can more easily adapt to or incorporate causes of nonstationarity into the model.

- **Refinement of neural encoding models.** Despite Bayesian models' impressive ability to peform robust model comparisons [29] and to make claims of causality in models [51], these advantages are not commonly implemented. To illustrate an example of how we might apply model comparison techniques, consider that our somatosensory experiments investigated the roles of *pulse width* and *amplitude* in sensory threshold experiments; in light of the fact that these variables' interactions the almost certainly violate assumptions of independence in the detection model, to improve the model, we might use Bayesian model-comparison measures to compare different hypotheses or model equations that specify how amplitude and pulse width each contribute to detection thresholds. As an example of how we might investigate causal models in a Bayesian analysis, consider how we might want to run experiments that help us determine how behavioral learning mechanisms e.g. classical conditioning, might contribute to fluctuations in threshold, as opposed to fluctuations you might see due to healing processes in the nerve, or inflammation tissue forming around the implant electrodes. By designing our experiments under a Bayesian framework we might be able to provide definitive (but probabilistic) answers to these questions about fundamental components of the model, which in effect might lay the foundation for more performant online adaptive procedures capable of executing tight feedback loops.

6.3 Software Tooling

Having provided extensive justification for Bayesian methods in neural engineering research, I now to turn to an examination of software tooling and methods that have enabled, and will continue to enable, more robust and performant computational models including but not limited to the types of models that are described above. Despite the fact that most scientists do not have the capacity to invest in learning modern best practices for software engineering and data management/processing, it becomes more and more essential that the sciences incorporate these tools and processes one way or another as our models and data sets become increasingly complex. This section presents a drive-by overview of critical

categories of software relevant to neuroengineering research[13]. These tools range from tools that are already see widespread usage in academia, tools that have not found widespread adoption but should still be accessible to academic researchers, and tools that likely require domain-specific knowledge in data engineering or software engineering to be able to utilize effectively. However, I refrain from positing which of these categories each tool might belong to. It should be noted that the landscape of software tooling is rapidly changing at the time of this writing, and the specific evaluations listed here will likely become outdated sooner than later.

6.3.1 All-in-One/Integrative Software

Software that integrates functionality from multiple external dependencies for domain-specific functionality. These libraries may or may not be available as installable software apart from the source code.

- **Automated Simulations to Characterize Electrical Nerve Thresholds (ASCENT).** ASCENT was developed out of the Grill lab at Duke University and was validated on a comprehensive example of vagus nerve stimulation. ASCENT integrates many components of peripheral nerve modeling based on the *hybrid model*. It integrates a few external dependencies (such as COMSOL and NEURON) in addition to some internally-developed modules. Key features include segmentation of histological images of peripheral nerve cross-section, parameterized simulation of anatomical characteristics of the nerve (e.g. number of fascicles or distribution of fiber sizes), preset but customizable COMSOL models of real PNI devices, compartment model activation simulation, stimulus waveform generation, FEM solving, and visualization of results. Perhaps the most comprehensive general-purpose "hybrid model" tool available for researchers is the ASCENT tool that serves as a connector between COMSOL and NEURON and helps researchers with common-sense defaults for some parameters and utilities for determining others. For example, it provides default (but customizable)

[13]Unless otherwise noted, I have at least some level of direct experience with the tools I describe in this section, either through academic settings, industry employment, hobby projects, or various trials and attempts at integrating tools with the research outlined in this dissertation. That means that this list is fairly Python-heavy.

material properties for the nerve and the implant, and provides utilities for determining the number and size of the fascicles in the nerve, generating COMSOL models from such parameters. It also provides a convenient interface for running simulations and post-processing the results. It is also free-and-open-source and well-documented, and is likely to continue to be actively maintained and developed. At the time of this writing, the authors of the tool have published two papers demonstrating how to apply the tool to VNS studies, which is of particular relevance to this dissertation [43, 7].

- **PyPNS.** This library offers a thin interface to NEURON functionality designed particularly for models of peripheral nerve. While this package is rarely updated, it can provide a soft introduction to the NEURON package for novice peripheral nerve researchers.

- **TxBDC Cuff Lab dashboard.** While this library has also not seen recent maintenance, it powers an impressive, fully-parameterizable, interactive dashboard for simulating peripheral nerve stimulation. Currently, an instance of the dashboard is hosted at `https://nervestimlab.utdallas.edu/`.

- **RatInABox.** *RatInABox* is a Python program that simulates a rat locomotion model and corresponding neural activity. For example, you can simulate neural activity of the kind you might observe in "place cells" in the hippocampus as a rodent explores a 2D environment. Impressively, RatInABox only relies on minimimal dependencies that are common to the Python "scientific stack" such as matplotlib and scipy. While the project was definitely developed with a rat locomotive model in mind, various components of the project are clearly useful in other experimental paradigms. Particularly, many components can be applied to the vagus nerve stimulation paradigm presented in Chapter 5, and we can apply the 2D-locomotion functions to eye tracking.

6.3.2 Probabilistic Computing

- **Stan** receives first mention in this section due to its unique status as a "domain-specific language" that is not tied to any particular programming language or operating system. However, most users interface with Stan code via bindings in Python and R, making execution of STAN code quite accessible. The language was specifically designed to simulate and evaluate Bayesian models. Stan documentation not only provides useful

instructions on how to use the library, but also contains a helpful introduction to Bayesian computation in general.

- **PyMC** is a powerful and mature Python package for Bayesian analyses powered by *Monte Carlo simulations* and is an excellent tool for beginners and experts alike to implement their own Bayesian models.

- **Bambi** is a novice-friendly interface to the PyMC library with a focus on a symbolic syntax and common model configurations. While bambi is nowhere near as customizable as PyMC, it can be a great way for beginners to quickly develop their own Bayesian model instances for commonly-implemented Bayesian models.

- **bibat** is a Python package that helps users create templates for Bayesian projects, inspired by Bayesian workflow. It also includes some Python utilities for executing stages of Bayesian workflows.

6.3.3 Data Visualization

Data visualization is an essential part of data analysis, particularly when it comes to data exploration. Developing one's knowledge of the various data visualization approaches and libraries can greatly increase researcher productivity by providing fast, interpretable feedback to data processing steps or model/simulation outcomes.

- **The HoloViz ecosystem** is a collection of various intercompatible data visualization tools which are listed at https://holoviz.org/. The most useful of these tools likely includes **Panel**, a highly performant and interactive dashboard framework that leverages "reactive"-style programming that makes it easier to visualize complex datasets using several representations and visualizations at once, and **HoloViews**, a framework that focuses on the semantics of data visualization (e.g. specifying dependent vs independent variables) instead of just the raw specifications of the visualizations (e.g. x axis vs y axis). The result of combining these packages is the ability to rapidly prototype higly complex and performant dashboards that can convey critical information much more quickly and clearly than existing solutions. These and other tools in the HoloViz ecosystem are powered by a Param, a library that provides syntactical methods to effectively declare and process "parameters" in your Python code which represent the

primary variables and controls of interesting in data processing and performant data visualization.

- **Plotly/Dash.** The Plotly visualization library and the Dash interactive dashboard framework provide simple and accessible APIs for the generation of nicely-formatted figures from structured datasets (specifically, the Plotly Express module is designed to generate a variety of plot types from "tidy"[113] Pandas DataFrames). Dash is an extremely accessible dashboarding framework and their documentation provides an excellent introduction to the concepts and mechanisms that power its style of interactive dashboarding. Dash serves the dashboard via a Flask service, and running it locally displays the dashboard on "localhost" in your browser (this structure makes it easier to deploy to production on the Internet if desired).

- **Quarto/Shiny.** Quarto applies the "reactive" programming style to notebook-centric development and is capable of publishing extremely professional-looking data reports. Instead of executing individual notebook cells and introducing invisible state to your analysis, Quarto notebooks run all cells at once so that the notebooks can be deterministically executed. They have first-class support for interactivity, including in the form of dashboard frameworks. The Shiny dashboard framework, originally developed for R but now fully available for Python, provides an alternative API to Dash but operates on similar principles. Quarto also supports the Observable framework for the rare data practitioner that is proficient in JavaScript; these types of programmers usually specialize in highly interactive animations and UIs.

- **Matplotlib** is perhaps the most commonly used Python visualization package, particulary among beginners. While it does not provide as clean of a syntax as some of the aforementioned data visualization packages, it is generally more configurable for plots that need special formatting spcifications, such as formatting specifications from scientific journals. Additionally, it is the most powerful Python library available if you wish to build performant custom data animations that can be displayed in the browser or saved as a .gif or .mp4 video.

- **ArViz** Arviz is a visualiztion package specifically designed to interpret and display results from Bayesian analyses executed in PyMC or Bambi.

6.3.4 Psychophysical Modeling

A full comparison of psychophysical modeling libraries is provided in Chapter 4.

6.3.5 Data Manipulation

Data manipulation and preprocessing is an inevitable component of any data analysis. Unfortunately for researchers who are not experienced or trained in these areas, particularly people unfamiliar with tabular and relational data models and their representation in SQL or DataFrame libraries, it is extremely easy to get bogged down in data cleaning and waste valuable time that could be spent building/tweaking models or reviewing data more frequently.

- **SQL/DuckDB.** First, a note on SQL: it is my experience that academic researchers, particularly those with less computational background, are more comfortable performing data analysis in prodecural, impearative programming languages such as Python, MATLAB, and R. However, it seems to me that most researchers would be able to perform analyses much more quickly with just a small committment to learning the basics of SQL. Some benefits of using SQL: 1. The limited number of keywords allowed in the language syntax, whose purpose mostly correlates with their meanings in English, simplifies initial learning of the language. 2. Only a few new concepts need to be understood to write effective SQL, namely the basics of tabular and relational data organization principles, how they relate to the most common transformation patterns such as GROUP BY and JOIN. 3. SQL is a *declarative* programming language that optimizes the way that certain steps in the data transformation process are performed *before* actually executing those steps. This has largely beneficial implications for researchers less experienced in proper data manipulation and management techniques, such as performing filtering on large table selections *before* joining that selection with another table; SQL will evaluate how you've defined the data you want via the query syntax and make that optimization for you. However, these optimizations can break down for more complex data pipelines if they haven't been carefully structured.

Having provided justification for the usage of SQL in data processing, I now explain why I strongly believe that DuckDB should be the first choice for analytics in the sciences. Unless you have specific infrastructural or contractual reasons to use a different SQL dialect, I strongly recommend DuckDB as your one-stop-shop for everything SQL. First, installation of DuckDB is very simple and most people can get started making queries in minutes if not seconds. Second, DuckDB comes with very performant and accessible data I/O utilities that include reading and writing from CSV or Parquet files. (As an aside, Parquet files are generally orders of magnitude more performant than CSV for I/O operations and should almost always be preferred to CSV unless you have specific requirements for CSV, such as a human-readable raw data file or a collaborator who only can only process data in Excel). Additionally, the database is implemented on disk as a single file, similarly to SQLite. While this means that researchers should take extra care to create backups and otherwise protect the integrity of the orginal (raw) representations of the data that were recorded, but they are afforded a vastly simpler way to interact with and port the database without relying on dedicated servers, network connections, or complicated authentication protocols. Just point DuckDB to the local file and you're ready to make queries. Finally, DuckDB is an excellent choice for researchers because it was specifically designed to be an *analytical* database solution, unlike traditional SQL dialects (including SQLite) that were designed to be *transactional* databases, optimizing for speed and integrity when inserting, modifying, or deleting entire rows of data. Instead, data analysis pipelines are often more concerned with the *columns* of the dataset, which represent a specific attribute of an entity across all (or many) instances of said entity (rows). Column-oriented databases such as DuckDB leverage these assumptions about *analytical* queries and workflows to represent the data on disk in a way that makes it *very* fast for analytical queries (e.g. averaging data from millions or billions of rows).

- **Pandas.** Pandas is by far the most commonly-used dataframe library in the Python ecosystem. While its overly permissive syntax leaves much to be desired, researchers coming from MATLAB should find that picking up Pandas' table-oriented syntax to be a natural transition. Additionally, its implemenation of the Apache Arrow backend starting with version 2.0 primarily addressed existing performance concerns about the library. Pandas is an excellent choice for researchers new to the Python data ecosystem.

- **Polars.** Researchers seeking a highly performant DataFrame library that sits more on the cutting edge of tooling may be interested in Polars, a DataFrame library written in the highly performant Rust programming language. Polars boasts a much clearer and simplified syntax when compared to Pandas, and imporantly features a *LazyFrame* API that executes queries similarly to the declarative approach used by SQL that usually makes it a more performant choice than Python for data querying.

- **Dask.** Finally, researchers with a clear need for paralellizable data processing steps should turn to Dask. In most of the important ways, Dask's syntax mirrors Pandas' sytnax, but its adantage comes with its built-in methods to parallelize processing across many cores on your laptop or machines in a cloud cluster. Additionally, Dask bundles with an extremely illustrative dashboard that can be monitored during the execution of Dask queries; in real time, it displays the queueing, scheduling, and execution of tasks as they are distributed across multiple nodes of computation (e.g. CPU cores). Interested readers are pointed to Dask's YouTube tutorial that introduces the library.

- **xarray.** Scientific data commonly adheres to hierarchical data patterns and multi-dimensional data scenarios. Unlike 2D DataFrames e.g. Pandas DataFrames, xarray data structures elegantly labels data in many dimensions and provides an API that leverages this data representation to provide a way to slice by different dimensions more easily, visualize the data from different dimensional perspectives, etc. The syntax is a bit tricky to internalize, but mastery of this framework can save significant effort when dealing with complex multidimensional datasets[14].

6.3.6 Neuron Simulation

- **NEURON.** NEURON is the most popular software package in existence for simulating neuronal activation patterns and can model networks of neurons in addition to single neuron responses. It is available as both a Python package and as a GUI program. It is widely used in neuroscience research (over 3100 citations on Google Scholar) and extensively documented. While it provides abstractions over many low-level biological and computational mechanistic decisions, it is also highly configurable if you find their sensible defaults unsatisfactory for your use case.

[14]See Thinking with Xarray

6.3.7 Finite Element Modeling Software (and Related Tooling)

- **COMSOL.** COMSOL is a commercial finite element modeling (FEM) tool that is commonly used in peripheral nerve interface research in addition to research across engineering disciplines in general. It provides an intuitive graphical interface for the computer-aided design (CAD) aspects of model-building, and includes many pre-set material properties and model equations that alleviate the need to construct such components from scratch. It also provides some helpful utilities for model parameterization and post-processing. Perhaps the most outstanding feature of COMSOL is the way that it seamlessly integrates various aspects of 3D modelings, such as handling mathematical calculations, 3D graphical rendering, mesh generation, etc. Other ecosystems for 3D modeling often require you to develop implementations of these components independently and require custom code to integrate the components. One significant downside of COMSOL is that it is commercial software. In contrast with "free and open source" (FOSS) software, COMSOL costs money and is not customizable if you would like to make changes to the source code or extend its functionality. I am tangentially aware COMSOL provides API access points for programmatic scripting, but I am unfamiliar with of the mechanics of the API and its general accessibility.

- **COMSOL alternatives.** While there is no software to my knowledge that can match COMSOL's capabilities one-for-one, alternative FOSS software for 3D modeling includes FreeCAD and Blender.

- **dolfinx (formerly FEniCS)** A performant and flexible PDE solver in Python that uses the finite element method. However, a decent amount of domain knowledge with respect to PDE solvers and the finite element method itself is likely a prerequisite for effective usage.

- **Mesh generators in Python.** There are several adequate Python packages available for 3D mesh generation, however gmsh is a common choice and is officially recommended by the developers of dolfinx.

6.3.8 Honorable Mentions to other Programming Languages

While almost all of the software tools presented in this chapter are part of the Python ecosystem, mostly because it is my strongest programming language and in part because it's by far the most leveraged programming language in the modern data science community, I must give honorable mentions to the developments in other programming languages, namely Julia and Rust, as well as give a nod to the historical significance of R.

R. There is a large community of very skilled scientists and statisticians who use R as their primary programming language for data analysis. This is certainly an acceptable choice, often particularly justified given specialized statistical libraries or subject domains. R is an extremely powerful language for statistical analysis and is more than capable of fulfilling a similar set of benefits as the other tool suggestions in this section.

Julia. Julia's status as a performant language for scientific programming is well-known, but my initial impressions of the language have led me to believe that its benefits are undersold. Core decisions made by the language, reinforced by the ethos of the community, clearly prioritize a disciplined and maintainable style of programming that make it a great choice for researchers who find themselves running into common frustrations with Python and are in invested in learning more powerful, scalable, and maintainable code architectures. I am looking forward to learning more Julia.

Rust. Rust is a relatively new programming language with a rapidly growing community. It provides not-too-unfamiliar syntax for experienced Python programmers, making it capable of implementing the same kinds of high-level APIs that have enabled Python's popularity, but it can also perform the role of a low-level systems language that is challenging C's ubuiquity as the go-to language for highly performant applications. Programs that are built with Rust are, more often than not, *blazingly fast*. For example, Rust has been used to create performant alternatives to critical Python developer tools– namely, the linter/code formatter Ruff (replacing flake8/black in the ecosystem), and the package manager uv (replacing pip in the ecosystem). These tools are so performant and easy to implement that they have been adopted by the Python community at unprecedented rates. While packages exist that are more tailored for scientific analysis and data processing, they are still quite behind what is available in Python. However, it is entirely within the realm of possibility that data tools built with Rust completely take over the ecosystem in the next 5 to 10 years.

6.4 A Bayesian Workflow Proposal to Guide Next Steps

Now that we have a general idea of the kinds of software tools that are available to us, what should we do with them?

The following Bayesian workflow proposal is adapted from the general Bayesian workflow suggestions established in Gelman et al. 2020[24]. It attempts to provide context on how one might implement the models we have discussed to perform follow-up analyses on our data sets or plan new experiments with methods that resemble the ones here.

6.4.1 Choosing an Initial Model

Choosing an initial model is simple in this instance: we may start with the Bayesian curve fits as is common in the field. As previously described, this is usually implemented via a generalized linear model with a sigmoid link function e.g. the logistic curve. Although we could theoretically rely on an existing psychophysics curve-fitting library to obtain these estimates, we should construct our own implementation to ensure that the results from the estimation are compatible with ensuing steps in the workflow.

This initial model should be implemented in PyMC. PyMC contains an accessible tutorial on the implementation of a GLM binomial regression model in their example documentation [15].

For the first iteration of the model, stimulus amplitude should be chosen as the value that is modulated along the abscissa.

6.4.2 Scaling and Transforming the Parameters

We have some choices to make when it comes to scaling and transforming our model parameters. According to Gelman et al.[24], our parameters should be in units that either enable informative priors or facilitate hierarchical modeling. In the context of the psychometric

[15] https://www.pymc.io/projects/examples/en/latest/generalized_linear_models/GLM-binomial-regression.html

curve, this might determine whether we use a "slope" or "spread" parameterization for the parameter that represents the sigmoid's rate of growth.

6.4.3 Prior Predictive Checking

Bayesian modeling requires the establishment of *prior* distributions that represent our knowledge of the distribution of the parameter space before data is acquired. In existing psychometric curve fitting libraries, priors are often broad and simply represent the range of allowable values (such as restricting the values of threshold to be a positive number). Prior predictive checks simulate experimental runs based on the priors we have established. We may then assess the model output, based on the priors alone, and make sure that the predicted outcomes make intuitive sense.

There is a golden opportunity on the table to develop more targeted specifications for priors in psychophysical data analysis; doing so would narrow the range of potential outcomes of model simulations by introducing stricter rules into the simulation's universe. For example, if we have a model that includes a random variable for a neuron's firing rate, instead of setting a prior distribution that samples a neuron's firing rate from all the real-valued numbers, we can introduce well-established and well-modeled concept of a refractory period of an action potential into the model so that the simulations don't attempt to sample values of firing rate that violate the refactory period.

If we decide to adjust model parameters at this step, to avoid "data snooping," we must sure that our adjustments come from data or scientific claims that preclude the data obtained in the experiments.

6.4.4 Running the Model

Because our experimental design is a natural use case for a common Bayesian GLM, running the default number of "warm up" and simulation samples provided by PyMC should be sufficient to avoid divergence or other model evalutation failures in early stages of model development, but a smaller number of iterations may be chosen as complexity is instroduced in order to maintain or facilitate speed of development and frequency of model iteration.

It is a good idea at this stage to run the canonical "checks" on model validity, such as the spread of residuals and posterior predictive checks, using pre-configured analysis in the Arviz library before moving on to more complex stages of model development. A few "hard-coded" values of "known" underlying parameter values might be carefully selected to provide further sanity checks and implement *simulation-based calibration* (SBC) if our posterior predictive checks do not feature maximum likelihood estimates around the hard-coded value. Such instances would likely necessitate adjustment of priors.

6.4.5 Trying Other Stimulus Dimensions

Once we are satisfied with the results of the above model, we should try a similar process using pulse width as the modulated value along the abscissa. Following that, we might try other parameterizations of the stimulus waveform such as the product of pulse width and amplitude that might align with established theories of peripheral nerve stimulation coding.

At this point, we might try to model various linear or non-linear combinations of multiple parameters from the original model, and adjust the model to accommodate multiple input dimensions. We can then try out different models of peripheral nerve encoding on for size and utilize canonical model-comparison methods to determine which model is more likely to be accurate given the priors and your data set.

6.4.6 Tuning the Model for a Single Session

In tandem with these methods, at this point we should explore many different angles of model adjustments; Gelman et al. [24] provides well-illustrated examples of the practical considerations involved in this process. Here we should take the time to thoughtfully consider what kinds of useful information we want to obtain from the model, and where certain "shortcuts" we take during modeling might be biasing the outcomes.

For example, we might consider replacing a simple binomial model with a "beta-binomial" model if we believe that there is evidence of non-stationarity within a session[88]. As another example, we might want to introduce variables from latent mechanisms such as fiber size distribution, estimates from simulation models such as ASCENT[67] or TouchSim[72], or

related biological hypothesis, and examine how the addition of these variables affect our model's behavior, such as simulation divergence, long runtimes, poor fit, prior and posterior checks, etc.

This stage of the process is the one that will likely ultimately require careful thought and a good deal of experience with probabilistic programming.

6.4.7 Introducing Hierarchy

When we are confident that our model for a single session is well-posed, then it will likely be a good idea to introduce hiearchy to the model, namely in the form of individuals or cohorts (such models would be stronger in non-monkey experiments), or the time dimension (implementing probability models and concepts appropriate for our time-series representation such as random walks, learning rates, asymptotal behavior, models that incorporate sudden perturbations, etc. Here, we can leverage repeated in-subject measurements to make assertions about within-subject variance). If our single-session model is robust, then the addition of hierarchy is very likely to make the explanatory power of our model and data much stronger. If our single-session model is weak, introducing hierarchy is very likely to lead to errant behavior during the model-fitting process, particularly in the form of diverging samples.

There are a good handful of studies that provide precedent for the application of Bayesian hierarchical models to psychophysical experiments; Houpt and Bittner is recommended as a model example which includes Stan source code [40]

6.5 Conclusion

The peripheral nerve provides an extremely useful system for implementing and measuring changes to our neural systems while avoiding a lot of the complexity that is natural to measurements or stimuluation taken directly from the brain. Research on the peripheral nerve, and specifically, a wholistic understanding of the neural mechanisms operating at the level of the periphery, are well-poised to deliver significant breakthroughs in neural interfaces and our understanding of our sensory systems in general. Techniques new and old may be

integrated to make the most out of our inquiries and experiments; my hope is that the work presented here can provide a humble signpost towards a more sophisticated and powerful approach to neural engineering and data processing.

References

[1] K. L. Akre and S. Johnsen. Psychophysics and the evolution of behavior. *Trends in Ecology & Evolution*, 29(5):291–300, May 2014.

[2] M. Alexandra Kredlow, R. J. Fenster, E. S. Laurent, K. J. Ressler, and E. A. Phelps. Prefrontal cortex, amygdala, and threat processing: implications for PTSD. *Neuropsychopharmacology*, 47(1):247–259, Jan. 2022. Publisher: Nature Publishing Group.

[3] S. Archibald, J. Shefner, C. Krarup, and R. Madison. Monkey median nerve repaired by nerve graft or collagen nerve guide tube. *Journal of Neuroscience*, 15(5):4109–4123, 1995. Publisher: Society for Neuroscience.

[4] J. E. Arle, K. W. Carlson, and L. Mei. Investigation of mechanisms of vagus nerve stimulation for seizure using finite element modeling. *Epilepsy Research*, 126:109–118, Oct. 2016.

[5] J. Badia, T. Boretius, D. Andreu, C. Azevedo-Coste, T. Stieglitz, and X. Navarro. Comparative analysis of transverse intrafascicular multichannel, longitudinal intrafascicular and multipolar cuff electrodes for the selective stimulation of nerve fascicles. *Journal of Neural Engineering*, 8(3):036023, June 2011.

[6] H.-R. Berthoud and W. L. Neuhuber. Functional and chemical anatomy of the afferent vagal system. *Autonomic Neuroscience: Basic and Clinical*, 85(1):1–17, Dec. 2000. Publisher: Elsevier.

[7] S. L. Blanz, E. D. Musselman, M. L. Settell, B. E. Knudsen, E. N. Nicolai, J. K. Trevathan, R. S. Verner, J. Begnaud, A. C. Skubal, A. J. Suminski, J. C. Williams, A. J. Shoffstall, W. M. Grill, N. A. Pelot, and K. A. Ludwig. Spatially selective stimulation of the pig vagus nerve to modulate target effect versus side effect. *Journal of Neural Engineering*, 20(1):016051, feb 2023.

[8] H. Bostock. The strength-duration relationship for excitation of myelinated nerve: computed dependence on membrane parameters. *The Journal of Physiology*, 341:59–74, Aug. 1983.

[9] D. H. Brainard. The Psychophysics Toolbox. *Spatial Vision*, 10(4):433–436, 1997.

[10] N. Brill, K. Polasek, E. Oby, C. Ethier, L. Miller, and D. Tyler. Nerve cuff stimulation and the effect of fascicular organization for hand grasp in nonhuman primates. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference*, 2009:1557–1560, 2009.

[11] F. Buchthal and A. Rosenfalck. Evoked action potentials and conduction velocity in human sensory nerves. *Brain Research*, 3(1):v–122, Nov. 1966.

[12] S. F. Cogan, K. A. Ludwig, C. G. Welle, and P. Takmakov. Tissue damage thresholds during therapeutic electrical stimulation. *Journal of Neural Engineering*, 13(2):021001, jan 2016.

[13] I. Cuberovic, A. Gill, L. J. Resnik, D. J. Tyler, and E. L. Graczyk. Learning of Artificial Sensation Through Long-Term Home Use of a Sensory-Enabled Prosthesis. *Frontiers in Neuroscience*, 13:853, 2019.

[14] M. Dali, O. Rossel, D. Andreu, L. Laporte, A. Hernández, J. Laforet, E. Marijon, A. Hagège, M. Clerc, C. Henry, and D. Guiraud. Model based optimal multipolar stimulation without a priori knowledge of nerve structure: application to vagus nerve stimulation. *Journal of Neural Engineering*, 15(4):046018, May 2018. Publisher: IOP Publishing.

[15] N. S. Davidovics, G. Y. Fridman, B. Chiang, and C. C. D. Santina. Effects of biphasic current pulse frequency, amplitude, duration, and interphase gap on eye movement responses to prosthetic electrical stimulation of the vestibular nerve. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(1):84–94, 2011.

[16] M. Davis, D. L. Walker, and K. M. Myers. Role of the Amygdala in Fear Extinction Measured with Potentiated Startle. *Annals of the New York Academy of Sciences*, 985(1):218–232, 2003. _eprint: https://nyaspubs.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1749-6632.2003.tb07084.x.

[17] G. D. Dawson. The relative excitability and conduction velocity of sensory and motor nerve fibres in man. *The Journal of Physiology*, 131(2):436–451, Feb. 1956.

[18] G. Dhillon and K. Horch. Direct Neural Sensory Feedback and Control of a Prosthetic Arm. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 13:468–72, Jan. 2006.

[19] S. Duvarci and D. Pare. Amygdala Microcircuits Controlling Learned Fear. *Neuron*, 82(5):966–980, June 2014. Publisher: Elsevier.

[20] J. O. Foley and F. S. DuBois. Quantitative studies of the vagus nerve in the cat. I. The ratio of sensory to motor fibers. *Journal of Comparative Neurology*, 67(1):49–67, 1937. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.900670104.

[21] I. Fründ, N. V. Haenel, and F. A. Wichmann. Inference for psychometric functions in the presence of nonstationary behavior. *Journal of Vision*, 11(6), May 2011.

[22] K. Fugleholm, H. Schmalbruch, and C. Krarup. Post reinnervation maturation of myelinated nerve fibers in the cat tibial nerve: chronic electrophysiological and morphometric studies. *Journal of the Peripheral Nervous System*, 5(2):82–95, 2000.

[23] J. L. Gaines, K. E. Finn, J. P. Slopsema, L. A. Heyboer, and K. H. Polasek. A model of motor and sensory axon activation in the median nerve using surface electrical stimulation. *Journal of Computational Neuroscience*, 45(1):29–43, Aug. 2018.

[24] A. Gelman, A. Vehtari, D. Simpson, C. C. Margossian, B. Carpenter, Y. Yao, L. Kennedy, J. Gabry, P.-C. Bürkner, and M. Modrák. Bayesian Workflow, Nov. 2020. arXiv:2011.01808 [stat].

[25] M. S. George, H. A. Sackeim, A. J. Rush, L. B. Marangell, Z. Nahas, M. M. Husain, S. Lisanby, T. Burt, J. Goldman, and J. C. Ballenger. Vagus nerve stimulation: a new tool for brain research and therapy. *Biological Psychiatry*, 47(4):287–295, Feb. 2000.

[26] M. S. George, H. E. Ward, P. T. Ninan, M. Pollack, Z. Nahas, B. Anderson, S. Kose, R. H. Howland, W. K. Goodman, and J. C. Ballenger. A pilot study of vagus nerve stimulation (VNS) for treatment-resistant anxiety disorders. *Brain Stimulation*, 1(2):112–121, 2008.

[27] W. Gerstner and W. M. Kistler. Mathematical formulations of Hebbian learning. *Biological Cybernetics*, 87(5):404–415, Dec. 2002.

[28] G. A. Gescheider. *Psychophysics: The Fundamentals*. Psychology Press, June 2013.

[29] J. Geweke. Bayesian Model Comparison and Validation. *The American Economic Review*, 97(2):60–64, 2007. Publisher: American Economic Association.

[30] U. Ghafoor, S. Kim, and K.-S. Hong. Selectivity and Longevity of Peripheral-Nerve and Machine Interfaces: A Review. *Frontiers in Neurorobotics*, 11, 2017.

[31] E. J. Gibson. Improvement in perceptual judgments as a function of controlled practice or training. *Psychological Bulletin*, 50(6):401–431, 1953.

[32] M. Gillespie, B. Jassal, R. Stephan, M. Milacic, K. Rothfels, A. Senff-Ribeiro, J. Griss, C. Sevilla, L. Matthews, C. Gong, C. Deng, T. Varusai, E. Ragueneau, Y. Haider, B. May, V. Shamovsky, J. Weiser, T. Brunson, N. Sanati, L. Beckman, X. Shao, A. Fabregat, K. Sidiropoulos, J. Murillo, G. Viteri, J. Cook, S. Shorser, G. Bader, E. Demir, C. Sander, R. Haw, G. Wu, L. Stein, H. Hermjakob, and P. D'Eustachio. The reactome pathway knowledgebase 2022. *Nucleic Acids Research*, 50(D1):D687–D692, Nov. 2021. _eprint: https://academic.oup.com/nar/article-pdf/50/D1/D687/42058295/gkab1028.pdf.

[33] P. H. Gorman and J. T. Mortimer. The effect of stimulus parameters on the recruitment characteristics of direct nerve stimulation. *IEEE Transactions on Biomedical Engineering*, BME-30(7):407–414, 1983.

[34] E. L. Graczyk, M. A. Schiefer, H. P. Saal, B. P. Delhaye, S. J. Bensmaia, and D. J. Tyler. The neural basis of perceived intensity in natural and artificial touch. *Science Translational Medicine*, 8(362):362ra142–362ra142, Oct. 2016. Publisher: American Association for the Advancement of Science Section: Research Article.

[35] P. C. Griffin, J. Khadake, K. S. LeMay, S. E. Lewis, S. Orchard, A. Pask, B. Pope, U. Roessner, K. Russell, T. Seemann, A. Treloar, S. Tyagi, J. H. Christiansen, S. Dayalan, S. Gladman, S. B. Hangartner, H. L. Hayden, W. W. Ho, G. Keeble-Gagnère, P. K. Korhonen, P. Neish, P. R. Prestes, M. F. Richardson, N. S. Watson-Haigh, K. L. Wyres, N. D. Young, and M. V. Schneider. Best practice data life cycle approaches for the life sciences. *F1000Research*, 6, June 2018.

[36] W. Grill and J. Mortimer. The effect of stimulus pulse duration on selectivity of neural stimulation. *IEEE Transactions on Biomedical Engineering*, 43(2):161–166, 1996.

[37] W. M. Grill and J. Thomas Mortimer. Electrical properties of implant encapsulation tissue. *Annals of Biomedical Engineering*, 22(1):23–33, Jan. 1994.

[38] L. Hermann. *Elements of Human Physiology*. Smith, Elder, 1875. Google-Books-ID: Lus2AQAAMAAJ.

[39] G. Hesselmann. Applying Linear Mixed Effects Models (LMMs) in Within-Participant Designs With Subjective Trial-Based Assessments of Awareness—a Caveat. *Frontiers in Psychology*, 9, May 2018.

[40] J. W. Houpt and J. L. Bittner. Analyzing thresholds and efficiency with hierarchical bayesian logistic regression. *Vision Research*, 148:49–58, 2018.

[41] J. Howells, L. Trevillion, H. Bostock, and D. Burke. The voltage dependence of Ih in human myelinated axons. *The Journal of Physiology*, 590(Pt 7):1625–1640, Apr. 2012.

[42] X. Huan and Y. M. Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232(1):288–317, Jan. 2013.

[43] W. J. Huffman, E. D. Musselman, N. A. Pelot, and W. M. Grill. Measuring and modeling the effects of vagus nerve stimulation on heart rate and laryngeal muscles. *Bioelectronic Medicine*, 9(1):3, Feb. 2023.

[44] D. R. Hulsey, J. R. Riley, K. W. Loerwald, R. L. Rennaker, M. P. Kilgard, and S. A. Hays. Parametric characterization of neural activity in the locus coeruleus in response to vagus nerve stimulation. *Experimental Neurology*, 289a(Supplement C):21–30, Mar. 2017.

[45] N. Jayaprakash, W. Song, V. Toth, A. Vardhan, T. Levy, J. Tomaio, K. Qanud, I. Mughrabi, Y.-C. Chang, M. Rob, A. Daytz, A. Abbas, Z. Nassrallah, B. T. Volpe, K. J. Tracey, Y. Al-Abed, T. Datta-Chaudhuri, L. Miller, M. F. Barbe, S. C. Lee,

T. P. Zanos, and S. Zanos. Organ- and function-specific anatomical organization of the vagus nerve supports fascicular vagus nerve stimulation.

[46] R. S. Johansson and A. B. Vallbo. Detection of tactile stimuli. Thresholds of afferent units related to psychophysical thresholds in the human hand. *The Journal of Physiology*, 297(1):405–422, 1979. _eprint: https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1979.sp013048.

[47] K. O. Johnson, S. S. Hsiao, and T. Yoshioka. Neural Coding and the Basic Law of Psychophysics. *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry*, 8(2):111–121, Apr. 2002.

[48] R. E. Kass and D. Steffey. Approximate Bayesian Inference in Conditionally Independent Hierarchical Models (Parametric Empirical Bayes Models). *Journal of the American Statistical Association*, 84(407):717–726, Sept. 1989. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01621459.1989.10478825.

[49] J. K. Kruschke and T. M. Liddell. The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*, 25(1):178–206, Feb. 2018.

[50] M. Kuss, F. Jäkel, and F. A. Wichmann. Bayesian inference for psychometric functions. *Journal of Vision*, 5(5):8–8, 2005.

[51] F. Li, P. Ding, and F. Mealli. Bayesian causal inference: a critical review. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 381(2247):20220153, Mar. 2023. Publisher: Royal Society.

[52] J. P. Lund, G.-D. Sun, Y. Lamarre, and T. P. Pons. Cortical Reorganization and Deafferentation in Adult Macaques. *Science*, 265(5171):546–548, 1994.

[53] N. A. Macmillan. Signal detection theory as data analysis method and psychological decision model. In *A handbook for data analysis in the behavioral sciences: Methodological issues*, pages 21–57. Lawrence Erlbaum Associates, Inc, Hillsdale, NJ, US, 1993.

[54] D. A. Magezi. Linear mixed-effects models for within-participant psychology experiments: an introductory tutorial and free, graphical user interface (LMMgui). *Frontiers in Psychology*, 6, 2015.

[55] R. Mandeville, S. Deshmukh, E. T. Tan, V. Kumar, B. Sanchez, A. S. Dowlatshahi, J. Luk, R. H. B. See, C. F. D. Leochico, J. A. Thum, S. Bazarek, B. Johnston, J. Brown, J. Wu, D. Sneag, and S. Rutkove. A scoping review of current and emerging techniques for evaluation of peripheral nerve health, degeneration and regeneration: part 2, non-invasive imaging. *Journal of Neural Engineering*, 20(4):041002, Aug. 2023.

[56] A. Maravita and A. Iriki. Tools for the body (schema). *Trends in Cognitive Sciences*, 8(2):79–86, Feb. 2004. Publisher: Elsevier.

[57] R. Marek, L. Xu, R. K. P. Sullivan, and P. Sah. Excitatory connections between the prelimbic and infralimbic medial prefrontal cortex show a role for the prelimbic cortex in fear extinction. *Nature Neuroscience*, 21(5):654–658, May 2018.

[58] C. C. McIntyre, A. G. Richardson, and W. M. Grill. Modeling the Excitability of Mammalian Nerve Fibers: Influence of Afterpotentials on the Recovery Cycle. *Journal of Neurophysiology*, 87(2):995–1006, Feb. 2002. Publisher: American Physiological Society.

[59] C. M. McKay and H. J. McDermott. The perceptual effects of current pulse duration in electrical stimulation of the auditory nerve. *The Journal of the Acoustical Society of America*, 106(2):998–1009, 1999.

[60] M. M. Merzenich, J. H. Kaas, J. T. Wall, M. Sur, R. J. Nelson, and D. J. Felleman. Progression of change following median nerve section in the cortical representation of the hand in areas 3b and 1 in adult owl and squirrel monkeys. *Neuroscience*, 10(3):639–665, Oct. 1983.

[61] M. M. Merzenich, R. J. Nelson, M. P. Stryker, M. S. Cynader, A. Schoppmann, and J. M. Zook. Somatosensory cortical map changes following digit amputation in adult monkeys. *Journal of Comparative Neurology*, 224(4):591–605, 1984.

[62] M. Mezzetti, C. P. Ryan, P. Balestrucci, F. Lacquaniti, and A. Moscatelli. Bayesian hierarchical models and prior elicitation for fitting psychometric functions. *Frontiers in Computational Neuroscience*, 17, 2023.

[63] J. O. Mills, A. Jalil, and P. E. Stanga. Electronic retinal implants and artificial vision: journey and present. *Eye*, 31(10):1383–1398, Oct. 2017. Publisher: Nature Publishing Group.

[64] I. Mogyoros, M. C. Kiernan, J.-M. Gracies, and D. Burke. The effect of stimulus duration on the latency of submaximal nerve volleys. *Muscle & Nerve*, 19(10):1354–1356, 1996.

[65] M. Moldovan and C. Krarup. Persistent abnormalities of membrane excitability in regenerated mature motor axons in cat. *The Journal of Physiology*, 560(3):795–806, 2004.

[66] M. A. Muniak, S. Ray, S. S. Hsiao, J. F. Dammann, and S. J. Bensmaia. The Neural Coding of Stimulus Intensity: Linking the Population Response of Mechanoreceptive Afferents with Psychophysical Behavior. *Journal of Neuroscience*, 27(43):11687–11699, Oct. 2007. Publisher: Society for Neuroscience, Section: Articles.

[67] E. D. Musselman, J. E. Cariello, W. M. Grill, and N. A. Pelot. ASCENT (Automated Simulations to Characterize Electrical Nerve Thresholds): A pipeline for sample-specific computational modeling of electrical stimulation of peripheral nerves. *PLOS Computational Biology*, 17(9):e1009285, Sept. 2021. Publisher: Public Library of Science.

[68] K. M. Myers and M. Davis. Mechanisms of fear extinction. *Molecular Psychiatry*, 12(2):120–150, Feb. 2007. Publisher: Nature Publishing Group.

[69] C. B. Nemeroff, H. S. Mayberg, S. E. Krahl, J. McNamara, A. Frazer, T. R. Henry, M. S. George, D. S. Charney, and S. K. Brannan. VNS Therapy in Treatment-Resistant Depression: Clinical Evidence and Putative Neurobiological Mechanisms. *Neuropsychopharmacology*, 31(7):1345–1355, July 2006.

[70] L. J. Noble, V. B. Meruva, S. A. Hays, R. L. Rennaker, M. P. Kilgard, and C. K. McIntyre. Vagus nerve stimulation promotes generalization of conditioned fear extinction and reduces anxiety in rats. *Brain Stimulation*, 12(1):9–18, 2019.

[71] D. Oberfeld and T. Franke. Evaluating the robustness of repeated measures analyses: The case of small sample sizes and nonnormal data. *Behavior Research Methods*, 45(3):792–812, Sept. 2013.

[72] E. V. Okorokova, Q. He, and S. J. Bensmaia. Biomimetic encoding model for restoring touch in bionic hands through a nerve interface. *Journal of Neural Engineering*, 15(6):066033, Oct. 2018. Publisher: IOP Publishing.

[73] D. H. O'Connor, L. Krubitzer, and S. Bensmaia. Of mice and monkeys: Somatosensory processing in two prominent animal models. *Progress in Neurobiology*, 201:102008, June 2021.

[74] E. L. Peck, C. J. Peck, and C. D. Salzman. Task-Dependent Spatial Selectivity in the Primate Amygdala. *Journal of Neuroscience*, 34(49):16220–16233, Dec. 2014. Publisher: Society for Neuroscience Section: Articles.

[75] J. Peirce, J. R. Gray, S. Simpson, M. MacAskill, R. Höchenberger, H. Sogo, E. Kastman, and J. K. Lindeløv. PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1):195–203, Feb. 2019.

[76] D. F. Peña, N. D. Engineer, and C. K. McIntyre. Rapid Remission of Conditioned Fear Expression with Extinction Training Paired with Vagus Nerve Stimulation. *Biological Psychiatry*, 73(11):1071–1077, June 2013.

[77] B. E. Pfingst. Changes over time in thresholds for electrical stimulation of the cochlea. *Hearing Research*, 50(1):225–236, Dec. 1990.

[78] T. P. Pons, P. E. Garraghty, A. K. Ommaya, J. H. Kaas, E. Taub, and M. Mishkin. Massive cortical reorganization after sensory deafferentation in adult macaques. *Science*, 252(5014):1857–1860, June 1991.

[79] N. Prins. Easy, bias-free Bayesian hierarchical modeling of the psychometric function using the Palamedes Toolbox. *Behavior Research Methods*, 2023.

[80] N. Prins and F. A. A. Kingdom. Applying the Model-Comparison Approach to Test Specific Research Hypotheses in Psychophysical Research Using the Palamedes Toolbox. *Frontiers in Psychology*, 9, 2018. Publisher: Frontiers.

[81] S. Raspopovic, F. M. Petrini, M. Zelechowski, and G. Valle. Framework for the development of neuroprostheses: From basic understanding by sciatic and median nerves models to bionic legs and hands. *Proceedings of the IEEE*, 105(1):34–49, 2017.

[82] J. C. A. Read. The place of human psychophysics in modern neuroscience. *Neuroscience*, 296:116–129, June 2015.

[83] K. Rehfeld, N. Marwan, J. Heitzig, and J. Kurths. Comparison of correlation analysis techniques for irregularly sampled time series. *Nonlinear Processes in Geophysics*, 18:389–404, June 2011.

[84] L. Resnik, G. Latlief, S. L. Klinger, N. Sasson, and L. S. Walters. Do users want to receive a DEKA Arm and why? Overall findings from the Veterans Affairs Study to optimize the DEKA Arm. *Prosthetics and Orthotics International*, 38(6):456–466, Dec. 2014.

[85] A. G. Rouse, J. J. Williams, J. J. Wheeler, and D. W. Moran. Spatial co-adaptation of cortical control columns in a micro-ECoG brain-computer interface. *Journal of Neural Engineering*, 13(5):056018, Oct. 2016.

[86] P. Rutecki. Anatomical, physiological, and theoretical basis for the antiepileptic effect of vagus nerve stimulation. *Epilepsia*, 31 Suppl 2:S1–6, 1990.

[87] T. C. Schlichenmeyer, E. R. Zellmer, H. Burton, W. Z. Ray, and D. W. Moran. Detection and discrimination of electrical stimuli from an upper limb cuff electrode in m. mulatta. *Journal of Neural Engineering*, 19(6):066009, nov 2022.

[88] H. H. Schütt, S. Harmeling, J. H. Macke, and F. A. Wichmann. Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research*, 122:105–123, May 2016.

[89] R. K. Shepherd and E. Javel. Electrical stimulation of the auditory nerve: Ii. effect of stimulus waveshape on single fibre response properties. *Hearing Research*, 130(1):171–188, 1999.

[90] M. Slugocki, A. B. Sekuler, and P. Bennett. BayesFit: A tool for modeling psychophysical data using Bayesian inference. *Journal of Open Research Software*, 7(1):2, Jan. 2019.

[91] D. Smith and C. Finley. Effects of electrode configuration on psychophysical strength-duration functions for single biphasic electrical stimuli in cats. *The Journal of the Acoustical Society of America*, 102:2228–37, Nov. 1997.

[92] R. R. Souza, M. B. Powers, R. L. Rennaker, C. K. McIntyre, S. A. Hays, and M. P. Kilgard. Timing of vagus nerve stimulation during fear extinction determines efficacy in a rat model of PTSD. *Scientific Reports*, 12(1):16526, Oct. 2022. Publisher: Nature Publishing Group.

[93] R. R. Souza, N. M. Robertson, C. K. McIntyre, R. L. Rennaker, S. A. Hays, and M. P. Kilgard. Vagus nerve stimulation enhances fear extinction as an inverted-U function of stimulation intensity. *Experimental Neurology*, 341:113718, July 2021.

[94] R. R. Souza, N. M. Robertson, D. T. Pruitt, P. A. Gonzales, S. A. Hays, R. L. Rennaker, M. P. Kilgard, and C. K. McIntyre. Vagus nerve stimulation reverses the extinction impairments in a model of ptsd with prolonged and repeated trauma. *Stress*, 22(4):509–520, 2019.

[95] W. C. Stebbins. *Animal Psychophysics: the design and conduct of sensory experiments*. Springer US, 1970.

[96] S. S. Stevens. *Psychophysics: Introduction to Its Perceptual, Neural and Social Prospects*. Routledge, Sept. 2017.

[97] J. Z. Sun, G. I. Wang, V. K. Goyal, and L. R. Varshney. A framework for Bayesian optimality of psychophysical laws. *Journal of Mathematical Psychology*, 56(6):495–501, Dec. 2012.

[98] D. W. Tan, M. A. Schiefer, M. W. Keith, J. R. Anderson, and D. J. Tyler. Stability and selectivity of a chronic, multi-contact cuff electrode for sensory stimulation in human amputees. *Journal of Neural Engineering*, 12(2):026002, Jan. 2015. Publisher: IOP Publishing.

[99] D. W. Tan, M. A. Schiefer, M. W. Keith, J. R. Anderson, J. Tyler, and D. J. Tyler. A neural interface provides long-term stable natural touch perception. *Science Translational Medicine*, 6(257):257ra138–257ra138, 2014. Publisher: American Association for the Advancement of Science.

[100] J. Tigerholm, A. H. Poulsen, O. K. Andersen, and C. D. Mørch. From Perception Threshold to Ion Channels—A Computational Study. *Biophysical Journal*, 117(2):281–295, 2019.

[101] E. Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, Sept. 2004.

[102] D. J. Tyler and D. M. Durand. Functionally selective peripheral nerve stimulation with a flat interface nerve electrode. *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, 10(4):294–303, Dec. 2002.

[103] J. Vanlier, C. A. Tiemann, P. A. J. Hilbers, and N. A. W. van Riel. A Bayesian approach to targeted experiment design. *Bioinformatics*, 28(8):1136–1142, Apr. 2012.

[104] C. Veraart, W. M. Grill, and J. T. Mortimer. Selective control of muscle activation with a multipolar nerve cuff electrode. *IEEE transactions on bio-medical engineering*, 40(7):640–653, July 1993.

[105] G. Wald. The Receptors of Human Color Vision. *Science*, 145(3636):1007–1016, Sept. 1964. Publisher: American Association for the Advancement of Science.

[106] J. T. Wall, D. J. Felleman, and J. H. Kaas. Recovery of Normal Topography in the Somatosensory Cortex of Monkeys After Nerve Crush and Regeneration. *Science*, 221(4612):771–773, Aug. 1983. Publisher: American Association for the Advancement of Science.

[107] J. T. Wall, J. H. Kaas, M. Sur, R. J. Nelson, D. J. Felleman, and M. M. Merzenich. Functional reorganization in somatosensory cortical areas 3b and 1 of adult monkeys after median nerve repair: possible relationships to sensory recovery in humans. *Journal of Neuroscience*, 6(1):218–233, Jan. 1986. Publisher: Society for Neuroscience Section: Articles.

[108] A. B. Watson. QUEST+: A general multidimensional Bayesian adaptive psychometric method. *Journal of Vision*, 17(3):10–10, 2017.

[109] S. Wendelken, D. M. Page, T. Davis, H. A. C. Wark, D. T. Kluger, C. Duncan, D. J. Warren, D. T. Hutchinson, and G. A. Clark. Restoration of motor control and proprioceptive and cutaneous sensation in humans with prior upper-limb amputation via multiple Utah Slanted Electrode Arrays (USEAs) implanted in residual peripheral arm nerves. *Journal of NeuroEngineering and Rehabilitation*, 14(1):121, Nov. 2017.

[110] L. Wernisch, T. Edwards, A. Berthon, O. Tessier-Lariviere, E. Sarkans, M. Stoukidi, P. Fortier-Poisson, M. Pinkney, M. Thornton, C. Hanley, S. Lee, J. Jennings, B. Appleton, P. Garsed, B. Patterson, W. Buttinger, S. Gonshaw, M. Jakopec, S. Shunmugam, J. Mamen, A. Tukiainen, G. Lajoie, O. Armitage, and E. Hewage. Online Bayesian optimization of vagus nerve stimulation. *Journal of Neural Engineering*, 21(2):026019, Apr. 2024. Publisher: IOP Publishing.

[111] N. Whittle, J. Fadok, K. P. MacPherson, R. Nguyen, P. Botta, S. B. E. Wolff, C. Müller, C. Herry, P. Tovote, A. Holmes, N. Singewald, A. Lüthi, and S. Ciocchi. Central amygdala micro-circuits mediate fear extinction. *Nature Communications*, 12(1):4156, July 2021. Number: 1 Publisher: Nature Publishing Group.

[112] F. A. Wichmann and N. J. Hill. The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, 63(8):1293–1313, Nov. 2001.

[113] H. Wickham. Tidy Data. *Journal of Statistical Software*, 59:1–23, Sept. 2014.

[114] B. S. Wilson, C. C. Finley, D. T. Lawson, R. D. Wolford, D. K. Eddington, and W. M. Rabinowitz. Better speech recognition with cochlear implants. *Nature*, 352(6332):236–238, July 1991.

[115] K. Woeppel, C. Hughes, A. J. Herrera, J. R. Eles, E. C. Tyler-Kabara, R. A. Gaunt, J. L. Collinger, and X. T. Cui. Explant Analysis of Utah Electrode Arrays Implanted in Human Cortex for Brain-Computer-Interfaces. *Frontiers in Bioengineering and Biotechnology*, 9, 2021.

[116] P. B. Yoo, N. B. Lubock, J. G. Hincapie, S. B. Ruble, J. J. Hamann, and W. M. Grill. High-resolution measurement of electrically-evoked vagus nerve activity in the anesthetized dog. *Journal of Neural Engineering*, 10(2):026003, jan 2013.

[117] T. Young. II. The Bakerian Lecture. On the theory of light and colours. *Philosophical Transactions of the Royal Society of London*, 92:12–48, Jan. 1802. Publisher: Royal Society.