PROBABILISTIC ENCODING AND FEATURE

SELECTIVITY IN THE SOMATOSENSORY

PATHWAY

A Dissertation
Presented to
The Academic Faculty

by

Clare Ann Gollnick

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Department of Biomedical Engineering

Georgia Institute of Technology
August 2014

PROBABILISTIC ENCODING AND FEATURE SELECTIVITY IN THE SOMATOSENSORY PATHWAY

Approved by:

Dr. Garrett B. Stanley, Advisor
Department of Biomedical Engineering
Georgia Institute of Technology

Dr. Ravi V. Bellamkonda, Co-Advisor
Department of Biomedical Engineering
Georgia Institute of Technology

Dr. Astrid A. Prinz
Department of Biology
Emory University

Dr. Peter Wenner
Department of Physiology
Emory University

Dr. Steve M. Potter
Department of Biomedical Engineering
Georgia Institute of Technology

Date Approved: June 10, 2014

ACKNOWLEDGMENTS

Science is not done in a vacuum. I would like to thank the many people who have made the work in this dissertation possible.

First, I must thank my primary advisor, Garrett Stanley. Garrett took me on as a student and supported me even when my project was not on an upswing. I could not have done this without your support of my unconventional, and at times inconvenient, ideas.

I would also like to thank all my lab mates, past and present, in both the Stanley Lab and the Bellamkonda lab. I have learned so much from each of you. In particular, Daniel Millard for always being willing to hear a new idea, or to let me use your data to test a hypothesis. I have also relied a great deal on the support from my friends in the Neurolab for both general support and spirited debate. To Amber Willett, Jenny Munson, Katie McNeeley, Sharon Norman, Ashley Allen, and Alex Ortiz: thank you for your support through the numerous ups and downs that come with the tumultuous research and discovery process.

Outside of lab, I have fantastic and supportive friends, some of whom are in the lab with me, but many who are not. I would like to thank Sarah Forte, Cindy Bower, and Andrea Salem for years of keeping my head in the real world while still putting up with my science musing. To the whole crew: all the great times throughout the past six years have made my time in Atlanta very special. I have to thank Christian Lease, specifically, for the many hours he let me talk about whiskers on the road trip to Andrea's wedding in

Detroit. Ideas can come at the strangest of times. To Becca Hamel, you may have left graduate school earlier than I did, but you made my first few years in Atlanta fantastic.

I would be remiss not to thank my boyfriend Prashanth Challa. Words will never be sufficient to express how much I needed your support during these six years. I look forward to many more. I would also like to thank my sisters Beth and Katie, who have known me for so many years and always pick up the phone to work through a problem. To my parents, thank you for all you have done for me.

Even to those I have failed to mention explicitly, thank you. Every single experience has helped shape my views, ideas, taught me new things and helped me grow as a scientist. I look forward to working with all of you in the future.

TABLE OF CONTENTS

vi

x

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Our sensory experiences are encoded in the patterns of activity of the neurons in our brain. While we know we are capable of sensing and responding to a constantly changing sensory environment, we often study neural activity by repeatedly presenting the same stimulus and analyzing the average neural response. It is not understood how the average neural response represents the dynamic neural activity that produces our perceptions. In this work, we use functional imaging of the rodent primary somatosensory cortex, specifically the whisker representations, and apply classic signal-detection methods to test the predictive power of the average neural response. Stimulus features such as intensity are thought to be perceptually separable from the average representation; however, we show that stimulus intensity cannot be reliably decoded from neural activity from only a single experience. Instead, stimulus intensity was encoded only across many experiences. We observed this probabilistic neural code in multiple classic sensory paradigms including complex temporal stimuli (pairs of whisker deflections) and multi-whisker stimuli. These data suggest a novel framework for the encoding of stimulus features in the presence of high-neural variability. Specifically we suggest that our brains can compensate for unreliability by encoding information redundantly across cortical space. This thesis predicts that a somatosensory stimulus is not encoded identically each time it is experienced; instead, our brains use multiple redundant pathways to create a reliable sensory percept.

CHAPTER 1   Introduction

1.1   Sensory Neural Processing

A central tenet of neuroscience is that all our sensory experiences are encoded directly in the activity of our neurons. Neurons communicate with each other through spikes of electrical activity known as action potentials. We study the patterns of action potentials in response to different sensory input (or stimuli) to try to understand how the neurons encode this essential sensory information.

The patterns of action potentials form the basis of what is called the neural code (Stanley, 2013). Understanding how neurons communicate will allow us not only to interpret signals that we observe in the brain – *to read the neural code*; but could also allow us to engineer devices to communicate with neurons, inputting information directly into the brain or *writing the neural code*.

However, neurons are noisy. When we record from a neuron or populations of neurons, even in response to identical stimuli, the patterns of neural activity are not always the same. The same sensory input may be encoded differently based on the arousal level (Ferezou et al., 2006; Petersen, 2007; Carandini et al., 2012), a behavioral goal (Wang et al., 2010; Ollerenshaw et al., 2014), or simply just due to random and unexplained biological variability. Differentiating randomness from meaningful differences in the

1

neural response is essential to understanding which type of neural activity is sufficient for information processing.

The context dependency and unexplained variability of neural activity makes interpreting neural code a difficult task. However, there are mathematical or experimental ways to tease apart which aspects of neural activity are necessary, which are sufficient, and which are not contributing to the creation of a stable perceptual representation of the external environment. In this work, signal-detection theory will be used to quantify information in the neural code. The central goal of this dissertation will be to measure and quantify the contribution of variability in specific experimental frameworks and to hypothesize how variability may contribute to context-dependent processing in our sensory systems.

1.1.1 Feature Selective Neurons: Evidence for Receptive Fields

The best working model for how the information could be encoded by patterns of action potentials was formulated following a seminal discovery made by Hubel and Wiesel in 1959 (Hubel and Wiesel, 1959). At the time, neuroscientists had already observed that different regions of the brain responded primarily to certain types of sensory stimuli. Specific areas of the cortex were known to respond to information from the eyes and others to pressure stimuli on the skin. Additionally, cells in the retina had been observed to be sensitive to dots specifically in a center-surround arrangement: light surrounded by the absence of light, or the reverse (Kuffler, 1953). Cortical neurons, however, had not been observed to respond to dots of light and responded minimally to diffuse light. Hubel and Wiesel were mapping which region of visual space activated neurons when they

2

observed that the cell they were recording from often fired an action potential (spiked) when they switched between slides. The cell was sensitive to something about the movement of the slide. This observation eventually developed into the broader concept of a feature-selective receptive field.
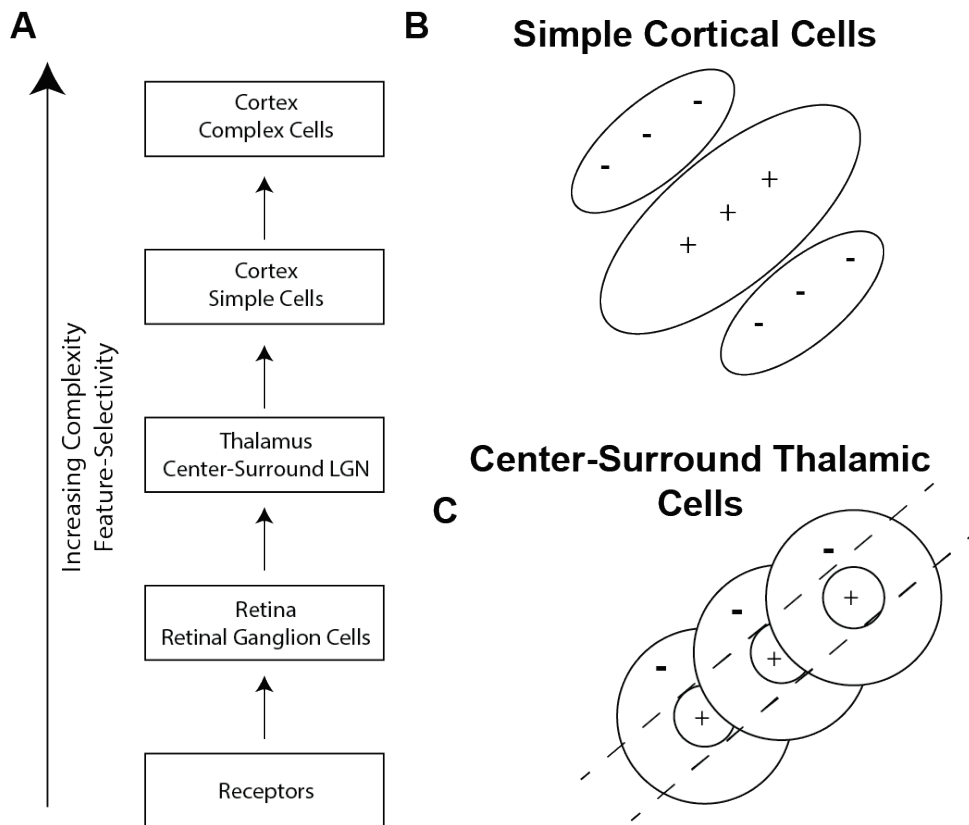
Figure 1-1 Hubel and Wiesel Model of Increased Complexity A) A hierarchical arrangement of neural structures in the visual pathway allows increased feature-selectivity to arise with each structure. B) A simple cortical cell feature-selective receptive field can be created from center-surround selective cells in the thalamus as shown in (C)

3

The feature of the stimulus that activated Hubel and Wiesel's neuron was the movement of lines at specific orientations, now widely known as 'orientation tuning'. Hubel and Wiesel also proposed that the feature-selective receptive cells in the visual cortex could be approximated as combinations of the type of receptive fields observed in the input to the cortex, the lateral geniculate nucleus of the thalamus (Figure 1-1) (Hubel and Wiesel, 1962). The result is a hierarchical model in which the specificity and selectivity of cells become progressively more complex and stimulus specific when processed through higher and higher order neural areas. With the discovery of feature-selective receptive fields and the model that more complex receptive fields could be formed as combinations of simpler receptive fields from input from previous neural areas, the field of sensory processing started to thrive.

1.2   Neural Processing in Somatosensation

1.2.1   Human Somatosensation

The somatosensation (touch) pathway shares many characteristics with human vision. Similar to vision, touch information travels through multiple hierarchical structures starting from the peripheral mechanoreceptors through the spinal cord, into the corresponding nuclei in the brainstem and thalamus before reaching the primary somatosensory cortex.

Somatosensory neurons respond selectively to one region of stimulus space (a specific region of the skin) known as the receptive field. Groups of neurons responding to similar or neighboring skin are arranged accordingly in cortical space, resulting in a topographic

4

map on the surface of the brain (Merzenich et al., 1978, 1983). Despite this one to one mapping of skin area to cortical area, size of the cortical space that is devoted to a specific body region is not correlated with the size of the area of skin but the number of nerves that innervate that area. As such, human body parts that we consider most sensitive to touch such as fingertips and lips represent a much higher percentage of cortical space than their small physical size would predict (Penfield and Boldrey, 1937). In general, skin regions that occupy a greater percentage of neural space (usually cortical area) than could be predicted from their physical size alone are also the most behaviorally relevant to the species.

Unlike vision however, there is not strong evidence for the existence of feature-selective receptive field properties in somatosensation, beyond the physical region of space to which a neuron is sensitive. There is no orientation-tuning equivalent, at least not in size or scale that is observed in the visual system (Hyvärinen and Poranen, 1978; Hsiao et al., 2002). The physical area of a skin that a neuron is responsive to increases with the hierarchical arrangement (Hyvärinen and Poranen, 1978). Neurons in the periphery respond to pressure stimuli in a small region of skin. Higher up in the processing pathway, the receptive field of cortical neurons can cover multiple fingers or regions of the hand (Sripati and Yoshioka, 2006). In this way, one could consider the receptive field complexity to be increasing; however, this is not related to specific properties of the stimulus such as orientation tuning in the visual cortex.

Despite lacking direct evidence for feature-selectively in the somatosensory system, the Hubel and Wiesel model of hierarchical complexity remains the best working model. A significant percentage of the research in the field of somatosensation, including in this dissertation, is centered on identifying and characterizing possible examples of feature-selectivity.

1.2.2 Rodent Vibrissa System as a Model System of Somatosensation

In this work, we study the neural code in the context of somatosensation. As with almost any biological experiment, it is not possible, or at least not ethical, to perform invasive experiments directly on the human system we seek to understand. Instead, we use animal models to approximate human somatosensation. A commonly studied and powerful animal model of somatosensation processing is the rodent vibrissa (whisker) system.

Neural processing of information from whiskers has many similarities to the neural processing of the human somatosensory system. While whiskers themselves are hairs, the follicle of the whisker is innervated with hundreds of mechanoreceptors (Rice et al., 1986; Ebara et al., 2002) that are similar to those observed in the human skin. These mechanoreceptors are sensitive to movement of the base of the whisker and transduce the energy from physical deformations into electrical neural activity. Whisker sensations are also processed through a similar hierarchy of neural structures from the peripheral whisker follicle to the brain stem, to the corresponding thalamic nucleus (ventral posteromedial nucleus, VPM), then to the primary somatosensory cortex and higher sensory areas (see Figure 1-2).

6

The cortical whisker representations are vastly overrepresented in the primary somatosensory cortex of rodents compared to their physical size (Land and Simons, 1985). This over-representation of cortical space is taken as evidence that the whiskers are exceptionally important to the sensory experience of the rodents on the same scale as lips and fingertips in humans. This is a reasonable assertion as rats and mice are nocturnal and therefore much less likely to depend on visual input than humans. Additionally, rodents are observed to use the whiskers to probe the environment (Carvell and Simons, 1990; Bermejo et al., 2002).

Figure 1-2 Barrel Cortex Introduction A) A schematic of the hierarchical pathway of the rodent vibrissae system. Whisker-related activity travels from the whisker pad to the brainstem to the thalamus and then to the cortical barrels. Figure adapted from Diamond et al., 2008) and reproduced here with permission. B) A cytochrome oxidase stain of tangential section through layer 4 of the cortex reveals the columnar and discrete arrangement of whisker barrels. Scale bar is 500um.

The physical size of the whisker representation makes the vibrissa system an appealing model system, but it has an additional experimentally helpful characteristic: it is discrete. Each whisker on the face maps directly to a specific region of the neural space. A single whisker representation is known as a barrelette in the brain stem, a barreloid in the VPM and most famously a barrel in the primary somatosensory cortex. These discrete columnar representations can be visualized in the cortex using a cytochrome oxidase stain as seen in Figure 1-2B. As whiskers are discrete physical structures, it is possible to activate one column unambiguously by the deflection of a specific whisker, without stimulation of peripheral receptors in the neighboring whisker follicle. Interestingly, the rest of the somatosensory cortex in both rodents and humans is also arranged into columnar structures like barrels, but these columns are much smaller and less prominent than the whisker representations and the input is not as easily separable (Mountcastle, 1997; Tommerdahl et al., 2010).

1.3   Sensory Processing in the Vibrissa System

1.3.1   Absence of Feature-Selectivity in the Rodent Barrel Cortex

In contrast to Hubel and Wiesel's groundbreaking discovery of orientation tuning for neurons in the primary visual cortex (but similar to the current understanding of human somatosensation), there are not feature-selective properties of neurons in the barrel pathway (although this is disputed). As mentioned above, a neuron is often sensitive primarily to a specific whisker (the primary whisker) and much less responsive to neighboring whiskers (Simons and Carvell, 1989; Higley and Contreras, 2005). This is

8

similar to a visual neuron being sensitive to a specific region in visual space. It is not equivalent to a neuron being specifically activated by something about the stimulus that occurs within that space, such as the direction of motion or orientation of a line.

Previously, some electrophysiologists have presented evidence for direction-tuning or angular-tuning in neurons in the barrel cortex (Simons and Carvell, 1989; Lee and Simons, 2004; Kerr et al., 2007); but this evidence is controversial and inconclusive. Importantly, recordings from pairs of neurons in the thalamus to the cortex show that connected neurons are not sensitive to the same direction, making it unclear how this information could be propagated along the pathway (Bruno et al., 2003). Others have postulated that a subset of neurons may be position-selective, acceleration (but not velocity) selective (Petersen et al., 2008), sensitive to more complex patterns of stimulation (Estebanez et al., 2012), or show frequency-specific activation (Ritt et al., 2008). However, even in these studies the percent of the neural population thought to be sensitive to any one feature is small (often 20% or less).

9
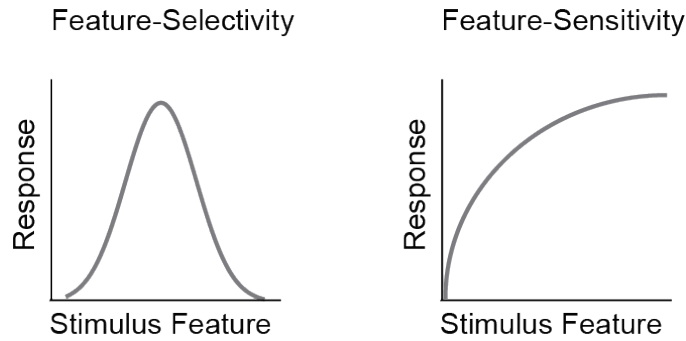
Figure 1-3 Feature-Selectivity versus Feature-Sensitivity Left: Feature selectivity requires a local maximum when some feature parameter is plotted along the x-axis and a neural response on the y-axis. Right: Feature-Sensitivity exists when no local maximum is observed however the neural response is modulated in some fashion by the feature parameter.

There is a subtle, but important, semantic distinction between feature-sensitive neurons and feature-selective neurons (Figure 1-3). Barrel cortex neurons are *velocity-sensitive*, as base measure of stimulus strength. If a whisker deflection is stronger (faster) then neurons fire more spikes on average (Simons, 1978; Wang et al., 2010). Importantly, this curve is monotonic. The neurons are not *velocity-selective* or *velocity-tuned,* responding maximally to a specific intermediate velocity and not to stronger velocities. This is in opposition to way that V1 neurons are specific to a particular orientation.

1.3.2 Behavioral Evidence for Whisker Information Processing

A large part of the confusion and controversy surrounding the lack of obvious feature-selectivity in the whisker pathway stems from a very practical problem. As human scientists, we do not have whiskers and consequently do not experience whisker sensations. We cannot immediately know which aspects of stimuli are meaningful or

10

even discriminable using whiskers alone. One of the central tenets of the study of sensory neural processing is that relevant perceptual information must be represented in some way in the activity of sensory neurons. With whiskers, we must first experimentally determine which features are perceptually relevant before searching the neural representations. The debate surrounding what type of information is available from whiskers has strong implications for development of theories about the corresponding neural code. As such, a careful consideration of recent behavioral paradigms is required.

Testing the perceptual capabilities of rodents is not a simple task. It was only recently that the experimental paradigms that allowed for careful and controlled analysis of whisker sensations were developed (Carandini and Churchland, 2013). Rodents, and most often rats, are clearly capable of performing complex behaviors using only whisker-related sensory information. Behavioral paradigms have shown that rodents are capable of responding more to deflections of increasing velocity (Stüttgen and Schwarz, 2008; Schwarz et al., 2010; Ollerenshaw et al., 2012, 2014). Additional studies show that rats are capable of basic object localization (O'Connor et al., 2010a, 2010b), texture discrimination (Wolfe et al., 2008; Jadhav et al., 2009; Diamond, 2010) and aperture discrimination (Krupa et al., 2001) all using information collected through the whiskers.

Intriguingly, while each of these experimental tasks demonstrates complex behavioral capabilities using information from whiskers, each paradigm can in fact be reduced to a simple detection task. For instance, in the experimental paradigm requiring object localization, rats were observed to move their whiskers only in the specific location of the

11

rewarded target, ignoring the other non-rewarded target positions (O'Connor et al., 2010a). While this did result in correct object localization behavioral performance, as the rats were able to report correctly whether the object was in the preferred location, the only behavioral strategy required was to detect or fail to detect an object in a specific searched location.

A rodent's ability to discriminate between different stimuli by adopting a strategy in which conditions are differentially detectable is particularly intriguing in the consideration of a recent texture discrimination task (Morita et al., 2011). Rats were trained to report whether a stimulus was textured (sandpaper) or smooth. The coarseness of the sandpaper determined the likelihood that the rodent correctly discriminated between the two stimuli, with the rodent responding correctly on more trials with coarser sandpaper (Wolfe et al., 2008; Jadhav et al., 2009; Morita et al., 2011). Interestingly, the ability to perform this type of discrimination was correlated with the presence of short high-velocity events, called 'slip-stick' events in which a whisker gets temporarily caught on the sandpaper, called a 'stick', followed by the release of that whisker, called a 'slip'. The discriminability between smooth and textured surfaces was also reduced to a detection task (Waiblinger et al., n.d.; Wolfe et al., 2008; Diamond, 2010).

The observation that all complex whisker-tasks reduce to a simple detection task greatly informed the work in this thesis. The key variable thought to determine the detectability of a whisker stimulus is the deflection velocity (Stüttgen and Schwarz, 2008; Ollerenshaw et al., 2014). I hypothesize that by understanding velocity encoding in the

12

whisker barrel pathway, we can understand what type of information is processed by whiskers and how it is encoded.

1.4   Experimental Approach

1.4.1   Application of Signal Detection Theory

Again, the central assertion of sensory neuroscience (and this work) is that features that are relevant to perception must be able to be reliably encoded (and thus decoded) from the patterns of neural activity. However, what does it mean to decode neural signals? We must define which feature of the signal is relevant, be capable of measuring that relevant feature, and consider the variability in this measurement due to noise.

The relevant neural feature is usually firing rate. Firing rate is assessed using a peri-stimulus time histogram (PSTH) collapsed from recording multiple trials of spiking activity from single or multiple neurons in the brain. These PSTHs are then used to estimate both the expected mean and variability in the neural representations of different stimuli. A common metric, known as d' (d-prime) is defined by the following equation:

$$d' = \frac{\mu_1 - \mu_2}{\sqrt{1/2(\sigma_1{}^2 + \sigma_2{}^2)}}$$

where μ_1, μ_2 and σ_1, σ_2 are the mean and standard deviations of the two distributions.

13

Figure 1-4 Signal Detection Theory d'   Distributions with a large d' are expected to be easily separable as there is minimal overlap. A small d' represents increased ambiguity and overlap of the distributions. The two signals are less discriminable.

An example of distributions with a large d' and small d' are shown in Figure 1-4. In the study of neural coding, distributions that are separable (a high d') would be interpreted to be reliably encoded and assumed to be perceptually different. Meanwhile distributions with a small d' would not be expected to be perceptually different.

We can also test a proposed encoding framework by classifying observed neural responses into stimulus categories and recording the performance. This type of approach is known as ideal observer analysis. The ideal observer analysis can be more widely applicable than the d' analysis as the classifier does not necessarily need to be Gaussian. An ideal observer classifies trials based a specific rule; often this rule is maximizing the Bayesian estimation of likelihood. The likelihood of an observed neural response, r, having resulted from any one of a possible set of stimuli (S_i) is calculated as:

14

$$P(S_i|r) = \frac{p(r|Si)p(S_i)}{p(r)}$$

The observer then classifies the observed response by maximizing the likelihood function, or for a discrete set of stimuli, choosing the stimulus that had the maximum likelihood of creating the observed response.

Using these signal-detection and classification frameworks, we develop and test the predictive power of different possible encoding schemes in neural networks.

1.4.2 Discrimination Using Single Trials as Opposed to Trial-Averages

In classic neural decoding frameworks, a central assumption is that the same stimulus will be encoded similarly each time it is presented. The neural representation varies based on context or state, but these differences would be predictable if the right variables are known. In this thesis, I used the same basic signal detection framework without this central assumption. I do not assume that distributions created from trial-average assessment of activity adequately represent single trials. Instead, I explicitly test this assumption using a maximum likelihood ideal observer analysis. Using an optical imaging technique known as voltage sensitive dye imaging (introduced in Chapter 2), I can record membrane potentials across a population in space with only one presentation of the stimulus. I can then test the discriminability of stimuli without making assumptions about the nature of single trial distributions.

15

1.5   Motivation and Thesis Organization

1.5.1   Thought Experiment: whiskers are not eyes

One unique characteristic of the whisker system as opposed to other sensory systems: *whiskers fall out*.

Eyes and retinas do not fall out. Ears and cochlea do not fall out. Whiskers are hairs. Similar to the hairs on our heads and bodies, whiskers go through recurring growth cycles that eventually result in an old whisker being replaced by a new whisker. While in many cases, the new whisker grows alongside the old whisker, it is common to see a rat missing a whisker, sometimes more than one (Young and Oliver, 1976).

While this may appear obvious, the neurobiological implications of this process should not be underestimated. First, if whiskers regularly fall out, regrow or change lengths on a day-to-day basis, one would also predict that the sensation from these changing whisker properties would alter the sensory percept. If each whisker were responsible for encoding of one specific property, say specific frequency information, then when that whisker falls out the animal would instantaneously lose all stimulus specific information associated with that frequency. As no human has directly experienced whisker sensations, it is possible that the sensory perception changes day-to-day; however, the sensory percepts we experience in other modalities do not seem so mutable.

Second, a missing whisker is a type of temporary sensory deprivation. We know from studies of sensory systems, including the barrel cortex, that sensory deprivation can cause

16

dramatic changes to the neural response (Wiesel and Hubel, 1963; Fox, 2002; Feldman and Brecht, 2005). Specifically, feature-selectivity is thought to develop as a direct result of constant, repeatable pairing of certain groups of stimuli. As the classic Hebbian saying goes: neurons that fire together, wire together. If the barrel cortex is going through random and recurring temporary sensory deprivation, it is logical to predict that meaningful neuron specific feature-selectivity could not be maintained within a whisker column.

With this thought experiment in mind, I developed the probability of activation hypothesis as a testable model for information coding when the sensors themselves are unreliable. If I were to engineer a system in which no individual input device could be expected to function identically (or even exist) on a given trial, the engineered system would need to be robust to this design specification. *Redundancy would be a central feature of this system*. I hypothesize that relevant stimulus information must be redundantly encoded, but still be sufficiently adaptable to allow for perceptual differences between stimuli. This will become a theme in the interpretation and proposed model of sensory coding in this work.

1.5.2   Organization of the Chapters

Despite the prominent role of the barrel cortex in the study of somatosensory processing for decades, there is still no conclusive evidence for feature-selectivity in receptive fields in the rodent vibrissa system. Hubel and Wiesel's pivotal discovery of orientation tuning in vision may in fact have led the whisker and somatosensory community astray. If the

17

absence of evidence of feature-selective neurons in the barrel cortex is interpreted as evidence of absence, it opens up the possibility for novel information processing schemes consistent with the unique properties of somatosensation. In this work, I propose a novel hypothesis to explain how stimulus information can be encoded such that complex behaviors can be performed without feature-specificity of individual neurons. I suggest that information can be encoded probabilistically across cortical space using multiple independent, but unreliable, whisker sensors. Initial evidence for the probability of activation hypothesis is presented in Chapter 2. The functional implications of the probability of activation hypothesis are then considered across time (Chapter 3) and space (Chapter 4). In Chapter 5, I present a framework where stimulus features including intensity and direction of motion could be encoded spatiotemporally as an array of unreliable detectors. In Chapter 6, I propose an accumulation of errors model and argue that probabilistic representations of stimulus strength not only possible, but also likely to occur as the result of a series of highly stochastic events.

18

CHAPTER 2   The Probability of Activation Hypothesis

2.1   Introduction

We use information derived from our sensory systems to interact dynamically with the external sensory world. Our senses help us avoid being hit by a car in a crosswalk or instantly recognize a face in a sea of strangers. Much of what we understand about neural processing of sensory stimuli has been derived from recordings of single or small populations of neurons averaged across repeated presentations of a sensory input. Trial-averaged neural activity, however, is not ethologically relevant; we do not often have the flexibility to replay an event or stimulus before reacting. Even when the average neural responses to different stimuli are significantly different, trial-to-trial variability can limit which stimulus parameters are effectively decoded (Abbott and Dayan, 1999; Averbeck et al., 2006; Beck et al., 2012). From a theoretical perspective the role of variability has been explored, particularly in the context of single neurons or the joint activity of pairs of neurons (Butts and Goldman, 2006; Churchland et al., 2010; Cohen and Kohn, 2011), but population-level variability remains poorly understood.

The rodent vibrissa (whisker) system is a powerful model system for detailed investigation of neural circuitry in early sensory pathways (Petersen, 2007), and is emerging as an important model system for behavior. Rodents are capable of choosing between different textures and patterns of whisker stimulation (Carvell and Simons,

1990; Wolfe et al., 2008; Morita et al., 2011) and of performing object detection and localization (O'Connor et al., 2010a, 2010b). Nevertheless, the specific stimulus features that are encoded in these tasks are not known. One hypothesis is that sporadic high-velocity 'slip-stick' events observed during texture discrimination as a whisker is moved across a surface are relevant stimulus features (Wolfe et al., 2008; Jadhav et al., 2009). Experiments in our laboratory and others have shown that both the behavioral detectability of whisker deflections (Stüttgen et al., 2006; Stüttgen and Schwarz, 2008; Ollerenshaw et al., 2012; Waiblinger et al., 2013) and the spike frequency in the thalamus and cortex (Simons, 1978; Pinto et al., 2000; Shoykhet et al., 2000a; Boloori et al., 2010; Wang et al., 2010) are modulated by short high-velocity whisker movements. However, single unit studies do not give direct information about population variability across-trials, which are critical for assessing the detectability of slip-stick events or the discriminability of different slip-stick events in different ethological contexts.

Here, we consider the encoding of whisker deflection velocity using in vivo voltage sensitive dye (VSD) imaging of the rat barrel cortex in an anesthetized preparation. VSD imaging measures population activity on a single trial with sufficient fidelity as to differentiate stimulus-evoked activity on a single stimulus presentation (Ollerenshaw et al., 2014). Consistent with previous VSD studies (Petersen et al., 2003a; Ollerenshaw et al., 2012), we show that the average amplitude of cortical activation increased with whisker deflection velocity. However, taking the perspective of an ideal observer of cortical activation, the single-trial cortical activity within a cortical column was not predictive of the velocity of the input. By examining the variability of the population

20

response, we find that the stimulus velocity instead modulated response reliability. We propose a new model for the encoding of stimulus strength in the presence of high trial-to-trial variability based on a probabilistic framework.

2.2 Methods

2.2.1 Animals

All procedures were approved by the Animal Care and Use Committee at the Georgia Institute of Technology. Data from six adult female Sprague-Dawley rats were using in this study (250g-330g). Data from four additional animals were excluded based on a predetermined criterion due to inadequate signal quality as reflected in the signal to noise ratio, predominantly caused by inadequate dye penetration.

2.2.2 Surgical Preparation

Animals were sedated with isoflurane (5%) and injected with sodium pentobarbital (50mg/kg, IP) for long-term anesthesia. A tail vein catheter was inserted and sodium pentobarbital (4.5mg/ml) was delivered continuously and adjusted to maintain an anesthetic depth at which no toe or tail pinch reflex was observed. Heart rate, oxygen saturation, respiratory rate, and temperature (37°C) were monitored continuously to ensure a constant level of anesthetic depth throughout the experiment. Once anesthetized, the animal was stabilized in a stereotaxic frame. A craniotomy was performed over the barrel cortex (stereotaxic coordinates: 0.5-4.0 mm caudal to the bregma, and 3.0-7.0 mm lateral to the midline), and the dura was removed. A dental-cement well was created

21

around the craniotomy to hold dye and saline. The surface was kept moist with sterile dye or saline throughout the entire experiment.

2.2.3 VSD imaging

Voltage-sensitive dye (VSD) imaging measures subthreshold activity from a population of neurons in layer 2/3 of the barrel cortex (Kleinfeld and Delaney, 1996). Voltage-sensitive dye (RH1691, 2mg/ml, Optical Imaging) was placed on the surface of the brain and mixed regularly with a micropipette every 5-10 minutes for 1.5 - 2 hours in order to allow sufficient time for diffusion into the cortex. Next, the unbound dye was removed by multiple washes with sterile saline. The cortex was illuminated with a 150W Halogen lamp passed through an excitation filter (621-643 nm). The camera was focused approximately 300μm from the surface of the brain. Images were recorded with a high-speed CCD camera (MiCOM2, SciMedia) as depicted in Figure 2-1. A 1x objective lens was combined with a 0.63X condenser lens, resulting in a total magnification of 1.6X. Forty to sixty trials for each stimulus condition were recorded for analysis. Image resolution was approximately 20μm/pixel.

2.2.4 Whisker Deflections

Controlled whisker deflections were delivered with a piezoelectric bending actuator (range of motion: 1 mm; band- width: 200 Hz; Polytec PI) attached to a glass pipette. The whiskers were trimmed to 15mm and the whisker tip was placed inside the pipette, which was positioned 10mm from the whisker base. The saw-tooth deflection consisted of an eight millisecond exponential rise phase, and an eight millisecond exponential decay

22

phase. The velocity of the whisker deflection was varied by changing the rostral-caudal distance the whisker was moved by the piezoelectric actuator. The duration of the deflection remained constant. The angle of deflection at the whisker base was calculated from the measured distances (distance from whisker pad, deflection distance) assuming triangular geometry. The six velocities (V_1 to V_6) were calibrated to be 75, 150, 300, 600, 900 and 1200˚/s, respectively, in the rostral-caudal plane. In most animals, noise trials in which no stimulus was delivered were also recorded. In data sets in which noise trials were not recorded, forty pre-stimulus frames were used to approximate noise trials. Dual whisker deflections were delivered using two individually calibrated piezoelectric actuators at equal velocity.

2.2.5 Image Analysis

All data analysis was done in custom written software in MATLAB (MathWorks, Natick, MA). For each trial, a background image (F_0) was created by averaging forty pre-stimulus frames (a 200ms interval). The stimulus evoked VSD signal was calculated as the percent change from this background image: $\%\Delta F/F = (F - F_0)/F_0 * 100$. Associated time series were calculated by averaging over a circular region of interest with a ten pixel radius ($\sim$200μm) from the center of activation in the onset frame (determined manually, usually 10ms or 15ms after stimulus delivery) from an average of all trials. The onset frame has previously been shown to be restricted to a single cortical column (Petersen et al., 2003a). The use of a spatially restricted ROI resulted in an increased signal to noise ratio and allowed the detection of even small activations that were stimulus-evoked. For calculation of the center of mass and image presentation only,

23

images were filtered with a spatial averaging filter (400μm X 400μm, approximately 1 barrel). In order to display response amplitude for all data sets on the same scale, single trial response amplitudes were normalized by the max single trial response amplitude observed in each data set. Average response amplitudes were normalized by the maximum average response.

2.2.6 Ideal Observer Classification by Stimulus Velocity

An ideal observer was tasked with classification of trials by response amplitudes (r) into six stimulus velocity categories V_1, V_2…V_6 (Wang et al., 2010). For each velocity, a probability distribution, $P(r|V_i)$, was estimated using bootstrapped estimates of the mean. The set of six velocity distributions formed the trial-average classifier. A given observed response amplitude, r, was sorted into the velocity class with the maximal likelihood estimator:

$$\hat{V}_{MLE} = \arg max \ _{V_i} \ P(r|V_i)$$

The result was a 6x6 performance matrix in which the [j,k] element was the frequency of assigning a single trial to V_k when the actual stimulus was V_j. The *actual classification performance* was defined as the frequency that a given trial was correctly classified, summed across stimuli, or simply the average of the diagonal of the performance matrix.

The *optimal classification performance* for a given classifier was limited only by the overlap between the velocity distributions. Conceptually, it was the frequency of

24

correctly classified trials if the response amplitudes were sampled directly from the classifier distributions:

$$Optimal\ Performance = 100 * \sum_{V_i} \int_{r \in R_i} P(r|V_i)\, P(V_i)\, dr$$

where R_i is the range of response amplitudes over which velocity, V_i, had the maximum probability and $P(V_i) = 1/6$, the relative frequency at which a given velocity stimulus was presented. Extreme errors were trials in which the assigned velocity class V_k differed from the actual stimulus class, V_j, such that $|k - j| \geq 2$. For simulations of multiple stimulus presentations, the maximum likelihood estimation was calculated from the effective response amplitude $\hat{r}$, which was the mean of n (1,3 or 10) randomly selected trials.

2.2.7 Response Reliability Analysis

Single trial time series were calculated by averaging over the same barrel-sized ROI used to create average time series described previously. Each trial was determined to be either a 'response' or 'no-response' trial based on the presence or absence of a measurable stimulus-locked change in amplitude using a template-matching algorithm. As the temporal response was extremely stereotyped across velocities, an average time series from 0ms to 250ms after stimulus delivery was used as a filter (or template). The filter included frames in which the activity had not yet reached the cortex, to ensure that changes in fluorescence were stimulus-locked. Individual trials were matched to the

25

template by point-wise multiplication (dot product). The sum was a single number, a match score, which quantified how well each single trial matched the template. Match score distributions from both noise frames and stimulus trials were calculated. A classification threshold was set at three standard deviations above mean match scores of noise trials. Trials with match scores above this threshold were considered response trials. Note that the use of a high threshold (a match score above the 99.7 percentile of the noise distribution) resulted in a greater number of false negatives (a misclassification of a true response as a no-response trial) while limiting the probability of false-positive misclassifications. The high threshold level was chosen to highlight the existence of large response amplitudes to small velocity stimuli; however, the results presented in this paper were consistent for all threshold levels.

2.2.8 Probability of Activation Simulations

We discussed two possible models that could generate the observed average VSD responses, the continuum model and the probability of activation model. As a demonstration of the unique characteristics of each model, we simulated data from the mean and standard deviation of the combined normalized distribution from all data sets. These simulations are meant to demonstrate the qualitative differences between the models, not to be quantitatively representative of any specific parameter of actual data. For the continuum model, distributions were assumed to be normally distributed about the observed mean with a standard deviation equal to the standard deviation of the noise distribution. For the probability of activation model, response trials were defined as any trial with response amplitude greater than three standard deviations above the noise

26

distribution mean. The response and no response means and standard deviations were estimated accordingly. The simulated distributions were randomly chosen from these distributions.

2.2.9 Multi-Whisker Spatial Classification

A spatial template-matching algorithm was used to classify single trials from simultaneous whisker deflection. This algorithm was adapted from a previously published method (Millard and Stanley, 2013). In each data set (n=2 animals), trials (data set 1, 100 trials; data set 2, 50 trials) were collected for individual whisker deflections and for both whiskers deflected simultaneously. Trial-average and time average (15-20ms) images from individual whisker deflections and simultaneous whisker deflections were used to create spatial templates (T_1, T_2 and T_3 respectively). A representative barrel map of thirty-two barrel columns was registered to the image based on known whisker locations (Wang et al., 2012). The average activation in each barrel was calculated, creating a one-dimensional representation of the image (1XN array). N is the number of barrel columns (32 maximum) that fall entirely within the VSD image in each dataset. No response trials were modeled as zero mean noise (template T_4) with a covariance (Σ) estimated from pre-stimulus noise frames. Single trials (*y*) were time-averaged and mapped to one dimension in the same manner as for the templates, then sorted by choosing the least-squares template, T_{LS}, that minimized the weighted mean squared error:

$$T_{LS} = \arg \min_{\substack{T_i \\ T_1, \ldots T_4}} \ (y - T_i)' \ \Sigma^{-1} \ (y - T_i)$$

27

Evidence for independence was assessed using a chi-squared test with one degree of freedom.

Figure 2-1 VSD System Schematic   A) A schematic of the voltage-sensitive dye imaging system. A piezoelectric actuator delivered controlled mechanical deflections of a whisker with variable velocity. The surface of the cortex was stained with voltage sensitive dye and images were recorded on a high-speed CCD camera (5ms resolution).

2.3   Results

Population responses to whisker deflections were studied using voltage sensitive dye imaging (VSD) as illustrated in Figure 2-1. Single punctate whisker deflections of

28

different velocities were delivered by a computer-controlled piezoelectric actuator in the rostral-caudal plane (see Methods). Images were recorded every five milliseconds and the fluorescence was quantified as percent change from background. Figure 2-2A shows trial-average images for six different velocity whisker deflections (V1-V6): 75°/s, 150°/s, 300°/s, 600°/s, 900°/s, 1200°/s which span the expected behavioral threshold for detection(Stüttgen and Schwarz, 2008; Ollerenshaw et al., 2012). The neural activity started locally over the primary barrel approximately 10ms after stimulus presentation and then spread into neighboring barrels (Figure 2-2B). Time series shown in Figure 2-2C were created by averaging over a circular region of interest approximately the size of one cortical barrel (400μm diameter) corresponding to the deflected whisker. Increasing the velocity of the stimulus increased the trial-average peak amplitude (or simply response amplitude) of the VSD response (Figure 2-2D) across all experiments (n=6 whiskers, 4 animals). These data are consistent with previous experimental findings (Petersen et al., 2003a; Ollerenshaw et al., 2012; Wang et al., 2012).

Figure 2-2 Representative Trial-Average VSD Data A) Representative cortical responses to whisker deflections of variable velocity (V1, lowest). Images were averaged over 60 trials. 1 mm scale bar applies to all images. B) Corresponding time series to the images seen in (A). Time series were calculated by averaging over a circular region of interest the size of one cortical column (~400μm diameter) centered over the barrel corresponding to the deflected whisker. The peak response frame determines the response amplitude. C) The trial-average response amplitude for each data set (dotted lines, n=6 whiskers, 4 animals) and the resulting mean (black line with error bars ± s.e.m.) . Trial-average response amplitude increased with velocity in all data sets.

2.3.1 Trial-average differences were not sufficient to allow for single trial velocity discrimination

We tested whether the relationship between peak cortical response amplitude and deflection velocity was sufficient to classify single trial responses. For each data set, we define a trial-average classifier based on a set of response distributions created from bootstrapped estimates of the mean responses at each of six velocities (V1 through V6). A representative example of a set of trial-average distributions is shown in Figure 2-3A. We used maximum likelihood estimation to assign an observed response to a particular velocity class (see Methods). If the single trials were perfectly modeled by the trial-average classifier, and thus drawn from these distributions, the theoretical performance would be limited only by the area of overlap of each of the distributions (referred to here as the optimal classifier performance). Optimal classification performance is visualized in matrices showing the frequency of trials predicted to be in each class (columns) for each actual input velocity (rows). A representative performance matrix is shown in Figure 2-3B. In all our data sets, the optimal performance matrix has a strong diagonal representing a high number of correctly classified trials (72.9% ± 10.4%, mean ± 1 SD)

31
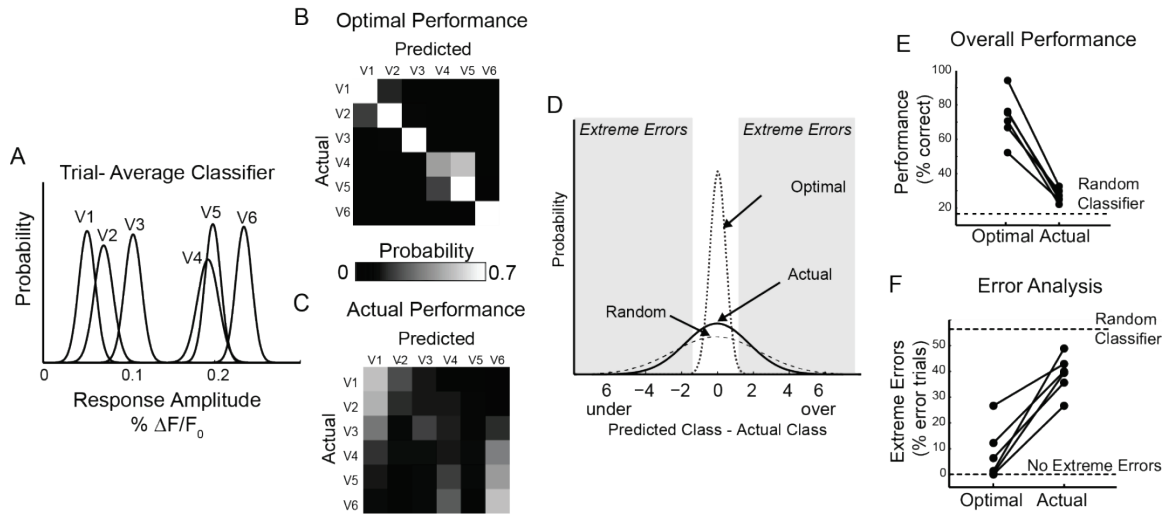
Figure 2-3 Ideal Observer Performance of Velocity Classification   A) Representative example of a trial-average classifier. B) The optimal performance of the classifier shown in (A) is limited only by the overlap of the velocity response distributions. Here, the probability of classifying an observation from velocity V_j as coming from V_k is represented graphically as element $[j, k]$ in the performance matrix. The presence of a strong diagonal represents a high number of correctly sorted trials. C) The actual performance matrix of the ideal observer given true single trial response amplitudes. D) Distribution of the errors associated with classification. Using the notation described above, the $= j - k$. E) The percent of trials correctly classified under both the optimal and actual conditions across all data sets. Optimal and actual performance from individual data sets are connected with a line. F) The percent of error trials that are extreme |j-k|≥2 in both the optimal and actual performance conditions.

Again taking the perspective of an ideal observer, we used the same decoding framework to classify observed single trials response amplitudes. The decoding performance was significantly less than optimal (paired t-test, n=6, p<0.001) and on average correctly classified 27.7% ± 3.7% of trials, which is only about 11% more than if the classification were purely random. The performance of each individual data set is shown in Figure 2-3E. It is not surprising that the classification was less than optimal, as bootstrapped

www.manaraa.com

estimates necessarily underestimate the true variability of the response; however, the actual performance of an ideal observer using this framework approached random classification, particularly when we considered the scale of the errors. In the example shown in Figure 2-3C, a trial from even lowest velocity (V1) was classified as one of the highest velocities (V5); similarly, a trial from the highest velocity was mistaken classified as the lowest (V1). To quantify this across data sets, we considered any trial that was misclassified by two velocity classes or more as an extreme error (Figure 2-3D). In the optimal case, the percent of errors predicted to be extreme was 7.8% with a standard deviation of 10.3%. The actual observed percentage of extreme errors was five times greater on average, 38.9% ± 7.4%. The difference is highly significant (paired t-test, n=6 whiskers from 4 animals, p<0.0001). The percent of misclassified trials that are expected to be extreme errors using a random classifier is 56.2%.

The poor performance of the ideal observer using a trial-average classifier suggests that the single trials are not well modeled by the trial-average response. In order to understand why, we considered single trial distributions directly. Figure 2-4A shows the time series associated with single trials from a representative data set. There is large variability in the response amplitude across all velocities. Consistent with the high number of extreme errors made when decoding with a trial-average classifier, the response amplitude from a single trial in response to a low velocity can equal the magnitude of the response from a trial observed in response to a high velocity. A jittered scatterplot shows the same data with a single data point representing the peak response amplitude frame (open circle) for each single trial at each velocity (Figure 2-4B). The mean response is indicated with a

33

solid black line. As previously observed, the mean increased with velocity; but a given

response amplitude can be the result of any velocity. The corresponding plot with all

single trials from all data sets (n=6 whiskers, 4 animals) is shown in Figure 2-4C. From

this representation of the data, it is clear that knowledge of the peak response amplitude is

not sufficient to classify individual trials into velocity categories.
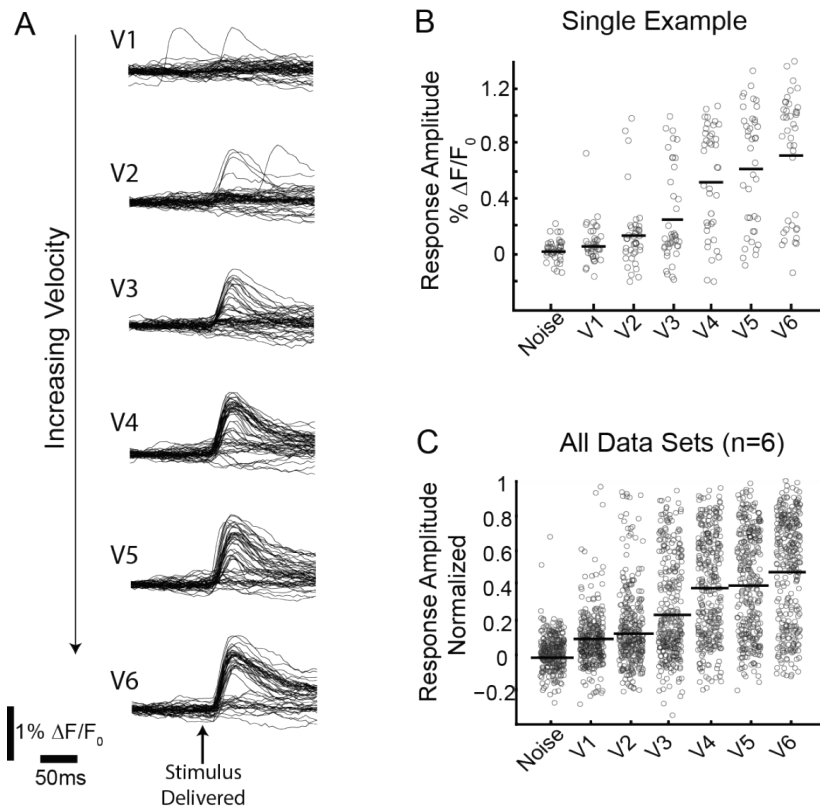
Figure 2-4 Single Trial Response Variability   A) All single trial VSD time series from a representative data set. V_1 is the lowest velocity while V_6 is the highest. B) The peak response amplitude from each single trial shown in (A) is represented as an open circle in the jittered scatter plot. C) All single trial response amplitudes from all six data sets are shown in a jittered scatterplot. Data sets are normalized such that the peak single trial from any data set has response amplitude of one.

34

2.3.2 Single-trial responses were well-described by a probability of activation model

It must be determined how, given such variability in the response amplitude, there are reliable differences in the trial-average amplitude between velocities across all animals. In all data sets, for a given velocity, we observed a subset of trials in which the stimulus-evoked response was not measurably increased from the pre-stimulus activity. These trials were distinctly different from the trials in which there was clearly an evoked response (Figure 2-5A). We sorted trials based on the presence or absence of a stimulus-evoked response on individual trials using a matched-filter algorithm (see Methods). We defined two categories: 1) 'response' trials in which post-stimulus fluorescence was distinguishable from noise and 2) 'no-response' trials in which it was not different from noise. An example set of sorted time series, as well as example corresponding single trial images, can be seen in Figure 2-5A. The magnitude of both response and no-response trials was variable (and not necessarily distinguishable by an absolute threshold), but the response trials all exhibited a clear, spatially specific activation over the primary barrel as well a stimulus-locked increase in fluorescence in the corresponding time series. No-response trials exhibited either non-specific background activity or no stimulus specific changes at all. In a small number of trials sorted as no-response trials, there was actually a very small spatially localized signal. We consider these trials to be response trials misclassified as no-response trials, but the change in activity was so small as to not be reliably distinguished from noise fluctuations (Figure 2-5B).

35
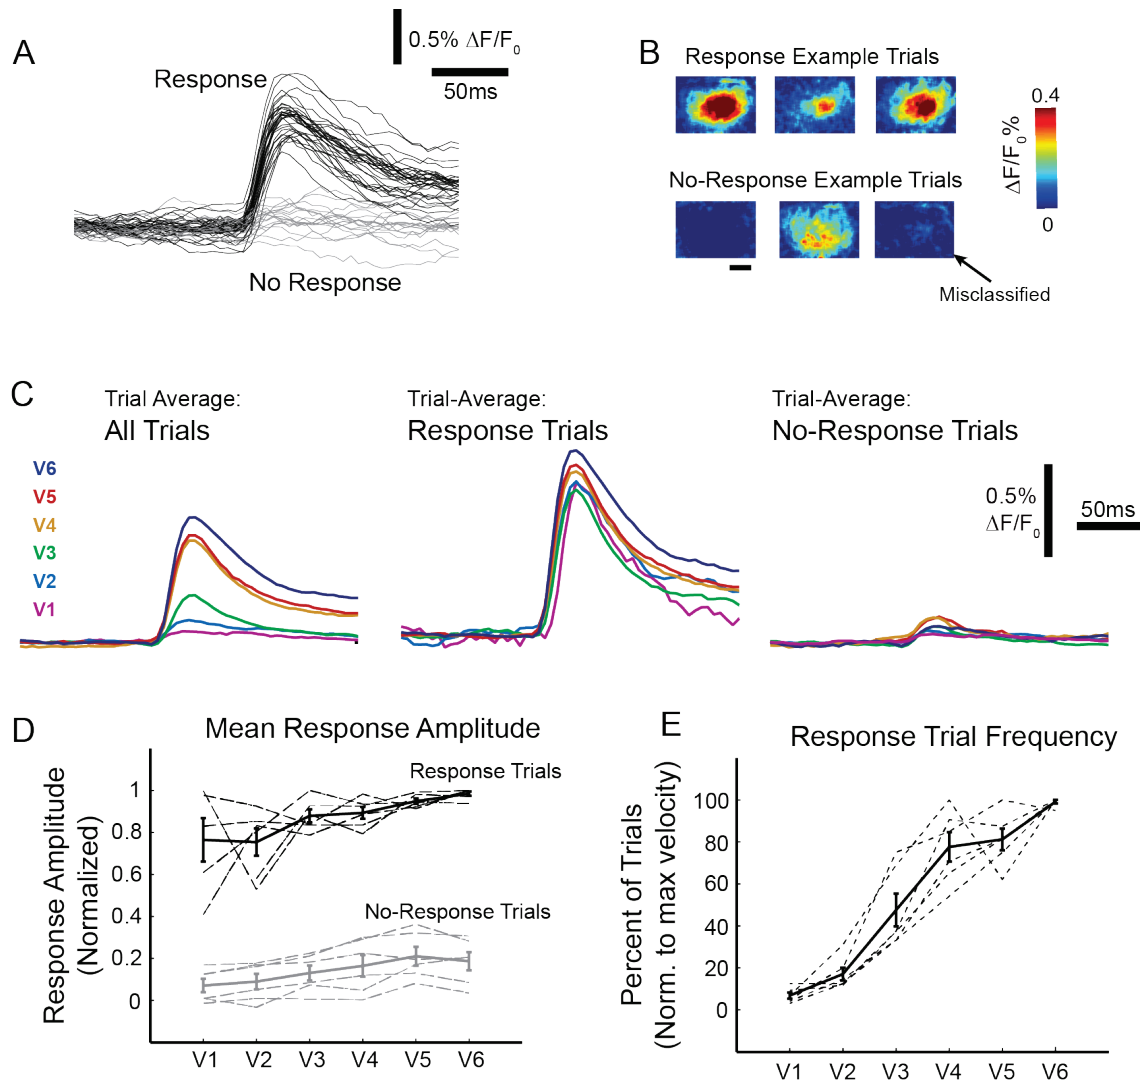
Figure 2-5 Response and No Response Dynamics of Single Trials A) Single trials were qualitatively and quantitatively separable into two groups: trials with stimulus evoked activity (*response trials,* black) and those without (*no-response*, gray) B) Example single trial images from trials shown in (A) demonstrate variability within the groups. Shown are three response trials with variable response amplitude (top). Three no-response trials (bottom) chosen to represent the types of trials that comprise this category, including a misclassified trial. C) When all trials were included in the average time series of a single data set, there were large differences between the velocities (left), but when considering the average of only response trials (middle) or only no-response trials (right) the differences were dramatically reduced. D) Results were consistent across all data sets (n=6 whiskers, 4 animals, dotted lines) for both the response (black) and no-response (gray) average responses. The mean ± s.e.m. for both groups is depicted with a solid line and error bars. Note that higher variability is associated with conditions with the fewest

number of trials (low velocities for response trials, high velocities for no response trials) E) The frequency of response trials increased with stimulus velocity for all data sets (dotted lines), as well on average (solid line, mean ± s.e.m.).

When considering response or no-response trials alone, the differences between the velocities were dramatically reduced. The ordered nature of the velocities was sometimes even lost: low velocities had larger mean response amplitudes than higher velocities, as seen in Figure 2-5C. The means of response and no-response groups for individual data sets are shown as dotted lines in Figure 2-5D. Once sorted, the absolute amplitude of response and no-response trials was no longer strongly modulated by the stimulus velocity. However, the number of trials sorted into the response category, which we call the response reliability, increased with stimulus strength (Figure 2-5E). This was true for every dataset collected. Note that the response reliability was considered relative to the velocity (not always the highest velocity) with the highest number of response trials for each data set. Importantly, no velocity in any dataset showed a perfect response rate. While the frequency of response trials always increased with velocity, the absolute response frequency was highly variable between data sets.

The presence of two distinct types of responses suggests a model for the encoding of the whisker deflection velocity that is distinctly different from what we had assumed in the original trial-average classifier. Without knowledge of the single trial variability, we had implicitly assumed a model in which the stimulus strength was encoded by the strength of the population response on a single trial (the trial-average classifier). We refer to this as

the continuum model as it proposes graded response with increasing deflection velocity (Figure 2-6, top). However, the same mean response profile can also be created from a discrete probability of activation model formed from two possible response categories, a response and no-response class (Figure 2-6, bottom). The differences in the mean response occur by changes in the number of observations sampled from each distribution.
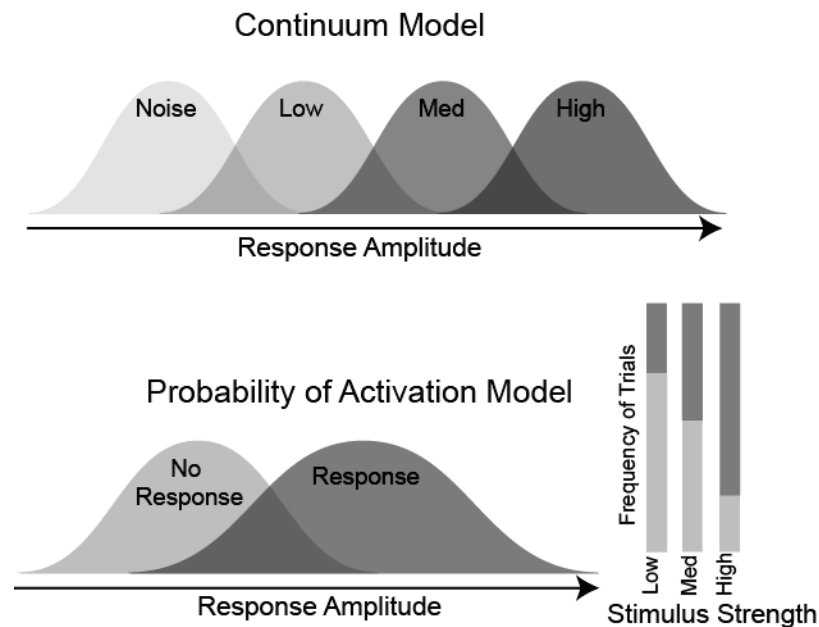
Figure 2-6 Two Possible Encoding Schemes The continuum model (top) has graded response amplitude with increasing stimulus strength (low, med, and high velocities). For the probability of activation model, (bottom) the frequency of response trials increases with increasing stimulus strength; however, the response and no response distributions were conserved across stimulus strength.

It is not possible to distinguish between the probability of activation model and the continuum model with only the trial-average response. For example, Figure 2-7 shows

two possible single trial distributions simulated from hypothetical realizations of both the continuum (Figure 2-7, top row) and probability of activation (Figure 2-7, bottom row) models. Both models result in the same trial average measurements (Figure 2-7, left column), making the source of the responses ambiguous. The primary differences between the two models are revealed in a more thorough examination of trial-to-trial variability. For the continuum model, in addition to increased mean response with velocity, three additional observations are evident. First, the minimum and maximum observed responses increased with the mean and with velocity (Figure 2-7B). Second, when pooled across velocities, the aggregate distribution of the responses is unimodal, reflecting the graduated increase in amplitudes with velocity (Figure 2-7C). Third, knowledge of the amplitude of a single trial response reduces the uncertainty about which of the stimuli could have produced that response. In contrast, the range of achievable responses does not change with increasing velocity in the probability of activation model (Figure 2-7H). When pooled, this distribution is bi-modal (Figure 2-7G). Although bi-modality is supporting evidence of a probability of activation model, it is not necessary, as it is also dependent on the overall difference between signal and noise distributions in a given dataset. Importantly, in the probability of activation model the single trial amplitude does not convey information, or reduce the uncertainty, about strength of the stimulus. Observed single trial distributions in Figure 2-7E and Figure 2-7F (same data as in Figure 2-4C) are more consistent with the probability of activation model.

Figure 2-7 Simulations of Single Trial Distributions A-C) Simulated continuum model distributions. A) Mean response amplitude versus velocity. B) A jittered scatter plot of the underlying single trials. C) A combined histogram of response amplitude of single trials from all velocities. The same analysis was done on the observed VSD single trials (D-F) and for simulated probability of activation model distributions (G-I). Notice that the range of achievable response amplitudes for each velocity is constrained in the continuum model, but not in the probability of activation model. The observed data is consistent with the probability of activation model.

2.3.3 Encoding stimulus strength within the probability of activation model required multiple independent sensors or events

The probability of activation model suggests that deflection is not encoded by the response amplitude within a single barrel. This is potentially troubling since velocity is likely an important parameter for sensory perception (Jadhav et al., 2009). Outside of controlled laboratory experiments, however, a rat would rarely perceive a stimulus, whether it is an object or wind stimulus, using only a single whisker. Additionally,

rodents are known to actively sense by repeatedly moving whiskers across objects, a behavior known as whisking (Carvell and Simons, 1990; Bermejo et al., 2002). Multiple observations of the same stimulus, whether on the same whisker over time or over an array of multiple whiskers, could provide an estimate of the strength of stimulus. For instance, a weak stimulus may activate two barrels out of ten barrels, or result in an observable activation in two out of ten whisks across an object. Similarly, a strong stimulus may activate all ten barrels or respond on every single whisk across an object. This qualitative hypothesis can easily be simulated by adapting the same maximum likelihood decoding framework used in Figure 2-3 to include multiple observations. We simulated having access to multiple observations by averaging one, three or ten single trials into a single response variable prior to classification. Overall accuracy of the classification increased with increasing number of observations of the stimulus. In Figure 2-8A, there are increasingly prominent diagonals in the performance matrices as the number of trials increased from one to three to ten observations. This was true across all data sets (Figure 2-8B).

41

Figure 2-8 Multi-Whisker Classification A) Example performance matrices for the optimal discrimination performance (top) and the performance matrix for an increasing number of observations (one, three, and ten observations). The same labels and scale bar apply to all performance matrices. B) Quantification of increased performance across all data sets (n=6 data sets, 4 animals). Bar graphs show mean ± s.e.m.

The key assumption in this simulation was that individual observations, whether across cortical space (multiple whiskers) or time (multiple whisks), are functionally independent. Temporal independence of stimuli is likely contingent on the inter-deflection interval and has been previously studied (Simons, 1985; Boloori et al., 2010). In the cortical space dimension, independence requires that it must be possible for one cortical column to responds while another does not at the same time. The functional independence of the response amplitude of neighboring cortical barrels within a single trial has not been previously considered. If the response/no response dynamics reported here were determined by a shared variable, potentially a non-specific global state

42

variable, both barrels would either respond or not respond on a given trial. In contrast, if the barrels were able to respond or not respond independent of the response of the neighboring barrel, the multiple observation frame work would be supported. To test this, we recorded trials in separate experiments in which two adjacent whiskers were deflected simultaneously. The trial-average VSD image from simultaneous deflection was spatially similar to summed responses of the two individual whiskers (Figure 2-9A). On some trials, even though both whiskers were physically deflected, the cortical response qualitatively resembled trials where only one whisker was deflected (Figure 2-9B). This suggests that two neighboring barrels can respond independently. Of course, it is difficult to prove that the two barrels are truly quantitatively independent. Independence depends on the exact frequencies with which we observe each type of trial, or more specifically on how well the observed frequencies match the probabilistic definition of independence: $P(A\&B)=P(A)P(B)$. To quantify frequency, we sorted the single trials in response to dual whisker deflection based on a spatial matched filter algorithm (see Methods) and recorded the number of trials that most resembled the spatial profile of four categories: each whisker individually (barrel 1 alone, barrel 2 alone), both barrels respond, and neither respond. The observed frequencies for two data sets (from two animals) are shown in

Table 1. The observed frequencies for both data sets were each consistent with a null hypothesis of independence assessed with a chi-squared test, as shown in the table. Taken together, this evidence suggests a probability of activation framework in which the strength of the sensory input is encoded not in a graded cortical response amplitude, but

43

instead in the probability of observing a response on a trial-by-trial basis; additionally, it appears that the individual cortical columns are independent enough to probabilistically encode strength or velocity information across the whisker array.
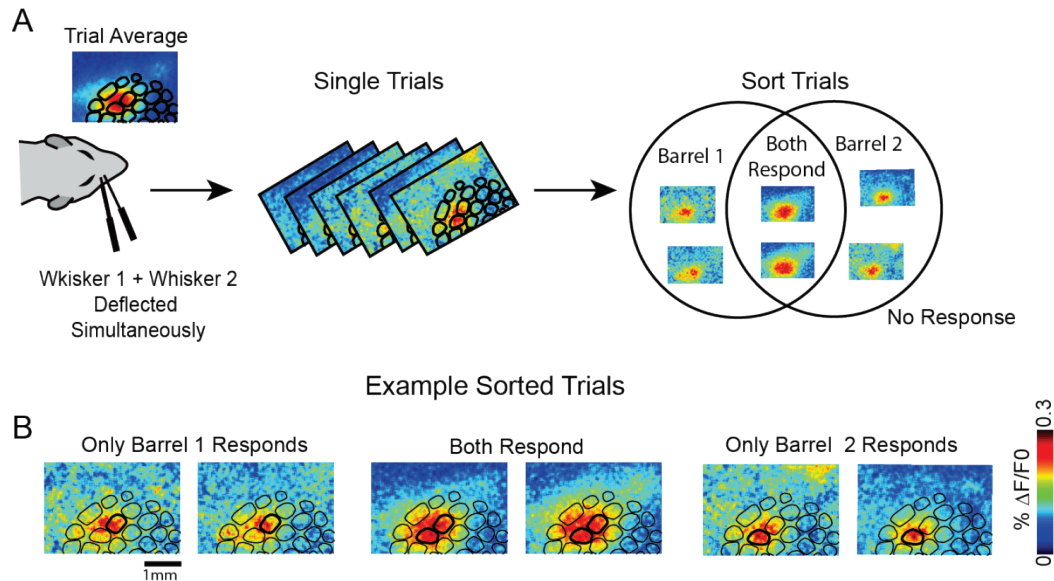
Figure 2-9 Barrels Can Respond Independently A) Schematic of the dual whisker experimental paradigm. Trial-average image from dual whisker deflections resembled a linear sum of two responding whisker barrels. Single trial responses exhibit variable response spatial profiles. Trials were sorted into four response categories: barrel 1 responds, barrel 2 responds, both barrels respond and no response. B) Example single trial images (15-20ms after stimulus presentation) from each of the three response categories. Scale bar (1mm) applies to all images.

2.4   Discussion

By examining the variability of the single trial VSD responses to whisker deflections of increasing velocity, we showed that the correlation between deflection velocity and mean VSD response amplitude results primarily from stimulus-dependent modulation of

response reliability, not the absolute response amplitude of individual trials. An ideal observer analysis showed that single trial response amplitudes were not sufficient to classify trials into velocity categories, but instead the observer often made large classification errors. While single trial response amplitudes were highly variable, there was structure in the single trial distributions. We observed that single trials from all velocities could be divided into two groups we refer to as response and no-response trials. Once sorted, differences between velocities were dramatically reduced. From these observations, we predict that a behaving animal would not be able to reliably discriminate stimulus velocities from the deflection of a single whisker. However, the spatial profile of single trial responses to the simultaneous stimulation of two adjacent whiskers resembled what would be expected if individual barrels acted as independent sensors of stimulus strength. This suggests that velocity discrimination would be possible with multiple independent samples of stimuli across multiple whiskers or repeated whisks of the stimulus.

Table 1 Evidence for Independence Frequencies are presented as probability estimates: the number of observed trials in each category divided by the total number of recorded trials.

	Observed Single Trial Spatial Profiles				Statistics	
	P(B1 not B2)	*P(B2 not B1)*	*P(B1 and B2)*	*P(NR)*	χ^2 (df=1)	p-value
Data Set 1	0.23	0.28	0.39	0.1	1.2385	0.265
Data Set 2	0.1	0.2	0.68	0.02	0.113	0.730

45

It is important to note that our data do not suggest an all-or-none type of bimodality, in which all response trials are identical. In contrast, the absolute response amplitude for all velocities was extremely variable, taking on any possible value: an anything-or-nothing response distribution. While some cortical responses from the lowest velocity were equal in magnitude to those from the highest, some responses from the highest velocities were just barely distinguishable from the noise distribution. This is an important distinction, suggesting that absolute amplitude of the stimulus is not reliably propagated through this sensory pathway. This variability is not an artifact from the VSD imaging technique itself or a result of changes in background activity as the level of variability observed with whisker driven input is in direct contrast to the highly repeatable and low-variability VSD responses observed in the same pathway in response to electrical microstimuluation of the thalamus (Millard et al., 2013).

It is important to note that VSD imaging is limited to activity in supragranular layers of cortex, primarily layer 2/3, due to physical limits of light scattering. We cannot rule out the possibility that this is a layer 2/3 specific response characteristic. This may even be likely as layer 2/3 is specifically characterized by minimal firing and high trial-to-trial variability (Petersen et al., 2003a; Kerr et al., 2007; Sato et al., 2007; Crochet et al., 2011). In these same studies, whisker deflections that do not cause spikes have been observed (Petersen et al., 2003b; Kerr et al., 2007; Sato et al., 2007). To our knowledge, the frequency of such trials has not been previously linked to stimulus strength.

46

Consistent with this interpretation, it is interesting to consider that in the primary input layer of the cortex, layer 4, neurons often respond to transient high-velocity events, such as the set of stimuli used in this study, with a single spike or no spike at all (Lee and Simons, 2004; Wang et al., 2010). Consequently, a trial-average firing rate is similar to an estimate of the reliability of an individual neuron across trials. Given that each individual neuron is operating close to the point of detection failure on a single trial, it is possible that a population of neurons within a single barrel may fail to sufficiently activate downstream processing, particularly at threshold velocities.

One of the biggest limitations of this study is the use of an anesthetized preparation. While it is impossible to know for certain how anesthesia may influence the observed effect, an attractive hypothesis is that the bimodality we observe might be related to different processing of stimuli in both the 'up' and 'down' states characteristic of an anesthetized cortex (Petersen et al., 2003b; Civillico and Contreras, 2012). It is possible that state fluctuations contribute to the overall variability of the VSD signal. Since we did not record any indicator of cortical state, we cannot specifically exclude this hypothesis. However, several observations suggest that a causal role for 'up' and 'down' states for the observed response reliability is unlikely. First, the frequency of no-response trials is stimulus specific. We are not controlling state and the stimuli are delivered in pseudo-random order. Each stimulus condition would be sampled from relatively equal occurrences of the 'up' and 'down' states and therefore state should affect all stimuli equally. Without a more complex non-linear relationship between state and response reliability, simple state fluctuations are not sufficient to explain the stimulus-dependent

47

differences. Second, the dual whisker deflection data suggest that one part of the cortex can respond while a neighboring region of cortex does not. If 'up' and 'down' states were determining this response/no-response characteristic, then one would have to conclude that neighboring whisker barrels could be in opposing states at the same time. While possible, this is inconsistent with existing literature (Petersen et al., 2003b). Finally, these results have been repeated with a fentanyl cocktail (fentanyl 5ug/kg, dexmedotodomine, 150ug/kg, midazolam 2mg/kg). Fentanyl, as an anesthetic, is thought to better represent a desynchronized cortical state similar to the awake brain (Simons and Carvell, 1989; Constantinople and Bruno, 2011). All velocity trends presented in this data are identical in this additional data set (Figure 2-10).
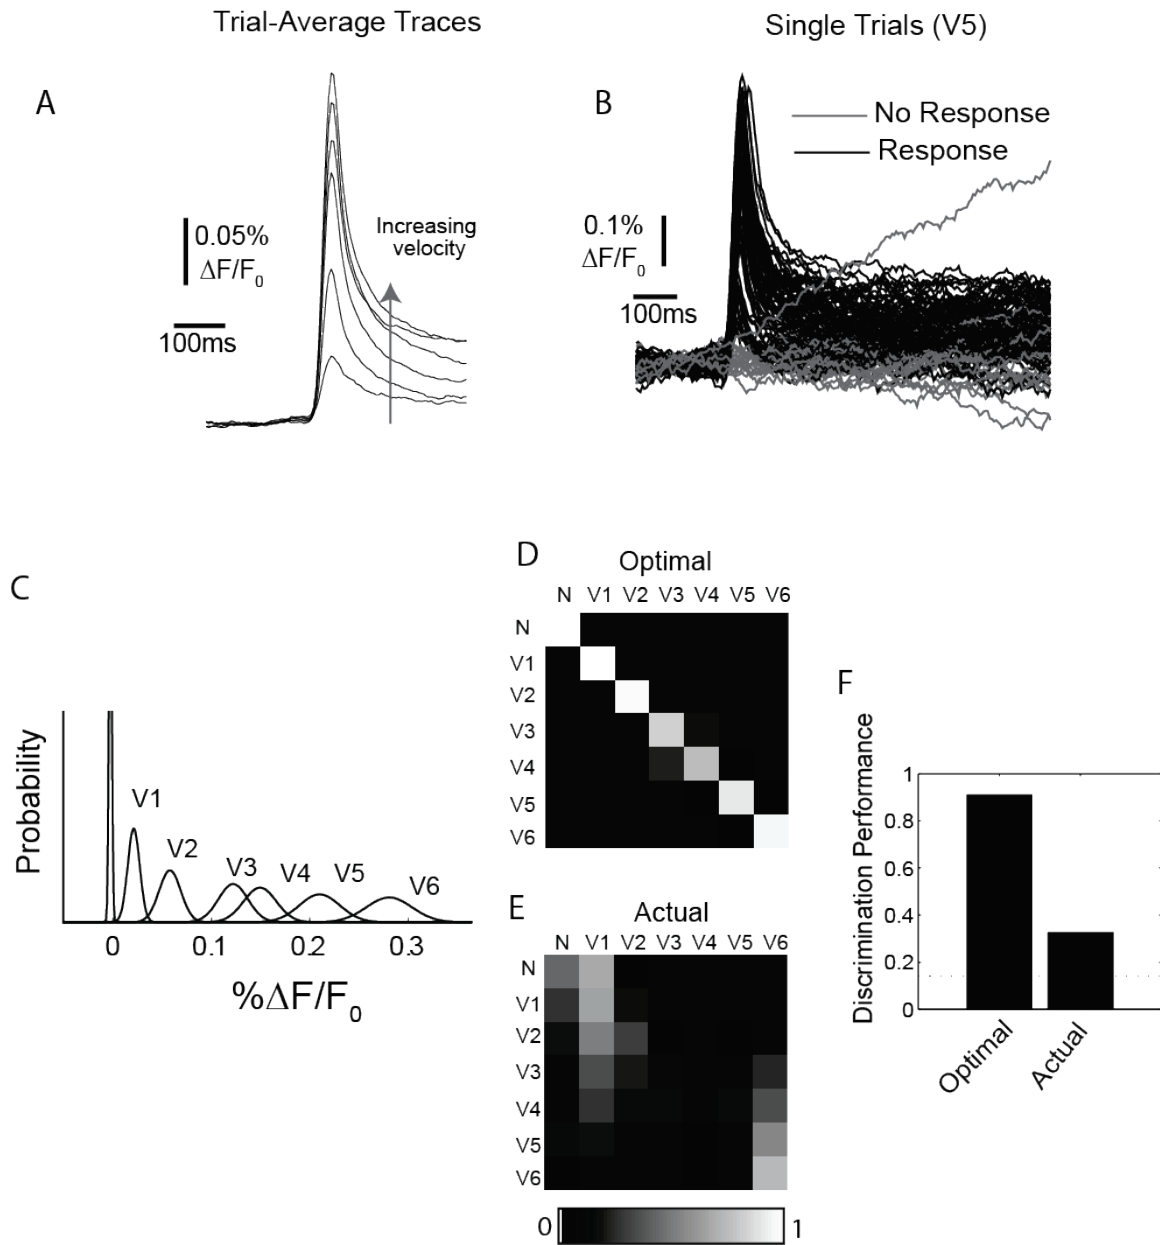
Figure 2-10 Ideal Observer Analysis under Fentanyl Anesthesia A) A replicated experiment using alternate anesthesia showed trial-average increases with velocity a (B) prominent response no response distributions. C) Trial-average classifier for this data set D) the optimal performance matrix and E) actual performance matrix for this data set. F) This data set also showed near chance levels of discriminability.

49

The results presented here allow an intriguingly simple interpretation of the canonical stimulus detection behavior results. Just as head-fixed awake behavioral experiments have shown that increased stimulus velocity increases the probability of detection (Stüttgen and Schwarz, 2008; Ollerenshaw et al., 2012), we observed a higher probability of response in layer 2/3 of the barrel cortex. In fact, the predicted detection threshold based on the probability of activation model presented in this study is between 200-300°/sec, consistent with existing behavior data from our lab and others (Stüttgen and Schwarz, 2008; Ollerenshaw et al., 2012). It is possible to interpret these behavioral results in the context of both the continuum and probability of activation models. However, to our knowledge there is no evidence that rodents can be trained to discriminate velocities on a single whisker. Additionally, the simplest interpretation of the continuum model, where detection occurs when the activity crosses an absolute threshold, cannot entirely account for the observed behavioral results and instead requires more complex frameworks such as an accumulation of evidence model (Ollerenshaw et al., 2012).

While an assertion of simplicity alone is not an argument for validity, the probability of activation model is also generally consistent with existing discrimination behavioral data when considered in conjunction with the 'slip-stick' hypothesis (Wolfe et al., 2008). This hypothesis asserts that texture discrimination can result from the animal perceiving differences in the number of high-velocity slip-stick events across the whisker array when whisking against a texture. Interestingly, the average slip-stick event velocity in this texture discrimination paradigm was 1,100°/sec (Wolfe et al., 2008), which is near the

50

maximum range of velocities tested here. The same stimulus would be predicted to be a threshold stimulus (determined behaviorally) in an adapted state (Ollerenshaw et al., 2014). Given this, we hypothesize that the discriminability of different textures could be regulated by altering the detectability of slip-stick events. The greatest texture discriminability would occur when the mean slip-stick event velocity occurs near a perceptual threshold.

It is tempting to speculate in which way these results may inform or support the study of sensory processing in other modalities. The instantaneous velocity of whisker deflection may be analogous to simplistic characterizations of the strength of inputs in other pathways, such as sound intensity in the auditory pathway or luminance/contrast in the visual pathway, both of which are primarily represented by an increased amplitude of neural response. Here, we observed that the absolute amplitude is not efficiently propagated from the sensory receptor into higher levels of cortical processing on a single presentation within a single cortical column. The analogous measures of stimulus strength in vision and audition are often not well perceived as an absolute. Instead perception is highly dependent on the stimulus context (Marks, 1994; Polley et al., 2006; Dixon et al., 2014). The results of this study predict that perception of stimulus strength in the barrel cortex requires context, specifically coordination across cortical columns. Whether cortical columns in vision or audition also function as independent, but possibly unreliable, sensors of a threshold stimuli remains to be investigated.

51

After examining both single trial amplitudes in response to the deflection of one whisker and single trial spatial profiles in response to the deflection of two whiskers, we propose a framework in which the strength of the sensory input is encoded not in a graded cortical response continuum, but instead probabilistically across the whisker array. Much of what we currently understand about the neural code has been derived from differences in trial-average responses; however, neither of the observations here could have been predicted from the trial-average response alone. Specific consideration of the variability of single trials allowed us to make additional predictions about how (or if) our brain is able to encode stimulus parameters in the presence of high variability.

CHAPTER 3   Probabilistic Encoding of Non-Linear Dynamics

3.1   Introduction

Natural whisker stimuli, even within the simplifying framework of the slip-stick hypothesis, are not just single whisker deflections. Whiskers are used actively by rats to probe stimuli over time in a process known as whisking (Carvell and Simons, 1990; Bermejo et al., 2002; Ollerenshaw et al., 2012). Extensive existing literature has shown that there are complex interactions between whisker deflections delivered separately in time.

Specifically, single-unit electrophysiology in the thalamus and cortex has shown that cortical neurons exhibit velocity-dependent suppression in response to paired whisker deflections (Simons, 1985; Simons and Carvell, 1989; Brumberg et al., 1996; Webber and Stanley, 2004; Boloori et al., 2010). This paired-pulse experimental paradigm consists of two whisker deflections given in series. The first pulse is known as the condition pulse, the second as the test pulse. In classic applications of this paradigm, the test pulse is suppressed by the presence of a condition pulse and the amount of suppression is determined by the intensity of the condition pulse (Boloori et al., 2010). In the whisker system, previously proposed mechanisms have assumed the continuum model where small velocity deflections result in lower response amplitudes and thereby

53

engage a smaller amount of inhibition(Boloori et al., 2010). Faster whisker deflections would result in more activation of the barrel cortex and therefore engage a greater inhibitory response. In context of the probability of activation hypothesis, in which single trial response amplitudes are not reflective of the stimulus velocity, it is not clear how condition pulses could modulate test pulse suppression.

We explicitly considered the non-linear temporal suppression dynamics in response to pairs of deflections in the context of the probability of activation hypothesis. We again used voltage sensitive dye imaging to study trial-to-trial variability of temporal dynamics. Trial-averaged VSD responses show velocity-dependent suppression consistent with previously reported dynamic trends. However, single trial responses are distinctly separable into three groups: 1) a response occurs on the condition pulse but not the test, 2) a response occurs on the test pulse but not on the condition, 3) no response to either pulse occurs. Importantly, there are no trials with responses to both pulses. Using a phenomenological model, we show that the probability of activation hypothesis can explain both the average observed trends and the structure of the single trial distributions.

3.2 Methods

3.2.1 Animals

All procedures were approved by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology. Data from 13 animals are presented in this study. Seven animals were newly tested for this study while data from the additional six animals were collected in the same experiment as data for previous studies.

3.2.2 Surgical Preparation and Voltage Sensitive Dye Imaging

All surgical preparations and imaging methods are identical to those described in Chapter 2. Briefly, a craniotomy, and in some cases a duratomy, was performed over the barrel cortex in animals anesthetized with sodium pentobarbital. Voltage sensitive dye RH1691 was applied to the surface of the cortex and allowed to diffuse into the brain for 1.5-2 hours. After saline washes, the activity in response to whisker deflections was recorded on a high-speed CCD camera with 5ms temporal resolution.

3.2.3 Whisker Stimuli

Controlled whisker deflections were delivered using a glass pipette attached to a piezoelectric actuator identically to previous studies. The stimuli in this study were a pairs of saw-tooth deflections of a specific velocity separated by a defined interval (ISI, inter-stimulus interval). The velocities tested were calibrated to be 75, 150, 300, 600, 900 and 1200°/sec. 10-60 trials were collected at each velocity, depending on available experimental time. For single trial analysis, trials were binned into low, med and high groups. For the equal condition-test paradigm the low, middle and high bins were 75,150°/sec; 300,600°/sec; 1200°/sec respectively. For the fixed test, the test pulse was 900°/sec. The condition velocities were also binned: 75,150°/sec (low); 300, 600°/sec (middle); 900, 1200°/sec. For five of the animals, we also collected data at 3 inter-stimulus intervals (ISIs): 100ms, 250ms and 500ms. For the two-whisker paired pulse studies, equivalent velocities of either 600 or 900 °/sec were delivered by separate piezo-electric control to both whiskers.

55

3.2.4 Image Analysis

Voltage sensitive dye images were processed as described in Chapter 2. Each image represents the percent change in fluorescence relative to a background image (an average of forty pre-stimulus frames). The analysis was done directly on single trial time series created by averaging the fluorescence signal in a region of interest with a radius of 10 pixels (200μm, approximately the size of one barrel) around the center of mass observed in the onset frame (15ms after stimulus delivery).

3.2.5 Trial-Average Analysis

Time series from each single trial (10-60 trials per data set) were averaged together, resulting in a single time series per stimulus condition. The temporal response to a whisker deflection was extremely stereotyped within each data set, but varied slightly between sets. As there is sometimes residual increased fluorescence (and presumably neural activity) from the first pulse at short inter-stimulus intervals, we defined the stimulus-evoked response amplitude using a differenced time series, not the absolute observed value. The response amplitude for each data set was defined as the measurement with the maximum change in fluorescence between 0 and 50ms on the time series from the highest velocity. This ensured that all changes in response amplitude represented evoked activity to the test pulse and not residual activity from the condition pulse. We defined a suppression index (SI) to measure suppression as in previous studies (Simons, 1985; Simons and Carvell, 1989; Boloori et al., 2010). Conceptually, the SI was the observed magnitude of a test-pulse relative to the expected size for that velocity deflection if delivered in isolation. Therefore, in the equal condition-test paradigm, the

56

suppression index is the ratio of the response amplitude of the test pulse over the condition pulse. In the fixed-test, the suppression index is the response to the test pulse divided by the response amplitude of a pulse of the same velocity (900°/sec) delivered without a first pulse.

3.2.6 Single Trial Analysis

Single trial analysis was performed on time series from single presentations of a stimulus. Similar to the trial average, response amplitudes were calculated based on the differenced time series. In differenced time series, the two frames with the greatest change in fluorescence identify the onset frame ($\Delta F_{max}=F_1-F_0$) where F_1 is the first frame with stimulus evoked activity and F_0 is the prior frame. Once identified, the same onset frame was used for all trials from a given animal. The average of the differenced time series for the onset frame and the two subsequent frames was considered the response amplitude. This reduces to the change in fluorescence from the absolute peak frame (F_4) to the pre-onset frame (F_0), in other words only the stimulus evoked change in fluorescence. This also ensured that even small amplitude response trials were different from noise. We then plotted the amplitude of the first pulse (condition pulse) versus the second pulse (test pulse) to observe differences in response probability across single trials.

3.2.7 Paired Pulse Probability of Activation Model

We built a phenomenological model of temporal dynamics to paired whisker deflections. Whisker deflection velocity defined the probability of a response on a given trial according to a sigmoidal function according to the following equation

57

$$P(V) = \left(1 - e^{-\left(\frac{V}{\alpha}\right)^{\beta}}\right)$$

where V is the deflection velocity, P is the probability of observing a response and α (354.17) and β (0.78) are free parameters governing the shape of the curve. We used the sigmoidal curve fits from an actual behavioral psychometric curve published previously (Ollerenshaw et al., 2014) where rats were trained to respond by licking to whisker deflections of variable velocity. Parameter α is the perceptual threshold and parameter β is the slope of the sigmoid. Our model is built directly from behavioral results, without any fitting to neural data. We assume that if a response occurs to the condition pulse, inhibitory circuitry is engaged such that there is no probability of response for the test pulse. If a response is not observed, the same probability of response curve is used to determine the probability of response to the test pulse. We assumed a binary model of response amplitudes, one for a response or zero for no-response. The probability of response to a test pulse of velocity, v_{test}, given a condition pulse of velocity, v_{cond} is:

$$p_{(v_{test}|v_{cond})} = P(v_{test})(1 - P(v_{cond}))$$

In the fixed test condition, $P(v_{test})$ is a constant as the test pulse velocity is fixed. In the equal condition-test paradigm, $v_{test} = v_{cond}$. In both cases, the predicted suppression index (SI) reduces to:

$$Suppression\ Index\ (SI) = \frac{p_{(v_{test}|v_{cond})}}{P(v_{test})} = 1 - P(v_{cond})$$

for all condition pulse velocities.

3.3 Results

In this study, we extended the probability of activation hypothesis to pulses delivered sequentially in time. We tested multiple variations of the traditional paired-pulse paradigms. First, we used a paired-pulse paradigm in which the condition and test pulses are of equal velocity and the inter-stimulus interval is fixed. Next, we kept the velocity of the test pulse constant, while varying the velocity of the condition pulse. We also varied the inter-stimulus interval between pulses. We propose a probabilistic model for these temporal dynamics that is consistent with both the trial-average observed responses and the single trial distributions.
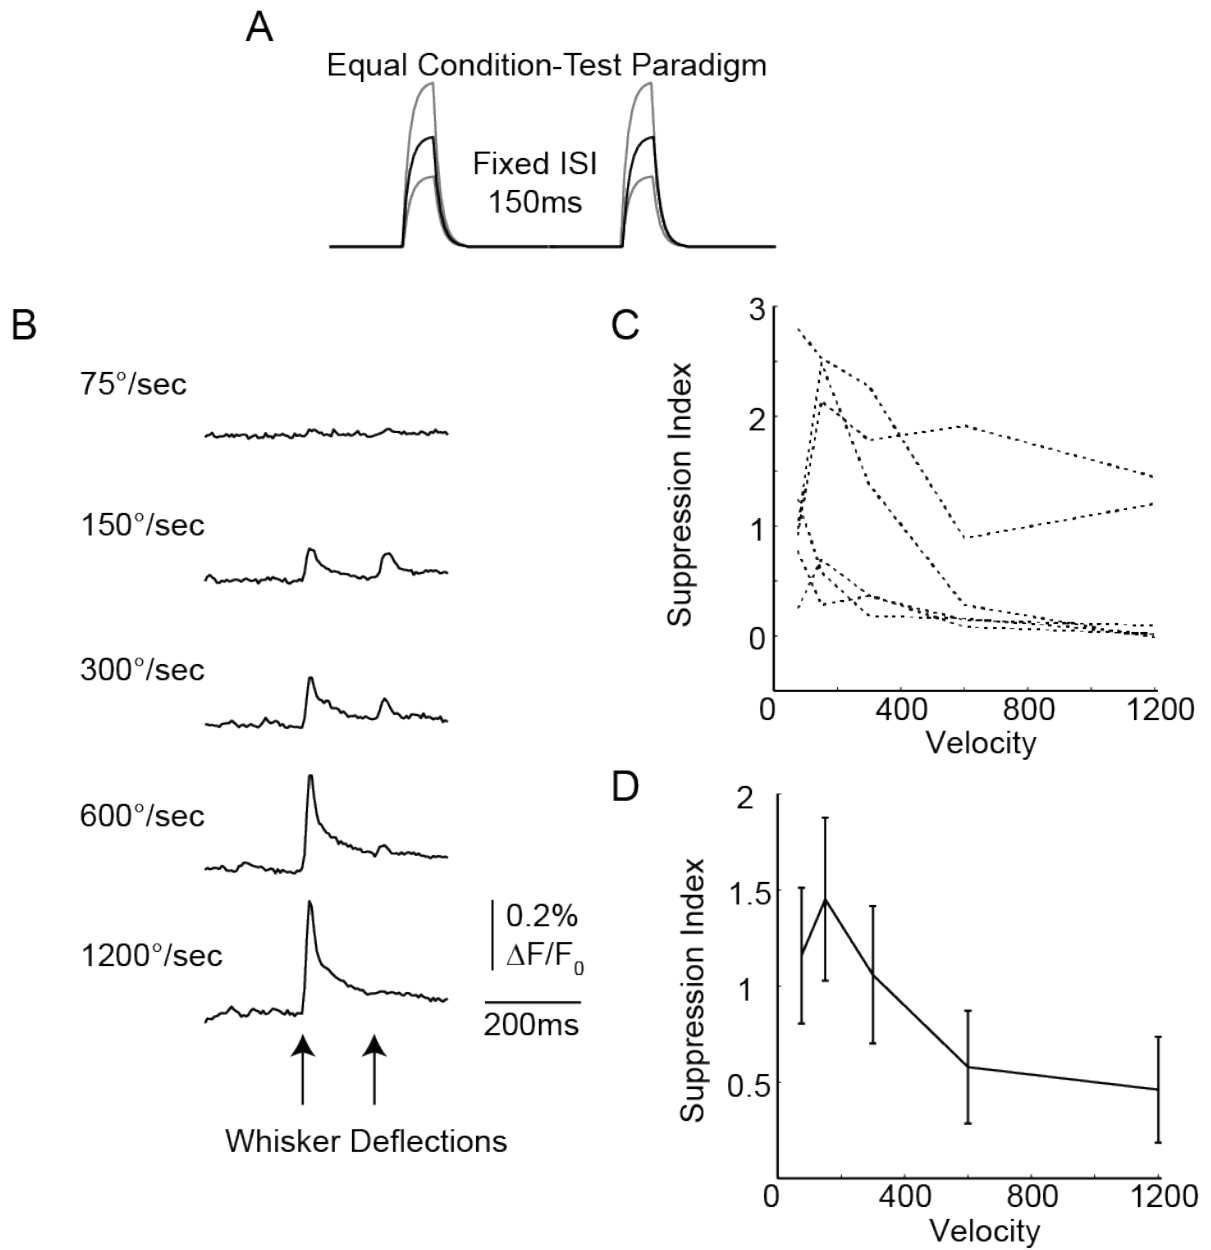
Figure 3-1 Equal Condition-Test Paired Pulse Paradigm A) Stimuli for the equal condition-test paired pulse paradigm consisted of two saw tooth whisker deflections of equal velocity separated by 150ms in time; B) Paired-pulses of five different velocities were delivered (at times denoted by arrows) and representative trial-averaged time series for one animal are shown. Notice that at lower velocities two clear equal sized pulses are visible while suppression to the second pulse is prominent at high velocities. C) Suppression index curves from six animals are shown. D) Mean and s.e.m. from the individual curves shown in C.

60

3.3.1 Temporal Dynamics are Velocity-Dependent

Trial-averaged paired-pulse dynamics observed using optical imaging in layer 2/3 were consistent with previously published single unit electrophysiological data collected in layer 4 of the cortex. We averaged the fluorescence over a region of interest centered on the primary barrel to create time series (see Methods). First, we considered the response to two paired whisker deflections of equal velocity, separated by a fixed inter-stimulus interval of 150ms (a schematic is shown in Figure 3-1A). The average response amplitude of the condition pulse increases with the deflection velocity (Figure 3-1B). The response amplitude of the second pulse, the test pulse, was inversely correlated with the condition pulse velocity indicating velocity-dependent suppression. Data from six individual animals is shown in Figure 3-1C with the mean and standard error shown in Figure 3-1D. In one data set, velocity dependent suppression was not observed. Instead, in this data set responses to all test pulses were approximately twice the condition pulse amplitude.

In a related but distinct stimulus protocol in an entirely different set of animals, we again varied the condition pulse velocity, but kept the test pulse at a fixed velocity (900°/sec). Similar to the equal condition-test paradigm, the fixed-test paradigm also showed velocity-dependent suppression (Figure 3-2). Since the test pulse velocity was constant, the response amplitudes to the test pulse shown in the representative trial-average time series can be compared directly (Figure 3-2B). Suppression was observed in 4 out of 5

61
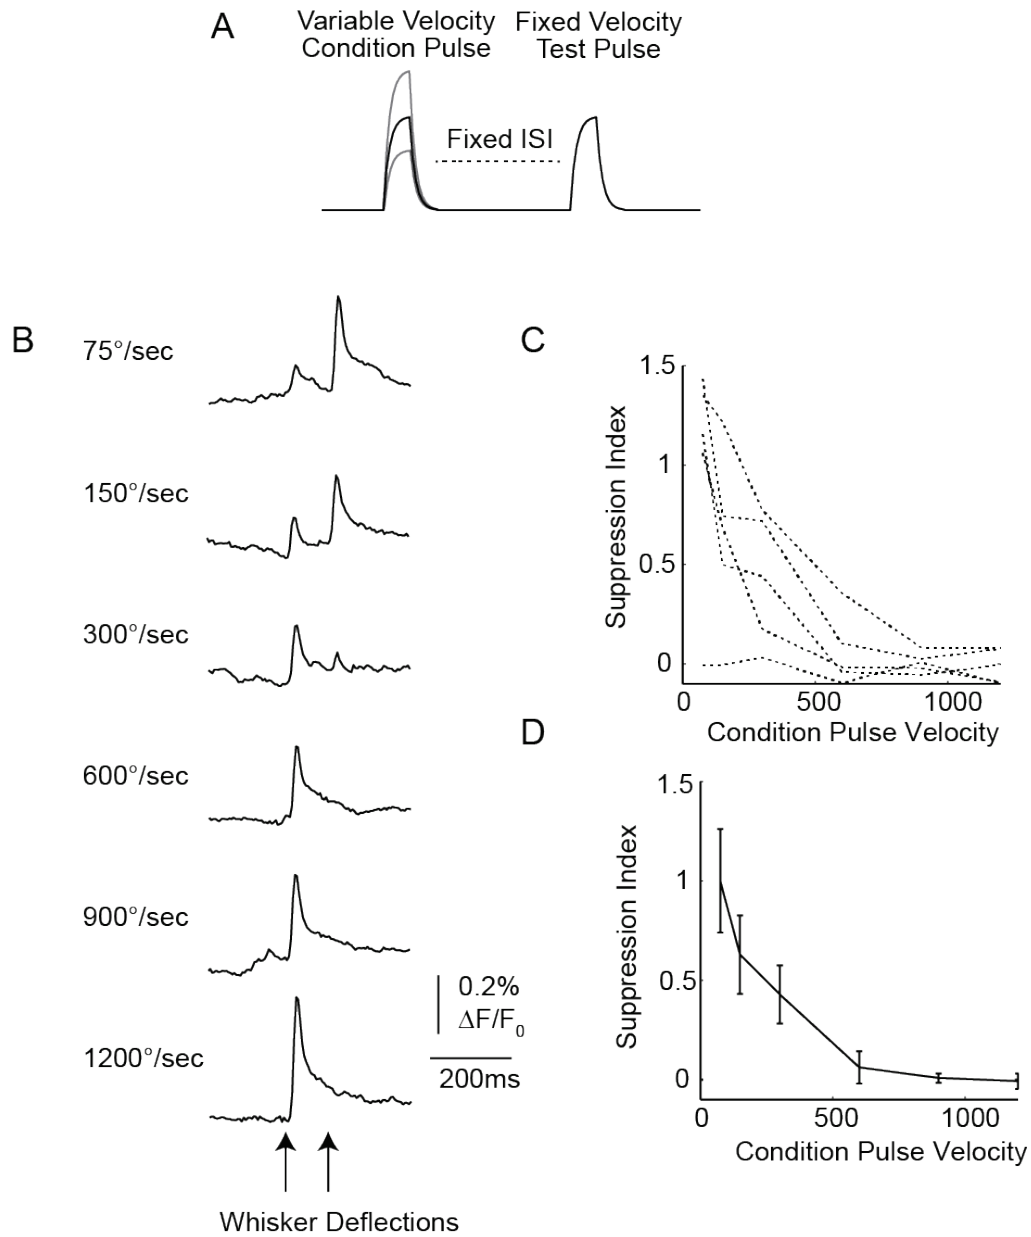
Figure 3-2 Fixed-Test Paired Pulse Paradigm A) Stimuli to the fixed test paradigm consistent of two saw tooth whisker deflections given at 100ms inter-stimulus interval (ISI). The condition pulse velocity is varied; the test pulse velocity is fixed at 900°/sec. B) Representative trial-average time series show that as the condition pulse grows, the response to the test pulse (although an identical stimulus) shrinks. C) Suppression indexes for both all animals. D) Mean and error bars (± s.e.m.) of data shown in C.

data sets (Figure 3-2C). In the fifth data set, no responses were observed to the second pulse with any condition pulse velocity. The average is summarized in Figure 3-2D.

In both of these classic paired pulse paradigms, the trial-average results recorded in layer 2/3 using VSD imaging matched the trends previously observed from single unit data in layer 4 of the cortex. This is the first time these trends have been repeated using a technique thought to primarily report sub-threshold activity.

3.3.2   Single Trial Distributions

As the trial-average responses are consistent with previous literature, we then considered the probability of activation dynamics of the single trials. A set of single trial time series in response to a pair of whisker deflections is shown in Figure 3-3.

63

Figure 3-3 Single Trial Paired-Pulse Time series Time series from VSD single trials are shown. Whiskers deflections are denoted by arrows. Single trial responses can easily be separated into trials that respond to the condition pulse but not the test pulse, or the test pulse but not the condition pulse.

We observed a dramatic bimodality. If a response occurs on the condition pulse, then a response was not observed on the test pulse. At times, test pulse responses equaled the magnitude of responses on other trials to the condition pulse, even if the velocity was different between the two conditions (consistent with the single pulse probability of activation hypothesis). This effect was true in all animals tested in both the fixed-test and equal condition-test paradigms. A summary (all single trials, all animals) is shown in

64

Figure 3-4. In the top row, we show all the single trial responses from the six animals in which we delivered the equal condition-test stimuli. The response amplitude (averaged over multiple frames as described the Methods) to the condition pulse is plotted on the x-axis and the response to the test pulse on the y-axis. The single-trials form an L-shaped distribution (Figure 3-4A). The absence of trials in the upper right quadrant indicated that trials with responses to both the condition and the test pulse were never observed. We also separated these trials by condition-test pulse velocity (low, mid and high). When the velocity of the condition pulse was low (Figure 3-4B), a group of no-response trials (points centered near the origin) was prominent. The response trials were sparse, and occurred at approximately equal number of trials to both the condition and the test pulse. This resulted in a suppression index of around one. As the velocity increased, we observed more responses to the condition pulse but fewer responses to the test pulse.

Equal Condition-Test Paradigm (150ms ISI) n=6 data sets

Fixed Test Paradigm (100ms ISI) n=4 data sets
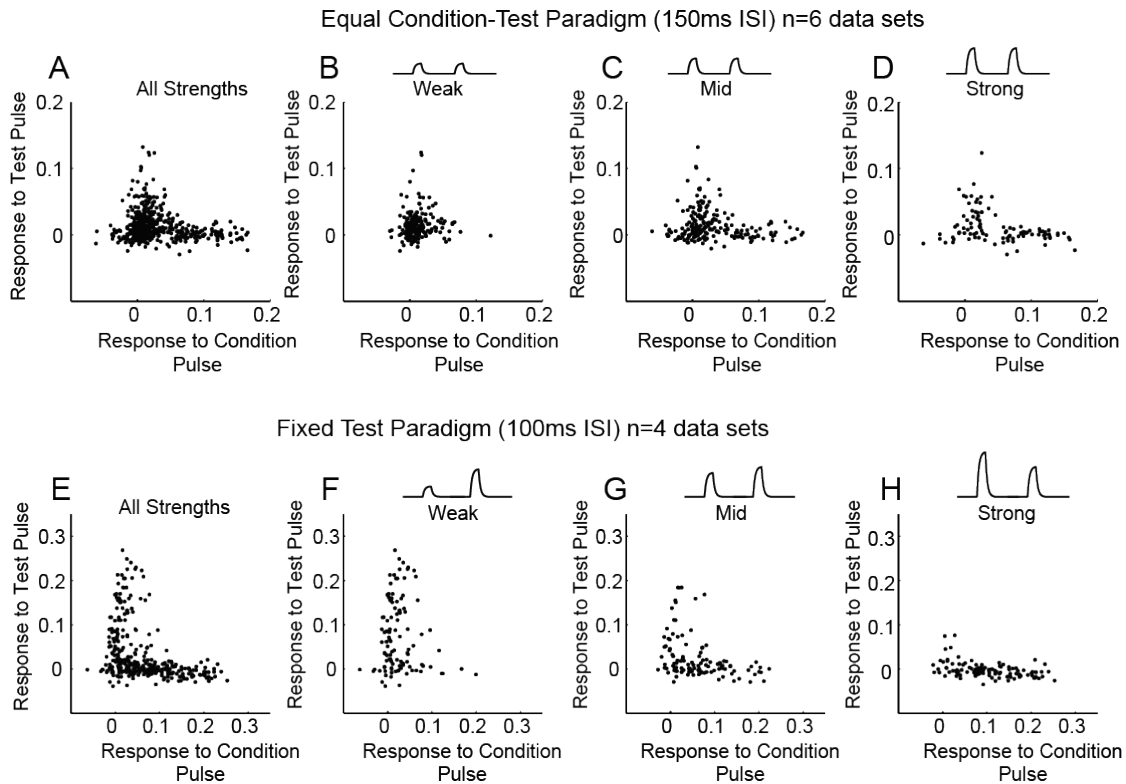
Figure 3-4 Single Trial Paired Pulse Distributions A) All single trials, all velocities, all animals from the equal condition-test paradigm are plotted as black dots. The response to the condition pulse is on the x-axis and the response to the test pulse is on the y-axis. No trials responded maximally to both the condition and test pulse resulting in an L-shaped distribution B) The same data from (A) separated by velocity. Low velocity responses were rare, but when they occurred, they did not occur on the same trial. C) Intermediate velocity stimuli had more responses to the condition pulse, and fewer responses to the test pulse D) High velocity pulses resulted in primarily responses to the conditional pulse with few responses to the test pulse. Responses to the test pulse only occurred on trials with no response on the first pulse. E) Similar to (A) but for the fixed-test paradigm F) A low velocity condition pulse resulted in responses primarily to the test pulse. However, if a response occurred in response to the small condition pulse, a response to the test pulse was not observed; G) Intermediate condition pulse velocities show approximately equal distributions of response to the test pulse and condition pulse. H) The high velocity condition pulse is capable of eliminating responses to the test pulse.

66

Similar plots are shown for five animals in the fixed-test paradigm in the bottom row of Figure 3-4. In this case, as the condition pulse velocity increased in strength from low to high, the structure of the single trial distributions shifted from primarily a vertical line (Figure 3-4F, responses to the second pulse only) to a horizontal line (responding to the first pulse only Figure 3-4H). All response amplitudes appeared achievable to all velocities; however, in this case there was a slight trend in the maximum single trial response from the low to the high condition velocities (Figure 3-4F to Figure 3-4H). This effect was not as dramatic as the changes to observed distribution bimodality.

3.3.3 Phenomenological Model of Velocity-Dependent Suppression

These dynamics could be explained by two assumptions: 1) response probability was modulated by stimulus intensity as predicted by the probability of activation hypothesis; 2) response trials, but not no-response trials, engaged a subsequent inhibitory response. We assumed that the probability of response curve as observed in layer 2/3 was the same as the probability of response curve in an awake behaving animal. To test this hypothesis, we used curve fits directly from behavioral detection data collected previously to model our neural data (Ollerenshaw et al., 2014). The probability of response curve, $P(V)$, is shown in Figure 3-5A. The exact same mapping from velocity to probability of response was used for the test pulse (the pulses are independent if no response is observed); however, trials that had previously responded to the condition pulse had a zero probability of response. Therefore, the probability of observing a response to a test pulse of velocity v_{test}, given a pre-pulse of velocity v_{cond}, is $P(v_{test})(1 - P(v_{cond}))$. In the fixed test condition, the probability of response for the test pulse is a constant. In the

67

equal condition-test paradigm, $(v_{test}) = P(v_{cond})$. Without performing any curve fitting, we mapped directly from behavior to neural data (see Methods). This model predicts the probability of response curves shown in Figure 3-5. Both the fixed test and equal condition-test paradigms started with identical condition pulses. The probability of response for the fixed test (900°/sec test velocity was used in this study, Figure 3-5B) and equal condition-velocity (Figure 3-5C) are shown in blue. As we assume the amplitude of all response trials is one and no-response trials is zero, the probability of response curve is also the expected trial average amplitude. In the fixed-test condition, our model predicts the trial-average response amplitude to decrease dramatically from a maximal response to zero, which was observed (Figure 3-2). For the equal condition-test paradigm, the model predicted that the absolute response to the test pulse would start near zero, increase slightly before beginning to decay again. This is consistent with the observed trend in the equal condition-test paradigm shown in Figure 3-1. Interestingly, despite the differences in these paradigms, the expected suppression index is not predicted to be different in these two conditions. This is a strong testable prediction.
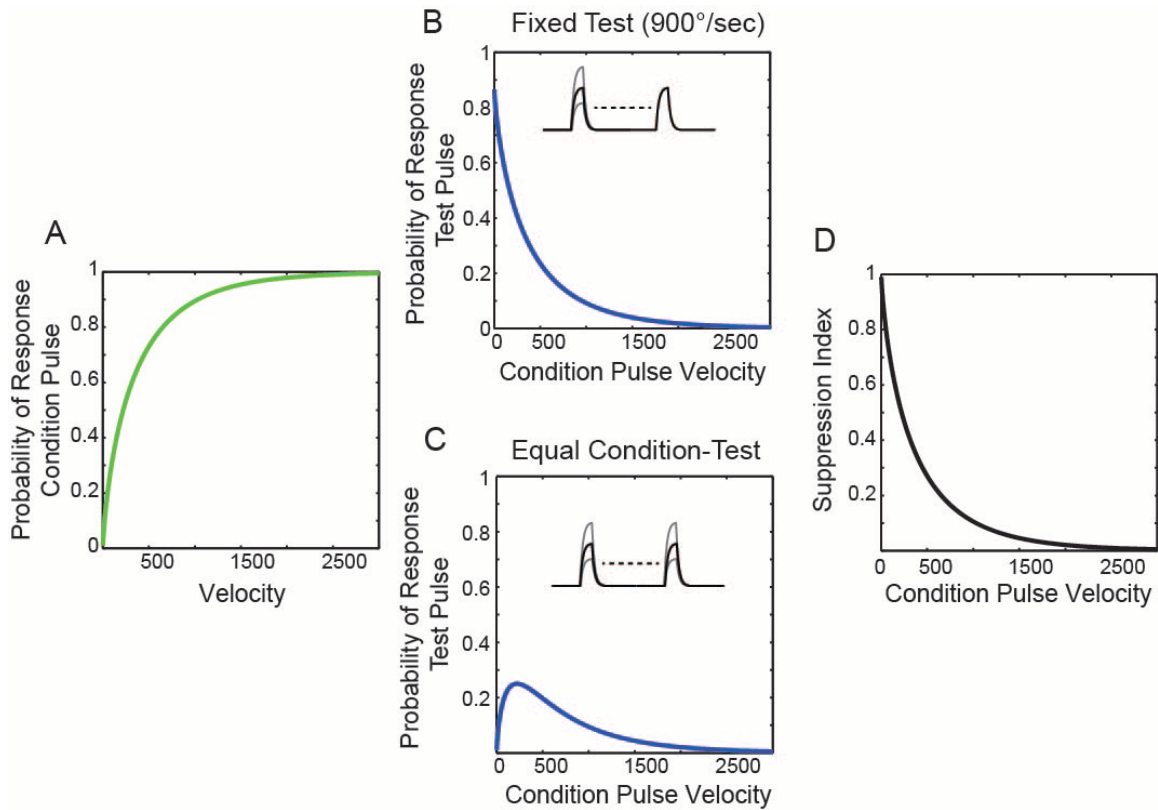
Figure 3-5 Deterministic Probabilistic Paired-Pulse Results A) A curve from an awake behaving detection paradigm published elsewhere (Ollerenshaw et al., 2014) was used to map from the a whisker deflection velocity to probability of a response trial. The condition pulse was identical for both paradigms. B) The predicted probability of response to a test pulse of 900°/sec is plotted for variable condition pulse velocities. C) The predicted probability of response for a test pulse in the equal condition-test paradigm given the behavior data. D) The expected suppression index for both paradigms is equivalent.

3.3.4   Model Derived From Behavioral Data Fits Observed Neural Results

We tested the suppression model, built from behavioral data directly, by plotting the model on top of the observed neural suppression curves. For reference, we plot again the trial-average responses from both paired-pulse paradigms in Figure 3-6 (shown in black). This time we also plot our phenomenological model predictions for the observed suppression index in blue. For the fixed-test paradigm (Figure 3-6A, B) there is a

69

Fixed Test Paradigm

Equal Test-Condition Paradigm
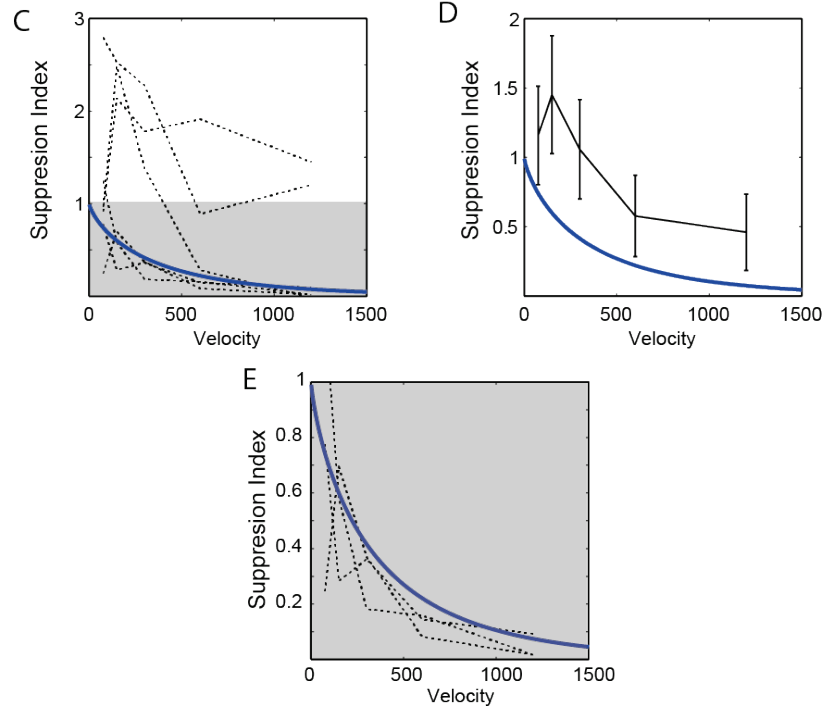
Figure 3-6 Probabilistic Model Fit to Data A) Individual animal SI curves from the fixed test (dotted black lines) are shown with the predicted SI curve built directly from behavior data (blue). B) Mean of data shown in (A). (C) and (D) are the same as (A) and (B) but for the equal condition-test paradigm. The predicted curve underestimates the observed response on average, but actually does a match 3 out of 6 curves that were observed fairly well. E) From (C), gray box zoomed in to show the three better fit curves.

70

reasonable agreement between the observed experimental curve and the predicted curve. In Figure 3-6A, we show that the model falls on top of 4 out the 5 data sets, except the data set with no observed test pulse responses.

We also show the model fit with the equal condition-test paradigm in Figure 3-6. The model captured the general trend of the data, but systematically underestimated the suppression index. By examining the single animal curves (Figure 3-6C) we saw that this curve actually fit about half the data extremely well (a zoomed in version of the gray box is shown in Figure 3-6E). The remaining three data sets were less suppressed than our model predicted, including the one data set that did not show suppression. For the other two data sets, the greatest error from our predicted fit occurred in the region of the graph with the lowest expected number of responses. In some animals, there were only 10-15 trials of data. For low-velocity stimuli, we expected to see at most 1-2 responses, but having no response trials was common. In this case, the act of dividing by near-zero response amplitude increased the variability of the suppression index. Some animals had suppression index values that moved from one to two simply because of one extra response trial.

3.3.5 Varying Inter-stimulus Interval

The results presented so far have only included analysis at one inter-stimulus interval. These suppression dynamics are known to be dependent on inter-stimulus interval as the circuitry returns to a baseline state (Simons and Carvell, 1989; Boloori and Stanley, 2006; Boloori et al., 2010). We tested the response dynamics in five animals (the same

five animals used for the fixed-test paradigm) at three different inter-stimulus intervals while keeping both the condition pulse and the test pulse at equal velocity. On average, allowing additional time in between whisker deflections resulted in less suppression of the test pulse. Representative time series can be seen in Figure 3-7B. The recoveries occurred at dramatically different rates in different animals (Figure 3-7C), with no animals reaching complete recovery by the longest ISI of 500ms. This recovery time is longer than reported in layer 4 of the cortex or thalamus (Simons, 1985; Simons and Carvell, 1989; Boloori et al., 2010).

Figure 3-7 Variable ISI Paired Pulse   A) Schematic of two whisker deflections of equal velocity given at a specified inter stimulus interval (ISI). B) The suppression indexes of 5 separate animals at increasing ISIs show great variability between animals C) representative trial-averaged time series from one animal show an incomplete recovery from suppression at 500ms ISI. D) The mean and s.e.m. of the data shown in (B). Single trial first pulse vs second pulse distributions at 100ms (E), 250ms (F) and 500ms (G) show some trials begin to response on both pulses

The single trial distributions show the bimodality at the 100ms interval (as shown previously). Beginning at 250ms ISI (Figure 3-7F) and then more dramatically at 500ms ISI (Figure 3-7G) we observed trials that responded to both the condition and test pulse. The bimodality in the data diminished with time. On trials with responses to both pulses, the response amplitude of the condition pulse was not obviously correlated with the response magnitude to the test pulse. This suggests that the inhibitory circuitry was not engaged relative to the absolute amplitude of the condition pulse. Again, we conclude that the absolute amplitude was not predictive of the subsequent dynamics; knowledge of only whether or not a response occurred was sufficient.

3.3.6 Two-Whisker Paired Pulse Paradigm

We have shown that a response to a condition pulse can prevent a response to a subsequent test pulse of the same whisker. However, we do not know if this suppressive effect is specific to a single whisker barrel. In Chapter 2, we showed cortical barrels appeared to be able to respond independently when deflections were delivered simultaneously, existing literature suggests that deflections of adjacent whiskers are suppressive if delivered preceding a deflection on a primary whisker (Simons and Carvell, 1989; Higley and Contreras, 2003, 2005; Civillico and Contreras, 2006). We tested this hypothesis using VSD. An adjacent whisker was defined as any whisker that was one row or arc removed from the primary whisker (Figure 3-8A). The test pulse in this paradigm was always a whisker deflection of constant high velocity. Response amplitude was averaged within a region of interest centered over the primary whisker

74

barrel. The condition pulse is delivered either on the same whisker (similar to previous paradigms) or on an adjacent whisker (Figure 3-8B).
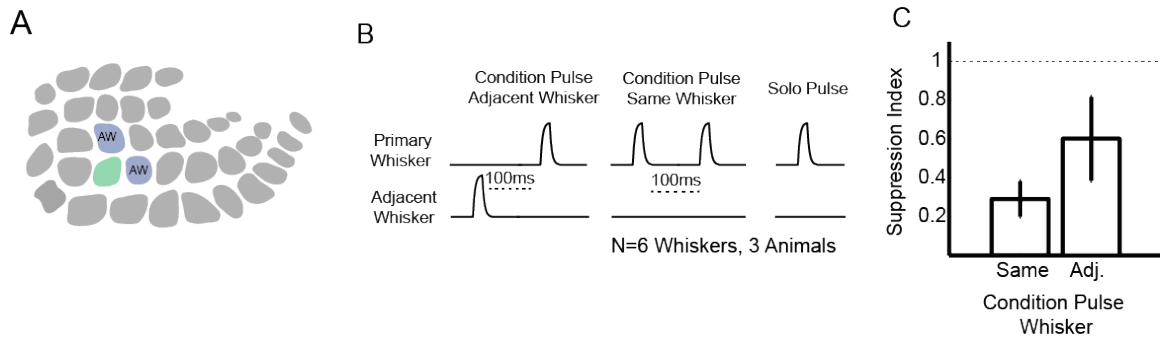
Figure 3-8 Two-Whisker Paired Pulse A) An adjacent whisker to the test whisker (green) is any whisker that shares a row or arc (blue). B) Three types of stimuli were collected, response to the whisker alone and the response with a condition pulse on either the same whisker or an adjacent whisker. C) A condition pulse on either the same whisker or neighboring whisker was sufficient to suppress the test pulse relative to the solo pulse. Equal response magnitudes would be an SI of one, shown in a dotted line.

We observed that a condition pulse on either the primary whisker or an adjacent whisker was capable of suppressing the test pulse. The test pulse was slightly more suppressed when the condition pulse was on the same whisker, however this effect was not significant (n=6 whiskers from 3 animals, paired t-test, p=0.19). This suggests that when a response occurs, the inhibitory response is not unique to the primary whisker.

3.4   Discussion

In summary, we have shown that the complex non-linear temporal suppression dynamics previously reported using single and multi-unit electrophysiology in layer 4 are also

75

observed in the trial-average VSD time series of layer 2/3. The suppressive effect of a condition pulse was determined by its intensity, with stronger whisker deflections suppressing subsequent deflections more. This suppression effect was strongest at short intervals and then recovered with time. The suppressive effect was not limited to the specific whisker barrel and was capable of suppressing neighboring whiskers.

The fact that the trial-average temporal dynamics were similar to previously reported studies was both a novel finding and a validation of previous work. VSD is an optical technique, and therefore like all optical techniques, is limited to the surface layer 1-3 of the cortex (Petersen et al., 2003b; Mateo et al., 2011). Additionally, it is thought that VSD primarily reflects sub-threshold activity. Previous studies on this topic were done in layer 4, the input layer from the thalamus (and some thalamic studies directly). Here, we show that these dynamics observed in layer 4 were also observed in its output layer.

Although the trial-average trends are identical to existing evidence, the previously proposed mechanisms to explain these temporal dynamics are insufficient to explain the single trial response distribution. Single trial population neural responses are not the same each time a stimulus is presented. Specifically, we observed that suppression of the test pulse required stimulus-evoked activity to the condition pulse on a trial-by-trial basis. We proposed a model where probabilistic trial-to-trial variability contributed to the trial – average suppressive dynamics. Trials that respond to the condition pulse cannot then respond to the test pulse; if no response occurs to the condition pulse, the test-pulse

76

response probability is unchanged. This model was sufficient to predict the observed neural responses.

Interestingly, we were able to predict the size and scale of neural response using detection data directly from an awake behaving rodent. The uniqueness and robustness of an effect that requires no fitting in order to map from behavior directly to an anesthetized recording cannot be overstated. This is strong evidence that the probabilistic effects we observe in layer 2/3 do in fact directly reflect the probabilistic performance and the perception of the animal.

We also saw differences in variability between different measurements that are predictable from a probabilistic framework only. In the continuum model there is no reason to expect that measurements from small velocities (when assessed with the same number of trials) should have a larger standard deviation (or standard error) than measurements at large velocities. If probabilistic, however, the observed variability depends on the expected number of true response trials which itself depends on the number of trials collected. In the case of the equal condition-test animals presented in this study, we used a small number of trials (10-15 per condition). The expected response probability is less than one. Because the true measurements are discrete, we could reasonably expect to observe zero, one, or two responses. Given a discrete event and small number of observed response trials, we predict higher variability (including shifts in magnitude from 0% to 200% of the true value) on small measurements more often than on measurements with a large number of expected response trials. This level of

77

variability and associated trends with stimulus strength was exactly what we observed in the data.

We further interpret the data and model presented in this chapter as evidence that the suppressive neural response to a whisker deflection can be considered stereotyped. Each time a whisker response occurs, it is variable in size, but functionally identical to all other responses. This assertion is still consistent with the complex non-linear dynamics of sequences of pulses. Non-linear dynamics can be created simply by modulating the reliability of the response alone. There were no robust and reliable differences on single trial responses to a condition or test pulse of equal velocity. Trends observed in the trial average are not necessarily represented on individual single trials.

There were some observations in the data that could not be accounted for within the probability of activation hypothesis. Specifically, very short inter-stimulus intervals (less than or equal to 100ms) were more suppressed than can be accounted for from observation of the layer 2/3 activity alone (data not shown). As highlighted in this chapter, there was one animal preparation that did not respond to any test pulse even when a fast whisker deflection follows a small condition pulse. This was inconsistent with the probability of activation hypothesis. However, if we were to examine this animal's data at longer inter-stimulus intervals, the probabilistic responses to the test pulse return (although the bimodality does not disappear at any length of tested ISI). This observation, the short-ISI problem, will come up again in the discussion of the proposed mechanism in Chapter 6.

78

In summary, we have shown that the non-linear temporal suppression dynamics are consistent with the probability of activation hypothesis. The neural representations of temporally complex stimuli are not identical each time a stimulus is presented. Even if on average the second pulse is smaller than the first, one cannot assume or predict if that relationship will hold true on any specific instance of the stimulus. This creates an interesting problem for the decoding of neural information in real time. This problem is addressed in a proposed framework for encoding and decoding of stimulus information the context of the probability of activation hypothesis in Chapter 5.

CHAPTER 4   Spatial Evidence for Probability of Activation Hypothesis

4.1   Introduction

The probability of activation hypothesis stems first from the observation that response amplitudes to whisker deflections were not predictive of stimulus intensity (Chapter 2). However, response amplitude was not the only measureable quality of the neural response. Velocity could also be encoded (and decoded) by other characteristics of the neural response. In this chapter, we specifically consider the discriminability of velocity stimulus by their representations in space.

If velocity information could be reliably decoded from single trials of whisker deflections based on the spatial extent of the cortical activation, then this would directly contradict the probability of activation hypothesis. Velocity-specific differences in spatial extent have been reported in literature in the barrel cortex (Wang et al., 2012). Stimulus intensity differences in the spatial spread have also been reported in other sensory modalities (Polley et al., 2006). As a result, it was of critical importance that the spatial dimension was considered within the context of this hypothesis.

As the response amplitude was not predictive of whisker deflection velocity, it was important that the method used to assess the differences in area of activation was not sensitive to covariance of response amplitude. Most common metrics of spatial

activation did not meet this criterion. There is a complete analysis of the ways in which covariance of amplitude can confound spatial data included as an appendix (see APPENDIX A). The methodology is not immediately relevant to the probability of activation hypothesis so only the relevant spatial analysis is presented in this chapter.

Here we show that the spatial spread of VSD response reflects known asymmetries in the anatomy, spreading more along whisker rows than arcs. Next, we analyze both the trial-average data and single trials separately in order to evaluate the reliability of changes in spatial spread. We found that using traditional spatial metrics on trial-average data, we could observe an increase in area with velocity; however, after controlling for the covariance of amplitude, we find no changes in spatial activation across velocities. Results in this chapter show that the spatial extent of the cortical circuitry engaged in response to whisker deflections is remarkably consistent and repeatable. This spatial analysis supports the primary assertion of the probability of activation hypothesis that whiskers can be viewed as stereotyped, but unreliable, detectors of stimulus intensity.

4.2   Methods

4.2.1   Surgical Preparation and Voltage Sensitive Dye Imaging

Data presented in this Chapter come from six animals including three of the same animals from the data in Chapter 2. The surgical procedures and imaging methods were identical to those described previously. For simplicity, we analyzed only four velocities (V1-V4; 150°/sec, 300°/sec, 600°/sec and 1200°/sec respectively), but the trends are also identical in all other velocities.

81

4.2.2 Image Analysis

Image analysis was performed identically to the previous chapters (2.2.5). Briefly, images represent fluorescent values as the percent change from background (pre-stimulus frames). Amplitudes were calculated as an average fluorescent value in a circular region of interest of approximately the size of a barrel (200um radius) centered on the center of mass of the onset frame. In this study, response amplitude is measured at the peak frame.
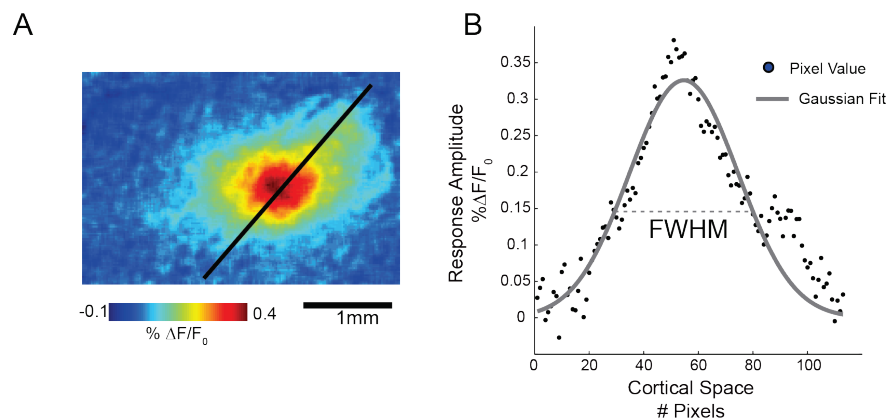
Figure 4-1 Area Schematic A) A representative VSD image shows localized signal B) Pixels values extracted along the black line shown in (A) show that the response is approximately Gaussian. Spatial spread can be quantified using the full width at half max (FWHM).

4.2.3 Defining Row and Arc Axis

All data analysis was done in custom software written in Matlab (MathWorks, Natick, MA). In order to register the anatomy of the whisker representation on a voltage sensitive dye image, we used the functional representation of two whiskers in the same row. Specifically, we calculated the center of activation of the onset frame (determined

manually) for each whisker. The two whisker centers defined the row axis. The arc axis then passed perpendicularly to the row axis and crossed through the center of activation of the whisker being analyzed. If a whisker in the same row was not available, we used a whisker in the same arc and defined the row axis perpendicular to the arc axis. Images were filtered with a spatial averaging filter (200μm x 200μm) prior to extracting the pixels that fall along the row or arc axis. These pixels formed a 1D representation of the spatial spread in this dimension. We fit a Gaussian to these image profiles and recorded the full width at half height (FWHM) as the metric of spatial spread.
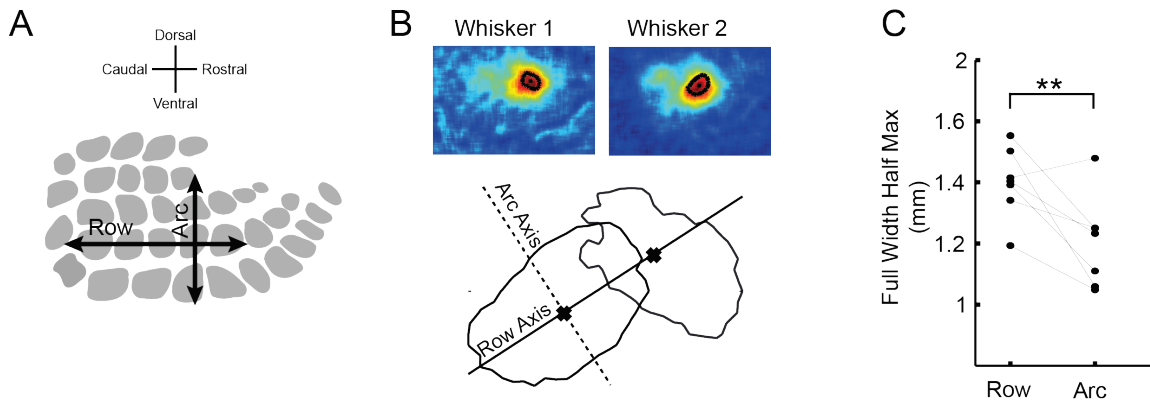
Figure 4-2 Row Arc Asymmetry A) Whisker rows are defined along the rostral-caudal axis of the whisker pad on the face and the arc is in the dorsal-ventral direction. B) A row axis was defined using the center of mass of two whiskers in the same row. The arc axis is perpendicular to the row axis. C) Spatial spread along the row axis is consistently larger than along the arc (6 out of 7 whiskers (from 7 animals), p=0.019; paired t-test, n=7).

83

4.3   Results

In Figure 4-1, we show a representative example of the neural response to a whisker deflection measured with voltage sensitive dye (VSD) imaging. The pixel intensities represent the magnitude of the fluorescent signal, previously shown to change linearly with the membrane potential of a neuron (Petersen et al., 2003a). Figure 4-1B shows the amplitudes measured in a cross section of the image, shown as a line in Figure 4-1B. The response to a whisker deflection is approximately Gaussian (as expected for light imaged scattered through tissue (Stallinga and Rieger, 2010). Given this shape, we used the full-width at half max (FWHM) of this Gaussian to quantify the spatial spread.

4.3.1   Responses to Whiskers Deflections Are Asymmetric

It has been observed in multiple studies that the spread of activity after a whisker deflection is asymmetric. Activation spreads farther along whisker rows than arcs (Petersen et al., 2003a; Lustig et al., 2013). We tested to see if the VSD data collected in this study is consistent with these previous observations. A schematic of the spatial arrangement of cortical barrels oriented to match the whisker map as seen from the side of the face is shown in Figure 4-2A. Whiskers in the same rostral-caudal direction are considered in the same row and whiskers in the dorsal-ventral axis are considered in the same arc.   We defined each axis functionally in VSD images by using the center of activation from two whiskers in the same row. The row axis was the line connecting these two points. The arc axis was perpendicular to this line through the center of mass of the whisker representation being evaluated. We then fit a one-dimensional Gaussian to the

84

pixel values along each of the two axes (similar what is shown in Figure 4-1B) and compared the full width at half max (FWHM) of these two Gaussians. This data is summarized in Figure 4-2C. In 6 out of 7 whiskers (from 6 animals), the FWHM was larger in the row axis than in the arc axis and this result was statistically significant (p=0.019; paired t-test, n=7).
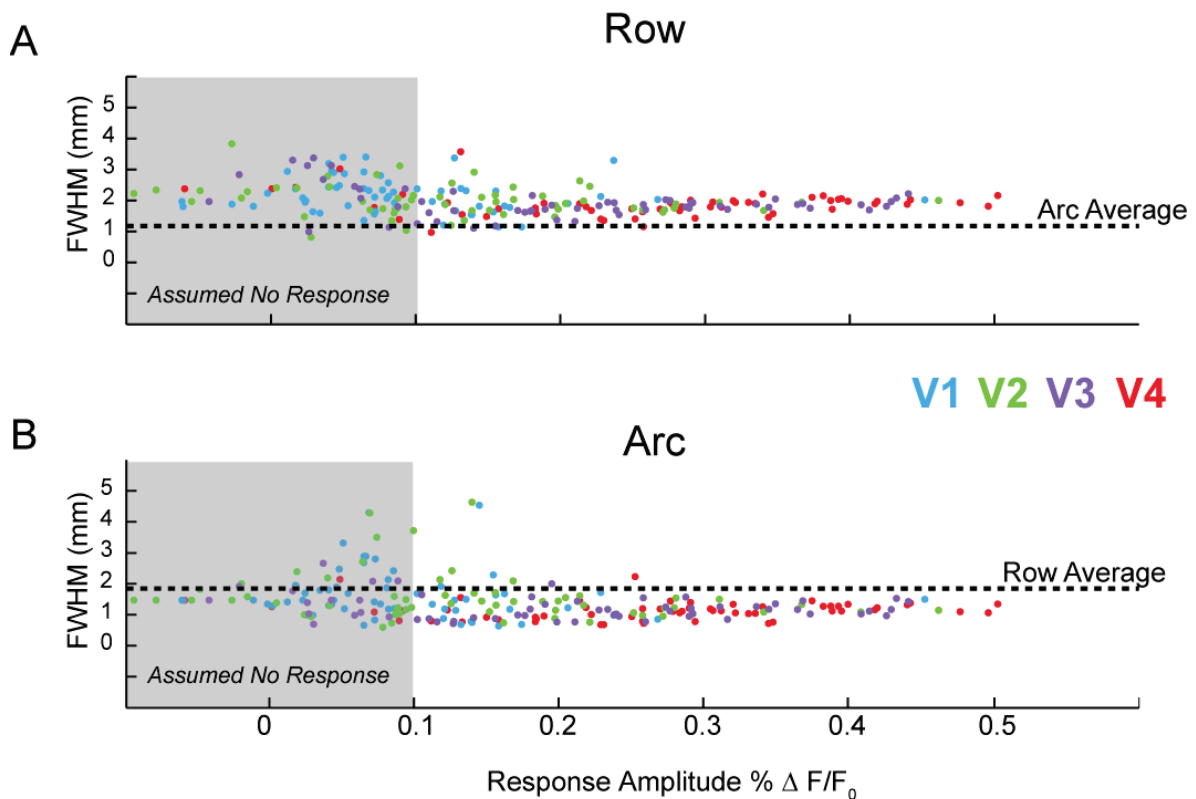
Figure 4-3 Single Trial Row-Arc Asymmetry   A) Full with at half max (FWHM) of Gaussian fits to the row axis from all single trials from one animal. Trials from different velocities are shown in different colors. Gaussian fits can fail for low amplitude responses, assumed to represent primarily no-response trials (gray box). As the amplitudes increase the variability decreases and single trials cluster around a single consistent point. The average arc response is shown (from B) is shown as a dotted line. The data is depicted similarly for the arc axis in B, showing the conserved asymmetry.

85

4.3.2 Spatial asymmetry is conserved across single trials

This asymmetry is also remarkably conserved on single trial representations. We performed the same analysis described above on images from single trial VSD. In Figure 4-3, we plot the peak response amplitude observed for each single trials versus the FWHM of a Gaussian fit along both the row (Figure 4-3A) and the arc (Figure 4-3B) axis. Trials are collected from four different velocities as indicated by color. There is a great deal of variability in the Gaussian fits to single trials, particularly at low response amplitudes. Low amplitude trials likely represent no-response trials. However, as the observed response amplitude increases, this variability dramatically decreases and a consistent spread is observed. The differences in the row and arc spreads are conserved across all single trials. The mean of the FWHM (spread) row and arc analysis is plotted as a dotted line in both Figure 4-3A and B. This asymmetry was not scaled or dependent on the absolute response amplitude. This was true across all data sets. In the one whisker in which arc spread was greater than the row spread, this same reversed trend was also reflected in the single trials.

86
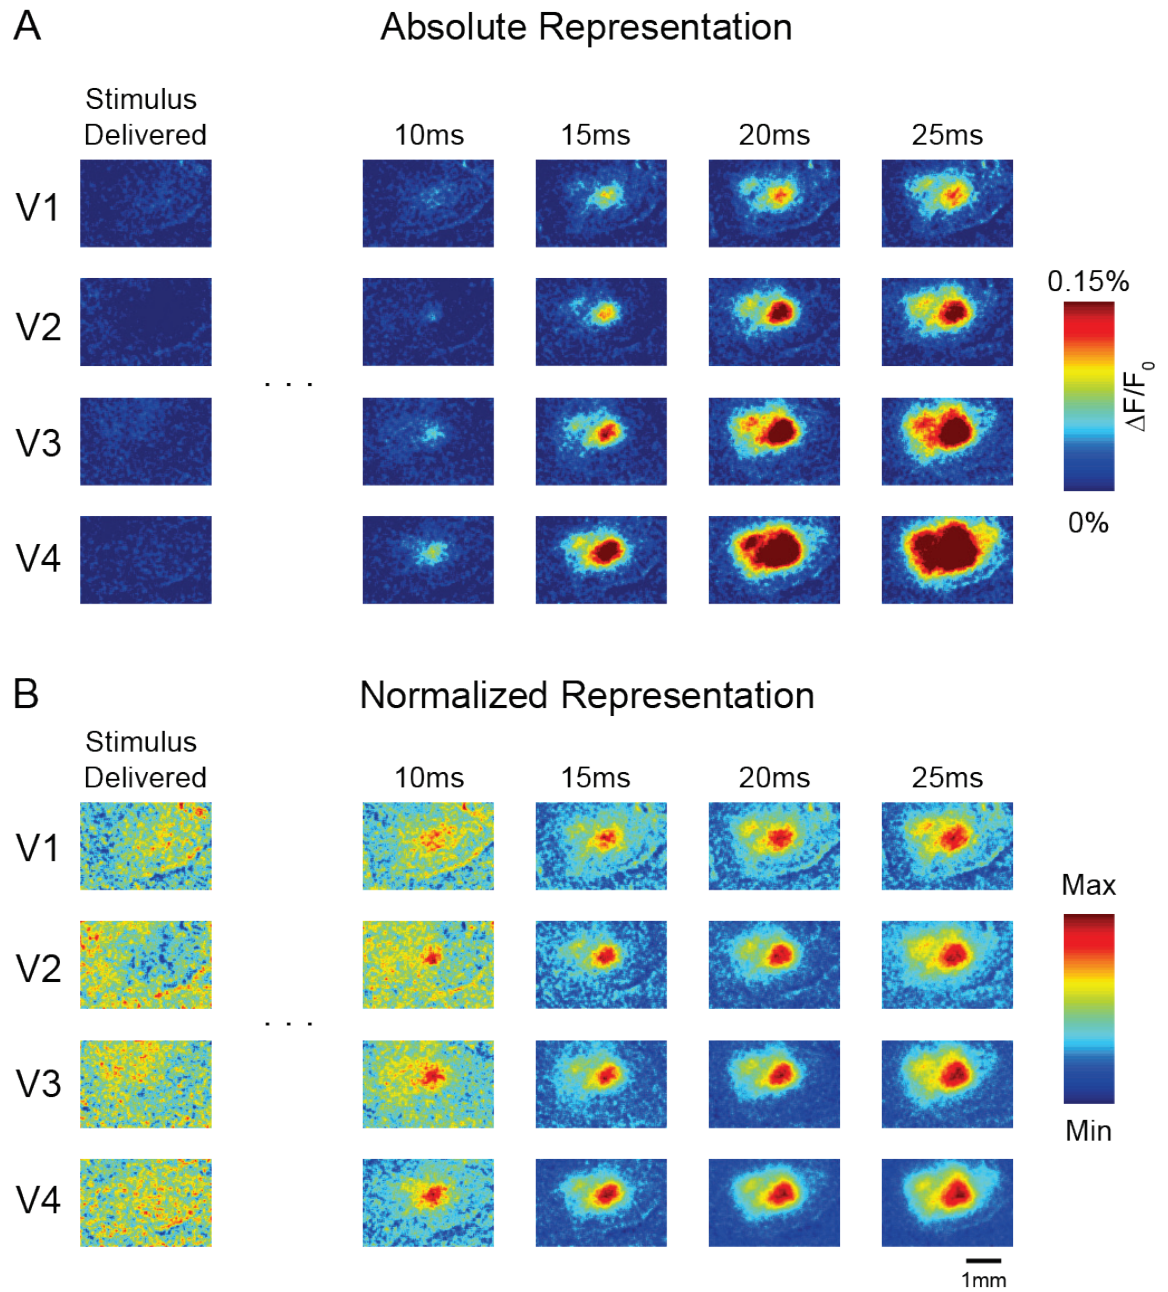
Figure 4-4 Normalized Color map A) Trial-average images are shown on an absolute color map for five frames and from four velocities. It appears in this representation that area increases with velocity. B) When each image is plotted such that it takes up the entire range of the color map (normalized) the differences between velocities disappear. However, differences that occur with time remain.

4.3.3 Trial-Average Response Amplitude Differences Interpreted as Different Spatial Activation

Previous correlations between spatial spread and velocity input were performed on trial-average images (Wang et al., 2012). We show trial average images to the four velocities in this study in Figure 4-4A. Qualitatively, it did appear as if the area of activation increased with velocity (down a column) and in time (across a row). If one were to count, for example, the number of red colored pixels in these images it would be simple to conclude that the area of activation increases with velocity. It is this qualitative assessment that led to the original spatial quantification methods that showed correlation between velocity and spatial spread. However, in the same data viewed on a color map in which each image is allowed to span the entire range, Figure 4-4B, it was easier to observe that increasing velocities are in fact scaled representations of other velocities.

Unfortunately, the Gaussian method used to fit the single trial is not robust to the dramatic changes in amplitude observed in trial-average data. As seen in the single trial data sets, Gaussian fits to very small responses often fail to represent the true spatial spread (see the *assumed no response* box in Figure 4-3) and report impossible FWHM values. Although this method is reliable with adequate signal to noise ratio, it can confound measurements at small velocities with low signal to noise ratios.

We developed a threshold metric for the quantification of spatial spread independent of response amplitude and it is described in detail in Appendix A. In Figure 4-5 we summarize these data across seven whiskers from six animals. We use two metrics of

88

spatial activation. One is a simple threshold at two standard deviations above pre-stimulus noise. This metric is equivalent to the metrics used in previous studies (Polley et al., 1999a, 2006; Wang et al., 2012). Using this method, there is an increase in spatial spread with velocity. However, this effect would be observed in all cases in which the amplitude also increases between experimental conditions even if the responses were just scaled versions (see Appendix A). We developed a relative metric to normalize the data and assess spatial spread independent of amplitude. Briefly, this metric normalized each image relative to both the peak and noise measurements in the image. Using this metric, we observed no amplitude-independent correlation between velocity and spatial spread, consistent with the trends observed in the single trial data.

4.4   Discussion

Here we demonstrated that area activated in response to a whisker deflection is asymmetric, spreading more along the rows than arcs. This asymmetry has been reported previously in literature and is thought to represent differences in the length of axonal projections that can be observed through axon tracing methodologies (Petersen et al., 2003a).

This asymmetry was also strikingly conserved for all single trials, independent of the response amplitude. Each single trial is a scaled version of every other single trial. Since VSD is thought primarily to reflect sub-threshold activity (Petersen et al., 2003b), it has been argued that an absolute threshold should be applied to the data to estimate area of spiking activity. This has actually been used to support some theories that directly

contradict this probability of activation hypothesis (Ollerenshaw et al., 2014). The single

trial data, however, does not support this interpretation.
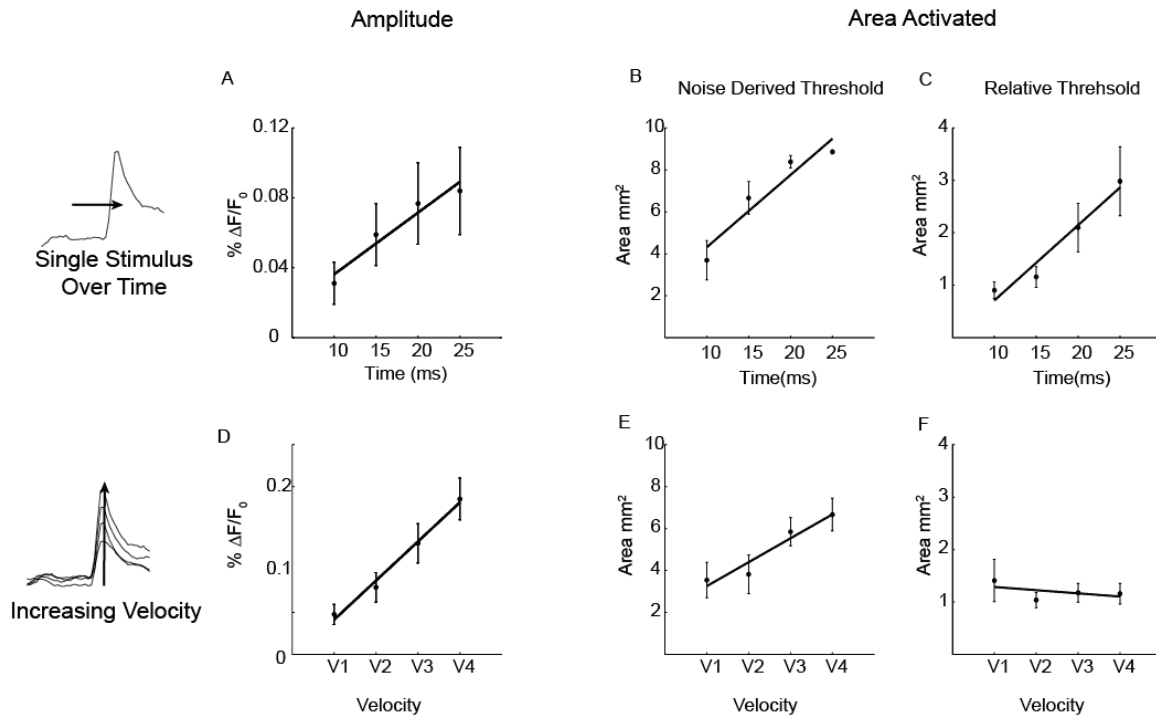
Figure 4-5 Amplitude-Dependent and Amplitude Independent Analysis   A) Response amplitude increases with frame number. B) Analysis with a noise-derived threshold shows that the area that is activated over time increases. C) Using a relative threshold, the area also increases over time. D) The response amplitude increases with velocity. E) Using a noise-derived threshold, the area appears to increase F) Using a relative threshold, the area does not increase with velocity.

First, a given VSD pixel likely does not include only cell somas. Because the physical

membrane surface area of a neuron is primarily determined by the size of its dendritic

and axonal arbor (Petersen et al., 2003a), it is likely that many pixels primarily represent

membrane potentials axons or dendrites that do not directly relate to spiking activity. If

90

the same neuron were to fire two spikes instead of one, the amplitude of the signal would increase across multiple pixels (the exposure of these frames is 5ms) and we would observe a greater area of activation. One neuron firing twice should not be interpreted as a change in area. Second, with the observation that single trial response amplitude is not reliably modulated by intensity, it seems troublesome to conclude that the area that is activated in every trial can double, triple or quadruple in size from one trial to the next.

The use of a normalized color map to represent spatial neural data will be important in other fields of neuroscience, specifically cortical plasticity and the study of reorganization of cortical maps. Most spatial data is analyzed using trial-average data. Absolute thresholds from trial average data may result in flawed conclusions if mapped to the single trial. The consequences of this observation are explored more fully in Appendix A.

We conclude that there is no change in spatial spread sufficient to separate or decode single trials into velocity categories using VSD images. Instead, individual responses to all whisker deflections can be considered scaled representations of every other response maintaining the distinct asymmetric spread along rows and arcs. This is consistent with the probability of activation hypothesis that considers all whisker responses to be stereotyped, but unreliable, activation of identical cortical circuitry.

91

CHAPTER 5   Emergence of feature-selectivity in a spatiotemporal neural code

5.1   Introduction

A central tenet of sensory neuroscience is that all our sensory experiences and perceptions are encoded directly in the activity of our neurons. At first glance, the data presented in this dissertation seem to either make this assertion untenable, or suggest that whiskers are not particularly useful as a sensory organ.

First, we have shown that from the perspective of an ideal observer of activity of a single whisker in layer 2/3, a response may or may not be detected following a whisker deflection. If a response is detected, the observer will have no predictive power about the intensity of the stimulus from either the response amplitude or the spatial extent of the response. Additionally, upon observation of a response, an observer will then have a reduced probability of detecting subsequent stimuli. This does not seem like an optimal starting point for reliable sensory encoding.

However, this is true only if an observer has access to information from a single whisker or region in space. Returning to the thought experiment that motivated this work, a powerful antidote to unreliability is redundancy. The whisker system is redundant. Most noticeably, there are many whiskers. Thirty-three whiskers exist in the primary barrel field alone (Land and Simons, 1985). Each whisker maps to a different cortical column.

If the whiskers are used as a unit, then an ideal observer with information from all barrels representing all whiskers could reduce uncertainty associated with the features of the stimulus. The unreliability of one whisker as a detector, instead of a detriment, becomes a central feature of a spatiotemporal neural code.

Many behavioral discrimination tasks require multiple whiskers. Performance of aperture discrimination degrades linearly as whiskers are trimmed (Krupa et al., 2001). A similar effect has been observed in texture discrimination (Carvell and Simons, 1995). However, one whisker is sufficient for a simple detection paradigm (although performance never reaches 100%) (Stüttgen et al., 2006; O'Connor et al., 2010a; Waiblinger et al., 2013; Ollerenshaw et al., 2014). Taken together, it appears that multiple whiskers are necessary for discrimination, but that one whisker is sufficient for detection.

In this chapter, I demonstrate using Monte Carlo simulations that stimulus features, specifically stimulus intensity and direction of motion, can be encoded using multiple whiskers within the probability of activation hypothesis. I model an array of detectors both in space and in time. First, I consider the encoding of stimulus intensity using whiskers deflected simultaneously, as model for wind stimuli. Wind stimuli (an air puff) have been observed to move all whiskers nearly instantaneously (even if they have directionality) (Ollerenshaw et al., 2012). Wind is a particularly relevant stimulus for rodents that are use olfaction as a primary sense. Probabilistic representation of stimulus intensity across space is shown to be sufficient to infer intensity information. Additionally, I show that this neural representation of intensity is not sensitive to the loss

93

of a whisker. Second, I simulate the response to a moving object interacting with the array such that arcs of whiskers are deflected in series. The predicted neural activity on average is identical to what has been reported in literature; however, this trial-average does not suggest that direction of motion could be reliably encoded. In contrast, there is sufficient patterned activity on a subset of single trials within the probability of activation framework that would allow for decoding of direction of motion.

5.2   Methods

5.2.1   Spatiotemporal Model

This spatiotemporal model was built in Matlab (MathWorks, Natick, MA). Whiskers are modeled as single pixels with probabilistic responses: zero for no response, one for response. Nine whiskers are simulated (a 3x3 array). The same mapping from whisker deflection velocity to probability of response is used as was used in Section 3.2.7. This curve fit was collected in awake behaving animals and was previously published (Ollerenshaw et al., 2014).

To add a temporal component to this model, we added a history component that lowers the probability response to subsequent pulses on a given trial only once a response is observed. The history function h(t) was a sigmoid of the form:

$$h(t) = \frac{1}{1 + e^{\frac{-(t-\beta)}{\alpha}}} - 1$$

94

Where β was the midpoint of the sigmoid and α was the slope of the sigmoid, t was the frame number following an observed response. One frame represented 10ms separation in time. Responses recovered to 50% of the original response probability at 10 frames (β=10, α=10). Whiskers in each arc were deflected every 2 frames. Distance between whiskers was equal to the maximum number of rows or columns crossed in order to reach the measured barrel, such that the distance between a whisker in position $[x, y]$ and whisker in position $[i, j]$ was $max(|x - i|, |y - j|)$. Suppression at a distance of [0 1 2] of the responding whisker of were scaled by 100%, 50%, and 40% respectively.

5.2.2 Detection and Discrimination Performance

Using the spatial temporal model described above, the expected number of responding whiskers and the associated distributions were estimated using 100 simulations of the spatial profile. The number of responding whiskers defined the response variable (r). A single responding whisker was sufficient for detection on a single trial. Discrimination performance was calculated identically to the optimal classifier performance described previously (2.2.6). For direction of motion, two responding whiskers in two different frames (separate points in time) resulted in the correct detection of motion. The percent-successfully decoded trials was determined by Monte Carlo simulations.

5.3   Results

5.3.1   Simultaneous Deflections of Multiple Whiskers

First we show that an array of whiskers responding independently can encode stimulus strength. In Figure 5-1, we show examples of possible spatial responses. The probability of response for each whisker was determined by the deflection velocity as shown in Figure 5-1B. Each column in Figure 5-1C shows five representative single

Figure 5-1 Distinct Spatial Representations of Identical Stimuli A) This model assumes nine whiskers are deflected simultaneously. Each barrel is represented by a pixel. B) The probability of response for each whisker (pixel) is determined by the probability of response curve. The simulated velocities are highlighted on the axis and with dotted lines. C) Five example spatial representations to each of the six velocities. Red pixels represent responding barrels, blue pixels represent no response.

trials for each of the six velocities. The same stimulus, in this case, nine whiskers deflected at identical velocities at the same time, was encoded with a different spatial profile on each trial.

The spatial response profiles are not unique to a given velocity; for example consider the example single trials from V4 and V5. However there is little chance of the minimum velocity (V1) causing the same spatial profile as the maximum (V6). We assume that stimulus intensity in this framework was encoded by the number of responding barrel columns. The expected distributions of responses to all nine whiskers are shown in Figure 5-2A. These distributions were similar to what was observed as trial-average classifiers in Chapter 2. The expected discrimination performance matrix is shown in Figure 5-2B. Given nine whiskers and the six velocities, the optimal classification performance of single trials was 48.5%. One whisker response was considered sufficient for detection of a stimulus on a trial (Ollerenshaw et al., 2014). The presence of a stimulus was detected on 93.7% of trials.

Returning to the thought experiment that motivated this work (1.5.1), we tested to see how robust this spatiotemporal neural code was to the loss of a whisker. This data is shown in Figure 5-2D. With 8 whiskers instead of 9, the discrimination performance decreased by less than 1% to 47.8% correct. Meanwhile, 91.9% of trials were detectable (2% less). We conclude that this coding framework provided reliable sensory information even if a whisker is lost.

98

Figure 5-2 Robust Performance with the loss of whisker A) The expected distributions for the number of responding whiskers out of nine for six velocities. These distributions are calculated from the single trials shown in Figure 5-1. B) The expected classification performance given these distributions. Classification occurs based on the number of responding whiskers (*r*). C) Expected response distributions and D) discrimination performance are not affected by the loss of a whisker (8 whiskers).

5.3.2   Direction of Motion

A stimulus feature that is thought to be important to sensation with whiskers is the direction of motion of an object. Specifically, this is thought to be important in distance based tasks such as aperture discrimination (Krupa et al., 2001). Previously, it has been

shown that if whisker arcs are deflected in series, there is only a response to first arc and all other responses are greatly suppressed (Drew and Feldman, 2007). This makes it unclear how directional information could be encoded across an array. We consider this observation in the context of probability of activation hypothesis.

We used a 3x3 array of whiskers that respond probabilistically as in the previous simulation. When a response occurs, a non-specific suppression of all adjacent whiskers follows for a period of time (see Methods). Using this model, we simulated deflections of whiskers one arc at a time. Assuming the whisker arrangement depicted in Figure 5-1A, this is equivalent to deflection all whiskers in the A column, then B, then C. This same protocol was previously tested in an anesthetized rodent preparation (Drew and Feldman, 2007). Like this paper, the probability of activation model predicts a strong response to the first arc on average, with a much smaller average response to all other arcs. From the trial-average response alone, it is unclear if direction information was encoded. A deflection of either direction could have caused just one arc to respond. However, more information could be observed in our probabilistic representations of single trials.
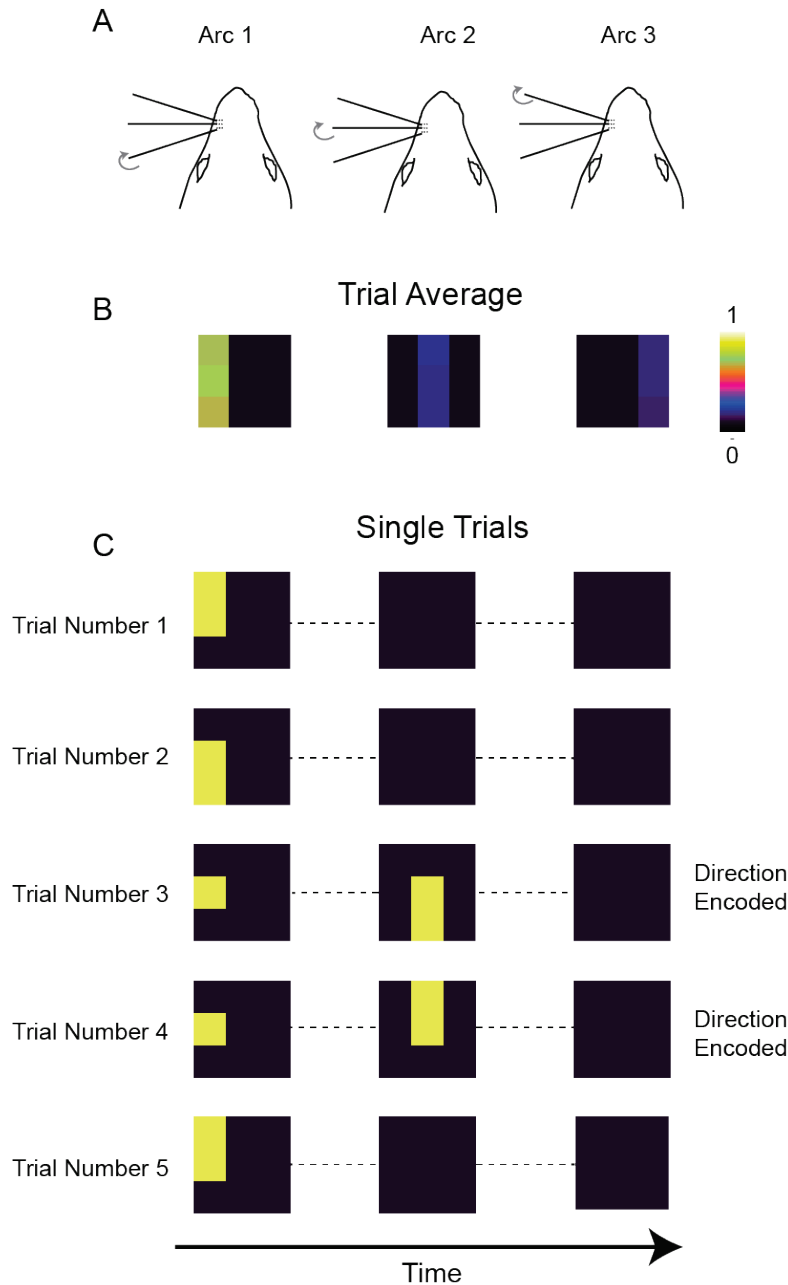
Figure 5-3 Probabilistic Encoding of Direction of Motion   A) Whiskers arcs are deflected in series (simulated responses) B) Trial average responses show a significant response only to the first arc. The second and third arc are greatly suppressed. C) Direction of motion is encoded on a subset of single trials within the probability of activation framework

101

Frames from single trials of this spatial temporal neural code are shown in Figure 5-3. Our model predicts high variability across both time and space in response to sequential deflections of whiskers in arcs; however, directionality can in fact be decoded, at least in some trials, in this framework. We observed that while most trials had at least one responding whisker in the first arc, not all trials responded to the first arc only. Often, at least one other response occurs somewhere in the whisker array such that the direction of motion of the deflections could be decoded. In this simulation, we had sufficient information to decode the direction of motion of the stimulus in 41% of the simulated trials.

5.4   Discussion

Here we have used simulations to show how a probabilistic and spatiotemporal neural code could encode stimulus features including intensity and direction of motion using multiple whiskers. We have shown that the number of responding whiskers was sufficient information for decoding of stimulus intensity correctly approximately half the time (chance is 16.6%). Additionally, we have shown that while the trial-average response to whisker arcs deflected in series did not appear to allow for the encoding of the direction of motion of an object, an array of probabilistic encoders provided this information more reliably than could be predicted from the trial-average alone.

Our simulations are consistent with evidence collected from electrophysiological recordings using multiple whiskers in literature. Specifically, we consider a paper published by Drew et al. in 2007. In this paper, the authors recorded the activity in a

102

single barrel while 9 whiskers were deflected simultaneously. They observed that the average response when all nine whiskers were deflected simultaneously was not different from when the primary whisker corresponding to the recorded barrel was deflected alone. This is in agreement with our instantaneous whisker deflection model.

This paper also tested the direction paradigm used in this simulation. Our model predictions are also consistent with the experimental observation presented in this paper. The trial-average response to arcs deflected in series showed that the first deflected arc responded while responses to all other arcs were suppressed. In the experimental paper, this was taken as evidence that a temporal code could not be used to encode direction; however, these authors did not have access to the single trials. In our simulations, we showed that a single second whisker response, sufficient to decode direction, occurred on a large subset of trials. We could not predict, however, which barrel would respond on a given trial. We predict that had the authors in the experimental paper had access to separable activity from multiple barrels, more reliable activity (although sparse) would have been observed to subsequent arc deflections.

The behavior of rodents is generally reported as response probabilities. If the animals perform above chance, then we conclude that animals are capable of performing tasks. Returning once again to our central tenet, all available perceptual information must be encoded in the neural representation. Performing reliably above chance is not evidence of a stable and reliable percept, only that sufficient information is available on a subset of

103

trials. When an animal can only perform a task 20% better than chance, then the necessary information may only be encoded 20% of the time.

Importantly, the simulations in this chapter have suggested that within the probability of activation framework different spatial neural representations can equivalently represent the same stimulus. This allows multiple solutions or pathways to encode the same stimulus information. This spatial type of redundancy increased the reliability of the neural code.

CHAPTER 6   Accumulation of Errors and Mechanistic Redundancy

6.1   Stochasticity, Independence, and Redundancy

Creating a stable sensory percept in the presence of variability and uncertainty is our brain's greatest challenge. Large-scale variability can be created from accumulation of small stochastic errors. At the lowest level, the biological hardware that performs neural computation is highly stochastic. It is impossible to predict with certainty whether a given ligand will bind with a given receptor as each of the basic processes have elements of randomness – the proximity of the molecules may be determined by a diffusive random walk, the binding kinetics by the instantaneous arrangement of the molecules. The success of even one neuron firing is dependent not on one stochastic event, but on the compounding effect of millions.

When individual events are stochastic, redundancy creates reliability. At the level of the hardware, we see that a cell does not rely on any single ligand binding to a single receptor, but on the collective activity of hundreds or thousands of such interactions. At a higher level, information is not encoded by the stochastic spiking activity of a single neuron, but by a population of neurons, such that the failure of any one single neuron provides a negligible impact on the actual information content available for perception (Averbeck et al., 2006). This assumption works well if neurons are independent;

however, if the input to these neurons is shared, even independent neurons cannot compensate for failures of the common input. A serial pathway operates such that success at one level depends on success in the previous. As such, errors due to stochastic variability are permanent. Errors can only accumulate with time.

In this final chapter, I propose an accumulation of errors model to explain the probabilistic dynamics observed in layer 2/3 of the barrel cortex. Like Hubel and Wiesel's vision model, the model features a series of hierarchical structures; unlike Hubel and Wiesel's model there is no increased in complexity or feature-selectivity of individual neurons at higher levels. Instead, response reliability is modulated by accumulation of random, unpredictable and corrosive encoding errors. Although I have no direct evidence of this model beyond the observation of supra-granular VSD dynamics presented in this thesis, it is intriguingly simple and remarkable as a predictor of neural dynamics. Importantly, it allows for a new type of redundancy in this pathway: mechanistic redundancy.

6.2 **Fundamental Observations**

Three observations were fundamental to the formulation of this conceptual model. The first was from a significant body of existing literature, the next two were observations from Chapters 2 and 3 of this dissertation.

First, the magnitude of the neural response decays as it progresses from lower levels of the pathway to the highest level of the cortex. Near the periphery, trigeminal afferents

and brainstem neurons spike robustly on each trial (Gottschaldt and Young, 1977; Minnery and Simons, 2003). This suggests that intensity-dependent differences in amplitude exist early in the pathway. At the next level, the VPM, single trial responses are less robust but still often result in multiple spikes (bursts) (Simons and Carvell, 1989). The maximum reliability of individual neurons across trials is approximately 80%, although neurons with 100% response reliability occur occasionally (Wang et al., 2010). At higher levels, layer 4 of cortex and then again to layer 2/3, the signal drops off dramatically. Layer 4 neurons respond to a stimulus with only a single spike (Simons and Carvell, 1989; Wang et al., 2010). Even the single spikes are less reliable, capping out at around only 60% of trials on average at a high intensity. At the top level, there is some disagreement if layer 2/3 neurons even spike at all (Sato et al., 2007; Sachidhanandam et al., 2013). At best, the spiking reliability of layer 2/3 neurons has been described as sparse (Jadhav et al., 2009). Importantly, all of these are assessments of the average neuron, not any given neuron.

Second, at some point in this sensory pathway the absolute amplitude is decoupled from the stimulus intensity. Interestingly, the single trial amplitudes are not systematically decreasing in strength like the average responses. Some single trial responses to the smallest stimuli equal the magnitude of the responses to the largest. Small responses, barely distinguishable from noise, were also observed in response to the fastest velocity. This modulation and decoupling of response amplitude was independent of stimulus intensity and appeared to allow either suppression or facilitation of single trials to occur.

107

Finally, the limitations and failures of the probability of activation hypothesis were informative. The probability of activation hypothesis is not capable of predicting or explaining the low response probability to short (less than 100ms) inter-stimulus intervals. A model in which the probability of observing a response to one whisker deflection depends only on its immediate history on a trial to trial basis (from Chapter 3) works well at inter-stimulus intervals (ISIs) of 100ms-200ms. However, if whisker deflections are given at shorter ISIs, the observed suppression is much greater than predicted by the number of observed responses in the final level of the pathway.

Figure 6-1 Short ISI Problem in One Animal A) One animal from the fixed-test condition showed no responses to the test pulse at 100ms ISI for any condition pulse velocity (data from all condition pulse velocities are shown) B) Responses to test pulses were observed at longer ISIs C) Complete recovery from the bimodality was not observed at the longest ISI recorded for this animal

In Chapter 3, I made a specific point of highlighting the variability of the temporal recovery between animals. I showed one animal that had no test pulse responses to any condition-pulse velocity at a 100ms ISI. At longer ISIs, the test-pulse responses returned

(Figure 6-1). In this animal, the responses at 100ms ISI were similar to the 50ms ISI responses in all other animals. There is a near zero probability of response to the test pulse at any condition pulse velocity if the ISI is short. The observation that short-ISIs were not well approximated by the probability of activation model suggests that a neural response to the first pulse and corresponding suppressive effects did exist at some point in the pathway, but were not observed at the last structure.

In summary, three key observations underpinned this model:

1) The spikes per stimulus and reliability decay with each successive level in the pathway
2) The decoupling of response amplitude and stimulus intensity can result in either suppression or facilitation of single trials
3) Short ISIs dynamics are not predicted by the number of response and no-response trials observed in layer 2/3 of the cortex

6.3   Accumulation of Errors Model

The accumulation of error model takes an agnostic but realistic approach to sensory encoding. Like Hubel and Wiesel's model, there is a series of hierarchical structures. At each structure, there is some potential for error as the signal is influenced by a random component. The influence of this random variable is unpredictable and irreversible. The majority of these random contributions are small and do not dramatically influence the propagation or integrity of the signal. However, occasionally a random contribution results in a substantial change to the signal. Whether the random element is small or large, the signal has been permanently changed as it propagates along the pathway.

109

Figure 6-2 Accumulation of Errors Schematic A schematic of the accumulation of errors model is shown above. Initial states are determined linearly with respect to stimulus intensity and then propagate through a series of three structures in a history-less manner. A random component is added at each structure and influences the strength state (S_n) prior to input into the next structure.

A schematic of the accumulation of errors model is shown in Figure 6-2. There is a linear relationship between stimulus intensity and an initial strength variable S_0. S_0 is then propagated through three neural structures in series (S_1, S_2, S_3). Each structure is identical. At each structure, there is an additive random component to the signal strength

and a non-linear threshold. If a signal on a given trial is affected by the random component such that it does not surpass the threshold level, it fails to propagate (Figure 6-3).

Figure 6-3 Probabilistic Strength State Propagations Within each structure, the initial strength is influenced by a random component (depicted as a series of random numbers). The new strengths are bounded such that if $S_n < S_{th}$ then $S_n = 0$. Once a trial reaches a strength state of zero, there is no probability of recovery.

This system is defined by five variables: 1) the standard deviation (σ_{noise}) of the random component (sampled from a normal distribution); 2) the mean of the random component (μ_{noise}); 3) the strength threshold, S_{th} at which a signal propagates; 4) number of structures (n); 5) a strength upper bound (S_{max}).

If $S_n < S_{th}$ at any point in the simulation, that trial has failed and has zero probability of recovery from this state. If $S_n > S_{max}$, then the signal is reset to S_{max}, to represent the maximum capacity of a biological system (all neurons active). Reliability of response

111

was considered as observer of activity in the final structure. Formally, this model can be considered a continuous-state Markov chain where the random component is:

$$\{\xi_t\} \sim N(\mu_{noise}, \sigma_{noise}{}^2)$$

The signal strength manipulation at each structure (n) is

$$S_{n+1} = h(S_n + \xi_{n+1})\mathbf{1}\{S_n > 0\}$$

where h(s) defines the absolute bounds on stimulus strength and the non-linearity:

$$h(s) := s\mathbf{1}\{S_{th} \leq s \leq S_{max}\} + S_{max}\mathbf{1}\{s > S_{max}\}$$

The indicator function $\mathbf{1}\{A\}$ evaluates to 1 for all s that exist in set A and zero otherwise. The model was solved using Monte Carlo simulations in Matlab (MathWorks, Natick, MA). At least 1000 simulations were included in each average observation.

6.4 Consistency with Fundamental Observations

The accumulation of error model is a conceptual model; the numerical values, including the units of S_n and associated constants have no units and no direct biophysical meaning. However, it was built to be broadly analogous to the whisker pathway. The three structures could be thought of as the brain stem, the thalamus and the cortex. Structure 3 represents the top level of this hierarchy (or the observation point), analogous to the

112
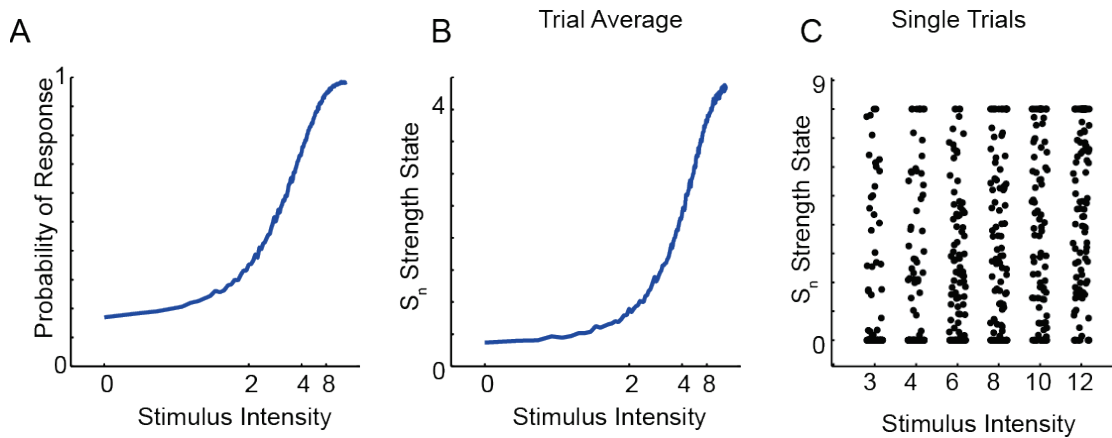
Figure 6-4 Markov Process Replicates Average and Single Trials Trends A) The probability of observing non-zero responses at the final structure for each original stimulus intensity B) The average stimulus strength state S_n for each stimulus intensity observed at the final structure increases with stimulus intensity C) The single trial distributions show decoupling of individual responses from the original stimulus intensity

cortical VSD imaging presented in this thesis. If any non-zero trial at Structure 3 is considered a response, a response reliability curve can be generated from the perspective of an observer of this final structure. An example of a response reliability curve from the accumulation of errors model is shown in Figure 6-4A. This curve is sigmoidal, consistent with both behavioral predictions and the probability of response data (Stüttgen and Schwarz, 2008; Ollerenshaw et al., 2012). The average strength state increases with stimulus intensity (until it plateaus) as seen in Figure 6-4B; however, there is no structure to the single trial distributions representing a decoupling of the observed neural strength and the stimulus intensity (Figure 6-4C) consistent with the observed single trial VSD

113

distributions.   There are a large number of trials at the imposed minimum and the imposed maximum, but again this is meant to be conceptual.

In the accumulation of errors model, the average signal intensity S_n, decreases with increasing structure number. Figure 6-5 shows the average strength state, S_n, after each structure. This is consistent with the first fundamental observation, the signal decays at each level of the pathway. This is also true to some extent at all other stimulus intensities, but is best observed in the plateau region.
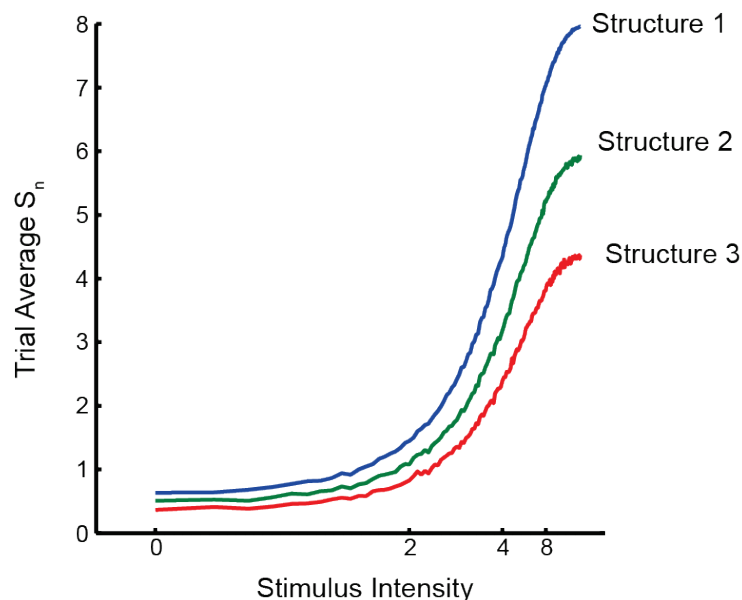
Figure 6-5 Average signal decays with each structure. The average S_n as observed at each structure in the pathway. The signal decays on average as it is processed sequentially.

There is no temporal component to the model (it is a future goal). However, the essential element to account for short ISI suppression exists. Conceptually, even though no

114

response was observed for most trials at a sub-threshold velocity at the last structure, these responses did exist in the pathway up until some point. Consider the sample path of a trial that successfully propagated to Structure 2 but did not propagate to Structure 3. An observer of level 3 would not be able to predict suppression to the subsequent pulse. However, the signal did reach Structure 1 at a sufficient strength to engage an inhibitory response, so suppression to subsequent inputs on this simulation would still be observed. A necessary assumption to account for short ISIs is that the inhibitory response at lower level structures is shorter than at higher structures. This assumption is consistent with thalamic versus cortical data in the whisker pathway (Simons and Carvell, 1989).

6.5   Regulation of Detection in a Hierarchical Markov Process

The most powerful aspect of the accumulation of errors model is the observation that a single modification, at any point in the pathway, does not affect the response probability of all stimulus intensities equally. Consider the three curves shown in Figure 6-6A where the expected probability of response given different propagation thresholds S_{th} (lowest, black) and then increasing (green and blue). The larger threshold makes it slightly more difficult to propagate between structures without failure. Even though the same bias exists on all trials, it specifically affects the probability of response of intermediate stimulus intensities. Figure 6-6B shows the difference in response probabilities from the low threshold to a higher threshold. The greatest change in response probability occurs at intermediate (threshold) stimuli, with minimal or no change extreme intensities.

115

Figure 6-6 Parameter Sweep of Accumulation of Errors Model A) Increasing the propagation threshold (S_{th}) increases the threshold B) This change differentially affects threshold velocities (gray box) C) Increasing the variability of the random component decreases the slope of the sigmoid D) Threshold stimuli are preferentially affected, either increasing or decreasing the probability of response; E-F) Adding additional structures changes the probability of response curve similarly to changing the propagation threshold

In fact, manipulation of any parameter of the model preferentially affects threshold stimuli (Figure 6-6). Every parameter except for the variability of the random component (σ_{noise}) modulates the curve such that the detection threshold shifts left or right on the graph. σ_{noise} does not affect the threshold, but the slope of the sigmoid (Figure 6-6C).

This non-linearity predicted by the accumulation of errors model exists in behaving animals and has been documented both in rats and humans in somatosensation (Stüttgen and Schwarz, 2008; Bari et al., 2013; Waiblinger et al., 2013; Ollerenshaw et al., 2014). I will consider the rodent example here specifically. In Figure 6-7, there are two curve fits from a behavioral experiment performed and published by Douglas Ollerenshaw in the Stanley Lab (Ollerenshaw et al., 2014). The raw data has been removed for clarity, but can be seen in the original paper. In this detection paradigm, a whisker was deflected at a given velocity as an animal responded by licking a waterspout. The probability of an animal correctly detecting the stimulus is determined by the whisker deflection velocity (Figure 6-7, green curve). In a subset of trials, an adapting stimulus (analogous to the condition pulse in Chapter 3) was delivered on the same whisker prior to the delivery of a probe. This data is shown in red in Figure 6-7A. The difference between the green curve and the red curve is plotted in Figure 6-7B. Given an equivalent perturbation across all stimulus intensities (an adapting stimulus), the response probabilities are not equally affected. The most dramatic effect is observed on threshold intensities (near a 50% detection probability).

Figure 6-7 Non-linear changes in detectability in awake animals A) Psychometric curves show the probability of a rat responding to a whisker deflection in a non-adapted state (green) and an adapted state (red). The gray boxes highlight stimulus intensities near the perceptual threshold B) The differences between the non-adapted curve and the adapted curve show that adaptive perturbation of the system preferentially affects threshold stimuli.

Consider this thought experiment in the context of the accumulation of errors model. A rat, given the same stimulus on two whiskers, responds 65% of the time to one whisker (S-) and 90% of the time on another whisker (S+) in some arbitrary initial state. In other words, the whiskers are not equally detectable. If a subtle handicap then perturbs the system such that it is marginally harder to propagate signals, we would predict that the animal would then respond less to all stimuli on average, but the change would be larger on the S- whisker as it was originally closer to threshold than to the S+ whisker. The ratio (S+/S-) would be predicted to increase. This is exactly what was observed in the adaption behavior task. Note that the author's interpretation of the effect was different as

118

they started with the assumption that all whiskers are equally detectable. However, the expected non-linear change in detectability from a whisker operating near threshold and one operating above threshold was sufficient to explain the increase in specificity in responses. A single perturbation differentially affecting the response probability of different stimuli would allow the creation of specificity in the neural code. This same logic could also apply to single neurons – a perturbation would be expected to affect less reliable neurons before more reliable neurons, at least until a neuron is so unreliable as to be on the other side of the effect. This effect would be increasingly true at each higher level in the pathway structures. More structures allow for more specificity in the response.

Previously, I have argued that the whisker system primarily functions as an array of unreliable detectors. Here, I argue that the detectability of stimuli can be selectively decreased or enhanced by regulating how easily a signal can propagate through the system. In all cases, threshold stimuli are more sensitive to perturbations than strong stimuli, allowing for the active modulation of the dynamic range without significant loss of function or feature-selectivity across the entire array of detectors.

6.6 **Always Sufficient, Never Necessary: the power of mechanistic redundancy**

Although the idea that our sensory systems may be inherently unreliable is existentially troubling, intellectual solace can be found in the robustness of the system as a whole. I have resisted providing any specific biophysical explanations for the source of the

119

random component of the accumulation of errors model. This is because it does not matter why or how the random process is random, only that it exists. Similarly, when the system is perturbed such that the probability of response is modulated, it does not matter what causes this perturbation. Any perturbation could be expected to influence the system as a whole in a predictable way.

The detectability in the behavior experiment introduced earlier was attributed to changes in thalamic synchrony (Wang et al., 2010; Ollerenshaw et al., 2014). In the accumulation of errors model, decreases in thalamic synchrony would make it more difficult for a signal to propagate on average and predict a decrease in the detectability. As such, changes to thalamic synchrony are consistent with the accumulation of errors model. The competing hypothesis, however, that adaptation simply results in a depression of synapses (Castro-Alamancos, 2004), would be equally supported by the accumulation of errors model.

Specific suppression of smaller stimuli has also been reported by the application of GABA agonists on the cortex (Kyriazi et al., 1996). An increase in inhibition would also be sufficient, given the accumulation of errors model, to predict non-linear decrease in detectability relative to original intensity. Additionally, trimming of adjacent whiskers has been shown to increase the trial-average response amplitude in a primary barrel (potentially an increase in response reliability) (Diamond et al., 1993, 1994; Rema et al., 1998). This can be blocked by the application of NMDA antagonists, which reduces firing rates and makes propagation more difficult (Rema et al., 1998). Although many

120

possible perturbations could have accounted for the original increase in response amplitude (or response probability) in this plasticity protocol, a systematic and non-biological level of inhibition resulting from blocking NMDA would be sufficient to reverse it. It would be sufficient, but not necessary to prevent the original effect.

Importantly, my model predicts that all of these effects could be observed in average firing rates, but not reliably from single trials. A given perturbation may reduce the response reliability on average, but that allows no predictive power for the sample path of any given trial.

One may interpret the flexibility of the accumulation of errors model as a weakness of the model. In fact, the opposite is true. The assertion that any one of a number of regulatory events occurring is *sufficient* to result in the desired functional output, but no one event is *necessary* is exactly like having multiple neurons act as a together as population: redundancy creates reliability. Mechanistic redundancy allows for increased functional (and presumably perceptual) reliability.

The idea that multiple possible underlying parameters could create the same functional output was first suggested and supported both computationally and experimentally in studies on the crustacean stomatogastric ganglion (Prinz et al., 2004; Gutierrez et al., 2013). In this system, it is clear that the functional output, not the tight regulation of a specific set of conduction parameters, is the key control variable. The accumulation of errors model is consistent, in fact, derived in part, from ideas presented and developed in this previous work. From this work, I hypothesize that many mechanistic solutions or

121

initial parameters can potentially explain a given observation. Tight regulation of individual parameters would not be necessary to achieve perceptual stability.

6.7   Conclusions

Redundancy is a common theme to all of neuroscience. Stochastic events are the key unit in essentially all neurological hardware from the level of diffusion in receptor-ligand interactions to the stochastic nature of a neuron spiking. Even with many neurons encoding the same stimulus, it is hard to imagine our neural code is uniquely immune to stochastic error. When stochastic processes are arranged in a hierarchical series of structures, accumulation of errors is inevitable. In fact, the most reliable aspect of the system is that information encoded in absolute values (the absolute response or firing rate) could never be reliable.

A system built of modulating the reliability of the signal would be mechanistically robust, as it does not require success *necessarily* on any particular regulatory mechanism in order to create a *sufficiently* regulated change in detectability of the signal. If our sensory systems do turn out to be mechanistically redundant, then we would be forced to reconsider what type of experimental design is evidence of what type of causality. It would become important not to mistake evidence of *sufficiency* as evidence of *necessity* or original cause.

CHAPTER 7   Discussion

7.1   A Summary

A central tenet of sensory neuroscience is that all our perceptual experiences are encoded directly in the activity of our neurons.

The tenet, however, does not compel the assumption that identical perceptual experiences are encoded by identical neural activity. It is also possible that equivalent percepts are the result of divergent neural representations.

In evidence of this hypothesis, I have shown, using the rodent vibrissa pathway as a model, that identical stimuli are not encoded identically on single trials in any of three tested dimensions: amplitude, time, or space. Given this fundamental unreliability, I hypothesize that using multiple whiskers (multiple columns) to redundantly encode stimulus information will result in a reliable spatiotemporal neural code. I propose an accumulation of errors model for creation of these probabilistic representations, which was characterized by small but numerous stochastic errors. Within this framework, I observed that modulation of detectability was a powerful method for control and regulation of neural output as any perturbation to the system did not affect all stimuli equally.

123

7.2 Limitations of the Trial-Average

When we compute trial-average measures in order to summarize, analyze or infer patterns of activity from a data set there is an implicit assumption: we assume that the average response is meaningful or capable of summarizing the underlying data in a useful way. This assumption, like all assumptions, should be tested whenever possible.

There are many examples in which the trial-average has been a valuable tool. For example, understanding the encoding of stimulus intensity in the whisker pathway has benefitted greatly from analysis of the trial-average. The trial-average model was sufficient to find that single units were sensitive to velocity (Simons and Carvell, 1989; Shoykhet et al., 2000b; Lee and Simons, 2004). Additionally, the non-linear temporal dynamics replicated in this thesis were all observed initially on the trial-average activity of single neurons (Simons and Carvell, 1989; Higley and Contreras, 2005; Boloori et al., 2010). It would be difficult, potentially impossible, to make these initial observations (and nearly every other summary of neural spiking patterns) without this simplifying assumption.

When we consider perception in the context of behavior, however, we do not assume that the brain computes trial-averages. We are capable of interacting in a dynamic sensory environment. Often, this means we only have one opportunity or experience with each stimulus. Once stimulus parameters that modulate the trial-average are identified, we must consider how this average response maps to perception. When analyzing single-unit electrophysiological data, it is common to assume that many similar neurons respond

124

independently to an input, thus motivating the trial-average as a model of a population code. This logic, however, does not account for the effects of variability of shared input. Although this has been considered in other contexts (Abbott and Dayan, 1999; Stocks, 2000; Wilke and Eurich, 2002; Moss, 2004; Harrison et al., 2005; McDonnell and Abbott, 2009). In this work, I proposed an accumulation of errors model for the development of probabilistic representations of stimulus intensity. The key inference from this model is not that individual populations of independent neurons are unreliable; in fact, each structure in the model introduces minimal error. However, when computations are performed in series, even very small inaccuracies can dramatically influence the propagation of the signal. Therefore, I predict that the validity of the trial-average neural activity as a model of the single trial experience will depend less on the independence of individual neurons within a structure or population and more on the potential for error within a hierarchical series of such structures. As a result, when we observe neural activity many steps removed from the initial encoding, we should attribute correspondingly less emphasis, and assert with less confidence, that the trial-average is representative of the single trial.

7.3 Cortical Columns and Relevance to Human Somatosensation

There many ways in which the rodent whisker system is not directly applicable to human somatosensation. Most of these I have considered in the original thought experiment in the introduction (1.5.1). However, I have not yet considered the most important way the barrel cortex is like the human primary somatosensory cortex: the prominence of cortical columns (Mountcastle, 1997; Tommerdahl et al., 2010).

Cortical columns exist in the human primary somatosensory cortex. These columns are characterized by local circuits that are highly connected vertically (between cortical layers) but with minimal horizontal connections (Tommerdahl et al., 1993; Mountcastle, 1997). Regions of skin that are highly relevant to behavior are over-represented in cortical space in all animals (Penfield and Boldrey, 1937; Mountcastle, 1997). These regions are also more densely innervated with peripheral receptors; but this increased cortical space may not just a trivial consequence of increased peripheral innervation. Given the organization of the barrel cortex and the work presented in this thesis, I hypothesize that the base unit of computation in somatosensation is the cortical column. As the size of cortical columns is largely conserved across sensory areas (Mountcastle, 1997), more cortical space directly implies more columns.

When more cortical columns redundantly encode sensory information, the probability of activation hypothesis predicts that the stimulus will be more detectable (have a lower perceptual threshold). Additionally, one would expect that sensory regions with more cortical columns would be capable of better fine scaled discrimination. Columns could respond in more unique combinations and provide a finer representation of stimulus features. Importantly, if many cortical columns perform the same computation, no one stochastic event can affect the entire representation of the stimulus. The stochastic events average out not over time (like in a trial-average) but over space (many separate columns).

126

The idea that cortical columns are the central unit of cortical organization is not new. This hypothesis was first proposed by Vernon Mountcastle around the same time of Hubel and Wiesel's discovery of feature-selectivity in vision (Mountcastle, 1957, 1997), and now championed by other researchers in human and primate somatosensation (Tommerdahl et al., 2010). The work in this thesis adds to this existing framework.

Mountcastle showed extensive evidence for columns across all multiple areas of the neocortex, including both motor and sensory areas. This framework lacks, however, a coherent reason for the columnar structure. I propose that columns are necessary because of the inherent unreliability of biological hardware, but could also key features of the neural code. For example, pressure stimuli on skin have been shown to activate a wide region of cortical space, but have a prominent focal point or center of mass (Mountcastle, 1957). If this signal is modulated such that it is harder to propagate, the accumulation of errors model predicts that the center of mass (columns that respond most reliably to a given stimulus) would not be greatly affected. However, the efficiency of columns that respond with intermediate probabilities could be dramatically reduced, creating spatial specificity. The same dynamic could contribute to classic examples of neuroplasticity (Merzenich et al., 1983). Since evidence for plasticity rests primarily on modulating the relative response amplitude across different columns, sensory deprivation may actually affect all columns equally. Given different initial conditions, however, the result is a non-linear change in the average cortical activated area.

127

The data in this dissertation strongly suggest that the cortical column that represents a whisker acts only as detector and does not encode specific stimulus features. Given this, I hypothesize that any one column in the human and primate somatosensory cortex can only function as a simple detector and does not encode additional stimulus-feature information. Instead, feature-selectivity arises by controlling the reliability of the column detectors such that specificity can be created across cortical space.

7.4   Interfacing with the Brain

One important reason for understanding the neural code is to develop technologies that can interact and communicate with the brain such as sensory prosthetics. In this context, the probability of activation hypothesis (if found to be more generally true) would change the design specifications for technologies that interface with the brain.

For one, the information recorded at a single point in space (with a single electrode) would always be more reliable if the interface is more peripheral or lower on the neural hierarchy.   This has already been observed, in fact, as peripheral interfaces such as targeted reinnervation (Kuiken, 2006) are having more initial clinical success.

A common current approach is to use a single electrode (Kennedy et al., 2000), or multi-electrode arrays (Hochberg et al., 2006), either invasively (Kennedy et al., 2000; Hochberg et al., 2006) or non-invasively (Millán et al., 2004). Each of these technologies are designed to specifically record from single neurons or units and infer specific stimulus features.   Given the probability of activation hypothesis, I predict that the

128

information available from a single electrode recording would be unreliable, potentially even unhelpful for the decoding of features of sensory stimuli. An interface design can be simplified if a spatially averaged stimulus is sufficient for decoding stimulus features.

For the reverse problem, engineering solutions for writing the neural code may also be simplified by the probability of activation hypothesis. Multiple redundant pathways, each of which is sufficient, allow flexibility in how information is delivered. Electrical stimulation of the brain is more reliable and repeatable than activity from peripheral receptors (Millard et al., 2013); therefore, it is possible to have greater control the reliability of this signal. The probability of activation hypothesis suggests that although the control of the spatial extent or spatial specificity of the input would be essential, it would not be necessary to have the ability to activate individual neurons.

As an overall assessment, the probability of activation hypothesis allows for simpler design specifications for brain-machine interface. If we do not assume that single neurons encode specific features, then we do not need to record activity of neurons individually or stimulate specifically. Instead, it would be more important to be able to control average activity across cortical space.

7.5 Future Directions

Many additional experiments could continue to validate, expand, or potentially falsify, the probability of activation hypothesis and the accumulation of error model presented in this thesis.

129

First, it is the unique characteristic of all scientific hypotheses that they are directly falsifiable. If it can be shown that rodents can discriminate between deflection velocity on single trials and on single whiskers, I will consider this hypothesis to be falsified. Specifically, this requires that rodents report the intensity of the stimulus with unique behaviors (multiple-alternative forced choice).

Short of direct falsification, some aspects of the experiments presented here weaken the strength of the evidence. The experiments in this work were all done in an anesthetized preparation. Despite the correspondence between the anesthetized data and the awake behaving data, I must concede that an anesthetized brain is not an awake brain. Specifically, most experiments in this work were performed under sodium pentobarbital, thought to act on the inhibitory neurotransmitter, GABA (Steinbach and Akk, 2001). With new technologies, it is becoming more and more feasible to study the brain in awake animals. The challenge is that many functional measures require averaging over many trials in order to create sufficient single to noise. Even multiunit recordings in response to the same stimulus exhibit extensive variability across single trials.

Some technologies, such as local field potential, do record amplitude information on single trials. However, it is unclear how to interpret absolute amplitude when the distance from the electrode and the electrode resistance dramatically affect a single neuron's contribution to the signal. As a result, I consider local field potential a spatially weighted average of activity and therefore not a valid technique for testing the proability of activation hypothesis. The most straightforward approach to testing the probability of

130

activation hypothesis in awake animals is to develop a recording technology that is similar to voltage sensitive dye imaging. There is promise for chronic awake recordings using genetically encoded voltage sensitive fluorescent protein (Akemann et al., 2010). If this methodology allows sufficient signal to noise such that estimates of response amplitudes on single trials are meaningful, the probability of activation hypothesis could be tested in awake animals similarly to the data in this dissertation.

Within an anesthetized preparation, it is also necessary to explicitly test the spatiotemporal coding framework I have proposed in Chapter 5. While studies done in anesthetized animals may or may not map directly to the awake brain, one can still discern some information about the nature of the neural circuitry. The spatial decoding problem becomes increasing complex as one is forced to differentiate between more whiskers. However, using the matched-filter algorithm developed in Chapter 2, it should be possible to deflect an array of whiskers and observe probabilistic dynamics. Sub-threshold velocities will allow for better observation of the probabilistic dynamics.

Finally, it would be interesting to test the probability of activation hypothesis in the encoding of skin pressure stimuli in other animals. Once this paradigm is moved outside the whisker system, however, some technical hurdles develop. The problem is two-fold: 1) a projections from a single column can spread into slightly into neighboring columns; 2) it is not possible to activate only a single column at the periphery. Unlike the peripheral whisker input, skin representations are known to be highly overlapping (Mountcastle, 1957). Even a single pin prick would activate both skin surface receptors

and deep tissue columns. Additionally, there is some amount of spatial spread even in the head of a pin and so we expect multiple columns to respond. In order to test this hypothesis, we would need to solve at least one of two difficult technical problems: 1) how to stimulate only a single column (as we did in the whisker system); or 2) how to record from only a single cortical column.

 One approach could be to use targeted genetic tools to specifically record from a single column. For instance, one may be able to use the voltage sensitive fluorescent protein technology to target a single column (or label columns in different colors). Even if the columns could not be differentiated at the periphery, the response dynamics of an individual column could be recorded uniquely.

7.6   Concluding Remarks

Human brains are as different from each other as the human personalities and experiences that they encode. As neuroscientists, we should embrace the idea that we do not need to be wired identically or perceive stimuli identically in order to co-exist. If we were to take the average of every definable characteristic to create a representative human, we would fail to model every individual human. No one person, no one experience is adequately modeled by the average. The neural code is likely not an exception to this general rule.

132

APPENDIX A   Covariance of Amplitude Confounds Metrics of

Spatial Activation

A.1   Introduction

Spatially organized cortical representations of stimulus space (i.e. topographic maps) or stimulus parameters (i.e. orientation tuning maps) are ubiquitous features across our sensory systems.   Anatomical evidence for topographic maps can be seen with histological stains (Land and Simons, 1985). Neurons that encode similar stimulus information tend to be close to each other in cortical space (Merzenich et al., 1983). The synaptic connections between these neurons reorganize when presented with patterned or coincident sensory inputs, particularly early in development (Wiesel and Hubel, 1963; Simons and Land, 1987).   These maps are also studied as functional entities.   The functional representation of topographic maps can be context-dependently modulated, by changes to brain state (Ferezou et al., 2006), the presence of an adapting stimulus (Ollerenshaw et al., 2014) or even stimulus-specific parameters (Sheth et al., 1998). Counter-intuitively, the methodology used to assess functional modification of topographic maps and the methodology used to assess structural plasticity of topographic maps is often the same (i.e. multi-electrode recordings, optical imaging etc) In this study, we consider some inferential problems in the interpretation of spatial data. We suggest that data collected as evidence for cortical plasticity and data collected as evidence for brain-state modulation of topographic maps may be more similar that previously thought.

133

Functional measures of cortical maps rely on the ability to distinguish between active and inactive cortical regions. The distinction between a cortical region that is active and one that is inactive (or comparably inactive) is question-dependent. A range of questions might arise. Is activation in a region different from spontaneous activity (Polley et al., 1999a)? Is the activation in one cortical region smaller than that observed prior to experimental manipulation (Wiesel and Hubel, 1963; Polley et al., 2004)? Is the amount of activation in the surround significantly less than the peak response with different stimulus conditions (Sheth et al., 1998; Ollerenshaw et al., 2014)? These questions are each distinctly different from each other, yet methodologically similar in a key respect: in each case, a direct comparison of response amplitudes, measured functionally, is used to infer that a change in the cortical area of activation has occurred. In fact, inference of changes of functional spatial activation requires corresponding changes in response amplitude. The two metrics are intrinsically linked.

The interdependence of response amplitude and area of activation creates a challenge if we seek to measure changes in area of activation separately from changes corresponding to response amplitude. Changes in the cortical area both functionally and anatomically occur almost exclusively with a covariance of the spatial spread of the signal. For example, using the examples above, absolute response amplitude decreases with sensory deprivation (Glazewski and Fox, 1996; Fox, 2002), increases with increased frequency of stimulation (Sheth et al., 1998) and can increase or decrease with different brain states (Ferezou et al., 2006).

134

Here, we test metrics used to quantify the topographic representations of whiskers in the barrel cortex of anesthetized rats. Using this knowledge, we develop a framework for the quantification of changes in functional maps across experimental conditions. We consider the sensitivity of metrics of spatial activation to the covariance response amplitude. We show that identical data analyzed with respect to different points of reference can report opposing spatial trends. We use voltage sensitive dye (VSD) imaging of neural responses to whisker deflections to demonstrate how the use of these metrics can confound interpretations of area of activation. Using a combination of absolute and relative metrics, we can differentiate between amplitude-independent and amplitude-dependent changes in spatial activation. Next, we consider the robustness of relative metrics in the presence of noise. We show that a relative threshold determined from a biased estimator of the point of reference is also sensitive to the covariance of amplitude. We conclude that covariance of response amplitude is systematically confounding spatial data across multiple fields of sensory neuroscience.

A.2   Methods

A.2.1   Threshold Analysis on Simulated Images

All simulations were performed in MATLAB (MathWorks, Natick, MA). Model images were two-dimensional radially symmetric Gaussians centered in a 50 x 50 pixel image. For Figure A.3, the amplitudes used for representation were 1, 2, and 3 for Image 1, Image 2 and Image 3, respectively. Each Gaussian had a standard deviation of $\sigma=80$ pixels. The images were analyzed with absolute thresholds of 0.4 (noise-derived and

135

peak-derived) and a relative threshold of 50%. For Figure A.4, we created five series with three images each. Image 1 was identical in all image series, with amplitude of one and standard deviation of $\sigma=80$. The standard deviation of the Gaussian in each image series covaried with the amplitude, as shown in Figure A.4A. Amplitudes were 1.1 for Image 2 and 1.2 in Image 3. As threshold severity increased, more and more pixels in an image are considered inactive. A small absolute peak-derived threshold was therefore considered a 'high severity' threshold and a small absolute noise derived threshold was considered 'low severity'. The threshold severity spanned the entire range of the image with the smallest amplitude. For the noise simulations, noise was added to each pixel by randomly choosing a value from a normal distribution with mean of zero and variance of 0.01.

A.2.2 Animals

All procedures were approved by the Animal Care and Use Committee at the Georgia Institute of Technology. Data from six female adult Sprague-Dawley rats (250-350g) were used in this study.

A.2.3 Surgical Preparation

Experimental procedures were similar to those utilized in our previous studies (Ollerenshaw et al., 2012, 2014; Wang et al., 2012; Millard et al., 2013). Briefly, animals were initially sedated with 5% isoflurane such that sodium pentobarbital (IP, 50mg/kg) could be injected for long-term anesthesia. A catheter was inserted into the tail vein and used to deliver additional sodium pentobarbital (I.V.) as needed during the duration of the

136

experiment. The anesthetic depth was kept such that the animals exhibited no pain reflex (toe pinch, or eye blink) but was light enough that neural activity could be observed. Heart rate, oxygen saturation, respiratory rate, and temperature (37°C) were monitored continuously to ensure a constant level of anesthetic depth. Animals were stabilized in a stereotaxic frame and a craniotomy and duratomy were performed over the barrel cortex in the left hemisphere (0.5-4 mm caudal to bregma, 3-7 mm lateral from the midline). A well was created using dental cement surrounding the craniotomy to hold the dye solution and keep sterile saline on the cortical surface throughout the experiment.

A.2.4 Voltage Sensitive Dye Imaging

Voltage-sensitive dye imaging measures neural activity, primarily sub-threshold changes in membrane voltage, from supra-granular layers of the cortex (Petersen et al., 2003a). A schematic of the imaging set-up is shown in Figure A.1A. The cortex was stained with the voltage sensitive dye (RH1691, 2mg/ml, Optical Imaging) from 1.5-2hours, until the cortex was visibly blue. Unbound dye was then rinsed off by multiple washes with sterile saline. The cortex was illuminated with a 150W Halogen lamp passed through an excitation filter (621-643 nm). Images were recorded with a high-speed CCD camera (MiCOM2, SciMedia). A 1x objective lens was combined with a 0.63X condenser lens, resulting in a total magnification of 1.6X. Twenty to sixty trials for each stimulus condition were recorded for analysis. The resultant image size was 183 x 124 pixels, at 20μm/pixel.

137

A.2.5 Image Analysis

All analysis was performed in custom written software in MATLAB (MathWorks, Natick, MA). Florescence was quantified as the percent change from background, % $\Delta F/F_0$ where background was defined as the average pre-stimulus activity (200ms, 40 frames) at the beginning of each trial. Unless otherwise noted, analyses were performed on trial-averaged images with a minimum of 20 trials. The barrel location of the deflected whisker was calculated as the center of mass of the fluorescence in the onset frame (chosen manually, typically 20ms after stimulus delivery). Time series and amplitude measurements were defined as the average fluorescence in a circular region of interest (10 pixel radius, ~200μm) centered over the center of mass. Images were filtered with a 400μm x400μm spatial averaging filter for display purposes only.

A.2.6 Area of Activation Analysis on Voltage Sensitive Dye Images

Spatial analysis was performed on unfiltered trial-averaged images (from 20 to 60 trials). Analysis was performed for the highest velocity for the time-dimension analysis. A time-averaged image (2 frames, 15-20ms) was used for velocity-dependent spatial analysis. The noise threshold was set at two standard deviations above the average noise fluorescence (a pre-stimulus VSD frame). The peak was the single maximum pixel in the image unless otherwise specified. The relative threshold was defined at half (50%) the peak metric minus the noise metric. For the pseudo-experiment of Figure A.9, a set of six single trials from each data set was chosen to be part of the 'high' group and the 'low' group. The single-trial response amplitudes were sorted and the six trials with the highest

138

amplitude were used in the 'high' group and the next six trials (6th -11th) were used as the 'low' group. The results held regardless of how 'high' and 'low' were defined, as long as there was an absolute change in amplitude between the groups.

A.3   Results

The results presented here pertain to spatial measures of neural activation, and are thus relevant for different modes of imaging (optical, fMRI, etc), electrode array recordings, or other measurement modalities that are designed to capture activation of neural tissue across space. For simplicity, we focus on optical imaging, which has been widely utilized for these purposes. In imaging, the spatial representation takes the form of a pixelated image, with the level of neural activation at a particular location in the tissue represented by the intensity in the image at the corresponding pixel. In Figure A.1, we have shown an example of an optical image that we collected in response to a whisker deflection using voltage sensitive dye (VSD) imaging. The response amplitude in voltage sensitive dye imaging is a measure of the relative fluorescence that linearly reflects the membrane potential of cells in the given cortical space represented by a pixel (Petersen et al., 2003a). A schematic of the optics of a voltage-sensitive dye imaging system is shown in Figure A.1A and a representative image recorded in response to a whisker deflection in Figure A.1B. Figure A.1C shows the pixel intensities (or response amplitudes) measured across the profile line as drawn in Figure A.1B. The image profile was approximately Gaussian. This is generally true in optical imaging experiments, as the point-spread function of light scattered through tissue has been shown to be well-approximated by a Gaussian function (Stallinga and Rieger, 2010). Although there is a clear qualitative

139

region of the image that has greater fluorescence representing stimulus-evoked activity, the individual pixel intensities taper off gradually and continuously. To quantify or compare the regions across of multiple images, boundaries must first be defined to identify *active* versus *inactive* pixels. While this is a common experimental problem, the solution is not unique. Multiple different widely used metrics have been developed and used in literature to quantify area of activation.
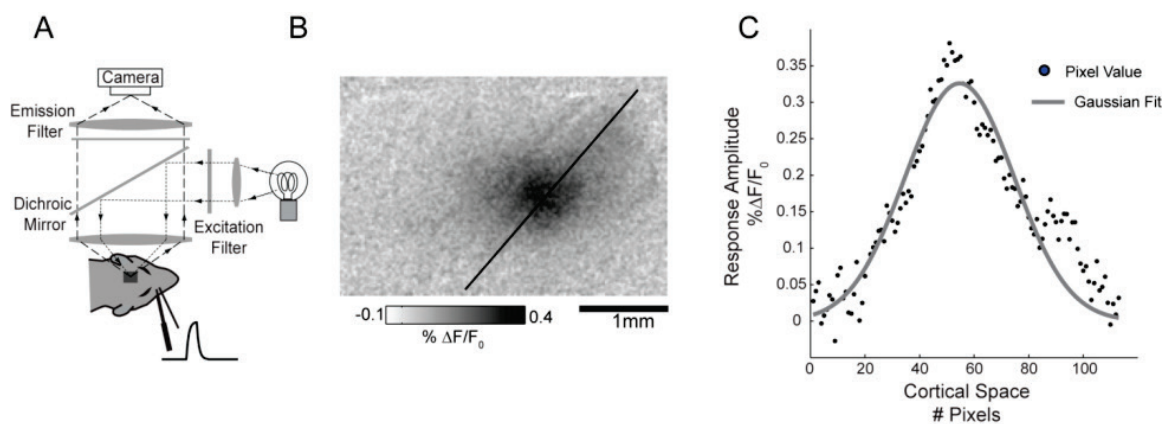
Figure A.1 Quantifying Area requires comparing response amplitudes A) A schematic of the VSD optical imaging set-up. Whiskers are deflected using a piezo-electric actuator. B) Representative VSD image collected following a whisker deflection. Pixels that lie along the black line are shown as dots in (C). A Gaussian fit to this image profile is shown.

A.2.7 Absolute versus Relative Metrics

We classified metrics of spatial activation by the point of reference. The point of reference is the metric used to register or normalize across images or experimental conditions. Most commonly, studies measure area of activation with respect to an

observed minimum or noise level. Some common examples of a noise-derived metric are 'statistically significant activation relative to pre-stimulus noise' or 'two standard deviations above noise'. Regions of the cortex that have response amplitudes higher than this point of reference are considered active, while regions that do not meet this level of activation are considered inactive. A schematic of a noise-derived metric is shown in Figure A.2B.

Figure A.2 Area Metrics Classified By Point of Reference A) Area is defined by the number of pixels above a threshold. Thresholds can be placed relative to different points of reference in the image. B) A noise-derived metric is determined relative the minimum observed value; C) A peak-derived metric is relative to the observed maximum. D) A relative threshold defined relative to the absolute size of the image (peak and noise metrics).

It is also possible to measure activation relative to the peak response (for example, series of concentric rings at specific distance from the peak). These metrics capture the sharpness of a cortical response by quantifying how fast the signal drops off from its

141

peak. This is a metric defined relative to the maximum observed response, or a peak-derived metric depicted in Figure A.2C.

Both of the examples of metrics described above are *absolute* metrics of spatial activation. Such metrics set the difference between active and inactive at an absolute and unchanging difference in response amplitudes between activity observed in some region and activity observed at the point of reference in every experimental condition. In some cases, this means specifically choosing a difference that is considered meaningful (i.e. a 2% change in signal amplitude) or by using statistical significance to define the absolute change in amplitude.

It is also possible to define a *relative* metric, where the distinction between active and inactive changes to reflect the size of the observed response in each experimental condition. For example, a 2% change would be required in an image with larger response amplitudes, but a 1% difference may be sufficient in images with smaller responses. A common example of a relative metric is a 50% contour, or area at half max, which pertains to raw image contours, or contours derived from Gaussian fits. In Figure A.2D we depict a type of relative threshold that is determined with respect to both the peak and noise measurements.

142
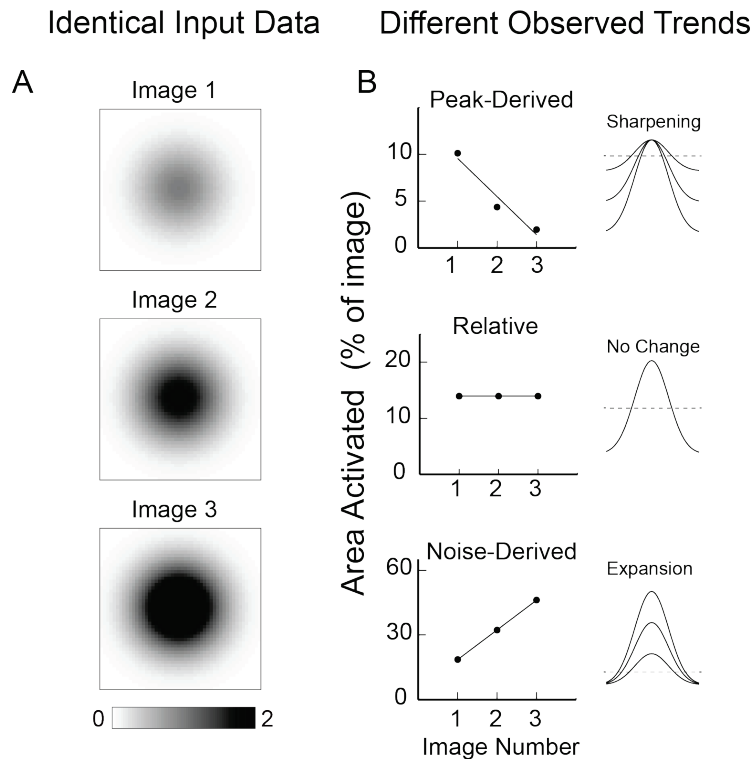
Figure A.3 Same Data Analyzed Three Ways   A) Three 2D Gaussian images that were scaled versions of each other are analyzed using one of each type of metric. B) Peak-Derived metric showed sharpening, Relative metric showed no change and the noise-derived metric showed an increase in area.

A.2.8   Same Data, Analyzed Three Ways

All of these methods seek to quantify the extent of cortical activation, but they are not equivalent. To demonstrate this, we analyzed the same data with a representative method from each category: noise-derived, peak-derived and a relative metric. Note that there are multiple ways to implement each class of metric. We used the simplest implementation, a threshold. As model images, we created radially symmetric two-dimensional Gaussians (see Methods).   By using a Gaussian, we can independently control the response

143

amplitude and the spatial spread defined by the standard deviation σ. We created three images in which we scaled the amplitude, while holding σ constant. Image 1 has the smallest amplitude and image 3 has the largest (Figure A.3A). Each image is exactly a scaled version of the other two. We then analyzed the images using a noise-derived threshold, a peak-derived threshold, and a relative threshold. The resulting trends are shown in Figure A.3B. Despite analyzing the same three images each time, we could infer three markedly different trends in area of activation between the three images (or experimental conditions). In the case of a peak-derived metric, images with higher absolute magnitude show faster drop-off from the peak, indicating a sharpening of the cortical response (Figure A.3B, top). In contrast, the same images analyzed with a noise-derived metric showed cortical expansion with increased amplitude, with a greater number of pixels surpassing the absolute difference in response amplitude necessary for activation (bottom). The relative metric, however, shows no change in area of activation (middle).

It is interesting, potentially troubling, that the three analyses of the same data result in three potentially contradicting interpretations. It is unclear if we should conclude that these images show cortical expansion, sharpening, or activation of the same area. This confound is caused by covariance of amplitude across the conditions or images.
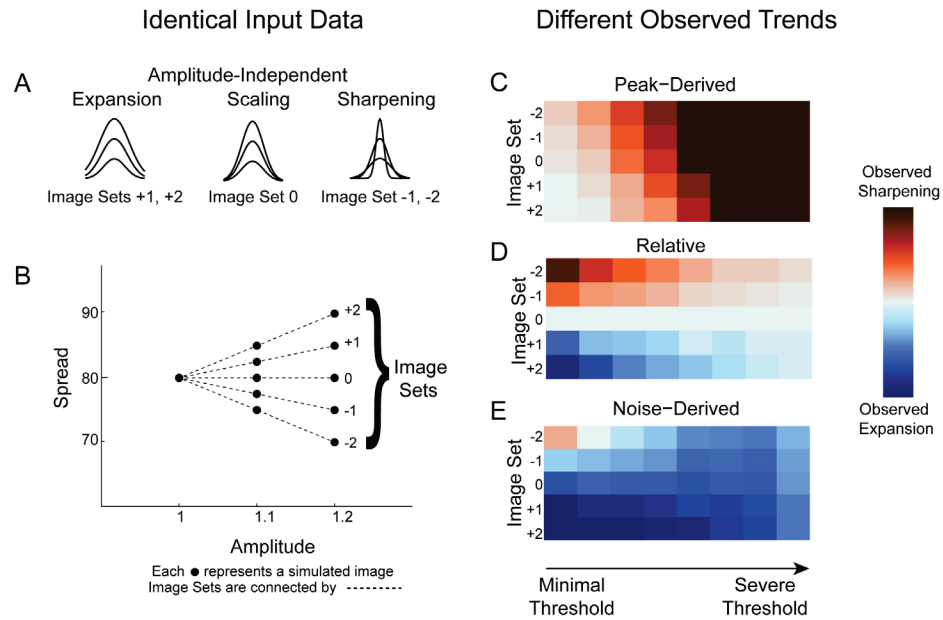
Figure A.4 Point of reference matters more than threshold severity   A) Fives sets of three 2D Gaussian images were created with the specific amplitude and spatial spread (sigma) metrics defined by the points in the graph (B). Image sets +1 and +2, sigma increases with amplitude (amplitude independent expansion), the reverse is true for image sets -1 and -2. C) Each of the five image sets analyzed with 8 peak-derived thresholds (C); 8 relative thresholds (D) and 8 noise-derived thresholds (E). The point of reference determined the direction of the trend, not the absolute threshold severity.

A.2.9   Point of Reference Determines Observed Trend: Same Data, Analyzed 24 Ways

It is common practice to assess area of activation at multiple threshold levels (Polley et al., 2004; Drew and Feldman, 2009) and widely acknowledged that analysis at different threshold levels can influence the magnitude of observed trends. We considered whether the contradictory results shown in Figure A.3 were sensitive to the absolute threshold value.

For this analysis, we created a more complex model data set of 2D Gaussian images and analyzed them 24 different ways (8 variable thresholds levels in reference to each of the 3 possible points of reference). The model data set included five image sets with three images in each. In Figure A.4A, we plot response amplitude versus sigma for our simulated data set. Each black dot represents a 2D Gaussian image with the corresponding amplitude and sigma values as labeled on the axis. The amplitudes of the three images in each group are not changed, only the sigma values. The total change in amplitude from Image 1 to Image 3 is 20%. Each image set (the three images connected by dotted lines) represented a different relationship between amplitude and spatial spread. Image set 0 is analogous to the data set analyzed previously in Figure 3, each image has the same sigma value but every pixel is scaled by the amplitude value. Image sets +1 and +2 represent expansion of the sigma independent of amplitude. Image sets -1 and -2 represent sharpening independent of amplitude (see schematics in Figure A.4A).

Figure A.4C-E show these five data sets analyzed across eight different threshold severities. Threshold severity conceptually represents the amount of the image that is considered active. If most pixels in an image are above the threshold then the threshold is minimal. If most pixels are below the threshold, it is considered severe. For example in a noise-derived metric, 0.5 standard deviations above noise is a minimal threshold. Five standard deviations above noise is a more severe threshold. There is no noise in these images, so the thresholds increase by equal absolute values from minimal to severe.

146

Figure A.5 Velocity and Temporal Images from Voltage Sensitive Dye Imaging   A) Representative VSD images from four velocities (V1, slowest to V4 fastest) at four frames (points in time). B) Time series created from each of the image series shown in A by averaging the fluorescence over the entire image.

While threshold severity did influence the magnitude of the observed trend, we observed that the point of reference, not the threshold severity, determined the direction of the trend (sharpening versus expansion). The relative method (Figure A.4D) in this case can be interpreted to report the amplitude-independent trend in the standard deviations, σ, of the images. The relative method reports that in image sets -1 and -2 sharpening is occurring (orange) and in image set +1 and +2 expansion is occurring (blue). As shown in Figure A.4C, we observed that a peak-derived metric reported spatial sharpening (orange) in nearly every case. In this way, the peak-derived metric was biased, as it reported sharpening with an increase in amplitude even if the sigma was increasing with image number. Conversely, we observed that a noise-derived metric (Figure A.4E) reported an expansion (blue) in nearly every combination of image set and threshold level.   In this way, the noise-derived threshold was also biased, as it reported an

147

expansion in the activated area with image number even if the standard deviations of the images were decreasing. We only observed trends against these biases if the threshold-severity was minimal (bottom left, Figure A.4C; top left, Figure A.4E).

A.2.10 Quantification of area of activation in response to whisker deflections using voltage sensitive dye imaging

Our simulations demonstrate that response amplitude confounds metrics of spatial activation. Here, we demonstrate this problem with an experimental example. We recorded the layer 2/3 response to whisker deflections using voltage-sensitive dye images of increasing velocity (V1-V4, 150°/sec, 300°/sec, 600°/sec and 1200°/sec respectively). Example trial-average image series in response to whisker deflections are shown in Figure A.5A. Time series were created by averaging the pixels in the image (Figure A.5B). Qualitatively, the area of activation appeared to increase both in the time dimension (across the rows from 15-25ms) and with increasing stimulus strength (down a column, V1 to V4).
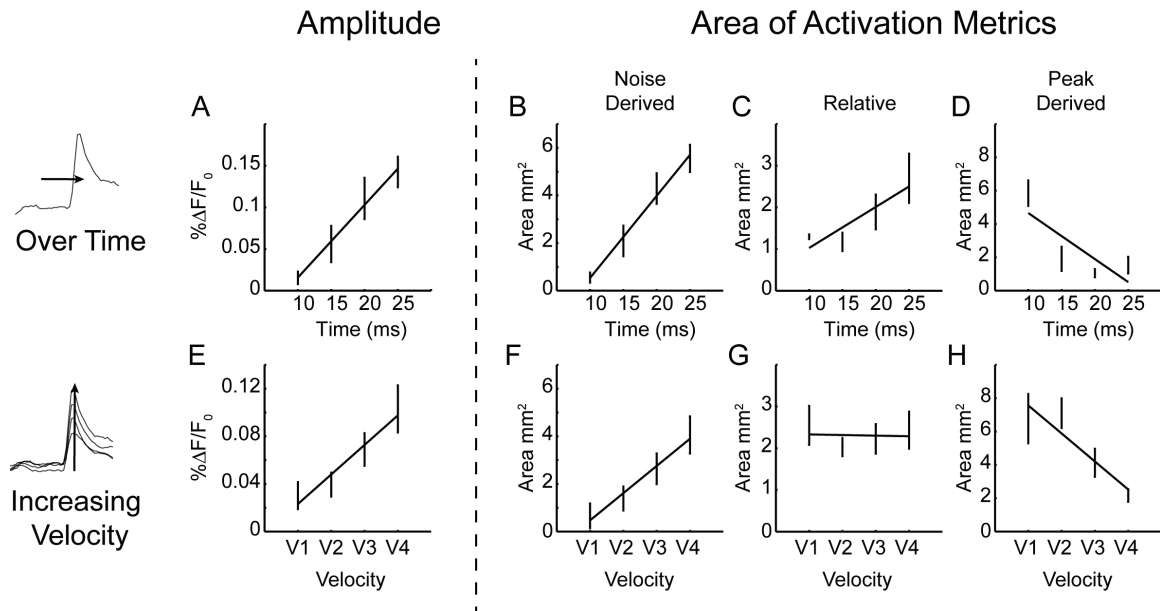
Figure A.6 Different area trends observed with different metrics A) Top row shows analysis with increasing frame number (in time). Amplitude increases with time from onset to peak. B) Noise-derived threshold shows an increase in area with time C) Relative threshold shows a less severe but clear increase in area with time D) Peak-derived threshold shows area decreasing with time. E) Similar to top row, but for one point in time for increasing velocity stimuli. F) Noise derived threshold shows an increase in area. G) Relative threshold shows no consistent change. H) Peak-derived threshold shows spatial sharpening.

We quantified amplitude and spatial spread (n=7 whiskers from 6 animals) in both dimensions: time and strength. First, we confirmed that the amplitude increased from 10ms to 25ms after stimulus delivery (Figure A.6A). Similarly, at a single time point (average of two frames, 15-20ms after deflection) the VSD amplitude also increases with whisker deflection velocity (Figure A.6E). Therefore, these VSD images represented an experimental example that paralleled our earlier simulations: we sought to quantify area

149

of activation in the face of a covariance of response amplitude across experimental conditions.

We quantified area of activation using a noise-derived threshold (two standard deviations above pre-stimulus noise) and a relative threshold (controlled for both noise and peak measurements) and a peak-derived threshold (5 noise standard deviations difference from the peak). In the time dimension, the inferred area of activation increased using both noise-derived threshold and the relative threshold (Figure A.6B,C). The peak-derived absolute metric showed a decrease in area (sharpening) as was predicted from our simulations (Figure A.6D). The peak-derived threshold does not have an obvious scientific interpretation in this experiment, so we considered the results of only the noise-derived and relative threshold. We conclude that the temporal increase in area of activation was *amplitude-independent*.

This was not true in the stimulus strength dimension. With a noise-derived threshold, the area appeared to increase with stimulus strength. However, with a relative threshold there was no significant correlation between whisker deflection velocity (strength) and area of activity. As expected, the peak-derived metric still showed spatial sharpening. The contradictory results of these three methods suggested, based on the results of our simulations, that VSD images from whisker deflections of strong velocities were simply scaled versions of deflections of weak velocities. The apparent increase in area was the result of the same physical space being more or less activated. In other words, the

150

observed change in area of activation in the noise-derived threshold was entirely *amplitude-dependent*.

Interestingly, if only a noise-derived threshold were employed to analyze these two cases, we would have concluded that both stimulus strength and time resulted in increased spatial spread. However, these two situations are distinctly different; using multiple metrics, we can definitively say that these two cases do not share a common mechanism.

A.2.11 A biased estimator of a point of reference confounds even relative metrics

Our simulations up until this point suggested that a relative threshold could differentiate between an *amplitude-independent* change in area of activation and an *amplitude-dependent* change. There are multiple ways to define a relative metric. In our previous example, we defined a relative threshold using two points of reference, the noise and the peak measurements.

Another common relative metric is the *area at half max* with the peak pixel used as the peak point of reference. A relative metric defined with respect to the peak pixel is in fact, also sensitive to the covariance of amplitude. To demonstrate this effect, we used the same set of images from our original simulation in Figure 3, shown again in Figure A.7A. We calculated activated area by the area at half-max metric both in the absence and in the presence of noise. When there is no noise, the result is identical to what was observed before (Figure A.7C in blue). There was no amplitude-independent change in area between the three images. However, if the noise was added to the images (as shown in Figure A.7B), prior to calculating the relative threshold level, this was no longer true. The

151

black lines in Figure A.7C show 10 separate instances of noise. In each case, the area at half-max method underestimates the true area (as determined by the noise-free image). However, this underestimation is more dramatic on average in Image 1 than in Image 3. Even though the noise added to each of the three images was identical, the effect the noise caused on the observed area was not.
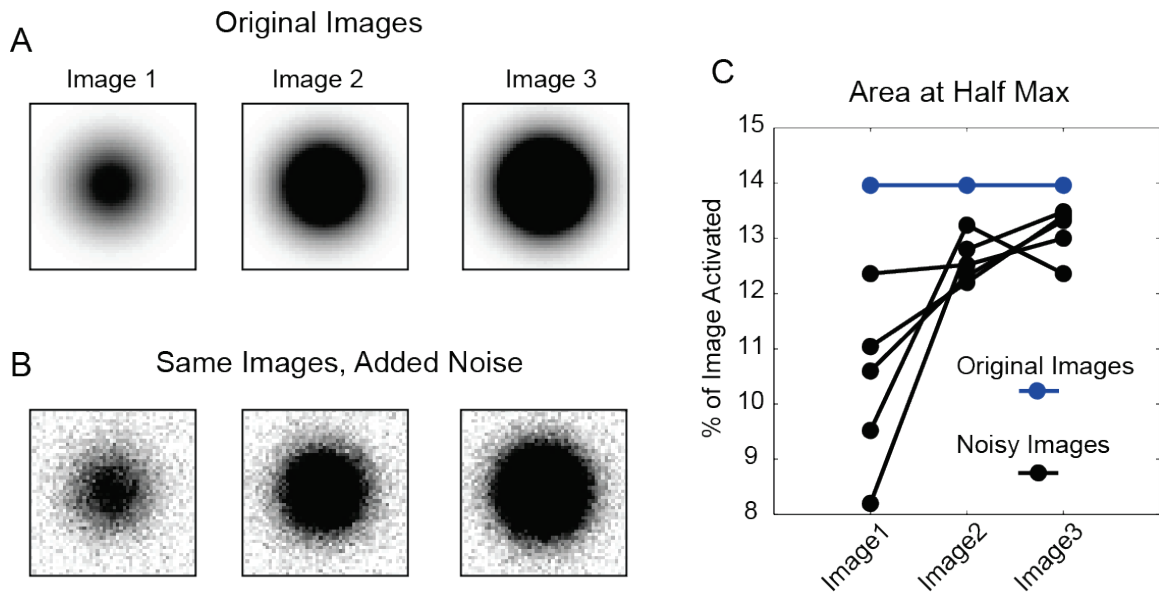
Figure A.7 Noise Biases Area at Half-Max Relative Metric Same images as in Figure A.3, without noise (A) and with noise (B). C) Each image analyzed by choosing the peak pixel and then setting a threshold at half the observed peak value. Images without noise show no change in area between images. Area of noisy images were systematically underestimated, with a larger effect on images with lower initial amplitude (image 1).

Noise in the images resulted in a systematic underestimation of the true area of activation because the peak pixel is a biased estimator of the true peak value. Choosing one max pixel from many possible pixels results in a selection bias for pixels that have

152

experienced a positive noise contribution. We demonstrated this effect using Monte Carlo simulations (Figure A.8). When the estimator of the measurement is systematically high, the threshold based off this estimator will be placed higher than intended, resulting in fewer pixels counted as active. This effect is grows as the signal to noise ratio decreases (Figure A.8).

We conclude that a relative metric that uses only the peak pixel as a point of reference is biased, and more likely to report that images with smaller amplitudes have proportionally smaller area of activation.
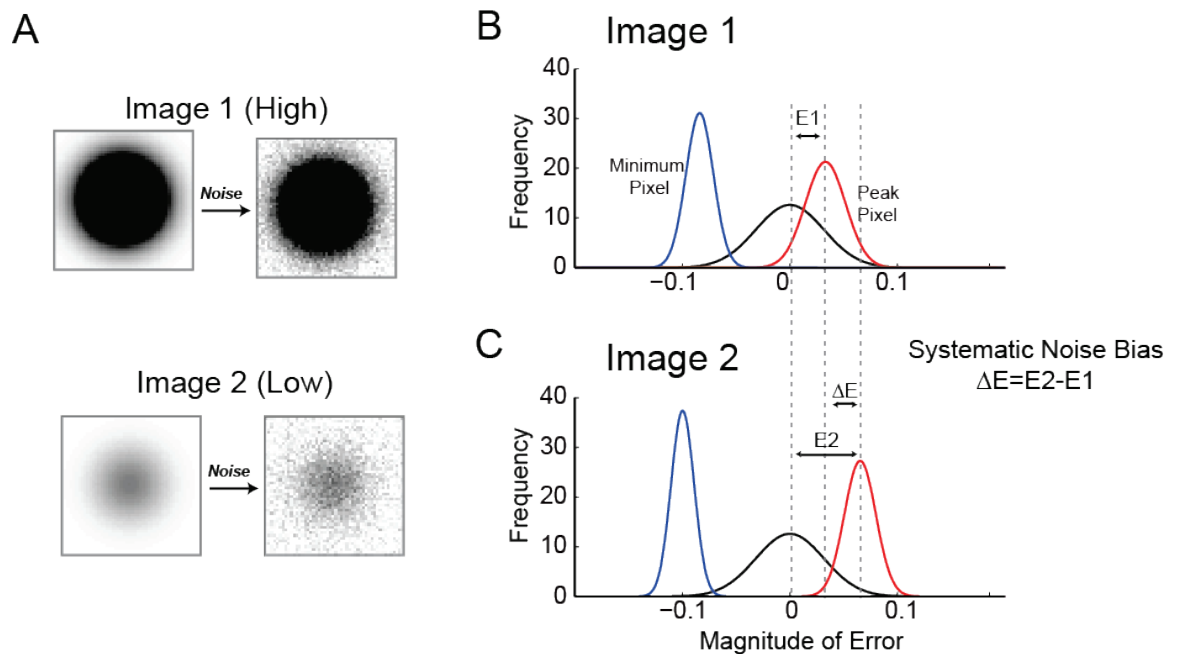
Figure A.8 Systematic Bias Created By Selecting the Peak Pixel A) Noise added to a standard 2D Gaussian with a high amplitude and low amplitude B) The error distribution associated with the peak pixel (red) is systematically more positive than the noise distribution added to the image as a whole (black distribution). If the minimum pixel were chosen instead, this bias would be negative (blue distribution). C) The expected contribution of noise when the original amplitude of the image is low is larger (E2) than

153

when the original image is bigger (E1). Thus the expected error also depends on the experimentally unknown pixel itself.

A.2.12 Demonstration of Shortcomings of the Relative Threshold

We created a pseudo-experiment in which we exploited the single trial variability of response amplitudes of voltage sensitive dye (VSD) imaging to create a situation analogous to the set of noise simulations presented above. A plot of the response amplitudes for 60 single trials of identical whisker deflections (the same velocity) within the same animal preparation is shown in Figure A.9A. The single trial response amplitude was variable, but it did not systematically vary across the experiment protocol. High amplitude responses occurred at trials early, middle and late in the protocol. This was true for all data sets used. For each dataset, we selected a subset of trials such that one group had a higher response amplitude ('high', red dots) on average than another group ('low', green dots) as shown in Figure A.9A (n=7 whiskers, 6 animals). Representative trial-averaged images from a 'high' group and a 'low' group are shown on an absolute scale in Figure A.9B. The enforced differences in amplitude are shown across animals in Figure A.9C. Qualitatively, in the absolute images, the 'high' group appears to have a greater area of activation than the 'low' group. However, we know that these images in fact represent responses from identical stimuli collected, interleaved in time.

154
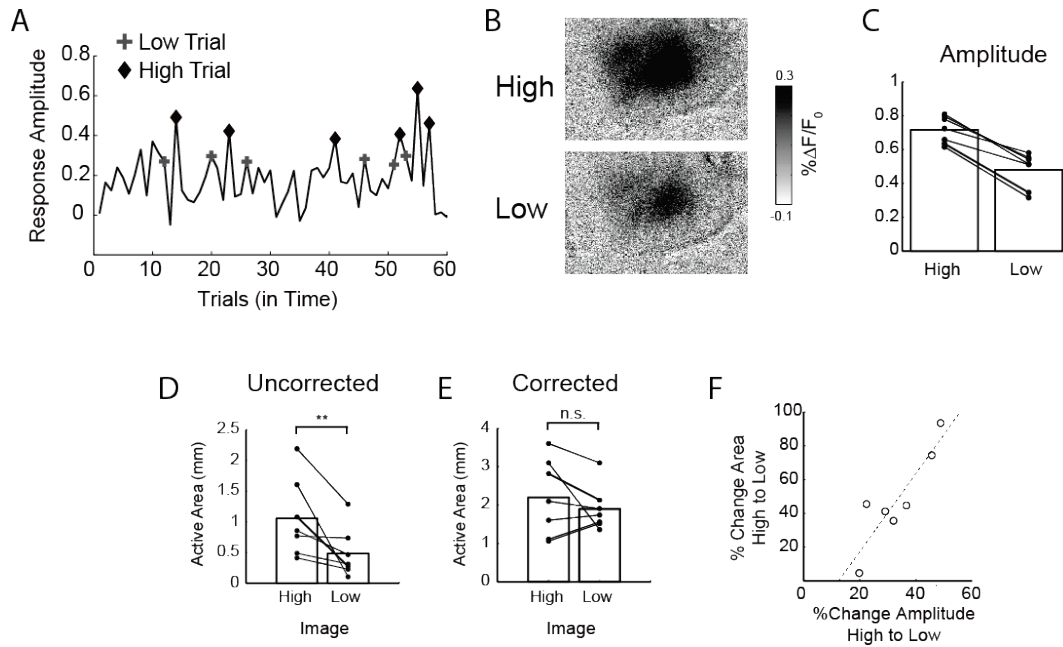
Figure A.9 Pseudo-experiment demonstrates area-confounds due to noise A) Single Trial response amplitude over time (trials) in a single data set. Trials included in the high group are denoted with a diamond, the low trials with a plus sign. B) Representative high and low images C) The pseudo-experiment required enforcing an amplitude covariance in all data sets, plotted here. Lines are individual data sets, the bar graph denotes the average. D) An uncorrected, peak-pixel derived relative metric shows a statistically significant decrease in area E) A corrected relative metric shows no statistically significant difference. F) This observed area differences (y-axis) are highly correlated with the enforced changed in amplitude (x axis) suggesting that the change in amplitude is sufficient to explain the change in area.

We then quantified the area of activation in these images using both an uncorrected and noise-corrected relative threshold. The uncorrected-relative threshold reported a 53% reduction in area and was statistically significant (p=0.028; paired t-test, n=7 whiskers from 6 animals). The measured area of activation in the 'high' trials was 1.06 mm ±0.24 (mean ± s.e.m.). The 'low' trials area of activation was 0.49 mm ±0.15 (mean ± s.e.m.).

155

This difference is larger than the average difference in peak amplitude (33%) we enforced in the artificial groups.

We then analyzed the area using a relative threshold determined with two points of reference (same as in Figure A.6). We corrected for the use of the biased peak pixel as an estimator by subtracting an expected noise contribution from the observed peak prior to setting a threshold. Using this corrected metric, the difference between the high and low groups was only 13%, and not statistically significant ($p=0.35$, paired t-test). The means are shown as bar graphs and the paired data points as lines in Figure A.9F.

No experimental variable can explain the statistically significant differences between the high and low groups in our pseudo-experiment; they were, to the best of our ability, representative of identical experimental conditions. It is possible that the observed variability in response amplitude reflected true differences in neural activity. Lower response amplitude trials may have a few less neurons firing and therefore decrease the area of activation. However, two pieces of evidence suggest that this is not a strong conclusion. First, use of the corrected 50% contour eliminated the statistical significance. Second, we plotted the change in amplitude from the high to low conditions against the observed change in area using the uncorrected area at half-max (Figure A.9E). There is a strong linear relationship, suggesting that the change in response amplitude is able to explain a significant portion of the observed area differences.

156

A.4   Discussion

We have shown that common metrics of spatial activation are confounded by changes in amplitude. Absolute metrics cannot differentiate between changes in area that are caused an increase in amplitude and those that are not. A relative metric must be used to infer amplitude independent changes in activated area. However, if a biased estimator such as the peak pixel is used to calculate the relative threshold, then even a relative metric is confounded by co-variance of amplitude between experimental conditions.

It is scientifically important to distinguish between amplitude-dependent and amplitude-independent changes in spatial activation. For one, distinguishing between amplitude-dependent and amplitude-independent changes allows insight into the mechanism of an effect. Cortical circuits can easily be activated differently such that we observe a change in the response amplitude, for instance with different strength stimuli (Shoykhet et al., 2000b; Lee and Simons, 2004; Boloori et al., 2010; Wang et al., 2010). In the rodent barrel cortex, it is known (and replicated in this study) that deflecting whiskers with greater velocity increases the observed response amplitude (Shoykhet et al., 2000b; Lee and Simons, 2004; Boloori et al., 2010; Wang et al., 2010) and this can be experimentally modulated by a number of means (Brumberg et al., 1996; Boloori and Stanley, 2006; Wang et al., 2010). Sensory deprivation by removal of whiskers (Glazewski and Fox, 1996; Fox, 2002) also results in a change of response amplitude that could be a general homeostatic response that occurs at any point in the pathway (suppressing all areas equally) or may be evidence of actual reorganization of cortical synapses. In order to provide evidence that cortical reorganization has occurred beyond a change in response

157

amplitude, one must show that the changes in spatial spread are different from what would be expected due to a change in response amplitude – to distinguish between an amplitude-dependent and amplitude-independent change in activated area.

To that point, the decrease in area we observed in low amplitude trials and high amplitude trials in our pseudo-experiment showed a change in spatial area comparable in both direction of effect and effect size to those observed in seminal barrel cortex plasticity studies (Polley et al., 1999a, 2004; Drew and Feldman, 2009). In these studies, the implementation of the area at half max relative to the peak pixel was interpreted as evidence that the effect was different from what would be expected if there were simply a decrease in response amplitude. It is important to reconsider whether these previous results are necessarily evidence of reorganization if they are also, given the bias described in this study, consistent with a change in response amplitude.

The analysis in this study has focused primarily on functional imaging data. However, the conceptual conclusions are not unique to imaging. Electrophysiological evidence also requires measured changes in response amplitude to infer changes in area of activation, and is much more likely to rely on the absolute metrics defined by statistical significance. An analogous design might be to measure response amplitude (average firing rate) at multiple points in space (one electrode is equivalent to a pixel). If the observed response amplitude changes across all electrodes from one experimental condition to another, then an application of an absolute threshold (statistical significance) will appear as if there were fewer pixels/electrodes activated. The failure to measure activation is not evidence

158

of its absence. We have demonstrated in our pseudo-experiment that is possible for the same stimulus, in the same experimental preparation to activate the cortical circuitry with different response amplitudes. We should be cautious with inferences that changes in response amplitudes are evidence of cortical reorganization.

The use of different types of spatial metrics has been discussed previously (Chen-Bee et al., 2000). Although the authors advocated for a different conclusion, the observation was the same: relative metrics show less dramatic changes in spatial activation than absolute metrics. In fact, in this previous consideration of analytical methods, the authors showed that representations of whiskers that had undergone classic plasticity paradigms did not appear different if quantified with relative thresholds. It is possible to identify additional pairs of studies with similar experimental protocols but with opposing conclusions consistent with the type of metric used For example, Polley et al., 1999b, used a noise-derived metric. Sheth et al, 1998 used a peak-derived metric. This confound is wide spread in the topographic map literature.

In summary, we have presented simulations and experimental examples of the ways covariance of amplitude confounds and biases metrics of area of activation. We showed that using both absolute and relative metrics, we could differentiate between amplitude-independent and amplitude-dependent effects on spatial spread. However, systematic contributions of noise can dramatically confound even relative metrics if a biased estimator is used. Finally, we showed that individual trials from the same animal and identical stimuli, collected interleaved in time, vary enough in response amplitude to

159

replicate changes in spatial spread previously interpreted to be evidence of reorganization of the barrel cortex topographic map. Taken together, we propose that changes in response amplitude are confounding interpretation of spatial plasticity data.

160

REFERENCES

Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. Neural Comput 11:91–101.

Akemann W, Mutoh H, Perron A, Rossier J, Knöpfel T (2010) Imaging brain electric signals with genetically targeted voltage-sensitive fluorescent proteins. Nat Methods 7:643–649.

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7:358–366.

Bari B a, Ollerenshaw DR, Millard DC, Wang Q, Stanley GB (2013) Behavioral and electrophysiological effects of cortical microstimulation parameters. PLoS One 8:e82170.

Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A (2012) Not noisy, just wrong: the role of suboptimal inference in behavioral variability. Neuron 74:30–39.

Bermejo R, Vyas A, Zeigler HP (2002) Topography of rodent whisking--I. Two-dimensional monitoring of whisker movements. Somatosens Mot Res 19:341–346.

Boloori A-R, Jenks R a, Desbordes G, Stanley GB (2010) Encoding and decoding cortical representations of tactile features in the vibrissa system. J Neurosci 30:9990–10005.

Boloori A-R, Stanley GB (2006) The dynamics of spatiotemporal response integration in the somatosensory cortex of the vibrissa system. J Neurosci 26:3767–3782.

Brumberg JC, Pinto DJ, Simons DJ (1996) Spatial gradients and inhibitory summation in the rat whisker barrel system. J Neurophysiol 76:130–140.

Bruno RM, Khatri V, Land PW, Simons DJ (2003) Thalamocortical angular tuning domains within individual barrels of rat somatosensory cortex. J Neurosci 23:9565–9574.

Butts DA, Goldman MS (2006) Tuning curves, neuronal variability, and sensory coding. PLoS Biol 4:e92.

Carandini M, Churchland AK (2013) Probing perceptual decisions in rodents. Nat Publ Gr 16:824–831.

Carandini M, Haider B, Ha M (2012) Inhibition dominates sensory responses in the awake cortex. :0–5.

Carvell GE, Simons DJ (1990) Biometric analyses of vibrissal tactile discrimination in the rat. J Neurosci 10:2638–2648.

Carvell GE, Simons DJ (1995) Task- and subject-related differences in sensorimotor behavior during active touch. Somatosens Mot Res 12:1–9.

Castro-Alamancos M a (2004) Absence of rapid sensory adaptation in neocortex during information processing states. Neuron 41:455–464.

Chen-Bee CH, Polley DB, Brett-Green B, Prakash N, Kwon MC, Frostig RD (2000) Visualizing and quantifying evoked cortical activity assessed with intrinsic signal imaging. J Neurosci Methods 97:157–173.

Churchland MM et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 13:369–378.

Civillico EF, Contreras D (2006) Integration of evoked responses in supragranular cortex studied with optical recordings in vivo. J Neurophysiol 96:336–351.

Civillico EF, Contreras D (2012) Spatiotemporal properties of sensory responses in vivo are strongly dependent on network context. Front Syst Neurosci 6:25.

Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci 14:811–819.

Constantinople CM, Bruno RM (2011) Effects and mechanisms of wakefulness on local cortical networks. Neuron 69:1061–1068.

Crochet S, Poulet JF a, Kremer Y, Petersen CCH (2011) Synaptic mechanisms underlying sparse coding of active touch. Neuron 69:1160–1175.

Diamond M, Huang W, Ebner F (1994) Laminar comparison of somatosensory cortical plasticity. Science (80-) 265:1885–1888.

Diamond ME (2010) Texture sensation through the fingertips and the whiskers. Curr Opin Neurobiol 20:319–327.

162

Diamond ME, Armstrong-James M, Ebner FF (1993) Experience-dependent plasticity in adult rat barrel cortex. Proc Natl Acad Sci U S A 90:2082–2086.

Diamond ME, von Heimendahl M, Knutsen PM, Kleinfeld D, Ahissar E (2008) "Where" and "what" in the whisker sensorimotor system. Nat Rev Neurosci 9:601–612.

Dixon E, Shapiro A, Lu Z-L (2014) Scale-invariance in brightness illusions implicates object-level visual processing. Sci Rep 4:3900.

Drew PJ, Feldman DE (2007) Representation of moving wavefronts of whisker deflection in rat somatosensory cortex. J Neurophysiol 98:1566–1580.

Drew PJ, Feldman DE (2009) Intrinsic signal imaging of deprivation-induced contraction of whisker representations in rat somatosensory cortex. Cereb Cortex 19:331–348.

Ebara S, Kumamoto K, Matsuura T, Mazurkiewicz JE, Rice FL (2002) Similarities and differences in the innervation of mystacial vibrissal follicle-sinus complexes in the rat and cat: a confocal microscopic study. J Comp Neurol 449:103–119.

Estebanez L, El Boustani S, Destexhe A, Shulz DE (2012) Correlated input reveals coexisting coding schemes in a sensory cortex. Nat Neurosci 15:1691–1699.

Feldman DE, Brecht M (2005) Map plasticity in somatosensory cortex. Science 310:810–815.

Ferezou I, Bolea S, Petersen CCH (2006) Visualizing the cortical representation of whisker touch: voltage-sensitive dye imaging in freely moving mice. Neuron 50:617–629.

Fox K (2002) Anatomical pathways and molecular mechanisms for plasticity in the barrel cortex. Neuroscience 111:799–814.

Glazewski S, Fox K (1996) Time course of experience-dependent synaptic potentiation and depression in barrel cortex of adolescent rats. J Neurophysiol 75:1714–1729.

Gottschaldt K, Young D (1977) Quantitative aspects of responses in trigeminal relay neurones and interneurones following mechanical stimulation of sinus hairs and skin in the cat. J Physiol:85–103.

Gutierrez GJ, O'Leary T, Marder E (2013) Multiple mechanisms switch an electrically coupled, synaptically inhibited neuron between competing rhythmic oscillators. Neuron 77:845–858.

163

Harrison LM, David O, Friston KJ (2005) Stochastic models of neuronal dynamics. Philos Trans R Soc Lond B Biol Sci 360:1075–1091.

Higley MJ, Contreras D (2003) Nonlinear integration of sensory responses in the rat barrel cortex: an intracellular study in vivo. J Neurosci 23:10190–10200.

Higley MJ, Contreras D (2005) Integration of synaptic responses to neighboring whiskers in rat barrel cortex in vivo. J Neurophysiol 93:1920–1934.

Hochberg LR, Serruya MD, Friehs GM, Mukand J a, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature 442:164–171.

Hsiao SS, Lane J, Fitzgerald P (2002) Representation of orientation in the somatosensory system. Behav Brain Res 135:93–103.

Hubel D, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol 160:106–154.

Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. J Physiol 148:574–591.

Hyvärinen J, Poranen A (1978) Movement-sensitive and direction and orientation-selective cutaneous receptive fields in the hand area of the post-central gyrus in monkeys. J Physiol 58:3567–3573.

Jadhav SP, Wolfe J, Feldman DE (2009) Sparse temporal coding of elementary tactile features during active whisker sensation. Nat Neurosci 12:792–800.

Kennedy PR, Bakay R a, Moore MM, Adams K, Goldwaithe J (2000) Direct control of a computer from the human central nervous system. IEEE Trans Rehabil Eng 8:198–202.

Kerr JND, de Kock CPJ, Greenberg DS, Bruno RM, Sakmann B, Helmchen F (2007) Spatial organization of neuronal population responses in layer 2/3 of rat barrel cortex. J Neurosci 27:13316–13328.

Kleinfeld D, Delaney KR (1996) Distributed representation of vibrissa movement in the upper layers of somatosensory cortex revealed with voltage-sensitive dyes. J Comp Neurol 375:89–108.

Krupa DJ, Matell MS, Brisben a J, Oliveira LM, Nicolelis M a (2001) Behavioral properties of the trigeminal somatosensory system in rats performing whisker-dependent tactile discriminations. J Neurosci 21:5752–5763.

Kuffler SW (1953) Discharge patterns and functional organization of mammalian retina. J Neurophysiol 16:37–68.

Kuiken T (2006) Targeted reinnervation for improved prosthetic function. Phys Med Rehabil Clin N Am 17:1–13.

Kyriazi HT, Carvell GE, Brumberg JC, Simons DJ (1996) Quantitative effects of GABA and bicuculline methiodide on receptive field properties of neurons in real and simulated whisker barrels. J Neurophysiol 75:547–560.

Land PW, Simons DJ (1985) Cytochrome oxidase staining in the rat SmI barrel cortex. J Comp Neurol 238:225–235.

Lee S-H, Simons DJ (2004) Angular tuning and velocity sensitivity in different neuron classes within layer 4 of rat barrel cortex. J Neurophysiol 91:223–229.

Lustig BR, Friedman RM, Winberry JE, Ebner FF, Roe AW (2013) Voltage-sensitive dye imaging reveals shifting spatiotemporal spread of whisker-induced activity in rat barrel cortex. J Neurophysiol 109:2382–2392.

Marks LE (1994) "Recalibrating" the auditory system: the perception of loudness. J Exp Psychol Hum Percept Perform 20:382–396.

Mateo C, Avermann M, Gentet LJ, Zhang F, Deisseroth K, Petersen CCH (2011) In vivo optogenetic stimulation of neocortical excitatory neurons drives brain-state-dependent inhibition. Curr Biol 21:1593–1602.

McDonnell MD, Abbott D (2009) What is stochastic resonance? Definitions, misconceptions, debates, and its relevance to biology. PLoS Comput Biol 5:e1000348.

Merzenich MM, Kaas JH, Sur M, Lin CS (1978) Double representation of the body surface within cytoarchitectonic areas 3b and 1 in "SI" in the owl monkey (Aotus trivirgatus). J Comp Neurol 181:41–73.

Merzenich MM, Kaas JH, Wall J, Nelson RJ, Sur M, Felleman D (1983) Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. Neuroscience 8:33–55.

Millán J del R, Renkens F, Mouriño J, Gerstner W (2004) Noninvasive brain-actuated control of a mobile robot by human EEG. IEEE Trans Biomed Eng 51:1026–1033.

165

Millard DC, Stanley GB (2013) Anatomically based Bayesian decoding of the cortical response to intracortical microstimulation. In: 2013 6th International IEEE/EMBS Conference on Neural Engineering (NER), pp 1457–1460. IEEE.

Millard DC, Wang Q, Gollnick CA, Stanley GB (2013) System identification of the nonlinear dynamics in the thalamocortical circuit in response to patterned thalamic microstimulation in vivo. J Neural Eng 10:066011.

Minnery BS, Simons DJ (2003) Response properties of whisker-associated trigeminothalamic neurons in rat nucleus principalis. J Neurophysiol 89:40–56.

Morita T, Kang H, Wolfe J, Jadhav SP, Feldman DE (2011) Psychometric curve and behavioral strategies for whisker-based texture discrimination in rats. PLoS One 6:e20437.

Moss F (2004) Stochastic resonance and sensory information processing: a tutorial and review of application. Clin Neurophysiol 115:267–281.

Mountcastle VB (1957) Modality and topographic properties of single neurons of cat's somatic sensory cortex. J Neurophysiol 20:408–434.

Mountcastle VB (1997) The columnar organization of the neocortex. Brain 120 (Pt 4:701–722.

O'Connor DH, Clack NG, Huber D, Komiyama T, Myers EW, Svoboda K (2010a) Vibrissa-based object localization in head-fixed mice. J Neurosci 30:1947–1967.

O'Connor DH, Peron SP, Huber D, Svoboda K (2010b) Neural activity in barrel cortex underlying vibrissa-based object localization in mice. Neuron 67:1048–1061.

Ollerenshaw DR, Bari B a, Millard DC, Orr LE, Wang Q, Stanley GB (2012) Detection of Tactile Inputs in the Rat Vibrissa Pathway. J Neurophysiol.

Ollerenshaw DR, Zheng HJV, Millard DC, Wang Q, Stanley GB (2014) The Adaptive Trade-Off between Detection and Discrimination in Cortical Representations and Behavior. Neuron 81:1152–1164.

Penfield W, Boldrey E (1937) Somatic Motor and Sensory Representation in the Cerebral Cortex of Man as Studied by Electrical Stimulation. Brain 60:389–443.

Petersen CCH (2007) The functional organization of the barrel cortex. Neuron 56:339–355.

Petersen CCH, Grinvald A, Sakmann B (2003a) Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured in vivo by voltage-sensitive dye imaging combined with whole-cell voltage recordings and neuron reconstructions. J Neurosci 23:1298–1309.

Petersen CCH, Hahn TTG, Mehta M, Grinvald A, Sakmann B (2003b) Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. Proc Natl Acad Sci U S A 100:13638–13643.

Petersen RS, Brambilla M, Bale MR, Alenda A, Panzeri S, Montemurro M a, Maravall M (2008) Diverse and temporally precise kinetic feature selectivity in the VPm thalamic nucleus. Neuron 60:890–903.

Pinto DJ, Brumberg JC, Simons DJ (2000) Circuit dynamics and coding strategies in rodent somatosensory cortex. J Neurophysiol 83:1158–1166.

Polley DB, Chen-Bee CH, Frostig RD (1999a) Two Directions of Plasticity in the Sensory-Deprived Adult Cortex. Neuron 24:623–637.

Polley DB, Chen-Bee CH, Frostig RD (1999b) Varying the degree of single-whisker stimulation differentially affects phases of intrinsic signals in rat barrel cortex. J Neurophysiol 81:692–701.

Polley DB, Kvasnák E, Frostig RD (2004) Naturalistic experience transforms sensory maps in the adult cortex of caged animals. Nature 429:67–71.

Polley DB, Steinberg EE, Merzenich MM (2006) Perceptual learning directs auditory cortical map reorganization through top-down influences. J Neurosci 26:4970–4982.

Prinz A a, Bucher D, Marder E (2004) Similar network activity from disparate circuit parameters. Nat Neurosci 7:1345–1352.

Rema V, Armstrong-James M, Ebner FF (1998) Experience-dependent plasticity of adult rat S1 cortex requires local NMDA receptor activation. J Neurosci 18:10196–10206.

Rice FL, Mance a, Munger BL (1986) A comparative light microscopic analysis of the sensory innervation of the mystacial pad. I. Innervation of vibrissal follicle-sinus complexes. J Comp Neurol 252:154–174.

Ritt JT, Andermann ML, Moore CI (2008) Embodied information processing: vibrissa mechanics and texture features shape micromotions in actively sensing rats. Neuron 57:599–613.

Sachidhanandam S, Sreenivasan V, Kyriakatos A, Kremer Y, Petersen CCH (2013) Membrane potential correlates of sensory perception in mouse barrel cortex. Nat Neurosci 16:1671–1677.

Sato TR, Gray NW, Mainen ZF, Svoboda K (2007) The functional microarchitecture of the mouse barrel cortex. PLoS Biol 5:e189.

Schwarz C, Hentschke H, Butovas S, Haiss F, Stüttgen MC, Gerdjikov T V, Bergner CG, Waiblinger C (2010) The head-fixed behaving rat--procedures and pitfalls. Somatosens Mot Res 27:131–148.

Sheth BR, Moore CI, Sur M (1998) Temporal modulation of spatial borders in rat barrel cortex. J Neurophysiol 79:464–470.

Shoykhet M, Doherty D, Simons DJ (2000a) Coding of deflection velocity and amplitude by whisker primary afferent neurons: implications for higher level processing. Somatosens Mot Res 17:171–180.

Shoykhet M, Doherty D, Simons DJ (2000b) Coding of deflection velocity and amplitude by whisker primary afferent neurons: implications for higher level processing. Somatosens Mot Res 17:171–180.

Simons DJ (1978) Response properties of vibrissa units in rat SI somatosensory neocortex. J Neurophysiol 41:798–820.

Simons DJ (1985) Temporal and spatial integration in the rat SI vibrissa cortex. J Neurophysiol 54:615–635.

Simons DJ, Carvell GE (1989) Thalamocortical response transformation in the rat vibrissa/barrel system. J Neurophysiol 61:311–330.

Simons DJ, Land PW (1987) Early experience of tactile stimulation influences organization of somatic sensory cortex. Nature 326:694–697.

Sripati A, Yoshioka T (2006) Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system. J … 26:2101–2114.

Stallinga S, Rieger B (2010) Accuracy of the gaussian point spread function model in 2D localization microscopy. Opt Express 18:24461–24476.

Stanley GB (2013) Reading and writing the neural code. Nat Neurosci 16:259–263.

Steinbach JH, Akk G (2001) Modulation of GABA(A) receptor channel gating by pentobarbital. J Physiol 537:715–733.

Stocks N (2000) Suprathreshold stochastic resonance in multilevel threshold systems. Phys Rev Lett 84:2310–2313.

Stüttgen MC, Rüter J, Schwarz C (2006) Two psychophysical channels of whisker deflection in rats align with two neuronal classes of primary afferents. J Neurosci 26:7933–7941.

Stüttgen MC, Schwarz C (2008) Psychophysical and neurometric detection performance under stimulus uncertainty. Nat Neurosci 11:1091–1099.

Tommerdahl M, Favorov O, Whitsel BL, Nakhle B, Gonchar Y a (1993) Minicolumnar activation patterns in cat and monkey SI cortex. Cereb Cortex 3:399–411.

Tommerdahl M, Favorov O V, Whitsel BL (2010) Dynamic representations of the somatosensory cortex. Neurosci Biobehav Rev 34:160–170.

Waiblinger C, Brugger D, Schwarz C (2013) Vibrotactile Discrimination in the Rat Whisker System is Based on Neuronal Coding of Instantaneous Kinematic Cues. Cereb Cortex.

Waiblinger C, Brugger D, Shephard CJ, Millard DC, Stanley GB, Schwarz C, Neurophysiology S, Reichardt W, Neuroscience I, Coulter WH, Engineering B, Tech G (n.d.) The role of kinematic , slip-like events ' in vibrotactile perception K , k I , i F , f. :1113.

Wang Q, Millard DC, Zheng HJ V, Stanley GB (2012) Voltage-sensitive dye imaging reveals improved topographic activation of cortex in response to manipulation of thalamic microstimulation parameters. J Neural Eng 9:026008.

Wang Q, Webber RM, Stanley GB (2010) Thalamic synchrony and the adaptive gating of information flow to cortex. Nat Neurosci 13:1534–1541.

Webber RM, Stanley GB (2004) Nonlinear encoding of tactile patterns in the barrel cortex. J Neurophysiol 91:2010–2022.

Wiesel TN, Hubel DH (1963) Single-cell responses in striate cortex of kiteens deprived of vision in one eye. J Neurophysiol 26:1003–1017.

Wilke SD, Eurich CW (2002) Representational accuracy of stochastic neural populations. Neural Comput 14:155–189.

Wolfe J, Hill DN, Pahlavan S, Drew PJ, Kleinfeld D, Feldman DE (2008) Texture coding in the rat whisker system: slip-stick versus differential resonance. PLoS Biol 6:e215.

Young RD, Oliver RF (1976) Morphological changes associated with the growth cycle of vibrissal follicles in the rat. J Embryol Exp Morphol 36:597–607.