INVESTIGATIONS INTO HUMAN VIBROTACTILE PERCEPTION

INVESTIGATIONS INTO HUMAN VIBROTACTILE PERCEPTION: PSYCHOPHYSICAL EXPERIMENTS AND BAYESIAN MODELLING

By ARINDAM BHATTACHARJEE, B.Sc., M.Sc., M.Sc.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree Doctor of Philosophy

McMaster University
DOCTOR OF PHILOSOPHY (2013) Hamilton, Ontario
(Psychology)

TITLE: Investigations into human vibrotactile perception: Psychophysical experiments and Bayesian modelling

AUTHOR: Arindam Bhattacharjee, B.Sc. (Gauhati University), M.Sc. (Allahabad University), M.Sc. (McMaster University)

SUPERVISOR: Dr. Daniel Goldreich

NUMBER OF PAGES: xv, 117

Abstract

A considerable amount of our everyday tactile experience requires interactions between textured surfaces and our fingertips. Such interactions elicit complex vibrations on our skin surface, which are encoded by the mechanosensitive afferents and conveyed to the brain where the perception of the textures emerges seemingly effortlessly. Intuitively, a fundamental question that may be asked is: "what features of the vibration stimuli are behaviourally relevant and what are the neural signatures of these features?" The goal of this thesis is to investigate these questions, which we have done using a combination of theoretical and experimental approaches.

Our theoretical approach (in Chapter 2) has been to create an ideal Bayesian perceptual observer that utilizes all the information available in a spike-rate based neural code and makes optimal inferences regarding the amplitude and the frequency of vibration stimuli. Our experimental approach has been to estimate the performance of human participants in vibrotactile detection (in Chapter 3), and in amplitude and frequency discrimination (in Chapter 4) tasks by using psychophysical procedures.

The results of these approaches suggest that the human perceptual observer, i.e. the human nervous system, probably uses a rate code to represent vibrotactile amplitude, but a non-rate code, such as a spike timing code, to represent vibrotactile frequency. Additionally, we conclude that humans are capable of inferring and separately perceiving the amplitude and frequency of vibrotactile stimuli; however, depending on experimental tasks, humans might also rely on a feature that combines the amplitude and frequency of vibrotactile stimuli.

I dedicate this thesis to:
Amalesh and Sarbani Bhattacharjee, my parents, for inspiring me to dream,
Sarah, my partner in life, for encouraging and supporting me through this dream,
and, Akshayan, my son, for giving me the opportunity to weave new dreams.

Preface

This is a model-driven thesis, which consists of five chapters. Chapter 1 provides relevant background information and an overview of the goal of this thesis. Chapter 2 provides complete description of the Bayesian ideal observer model and presents simulation results that motivated the empirical studies reported in chapters 3 and 4. Chapter 5 summarizes the general findings of this thesis and discusses the implications of this body of work. None of the chapters has yet been published in any journal; however, Chapters 2, 3, and 4 have been prepared with the intention to submit as manuscripts for peer-reviewed publications.

The research work presented in this thesis was supported by a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant awarded to Dr. Daniel Goldreich. This research was also supported by scholarships and bursaries awarded to me by the Department of Psychology, Neuroscience & Behaviour (McMaster University).

Acknowledgments

As I write the last section for my thesis, I am thinking of all the people who have contributed in so many different ways to make this thesis a reality. Needless to say that one of the biggest contributions was from my graduate advisor, Dr. Daniel Goldreich and I cannot thank him enough for being such an amazing mentor, colleague, and a friend! Not only you have helped me in every aspect of my research, during past 6 years you have also trained me in English grammar – a skill that is essential for everyone, but particularly for aspiring scientists; this training has been very crucial during writing this thesis. I express my heartfelt gratitude to you. While this thesis has come to an end, I hope this is a new beginning of many years of collaboration.

I would also like to thank my committee members, Dr. Bruce Milliken and Dr. David Shore. Bruce, thank you for the encouraging words at the end of every committee meeting, and David, thank you for those difficult questions that helped us improve the model after every committee meeting.

One of the best things about working in this lab is the people of this lab. I thank each and every member of the lab, but I particularly thank all the graduate students: Jonathan Tong, Luxi Li, Onkar Marway (honorary), Ryan Peters and of course my very good friend Mike Wong with whom I have spent 5 of the 6 years in this lab. Mike, I have enjoyed all the moments we have spent together chatting about food and spices, coffee, grammar, research, and life in general; thank you for all those moments. Jon, I'll miss our coffee runs and our long useless conversations about science, religion, and politics. Onkar, your help has been very critical for this thesis; I have no words to thank you for all your help!

I cannot think of this department without thinking of Donna Waxman, Nancy Riddell, Sally Presutti, Wendy Selbie, and Gary Weatherill. I thank you all!

This thesis would not have been possible without the love and support of our lovely little family – my partner Sarah and all our children, Reine, Pery, and Akshayan! Sarah, my love, thank you for proof reading various versions of every chapter! And, even a bigger thank you for the beautiful food and treats that you make every day. The last few months have been particularly crazy and I cannot thank you enough for being so supportive and understanding.

Last but not the least, I thank my mum and dad, Sarbani and Amalesh Bhattacharjee; without your love, support, encouragement, sacrifices, and blessings nothing would have been possible!

Table of Contents

List of figures

List of Tables

List of all Abbreviations

2IFC	2-Interval Forced Choice
AD	Amplitude Discrimination
AD_{vF}	Amplitude Discrimination with variable Frequency
BF	Bayes Factor
BFL_{ntFAD}	Best Fit Line for non-target Feature Affected Discrimination
CI	Confidence Interval
FD	Frequency Discrimination
FD_{vA}	Frequency discrimination with variable Amplitude
MCS	Method of Constant Stimuli
ntFAD	non-target Feature Affected Discrimination
PC	Pacinian Corpuscles
PDF	Probability Distribution Function
RA	Rapidly Adapting
SA1	Slowly Adapting type-1
SA2	Slowly Adapting type-2
SOA	Stimulus Onset Asynchrony
TD	Threshold discrimination
TSC	Temporal Summation Curve

List of all Symbols

a	Amplitude
d	Distance
f	Frequency
t	Stimulus period
r	Single afferent or neuron response
$\{r\}$	Population response
μ_f	Expected stimulus-evoked spike count
ψ	Psychometric function
θ	Threshold parameter
β	Slope parameter of psychometric function
δ	Lapse rate parameter
l	Stimulus level
cr_l	Correct response at stimulus level l
icr_l	Incorrect response at stimulus level l
s	Summation constant
τ	Exponential decay constant
m_{ntFAD}	Slope due to non-target Feature Affected Discrimination
c	Intercept

Declaration of Academic Achievement

Chapter 2.

This chapter is a preliminary version of a manuscript to be co-authored with my graduate advisor, Dr. Daniel Goldreich. My contributions in this chapter involve conceptualizing, designing, programming, running simulations, performing statistical analyses, preparing figures, and writing. Dr. Daniel Goldreich contributed by providing guidance at every stage, programming, and editing the manuscript.

Chapter 3.

This chapter is a preliminary version of a manuscript to be co-authored with Mobolaji Adeolu, Onkar Marway, and Dr. Daniel Goldreich. Mobolaji Adeolu and Onkar Marway were undergraduate thesis students (academic years: 2011, 2012). Dr. Daniel Goldreich and I conceived the study and the experimental design. Mobolaji Adeolu, Onkar Marway, and I collected the data. My other contributions in this chapter involve performing statistical analyses, preparing figures, and writing. Mobolaji Adeolu completed an undergraduate thesis based on the data reported in experiments 1, 2, and 3. Dr. Daniel Goldreich contributed by providing guidance at every stage, especially on performing statistical analyses, and editing the manuscript.

Chapter 4.

This chapter is a preliminary version of a manuscript to be co-authored with Onkar Marway, and Dr. Daniel Goldreich. Dr. Daniel Goldreich and I conceived the study and the experimental design. Onkar Marway, and I collected the data. My other contributions in this chapter involve programming, running simulations, performing statistical analyses, preparing figures, and writing. Dr. Daniel Goldreich contributed by providing guidance at every stage, especially on programming, performing statistical analyses, and editing the manuscript.

CHAPTER 1: GENERAL INTRODUCTION

1.1 Perception

Perception materializes as the brain interprets stimulus-evoked sensory data based on its internal hypotheses about the external world. This implies that perceptual accuracy critically depends on – a) the fidelity of the sensorineural response to the stimulus, and b) the efficiency of the brain's interpretation or inference based on the sensorineural data. Note that when sensory data reaches the brain, it is obligated to process this data and determine the feature(s) of the stimulus that might have evoked the sensorineural data. Therefore, the quality of perception is primarily affected by the fidelity of the stimulus transformation process that evokes the sensory data, and by the efficiency of the brain in deriving the stimulus information from the acquired sensory data.

The stimulus transformation process in the periphery is the first stage of encoding. Here, various factors can affect the fidelity of the sensorineural response. For example: intuitively, receptor density should affect the fidelity with which a stimulus is represented in the sensory data. Recently, Peters et al. (2009) demonstrated that people with smaller fingers, who presumably have more densely packed receptors than people with larger fingers, are better able to discern the spatial details of passively presented textured surfaces to their fingertip. There is no obvious reason to believe that humans with larger fingers, compared to those with smaller fingers, are less efficient in deciphering tactile stimuli yet there exists a difference in performance, which corroborates the effect of receptor density and demonstrates that the efficiency of encoding is critical for perception.

Once the stimulus has been encoded, the efficiency with which the brain deciphers (i.e., "decodes") the relevant stimulus information from the sensorineural data will ultimately determine the accuracy of perception. For example, studies have consistently shown that blind humans perform better than sighted humans on a variety of tactile tasks (Goldreich and Kanics, 2003; Bhattacharjee et al., 2010; Wong et al., 2011; see for review Wong, 2013); however, there is no reason to believe that the fingertips of a blind person are any different from those of a sighted person. Presumably, cortical plasticity has enhanced the decoding capability in blind humans; this enhancement in decoding efficiency increases perceptual accuracy.

In a recent study, Wong et al. (2013) trained sighted participants to discriminate orientation of fine textures presented passively to the participants' fingertip. The authors observed that after training the participants' spatial acuity enhanced; however,

intriguingly, this enhancement was limited by the participants' finger size, i.e. the proxy for their receptor density. This can be interpreted as an indication that the perceptual decoder could be trained to increase its efficiency; however, perceptual accuracy is ultimately limited by the fidelity of the encoding process.

Another recent study beautifully demonstrates how the efficiency of the encoding and the decoding interplay in perception, which is reflected in the overall performance in psychophysical tasks. Peters and Goldreich (submitted) following from Peters et al. (2009) predicted that if humans with smaller finger, compared to those with larger fingers, perform superiorly in tactile spatial acuity tasks, then children (who tend to have smaller fingers than adults) should have much better spatial acuity than adults. Surprisingly, Peters and Goldreich (submitted) did not find any difference in acuity between the two groups. To explain the results they reasoned that whereas younger children have smaller fingers, their central nervous system is still undergoing development and has not reached adult level maturation. The authors observed that, when they controlled for finger size, the performance of the children in the spatial acuity task enhanced with age, which validates their reasoning. Therefore, despite acquiring presumably very high fidelity sensory information the decoder in the younger children could not attain the efficiency acquired by the adults. As children grow to adulthood, the receptor density decreases and decoding efficiency increases, two processes that counteract each other; consequently the overall estimates of spatial acuity remain approximately constant. Later in life, aging presumably causes loss of receptors but the efficiency of the perceptual decoder stays relatively stable; as a result, tactile acuity declines with aging in adulthood (Gescheider et al., 1994; Goldreich and Kanics, 2003; Bhattacharjee et al., 2010; Wong et al., 2011).

Evidently, the failure to observe a quantitative difference in performance between the children and the adults in Peters and Goldreich (submitted) suggests that psychophysical investigations of human perceptual tasks are only capable of quantifying our general ability to perceive different stimuli, and they cannot separately quantify the efficiency of the individual aforementioned steps of perception, i.e. the efficiency of the encoder and of the decoder. This limitation of psychophysical experimentation is overcome by the use of ideal observer models, which do allow us to separately study the encoding and decoding processes.

Ideal observer analysis: This is an approach to determine the physical limits of perception. Central to this idea is an observer that ideally (i.e., optimally) decodes the sensory information – hence, the "ideal observer". Conceptually, this analytical tool allows us to ask a simple, and yet fundamental, question in sensory neuroscience: Given full knowledge of the stimulus-evoked neural response (and the means by which that response is produced), what is the best performance that is possible? (Geisler, 2003). Although this concept has been known for a while, the application of this analytical tool became more popular during the 1980s (see Geisler, 1987). Because sensory information gets transformed at each level of sensory processing, from the activation of the receptors

to cortex, Geisler (1989) has demonstrated that this analytical tool – which the author called sequential ideal-observer analysis – can be applied at each level of sensory processing to determine the best performance and to quantify the progressive loss of information as it moves centrally through the nervous system.

1.2 Tactile perception

To study perception, I investigated the sense of touch as a model system. More specifically, I studied humans' ability to perceive vibration presented on the glabrous part of the hand. During movement of our fingers over textured surfaces, or when textured surfaces are moved against our fingertip, the interaction between the fingertip and the textured surface generates complex vibrations (Bensmaïa and Hollins, 2003; Weber et al., 2013). Interestingly, the fingerprint also contributes in this interaction. Scheibert et al. (2009) recently proposed that the fingerprint might work as a band-pass filter to accentuate or attenuate certain frequencies that are generated by the fingertip-texture movement. Presumably, even when the finger is scanned over a flat surface, vibrations are thereby induced as a result of the fingerprint. Therefore, to understand texture perception or even the perception of smooth surfaces via the moving finger, we have to understand an even more fundamental question: how do we perceive vibrations?

There are four different tactile afferents that can be activated by vibrotactile stimuli of different frequencies (see for description: Johnson et al., 2000; Mountcastle, 2005; Gescheider et al., 2009). Stimuli with very low frequencies (0.5 to 5 Hz) activate the slowly-adapting type 1 (SA1) afferents, which terminate in the Merkel receptors. Functionally, these afferents and the corresponding receptors are activated by static touch and are involved in tactile spatial tasks. Stimuli with frequencies ranging from 5Hz to 50Hz activate the rapidly-adapting type (RA) afferents, which terminate in the Meissner's corpuscles. These afferents, functionally, are activated by micro-slippage or minute skin motion. Interestingly, the speed with which proficient Braille readers move their finger over Braille text suggests that they might be recruiting these afferents during Braille reading (Davidson, 1992; Bhattacharjee et al., 2010). Pacinian (PC) afferents, which terminate in the Pacinian corpuscles, are activated by stimuli with frequencies from 50Hz to 500 Hz (Bolanowski et al., 1988) or higher, and the best frequencies are around 250Hz to 300Hz. For instance, these afferents respond to vibrations produced by scanning a hard tool across a textured surface. The slow-adapting type 2 (or SA2) afferents apparently do not produce a conscious tactile percept (Ochoa and Torebjörk, 1983; Johnson, 2000). These afferents, which terminate in the Ruffini endings, signal skin stretch that may be important for proporioceptive feedback; they are activated at high frequencies around 100Hz to 500Hz but only with very small vibratory probes (Verrillo and Bolanowski, 1986).

Hollins and Risner (2000) argued that humans discriminate coarse texture, when presented passively, probably by using the sensory information in the SA1 afferents; however, to perceive fine textures finger movement might be necessary, which transforms the texture information into vibrotactile stimuli. Interestingly, both RA and PC afferents are sensitive to vibrotactile stimuli; therefore, to determine whether the RA1 afferents or the PC afferents are more informative for texture discrimination, Hollins and colleagues, conducted texture discrimination tasks on human participants using an adaptation paradigm (Hollins et al., 2001). In different experimental conditions, the authors selectively adapted the PC (with a 250 Hz vibration) and the RA afferents (with a 10 Hz vibration), and observed that human participants' performance worsened only when PC afferents were adapted, which led Hollins et al. (2001) to conclude that PC afferents are more important for texture discrimination.

Whereas Hollins et al. (2001) have suggested that the sensory information encoded in the PC afferents are more important, results from Gamzu and Ahissar (2001) suggest that human participants might prefer to encode sensory information by activating the RA afferents. Gamzu and Ahissar (2001) asked participants to discriminate textures with spatial gratings, i.e. the textures had grooves and ridges that varied in spatial frequency. The participants freely scanned the textures with their fingertip while wearing a glove with a tiny probe attached near the tip of the index finger such that only the probe touched the textured surface. Interestingly, the participants adjusted their scanning speed depending on the spatial frequency of the stimulus pieces such that the movement elicited temporal frequencies that primarily activate the RA afferents.

1.3 Which features of vibration are behaviourally relevant?

Miyaoka et al. (1999) investigated whether human participants utilized amplitude information of texture particles during texture discrimination. The investigators first estimated the discriminability of textured surfaces that varied in texture particle size; next, they reasoned that if participants were utilizing amplitude information of the particles then the participants' texture discrimination performance should match a ridge-height discrimination task (similar to an amplitude discrimination task). The authors found that the participants' performance on both tasks were similar, which led them to conclude that to discriminate between fine-surface textures, humans might use amplitude information of texture particles.

Results from Gamzu and Ahissar (2001) suggest those temporal cues generated by movement-induced vibrations are behaviourally important for texture perception. Similarly, Cascio and Sathian (2001) tested humans on roughness discrimination experiments by using gratings that varied in groove and ridge width, which were

passively moved below human participants' index fingertip, and observed that the participants' performance in roughness discrimination declined when the cue based on the temporal frequency was eliminated. These studies suggest that frequency of vibrations is one of the relevant cues. Interestingly, Bensmaïa and Hollins (2003) argued that instead of the frequency cues humans use intensity cues for fine texture discrimination.

In fact, studies that did not investigate texture discrimination but investigated vibrotactile perception, found that humans use intensity cues in human frequency discrimination tasks (Goff, 1967; LaMotte and Mountcastle, 1975). More recently, Harris et al. (2006) proposed that humans encode the product of the amplitude and frequency of the stimulus (i.e., stimulus energy) as the cue to perform frequency discrimination tasks.

All these studies considered together suggest that depending on the task, humans might use different features (amplitude, frequency, intensity, and/or energy) for vibrotactile stimulus discrimination tasks.

1.4 What is the neural code?

Central to the discussion of perception is, not only which stimulus features are encoded, but how are they encoded.

Afferent level coding: In a landmark study, Talbot et al. (1968) showed that two different afferent types are involved in the perception of low (flutter-vibration) and high (vibration) frequency vibratory stimuli, which were RA and Pacinian afferents, respectively. Additionally, the authors observed that as they linearly increased the amplitude of a low (flutter) frequency vibration, the human participants' estimation of the subjective magnitude also increased linearly. However, neural recordings obtained from monkeys showed that the firing rate of each RA afferent surprisingly did not follow the behavioural trend. In each RA afferent, as the amplitude of the stimulus was increased, the firing rate increased but only up to a certain amplitude value beyond which it responded (or entrained) at a rate of 1 spike/stimulus-cycle, and not until the stimulus amplitude crossed some other high critical amplitude value did the firing rate start to increase again. This was a perplexing result because there existed an apparent plateau where the change in amplitude had no effect on firing rate of the RA afferent yet the subjective magnitude matched the increase in stimulus amplitude. Based this observation, Johnson (1974) hypothesized that the perceived intensity depends on the activity of the population of RA afferents rather than the activity of any single afferent. To this end, Johnson (1974) proposed several candidate neural codes and tested whether each neural code could reliably describe the behavioural results. The author reported that not a single but a collection of neural codes probably determines perceived intensity, such as the total population activity, and various codes related to the total number of active fibers.

Recently, Bensmaïa and colleagues (Muniak et al., 2007) extended Johnson's (1974) study by recording neural responses from RA, PC, and SA1 afferents and by using a range of stimulus types, such as diharmonic, noise stimuli, and sinusoidal stimuli. After testing eight plausible hypotheses, the authors concluded that perceived intensity depends on the firing rates in each afferent underneath or near the probe and the activity evoked in RA, PC, and SA1 afferents (or "weighed by afferent type"; Muniak et al., 2007).

More recently, Güçlü and Dinçer (2013) investigated the neural codes that might explain human sensation magnitude exclusively in the RA afferents. To accomplish this, the authors estimated the participants' sensation magnitude after adapting the PC afferents (which might contribute to the sensation magnitude, according to Muniak et al., 2007), and compared simulation results based on five hypothetical neural codes. The authors found that the human data matched the simulation results when they considered the neural codes that were based on the 'number active fibers', 'the distribution of spike count', and 'the total spike count' (Güçlü and Dinçer, 2013).

Whereas the above mentioned neural codes put forth to explain perceived intensity were based on spike counts, the strong entrainment seen in the afferent recordings suggests that stimulus frequency might be encoded not by spike count but by the temporal regularity of spiking activity (Talbot et al., 1968; Johnson, 1974). Whereas Güçlü and Dinçer (2013) reported that the neural code based on interspike interval did not match the human sensation magnitude, Whitsel et al (2000) strongly argued that the entrainment in the RA afferents code the frequency of the stimulus.

Cortical level coding: In a seminal work Mountcastle et al. (1969), tested whether firing rates of the RA-like neurons (cortical neurons that show characteristics of RA afferents) modulate with changes in the stimulus frequency. However, these neurons entrained to the stimulus frequency, particularly between 20Hz and 40Hz, which led Mountcastle et al. (1969) to propose the periodicity hypothesis for encoding stimulus frequency.

Relatively recently, Whitsel et al. (2001) reinvestigated the cortical coding strategy for vibrotactile stimulus frequency by recording from the RA afferent as well as RA-like cortical neurons. The authors observed that whereas the mean firing rate in the RA afferents increased monotonically, the mean firing rate in the RA-like neurons did not change with increasing frequency. Whitsel et al. (2001) illustrated a convincing comparison between human data acquired from other frequency discrimination study (for example, Goff, 1967; Mountcastle et al., 1969, 1990) and the predicted frequency discrimination threshold based on the entrainment neural code showed remarkable resemblance.

In an interesting study, Torebjörk et al. (1987) using the microstimulation technique (where small currents are injected into the tactile afferents) showed activation of single afferents could elicit sensation similar to those felt during mechanical stimulation.

Intriguingly, participants reported ‘tapping on the skin’, ‘flutter sensation’, and ‘buzzing vibration’, as the authors increased the stimulation frequency from low, to 20Hz – 50Hz, and to frequencies above 50Hz, respectively.

Whereas these aforementioned studies provide evidence for the periodicity or temporal code, Romo and colleagues, using various experimental and analytic techniques, have been consistently providing evidence for spike count or rate code as a neural code for stimulus frequency (see for review, Romo and Salinas, 2001, 2003).

Romo et al. (1998) conducted frequency discrimination experiments on monkeys in which they presented mechanical stimulation on the monkeys’ fingertip, and / or microelectrical stimulus to the fingertip representation in the primary somatosensory cortex. The authors reported that regardless of the stimulation technique, the performance by the monkeys was indistinguishable and highly reliable. Furthermore, to test whether frequency discrimination requires periodic activation of relevant cortical neurons, Romo et al. (1998) created aperiodic stimuli by keeping the total duration of the stimuli similar to the periodic stimuli, but jittered the inter-pulse intervals. Once again, irrespectively of the stimulation technique, the frequency discrimination results suggested that the monkeys were able to extract the mean frequency from the aperiodic stimuli and thereby perform the frequency discrimination tasks.

In Salinas et al. (2000), Romo and his colleagues tested monkeys’ ability to perform frequency discrimination task where the stimuli were periodic and / or aperiodic. The authors reported that the behavioural results were similar to Romo et al. (1998). In addition, using an information theoretic approach, the authors calculated the amount of information that was available in the stimulus evoked neural responses based on periodicity and on firing rate. Salinas et al. (2000) observed that the neural code based on periodicity had higher information content than the code based on firing rate; however, if the monkeys had used the periodicity code, the performance of the monkeys should have been at least 3 times better than the observed performance. Based on this result the authors suggested that the monkeys were most likely using spike rate as the neural code for frequency discrimination.

In Hernández et al. (2000), Romo and colleagues conducted experiments similar to the Salinas et al. (2000) study. However, in the current study, the authors created neurometric functions based on periodicity and firing rate as two different measures, and compared the psychometric function of the monkeys to the derived neurometric functions. Like Hernández et al. (2000), Salinas et al. (2000) observed that the behavioural performance predicted spike rate as the neural code, i.e. the neurometric function derived from the spike rate as the decision criterion was more similar to the psychometric function than was the periodicity based neurometric function. Because the periodicity based neurometric function predicted much lower thresholds than the spike rate based neurometric function, Romo and colleagues once again suggested that the monkeys must be using a neural code based on spike rate to perform frequency discrimination tasks.

Interestingly, whereas this debate regarding the neural code exists in the somatosensory field, conceptually the distinction between rate code and temporal code is not absolute. The idea of the rate code strictly depends on a window within which the number of spikes are counted, and if we consider a small enough window that allows only a single spike to occur then the rate code becomes similar to a temporal code (Rieke et al., 1999). Indeed, it can be difficult to clearly define and distinguish the terms temporal and rate coding. One proposal, put forth by Dayan and Abbott (Dayan and Abbott, 2001), is that a temporal code is one in which information is carried by variations in spike rate that occur more rapidly than variations in the stimulus itself. In this thesis, however, we use the term "temporal code" as it is commonly used in the somatosensory literature, to mean that information is carried by the stimulus-driven time-varying responses of neurons, rather than by the average neuronal firing rates within a larger temporal window.

1.5 Thesis overview

Both neural coding strategies – periodicity coding and spike-rate coding, seem plausible in light of previous investigations. Whereas most of the aforementioned studies have linked encoding strategies to humans' or monkeys' perceptual behaviour during sensory testing, note that the thresholds obtained at the end of psychophysical experiments are quantified estimates of perceptual encoding and decoding processes. For example, Romo and colleagues claim that periodicity in the stimulus-evoked neural responses have high information content (Salinas et al., 2000) and if humans' or monkeys' were utilizing that encoding strategy for frequency discrimination then the animals' threshold should have been much better than that observed in experiments. Salinas et al. (2000) have demonstrated that based on the periodicity of neural responses, the nervous system could conceivably perceive the temporal features of the vibrotactile stimuli with very high fidelity. However, it is quite plausible that humans and monkeys are suboptimal (Putzeys et al., 2012) in the sense that they do not take advantage, during the decoding process, of all information in the neural response. Indeed, Salinas et al. (2000) show that the periodicity code contains much more precise information than is evident in monkey's performance, and suggest instead that monkey use a rate code for frequency discrimination. To further investigate whether this possibility applies as well to human observers, in Chapter 2, we create a Bayesian ideal observer model that optimally decodes rate code based simulated neural responses to vibrotactile stimuli, and in Chapters 3 and 4 we compare human behavior to the predictions resulting from the Bayesian observer.

In chapter 2, we show that our Bayesian model has two parts, a generative module or the encoder, and an inference module or the decoder. The encoder implements the stimulus-response functions reported in Johnson (1974) and Freeman and Johnson (1982), using

the afferent density of human glabrous skin reported in Johansson and Vallbo (1979). The decoder infers the amplitude and frequency feature values. Following the logic of Geisler (1989), we first fed the model the neural responses of RA afferents; this allowed us to determine the optimal performance based on the full information content contained in the spike rates of the first neurons that respond to the vibrotactile stimulus. Next, we added Poisson variability to the afferent responses in order to determine the degree of information loss due to this feature of cortical neuronal firing rates. We quantified the model's performance on common vibrotactile tasks that are conducted on humans, namely, vibrotactile threshold detection, amplitude discrimination, and frequency discrimination tasks.

In chapter 3, we tested the model's predictions that spatial and temporal summation should exist in the human RA afferent system – the phenomena that increasing the number of activated afferents, and increasing the duration of stimulus presentation, reduces the threshold (i.e. the amplitude required to detect the stimulus). Our results support our prediction and suggest that spatial and temporal summation exist in RA afferent system.

In chapter 4, we compared the performance of our rate code based ideal observer model against human performance on a range of amplitude and frequency discrimination tasks. We examined three stimulus features – amplitude, frequency, and the product of amplitude and frequency (i.e. stimulus energy) – that the humans might use to discriminate vibrotactile stimuli. Because vibrotactile stimuli have two concomitant features – amplitude and frequency – we investigated the effect of each feature on tasks requiring the discrimination of the other. We also investigated whether humans can perceive separately the amplitude and frequency of a vibrotactile stimulus. We obtained psychophysical results that suggest that humans do have access to the individual stimulus features, and the non-target parameter does not affect discrimination performance unless the non-target parameter is surreptitiously changed during the experiment.

Collectively, the results reported in this thesis suggest that humans utilize periodicity information while performing the frequency discrimination tasks, but use spike-rate code to perform the amplitude discrimination tasks.

1.6 References

Bensmaïa SJ, Hollins M (2003) The vibrations of texture. Somatosens Mot Res 20(1):33-43.

Bhattacharjee A, Ye AJ, Lisak JA, Vargas MG, Goldreich D (2010) Vibrotactile masking experiments reveal accelerated somatosensory processing in congenitally blind braille readers. J Neurosci 30:14288-14298.

Dayan P, Abbott LF (2001) Theoretical neuroscience: computational and mathematical modeling of neural systems. Cambridge: MIT Press.

Davidson PW, Apelle S, Haber RN (1992) Haptic scanning of Braille cells by low- and high-proficiency blind readers. Res Dev Disabil 13:99-111.

Freeman AW, Johnson KO (1982) A model accounting for effects of vibratory amplitude on responses of cutaneous mechanoreceptors in macaque monkey. J Physiol 323:43-64.

Geisler WS (1987) Ideal-observer analysis of visual discrimination. In Frontiers of Visual Science: Proceedings of the 1985 Symposium National academy press, Washington.

Geisler WS (1989) Sequential ideal-observer analysis of visual discriminations. Psychol Rev 96:267-314.

Geisler WS (2003) Ideal observer analysis. In: L Chalupa and J Werner (eds.), The visual neuroscience. Boston: MIT press, 825 – 837.

Gescheider GA, Bolanowski SJ, Hall KL, Hoffman KE, Verrillo RT (1994) The effects of aging on information-processing channels in the sense of touch: I. Absolute sensitivity. Somatosens Mot Res 11:345-357.

Gescheider GA, Wright JH, Verrillo RT (2009) Information-Processing Channels in the Tactile Sensory System: A Psychophysical and Physiological Analysis. New York: Psychology Press.

Goff GD (1967) Differential discrimination of frequency of cutaneous mechanical vibration. J Exp Psychol 74(2):294-299.

Goldreich D (2007) A Bayesian Perceptual Model Replicates the Cutaneous Rabbit and Other Tactile Spatiotemporal Illusions. PLoS ONE 2(3): e333. doi:10.1371/journal.pone.0000333

Goff GD (1967) Differential discrimination of frequency of cutaneous mechanical vibration. J Exp Psychol 74(2):294-9.

Güçlü B, Dinçer SM (2013) Neural coding in the Non-Pacinian I tactile channel: a psychophysical and simulation study of magnitude estimation. Somatosens Mot Res 30:1-15.

Harris JA, Arabzadeh E, Fairhall AL, Benito C, Diamond ME (2006) Factors affecting frequency discrimination of vibrotactile stimuli: implications for cortical encoding. PLoS One 1:e100.

Hernández A, Zainos A, Romo R (2000) Neuronal correlates of sensory discrimination in the somatosensory cortex. Proc Natl Acad Sci U S A 97:6191-6196.

Hollins M, Bensmaïa SJ, Washburn S (2001) Vibrotactile adaptation impairs discrimination of fine, but not coarse, textures. Somatosens Mot Res 18(4):253-62.

Hollins M, Risner SR (2000) Evidence for the duplex theory of tactile texture perception. Percept Psychophys 62(4):695-705.

Johnson KO (1974) Reconstruction of population response to a vibratory stimulus in quickly adapting mechanoreceptive afferent fiber population innervating glabrous skin of the monkey. J Neurophysiol 37:48-72.

Johnson KO, Yoshioka T, Vega-Bermudez F (2000) Tactile functions of mechanoreceptive afferents innervating the hand. J Clin Neurophysiol 17(6):539-58.

Johansson RS, Vallbo AB (1979) Detection of tactile stimuli. Thresholds of afferent units related to psychophysical thresholds in the human hand. J Physiol 297(0):405-22.

LaMotte RH, Mountcastle VB (1975) Capacities of humans and monkeys to discriminate vibratory stimuli of different frequency and amplitude: a correlation between neural events and psychological measurements. J Neurophysiol 38(3):539-59.

Miyaoka T, Mano T, Ohka M (1999) Mechanisms of fine-surface-texture discrimination in human tactile sensation. J Acoust Soc Am 105(4):2485-92.

Mountcastle VB (2005) The sensory hand: neural mechanisms of somatic sensation. Harvard University Press.

Mountcastle VB, Steinmetz MA, Romo R (1990) Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. J Neurosci 10(9):3032-44.

Mountcastle VB, Talbot WH, Sakata H, Hyvärinen J (1969) Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. J Neurophysiol 32(3):452-84.

Muniak MA, Ray S, Hsiao SS, Dammann JF, Bensmaïa SJ (2007) The neural coding of stimulus intensity: linking the population response of mechanoreceptive afferents with psychophysical behavior. J Neurosci 27:11687-11699.

Ochoa J, Torebjörk E (1983) Sensations evoked by intraneural microstimulation of single mechanoreceptor units innervating the human hand. J Physiol 342:633-654.

Peters RM, Goldreich D (submitted) Tactile spatial acuity in childhood: effects of age and fingertip size.

Peters RM, Hackeman E, Goldreich D (2009) Diminutive digits discern delicate details: fingertip size and the sex difference in tactile spatial acuity. J Neurosci 29:15756-15761.

Putzeys T, Bethge M, Wichmann F, Wagemans J, Goris R (2012) A new perceptual bias reveals suboptimal population decoding of sensory responses. PLoS Comput Biol 8(4):e1002453.

Rieke F, Warland D, de de Ruyter van Steveninck R, Bialek W (1999) Spikes: exploring the neural code. London: MIT Press.

Romo R, Hernández A, Zainos A, Salinas E (1998) Somatosensory discrimination based on cortical microstimulation. Nature 392:387-390.

Romo R, Salinas E (2001) Touch and go: decision-making mechanisms in somatosensation. Annu Rev Neurosci 24:107-137.

Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. Nat Rev Neurosci 4:203-18.

Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. J Neurosci 20:5503-5515.

Scheibert J, Leurent S, Prevost A, Debrégeas G (2009) The role of fingerprints in the coding of tactile information probed with a biomimetic sensor. Science 323:1503-1506.

Talbot WH, Darian-Smith I, Kornhuber HH, Mountcastle VB (1968) The sense of flutter-vibration: comparison of the human capacity with response patterns of mechanoreceptive afferents from the monkey hand. J Neurophysiol 31(2):301-34.

Torebjörk HE, Vallbo AB, Ochoa JL (1987) Intraneural microstimulation in man. Its relation to specificity of tactile sensations. Brain 110:1509-1529.

Verrillo RT, Bolanowski SJ Jr (1986) The effects of skin temperature on the psychophysical responses to vibration on glabrous and hairy skin. J Acoust Soc Am 80:528-532.

Weber AI, Saal HP, Lieber JD, Cheng JW, Manfredi LR, Dammann JF 3rd, Bensmaïa SJ (2013) Spatial and temporal codes mediate the tactile perception of natural textures. Proc Natl Acad Sci U S A 110:17107-17112.

Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. Nat Neurosci 5:598-604.

Whitsel BL, Kelly EF, Delemos KA, Xu M, Quibrera PM (2000) Stability of rapidly adapting afferent entrainment vs responsivity. Somatosens Mot Res 17:13-31.

Whitsel BL, Kelly EF, Xu M, Tommerdahl M, Quibrera M (2001) Frequency-dependent response of SI RA-class neurons to vibrotactile stimulation of the receptive field. Somatosens Mot Res 18:263-285.

Wong M (2013) What drives tactile spatial acuity enhancement in the blind? Open Access Dissertations and Theses. http://digitalcommons.mcmaster.ca/opendissertations/7343

Wong M, Gnanakumaran V, Goldreich D (2011) Tactile spatial acuity enhancement in blindness: evidence for experience-dependent mechanisms. J Neurosci 31:7028-7037.

Wong M, Peters RM, Goldreich D (2013) A physical constraint on perceptual learning: tactile spatial acuity improves with training to a limit set by finger size. J Neurosci 33:9345-9352.

CHAPTER 2: PERFORMANCE OF A BAYESIAN IDEAL-OBSEVER MODEL ON VIBROTACTILE DETECTION AND DISCRIMINATION TASKS

2.1 Introduction

Our ability to perceive the world appears deceptively simple! However, during this process our brain constantly encounters a challenging problem: how to efficiently infer the information about the world from the stimulus-evoked sensorineural data. Whereas the neural responses generated by the sensory system are informative about the sensory event, the stochastic nature of the encoding process introduces ambiguity in the stimulus-evoked sensory data. An identical stimulus elicits a range of neural responses, and different stimuli can evoke the same neural responses. In short, there is not a one-to-one mapping from stimulus to neural responses, or from response back to stimulus. Under such conditions, the brain must somehow decode the neural response to infer the stimulus that caused it. How does the brain accomplish this difficult task?

When faced with uncertainty, the ideal strategy for the brain is to generate probabilistic inferences about the stimulus that might have evoked the acquired sensory data (Gold and Shadlen, 2007; Goldreich and Tong, 2013). Interestingly, over the past decade several studies have provided convincing evidence that the brain performs probabilistic inference (Pouget et al 2000; Yang and Shadlen, 2007). Furthermore, studies have also shown that humans' performance in perceptual tasks could be expressed in a Bayesian framework (Ernst and Banks, 2002; Geisler and Kersten, 2002; Knill and Saunders, 2003; Knill and Pouget, 2004; Ma et al., 2006; Goldreich, 2007; Kording, 2007; Wozny et al., 2008). In the current paper, we present a Bayesian ideal-observer model that decodes neural activity elicited by vibrotactile stimuli, i.e. vibrations presented on skin surfaces.

Ideal-observer models, which are often described in a Bayesian decision theoretic framework, are analytical tools that are applied to determine the optimal performance possible in a given perceptual task (Geisler, 2003). Because of the neural response variability introduced during stimulus encoding, ideal-observer models (like humans) make errors in perceptual tasks. However, given the amount of information available to the ideal-observer, it decodes the neural responses most efficiently. Therefore, the optimality in the ideal-observers' performance sets a reliable benchmark against which to compare human performance (Green and Swets, 1966). Furthermore, ideal observer

models can be used to investigate the effects of biologically plausible constraints on performance accuracy; for example, Bayesian models can be used to assess the effect of different values of receptor density, or the effect of different degrees of neural response variability, on perceptual tasks.

Four different types of mechanoreceptive afferents – slowly adapting 1 (SA1), slowly adapting 2 (SA2), rapidly adapting (RA), and Pacinian (PC) – are present in the human skin and carry vibrotactile information to the brain (Bolanowski et al. 1988). The RA afferents are most sensitive to intermediate frequencies, responding best to stimuli between 30 and 50 Hz, whereas SA1 afferents are sensitive to low frequency vibrotactile stimulations (0.4 to 2 Hz), and PC afferents are most sensitive to high frequency vibrotactile stimuli (200-300 Hz) (Gescheider et al., 2004). The SA2 afferents are only activated using small diameter stimulators, and respond to frequencies above 100 Hz (Verrillo and Bolanowski, 1986).

Although all four types of mechanoreceptive afferents may differentially contribute to vibrotactile perception in daily life (Muniak et al., 2007; Mackevicius et al., 2013), only the RA afferents respond the best to the vibrotactile stimuli in mid frequency range (Talbot et al., 1968; Johnson et al., 2000) that we have used to test humans in the studies described in this thesis. Therefore, we programmed our Bayesian model to decode only neural activity observed in RA afferents (Johnson, 1974, Freeman and Johnson, 1982). By optimally decoding the sensorineural data generated only from RA afferents, our ideal-observer model provides us with an estimate of the best performance achievable based on the spike rates of RA afferents alone.

Periodic vibrations have two parameters - amplitude and frequency; a key issue for neuroscientists is to identify the neural code that represents these parameters. Vibrotactile amplitude or intensity is thought to be encoded by the number of spikes that occur within a certain time period, i.e. a spike rate code (Bensmaïa, 2008; Tommerdahl et al., 2010); however, the neural code for vibrotactile frequency is unclear (see Johnson et al., 2000 for a review). Some studies suggest that vibrotactile frequency is represented by a temporal code. According to this hypothesis, the regularity or the periodicity of the spiking activity in neurons, which is phase-locked to the stimulus periodicity, encodes vibrotactile frequency (Mountcastle et al. 1969, 1990). However, other studies argue that, as for vibrotactile amplitude, vibrotactile frequency is represented by a spike rate code (Hernandez et al 2000; Salinas et al 2000; Luna et al., 2005). In the current study, we have constructed a Bayesian ideal observer that uses spike rate as the neural code for both vibrotactile amplitude and frequency. We do not discard the possibility of a temporal code for vibrotactile frequency; rather, our strategy is to learn how well an ideal observer makes perceptual decisions based only on the rate code.

Johnson (1974) recorded neural activity from the RA afferents of macaque monkeys and measured the afferents' response to vibrations at different amplitudes. Studying the same afferent type in the same primate species, Freeman and Johnson (1982) investigated the

changes in neural activity at different vibrotactile frequencies between 5 Hz and 80 Hz. By utilizing the information reported by Johnson (1974) and Freeman and Johnson (1982), we simulated neural responses to vibrotactile stimuli that activate the RA afferents. In the current study, because our simulated neural responses were derived from the recordings done on monkeys, we are making an assumption that these responses to identical vibrotactile stimuli are representative of those expected in human afferents. Interestingly, two different studies on vibrotactile detection and discrimination tasks have provided evidence that support the plausibility of this assumption (Mountcastle et al., 1972; LaMotte and Mountcastle, 1975). Mountcastle and colleagues, in both studies, observed that on identical stimuli the human participants and the macaque monkeys performed very similarly, and that the psychophysical thresholds of both primate species were comparable in vibrotactile detection and frequency discrimination tasks. Therefore, we believe it is valid to compare the performance of our ideal-observer model, which decodes simulated neural activity based on recordings in monkeys, to that of humans.

To perform ideal-observer analysis, we generated neural data based on the stimulus-response functions characterized by Johnson (1974); however, these stimulus-response functions do not provide information about neural response variability. As we mentioned above, neural response variability is typical in all sensory systems – identical stimuli rarely evoke identical neural responses. Interestingly, different degrees of variability exist in the somatosensory pathway. For example, the variability in firing rates of peripheral afferents is lower than that in primary cortical areas. To realistically simulate the responses of RA afferents, we incorporated RA afferent neural variability (Vega-Bermudez and Johnson, 1999) into the stimulus-response function (Johnson, 1974). To simulate cortical responses, we added firing rate variability characteristic of primary somatosensory cortex (Sripati et al., 2006). Apart from this variability, the brain also exhibits neural activity in the absence of any stimulus, which is an irregular and low spiking activity per unit of time, known as spontaneous noise (Mountcastle et al., 1969; de Lafuente and Romo, 2006; Vazquez et al., 2013). Therefore, in our cortical simulations, we also included spontaneous firing activity.

Having constructed our ideal observer as described, we tested its performance on three vibrotactile tasks – threshold detection (TD), amplitude discrimination (AD), and frequency discrimination (FD). By separately feeding the ideal observer peripheral and cortical noise-affected stimulus responses, we quantified its performance in all three tasks at two different decoding levels within the somatosensory pathway – an approach known as sequential ideal-observer analysis (Geisler, 1989). This allowed us to determine the best performance an observer can attain in each of the above-mentioned tasks, using neural spike rates representative of either peripheral or cortical neurons.

2.2 Methods

In this section, first we explain the steps for creating our ideal observer model, and next we describe the steps for determining the performance of the ideal observer model on simulated vibrotactile perceptual tasks. Our ideal observer model has two components – a generative component (or the encoding model), which simulates the sensorineural data evoked by vibrotactile stimuli of different amplitudes and frequencies, and a decoding component that optimally decodes the simulated stimulus-evoked neural responses to infer the task relevant parameter(s) of the vibrotactile stimuli. In the following subsection we describe in detail how we generated the stimulus-evoked responses.

2.2.1 Generative model

Because vibrotactile stimuli of mid-frequency range (5 Hz to 50 Hz) primarily activate the RA afferents (Bolanowski et al., 1988; Gescheider et al., 2009; Mountcastle, 2005), we simulated only the RA afferent responses by using the information about the stimulus response functions reported in Johnson (1974), and in Freeman and Johnson (1982). We have skipped the transduction level of encoding, which is at the level of the receptors, and simulated the neural responses that are represented in the RA afferents and cortical neurons.

2.2.1.1 Simulating the afferent response

Each RA afferent branches and terminates in Meissner's corpuscles that are activated when vibrations are presented on the skin surface. Meissner's corpuscles convert the vibratory stimuli into neural impulses, which are relayed by the RA afferents to the central nervous system (Mountcastle, 2005). Studies have shown that the neural activity in the RA afferents is affected by the frequency and amplitude of the vibration stimuli (Talbot et al., 1968; Johnson, 1974; Freeman & Johnson, 1982; Muniak et al., 2007). Talbot et al. observed, using vibrotactile stimuli of 40Hz and a range of stimulus amplitudes (0-200μm), that a linear increase in vibration amplitude did not increase the firing rate linearly in any recorded RA afferent; the responses of each RA afferent, instead, depended critically on certain threshold amplitude values and showed step-like functions (see Fig. 13 in Talbot et al., 1968). Johnson (1974), using stimuli of the same frequency but with a wider range of amplitudes (0-1000μm) compared to the range used by Talbot et al. (1968), observed the same phenomenon. In both studies the investigators computed the stimulus-response function by dividing the total number of spikes they acquired in each stimulus presentation by the total number of stimulus cycles that occurred during those corresponding stimulus durations. See figure 2.1, which shows the stimulus-response function of a simulated RA afferent. Separately, these investigators observed that in each stimulus cycle, i.e. per 25ms for a 40Hz stimulus, on average each RA afferent began to respond with a few spikes per second only when the stimulus

amplitude exceeded a critical threshold amplitude, which Johnson (1974) termed as I_0 (or absolute threshold amplitude); beyond I_0 the number of spikes per cycle increased linearly until the stimulus amplitude exceeded a second critical threshold, which Johnson (1974) termed as I_1 (or entrainment threshold). At stimulus amplitudes above I_1 the afferent responded periodically at 1-spike per stimulus cycle, and any further increase in amplitude did not affect the afferent's response until the amplitude was increased to exceed another critical threshold called I_2 (from Johnson, 1974). When Johnson increased the stimulus amplitude beyond I_2, the number of spikes per cycle increased linearly to 2 until the stimulus amplitude reached the final critical threshold, which Johnson (1974) termed it as I_3 (or the doubling threshold). At stimulus amplitudes exceeding I_3, the afferent entrained at 2-spikes per cycle without showing any effect of change in stimulus amplitude (Fig. 2.1; also see figure 3 in Johnson, 1974). Johnson (1974) formulated this stimulus-response function into a series of equations to predict the stimulus-evoked spike count of an RA afferent (see equation 1).

For sinusoidal vibrations of frequency *f* (cycles per sec) and amplitude *a* (microns), the equations for the mean stimulus-evoked spike count of an RA afferent are (adapted from Johnson, 1974):

$$\mu_f(a) = \begin{cases} 0 & 0 < a < I_0 \\ n(a - I_0)/(I_1 - I_0) & I_0 < a < I_1 \\ n & I_1 < a < I_2 \\ n(1 + (a - I_2)/(I_3 - I_2)) & I_2 < a < I_3 \\ 2n & I_3 < a \end{cases} \qquad \text{(Equation 1)}$$

Here, *n* is the number of sinusoidal cycles, which depends on *f* and the stimulus duration, *t*:

$$n = (t)(f) \qquad \text{(Equation 2)}$$

How to simulate the "critical threshold points": I_0, I_1, I_2, and I_3? Interestingly, Johnson (1974) observed that a) the critical thresholds for each afferent are not random, rather there is a strict pattern where $I_0 < I_1 < I_2 < I_3$, and b) the I_0 and I_1 (and similarly I_2 and I_3) values are not independent of each other. This implies that to assign critical threshold amplitude values to our simulated afferents, we cannot independently sample the I_0 and the I_1 values from their respective population distributions (which is also true for I_2 and I_3).

Johnson (1974) reported that the ratios I_0/I_1 and I_2/I_3 vary between 0 and 1 (i.e. very shallow to very steep slopes), and these ratios are normally distributed with a mean of 0.32 and 0.62, respectively (see Fig. 2.2; also see Fig. 13 in Johnson, 1974). Furthermore, he showed that I_0/I_1 and I_2/I_3 are independent of the entrainment (i.e. I_1) and the doubling (i.e. I_3) thresholds, respectively (Johnson, 1974). Therefore, to determine an I_0-I_1 pair for

each RA afferent, we first sampled I_1 values from the I_1 distribution that we recreated from Johnson (1974) (see Fig. 2.3). Next, we sampled an I_0/I_1 ratio from the normal distribution shown in figure 2.2. Finally, by multiplying the sampled I_1 and I_0/I_1 values, we derived the value of I_0. Using the identical procedure we derived the I_2 and I_3 values for each RA afferent.

Because Johnson (1974) only used 40Hz stimuli, the derived critical thresholds might be true only when the stimulus frequency is 40Hz; what happens when the frequency of the stimulus changes? Note that in equation 1 between I_1 and I_2, as well as beyond I_3 (i.e. the entrainment periods), the afferents entrain at 1 (or 2 in case of $I_3 < a$) spike(s) per stimulus cycle; this suggests that changes in stimulus frequency will correspondingly modify the firing rate of the afferents by the same factor. However, this might pose a problem for the observer. For instance, according to equation 1, a 20Hz stimulus at amplitudes above I_3 will elicit 40 spikes in a second, and a 40Hz stimulus at any amplitude between I_1 and I_2 will also elicit 40 spikes in a second. How can an observer know whether it was a 20Hz or a 40Hz stimulus? Do the critical thresholds vary with frequency?

Stimulus frequency affects critical threshold points. Freeman and Johnson (1982) observed that the absolute threshold (I_0), the entrainment threshold (I_1), and the ratio I_0/I_1 all change with frequency (see Figs. 2.4 and 2.5, which are recreated from figures 2.7 and 2.5, respectively of Freeman and Johnson, 1982). Similar to deriving the I_0-I_1 pairs for the 40Hz vibrations mentioned in the subsection above, for all the frequencies between 5Hz and 50Hz, we first derived I_1 and then the I_0/I_1 ratio using information from Freeman and Johnson (1982), and by combining these values, we derived I_0.

To derive the I_0-I_1 pairs at each relevant frequency (i.e., 5-50Hz) for all the afferents in our model, we followed a 4-step procedure. First, because Freeman and Johnson (1982) only reported the mean values of I_1 at different frequencies and not the standard deviations (see fig. 2.4), we had to determine the factor by which the mean and the standard deviation of the 40Hz I_1 and all the other frequencies between 5 and 50 Hz differed. To calculate this "I_1 translation factor", we divided the mean values of I_1 at all the frequencies between 5 and 50Hz by the mean of 40Hz I_1. Second, we sampled an I_1 value from the 40Hz I_1 distribution (from Johnson, 1974) and depending on the frequency we multiplied it by the corresponding I_1 translation factor. Third, we sampled an I_0/I_1 ratio for that frequency from figure 2.5. As shown in figure 2.5A, Freeman and Johnson (1982) investigated only a small number of frequencies. To sample from intervening frequencies, we interpolated between two I_0/I_1 ratios sampled from figure 2.5B. In the fourth and the final step, we combined the I_1 and I_0/I_1 values to calculate the new I_0 values, as in the final step mentioned in the previous subsection.

Freeman and Johnson (1982) did not investigate the distribution of I_3, or the ratio of I_2/I_3 at different frequencies. Therefore, we did not update the I_2/I_3 or the I_3 distribution in our model; however, to keep the distance (in amplitude points) constant between I_1 to I_2 and

I_1 to I_3, we shifted the mean of each distribution by the same amount (in amplitude) as was predicted for the 40Hz stimulus (Johnson, 1974). It is likely that, with an increase in frequency the amplitude required to reach I_3 threshold might be higher than that at lower frequencies, and the I_2/I_3 slope might be shallower. Nevertheless, this will not pose a problem in the current study because the maximum amplitude tested here is 100μm (i.e., none of the RA afferents should reach the amplitude to entrain at 2-spikes per cycle). Interestingly, Güçlü and Bolanowski (2003a) reported I_2 and I_3 values in cat RA afferents lower than those reported by Johnson (1974) in monkey RA afferents. Furthermore, the authors observed that in cat RA afferents the ratios of I_0/I_1 and I_2/I_3 are similar (Güçlü and Bolanowski, 2003a,b). We have chosen to base our generative model on the available monkey RA data, under the assumption that it is more likely to be representative of human afferents.

Because our sampling processes for I_1 and I_2 were independent from each other, there was a very small chance of obtaining samples where $I_1 > I_2$ (a situation that is biologically impossible). To assess this possibility, we drew 100,000 samples of I_1 and I_2 and confirmed that I_1 was never greater than I_2.

How do the afferents respond that are away from the centre of the probe? Johnson (1974), using a 1mm stimulus-probe radius, studied a population of RA afferents and observed that the intensity of the stimulus attenuates with increasing distance *d*, beyond 2 mm (probe radius + 1mm; see Muniak et al., 2007) from the centre of the stimulus-probe. We calculated the effective amplitude at a distance *d* by implementing equation 3 shown below:

$$a(d) = \begin{cases} a & d \le (\text{probe radius} + 1\text{mm}) \\ \left(\dfrac{\text{probe radius} + 1\text{mm}}{d} \right)^{1.9} \cdot a & d > (\text{probe radius} + 1\text{mm}) \end{cases} \quad \text{(Equation 3)}$$

Interestingly, the critical threshold points (I_0, I_1, I_2, I_3) are characteristics of the afferents and are invariant to the distance from the centre of the stimulus-probe. The afferents respond to the effective stimulus amplitude as if the receptive field hotspot received a stimulus at reduced amplitude, which is similar to the distance-attenuated amplitude. Thus, different receptive fields equidistant from the centre of the probe will have very similar afferent response. Using this information we generated the population response by estimating the firing rate at the effective stimulus amplitude for each RA afferent.

The density of the RA afferents on our fingertip and thenar eminence are 141units/cm^2 and 25units/cm^2, respectively (Johansson and Vallbo, 1979). The afferent density probably varies in the proximodistal axis of the fingertip where the afferents are denser near the distal tip of the finger than near the first interphalangeal crease i.e. the proximal part of the distal phalanx (Güçlü and Bolanowski, 2002). However, in the current study

we made a simplifying assumption that the afferents are regularly distributed in a grid like pattern where each afferent is equidistant from its neighboring afferent by 0.84mm and 2.5mm on simulated fingertip and thenar eminence respectively. For all the simulations we considered a 10mm-by-10mm grid of receptors or cortical neurons, and probe radius of 5.5mm.

Vibrotactile stimulus evoked RA afferent responses. We followed all the steps mentioned in the above subsections and recreated the stimulus-response function at all the frequencies between 5Hz and 50Hz for each afferent that we considered in any experiment. To generate the number of spikes for a stimulus of a certain amplitude and frequency, we first referred to each afferent's frequency relevant stimulus-response curve and depending on the effective amplitude (because effective amplitude changes with distance from the probe: see subsection above) we read out the expected number of spikes per stimulus cycle (see Fig. 2.1).

Equations (1) give the mean stimulus-evoked spike count for an RA afferent stimulated with a vibration of effective amplitude a. The mean spikes per cycle are not necessarily integers. For example, a neuron might be expected to fire 0.6 spikes per cycle. We interpret fractional mean spikes per cycles as probabilities: a neuron expected to yield 0.6 spikes/cycle has a 60% probability of firing a spike on any given cycle. Therefore, we drew binomial samples for each afferent, i, such that the probability, $P(r_i)$, of obtaining r_i stimulus-evoked spikes was:

$$P(r_i) = 0 \qquad \text{when, } 0 < a < I_0 \qquad \text{(Equation 4)}$$

$$P(r_i) = \binom{n}{r_i} \left(\frac{\mu_f(a)_i}{n}\right)^{r_i} \left(1 - \frac{\mu_f(a)_i}{n}\right)^{n-r_i} \qquad \text{when, } I_0 < a < I_1 \qquad \text{(Equation 5)}$$

$$P(r_i) = \begin{cases} 1 & \text{if } r_i = n \\ 0 & \text{if } r_i \neq n \end{cases} \qquad \text{when, } I_1 < a < I_2 \qquad \text{(Equation 6)}$$

$$P(r_i) = n + \binom{n}{r_i} \left(\frac{\mu_f(a)_i}{n} - 1\right)^{r_i} \left(2 - \frac{\mu_f(a)_i}{n}\right)^{n-r_i} \qquad \text{when, } I_2 < a < I_3 \qquad \text{(Equation 7)}$$

$$P(r_i) = \begin{cases} 1 & \text{if } r_i = 2n \\ 0 & \text{if } r_i \neq 2n \end{cases} \qquad \text{when, } I_3 < a \qquad \text{(Equation 8)}$$

To simulate the population response, we repeated this process for every afferent. Note that here we did not implement any convergence of afferents to a higher order neuron (see Güçlü and Bolanowski, 2000). We made the assumption that the N afferents are conditionally independent, given the stimulus, so the probability of the population response $\{r\}$ is simply the product:

$$P(\{r\}) = \prod_{i=1}^{N} P(r_i) \qquad \text{(Equation 9)}$$

2.2.1.2 Modeling cortical response

Compared to primary afferents, cortical somatosensory neurons show much greater response variability, and they show spontaneous activity in the absence of any stimulus. Each of these features reduces the signal-to-noise ratio in the cortex, in comparison to that in the periphery. To investigate the perceptual effects of these characteristics of cortical activity, we transformed each RA afferent into an idealized cortical neuron. For this purpose, we used a one-to-one mapping from peripheral afferent to idealized cortical neuron; we do not attempt here to model the convergence and divergence of inputs from the periphery to the cortex.

The stochastic variability that occurs with the stimulus-evoked responses in the primary somatosensory system can be characterized as a Poisson distribution (Sripati et al. 2006). The unique feature of such a distribution is that the variability is equal to the mean of the distribution. Therefore, to represent cortical neurons, we introduced Poisson noise into our model neurons. Regarding cortical spontaneous activity, Mountcastle and colleagues (1969) observed that the average neuronal firing rate in the absence of any stimulus was 5-20 spikes per second in monkey primary somatosensory cortex. More recently, Romo and colleagues (de Lafuente and Romo, 2005; Vazquez et al., 2013) reported a wider range (0 to 40 spikes per sec) with a median of 10 spikes per second (Vazquez et al., 2013). To understand the effect of spontaneous noise in the perceptual tasks, we parametrically varied the number of spontaneous-spikes considered during the simulations.

To simulate the total expected spike count for each cortical neuron *i*, we summed the expected stimulus-evoked spike count, $\mu_f(a)_i$, with the spontaneous spike count resulting from a spontaneous rate of *s* spikes/sec. We then sampled from a Poisson distribution to obtain the number of spikes during the stimulus period, *t*:

$$P(r_i) = \frac{\exp\left(-\left(\mu_f(a)_i + (s)(t)\right)\right)\left(\mu_f(a)_i + (s)(t)\right)^{r_i}}{r_i!} \qquad \text{(Equation 10)}$$

By applying equation 9, we simulated neural response of a population of cortical neurons.

As a preliminary attempt at creating a generative model, in the current study, we assumed that each afferent feeds directly to a cortical neuron. This approach, albeit simple, is probably generating neural responses with fidelity higher than real stimulus-evoked cortical neural responses in a monkey or a human participant, which means the ideal observer is getting more information than it would have gotten otherwise.

2.2.1.3 Modeling simulated subjects

In the current study we created a population of RA afferents that represents the average density on human fingertip, and each RA afferent's sensitivity thresholds (I_0, I_1, I_2, and I_3) are different, as explained in subsections above (see section 2.2.1.1). Because of the sampling process associated with simulating the RA afferents, it is possible to generate a new set of RA afferent population for each trail of a perceptual experiment; however, in reality the population of RA afferents on human skin does not change during multiple trials or during different experiments. Therefore, we only sampled once per simulated subject and created a population of RA afferents to match the densities of the fingertip and the thenar eminence.

2.2.2 Decoding model

After generating peripheral and cortical population responses to various mid-frequency vibrotactile stimuli, we programmed the ideal observer (i.e. the decoding model) to perform Bayesian inference on the simulated sensorineural data and to perceive the stimulus features necessary to perform the simulated vibrotactile perceptual tasks – vibrotactile threshold detection (TD) task, amplitude discrimination (AD) task, and frequency discrimination (FD).

We implemented Bayesian model comparison for the detection task (TD), and Bayesian parameter estimation for vibrotactile discrimination tasks (AD and FD). Bayesian model comparison is a procedure in which the observer evaluates the probability of the acquired sensorineural data in light of two models (e.g., the stimulus was present or it was absent), and then compares these probabilities to identify the winning model (see subsection 2.2.3.1, TD task for equations). Bayesian parameter estimation is a procedure in which the observer estimates the probability of stimulus parameter values (e.g., the probability of different values of stimulus amplitude in microns) given the sensorineural (see subsection 2.2.3.2, AD and FD tasks for equations). In all simulations the ideal observer knew its own receptive fields and knew the sources of variability, i.e., the ideal observer knew whether the sensorineural data were generated from the RA afferent population or from the cortical neuronal population, and it knew the mean spontaneous firing rate of the cortical neurons.

2.2.2.1 Vibrotactile detection task

In this 2-interval forced choice (2IFC) task, using the method of constant stimuli (MCS), we randomly presented a stimulus in one of two intervals. We started the simulation at the stimulus strength of 1 μm and after every 100 trials we increased the stimulus amplitude by 1 μm. The ideal observer treated each interval separately and performed Bayesian

model comparison, which is quantified as a likelihood ratio of two models (known as Bayes factor):

$$BF = \frac{P(\{r\} \mid M_1)}{P(\{r\} \mid M_0)} \quad \text{(Equation 11)}$$

In the denominator of the ratio, M_0 is the model that hypothesized that the stimulus amplitude was zero, which is equal to the probability calculated in equation 4, (i.e. the condition where $0 < a < I_0$):

$$P(\{r\} \mid M_0) = P(\{r\} \mid a = 0) \quad \text{(Equation 12)}$$

And, in the numerator of the ratio, M_1 is the model that hypothesized that the stimulus amplitude is a nonzero value from 1 to 100 microns:

$$P(\{r\} \mid M_1) = \sum_{a=1\mu m}^{100\mu m} P(\{r\} \mid a) \cdot P(a \mid M_1) \quad \text{(Equation 13)}$$

Thus, the ideal observer evaluated whether the interval under consideration had a stimulus and returned its confidence about that decision as a Bayes factor. Hence, for each trial the ideal observer returned two Bayes factor values and chose the interval with the higher Bayes factor as the interval with the stimulus. Finally we evaluated the ideal observer's decisions to find the number of correct responses at all stimulus amplitudes, and estimated the ideal observer's threshold as shown in equation 16 and 17.

2.2.2.2 Vibrotactile discrimination tasks

The implementation of vibrotactile AD and FD tasks was procedurally similar; we varied the target parameter (*a* for AD task, and *f* for FD task) while keeping the non-target parameter identical in both intervals of a trial in these 2IFC tasks. For each block of 100 trials, we kept the standard stimulus fixed and progressively decreased the value of the comparison stimulus to below that of the standard stimulus. Thus, we progressively increased the difference between the target parameters, making the task progressively easier. We started with a difference of 1 μm for the AD task (1 Hz for FD task) and increased the difference by 1 unit after every 100 trials. As in the detection task, we analyzed the ideal observer's responses to determine its threshold on the discrimination tasks.

Vibrotactile amplitude discrimination task: In the AD tasks we presented two vibrotactile stimuli representing two intervals; in both intervals the *f* of the stimuli was identical, however the *a* was manipulated experimentally. In this 2IFC task, during each interval of a trial the ideal observer performed Bayesian parameter estimation to identify the most likely *a* value given the sensorineural data:

$$P(a \mid \{r\}) \propto P(\{r\} \mid a) \cdot P(a) \quad \text{(Equation 14)}$$

We considered the mode of $P(a|\{r\})$ as the perceived a for each interval. Note that $P(a)$ is a constant (i.e. uniform prior over amplitude); therefore, the mode of $P(a|\{r\})$ is the same as the maximum likelihood estimate. The simulated subjects chose the interval that yielded higher perceived a as the response for each trial.

Vibrotactile frequency discrimination task: In the FD tasks, the stimulus f in one interval was higher (standard stimulus) than in the other (comparison stimulus), but the a was identical in the two intervals. The ideal observer performed Bayesian parameter estimation in each interval to identify the most likely f value given the sensorineural data:

$$P\left(f \mid \{r\}\right) \propto P\left(\{r\} \mid f\right) \cdot P\left(f\right) \qquad \text{(Equation 15)}$$

As in the AD task, we considered the mode of $P(f|\{r\})$ as the perceived frequency; note that $P(f)$ is a constant (i.e. uniform prior over frequency). Hence, for each trial we received two perceived f values from the ideal observer, which we compared and chose the interval with higher "perceived" f as the response in each trial.

2.2.3 Estimation of thresholds

To quantify the performance of the simulated subjects in the detection and discrimination tasks, we fit psychometric functions relating the probability of correct response, $P(c)$, to stimulus level, l (where l could refer to stimulus amplitude, or change in amplitude or frequency, depending on the task). Like Kontsevich and Tyler (1999) we used cumulative normal psychometric functions, parameterized by θ (the threshold parameter) and β (the slope parameter of the psychometric function), such that:

$$P\left(c_l \mid \theta,\beta\right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{1}{\sqrt{2}}\left(\frac{l}{\theta}\right)^{\beta}} \exp\left(\frac{-x^2}{2}\right) dx \qquad \text{(Equation 16)}$$

where, $d' = \left(\frac{l}{\theta}\right)^{\beta}$ such that d' = 1 when $l=\theta$.

We assumed a uniform prior over the parameters θ and β, and found the best-estimate, $\hat{\theta}$, for each subject's threshold as the posterior mode resulting from Bayesian binomial parameter estimation:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \sum_{\beta} \left[\prod_{l} P\left(c_l \mid \theta,\beta\right)^{cr_l} \cdot \left(1 - P\left(c_l \mid \theta,\beta\right)\right)^{100-cr_l} \right] \qquad \text{(Equation 17)}$$

where, cr_l refers to the number of correct responses of the subject at stimulus level l.

2.2.4 Data extraction, simulations and statistics

For the generative portion of the ideal observer analyses, we had to rely on previously published neurophysiological studies; therefore, we extracted relevant data from different studies using GraphClick (2008, Arizona Software). We performed all the simulations using LabVIEW 11 (National Instruments). Any statistics that we have reported were done using SPSS version 20 (IBM).

2.3 Results

Before implementing the tasks, we generated 10 simulated subjects (see section 2.2.1.3 above). Each simulated subject was tested on 5 blocks in each perceptual task; within each block we presented 100 trials per stimulus level. For a given stimulus duration the number of stimulus cycles varies with frequency; therefore, for all the frequency discrimination tasks we kept the number of stimulus cycles constant.

2.3.1 Effect of neural noise: afferent variability, cortical variability, and spontaneous noise

To investigate the effect of neural noise on decoding efficiency, we conducted detection (TD) and discrimination (AD and FD) tasks where the ideal observer separately decoded RA noise and cortical noise (with and without spontaneous noise) affected sensory data. In this current set of experiments, we set the number of spontaneous spikes to zero.

In the TD task, the neural noise (RA1 versus cortical) did not affect the performance of the ideal observer ($t = –1.633$, $p = 0.137$) (see Fig. 2.6A). In the discrimination tasks, an ANOVA with neural noise and task as within subject factors revealed a main effect of task [$F(1,9) = 31.28$; $p < 0.001$] and noise [$F(1,9) = 5.51$; $p = 0.04$] (see Fig. 2.6B). Although the mean discrimination thresholds were slightly higher for the cortical noise affected data, pairwise comparisons with Bonferroni corrections revealed that within each discrimination task, there was no significant effect of the level of sensory noise (p=0.05, and p=0.14, for AD and FD tasks, respectively).

Next, to test the influence of spontaneous noise in cortical variability and how that affects inference, we added spontaneous spikes to the stimulus-evoked neural responses during the cortical encoding of simulated vibrotactile stimuli (see section 2.2.1.2). Addition of spontaneous spikes worsened the performance of the ideal observer in all the tasks (see Fig. 2.7 and 2.8).

An ANOVA with spontaneous rate (0, 10, 20, 30 spikes/s) as the only variable revealed its main effect in the TD task [$F(3,27) = 308.17$; $p < 0.001$], and pairwise comparisons with Bonferroni corrections revealed that the mean thresholds at each spontaneous rate were significantly different ($p < 0.01$) (see Fig. 2.7). In discrimination tasks, an ANOVA with two within subject factors – spontaneous rate (0, 10, 20, 30 spikes/s) and task (AD, FD), revealed a main effect of both factors (task: [$F(3,27) = 1562.63$; $p < 0.001$], spontaneous rate: [$F(3,27) = 3466.37$; $p < 0.001$]) (Fig. 2.8). The addition of spontaneous spikes particularly affected the ideal observer's performance in the FD task (Fig. 2.8B). Pairwise comparisons with Bonferroni corrections revealed that within each task an increase in spontaneous rate increases the threshold to discriminate between two vibrations ($p < 0.05$ for all comparison except for FD, 20 spikes/s vs 30 spikes/s where p=0.07).

2.3.2 Effect of receptor density and stimulus duration

To investigate whether receptor density affects the ideal observer's performance on detection and discrimination tasks we estimated the ideal observer's performance at two receptor densities that match the density in human fingertip (141 RA units/cm^2) and thenar eminence (25 RA units/cm^2). For all cortical level simulations we added the median spontaneous rate (i.e. 10 spikes/sec, see Vazquez et al., 2013) to the stimulus-evoked neural responses during the encoding stage. In all three tasks, irrespective of the decoding level (afferent or cortical), receptor density affected the performance of the ideal observer (see Fig. 2.9 and 2.10).

Next, we investigated whether stimulus duration (i.e. number of stimulus cycles) affect the ideal observer's performance. We tested the model only on the fingertip receptor density, and for all cortical level simulations we added the median spontaneous rate (i.e. 10 spikes/sec, see Vazquez et al., 2013) to the stimulus-evoked neural responses during the encoding stage. In all tasks, it is clearly noticeable that the performance of the ideal observer improves with increasing stimulus cycles (see Fig. 2.11 and 2.12). Interestingly, with increasing stimulus cycles, the performance of the cortical ideal observer in the AD task approached the RA performance level.

2.3.3 Effect of stimulus frequency on TD task

Irrespective of the stimulus frequencies, 20Hz or 40Hz, the ideal observer's threshold decreased with increasing stimulus duration (see Fig. 2.13). Also, the ideal observer's detection thresholds at 40Hz vibrations were lower than at 20Hz vibrations.

We performed two ANOVAs, one for each decoding levels, with number of taps (1, 2, 3), and stimulus frequency as within subject factors. The ANOVA on the afferent level

decoding thresholds revealed a significant main effect of the number of taps [$F(2,18) = 211.66$; $p < 0.001$] and a marginal main effect of stimulus frequency [$F(1,9) = 4.47$; $p = 0.06$]. Pairwise comparisons with Bonferroni corrections revealed that at each tap number condition the 40Hz stimuli had a threshold that was significantly lower than the 20Hz stimuli ($p < 0.05$), except when the stimulus contained a single cycle (p=0.21) (see Fig. 2.13A). The ANOVA on the cortical level decoding thresholds revealed a significant main effect of the number of taps [$F(2,18) = 227.39$; $p < 0.001$] and stimulus frequency [$F(1,9) = 7.72$; $p < 0.05$]. Like afferent level decoding threshold differences, on the cortical level decoding threshold pairwise comparisons with Bonferroni corrections revealed that at each tap number condition the 40Hz stimuli had a threshold that was significantly lower than the 20Hz stimuli except when the stimulus contained a single cycle (1 tap: p=0.1; 2 taps: p=0.01; 3 taps: p=0.03) (see Fig. 2.13B).

2.3.4 Effect of standard stimulus in discrimination tasks (effect of target and non-target features)

To investigate the effect of changing the target feature (i.e. amplitude in AD, and frequency in FD tasks) and the non-target feature (i.e. frequency in AD, and amplitude in FD tasks) in discrimination tasks, we ran a series of simulations. We observed that the difference required to reliably discriminate the target feature in two vibrations grows with increase in the value of the target feature (see Fig. 2.14A and C). For example, the ideal observer required a bigger difference in amplitude between the standard and the comparison vibrations in the AD task when we increased the amplitude of the standard stimulus from 20µm to 40µm. Interestingly, the ideal observer's performance at discriminating the target feature was affected very little by changing the non-target feature in different blocks. For instance, we estimated the ideal observer's AD threshold at 20Hz and at 40Hz vibrations, and observed no change in its AD threshold (see Fig. 2.14B). However, increasing the amplitude from 20µm to 40µm in the FD tasks slightly worsened the performance of the ideal observer (see Fig. 2.14D).

Finally, to explore whether humans can actually infer the stimulus features during vibrotactile discrimination task, we ran simulations with identical combination of amplitude and frequency. We reasoned that if humans are melding the features then two tasks should reveal identical performances; however, because the ideal observer optimally infers the features, similarity between the trends in human performance and that of the ideal observer will provide evidence whether humans are inferring the features of vibrotactile stimuli. The simulation results show there the discrimination thresholds for AD and FD tasks should be different.

2.4 Discussion

In the current study, we implemented a simple generative model based on binomial probabilities to simulate the total number of spikes evoked by vibrotactile stimuli of different amplitudes and frequencies. Whereas a sophisticated generative model has been proposed by Güçlü and Bolanowski (2004), which would have allowed us to create spike trains, the motivation of the current study was to explore the information content carried by a simple rate code, in order to determine whether this code is sufficient to perform different vibrotactile perceptual tasks.

Thus we created a Bayesian ideal observer model to estimate the upper bounds of the performance efficiency that any observer can achieve from decoding spike rates on three different vibrotactile tasks. In each task, the ideal observer optimally inferred the stimulus parameters from the simulated neural responses evoked in the RA afferents and in the cortical neurons. Note that, in the current study our Bayesian ideal observer considered uniform prior probability distributions over stimulus amplitudes and frequencies, which means that its percepts were the same as those that would result from maximum likelihood estimation. In other words, simple decision rules based on signal detection theory would have generated the same results that we obtained by implementing Bayesian estimation. However, the benefits of using Bayesian estimation include systematic examination of the information that the perceptual observer might have while performing vibrotactile tasks, and of the effects of assumptions made by the observer. For example, during a 2IFC FD task human participants might not know the exact values of the amplitudes in the two intervals, but they might assume that the amplitudes are identical. Implementation of such scenarios is only possible by using Bayesian estimation where the ideal observer could effectively marginalize over the amplitude feature to determine the frequency of each interval. Our Bayesian observer can readily be extended to allow for these interesting decoding scenarios.

Effect of neural noise – the quality of sensorineural data affects perceptual inference: The fidelity of the stimulus-evoked neural responses encoded by the nervous system affects an observer's performance on any perceptual task. For example, the ideal observer performed better at the afferent than at the cortical level. We used the identical approach to test the ideal observer's performance at both levels, and by definition the ideal observer takes all possible factors into consideration to achieve the highest level of accuracy. Therefore, the result of the ideal observer analysis demonstrated that the quality of the sensorineural data is better at the somatosensory afferent level than at the cortical level, presumably due to less response variability in the RA afferents than in the primary somatosensory areas (Vega-Bermudez and Johnson, 1999; Sripati et al., 2006), which results in a loss of information from the afferent to the cortical level. The spontaneous neural activity in the cortex probably adds to the noisiness of the cortical sensorineural data and contaminates the sensory information. Interestingly, when we restricted the contribution of the spontaneous activity to zero, the ideal observer's performance at the

afferent and cortical level was comparable in all the tasks (see Fig. 2.6). This suggests that the cortical response variability is less influenced by the Poisson noise alone, and is more affected by a combined process of spontaneous spiking activity and Poisson noise. In a Poisson process the variability of a sampled value depends on the mean or the expected value of the distribution, i.e. if the expected value is high then the variance of the sampled values will be correspondingly high. This dependence of the sampled value on the expected value predicts the apparent reason for a detrimental effect of spontaneous noise. For example, even if the stimulus-evoked expected number of spikes is relatively small, adding a large number of spontaneous spikes will increase the new-expected number of spikes or the mean of the Poisson distribution, which introduces a bigger variability in the sampled number of spikes compared to a condition where the number of spontaneous noise-generated spikes is zero.

Effect of receptor density and stimulus duration – the quantity of sensorineural data affects perceptual inference: Goldreich (2007), using a Bayesian model, demonstrated that the accuracy in perceptual inference depends on the density of tactile sensory receptors. Interestingly, Goldreich and colleagues have provided evidence in different studies that tactile spatial acuity is constrained by finger size (Peters et al., 2009; Wong et al., 2013). Assuming that the density of slowly adapting afferents that carry the information about tactile spatial details (Johnson et al., 2000) decreases with increasing finger size, these studies support the idea that the number of afferents activated affects the quantity of sensorineural data available to the perceptual observer, which in turn influences the accuracy of any perceptual inference. Similarly, in the current study, a decrease in afferent density increased both detection and discrimination thresholds, i.e. worsened the performance of the ideal observer (see Fig. 2.9 and 2.10).

Similar to the effect of receptor density, we hypothesized that changes in stimulus duration would affect the quantity of sensorineural data and consequently the accuracy of inference of the ideal observer. Our simulation results support the hypothesis that stimulus duration affects the performance of the ideal observer in detection and discrimination tasks: the ideal observer performed better on long than on short duration stimuli (see Fig. 2.11, 2.12, and 2.13). Therefore, we predict that in a behavioural study, by changing the quantity of sensorineural data provided to humans, we will be able to affect their performance on tactile tasks (see Chapter 3).

Effect of stimulus features in detection and discrimination tasks: The neural activity at the afferent level strictly depends on the critical amplitudes (I_0, I_1, I_2, and I_3), and it changes with frequency. Therefore, we predicted that changes in frequency in TD and AD tasks, and amplitude in FD tasks, would affect the performance of the ideal observer. In the TD tasks we assume that the threshold values rely on the afferents' I_0 values, which decreases with increasing frequency (Freeman and Johnson, 1982). However, when we compared the ideal observer's detection ability at 20Hz to that at 40Hz, we observed that the detection threshold was statistically similar at stimulus durations that contained a single tap. Interestingly, as we increased the number of taps, the difference between the

detection threshold at these two frequencies became statistically significant, which suggests that probably the signal-to-noise ratio increases with the increase in the number of taps.

We observed that the mean difference in amplitude (in the AD task) or frequency (in the FD task) increases with the increase in the target feature value, which might be due to the Poisson variability. Interestingly, there was very little effect of manipulation of the non-target feature in AD and FD tasks, which reflects that the ideal observer optimally infers the target feature before performing the discrimination task.

Encoding of vibrotactile stimulus features: The two key features of any vibrotactile stimuli are amplitude and frequency; however, it is unclear whether humans can perceive these features separately to perform corresponding discrimination tasks, i.e. vibrotactile AD and FD tasks. Recently, Harris et al. (2006) suggested that during the encoding of any vibrotactile stimulus, humans are unable to perceive the amplitude and the frequency features separately. According to the authors, humans meld these features together into a single feature that can be characterized as the product of the amplitude and the frequency, which is similar to the energy of the vibrotactile stimulus. The current study shows, however, that an ideal observer is capable of independently discriminating vibrotactile frequencies and amplitudes. Because the ideal observer analyses reflect the information content in the generative or the encoding stages, should there be qualitative similarity in performance between the ideal observer and the human participant's performance, we can conclude that humans have access to both amplitude and frequency information. This interpretation is critical for our understanding of vibrotactile perception (see Chapter 4).

2.4.1 What is the neural representation of vibrotactile features?

Studies suggest that there are two plausible neural codes that the brain might utilize to encode vibrotactile stimulus features: a) the rate code, i.e. the total number of spikes per certain duration, and b) the temporal code, or the inter-spike interval – the duration between the spikes that are phase-locked to the periodicity of the stimulus (Romo and Salinas, 2003; Romo et al., 2012). Whereas there is a unanimous agreement that stimulus amplitude is encoded in a rate code (Bensmaïa, 2008; Tommerdahl et al., 2010), the literature is contradictory on the neural code for stimulus frequency. Some studies have argued that frequency is encoded in a rate code (Luna et al., 2005); however, others have argued in favour of a temporal code (Mountcastle et al., 1969, 1990; Mackevicius et al., 2013).

Mountcastle and colleagues have reported that the periodicity in spike trains elicited in RA afferents (Talbot et al., 1968) and in the cortex (LaMotte and Mountcastle, 1975; Mountcastle et al., 1969) of monkeys phase locks very precisely to the stimulus frequency, providing a temporal code for frequency discrimination tasks. However, Romo and colleagues showed through a series of studies that a rate code was sufficient for

monkeys to perform frequency discrimination tasks (Hernandez et al., 2000; Salinas et al., 2000; Romo and Salinas, 2003). To further demonstrate that periodicity is not a necessary condition to perform frequency discrimination, Romo and colleagues (Salinas et al., 2000) jittered the inter-pulse interval of sinusoidal stimuli, which eliminated the timing information present in the stimulus. The authors reasoned that if periodicity is essential for frequency discrimination, then scrambling the periodicity by varying the inter-pulse interval of the stimulus would disrupt the monkeys' ability to do the task. However, in spite of a lack of temporal regularity, the experimental monkeys discriminated the aperiodic stimuli just as well as they discriminated the periodic stimuli. Thus, Romo and colleagues concluded that there is sufficient information in a rate code to discriminate between the frequencies of two vibrotactile stimuli. But interestingly, Harris et al. (2006) tested human participants' ability to discriminate aperiodic stimuli, and the authors reported that participants' performance in this task was worse than for periodic stimuli. This led Harris et al. (2006) to conclude that humans must be utilizing temporal information to discriminate the frequency of different vibrations. Note that in Harris et al. (2006), the human participants were capable of discriminating two aperiodic stimuli, albeit only with a bigger difference between the stimuli compared to the periodic stimuli. Therefore, this suggests that humans probably have the capability to access either the rate code or the temporal code, depending on the difficulty of the task. Interestingly, Hernandez et al (2000) recorded neuronal activity from area 3b and 1 (primary somatosensory areas) while the monkeys were performing a frequency discrimination task; the authors observed that out of 188 RA neurons that responded to stimulus frequencies, 139 showed periodicity as the neural code, 72 modulated their firing rate with changing stimulus frequency, and 23 neurons responded to both periodicity and firing rate codes.

In the current study, the ideal observer did not have any information about the stimulus inter-pulse interval or inter-spike durations; its inferences about stimulus frequency were strictly based on a rate code. Thus we conclude that there is sufficient information in a rate code for a perceptual observer to perform vibrotactile frequency discrimination tasks, along with amplitude discrimination and threshold detection tasks.

2.4.2 Assumptions

Although ideal observers are useful theoretical tools to examine any perceptual system's decoding efficiency, finding and applying the right generative model is a strong limitation in this approach. In the current study, neurophysiological evidence from Johnson (1974) and from Freeman and Johnson (1982) provided us invaluable information to create an informative generative model for the afferent system; however, there is a need for studies in the monkey RA afferent system like the ones conducted on the cat RA afferent system (see Güçlü and Bolanowski, 2003a,b). For example, in the current study we assumed that the distribution of I_0/I_1 in all frequencies between 5Hz to 50Hz is identical to that at 40Hz

reported in Johnson (1974), which might not be true as shown by Güçlü and Bolanowski (2003b) in cat RA afferent recordings.

Next, due to lack of complete information about a cortical generative model, we had to extend our afferent model to a cortical model with Poisson process. RA afferent responses have the characteristic rise (I_0 to I_1, I_2 to I_3) and plateau (I_1 to I_2, and I_3 onwards) of responses (see Fig. 2.1); however, in cortical neurons the spiking activity monotonically increases with increasing stimulus amplitude (Vazquez et al., 2013) before saturating at a certain firing rate (Mountcastle et al., 1969). We acknowledge that the simulated cortical representation of the vibrotactile stimulus is unrealistic; however, it informs us about the effect of Poisson variability in the cortex, and the variability introduced by the spontaneous spikes. Because the architecture of the RA and cortical populations are identical in our model, a comparison of the ideal observer's performance at these two levels interestingly reveals how much information is lost from the periphery to the cortex due to Poisson noise and spontaneous activity alone.

In the current model, we assumed that the responses of different neurons are conditionally independent given the stimulus, which is very likely untrue for cortical neurons (Zohary et al 1994); however, it is a common assumption that is made for maintaining the tractability and simplicity of models (Jazayeri & Movshon, 2006; Gold & Shadlen, 2007). For example, within vibrotactile studies, to characterize the behaviour of thalamic and primary somatosensory neurons during vibrotactile detection tasks, Romo and colleagues simulated hypothetical neurons and assumed that the neuronal responses are conditionally independent (Vazquez et al paper 2013), even though Romo et al. (2003) reported mean noise correlation coefficients of 0.16 in primary somatosensory cortical neurons (de Lafuente and Romo, 2005).

In the current study, we have assumed that tactile perceptual decisions are based on the activity of primary somatosensory cortical neurons. This assumption, though probably unrealistic, does provide us with valuable knowledge about the amount of information available in S1 that might be sufficient to perform different perceptual tasks. Romo and colleagues found that though S1 neurons respond consistently to changes in stimulus features, this neural activity does not co-vary with the monkeys' trial-to-trial behavioural responses. Instead, the authors reported that the neural activity pattern of the prefrontal cortical areas more strongly predict the monkeys' behavioural response patterns (Herenandez et al., 2002; de Lafuente and Romo, 2006, Romo et al., 2004, 2012). However, we do not have sufficient neurophysiological information to model the cortical areas downstream of S1. Therefore, we assumed that the ideal observer makes decisions based on the activity in the primary somatosensory cortical areas.

2.4.3 Predictions and future experiments

The primary objective of creating this ideal observer model was not to propose a model of the brain but to generate testable predictions. We have tested some of these predictions in chapters 3 and 4; the remaining predictions suggest possible future experiments. Within the modeling paradigm, future experiments should include controlled sub-optimization of the model and comparing the obtained results to human performances; this will enable us to identify how information is lost or transformed to generate behaviour related to vibrotactile perception. For example, in the current study the ideal observer had knowledge of its average spontaneous noise rate; however, we can imagine a suboptimal observer that does not have this information.

The simulation results in the current study demonstrated that the ideal observer's performance improved with increasing amounts of information; therefore, we predict that human participants' performance on vibrotactile detection and discrimination tasks should improve with increasing stimulus duration (Fig. 2.11, 2.12, and 2.13) – a phenomenon known as temporal summation (Gescheider et al., 2009). Because the RA afferent density is higher in the fingertip than in the thenar eminence, we also predict that the detection threshold estimated at the fingertip will be lower than that at the thenar eminence – a phenomenon known as spatial summation (Gescheider et al., 2009). Similar to detection thresholds, simulation results predict that the participants' discrimination threshold must change with body site (see Fig. 2.9 and 2.10).

The AD and FD simulation results, using the identical combination of amplitude and frequency, predict that human participants should have different AD and FD threshold values even for identical amplitude and frequency combinations (see Fig. 2.15).

2.5 Conclusion

In conclusion, we have shown that the stimulus-evoked spiking rate code has sufficient information to allow an ideal observer to perform three different vibrotactile tasks. By utilizing the strength of a Bayesian ideal observer model, i.e. the ability to perform optimal decoding, we have quantified the best possible performance achievable by a rate-code-based perceptual observer. Our simulation results suggest that human performance in vibrotactile tasks should be affected by receptor density and stimulus duration; these predictions are tested empirically in the following chapter.

2.6 References

Bensmaïa S (2002) A transduction model of the Meissner corpuscle. Math Biosci 176:203-17.

Bensmaïa SJ (2008) Tactile intensity and population codes. Behav Brain Res 190:165-73.

Bolanowski SJ Jr, Gescheider GA, Verrillo RT, Checkosky CM (1988) Four channels mediate the mechanical aspects of touch. J Acoust Soc Am 84(5):1680-94.

Brody CD, Hernández A, Zainos A, Romo R (2003) Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. Cereb Cortex 13:1196-207.

de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. Nat Neurosci. 8(12):1698-703.

de Lafuente V, Romo R (2006) Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc Natl Acad Sci U S A 103:14266-71.

Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-33.

Freeman AW, Johnson KO (1982) A model accounting for effects of vibratory amplitude on responses of cutaneous mechanoreceptors in macaque monkey. J Physiol 323:43-64.

Geisler WS (1989) Sequential ideal-observer analysis of visual discriminations. Psychol Rev 96:267-314.

Geisler, W. S. (2003) Ideal Observer analysis. In: L. Chalupa and J. Werner (Eds.), The Visual Neurosciences. Boston: MIT press, 825-837.

Geisler WS, Kersten D (2002) Illusions, perception and Bayes. Nat Neurosci 5:508-10.

Gescheider GA, Bolanowski SJ, Verrillo RT (2004) Some characteristics of tactile channels. Behav Brain Res 148(1-2):35-40.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30:535-74.

Goldreich D (2007) A Bayesian perceptual model replicates the cutaneous rabbit and other tactile spatiotemporal illusions. PLoS One 2:e333.

Goldreich D (2009) An ideal observer for passive tactile spatial perception. Front. Syst. Neurosci. Conference Abstract: Computational and systems neuroscience 2009. doi: 10.3389/conf.neuro.06.2009.03.058.

Goldreich D, Tong J (2013) Prediction, postdiction, and perceptual length contraction: a bayesian low-speed prior captures the cutaneous rabbit and related illusions. Front Psychol. 4:221.

Green DM, Swets JA (1966) Signal detection and psychophysics. New York: Wiley.

Güçlü B, Bolanowski SJ (2002) Modeling population responses of rapidly-adapting mechanoreceptive fibers. J Comput Neurosci 12:201-18.

Güçlü B, Bolanowski SJ (2003a) Distribution of the intensity-characteristic parameters of cat rapidly adapting mechanoreceptive fibers. Somatosens Mot Res 20:149-55.

Güçlü B, Bolanowski SJ (2003b) Frequency responses of cat rapidly adapting mechanoreceptive fibers. Somatosens Mot Res 20:249-63.

Harris JA, Arabzadeh E, Fairhall AL, Benito C, Diamond ME (2006) Factors affecting frequency discrimination of vibrotactile stimuli: implications for cortical encoding. PLoS One 1:e100.

Hernández A, Zainos A, Romo R (2000) Neuronal correlates of sensory discrimination in the somatosensory cortex. Proc Natl Acad Sci U S A 97:6191-6.

Hernández A, Zainos A, Romo R (2002) Temporal evolution of a decision-making process in medial premotor cortex. Neuron 33:959-72.

Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. Nat Neurosci. 9(5):690-6.

Johansson RS, Birznieks I (2004) First spikes in ensembles of human tactile afferents code complex spatial fingertip events. Nat Neurosci 7(2):170-7.

Johansson RS, Vallbo AB (1979) Detection of tactile stimuli. Thresholds of afferent units related to psychophysical thresholds in the human hand. J Physiol 297(0):405-22.

Johnson KO (1974) Reconstruction of population response to a vibratory stimulus in quickly adapting mechanoreceptive afferent fiber population innervating glabrous skin of the monkey. J Neurophysiol 37:48-72.

Johnson KO, Yoshioka T, Vega-Bermudez F (2000) Tactile functions of mechanoreceptive afferents innervating the hand. J Clin Neurophysiol 17(6):539-58.

Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci 27:712-9.

Knill DC, Saunders JA (2003) Do humans optimally integrate stereo and texture information for judgments of surface slant? Vision Res 43:2539-58.

Körding K (2007) Decision theory: what "should" the nervous system do? Science 318:606-10.

Kontsevich LL, Tyler CW (1999) Bayesian adaptive estimation of psychometric slope and threshold. Vision Res 39:2729-37.

LaMotte RH, Mountcastle VB (1975) Capacities of humans and monkeys to discriminate vibratory stimuli of different frequency and amplitude: a correlation between neural events and psychological measurements. J Neurophysiol 38(3):539-59.

Luna R, Hernández A, Brody CD, Romo R (2005) Neural codes for perceptual discrimination in primary somatosensory cortex. Nat Neurosci 8(9):1210-9.

Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9:1432-8.

Mackevicius EL, Best MD, Saal HP, Bensmaïa SJ (2012) Millisecond precision spike timing shapes tactile perception. J Neurosci 32:15309-17.

Mountcastle VB (2005) The sensory hand: neural mechanisms of somatic sensation. Harvard University Press.

Mountcastle VB, LaMotte RH, Carli G (1972) Detection thresholds for stimuli in humans and monkeys: comparison with threshold events in mechanoreceptive afferent nerve fibers innervating the monkey hand. J Neurophysiol 35:122-36.

Mountcastle VB, Steinmetz MA, Romo R (1990) Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. J Neurosci 10(9):3032-44.

Mountcastle VB, Talbot WH, Sakata H, Hyvärinen J (1969) Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. J Neurophysiol 32(3):452-84.

Muniak MA, Ray S, Hsiao SS, Dammann JF, Bensmaïa SJ (2007) The neural coding of stimulus intensity: linking the population response of mechanoreceptive afferents with psychophysical behavior. J Neurosci 27:11687-99.

Peters RM, Hackeman E, Goldreich D (2009) Diminutive digits discern delicate details: fingertip size and the sex difference in tactile spatial acuity. J Neurosci. 29(50):15756-61.

Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. Nat Rev Neurosci. 1(2):125-32.

Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. Nature 399:470-3.

Romo R, Hernández A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. Neuron 41:165-73.

Romo R, Hernández A, Zainos A, Salinas E (2003) Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. Neuron. 38(4):649-57.

Romo R, Lemus L, de Lafuente V (2012) Sense, memory, and decision-making in the somatosensory cortical network. Curr Opin Neurobiol. 22:914-9.

Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. Nat Rev Neurosci 4:203-18.

Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. J Neurosci 20:5503-15.

Sripati AP, Yoshioka T, Denchev P, Hsiao SS, Johnson KO (2006) Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system. J Neurosci 26:2101-14.

Talbot WH, Darian-Smith I, Kornhuber HH, Mountcastle VB (1968) The sense of flutter-vibration: comparison of the human capacity with response patterns of mechanoreceptive afferents from the monkey hand. J Neurophysiol 31(2):301-34.

Tommerdahl M, Favorov OV, Whitsel BL (2010) Dynamic representations of the somatosensory cortex. Neurosci Biobehav Rev 34:160-70.

Vazquez Y, Salinas E, Romo R (2013) Transformation of the neural code for tactile detection from thalamus to cortex. Proc Natl Acad Sci U S A. 110(28):E2635-44.

Vega-Bermudez F, Johnson KO (1999) SA1 and RA receptive fields, response variability, and population responses mapped with a probe array. J Neurophysiol 81:2701-10.

Verrillo RT, Bolanowski SJ (1986) The effects of skin temperature on the psychophysical responses to vibration on glabrous and hairy skin. J Acoust Soc Am 80(2):528-32.

Wong M, Peters RM, Goldreich D (2013) A physical constraint on perceptual learning: tactile spatial acuity improves with training to a limit set by finger size. J Neurosci. 33(22):9345-52.

Wozny DR, Beierholm UR, Shams L (2008) Human trimodal perception follows optimal statistical inference. J Vis 8:24,1-11.

Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. Nature 447:1075-80.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature. 370(6485):140-3.

2.7 Figures and figure captions

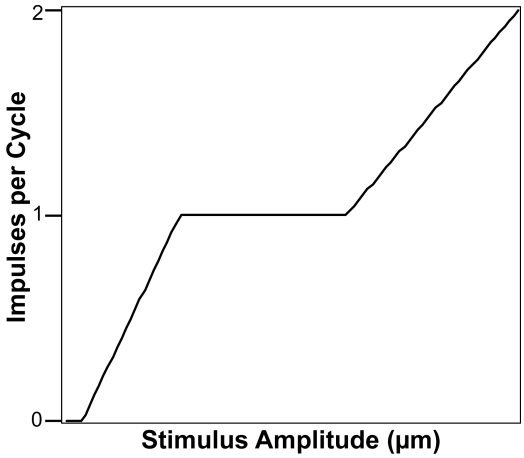

Figure 2.1: Typical stimulus response function of an RA afferent.

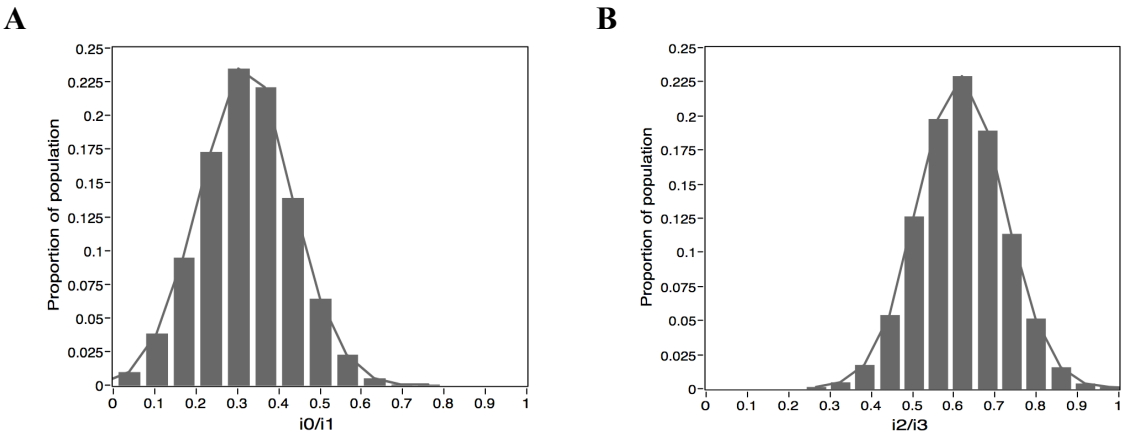

Figure 2.2: Distribution of ratios A) I_0/I_1 and B) I_2/I_3

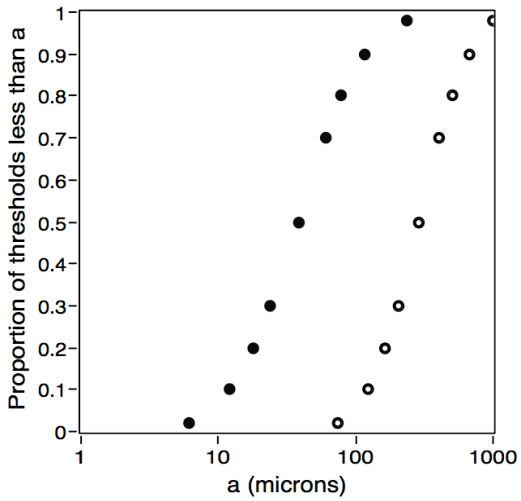

Figure 2.3: Cumulative distribution of I_1 and I_3, recreated from Johnson, 1974 figure 11.

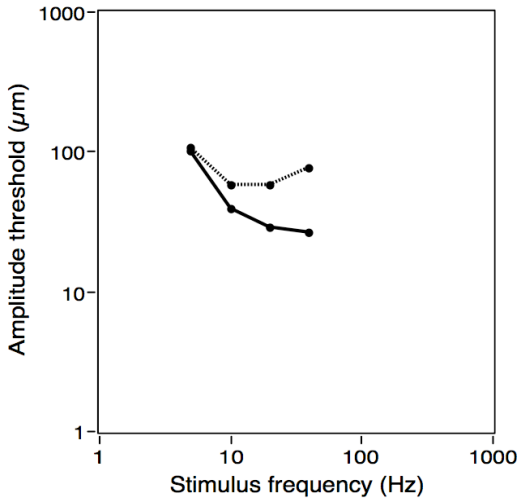

Figure 2.4: Absolute (I_0, solid curve) and entrainment (I_1, dashed curve) thresholds as a function of increasing frequency.

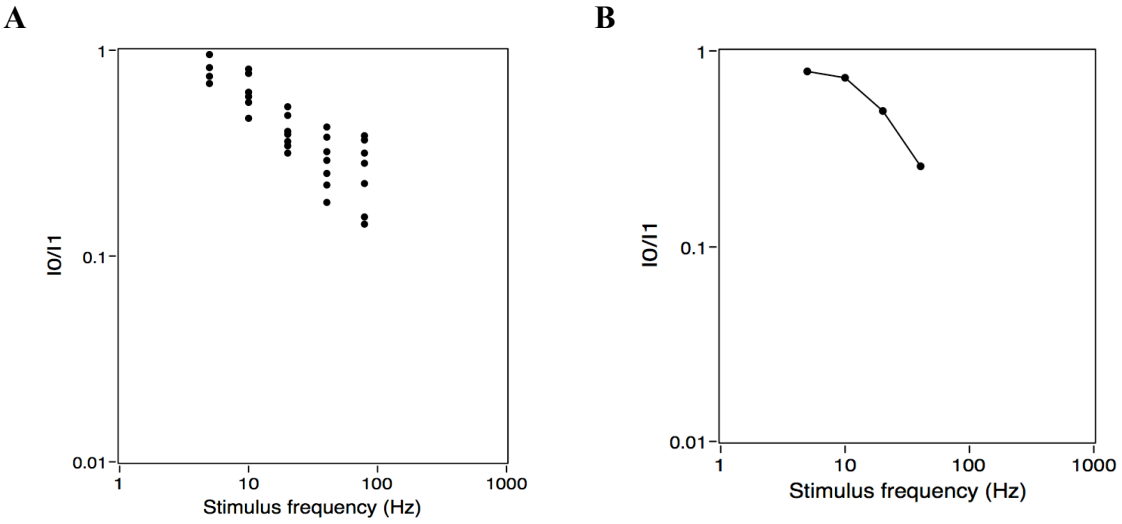

Figure 2.5: A) Ratio of I_0/I_1 as a function of stimulus frequency (Hz). B) Recreation of figure 5 Freeman and Johnson, 1982.

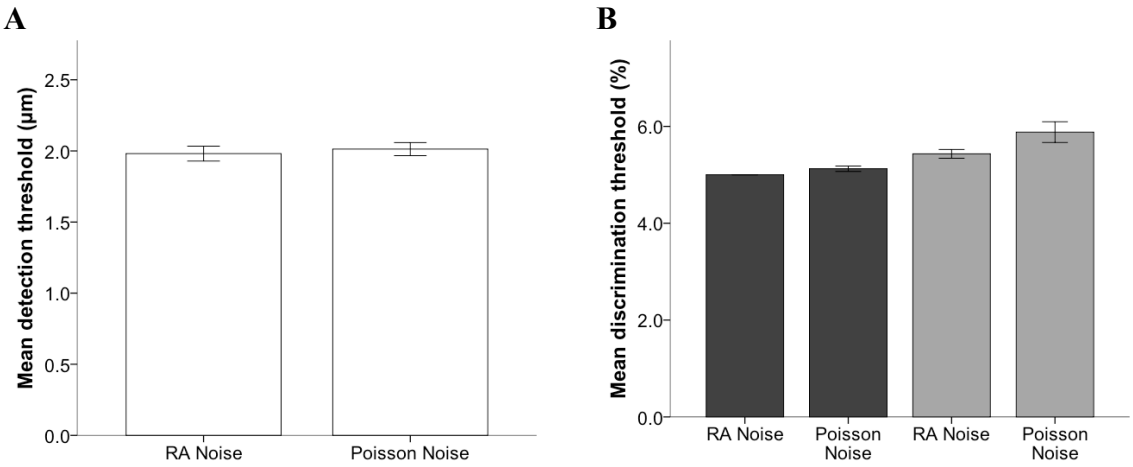

Figure 2.6: An average performance of 10 simulated subjects in TD, AD, and FD tasks, where the sensory data were affected by RA and Cortical noise. Because we set the spontaneous rate to 0 spikes/s, the cortical sensory data is only affected by the Poisson variability. The amplitude refers to peak-to-peak values. The error bars represent ±1SE. A) The simulated subjects detected a 100ms stimulus at 20Hz in a 2IFC task. B) The black and grey bars represent the simulated subject's performance in AD and FD task, respectively. In both tasks the standard stimulus was a 20Hz 20µm vibration, and the bars represent the percent difference in amplitude (for AD task) and frequency (for FD task) of the comparison stimulus that the simulated subjects required to reliably discriminate the vibrations in the 2IFC tasks.

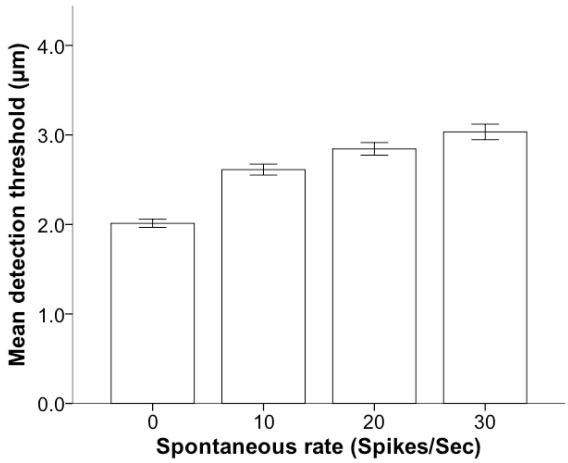

Figure 2.7: An average performance of 10 simulated subjects in TD task. The number of spontaneous spikes added to the stimulus-evoked cortical neural responses depended on the stimulus duration that was 100ms. Each bar represents the average threshold required by the simulated subjects to detect a 20Hz stimulus in a 2IFC task at four different spontaneous rates (spikes/sec). The amplitude refers to peak-to-peak values. The error bars represent ±1SE.

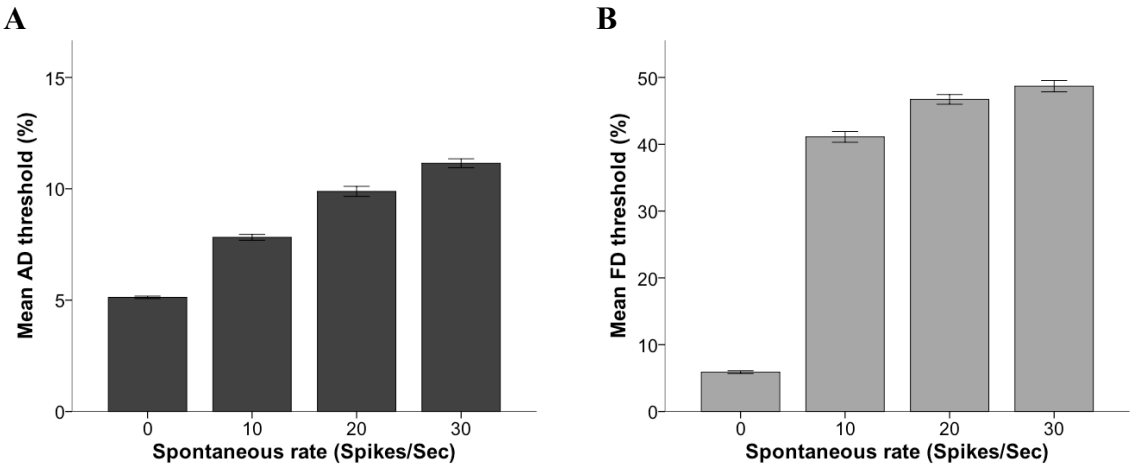

Figure 2.8: An average performance of 10 simulated subjects in AD (panel A) and FD (panel B) tasks. The number of spontaneous spikes added to the stimulus-evoked cortical neural responses depended on the stimulus duration. The amplitude refers to peak-to-peak values. In AD and FD tasks the standard stimulus was a 20Hz 20µm vibration. The bars (black for AD, and grey for FD) represent the percent difference in amplitude (for AD task) and frequency (for FD task) of the comparison stimulus that the simulated subjects required to reliably discriminate the vibrations in the 2IFC tasks at four different spontaneous rates (spikes/sec). The error bars represent ±1SE.

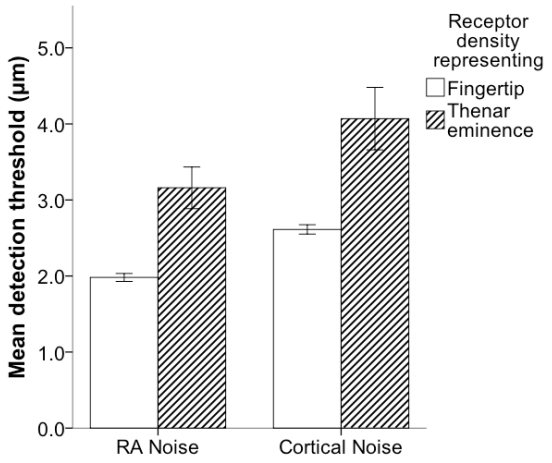

Figure 2.9: Each bar represents the average performance of 10 simulated subjects in TD task (at 20Hz stimulus frequency) conducted at afferent densities that represent the human fingertip (white bar) and thenar eminence (hatched bar). The spontaneous rate was 10 spikes/s for cortical noise simulations, which was equal to the stimulus duration of 100ms. The amplitude refers to peak-to-peak values. The error bars represent ±1SE.

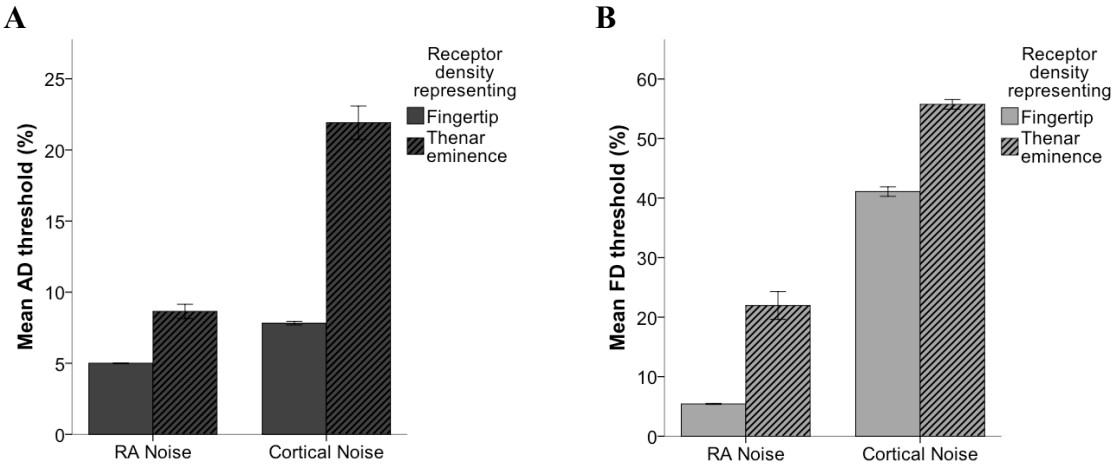

Figure 2.10: An average performance of 10 simulated subjects in AD (panel A) and FD (panel B) tasks conducted at fingertip (clear bars) and thenar eminence (hatched bars) densities. The spontaneous rate was 10 spikes/s for cortical noise simulations, and the number of spontaneous spikes depended on the stimulus duration. The amplitude refers to peak-to-peak values. In AD and FD tasks the standard stimulus was a 20Hz 20μm vibration, and the bars (black for AD, and grey for FD) represent the percent difference in amplitude (for AD task) and frequency (for FD task) of the comparison stimulus that the simulated subjects required to reliably discriminate the vibrations in the 2IFC tasks. The error bars represent ±1SE.

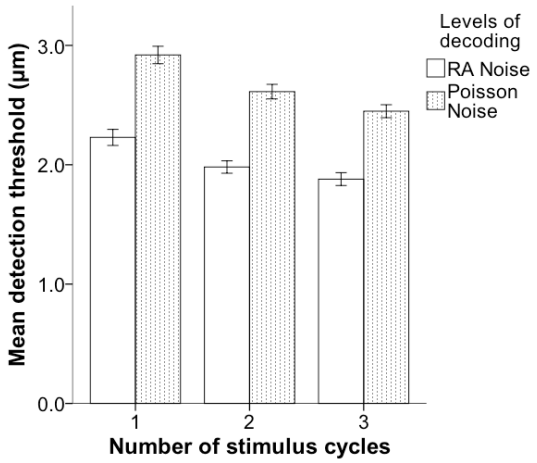

Figure 2.11: Each bar represents the average performance of 10 simulated subjects in TD task (at 20Hz stimulus frequency) conducted at two levels of decoding (afferent level: clear bars, cortical level: stippled bars). In separate blocks, we increased the stimulus duration by increasing the number of stimulus cycles. The spontaneous rate was 10 spikes/s for cortical noise simulations and the stimulus duration determined the number of spontaneous spikes. The amplitude refers to peak-to-peak values. The error bars represent ±1SE.

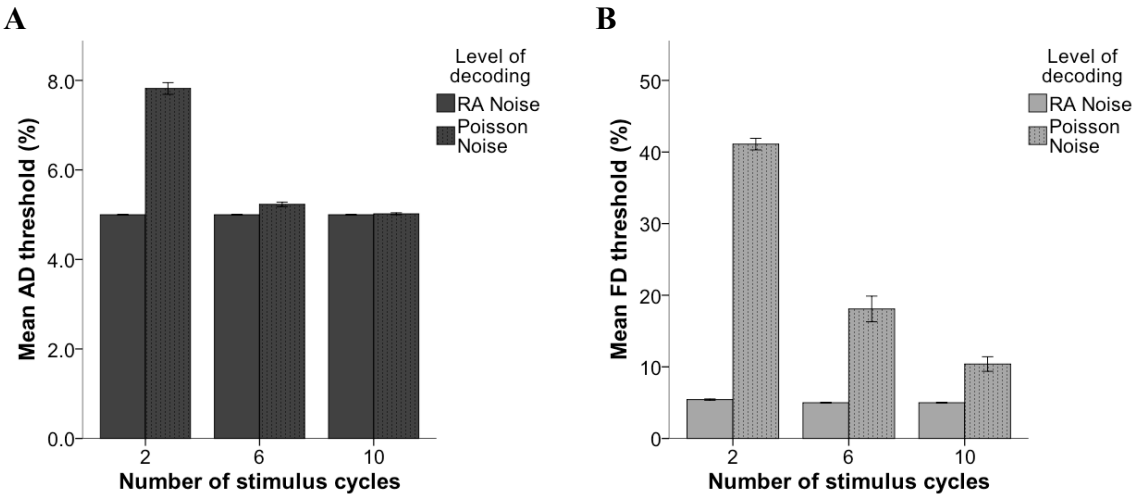

Figure 2.12: Each bar represents the average performance of 10 simulated subjects in AD (panel A) and FD (panel B) tasks conducted at different stimulus durations, which we manipulated by changing the number of stimulus cycles. We conducted the tasks at two levels of decoding (afferent level: clear bars, cortical level: stippled bars). The spontaneous rate was 10 spikes/s for cortical noise simulations and the stimulus duration determined the number of spontaneous spikes. The amplitude refers to peak-to-peak values. In AD and FD tasks the standard stimulus was a 20Hz 20μm vibration, and the bars (black for AD, and grey for FD) represent the percent difference in amplitude (for AD task) and frequency (for FD task) of the comparison stimulus that the simulated subjects required to reliably discriminate the vibrations in the 2IFC tasks. The error bars represent ±1SE.

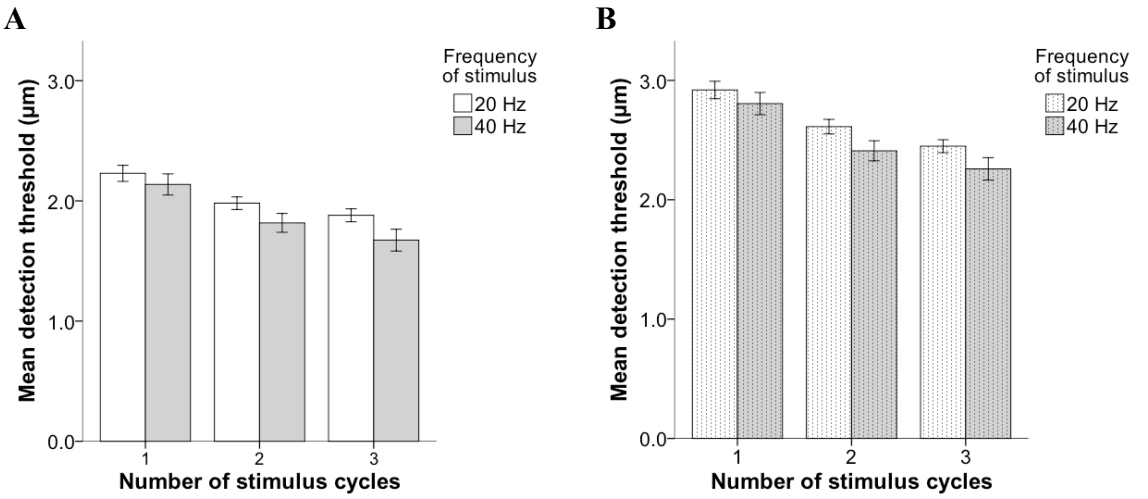

Figure 2.13: Each bar represents the average performance of 10 simulated subjects in TD task at two stimulus frequencies (20Hz – white bars, and 40Hz – grey bars). The ideal observer's performance at two decoding levels are shown in panel A (afferent level) and panel B with stippled bars (cortical level). The spontaneous rate was 10 spikes/s for cortical noise simulations and the stimulus duration determined the number of spontaneous spikes. The amplitude refers to peak-to-peak values. The error bars represent ±1SE.

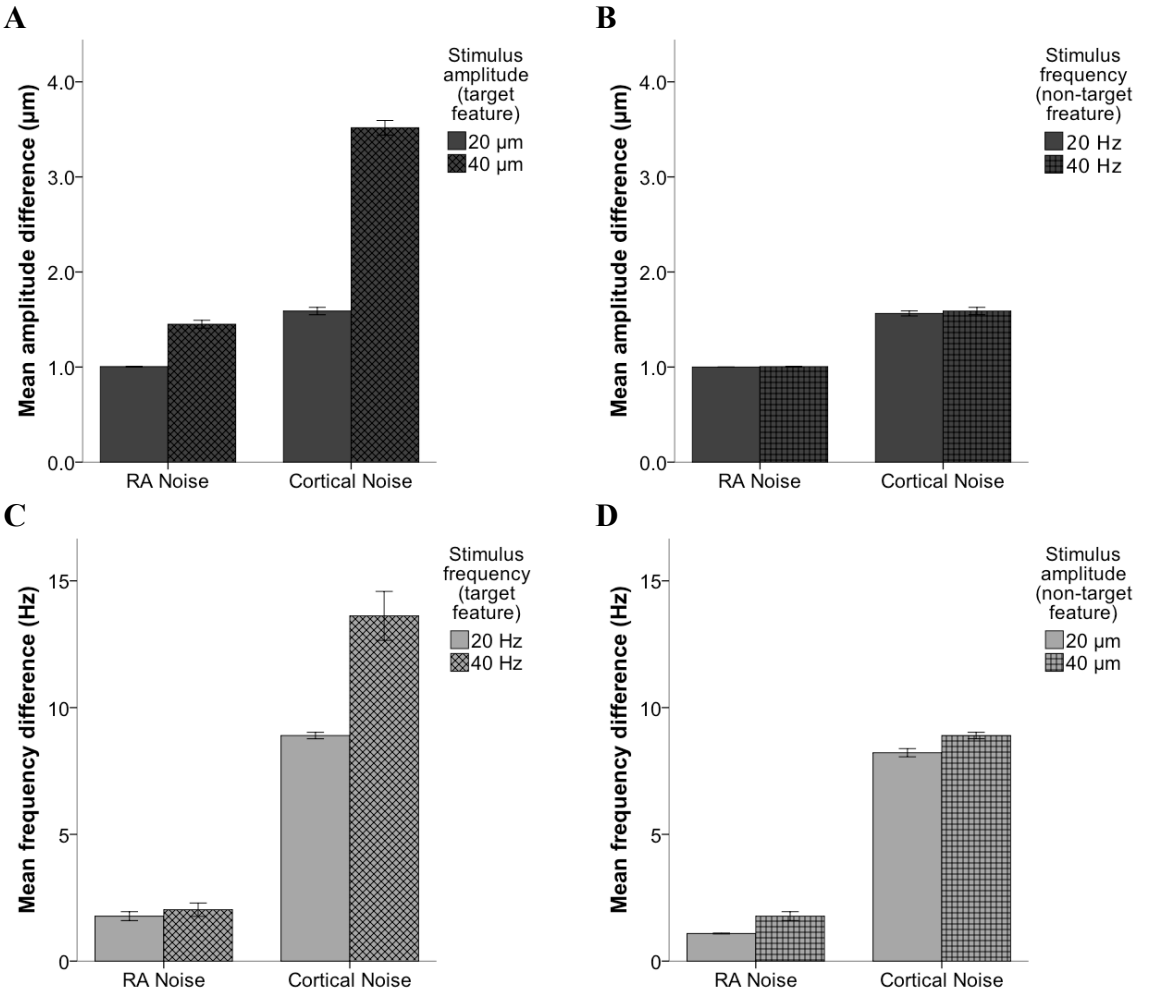

Figure 2.14: The average discrimination threshold of 10 simulated subjects at two decoding levels in AD tasks, in which A) the standard stimulus amplitudes were 20µm (clear black bars) and 40µm (hatched black bars), and frequency was 40Hz in all the blocks; B) the standard stimulus amplitude was 20µm in all the blocks, but the frequency (i.e. the non-target feature) was 20Hz (clear bars) and 40Hz (checkered bars). Performance of the simulated subjects at two decoding levels in FD tasks, in which C) the standard stimulus frequency were 20Hz (clear grey bars) and 40Hz (hatched grey bars), and amplitude was 40µm in all the blocks; D) the standard stimulus frequency was 20Hz in all the blocks, but the amplitude (i.e. the non-target feature) was 20µm (clear bars) and 40µm (checkered bars). The spontaneous rate was 10 spikes/s for cortical noise simulations and the stimulus duration determined the number of spontaneous spikes. The amplitude refers to peak-to-peak values. The error bars represent ±1SE.

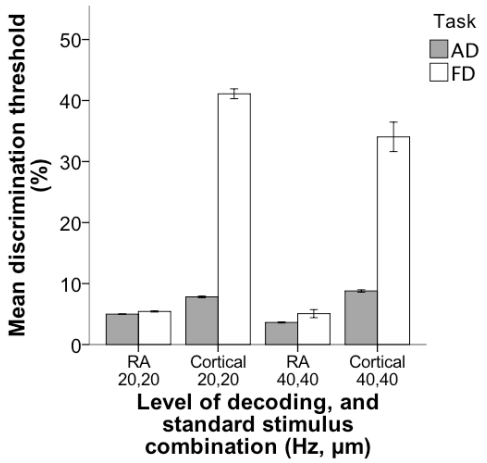

Figure 2.15: The average discrimination threshold of 10 simulated subjects at two decoding levels in AD (black bars) and FD (grey bars) tasks, where the standard stimuli were identical. The amplitude refers to peak-to-peak values. The error bars represent ±1SE.

CHAPTER 3: SPATIAL AND TEMPORAL SUMMATION IN THE RA CHANNEL

3.1 Introduction

Perception is a process of making inferences about the stimuli that are presented to the sensory system (Goldreich, 2007). Sensory receptors convert stimulus features into nerve impulses, which are then relayed to the brain. Although this neural encoding process might be highly efficient, repeated presentations of identical stimuli rarely evoke identical cortical responses. As a consequence, a particular response could have been produced by any one of several stimuli. The brain can therefore at best infer probabilistically the most likely stimulus that may have evoked the neural data. Under these conditions, the accuracy of a perceptual inference depends strongly on the amount of sensory evidence that the nervous system has acquired (Shadlen and Newsome, 1996; Yang and Shadlen, 2007). In this chapter, we investigate two features of the nervous system that can influence the amount of evidence it can acquire: 1) the ability of the nervous system to accumulate information over time, and 2) the ability of the nervous system to accumulate information over space (i.e., from pools of sensory receptors).

A common and procedurally simple experiment used to investigate human's sensitivity to vibration is the "vibrotactile threshold detection task". In a 2-interval forced-choice (2IFC) protocol, we present a vibrotactile stimulus in one of two intervals, and the participants report the interval in which they perceived the stimulus. By using this procedure, we estimate the minimum amplitude at which the participants reliably detect the vibrotactile stimulus, i.e. we determine the detection threshold of the participants. Presumably, if the nervous system is capable of accumulating evidence over time, then sinusoidal stimuli applied for longer periods of time should be more easily detectable than the same stimuli applied for a shorter period of time – a phenomenon known as temporal summation. Similarly, if the nervous system is capable of accumulating responses from a pool of sensory receptors, then stimuli applied to skin regions of high receptor density should be more easily detectable than the same stimuli applied to skin regions of low receptor density – a phenomenon known as spatial summation. In both phenomena, an increase in detectability is quantified as a decrease in detection threshold.

A vibratory stimulus delivered to the skin surface primarily activates two afferent systems or psychophysical channels: 1) rapidly adapting type-1 afferents (RA channel), which are effectively activated by vibrotactile frequencies ranging from 5 Hz to 50 Hz, and 2) rapidly adapting type-2 afferents (or Pacinian channel), which are effectively activated by

vibrotactile frequencies ranging from 50 Hz to 300 Hz (Talbot et al., 1968; Bolanowski et al., 1988; Gescheider et al., 2001). Several studies have investigated temporal and spatial summation in vibrotactile perception (for review, see Gescheider et al., 2009). To test for spatial summation, the investigators varied the number of receptors activated by a) changing the size of the stimulus probe, b) testing different body locations of varying receptor density, and/or c) testing different groups of participants who might have different receptor densities, e.g., young and old participants (Gescheider et al., 1994; Goble et al., 1996). By implementing one of these strategies (or some combination thereof) several studies have reported that spatial summation exists in the Pacinian channel (Verrillo, 1963; Gescheider et al., 1994, 2002, 2005; Goble et al., 1996; Whitehouse et al., 2006; Morioka et al., 2008; for review see Gescheider et al., 2009). By contrast, the results for the RA channel are unclear. While several studies have reported the absence of spatial summation in the RA channel (Gescheider et al., 1994; Goble et al., 1996), others have reported that spatial summation does exist in the RA channel (e.g., Whitehouse et al., 2006; Morioka et al., 2008). To resolve this discrepancy, we investigated whether spatial summation exists in the RA channel by estimating vibrotactile detection threshold in the fingertip and thenar eminence; the afferent density in fingertip is 141 units/cm^2, which is approximately six times that in thenar eminence (25 units/cm^2) (Johansson and Vallbo, 1979a).

As with spatial summation, studies have consistently shown that temporal summation occurs in the Pacinian channel (Verrillo, 1965; Gescheider et al., 1999, 2002). However, it is unclear whether temporal summation occurs in the RA channel too. Whereas, results from a few studies indicate that temporal summation does not exist in the RA channel (Verrillo, 1965; Gescheider and Joelson, 1983), a few other studies reported evidence suggesting that temporal summation does exist in the RA channel (Green, 1976; Hämäläinen et al., 1981). To investigate whether temporal summation exists in the RA channel we conducted a series of five experiments to estimate vibrotactile detection thresholds with stimuli of different durations.

3.2 Methods

3.2.1 Participants

We tested a total of 75 normally sighted (18.62 – 30.25 years of age; median age, 20.97; 35 male, 40 female) participants. Based on a handedness survey (modified from Oldfield, 1971), 65 of our participants were right-hand dominant. We ensured that participants (based on their self-report) did not have any nervous system disorders, learning disabilities, hearing impairment, and injuries or calluses on their index fingertip. Additionally, we screened participants for dyslexia as it adversely affects tactile acuity

(Grant et al., 1999), and for diabetes because it causes peripheral neuropathy and slows down action potential conduction (Hyllienmark et al., 1995). The Research Ethics Board of McMaster University approved all procedures of this study. The participants signed an informed consent form before starting the experiment, and were paid in cash or course-credit for their time in the laboratory.

3.2.2 Vibrotactile stimulation

The experimental setup is identical to that reported in Bhattacharjee et al. (2010). Using LabVIEW 6.1 (National Instruments) running in a Macintosh G3 computer (Apple), we generated vibrotactile stimuli by passing voltage waveforms (32 kilo samples per second at 12 bit resolution) through an analog output channel (National Instruments PCI-MIO-16E-1 board) to a power amplifier (Bruel & Kjaer type 2718) that displaced a flat surface circular glass probe (11 mm diameter) attached to a precision mini-shaker (Bruel & Kjaer 4810). The probe displacement amplitudes (in microns) that we report in this study should be considered as nominal values because the on-line displacement profile was unavailable (for technical reasons); however, we precalibrated the system with a charge accelerometer (Bruel & Kjaer type 4381) threaded onto the mini-shaker.

To ensure that the participant's tested body region (fingertip or thenar eminence) maintained a light but steady force against the probe surface, we attached a load cell transducer (Daytronic model 434AM-250G) that triggered a warning (announced to the participant by a computer voice) if the forces were < 20 g or > 50 g; a force warning resulted in exclusion of the trial. We considered 20-50 g as an acceptable window because it was convenient for most participants and variation in baseline force within this range presumably exerts little effect on the responses of the relevant peripheral afferents (Johnson and Lamb, 1981).

The mini-shaker was housed in a medium-density fiberboard box (L-29, W-29, H-14 cm) with a circular opening (13 mm diameter) in the top surface through which the probe contacted the skin. The probe surface was level with the top surface of the box with a 1 mm gap separating the probe edge and the edge of the opening. During testing the forearm rested in prone position, level with the fiberboard box, on a foam pad (H-14cm). Depending on the experiment, the probe contacted the distal pad of the index finger or the thenar eminence (the palm region at the base of the thumb) while the palm and/or the (remaining) fingers rested on top of the box.

3.2.3 Procedure

The overall perceptual task and the design were identical in all five experiments. Before starting the experiment, the experimenter familiarized the participants with the equipment

and the response unit. The experimenter then conducted a "force practice" task in order to acquaint the participants with the amount of force they were allowed to exert on the vibrotactile probe (i.e. between 20-50 g, see section 3.2.2). Next, the experimenter explained the task to the participants and answered any questions the participants had about the task. To assure the experimenter that the participants fully understood the task, the participants had to repeat the task instructions back to the experimenter. The testing session began when the experimenter was satisfied that the participants had fully understood the task instructions.

We tested participants on several blocks (10 in experiments 1, 3, and 5, 11 in experiment 2, 22 in experiment 4) to estimate their vibrotactile detection threshold at two frequencies (40 Hz in experiments 1, 2, and 5, 20 Hz in experiment 3, 20 Hz and 40 Hz in experiment 4). Every block consisted of 40 test trials and at least 10 practice trials. During practice trials, the participants received auditory feedback. In every experiment we implemented the 2IFC protocol, i.e. in each trial, a vibratory stimulus was presented randomly in one of two intervals announced by auditory tones; the participants, by pressing a response button, reported the interval in which they detected the stimulus. To estimate each participant's detection threshold, we adaptively adjusted (see section 3.2.4) the amplitude of the vibrotactile stimulus to values between 1μm and 100μm. Specifics of each experiment are given in the following subsections.

Experiment 1: To test for spatial summation, in 5 out of 10 blocks we delivered 40 Hz vibrotactile stimuli to the distal pad of the index finger, and in the remaining 5 blocks we delivered stimuli to the thenar eminence. To test for temporal summation, we delivered stimuli at 5 different durations (represented by the number of taps at 40 Hz), which were fixed within each block. The stimulus durations were 12.5ms, 237.5ms, 487.5ms, 737.5ms, and 987.5ms, corresponding to 1, 10, 20, 30, and 40 taps, respectively. The starting duration was counterbalanced such that the participants either started at the 40-taps block or at the 1-tap stimulus duration block. Four different combinations allowed us to counterbalance for the duration of the stimulus in the first block and for the stimulation location; 24 participants (6 participants X 4 conditions) were randomly assigned to each one of the 4 combinations. Before starting the first block on a new stimulation location (fingertip or thenar eminence), the participants performed the "force practice" task (see section 3.2.2 and 3.2.3). The participants received 20 practice trials on the first block at both stimulation locations.

Experiment 2: Using 40 Hz vibrations, we further tested for temporal summation on the index fingertip. We conducted 11 experimental blocks at 7 different tap number conditions: 2 blocks each at 1, 2, 3, and 4 taps, which corresponded to 12.5ms, 34.5ms, 62.5ms, and 87.5ms respectively, and 1 block each at 5, 10, and 20 taps, which corresponded to 112.5ms, 237.5ms, and 487.5ms respectively. Each participant was first tested sequentially on blocks with 20, 10, and 5 taps and then randomly tested twice on blocks with 1, 2, 3, and 4 taps. During the 20-taps block, the participants were presented with 20 practice trials.

Experiment 3: We tested for temporal summation on 20 Hz vibrations presented to the index fingertip. Each participant completed 10 experimental blocks, which consisted of 10 different tap numbers: 1, 2, 3, 4, 5, 6, 8, 10, 15, and 20 taps; these tap numbers corresponded to stimulus durations: 25ms, 75ms, 125ms, 175ms, 225ms, 275ms, 375ms, 475ms, 725ms, and 975ms, respectively. In the first 3 of 10 blocks, each participant was tested sequentially on 20-, 15-, and 10-taps condition, and then on the remaining 7 blocks, the participants were tested on a unique random order of the tap numbers enumerated above. Similar to experiment 2, during the first block, the participants were presented with 20 practice trials.

To study the effect of stimulus duration and frequency on vibrotactile detection thresholds, we chose to control for stimulus duration by manipulating the number of stimulus cycles or pulses because at any fixed duration, doubling the stimulus frequency would double the number of stimulus cycles. Freeman and Johnson (1982) showed that increasing the frequency of the stimulus lowers the absolute threshold (i.e. the amplitude at which the afferents starts to elicit neural responses) of all the RA afferents. Moreover, if every RA afferent at certain amplitude has a certain probability of responding, then an increase in the number of stimulus cycles will increase the possibility that the afferent will respond at least in one of several stimulus cycles (even though the probability of generating a neural response of that afferent in a single cycle may be very low); this phenomenon is known as *probability summation*. Therefore, by keeping the stimulus duration constant at different frequencies we would have delivered different number of taps, and consequently, we would not have been able to determine whether the decline in detection threshold with the increase in the stimulus frequency was due to a) an increase in the number of stimulus cycles or pulses, or b) difference in the sensitivity of the RA afferents.

Experiment 4: In experiments 1, 2, and 3, the stimuli were a combination of sinusoids and half-cycle sinusoids: for the single tap and the last tap in a series of taps or stimulus cycles, we presented a half-cycle pulse at the target frequency. Interestingly, at identical stimulus amplitudes, the effective amplitude of a full sinusoid (i.e. stimulus with 2 or more taps) is twice that of a half-cycle sinusoid (i.e. stimulus with 1 tap) (see Fig. 3.1 a and b). This suggests that if participants reliably detect a 1-tap stimulus at *a* microns, then they might be able to detect a multi-tap stimulus at *a*/2 microns (see Fig. 3.1c). To avoid this potential confound, in the current experiment we used pulsatile stimuli where all multi-tap stimuli consisted of a series of half-wave rectified half-cycle sinusoid taps (see Fig. 3.1d).

Every participant attended two separate counterbalanced sessions (11 blocks per session) in which we tested for temporal summation on 20 Hz and 40 Hz vibrations using pulsatile stimuli presented to the index fingertip. On each testing session, as in experiment 2, we conducted 11 experimental blocks at 7 different tap number conditions: every participant was first tested sequentially on blocks with 20, 10, and 5 taps and then randomly tested

twice on blocks with 1, 2, 3, and 4 taps. The stimulus duration depended on the stimulus frequency. During the first testing block of both sessions, the participants were presented with 20 practice trials.

Experiment 5: In all the experiments of this study the intervals within each trial were announced by 25ms auditory tones, and the SOA between the auditory tone and the vibrotactile stimuli was 0ms. To investigate whether the auditory tones that defined the intervals were suppressing the tactile stimulus-evoked responses (i.e. crossmodal masking; see Gescheider and Niblette, 1976), we tested each participant on two SOA durations (0ms and 500ms) in separate (counterbalanced) blocks.

Within each SOA duration condition, we tested for temporal summation using pulsatile stimuli presented at 40 Hz on the index fingertip. We conducted 5 experimental blocks (per SOA condition) at 3 different tap number conditions. The participants were always tested first on a block with 10 taps, and then randomly tested twice on blocks with 1 and 4 taps. During the 10-taps block, the participants were presented with 20 practice trials.

3.2.4 Bayesian adaptive method

To choose the most informative stimulus levels during the experiment, and to estimate each participant's psychometric function ($P_c(l)$, i.e. the probability of a correct response as a function of stimulus level, l), we implemented a modified version of the ψ method (Konstevich and Tyler, 1999). We modeled each participant's psychometric function as a mixture of a cumulative normal function and a lapse rate term (δ):

$$P(c_l \mid \theta,\beta,\delta) = \frac{\delta}{2} + (1-\delta)\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\frac{d'}{\sqrt{2}}} \exp\left(-\frac{x^2}{2}\right)dx$$

$$\text{where, } d' = \left(\frac{l}{\theta}\right)^{\beta} \text{ such that } d' = 1 \text{ when } l = \theta.$$

Thus, the θ-parameter is the threshold stimulus level, which corresponds to 76% correct response probability (see Fig. 3.2c), and the β-parameter is the slope of the psychometric function.

We treated all three psychometric function parameters as unknown parameters and began the algorithm with a uniform prior probability distribution over a wide range of parameter values – threshold (θ: 1 to 100μm), slope (β: 0.01 to 15), and lapse rate (δ: 0.01 to 0.1). Thus the algorithm generated several thousand psychometric functions using all possible combinations of the above-mentioned parameter values. After each trial the algorithm updated the likelihood of the psychometric functions and calculated the expected information gain associated with each stimulus level; it then chose for the next stimulus

level the one that predicted the highest information gain. At the end of the testing block, the joint probability distribution over psychometric function parameters was marginalized over β and δ to obtain the posterior probability density function (PDF) for each participant's θ-parameter (see Fig. 3.2b).

3.2.5 Derived performance measures

We parameterized a Temporal Summation Curve (TSC) to describe the decline in a participant's threshold as a function of tap number, n:

$$\theta(n) = (\theta_1 - b)s^{n-1} + b$$

$$\text{where, } s = e^{\left(\frac{-1}{\tau}\right)}.$$

Here, s is the summation constant, τ is the exponential decay constant, θ_1 is the threshold for a 1-tap stimulus, and b is the asymptotic threshold approached as the number of taps increases towards infinity. In the equation above, s is the magnitude of the reduction in detection threshold caused by each additional tap; hence, $s = 1$ implies no proportional difference between threshold at n and at n-1 taps.

We found the TSC that best fit each participant's performance (correct and incorrect answers at each tested stimulus level $\{cr_l, ic_l\}$) across all testing blocks. To do so, we considered TSCs with summation constants ranging from $s = 1$ to 1000 taps, θ_1 ranging from 1 to 100 microns, and b ranging from 0 to θ_1-1 microns. We then determined the maximum likelihood TSC:

$$\hat{TSC} = \underset{s,\theta_1,b}{\operatorname{argmax}} \left[\prod_n P(\{cr_l, ic_l\}_n \mid s, \theta_1, b) \right]$$

3.2.6 Statistical analysis

Participant response consistency: Based on a participant's performance i.e. correct and incorrect responses at different stimulus amplitudes (Fig. 3.2a), the Bayesian adaptive method estimates the detection threshold for that participant. Specifically, the procedure calculates the probabilities that certain stimulus amplitude is the actual threshold of the participant, which is represented as a distribution known as the posterior Probability Distribution Function (PDF) (Fig. 3.2c); the mode of the posterior PDF is the most probable value of the participant's threshold. The posterior PDF also reflects the participant's attentiveness and response consistency, which can be quantified as the 95% confidence interval (CI). For example, Figure 3.3 shows performance of a participant in

two different experimental conditions; compared to the grey curve, the black curve represents more consistent performance. Because inconsistent performance leads to poor estimate of a participant's threshold, we calculated the shortest 95% CI in each block for every participant and only considered the blocks where the shortest 95% CI was ≤ 10μm.

We considered the derived performance measures and threshold estimates as dependent variables, and performed *t*-tests and analyses of variance (ANOVAs) using SPSS 20 (IBM) for Macintosh, with a significance criterion of 0.05. In the results section, all reported Bonferroni-corrected comparisons are *p* values corrected for multiple comparisons.

3.3 Results

3.3.1 Detection threshold significantly varies with stimulation location

To test for spatial summation in the RA channel, in experiment 1 we estimated participants' ability to detect 40 Hz stimuli of varying tap numbers delivered to the fingertip and the thenar eminence. An ANOVA with stimulation location (fingertip, thenar eminence) and number of taps (1, 10, 20, 30, 40) as within-subject factors, and sex (male, female) as a between-subject factor revealed a main effect of stimulation location [$F(1,11) = 14.67$; $p < 0.01$; $\eta^2 = 0.571$] and number of taps [$F(4,44) = 34.44$; $p < 0.001$; $\eta^2 = 0.758$], but no main effect of sex ($p = 0.98$). For conditions with more than one tap, Bonferroni corrected comparisons of thresholds between thenar eminence and fingertip revealed significantly higher thresholds at the thenar eminence than at the fingertip ($p<0.05$). For the 1-tap condition, the thenar eminence threshold was also higher than the fingertip threshold, although this effect was not statistically significant ($p = 0.066$) (see Fig. 3.4a).

3.3.2 Detection threshold significantly varies with the number of taps

To test for temporal summation, we performed two-tailed *t*-tests on the summation component *s* of the derived measure TSC, and determined whether the mean *s*-value is significantly different from 1 (where $s = 1$ suggests no temporal summation). In all 5 experiments, the mean *s*-value was significantly different from 1, suggesting that the detection thresholds decreased with increasing number of taps (see Fig. 3.5). In all the analyses on *s* values, we did not consider sex as a between-subject factor because

ANOVAs, performed on the threshold data of all 5 experiments, with sex as a between-subject factor did not reveal a significant interaction of sex by number of taps.

In experiment 1, participants showed temporal summation at both the thenar eminence and the fingertip (one-sample *t*-test: on *s*-values, $p<0.001$). A paired samples *t*-test showed that temporal summation was not different between the two stimulation locations (2-tailed, $p=0.159$) (see Fig. 3.4b; Fig. 3.5 bar 1, 2). Interestingly, whereas pairwise comparisons with Bonferroni corrections revealed that the participants' detection thresholds significantly decreased as we increased the number of taps from 1 to 10 ($p<0.01$), further increase in the number of taps did not have any favourable effect on the detection thresholds ($p=1$), i.e. the detection thresholds at 10, 20, 30, and 40 taps did not differ from one another.

Because experiment 1 did not include tap numbers between 1 and 10, in experiment 2 we investigated whether, and how, gradually increasing the number of taps affect detection thresholds. We tested participants' detection thresholds for 40 Hz vibrations presented as 1-, 2-, 3-, 4-, 5-, 10-, and 20-tap stimuli. As evidence for temporal summation, a one-sample t-test showed that the *s*-values were significantly different from 1 ($p<0.001$) (see Fig. 3.6b; Fig. 3.5 bar 3). Interestingly, an ANOVA with number of taps as the within-subject factor and sex as between-subject factor revealed a main effect of number of taps [$F(6,36) = 21.06$; $p < 0.001$; $\eta^2 = 0.778$] and a main effect of sex ($p=0.671$). Further pairwise comparisons with Bonferroni corrections revealed that the participants' detection thresholds significantly decreased as we increased the number of taps from 1 to 2 ($p<0.01$); however, the detection thresholds at 2 taps and above were not significantly different from one another (p ranges between 0.35 and 1) (see Fig. 3.6a).

Next we investigated whether we would see effects similar to experiment 2 if we used a stimulus frequency of 20 Hz. In experiment 3, we tested participants' detection thresholds at tap numbers 1, 2, 3, 4, 5, 6, 8, 10, 15, and 20 using 20 Hz vibrotactile stimuli. A one-sampled *t*-test on *s*-values provided evidence for temporal summation ($p<0.001$) (See Fig. 3.7b; Fig. 3.5 bar 4). However, an ANOVA with number of taps as the within-subject factor and sex as the between-subject factor revealed no main effect of number of taps ($p=0.176$) or sex ($p=0.918$) (see Fig. 3.7a).

In experiment 4, we estimated the detection threshold for 20 Hz and 40 Hz pulsatile stimuli on the same group of participants. At both frequency conditions, increasing the number of taps of the stimuli reduced the detection threshold. The *s*-values in both conditions were significantly different from 1 (20 Hz, 40 Hz: $p<0.01$) and they did not differ between conditions ($p=0.88$) (see Fig. 3.8b; Fig. 3.5 bar 5, 6).

In experiment 5, we estimated detection thresholds when the interval-identifying auditory stimuli were presented either simultaneously (SOA: 0ms) or 525ms before the 40 Hz vibrotactile stimuli. The *s* values in both conditions were significantly different from 1

(SOA 0ms: $p<0.05$, SOA 525ms: $p<0.01$) and they did not differ between conditions ($p=0.74$) (see Fig. 3.9b; Fig. 3.5 bar 7,8).

3.3.3 Effect of frequency on detection threshold

In experiments 2 and 3, we tested participants' ability to detect 40 Hz and 20 Hz vibrotactile stimuli, respectively. However, these experiments were conducted on two different groups of participants. Therefore, to investigate the effect of frequency on detection threshold, in experiment 4 we estimated detection threshold for 20 Hz and 40 Hz pulsatile stimuli on the same group of participants.

An ANOVA with frequency (20 Hz, 40 Hz) and number of taps (1, 2, 3, 4, 5, 10, 20) as within-subject factors and sex as a between-subject factor revealed a main effect of frequency [$F(1,8) = 17.37$; $p < 0.01$; $\eta^2 = 0.685$] and number of taps [$F(6,48) = 10.99$; $p < 0.001$; $\eta^2 = 0.579$] but no main effect of sex ($p=0.321$) (see Fig. 3.8a). The ANOVA failed to show a significant interaction between frequency and number of taps ($p=0.362$), which suggests that the effect of number of taps did not differ in the two frequency conditions. A further examination into the effect of frequency at each tap number condition revealed that the thresholds at 20 Hz were always higher than those at 40 Hz; however, Bonferroni corrected pairwise comparisons showed that this difference was statistically significant only for conditions where the stimulus contained 2, 3, 5, and 20 taps ($p<0.05$), and not for conditions where the stimulus contained 1 ($p=0.08$), 4 ($p=0.09$), and 10 ($p=0.25$) taps.

3.3.4 No evidence for crossmodal masking

In the current study, because the intervals of each trial in all our tasks are identified by auditory stimuli, it might have been possible that the auditory stimuli were perceptually masking the single tap vibrotactile stimuli. To investigate, we tested participants' ability to detect 1, 4, and 10 tap 40 Hz pulsatile stimuli at two different conditions where the interval-indicating auditory stimuli either appeared with (SOA: 0ms) or before (SOA: 525ms) the vibrotactile stimuli.

An ANOVA with SOA (0ms, 525ms) and number of taps (1, 4, 10) as within-subject factors and sex as a between-subject factor revealed only a main effect of number of taps [$F(2,16) = 12.75$; $p < 0.001$; $\eta^2 = 0.614$] and no main effect of SOA ($p=0.358$) or sex ($p=0.354$), which suggests that the detection thresholds were not affected by the auditory stimuli (see Fig. 3.9a). If auditory stimuli were masking the responses of vibratory stimuli, particularly for the block with 1 tap stimuli, we would expect the detection threshold at 0 ms SOA to be higher than that at 525 ms SOA. On the contrary, a pairwise comparison with Bonferroni correction revealed that the average detection threshold of 1-

tap stimuli at 0 ms SOA was significantly lower than that at 525 ms SOA (p=0.02), whereas there was no statistical difference at the other tap conditions (4 taps: p=0.65; 10 taps: p=0.69).

3.4 Discussion

In this study we tested whether the RA channel is capable of spatial and temporal summation. Motivated by the simulation results from the ideal observer model (described in chapter 2), which showed differential performance due to changes in receptor density and stimulus duration, we conducted five experiments on human participants. Our data provide evidence for spatial and temporal summation in this channel.

3.4.1 Effect of stimulation location: Evidence for spatial summation in the RA channel

Decline in detection threshold with concurrent increase in stimulated area is defined as spatial summation (Gescheider et al., 2009). Alternatively, differences in detection sensitivity due to differences in receptor density could also be interpreted as spatial summation (Gescheider et al., 2002).

Here, using the same equipment and testing protocol, we estimated participants' ability to detect 40 Hz stimuli presented to their fingertip and to their thenar eminence. The thresholds were higher at the thenar eminence than at the fingertip in all stimulus durations tested. Because the RA afferent density at the fingertip is higher (by ~6 times) than at the thenar eminence (Johansson and Vallbo, 1979a), we interpret our results, albeit cautiously, as evidence for spatial summation.

Gescheider et al. (2002) (using a 0.72 cm^2 circular probe to present 5 Hz stimuli) did not find any difference in detection thresholds between the index fingertip and the thenar eminence. In the current study, we used a 0.95 cm^2 circular probe to present 40 Hz stimulus and found an effect of stimulation location on detection thresholds. It is possible that, by using 40 Hz stimulus, we may have inadvertently activated the Pacinian channel (Gescheider and Joelson, 1983), which is known to show spatial summation (see for review Gescheider et al. 2009). Interestingly, Gescheider et al. (2002) reportedly failed to find a difference in detection thresholds at the fingertip and thenar eminence even for high frequency stimuli that activate the Pacinian channel, which is known to exhibit spatial summation (see for review Gescheider et al. 2009).

Several investigators using different contactor sizes (probe diameters 2 and 8 mm, Kekoni et al., 1989; 6 and 10 mm, Morioka and Griffin, 2005; 1, 3, 6, and 10 mm, Gu and Griffin,

2013) reported lower detection thresholds at the fingertip than at the thenar eminence in human participants; however, none of these studies found an effect of contactor size, i.e. contactors with larger probe diameter did not yield lower detection thresholds at each tested location, which is contrary to the definition of spatial summation.

Morioka et al. (2008) tested four different stimulation locations (fingertip, volar forearm, large toe, and heel) using 1 and 6 mm diameter stimulus probes. The authors reported that the detection thresholds were lowest on the fingertip, which has the highest density of afferents among the tested locations, and that the threshold at the volar forearm was the highest, which lacks Meissner's corpuscles (the receptors that activate the RA afferents in the glabrous skin) (Vallbo et al., 1995). Morioka et al's (2008) observations occurred for both stimulus probes; however, at all stimulated locations there was no effect of probe diameter on the detection thresholds at the RA effective frequencies.

Verrillo (1963) did not test detection thresholds at the fingertip; however, the author used a range of contactor sizes (probe areas: 0.005, 0.02, 0.08, 0.32, 1.3, 2.9, and 5.1 cm^2) to present RA channel activating frequencies at the thenar eminence and did not observe any change in detection threshold due to changes in contactor probe areas.

Why do different studies consistently show an effect of stimulation location, but no effect of contactor size? If afferent density affects detection thresholds then changing stimulation location or contactor size should vary the number of activated afferents. Could the sensitivity of the afferents vary with location but when different contactor sizes are used on the same location the most sensitive afferents are always inadvertently activated? Johansson and Vallbo (1979b) studied psychophysical detection thresholds and the RA afferent sensitivities by directly recording neural responses from the afferents. They identified regions of high and low psychophysical thresholds on the glabrous skins of the human hand and compared the afferent thresholds (i.e., the minimum amplitude required to elicit neural responses) of those skin areas. The authors did not find any difference in afferents' sensitivity at the high and low psychophysical threshold regions. Thus, Johansson and Vallbo's (1979b) study suggests that the difference in psychophysical thresholds might be due to other factors and not due to any sensitivity differences in RA afferents at different skin regions.

We speculate that the biomechanics of the skin at different stimulation locations might influence vibratory sensitivity. For example, Morioka et al. (2008) observed higher detection thresholds at the heel than at the large toe of human participants, although the density of afferents is presumably uniform throughout the plantar region of the foot (Kennedy and Inglis, 2002). The Meissner's corpuscles that transform the vibratory stimulus into neural impulses are found in the papillary ridges of the dermis that extend into the epidermis of the glabrous skin (Nolano et al., 2003; Verrillo and Bolanowski, 2009). Recently, an ultrasound study estimated the thickness of the skin on the plantar region of the foot and reported that the epidermis at the heel is thicker than that at the

pulp of the large toe (Chao et al., 2011), which might explain why Morioka et al. (2008) observed higher detection thresholds at the heel than at the large toe.

Interestingly, Whitehouse et al. (2006) tested five different locations on the distal phalanx of the right middle finger starting from the most distal part to the more proximal part near the interphalangeal crease representing high density to lower density of receptors respectively. The investigators, using 1 and 6 mm diameter probe and stimulus frequencies that activate the RA channel, found a significant effect of location, with the lowest threshold at the highest density region. Nevertheless, at each stimulated location the authors of this study, like the above-mentioned studies, did not find a significant difference in thresholds due to differences in contactor size. Whitehouse et al. (2006) did, however, report that the effect of stimulation location was stronger for the 1 mm diameter probe than for the 6 mm diameter probe, presumably due to overlap in stimulation location with the bigger diameter probe. The 6 mm diameter probe probably activated the most sensitive afferents at the stimulated location, as well as some afferents at the adjoining stimulation locations.

In light of these findings, we speculate, in addition to skin mechanics, that two further factors influence vibrotactile detection – a) the activation of the most sensitive afferent under the stimulus probe, and b) the activation of a small but critical number of mechanosensitive afferents. Perhaps the same most sensitive afferents are activated when a stimulus is delivered to one high-receptor-density location with probes of varying diameter; however, when the stimulation location is changed to one with sparser afferent density, the activation of a critical number of afferents is not reached. However, it is still unclear why a sufficiently large probe would not be effective at activating the critical number of receptors in a sparsely innervated skin area.

Aging causes loss of receptors (Cauna, 1965; Bolton et al., 1966; Bruce, 1980), and if the density of receptors reduce due to aging this would presumably affect the threshold of vibrotactile detection (Gescheider et al. 1994; Goble et al. 1996). Gescheider et al. (1994), using different contactor sizes and a range of frequencies, concluded that aging adversely affects detection threshold with younger participants outperforming older participants. The difference in thresholds between the two age groups were, however, smaller for low-medium frequency vibrations. Goble et al. (1996) found consistent differences in thresholds between young and old adults for a range of frequencies covering low-medium to high frequencies. The authors suggested that loss of receptors, among other factors, was a parsimonious explanation for the observed differences in thresholds between age groups.

In the current study, detection thresholds of male and female participants in all the experiments were almost identical. Verrillo (1979), and more recently, Seah and Griffin, (2008) did not find a significant difference in detection thresholds in male and female participants. Bhattacharjee et al. (2010), however, reported a significant main effect of sex on the detection task; the investigators tested the ability of blind and sighted male and

female participants of a wide range of ages (19-80 years) to detect a single tap in a 2IFC task and found that female outperformed male participants. Interestingly, Gescheider et al. (1994) reported an effect of participants' sex on vibrotactile detection tasks – female outperformed male participants; however, the authors found that the difference in the detection thresholds of male and female participants were statistically significant only for the older participants but not for the young participants. An interpretation of histological data from Dillon et al. (2001) and Nolano et al. (2003) suggest that the ratio of density of Meissner's corpuscles in female to male participants are ~1.5 and ~1.6, respectively. Therefore, an effect of difference in receptor density between male and female participants may not reflect on the psychophysical thresholds; note that the fingertip surface area of male also tend to be larger than that of female participants (Dillon et al., 2001; Nolano et al., 2003; Peters et al., 2009).

3.4.2 Effect of stimulus duration: Evidence for temporal summation in the RA channel

Temporal summation may be defined as a decrease in the detection threshold due to an increase in stimulus duration (Gescheider et al., 2009). In five experiments we changed the duration of the stimuli by changing the number of taps at 20 Hz and 40 Hz stimulus frequencies, and the data consistently provide evidence for temporal summation in the RA channel. Whereas different studies have concluded that temporal summation does not exist in the RA channel (Verrillo, 1965; Gescheider and Joelson, 1983; see Gescheider et al., 2009 for review), other studies (including the current study) have shown evidence supporting temporal summation in this channel (Green, 1976; Hämäläinen et al., 1981).

Verrillo (1965), using 100 Hz, 200 Hz, and 500 Hz vibrations (that primarily activate the Pacinian channel) at different stimulator sizes, tested human participants' ability to detect vibratory stimuli varying in number of taps or pulses, and showed that temporal summation is only possible when stimuli are presented with big surface contactors. According to Verrillo (1965) at contactor sizes ≤ 0.02 cm^2 there was no effect of pulse number on the detection thresholds because smaller contactors (≤ 0.02cm^2) are unable to activate the Pacinian channels (which shows temporal summation) and presumably instead activate the RA channel or the slowly adapting type 2 channel (which do not show temporal summation) (Bolanowski et al., 1988). However, Green (1976) pointed to the absence of electrophysiological studies that tested for the effect of contactor size on afferent responses; thus, the author argued that to determine whether the RA channel demonstrates temporal summation, the effect of the number of pulses on detection thresholds at low frequencies should be tested directly, and if temporal summation is still not present even with a larger probe then it could be concluded that the RA channel fails to show temporal summation. To investigate, Green (1976) delivered RA activating frequencies (25 Hz, 40 Hz) with a large contactor (diameter: 12.96 mm), and reported data that show temporal summation in the RA channel. However, the author was unsure

about the mechanism for observing the effect of pulse number on detection thresholds and reasoned that the effect could be due to: a) that temporal summation actually exists in the RA channel when activated by large sized contactors, or b) that there are two psychophysical channels and possibly both channels are getting activated; the RA channel might be the nonsummating channel, and the other summating channel might be causing the temporal summation effects.

Hämäläinen et al. (1981) investigated human participants' ability to detect 20 Hz vibrations at three different stimulus durations (50, 150, and 300 ms) and observed that participants' detection threshold declined with the increase in stimulus duration. However, Gescheider and Joelson (1983) tested detection thresholds in human participants on a range of frequencies including 25 and 40 Hz, and reported absence of temporal summation in these frequencies.

Should contactor size matter for temporal summation? Green (1976) used a contactor size of 12.96 mm diameter and we used 11 mm diameter. Both studies used a large contactor size and reported temporal summation, which follows one of the predictions of Green (1976) that large contactor size will cause temporal summation in the RA channel. However, Gescheider and Joelson (1983) used a contactor size of 19.54 mm diameter, which is larger than both of the above-mentioned studies, and yet did not observe temporal summation for 25 Hz and 40 Hz frequency vibrations. Interestingly, Hämäläinen et al. (1981) using a contactor size of 1.99 mm diameter observed temporal summation. Therefore, the difference in results between our study and that of Gescheider and Joelson's (1983) study is certainly not due to contactor size.

3.4.3 Our finding of temporal summation is not an artifact of cross-modal masking, or of activation of the Pacinian channel

Our analyses of behavioural data from all the experiments showed temporal summation in the RA channel; it is important to note that our participants performed particularly poorly when their objective was to detect the presence of a single tap in one of two intervals. In the current study, to demarcate two intervals we used auditory tones and the stimulus-onset-asynchrony (SOA) between the tactile stimulus and the auditory tone was set to 0 ms. Because the duration of the auditory beeps were 25 ms, the 40 Hz single tap (12.5 ms duration) and 20 Hz (25 ms duration) single tap stimuli may have been partially masked by the auditory tone. Thus, we wondered whether increased detection threshold for a single tap stimulus could be due to cross-modal masking i.e. whether the auditory stimulus suppressed the tactile stimulus-evoked responses. To investigate, in experiment 5 we compared participants' performance in two conditions in separate blocks where we presented the auditory stimulus 525 ms or 0 ms before the onset of the tactile stimulus. We did not observe any difference in the participants' performance. We conclude that

participants' higher detection threshold at the single tap condition is not an artifact of the interval-identifying stimulus.

Although both 20 and 40 Hz vibrations primarily activate the RA afferents (Gescheider et al., 2009; Mountcastle et al., 2005), Gescheider and Joelson (1983) suggested that 40 Hz vibrations might also activate the Pacinian channels but 20 Hz vibrations at threshold amplitudes should not activate the Pacinian channel. Because the Pacinian channel demonstrates both temporal and spatial summation, low frequency vibratory stimuli inadvertently activating the Pacinian channel could lead to erroneous conclusions about temporal and spatial summation in the low frequency channel. In experiment 1 and 2, we estimated participants' detection threshold for 40 Hz stimuli. Therefore, to verify whether temporal summation exists in the RA channel we tested participants' detection threshold at 20 Hz vibratory stimuli in experiment 3, and observed that increasing the number of taps decreased detection thresholds. To further confirm that in the current study both 20 and 40 Hz vibrations are activating the RA channel, we conducted experiment 4 where within the same group of participants we measured detection thresholds at both 20 and 40 Hz. The analyses of the summation factor *s* showed strong temporal summation effects, and the values of *s* at the two frequencies were almost identical, thus confirming the presence of temporal summation in the RA channel.

3.5 Conclusion

In Chapter 2 we simulated vibrotactile detection task using the ideal observer model, and found that afferent density and stimulus duration affected its performance on the detection task. In the current chapter, we tested whether human performance showed similar effects. In experiment 1, by testing the fingertip and the thenar eminence, we showed that detection threshold varies with RA afferent density. In experiment 1 to 5, we showed that detection threshold varies with stimulus duration. Further experiments are needed to conclusively determine whether spatial summation exists in the RA channel and the current effect of stimulus location is not due to accidental activation of the Pacinian channel, or due to differences in skin biomechanics of the fingertip and the thenar eminence. Future experiments are also needed to verify the existence of temporal summation by estimating the detection threshold in human participants with varying stimulus duration delivered through stimulus contactors of different sizes.

3.6 References

Bhattacharjee A, Ye AJ, Lisak JA, Vargas MG, Goldreich D (2010) Vibrotactile masking experiments reveal accelerated somatosensory processing in congenitally blind braille readers. J Neurosci 30:14288-14298.

Bolanowski SJ Jr, Gescheider GA, Verrillo RT, Checkosky CM (1988) Four channels mediate the mechanical aspects of touch. J Acoust Soc Am 84:1680-1694.

Bolton CF, Winkelmann RK, Dyck PJ (1966) A quantitative study of Meissner's corpuscles in man. Neurology 16:1–9.

Bruce MF (1980) The relation of tactile thresholds to histology in the fingers of elderly people. J Neurol Neurosurg Psychiatry 43:730 –734.

Cauna N (1965) The effects of aging on the receptor organs of the human dermis. In: Advances in biology of skin (Montagna W, ed), pp 63–96. New York: Pergamon.

Chao CY, Zheng YP, Cheing GL (2011) Epidermal thickness and biomechanical properties of plantar tissues in diabetic foot. Ultrasound Med Biol 37:1029-1038.

Dillon YK, Haynes J, Henneberg M (2001) The relationship of the number of Meissner's corpuscles to dermatoglyphic characters and finger size. J Anat 199:577-584.

Freeman AW, Johnson KO (1982) A model accounting for effects of vibratory amplitude on responses of cutaneous mechanoreceptors in macaque monkey. J Physiol 323:43-64.

Gescheider GA, Berryhill ME, Verrillo RT, Bolanowski SJ (1999) Vibrotactile temporal summation: probability summation or neural integration? Somatosens Mot Res 16:229-242.

Gescheider GA, Bolanowski SJ, Hall KL, Hoffman KE, Verrillo RT (1994) The effects of aging on information-processing channels in the sense of touch: I. Absolute sensitivity. Somatosens Mot Res 11:345–357.

Gescheider GA, Bolanowski SJ, Hardick KR (2001) The frequency selectivity of information-processing channels in the tactile sensory system. Somatosens Mot Res 18:191-201.

Gescheider GA, Bolanowski SJ, PopeJV, Verrillo RT (2002) A four-channel analysis of the tactile sensitivity of the fingertip: frequency selectivity, spatial summation, and temporal summation. Somatosens Mot Res 19:114 –124.

Gescheider GA, Güçlü B, Sexton JL, Karalunas S, Fontana A (2005) Spatial summation in the tactile sensory system: probability summation and neural integration. Somatosens Mot Res 22:255-268.

Gescheider GA, Joelson JM (1983) Vibrotactile temporal summation for threshold and suprathreshold levels of stimulation. Percept Psychophys 33:156-162.

Gescheider GA, Niblette RK (1967) Cross-modality masking for touch and hearing. J Exp Psychol 74:313-320.

Gescheider GA, Wright JH, Verrillo RT (2009) Information-Processing Channels in the Tactile Sensory System: A Psychophysical and Physiological Analysis. New York: Psychology Press.

Goble AK, Collins AA, Cholewiak RW (1996) Vibrotactile threshold in young and old observers: the effects of spatial summation and the presence of a rigid surround. J Acoust Soc Am 99:2256-2269.

Goldreich D (2007) A Bayesian Perceptual Model Replicates the Cutaneous Rabbit and Other Tactile Spatiotemporal Illusions. PLoS ONE 2(3): e333. doi:10.1371/journal.pone.0000333

Grant AC,Zangaladze A,Thiagarajah MC,Sathian K (1999) Tactile perception in developmental dyslexia: a psychophysical study using gratings. Neuropsychologia 37:1201–1211.

Green BG (1976) Vibrotactile temporal summation: effect of frequency. Sens Processes 1:138-149.

Gu C, Griffin MJ (2013) Spatial summation of vibrotactile sensations at the foot. Med Eng Phys 35:1221-1227.

Hämäläinen H, Pertovaara A, Soininen K, Järvilehto T (1981) Is there low frequency vibrotactile temporal summation? Scand J Psychol. 1981;22(3):203-206.

Hyllienmark L, Brismar T, Ludvigsson J (1995) Subclinical nerve dysfunction in children and adolescents with IDDM. Diabetologia 38:685– 692.

Johansson RS, Vallbo AB (1979a) Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin. J Physiol 286:283-300.

Johansson RS, Vallbo AB (1979b) Detection of tactile stimuli. Thresholds of afferent units related to psychophysical thresholds in the human hand. J Physiol 297:405-22.

Johnson KO (1974) Reconstruction of population response to a vibratory stimulus in quickly adapting mechanoreceptive afferent fiber population innervating glabrous skin of the monkey. J Neurophysiol 37:48-72.

Johnson KO, Lamb GD (1981) Neural mechanisms of spatial tactile discrimination: neural patterns evoked by Braille-like dot patterns in the monkey. J Physiol 310:117–144.

Kekoni J, Hämäläinen H, Rautio J, Tukeva T (1989) Mechanical sensibility of the sole of the foot determined with vibratory stimuli of varying frequency. Exp Brain Res 78:419-424.

Kennedy PM, Inglis JT (2002) Distribution and behaviour of glabrous cutaneous receptors in the human foot sole. J Physiol 538:995-1002.

Morioka M, Griffin MJ (2005) Thresholds for the perception of hand-transmitted vibration: dependence on contact area and contact location. Somatosens Mot Res 22:281-297.

Morioka M, Whitehouse DJ, Griffin MJ (2008) Vibrotactile thresholds at the fingertip, volar forearm, large toe, and heel. Somatosens Mot Res 25:101-112.

Mountcastle VB (2005) The sensory hand: neural mechanisms of somatic sensation. Harvard University Press.

Mountcastle VB, Talbot WH, Sakata H, Hyvärinen J (1969) Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. J Neurophysiol 32:452-484.

Nolano M, Provitera V, Crisci C, Stancanelli A, Wendelschafer-Crabb G, Kennedy WR, Santoro L (2003) Quantification of myelinated endings and mechanoreceptors in human digital skin. Ann Neurol 54:197–205.

Oldfield RC (1971) The assessment and analysis of handedness: the Edinburgh inventory. Neuropsychologia 9:97–113.

Peters RM, Hackeman E, Goldreich D (2009) Diminutive digits discern delicate details: fingertip size and the sex difference in tactile spatial acuity. J Neurosci 29:15756-15761.

Seah SA, Griffin MJ (2008) Normal values for thermotactile and vibrotactile thresholds in males and females. Int Arch Occup Environ Health 81:535-543.

Shadlen MN, Newsome WT (1996) Motion perception: seeing and deciding. Proc Natl Acad Sci 93:628-633.

Sripati AP, Yoshioka T, Denchev P, Hsiao SS, Johnson KO (2006) Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system. J Neurosci 26:2101-2114.

Talbot WH, Darian-Smith I, Kornhuber HH, Mountcastle VB (1968) The sense of flutter-vibration: comparison of the human capacity with re- sponse patterns of mechanoreceptive afferents from the monkey hand. J Neurophysiol 31:301–334.

Vallbo AB, Olausson H, Wessberg J, Kakuda N (1995) Receptive field characteristics of tactile units with myelinated afferents in hairy skin of human subjects. J Physiol 483:783-795.

Verrillo RT (1963) Effect of Contactor Area on the Vibrotactile Threshold. J Acoust Soc Am 35:1962-1966.

Verrillo RT (1965) Temporal summation in vibrotactile sensitivity. J Acoust Soc Am 37:843-846.

Verrillo RT (1979) Comparison of vibrotactile threshold and suprathreshold responses in men and women. Percept Psychophys 26:20–24.

Verrillo RT, Bolanowski S J (2009) Tactile responses to vibration. In: Handbook of signal processing in acoustics (Havelock D, Kuwano S, Vorlander M, ed), pp1185-1213. New York: Springer.

Whitehouse DJ, Morioka M, Griffin MJ (2006) Effect of contact location on vibrotactile thresholds at the fingertip. Somatosens Mot Res 23:73-81.

Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. Nature 447:1075-1080.

3.7 Figures and figure captions

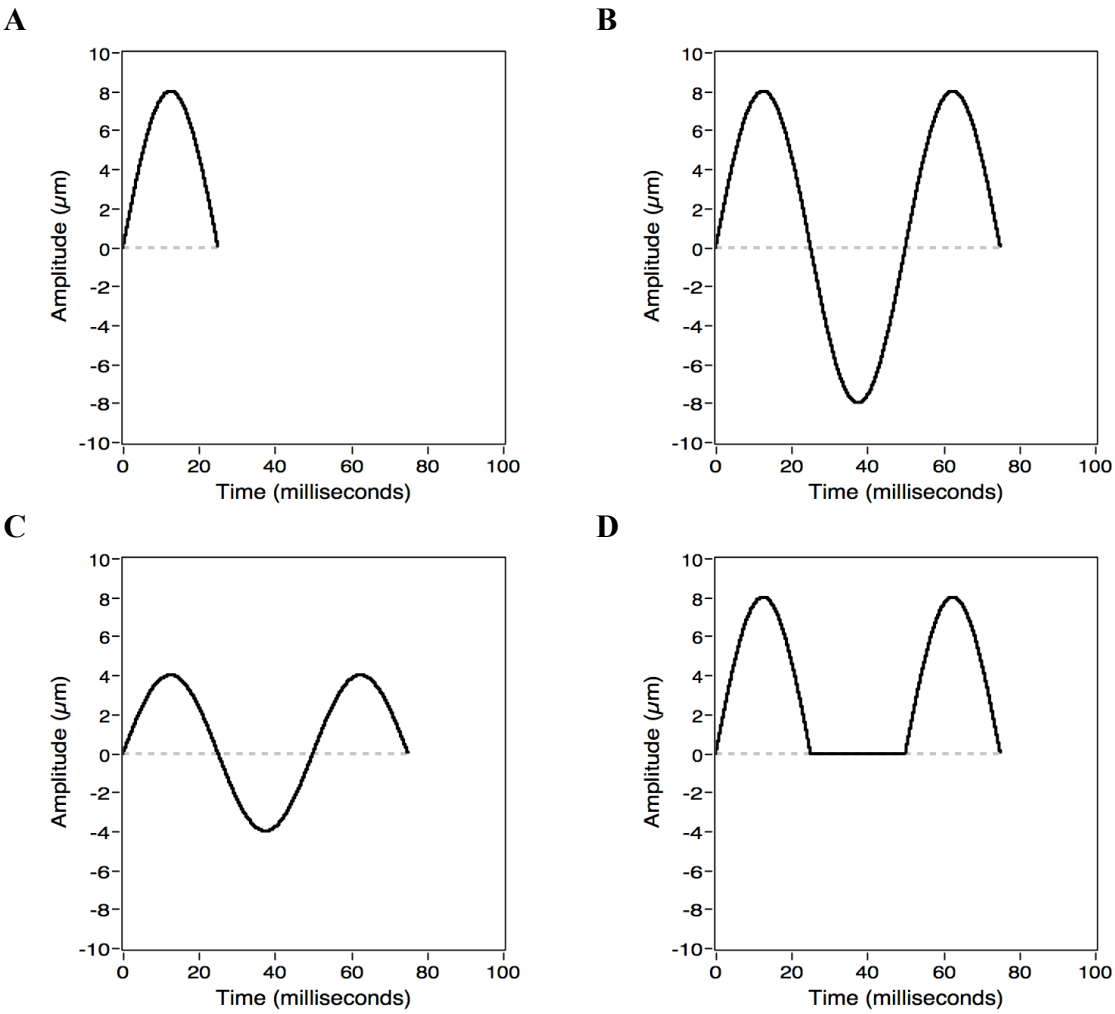

Figure 3.1: 20 Hz vibrotactile stimuli shown at two different amplitudes. Note that the amplitude of the single tap (panel A) is same as that of the two tap (panel B) stimulus; however, the effective (peak-to-peak) amplitude of the 2-tap stimulus (panel B) is twice that of 1-tap stimulus. Panel C shows a 2-tap stimulus with an amplitude of 4μm but has an effective amplitude of 8μm, which is same as the amplitude of the single tap stimulus (panel A). To resolve this discrepancy, in experiment 4 and 5 we used pulsatile stimuli for all multi-tap stimuli (panel D). Irrespective of the number of taps the effective amplitude was equal to the waveform amplitude of the stimulus.

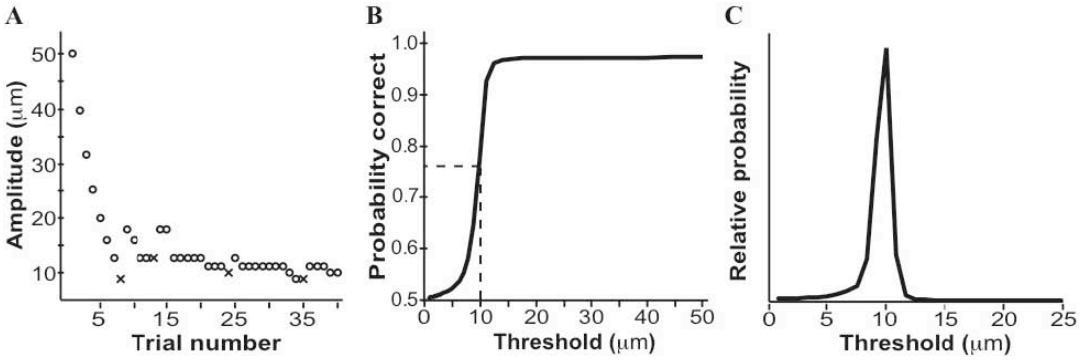

Figure 3.2: a) Participant's performance (o = correct, x = incorrect), b) participant's psychometric function, c) posterior probability distribution function of threshold (θ) parameter.

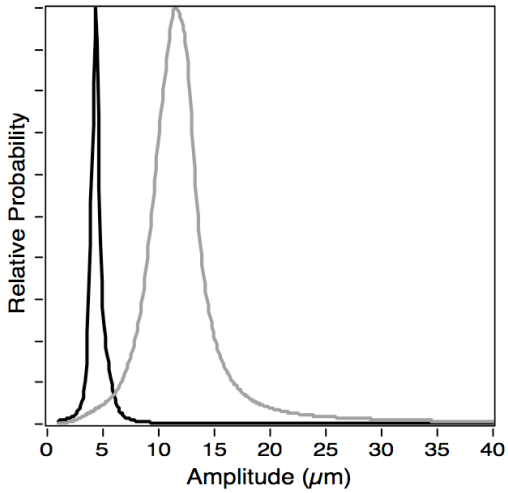

Figure 3.3: Shows a participant's posterior probability distribution function of the threshold (θ) parameter on two different conditions. When compared to the black curve, the grey curve represents poor performance by the participant. The width of the probability distribution function affects our confidence on the participant's threshold.

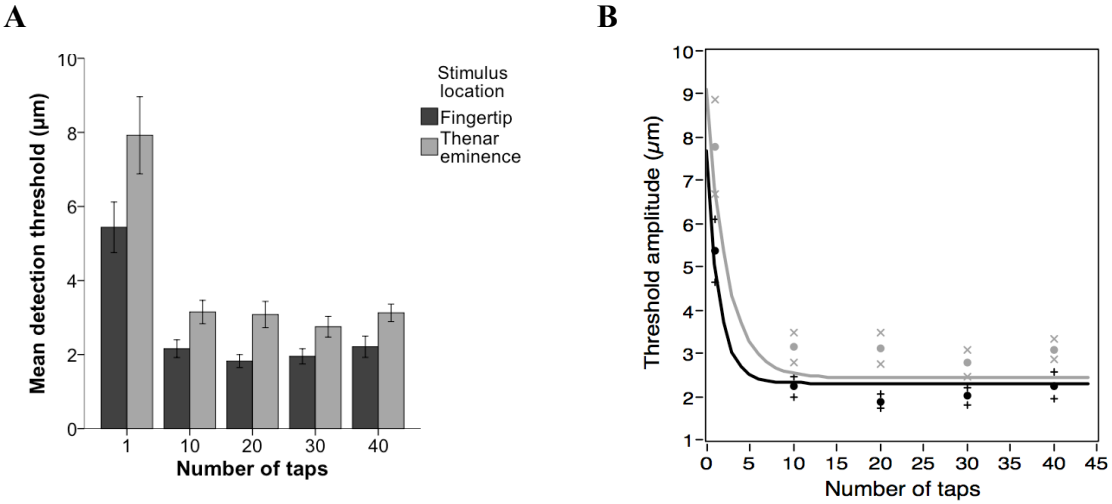

Figure 3.4: Panels A and B, show the mean threshold of participants on experiment 1 where 40 Hz vibrations of 1, 10, 20, 30, and 40 taps (x-axis) were delivered to their fingertip (black bars in A, black dots in B) and thenar eminence (grey bars in A, grey dots in B). The error bars in panel A, the black pluses (+) and the grey crosses (x) flanking the mean thresholds in panel B represent ±1 SE. In panel B, along with the mean threshold values (also shown in A), we show the mean TSCs for fingertip (black curve) and thenar eminence (grey curve). To calculate each TSC, we computed the average of all the participants' s, θ_1, b values in each condition.

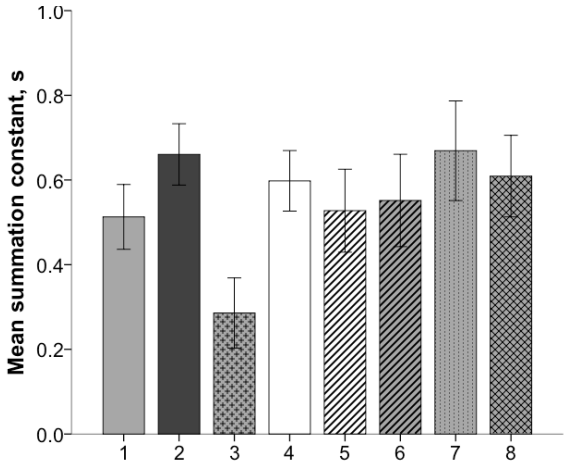

Figure 3.5: Shows the mean summation constant, i.e. the summation or *s* component of the TSCs, for all 5 experiments. The numbers in the x-axis represents every condition tested in this study. Bar 1 and 2, refers to fingertip and thenar eminence *s*-values in experiment 1, respectively. Bar 3, refers to experiment 2. Bar 4, refers to experiment 3. Bar 5 and 6 refer to 20 Hz and 40 Hz stimulus frequency *s*-values in experiment 4. Bar 7 and 8 refer to 0 ms and 500 ms SOA *s*-values in experiment 5.

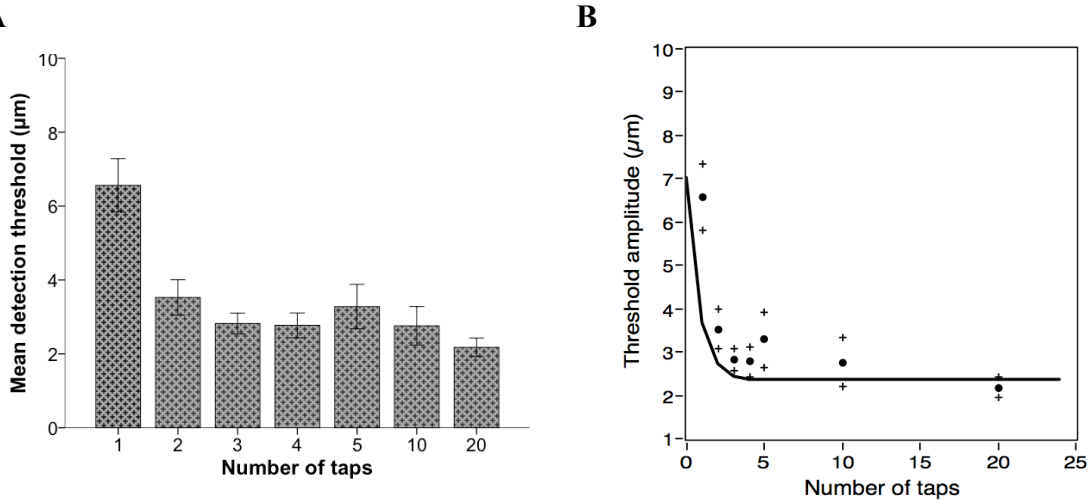

Figure 3.6: Panel A and B, shows the mean threshold of participants on experiment 2 where 40 Hz vibrations of 1, 2, 3, 4, 5, 10, and 20 taps (x-axis) were delivered to their fingertip (bars in A, dots in B). The error bars in panel A, and the black pluses (+) flanking the mean thresholds (dots) in panel B represent ±1 SE. In panel B, along with the mean threshold values (also shown in A), we show the mean TSC, which we computed from the average of all the participants' *s*, θ_1, *b* values in each condition.

A **B**

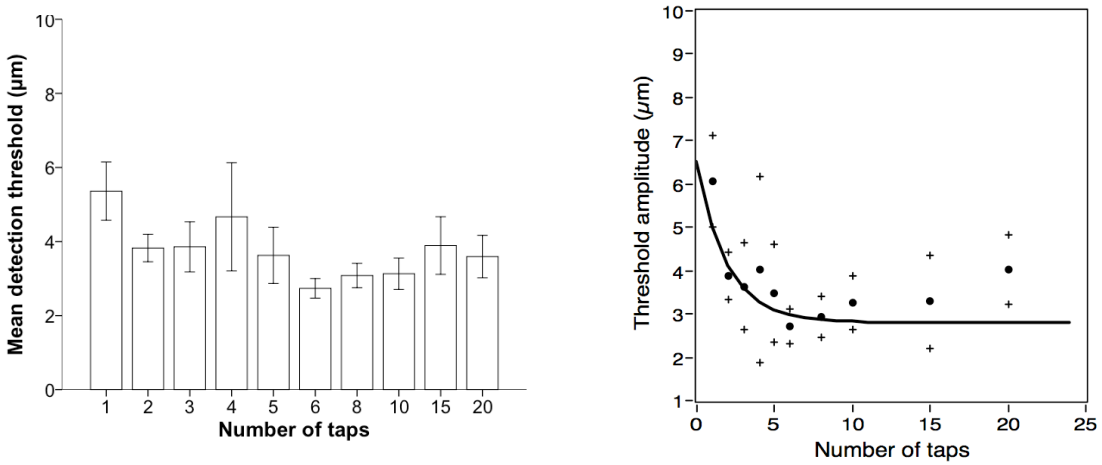

Figure 3.7: Panel A and B, shows the mean threshold of participants on experiment 3 where 20 Hz vibrations of 1, 2, 3, 4, 5, 6, 8, 10, 15, and 20 taps (x-axis) were delivered to their fingertip (bars in A, dots in B). The error bars in panel A, and the black pluses (+) flanking the mean thresholds (dots) in panel B represent ±1 SE. In panel B, along with the mean threshold values (also shown in A), we show the mean TSC, which we computed from the average of all the participants' s, θ_1, b values in each condition.

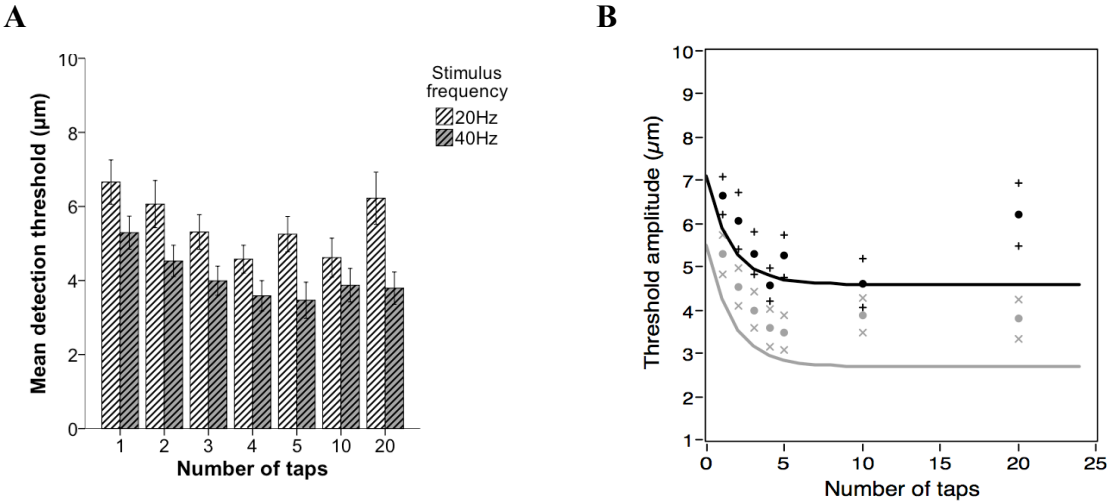

Figure 3.8: Panel A and B, shows the mean threshold of participants on experiment 4 where 1-, 2-, 3-, 4-, 5-, 10-, and 20-tap (x-axis) vibrations were presented to their fingertip at two different frequencies - 20 Hz (hatched bars with clear background in panel A, and black dots in panel B) and 40 Hz (hatched bars with grey background in panel A, and grey dots in panel B). The error bars in panel A, the black pluses (+) and the grey crosses (x) flanking the mean thresholds in panel B represent ±1 SE. In panel B, along with the mean threshold values (also shown in A), we show the mean TSCs for 20 Hz (black curve) and 40 Hz (grey curve) vibration conditions. To calculate each TSC, we computed the average of all the participants' s, θ_1, b values in each condition.

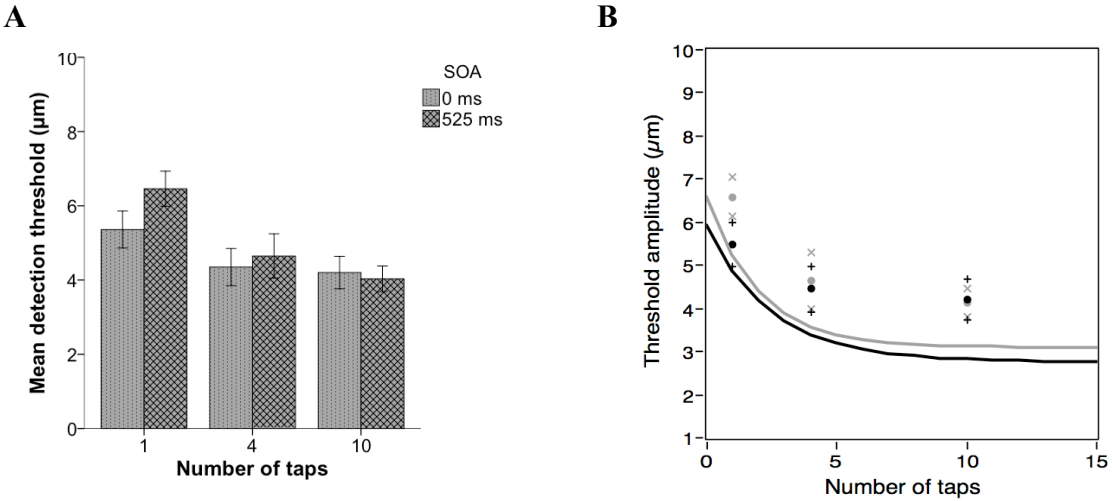

Figure 3.9: Panel A and B, shows the mean threshold of participants on experiment 5 where 1-, 4-, and 10-tap (x-axis) vibrations were presented to their fingertip at two different beep and vibration SOAs – 0ms (stippled bars in A, and black dots in B) and 525ms (plaid bars in A, and grey dots in B). The error bars in panel A, the black pluses (+) and the grey crosses (x) flanking the mean thresholds in panel B represent ±1 SE. In panel B, along with the mean threshold values (also shown in A), we show the mean TSCs for 0ms SOA (black curve) and 525ms SOA (grey curve) conditions. To calculate each TSC, we computed the average of all the participants' s, θ_1, b values in each condition.

CHAPTER 4: EFFECT OF TARGET AND NON-TARGET FEATURE MANIPULATION IN AMPLITUDE AND FREQUENCY DISCRIMINATION TASKS

4.1 Introduction

A vibrotactile stimulus presented on the skin surface evokes neural responses, which the brain has to decode to determine the stimulus features. In this study, we asked what features of a vibrotactile stimulus are encoded, and what are the neural signatures (i.e. neural codes) of those features?

Two obvious choices for the vibrotactile stimulus features are amplitude and frequency. There is unanimous agreement that the brain encodes vibrotactile amplitude in neural firing rate, i.e. number of spikes per second (Bensmaïa, 2008; Simons et al., 2005; Tommerdahl et al., 2010; Güçlü and Dinçer, 2013; Harvey et al., 2013). However, the neural code for vibrotactile frequency is unclear. On the one hand, Romo and colleagues have argued that neural firing rate adequately encodes vibrotactile frequency (Hernandez et al 2000; Salinas et al 2000; Luna et al., 2005; see for review Romo and Salinas, 2003; Romo and de Lafuente, 2013). On the other hand, studies have shown that the activity of neurons that respond to vibrotactile stimuli are phase-locked to the frequency of the stimulus, which suggests that the stimulus frequency is encoded in the regularity of the neural activity (Mountcastle et al. 1969, 1990; Whitsel et al. 2001).

Interestingly, because amplitude and frequency are two coexisting features of a vibrotactile stimulus, it is possible that humans might meld these features into a single feature to perform vibrotactile perceptual tasks (Harris et al., 2006). Neurophysiological studies, however, suggest that the frequency affects the sensitivity of the tactile mechanosensitive afferents (Freeman and Johnson, 1982) and the amplitude modulates the number of spikes (Johnson, 1974). Our Bayesian ideal observer model utilizes this information to infer the amplitude and frequency of a vibrotactile stimulus (see Chapter 2). Therefore, to examine whether humans can separately perceive the amplitude and frequency features of vibrotactile stimuli, and to make an educated speculation about the neural code that the brain might utilize during vibrotactile tasks, we compared the performance of human participants to that of the ideal observer.

To probe the ideal observer model's accuracy of inference, we could present a single stimulus and record the inferred feature values (e.g., the model's inferred frequency in Hz and amplitude in μm); however, because it is infeasible for human participants to report the feature values of a vibrotactile stimulus, we could not implement this straightforward approach to determine the participants' ability to infer the stimulus features. Therefore, in the current study we conducted a series of vibrotactile discrimination tasks, namely – amplitude discrimination (AD) and frequency discrimination (FD) tasks.

Ideally, to perform a vibrotactile discrimination task the participant should infer the amplitude and frequency values from the stimulus-evoked neural responses in each interval and then, depending on the task (AD or FD), choose the interval with the higher target feature (A or F, respectively) value as the interval of choice. Recently, Harris et al. (2006) conducted two different versions of an FD experiment: one in which the amplitude in the two intervals was identical, and the other in which the authors adjusted the amplitude of both intervals such that the total energy (a x f) in the two intervals was identical. The authors observed that the participants' FD threshold was higher on the identical-energy than on the identical-amplitude version of the task. Because Harris et al. (2006) eliminated the energy cue by matching the total energy in the intervals, they concluded that human participants encode the product of frequency and amplitude (i.e. the energy) of the vibration, and use that as a feature to perform FD tasks. Note that, although the participants showed higher threshold in the equal energy version compared to the identical amplitude version, the participants were still able to perform the task. How could the participants do the task if vibrotactile discrimination requires the energy cue and the experimenters eliminated that cue?

Similar to energy difference as a potential cue, different studies have suggested that humans might utilize the intensity difference as a cue to perform vibrotactile FD tasks (Goff, 1967; LaMotte and Mountcastle, 1975; Mountcastle and Romo, 1990). Therefore, to match the "subjective intensity" in the two intervals, which presumably matches the number of spikes evoked in the intervals (Johnson, 1974; Mountcastle, 2005), the investigators adjusted the stimulus amplitude in two intervals by delivering high frequency intervals at amplitudes lower than the amplitudes of the low frequency intervals. Despite eliminating the intensity cue, the participants were able to perform the discrimination task and the investigators could estimate the participants' threshold values, which suggest that the participants were probably inferring the frequency information in each interval, or the participants might have utilized the neural code related to the periodicity of the neural activity.

Interestingly, in another experiment of the same study, Harris et al. (2006) showed that in an FD task the difference in amplitude (i.e. the non-target feature) of the vibrotactile stimulus in two intervals affects the FD threshold of human participants. During the task, the participants received two vibrations (corresponding to two intervals of a trial) at different frequencies and the amplitude of these vibrations was either identical or different such that in some "different" trials the interval containing the high frequency

stimulus was presented at amplitudes higher than that of the low frequency stimulus interval (i.e. congruent condition); in other trials, the low frequency stimulus was presented at amplitudes higher than that of the high frequency stimulus interval (i.e. incongruent condition) (see Table 4.1). The participants were unaware of this manipulation and were instructed to report the interval with the high frequency stimulus. Harris et al. (2006) observed that in the congruent condition, the participants correctly chose the high amplitude high frequency interval as the response; however, in the incongruent condition, the participants incorrectly chose the high amplitude low frequency interval as the response. Therefore, the amplitude manipulation affected the participants' performance such that their FD thresholds were biased by the interval with the higher amplitude. This finding suggests that the participants may not have the ability to infer the amplitude and frequency features from the stimulus evoked neural responses; instead, they might have been choosing the interval with the higher number of spikes (assuming vibrotactile stimulus features are encoded in a rate code) as the response in each trial.

Is it possible that humans are capable of making inference about the stimulus features, but rely on the energy or the intensity cue because these cues are always consistent with the change in frequency in FD tasks? This could explain the findings of Harris et al. (2006) where the authors secretly modified the amplitude feature of the stimulus during the FD task: The naïve participants assumed that the amplitude in the two intervals was equal, which would imply that the difference in the neural responses evoked by the stimuli in the two intervals is due to the manipulation of the target feature of the vibration, i.e. the frequency of the stimulus. Hence, the participants were deceived by this secret manipulation!

We predicted that if participants were aware of this manipulation, then during the discrimination task they would infer the target feature from the stimuli, and the manipulation would no longer affect their discrimination performance. However, if participants only use the energy feature to discriminate between vibrations, they would be unable to detect the manipulation, and furthermore their performance on AD and FD task would be identical.

To investigate this, we conducted three different vibrotactile discrimination experiments. In the first experiment we tested participants' ability to discriminate vibrotactile stimuli separately based on the amplitude and frequency. For both AD and FD tasks, we chose the standard-stimulus amplitude and frequency combinations such that two experimental blocks had the same energy. We reasoned that if participants use only the energy cue to discriminate between vibrations, the experimental blocks with identical energy, irrespective of the task, should yield identical discrimination thresholds. In the second experiment, we tested whether surreptitiously changing the non-target feature affects the performance of the participants as reported by Harris et al. (2006). In the third experiment we tested whether participants can selectively identify the feature that changes in each trial. To test our prediction that participants' assumption of equal non-target feature in

both intervals of every trial leads to the effect that was reported by Harris et al. (2006), we retested participants on the second experimental paradigm after they have been trained in a paradigm identical to the third experiment. We also simulated the first and the second experiment using the ideal observer model and compared the human performance to the ideal observer's performance to gain insight into stimulus feature inference. The data obtained in the current study support our prediction that human participants are capable of inferring the stimulus features; however, they primarily rely on the energy cue to discriminate vibrotactile stimuli.

4.2 Methods

4.2.1 Behavioural methods

4.2.1.1 Participants

Two groups of normally sighted neurologically healthy naïve participants were tested on different vibrotactile discrimination tasks. One group participated only in experiment 1, and the other group participated in experiments 2 and 3. We ensured that the participants did not have any injuries or calluses on their tested fingertip. Also, all participants were screened (based on their self-report) for diabetes, dyslexia, learning disabilities, hearing impairment, or any nervous system disorders. Diabetes causes peripheral neuropathy and slows action potential conduction (Hyllienmark et al., 1995), and dyslexia adversely affects tactile acuity (Grant et al., 1999). The McMaster University Research Ethics Board approved all procedures implemented in this study. All participants signed the informed consent form, and were remunerated (cash or course-credit) for their participation.

4.2.1.2 Vibrotactile stimulation

Because this section is identical to that mentioned in the previous chapter, please refer to section 3.2.2 (in Chapter 3). In all the experiments, the stimulus duration was set to 10 taps (i.e., 9.5 cycles) at respective stimulus frequencies.

4.2.1.3 General procedure

Procedures specific to each experiment are described in the corresponding subsections below; however, certain procedural steps that were identical for all the experiments are described here. First, we acquainted the participants with the equipment and the response unit. Next we conducted a "force practice" task, which familiarized the participants with

the permissible level of forces (between 20-50 g) that they could exert on the vibrotactile probe (during the experiment, forces beyond the permissible levels lead to a force warning and exclusion of the trial). Finally, we explained each task to the participants and answered their questions (if any) about the task. The participants repeated the instructions back to us to demonstrate full comprehension of the task instructions. We started the testing protocol once we were satisfied that the participants had fully understood the instructions. In all experiments, the stimuli were delivered to the participants' distal pad of their left index finger.

4.2.1.4 Experiment 1

To investigate whether humans use energy as the discrimination-feature to perform vibrotactile amplitude (AD) and frequency discrimination (FD) tasks, we tested the group-one participants (total 24: 9 male and 15 female; age range: 18.42 – 23.14 years; median age: 21.06 years), of which 17 were right-hand dominant (based on questions modified from the Edinburgh Handedness Inventory; Oldfield, 1971).

Perceptual tasks: All participants were tested on a total of 10 testing blocks. Because participants' ability to detect vibrotactile stimuli presumably affects their performance in the FD task (LaMotte and Mountcastle, 1975), first we conducted two blocks of a vibrotactile threshold detection (TD) task. Next, we conducted 4 blocks (each) of AD and FD tasks. To counterbalance, we tested the odd numbered participants first on FD and then on AD tasks, and reversed the sequence for the even numbered participants.

TD task: In two separate 40-trial blocks we tested each participant's ability to detect a 10-tap vibration delivered at 20Hz and 40Hz stimulus frequencies. The odd numbered participants were tested first in the 40Hz and then in the 20Hz block, and the reverse was true for the even numbered participants. During each trial, we presented the stimulus in 1 of 2 intervals identified by auditory beeps, and using the response unit the participants reported the interval in which they perceived the stimulus. We determined each participant's detection threshold by adaptively adjusting the amplitude of the stimulus ranging between 1μm and 100μm. In each block, prior to the testing phase every participant completed 20 practice trials during which the participant received an auditory feedback after each response.

AD task: In four different 50-trial blocks we tested each participant's ability to discriminate two 10-tap vibrations (representing each interval of a trial) delivered at 20Hz and 40Hz (2 blocks each) stimulus frequencies. For every trial, in 1 of 2 intervals we presented the standard stimulus at 40μm or 80μm (depending on the testing condition) stimulus amplitude, and in the other we presented the comparison stimulus at amplitudes that are a certain percentage lower than the standard stimulus. To estimate each participant's AD threshold, we adaptively adjusted the difference in amplitude of the stimulus ranging between 1% and 100%. The stimulus frequency was identical in the two intervals. We created 4 sets of standard stimulus (in 4 separate blocks), shown in Table

4.2; two of which (20Hz-80μm, 40Hz-40μm) were iso-energy (i.e. equal $a \times f$) combinations.

FD task: The procedure and the number of trials for the FD task were identical to that of the AD task; here we manipulated the frequency in lieu of the amplitude feature of the vibrotactile stimulus. In separate blocks, we estimated each participant's FD threshold at 20Hz and 40Hz stimulus frequencies where the amplitude of the standard stimulus was either 40μm or 80μm (depending on the testing condition). Thus we tested each participant on all 4 sets of standard stimulus conditions (as shown in Table 4.2).

Psychophysical method: To efficiently choose the stimulus levels, and to quantify the participants' performance, we used the Bayesian adaptive method as described in Chapter 2 (section 2.2.4). In this method, which is a modified version of the ψ method (Konstevich and Tyler, 1999), we considered each participants' psychometric function ($P_c(l)$, i.e. the probability of a correct response as a function of stimulus level, l) as a mixture of a cumulative normal function and a lapse rate term (δ):

$$P\left(c_l \mid \theta, \beta, \delta\right) = \frac{\delta}{2} + \left(1 - \delta\right)\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\frac{d'}{\sqrt{2}}} \exp\left(-\frac{x^2}{2}\right)dx$$

$$\text{where, } d' = \left(\frac{l}{\theta}\right)^{\beta} \text{ such that } d' = 1 \text{ when } l = \theta.$$

Therefore, the threshold stimulus level is the θ-parameter, and it corresponds to 76% correct response probability (see Fig. 4.1). For the TD task, the θ-parameter refers to the stimulus amplitude values in microns, whereas for the discrimination tasks, this parameter refers to the percent difference of the target feature of the standard stimulus (i.e. standard stimulus' amplitude value for AD, and frequency value for FD task). The slope of the psychometric function is represented by the β-parameter.

We considered all the psychometric function parameters (θ, β, and δ) as unknown parameters and initiated the algorithm with a uniform prior probability distribution over a wide range of parameter values – threshold (θ: 1 to 100μm, or 1% to 100%), slope (β: 1 to 15), and lapse rate (δ: 0.01 to 0.1). Using different combinations of these parameter values, the algorithm generated several thousand psychometric functions, and after each trial the algorithm updated the likelihood of the psychometric functions. During the testing session, based on the participant's response after every trial, the algorithm estimated the expected information gain associated with each stimulus level; the stimulus level that predicted the highest gain value was chosen as the stimulus level for the next trial. At the end of each experimental block, this iterative process yielded a psychometric function that most closely represented the participant's performance. Because each psychometric function is a combination of three parameters, we marginalized over the β- and the δ-parameters to generate the posterior probability density function (PDF) for each participant's θ-parameter (see Fig. 4.1C).

4.2.1.5 Experiment 2

To investigate whether the participants' AD and FD performances are affected by surreptitious manipulation of the non-target (*f* and *a*, respectively) parameters, we tested the group-2 participants (7 male and 5 female; age range: 19.69 – 23.29 years; median age: 21.94 years; 11 right-hand dominant) in a series of experimental conditions.

Perceptual tasks: Whereas all participants of group-2 were tested on 3 common tasks (a block each of TD, AD, and FD task), 6 odd-numbered participants were further tested on 6 blocks of AD_{vF} (a variation of the AD) task, and similarly, 6 even-numbered participants were further tested on 6 blocks of FD_{vA} (a variation of the FD) task. Thus, every participant of this group was tested on 9 different experimental blocks.

TD task: For the reason mentioned in section 2.2.1.4 (subsection, perceptual tasks), we estimated each participant's detection threshold. The current testing protocol, procedure, and the number of trials were identical to those implemented in experiment 1 except the stimulus frequency was set to 30Hz.

AD task: Similar to the AD task in experiment 1, we estimated each participant's ability to discriminate two vibrations. The amplitude of the standard stimulus was 30μm and it was always higher than that of the comparison stimulus; the frequency of both stimuli representing each interval of the trials was 30Hz. The current testing protocol, procedure, and the number of trials were identical to that implemented in experiment 1.

FD task: The FD task implemented here (including the procedure, testing protocol, and the number of trials) is identical to that mentioned in experiment 1. We estimated each participant's ability to discriminate two vibrations based on the difference in frequencies. The frequency of the standard stimulus was 30Hz and it was always higher than that of the comparison stimulus; the amplitude of both stimuli was 30μm.

AD_{vF} (amplitude discrimination with variable frequency) task: In this 2-interval forced choice (2IFC) task, both vibrations representing two intervals of every trial contained 10 taps. Conceptually the AD_{vF} task is identical to the AD task mentioned above; however, in the AD_{vF} task the frequency feature of the vibrations (i.e. the non-target feature) in each interval were secretly manipulated. In 6 different conditions, the frequency difference between intervals were 2Hz, 4Hz, 6Hz, 8Hz, 10Hz, and 12Hz, which corresponded to a difference of 6.66%, 13.33%, 20%, 26.66%, 33.33%, and 40% of 30Hz. This difference in frequencies between intervals created two sets of stimuli: congruent (+) and incongruent (–) (see Table 4.1). In the congruent set of stimuli, the standard stimulus (which was always higher in amplitude than the comparison stimulus) was present at a higher frequency than the frequency of the comparison stimulus: for example, at +20% frequency difference between to intervals, the interval containing the standard stimulus was delivered at 33Hz and the comparison stimulus was delivered at

27Hz. Conversely, in the incongruent set of stimuli, the standard stimulus was presented at a lower frequency than that of the comparison stimulus; for example, at –20% frequency difference between to intervals, the interval containing the standard stimulus was delivered at 27Hz and the comparison stimulus was delivered at 33Hz. Each experimental block consisted of 100 trials; an equal number of trials were randomly chosen from congruent and incongruent set of stimuli. To estimate the participants' AD_{vF} threshold at each non-target-difference condition, we adaptively adjusted the difference in amplitude of the stimulus ranging between 1% and 100% for each set of stimuli (i.e. congruent and incongruent).

FD_{vA} (frequency discrimination with variable amplitude) task: Procedurally, the FD_{vA} task is identical to the AD_{vF} task, except we secretly manipulated the amplitude feature of the standard and comparison stimuli. We tested each participant on 6 different conditions where the amplitude difference in both intervals were 2µm, 4µm, 6µm, 8µm, 10µm, and 12µm, which corresponded to a difference of 6.66%, 13.33%, 20%, 26.66%, 33.33%, and 40% of 30µm. Similar to the estimation of FD thresholds, we determined the participants' FD_{vA} threshold at each non-target-difference condition by adaptively adjusting the difference in frequency of the stimulus ranging between 1% and 100% for each set of stimuli (i.e. congruent and incongruent) (see Table 4.1).

Psychophysical method: To conduct the experiment and to estimate each participant's TD, AD, FD, and AD_{vF} or FD_{vA} thresholds, we implemented the Bayesian adaptive method mentioned in the previous experiment. In case of AD_{vF} and FD_{vA} tasks, although we secretly manipulated the non-target feature of the standard and comparison stimuli, procedurally we treated these experimental conditions identically to the AD and FD conditions. We reasoned that if the non-target feature manipulation affected the participants' response behaviour, the adaptive nature of the psychophysical method would accordingly adjust the stimulus levels of the target feature such that the participants' performance is reflected in their discrimination thresholds. For example, if the non-target-feature manipulation helped the participants to discriminate between the standard and the comparison stimuli, the participants would respond correctly, and in turn our adaptive algorithm would reduce the difference in the stimulus level of the target feature between the standard and the comparison stimuli in an attempt to increase the level of task difficulty. Conversely, if the non-target-feature manipulation hindered the participants' ability to discriminate between the standard and the comparison stimuli, the participants would respond incorrectly, which would cause the adaptive algorithm to increase the difference between the standard and the comparison stimuli and alleviate the difficulty of the task.

4.2.1.6 Experiment 3

The goal of this experiment was to investigate whether humans can identify the features (amplitude and frequency) of a vibrotactile stimulus. To this end, we tested all the group-2 participants who were also tested in experiment 2 in a previous testing session.

Perceptual task: In each trial we presented two 10-tap vibrations where the amplitude-frequency combination of the first vibration was set to 30μm-30Hz. The amplitude-frequency combination of the second vibration was such that its energy ($a \times f$) was always higher than that of the first vibration by 20%. Hence, two possible combinations for the second vibrations were – a) high-amplitude trials, where the amplitude and frequency combination was 36μm and 30Hz, respectively, and b) high-frequency trials, where the combination changed to 30μm and 36Hz. We asked the participants to report, by pressing respective buttons of the response unit, the feature (amplitude or frequency) that they identified was higher in the second compared to the first interval.

Before conducting the testing phase, we presented 40 practice trials that included 20 trials of both types (high-amplitude and high-frequency). In the practice trials the second vibration had 46.66% higher energy than the first vibration. Participants were provided feedback after every response.

Psychophysical method: To quantify the participants' ability to detect the stimulus feature that was higher in the second vibration, we randomly presented 50 trials each where the second vibration either had higher amplitude or higher frequency than that of the first vibration.

4.2.1.7 Training regimen and test of effectiveness of training

After determining participants' ability to identify the features (amplitude and frequency) of a vibrotactile stimulus when the energy difference between the first and the second vibration was 20%, we initiated a training regimen using a procedure identical to experiment 3. The training regimen consisted of 24 blocks of 100 trials each, which we conducted in two sessions on consecutive days. Similar to experiment 3, all training blocks consisted of 50 high-amplitude and high-frequency trials each. The participants received feedback after every trial.

During the training regimen, the participants' percentage of correct responses in each block determined the energy difference for the next training block. The energy difference between the two vibrations in each trial increased (making the task easier), decreased (making the task harder), or stayed the same, if the participants' percent correct was $<70\%$, $\geq 85\%$, or between 70% and 85%, respectively. To start the training regimen we applied this criterion after calculating the participant's percent-correct in experiment 3 (i.e. participant's first performance with an energy difference of 20%).

After conducting the training regimen (i.e. all 24 blocks), we tested the effectiveness of the training regimen by conducting – a) experiment 3 once again, but without the practice session, and b) experiment 2 on the following day.

4.2.1.8 Data analyses

We performed statistical tests (ANOVA, *t*-test), using SPSS v20 (IBM) for Macintosh with an alpha level of 0.05, on different performance measures.

Performance measures: To analyze the performance of the participants in experiment 1, we used two measures – 1) the mode of the posterior PDF of the θ-parameter, and 2) the energy measure. The posterior PDF of the θ-parameter represents the probability that each stimulus level is the threshold of a participant; we took the mode of the posterior PDF as the participant's threshold; this is our best-estimate of the participant's detection amplitude (μm) for the TD task, and discriminable percent difference of target feature for AD and FD tasks. As shown in Table 4.2, two standard stimuli (20Hz-80μm, 40Hz-40μm) have equal energy (1600 energy units); therefore, to compare each participant's discrimination performance for these two iso-energy combinations in AD as well as FD tasks, we calculated the participant's discrimination threshold in energy terms by multiplying the threshold estimate with the non-target parameter. This operation was necessary because both the target and the non-target features were different in these iso-energy combinations.

To analyze the discrimination performance of the human participants in experiment 2 (i.e. the "non-target Feature Affected Discrimination" or ntFAD performance), both before and after the training regimen, we derived the Best-Fit Line (BFL_{ntFAD}), which represented each participant's performance change from the most incongruent (i.e. –40%) to most congruent (+40%) combination of target and non-target features of the two vibrations in each trial of all the AD_{vF} and FD_{vA} blocks. In other words, we parameterized a BFL_{ntFAD} as a function of non-target feature change, x_{ntFAD}:

$$\theta_{\text{ntFAD}} = m_{\text{ntFAD}} \times x_{\text{ntFAD}} + c$$

Here, m_{ntFAD} is the ntFAD slope, c is the intercept, and θ_{ntFAD} is the discrimination threshold due to the manipulation.

We estimated the BFL_{ntFAD} that best fit each participant's performance (correct and incorrect answers at each tested stimulus level $\{cr_l, ic_l\}$) across all relevant discrimination blocks (i.e. AD for AD_{vF} and FD for FD_{vA}). To do so, we considered BFL_{ntFAD} with ntFAD slope values ranging from $m_{ntFAD} = 0$ to 1, and c ranging from 1 to 100 microns. We then determined the maximum likelihood BFL_{ntFAD}:

$$\hat{BFL}_{\text{ntFAD}} = \underset{m_{\text{ntFAD}},c}{\arg\max}\left[\prod_{x_{\text{ntFAD}}} P\left(\left\{cr_l, ic_l\right\}_{x_{\text{ntFAD}}} \mid m_{\text{ntFAD}}, c\right)\right]$$

To derive the best-fit line, instead of using a point estimate from each participants posterior PDF of the θ-parameter, we believe that by utilizing the participants' every

correct and incorrect response provided us a reliable estimate of each participants' BFL_{ntFAD} because the current approach weights the testing blocks during which the participants were attentive and responded most consistently.

We analyzed the participants' ability to detect the changing features of a vibration in experiment 3 by using the percentage of correct responses as the dependent variable.

4.2.2 Ideal observer simulations for experiment 1 and 2

In chapter 2, we demonstrated that the ideal observer can perform AD and FD tasks. To conduct the AD and FD tasks, we created 10 simulated subjects, and to conduct the AD_{vF} and FD_{vA} tasks, we created 6 simulated subjects. To implement AD_{vF} and FD_{vA} tasks that simulated pre-training performance, we changed the non-target feature (i.e. frequency and amplitude, respectively) values in the two intervals of each trial only in the generative section (i.e., the encoding section) of the model; however, during stimulus decoding we set the non-target feature value to 30Hz or 30µm (depending on the task) for both intervals. Therefore, the simulated subjected "assumed" that the non-target feature was identical in both intervals. To simulate the post-training AD_{vF} and FD_{vA} tasks, we provided the encoder and the decoder with identical non-target feature values; this replicated the behavioural condition in which the human subjects knew that the non-target feature values were not the same in the two intervals. To determine the BFL_{ntFAD} and ntFAD slope, we used the linear fitting algorithm provided in the LabVIEW software package.

4.3 Results

4.3.1 Experiment 1

LaMotte and Mountcastle (1975) observed that during FD tasks stimulus amplitudes close to (approximately < 3 times) the detection threshold adversely affect human and non-human primates' performances. Therefore, to determine whether the lowest stimulus amplitude (i.e. 40µm) of the standard stimuli that we delivered during the FD task was at least 3-times the participants' detection threshold, we conducted TD tasks at 20Hz and 40Hz stimulus frequencies. The mean detection thresholds for 20Hz and 40Hz stimuli were 3.6µm (SD: 1.2) and 2.2µm (SD: 0.91), respectively. Therefore, for all participants the 40µm stimulus amplitude was higher than 3-times their detection threshold.

4.3.1.1 Effect of target and non-target feature manipulation

To examine the effect of increasing the target and the non-target features of the standard-stimulus in the AD and FD tasks, we performed two separate full factorial ANOVAs with standard-stimulus target feature (amplitude for AD: 40µm, 80µm; frequency for FD: 20Hz, 40Hz) and non-target feature (frequency for AD: 20Hz, 40Hz; amplitude for FD: 40µm, 80µm;) as within-subject factors.

AD task: The ANOVA on the AD absolute thresholds revealed a significant main effect of standard-stimulus amplitude [$F(1,23) = 109.27$; $p < 0.001$; $\eta^2 = 0.826$] but not of standard-stimulus frequency ($p = 0.99$), and there was no significant interaction between these factors ($p = 0.75$) (see Fig. 4.2A). Bonferroni corrected comparisons of AD thresholds showed that irrespective of stimulus frequency (20Hz or 40Hz), increasing standard-stimulus amplitude significantly increases the absolute AD thresholds ($p < 0.001$ at both stimulus frequencies). A separate ANOVA with the above-mentioned factors on the percent-difference thresholds revealed no difference in thresholds, which was ~20% for all standard-stimulus conditions (see Fig. 4.2A).

FD task: The ANOVA on the FD absolute thresholds revealed a significant main effect of standard-stimulus frequency [$F(1,23) = 68.00$; $p < 0.001$; $\eta^2 = 0.747$] but no main effect of amplitude ($p = 0.1$), and the interaction was also not significant ($p = 0.1$) (see Fig. 4.3A). To determine whether increasing the standard-stimulus frequency (i.e. the target feature) significantly increased the FD absolute thresholds at both stimulus amplitudes, we performed Bonferroni corrected comparisons of FD thresholds at each stimulus amplitudes, which revealed that irrespective of the amplitude of the stimuli there was a significant increase in threshold with increasing standard-stimulus frequency. An ANOVA on the FD percent-difference thresholds, with the same between-subject factors as mentioned above, revealed a significant main effect of standard-stimulus frequency [$F(1,23) = 5.27$; $p < 0.05$; $\eta^2 = 0.186$] such that discrimination thresholds at the 40Hz were better than those at 20Hz stimulus by ~2Hz (see Fig. 4.3A).

4.3.1.2 Performance on iso-energy standard-stimulus combinations

In the current experiment we consider two iso-energy combinations – 20Hz-80µm, and 40Hz-40µm, which corresponds to 1600 energy (i.e., $a \times f$) units. Although AD and FD tasks presumably require discrimination of different features of the vibrations, converting the threshold to energy units allow us to compare different tasks and standard-stimulus combinations. Therefore, to examine whether humans compare vibrations using the energy feature, we performed a full-factorial ANOVA with iso-energy combination (20Hz-80µm, and 40Hz-40µm) and task (AD and FD) as the within-subject factors. The analysis revealed no main effect of iso-energy combination ($p = 0.83$) but a significant main effect of task [$F(1,23) = 4.37$; $p = 0.048$; $\eta^2 = 0.16$] (see Fig. 4.4); participants required marginally higher energy to discriminate vibrations based on the amplitude rather than frequency of those vibrations.

4.3.1.3 Test for performance difference due to type of discrimination task

To further examine the effect of task on vibrotactile discrimination performance, we compared the percent-difference threshold for the AD and FD tasks where the standard-stimulus was 40Hz-40µm. The threshold for FD task was on average ~4% lower than that of the AD task, and a paired samples *t*-test shows a trend towards statistical difference in tasks ($p = 0.07$) (see Fig. 4.5A).

4.3.2 Experiment 2

Similar to experiment 1, to examine whether the stimulus amplitude for frequency discrimination (on average 30µm) was at least 3-times that of participants' detection threshold, we tested participants' ability to detect a 30Hz stimulus in a TD task. The mean detection threshold was 3.4µm (SD: 1.23), and for all participants, the 30µm stimulus amplitude was higher than 3-times their detection threshold.

4.3.2.1 Secretly changing non-target feature affects discrimination threshold

While group-2 participants were performing respective discrimination tasks we surreptitiously manipulated the non-target feature (frequency for AD_{vF}, and amplitude for FD_{vA} sub-groups) such that we obtained congruent and incongruent stimulus combination (see 2.1.5, AD_{vF}, FD_{vA} tasks). To characterize participants' performance in all these conditions, we derived the ntFAD slope (see 2.1.8, performance measure) for each participant. We predicted that if participants are unaffected by the secret manipulation then the slope should be zero; however, the mean ntFAD slope calculated from the data obtained from the group-2 participants suggests that the secret manipulation of the non-target feature affected their discrimination threshold (see Table 4.3). The incongruent condition increased whereas the congruent condition decreased the discrimination threshold in all participants; a one-sample *t*-test verified that the mean ntFAD slope was significantly different from 0 ($t_{(11)} = -4.215$, $p < 0.01$). Separate one-sample *t*-tests on AD_{vF} and FD_{vA} sub-groups confirmed the overall effect (AD_{vF} sub-group: $t_{(5)} = -3.995$, $p = 0.01$; FD_{vA} sub-group: $t_{(5)} = -2.602$, $p = 0.048$). Interestingly, the effect of non-target feature manipulation was stronger in the AD_{vF} compared to the FD_{vA} sub-group (ntFAD slope difference for $AD_{vF} - FD_{vA}$: –0.199), and an independent samples *t*-test shows a trend towards statistical significance ($p = 0.07$) (see Fig. 4.6A).

4.3.2.2 Test for performance difference due to type of discrimination task

In the current experiment we also estimated participants' discrimination threshold in AD and FD tasks where the non-target feature was identical in both intervals; the frequency-amplitude combination of the standard stimulus in both tasks was 30Hz-30μm. To examine whether the type of task determined the discrimination thresholds, we compared the participants' percent discrimination threshold for AD and FD task. Interestingly, the thresholds were almost identical (AD mean: 19.04%, FD mean: 19.82%) (see Fig. 4.5B).

4.3.3 Experiment 3

All group-2 participants were also tested in this experiment. Here we tested whether participants could identify the feature that was increased in the second interval of each trial by 6μm or 6Hz (i.e. 20% increase in energy) depending on "high-amplitude" or "high-frequency" trial respectively. We predicted that if participants cannot detect the increased feature then the percentage of correct responses should be 50; however, a one-sample *t*-test on the participants' percentage of correct responses revealed a mean of 72.6 (SD: 6.98), which was significantly different from 50 ($t_{(11)} = 11.194$, $p < 0.001$).

4.3.4 Effect of training regimen

To determine the effectiveness of the training regimen, we re-tested participants on experiment 3. A paired-samples *t*-test on participants' pre- and post-training performance showed an increase in percentage of correct responses after training by 4.5; this difference shows a trend towards statistically significant ($p = 0.069$) (see Fig. 4.7).

After conducting the training regimen we also re-tested participants on all the tasks in experiment 2. First, to re-examine whether the performance of participants on AD and FD tasks (where the non-target feature was identical in both intervals) after the training differed due to type of discrimination task, we performed an ANOVA with training (pre, and post) and task (AD, FD) as two within-subject factors. The analysis revealed a main effect of training [$F(1,20) = 17.77$; $p < 0.01$; $\eta^2 = 0.618$] but no main effect of task ($p = 0.44$) and no significant interaction between the factors ($p = 0.16$). Whereas the pre-training difference in AD and FD performance was negligible (see section 4.3.2.2), after training the FD threshold was ~3% lower than the AD threshold; however, this difference was not significant (Bonferroni corrected $p = 0.13$) (see Fig. 4.8).

Finally, to test whether the ntFAD slopes changed after training, we performed a paired-samples *t*-test on the participants' pre- and post- ntFAD slopes and the test revealed no significant differences in the slopes ($p = 0.47$) (see Fig. 4.9A). The non-target feature was different for both discrimination tasks conducted on two sub-groups of participants;

therefore, to explore whether there was a differential effect of training within each tasks, we performed an ANOVA with training as a within-subjects factor (pre, and post) and task sub-group (AD_{vF}, FD_{vA}) as a between-subjects factor. The analyses revealed no main effect of training ($p = 0.46$) or condition ($p = 0.16$), and there was no significant interaction ($p = 0.24$) (see Fig. 4.9B and 4.10A).

4.3.5 Simulation results

To determine the ideal performance attainable in the AD and FD tasks with a spike-count code, we simulated experiment 1 and 2 using the Bayesian ideal observer (see Chapter 2).

4.3.5.1 Experiment-1 simulation results

Similar to human participants we ran all the conditions on the simulated subjects (see Chapter 2 section 2.2.1.3) and analyzed the absolute AD and FD thresholds in two different ANOVAs.

AD task: The overall trends in performance of the model as a function of amplitude and frequency were similar to those of the human, though the model gave lower thresholds. An ANOVA with standard-stimulus amplitude (40μm, 80μm), and standard-stimulus frequency (20Hz, 40Hz) as between-subjects factor revealed a significant main effect of stimulus amplitude [$F(1,9) = 173.07$; $p < 0.001$; $\eta^2 = 0.951$] and of stimulus frequency [$F(1,9) = 16.63$; $p < 0.01$; $\eta^2 = 0.649$], which is the non-target feature in this task. The interaction of these two factors was also not significant ($p = 0.85$) (see Fig. 4.2B). Similar to human participants, Bonferroni corrected comparison of the simulation AD results showed an increase in threshold with an increase in the stimulus amplitude in both frequency conditions ($p < 0.001$).

FD task: Unlike in AD, the model's trends in the model's performance as a function of amplitude and frequency differed markedly from those of humans. An ANOVA with standard-stimulus frequency (20Hz, 40Hz), and standard-stimulus amplitude (40μm, 80μm) as between-subjects factor revealed a significant main effect of stimulus frequency [$F(1,9) = 5.67$; $p < 0.05$; $\eta^2 = 0.386$] and of stimulus amplitude [$F(1,9) = 62.32$; $p < 0.001$; $\eta^2 = 0.874$]; however, there was no significant interaction between these factors ($p = 0.47$) (see Fig. 4.3B). Interestingly, the ideal observer performed relatively better at the 40Hz compared to the 20Hz standard-stimulus frequency condition. Remarkably, the model performance was clearly worse than the human's at 20Hz-80μm.

Test for performance difference due to type of discrimination task: To test whether the model's percent discrimination threshold for AD and FD tasks at 40Hz-40μm condition are statistically equal, we performed a paired samples *t*-test, which showed a significant

difference in threshold between the two tasks ($t_{(9)} = –6.67$, $p < 0.001$); the FD threshold was higher than the AD threshold by ~6% (see Fig. 4.5C).

4.3.5.2 Experiment-2 simulation results

We simulated experiment-2 by presenting the ideal observer with all the stimulus combinations that we used for human participants. The encoding portion of the model (see Chapter 2 for details) generated the neural responses according to the stimulus combinations; however, the decoding portion "assumed" that the non-target feature was identical in both intervals, which we fixed at 30Hz or 30μm depending on the task (AD_{vF} or FD_{vA}, respectively). To replicate post-training simulations (see section 4.3.4 above), we provided the decoding portion of the model with the complete stimulus information.

We conducted AD_{vF} and FD_{vA} on the same simulated subjects, which allowed us to perform an ANOVA on the ntFAD slope values with the decoder's knowledge (unaware, aware) of the manipulation of the non-target feature and type of task (AD_{vF}, FD_{vA}) as two within-subject factors. The analysis revealed a significant main effect of decoder's knowledge about the non-target feature manipulation [$F(1,5) = 126.53$; $p < 0.001$; $\eta^2 = 0.962$] and a main effect of the type of task [$F(1,5) = 7.76$; $p < 0.05$; $\eta^2 = 0.608$]; the interaction between these factors was also statistically significant [$F(1,5) = 18.03$; $p < 0.01$; $\eta^2 = 0.783$] (see Figure 4.10B). Bonferroni corrected pairwise comparison of slopes from both tasks revealed a significant difference in AD_{vF} and FD_{vA} slopes that resulted from the "pre-training" simulations ($p < 0.001$); there was no difference between "post-training" AD_{vF} and FD_{vA} slopes ($p = 0.59$). Most interestingly, the "pre-training" simulations revealed positive (see Table 4.4) rather than negative slopes, which we obtained from human participants (see Table 4.3).

4.4 Discussion

Can humans infer the features of vibrotactile stimuli? We investigated this question in three different experiments involving several discrimination tasks. The features that we considered in this study were amplitude, frequency, and the product of these two features, which is referred to as energy. The data suggest that human participants are capable of making inferences about the stimulus features; however, they might mostly rely on a feature that combines the amplitude and frequency, for example, the energy feature. We compared the human data to the performance of a spike-count-based ideal observer and found certainly similarities but several intriguing discrepancies, suggestive of the possibility that humans may make use of non-spike-count data (e.g., spike timing) on the FD task.

4.4.1 An increase in the standard-stimulus target feature value increases discrimination threshold

AD task: In separate conditions, we tested participants' ability to discriminate vibrations at 40µm and 80µm standard-stimulus amplitudes, and observed that doubling the standard-stimulus amplitude approximately doubled the participants' absolute AD threshold. However, when the thresholds were expressed in percent difference, there was no difference between the conditions, which supports Weber's law. We found that the AD threshold of ~20% is consistent with previous studies; some of those are: Bhattacharjee et al. (2010), using 20Hz vibrations and 100µm standard-stimulus amplitude, tested blind and sighted participants on an AD task and reported a threshold range of ~20% – 30%. Francisco et al. (2008) used 25Hz vibrations at various standard-stimulus amplitude values to determine participants' Weber fraction; for the amplitude values 50µm and 100µm (which are closer to the amplitudes we considered in our current study) the authors reported a discrimination threshold of ~30% and ~23%, respectively. Goble and Hollins (1993), using a standard stimulus of 25Hz and ~20µm, observed ~16% threshold in human participants. Gescheider et al. (1990) using 25Hz, and later Güçlü (2007) using 40Hz vibrations, tested participants on at a range of standard-stimulus amplitude and reported Weber fractions of 19% and 32%, respectively.

FD task: We conducted four conditions of the FD task, which resulted in threshold values that range between ~13% to ~17%. The absolute threshold of the participants increased as we doubled the standard-stimulus frequency from 20Hz to 40Hz; however, unlike the AD task, the FD threshold did not double. We observed that the threshold for the 20Hz stimulus was higher than that for the 40Hz stimulus, a phenomenon known as the near-miss to Weber's law (Gescheider et al., 1990). Goff (1967) matched subjective intensity of the stimuli and the Weber fraction reported by the author suggests near-miss only at a lower but not at higher amplitude. Because participants presumably use an intensity cue to perform the FD task, several studies matched the subjective intensity of both intervals in every trial during the task. Relative recently, Mountcastle et al. (1990), tested human and non-human primates on several FD tasks that included the conditions where the standard-stimulus frequency was 20Hz and 40Hz, and reported thresholds at 17% and 7%, respectively showing near-miss to Weber's law. These authors also matched the subjective intensity of stimulus in each interval. Goble and Hollins (1994) determined the threshold of their participants at 25Hz standard-stimulus frequency and reported Weber fractions between 0.14 and 0.23, i.e. the threshold range is 14% to 23%. Tommerdahl et al. (2005), using 25Hz standard-stimulus frequency, reported an average Weber fraction of 0.38, which is quite high at 38% threshold. Whereas most of the studies mentioned here reportedly matched the subjective intensity of the stimulus, Harris et al. (2006) matched the energy of the stimuli in each trial. The authors using 32Hz standard-stimulus frequency reported a threshold of ~20%; however, when the authors matched the stimulus amplitude (280µm) in both intervals of a 2IFC FD task, the threshold dropped to ~15%. Therefore, all these studies that used frequencies between 20Hz and 50Hz generally

reported thresholds within a small range between 10% and 25%, which include studies that matched energy or intensity or amplitude in both intervals of every trial. This raises the importance of the amplitude, i.e. the non-target feature, in FD tasks.

4.4.2 Effect of non-target feature in vibrotactile discrimination tasks

When the non-target feature is identical in both intervals: In all the AD or the FD tasks, we did not find any effect of the non-target feature. Gescheider et al. (1990), using 25Hz and 250Hz stimulus frequencies, conducted AD task on a range of standard-stimulus amplitudes and reported that there was no effect of the stimulus frequency on the AD thresholds. Craig (1974) tested participants on an AD task with the stimulus frequency of 160Hz and observed a discrimination threshold of ~23%, which is similar to all the thresholds mentioned above (see section 4.4.1, AD task). However, Goff (1967) conducted FD tasks, at two different amplitude levels, on a range of standard-stimulus frequencies between 25Hz and 200Hz. The author observed that at all frequencies, the FD threshold was lower in the high than in the low amplitude conditions. Whereas Goff (1967) used supra-threshold amplitudes, Kuroki et al. (2013) tested participants on several FD tasks where the standard-stimulus frequency ranged from 15Hz to 240Hz, and the stimulus amplitude was 2 and 6.3 times the participant's detection threshold. The authors reported that increasing stimulus amplitude lowered the FD threshold at all tested frequencies, which provides evidence for an "atonal interval". Nonetheless, increasing amplitude enhanced the performance of the participants in the FD task. However, this enhancement in threshold with an increase in the non-target feature value was limited to the FD tasks only.

When the non-target feature is not identical in both intervals: To the best of our knowledge, Harris et al. (2006) is the only other study that systematically examined the effect of a secret manipulation of the non-target feature on a discrimination task. Because Harris et al. (2006) only investigated the effect of secret manipulation of amplitude in an FD task, we will compare our FD_{vA} results to that of Harris et al.'s (2006) study. Harris et al. (2006) tested 6 participants and showed a very strong effect of amplitude manipulation with an average slope of –0.42 (SD: 0.27), which is much steeper than the slope (mean: –0.14, SD: 0.13) we obtained. Note that, on average AD thresholds vary around ~20% (see section 4.4.1, AD task), which implies that Harris et al.'s participants might not have detected the change even if they were informed about the manipulation. In the current study, we extended the amplitude difference to 40% of the mean amplitude (i.e. 30μm), which might have revealed the secret about the manipulation and resulted in a shallower slope than that from the Harris et al. (2006) study. Moreover, the amplitude that we used were comparatively much lower than that delivered by Harris et al., which suggests that the subjective intensity of their stimulus was much higher than that of our stimulus

(Verrillo, 1969). This might have affected their participants' discrimination performance more strongly.

To complement the FD_{vA} task, we also conducted the AD_{vF} task where we secretly manipulated the frequency of both intervals within each trial. We reasoned that if the participants use energy as the discrimination feature then the slopes of FD_{vA} and AD_{vF} should be identical. We observed that the AD_{vF} slopes were relatively steeper than the FD_{vA} (see Fig. 4.9B).

4.4.3 Do humans use energy to discriminate vibrations?

The results are puzzling because different lines of evidence suggest contradictory features or mechanisms for vibrotactile discrimination. For example, there was no effect of the non-target feature on the discrimination tasks with identical non-target feature in both intervals (Figs. 4.2a, 4.3a). Presumably, if participants were using energy as the discrimination feature, then an increase in the energy (by doubling the non-target feature) should have increased the discrimination threshold. However, within a particular discrimination task (AD or FD) there is no difference in thresholds of standard-stimulus with iso-energy combinations (20Hz-80µm and 40Hz-40µm) (Fig. 4.4), which indicates that humans might be using energy as a feature. Interestingly, the ANOVA that revealed no difference in thresholds obtained from the iso-energy discrimination tasks, also revealed a statistically significant effect of the type of discrimination task (i.e. AD vs FD). Once again, if participants were using the energy feature then irrespective of the task the thresholds should have been similar. For instance, comparison of the relative percent difference in threshold in AD and FD tasks at 30Hz-30µm standard stimulus showed almost identical threshold value (Fig. 4.5b). Although statistically not significant (p = 0.07), a comparison of the relative percent difference in threshold in AD and FD tasks at 40Hz-40µm standard stimulus showed lower threshold for the FD than for the AD task, which is inconsistent with the aforementioned assumption. Interestingly, Goff (1967) conducted FD task and observed that participants' responses did not match the availability of the energy cues, which argues against the energy discrimination argument. Finally, the AD_{vF} and FD_{vA} tasks yielded different slopes, which suggest that the two groups of participants might have been using different discrimination features. Thus, it is unclear whether humans exclusively use energy as a discrimination feature.

4.4.4 Do humans have access to the amplitude and frequency features of vibrotactile stimuli?

The result from experiment 3 strongly suggests that humans are capable of inferring the stimulus amplitude and frequency (Fig. 4.7). Interestingly, all the participant's performance was above chance even before we conducted the training regimen, which

demonstrates that consistently every participant could identify the feature that had a higher value in the second interval of each trial in the 2IFC task.

We interpreted the results reported by Harris et al. (2006), where the authors secretly manipulated the amplitude during the FD task, as an indication that humans were unaware of the amplitude manipulation, and we predicted that if the participants are adequately informed about the manipulation of the non-target feature, we should see evidence for inference making during the AD_{vF} or FD_{vA} task, i.e. reduced effect of the non-target manipulation as quantified by the ntFAD slopes. We assumed that the most effective approach to inform participants about the manipulation would be to train them to identify the stimulus features. In the current study, we found that whereas the post-training ntFAD slopes were lower than that of the pre-training, the difference was not significant in either task. The negative slope implies that the incongruent conditions are strongly affecting the performance of the humans in the AD_{vF} or FD_{vA} task, and that human participants are probably utilizing the energy feature to perform the discrimination tasks. Interestingly, post-training performance revealed a trend towards shallower slope, which implies that the human participants were inferring the target feature to make the appropriate comparison; however, their inference was probably suboptimal, which might be the reason that the pre-training and post-training slopes were not statistically different.

Note that the performance of the participants in experiment 3 before we conducted the training regimen was quite high (an indication that humans can perceive amplitude and frequency separately), and the post-training ntFAD slopes became shallower but not statistically significantly different from the pre-training ntFAD slopes. Therefore, it might be possible that during the post-training session just being aware of the non-target feature manipulation and not the training *per se* caused the shallowness of the slope, i.e. the participants realized that they should infer the features rather than use energy as a cue to perform in those blocks. Nevertheless, in light of these results we speculate humans can infer the amplitude and frequency features of a vibrotactile stimulus.

4.4.5 Comparison of human and simulation results

We simulated experiment 1 and 2 using the ideal observer model, which (by virtue of its architecture) makes inferences only about the amplitude and frequency feature of the stimulus. Whereas some of the simulation results matched the human performance, we also observed 4 main differences in simulation and human performance results.

First difference: Increasing non-target stimulus affects ideal observer but not human participants. We tested the ideal observer on all the stimulus condition that we conducted in experiment 1. The simulation results revealed that the ideal observer's performance enhanced in the AD, whereas worsened in the FD task, when the non-target feature was doubled. However, the non-target feature in any of the discrimination tasks did not affect the human participants. This discrepancy might have originated from the approach we

implemented to generate cortical neural responses – we assumed that the vibrotactile stimulus-evoked neural responses that are generated in the periphery pass through a Poisson sampler (see Chapter 2, section 2.2.1.2). Because fano factor (i.e. the ratio of the mean to the variance) of Poisson distribution is 1, increasing the mean increases the variance, which reflects in the sampling process. During an AD task, when we increased the frequency, (rather counter-intuitively) the firing rate in each afferent decreased, which was fed through the Poisson sampler. This in turn reduced the variability of the population response of each simulated vibration corresponding to each interval of a trial and the ideal observer predicted the amplitude more accurately, which reflected on the simulation results. On the contrary, during the FD task, we increased the amplitude of the simulated vibration, which (as expected) increased the neural responses in each afferent. To simulate the cortical responses, we passed the afferent responses through the Poisson sampler and (as mentioned above) the sampling variability increased, which affected the inference accuracy of the ideal observer. This in turn affected the simulation results.

Second difference: The ntFAD slopes of AD_{vF} and FD_{vA} tasks in human participants showed negative values, which we predicted based on Harris et al. (2006) results, and based on the premise that human participants are unaware of the secret manipulation. Antithetical to human performance, the slope of AD_{vF} and FD_{vA} tasks in the ideal observer showed positive values. Similar to the reason mentioned above, during the congruent trials (i.e. the high amplitude interval was delivered at high frequency) the high amplitude interval elicited fewer spikes than the interval with low frequency and amplitude. Because the ideal observer was "unaware" of the frequency manipulation, it incorrectly inferred that a low amplitude stimulus must have been presented; however, during decoding the low amplitude interval, which also had the low frequency, the generative model elicited more spikes than the high amplitude interval. The ideal observer inferred that only a high amplitude stimulus could elicit this high number of spikes, which lead to the incorrect choice of interval as the response. The simulated subject erroneously chose the low amplitude interval as the "perceived high amplitude" interval. Conversely, during the simulation of incongruent trials, the high amplitude stimulus was presented at a low stimulus frequency. Because the low frequency interval evoked more spikes, the ideal observer inferred this interval as the "perceived high amplitude" interval, which was also the correct answer. Therefore, the thresholds for the incongruent blocks of trials were lower than the congruent trials, which created positive-value rather negative-value slopes. Why did the FD_{vA} task showed positive slopes? Because the ideal observer has all the information including the precise response characteristics of each afferent, the ideal observer "knows" that a high frequency stimulus evokes fewer spikes compared to a low frequency stimulus. During congruent conditions of FD_{vA} task, high frequency stimuli were delivered at high stimulus amplitudes. The high stimulus amplitude evoked more spikes in the high frequency interval than the low frequency interval. Based on its knowledge, the ideal observer incorrectly inferred that the stimulus must have been presented at a low frequency. Conversely, the low frequency low amplitude interval evoked fewer spikes, which is consistent with the ideal observer's knowledge of "high perceived frequency". Thus, the simulated subject incorrectly

chooses the "high perceived frequency" as the response, which leads to increase in threshold for the congruent trials. In the incongruent trials, because the low amplitude and high frequency stimulus pair evokes few spikes, the ideal observer correctly infers that interval as the "high perceived frequency", which leads to correct responses and lower thresholds. Interestingly, human participants do not show this behaviour, which suggests that human participants do not have access to the afferent responses.

Third difference: Before performing the simulations that represented the post-training tests, we predicted that the ideal observer by "knowing" about the manipulation of the non-target feature it would correctly infer the interval with high target feature stimulus and veridically perform the AD_{vF} and FD_{vA} tasks. Our simulation results supported this prediction; however, the human performance did not.

Fourth difference: Overall, the result of all the simulation tasks showed a higher threshold for the FD than for the AD tasks. However, a scrutiny of the human AD and FD thresholds show a higher threshold for the AD than for the FD tasks. More importantly, the ideal observer had higher thresholds than did the human on the FD task under certain conditions. Because an ideal observer, by definition, has all the information to perform a task most efficiently, the thresholds achieve by the ideal observer serves as the best possible performance. The current discrepancy suggest that the human participants had more information that did the ideal observer, which lead the human participants to attain a higher level of performance compared to the ideal observer.

4.5 Conclusion: Speculations on the neural code

The ideal observer inferred both features of the vibrotactile stimulus – amplitude and frequency – using the stimulus evoked number of spikes. Quantitatively, the AD thresholds of the ideal observer were better than human thresholds on the same task; however, the FD thresholds of the ideal observer were much higher compared to those of the human participants. We speculate that humans probably use both temporal and rate codes to perform the FD tasks. If human participants were exclusively using the rate code, their performance should have been worse than the ideal observer (as we see for the AD tasks). However, the results of the experiment in which we secretly manipulated the non-target feature suggest that the human participants are unable to utilize the temporal code exclusively; we argue that exclusive use of temporal code should eliminate the effect of amplitude in the FD_{vA} task because the temporal structure of the stimulus was identical in both, congruent and incongruent, conditions.

4.6 References

Bhattacharjee A, Ye AJ, Lisak JA, Vargas MG, Goldreich D (2010) Vibrotactile masking experiments reveal accelerated somatosensory processing in congenitally blind braille readers. J Neurosci 30:14288-14298.

Bensmaïa SJ (2008) Tactile intensity and population codes. Behav Brain Res 190:165-173.

Craig JC (1974) Vibrotactile difference thresholds for intensity and the effect of a masking stimulus. Percept Psychophys 15:123–127.

Francisco E, Tannan V, Zhang Z, Holden J, Tommerdahl M (2008) Vibrotactile amplitude discrimination capacity parallels magnitude changes in somatosensory cortex and follows Weber's Law. Exp Brain Res 191:49-56.

Freeman AW, Johnson KO (1982) A model accounting for effects of vibratory amplitude on responses of cutaneous mechanoreceptors in macaque monkey. J Physiol 323:43-64.

Gescheider GA, Bolanowski SJ Jr, Verrillo RT, Arpajian DJ, Ryan TF (1990) Vibrotactile intensity discrimination measured by three methods. J Acoust Soc Am 87:330-338.

Goble AK, Hollins M (1993) Vibrotactile adaptation enhances frequency discrimination. J Acoust Soc Am 96:771-780.

Goble AK, Hollins M (1994) Vibrotactile adaptation enhances frequency discrimination. J Acoust Soc Am 96:771-780.

Goff GD (1967) Differential discrimination of frequency of cutaneous mechanical vibration. J Exp Psychol 74:294-299.

Grant AC,Zangaladze A,Thiagarajah MC,Sathian K (1999) Tactile perception in developmental dyslexia: a psychophysical study using gratings. Neuropsychologia 37:1201–1211.

Güçlü B (2007) Deviation from Weber's law in the non-Pacinian I tactile channel: a psychophysical and simulation study of intensity discrimination. Neural Comput 19:2638-2664.

Güçlü B, Dinçer SM (2013) Neural coding in the Non-Pacinian I tactile channel: a psychophysical and simulation study of magnitude estimation. Somatosens Mot Res 30:1-15.

Harris JA, Arabzadeh E, Fairhall AL, Benito C, Diamond ME (2006) Factors affecting frequency discrimination of vibrotactile stimuli: implications for cortical encoding. PLoS One 1:e100.

Harvey MA, Saal HP, Dammann JF 3rd, Bensmaïa SJ (2013) Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. PLoS Biol 11:e1001558.

Hernández A, Zainos A, Romo R (2000) Neuronal correlates of sensory discrimination in the somatosensory cortex. Proc Natl Acad Sci U S A 97:6191-6196.

Hyllienmark L, Brismar T, Ludvigsson J (1995) Subclinical nerve dysfunction in children and adolescents with IDDM. Diabetologia 38:685– 692.

Johnson KO (1974) Reconstruction of population response to a vibratory stimulus in quickly adapting mechanoreceptive afferent fiber population innervating glabrous skin of the monkey. J Neurophysiol 37:48-72.

Kontsevich LL, Tyler CW (1999) Bayesian adaptive estimation of psychometric slope and threshold. Vision Res 39:2729-2737.

Kuroki S, Watanabe J, Nishida S (2013) Contribution of within- and cross-channel information to vibrotactile frequency discrimination. Brain Res 1529:46-55.

LaMotte RH, Mountcastle VB (1975) Capacities of humans and monkeys to discriminate vibratory stimuli of different frequency and amplitude: a correlation between neural events and psychological measurements. J Neurophysiol 38(3):539-559.

Luna R, Hernández A, Brody CD, Romo R (2005) Neural codes for perceptual discrimination in primary somatosensory cortex. Nat Neurosci 8(9):1210-1219.

Mountcastle VB (2005) The sensory hand: neural mechanisms of somatic sensation. Harvard University Press.

Mountcastle VB, Steinmetz MA, Romo R (1990) Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. J Neurosci 10(9):3032-3044.

Mountcastle VB, Talbot WH, Sakata H, Hyvärinen J (1969) Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. J Neurophysiol 32(3):452-484.

Oldfield RC (1971) The assessment and analysis of handedness: the Edinburgh inventory. Neuropsychologia 9:97–113.

Romo R, de Lafuente V (2013) Conversion of sensory signals into perceptual decisions. Prog Neurobiol 103:41-75.

Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. Nat Rev Neurosci 4:203-218.

Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. J Neurosci 20:5503-5515.

Simons SB, Tannan V, Chiu J, Favorov OV, Whitsel BL, Tommerdahl M (2005) Amplitude-dependency of response of SI cortex to flutter stimulation. BMC Neurosci 6:43-57.

Tommerdahl M, Hester KD, Felix ER, Hollins M, Favorov OV, Quibrera PM, Whitsel BL (2005) Human vibrotactile frequency discriminative capacity after adaptation to 25 Hz or 200 Hz stimulation. Brain Res 1057:1-9.

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

Verrillo R T, Fraioli AJ, Smith RL (1969) Sensation magnitude of vibrotactile stimuli. Percept Psychophys 6:366-372.

Whitsel BL, Kelly EF, Xu M, Tommerdahl M, Quibrera M (2001) Frequency-dependent response of SI RA-class neurons to vibrotactile stimulation of the receptive field. Somatosens Mot Res 18:263-285.

4.7 Tables and table captions

		Amplitude: High	Amplitude: Low
Frequency	High	+	–
	Low	–	+

Table 4-1: Congruent (+) and incongruent (–) combinations of amplitude and frequency.

		Amplitude: 40µm	Amplitude: 80µm
Frequency	20Hz		
	40Hz		

Table 4-2: All possible amplitude and frequency combinations of the standard-stimulus that we delivered in the AD and FD tasks. The grey cells represent iso-energy amplitude-frequency combinations.

Participant Number	Pre-training Slope	Post-training Slope
1	-0.27	-0.256
2	-0.181	0
3	-0.058	-0.067
4	-0.329	-0.3
5	-0.321	-0.244
6	0	0
7	-0.215	-0.188
8	0	-0.102
9	-0.536	-0.318
10	-0.103	-0.591
11	-0.608	-0.188
12	-0.201	0

Table 4-3: Pre- and post- training slopes obtained from experiment 2. The even numbered (grey cells) participants performed the FD_{vA} and the remaining performed the AD_{vF} task.

Participant Number	Pre-training AD_{vF} Slope	Post-training AD_{vF} Slope	Pre-training FD_{vA} Slope	Post-training FD_{vA} Slope
1	0.412	-0.032	0.8	-0.074
2	0.418	-0.006	0.969	0.231
3	0.47	0.001	0.899	-0.43
4	0.407	0	0.707	0.036
5	0.442	0.017	0.744	-0.064
6	0.436	0.016	0.61	-0.011

Table 4-4: Pre- and post- training slopes obtained from simulated AD_{vF} and FD_{vA} tasks.

4.8 Figures and figure captions

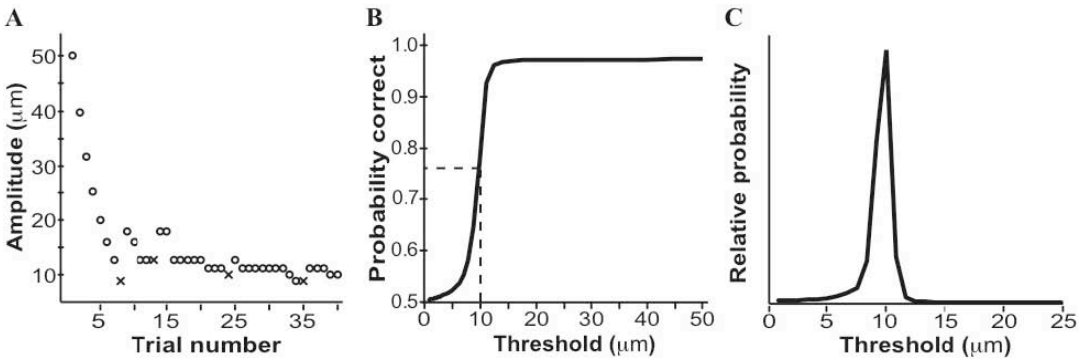

Figure 4.1: a) Participant's performance (o = correct, x = incorrect), b) participant's psychometric function, c) posterior probability distribution function of threshold (θ) parameter.

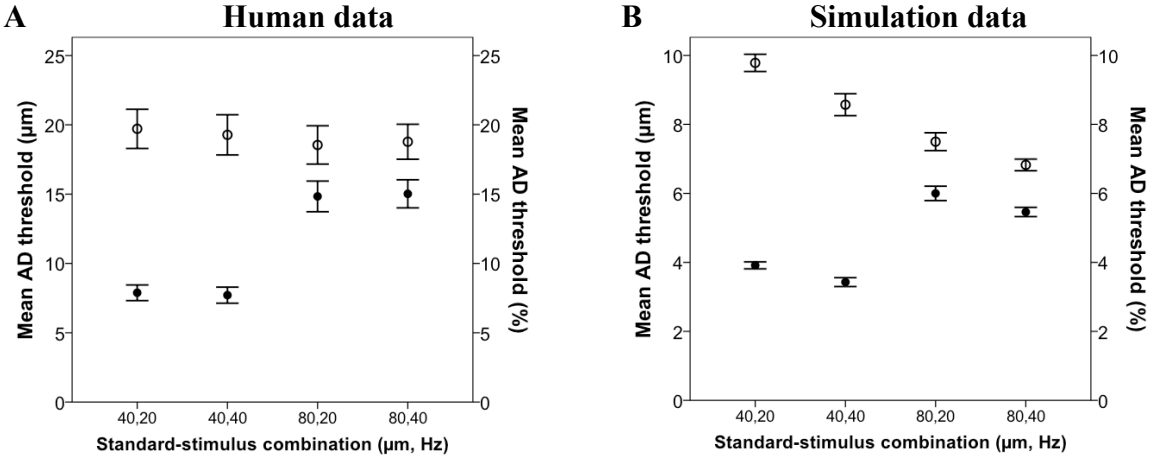

Figure 4.2: Amplitude discrimination performance estimated in human participants (panel A) and simulated subjects (panel B) are shown in this figure. In all simulations we presented 10 stimulus cycles and the spontaneous spike-rate was set to 10 spikes/sec (see Chapter 2 for detail). The open circles represent the mean percent difference threshold, which refers to the right y-axis, and the filled circles represent the mean absolute difference threshold (in microns), which refers to the left y-axis. The error bar represents ±1SE.

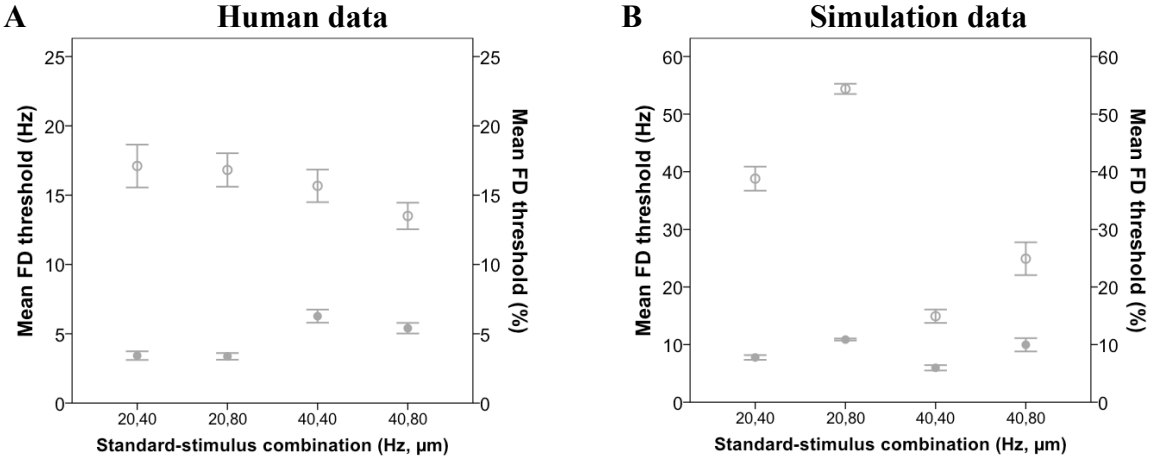

Figure 4.3: Frequency discrimination performance estimated in human participants (panel A) and simulated subjects (panel B) are shown in this figure. In all simulations we presented 10 stimulus cycles and the spontaneous spike-rate was set to 10 spikes/sec (see Chapter 2 for detail). The open circles represent the mean percent difference threshold, which refers to the right y-axis, and the filled circles represent the mean absolute difference threshold (in Hertz), which refers to the left y-axis. The error bar represents ±1SE.

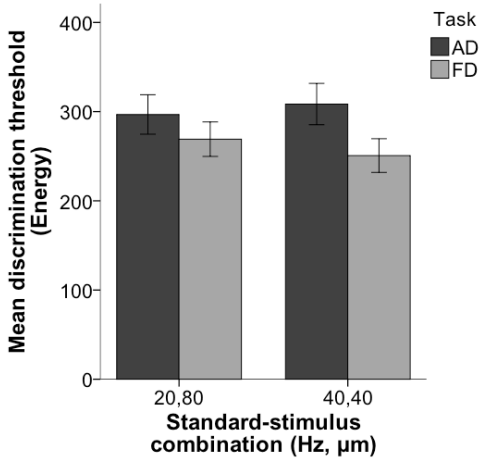

Figure 4.4: Performance of human participants on two different combinations of standard-stimulus where the energy ($a \times f$) is equal. The black bars represent AD threshold and grey bars represent FD threshold. The error bar represents ±1SE.

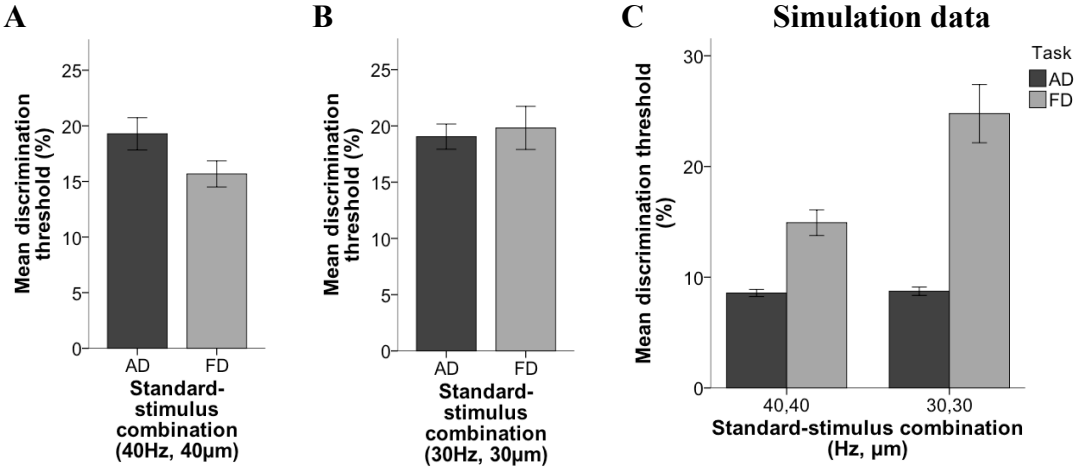

Figure 4.5: Comparison of thresholds that were estimated in AD and FD tasks when the standard stimulus combination was identical. Panels A and B show human participants' and panel C shows the performance of simulated subjects. In all simulations we presented 10 stimulus cycles and the spontaneous spike-rate was set to 10 spikes/sec (see Chapter 2 for detail). In all panels, the black bar indicates the AD threshold and the grey bar represents the FD threshold. The error bar represents ±1SE.

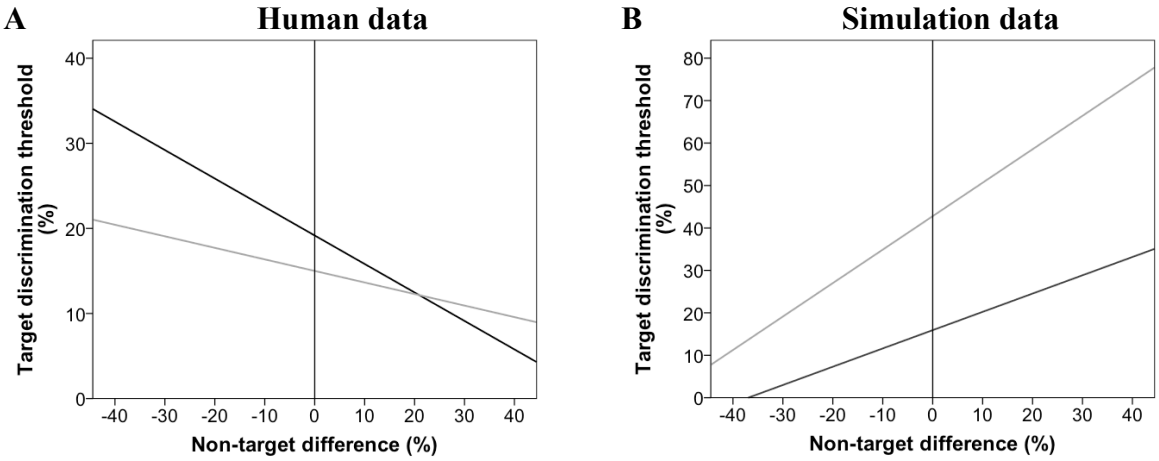

Figure 4.6: Panels A and B, show the AD_{vF} (black curve) and FD_{vA} (grey curve) best-fit line (BFL_{ntFAD}) calculated by averaging the ntFAD slopes and intercepts from 12 (6 in each subgroup) human participants (panel A) and 6 simulated subjects (panel B). In all simulations we presented 4 stimulus cycles and the spontaneous spike-rate was set to 10 spikes/sec (see Chapter 2 for detail). Note that the scale of y-axis in both panels is different. The non-target difference refers to the systematic manipulation of amplitude and frequency during the AD_{vF} and FD_{vA} tasks, respectively.

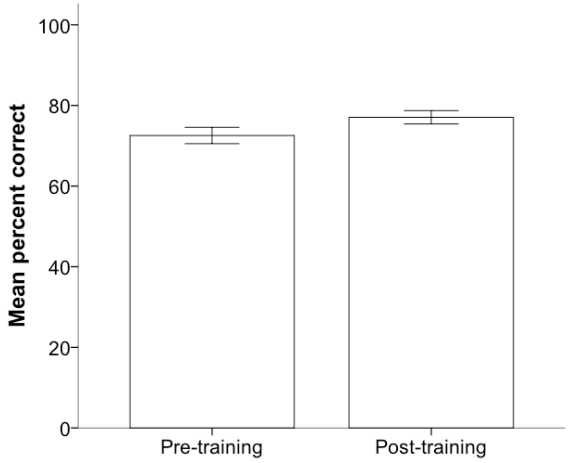

Figure 4.7: Shows the performance of the participants in experiment 3 in which the participants had to identify the feature (amplitude, or frequency) was higher in the second interval of the 2IFC task. Note that the performance of the participants before the training regimen was conducted is already higher than chance performance (50%).

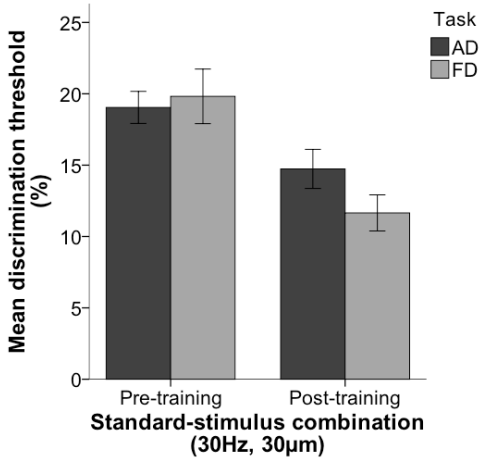

Figure 4.8: Comparison of thresholds that were estimated AD and FD tasks where the standard stimulus combination is identical before and after implementing the training regimen. The pre-training performance is also shown in Fig. 5C. The black bar indicates the AD threshold and the grey bar represents the FD threshold. The overall enhancement in discrimination threshold after training might also be due to practice effect. The error bar represents ±1SE.

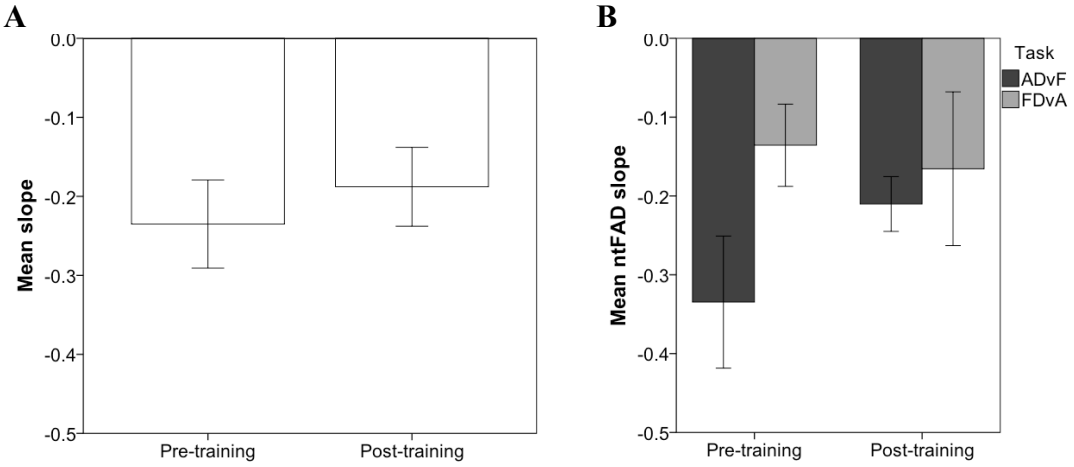

Figure 4.9: The graph in panel A shows the mean ntFAD slope of all human participants' performance before and after the training regimen was conducted. The graph in panel B shows the AD_{vF} and FD_{vA} data, which is also shown in panel A, separated in the corresponding subgroups.

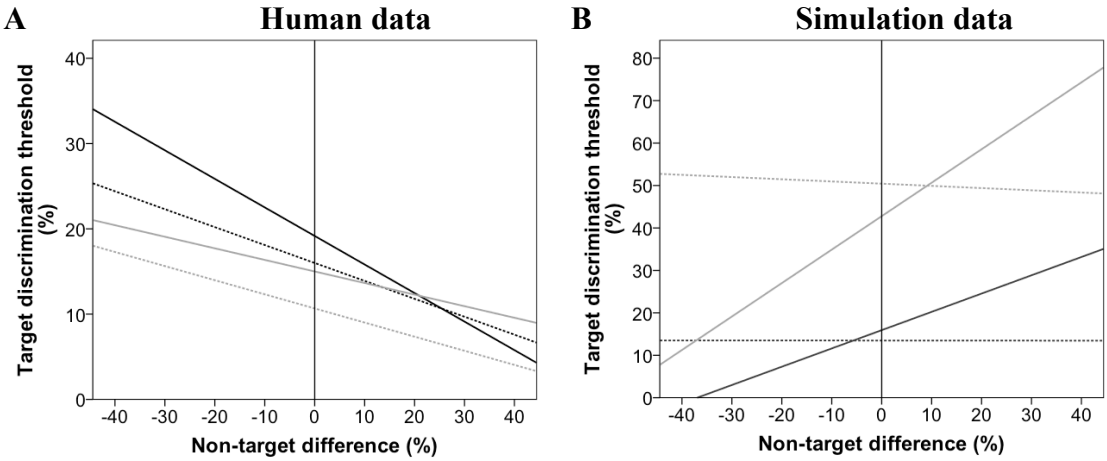

Figure 4.10: Shows the effect of training or "awareness of manipulation" on the mean best-fit line (BFL_{ntFAD}) calculated by averaging the ntFAD slopes and intercepts ntFAD slopes that were estimated from the human participants' (panel A) and simulated subjects' (panel B) performances in the AD_{vF} and FD_{vA} tasks. In all simulations we presented 4 stimulus cycles and the spontaneous spike-rate was set to 10 spikes/sec (see Chapter 2 for detail). The solid curves in both panels represent pre-training performance, which is also shown in Fig. 6. The dash-curve in both panels represents post-training performance. The black curve in both panels refers to the participants' and the simulated subject's mean BFL_{ntFAD} in the AD_{vF} task, and the grey curve refers to that of FD_{vA} task. The non-target difference refers to the systematic manipulation of amplitude and frequency during the AD_{vF} and FD_{vA} tasks, respectively.

CHAPTER 5: GENERAL DISCUSSION

Every vibrotactile stimulus has two features – amplitude and frequency. There is a general agreement that the neural code (i.e. the neural representation) for stimulus amplitude is the rate code; however, the neural code for stimulus frequency is unclear – whereas evidence from some studies support the periodicity code (i.e. a temporal code), other studies support the rate code. In this thesis, by using computational and psychophysical techniques, I investigated the neural code for these features, and the importance of and interaction between these features in vibrotactile detection and discrimination tasks. The results are consistent with the rate code for amplitude tasks and the periodicity code for frequency tasks.

5.1 Summary of studies

In chapter 2, to investigate whether human performance on amplitude and frequency tasks is consistent with optimal inference based on spike rate, we first created a spike-rate-based ideal observer model. We implemented detection and discrimination (amplitude, frequency) tasks for three reasons: a) to quantify the loss of information due to Poisson noise, which represented the cortical noise in the model, b) to determine whether the amount of neural evidence (in terms of receptor density, stimulus duration, etc.) affects perception, and c) to quantify how well the ideal observer performed in frequency discrimination task using the rate code. The simulation results provide behavioural trends that we tested in chapters 3 and 4.

In chapter 3, we conducted several behavioural experiments on young adults, investigating the effect of receptor density and stimulus duration on vibrotactile detection. Intuitively as well as based on the simulation results from chapter 2, perceptual performance should reflect the quantity of sensory evidence. Therefore, increases in receptor density and / or stimulus duration should result in lower thresholds, i.e., spatial summation and temporal summation. To test for spatial summation, in 1 experiment we estimated human participants' threshold for detecting a 40Hz stimulus presented on their fingertip and thenar eminence because in humans, the RA afferent density in the thenar eminence is ~6 times lower than that in the fingertip. To test for temporal summation, in 5 different experiments, we estimated human participants' threshold for detecting 20Hz and 40Hz vibrations presented for a range of stimulus durations. The results supported our predictions and validated the use of rate code to represent stimulus amplitude. These results are in accord with those of some previous studies (Green, 1976; Hämäläinen et al., 1981) but in apparent disagreement with those of other (Verrillo, 1965; Gescheider and Joelson, 1983) (See discussion in Ch. 3).

In chapter 4, we tested human participants on a range of discrimination tasks to determine whether – a) human participants combined the vibrotactile stimulus features (amplitude and frequency) during discrimination tasks, and b) a rate code can explain frequency discrimination performance. We replicated Harris et al.'s (2006) experiment by surreptitiously manipulating the amplitude (i.e. the non-target) feature values during a frequency discrimination task. To complement Harris et al. (2006), we surreptitiously manipulated the frequency (i.e. the non-target) feature values during an amplitude discrimination task. We quantified the effect using a single slope parameter (ntFAD), and found that the frequency manipulation had a stronger effect on amplitude discrimination than vice-versa. This suggests, consistently with Harris et al. (2006), that the non-target feature influences performance, but also that humans do not combine the frequency and amplitude features precisely as suggested by Harris et al. (2006). To further investigate this, we conducted a novel feature identification task in which we asked the participants, using a 2-interval forced choice protocol, to identify the feature that increased in the second interval. Evidently, if participants can identify the features, then they certainly do not combine amplitude and frequency into a single discrimination feature. The comparisons between the human participants' and the ideal observer's performances in amplitude discrimination tasks further support the use of a rate code to represent stimulus amplitude; however, similar comparisons in frequency discrimination tasks imply that humans probably use a periodicity code to represent stimulus frequency.

5.2 Neural code revisited

Previous studies have consistently shown that stimulus amplitude is represented by a rate code (Bensmaïa, 2008; Simons et al., 2005; Tommerdahl et al., 2010; Güçlü and Dinçer, 2013; Harvey et al., 2013). The trends that we found in all the simulated tasks related to stimulus amplitude are also present in the human data, which convinces us that vibrotactile amplitude information is represented by a rate code.

However, there exists a debate about whether stimulus frequency is represented by a periodicity code or a rate code. Consistently, neurophysiological studies have provided evidence in support of a periodicity code (Moutcastle et al., 1969, 1990; Whitsel et al., 2001; Harvey et al., 2013). Complementing these neurophysiological studies, three lines of behavioural evidence from this thesis suggest that humans represent frequency either with a periodicity code, or at the least with a neural code that is certainly not a rate code: a) whereas our rate-code-based ideal observer outperformed humans for amplitude discrimination tasks, humans outperformed the model in frequency discrimination tasks, b) humans could identify the stimulus feature that had a higher value in the second interval of each trial, and c) there was less effect of secret manipulation of amplitude in

the frequency discrimination task than of secret manipulation of frequency in the amplitude discrimination task.

In contrast, Romo and colleagues have provided intriguing evidence in support of a rate code for frequency (Romo et al., 1998; Hernández et al., 2000; Salinas et al., 2000). Note that Romo and colleagues do show evidence for temporal response characteristics in cortical neural response (Hernández et al., 2000; Salinas et al., 2000). Of course, observing temporal signatures of cortical neural responses does not imply that experimental animals are utilizing that information. Nonetheless, the temporal information is available at the cortex.

It is not clear why Romo and colleagues have consistently reported the rate code as a representation for stimulus frequency. I speculate that their frequency discrimination tasks may have procedurally become amplitude discrimination tasks. In Hernández et al. (2000), Salinas et al. (2000), and Luna et al. (2005), Romo and colleagues used an identical waveform frequency for each pulse (reportedly 20ms duration, which converts to 50Hz) and manipulated the duration between pulses to create stimuli of different frequencies. For example, to create a 500ms duration 20Hz stimulus, the authors presented 11 pulses (20ms-50Hz waveform frequency) with an inter-pulse duration of 50ms (Salinas et al., 2000). Because the RA afferents' sensitivity (I_0) and entrainment (I_1) thresholds are sensitive to the waveform frequency (see chapter 2, Freeman and Johnson, 1982) and not the temporal frequency, the number of spikes evoked per stimulus cycle does not change in each interval unless the amplitude of each stimulus cycle is changed. Next the authors matched the subjective intensity by changing the amplitude of the stimuli, which will change the number of spikes evoked in each interval (Johnson, 1974). During training the authors presented bigger difference in frequencies, which implies the amplitude difference also increased concurrently. Furthermore, during the training, the authors presented visual cues to direct the monkeys to choose the correct "high temporal frequency low amplitude" interval and correct responses were rewarded. Therefore, it is possible that the monkeys associated juice reward with the low amplitude stimulus, effectively changing the intended "frequency discrimination" task into an "amplitude discrimination" task. This thesis shows that amplitude can be reliably represented as a rate code.

5.3 Future experiments: investigating biological constraints

To the best of our knowledge, the Bayesian ideal observer model presented in this thesis is the first of its kind to estimate optimal vibrotactile perception. Note that this model does not aspire to be a model of the brain; rather it is a tool with which we can estimate the efficiency of the human perceptual system. We found that human vibrotactile

amplitude discrimination is inefficient, in the sense that it quantitatively falls short of the performance of the ideal observer. Therefore, for future modelling experiments, a fascinating next step would be to investigate systematically the sources of sub-optimality in vibrotactile perception via the incorporation of plausible biological constraints.

A strong advantage of Bayesian observer models is that they allow the incorporation of realistic biological constraints, and indeed we have already begun this process in this thesis. For example, neuronal response variability is a biological constraint: We found that the ideal observer's performance in amplitude discrimination was superior to that of humans, which suggests that during perceptual processing there is loss of information in humans. This is typically the case in ideal observer comparisons to human performance (Geisler, 1989, 2011). Among other reasons, the loss of information could be due to the stochastic nature of neural information processing. Variability is intrinsically incorporated in successive stages of processing. For example, the RA afferent variability can be characterized as a Gaussian with a small standard deviations (Vega-Bermudez and Johnson, 1999); however, in RA-like cortical neurons the variability increases and is characterized by the Poisson or Poisson-like distribution (Sripati et al., 2006). Therefore, incorporating Poisson variability into the RA afferent activity, we quantified the decline in performance due exclusively to Poisson variability.

Spontaneous neural activity as a biological constraint: The presence of spontaneous spiking activity in the cortex interferes with the stimulus-evoked neural responses, i.e. it reduces the fidelity of the neural signal. The Bayesian observer model presented in this thesis "knows" the sources of noise or variability (including spontaneous noise), which makes it ideal; however, it is possible that our brain, i.e., our perceptual observer, does not know about the spontaneously occurring spikes. How, and by how much, performance would degrade if the Bayesian observer were unaware of the spontaneous noise spikes could be explored in future investigations.

Another biological constraint concerns noise correlations among cortical neurons. In the Bayesian ideal observer model presented in this thesis, the cortical neuronal responses are conditionally independent given the stimulus; this assumption of conditional independence is a common practice used to keep models computationally tractable (Jazayeri and Movshon, 2006); however this is possibly not true in human cortical neurons (Zohary et al., 1994). Therefore, such assumptions might oversimplify the model, which could lead to lower thresholds in relevant perceptual tasks. Future modelling work could incorporate dependencies among the simulated cortical neurons.

By comparing the Bayesian ideal observer model results to the human results we learned that humans are not using a rate code to represent frequency, and that humans are sub-optimal at using the rate code to represent amplitude. To complement the investigation into the sources of sub-optimality in human vibrotactile perception, another important modelling mission should be to create an ideal observer model that utilizes the periodicity code for frequency. Presumably, owing to several biological constraints, when compared

to the periodicity-based ideal observer model, humans would be inferior at utilizing all the periodicity related information to optimally perform in frequency tasks. Therefore, a systematic exploration of the sources of sub-optimality, using a periodicity-based ideal observer, would enhance our understanding of vibrotactile perception.

5.4 Conclusion

In conclusion, this thesis presents an ideal observer model for the RA afferent population with which we tested whether a rate code model faithfully represents the features of vibrotactile stimuli. The ideal observer analyses along with the behavioural results strongly suggest that vibrotactile frequency is represented as a periodicity code but that amplitude is represented by a rate code. Moreover, we show that humans can separately perceive the amplitude and the frequency features of vibrotactile stimuli.

5.5 References

Bensmaïa SJ, Hollins M (2003) The vibrations of texture. Somatosens Mot Res 20(1):33-43.

Bhattacharjee A, Ye AJ, Lisak JA, Vargas MG, Goldreich D (2010) Vibrotactile masking experiments reveal accelerated somatosensory processing in congenitally blind braille readers. J Neurosci 30:14288-14298.

Freeman AW, Johnson KO (1982) A model accounting for effects of vibratory amplitude on responses of cutaneous mechanoreceptors in macaque monkey. J Physiol 323:43-64.

Geisler WS (1989) Sequential ideal-observer analysis of visual discriminations. Psychol Rev 96:267-314.

Geisler WS (2011) Contributions of ideal observer theory to vision research. Vision Res 51:771-781.

Gescheider GA, Joelson JM (1983) Vibrotactile temporal summation for threshold and suprathreshold levels of stimulation. Percept Psychophys 33:156-162.

Green BG (1976) Vibrotactile temporal summation: effect of frequency. Sens Processes 1:138-149.

Güçlü B, Dinçer SM (2013) Neural coding in the Non-Pacinian I tactile channel: a psychophysical and simulation study of magnitude estimation. Somatosens Mot Res 30:1-15.

Hämäläinen H, Pertovaara A, Soininen K, Järvilehto T (1981) Is there low frequency vibrotactile temporal summation? Scand J Psychol. 1981;22(3):203-206.

Harris JA, Arabzadeh E, Fairhall AL, Benito C, Diamond ME (2006) Factors affecting frequency discrimination of vibrotactile stimuli: implications for cortical encoding. PLoS One 1:e100.

Harvey MA, Saal HP, Dammann JF 3rd, Bensmaïa SJ (2013) Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. PLoS Biol 11:e1001558.

Hernández A, Zainos A, Romo R (2000) Neuronal correlates of sensory discrimination in the somatosensory cortex. Proc Natl Acad Sci U S A 97:6191-6196.

Hollins M, Bensmaïa SJ, Washburn S (2001) Vibrotactile adaptation impairs discrimination of fine, but not coarse, textures. Somatosens Mot Res 18(4):253-262.

Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. Nat Neurosci. 9(5):690-696.

Johnson KO (1974) Reconstruction of population response to a vibratory stimulus in quickly adapting mechanoreceptive afferent fiber population innervating glabrous skin of the monkey. J Neurophysiol 37:48-72.

Luna R, Hernández A, Brody CD, Romo R (2005) Neural codes for perceptual discrimination in primary somatosensory cortex. Nat Neurosci 8(9):1210-1219.

Romo R, Hernández A, Zainos A, Salinas E (1998) Somatosensory discrimination based on cortical microstimulation. Nature 392:387-390.

Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. J Neurosci 20:5503-5515.

Sripati AP, Yoshioka T, Denchev P, Hsiao SS, Johnson KO (2006) Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system. J Neurosci 26:2101-2114.

Simons SB, Tannan V, Chiu J, Favorov OV, Whitsel BL, Tommerdahl M (2005) Amplitude-dependency of response of SI cortex to flutter stimulation. BMC Neurosci 6:43-57.

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

Vega-Bermudez F, Johnson KO (1999) SA1 and RA receptive fields, response variability, and population responses mapped with a probe array. J Neurophysiol 81:2701-2710.

Verrillo RT (1965) Temporal summation in vibrotactile sensitivity. J Acoust Soc Am 37:843-846.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature. 370(6485):140-143.