An Information Theoretic Approach to Tactile Encoding and Discrimination

Hannes Philipp Saal

Doctor of Philosophy
Institute of Perception, Action and Behaviour
School of Informatics
University of Edinburgh

2011

Abstract

The human sense of touch integrates feedback from a multitude of touch receptors, but how this information is represented in the neural responses such that it can be extracted quickly and reliably is still largely an open question. At the same time, dexterous robots equipped with touch sensors are becoming more common, necessitating better methods for representing sequentially updated information and new control strategies that aid in extracting relevant features for object manipulation from the data. This thesis uses information theoretic methods for two main aims: First, the neural code for tactile processing in humans is analyzed with respect to how much information is transmitted about tactile features. Second, machine learning approaches are used in order to influence both what data is gathered by a robot and how it is represented by maximizing information theoretic quantities.

The first part of this thesis contains an information theoretic analysis of data recorded from primary tactile neurons in the human peripheral somatosensory system. We examine the differences in information content of two coding schemes, namely spike timing and spike counts, along with their spatial and temporal characteristics. It is found that estimates of the neurons' information content based on the precise timing of spikes are considerably larger than for spikes counts. Moreover, the information estimated based on the timing of the very first elicited spike is at least as high as that provided by spike counts, but in many cases considerably higher. This suggests that first spike latencies can serve as a powerful mechanism to transmit information quickly. However, in natural object manipulation tasks, different tactile impressions follow each other quickly, so we asked whether the hysteretic properties of the human fingertip affect neural responses and information transmission. We find that past stimuli affect both the precise timing of spikes and spike counts of peripheral tactile neurons, resulting in increased neural noise and decreased information about ongoing stimuli. Interestingly, the first spike latencies of a subset of afferents convey information primarily about past stimulation, hinting at a mechanism to resolve ambiguity resulting from mechanical skin properties.

The second part of this thesis focuses on using machine learning approaches in a robotics context in order to influence both what data is gathered and how it is represented by maximizing information theoretic quantities. During robotic object manipulation, often not all relevant object features are known, but have to be acquired from sensor data. Touch is an inherently active process and the question arises of how

to best control the robot's movements so as to maximize incoming information about the features of interest. To this end, we develop a framework that uses active learning to help with the sequential gathering of data samples by finding highly informative actions. The viability of this approach is demonstrated on a robotic hand-arm setup, where the task involves shaking bottles of different liquids in order to determine the liquid's viscosity from tactile feedback only. The shaking frequency and the rotation angle of shaking are optimized online. Additionally, we consider the problem of how to better represent complex probability distributions that are sequentially updated, as approaches for minimizing uncertainty depend on an accurate representation of that uncertainty. A mixture of Gaussians representation is proposed and optimized using a deterministic sampling approach. We show how our method improves on similar approaches and demonstrate its usefulness in active learning scenarios.

The results presented in this thesis highlight how information theory can provide a principled approach for both investigating how much information is contained in sensory data and suggesting ways for optimization, either by using better representations or actively influencing the environment.

Acknowledgements

I would like to thank my supervisors Sethu Vijayakumar, Roland Johansson, and Peggy Series who were always happy to help me along the way. I am also grateful to Jo-Anne Ting and Nicolas Heess, with whom I collaborated on individual projects. Stefan Klanke helped getting the robot arm under control. Adrianna Teriakidis, Jun Nakanishi, and Nicolas Heess undertook the job of proofreading. Finally, I am grateful to my parents and my sister Lena for always believing in me. Thank you all so much.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Hannes Philipp Saal*)

Table of Contents

Chapter 1

Introduction

1.1 Motivation and aims

Imagine searching your coat pocket for your car keys. Your fingers get hold of a certain object and thousands of touch receptors in your fingertips start sending information to your brain about various properties associated with this object: how hard or soft it is, its texture, and properties of shape like curvature. Based on this feedback, you might disregard an object or explore it further if your expectations of how your keys should feel match what you actually perceive.

Humans are experts at such tasks in that they can rapidly and accurately extract information about object properties from sensory information relayed by mechanoreceptors. Moreover, they are also very good in utilizing this information to guide their exploration of an unknown object, so as to classify it as fast as possible.

Robotic systems, on the other hand, still do not come close to mimicking the human hand and arm, neither in dexterity nor in tactile capabilities. However, in order to prove useful in environments shaped for humans, robotic systems will need a sense of touch to be able to manipulate objects with unknown dynamical properties successfully. While touch sensor technology is getting more and more sophisticated, algorithms that efficiently extract tactile information from sensor feedback and guide the robot in exploring its environment through touch are still very much at the early stages of what is possible in terms of their performance.

The aim of this thesis is to apply information theoretic methods to elucidate tactile

processing in humans, and to design better algorithms for extracting information from tactile data in robots.

1.2 Thesis outline

In the following, we provide a quick chapter-by-chapter outline of the thesis.

Chapter 2: Tactile information processing in humans—from spikes to discrimination

This chapter provides an overview of tactile processing in humans. It also discusses the importance of spike timing and what is known about the mechanical properties of the fingertip. Finally, we mention some common approaches for estimating information theoretic quantities from electrophysiological data.

Chapter 3: Information content of different coding schemes

In this chapter, we explore how rich tactile information can be transmitted both rapidly and accurately by the human peripheral nervous system. This is done by examining different coding schemes and measuring the amount of information transmitted by each of them.

Main contributions:

- A quantitative information theoretic analysis showing that the information in first spike latencies about tactile features relevant for grasping is at least as high as that in spike counts.
- An examination of spatial properties revealing how different tactile features (object curvature and force direction) are represented by different neuronal subgroups depending on the receptors' position on the fingertip.
- An analysis showing information between temporal and rate codes to be mainly independent, rather than redundant or synergistic.

Publications:

- Saal HP, Vijayakumar S, and Johansson RS. *Information about complex fingertip parameters in individual human tactile afferent neurons*. J Neurosci (2009) vol. 29 (25) pp. 8022-31.

Chapter 4: Encoding and exploiting the effects of hysteresis

After having analyzed the responses of tactile neurons under ideal conditions, we turn our attention to more natural settings. We explore how the mechanical properties of the fingertip lead to a degradation of the transmitted information during sequential contact events, and how the responses from a subset of peripheral neurons might transmit a signal to remedy this problem.

Main contributions:

- An analysis showing that both fingertip skin mechanics and neural responses are affected by repeated application of different tactile stimuli.
- An information theoretic analysis that suggests that some peripheral neurons signal properties relating to past stimulation in the timing of their first spikes even as they are responding to a new stimulus.

Publications:

- Saal HP, Vijayakumar S, and Johansson RS. *Signals in ensembles of human tactile afferents code past and current fingertip stimuli*. In preparation.

Chapter 5: Haptic discrimination in robotics—from sensors to actions

In this chapter we present a survey of current approaches to tactile processing in robotics. Furthermore, we discuss an approach to mapping sequentially arriving sensory observations to object properties: the nonlinear filtering paradigm. Finally, we introduce active learning, which presents a principled way of influencing incoming data such that it is most informative about object properties of interest.

Chapter 6: Active learning for tactile processing in robotics

We examine how a robot's actions should be controlled so as to maximize the informativeness of sequentially arriving sensor data about relevant object features and present

results on a robotics hand-arm setup.

Main contributions:

- A machine-learning approach that combines active learning with nonlinear filtering.
- A robotics application of haptic discrimination to dynamically estimate the viscosity of different liquids from tactile feedback by shaking them: shaking parameters are optimized online in order to maximize expected information gain.

Publications:

- Saal HP, Ting J, and Vijayakumar S. *Active sequential learning with tactile feedback*. 13th International Conference on Artificial Intelligence and Statistics (AISTATS), Chia Laguna, Italy, JMLR: W&CP 9 (2010).
- Saal HP, Ting J, and Vijayakumar S. *Active estimation of object dynamics parameters with tactile sensors*. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Taipei, Taiwan (2010) pp. 916 - 921.

Chapter 7: Representing multimodal distributions for sequential observations

In this chapter, we are concerned with the problem of how to represent complex and possibly multimodal probabiltiy distributions, which are common in tactile processing, for example in tactile search tasks. We also show that representing uncertainty accurately is important for active learning scenarios, such as those discussed in the previous chapter.

Main contributions:

- A novel approach to nonlinear filtering, which maintains the state distribution as a mixture of Gaussians using deterministic sampling.

Publications:

- Saal HP, Heess NMO, and Vijayakumar S. *Multimodal nonlinear filtering using Gauss-Hermite Quadrature*. European Conference on Machine Learning (ECML), Athens, Greece (2011). In press.

Chapter 8: Conclusion

We conclude the thesis by discussing some potential future avenues of research in tactile processing, both in neuroscience and robotics. We go on to discuss the merit of using information theory, and how the fields of neuroscience and robotics can inform and benefit from each other.

Chapter 2

Tactile information processing in humans—from spikes to discrimination

2.1 Introduction

In this chapter we will give an overview of what is currently known about what and how tactile features are transmitted in the human tactile system.

When grasping an object, thousands of mechanoreceptors embedded in the human fingertip respond and start transmitting information about object properties that are important for grasp stability. Based on this feedback, any necessary changes to grasp forces are executed rapidly so as to prevent objects from slipping or breaking due to overgrasping. This happens reliably even though in real-world object manipulation tasks the same force applied in quick succession onto the fingertip will affect indentation differently due to the mechanical properties of the different fingertip tissues.

In our overview, we will first review the different stages of tactile processing in Sec. 2.2, with an emphasis on the signals that are picked up and transmitted by peripheral neurons. We then ask how tactile information could be encoded in spike trains in order to be transmitted as quickly as possible. To this end, we discuss the importance of spike timing, with a focus on the somatosensory system, in Sec. 2.3. We will describe our own investigation of this matter in Chapter 3, where we quantitatively compare differ-

ent potential coding schemes. Next, the mechanical properties of the fingertip along with their effects on neural responses and perception are discussed in Sec. 2.4. Our own analysis into this matter is the topic of Chapter 4. Our analyses are mainly based on information theoretic concepts. We will present a short overview of information theoretic methods as they are used in the neurosciences in Sec. 2.5, which concludes this chapter.

2.2 The stages of tactile processing in humans

Humans show amazing dexterity when interacting with and manipulating objects using their hands and fingers. This ranges from seemingly simple tasks like picking up objects to more complex ones, like tying a shoelace, which require very accurate spatial and temporal movement control. To handle such tasks successfully, it is not only important to recognize the occurrence of certain events like making or breaking contact with objects, but also to correctly estimate object properties like weight, curvature and other parameters relating to friction and object elasticity. For example, when grasping objects, the grip force has to be adjusted for object properties so as to avoid slips (Johansson and Westling, 1984).

Usually, input from a number of senses is integrated in order to generate the necessary motor commands. This sensory data includes feedback from visual, auditory, and proprioceptive senses, together with predictions from internal models (Johansson and Flanagan, 2009). Feedback from the tactile system plays a very important part in this, especially in the absence of prior information or input from other senses (reviewed in Witney et al., 2004). For example, the ability to light a match is not impaired much when subjects are blindfolded, but the same task without tactile input takes considerably longer and can pose a serious challenge (R. S. Johansson, personal communication). To provide the necessary information, each human fingertip is innervated by around 2,000 tactile afferents (Johansson and Vallbo, 1979), which is a much higher density than in most of the rest of the body.

In the following section, we review the most relevant literature on how tactile information is picked up by mechanoreceptors located within the skin and transferred to the central nervous system. We will focus on human data, but also refer to studies in monkeys and other animals if they appear to be relevant.

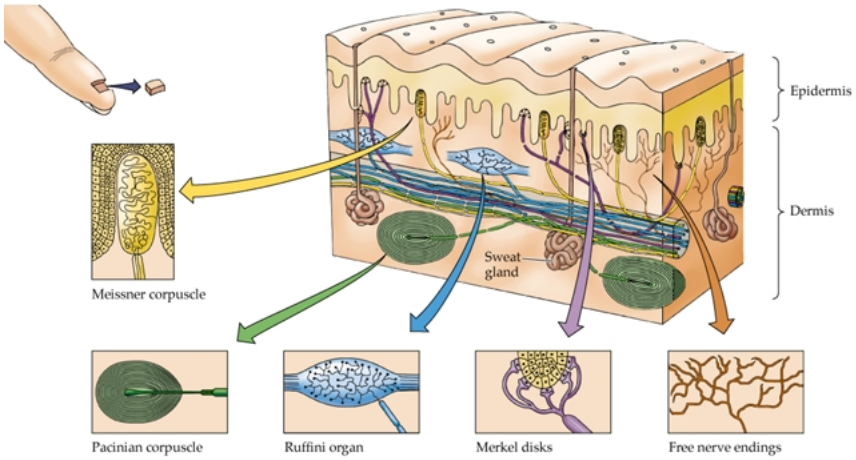

Figure 2.1: Mechanoreceptors and their locations in human glabrous skin. Figure taken from Purves et al. (2001).

2.2.1 Mechanoreceptors and tactile afferents

Four major classes of mechanoreceptors have been identified in human glabrous (hairless) skin (see Fig. 2.1): Meissner's corpuscles and Merkel disk receptors lie in the superficial layer of the skin, i.e. directly below the epidermis, and are connected to FA-I (fast adapting type I) and SA-I (slowly adapting type I) primary tactile afferents, respectively; Pacinian corpuscles and Ruffini endings are located deep within the dermis or even in subcutaneous regions and are innervated by FA-II and SA-II afferents, respectively. The relationship between mechanoreceptors and tactile afferents is not one-to-one, i.e. a single mechanoreceptor can be connected to several afferents and vice versa (Paré et al., 2002). Whereas slowly adapting afferents show ongoing activity during stimulus presentation, fast adapting ones only signal stimulus on- and offset, thus responding only to changes in the external stimulus (see Johansson and Flanagan, 2009, for a review). The density of each of these receptor types varies with location. In the human fingertip there are an estimated 140 FA-I afferents and 70 SA-I afferents per cm^2 of skin area but only around 22 FA-II and 12 SA-II afferents per cm^2 (Johansson and Vallbo, 1979). While many aspects of mechanotransduction are still poorly understood, there has been much progress in recent years in identifying and characterizing the molecular mechanisms behind it (see Delmas et al., 2011, for a review).

2.2.2 Response properties of primary tactile afferents

In most studies, response properties of tactile afferents have been characterized in terms of their firing rates. We will provide a summary of the main results on firing rates below and will focus on some recent results related to spike timing in Sec. 2.3.

It is thought that the different classes of tactile afferents code for distinct tactile features (see Johnson, 2001, for a review). While this view might not be as clear-cut as argued by some, there are undoubtedly differences in what types of stimuli are most effective for different afferent classes.

SA-I afferents are thought to be responsible for encoding object shape and fine-grained texture discrimination (Johnson and Hsiao, 1992), and they are also highly sensitive to edges and gratings (Bensmaia et al., 2006). FA-I afferents seem to be involved in motion detection (Friedman et al., 2002), stimulus localization, and especially the detection of sudden slip during object handling (Macefield et al., 1996). As for type II afferents, FA-II afferents are mainly sensitive to vibratory stimuli, and SA-II afferents seem to be involved in the detection of skin stretch. Consistent with this view, information about the location and movement of small surface features seems to only be represented in type I afferents (Blake et al., 1997; Phillips et al., 1992; LaMotte and Srinivasan, 1996).

Most of the aforementioned studies have been carried out using stimuli that focus on specific tactile features (e.g. pointy probes or braille dots), thus sometimes making it hard to compare and unify their results. In an attempt to provide a general view as to how type I tactile afferents' receptive fields (RFs) are structured, Sripati et al. (2006b) used tactile white noise stimuli to determine afferents' RFs using spike-triggered averages. They found mainly circular receptive fields, along with center-surround suppression (see also Vega-Bermudez and Johnson, 1999), quite similar to the early stages of the visual system.

In the studies mentioned so far, only afferents on the flat part of the fingertip were examined. Moreover, applied forces were generally low, operating around the afferents' individual thresholds. However, when grasping or manipulating objects in natural settings, the forces acting on the fingertip are usually much bigger (on the order of several Newtons), and lead to widespread deformation of the fingertip (see Sec. 2.4). As a result, afferents all over the fingertip respond to these stimuli (Bisley et al., 2000; Birznieks et al., 2001). During grasping, grip forces have to be adjusted

to match loading forces, and for this it is crucial to accurately estimate object properties at the contact point. A number of object and contact properties have been shown to be picked up by (mainly type I) mechanoreceptors for stimuli mimicking grasping conditions. Some examples are force loading direction (Birznieks et al., 2001; Wheat et al., 2010), object curvature (Goodwin et al., 1995, 1997; Jenmalm et al., 2003), and tangential torque (Birznieks et al., 2010). These features excite both SA-I and FA-I afferents, with only minor differences: in general, SA-I afferents seem to be less noisy, while FA-I afferents tend to respond more quickly (Johansson and Birznieks, 2004; Birznieks et al., 2010).

As multi-unit recordings from primary tactile afferents are very difficult, most studies focus on the responses of single neurons. A number of factors have to be considered when extrapolating results from single afferents to a full population. First, their responses are correlated, not only because they respond to the same stimulus (signal correlations), but also by other factors (noise correlations), such as heart rate and breathing (Macefield, 2003) that affect neighbouring afferents in a similar way. Moreover, some properties like conduction velocity (Kakuda, 1992), stimulation thresholds, and noise levels vary considerably between neurons and have to be taken into account when attempting to read out population responses (Goodwin and Wheat, 1999). Given known response statistics as well as estimates of factors affecting population responses, a number of models have been proposed for synthesizing information from primary afferents so that it can be fed into further processing stages (see Goodwin and Wheat, 2004, for an overview). An open question is how to combine them all into a single read-out model.

There have also been approaches to modelling the response properties of primary tactile afferents. Recently, a number of models predicting neural responses from fingertip indentations, force impacts, or vibratory stimuli have been developed (Sripati et al., 2006a; Kim et al., 2010; Rivest and Gerling, 2010; Lesniak and Gerling, 2009; Russell et al., 2009). These models serve as tools to test our understanding of how neural responses are related to the mechanical properties of the skin (see Sec. 2.4) and mechanotransduction. Furthermore, they constitute important steps towards interfacing touch sensor feedback from neuroprosthetic devices with severed tactile afferents in amputees.

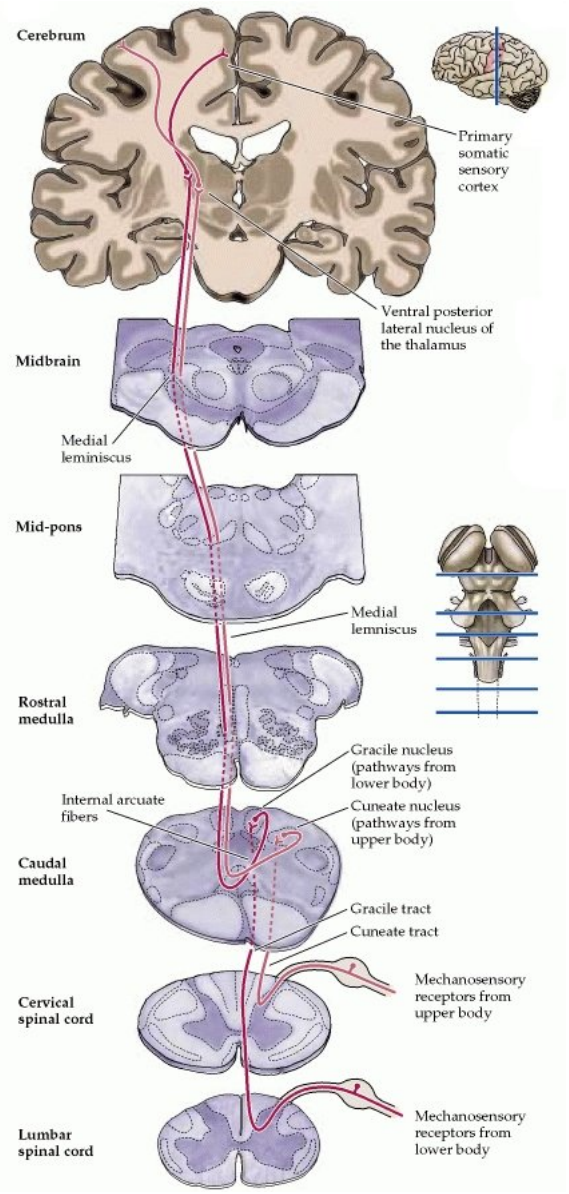

Figure 2.2: The tactile pathway from primary afferents to the cortex. Figure adapted from Purves et al. (2001).

2.2.3 Further processing stages and cortex

So far, we have concentrated on primary tactile afferents. In the following, we will present a short overview of the (slightly simplified) pathway from first-order neurons to the cortex (see Jones and Lederman, 2006, for more details). The cell bodies of the primary afferents lie in the dorsal root ganglion. Dorsal root ganglion cells enter the spinal cord on the dorsal side and then project upwards onto secondary neurons in the

cuneate nucleus, which is part of the medulla. Secondary neurons then project from there onto neurons in the ventral posterior lateral part of the thalamus. This pathway is known as the dorsal column medial lemniscal system. Finally, third-order neurons project from the thalamus to the primary somatosensory cortex. Fig. 2.2 provides a graphical overview of the whole pathway.

Studies have shown that there is widespread convergence and divergence of tactile neurons in the cuneate nucleus (see Johansson and Flanagan, 2009, for more details), which hints at this area being an important first step in consolidating the potentially highly redundant information from numerous primary afferents. In the thalamus, modulation of the tactile signal by other systems, for example by the primary motor cortex, has been observed in rats, which could be a first step in shaping the sensory signal according to an animal's current actions (Lee et al., 2008). Other studies on the whisker pathway in rats found that dynamic parameters of whisker movement are relayed with very high temporal precision in the thalamus (e.g. Montemurro et al., 2007a; Petersen et al., 2008).

The primary somatosensory cortex (S1, Brodmann area 3 and especially 3b for tactile processing) is arranged topographically, with different fingers represented independently (Iwamura et al., 1983; Sanchez-Panchuelo et al., 2010). However, it has been shown that there seems to be widespread integration, with single tactile stimuli being able to elicit widespread responses across S1 (Reed et al., 2008). In contrast to peripheral receptive fields, cortical RFs are often elongated, and appear well fit by Gabor filters (Bensmaia et al., 2008) and exhibit regions of high inhibition (DiCarlo et al., 1998; Sripati et al., 2006b). It has also been argued that in some respects, the processing of touch in the cortex is similar to cortical visual processing. For example, motion across the skin—which can be ambiguous to resolve due to the aperture problem—is integrated in the tactile system in a similar fashion as in the visual system (Pei et al., 2008, 2011). Additionally, neurons in both visual area V4 and somatosensory area S2 have been found to code for curvature direction, which might indicate a common underlying mechanism to represent shape or facilitate multisensory integration (Yau et al., 2009).

2.2.4 Perception

Perceptual studies have confirmed that many tactile features that are represented in neural responses are indeed perceived by human subjects. For example, subjects are readily able to distinguish multiple force stimuli in various directions (Panarese and Edin, 2011), object curvature (Jenmalm and Johansson, 1997; Goodwin and Wheat, 1992), object position (Wheat et al., 1995), stimulus intensity (Muniak et al., 2007), and skin indentation depth (Mei et al., 1983).

It has been found that even single action potentials in FA-I afferents elicit a perception, while for other afferent types, more vigorous stimulation may be required (Vallbo and Johansson, 1984). Additionally, different types of afferents elicit different kinds of percepts, which further hints to a functional distinction between them (cf. Sec. 2.2.2).

Lederman and Klatzky (1997) have studied when different tactile features become available to subjects, which might give insights into how these features are processed in the central nervous system. They found that coarse features related to stimulus intensity become available early on, while features related to object shape and form took longer.

2.3 Spike timing and temporal codes

A central question in neuroscience is how information in general and external stimuli in particular are encoded in the responses of neurons. Since Adrian's seminal work in the 1920s (Adrian, 1928), a lot of studies have focused on firing rates, i.e. the number of action potentials (spikes) within a given time frame. In modelling approaches, the exact timing of spikes is often assumed to carry no information, for example Poisson spike generators are commonly used. However, in more recent work this assumption has been increasingly challenged. In the following, we will present a short overview of the mounting evidence for the importance of spike timing, with an emphasis on the tactile system.

2.3.1 Spike timing in somatosensory and other systems

Information about various tactile stimuli is rapidly extracted and processed by the central nervous system. When grasping an object, contact features like friction, force

direction and surface shape are reflected in reactive fingertip forces as early as at 100 ms after initial contact (Johansson and Westling, 1984, 1987; Jenmalm and Johansson, 1997). These motor responses reflect complex stimulus features, which are processed in cortex to guide fine-tuned motor commands. Taking into account neuronal conduction velocities (Kakuda, 1992), the speed at which the motor system is able to react does not leave much time at each individual processing stage. It has been estimated that there is only enough time at each processing stage for a single spike per neuron (Johansson and Birznieks, 2004).

While it has been shown that first spike latencies in human primary tactile afferents can be used to discriminate different tactile features (Johansson and Birznieks, 2004), so far it is unknown how firing rates and spike latencies compare in terms of the amount of information that is transmitted. We will attempt to tackle this question in Chapter 3.

Another important question is whether and how subsequent processing stages could take advantage of spike timing information. A first relay station in the somatosensory pathway is the spinal cord and the cuneate nucleus in the medulla, where it has been shown that single action potentials are relayed reliably and with high timing precision in cats (Vickery et al., 1994). Moreover, recordings in rat thalamus have found spike timing, and especially first spike latencies to carry at least as much if not more information than firing rates (Foffani et al., 2009). Finally, recordings from the somatosensory cortex in rodents have shown that spike timing does contain more information than spike counts, both when applying stimuli to whiskers (Panzeri et al., 2001; Petersen et al., 2001, 2002) and forelimbs (Foffani et al., 2004). However, so far there are no studies examining spike timing in the somatosensory cortex of humans or monkeys.

Information in spike timing has been found in neurons throughout the nervous system, for example in auditory (Furukawa et al., 2000; Kayser et al., 2010), visual (Gawne et al., 1996; Reich et al., 2001; Nemenman et al., 2008; Gollisch and Meister, 2008; Jacobs et al., 2009) and olfactory (Junek et al., 2010) circuits (see also Heil, 2004; Stein et al., 2005; Van Rullen et al., 2005, for general reviews). Moreover, it has been shown that small milli-second delays can be behaviourally relevant and drive decisions (Yang et al., 2008). On the other hand, a recent study has shown that even small differences in spike timing can have big effects on subsequent neuronal populations, while affecting firing rates only marginally (London et al., 2010). It was argued that spike timing might be important in the initial feedforward pass of incoming sensory information, but much less so in ongoing cortical feedback loops.

2.3.2 Decoding spike times

One important question is how the information in spike timing could be "read out" by subsequent processing stages. There have been a number of modelling attempts as well as some first experimental results elucidating this issue. First, there is the well-known tempotron model (Gütig and Sompolinsky, 2006), which illustrates how a single neuron could learn to distinguish different spike patterns. There has also been research into whether spike timing dependent plasticity (STDP) might play a role, and computational modelling has indicated that STDP enables neurons to learn to respond to specific temporal input patterns. Moreover, such patterns could be recognized ever faster as learning progresses, as neurons identify and become sensitive to the very first information-bearing spikes in each distinct pattern (Guyonneau et al., 2005; Masquelier et al., 2008). This could provide a mechanism by which temporal patterns could be learned and the information contained in them passed on to other cortical areas.

Additionally, bigger hierarchical models have been developed where the influence of spike timing can be studied on larger sets of neurons spread out over different areas. For example, one model of the visual hierarchy showed how spike timing can enable rapid object classification by acting on only the very first spikes elicited by each processing stage (Thorpe et al., 2001).

Finally, there has also been some work in trying to determine how various features could be encoded by different neural response codes and how these codes could interact. Such multiplexed codes would contain information on different levels (Panzeri et al., 2010). For example, complementary stimulus features could be transmitted in parallel by codes operating with both spike counts and relative spike timing.

2.4 Human fingertip mechanics

2.4.1 Viscoelastic properties and mechanical modelling

The human fingertip exhibits complex mechanical properties, which become evident during grasping and object manipulation when the fingertip undergoes considerable deformation due to the forces acting upon it. After deformation, seconds to minutes of inactivity are required for recovery to the resting state, a phenomenon called hysteresis (Pubols, 1982a,b; Pawluk and Howe, 1999; Serina et al., 1997). The most common

description of the mechanical properties of the fingertip is in terms of models of viscoelastic materials (Pawluk and Howe, 1999; Serina et al., 1997). This means that the amount of indentation of the skin does not only depend on the amount of force that is present, but also on the stimulation history.

Most studies so far have looked into the mechanical properties of the fingertip in the direction orthogonal to the skin surface (Jindrich et al., 2003). It was found that fingertip stiffness is nonlinear and increases with indentation depth. High hysteresis was also reported. Other studies have examined the mechanical properties of skin in the direction tangential to the skin surface (Nakazawa et al., 2000; Pataky et al., 2005). Stiffness has been found to be constant and therefore independent of force impact (but see our results in Chapter 4). Interestingly, anisotropic effects have been found with stiffness and hysteresis values varying between different tangential directions.

Modelling the mechanical behaviour of the fingertip is complicated by different properties of the different tissue layers that are involved and their arrangement in three dimensions. Some of the recently developed models allowed for the calculation of correlations between stress/strain distributions within the fingertip tissue and the response properties of mechanoreceptors (Dandekar et al., 2003; Wu et al., 2006a; Sripati et al., 2006b). However, so far these studies have been limited to firing rates of tactile afferents and ignored spike timing (see also Sec. 2.2.1 above). Moreover, only the responses of mechanoreceptors located at the flat part of the fingertip or deep within tissue have been correlated with stress/strain patterns within the skin. This is likely to change once more sophisticated models are developed that also take into account more complex stimuli (see Wu et al., 2006b, for a recent model of how stress/strain patterns change depending on force direction).

2.4.2 Effects on neural response and perception

Tactile indentations delivered in the direction orthogonal to the skin surface and following each other in quick succession may exert marked effects on the firing rates of tactile afferents (e.g. Pubols, 1982b; Werner and Mountcastle, 1965; Lindblom, 1965). This manifests itself as a reduction in firing rate for subsequent indentations. It has been argued that this response attenuation is not due to properties intrinsic to mechanoreceptors (e.g. rate adaptation), but rather due to the mechanical properties of the skin, as changes in firing rates exhibit significant correlations with changes in the

effective stimulation (Pubols, 1982b): skin hysteresis resulting from previous stimulation causes any subsequent stimuli to have less impact on the skin, as the skin has not yet reverted back to its original resting state and due to nonlinear stiffness of the fingertip the resulting skin indentation will be smaller.

However, these effects have not been systematically studied in more complex settings (for example when stimuli are not delivered orthogonally to the skin surface) and it is unknown whether they affect spike timing or only firing rates. Most experimental studies deliberately limit the effect of past stimuli by either allowing for sufficient time between successive stimuli such that the skin can revert back to its resting state or using only small forces close to the mechanoreceptors' thresholds so that hysteresis is minimal. However, in real-world object manipulation and grasping tasks, forces leading to considerable skin deformations generally follow each other in quick succession (Kunesch et al., 1989) and any confounding effects of past stimulation on neural responses would limit the accuracy at which the tactile system operates (Pubols, 1982a; Beitel et al., 1977; Kenton et al., 1971). We will address some of these questions in Chapter 4, where we analyze how neural responses change when stimuli are applied in quick succession. We will also examine a possible mechanism of how confounding effects due to fingertip hysteresis could be mitigated by the central nervous system.

As for whether the mechanical properties of the fingertip also affect perception, there is preliminary evidence that such history effects are visible in psychophysical experiments and affect the performance of subjects in a discrimination task (Singh and Fearing, 1998). Viscoelasticity has to be taken into account by any controller trying to achieve stable grasps, as contact points and forces need to change depending on the indentation history of the fingertip (Pataky et al., 2005). Whether such effects also influence perception is currently unknown.

2.5 Information theory applied to neurophysiological data

After discussing tactile processing, different coding schemes, and human fingertip mechanics, we now turn our attention to methods that allow us to estimate the amount of information that is transmitted by neural responses (e.g. spike trains or specific features of spike trains) about tactile stimuli acting on the skin. This allows for a principled comparison of different coding schemes (e.g. firing rates and spike latencies). It also gives an indication of what kinds of tactile features are reflected in the neural response

at all. We will first discuss the basics of how to estimate mutual information from electrophysiological data, and then discuss how to remedy some common problems that might bias the results.

2.5.1 Mutual information and entropy

Commonly in neuroscientific experiments, stimuli (s) are presented to human or animal subjects, while neural responses (r) are measured from the peripheral or central nervous system. The stimuli can range from visual gratings with different bandwidths, to tactile indentations at various forces. These stimuli are then varied with respect to a particular parameter that one is interested in. Responses can either be gathered directly from single or multiple neurons as spike trains, or by less invasive methods such as EEG or fMRI. One then often wants to quantify how much the stimulus influences the neural response. Moreover, one might be interested in specific aspects of the response, such as whether a given stimulus is represented in the measured spike counts or in the timing of single spikes. Such analyses provide insight into the neural code by allowing us to compare different coding schemes. Some of these schemes might be ruled out as viable candidates, if their information capacity is below that needed to explain human or animal performance in behavioural tasks.

A principled method of characterizing the relationship between stimuli and responses is given by information theory in the form of mutual information (Shannon, 1948, refered to as "rate of transmission" therein). This measure indicates how informative the recorded responses are about the presented stimulus, and is related to the accuracy with which one can classify a given stimulus from its response. It provides a general way to measure how much information one random variable provides about another one, without having to rely on classifiers which might rely on unrealistic assumptions or otherwise perform sub-optimally. For discrete probability distributions mutual information is calculated according to the following equation:

$$I = \sum_{s,r} p(s,r) \log \frac{p(s,r)}{p(s)p(r)} \tag{2.1}$$

$$= H(R) - H(R|S) \tag{2.2}$$

where $p(r)$ is the probability that a given response r is recorded, $p(s)$ is the probability that stimulus s occurs (as defined by the experimenter), and $p(r,s)$ is the joint distribution that a stimulus s elicits a response r. Furthermore, $H(R)$ is the entropy of the

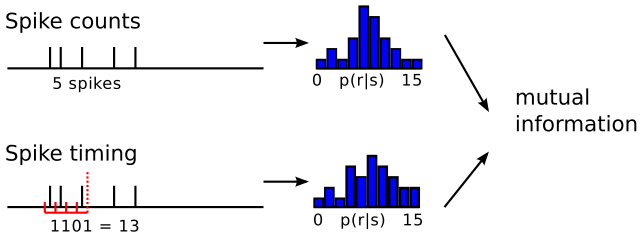

Figure 2.3: Illustration of how relative frequency distributions are estimated from neural spike trains.

response as assessed from the response variability across all stimuli, and represents the overall information contained in the neural response. $H(R|S)$ is the conditional entropy of the response given the stimulus, and represents noise in the system, i.e. information in the afferent signal that is not about the nominal stimulus parameter but may include information about the stimulus not accounted for by the model used, or measurement related or neural noise.

The entropy terms are calculated as follows:

$$H(R) = -\sum_{r} p(r) \log p(r) \tag{2.3}$$

$$H(R|S) = -\sum_{s} p(s) \sum_{r} p(r|s) \log p(r|s) \tag{2.4}$$

where $p(r|s)$ is the conditional probability that the stimulus s elicits response r. The probability distribution of the response given the data $p(r|s)$ and the marginal probability of the data $p(r)$ are estimated from recorded data. If stimulus and response spaces are discrete, one can use the relative frequencies as stand-ins for the true distributions and evaluation of the above equations is straight-forward. Fig. 2.3 illustrates how relative frequency distributions are commonly constructed for two coding schemes, namely spike counts (firing rates) and spike timing (temporal spiking patterns).

2.5.2 Avoiding bias due to low sample sizes

Often the amount of data that is gathered is relatively small (e.g. the number of trials collected is on the order to possible responses r or smaller), due to technical or time limitations. In these cases, the relative frequencies estimated from the response data r might not reflect the true probabilities very well. This can lead to considerable bias in estimating entropy terms. As this effect leads to an underestimation of entropy terms

and $H(R|S)$ is more affected by this than $H(R)$, overall the estimate of the mutual information will be biased upwards. That is, scarce data can lead to an overestimation of the amount of information that is transmitted (see Paninski, 2003; Quian Quiroga and Panzeri, 2009, for reviews).

There are several different approaches for dealing with the bias problem and ensuring that the estimates are robust. We will give a short overview of the most commonly used ones (see e.g. Panzeri et al., 2007, for a more extensive review).

First, some approaches try to estimate the bias so that it can subsequently be subtracted from the information estimate to arrive at corrected mutual information values. A simple approach is to shuffle the data with respect to the different stimuli, and then use the estimated information values from this new data set as an approximation of the bias. As the responses and stimuli are randomly paired, the mutual information should be close to zero and any values above that can be considered as stemming from bias. However, this method has been shown to be rather inaccurate. A better approach is to use a series expansion of the analytical estimate of the average bias (Treves and Panzeri, 1995; Panzeri and Treves, 1996), as this method takes some variables directly into account (such as the number of relevant responses), which are important for the bias term. Additionally, this method uses a Bayesian estimate of the number of actually occuring, distinct responses (relevant responses) in the data set (which could be smaller than the number of potential responses) to further improve the information estimate. This method has been shown to perform well, as long as the number of trials is at least as high as the number of distinct responses. More details on this method can be found in Sec. 3.2.2.2.

Second, other approaches try to directly arrive at an unbiased estimate, by using a fully Bayesian approach (Nemenman et al., 2002, 2004). In this so-called NSB (Nemenman-Shafee-Bialek) method, a (rather) flat prior distribution is put on the entropy estimate in order to ensure that the results are not biased. The authors of this method also show how the traditional estimates using relative frequencies bias the results by (implicitly) using too narrow a prior. While the NSB method is computationally rather expensive, it has been shown to work very well, especially when the number of different stimulus conditions is high.

A final set of methods calculate a lower bound to the mutual information from the performance of a classifier on the data set. The most common approach when dealing with neuronal spike trains is the metric space method (Victor and Purpura, 1996, 1997).

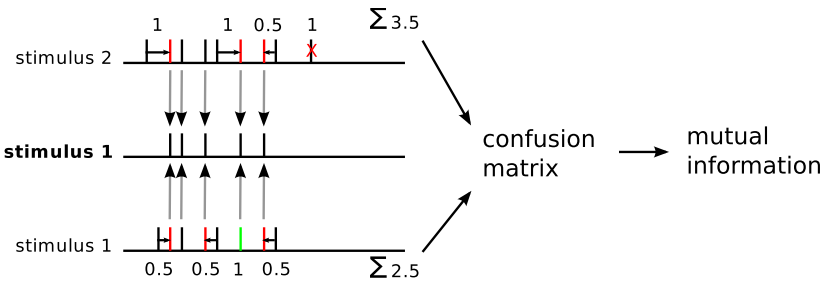

Figure 2.4: Illustration of the metric space method. Distances between every pair of spike trains are calculated, which are used to build a confusion matrix from which an estimate of the mutual information is calculated.

In this method, the classifier relies on spike train distances, which for any pair of spike trains measure how "far apart" these two responses are. This is done by defining a set of transformation rules for how to transform one spike train into another one: Single spikes can be inserted, deleted or moved in time. Each rule application incurs a certain cost and the overall distance of one spike train from another one is calculated as the minimal cost that has to be incurred to transform one spike train into the other one (see Fig. 2.4 for an illustration). A confusion matrix is then computed from the pairwise distances by assigning a spike train to the stimulus condition with the lowest average spike train distance. Finally, an estimate of the mutual information is calculated from the confusion matrix. This method is also described in Sec. 3.2.2.3, and results using this method are presented in Secs. 3.3.1 and 4.3.3. Alternative spike train distances have been developed (e.g. van Rossum, 2001), which are easier to calculate, and there are extensions for multi-neuronal recordings as well (Aronov et al., 2003).

In this overview, we have focused on approaches that are currently used within the computational neuroscience community, but we would like to note that there is a variety of other estimators for entropy and mutual information (see e.g. Hausser and Strimmer, 2009, for a recent approach) that are of potential use.

Chapter 3

Information content of different coding schemes

3.1 Introduction

In this chapter, we ask to what extent human primary tactile afferents carry information about complex tactile features in their spike times and their firing rates. As described in Sec. 2.3, spike timing has been shown to carry information in various neural circuits, sometimes in excess of the information contained in firing rates. In spite of this, most studies on the encoding of information in peripheral tactile afferents have focused solely on firing rates estimated from spike counts of individual neurons within a defined period (for reviews see Torebjörk et al., 1987; Johnson and Hsiao, 1992; Craig and Rollman, 1999; Johnson et al., 2000; Goodwin and Wheat, 2004). Recent findings in humans suggest, however, that information about complex mechanical fingertip events are conveyed by the relative timing of impulses in individual members of ensembles of afferents and that this code can operate fast enough to account for the use of tactile signals in natural manipulation tasks (Johansson and Birznieks, 2004). Specifically, the relative timing of the first spikes elicited when contact with objects is made provides information about the shape of the contacted surface as well as force direction. That is, changes in either stimulus parameter can differentially alter the latency of the first spike of individual afferents with little systematic effects on the distribution of first-spike latencies within the population as a whole. Information about curvature of objects contacted by the human fingertip and the direction of the contact

force can be conveyed both by firing rates assessed by spike counts (Goodwin et al., 1997; Birznieks et al., 2001; Jenmalm et al., 2003) and the precise timing of spikes in populations of tactile afferents (Johansson and Birznieks, 2004), but it is unclear how these two coding schemes compare regarding the amount of information conveyed. A related question is whether spike counts and spike timing transmit the same (redundant) or complementary (independent) information, and whether additional information is carried jointly by both codes (synergistic information) (see Nelken et al., 2005; Mikula and Niebur, 2005). Finally, little is known about whether information about fingertip stimuli conveyed by different coding schemes depends on the position where the afferent terminates in the fingertip in relation to the stimulation site.

In this chapter, we employ information-theoretic measures to compare information transmitted in individual afferents by spike counts and spike timing, with a special emphasis on the very first spikes elicited by each afferent after stimulus onset. We find that spike timing contains more information about tactile stimuli than spike counts, and that the latencies of each afferent's first spike transmit at least as much information on average as spike counts, and sometimes much more. Additionally, we show that the information conveyed by the two different coding schemes seems to be mainly independent.

3.2 Methods

3.2.1 Data collection and experimental setup

The data used in this study was collected at the Laboratory of Dexterous Manipulation, Umeå University, Sweden, by Roland Johansson and collaborators. I did not take part in the design of the experiments or the data collection. The following Secs. 3.2.1.1–3.2.1.3 give a quick overview of the experimental setup insofar as it is relevant for the subsequent analysis (please refer to Saal et al., 2009, for more details).

3.2.1.1 Microneurography

Neural data was recorded from thirty-three human volunteers using microneurography (see Gandevia and Hales, 1997, for a review). In this procedure, a very thin electrode is inserted into the median nerve between a subject's shoulder and elbow. Neural sig-

nals were recorded at 12.8 kHz and filtered (bandwith 0.5–5 kHz). To ensure that only activity from single afferents was considered, off-line spike sorting was used (both by computational analysis and manual inspection of wave-forms). Receptive field outlines and centres ("hot spots") were found using von Frey hairs (nylon filaments) that were systematically stroked across the fingertip. Location of each individual receptive field centre was mapped onto a "generic" fingertip, as described by Birznieks et al. (2001). Afferents were classified according to their firing patterns and which stimuli were able to elicit responses (Vallbo and Johansson, 1984). In this study we were only concerned with fast-adapting type I (FA-I) and slowly-adapting type I (SA-I). Briefly, FA afferents respond only to changes in skin deformation, whereas the SA afferents show an ongoing response during periods of static skin deformation (see Fig. 3.1A, middle panel, and also Sec. 2.2.1).

3.2.1.2 Stimulation protocol

Stimuli were delivered using a custom-made device that allowed the probe to be moved in three dimensions (Birznieks et al., 2001). Three linear motors were connected to the probe via stiff rods to achieve this. A six-axis force/torque transducer was attached to the probe for measuring the applied forces. The part of the probe that was in contact with the fingertip skin could be exchanged to allow for surfaces of different curvatures (Jenmalm et al., 2003). All surfaces were covered with silicon carbide grains in order to increase friction between them and the fingertip skin so as to prevent accidental slips of the probe.

A single trial consisted of an inter-stimulus interval (250 ms); a protraction period (125 ms) in which the normal force (as well as the tangential force, if present) was ramped up following a sinusoidal function; a plateau phase (250 ms) at which the force acting on the skin was kept constant; and a retraction period (125 ms) in which the force was ramped down (see Fig. 3.1A, middle panel). A minimum normal force of 0.2 N was maintained during the inter-stimulus intervals to ensure that the probe was in contact with the skin at all times. Maximum normal force during the plateau phase was 4 N. Some trials had an added radial, distal, ulnar, or proximal component, resulting in an angle of 20° relative to the vector orthogonal to the skin surface. In these trials, the tangential force reached a maximum of 1.4 N and changed in synchrony with the normal force. Time courses and force magnitudes were chosen so as to simulate load and grip forces during grasping of an object weighing $\sim$280 gr (see Johansson and

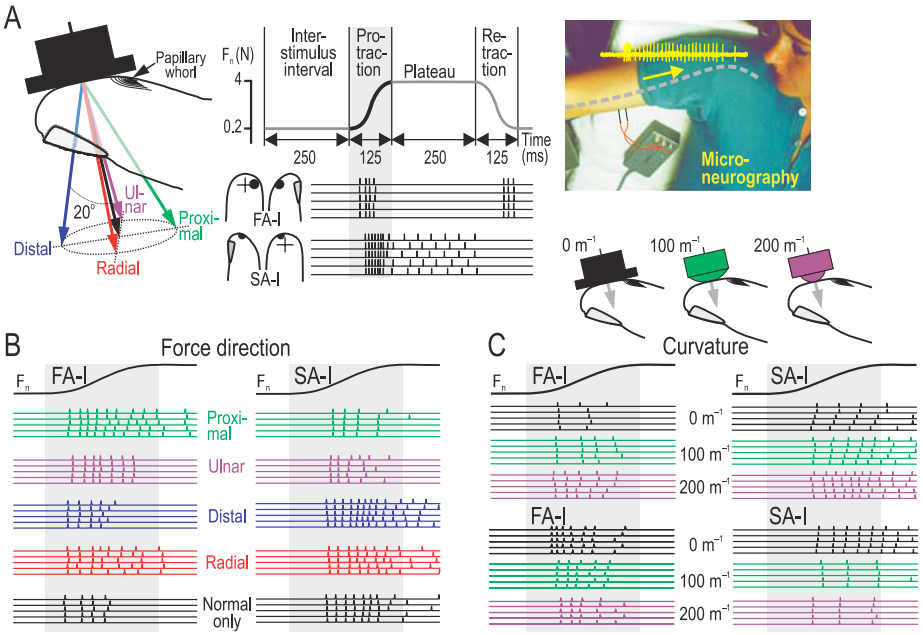

Figure 3.1: Responses to fingertip force direction and contact surface shape by human tactile afferents under conditions representative for object manipulation tasks. **A.** The fingertip was stimulated with objects of different surface curvature (bottom right) in five different directions (left). The time course of force application, illustrated with the normal force component, is shown in the middle panel. Responses from a FA-I and a SA-I afferent to repeated stimuli with normal force only ($n = 5$) are displayed below. Microneurographic recordings were obtained from the median nerve (right). **B.** Impulse ensembles in a FA-I and a SA-I afferent to repeated force stimuli ($n = 5$) applied in different directions (color coded). **C.** Impulse ensembles in two FA-I and two SA-I afferent to repeated stimuli ($n = 5$) with normal forces applied with a flat surface and with two spherically curved surfaces (5 and 10 mm radius). Figure prepared by Roland Johansson (taken from Saal et al., 2009).

Westling, 1984; Westling and Johansson, 1984).

A stimulation sequence consisted of 50 force stimuli, 10 in each of the five force directions. The first five repetitions of each stimulus direction were in a fixed order (*D, U, P, N, R*), while for the rest of the sequence, stimuli followed each other in pseudo-random order (see also the description of the regular and irregular series in Sec. 4.2.1). Such sequences were applied with each of three different stimulus surfaces (Fig. 3.1A, bottom right panel). One was flat (curvature 0 m^{-1}), and the other two were spherically convex with radii of 10 and 5 mm (corresponding to curvatures of 100 and 200 m^{-1}),

respectively. These curvatures spanned the range of what humans comfortably deal with at force levels such as those used in this study, i.e. objects with curvatures of the same order of magnitude as those of the fingertips (see Jenmalm et al., 2003). A single stimulation sequence containing trials in the different force directions as described above was delivered with a fixed surface curvature, while the curvature was exchanged between sequences (see Jenmalm et al., 2003, for details). Depending on the stability of the recording, 1–3 stimulation sequences were executed with each surface curvature.

3.2.1.3 Afferent sample

We analyzed 43 FA-I and 49 SA-I afferents that were included in a sample of in total 196 tactile afferents with receptive fields distributed over the entire distal segment of the finger (Birznieks et al., 2001). We focused on these types of afferents because their density is very high in the fingertip (see Sec. 2.2.1) and earlier studies have shown their involvement in signalling force direction and object curvature (Birznieks et al., 2001; Jenmalm et al., 2003). To minimize potential bias effects introduced by a low number of stimulations we required that afferents had been exposed to at least 150 stimulations (corresponding to 3 stimulation sequences, one with each surface curvature), and we pooled data from trials with all curvatures when we examined information about force direction and from all force directions when information about curvature was examined. Thus, between 30 and 90 trials were available for each force direction, and between 50 and 150 for each curvature. We also analyzed the subset of afferents for which more than 300 trials in total were available ($\geq$2 stimulation sequences $\times$ 3 curvatures) to calculate interaction effects (see results in Sec. 3.3.2) without the need for pooling.

3.2.2 Information theoretic analysis

3.2.2.1 Choice of response window

All our analyses focused on the protraction phase of the force stimuli, i.e. the first 125 ms after the onset of the force increase. This time window included most spikes of the FA-I afferents elicited by the stimuli since FA-I afferents mostly respond to only deformation changes of the fingertip (see FA-I in middle panel of Fig. 3.1A). Because of their sensitivity to static skin deformations, the SA-I afferents could also fire during

the plateau phase of the force stimulations, but at lower rates than during the force protraction phase (see SA-I in middle panel of Fig. 3.1A). Figure 3.1B shows examples of responses in one FA-I and SA-I afferent to stimuli in the five force directions with the flat contact surface, while Fig. 3.1C shows examples of responses in two FA-I and two SA-I afferent to normal force stimuli with each of the three surface shapes. As previously shown, the number of impulses elicited during the protraction phase was related to both shape and force direction (Birznieks et al., 2001; Jenmalm et al., 2003). Moreover, the timing of spikes in ensembles of afferents of each type also changed depending on both shape and force direction (Johansson and Birznieks, 2004). Changes in either of these parameters could differentially influence both the spike counts and the spike timing amongst afferents without having significant effects on the distribution of these response measures within the afferent populations as a whole (Johansson and Birznieks, 2004). The first spikes elicited by the force stimuli could pass the recording electrode as soon as 35 ms after stimulus onset, but many afferents responded later. The overall mean of first-spike latencies across afferents and stimulation conditions was 55 ms (for distributions of latencies of the first and second spike see Figs. 3 and 4 in Johansson and Birznieks, 2004). Several factors contributed to the delay between the onset of the stimulation and the appearance of the action potentials at the recording electrode and to the variation of the first-spike latencies between afferents. These include nerve conduction delays, which vary across afferents (Kakuda, 1992) and the afferents threshold sensitivity to the force stimulation, which also varies across afferents, combined with the fact that the force initially increased very slowly because of the sinusoidal waveform of the stimulus.

3.2.2.2 Information in spike counts

We assessed mutual information about the stimulus parameters (force direction and surface curvature) by counting the number of spikes recorded from a single afferent during the force protraction phases (125 ms fixed window). Because this number ranged from 0 to 12 (mean: 2.2; median: 2) and from 0 to 13 (2.6; 2) amongst stimulation for the FA-I and SA-I afferents, respectively, the corresponding sizes of the response space was 13 and 14. From these counts we calculated relative frequencies and then calculated the mutual information using the standard formula (see Sec. 2.5.1 for mathematical details). Since sparse data can lead to an upward bias of estimates of mutual information, this estimate was corrected by subtracting the bias, B, estimated

as follows (Panzeri and Treves, 1996):

$$B = \frac{1}{2N\log(2)}\left(\sum_s R_s - R - (S-1)\right) \tag{3.1}$$

where N is the sample size (all stimuli delivered), R is the size of the response space that was not empty (relevant responses, cf. Sec. 2.5.2), R_s is the number of response bins that were not empty conditioned on the stimulus, and S is the number of different stimuli. This estimate of bias is considered correct provided that the number of trials for each stimulus condition is larger than the number of possible distinct responses, i.e. the size of the response space (Panzeri et al., 2001). This condition was fulfilled in the present study. In addition, we assessed the validity of our bias-compensation methods by estimating the parameters of a Markov model of spike generation (cf. Montemurro et al., 2007b) applied to signals in 10 FA-I and 12 SA-I afferents for which more than 300 trials were available. Generation of spike trains from these models and calculations of the remaining bias for sample sizes representative for afferents with the least data indicated negligible differences ($<$0.01 bit for both afferent types and for both stimulus parameters).

3.2.2.3 Information in precise timing of spikes

We used the metric-space method (see Sec. 2.5.2 and Victor and Purpura, 1996, 1997) to estimate mutual information contained in the precise timing of all spikes elicited during the force protraction phase. Because of the potentially large response space, use of other methods for estimating mutual information might have resulted in large biases (Panzeri et al., 2007). The metric-space method calculates a lower bound on the mutual information and permits assessment of the temporal precision by which spike timing can contribute to stimulus classification. The method provides a measure of difference, or distance, between spike-trains that reflects the minimal cost associated with transforming one spike train into the other. Single spikes are inserted and deleted at a cost of unity and shifting a spike in time incurs a certain cost q per unit of time. For each afferent we first computed the distance between each spike train recorded and every other spike train. We then computed the average distances between every given spike train r and all other spike trains r' belonging to a certain stimulus condition s (i.e. surface curvature or force direction), denoted by $d(r, r_s)$, as follows:

$$d(r, r_s) = \left\langle d(r, r')^z \right\rangle^{\frac{1}{z}} \tag{3.2}$$

where $d(r,r')$ indicates the distance between spike trains r and r', angle brackets denote the average with respect to all r observed under condition s, and z, set to -2 in this study, refers to the contribution of outliers to the average distance[1]. A lower bound on the mutual information was then computed from a confusion matrix that indicated the condition from which a spike train originated and the condition to which it was most similar, assessed by the average distance computed as described above. To estimate the bias of the mutual information terms, we reassigned spike trains to random conditions and computed the mutual information with the metric space method. The bias term used was the average of the outcome of 20 such operations. By varying the cost parameter q, we assessed the temporal resolution required for obtaining maximum mutual information for each afferent and stimulation condition. We used eight different q values that ranged from 8 to 1024 derived from powers of 2. This corresponds to a temporal precision ranging from a 125 ms window up to a 1 ms window. In Sec. 3.3, we report mutual information as obtained using an optimal temporal resolution for each afferent. To cross-validate the outcomes of the metric space method and the direct method used with spike counts, we also computed mutual information at $q = 0$ where the cost of moving spikes in time is zero and the distance between two spike trains is just the difference in spike counts. It was found that both methods provided consistent results: Out of the 43 FA-I afferents analyzed, the information about force direction and curvature estimated by the metric-space analysis exceeded that computed from spike counting in only 2 and 4 afferents, respectively. The corresponding numbers for the 49 SA-I analyzed were 1 afferent and 3 afferents, respectively. Mutual information values for force direction calculated using the metric-space method were on average 0.13 bits (standard deviation: 0.10 bits) and 0.15 (0.07) bits lower than the values obtained from the direct methods for FA-I and SA-I afferents, respectively. For curvature, the corresponding reduction was 0.04 (0.13) and 0.07 (0.13) bits. As a percentage of the absolute values, the mutual information calculated using the direct method was 134% and 148% of that calculated by using the lower bound for force direction and curvature, respectively.

[1] Our value of z is a common choice and we ensured that choosing other values did not considerably affect the results.

3.2.2.4 Information in first spikes

To assess the mutual information contained in the timing of the first spikes in individual afferents, we divided the force protraction phase into 2 ms long consecutive bins and allocated the first spike elicited by each stimulus into the appropriate bin. From this, we calculated the mutual information with bias correction as detailed in Sec. 3.2.2.2. We chose 2 ms bins because this bin size approximately matched the temporal resolution for maximal information as estimated by the metric space method when applied to the first spikes only. Furthermore, for nearly all afferents, the variability in the timing of first spikes across all stimulus conditions was contained within a 30 ms window (see Johansson and Birznieks, 2004), which resulted in similar sizes of the response space as used with spike counts. Indeed, control analyses showed that the mutual information computed with an appropriately positioned 30 ms window matched that obtained with the procedure we used. That is, the remaining bias, computed as described for spike counts, was negligible ($<$0.02 bits for FA-I afferents and $<$0.01 bits for SA-I afferents).

3.2.2.5 Information in spike presence

The information in spike presence refers to information that is transmitted by the mere presence of a neural response to a given stimulus, as opposed to no response (i.e. no spiking). It was computed as in Sec. 3.2.2.2, however with a binary response space, i.e. '0' for no spiking activity, and '1' if there were any number of spikes present.

3.2.2.6 Redundant and synergistic information contained in spike counts and first spike latencies

To assess if information transmitted by firing rates and spike timing is redundant, independent, or synergistic we compared information available when both decoding schemes were used simultaneously with that contained in timing of first spikes and in spike counts independently according to the following equation:

$$Syn = I(R_c, R_f; S) - I(R_c; S) - I(R_f; S) \tag{3.3}$$

where $I(R_c, R_f; S)$ represents the joint mutual information between the stimulus and both coding schemes (spike counts (c) and first spike latencies (f)) and $I(R_c; S)$ and $I(R_f; S)$ stand for the mutual information about the stimulus with spike counts and first

spike latencies, respectively. Positive values of *Syn* indicate presence of synergistic information, where information contained in the two coding schemes jointly exceeds the sum of the information contained in each coding scheme if considered independently. Accordingly, negative values of *Syn* indicate that the two coding schemes provide redundant information. A value of zero means that both coding schemes contribute information independently. To minimize potential bias in the estimates of *Syn* related to the small sample sizes, we computed the joint information by using the recently introduced shuffling method (Montemurro et al., 2007b; Panzeri et al., 2007)[2]. In addition, to correct for confounding effects due to the low number of different stimuli used, we also compared the distribution of *Syn* values based on the original data with a distribution generated under the assumption of independence between spike counts and latencies, where spike counts and latencies were randomly reassigned across afferents.

3.3 Results

The Results are divided into four sections: First, in Sec. 3.3.1 we compare the mutual information between stimulus and response estimated from the number of spikes elicited during the protraction phase of the force stimuli and from the precise timing of all spikes elicited during the same 125 ms epoch. Second, in Sec. 3.3.2 we compare information in spike counts with that contained in the timing of the first spikes elicited by the force stimuli. Next, in Sec. 3.3.3 we examine for each coding scheme how information about the direction of contact force and the curvature of the contact surface in single afferents might depend on where they terminate in the fingertip skin relative to the site of stimulation. Finally, we analyze to which degree first spikes and spike counts may convey redundant, independent or synergistic information about the stimuli in Sec. 3.3.4.

Preliminary versions of the results in Secs. 3.3.1 and 3.3.2 have appeared in the author's MSc thesis (Saal, 2007).

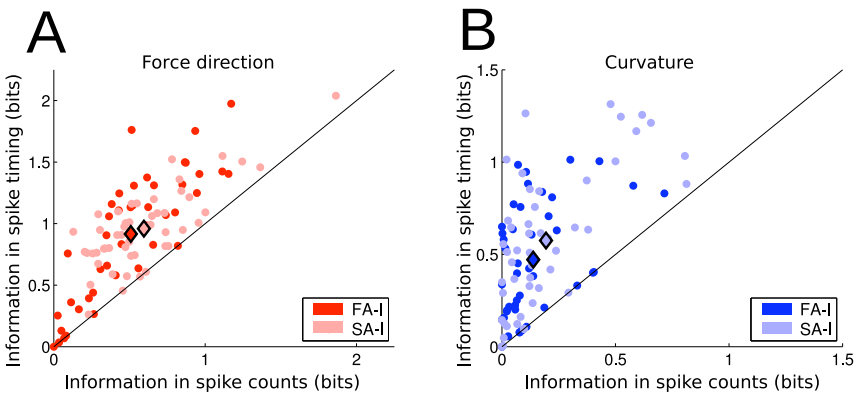

Figure 3.2: Amount of mutual information provided in the responses of FA-I and SA-I afferents. **A.** Force direction. **B.** Curvature. Information in bits in spike counts versus maximal information in spike timing. Each dot represents data from an individual afferent. Points on the diagonal line indicate that spike count and temporal information are equal. Dots framed in black denote population means.

3.3.1 Spike times versus spike counts

As expected, the information obtained from the timing of all spikes was found to be substantially greater than that obtained from spike counts. This applied to both FA-I and SA-I afferents and to information about both force direction (Fig. 3.2A) and surface curvature (Fig. 3.2B). In order not to introduce any bias due to low sample sizes into our estimates of the mutual information, we calculated a lower bound on the information in spike timing and compared this lower bound with a direct estimate of the information in spike counts. On average, the estimated information content for individual FA-I afferents was 0.51 $\pm$ 0.33 bits (mean $\pm$ 1 SD) for force direction as assessed by spike counts and 0.92 $\pm$ 0.52 bits (lower bound) if spike timing was taken into account. The corresponding data for the SA-I afferents were 0.60 $\pm$ 0.33 bits for spike counts and 0.96 $\pm$ 0.32 bits as the lower bound for spike timing. Thus, considering spike timing provided an average information gain of at least 0.41 bits (80%) for the FA-I afferents and 0.36 bits (60%) for the SA-I afferents as compared to spike counts. For surface curvature, spike timing conveyed information by on average 0.47 bits and 0.57 bits (lower bounds) for the FA-Is and the SA-Is, respectively, which cor-

[2]The shuffling method has been found to overcompensate slightly, i.e. the estimated bias might be higher than the actual bias, which should be taken into account when interpreting our results (see Sec. 3.3.4).

responded to 336% and 300% of the information about curvature in spike counts (0.14 $\pm$ 0.16 bits for the FA-I and 0.19 $\pm$ 0.22 bits for the SA-I afferents). According to a two-way ANOVA, there was a significant effect of spike count vs. temporal coding ($F(1,181) > 47.04$; $p < 0.001$ in both cases), but not between different afferent types ($F(1,181) < 3.73$; $p > 0.05$). We were also interested in the temporal resolution at which information was encoded in tactile afferents. We found that the median temporal resolution at which afferents transmitted maximal information was 16 ms and a two-way mixed design ANOVA failed to indicate an effect on this of both afferent type (FA-I, SA-I; $p = 0.53$) and stimulus feature (curvature and force direction; $p = 0.24$; treated as repeated measures). However, we noted a marked variability amongst afferents; the lower and upper quartiles for the FA-Is were 32 ms and 8 ms and those of the SA-Is were 64 ms and 8 ms. Estimates of the mutual information for temporal resolutions other than the optimal ones resulted in fairly small changes. Different spikes within single spike trains may, however, contain information at different temporal resolutions (see analysis of first spikes in Sec. 3.3.2 below).

Based on a lower bound on the mutual information for spike timing, we conclude that the tactile afferents examined (FA-I and SA-I) on average convey at least 60% more information about force direction and 200% or more information about surface curvature than spike counts do.

3.3.2 First spikes versus spike counts

The information about force direction provided by timing of the first spikes, irrespective of when they occurred during the force protraction period, was 0.48 $\pm$ 0.32 bits (mean $\pm$ 1SD) for the FA-Is and 0.59 $\pm$ 0.36 bits for the SA-Is. The corresponding numbers for surface curvature were 0.32 $\pm$ 0.23 bits and 0.45 $\pm$ 0.34. The estimate of the information conveyed by the first spikes in the two types of afferents did not differ significantly regarding force direction ($p = 0.13$; two-sample Kolmogorov-Smirnov test) but did so for surface curvature ($p = 0.04$). We observed no simple relationship between the time during the protraction phase when an individual afferent reached maximum information and the information conveyed by the first spike (Figs. 3.3C and D). Tests of correlations between the amount of information in the first spikes and this time indicated significance only for force direction in the SA-I afferents ($r = 0.48$; $p < 0.001$). The amount of information was smaller for afferents reaching maximum information early as compared to late (Fig. 3.3C). Afferents that transmitted a large

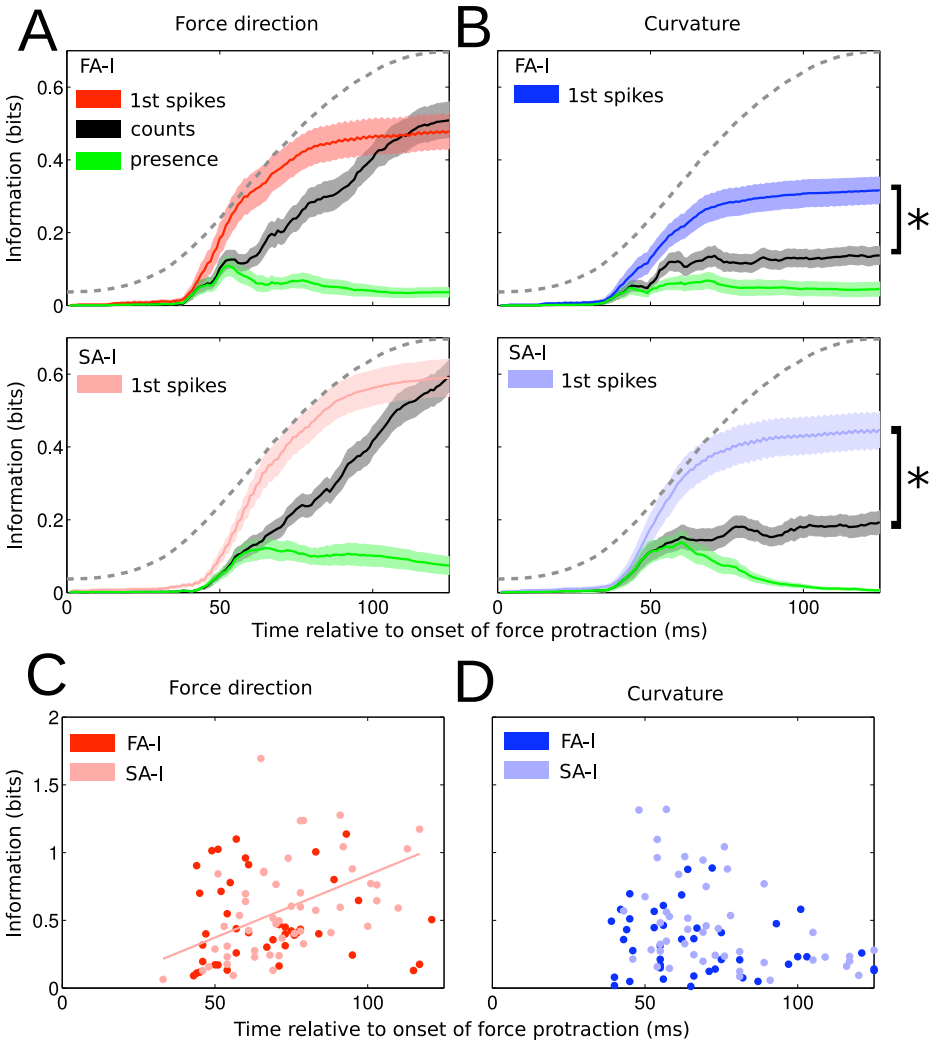

Figure 3.3: Information in bits for FA-I and SA-I afferents about force direction and curvature provided by spike counts and by timing of first spikes. **A and B.** Information about force direction and surface curvature, respectively, estimated from spike counts (black line) and from first (red or blue) spike latencies and displayed as a function of time during the force protraction phase. Information in spike presence over time is displayed by the green line. Data averaged across all afferents of each type. Shaded areas correspond to $\pm$ 1 standard error of the mean value given by the solid line. The dashed line denotes the profile of the normal force increase from 0.2 N to 4 N during the force protraction phase. **C and D.** Information in 1st spike latencies of individual FA-I and SA-I afferents about force direction and surface curvature plotted against the time when 90% of each afferent's maximum information was reached. Linear regression lines shown for data exhibiting a significant correlation.

amount of information in their first spikes also tended to transmit high information in their firing rates. That is, with each stimulation parameter (force direction and curvature) and for both afferent types (FA-I and SA-I) there was a significant positive correlation across afferents between information contained in spike counts and in first spike latencies ($p < 0.002$ in all four instances). For the FA-I afferents, the correlation coefficient was 0.50 for force direction and 0.58 for curvature. The corresponding values were 0.55 and 0.45 for the SA-I afferents.

We analyzed when information about force direction and curvature gradually became available in spike counts and in the timing of the first spike during the progression of the force protraction phase: we computed the information contained in the two encoding schemes during a series of time windows that all began at the onset of the force protraction phase and had durations that increased stepwise by 1 ms from zero to 125 ms. For each window we averaged information across all afferents in our sample irrespective of whether they had responded or not. Before an afferent responded (in any of the trials recorded), the transmitted information was considered to be zero. Figures 3.3A and B plot information about force direction and curvature provided up to the time of the end of each window (abscissa value). Overall, the timing of the first spikes conveyed more information than spike counts during the force protraction phase. For surface curvature (Fig. 3.3B), the first spike contained about twice as much information as spike counts even at the end of the protraction phase for both the FA-I afferents (229%; $p < 0.001$) and the SA-I afferents (237%; $p < 0.001$). With force direction the information contained in spike counts matched that in the first spike at the end of the protraction phase (Fig. 3.3A; $p = 0.55$ and $p = 0.90$ for FA-I and SA-I afferents, respectively; paired t-test). From Figures 3.3A and B it can be seen that the initial increase of information in first spike latencies and spike counts happens at around the same time. This early information is based on whether an afferent has already fired or not (information in spike presence), which was available from both coding schemes and is indicated in the plots in green.

All above analyses concerning afferent information about force direction were based on data pooled across surface curvatures and the analyses concerning surface curvature on data pooled across force directions. To check for interaction effects between force direction and curvature, we compared mutual information calculated with and without pooling data across conditions for a subset of afferents for which a sufficient number of trials was available. We found that pooling decreased the information available for

both spike counts and first spike latencies by roughly the same amount, suggesting that the influence of interaction effects was similar for both coding schemes. For force direction and FA-I afferents ($N = 10$), on average the mutual information decreased by 0.13 bits and 0.18 bits for spike counts and first spike latencies ($p = 0.38$; Wilcoxon signed rank test), respectively. With SA-I afferents ($N = 12$) the corresponding decreases were 0.24 bits and 0.18 bits ($p = 0.42$). For curvature and FA-I afferents ($N = 14$), the information loss was 0.14 bits and 0.22 bits for spike counts and first spike latencies, respectively (p = 0.09), and for SA-I afferents ($N = 15$) the loss was 0.24 bits for both coding schemes ($p = 0.76$).

We conclude that the first spikes provide significant information about our tactile stimuli earlier than spike counts do. Moreover, they appear to provide more than twice the information about curvature present in spike counts over the entire force protraction phase and similar amount of information about force direction as present in spike counts over the entire protraction phase.

3.3.3 Influence of afferent termination site

It has been suggested that afferents terminating at different locations with reference to the primary site of stimulation of a fingertip might preferentially contribute information about different aspects of tactile stimuli (Bisley et al., 2000; Birznieks et al., 2001; Jenmalm et al., 2003). We examined this issue by analyzing the information transmitted about curvature and force direction in timing of first spikes and in spike counts over the entire protraction period (see Fig. 3.4). The color of the circles on the contours of the fingertips in Fig. 3.4 illustrates the relative contribution by each afferent to information about surface curvature and force direction. That is, the circles are coloured in proportion to the information estimated for curvature (blue) and force direction (red), respectively. The area of the circles is proportional to the total information transmitted, computed as the sum of the information about curvature and force direction. Afferents terminating closer to the primary site of stimulation tended to convey more information about curvature than afferents terminating more remotely. This applied to either coding scheme and to both the FA-I and the SA-I afferents (FA-I spike counts: $r = -0.31$, $p = 0.045$; FA-I first spike latencies: $r = -0.49$, $p = 0.001$; SA-I spike counts: -0.41, $p = 0.003$; SA-I first spike latencies: $r = -0.38$, $p = 0.009$). Concerning force direction, we found no correlation with distance except in first spikes latencies of FA-I afferents, which tended to convey more information with greater distance ($r = 0.56$,

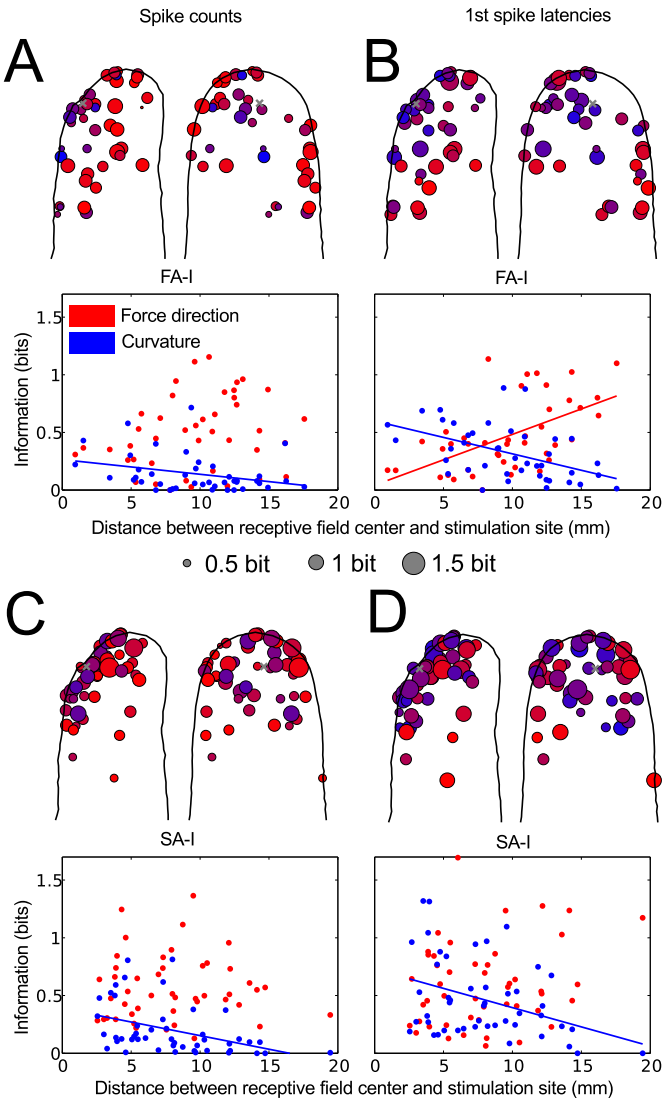

Figure 3.4: Information in bits about force direction and surface curvature for individual afferents plotted on the contours of the generic fingertip. Each circle represents data for a single afferent and the center of the circle corresponds to the hot spot of the afferent's receptive field. Grey crosses denote average stimulus contact points. The side-view of the fingertip includes afferents located on either side of the finger. The area of a circle is proportional to the amount of information that is transmitted by the afferent and the color reflects the stimulus parameter this information relates to, with blue standing for curvature, red for force direction and mixtures for both. The scatter diagrams show for each afferent amount of information conveyed plotted against the distance between the center of the afferent's receptive field and the primary site contact. Red and blue dots (and linear regression lines) represent information about force direction and curvature, respectively. **A and B.** FA-I afferents. **B and C.** SA-I afferents.

$p < 0.001$; Fig. 3.4B).

In conclusion, afferents terminating all over the fingertip seem to be important in conveying tactile features and there exists a distinct spatial relationship between the location on the fingertip where an afferent terminates and the type of feature that it is conveying information about.

3.3.4 Redundancy/synergy between different coding schemes

Finally, we addressed the question of whether individual afferents transmit the same (i.e. redundant) or complementary (i.e. independent) information, or whether both coding schemes (i.e. spike counts and first spike latencies) might even jointly carry information, that is synergistic information which would only be available from the combination of both coding schemes. We did this by calculating a measure of shared information between the two decoding schemes (see Sec. 3.2.2.6). Positive values of this measure indicate that the two coding schemes contain more information if operating jointly as compared to independently whereas negative values indicate that the coding schemes would provide redundant information. For force direction, this measure was on average -0.07 ± 0.23 and -0.13 ± 0.38 bits for the FA-I and SA-I afferents, respectively. The corresponding values for curvature were -0.06 ± 0.16 and -0.09 ± 0.18 bits. This means that for FA-I and SA-I afferents, on average just 8% and 10% of the information about force direction, respectively, represented redundant information. The corresponding values for curvature were 22% and 15%.

Most of the slight redundancy observed above can be explained by noticing that both coding schemes carry information about whether an afferent had fired during a given trial or not (see results on information in spike presence in Sec. 3.3.2). This means that there is necessarily some redundancy between both coding schemes, which might be sufficient to explain the low average redundancy reported here. We found that information in spike presence amounted to between 0.01 to 0.06 bits. Adjusting the syngery values for this brings both coding schemes very close to independence. The synergy measure showed considerable variance across afferents, where the signals in some afferents contained redundant information while others contained synergistic information (Fig. 3.5D, green lines). For both afferent types, redundancy was more common for force direction than surface curvature. The distances between the afferents' terminations in the fingertip and the primary stimulation site explained a small

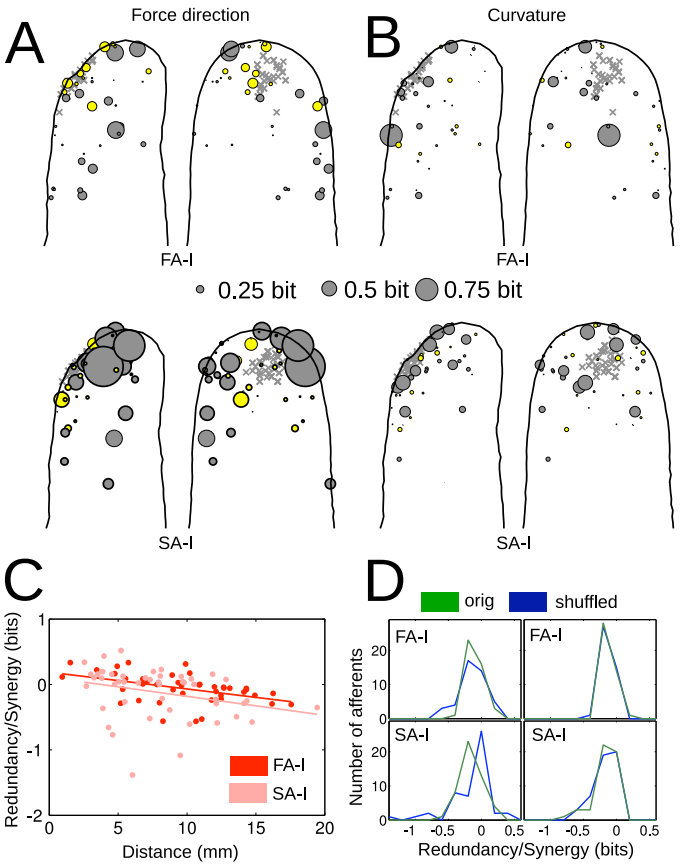

Figure 3.5: Information about force direction and surface curvature jointly represented in spike counts and first spike latencies of individual afferents and amount of synergistic vs. redundant information between the two coding schemes. **A and B.** Synergy (yellow circles) and redundancy (grey circles) in individual FA-I (top) and SA-I (bottom) afferents about force direction and surface curvature, respectively, plotted on the contours of the generic fingertip, where the diameter of the circles corresponds to the amount of overall information. Grey crosses denote stimulus contact points for all trials. **C.** Synergy/redundancy in the two coding schemes for force direction as a function of distance between the center of the afferent's cutaneous receptive field and the primary site of contact. Each symbol represents data from a single afferent. Positive and negative values imply synergistic and redundant information, respectively, given as the fraction of the information jointly contained in both coding schemes. **D.** Synergy/redundancy distributions for the original data (green) and shuffled data (blue) for FA-I (top) and SA-I (bottom) afferents, and for force direction (left) and curvature (right).

amount of the variance in the synergy measure that concerned force direction. Afferents that terminated close to the stimulation site showed positive synergistic effects between the two encoding schemes while afferents terminating farther away showed a gradual increase in redundancy (Figs. 3.5A, B, and C). For curvature, there were no reliable relationships in this respect. Finally, the bias correction method we used is prone to slightly underestimate the true mutual information (see Sec. 3.2.2.6 for more details). This might have resulted in a slight overestimation of redundacy in the current analysis.

Last, we examined whether the variance of the synergy measure might be explained by the fact that the small stimulus sets used (five force detections and three surface curvatures) might have introduced spurious synergy estimates between coding schemes. We compared the synergy measures computed from the recorded neural data with measures computed from a shuffled distribution generated under the assumption that both coding schemes acted independently (see Sec. 3.2.2.6; green and blue line histograms in Fig. 3.5D). Two-sample t-tests failed to indicate significant differences in the mean values of the synergy measures computed for original and shuffled data ($p > 0.36$ for force direction and curvature with both afferent types). Moreover, Kolmogorov-Smirnov tests did not find any significant differences between the original and the shuffled distributions ($p > 0.40$) apart from those for force direction in the SA-I afferents ($p = 0.03$).

Taken together, our results suggest that coding of surface curvature and force direction by tactile afferents by spike timing and spike counts contributes principally complementary information. We did not find evidence for either strong synergy or redundancy between both coding schemes over the whole population of afferents in our dataset.

3.4 Discussion

To our knowledge, the study presented in this chapter is the first to employ information theory for assessing information represented in tactile afferents supplying glabrous skin, i.e. body areas that actively contact objects. We found that the information contained in spike timing about the shape of the surfaces contacting the fingertip was at least 3 times that conveyed by firing rates, and for fingertip forces it was more than 1.6 times larger. Moreover, the timing of only the first spikes elicited in tactile afferents by our stimuli appeared to transmit the same amount of information about force direc-

tion as was transmitted by firing rates, and transmitted considerably more information about object curvature. Furthermore, our results indicate that coding of surface curvature and force direction in tactile afferents by means of spike timing and spike counts overall contribute complementary information. Given that multiple parallel pathways use tactile afferent information, it is possible that different processes and tasks use different codes. For example, relative spike timing may primarily support fast stimulus classification in the control of action. Firing rates, in contrast, might preferentially support perceptual mechanisms that operate under less time pressure.

Our results extend an earlier study (Johansson and Birznieks, 2004) by explicitly quantifying the amount of information in spike counts and first spike latencies and by providing a lower bound on the overall information provided by spike timing within the protraction period. Some effects of afferent location on the neural responses have been examined for spike counts before (e.g. Birznieks et al., 2001; Jenmalm et al., 2003), and here we extend these results to include effects on first spike latencies. Finally, while it had been observed earlier that spike counts and first spike latencies do not seem to be correlated within individual afferents (i.e. a higher firing rate does not imply a lower first spike latency, see Johansson and Birznieks, 2004), we provide a more thorough treatment and take into account spike presence and afferent location.

3.4.1 Information theoretic analysis

There are two main concerns, which generally arise in information-theoretic studies: First, the selection of an adequate stimulus set, and second, the elimination of bias introduced by small sample sizes. Stimulus parameters were chosen such that they reflected fingertip stimuli that commonly arise during natural object manipulation tasks, but we used only three curvatures and five force directions. We argue however that we made a reasonable trade-off between the size of the stimulus set and the number of trials that we could obtain with each type of stimulation given the time limitation for recording signals in single human afferents. We used three methods to ensure that low sample sizes did not affect the validity of our results: First, we used established bias-reduction methods and we confirmed the effectiveness of these methods in simulation and by cross-validating with alternative methods. This ensured that any potential remaining bias would be considerably smaller than the effect sizes we report. Second, we calculated a reliable lower bound on the mutual information in spike timing and compared this bound with a direct estimation of the information in spike counts. This

resulted in a robust lower bound on the effect size. Third, to increase the effective sample size we pooled data from trials with all curvatures when analyzing information about force direction and vice versa. We verified that pooling did not affect the comparison between different coding schemes (see end of Sec. 3.3.2).

Werner and Mountcastle (1965) quantified information about the amplitude of skin indentations in spike counts in slowly-adapting afferents innervating the hairy skin of the hind legs of cats and monkeys. Both the spike counts and the information transmitted (up to 2 bits) during the initial 100 ms of the stimulation were higher than what we report. That study was however designed to maximize the number of spikes elicited per stimulus. To that end, the skin was shaved and, for each afferent, the spot of maximum sensitivity was subjected to pointed skin-indentations. In contrast, our approach was to obtain a representative picture of information conveyed in the population of afferents by analyzing information in all responsive afferents while stimuli were delivered to the same site irrespective of the location of the afferents terminals in the skin. Hence, many of the afferents analyzed were just weakly excited. Furthermore, we assessed information about more complex stimulus features that should be harder to discriminate than stimulus intensity (which was kept about constant). Despite that, we found that some afferents could be highly informative (see Secs. 3.3.1 and 3.3.2).

3.4.2 Latency coding in sensory systems

While we provided a lower bound on the overall information contained in spike timing during the protraction period, we focused specifically on the information in the latencies of the very first spikes elicited by each new stimulus. This is because such latency codes can transmit information very quickly, as there is no need to wait for subsequent spikes in the read-out stage. Latency codes have been considered in other sensory systems as well. For example Gollisch and Meister (2008) considered first spike latency coding in the retina. In that study, the authors focused on relative latencies between pairs of neurons. In our analysis we didn't include relative latencies between different afferents, as we only had data from single-unit recordings available. In the somatosensory system, spike timing has been found to be highly informative about stimulus properties in the rat thalamus (Foffani et al., 2009). It was found that first spike latencies carried a considerable amount of that information. This analysis agrees with findings in the rat somatosensory (barrel) cortex, where spike timing offered additional information over spike counts in encoding stimulus location. Again,

the very first spikes were found to be most informative in this analysis. Studies in auditory (Furukawa and Middlebrooks, 2002) and visual (Gawne et al., 1996) cortex have also emphasized the importance of spike latency. In our analysis, we found that the information in the first spike latencies was around 60% of that in overall spike timing. However, we only calculated a lower bound on the information in spike timing and comparisons for spike counts showed that this bound might not be particularly tight (see Sec. 3.2.2.3). Still, the information in first spike latency covers a considerable fraction of the overall information.

3.4.3 Complementary information in spike counts and first spike latencies

Although we found significant positive correlation between the information transmitted by precise spike timing and firing rates for both curvature and force direction, the two coding schemes largely conveyed independent information about both force direction and curvature. As such, these results agree with previous analyses indicating a lack of simple and consistent relationship between spike latencies and firing rate in tactile afferents (Johansson and Birznieks, 2004). Indeed, the complex non-linear viscoelastic deformation properties of the fingertip caused by its geometry and composite material properties makes the afferent encoding of mechanical fingertip events highly intricate (Birznieks et al., 2001; Jenmalm et al., 2003; Goodwin and Wheat, 2004). Researchers have attempted to model the fingertip mechanically while incorporating its composite material properties with the goal of predicting responses of populations of tactile afferents to arbitrary fingertip stimuli (Srinivasan and Dandekar, 1996; Maeno et al., 2004; Serina et al., 1998; Maeno et al., 1998; Nakazawa et al., 2000; Dandekar et al., 2003; Wu et al., 2006c). However, no model yet possesses the level of realism that satisfies this goal.

The mechanics of the fingertip presumably also accounted for the relationship between information conveyed about object curvature and the distance from the site of stimulation on the fingertip and afferents' termination in the skin. In agreement with observations in monkeys suggesting that SA-I afferents terminating at the sides of the fingertip are less affected by changes in curvature than afferents terminating close to the stimulation site (Bisley et al., 2000), we observed that afferents located close to the primary stimulation site transmitted higher information about object curvature than afferents located farther away in both firing rates and first spikes. Interestingly, we found the

opposite relationship for information about force direction conveyed by the first spikes in FA-I afferents, where afferents terminating on the sides and end of the fingertip tended to convey more information about force direction than afferents located close to the stimulation site.

The intricacy of the mechanical properties of the fingertip might also have accounted for the apparent difference between force direction and curvature concerning information contained in spike counts vs. first spikes. Besides the fact that the first spikes conveyed significant information about both stimulus parameters considerably faster than firing rates, the information about force direction in spike counts increased throughout the force protraction period, while the amount of information about surface curvature in spike counts saturated comparatively early during the force protraction period. Hence, later spikes appeared to contribute little information about curvature.

The SA-I and FA-I afferents showed overall similar properties regarding information content across the coding schemes. One likely explanation for this similarity is that the force stimulus during the protraction phase followed a time course that could excite both types of afferents with similar efficiency: its frequency content was 4 Hz. A major difference between the FA-I and SA-I systems is their sensitivity, with the FA-Is being most sensitive to relatively high frequencies (5–50 Hz) and the SA-Is to lower frequencies content (less than 5 Hz) (see e.g. Johansson et al., 1982).

3.4.4 The importance of spike timing in the early somatosensory system

Neural mechanisms that support fast reception and processing of sensory information are critical for control of actions. In object manipulation tasks, information about discrete mechanical fingertip events is typically expressed in the control of fingertip actions within 100 ms or less (for a review see Johansson and Flanagan, 2009). This requires that significant information is transmitted even when most afferents recruited have only had time to fire one spike; this, in turn, implies that the precise timing of the first impulses elicited in the afferent population convey information, which we conclude to be the case. It is improbable that the CNS can predict contact events with millisecond precision in manipulation tasks so as to calculate spike latencies with stimulus onset exactly. Therefore, rather than using absolute spike latencies, the relative timing of impulses in ensembles of afferents most probably convey the relevant information (Johansson and Birznieks, 2004). How the central nervous system can de-

code such information is not known, but various mechanisms have been proposed (see e.g. Johansson and Flanagan, 2009). Modelling studies have shown that spike-timing dependent plasticity (STDP) could be one crucial mechanism that allows neurons to learn temporal patterns (Masquelier et al., 2008).

Neural codes operating on latency differences between pairs of neurons might convey substantially more information than codes focusing on latencies in single neurons alone, if noise is correlated across neurons (Gollisch and Meister, 2008). We could not specifically address this possibility, as our experimental technique was limited to single-cell recordings. However, this assumption seems a very plausible one, as a substantial inter-trial variability in tactile afferent responses to the same stimulus is accounted for by non-stationary physiological conditions that globally affect the responsiveness of afferent populations. This includes volume changes of the fingertip with the pulse pressure wave and the respiratory cycle (Macefield, 2003) and, in particular, changes in viscoelastic properties of the fingertip related to previous tactile stimuli (Pubols, 1982a; Birznieks et al., 2001). Because of such correlations, the relative timing of spikes in pairs of tactile afferents is likely to fluctuate less than expected from the noise in individual afferent latencies, and thereby leading to even higher information rates. Thus, based on our analyses and these considerations, it seems suggestive that codes based on the timing of first spikes may provide an explanation for the rapid use of tactile information, as for example shown in object-manipulation tasks.

Chapter 4

Encoding and exploiting the effects of hysteresis

4.1 Introduction

In this chapter, we focus on how the mechanical properties of the fingertip affect neural responses when tactile stimuli follow each other rapidly. As described in Sec. 2.4, the human fingertip undergoes considerable deformation during grasping and object manipulation (see Fig. 4.1) and only slowly reverts back to its resting state. Most previous studies dealing with tactile afferent mechanisms have used protocols that deliberately reduce or standardize possible influences of earlier tactile stimuli on the responses recorded to the test stimuli. However, in real-world object manipulation and haptic tasks, forces leading to skin deformations generally follow each other in rapid succession (Kunesch et al., 1989) and any confounding effects of past stimulation on neural responses would limit the accuracy at which the tactile system operates (Pubols, 1982b; Beitel et al., 1977; Kenton et al., 1971). For example, stimulation of a fingertip with a given force perpendicular to the skin surface would result in a lower rate of deformation if the fingertip is not yet fully recovered from a previous deformation by a perpendicular force. These differences in the mechanical response of the fingertip tissues would lead to differences in the responses of the tactile afferents, and therefore might impair the accuracy in estimating relevant mechanical events at the fingertip. Hence, to correctly estimate impending forces based on the afferent signals, the CNS would also need to know the state of the fingertip at the commencement of

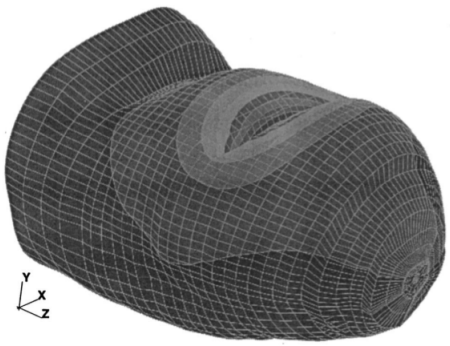

Figure 4.1: Simulation of widespread displacement of fingertip tissue in response to indentation by a wedge-shaped object. Lighter shades of grey indicate displacement. Figure taken from Dandekar et al. (2003).

the stimulus.

There are several potential mechanisms for mitigating the effects of skin hysteresis. First, memory mechanisms within the CNS might keep track of a recent history of skin stimulation, similar to the sensorimotor memory which keeps track of recent grip forces (Johansson and Westling, 1984; Quaney et al., 2003). Second, SA-II afferents that respond continuously even in the absence of skin stimulation might provide a constant signal of the mechanical state of the fingertip to the CNS (see Discussion in Birznieks et al., 2009). Last, the state might be encoded in the pattern of responses in ensembles of tactile afferents that respond when a new stimulus is applied to the skin. In this chapter, we focused on the last possibility.

To this end, we investigated the effects of past stimuli on the spike counts (firing rate) and the precise timing of the first spike of primary tactile afferents innervating the human fingertips when stimulated with forces in various directions applied in quick succession. We report that varying the past stimuli affects both spike timing and spike counts in FA-I and SA-I afferents, apparently increasing the noise in the afferent response as compared to when the influence of past stimuli is kept constant for each trial. However, we found that the precise timing of the first spike of a subset of afferents encodes the state of the skin at the onset of stimulation and thereby provides information related to past stimuli. These signals might be used to estimate the current state of the fingertip so as to uphold the accuracy of estimating important tactile events.

We used first spike latencies and spike counts as response measures in our analysis

of information about the tactile stimuli encoded in the tactile afferents, since these measures provide complementary early and late information on tactile events (see Johansson and Birznieks, 2004, and Sec. 3.4.3 in this thesis).

4.2 Methods

4.2.1 Data collection

The overall data set used in this chapter of the thesis is the same at the one used in Chapter 3. Because we selected the data based on different criteria, the number and identity of the afferents included in the analysis in this chapter are different from Chapter 3. The experimental setup has already been described in Secs. 3.2.1.1–3.2.1.3. In the following, we note some relevant differences.

In this chapter, we only use data from sequences in which a flat probe was used as stimulus. These probes were pushed onto the fingertip skin in one of five different force directions: normal (N) to the skin surface, or with an additional tangential force component in the radial (R), distal (D), ulnar (U) or proximal (P) direction (see Fig. 4.2A). As described in Sec. 3.2.1.2, each stimulation series consisted of 50 individual trials, i.e. 10 per force direction. The temporal ordering of the stimuli within a sequence differed: The first 25 stimuli made up the regular sequence, where stimuli followed each other in a fixed order, with each stimulus being presented 5 times. After the regular sequence, the irregular sequence was delivered to explore the effect of changes in the immediate stimulation history: the above five stimuli (R, D, U, P, N) were presented five times each in such a way that each stimulus was preceded once by each of the five stimuli (see Fig. 4.2B).

4.2.2 Skin position and fingertip hysteresis

In order to determine the effects of hysteresis and other viscoelastic properties on the transformation of the fingertip when stimulated by controlled forces, we measured the time-varying 3D position of the stimulus surface (see Birznieks et al., 2001) which reflected the position of the skin at the primary site of stimulation. In short, the stimulus position was controlled by three linear motors that each were connected to the stimulation probe by stiff rods. Linear optical position transducers were attached to each of

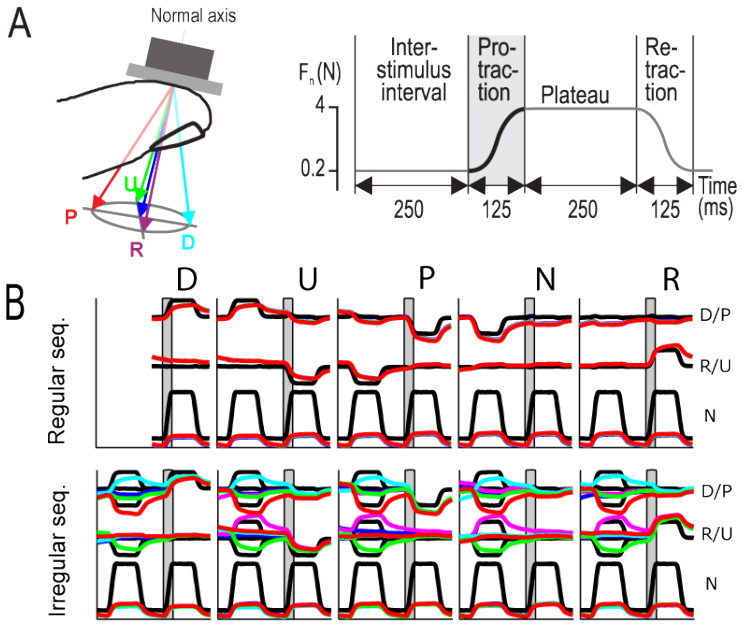

Figure 4.2: **A.** Experimental setup: Flat objects are pressed into the fingertip in the direction normal to the skin surface (N, dark blue) or with an additional lateral force component in the distal (D, light blue), radial (R, magenta), proximal (P, red), or ulnar (U, green) direction. A single trial consists of several phases: the inter-stimulus interval, where normal forces are kept at 0.2 N, lasting for 250 ms; protraction and retraction phases, where normal and tangential forces are ramped up or down, respectively, lasting for 125 ms each; and a plateau period where the normal force is kept constant at 4 N, lasting for 250 ms. **B.** Forces (black traces) acting on the fingertip and resultant skin positions (colored traces) along the proximal-distal (P/D), radial-ulnar (R/U), and normal (N) axis for five trials per stimulus direction, along with preceding trials. Trials are shown in full, with the protraction period shaded in grey. Stimuli are either presented in a fixed order (regular sequence, top panel), or varying systematically (irregular sequence, bottom panel). For the irregular sequence, trials are color-coded according to the stimulus direction of the preceding trial.

the motors, which allowed the varying 3D position to be measured with a resolution of $< 0.6\mu m$. Although these measurements do not reflect the complex changes in the distribution of stresses and strains within the fingertip, they still serve to estimate tissue hysteresis and stiffness (e.g. Pawluk and Howe, 1999; Pataky et al., 2005). Because a constant background force of 0.2 N orthogonal to the stimulus surface was maintained throughout the experiment and the surface material was chosen to have high friction,

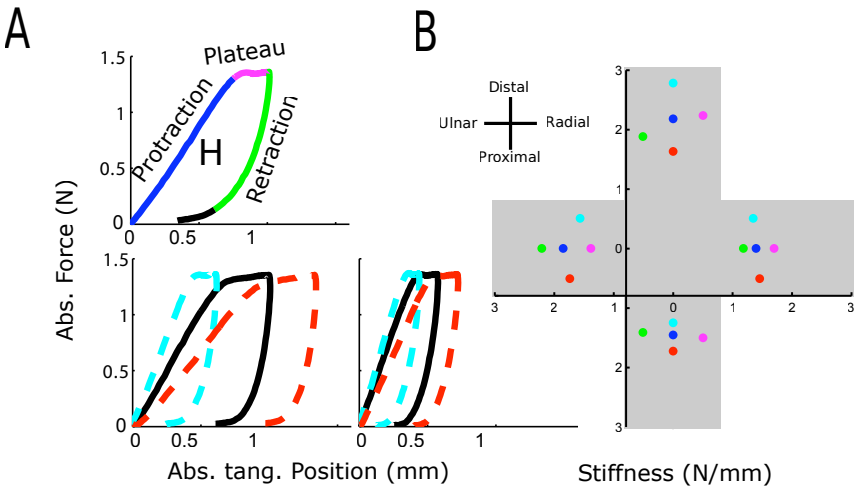

Figure 4.3: **A.** Top panel: Absolute position plotted against absolute force in the tangential plane for a single exemplary trial. The different stimulation phases (cf. Fig. 4.2A) are color-coded, with the protraction period in blue, the plateau period in magenta, the retraction period in green, and the inter-stimulus-interval in black. Absolute positions are calculated with respect to the tangential skin position at the beginning of the protraction period. The area enclosed by the curve (denoted by H) corresponds to the mechanical hysteresis. The slope of the curve during the protraction period corresponds to the stiffness of the fingertip during loading. Bottom panels: Position-force curves for exemplary trials from two different fingertips. Black curves denote a preceding stimulus without a tangential force component (N), while light blue and red curves stand for preceding stimuli that were either in the same or in the opposite direction of the current stimulus direction, respectively. **B.** Tangential stiffness values for four different force directions depending on previous stimulation direction. The color of each dot corresponds to the previous force direction (light blue: distal, magenta: radial, green: ulnar, red: proximal, dark blue: normal only; same as in Fig. 4.2). Stiffness values for the distal-proximal line are shown on the vertical axis, and values for the radial-ulnar line are shown on the horizontal axis. Displacement of some dots orthogonally to these axes is for clarification only.

little movement of the probe relative to the primary site of contact with the skin could happen.

The two mechanical variables that we calculated were stiffness and hysteresis. Stiff-

ness is defined as the amount of force that has to be expedited in order to effect a certain unit of change in the position of the fingertip skin. As we measured both force as well as position, stiffness simply corresponds to the slope of the force-position curve. Hysteresis refers to the amount of "lag" that the skin exhibits when reverting back to its resting state after indentation. When plotting force against position for a force ramp-up and a force ramp-down, these curves will be different and the area between them is referred to as hysteresis.

In order to measure the viscoelastic properties of the human fingertip in the tangential plane, we calculated stiffness and hysteresis from the time-varying force and position data. To this end, we used linear regression to fit a straight line to the force-position data from the protraction period (i.e. the blue curve in Fig. 4.3A, upper panel). The slope of this line then reflects the stiffness (in Nmm^{-1}) of the fingertip skin in the tangential direction during the protraction period. Generally, we found the linear approximation to hold very well, in agreement with other studies (Nakazawa et al., 2000; Pataky et al., 2005). The area enclosed between the protraction, plateau, and retraction curves corresponds to the hysteresis (indicated by H in Fig. 4.3A). We calculated stiffness and hysteresis values for different force directions and also analyzed whether past stimuli had an effect on these variables. For all measurements, absolute force and position are taken with respect to their values at the beginning of the protraction period; as trials followed each other quickly this might not be the fingertip's resting position, but instead depend on the previous force direction. In total, data from 70 unique fingertips (23 index, 31 middle, and 16 ring fingers) was included in this analysis.

4.2.3 Information theoretic analysis and classification

In order to determine how sensitive an afferent's response was to the current, previous, or future force direction, we calculated a lower bound on the information transmitted by a given afferent. We used the metric space analysis (Victor and Purpura, 1996, 1997), which defines a classifier on the neural response data, calculates a confusion matrix from several runs of the classifier, and finally computes a lower bound on the mutual information from the confusion matrix (see full explanation in Sec. 3.2.2.3). We chose the metric-space method, because our separate treatment of the regular and irregular series in this chapter resulted in a low number of overall trials and the simulations described in Sec. 3.2.2.2 suggested that the bias would be unacceptably high using the direct method as described in Secs. 3.2.2.2 and 3.2.2.4. In our treatment of

spike time latencies, the temporal resolution was fixed at 4 ms, as this was shown to maximize information transmission for tactile afferents (see Sec. 3.3.1). However, the exact setting of this parameter did not substantially affect our findings. We assessed two coding schemes, spike counts and first spike latencies. For spike counts, we used the full neural response during the protraction period wit the "cost" parameter q set to zero (cf. Sec. 3.2.2.3), while for first spike latencies, we calculated spike distances from the very first spike of each afferent only, with the cost parameter set according to the optimal temporal resolution as described above.

To estimate the bias of the estimate of the mutual information, we reassigned neural responses from individual trials to random conditions and recalculated the mutual information. The bias term was set as the average of the outcome of 10 such assignments and subtracted from the estimate of the mutual information.

In a further analysis, we asked whether we could predict current, previous, or future force directions from trials taken from the irregular series when the classifier was trained with data from the regular series. For this, we randomly selected individual trials from the irregular series for a given force direction. For each afferent from this set we then computed the average distance between this trial and all trials from all other conditions from the regular data set. The chosen sample was then classified as originating from the force direction with the smallest average distance. When trying to classify previous or future force directions, trials from the regular series were grouped according to the previous or future force direction respectively.

4.3 Results

4.3.1 Viscoelastic properties of the fingertip

In agreement with earlier reports (Birznieks et al., 2001; Jenmalm et al., 2003; Nakazawa et al., 2000; Pataky et al., 2005; Wang and Hayward, 2007), we found marked anisotropic properties of fingertip mechanics, when calculating stiffness and hysteresis for different tangential force directions (see Sec. 4.2.2 and Fig. 4.3A, upper panel). Stiffness in the tangential directions during the protraction period was on average 1.4 Nmm^{-1} ($\pm$ 0.4 Nmm^{-1} standard deviation) in the radial, 1.8 ($\pm$ 0.5) Nmm^{-1} in the ulnar, 2.2 ($\pm$ 0.6) Nmm^{-1} in the distal, and 1.4 ($\pm$ 0.5) Nmm^{-1} in the proximal direction. A one-way analysis of variance (ANOVA) found highly significant differences between

the directions ($p \ll 0.01$), and post-hoc paired comparisons showed that while radial and proximal stiffness did not differ significantly, all other pairs did ($p < 0.05$, Tukey correction). Similarly, hysteresis differences between force directions were highly significant ($p \ll 0.01$) for the same pairs as for stiffness. Average hysteresis was 66% in the radial, 69% in the ulnar, 74% in the distal, and 66% in the proximal direction. There was considerable variation in stiffness and hysteresis between subjects, as well as between different fingers for individual subjects.

Fig. 4.2B shows the forces (black lines) and resultant skin positions (colored lines) for trials in each of the five force direction along with their respective preceding stimuli for both the regular (top panel) and irregular sequence (bottom). Shaded areas correspond to the protraction period. Fig. 4.4A shows skin positions (top row) and forces (bottom row) in the tangential plane plotted over time for both the regular and irregular series, with the force protraction periods color coded. The position signal revealed that past force directions affected the transformation of the fingertip in response to the current force stimulus. The trajectory of the probe position was very repeatable over several trials for the regular sequence (Fig. 4.2B, upper panel, and Fig. 4.4A, left) in which the order of stimuli was fixed. In contrast, there was much variation in probe trajectory during the irregular sequence (Fig. 4.2B, lower panel, Fig. 4.4A, right) in which the order of stimuli was systematically varied (see Sec. 4.2.2), even though the forces exerted onto the fingertip during individual trials did not differ from those during the regular sequence.

Fig. 4.4B shows, for the irregular series, absolute probe positions, and forces in the plane tangential to the contact surface at the beginning of the protraction phase. Different colors now refer to different preceding stimuli (i.e. the stimulus that was presented directly before the current one), while different marker shapes indicate the currently applied force directions. The grey traces indicate the paths during the first 85 ms of the protraction period, whereas fully colored markers indicate the probe positions at the start of the protraction period. As can be seen, initial probe positions appear clustered together based on the previous stimulation direction, as the fingertip has not yet fully recovered back to its resting position from the previous stimulation.

We found that past stimuli also affected stiffness and hysteresis values for the subsequent trial. Generally, stiffness increased if the previous trial had had the same tangential force direction, and decreased if the previous trial had been in the opposite direction (see Fig. 4.3A, bottom panels, for examples). These changes not only

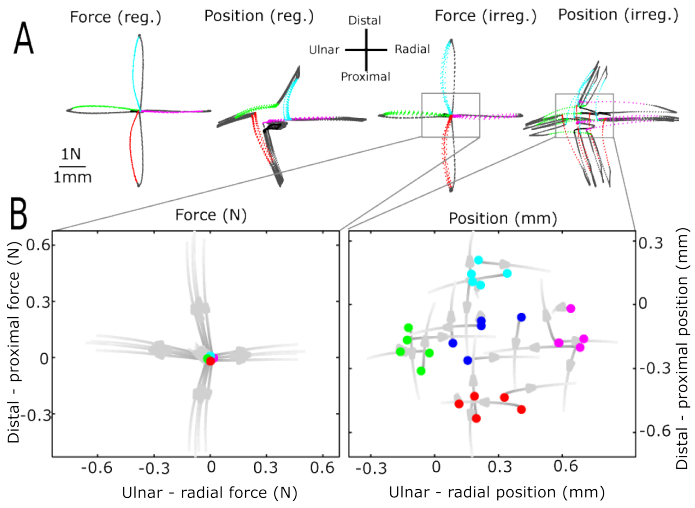

Figure 4.4: **A.** Forces and skin position traces in the tangential plane on the fingertip for five repetitions of each of the five stimuli shown for the regular series (left panels) and the irregular series (right panels). A dot is plotted every 5 ms. Coloured dots indicate the protraction period for stimuli in the five force directions. **B.** Forces (left panel) and skin positions (right panel) for the irregular series at the start of the protraction period, color-coded according to the preceding stimulus. Grey traces denote the forces and positions during the first 85 ms after stimulus onset, respectively.

affected the very beginning of the trial, but are still evident later on in the trial, for example during the plateau period. A graphical depiction of stiffness values during the protraction period for various previous stimuli is shown in Fig. 4.3B. It can be seen that earlier stimulation in a direction opposite to the subsequent force direction influences stiffness in that direction considerably. Finally, hysteresis was also affected by stimulation history: It increased between 3 and 6 percentage points when the earlier stimulation was in the opposite direction and decreased by around 5 to 8 percentage points when the preceding trial had used the same force direction.

In summary, past stimuli have both immediate and long-lasting effects on fingertip mechanics, resulting in marked effects on skin position trajectories during delivery of a subsequent stimulus. In the next section, we analyze whether the neural response was affected.

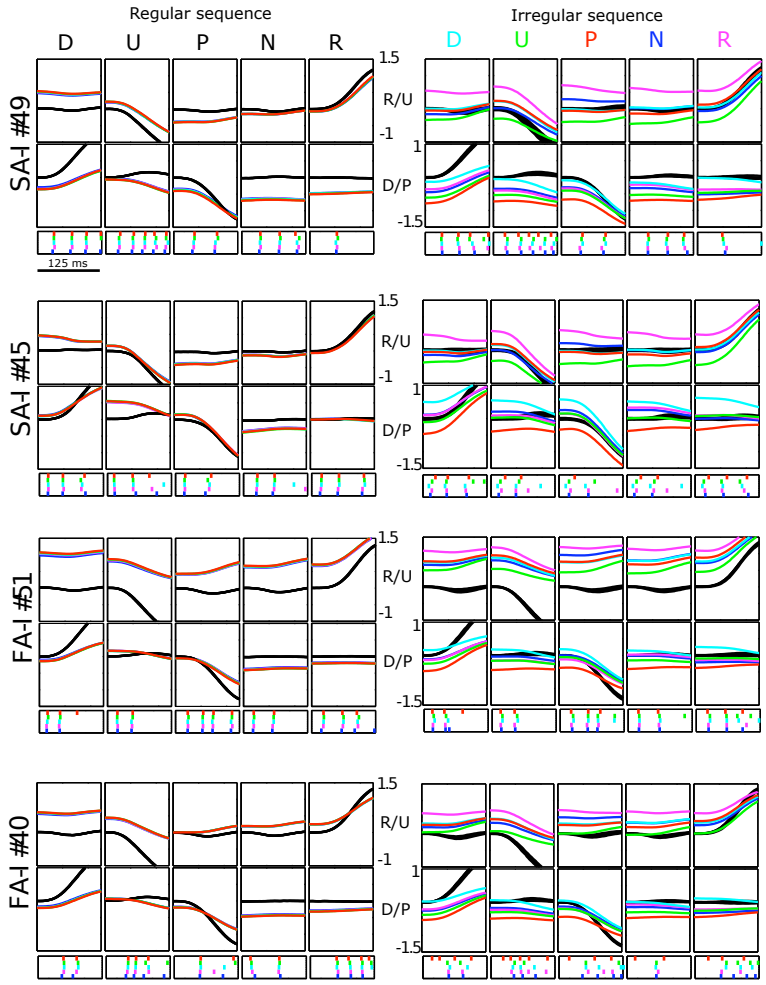

Figure 4.5: Forces (black lines) and skin positions (colored lines) along tangential axes (upper panels: radial-ulnar, middle panels: proximal-distal) as well as corresponding neural responses (colored ticks, bottom panels) during the protraction period for selected afferents, using data from the regular sequence (left column) or the irregular sequence (right column). Five repetitions per stimulus direction are shown. In the irregular sequence, skin positions and neural responses are colored-coded according to the stimulus direction of the trial directly preceding the displayed trial.

4.3.2 Effect of stimulation history on neural responses

We analyzed the effects of stimulation history on the tactile afferent responses during the force protraction phase only, as the initial neural response has been shown to contain rich information about stimulus direction (see results in Chapter 3 and Birznieks

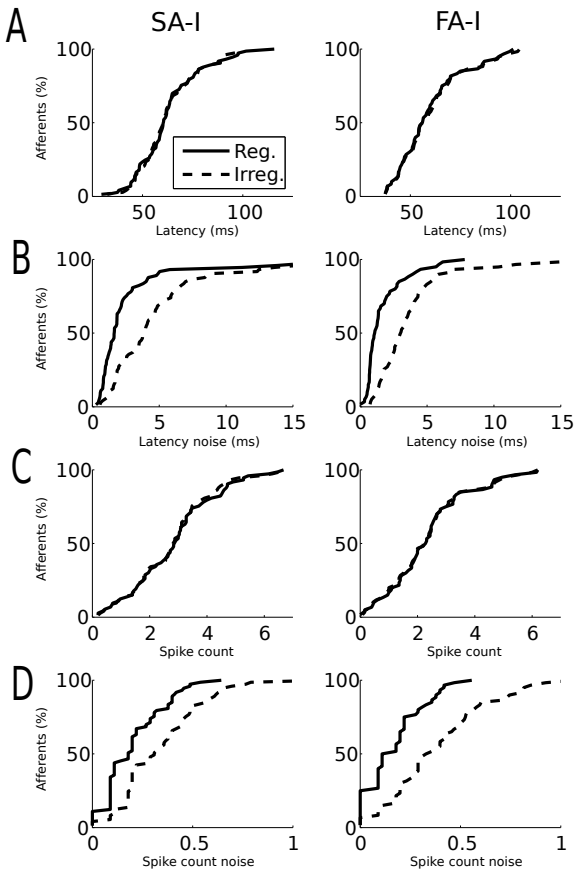

Figure 4.6: Cumulative distributions of neural sensitivity and noise for trials from the regular (continuous lines) and irregular (dashed lines) sequences. **A.** First spike latencies (averaged over all trials and stimulus directions). **B.** First spike latency noise, corresponding to the standard deviation of the first spike latencies (averaged over all trials and stimulus directions). **C.** Spike counts (averaged). **D.** Spike count noise (averaged).

et al., 2001; Johansson and Birznieks, 2004) as well as other important object properties (Jenmalm et al., 2003; Birznieks et al., 2010) and plays an important role during object manipulation (see Johansson and Flanagan, 2009, for a review).

Fig. 4.5 shows skin positions along the radial-ulnar and distal-proximal axis as well as neural responses to all five force-directions (N, R, U, D, P) during the protraction period of some SA-I and FA-I afferents for the regular sequence (left-hand columns) and the irregular series (right-hand columns). Some afferents' responses were just modestly affected by stimulation history (afferents SA-I #49 and FA-I #51), while

for others the stimulus history appeared to more dramatically affect the neural responses (afferents SA-I #45 and FA-I #40). The skin positions along the radial-ulnar and proximal-distal axes show considerable effects of past stimulation throughout the whole protraction period, as is evident from the data from the irregular sequence.

In the following, we quantify the effects of stimulus history on the timing of the first spikes and spike counts, to get an idea of how these response measures are affected. The first spike latency averaged over all afferents and force directions was identical in the regular and irregular sequences for FA-I afferents (60 ms after stimulus onset) and SA-I (63 ms). See Fig. 4.6A for cumulative distributions of first spike latencies. Expressed as the median of the standard deviations of the first spike latencies computed for afferent and force direction, the noise in spike timing of the FA-I afferents was 3.1 ms in the irregular series as compared with 1.1 ms during the regular series ($p \ll 0.01$; Wilcoxon rank-sum tests). The corresponding numbers for the SA-I afferents were 3.9 ms and 1.7 ms, respectively ($p \ll 0.01$, see Fig 4B).

The mean number of spikes elicited during the force protraction phase in the FA-I afferents (2.4) and SA-I afferents (2.9) did not differ significantly between the irregular and regular sequence (Fig. 4.6C). However, the median noise in the spike count increased from 0.14 to 0.34 for FA-I and from 0.18 to 0.31 for SA-I afferents (Fig. 4.6D; $p \ll 0.01$ in both cases).

We conclude that past stimulation in the tangential plane can affect both spike timing and firing rates of afferents. While the responsiveness of the afferents, i.e. the mean first spike latencies and spike counts, did not change when averaged over the whole population, variation of the past stimulus increased the variance in responses to a given stimulation for both spike timing and spike counts.

4.3.3 Responsiveness of individual afferents to previous and current stimulus

Given that responses of tactile afferents were affected by the stimulation history, we asked how the force direction of the previous stimulation influenced the amount of information conveyed about the current force direction and whether the afferent responses might contain information about previous force directions. In short, we calculated a lower bound on the mutual information transmitted about the stimulus by assessing the performance of a classifier that assigned neural responses to force direc-

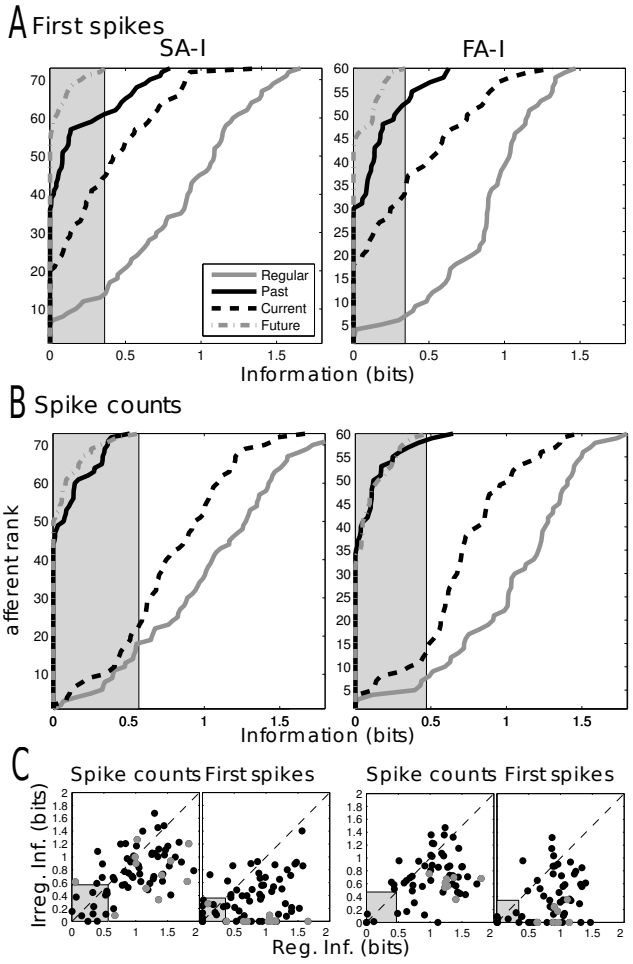

Figure 4.7: **A and B.** Mutual information (lower bound) contained in first spike latencies (A) or spike counts (B) of SA-I (left column) and FA-I (right column) afferents for current stimulus (dashed black line), previous stimulus (continuous black line) and future stimulus (dash-dotted grey line) from the irregular sequence and the current stimulus from the regular series (continuous grey line). To avoid confusion, each individual line is plotted in order of ascending information, i.e ranked, with the order determined for each condition separately. Thus, horizontal sections through this plot do not correspond to values from the same afferent. Grey boxes denote areas that are equal or smaller in information to the maximal information obtained for the future condition. Afferents passing the threshold given by this box are classified as belonging to the *C* (current stimulus) and *P* (previous stimulus) sets, respectively. **C.** Mutual information calculated from the regular sequence plotted against information in the irregular sequence for spike counts (left panels in each column) and first spike latencies (right panels). Afferents belonging to the *P* set are shown in grey, all other afferents in black.

tions (see Sec. 4.2.3). We ran this classifier to find out how well we could determine current, previous, and (as a control[1]) future force directions from a given afferent's neural responses during the protraction period. Values were computed for two different coding schemes: Spike counts and first spike latencies only.

Trials from the irregular sequence were grouped in three different ways: According to the current force direction (current grouping), the preceding force direction (past grouping), or the upcoming force direction (future grouping). For example, when using the past grouping, all trials that share the same preceding stimulus were grouped together, irrespective of the current force direction. By examining information values calculated from different groupings of the irregular sequence, we determined whether individual afferents carry information about ongoing stimulation, previous stimuli, or (as a control) about future stimulation.

As a comparison, we also examined the information contained about the current force direction in neural responses obtained in the regular sequence. In this data set past, current, and future groupings are identical as all trials for a given current force direction shared the same preceding and subsequent force direction.

The graphs in Fig. 4.7A and 4.7B show mutual information present in first spike timing and spike counts per grouping for all afferents of each type. The maximal information about the future stimulus (indicated by a light grey box in Fig. 4.7) acts as a threshold: any values falling below are deemed resulting from noise. We used this rather conservative estimate to ensure the exclusion of spurious results. As expected from the higher noise, the information about the force direction in the current stimulus was overall lower for the irregular series (black dashed curves in Figs. 4.7A and 4.7B) compared to the regular series (solid gray curves, see also scatter plots in Fig. 4.7C).

However, the afferents provided not only information about the current stimulus, but some afferents also provided information about the force direction in the previous stimulus when considering first spike timing (solid black curves). For data from the irregular series, we examined which afferents exhibited information above threshold about the *current* force direction and will refer to this group of afferents as the *C* set. Similarly, those afferents transmitting information above threshold about the *previous* force direction make up the *P* set. Using this procedure on first spike latencies, we found that 29 (i.e. 40% of all) SA-I afferents belonged to the *C* set and 12 (i.e. 16%

[1] As the sequence was not random, some amount of information about the future stimulus is available. This makes the control condition slightly more conservative than necessary.

of those afferents) to the *P* set. Notably, there was no overlap between these two sets. That is, no afferents transmitted high information about both the current and the previous stimulus (see also Fig. 4.7C, left panels in each column). Similar results were obtained for the FA-I afferents, where 26 (i.e. 43% of all FA-I afferents) and 8 (i.e. 13% of the total) belonged to the *C* and *P* sets, respectively. Only one FA-I afferent belonged to both sets. Considering spike counts, during the irregular series most afferents were informative with respect to the current force direction, with 51 (i.e. 70% of all) SA-I and 46 (i.e. 77% of all) FA-I afferents falling into this group. No SA-I afferent and only two FA-I afferents seemed to provide information about the previous direction, with one of them also responding above threshold to the current direction.

There was a highly significant reduction in the information in spike counts between the regular and irregular sequence for FA-I afferents belonging to the *P* set (t test, $p < 0.01$), but not so for equivalent SA-I afferents ($p = 0.06$, see also Fig. 4.7C, right panels in each column). Despite this reduction, for the afferents belonging to the *P* set, 8 (out of 12) SA-I afferents and 7 (out of 8) FA-I afferents also transmitted information above threshold about the direction of the current stimulus in their spike counts. Thus, some afferents signaled the past force direction in their first spike latencies, and information about ongoing stimulation in their spike counts.

We also examined whether afferents belonging to the *P* set would tend to respond earlier than other afferents. We found significant differences in response onset between *P* and *C* set afferents for both SA-I afferents (*P* set mean: 58 ms, *C* set mean: 69 ms, $p = 0.01$) and FA-I afferents (53 ms vs. 67 ms, $p = 0.05$). We did not find any differences between where the afferents terminated on the fingertip and whether they belonged to the *P* or *C* set.

Hysteresis and other viscoelastic properties of the fingertips vary between subjects (see Sec. 4.3.1 and Pataky et al., 2005). As we used single afferent recordings from different subjects, it could be that the afferents in the *P* and *C* sets simply reflected the different mechanical properties exhibited by the subjects. For example, some of the subjects' fingertips might show especially high hysteresis and therefore be more likely to cause afferents to reflect past force directions, while others might show low hysteresis and therefore be more likely to respond to the current stimulus. If this were the case, it would be wrong to make any projections about how a population of afferents within a single fingertip would respond to a sequence of stimuli. In order to test this, we

calculated the stiffness and hysteresis associated with each afferent's fingertip from the recorded position signals (see Sec. 4.2.2) and then examined whether the stiffness or hysteresis levels differed between the *P* and *C* sets. Wilcoxon rank-sum tests did not show a significant difference between the medians of both distributions for either SA-I or FA-I afferents. Additional evidence comes from the fact that in some instances (in 10 out of the 70 distinct fingertips), there are both *P* and *C* set afferents on the same fingertip[2].

In summary, we found that a large number of afferents convey information about the current stimulus in both spike counts and first spike latencies. Additionally, a subset of afferents signal information related to the previous stimulus in their first spike latencies.

4.3.4 Comparison of latency tuning curves for regular and irregular series

To analyze further the view that some afferents can signal the past stimulus in their first spike latencies, for each afferent we inspected latency tuning curves showing the average response to each of the different force directions along with the standard deviation computed across all stimulus repetitions as an estimate of the expected amount of noise. We compared tuning curves constructed from the irregular data, grouped according to either the previous, current, or (again as a control) future force direction, with tuning curves generated from the regular series. If afferents in the *C* and *P* sets respond predominately to the current and previous force direction, respectively, there would be a rather close match between the tuning curves, along with rather low overall noise values.

Fig. 4.8 shows latency tuning curves computed from the regular (black curves) and irregular series (gray curves) for different subsets of afferents and stimulus groupings. The tuning curves for the current stimulus for the *C* set of afferents matched the corresponding curves computed from the regular series very well (Fig. 4.8A). Likewise, tuning curves for the *P* set match those referring to the previous stimulus obtained from the regular series (Fig. 4.8B). In contrast, the fit is much worse for afferents from the *P* set when compared to tuning curves for both the current and the future force direction (Figs. 4.8C and 4.8D). In these cases, the tuning curves from the irregular

[2]For the majority of fingertips, data from only one or two afferents are available.

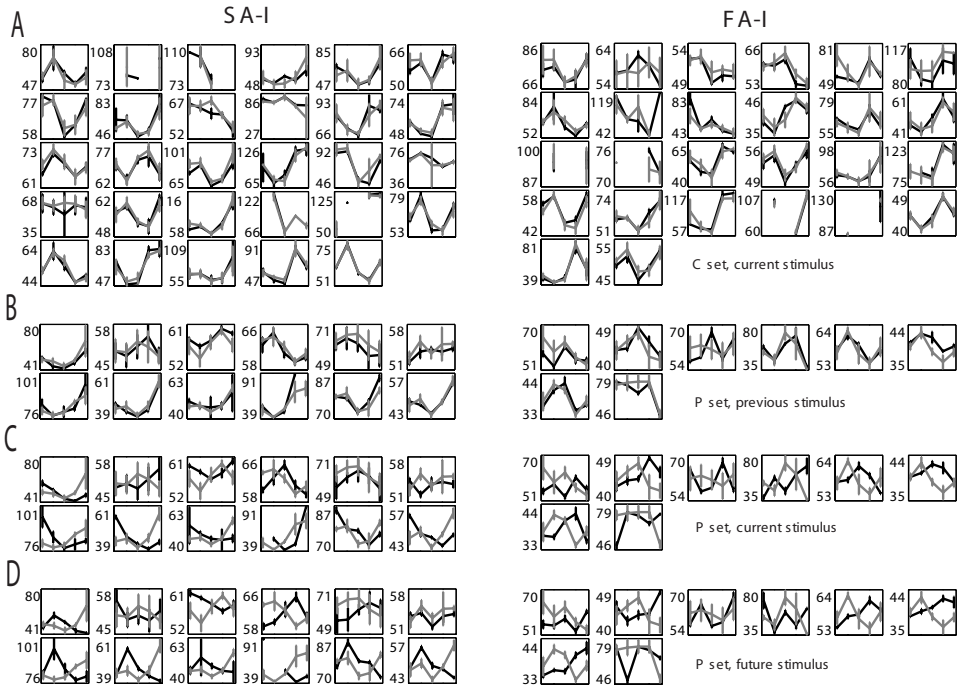

Figure 4.8: Latency tuning curves for different afferent groups and with different stimulus groupings calculated from the irregular (black line) and regular (grey line) sequences. The ordinate shows average first spike latencies after stimulus onset (i.e. start of for protraction) along with error bars denoting one standard deviation (over five trials) for the different force directions (N, D, U, P, R) given by the abscissa. **A.** Tuning curves with respect to the current force direction for afferents in the *C* set. **B.** Previous force direction for afferents in the *P* set. **C.** Current force direction for afferents in the *P* set. **D.** Future force direction for afferents in the *P* set.

data set show much higher variance and generally do not match the tuning curves from the regular data set very well.

These results are also consistent with the noise in the first spike latency distributions: The median noise within the first spike latencies of the *C* set when grouped with respect to the current stimulus it was 3.0 ms and 2.9 ms for SA-I and FA-I afferents, respectively. The corresponding values for the *P* set grouped according to the previous stimulus was 1.7 ms and 1.9 ms. However, when the *P* set was grouped according to the current stimulus, the noise increased to 4.0 ms and 2.2 ms, respectively.

Finally, we looked at how well we could classify the force direction of single trial responses from the irregular data set for groups of afferents, with a classifier trained only

on data from the regular series (see Sec. 4.2.3). We found that we could classify the current force direction with very high accuracy when using all the afferents from the *C* set: Correct classification was possible 98% of the time for the SA-I afferents and 92% of the time for the FA-I afferents. The accuracy when classifying the previous stimulus from afferents in the *P* set reached 71% and 47% for SA-I and FA-I afferents respectively, which is well above the chance level of 20%. It should be noted that classification accuracy also depends on the size of the neural ensemble, which was much smaller for the *P* set. Finally, when trying to classify the current or future stimulus from afferents in the *P* set, accuracy rates were as low as around 1%. The fact that the training data for the classifier comes from the regular series, which makes it impossible to distinguish whether a response is caused by the current or previous stimulus, explains why performance is considerably lower than chance in these cases. That is, the low accuracy is merely an artifact of the data set used.

4.4 Discussion

During object manipulation and haptic tasks, forces acting upon the fingertip follow each other in rapid succession. The tactile afferent responses, driven by the resultant skin deformations, have to be processed quickly and reliably. However, the viscoelastic properties of the skin and underlying tissues render the skin deformations dependent on earlier deformations, potentially disrupting the accurate estimation of impending forces. Therefore, the tactile system must have developed a way to cope with effects due to hysteresis and other viscoelastic properties of the skin in order to estimate the magnitude and direction of indenting forces correctly. However, in most physiological studies concerning tactile mechanisms, the non-stationary properties of the fingertip mechanics are rarely considered to play a role in determining how afferents respond to a mechanical stimulus; the unwritten assumption being that the skin provides high fidelity transmission of the stimulus to the underlying receptors. The purpose of the present study was to examine this assumption experimentally.

We found that stiffness and hysteresis of the fingertip in the tangential plane depended on the stimulation history. Moreover, stimulation history affected both spike counts and first spike latencies, resulting in increased neural noise[3]. While most afferents

[3]Of course such "noise" from the view of our analysis might still contain useful information about parameters not included in this analysis.

continued to signal the ongoing stimulus (collectively referred to as the *C* set), some afferents (the *P* set) transmit high information about past stimulation in their first spike latencies. An analysis of latency tuning curves confirms that those afferents' responses were influenced mainly by past stimulation. This might be explained by afferents responding to the state of the skin at the start of stimulation, before the response is influenced significantly by the tangential components. Such a mechanism might provide an "offset signal" at initial touch that could be used to mitigate confounding effects due to hysteresis on the neural responses.

4.4.1 Viscoelastic properties of the fingertip

We found marked anisotropic effects, in that the stiffness and hysteresis of the fingertip in the tangential direction varied with direction. In general, the stiffness values we report are somewhat higher than the values reported before (Pataky et al., 2005), which is most likely due to the fact that in our experiments the fingertip was fixated by gluing the nail to a cast, resulting in less rotational skin movement around the distal-proximal axis than when retaining the distal interphalangeal joint.

Interestingly, we found marked effects of the previous force direction on the tangential fingertip stiffness and hysteresis. Tangential force and position during loading (i.e. the protraction period) were linearly related, as reported before (Pataky et al., 2005), however the slope of this line, which represents the tangential stiffness, changed depending on the stimulus direction of the preceding trial. These effects might be caused by ongoing structural changes within the fingertip, while it is restoring to its resting state.

Effects of the preceding stimulus direction were visible throughout the protraction period and even continued thereafter: The fingertip stiffness during the plateau period (where mechanical creep was observed) was still affected by the previous force direction. These long-lasting effects are probably the reason why both the timing of the first spike as well as spike counts (i.e. firing rates) were affected by stimulation history: Both coding schemes exhibited higher noise in the irregular than in the regular stimulus sequence.

4.4.2 Afferents signaling past and current stimulation

Due to the ongoing effects of past stimuli throughout the delivery of subsequent stimuli, skin positions as well as strains and stresses within the fingertip would appear more noisy, which in turn would affect the neural response. Indeed, our information theoretic analysis showed that most afferents continued to transmit information about the ongoing stimulus, even though the amount of information was lower due to the increased noise. However, we also found that some afferents signaled the past stimulus direction in their first spike latencies, without conveying any information about the ongoing stimulation.

It is usually assumed that tactile afferents' firing rates are mainly driven by changes in tissue strain (Pubols and Pubols, 1976), although effects on spike timing are largely unknown. It is conceivable that different properties of the neural responses, e.g. first spike latencies and spike counts, might respond to different mechanical factors. In previous work it was found that correlation between latency and spike count tuning curves was low (Johansson and Birznieks, 2004) and that the information transmitted by these coding schemes appeared to be mainly independent (see Sec. 3.3.4). Furthermore, first spike latencies have greater chances to be affected by the previous stimulus (or the initial skin state at the current stimulus) than firing rates that are averaged over the entire protraction phase. Indeed we found that the response onset for afferents in the *P* set was 10–15 ms earlier than for afferents in the *C* set. As we showed, at the beginning of a new stimulus, skin positions were clustered according to the previous stimulus, as the fingertip had not fully reverted back to its resting state. It is possible that these differences in absolute position are picked up by afferents in the *P* set. It is important to note however, that the initial state by itself would not have caused the tactile afferents to respond: Without new stimulation of the skin, these afferents would not have fired. More detailed mechanical models of the fingertip than the ones currently available might be able to shed some light on the question of what mechanical factors could be involved in these responses. Our conservative estimate was that around 13% to 16% of SA-I and FA-I afferents belong to the *P* set. Projecting from this number to the whole fingertip which is innervated by around 2,000 afferents (Johansson and Vallbo, 1979), with the majority of them being either FA-I or SA-I afferents, results in a considerable number of afferents involved in signaling past stimulation.

4.4.3 Disambiguating neural responses under the presence of hysteresis

The question remains how the signals in the first spike latencies of the *P* set could be exploited in order to mitigate the effects of hysteresis. We argue that those afferents transmit an "offset signal" indicating the state of the skin at the onset of a new stimulus. It seems plausible that the different tuning properties in the *P* and *C* sets are entirely dependent on strain changes in the skin and other external factors, and are not intrinsic to those neurons. Therefore, membership of afferents in the *C* and *P* sets are subject to change for different stimulus locations and possibly other factors. How then, could the brain make sense of the incoming information? A first clue would be that afferents in the *P* set on average respond slightly earlier than other afferents, which might be exploited by higher-level read-out mechanisms. Furthermore, tuning properties of tactile afferents to tangential forces are complicated already, and the brain is able to make sense of them. For example, both latency and firing rate tuning curves can differ substantially even for afferents terminating close to each other (Birznieks et al., 2001; Wheat et al., 2010; Johansson and Birznieks, 2004).

There might be other mechanisms at work to mitigate the effects of hysteresis and it is possible that some of these work in parallel. For example short-term memory mechanisms that keep track of recent force impacts could help with predictively scaling grip forces, without requiring feedback first. Another mechanism that has been proposed is that SA-II afferents located close to the nail could send a continuous feedback signal about the viscoelastic state of the fingertip (Birznieks et al., 2009). Such a signal would be useful for continuously monitoring the state of the fingertip. The mechanism proposed in this paper would be complementary: The very first volley of neural responses after grasping an object would contain information about the state of the fingertip at contact, so as to be able to correctly estimate impending forces.

Finally, we found that spike counts were less affected by hysteresis effects than first spike latencies. Thus, it is also possible that any ambiguous force estimates due to hysteresis effects could be corrected using the information in spike counts. In such a scenario, the ambuity wouldn't be resolved immediately (contrary to the options above), as several spikes per afferent would be needed.

4.4.4 Effects on task performance and perception

Not much is known about whether and how effects associated with the mechanical properties of the fingertip such as hysteresis affect perception or motor control mechanisms (see also Sec. 2.4.2). Earlier studies on perception of tangential forces did not systematically look at whether the viscoelastic properties of the fingertip affect perception (Paré et al., 2002; Wheat et al., 2004). One study found that the detection of ridges is impaired when a smaller ridge follows a bigger one in quick succession (Singh and Fearing, 1998). However, accurate perception might just be impossible in such a setup, even if additional information about the state of the fingertip was available. In our study, stimuli are chosen so as to mimic the forces occurring during object manipulations. In object handling and manipulation tasks, accurate determination of forces and their direction is essential for establishing and reacting to contact events and it has been shown that cutaneous afferents play an important part in providing such feedback (see Witney et al., 2004, for a review).

4.4.5 Tactile feedback under natural conditions

As we have seen in the past two chapters, it is important to create experimental conditions that are close to the ones that would be naturally encountered, as some important features or constraints on tactile processing might only become evident in such realistic settings. For example, we have argued that fingertip mechanics can affect the neural responses of primary tactile afferents in important ways: mechanical properties of the fingertip tissues are ultimately responsible for the spatial distribution of information about tactile stimuli (see Chapter 3), and also need to be taken into account when different stimuli follow each other quickly, as happens during natural object exploration (as we saw in this chapter).

Another interesting aspect of this problem is the coordination of finger movements during grasping and object exploration. The tactile sense is necessarily an *active* sense, in that finger (or other limb) movements are required in order to obtain tactile feedback. One can then pose the question of how these movements should be controlled in order to obtain relevant tactile feedback for the task at hand. One example of this would be a tactile object classification task, where the aim would be to direct movements so as to collect tactile feedback data that is most suitable for discrimination between different object classes. We will describe computational approaches to dealing with

such scenarios in a robotics setup over the next three chapters (with a review in Sec. 5.4 and our own results in Chapter 6), and will discuss the problem from a neuroscientific point of view in the general conclusion in Chapter 8.

Chapter 5

Haptic discrimination in robotics—from sensors to actions

5.1 Introduction

In previous chapters we examined what kind of tactile features are represented in a population of human peripheral tactile neurons and how these are encoded. We used information theory in order to arrive at a quantitative analysis of the information transmitted by these afferents. In this chapter, we turn our attention to tactile processing in robotics. We will discuss how to represent sequentially arriving data efficiently and how to maximize incoming information by controlling a robot's actions during the collection of this data. We will begin with a short survey of tactile sensors along with robotic applications

5.2 Tactile sensing in robotics

5.2.1 Sensors and hardware

For much of its existence, robotics research has focused on industrial machines operating in static environments. More recently though, interest has shifted towards robots that are able to execute tasks in unstructured environments, that can work safely alongside humans, and that can manipulate objects originally designed for human beings. These new requirements led to the design of robotic hands that try to mimic the human

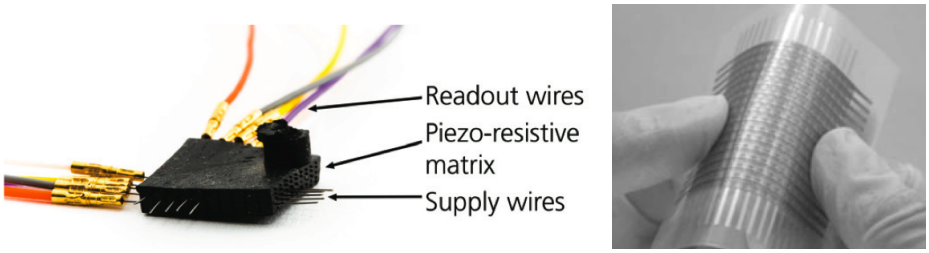

(A) (B)

Figure 5.1: Examples of flexible tactile sensors. **A.** Prototype of a viscoelastic, flexible, piezo-resistive sensor matrix. Figure taken from Strohmayr et al. (2010). **B.** Flexible tactile matrix based on organic transistors. Figure adapted from Someya et al. (2004).

hand in being multi-fingered and comparable in size. The key motivation was that such hands would be best suited for object grasping and dexterous manipulation in environments designed for humans (see e.g. Borst et al., 2003, for a modern approach). Moreover, soft fingertips exhibiting similar mechanical properties as human fingertips have been increasingly used. This is because compliant surfaces offer superior properties in grasping stability and slip prevention (Lee and Nicholls, 1999). The move to unstructured environments also emphasizes the need for sophisticated tactile sensing abilities in order to cope with the increased task complexity. Tactile sensing facilitates task execution by providing feedback directly from the site of contact between robot and object, information that cannot be gathered solely from camera images or force/torque sensors embedded in joints. Such features include the state of contact (e.g. detecting when making and breaking contact), slip detection, or texture recognition, all of which are important for grasping and object manipulation. Other sensor technologies like cameras or laser scanners might suffer from occlusion, reflections, bad lighting or a fixed viewpoint, and therefore might only be able to supply partial or degraded views of the environment.

There is a multitude of different touch sensor technologies (see Lee and Nicholls, 1999; Tegin and Wikander, 2005; Dahiya et al., 2010, for reviews). The most common techniques involve capacitive and resistive sensors, i.e. material that changes their electrical properties when physical force is exerted upon them. All methods come with their individual benefits and disadvantages, differing in robustness, spatial resolution and sensitivity among many other factors, which explains why no single technology has become standard yet.

Some combined designs have been developed that integrate tactile sensors directly into robotic hands or fingers (e.g. Edin et al., 2008; Wettels et al., 2008). These aproaches allow for tighter integration, but result in specialized hardware that is harder to combine with existing systems. Full-fledged robotic hands with integrated touch sensors have also been developed (e.g. Hosoda et al., 2006).

One noteworthy recent development is the aim to extend the area usually covered by tactile sensors—fingertips and maybe a small number of other regions on the hand—and develop big, soft, and flexible sensors. These would then make up an "artificial skin", covering robotic limbs and joints, and provide tactile feedback from all over the robot's surface (see e.g. Papakostas et al., 2002; Someya et al., 2004; Strohmayr et al., 2010, for approaches to flexible sensors). These and other developments, e.g. higher spatial resolution or sensing of 3D forces rather than just those orthogonal to the sensor surface, massively increase the amount of data and therefore call for sophisticated processing techniques.

5.2.2 Robotic applications involving tactile sensing

Data from touch sensors is often high-dimensional due to the large number of small sensing elements that are typically used. It is also highly redundant, as individual sensors are close to each other and tactile stimuli often affect several sensors, either because objects are much bigger than individual sensors or because robotic parts (e.g. grippers) that are in direct contact with objects are coated in soft material, so that even small tactile impressions spread through the material and activate several sensors.

While there is a multitude of public data sets in the vision community, which can be used to test and compare machine vision algorithms, only a single database containing tactile feedback gathered during exploration of a variety of objects has been released so far (Schöpfer et al., 2007). This may have to do with the fact that the hardware is not yet as mature and is much more diverse than in the vision community, leading to a high number of rather specialized algorithms that rely on specific hardware or certain kinds of input data.

In recent years, machine learning techniques for extracting tactile features have become more widespread. This is due to sensors becoming more sophisticated and being able to transmit much higher-dimensional data, but also due to increased demand for more complex tactile features and tasks. For example, setting up hand-crafted models

that determine simple shape information from probes equipped with contact sensors is laborious but still possible. Devising models by hand for classifying objects from high-dimensional tactile data, on the other hand, is very difficult and needs to rely on at least some simple machine learning techniques. Quite a few of the techniques employed aim to automatically extract relevant features from the tactile data by using (mainly unsupervised) dimensionality reduction methods. For example, in one study multi-dimensional scaling is used in order to generate a "somatosensory map", which is then used to classify different types of tactile interactions, like hugging or patting (Noda et al., 2008). Other methods include (hierarchical) self-organizing maps, which have been used for haptic shape perception and dynamic object discrimination (Johnsson and Balkenius, 2007; Suzuki et al., 2006), unsupervised clustering for object discrimination (Schneider et al., 2009), and neural networks for grasp control (Ascari et al., 2009). The majority of these methods try to automatically reduce high-dimensional tactile data to a small number of relevant features in order to discriminate between different objects. One of the limitations of current studies is that the number of objects is usually rather small (often ten or less items).

There are also approaches that try to determine tactile features (like ridges or edges of objects) by defining and using a set of exploratory actions (e.g. Okamura and Cutkosky, 2001; Allen and Michelman, 1990). This work is partly inspired by psychological studies on humans and how they identify relevant tactile features (see Lederman and Klatzky, 1993, for a review, and our discussion in Sec. 8.2.2). Mostly, such methods have focused on pre-determined strategies rather than on machine learning techniques for either feature extraction or optimization of the exploration process. Similar strategies have also been used in order to help with grasping different objects (Natale and Torres-Jara, 2006). In these strategies, individual actions are not optimized for their usefulness (i.e. how informative they are about the feature in question), but rather chosen by the programmer either to mimic human strategies or because they simply seem to work. In Sec. 5.4.2 below, we review some approaches that try to automatically determine suitable strategies for feature extraction by maximizing the amount of information that is transmitted. Our own approach to this problem will be presented in Chapter 6.

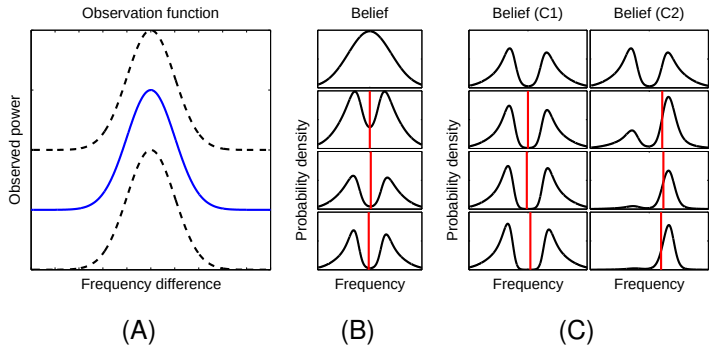

Figure 5.2: Illustration of the viscosity estimation problem. **A.** Mean (blue) and standard deviation (black) of the observation function, plotting the difference between the liquid's resonance frequency and the robot's shaking frequency against observed power in the frequency spectrum. The closer the shaking frequency is to a liquid's resonance frequency, the higher is the observed power. **B.** After each new observation, the probability distribution over the liquid's resonance frequency (and therefore its viscosity) is updated (time running from top to bottom panel). Red lines indicate the robot's shaking frequency at which the observation was obtained. **C.** Two examples (C1 and C2) of how the distribution might evolve given different shaking frequencies being selected by the robot.

5.2.3 Integrating perception and action

As we have seen in the previous section, it is often desirable to learn sensor models that connect object attributes with tactile feedback. Moreover, due to the nature of tactile exploration, data arrives sequentially and often decisions have to be made about how the robot should move its sensing devices in order to collect relevant data.

As a motivating example, let us consider a robot trying to determine the type of various liquids, for example water, vegetable oil, or peanut butter. These liquids differ in viscosity (i.e. their "thickness"). The dynamical behaviour of each container is determined by the liquid inside, which means that shaking them should result in different temporal force profiles being recorded by the touch sensors mounted on the robot's fingertips. These responses also differ depending on the frequency that the bottle is shaken with and other movement parameters. For example, when shaking close to the resonance frequency of a liquid, we would expect to observe a high amplitude in the power spectrum of the recorded tactile responses, while observing less power

when shaking at some other frequency. The smaller the difference in shaking frequency and the liquid's resonance frequency (which is determined by its viscosity), the higher the observed power in the frequency spectrum of the tactile sensor feedback (see Fig. 5.2A). In our simple example, we have modelled this relationship (the observation function) qualitatively using a squared exponential function. In general, the observation function could be known analytically (e.g. from a physical model), if the problem is simple enough, or learned from data.

We can now use the observation model to determine the viscosity of an unknown liquid. To do this within the Bayesian framework that we are considering, we place a broad prior distribution over viscosity space that reflects our initial belief of the viscosity of the unknown liquid. We then shake the bottle at some frequency, record the force profile on the tactile sensors, and update our belief of the viscosity using the newly gathered data and the observation function. We repeat this process until the viscosity belief becomes accurate enough for our purposes. Fig. 5.2B shows the robot's simulated estimate of the liquid's viscosity over time as new data arrives and influences the robot's belief state. We can see that the distribution quickly turns multimodal. This is a common occurrence for nonlinear observation models such as the one displayed in Fig. 5.2A. In Sec. 5.3 we introduce the problem of representing sequentially arriving data and discuss some approaches of how to deal with multimodality. In Chapter 7 we then present our own approach.

Another problem highlighted by our example is how to select the shaking frequency so as to estimate the viscosity of the unknown liquid as quickly as possible. Fig. 5.2C shows the effects of two different sequences of shaking frequencies (starting from the belief state that was reached at the end of Fig. 5.2B) and we can see that these influence the belief state quite differently: in the left column, where the shaking frequency is not close to the actual resonance frequency of the liquid, the distribution hardly changes, implying the data gathered was not very informative; however, in the right column, where the shaking frequency is close to the actual resonance frequency, one of the two modes is eliminated, suggesting that the data collected at this shaking frequency was much more useful. We discuss this problem—generally referred to as active learning—in Sec. 5.4, and then present our own approach in Chapter 6. In the following we will give an overview of sequential models and the active learning problem along with relevant robotics applications.

5.3 Dealing with sequential data

5.3.1 Common robotics approaches

Dealing with data that is arriving sequentially is a common problem in robotics, as all sensing devices, whether they are visual, tactile, or of another kind, will gather new data frames one after another, often at a fixed frequency. Integrating a new data frame with a robot's prior knowledge and with what has been learned from earlier data frames is thus an important problem. This especially holds true for data from touch sensors that often only provide a partial view of the relevant information due to the "local" nature of their feedback, so that multiple successive data frames have to be combined into a more complex representation. For example, it might not be possible to fully determine the shape of an object from just a single grasp. Instead, several sequential contact events at different locations on the object's surface might be required. The underlying representational framework must therefore be powerful enough to allow for easy and accurate incorporation of new data points, yet computationally still be tractable.

Several different approaches to this problem have been proposed, and we will mention the main ones in the following section along with relevant example involving tactile sensing, where possible. First, the problem can be simplified considerably by ignoring the sequential nature of the data and treating a set of data frames collected within a specified time span as a single data item (e.g. Takamuku et al., 2008). This has the obvious advantage that the data representation does not have to be updated sequentially. However, one big drawback is that whatever data is gathered cannot influence the robot's actions immediately. Thus, data gathering can be very time-consuming and ineffective, so that usually truly sequential methods are preferred. In the simplest case, the parameter of interest is directly observable. For example, when exploring shapes, coordinates of positions from an object's surface can be gathered by touch sensors through inverse kinematics and then incorporated into an object modelling framework (e.g. Dragiev et al., 2011). In such cases, both parametric and non-parametric methods for online learning can be used. However, in a lot of cases, object features of interest cannot be directly accessed. In such cases, observations depend on the underlying feature through some (often nonlinear) function. In some cases the function that maps from the underlying state to observations is known in analytical form, otherwise it can be learned from training data using regression techniques. For example, liquids of

various viscosities differ in their impact on tactile sensors when being shaken (see our earlier example in Sec. 5.2.3). What is needed in such cases is a way of inverting the function so that it maps from observations to the state[1]. Often direct calculation of an inverse function is intractable, but it might be possible to learn such a function from data. For example, (recurrent) neural networks have been used in order to classify objects from data obtained through dynamics movements (Ogata et al., 2005). Other approaches collect sets ("bags") of features over time that eventually disambiguate different objects (Schneider et al., 2009). All of these approaches can be interpreted as providing a direct mapping from observations to object features or identity. Finally, using fully probabilistic frameworks, sequential Bayesian methods maintain a probability distribution over the state which is updated at each step using Bayes' rule (e.g. Petrovskaya et al., 2006). These approaches rely on an observation function, which is usually more readily available than an inverse function. The fully probabilistic framework also works in cases were an inverse function may not exist and the resulting probability distributions over features become multimodal.

Often, the feature that one is interested in is not stationary, but dynamic, i.e. it changes over time. For example, a service robot mixing cocktails might have to take into account weight changes of bottles and glasses as drinks are poured. In such cases, changes in the state from one step to another are defined by a dynamics model. In the context of sequential Bayesian updating, such problems are known as filtering problems and will be discussed in the next section.

5.3.2 Nonlinear filtering for robotics

5.3.2.1 Formal problem setup

In filtering problems, observations $\mathbf{y}$ are received sequentially (at discrete time steps $t, t+1, \ldots$). These are are stochastically dependent on an underlying state θ, which cannot be observed directly. Additionally, the state evolves over time according to some underlying dynamics. The goal of filtering is to maintain a probability distribution over the state space and to update this distribution iteratively as new observations arrive and the state dynamics unfold.

The problem can be formalized as follows (see Fig. 5.3 for the corresponding graphi-

[1] Note that an inverse function might not exist, if the observation function is nonlinear.

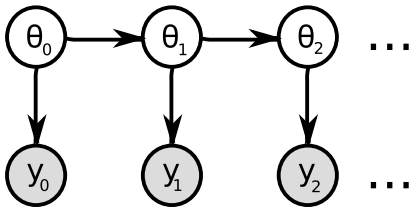

Figure 5.3: Gaphical model of the filtering problem. Time progresses from left to right, observed variables are shown in shaded grey.

cal model): both the state dynamics and observation models are known in the form of conditional probability distributions: $p(\theta_{t+1}|\theta_t)$ for the dynamics model and $p(\mathbf{y}_t|\theta_t)$ for the observation model. Often, these distributions are assumed to be Gaussian with a fixed noise covariance and a mean that depends through a dynamics function $g(\theta)$ and an observation function $f(\theta)$, respectively, on the state θ. Given the current belief of the state at time step t (after having received observations up to time step t), represented as the probability distribution $p(\theta_t|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0)$, we would like to calculate the updated belief for the next time step $t+1$, taking into account both the state dynamics, as well as the newly arrived observation $\mathbf{y}_{t+1}$. The update can be divided into a dynamics and an observation update. The dynamics update is calculated as follows:

$$p(\theta_{t+1}|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0) = \int \mathrm{d}\theta_t \; p(\theta_{t+1}|\theta_t) p(\theta_t|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0) \tag{5.1}$$

Next, the new observation $\mathbf{y}_{t+1}$ is incorporated in the observation update:

$$p(\theta_{t+1}|\mathbf{y}_{t+1}, \mathbf{y}_t, \ldots, \mathbf{y}_0) = \frac{p(\mathbf{y}_{t+1}|\theta_{t+1}) p(\theta_{t+1}|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0)}{p(\mathbf{y}_{t+1}|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0)} \tag{5.2}$$

where $p(\mathbf{y}_{t+1}|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0) = \int \mathrm{d}\theta_{t+1} \; p(\mathbf{y}_{t+1}|\theta_{t+1}) p(\theta_{t+1}|\mathbf{y}_t, \mathbf{y}_{t-1}, \ldots, \mathbf{y}_0)$.

In some cases, the dynamics and observation updates in Eqs. (5.1) and (5.2) can be calculated analytically (see Sec. 5.3.2.2 below), but in others approximations have to be used.

In the following section, we will describe several approaches to dealing with the filtering problem. We will concentrate on the observation update, as in Chapters 6 and 7, for simplicity, we only consider problems where the underlying state is constant[2]. In the following, we discuss different approaches that been brought forward to deal with the filtering problem. First, we discuss filters that maintain a unimodal Gaussian over the state distribution in Sec. 5.3.2.2. In Sec. 5.3.2.3 we then review extensions

[2]That is, we assume that the state dynamics correspond to a delta function $p(\theta_{t+1}|\theta_t) = \delta(\theta_{t+1} - \theta_t)$

that use mixtures of Gaussians for that purpose. Finally, we quickly mention Monte-Carlo approaches that represent the state distribution using a set of weighted samples in Sec. 5.3.2.4.

5.3.2.2 Unimodal Gaussian filters

The filtering problem can be solved exactly when it is in the linear-Gaussian form, i.e. the observation and dynamics functions depend linearly on the state and all probability distributions involved are normal, resulting in the well-known Kalman filter equations (Kalman, 1960). These equations describe the mean μ_{t+1} and covariance Σ_{t+1} of the Gaussian representing the state distribution $p(\theta_{t+1}|\mathbf{y}_{t+1}, \mathbf{y}_t, \ldots, \mathbf{y}_0)$:

$$\mu_{t+1} = \mu_t + \mathbf{C}_{t+1}^T \mathbf{S}_{t+1}^{-1} (\mathbf{y}_{t+1} - \mathbf{m}_{t+1}) \tag{5.3}$$

$$\Sigma_{t+1} = \Sigma_t - \mathbf{C}_{t+1}^T \mathbf{S}_{t+1}^{-1} \mathbf{C}_{t+1} \tag{5.4}$$

where $\mathbf{m}_{t+1}$ and $\mathbf{S}_{t+1}$ are the marginal mean and covariance of the observation, respectively; $\mathbf{C}_{t+1}$ is the cross-covariance between hidden state and observations; and $\mathbf{y}_{t+1}$ is the newly received observation. The marginal terms are calculated as follows[3]:

$$\mathbf{m} = \int \mathrm{d}\theta \; p(\theta) f(\theta) \tag{5.5}$$

$$\mathbf{S} = \int \mathrm{d}\theta \; p(\theta) f(\theta) f(\theta)^T \tag{5.6}$$

$$\mathbf{C} = \int \mathrm{d}\theta \; p(\theta) \theta f(\theta)^T \tag{5.7}$$

If $f(\cdot)$ is linear, then the integrals involved can be calculated analytically.

In many interesting cases, however, the observation function is nonlinear, necessitating the use of approximative methods. A popular type of approach deals with this problem by linearizing the observation function, such that the Kalman filter solution is applicable (see Lefebvre et al., 2004, for an overview) and the exact state distribution is approximated as a single Gaussian distribution at each time step. There are different approaches for calculating the linearization. The most well-known of these methods is the extended Kalman filter (EKF) which at each time step linearizes the observation function with a first-order Taylor expansion at the mean of the current belief state. This work has been extended to filters which try to incorporate higher orders of the Taylor expansion, for example in the truncated second order filter (Henriksen, 1982). A similar approach uses Stirling's interpolation formula for polynomial approximation

[3] We have dropped the time index for clarity.

instead of the Taylor expansion (Norgaard et al., 2000). More recently, a new class of filters has emerged that uses statistical linearization, which means that the observation function is linearized using the full state distribution, rather than just its mean. Such calculations are often intractable and deterministic sampling schemes have been employed in order to approximate the required integrals. These schemes rely on the unscented transform (Julier and Uhlmann, 2004) or Gauss-Hermite quadrature (e.g. Ito and Xiong, 2000). Filters using such methods have been found to be less computationally expensive than those using higher-order Taylor expansions while performing at least as well. If the observation function has a special form, e.g. a Gaussian Process or radial basis function network, analytical solutions to the statistical linearization problem are possible and have been used in Gaussian filters (Deisenroth et al., 2009).

Besides Kalman-like filters, there are several other approaches that represent the state distribution as a single Gaussian. For example, the iterated extended Kalman filter (Bell and Cathey, 1993) and the iterated unscented filter (Sibley et al., 2006) successively apply Kalman-like updates until convergence to a new Gaussian state distribution. Finally, other approaches try to approximate the true posterior distribution directly with a single Gaussian, using deterministic sampling to calculate the integrals involved (Kushner and Budhiraja, 2000). However, this requires the approximation of integrals in observation space, which is often high-dimensional.

5.3.2.3 Gaussian mixture filters

In cases where the observation function is highly nonlinear, the state distribution can have complex shapes or even become multimodal and is thus not well represented using just a single Gaussian distribution. One solution to this problem is to represent the state distribution as a mixture of Gaussians. Increasing the number of components in this mixture would allow this distribution to become arbitrarily close to the actual posterior distribution, provided that there is a way to correctly determine the individual components' optimal means and covariances, and the mixture weights.

The classic approach to Gaussian mixture filtering uses a weighted sum of extended Kalman filters (EKFs) running in parallel (Alspach and Sorenson, 1972). Newer approaches replace the EKFs with linear filters using deterministic sampling (Ito and Xiong, 2000; Arasaratnam et al., 2007). In all these cases several unimodal filters run independently of each other. While this is computationally very efficient, it also leads

to inferior representation of the posterior distribution, as interactions between the individual components are neglected. Different ways to compute the posterior mixture weights have been proposed, in an attempt to decrease the distance between the true posterior and the mixture approximation, but the impact on performance seems small (Ito and Xiong, 2000). In Chapter 7, we will present an approach that attempts to adapt *all* Gaussian mixture parameters (i.e. means, covariances, and mixture weights) *jointly* so as to fit the true posterior distribution as closely as possible.

5.3.2.4 Sequential Monte-Carlo methods

Another class of methods that has been developed to specifically deal with the problem of highly nonlinear observation functions and the resulting multimodal state distributions are sequential Monte-Carlo methods or particle filters (Doucet et al., 2001). These methods maintain a collection of samples θ_p with associated weights w_p that together represent a probability distribution:

$$p(\theta) = \sum_p w_p \delta(\theta - \theta_p) \tag{5.8}$$

For the observation step, the weights w_p are set to be proportional to the likelihood of the new observation. As this procedure can lead to an impoverished representation over time (e.g. lots of samples with tiny weights and a few with high weights), resampling schemes are usually employed that periodically draw a "fresh" set of samples from the current state distribution. While particle filter are potentially very good at representing complex distributions, these types of methods run into problems in very high-dimensional spaces, as in such cases an untenable number of particles would be needed to model the resulting probability distributions accurately.

Particle filters have frequently been used in robotics, especially in localization or tracking problems (e.g. Isard and Blake, 1996). There are currently a lot of applications to tactile sensing. One noteworthy example comes from a study where the authors estimated the pose of objects from sequential tactile "probings" and modelled their highly complex state distributions using sampling methods (Petrovskaya et al., 2006).

5.4 Active learning in robotics

In this section, we discuss active learning, i.e. the problem of how to select actions or movement parameters that influence data gathering, such that future observations are expected to be as informative as possible about a variable of interest. We first present the mathematical background and then present a survey of some applications of this framework to problems in robotics.

A large number of different scenarios is covered by the term "active learning" and the connections between those might not always be obvious at first sight (see Settles, 2010, for a broad review). For example, active learning techniques have been used in order to request labels for a small number of data point from a bigger data set so as to maximize the information about class boundaries. We will focus on a rather narrow definition of active learning and on applications that are relevant for robotics in the following sections.

5.4.1 Theoretical background

In most classical problems involving learning from data the learner stays passive in that it does not have any control over what data is provided. In many situations however, active feedback from the learner is not only possible, but even required. For example, if a robot is to gather data using a touch sensor when locating a specific object, it has to make a decision about how to move its sensor in order to make contact with objects, and explore the environment for specific and relevant features distinguishing this object from other objects. One can now pose the question how the robot should move its touch sensor in order to gather the most *informative* data about the object or feature of interest. We will formalize this question in the following.

We assume the existence of an observation model, which is given as a conditional distribution $p(\mathbf{y}|\theta,\mathbf{x})$, which gives the probability of an observation $\mathbf{y}$ (e.g. data from a touch sensor), given a parameter of interest θ (e.g. liquid viscosity or object identity) and an action $\mathbf{x}$ (e.g. shaking at a specific frequency). The problem is now how to choose a specific $\mathbf{x}^*$ such that the future observation $\mathbf{y}$ is expected to be maximally informative about the unknown parameter θ as compared to all other actions that could be executed (see Fig. 5.4A for the corresponding graphical model). Often, active learning problems are sequential, and can, for example, be combined with the filtering ap-

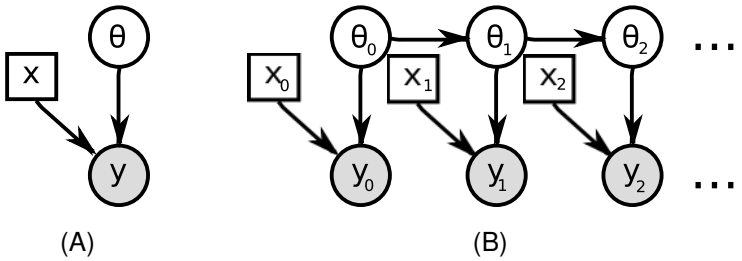

Figure 5.4: Graphical model of the active learning problem. **A.** Single step problem. **B.** Sequential problem, where a new action $\mathbf{x}$ has to selected at every time step t.

proaches presented in the previous section (as illustrated in Fig. 5.4B).

The classical approach to the active learning problem is to take a frequentist view and use maximum likelihood to estimate the underlying parameter and Fisher information as a measure of the informativeness of different actions. This method is known under the name Optimal Experimental Design (OED; see Silvey, 1980, for a comprehensive introduction). Usually, OED has been used when the model is linear in θ, as in this case the optimal action is independent of the true value of θ. In the linear-Gaussian case, the Bayesian criterion can be reduced to the Fisher information (Chaloner and Verdinelli, 1995) or the variance of the estimator (MacKay, 1992; Cohn et al., 1996). For nonlinear models, the optimal action depends on the current estimate of θ, and although the classical OED approach has been extended in this respect (Chaudhuri and Mykland, 1993), the most principled approach is to use the Bayesian approach and sequentially update the posterior $p(\theta|\mathbf{x},\mathbf{y})$ as new data comes in.

In such a Bayesian point of view, a prior over θ is introduced and the informativeness of an action can be quantified by the expected information gain, i.e. the expected KL divergence between the posterior and the prior of θ, which corresponds to the mutual information between θ and $\mathbf{y}$ (Lindley, 1956; Luttrell, 1985; Paninski, 2005):

$$\mathbf{I}(\theta;\mathbf{y}|\mathbf{x}) = \iint \mathrm{d}\mathbf{y}\,\mathrm{d}\theta\; p(\theta,\mathbf{y}|\mathbf{x}) \log \frac{p(\theta,\mathbf{y}|\mathbf{x})}{p(\theta)p(\mathbf{y}|\mathbf{x})} \tag{5.9}$$

$$= \iint \mathrm{d}\mathbf{y}\,\mathrm{d}\theta\; p(\mathbf{y}|\mathbf{x})p(\theta|\mathbf{y}) \log \frac{p(\theta|\mathbf{y})}{p(\theta)} \tag{5.10}$$

$$= E_{p(\mathbf{y}|\mathbf{x})}\mathrm{KL}\left[p(\theta|\mathbf{y})||p(\theta)\right] \tag{5.11}$$

The main downside of using mutual information as the criterion to rate the informativeness of certain actions is its high computational demand. Even rather simple models turn out to be analytically intractable, therefore one has to resort to numerical approx-

imations or analytical bounds. An approach that is often taken is to discretize the problem into a small number of states and observations such that the the entropy terms can still be calculated in reasonable time. Alternatively, one can use linear-Gaussian models (e.g. Grocholsky et al., 2003; Lewi et al., 2007), but these are either restricted to linear dynamics or to linearized approximations of non-linear dynamics. Moreover, Gaussian approximations do not capture multi-modality of the posterior, as happens when incoming data is inherently ambiguous. In such cases, particle filters have been used, which can capture both non-linear dynamics and arbitrary distributions (Hoffmann et al., 2006). However, if the state-space is high-dimensional a prohibitively large number of samples is needed. This problem is even more prnounced when measures like mutual information are to be computed from these samples, as these measures suffer considerably from under-sampling. Finally, another approach that has been used is to generate a small number of candidate actions and then choose the one with the highest expected information gain. Such a model has been used in a study that generated internal models of the current configuration of the robot (e.g. for detecting missing limbs or other reconfigurations) by performing discriminatory actions (Bongard et al., 2006).

5.4.2 Applications in robotics

Active learning has been applied to a number of robotics problems, where it has proven useful, especially in problems that involve cluttered and unstructured environments. In such scenarios, data about object features that are relevant for manipulation are often only partially observable and data from vision, tactile sensors, and other sensing technologies is noisy and incomplete. Moreover, gathering data can be expensive, both in terms of the time that is needed as well as in energy requirements. State-of-the-art robots also have a high number of degrees of freedom, resulting in a seemingly infinite number of different movements that can be executed. Thus, knowledge about which movements are going to result in informative data from touch sensors or video cameras is an important problem. Older approaches often used pre-determined strategies (see Sec. 5.2.2) rather than online optimization, but in more recent applications active learning approaches have started to become more widespread.

One area where active learning has proven useful is in determining object position using tactile data. Often, camera images can give rough position information, but are not accurate enough to allow for object manipulation (like grasping), or suffer from

occlusion. A recent study used partial camera views in order to determine regions occupied by objects and then used tactile sensing to refine that estimate (Bohg et al., 2010). In this example, the search trajectories were optimized in order to minimize the residual uncertainty of the occupancy model. In another application, objects were classified from a number of sequential tactile impressions resulting from different grasping actions (Schneider et al., 2009). Each of these actions was optimized to maximize information about object class. There are also applications to improving the grasping of objects. These either try to establish good grasping locations (Kroemer et al., 2009) or try to improve object localization by determining exploratory actions that improve knowledge about object position (Hsiao et al., 2010).

Vision is another robotics area where active learning has been proven useful. It has been used successfully in image classification where the classifier can specify a viewpoint from which the next picture should be taken (Schiele and Crowley, 1998; Paletta et al., 2000; Callari and Ferrie, 2011; Denzler and Brown, 2002). For example, one of the studies (Denzler and Brown, 2002) optimized the control of camera parameters such as zoom and pan while sequentially taking pictures in order to discriminate between different objects. Some of these objects differed only in minor details, so a high zoom level was needed in order to distinguish them, while other objects could be classified more easily using broader views. More recent work has extended this approach to include pose estimation (Ma and Burdick, 2009).

Finally, some techniques related to active learning have been used for robot control. For example, some studies have used reinforcement learning (Paletta and Pinz, 2000) or evolutionary algorithms (Tuci et al., 2009) in order to learn control policies that allow for quick discrimination of objects or parameter discrimination. Other approaches combine learning about the model and using it, leading to the well-known problem of how to trade off exploration and exploitation. One study developed a method that did path planning, while improving its position estimate simultaneously (Martinez-Cantin et al., 2009).

All of these applications highlight the benefits of using active learning, both in terms of increased speed of convergence as well as better parameter estimates. Moreover, rather than having to be hand-crafted, strategies are adapted so as to be maximally informative.

Chapter 6

Active learning for tactile processing in robotics

6.1 Introduction

In this chapter, we will present our own approach to the active learning problem. Specifically, we will try to maximize tactile feedback information gathered by having a robot perform movements that are maximally discriminative with respect to an object feature we are interested in and our current probabilistic belief of the value of this feature.

In application domains such as service robotics, haptics, surgical robotics, and human-robot interaction, sophisticated dexterous manipulation of objects can be crucial. In the absence of visual information, force and tactile sensory data are typically the only feedback available. Since not all the dynamic properties of objects will be known beforehand, the robot may have to estimate some of these features "on the fly". For example, when handling open containers filled with fluids, the viscosity of the liquid constrains the kinds of movement a robot can make without any spillage; while containers with highly viscous fluids like jam can be moved quickly, cautious and slower movements should be used for fluids of lower viscosity like soft drinks. Quick and reliable estimation of such parameters is, thus, important for bringing robots into real-world environments. Importantly, viscosity is a dynamical feature in the sense that the container has to be moved for viscous properties to manifest themselves.

To estimate an object's dynamic properties on the fly, the robot needs to be equipped

with a sensor model that maps both the object parameter of interest and the robot's current state to sensor readings. Assuming the sensor model is known (either learned from data or modelled using physics-based models), the robot is faced with the challenge of determining a sequence of actions to perform (e.g. shaking movements at certain frequencies) in order to gather maximally informative observations. The smallest set of informative observations is desired so that the object feature of interest can be estimated as quickly as possible. We are interested in application domains where observations and actions are potentially high-dimensional and continuous and where decisions need to be made quickly, often close to real-time.

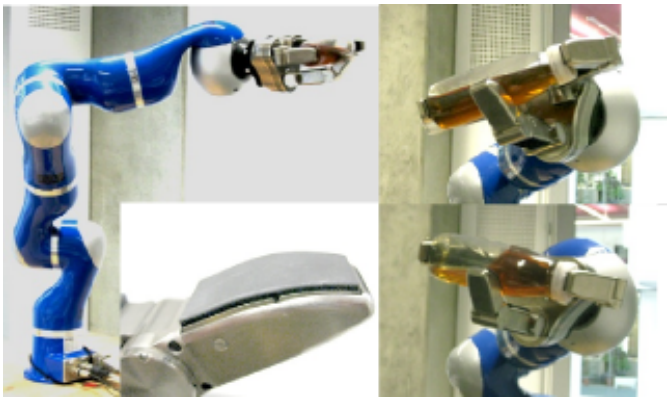

Figure 6.1: Left: Schunk SDH 7-DOF hand mounted on DLR 7-DOF robotic arm. Inset: Touch sensor array on robotic finger. Right: Liquid-filled bottles are shaken at different frequencies and rotation angles. Top right: Positive rotation angle, with bottle tilted upwards. Bottom right: Negative angle, leading to downwards tilt.

In our approach, we use active learning in order to speed up the sequential estimation of the viscosity of various liquids. Filled bottles are gripped and then shaken by a robotic hand-arm system, shown in Fig. 6.1, while tactile feedback from touch sensor arrays mounted on the robotic fingers is used to determine the viscosity of the liquids. Crucially, we optimize the robot's actions (shaking frequencies and rotation angle, as shown in Fig. 6.1) at each time step in order to maximize the informativeness of the next incoming observation.

We describe our proposed framework in Sec. 6.2 before presenting experimental results in Sec. 6.3. Evaluations are shown for both simulations on toy problems as well as a liquid viscosity estimation task using the DLR-Schunk robotic hand-arm system (see Fig. 6.1).

6.2 Methods

6.2.1 Problem formulation

We assume the following notation: $\theta \in \mathbb{R}^{d_\theta}$ is the object feature of interest that we wish to estimate (here, viscosity of a liquid). $\mathbf{x} \in \mathbb{R}^{d_\mathbf{x}}$ is a vector of action parameters (e.g. shaking frequency and rotation angle). $\mathbf{y} \in \mathbb{R}^{d_\mathbf{y}}$ is observed sensory data. We assume that observations $\mathbf{y}$ are a (nonlinear) function of both the actions $\mathbf{x}$ and state parameters θ, thus we get for each output dimension m:

$$y_m = f_m(\mathbf{x}, \theta) + \varepsilon_m \tag{6.1}$$

where $\{\mathbf{x}, \theta, \mathbf{y}\}$ are all continuous, and ε_m is Gaussian observation noise. Both $\mathbf{x}$ and $\mathbf{y}$ can be high-dimensional.

Assuming that the observation model in Eq. (6.1) is learnt accurately in an offline training phase, we are faced with the task of estimating the particular value of θ^* (i.e. the "true" value of the parameter) at run/test time. As data samples are sequentially observed/collected, our estimate of θ^*, denoted as $p(\theta_{t+1}|\mathbf{y}_{t\ldots0}, \mathbf{x}_{t\ldots0})$, is recursively updated whenever an action $\mathbf{x}_{t+1}$ has been taken and the corresponding observations $\mathbf{y}_{t+1}$ received, to yield $p(\theta_{t+1}|\mathbf{y}_{t+1\ldots0}, \mathbf{x}_{t+1\ldots0})$. For more details see the background in Sec. 5.3.2. We note that in our current setup the value of θ^* is regarded as constant, but it is straightforward to extend our approach to the general filtering problem[1].

In order to incorporate active learning, we are additionally interested in determining the optimal actions $\mathbf{x}_t^*$ to take at each time step so that the mutual information (MI) between the expected next observation $\mathbf{y}_t$ and the current belief $p(\theta_t|\mathbf{y}_{t-1\ldots0}, \mathbf{x}_{t-1\ldots0})$, denoted by $\mathbf{I}(\theta_t; \mathbf{y}_t|\mathbf{x}_t)$, is maximized, i.e., $\mathbf{x}_t^* = \arg\max_{\mathbf{x}_t \in \mathbf{X}} \mathbf{I}(\theta_t; \mathbf{y}_t|\mathbf{x}_t)$. The mutual information is defined as:

$$\mathbf{I}(\theta_t; \mathbf{y}_t|\mathbf{x}_t) = \iint \mathrm{d}\mathbf{y}_t \mathrm{d}\theta_t \, p(\theta_t, \mathbf{y}_t|\mathbf{y}_{t-1\ldots0}, \mathbf{x}_{t\ldots0}) \log \frac{p(\theta_t, \mathbf{y}_t|\mathbf{y}_{t-1\ldots0}, \mathbf{x}_{t\ldots0})}{p(\theta_t|\mathbf{y}_{t-1\ldots0}, \mathbf{x}_{t-1\ldots0}) p(\mathbf{y}_t|\mathbf{x}_t)} \tag{6.2}$$

See Sec. 5.4 for a more thorough overview of the active learning problem, along with pointers to the relevant literature.

To provide a more concise overview of the full algorithm, all necessary steps are outlined in Algorithm 6.1.

[1] That is, the dynamics update is simply $p(\theta_{t+1}|\mathbf{y}_{t\ldots0}, \mathbf{x}_{t\ldots0}) = p(\theta_t|\mathbf{y}_{t\ldots0}, \mathbf{x}_{t\ldots0})$.

Algorithm 6.1 Pseudo code detailing our sequential active learning approach. The algorithm can either be run for a fixed number of time steps T, or until $p(\theta)$ has converged.

learn sensor model $p(\mathbf{y}|\theta,\mathbf{x})$ from training data set
set prior distribution $p(\theta_0)$
for $t = 1 \rightarrow T$ **do**
  select action $\mathbf{x}_t^*$ by optimizing $\mathbf{I}(\theta_t;\mathbf{y}_t|\mathbf{x}_t)$
  execute $\mathbf{x}_t^*$ and gather sensor feedback $\mathbf{y}_t$
  calculate $p(\theta_t|\mathbf{y}_{t\ldots0},\mathbf{x}_{t\ldots0})$ to reflect new observation $\mathbf{y}_t$
end for

In the following sections, we take a first step towards tackling the problems described above. First, in Sec. 6.2.2 we describe how the sensor model is represented and learned. Then, we propose our active sequential updating approach, which is based on a Gaussian approximation in Sec. 6.2.3. Finally, we examine the computational complexity of the proposed methods and describe implementation details in Sec. 6.2.4.

6.2.2 Learning the sensor model

In this section, we describe how we train a regression model on a data set, so as to approximate the true function $f(\cdot)$, given in Eq. (6.1), as closely as possible. We assume that the data consist of a set of independent samples $D = \left\{\mathbf{x}^i_{\text{train}}, \mathbf{y}^i_{\text{train}}, \theta^i_{\text{train}}\right\}_{i=1}^N$, which are gathered offline in an initial training phase.

We first describe a sensor model that is based on Gaussian Processes. This regression model is straightforward to optimize and can be used to find the predictive mean and variance for any given input, with the variance depending on the distance of other training data. Then, we will focus on an extension of this model based on sparse Gaussian Processes that is capable of handling many more data points and is therefore of more use in real-world applications.

6.2.2.1 Gaussian Process model

To learn the sensor model, we use a Gaussian process (GP; Williams and Rasmussen, 1995; Rasmussen and Williams, 2006) to approximate the nonlinear function $f(\cdot)$ in Eq. (6.1). GPs rely on a kernel function to determine the correlations between different

input points. A kernel matrix Λ containing the covariance between each pair of input data points $\mathbf{z}_p$ and $\mathbf{z}_q$ is computed from this kernel function, so that $\Lambda_{pq} = k(\mathbf{z}_p, \mathbf{z}_q)$. We learn a GP for each output dimension $m = 1, .., d_\mathbf{y}$. We use a squared exponential kernel function of the following form to compute the covariance:

$$k_m(\mathbf{z}_p, \mathbf{z}_q) = \alpha_m^2 \exp\left\{\frac{1}{2}(\mathbf{z}_p - \mathbf{z}_q)^T \mathbf{H}_m^{-1}(\mathbf{z}_p - \mathbf{z}_q)\right\} + \sigma_m^2 \delta_{pq} \tag{6.3}$$

where $\mathbf{z} \in \mathbb{R}^{d_\theta + d_\mathbf{x}}$ is the vector $\begin{bmatrix}\theta^T & \mathbf{x}^T\end{bmatrix}^T$; $\mathbf{H}_m = \begin{pmatrix} \mathbf{H}_m^\theta & 0 \\ 0 & \mathbf{H}_m^\mathbf{x} \end{pmatrix}$, $\mathbf{H}_m^\theta$ and $\mathbf{H}_m^\mathbf{x}$ are diagonal matrices; α_m^2 is a scaling parameter; and σ_m^2 denotes the variance of additive noise. The set of hyperparameters Γ_m to be optimized for kernel function k_m is $\Gamma_m = \left\{\alpha_m^2, \sigma_m^2, \mathbf{H}_m^\theta, \mathbf{H}_m^\mathbf{x}\right\}$.

Given the Kernel matrix and a query point $\tilde{\mathbf{z}} = [\tilde{\theta}^T \;\; \tilde{\mathbf{x}}^T]^T$, the predictive mean m_m^* and variance v_m^* can then be calculated as follows, for each dimension m:

$$m_m^* = \mathbf{k}_m(\tilde{\mathbf{z}})^T \Lambda_m^{-1} \mathbf{y}_m \tag{6.4}$$

$$v_m^* = \sigma_m^2 - \mathbf{k}_m^T \Lambda_m^{-1} \mathbf{k}_m, \tag{6.5}$$

where $\mathbf{k}_m(\tilde{\mathbf{z}})$ is a row vector, where each entry contains the covariance between the respective point from the training data set and $\tilde{\mathbf{z}}$.

6.2.2.2 Sparse Gaussian Process model

If the number of training points is large, the Gaussian Process Kernel matrix becomes very big, slowing down learning as well as calculating updates in a filtering context (see Sec. 6.2.4, and Deisenroth et al., 2009). To remedy this problem, sparse extensions of GPs have been developed that rely on a small number of pseudo-inputs to speed up computations. We use one of these methods (Snelson and Ghahramani, 2006), where the location of the pseudo-inputs are optimized so as to maximize the likelihood of the model. An additional benefit of this approach is that it can incorporate heteroscedastic noise, albeit in limited form. Other approaches for modelling sparse GPs have been developed, and could be substituted here (see Quiñonera Candela and Rasmussen, 2005, for an overview).

In this model, we again compute the mapping from object features θ and action parameters $\mathbf{x}$ to a probability distribution over sensor readings $\mathbf{y}$. In particular, the predictive

distribution over $\mathbf{y}$ for a test point $\tilde{\mathbf{z}}$ is Gaussian with predictive mean m_m^* and predictive variance v_m^* is now expressed as follows:

$$m_m^* = \mathbf{k}_m(\tilde{\mathbf{z}})^T \mathbf{W}_m^{-1} \mathbf{K}_m^{\uparrow} \Lambda_m^{-1} \mathbf{y}_m \tag{6.6}$$

$$v_m^* = \sigma_m^2 - \mathbf{k}_m(\tilde{\mathbf{z}})^T (\mathbf{K}_m^{-1} - \mathbf{W}_m^{-1}) \mathbf{k}_m(\tilde{\mathbf{z}}) \tag{6.7}$$

where $\mathbf{W}_m = \mathbf{K}_m + \mathbf{K}_m^{\uparrow} \Lambda_m^{-1} \mathbf{K}_m^{\downarrow}$. Here, Λ_m is the kernel matrix calculated from the training inputs, $\mathbf{K}_m$ is the kernel matrix for the pseudoinputs and $\mathbf{K}_m^{\uparrow}$ and $\mathbf{K}_m^{\downarrow}$ are $N \times U$ and $U \times N$ matrices, respectively, containing the correlations between pseudo-input and training inputs, where U is the number of pseudo-inputs and N is the number of training inputs.

The GPs are learned from training data by first optimizing the hyperparameters Γ_m of each kernel function k_m while keeping the location of the pseudo-inputs fixed, and then optimizing the position of the pseudo-inputs. Optimization is done by gradient ascent on the model likelihood (Snelson and Ghahramani, 2006). To prevent overfitting, we leave out some of the initial training data gathered to use as a validation set. Once the estimated error on this validation set starts to increase, we stop optimization.

6.2.3 Information theoretic active belief update

Once the sensor model is learned offline, we can use it in order to perform online estimation of the dynamical parameter we are interested in. That is, we would like to determine suitable actions as well as update the belief over the parameter given sensor feedback when executing these actions.

In this section, we describe an update method that represents the distribution $p(\theta_t)$ using a single Gaussian $\mathcal{N}(\theta|\mu_t, \Sigma_t)$. We follow the general results by Deisenroth et al. (2009), but extend their method by adding an active learning component.

In this approach, the posterior mean μ_{t+1} and posterior covariance Σ_{t+1} are computed with Kalman-filter like updates (see overview in Sec. 5.3.2.2):

$$\mu_{t+1} = \mu_t + \mathbf{C}_{t+1}^T \mathbf{S}_{t+1}^{-1} (\mathbf{y}_{t+1}^{obs} - \mathbf{m}_{t+1}) \tag{6.8}$$

$$\Sigma_{t+1} = \Sigma_t - \mathbf{C}_{t+1}^T \mathbf{S}_{t+1}^{-1} \mathbf{C}_{t+1} \tag{6.9}$$

where μ_t is the posterior mean of θ at time step t and Σ_t is the corresponding posterior covariance. $\mathbf{y}^{obs}$ is the new sensory data. The terms $\mathbf{C}_{t+1}$, $\mathbf{S}_{t+1}$, and $\mathbf{m}_{t+1}$ correspond to the cross-covariance between $p(\theta_{t+1}|\mathbf{y}_{t\ldots 0}, \mathbf{x}_{t\ldots 0})$ and $p(\mathbf{y}_{t+1}|\mathbf{x}_{t+1})$, and the marginal

mean and variance of $p(\mathbf{y}_{t+1}|\mathbf{x}_{t+1})$, respectively (see Eqs. (6.11)–(6.13) below and Sec. 5.3.2.2).

We introduce an active component to the selection of action $\mathbf{x}_{t+1}$ at time step t by maximizing the mutual information $\mathbf{I}(\theta_{t+1};\mathbf{y}_{t+1}|\mathbf{x}_{t+1})$, as given in Eq. (6.2), between the belief over θ_{t+1} and expected observations $\mathbf{y}_{t+1}$, conditioned on the next action $\mathbf{x}_{t+1}$. That is to say, the optimal action $\mathbf{x}^*_{t+1}$ can be found with the following equation:

$$\begin{aligned}\mathbf{x}^*_{t+1} &= \text{argmax}_{\mathbf{x}_{t+1}}\, \mathbf{I}(\theta_{t+1};\mathbf{y}_{t+1}|\mathbf{x}_{t+1}) \\ &= \text{argmax}_{\mathbf{x}_{t+1}}\, |\mathbf{C}_{t+1}(\mathbf{x}_{t+1})\mathbf{S}_{t+1}(\mathbf{x}_{t+1})\mathbf{C}_{t+1}(\mathbf{x}_{t+1})^T|\end{aligned} \tag{6.10}$$

where the above expression is maximized by performing gradient ascent.

In the rest of this section, we drop the time index (subscript t) on $\mathbf{C}$, $\mathbf{S}$, and $\mathbf{m}$. Since the sensor model in Eq. (6.1) is non-linear, these quantities have to be approximated, as is done, for example, in the unscented Kalman filter (Julier and Uhlmann, 2004). However, it has been shown that these quantities can be calculated analytically if $p(\mathbf{y}_{t+1}|\mathbf{x}_{t+1})$ is approximated with a Gaussian distribution (Girard et al., 2003; Deisenroth et al., 2009):

$$\begin{aligned} m_m &= \int \mathrm{d}\theta f_m(\theta,\mathbf{x})p(\theta) \\ &= \mathbf{q}_m(\mathbf{x})^T\mathbf{a}_m \end{aligned} \tag{6.11}$$

$$\begin{aligned} S_{mn} &= \int \mathrm{d}\theta p(\theta) f_m(\theta,\mathbf{x}) f_n(\theta,\mathbf{x}) - m_m(\mathbf{x})m_n(\mathbf{x}) \\ &= \mathbf{a}_m^T\mathbf{Q}_{mn}(\mathbf{x})\mathbf{a}_m - m_m m_n + \delta(m-n)\left(\alpha_m^2 - \text{tr}(\mathbf{B}_m\mathbf{Q}_{mm})\right) \end{aligned} \tag{6.12}$$

$$\begin{aligned} C_{mn} &= \int \mathrm{d}\theta f_m(\theta,\mathbf{x})\theta p(\theta) \\ &= \mathbf{Z}_m^T(\mathbf{x})\mathbf{a}_m - \mu_n m_m \end{aligned} \tag{6.13}$$

where m_m is the m-th coefficient of $\mathbf{m}$; S_{mn} is the (m,n)-th entry of the matrix $\mathbf{S}$; C_{mn} is the (m,n)-th entry of the matrix $\mathbf{C}$. Moreover, $\mathbf{a}_m = \Lambda_m^{-1}\mathbf{y}_m$ or $\mathbf{a}_m = \mathbf{W}_m^{-1}\mathbf{K}_m^{\uparrow}\Lambda_m^{-1}\mathbf{y}_m$, depending on whether standard or sparse GPs are used for the sensor model (see Sec. 6.2.2). Likewise, $\mathbf{B}_m = \Lambda_m^{-1}$ or $\mathbf{B}_m = \mathbf{K}_m^{-1} - \mathbf{W}_m^{-1}$, respectively.

We note that $\mathbf{q}_m$ and $\mathbf{Q}_{mn}$ can be written as Schur products of two factors:

$$\mathbf{q}_m = \mathbf{q}_m^{\theta} \circ \mathbf{q}_m^{x} \quad \text{and} \quad \mathbf{Q}_{mn} = \mathbf{Q}_{mn}^{\theta} \circ \mathbf{q}_m^{x}(\mathbf{q}_n^{x})^T \tag{6.14}$$

Detailed equations for $\mathbf{q}_m$ and $\mathbf{Q}_{mn}$ as well as derivatives can be found in Appendix A.1.

6.2.4 Computational complexity

During run time, the computational complexity of the analytical Gaussian update is $O(d_{\mathbf{y}}^3) + O((d_\theta + d_{\mathbf{x}})d_{\mathbf{y}}^2 N^2)$ at each time step (see also Deisenroth et al., 2009). The additional computational burden due to the information maximization in the active case is linear in the number of steps the gradient ascent takes. Using a sparse GP (Snelson and Ghahramani, 2006) reduces the computational complexity at each time step, where now $N = M$, i.e. the number of pseudo-inputs that is used.

It is worth noting that during optimization each successive step along the gradient improves the informativeness of the selected action. Therefore the optimization does not have to be run until convergence for time-critical applications, while still providing benefits compared to the passive case.

6.3 Results

In the following two sections, we present experimental evaluations of our active learning scheme. We first show results on synthetic data in Sec. 6.3.1 to analyze the benefits of our approach for different sensor models and noise conditions. Then, in Sec. 6.3.2, we present results obtained from a real-world task on a robotic system.

6.3.1 Synthetic data

The simulations presented in this section along with the figures were prepared by Jo-Anne Ting. They are included here for illustration and completeness.

Table 6.1: Converged average predictive squared error values ($\times 10^{-3}$) after 200 time steps for different belief updating approaches, averaged over 10 trials, for functions (i)–(iii).

	Function (i)		Function (ii)		Function (iii)	
Observation noise	low	high	low	high	low	high
Gaussian/random	0.1	1.7	-	-	-	-
Gaussian/active	0.1	0.5	1.8	-	-	-
MC/random	0.4	1.8	1.3	6.2	3.5	83.3

We applied the active belief update to synthetic data sets, for the purpose of estimating a particular value of θ^* at runtime. We compared these results with the passive version (i.e. random actions) and with a sampling-based approach (see Sec. 5.3.2.4) that also used random actions. All code was implemented in Matlab; for fitting Gaussian processes we used the GPML toolbox, version 2007-07-25[2].

We considered the following three observation models, which cover a range of nonlinearities:

i. $y = \exp(-(\theta - x)^2)$,

ii. $y = \cos(\theta)\sin(\theta x)$,

iii. $y = \exp(-(\theta_1 - x_1)^2) + \exp(-0.5(\theta_2 - x_2)^2)$.

Functions (i) and (ii) both assume scalar θ and x, while in function (iii), both θ and $\mathbf{x}$ are 2-dimensional. We generated data sets, each having 500 samples, from functions (i)–(iii) with either a small amount or a larger amount of observation noise[3]. We used a Gaussian Process with $N = 500$ inputs during training to learn the observation model over $\mathbf{y}$. The sampling-based belief updates used a population of $P = 500$ samples. For the MC version, we also applied a small amount of noise after each time step (corresponding to an identity dynamics model with added noise), in order to avoid divergence due to particle "stacking".

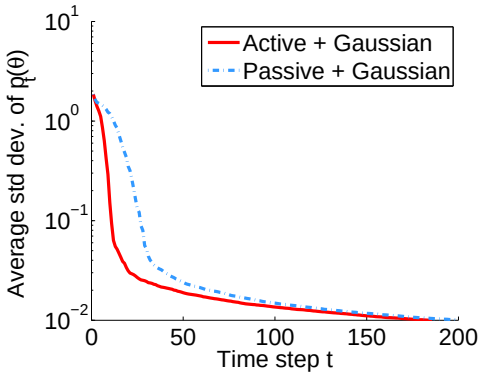

Figure 6.2: Average standard deviations of posterior $p(\theta_t)$ over all time steps t, averaged over 10 trials for function (i) with high observation noise.

[2] available at `http://www.gaussianprocess.org/gpml/code/matlab/doc/`

[3] We parameterize noise with the coefficient of determination, $r^2 = (\sigma_y^2 - \sigma_{res}^2)/\sigma_y^2$, where σ_{res}^2 is the variance of the residual error. We added noise scaled to the variance of the noiseless $\bar{\mathbf{y}}$ such that $\sigma_{noise}^2 = c\sigma_{\bar{\mathbf{y}}}^2$, where $c = 1/r^2 - 1$. We set $r^2 = 0.995$ and $r^2 = 0.95$ for the low and higher noise cases, respectively.

Table 6.1 lists the three different frameworks evaluated: a) Gaussian belief updates with random actions, b) active Gaussian belief updates, and c) MC sampling-based belief updates with random actions. The table also shows the converged average predictive squared error (MSE) (using the posterior mean of θ) after 200 time steps for the four frameworks, evaluated on functions (i)–(iii). Results are averaged over 10 trials for each combination of framework and observation. Function (i) appears to be relatively easy (with all methods attaining low predictive errors), while functions (ii) and (iii) are harder, being highly nonlinear in θ. In this scenario, analytical Gaussian belief updating fails to capture the multimodal distribution of $p(\theta_t)$, and often fails to converge correctly. In such cases, the MSE would be meaningless and was therefore omitted[4]. In chapter 7, we will discuss an approach that deals with multimodal posterior distributions without relying on Monte-Carlo sampling. The results also show that, as expected, a higher observation noise results in slower convergence. As the results show, the active method results in a considerably lower error than the respective passive counterpart in the majority of cases that we considered.

Fig. 6.2 shows the standard deviation of the posterior $p(\theta_t)$, averaged over all 10 trials, for Gaussian updating using function (i). These figures demonstrate that adding an active component to sequential data selection/collection appears to speed up convergence to θ^*.

6.3.2 Evaluation on robotic system

We also tested our proposed framework on a real robotic system using a viscosity estimation problem from tactile data. We first describe the setup along with implementation issues and data pre-processing in Sec. 6.3.2.1. Then, we present results on a 1D action space in Sec. 6.3.2.2 and on a more challenging 2D space in Sec. 6.3.2.3.

6.3.2.1 Setup and preprocessing

We examined the problem of viscosity estimation from tactile sensory data using the robotic anthropomorphic hand-arm system in Fig. 6.1. The setup consisted of the

[4]Even though reported results are averaged over 10 trials, the range of average predictive error values hold for Gaussian belief updates when averaged over 100 trials.

(A) (B)

Figure 6.3: **A.** Robotic finger with two tactile arrays shown in black. **B.** Close-up of bottle being gripped and shaken by the robot.

KUKA light weight arm (LWR)[5] on which a Schunk Dexterous Hand 2 (SDH)[6] was mounted. The hand exhibits 3 fingers, and 7 degrees of freedom in total, with 2 joints per finger plus an additional rotational joint that allows for rotating two of the fingers simultaneously at the base. The fingers are equipped with 6 tactile arrays in total (2 per finger), each of which contains 486 texels. They are manufactured by Weiss Robotics[7]. In the current setup a 12 bit value corresponding to the activation of each of the texels is transmitted at a sampling frequency of 30 Hz. Fig. 6.3 shows a single finger as well as a bottle being gripped for shaking.

The robot's task was to determine the viscosity θ of various liquids in bottles by shaking the containers using different action parameters $\mathbf{x}$. We evaluated the robot's performance on this task using either a one-dimensional action space, where we only optimized shaking frequency, or a two-dimensional action space, where we additionally optimized the rotation angle of the shaking movement (that is, the angle of the last arm joint at one of the end-points of the shaking trajectory). At a rotation angle of 0 degrees, the bottle is held horizontally while being shaken. For positive or negative rotation angles, the bottle is tilted upwards or downwards, respectively. Sensor data $\mathbf{y}$ is measured from tactile arrays mounted on the robot's fingers, while the bottle is being shaken, and is preprocessed as described in the next paragraphs.

To gather training data, we took plastic bottles[8] containing three different liquids and

[5] http://www.kuka-robotics.com/en/products/addons/lwr/

[6] http://www.schunk-modular-robotics.com/left-navigation/service-robotics/components/actuators/robotic-hands/sdh.html

[7] http://www.weiss-robotics.de/en/

[8] The bottles were standard super-market bought *Innocent pure fruit smoothies (strawberries and bananas)* with a capacity of 250 ml that had been emptied and then re-filled with one of the four test liquids.

recorded the tactile responses while shaking the bottles under a varied set of action parameters. The three liquids had viscosities of 1 cst[9] (water), 120 cst (motor oil), and 1200 cst (glycerine). These values were transformed to $\log_{10}$ space[10], yielding values of 0, 2.07 and 3.07, respectively. The bottles used for the three liquids had identical shape, and the content was matched for weight (160g). Bottles were gripped using a force-controlled strategy, after which finger joint positions were held constant throughout the shaking motions. A typical tactile response during shaking movements would involve time series data for around 20 active texels. When gathering training data for the 1D action task, bottles were shaken at a range of frequencies from 0.3 to 1.1 Hz in 0.1 Hz increments for 5 s each, while tactile time series were recorded. The shaking amplitude was fixed at 10 cm. For the 2D task, apart from the shaking frequency rotation angles were varied from -25 to 25 degrees. In that case, training data came both from a fixed grid over action space as well as randomly chosen actions.

We preprocessed the tactile data in the following way. First, the time series data containing the responses of all active texels recorded during the shaking movements were projected onto their principal component, in order to achieve spatial invariance with respect to the most responsive texels. Then, we calculated the Fourier transform of the resulting time series and normalized the power spectrum between the frequencies of 1.35 Hz and 2.9 Hz (for $d_y = 11$, 2D task) or 0.30 Hz and 3.64 Hz (for $d_y = 23$, 1D task) in order to remove any effects due to variations in individual grip strengths and locations. The time series data were re-normalized at each time step in order to counteract drift of the tactile responses over time that were due to small changes in bottle and finger positions as shaking progressed.

We then fitted individual Gaussian Processes to each Fourier component (11–23 in total). The resulting model maps the joint space of viscosity and shaking frequency $[\boldsymbol{\theta}\ \mathbf{x}]^T$ (2D or 3D) to the preprocessed tactile space $\mathbf{y}$ (23D or 11D). For the 1D task, we used a standard GP model (see Sec. 6.2.2.1). 270 training points were used for this in total. For the 2D task, we used sparse Gaussian Processes (see Sec. 6.2.2.2) and collected 5070 training points in total and used 297 pseudo-inputs[11]. The pseudo-input were initialized over a grid over action parameter space (instead of random initialization) before being optimized. For optimization we used a modified version of Ed

[9] centi-stokes (cst) is the commonly used measure of viscosity.

[10] Viscosity is often displayed on a logarithmic scale.

[11] Because the pseudo-inputs were initialized over a grid instead of randomly, we used a grid, discretized over 11 shaking angles, 9 shaking frequencies and 3 viscosity values, to give $11 \times 9 \times 3 = 297$ pseudo-input values.

Snelson's sparse GP Matlab code[12].

At the start of run/test time, we placed an initial broad Gaussian prior over the viscosity space and ran the experiment for 20 time steps. At the very first time step, a shaking frequency of 0.5 Hz (at a 0 degree rotation angle for the 2D task) was always chosen since the sensor values had to be normalized to account for slight variations in grip force and grip location. As in the training part, we recorded tactile responses for 5 s. When switching action parameters, we allowed the liquid to settle into the new shaking pattern for 3 s, before starting to record new tactile observations. Each shaking behavior/trajectory (corresponding to a set of action parameter values) needed to be completed before a new set of actions could be performed (to avoid large jerk movements by the hand-arm system). On average, it took more than 500 ms for a specified shaking trajectory to complete. Thus, if our active learning optimization procedure could complete within 500 ms in each step, its effect on the overall runtime would be negligible.

All code for active learning and filtering was implemented in Python, using the SciPy framework[13]. For gradient ascent we used the `fmin_bfgs` function from the SciPy optimization package. Low-level code for robotic hand and arm communication was implemented in C(++). Some routines for control of the robotic arm were provided by Stefan Klanke.

6.3.2.2 Results on 1D action space

First, we examined the effect of active learning for the task, where only shaking frequency was adapted according to expected information gain.

Instead of doing classification over a fixed number of liquids, our model is continuous over the viscosity domain[14] and, therefore, is capable of exhibiting generalization properties. In order to test how well our system generalizes, we introduced a fourth liquid which was a water-glycerine mix with a viscosity of 30 cst (1.47 in log space).

For Gaussian-based updates, the predictive mean squared error—averaged over all four liquids and 3 trials per liquid—for the passive case after 10 time steps was 0.37, compared to 0.07 for the active case. Estimation worked equally well for the liquids used in

[12] `http://www.gatsby.ucl.ac.uk/~snelson/`
[13] `http://www.scipy.org/`
[14] In fact, all random variables in the model are treated as continuous.

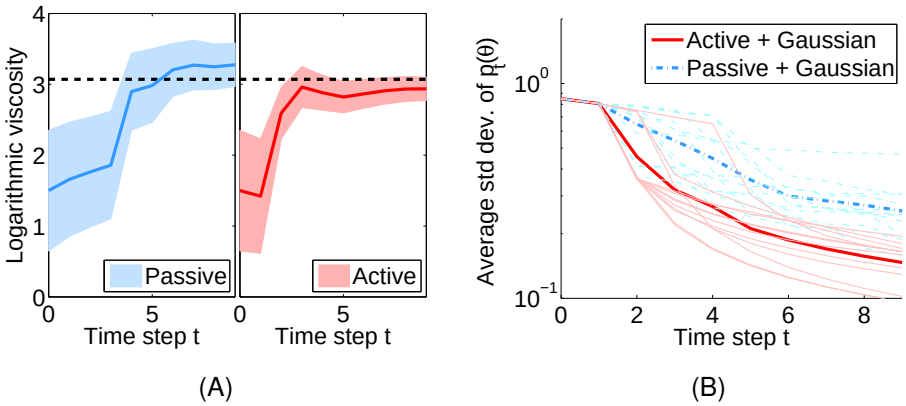

Figure 6.4: **A.** Single trials for passive (left) and active (right) Gaussian belief update of viscosity over all time steps, where the liquid tested on was glycerine (target indicated as dashed line). **B.** Standard deviation of posterior viscosity over time (thick lines denote averages) for passive (blue) and active (red) Gaussian update for all trials over the 4 tested liquids.

training as well for the newly introduced liquid. The posterior standard deviation after 10 steps was 0.26 for the passive case, compared to 0.15 for the active case. Figure 6.4 confirms this result, showing a faster convergence in the active case (see example runs in Fig. 6.4A and posterior variance for all trials in Fig. 6.4B). For the active learning strategy, the execution time of optimization was below 500 ms at each time step in all cases. This ensured that the optimization procedure did not affect overall run-time.

6.3.2.3 Results on 2D action space

In a further set of experiments we tried the active learning paradigm on a more difficult problem, where the actions space was now two-dimensional. Additionally, we wished to obtain more insight into which kind of strategy would be chosen by the information maximization algorithm.

As explained before, we used sparse Gaussian Processes in order to model the observation function for the 2D task. Fig. 6.5 shows slices of the learned GP's mean functions for two different output frequencies (learned from the preprocessed tactile feedback time series data) for both water (left column) and glycerine (right column). It can be seen that most of the power is concentrated at a frequency that is double the shaking frequency (vertical red bands). However, other frequency bands also contribute power

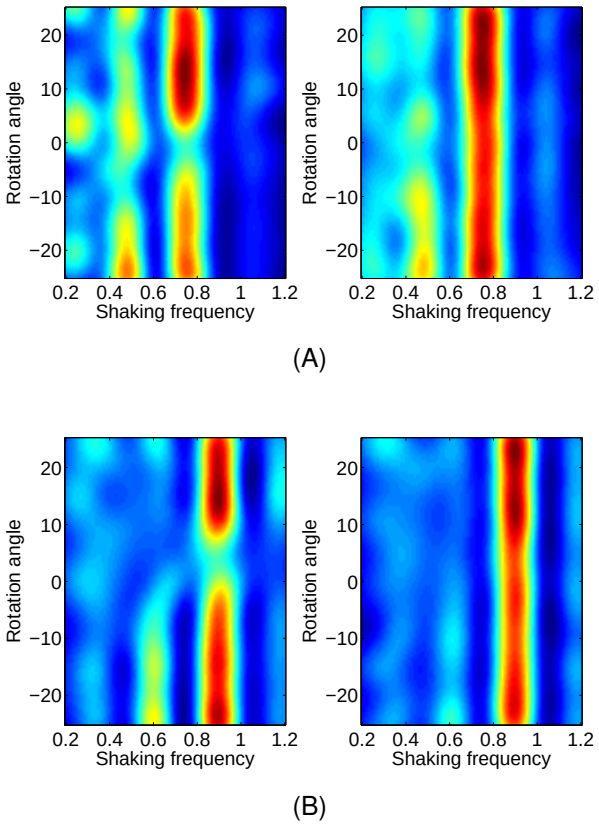

Figure 6.5: Examples of learned GP mean functions over power spectrum amplitudes at different observed frequencies. **A.** Observed power at frequency 1.5 Hz. **B.** Observed power at frequency 1.8 Hz. The left column shows water (low viscosity) and the right one glycerine (high viscosity), respectively.

(yellow bumps throughout plot). The specific relationship between action parameters and observed power spectrum most likely depends on the exact shape of the bottle as well as other parameters.

As can be seen from the plots, at certain shaking angles, the amplitudes between water and glycerine filled bottles differ significantly—at these angles, the two liquids could be distinguished more easily. It should be noted though that the plots only show the mean function and not the corresponding variance, which also affects discriminability of liquids. The variance changes depending on the density of training data in the surrounding space. Moreover, in sparse GPs, some characteristics of heteroscedastic noise can also be modeled by shifting the locations of pseudo-inputs. We did observe that signal-dependent noise in our data set in regions of higher amplitude generally

showed higher variance.

For the 2D task, we compared four different strategies for determining optimal actions at each time step:

i. *Random strategy*: Actions are selected randomly from a uniform distribution over the action space (shaking frequencies and rotation angles from the range that was encountered during training) that was explored in the initial training phase.

ii. *Grid strategy:* A grid is placed to cover the action space uniformly and actions from this grid are executed sequentially in a fixed order. The idea behind this is to make sure that both high and low frequencies, as well as a range of rotation angles, are used.

iii. *Frequency strategy:* The shaking frequency starts low and is increased over time while the bottle is held constant at zero rotation angle.

iv. *Active learning strategy:* The informativeness of each subsequent action is maximized by gradient ascent on the current information landscape.

During run-time, once a new observation has arrived, the next action parameters have to be selected as quickly as possible. We used several strategies to ensure that this happened quickly in the active learning scenario, such that it would not affect the overall runtime of the algorithm. In more detail, we implemented the following measures:

- During gradient ascent, only the terms dependent on $\mathbf{x}$ need to be recalculated (that is, $\mathbf{q}_m^x$, in Eq. (6.14)), thereby speeding up the calculations.

- We use thresholding for the calculations involving the kernel matrix, excluding entries which only have a minor impact during optimization.

- The optimization is run in parallel in several threads from different starting points, and the action from the run with the highest information gain is selected at each step. In our case, we start the optimization of one thread from the action that was chosen for the last time step, and initialize the others at randomly chosen starting positions. We usually ran 3–5 threads in parallel on several networked machines.

- We place an upper limit on the run-time of the optimization at each time step (500 ms) and stop when this upper limit has been reached. Since we perform

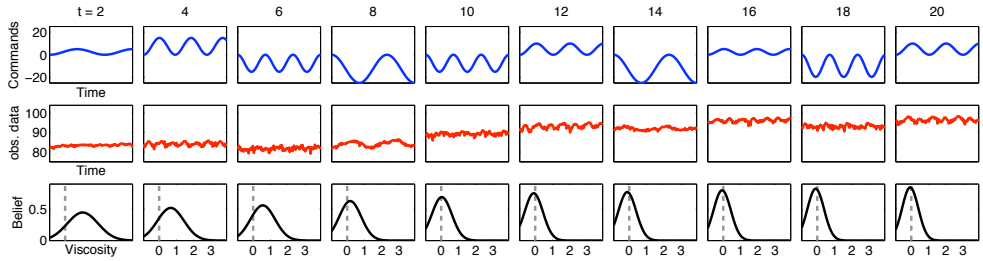

(A)

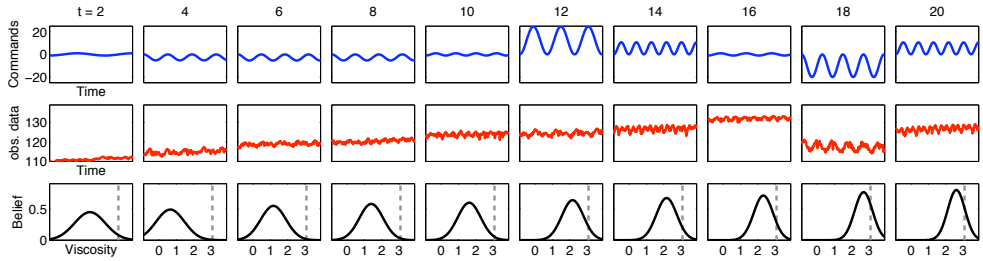

(B)

Figure 6.6: Sample runs with two different liquids. **A.** Water. **B.** Glycerine. Each individual column corresponds to a single time step (lasting 5 s). Shown are the commanded joint angle of the last arm joint (top plot, where the abscissa indicates time and shows commands taken over 5 s), average tactile observations (middle, where the abscissa again indicates time and shows sensor data observed over a period of 5 s) and belief over log viscosity space (bottom, where the x axis represents log viscosity values ranging from -1 to 4, and dashed lines denote the actual viscosity of the tested liquid) at each time step. Note that only every other time step t is shown in the figure due to space constraints. In the both examples, the belief starts out broad and over time converges to the true viscosity of the liquid that was tested.

gradient ascent, this ensures that whenever we stop the optimization, we will have improved on the informativeness of the next action.

- We share the Gaussian Process hyperparameters Γ across all M output dimensions, such that kernel matrices could be re-used for calculations in the different output dimensions and computation time was greatly reduced. This was possible in our case as the length scales and other parameters in the different dimensions all proved to be rather similar. When fitting the Gaussian Processes, we therefore only optimized a single set of hyperparameters and all output dimensions

concurrently.

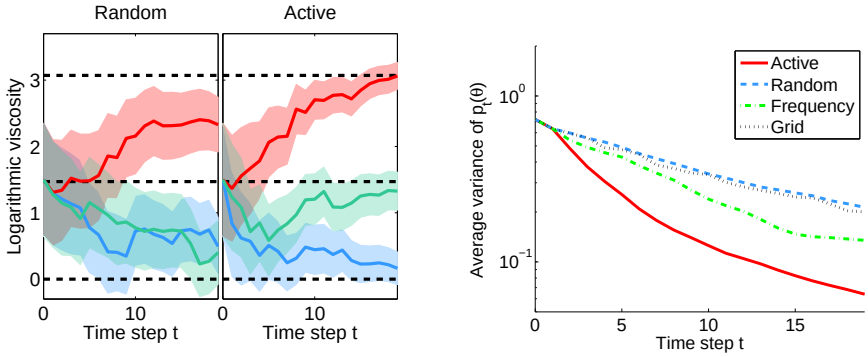

Figure 6.7: **A.** Sample runs for the random (left) and active learning strategies (right), shown over 20 time steps. The posterior over viscosity is shown, over time, for three different liquids: water (blue), a water-glycerine mix (green), and glycerine (red) with the liquids' true viscosities indicated by the dashed lines. Solid lines refer to posterior means, while shaded regions indicate posterior standard deviations. **B.** Posterior variance of viscosity (on log scale) over time for active learning (red), random (blue), frequency (green) and grid (black) strategies.

Two step-by-step sample runs of the active learning strategy are shown in Fig. 6.6, for water and glycerine-filled bottles, respectively. Further examples are shown in the accompanying video (see attached CD-R).

We performed six trials for each of the four liquids (the three liquids from the training set plus the new liquid as described in Sec. 6.3.2.2), using both the random and active learning strategies, and three trials for both the frequency and grid strategies. Fig. 6.7A shows sample test runs for the random (left) and active learning (right) strategies applied to three different viscosities (water, water-glycerine mix, and glycerine). Estimation worked equally well for the liquids used in the training phase as for the newly introduced (water-glycerine mix) liquid.

The mean squared error—averaged over all four liquids and trials—for the active learning strategy after 20 time steps was 0.48. The frequency strategy was close with 0.52, while the other strategies had a mean squared error of 0.72 and greater. Additionally, the posterior variance after 20 steps was considerably lower for the active learning strategy than for any of the other strategies: Fig. 6.7B shows the average posterior variance of θ over time, indicating that the active learning strategy leads to the fastest

convergence. The figure shows that the active strategy is able to reduce uncertainty about the true viscosity in just 7-10 steps to the same level that the other strategies take 20 steps to achieve.

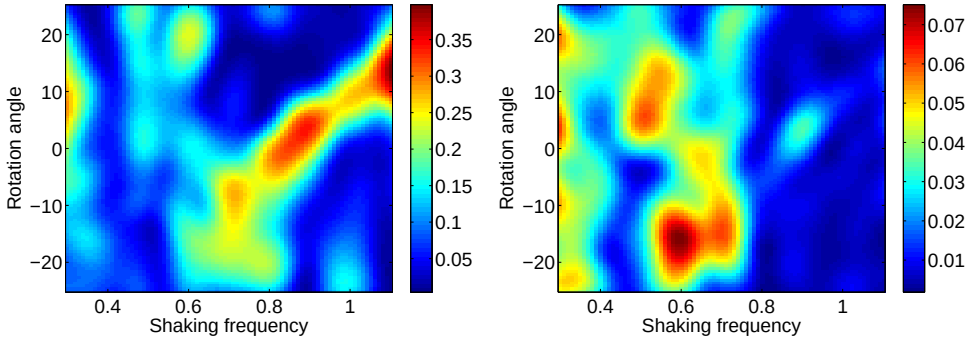

Figure 6.8: Information landscapes for two different probability distributions $p(\theta)$. **A.** broad distribution over viscosity space ($\mu = 1.5$ and $\Sigma = 0.7$). **B.** Narrow distribution in a low viscosity region ($\mu = 0.3$ and $\Sigma = 0.2$). Increasing shaking frequencies are shown on the abscissa, different rotation angles on the ordinate. Blue regions have low information, while yellow and red regions are medium and highly informative, respectively.

We also examined the strategy that was implicitly chosen by the active controller in more detail. Generally, it is not straightforward to come up with information maximizing strategies by hand, as even small changes in the sensor model can result in rather large changes in informativeness of different actions parameters. Moreover, in problems where there is no analytical model of all the effects involved, intuitions about what might be effective strategies can be misleading. Another challenge is that in nonlinear problems, the informativeness of different actions depends on the current belief state, and thus pre-computing the actions ahead of an observation sequence is not possible.

To illustrate this, Fig. 6.8 shows information landscapes for two different prior distributions over the viscosity. As can be seen in Fig. 6.8A, when uncertainty about the true viscosity is high, relatively high shaking frequencies are preferred, along with a positive rotating angle (resulting in an upwards tilted bottle) that increases with shaking frequency. On the other hand, when the distribution over θ is restricted to a narrow low viscosity region, low shaking frequencies become much more informative (as Fig. 6.8B shows). This trend can also be seen in Fig. 6.6, where towards the

end of each run, low shaking frequencies are selected for the water-filled bottle (see Fig. 6.6A), while high shaking frequencies are chosen for the glycerine-filled bottle (see Fig. 6.6B).

Since informative regions of the action space can vary drastically, depending on the current belief, and highly informative regions tend to be sparse, any strategy relying on chance to encounter such informative regions would be expected to fail. Moreover, regions of high information may be located very close to parameter regions containing very little information. As a result, small differences in action values can have a huge effect on how informative resulting observations will be.

6.4 Discussion

In this chapter, we have presented an active framework that exploits a learned sensor model in order to determine dynamics parameters of objects. Our model was derived based on a (sparse) Gaussian Process sensor model. We evaluated our framework on a real robotic hand-arm system that is able to determine the viscosity of liquids based on touch sensor data by shaking bottles at different frequencies and angles. Actions are optimized at each time step based on the expected informativeness of future observations. We demonstrated that the active learning strategy performs better than other simple strategies and is not slowed down by protracted calculations.

6.4.1 Estimation of object dynamics parameters

Previous work on dynamic touch uses fixed end-effector motions and discriminates between objects using object dynamics (Suzuki et al., 2006). In contrast, we focus on the harder problem of regression instead of classification and dynamically select actions that maximize the informativeness of incoming observations.

There exists prior work on dynamic estimation of object properties. As for these problems, the robot has to move, they are inherently active. However, for the most part, control strategies were not optimized in order to maximize the informativeness of incoming samples, as we will see in the following. Early studies have looked into estimating the weight, center of mass, and modes of inertia to loads attached to a robot arm (Atkeson et al., 1985). These quantities can be estimated from rigid body

dynamics equations, and observations were collected by applying specific, parameterized trajectories that were chosen by the experimenters. Other properties like object diameter, weight or stiffness—which are difficult or impossible to determine from static grasps—have also been estimated with dynamic movements. Another study used shaking movements in order to discriminate between objects, but movement parameter values (like shaking frequency) were hand-tuned by experimenters and were reported to affect the quality of object discrimination (Suzuki et al., 2006; Takamuku et al., 2008). An active strategy to select the movement parameters would clearly be beneficial in this case. While some of the studies mentioned above deal with estimation of continuous-valued object parameters, most focus on the task of discrimination within a discrete set of objects.

6.4.2 Alternative approaches

As mentioned in Sec. 5.4.2, reinforcement learning (Paletta and Pinz, 2000) and evolutionary algorithms (Tuci et al., 2009) have been applied to the problem of learning control policies that allow for quick discrimination of objects or parameter discrimination. These methods differ from our approach since we determine optimal actions to be taken in an online manner (i.e., during run-time). While this means that the optimal action to be performed needs to be computed at each time step, our approach makes it easy to add more training data (i.e., to gather more observations) without having to re-learn the policy (i.e., the sequence of actions to be taken). In contrast to past work that optimize actions and perceptions together, we de-couple the optimization (learning the sensor model and optimal actions in separate phases) so that discrimination can be done even when sub-optimal actions are taken.

6.4.3 Possible extensions

An interesting extension of this work would be to incorporate active learning concepts to the collection of the training data (in addition to run-time data). As the dimensionality of the action space increases, the amount of training data samples grows very quickly, and collecting training data can be very time-consuming and involved. One can imagine gathering data only from regions that would contribute to a sufficiently rich training set so that the sensor model can be properly learned in a reasonable amount of time. Finally, as became evident from our evaluation on synthetic data, uni-

modal Gaussian posterior distributions are too restrictive when the actual distribution is multimodal. In the next chapter, we will explore how to represent such distributions with mixtures of Gaussians and evaluate our approach in an active learning setting.

Chapter 7

Representing multimodal distributions for sequential observations

7.1 Introduction

As we have seen in the previous chapter (see Sec. 6.3.1), maintaining a single Gaussian distribution to represent the posterior distribution in filtering problems can result in divergence of the estimate and generally poor performance. This problem can be remedied by relying on Monte-Carlo methods instead, but these can be rather slow when the state space is high-dimensional. This kind of problem is especially important in tactile processing, where each new tactile object contact or exploratory procedure might reveal only relatively little information, which then needs to be integrated with successive data points. In such sequential tasks, the sensory observations have a stochastic dependence on an unobserved underlying state (e.g. object features that one is interested in like pose or dynamical attributes such as size or viscosity) and the goal is to infer the posterior distribution over the state at a particular time step given all observations that have been received up to that time step. In settings where the observation function is nonlinear the posterior state distribution can take on complex shapes and even become multimodal. For example, when trying to determine the pose of an object from just a single or small number of contact points lying on the object's surface, several different potential object configurations often fit the data

(Petrovskaya et al., 2006). In such cases, the posterior distribution might comprise a relatively small number of reasonably well isolated modes, each of which could be modeled well by a single Gaussian distribution. Properly representing posterior state distributions, including their multimodality, is especially important when using active learning methods since the uncertainty captured by the posterior is used directly in the decision of how to query the next observation in order to resolve ambiguities in the state as quickly as possible.

In this chapter, we present a novel approach to the problem of mixture filtering that takes inspiration from variational approaches to approximate inference and combine this with a deterministic sampling approach: We assume that the prior (e.g. the filtering distribution from the previous time step) is given as a mixture of Gaussians (MoG). Due to this MoG prior and the nonlinear observation function the posterior distribution over the state given a new observation can have a complex shape. We therefore approximate the new posterior distribution again as a MoG distribution. We optimize this approximate posterior distribution by approximately minimizing the Kullback-Leibler (KL) divergence between the true updated state distribution and the approximate MoG representation. Exact minimization of the KL divergence is intractable. Our approximate minimization relies on a deterministic sampling approach, Gauss-Hermite Quadrature, which evaluates integrals by evaluating the integrand at suitably chosen sample points, and we describe a novel way to re-formulate the required integrals in order to optimize the accuracy of the method for the problem at hand.

This chapter is based on joint work with Nicolas Heess, who worked on the variational formulation, especially in Secs. 7.2.1 and 7.2.2, and also contributed some Matlab code for running simulations.

7.2 Methods

7.2.1 Problem statement

The general filtering problem is described in Sec. 5.3.2.1. Here, we ignore the time update and instead focus on the measurement update.[1] Similar to the approaches discussed in Sec. 5.3.2.3, we will be representing the prior and the approximate posterior

[1] This corresponds to a static target. Including a dynamics model is straightforward and time updates could be done as in other mixture filters by propagating each mixture component independently.

as a MoG distribution. Furthermore, we assume that the observation likelihood is a Gaussian with fixed covariance, but with a mean that depends on θ via a nonlinear function $f(\cdot)$[2] which we choose to represent as a radial basis function (RBF) network. We choose this form because it allows us to treat terms arising from the likelihood analytically (cf. Sec. 7.2.2 below) and at the same time is general enough to approximate any nonlinear function to arbitrary accuracy (Park and Sandberg, 1991). Alternative formulations, e.g. using a Gaussian process representation are also conceivable (Deisenroth et al., 2009). Moreover, a lot of filtering problems in robotics are complex enough so that mathematical models for observation function are not available and they have to be learned from data instead (see also Chapter 6). In cases where the observation function is given in analytical form, the required expected values can instead be estimated using linearization methods from commonly used unimodal filters, like the extended Kalman filter or the unscented filter (see Sec. 5.3.2.2). Thus, at any time step we are faced with the following problem: Given

$$p(\theta) = \sum_n \gamma_n p_n(\theta) \tag{7.1}$$

$$p_n(\theta) = \mathcal{N}(\theta|\nu_n, \Theta_n) \tag{7.2}$$

$$p(\mathbf{y}|\theta) = \mathcal{N}(\mathbf{y}|f(\theta), \Sigma_y) \tag{7.3}$$

$$f(\theta) = \sum_j c_j k(\theta, \mathbf{m}_j) \tag{7.4}$$

$$k(\theta, \mathbf{m}_j) = \exp\left\{-\frac{1}{2}(\theta - \mathbf{m}_j)^T \mathbf{S}^{-1}(\theta - \mathbf{m}_j)\right\} \tag{7.5}$$

our goal is to approximate the state posterior $p(\theta|\mathbf{y})$ with a mixture of Gaussians

$$q_{\text{mix}}(\theta) = \sum_m \alpha_m q_m(\theta) \tag{7.6}$$

$$q_m(\theta) = \mathcal{N}(\theta|\mu_m, \Sigma_m). \tag{7.7}$$

7.2.2 Fitting a Gaussian mixture to the posterior

Given a new observation $\mathbf{y}$, and the current prior $p(\theta)$ we need to optimize the parameters of $q(\theta)$ (i.e. means, covariance and weights of the individual mixture components) so as to match the actual posterior distribution $p(\theta|\mathbf{y})$ as closely as possible. The variational approach requires the Kullback-Leibler (KL) divergence between the approximate posterior $q(\theta)$ and the true posterior to be minimized in order to arrive at

[2] Throughout this chapter, we treat $f(\cdot)$ as one-dimensional for simplicity, but a multi-dimensional extension is straightforward.

the approximated distribution $q^*(\theta)$:

$$q^* = \arg\min_q \mathrm{KL}\left[q||p\right] \tag{7.8}$$

$$\mathrm{KL}\left[q||p\right] = \int \mathrm{d}\theta\; q(\theta) \log \frac{q(\theta)}{p(\theta|\mathbf{y})} \tag{7.9}$$

In our case $\mathrm{KL}\left[q_{\mathrm{mix}}||p\right]$ can be broken down as follows:

$$\mathrm{KL}\left[q_{\mathrm{mix}}||p\right] = \int \mathrm{d}\theta\; q_{\mathrm{mix}}(\theta) \log \frac{q_{\mathrm{mix}}(\theta)}{p(\theta|\mathbf{y})} \tag{7.10}$$

$$= -\mathrm{H}\left[q_{\mathrm{mix}}\right] - \sum_m \alpha_m \mathrm{E}_{q_m}\left[\log p(\theta)\right]$$

$$- \sum_m \alpha_m \mathrm{E}_{q_m}\left[\log p(\mathbf{y}|\theta)\right] + \mathrm{const} \tag{7.11}$$

where $\mathrm{H}\left[q\right] = -\int \mathrm{d}\theta\; q(\theta)\log q(\theta)$ is the differential entropy of q and the expectations $\mathrm{E}_{q_m}\left[\cdot\right]$ are taken with respect to the Gaussian components q_m of the mixture posterior.

In order to obtain the approximate posterior this expression needs to be minimized with respect to the parameters of q_{mix}. With the choices made above for prior, likelihood, and approximate posterior in Eqs. (7.1)–(7.7), we can exactly compute the third term in Eq. (7.11) (see Appendix A.2) but the first and second term are not tractable since they involve integrals over log-sums. We approximate these intractable integrals in Eq. (7.11) using quadrature methods as described in Sec. 7.2.3.

Mixture distributions as approximate posterior distributions have been considered previously in the literature on variational inference (Jaakola and Jordan, 1998; Lawrence and Azzouzi, 1999; Bouchard and Zoeter, 2009), often with the aim of ensuring bounds on the likelihood. Here, we are interested in a fast method for approximating the relevant integrals, and therefore resort to Gauss-Hermite quadrature, as detailed in the next section. A more thorough threatment of the variational literature can be found in Saal et al. (2011).

7.2.3 Gauss-Hermite Quadrature

Gauss-Hermite quadrature (Abramowitz and Stegun, 1972) approximates d-dimensional integrals by deterministically selecting sample points from a weight function—in this case a Gaussian $\mathcal{N}(\theta|\mu_n,\Sigma_n)$—and then computing a weighted sum of the function values at those sample points:

$$\int \mathrm{d}\theta\; \mathcal{N}(\theta|\mu_m,\Sigma_m) f(\theta) = \mathrm{E}_{q_m}\left[f(\theta)\right] \approx \pi^{-\frac{d}{2}} \sum_{\mathbf{h}} w_{\mathbf{h}} f(\mathbf{z}_{\mathbf{h}}) \tag{7.12}$$

where $w_{\mathbf{h}} = \prod_d w_{\mathbf{h}(d)}$, i.e. the overall weights are determined as the product of the individual single dimension weights. $\mathbf{z_h}$ are the transformed sample points $\mathbf{z_h} = \mathbf{L}_m \theta_{\mathbf{h}} \sqrt{2} + \mu_m$, where $\mathbf{L}_m \mathbf{L}_m^T = {\Sigma_m}^3$. Here, we set $\mathbf{L}_m$ to be the Cholesky factor, but any matrix square root decomposition of Σ_m could be used (cf. Arasaratnam et al., 2007). The sample points and corresponding weights are given by the roots of the Hermite polynomial and can be calculated offline and stored. Derivatives of the resulting approximation are straightforward to calculate. In our setup, the function $f(\cdot)$ is given as an RBF and thus the integral has the following form:

$$\mathrm{E}_{q_m}[f] = \mathrm{E}_{q_m}\left[\log \sum_n \alpha_n \mathcal{N}(\theta|\mu_n, \Sigma_n)\right] \tag{7.13}$$

Gauss-Hermite quadrature could be used to approximate this integral directly (see Goldberger et al., 2003, for such an approach), but we found that the estimate can become highly inaccurate if the variances of the individual mixture components differ considerably and only a small number of sample points is used. This can then lead to a divergence of the optimization of the KL-divergence. Instead of increasing the number of sample points, which quickly becomes untenable in high dimensions, we instead rewrite the integral as a sum, which allows us to approximate each term of this sum individually. This should allow for higher accuracy, as we can optimize the sample points for each term separately. Thus, we write $\mathrm{E}_{q_m}[f]$ as:

$$\mathrm{E}_{q_m}[f] = \mathrm{E}_{q_m}[\log \alpha_1 \mathcal{N}_1] + \sum_{n=2}^{N} \mathrm{E}_{q_m}\left[\log\left(1 + \frac{\alpha_n \mathcal{N}_n}{\sum_{k=1}^{n-1} \alpha_k \mathcal{N}_k}\right)\right] \tag{7.14}$$

where we use the abbreviation $\mathcal{N}_n$ to stand in for the longer $\mathcal{N}(\theta|\mu_n, \Sigma_n)$. The first component of this sum can be calculated analytically, while the remaining ones have to be approximated. We use different weighting functions for each of these terms. For each integral, we approximate it in the standard way as described above, if the variance of $\mathcal{N}_m$ is smaller than the variance of $\mathcal{N}_n$. Otherwise, we rewrite the integral and choose $\mathcal{N}_n$ as the weighting function:

$$\underbrace{\mathrm{E}_{q_m}\left[\log\left(1 + \frac{\alpha_n \mathcal{N}_n}{\sum_{k=1}^{n-1} \alpha_k \mathcal{N}_k}\right)\right]}_{\text{Integral over } q_m} = \underbrace{\mathrm{E}_{q_n}\left[\frac{\mathcal{N}_m}{\mathcal{N}_n} \log\left(1 + \frac{\alpha_n \mathcal{N}_n}{\sum_{k=1}^{n-1} \alpha_k \mathcal{N}_k}\right)\right]}_{\text{Integral over } q_n} \tag{7.15}$$

As illustrated in Fig. 7.1 this makes it more likely that the sample points will capture the region that is interesting for integration, as the mean and variance of the new

[3] One dimensional weights $w_{\mathbf{h}(d)}$ and sample points $\theta_{\mathbf{h}(d)}$ are computed offline or taken from appropriate tables, e.g. `http://www.efunda.com/math/num_integration/findgausshermite.cfm`

weighting function should be closer to the mode and log curvature of the integrand, and thereby improving accuracy (Liu and Pierce, 1994). In the multivariate case, we either pick the component with the lowest covariance determinant, or (when restricting ourselves to diagonal covariance matrices) treat each dimension independently.

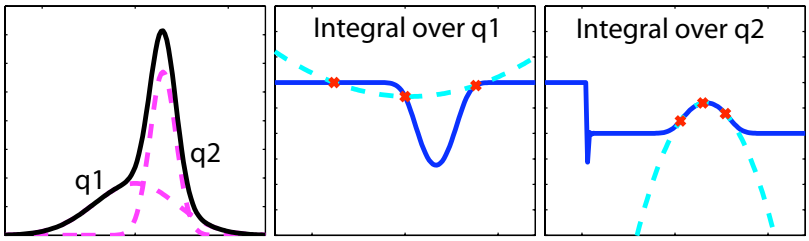

Figure 7.1: Illustration of the improved Gauss-Hermite method. Left: Mixture of two Gaussians, q_1 (left) and q_2 (right). Middle: Sample points from q_1 (red), integrand (dark blue) and implicit polynomial fit to integrand by quadrature method (light blue) when integrating according to the left-hand term in Eq. (7.15). Right: Improved fit in relevant region (around sample points) when integrating according to the right-hand term in Eq. (7.15), with sample points taken from q_2.

7.3 Results

7.3.1 Problem setup

In order to test the performance of our proposed variational approach, we examined how well it could fit complex (i.e. multimodal or skewed) state distributions when compared with other approaches. To highlight the importance of such representations, we additionally tested whether an improved posterior representation would help in a localization task with ambiguous observations, while using active learning in order to speed up convergence.

Depending on task difficulty, we compared our approach with a number of other options: first, a linear mixture filter[4] using the same number of components as the variational mixture filter; second, linear filters using a higher number of components (3^d or 7^d)[5]; and finally, a variational filter using just a single component, in order to examine

[4]Our linear mixture filter implementation is a bank of independent filters that are updated independently. Note that because of the particular form of the nonlinear likelihood (RBF, cf. Eqs. (7.4), (7.5)) all required expectations can be computed analytically (see Appendix A.2).

[5]d denotes the dimensionality of the state distribution.

how well a unimodal approach would perform[6].

The problem was set up as a localization task where the position of a (stationary) target at location θ^* had to be estimated by probing search locations $\mathbf{x}$ sequentially and receiving observations $\mathbf{y}$ that depended on both the probe and target locations: $\mathbf{y} = f(\mathbf{x} - \theta^*)$. Each run started with a wide Gaussian prior (zero-mean with an isotropic variance of 9), reflecting the fact that the location of the target was unknown. The observation function $f(\cdot)$ was set up as an RBF consisting of a small number of individual squared exponential components. In each new run, the location of these components with respect to the target location was sampled from a uniform distribution over a hyper-cube of length 8 centered at the origin. The number and kernel width varied with the dimensionality of the problem: We used 3 kernels with a width of 1.5 for the 4D problem, 4 kernels with width 1.5 for 6D, 5 kernels with width 3 for 8D, and 5 kernels with width 7 for 10D. We used two different types of observations: Ambiguous and Infomax observations. For ambiguous ones, search locations $\mathbf{x}$ were fixed such that observations always came from the mode of an individual RBF component, which frequently resulted in the posterior becoming multimodal. This type of observations was used to test how well the different methods were able to model multimodal state distributions (see results in Sec. 7.3.2). For Infomax observations, the search locations $\mathbf{x}$ were optimized via active learning such that they resulted in the biggest information gain about the position of the target (see Appendix A.3 for mathematical details and Sec. 7.3.3 below for results). These observations should allow the posterior to quickly converge onto the correct target. That is, knowing the observation function as well as the current (possibly multimodal) state distribution allows selecting search locations that disambiguate between the different potential target positions effectively.

For the cases where the linear filter used a higher number of components than the variational one, its prior was initialized to match the original prior as closely as possible by placing components on a grid and adjusting their weights so that the resulting mixture distribution matched the original broad Gaussian prior distribution. We also tried other initializations, e.g. randomly sampling components from the original prior, and found that different initializations did not influence the results much. For the variational approach, we used 3 quadrature points per dimension in all examples, leading to 3^d samples in total. For problems where $d > 2$, we restricted the variational method to

[6]The unimodal variational filter still needs to be optimized iteratively, but since the posterior distribution only consists of a single Gaussian, the KL divergence and its gradient can be evaluated analytically, so no numerical quadrature is needed and optimization is usually much quicker.

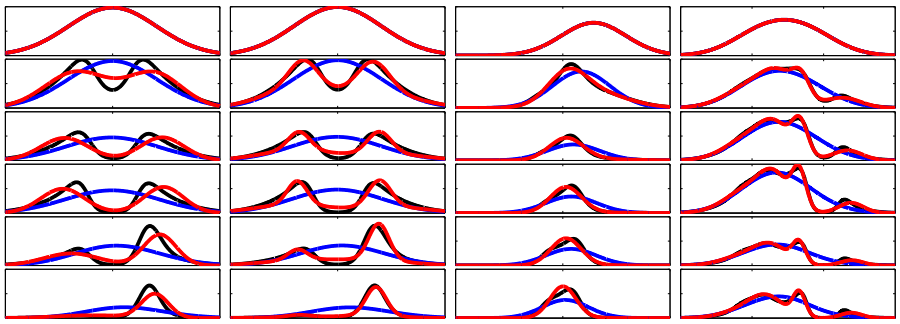

Figure 7.2: Comparison of true posterior (black lines), variational (red) and "bank of filters" (blue) approximations for different observation functions. Each row represents a new observation, with successive time steps ordered from top to bottom, while each of the four columns stands for a different problem: 1. Bimodal posterior, approximated with two components. 2. Same posterior as in (1), approximated with 4 components, 3. Skewed posterior, 2 components, 4. Highly multimodal posterior, 4 components.

diagonal covariance matrices; components in linear filters always maintained full covariance matrices, however. All algorithms were implemented in Matlab, using some functions from the Lightspeed toolbox[7]. For gradient descent we used the scaled conjugate gradient implementation provided by the Netlab toolbox[8].

7.3.2 Representation of multimodal posterior distributions

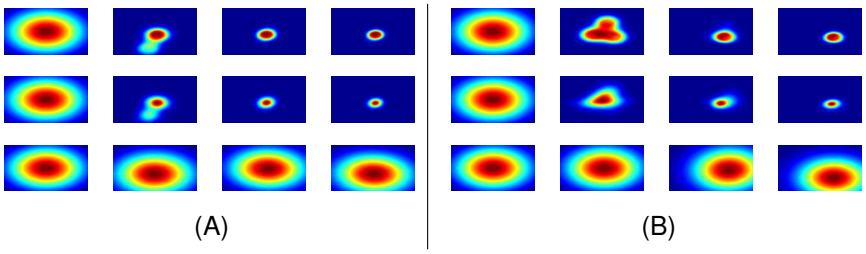

Figure 7.3: **A and B.** Two example runs with time increasing from left to right (4 steps each). Top row: Actual posterior (calculated numerically). Middle row: Variational mixture approximation (3 components). Bottom row: Linear mixture approximations (3 components). All methods start with the same prior (left-most columns) and receive the same observations.

[7] http://research.microsoft.com/en-us/um/people/minka/software/lightspeed/
[8] http://www1.aston.ac.uk/eas/research/groups/ncrg/resources/netlab/

First, we tested how well our approach could represent skewed and multimodal posterior state distributions of different dimensionalities. Some examples for different setups in 1D and 2D are shown in Figs. 7.2 and 7.3. As can be seen, the variational approximation generally approximates the posterior well and correctly finds and represents the major modes. The quality of approximation evidently improves with the number of mixture components that is used (see second column in Fig. 7.2). Moreover, skewed distributions can be fitted well, by using several mixture components (see e.g. the third column in Fig. 7.2). Sometimes the variational approximation covers several posterior modes with a single component, which also happens when there are more posterior modes than mixture components (see right panel in Fig. 7.3). The "bank of filters" method on the other hand runs into problems if a unimodal prior splits into a posterior consisting of several components, and often retains an excessively high variance.

We also ran a more exhaustive test on low-dimensional problems (1–2D), where we kept the observation function constant, but systematically varied the mean and covariance of the prior distribution with respect to the observation function. This introduced a big range of different nonlinearities for the methods to encounter. We ran the algorithms on several different observation functions by varying the number and locations of RBF components. For each individual trial, we (numerically) calculated the KL divergence between the distributions approximated by the mixture methods and the true state posterior. We found that, generally, our algorithm was at least as good as the linear mixture filter but often dramatically better, usually when the posterior state distribution became either considerably skewed or multimodal.

In a further set of tests, we examined how well the different mixture methods were able to capture high-dimensional complex posterior distributions. In order to quantify the fit of the different representations, we calculated differences in the KL divergences between the different mixture approximations and the true posterior distributions over time: we used Monte-Carlo integration in order to arrive at an estimate of the KL divergence (up to a constant). In these tests we only presented ambiguous samples as we were interested in a complex posterior shape. Results for tasks in 4D (with either a single or three observation function modes) and 8D (using 5 modes) can be seen in Fig. 7.4. We ran this task for 100 (4D) or 25 (8D) random configurations of the observation function and target. We noticed several interesting effects. First, a linear filter using the same number of components as the variational filter consistently performs

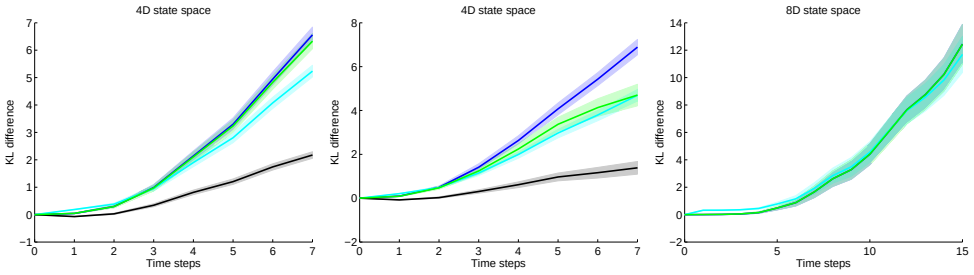

Figure 7.4: Differences between the KL divergence between respective posterior approximation and the true posterior, and the KL divergence between the variational filter and the true posterior as a function of the number of observations. Positive values indicate that the variational mixture filter was closer to the true posterior (in terms of the KL divergence), while negative values denote the respective other filter being a better fit. Left: 4D problem with observation function consisting of a single bump. Middle: 4D problem with three observation function modes. Right: 8D problem with 5 modes. Blue line: Average difference between linear and variational mixture filters (both using 4 mixture components). Green line: Difference between unimodal (1 component) and variational filter). Light blue and black lines (4D only): Same for linear filters using 81 and 2401 components, respectively. Shaded regions indicate standard error of the mean.

worse, independent of whether the posterior state distribution is skewed (left panel) or becomes multimodal (middle panel). Second, adding more components to the linear filter improves the difference in KL divergence. However, even with a very large number of components (2401), performance is generally much worse than the variational mixture filter. Also, the number of components that would be needed quickly becomes untenable in higher dimensions ($d > 4$). For example, using 7 mixture components per dimension would have required more than 5 million components in 8 dimensions, which exceeds the memory limitations of our setup. Finally, a unimodal variational filter cannot represent the complex posterior shapes properly. We noticed that most of the time, the unimodal version tends to cover all of the posterior modes, although in some cases it could "fall" into a single posterior mode, leaving other ones uncovered.

7.3.3 Convergence when using active learning

In another set of simulations, we used active learning in order to quickly resolve the uncertainty introduced by ambiguous observations. For this, we optimized successive search locations with respect to their informativeness about the target location (see Appendix A.3 for mathematical details, and also Chapter 6 for a much more thorough treatment of active learning). This was to highlight the importance of multimodal representations in a practical scenario. In these tests we first presented a number of ambiguous samples (3–5), causing the posterior distribution to acquire a complex shape and possibly become multimodal. We then iterated between optimizing the next search location and updating the state distribution after receiving an observation from that location. This optimization crucially depended on the prior state distribution at the current time step. Using active learning should quickly resolve any ambiguities in the state distributions and lead to quick convergence of the posterior distribution to the actual target. Methods better at representing multimodal posteriors should converge more quickly as they should be better at correctly representing the uncertainty in the state space. In this part of the analysis, we did not examine linear filters with a bigger number of components than the variational filter, as the runtime of our active learning framework is quadratic in the number of Gaussian mixture components (see Appendix A.3), and therefore prohibitively slow to use with a lot of individual components.

We examined how well the different algorithms converged onto the target location for both 6D and 10D problems. Fig. 7.5 shows both the mean squared error (MSE) as well as the root covariance determinant over time[9]. Additionally we plot the log likelihood of the actual target over time. Convergence is indicated by both decreasing error and root covariance determinant, while the log likelihood of the target should increase over time. We found that the variational mixture filter converged well towards the actual target over time, while both the linear and unimodal filters seemed to stall. This means that the active learning component could exploit the multimodal representation of the variational approach and resolve the ambiguity about the target location. The representation of the state distribution by the other methods, however, was not sufficient to allow for effective target localization.

As the variational mixture filter is computationally more involved, it is slower in updating the posterior state distribution after receiving an observation due to the numerical

[9]The MSE and mean root covariance determinant were calculated with respect to the overall mean and covariance of the approximate mixture distribution.

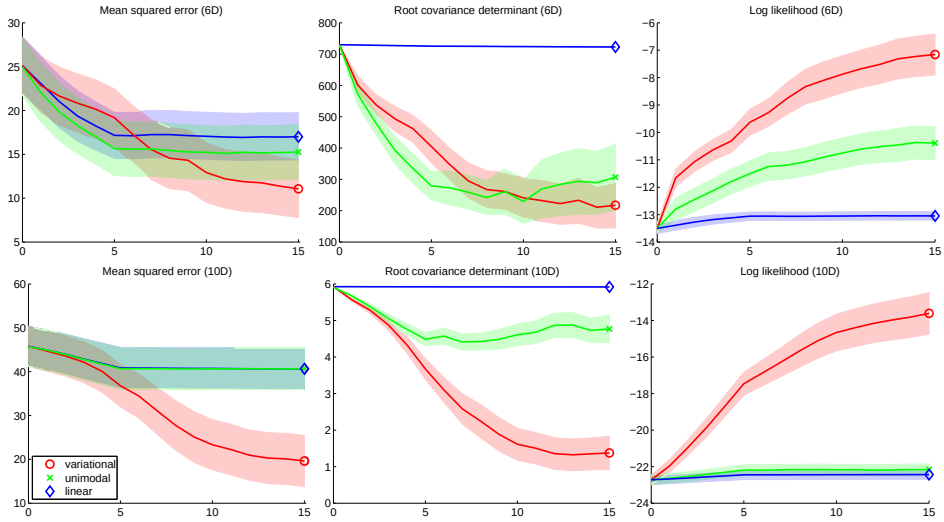

Figure 7.5: Convergence results for different algorithms on a 6D (top row of panels) and 10D (bottom row) problem, respectively, with the linear and variational algorithm using 3 components each. Left panels: Average mean squared error over 15 or 20 time steps, respectively, for 25 runs of the variational (red), linear (blue), and unimodal variational (green) algorithms. Shaded regions indicate standard error of the mean. Both multimodal approaches used 4 mixture components. Middle panels: Average covariance determinant over time. Right panels: Log likelihood of the actual target over time.

optimization involved. We asked whether these delays would have any effects on performance. We therefore set a fixed time span (30–90 seconds), during which each of the methods would iteratively determine the next sample point using active learning, then receive an observation and update its posterior state distribution. The time needed for both optimization of the next search location (i.e. the active learning part) and updating the posterior distributions counted towards each method's time budget. Thus, the faster an algorithm updated its state representation and the lower its number of mixture components, the more observations it was able to request. Fig. 7.6 plots the number of observations that was used by each method against the log likelihood at the end of the time span for a 3D (left) and a 6D problem (right). In this plot, a marker in the left, upper corner indicates that the respective method was only able to request a small number of observations, but achieving a high log likelihood. Markers in the right, lower corner, however, would indicate that an algorithm requested a high number of observations but failed to increase the log likelihood considerably. We found that

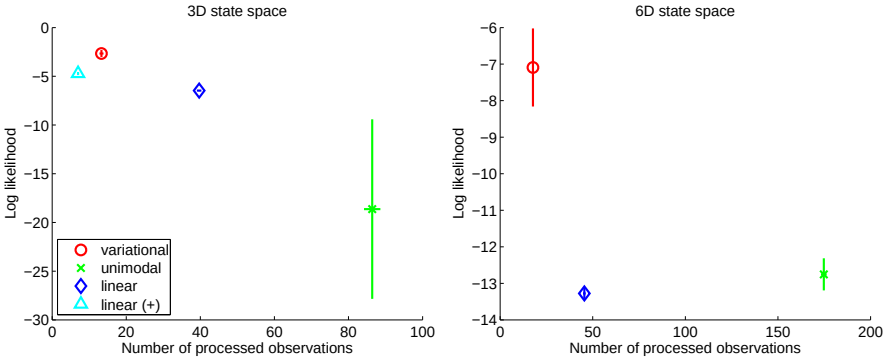

Figure 7.6: Log likelihood of target after updates plotted against number of processed observations. Markers denote mean (over 25 runs each) and vertical lines indicate standard errors of the mean. Left: 3D state space. Right: 6D state space. A fixed time budget of 30 and 90 seconds, respectively, was imposed per algorithm. Light blue marker (left plot only): Linear filter with 27 components. Red: Variational mixture filter (3 components). Blue: Linear filter (3 components). Green: Unimodal variational filter.

the variational mixture filter does well in both tasks, by increasing the log likelihood more than other methods, despite being relatively slow and therefore only processing a small number of observations. The linear filter with 27 components also does well in the 3D example, but performs even slower due to increased computational demands in the active learning stage. The variational unimodal filter is generally the fastest, but does not perform well. The extremely high standard error observed is due to its mode-seeking behavior, which caused it to model only a single posterior mode, which often turned out to be the "wrong" one.

7.4 Discussion

In this chapter we have proposed a novel approach to nonlinear filtering in which the approximate posterior distribution over the state is maintained as a mixture of Gaussians. Using a MoG to approximately represent the posterior makes it possible to capture complex shapes of the true state distribution such as skewness or multimodality which often arise when the observation function is nonlinear. Unlike previous approaches to mixture filtering we do not maintain a set of independent Gaussian components but take interactions between the mixture components into account when optimizing the approximate representation of the posterior distribution given a new

observation. This requires the calculation of expectations over log-sums, which cannot be done analytically, and we propose to approximate these terms using quadrature methods. We find that optimizing the mixture representation directly captures the true shape of the posterior much better (in terms of the KL divergence) than a bank of independent linear filters, even when allowing many more components. We demonstrate the impact of this improved representation in a localization task where active learning is used in order to directly resolve the uncertainty in the posterior distribution resulting from ambiguous samples: The proposed approach converges considerably faster and more reliably than the alternative filtering approaches (which frequently get stuck at noticeable distance from the target). Importantly, faster convergence is achieved not just when measured as a function of observations but also in terms of overall compute time, despite the fact that our filtering approach is computationally more expensive than the alternatives: The additional computational complexity in processing individual observations by the filter is more than compensated for by the noticeably faster convergence per observation as demonstrated in the experiments with limited overall run-time.

7.4.1 Alternative approaches

We have focused on a mixture of Gaussians representation in order to capture multimodal or skewed posterior distributions. Sampling methods like the particle filter have been proposed as an alternative to traditional filtering methods when dealing with nonlinear observation functions and the resulting multimodal state distributions (see Sec. 5.3.2.4). While a sample based representation can in principle approximate distributions of any shape (including multimodal distributions), the number of particles required increases exponentially so that they become impractical in high-dimensional spaces. Furthermore, for many applications a compact representation of the posterior e.g. in terms of a small number of mixture components is crucial: For instance, in the context of the application discussed in this paper, active learning, the sample-based methods that have been proposed so far are very slow even with a relatively small number of particles, which renders them feasible in very low-dimensional spaces only (Hoffmann et al., 2006, and see also Chapter 6 in this thesis).

7.4.2 Possible extensions

There are several interesting directions for future work. Firstly, our method uses an observation function that is represented as an RBF network. A relatively straightforward extension would be to use a Gaussian Process representation instead (Deisenroth et al., 2009). In cases where the observation function is not learned from data, but its analytic form is given directly, it might not be possible to calculate expectations over this function. In such cases, expectations can be estimated using methods from unimodal Gaussian filters, such as linearization by computing derivatives, as in the extended Kalman filter, or statistical linearization, as in the unscented filter (Julier and Uhlmann, 2004).

Secondly, an important practical consideration is the number of mixture components used to represent the posterior. In the experiments described above this number was fixed from the start of the algorithm. Using too many components does not negatively impact the quality of the posterior representation, but it does slow down the algorithm. Using as few components as possible is therefore desirable in order to achieve fast convergence in terms of compute time. On the other hand, choosing the number of components too small will leave some part of the state space unrepresented or requires a single component to cover multiple posterior modes, which will result in a worse representation and in the extreme case can lead to similar problems as for a unimodal filtering approach. Choosing the number of components optimally is an open question so far. One interesting approach would be to dynamically adjust the number of components, adding new components in each observation step, and then merging the most similar ones.

7.4.3 Active learning

In this chapter we have focused on developing a state representation that is rich enough to capture multimodality and other complex shapes in an accurate way. One of the examples that we used to highlight the importance of such representations is the use of active learning in order to resolve ambiguous information. For such problems it is important to properly represent uncertainty, so that active learning can then be used to reduce any remaining ambiguity as quickly as possible.

Chapter 8

Conclusion

8.1 Summary

In this thesis, we have applied information theoretic measures to tactile processing, considering both the human and robotic tactile systems. In Chapters 3 and 4 we analyzed neural responses from primary afferents in the human peripheral nervous system with respect to the information content they provide about tactile features important for grasping. To this end, we presented an in-depth analysis of two coding schemes—spike counts and first spike latencies—in Chapter 3. In Chapter 4 we then considered a more natural experimental setting and analyzed what role fingertip mechanics play in coding stimulus properties. In both chapters, we estimated the mutual information about object/grasp properties that is carried by spike trains, whenever possible directly and in other cases by relying on lower bounds. In these chapters, we were mainly interested in the information capacity of the human tactile system and the coding schemes involved in relaying information gathered by mechanoreceptors to the brain as quickly as possible. In our robotics experiments we then expanded our view and considered the sense of touch as an "active" modality. Most feedback will be gathered during active manipulation, grasping, or exploration of objects and therefore the "quality" of this data will depend on the exact movements that are executed by a robot. We phrased this problem in the context of active learning: How to select movement parameters such that the information about a parameter of interest is maximized? In Chapter 6, we developed our own approach to this problem and presented results obtained on a robotic hand-arm setup. Finally, in Chapter 7, we extended our representation to deal with complex and multimodal probability distributions, which are especially common

when processing tactile data.

8.2 Closing the circle

In what way can the examination of the human tactile systems help with solving robotic and engineering questions? How can machine learning and robotics approaches provide insight into the principles underlying information processing in the human brain? In the following, we will explore some fruitful connections between neuroscience and machine learning / robotics.

8.2.1 From neuroscience to robotics...

One straightforward approach where neuroscientific results can be transferred into the robotics domain is the examination of mechanical properties of the fingertip along with mechanoreceptor placement and function in order to mimic properties that are advantageous for sensing or grasping capabilities. For example, the soft coating that is often used on robotic fingertips helps with object stabilization during grasping and prevents slippage, just as the soft fingertip tissue does in humans. Another example comes from a recent study that argues that fingerprints cause skin vibrations when in contact with textures and the range and depth of the frequencies that are elicited seem tuned for being picked up by Pacinian corpuscles located within deeper areas of the human skin (Scheibert et al., 2009). Designs inspired by this principle might be of potential use in robotic tactile sensor development. For example, one recent application that has been developed trying to replicate this function is the BioTac sensor by SynTouch[1] (see also Wettels et al., 2008). Its goal is to use insights from the human tactile system to develop a sensorized artificial fingertip that comes close to the functional properties of the human one (e.g. in texture recognition or temperature sensing capabilities). This proves a good example where biological insights have contributed directly to a state-of-the-art robotics sensing device.

One area of research that is important for these kinds of developments is examining what kinds of external object properties are picked up during grasping and relayed to the brain. This might provide hints into what kinds of features might be important for robotic touch sensors to pick up in order to aid robots perform human-like tasks.

[1] http://www.syntouchllc.com

For example, the development of robotic touch sensors puts considerable emphasis on increasing the spatial resolution, yet a high spatial resolution might actually not be that important for most haptic tasks that humans can perform with ease (with the two-point discrimination threshold being around 2mm on the human fingertip). Increasing the tactile resolution of robotic sensors, which is still an important selling point, might therefore not be of that much help in improving a robot's haptic abilities. Contrarily, numerous object and contact features are transmitted by human tactile afferents that are currently very hard for tactile sensors to pick up: object curvature and the angle of force impact, as we examined in Chapter 3 are only two examples. Other studies have looked at more complex features like tangential torque, which arises from picking up an object and then re-aligning the hand position with respect to the ground (Birznieks et al., 2010). Importantly, such features and tactile feedback in general have been shown to be critical for grasp stability (e.g. Jenmalm et al., 2000; Augurelle et al., 2003).

An additional important connection between somatosensory neuroscience and robotics is the field of neuroprosthetics. Currently, amputees who are fitted with prosthetic limbs do not receive any (or only very limited) feedback from these devices, which makes it difficult for them to correctly scale grasp forces, for example. One the goals of current research into this area is to equip prosthetic limbs with touch sensors and then relay any data that is gathered by these directly to the brain via connections between the prosthetic hardware and any remaining peripheral afferents. Other approaches try to target the brain directly by implanting electrode arrays into the somatosensory cortex. Knowledge about what kinds of tactile features are picked up and how they are encoded are therefore important to ensure that these artifically generated signals come as close as possible to the kind of feedback that the intact tactile system would provide (see e.g. Russell et al., 2009, for a model that maps fingertip indentation to neural responses).

8.2.2 ...and vice versa

In Chapter 6 we were concerned with strategies for robotic movement control that allow for maximizing the amount of information that is gathered about a parameter of interest. Importantly, simply imitating human strategies would not have worked, as those are (presumably) optimised for the human tactile system, which differs considerably from robotic touch sensors. Rather, we considered a framework that is general enough so as to work with arbitrary sensor models and is therefore independent of the

specific hardware that is used.

We can now ask the question whether humans use similar approaches that maximize informativeness in haptic tasks, or whether they rely on simpler but less effective strategies. A number of studies have looked into what movements were selected during exploration of objects in classification tasks. The importance of the finger and hand movements for haptic object recognition was first emphasized by Gibson (1962), who conjectured that appropriate movements are selected to increase the information one obtains about the object. The first psychophysical evidence was brought forward by Klatzky and Lederman who showed that when humans try to recognize objects by haptic interaction alone, they tend to select stereotyped movements from a set of exploratory actions (see Lederman and Klatzky, 1993, for a review). Among other things they showed that these exploratory actions differ in the amount of information they provide about different object features (e.g. softness, texture) and that some of them conveyed information about a broader feature range than others. Specifically, they found that during early stages of object recognitions, exploratory actions that could deliver information about a broad spectrum of tactile features were favoured, while later on more specific exploratory actions were chosen (Lederman and Klatzky, 1990). Finally, it seems that these actions are selected specifically in order to minimize posterior variance (i.e. uncertainty) about object class membership (Lederman and Klatzky, 1987). More recently, it has been shown that fingertip forces are adapted during tactile exploration depending on what type of target that has to be located (Smith et al., 2002). How such exploratory actions are learned, and whether they are optimized during exploration or whether action sequences useful for feature detection are learned is still an open question. Likewise, computational models for how such tasks could be achieved in neural systems are also scarce or non-existent.

The question whether humans use active learning strategies has also been investigated in a number of other modalities and systems and some computational modelling has been done in some of those. First, and closely related to haptics, active learning has also received some attention in the human motor control community. The question is whether humans select actions that can be considered optimal when learning a new motor skill. It has been found that while those participants that selected their training actions according to what an optimal active learner would have done were better than participants who did not follow such a strategy, there was no evidence that on average humans are more inclined to follow the optimal strategy than any other one

(Huang et al., 2008). Second, the idea that actions are chosen so as to maximize incoming information has also been applied to the visual system: Some studies took the approach that instead of specifying specific features that determine the saliency of an image region (see Itti and Koch, 2001, for a review), one can define saliency in terms of the self-information that an image patch provides, i.e. a measure of how surprising a given area is given the rest of the image statistics (Bruce and Tsotsos, 2005; Lee and Yu, 2000). This would explain fixation points being at locations that are highly informative about the underlying image. This idea is also corroborated by other studies showing that in a visual search task, fixation points are chosen according to how informative they are about the location of the target (Najemnik and Geisler, 2005, 2008). Finally, there has also been work trying to determine whether humans maximize informativeness when learning about new concepts that belong to different classes. In such cases it is an advantageous strategy to request help for individual cases that are close to border between different concepts. For example in one study, participants had to learn different classes objects that differed in how "spiky" objects belonging to different classes looked. They found that humans perform better when they can choose where to sample than in a conditions where samples are presented to them randomly. However, they did not achieve the theoretical optimal performance (Castro et al., 2008).

In summary, while there have been some approaches that use active learning in order to ask questions about optimality of human behaviour, studies in the field of haptics are still lacking. More experimental results are needed which together with computational concepts from machine learning should provide more insight into this problem.

8.3 Open problems

There remains a large number of open problems and unanswered questions related to tactile processing and the wider area of haptic interaction. In the following, we will name a few which are particularly relevant to the issues discussed in this thesis.

First, while there is a multitude of tactile features that are coded for in peripheral afferents, it is still not clear how these are integrated and processed in cortex. Recent studies have identified features present in cortex that do not seem to be available from single peripheral afferents (Yau et al., 2009), but how these arise is far less understood than, for example, in the visual system. Multi-unit recordings together with information-

theoretic methods allow for examining correlations (e.g. Reed et al., 2008) between individual neurons and analyzing different coding schemes and their interplay, as has been done in the visual system (e.g. Montani et al., 2007; Panzeri et al., 2010). So far, there do not seem to be any information-theoretic analyses of recordings from somatosensory cortex in monkeys for example, even though there is a multitude of such studies from the barrel cortex in rats (e.g. Arabzadeh et al., 2006; Panzeri et al., 2001).

Another interesting problem is that of sensorimotor interaction, i.e. how the tactile features that are extracted during object exploration and grasping influence motor actions and vice versa. Such interactions presumably take place on several levels along the sensory/motor hierarchy. For example, it has been shown that the motor cortex can directly modulate neural activity in the rat thalamus (Lee et al., 2008).

A related issue is the active learning problem that we discussed in Chapter 6. As mentioned in the previous section, there are currently no computational approaches towards human object exploration and classification in terms of active learning.

On the robotics side, active learning is one aspect that must be mastered on the road to autonomous learning and task-driven systems. A notable problem is how to solve the active learning problem, when the model is only partly known and has to be simultaneously estimated from data that is collected online: this is also known as the exploration/exploitation problem. Tractable algorithms that are useful in robotics are currently being developed (e.g. Simpkins et al., 2008). This problem is also looked at in the context of reinforcement learning and Markov decision processes (e.g. Doshi et al., 2008). Finally, estimating a certain relevant parameter is often only part of a bigger goal, for example when this information is needed to perform a specific task, a problem that is often called dual control (see Filatov and Unbehauen, 2000, for a review from a control theory perspective). Developing tractable algorithms for these problems that work in high-dimensional spaces will be a major challenge in the future.

Such advances will also need more powerful models for data representation and integration. Recently, a variety of hybrid models (e.g. Sudderth et al., 2010) have been developed that combine aspects of Monte-Carlo approaches and more traditional (semi-) parametric models. Another important problem, especially in robotics, is that often spaces are non-Euclidean. This might be caused, for example, by joint limits in the robot's kinematics. Approaches that work in generalized spaces, e.g. reproducing kernel Hilbert spaces (RKHS), can be a big benefit in such situations (see Song et al., 2009, for an RKHS approach to the filtering problem).

Tackling these open questions will provide further insights into how human tactile processing is organized and will further guide engineering efforts to improve robotic and neuroprosthetic devices. Similarly, more advanced mathematical models should open up bigger areas of neuroscience for computational analysis.

Appendix A

Additional active learning and filtering equations

A.1 Expansion of terms for the Gaussian active update

The terms $\mathbf{q}^{\theta}_m$, $\mathbf{q}^{x}_m$, $\mathbf{Q}^{\theta}_{mn}$, and $\mathbf{Z}_m$, as introduced in Eqs. (6.11)–(6.13) in Sec. 6.2.3, can be expanded as follows:

$$(\mathbf{q}^{\theta}_m)_i = \alpha_m^2 |\Sigma_m(\mathbf{H}^{\theta}_m)^{-1}+\mathbf{I}|^{\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mu-\tau^i)^T(\Sigma_m+\mathbf{H}^{\theta}_m)^{-1}(\mu-\tau^i)\right\} \tag{A.1}$$

$$(\mathbf{q}^{x}_m)_i = \exp\left\{-\frac{1}{2}(\mathbf{x}-\chi^i)^T(\mathbf{H}^{\mathbf{x}}_m)^{-1}(\mathbf{x}-\chi^i)\right\} \tag{A.2}$$

$$\begin{aligned}(\mathbf{Q}^{\theta}_{mn})_{ij} = \alpha_m^2\alpha_n^2 |\left\{(\mathbf{H}^{\theta}_m)^{-1}+(\mathbf{H}^{\theta}_n)^{-1}\right\}\Sigma+\mathbf{I}|^{\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mathbf{z}^{mj}-\mathbf{z}^{nj})^T\mathbf{U}^{-1}_{mn}(\mathbf{z}^{mi}-\mathbf{z}^{nj})\right\} \\ \exp\left\{-\frac{1}{2}(\tau^i-\tau^j)^T(\mathbf{H}^{\theta}_m+\mathbf{H}^{\theta}_n)^{-1}(\tau^i-\tau^j)\right\}\end{aligned} \tag{A.3}$$

$$(\mathbf{Z}_m)_{\bullet i} = (\mathbf{q}_m)_i\left\{\Sigma^{-1}+(\mathbf{H}^{\theta}_m)^{-1}\right\}^{-1}(\Sigma^{-1}\mu+(\mathbf{H}^{\theta}_m)^{-1}\tau^i) \tag{A.4}$$

where $\mathbf{z}^{mi} = \mathbf{H}^{\theta}_m(\tau^i-\mu)$, $\mathbf{R}_{mn} = (\mathbf{H}^{\theta}_m+\mathbf{H}^{\theta}_n)^{-1}+\Sigma$ and $\mathbf{U}_{mn} = (\mathbf{H}^{\theta}_m+\mathbf{H}^{\theta}_n)^{-1}\mathbf{R}^{-1}\Sigma(\mathbf{H}^{\theta}_m+\mathbf{H}^{\theta}_n)^{-1}$. $(\mathbf{Z}_m)_{\bullet i}$ denotes the i-th column vector of matrix $\mathbf{Z}_m$.

The gradients are then calculated as below:

$$\frac{\partial(\mathbf{q}^{x}_m)_i}{\partial\mathbf{x}} = (\mathbf{q}_m)_i(\mathbf{H}^{\mathbf{x}}_m)^{-1}(\mathbf{x}-\chi^i) \tag{A.5}$$

$$\frac{\partial(\mathbf{Z}_m)_{ij}}{\partial\mathbf{x}} = \frac{\partial(\mathbf{q}_m)_i(\mathbf{x})}{\partial\mathbf{x}}(\Sigma^{-1}+(\mathbf{H}^{\theta}_m)^{-1})^{-1}(\Sigma^{-1}\mu+(\mathbf{H}^{\theta}_m)^{-1}\tau) \tag{A.6}$$

A.2 Expectations over RBF networks

In Chapter 7, we are faced with computing expectations over the observation function $f(\cdot)$, both in our approach and in the linear mixture filter. Usually, expectations over the nonlinear observation function $f(\cdot)$ are intractable and have to be approximated, however if $f(\cdot)$ is represented as a RBF network as in our case, or as a Gaussian Process (Deisenroth et al., 2009), then these expectations can be calculated analytically (e.g. Girard et al., 2003) as follows:

$$\mathrm{E}_{q_{mix}}[f(\mathbf{x})] = \sum_m \sum_j \mathrm{E}_{q_m}[c_j k(\mathbf{x}, \mathbf{m}_j)] \tag{A.7}$$

$$\mathrm{E}_{q_m}[c_j k(\mathbf{x}, \mathbf{m}_j)] = c_j |\mathbf{S}^{-1}\Sigma_m + \mathbf{I}|^{-\frac{1}{2}} \cdot \tag{A.8}$$

$$\exp\left\{-0.5(\mu_m - \mathbf{m}_j)^T (\Sigma_m + \mathbf{S})^{-1}(\mu_m - \mathbf{m}_j)\right\} \tag{A.9}$$

$$\mathrm{E}_{q_{mix}}\left[f(\mathbf{x})^2\right] = \sum_m \sum_i \sum_j \mathrm{E}_{q_m}[c_i k(\mathbf{x}, \mathbf{m}_i) c_j k(\mathbf{x}, \mathbf{m}_j)] \tag{A.10}$$

$$\mathrm{E}_{q_m}[c_i k(\mathbf{x}, \mathbf{m}_i) c_j k(\mathbf{x}, \mathbf{m}_j)] = c_i c_j |2\mathbf{S}^{-1}\Sigma_m + \mathbf{I}|^{-\frac{1}{2}} \cdot \tag{A.11}$$

$$\exp\left\{-0.5(\mu_m - \hat{\mathbf{m}}_{ij})^T (\Sigma_m + 0.5\mathbf{S})^{-1}(\mu_m - \hat{\mathbf{m}}_{ij})\right\} \tag{A.12}$$

where $\hat{\mathbf{m}}_{ij} = \frac{\mathbf{m}_i + \mathbf{m}_j}{2}$.

A.3 Quadratic mutual information for MoGs

In the active learning scenario described in Sec. 7.3.1, our aim is to choose a search location θ^*, which maximizes the mutual information between the current state distribution and the expected observation. As the mutual information cannot be calculated analytically when distributions are represented as mixtures of Gaussians, we instead optimize a surrogate measure, called quadratic mutual information (QMI) that has been originally proposed for clustering (Torkkola, 2003):

$$I_{QMI}(X;Y|\Theta) = \iint \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}\,(p(\mathbf{x},\mathbf{y}|\theta) - p(\mathbf{x})p(\mathbf{y}|\theta))^2 \tag{A.13}$$

$$= \iint \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}\; p(\mathbf{x},\mathbf{y}|\theta)^2 + \iint \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}\; p(\mathbf{x})^2 p(\mathbf{y}|\theta)^2 - 2\iint \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}\; p(\mathbf{x},\mathbf{y}|\theta)p(\mathbf{x})p(\mathbf{y}|\theta) \tag{A.14}$$

Each of the integrals involved can now be calculated analytically in a similar fashion as described in Appendix A.2. At each step, I_{QMI} is optimized by gradient ascent with

respect to the new search location θ. The computational complexity of this approach is quadratic in the number of mixture components. For a more thorough treatment of the active learning problem, see Chapter 6.

Bibliography

Abramowitz, M. and Stegun, I. A. (1972). *Handbook of mathematical functions with formulas, graphs, and mathematical tables*. U.S. Dept. of Commerce.

Adrian, E. D. (1928). *The basis of sensation*. New York: W. W. Norton.

Allen, P. K. and Michelman, P. (1990). Acquisition and interpretation of 3-D sensor data from touch. *IEEE Trans Robot Autom*, 6(4):397–404.

Alspach, D. and Sorenson, H. (1972). Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Trans Autom Control*, 17(4):439–448.

Arabzadeh, E., Panzeri, S., and Diamond, M. E. (2006). Deciphering the spike train of a sensory neuron: Counts and temporal patterns in the rat whisker pathway. *J. Neurosci.*, 26(36):9216–9226.

Arasaratnam, I., Haykin, S., and Elliott, R. J. (2007). Discrete-time nonlinear filtering algorithms using Gauss-Hermite quadrature. *Proc. IEEE*, 95(5):953–977.

Aronov, D., Reich, D. S., Mechler, F., and Victor, J. D. (2003). Neural coding of spatial phase in V1 of macaque monkey. *J Neurophysiol*, 89(6):3304–27.

Ascari, L., Bertocchi, U., Corradi, P., Laschi, C., and Dario, P. (2009). Bio-inspired grasp control in a robotic hand with massive sensorial input. *Biol Cybern*, 100(2):109–28.

Atkeson, C. G., An, C. H., and Hollerbach, J. M. (1985). Rigid body load identification for manipulators. In *Proc. 24th Conf. on Decision and Control*, volume 24, pages 996–1002.

Augurelle, A. S., Smith, A. M., Lejeune, T., and Thonnard, J. L. (2003). Importance of cutaneous feedback in maintaining a secure grip during manipulation of hand-held objects. *J Neurophysiol*, 89(2):665–71.

Beitel, R. E., Gibson, J. M., and Welker, W. I. (1977). Functional development of mechanoreceptive neurons innervating the glabrous skin in postnatal kittens. *Brain Res*, 129(2):213–26.

Bell, B. and Cathey, F. (1993). The iterated Kalman filter update as a Gauss-Newton method. *IEEE Trans Autom Control*, 38(2):294–297.

Bensmaia, S. J., Craig, J. C., Yoshioka, T., and Johnson, K. O. (2006). SA1 and RA afferent responses to static and vibrating gratings. *J Neurophysiol*, 95(3):1771–82.

Bensmaia, S. J., Denchev, P. V., Dammann, J. F., Craig, J. C., and Hsiao, S. S. (2008). The representation of stimulus orientation in the early stages of somatosensory processing. *J Neurosci*, 28(3):776–86.

Birznieks, I., Jenmalm, P., Goodwin, A. W., and Johansson, R. S. (2001). Encoding of direction of fingertip forces by human tactile afferents. *J Neurosci*, 21(20):8222–8237.

Birznieks, I., Macefield, V. G., Westling, G., and Johansson, R. S. (2009). Slowly adapting mechanoreceptors in the borders of the human fingernail encode fingertip forces. *J Neurosci*, 29(29):9370–9.

Birznieks, I., Wheat, H. E., Redmond, S. J., Salo, L. M., Lovell, N. H., and Goodwin, A. W. (2010). Encoding of tangential torque in responses of tactile afferent fibres innervating the fingerpad of the monkey. *J Physiol (Lond)*, 588(Pt 7):1057–1072.

Bisley, J. W., Goodwin, A. W., and Wheat, H. E. (2000). Slowly adapting type I afferents from the sides and end of the finger respond to stimuli on the center of the fingerpad. *J Neurophysiol*, 84(1):57–64.

Blake, D. T., Johnson, K. O., and Hsiao, S. S. (1997). Monkey cutaenous SAI and RA responses to raised and depressed scanned patterns: effects of width, height, orientation, and a raised surround. *J Neurophysiol*, 78(5):2503–17.

Bohg, J., Johnson-Roberson, M., Björkman, M., and Kragic, D. (2010). Strategies for multi-modal scene exploration. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 4509–15.

Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121.

Borst, C., Fischer, M., Haidacher, S., Liu, H., and Hirzinger, G. (2003). DLR hand II: experiments and experience with an anthropomorphic hand. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 702–707.

Bouchard, G. and Zoeter, O. (2009). Split variational inference. In *Proc. Int. Conf. on Machine Learning (ICML)*, pages 57–64.

Bruce, N. and Tsotsos, J. (2005). Saliency based on information maximization. In *Proc. Advances in Neural Information Processing Systems (NIPS)*.

Callari, F. G. and Ferrie, F. P. (2011). Active object recognition: Looking for differences. *Int J Comput Vision*, 43(3):189–204.

Castro, R., Kalish, C., Nowak, R., Qian, R., Rogers, T., and Zhu, X. (2008). Human active learning. In *Proc. Advances in Neural Information Processing Systems (NIPS)*.

Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Stat Sci*, 10(3):273–304.

Chaudhuri, P. and Mykland, P. (1993). Nonlinear experiments: Optimal design and inference based on likelihood. *J Am Stat Assoc*, 88(422):538–546.

Cohn, D. A., Ghahramani, Z., and Jordan, M. I. (1996). Active learning with statistical models. *J Artif Intell Res*, 4:129–145.

Craig, J. C. and Rollman, G. B. (1999). Somesthesis. *Annu Rev Psychol*, 50:305–31.

Dahiya, R. S., Metta, G., Valle, M., and Sandini, G. (2010). Tactile sensing—from humans to humanoids. *IEEE T Robot*, 26(1):1–20.

Dandekar, K., Raju, B. I., and Srinivasan, M. A. (2003). 3-D finite-element model of human and monkey fingertips to investigate the mechanics of tactile sense. *J Biomed Eng*, 125(5):682–91.

Deisenroth, M. P., Huber, M. F., and Hanebeck, U. D. (2009). Analytic moment-based Gaussian process filtering. In *Proc. Int. Conf. on Machine Learning (ICML)*.

Delmas, P., Hao, J., and Rodat-Despoix, L. (2011). Molecular mechanisms of mechanotransduction in mammalian sensory neurons. *Nat Rev Neurosci*, 12(3):139–153.

Denzler, J. and Brown, C. M. (2002). Information theoretic sensor data selection for active object recognition and state estimation. *IEEE T Pattern Anal*, 24(2):145–157.

DiCarlo, J. J., Johnson, K. O., and Hsiao, S. S. (1998). Structure of receptive fields in area 3b of the primary somatosensory cortex in the alert monkey. *J Neurosci*, 18(7):2626–45.

Doshi, F., Pineau, J., and Roy, N. (2008). Reinforcement learning with limited reinforcement: using Bayes risk for active learning in POMDPs. In *Proc. Int. Conf. on Machine Learning (ICML)*.

Doucet, A., de Freitas, N., and Gordon, N. (2001). *Sequential Monte Carlo Methods in Practice*. Springer.

Dragiev, S., Toussaint, M., and Gienger, M. (2011). Gaussian Process implicit surfaces for shape estimation and grasping. In *Proc. IEEE Int. Conf. on Robotics and Automation*.

Edin, B. B., Ascari, L., Beccai, L., Roccella, S., Cabibihan, J. J., and Carrozza, M. C. (2008). Bio-inspired sensorization of a biomechatronic robot hand for the grasp-and-lift task. *Brain Res Bull*, 75(6):785–95.

Filatov, N. M. and Unbehauen, H. (2000). Survey of adaptive dual control methods. *IEE P-Contr Theor Ap*, 147(1):118–128.

Foffani, G., Morales-Botello, M. L., and Aguilar, J. (2009). Spike timing, spike count, and temporal information for the discrimination of tactile stimuli in the rat ventrobasal complex. *J Neurosci*, 29(18):5964–73.

Foffani, G., Tutunculer, B., and Moxon, K. A. (2004). Role of spike timing in the forelimb somatosensory cortex of the rat. *J Neurosci*, 24(33):7266–71.

Friedman, R. M., Khalsa, P. S., Greenquist, K. W., and LaMotte, R. H. (2002). Neural coding of the location and direction of a moving object by a spatially distributed population of mechanoreceptors. *J Neurosci*, 22(21):9556–66.

Furukawa, S. and Middlebrooks, J. C. (2002). Cortical representation of auditory space: information-bearing features of spikes patterns. *J Neurophysiol*, 87(4):1749–1762.

Furukawa, S., Xu, L., and Middlebrooks, J. C. (2000). Coding of sound-source location by ensembles of cortical neurons. *J Neurosci*, 20:1216–1228.

Gandevia, S. C. and Hales, J. P. (1997). The methodology and scope of human microneurography. *J Neurosci Methods*, 74(2):123–136.

Gawne, T. J., Kjaer, T. W., and Richmond, B. J. (1996). Latency: another potential code for feature binding in striate cortex. *J Neurophysiol*, 76(2):1356–60.

Gibson, J. J. (1962). Observations on active touch. *Psychol Rev*, 69:477–91.

Girard, A., Rasmussen, C. E., Candela, J. Q., and Murray-Smith, R. (2003). Gaussian process priors with uncertain inputs – application to multiple-step ahead time series forecasting. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pages 545–552. MIT Press.

Goldberger, J., Gordon, S., and Greenspan, H. (2003). An efficient image similarity measure based on approximations of KL-divergence between two Gaussian mixtures. In *Proc. IEEE Int. Conf on Computer Vision (ICCV)*, volume 1, pages 487–493.

Gollisch, T. and Meister, M. (2008). Rapid neural coding in the retina with relative spike latencies. *Science*, 319(5866):1108–11.

Goodwin, A. W., Browning, A. S., and Wheat, H. E. (1995). Representation of curved surfaces in responses of mechanoreceptive afferent fibers innervating the monkey's fingerpad. *J Neurosci*, 15(1 Pt 2):798–810.

Goodwin, A. W., Macefield, V. G., and Bisley, J. W. (1997). Encoding of object curvature by tactile afferents from human fingers. *J Neurophysiol*, 78(6):2881–8.

Goodwin, A. W. and Wheat, H. E. (1992). Human tactile discrimination of curvature when contact area with the skin remains constant. *Exp Brain Res*, 88(2):447–50.

Goodwin, A. W. and Wheat, H. E. (1999). Effects of nonuniform fiber sensitivity, innervation geometry, and noise on information relayed by a population of slowly adapting type I primary afferents from the fingerpad. *J Neurosci*, 19(18):8057–8070.

Goodwin, A. W. and Wheat, H. E. (2004). Sensory signals in neural populations underlying tactile perception and manipulation. *Annu Rev Neurosci*, 27:53–77.

Grocholsky, B., Makarenko, A., and Durrant-Whyte, H. (2003). Information-theoretic coordinated control of multiple sensor platforms. In *Proc. IEEE Int. Conf. Robotics and Automation*, volume 1, pages 1521–26.

Gütig, R. and Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing based decisions. *Nat Neurosci*, 9(3):420–428.

Guyonneau, R., VanRullen, R., and Thorpe, S. J. (2005). Neurons tune to the earliest spikes through stdp. *Neural Comput*, 17(4):859–879.

Hausser, J. and Strimmer, K. (2009). Entropy inference and the James-Stein estimator, with application to nonlinear gene association networks. *J Mach Learn Res*, 10:1469–84.

Heil, P. (2004). First-spike latency of auditory neurons revisited. *Curr Opin Neurobiol*, 14:461–467.

Henriksen, R. (1982). The truncated second-order nonlinear filter revisited. *IEEE Trans. Automatic Control*, 27(1):247–51.

Hoffmann, G., Waslander, S., and Tomlin, C. (2006). Mutual information methods with particle filters for mobile sensor network control. In *Proc. IEEE Conf. Decision and Control*, pages 1019–1024.

Hosoda, K., Tad, Y., and Asada, M. (2006). Anthropomorphic robotic soft fingertip with randomly distributed receptors. *Robot Auton Syst*, 54(2):104–9.

Hsiao, K., Kaelbling, L., and Lozano-Perez, T. (2010). Task-driven tactile exploration. In *Proc. Robotics: Science and Systems*.

Huang, V. S., Shadmehr, R., and Diedrichsen, J. (2008). Active learning: learning a motor skill without a coach. *J Neurophysiol*, 100(2):879–87.

Isard, M. and Blake, A. (1996). Contour tracking by stochastic propagation of conditional density. In *Proc. Europ. Conf. on Computer Vision (ECCV)*, pages 343–356.

Ito, K. and Xiong, K. (2000). Gaussian filters for nonlinear filtering problems. *IEEE Trans Autom Control*, 45(5):910–927.

Itti, L. and Koch, C. (2001). Computational modelling of visual attention. *Nat Rev Neurosci*, 2(3):194–203.

Iwamura, Y., Tanaka, M., Sakamoto, M., and Hikosaka, O. (1983). Functional subdivisions representing different finger regions in area 3 of the first somatosensory cortex of the conscious monkey. *Exp Brain Res*, 51(3):315–326.

Jaakola, T. S. and Jordan, M. I. (1998). Improving the mean field approximation via the use of mixture distributions. In Jordan, M. I., editor, *Learning in graphical models*, pages 163–173. Kluwer Academic Publishers.

Jacobs, A. L., Fridman, G., Douglas, R. M., Alam, N. M., Latham, P. E., Prusky, G. T., and Nirenberg, S. (2009). Ruling out and ruling in neural codes. *P Natl Acad Sci Usa*, 106(14):5936–41.

Jenmalm, P., Birznieks, I., Goodwin, A. W., and Johansson, R. S. (2003). Influence of object shape on responses of human tactile afferents under conditions characteristic of manipulation. *Eur J Neurosci*, 18(1):164–176.

Jenmalm, P., Dahlstedt, S., and Johansson, R. S. (2000). Visual and tactile information about object-curvature control fingertip forces and grasp kinematics in human dexterous manipulation. *J Neurophysiol*, 84(6):2984–97.

Jenmalm, P. and Johansson, R. S. (1997). Visual and somatosensory information about object shape control manipulative fingertip forces. *J Neurosci*, 17(11):4486–99.

Jindrich, D. L., Zhou, Y., Becker, T., and Dennerlein, J. T. (2003). Non-linear viscoelastic models predict fingertip pulp force-displacement characteristics during voluntary tapping. *J Biomech*, 36(4):497–503.

Johansson, R. S. and Birznieks, I. (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat Neurosci*, 7(2):170–177.

Johansson, R. S. and Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nat Rev Neurosci*, 10(5):345–59.

Johansson, R. S., Landström, U., and Lundström, R. (1982). Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements. *Brain Res*, 244(1):17–25.

Johansson, R. S. and Vallbo, A. B. (1979). Detection of tactile stimuli: Thresholds of afferent units related to psychophysical thresholds in the human hand. *J Physiol (Lond)*, 297:405–422.

Johansson, R. S. and Westling, G. (1984). Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Exp Brain Res*, 56(3):550–64.

Johansson, R. S. and Westling, G. (1987). Signals in tactile afferents from the fingers eliciting adaptive motor responses during precision grip. *Exp Brain Res*, 66(1):141–54.

Johnson, K. O. (2001). The roles and functions of cutaneous mechanoreceptors. *Curr Opin Neurobiol*, 11:455–461.

Johnson, K. O. and Hsiao, S. S. (1992). Neural mechanisms of tactual form and texture perception. *Annu Rev Neurosci*, 15:227–50.

Johnson, K. O., Yoshioka, T., and Vega-Bermudez, F. (2000). Tactile functions of mechanoreceptive afferents innervating the hand. *J Clin Neurophysiol*, 17(6):539–558.

Johnsson, M. and Balkenius, C. (2007). Neural network models of haptic shape perception. *Robot Auton Syst*, 55(9):720–27.

Jones, L. A. and Lederman, S. J. (2006). *Human hand function*. Oxford University Press.

Julier, S. and Uhlmann, J. (2004). Unscented filtering and nonlinear estimation. *Proc. IEEE*, 92(3):401–422.

Junek, S., Kludt, E., Wolf, F., and Schild, D. (2010). Olfactory coding with patterns of response latencies. *Neuron*, 67(5):872–84.

Kakuda, N. (1992). Conduction velocity of low-threshold mechanoreceptive afferent fibers in the glabrous and hairy skin of human hands measured with microneurography and spike-triggered averaging. *Neurosci Res*, 15:179–188.

Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Trans ASME J Basic Eng*, 82:35–45.

Kayser, C., Logothetis, N. K., and Panzeri, S. (2010). Millisecond encoding precision of auditory cortex neurons. *P Natl Acad Sci USA*, 107(39):16976–81.

Kenton, B., Kruger, L., and Woo, M. (1971). Two classes of slowly adapting mechanoreceptor fibre in reptile cutaneous nerve. *J Physiol (Lond)*, 212(1):21–44.

Kim, S. S., Sripati, A. P., and Bensmaia, S. J. (2010). Predicting the timing of spikes evoked by tactile stimulation of the hand. *J Neurophysiol*, 104(3):1484–96.

Kroemer, O., Detry, R., Piater, J., and Peters, J. (2009). Active learning using mean shift optimization for robot grasping. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 2610–5.

Kunesch, E., Binkofski, F., and Freund, H. J. (1989). Invariant temporal characteristics of manipulative hand movements. *Exp Brain Res*, 78(3):539–46.

Kushner, H. J. and Budhiraja, A. S. (2000). A nonlinear filtering algorithm based on an approximation of the conditional distribution. *IEEE T Automat Contr*, 45(3):580–585.

LaMotte, R. H. and Srinivasan, M. A. (1996). Neural encoding of shape: responses of cutaneous mechanoreceptors to a wavy surface stroked across the monkey fingerpad. *J Neurophysiol*, 76(6):3787–97.

Lawrence, N. D. and Azzouzi, M. (1999). A variational Bayesian committee of neural networks. Technical report, University of Cambridge, UK.

Lederman, S. J. and Klatzky, R. L. (1987). Hand movements: a window into haptic object recognition. *Cognitive Psychol*, 19(3):342–68.

Lederman, S. J. and Klatzky, R. L. (1990). Haptic classification of common objects: knowledge-driven exploration. *Cognitive Psychol*, 22(4):421–59.

Lederman, S. J. and Klatzky, R. L. (1993). Extracting object properties through haptic exploration. *Acta Psychol*, 84(1):29–40.

Lederman, S. J. and Klatzky, R. L. (1997). Relative availability of surface and object properties during early haptic processing. *J Exp Psychol Human*, 23(6):1680–707.

Lee, M. H. and Nicholls, H. R. (1999). Tactile sensing for mechatronics—a state of the art survey. *Mechatronics*, 9:1–31.

Lee, S., Carvell, G. E., and J., S. D. (2008). Motor modulation of afferent somatosensory circuits. *Nat Neurosci*, 11(12):1430–8.

Lee, T. S. and Yu, S. X. (2000). An information-theoretic framework for understanding saccadic eye movements. In *Proc. Advances in Neural Information Processing Systems (NIPS)*.

Lefebvre, T., Bruyninckx, H., and de Schutter, J. (2004). Kalman filters for non-linear systems: a comparison of performance. *Int J Control*, 77(7):639–53.

Lesniak, D. R. and Gerling, G. J. (2009). Predicting SA-I mechanoreceptor spike times with a skin-neuron model. *Math Biosci*, 220(1):15–23.

Lewi, J., Butera, R., and Paninski, L. (2007). Real-time adaptive information-theoretic optimization of neurophysiology experiments. In *Proc. Advances in Neural Information Processing Systems (NIPS)*.

Lindblom, U. (1965). Properties of touch receptors in distal glabrous skin of the monkey. *J Neurophysiol*, 28(5):966–85.

Lindley, D. (1956). On a measures of the information provided by an experiment. *Ann Math Stat*, 27(4):986–1005.

Liu, Q. and Pierce, D. A. (1994). A note on Gauss-Hermite quadrature. *Biometrika*, 81(3):624–629.

London, M., Roth, A., Beeren, L., Häusser, M., and Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302):123–7.

Luttrell, S. (1985). The use of transinformation in the design of data sampling schemes for inverse problems. *Inverse Probl*, 1(3):199–218.

Ma, J. and Burdick, J. (2009). Sensor planning for object pose estimation and identification. In *Proc. IEEE Int. Workshop on Robotic and Sensors Environments (ROSE)*, pages 69–74.

Macefield, V. G. (2003). Cardiovascular and respiratory modulation of tactile afferents in the human finger pad. *Exp Physiol*, 88(5):617–25.

Macefield, V. G., Häger-Ross, C., and S., J. R. (1996). Control of grip force during restraint of an object held between finger and thumb: responses of cutaneous afferents from the digits. *Exp Brain Res*, 108.

MacKay, D. J. C. (1992). Information-based objective functions for active data selection. *Neural Comp*, 4(4):590–604.

Maeno, T., Kawamura, T., and Cheng, S. (2004). Friction estimation by pressing an elastic finger-shaped sensor against a surface. *IEEE T Robotic Autom*, 20(2):222–8.

Maeno, T., Kobayashi, K., and Yamazaki, N. (1998). Relationship between the structure of human finger tisue and the location of tactile receptors. *JSME Int J C-Mech Sy*, 41:94–100.

Martinez-Cantin, R., de Freitas, N., Brochu, E., Castellanos, J., and Doucet, A. (2009). A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Auton Robot*, 27:93–103.

Masquelier, T., Guyonneau, R., and Thorpe, S. J. (2008). Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS ONE*, 3(1):e1377.

Mei, J., Tuckett, R. P., Poulos, D. A., Horch, K. W., Wei, J. Y., and Burgess, P. R. (1983). The neural signal for skin indentation depth. II. steady indentations. *J Neurosci*, 3(12):2652–9.

Mikula, S. and Niebur, E. (2005). Rate and synchrony in feedforward networks of coincidence detectors: analytical solution. *Neural Comput*, 17(4):881–902.

Montani, F., Kohn, A., Smith, M. A., and Schultz, S. R. (2007). The role of correlations in direction and contrast coding in the primary visual cortex. *J Neurosci*, 27(9):2338–48.

Montemurro, M. A., Panzeri, S., Maravali, M., Alenda, A., Bale, M., Brambilla, M., and Petersen, R. S. (2007a). Role of precise spikes timing in coding of dynamic vibrissa stimuli in somatosensory thalamus. *J Neurophysiol*, 98(4):1871–82.

Montemurro, M. A., Senatore, R., and Panzeri, S. (2007b). Tight data-robust bounds to mutual information combining shuffling and model selection techniques. *Neural Comput*, 19:2913–2957.

Muniak, M. A., Ray, S., Hsiao, S. S., Dammann, J. F., and Bensmaia, S. J. (2007). The neural coding of stimulus intensity: linking the population response of mechanoreceptive afferents with psychophysical behavior. *J Neurosci*, 27(43):11687–99.

Najemnik, J. and Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434(7031):387–91.

Najemnik, J. and Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *J Vis*, 8(3):4.1–14.

Nakazawa, N., Ikeura, R., and Inooka, H. (2000). Characteristics of human fingertips in the shearing direction. *Biol Cybern*, 82(3):207–14.

Natale, L. and Torres-Jara, E. (2006). A sensitive approach to grasping. In *Proc. 6th Int. Workshop on Epigenetic Robotics*, pages 87–94.

Nelken, I., Chechik, G., Mrsic-Flogel, T. D., King, A. J., and Schnupp, J. W. (2005). Encoding stimulus information by spike numbers and mean response time in primary auditory cortex. *J Comput Neurosci*, 19(2):199–221.

Nemenman, I., Bialek, W., and de Ruyter van Steveninck, R. (2004). Entropy and information in neural spike trains: Progress on the sampling bias. *Phys Rev E*, 69:056111.

Nemenman, I., Lewen, G. D., Bialek, W., and de Ruyter van Steveninck, R. R. (2008). Neural coding of natural stimuli: Information at sub-millisecond resolution. *PLoS Comput Biol*, 4(3):e1000025.

Nemenman, I., Shafee, F., and Bialek, W. (2002). Entropy and inference, revisited. In *Proc. Advances in Neural Information Processing Systems (NIPS)*.

Noda, T., Miyashita, T., Ishiguro, H., and Hagita, N. (2008). Super-flexible skin sensors embedded on the whole body, self-organizing based on haptic interactions. In *Proc. Robotics: Science and Systems (RSS)*.

Norgaard, M., Poulsen, N. K., and Ravn, O. (2000). New developments in state estimation for nonlinear systems. *Automatica*, 36(11):1627–1638.

Ogata, T., Ohba, H., Komatani, K., and Okuno, H. G. (2005). Extracting multi-modal dynamics of objects using RNNPB. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 966–971.

Okamura, A. M. and Cutkosky, M. R. (2001). Feature detection for haptic exploration with robotic fingers. *Int J Robot Res*, 20(12):925–938.

Paletta, L. and Pinz, A. (2000). Active object recognition by view integration and reinforcement learning. *Robot Auton Syst*, 31(1–2):71–86.

Paletta, L., Prantl, M., and Pinz, A. (2000). Learning temporal context in active object recognition using bayesian analysis. In *Proc. Int. Conf. Pattern Recognition*, pages 695–699.

Panarese, A. and Edin, B. B. (2011). Human ability to discriminate direction of 3d force stimuli applied to the finger pad. *J Neurophysiol*, 105(2):541–547.

Paninski, L. (2003). Estimation of entropy and mutual information. *Neural Comput*, 15:1191–1253.

Paninski, L. (2005). Asymptotic theory of information-theoretic experimental design. *Neural Comput*, 17(7):1480–507.

Panzeri, S., Brunel, N., Logothetis, N. K., and Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends Neurosci*, 33(3):111–20.

Panzeri, S., Petersen, R. S., Schultz, S. R., Lebedev, M., and Diamond, M. E. (2001). The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron*, 29(3):769–777.

Panzeri, S., Senatore, R., Montemurro, M. A., and Petersen, R. S. (2007). Correcting for the sampling bias problem in spike train information measures. *J Neurophysiol*, 98:1064–1072.

Panzeri, S. and Treves, A. (1996). Analytical estimates of limited sampling bias in different information measures. *Network: Comput Neural Syst*, 7:87–107.

Papakostas, T., Lima, J., and Lowe, M. (2002). A large area force sensor for smart skin applications. In *Proc. IEEE Sensors*, volume 2, pages 1620–24.

Paré, M., Smith, A. M., and Rice, F. L. (2002). Distributions and terminal arborizations of cutaneous mechanoreceptors in the glabrous finger pads of the monkey. *J Comp Neurol*, 445(4):347–359.

Park, J. and Sandberg, I. W. (1991). Universal approximation using radial-basis-function networks. *Neural Comput.*, 3:246–257.

Pataky, T. C., Latash, M. L., and Zatsiorsky, V. M. (2005). Viscoelastic response of the finger pad to incremental tangential displacements. *J Biomech*, 38:1441–49.

Pawluk, D. T. and Howe, R. D. (1999). Dynamics lumped element response of the human fingerpad. *J Biomech Eng—T ASME*, 121(2):178–183.

Pei, Y. C., Hsiao, S. S., and Bensmaia, S. J. (2008). The tactile integration of local motion cues is analogous to its visual counterpart. *P Natl Acad Sci USA*, 105(23):8130–5.

Pei, Y. C., Hsiao, S. S., Craig, J. C., and Bensmaia, S. J. (2011). Neural mechanisms of tactile motion integration in somatosensory cortex. *Neuron*, 69(3):536–47.

Petersen, R. S., Brambilla, M., bale, M. R., Alenda, A., Panzeri, S., Montemurro, M. A., and Maravall, M. (2008). Diverse and temporally precise kinetic feature selectivity in the VPm thalamic nucleus. *Neuron*, 60(5):890–903.

Petersen, R. S., Panzeri, S., and Diamond, M. E. (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron*, 32(3):503–514.

Petersen, R. S., Panzeri, S., and Diamond, M. E. (2002). The role of individual spikes and spike patterns in population coding of stimulus location in rat somatosensory cortex. *Biosystems*, 67:187–193.

Petrovskaya, A., Khatib, O., Thrun, S., and Ng, A. Y. (2006). Bayesian estimation for autonomous object manipulation based on tactile sensors. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 707–714.

Phillips, J. R., Johansson, R. S., and Johnson, K. O. (1992). Responses of human mechanoreceptive afferents to embossed dot arrays scanned across fingerpad skin. *J Neurosci*, 12(3):827–39.

Pubols, B. H. (1982a). Factors affecting cutaneous mechanoreceptor response. I. constant-force versu constant-displacement stimulation. *J Neurophysiol*, 47(3):515–529.

Pubols, B. H. (1982b). Factors affecting cutaneous mechanoreceptor response. II. changes in mechanical properties of skin with repeated stimulation. *J Neurohysiol*, 47(3):530–42.

Pubols, B. H. and Pubols, L. M. (1976). Coding of mechanical stimulus velocity and indentation depth by squirrel monkey and raccoon glabrous skin mechanoreceptors. *J Neurophysiol*, 39(4):773–787.

Purves, D., Augustine, G. J., Fitzpatrick, D., Katz, L. C., LaMantia, A., McNamara, J. O., and Williams, S. M., editors (2001). *Neuroscience*. Sinauer Associates, 2nd edition.

Quaney, B. M., Rotella, D. L., Peterson, C., and Cole, K. J. (2003). Sensorimotor memory for fingertip forces: evidence for a task-independent motor memory. *J Neurosci*, 23(5):1981–6.

Quiñonera Candela, J. Q. and Rasmussen, C. E. (2005). A unifying view of sparse approximate Gaussian process regression. *J Mach Learn Res*, 6:1939–59.

Quian Quiroga, R. and Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nat Rev Neurosci*, 10(3):173–85.

Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT press.

Reed, J. L., Pouget, P., Qi, H. X., Zhou, Z., Bernard, M. R., Burish, M. J., Haitas, J., Bonds, A. B., and Kaas, J. H. (2008). Widespread spatial integration in primary somatosensory cortex. *P Natl Acad Sci USA*, 105(29):10233–7.

Reich, D. S., Mechler, F., and D., V. J. (2001). Independent and redundant information in nearby cortical neurons. *Science*, 294(5551):2566–8.

Rivest, I. R. and Gerling, G. J. (2010). Evaluating populations of tactile sensors for curvature discrimination. In *IEEE Haptics Symposium*, pages 59–62, Waltham, Massachussetts, USA.

Russell, A. F., Armiger, R. S., Vogelstein, R. J., Bensmaia, S. J., and Etienne-Cummings, R. (2009). Real-time implementation of biofidelic SA1 model for tactile feedback. In *Ann. Int. Conf. of the IEEE Engineering in Medicine and Biology Society (EMBS)*, pages 185–188.

Saal, H. P. (2007). Tactile feature discrimination under natural conditions. Master's thesis, University of Edinburgh.

Saal, H. P., Heess, N. M. O., and Vijayakumar, S. (2011). Multimodal nonlinear filtering using Gauss-Hermite quadrature. In *Proc. European Conference on Machine Learning (ECML)*.

Saal, H. P., Vijayakumar, S., and Johansson, R. S. (2009). Information about complex fingertip parameters in individual human tactile afferent neurons. *J Neurosci*, 29(25):8022–31.

Sanchez-Panchuelo, R. M., Francis, S., Bowtell, R., and Schluppeck, D. (2010). Mapping human somatosensory cortex in individual subjects with 7T functional MRI. *J Neurophysiol*, 103(5):2544–2556.

Scheibert, J., Leurent, S., Prevost, A., and Debregas, G. (2009). The role of fingerprints in the coding of tactile information probed with a biomimetic sensor. *Science*, 323(5920):1503–6.

Schiele, B. and Crowley, J. (1998). Transinformation for active object recognition. In *Proc. 6th Int. Conf. Computer Vision*, pages 249–254.

Schneider, A., Sturm, J., Stachniss, C., Reisert, M., Burkhardt, H., and Burgard, W. (2009). Object identification with tactile sensors using bag-of-features. In *Proc. Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 243–8.

Schöpfer, M., Ritter, H., and Heidemann, G. (2007). Acquistion and application of a tactile database. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 1517–1522.

Serina, E. R., Mockensturm, E., Mote, C. D., and Rempel, D. (1998). A structural model of the forced compression of the fingertip pulp. *J Biomech*, 31(7):639–46.

Serina, E. R., Mote, C. D., and Rempel, D. (1997). Force response of the fingertip pulp

to repeated compression—effects of loading rate, loading angle and anthropometry. *J Biomech*, 30(10):1035–40.

Settles, B. (2010). Active learning literature survey. Technical Report 1648, University of Wisconsin-Madison, Madison, WI, USA.

Shannon, C. (1948). A mathematical theory of communication. *Bell Syst Tech J*, 27(3):379–423.

Sibley, G., Sukhatme, G., and Matthies, L. (2006). The iterated sigma point Kalman filter with applications to long range stereo. In *Proc. Robotics: Science and Systems (RSS)*.

Silvey, S. D. (1980). *Optimal design*. Chapman and Hall.

Simpkins, A., de Callafon, R., and Todorov, E. (2008). Optimal trade-off between exploration and exploitation. In *Proc. American Control Conf.*

Singh, U. and Fearing, R. S. (1998). Tactile after-images from static contact. In *Proc. ASME Dynamic Systems, Control Division*, volume 64, pages 163–170.

Smith, A. M., Gosselin, G., and Houde, B. (2002). Deployment of fingertip forces in tactile exploration. *Exp Brain Res*, 147(2):209–18.

Snelson, E. and Ghahramani, Z. (2006). Sparse Gaussian processes using pseudo-inputs. In *Proc. Advances in Neural Information Processing Systems (NIPS)*.

Someya, T., Sekitani, T., Iba, S., Kato, Y., Kawaguchi, H., and Sakurai, T. (2004). A large-area, flexible pressure sensor matrix with organix field-effect transistors for artificial skin applications. *P Natl Acad Sci USA*, 101:9966–9970.

Song, L., Huang, J., Smola, A., and Fukumizu, K. (2009). Hilbert space embedding of conditional distributions with applications to dynamical systems. In *Proc. Int. Conf. on Machine Learning (ICML)*.

Srinivasan, M. A. and Dandekar, K. (1996). An investigation of the mechanics of tactile sense using two-dimensional models of the primate fingertip. *J Biomed Eng*, 118(1):48–55.

Sripati, A. P., Bensmaia, S. J., and Johnson, K. O. (2006a). A continuum mechanical model of mechanoreceptive afferent responses to indented spatial patterns. *J Neurophysiol*, 95(6):3852–64.

Sripati, A. P., Yoshioka, T., Denchev, P. V., Hsiao, S. S., and Johnson, K. O. (2006b). Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system. *J Neurosci*, 26(7):2101–14.

Stein, R. B., Gossen, R. E., and Jones, K. E. (2005). Neuronal variability: noise or part of the signal? *Nat Rev Neurosci*, 6(5):389–397.

Strohmayr, M. W., Saal, H. P., Potdar, A. H., and van der Smagt, P. (2010). The DLR touch sensor I: A flexible tactile sensor for robotic hands based on a crossed-wire approach. In *Proc. 2010 IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, pages 897–903.

Sudderth, E. B., Ihler, A. T., Isard, M., Freeman, W. T., and Willsky, A. S. (2010). Nonparametric belief propagation. *Commun ACM*, 53(10):95–103.

Suzuki, M., Noda, K., Suga, Y., Ogata, T., and Sugano, S. (2006). Dynamic perception after visually guided grasping by a human-like autonomous robot. *Adv Robotics*, 20(2):233–254.

Takamuku, S., Hosoda, K., and Asada, M. (2008). Object category acquisition by dynamic touch. *Adv Robotics*, 22(10):1143–1154.

Tegin, J. and Wikander, J. (2005). Tactile sensing in intelligent robotic manipulation—a review. *Ind Robot*, 32(1):64–70.

Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Netw*, 14(6-7):715–725.

Torebjörk, H. E., Vallbo, A. B., and Ochoa, J. L. (1987). Intraneural microstimulation in man. Its relation to specificity of tactile sensations. *Brain*, 110(Pt 6):1509–29.

Torkkola, K. (2003). Feature extraction by non-parametric mutual information maximization. *J Mach Learn Res*, 3.

Treves, A. and Panzeri, S. (1995). The upward bias in measures of information derived from limited data samples. *Neural Comput*, 7(2):399–407.

Tuci, E., Massera, G., and Nolfi, S. (2009). Active categorical perception in an evolved anthropomorphic robotic arm. In *Proc. IEEE Congress on Evolutionary Computation*, pages 31–38.

Vallbo, A. B. and Johansson, R. S. (1984). Properties of cutaneous mechanoreceptors in the human hand related to touch sensation. *Hum Neurobiol*, 3(1):3–14.

van Rossum, M. C. (2001). A novel spike distance. *Neural Comput*, 13(4):751–63.

Van Rullen, R., Guyonneau, R., and Thorpe, S. J. (2005). Spike times make sense. *Trends Neurosci*, 28(1):1–4.

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

Vickery, R. M., Gynther, B. D., and Rowe, M. J. (1994). Synaptic transmission between single slowly adapting type i fibres and their cuneate target neurones in cat. *J Physiol (Lond)*, 474(3):379–92.

Victor, J. D. and Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: A metric-space analysis. *J Neurophysiol*, 76(2):1310–1326.

Victor, J. D. and Purpura, K. P. (1997). Metric-space analysis of spike trains: theory, algorithms and application. *Network: Comput Neural Syst*, 8:127–164.

Wang, Q. and Hayward, V. (2007). In vivo biomechanics of the fingerpad skin under local tangential traction. *J Biomech*, 40(4):851–60.

Werner, G. and Mountcastle, V. B. (1965). Neural activity in mechanoreceptive cutaneous afferents: Stimulus-response relations, weber functions, and information transmission. *J Neurophysiol*, 28:359–97.

Westling, G. and Johansson, R. S. (1984). Factors influencing the force control during precision grip. *Exp Brain Res*, 53(2):277–84.

Wettels, N., Santos, V. J., Johansson, R. S., and Loeb, G. E. (2008). Biomimetic tactile sensor array. *Adv Robotics*, 22(8):829–849.

Wheat, H. E., Goodwin, A. W., and Browning, A. S. (1995). Tatcile resolution: Peripheral neural mechanisms underlying the human capacity to determine positions of objects contacting the fingerpad. *J Neurosci*, 15(8):5582–95.

Wheat, H. E., Salo, L. M., and Goodwin, A. W. (2004). Human ability to scale and discriminate forces typical of those occurring during grasp and manipulation. *J Neurosci*, 24(13):3394–401.

Wheat, H. E., Salo, L. M., and Goodwin, A. W. (2010). Cutaneous afferents from the monkeys fingers: responses to tangential and normal forces. *J Neurophysiol*, 103(2):950–961.

Williams, C. K. I. and Rasmussen, C. E. (1995). Gaussian processes for regression. In *Proc. Advances in Neural Information Processing Systems (NIPS)*. MIT Press.

Witney, A. G., Wing, A., Thonnard, J. L., and Smith, A. M. (2004). The cutaneous contribution to adaptive precision grip. *Trends Neurosci*, 27(10):637–43.

Wu, J. Z., Krajnak, K., Welcome, D. E., and Dong, R. G. (2006a). Analysis of the dynamic strains in a fingertip exposed to vibrations: Correlation to the mechanical stimuli on mechanoreceptors. *J Biomech*, 39(13):2445–56.

Wu, J. Z., Krajnak, K., Welcome, D. E., and Dong, R. G. (2006b). The stimulus location and static compression affect the biodynamic interaction between the fingertip and probe in the vibrotactile test: 3d finite element simulations. In *Inter-noise 06*.

Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., and Black, M. J. (2006c). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Comput*, 18:80–118.

Yang, Y., DeWeese, M. R., Otazu, G. H., and Zador, A. M. (2008). Millisecond-scale differences in neural activity in auditory cortex can drive decisions. *Nat Neurosci*, 11(11):1262–3.

Yau, J. M., Pasupathy, A., Fitzgerald, P. J., Hsiao, S. S., and Connor, C. E. (2009). Analogous intermediate shape coding in vision and touch. *P Natl Acad Sci USA*, 106(38):16457–62.