



This is a repository copy of *Biomimetic robots as scientific models: a view from the whisker tip*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/107024/>

Version: Accepted Version

Book Section:

Prescott, A.J., Mitchinson, B., Pearson, M.J. et al. (1 more author) (2011) Biomimetic robots as scientific models: a view from the whisker tip. In: Krichmar, J. and Wagatsuma, H., (eds.) *Neuromorphic and Brain-Based Robots*. Cambridge University Press , pp. 23-57. ISBN 9780521768788

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Biomimetic robots as scientific models: A
view from the whisker tip

Ben Mitchinson, Martin J. Pearson, Anthony G. Pipe, Tony J.
Prescott

Abstract

Why build robot models of animals and their nervous systems? One answer is that in building a robot model of a target organism, that mimics sufficiently some aspects of that animal's body, brain and behaviour, we can expect to learn a good deal about the original creature. Synthesis (engineering) is quite different from analysis (reverse-engineering), often easier, and teaches fascinating lessons (Braitenberg, 1986). Another answer is that a robot model should allow us to conduct experiments, that will help us better understand the biological system, and that would be impossible or at least much more difficult to perform in the original animal (Rosenblueth and Wiener, 1945). In this chapter our target organism is the rat and our specific focus is on the sophisticated tactile sensory system provided by that animal's facial whiskers (vibrissae). Neurobiology shows us that the brain nuclei and circuits that process vibrissal touch signals, and that control the positioning and movement of the whiskers, form a neural architecture that is a good model of how the mammalian brain, in general, co-ordinates sensing with action. Thus, by building a robot whisker system, we can take a significant step towards building the first robot 'mammal'. Following a short review of relevant rat biology, this chapter will describe the design and development of two whiskered robot platforms—Whiskerbot and SCRATCHbot—that we have constructed in order to better understand the rat whisker system, and to test hypotheses about whisker control and vibrissal sensing in a physical brain-based device. We provide a description of each platform, including mechanical, electronic and software components, discussing, in relation to each component, the design constraints we sought to meet and the trade-offs made between biomimetic ideals and engineering practicalities. Some results obtained using each platform are described together with a brief outline of future development plans. Finally, we discuss the use of biomimetic robots as scientific models and consider, using the example of whiskered robots, what contribution robotics can make to the brain and behavioural sciences.

0.1 Introduction

Rats are endowed with prominent facial whiskers (Figure 0.1) which they use to explore the environment immediately surrounding their head. This tactile sense is generally considered to be primary in rat in the way vision




Figure 0.1 Macro- and microvibrissae. Left inset shows a close-up of the microvibrissal region centered around the upper lip (see outlined region on main image), the right inset shows the microvibrissae being used to investigate a coin lying on the floor. Note the regular grid-like organisation of the actuated macrovibrissae.

is primary in primates—to the untrained eye the behaviour of blind rats can appear indistinguishable from that of sighted animals.

One group of whiskers are the long ‘macrovibrissae’ that are arranged into a regular grid of rows and columns set into the ‘mystacial pads’ on each side of the snout. These are moved back and forth, when the animal is actively sensing its environment, in a behaviour known as ‘whisking’. A second group, shorter and less regularly organised, is distributed over the front and underside of the snout, and is referred to as the ‘microvibrissae’. These whiskers do not have a musculature. This physical and functional dichotomy is reflected in the different uses to which the animal seems to put the two groups. The primary role of the macrovibrissae appears to be locating environmental features, whilst fine investigation is performed, in large part, by the microvibrissae, in concert with the other sensory apparatus around the front of the snout (the teeth, lips, tongue and nose) (Welker, 1964; Brecht et al., 1997).

The whisker system has become very popular as a model sensory system in neuroscience owing to its discrete organisation from the sensory apparatus (the whisker shaft) all the way to the sensory cortex (Petersen, 2007), its ease of manipulation and, not least, its presence in the laboratory rat. Our approach to this system begins with neuroethology, wherein we study neural systems holistically, including the observation of natural behaviour as well as comparative and evolutionary data, and leading to computational models. We then expose these models to the complexities of real-world operation, and the demands of functional robotics, revealing shortcomings that do not manifest in simulation. This engineering process feeds back, raising questions that are not raised (or addressed) by current biological data, and guiding us in the design of future biological experiments. Along the way, we hope to show that whiskers can be a useful robotic sensory system.

Below, we briefly review the neuroethology of the rat whisker system, including results from behavioural experiments conducted in our own laboratory. We then go on to describe two robotic platforms that we have developed, ‘Whiskerbot’ and ‘SCRATCHbot’. Whiskerbot was our first attempt to develop a biomimetic model of the rat whisker system and was, consequently, the more primitive mechanically. Investigations focussed on embedded models of neural processing as well as our first efforts to build and control artificial whiskers. SCRATCHbot¹ is our latest platform and still a work-in-progress. Compared to Whiskerbot, the new robot is more refined mechanically and electrically and has more degrees of freedom. The sensor transduction and whisker actuation mechanisms have been redesigned based on insights drawn from our earlier attempt. SCRATCHbot’s control system uses a mixture of neural-like and arithmetic computation with a focus on modelling motor control and sensory processing at a relatively abstract level. Its whisker system is also much closer to being a practical artificial sensory system, with possible applications in autonomous robotics.

Our work builds on, and was inspired by, a large number of previous research efforts in robotic tactile sensing systems, including, but not limited to, other whiskered robots that have been developed (Russell and Wijaya, 2005; Fend et al., 2004; Kim and Möller, 2004; Seth et al., 2004a; Solomon and Hartmann, 2006). We have recently provided an extensive review of artificial whisker systems (Prescott et al., 2009) therefore related projects are mentioned below only where they are of

¹ The name of the robot is derived from the acronym Spatial Cognition and Representation through Active TouCH.

direct relevance to the design decisions that we made in developing our own robot platforms. Our wider goal, through this review, is to describe the development of a research programme in neuromorphic robots, including the trade-offs made between accurate biomimicry and the need to engineer functioning systems at reasonable cost. We aim to show that, despite these constraints, the robotics provides insights to the biology; we will illustrate this using examples, and discuss the matter at the end of the chapter.

0.2 Neuroethology of the rat vibrissal system

0.2.1 Morphology, sensory transduction, and whisker actuation

Rat macrovibrissae are made of keratin, are tapered from base to tip, are curved (Figure 0.1), and are typically between 20 and 50 mm long (length varying regularly with location on the face) (Brecht et al., 1997). Their frequency response and other mechanical characteristics have been quantified both *in vivo* and *ex vivo* (Hartmann et al., 2003; Neimark et al., 2003), and mechanical response seems to play a key role in signal transduction (Lottem and Azouz, 2009).

Each macrovibrissa is mounted in a modified hair follicle, a roughly ellipsoidal capsule around 1mm in diameter and 3mm long (Rice et al., 1986), which is responsible for transducing mechanical signals into neural signals. A rich variety of mechanical signals are transduced—around 150-200 sensory nerves serve each follicle, and seven or more anatomically distinct classes of ‘mechanoreceptor’ (the mediators of biological tactile transduction) are found distributed throughout the follicle (Ebara et al., 2002). Amongst this range of signals, transverse whisker deflections have been the most studied and are known to generate strong signals in a large proportion of sensory cells (Lichtenstein et al., 1990; Shoykhet et al., 2000) (‘deflection cells’). Cells that transduce something related to whisker angular position (Szwed et al., 2003) (‘angle cells’) and longitudinal deflections (Stüttgen et al., 2008) have also been observed.

The principal, and first-described, component of whisker kinematics is the periodic, forward-backward (anterior-posterior, AP) motion of all macrovibrissae together (Welker, 1964; Zucker and Welker, 1969), a behaviour known as ‘whisking’. A smaller, synchronised, up-down (dorsal-ventral, DV) component to this motion (Bermejo et al., 2002) has been

identified (that is, a typical ‘whisk’ is reminiscent of a ‘rowing’ action), as has a torsional rotation of the shaft during the whisk cycle (Knutsen et al., 2008). Furthermore, the whisker columns move at somewhat different speeds with the net effect that the angular separation, or spread, between the whiskers varies significantly within each whisk cycle (Grant et al., 2009). Finally, the whiskers do not always move in concert on the two sides of the face (Sachdev et al., 2003; Mitchinson et al., 2007), and the mystacial pad moves substantially during whisking (Hill et al., 2008). Nonetheless, AP motion of all whiskers together describes a large proportion of overall whisker motion (Grant et al., 2009).

The ‘intrinsic’ muscles are found under the skin of the pad, wrap around each follicle, and are anchored to the skin and/or to neighbouring follicles (Dörfl, 1982). These drive ‘protraction’ (forward angular motion) of whiskers individually, by rotating the follicle around a lower pivot point beneath the skin (Dörfl, 1982; Wineski, 1985). Whisker ‘retraction’ (rearward angular motion) is partly passive, due to the elastic properties of the skin, and partly active, driven by the ‘extrinsic’ muscles to the rear of the pad (Carvell et al., 1991). These muscles pull the pad backward, causing all the follicles to rotate around an upper pivot point (Berg and Kleinfeld, 2003). A more recent study reports a contribution to protraction from another set of extrinsic muscles forward of the pad (Hill et al., 2008).

0.2.2 Whisker motion and active sensing behaviour

Rats generally whisk when they are exploring an environment or attempting most forms of tactile discrimination. Studies of neural responses to ‘passive’ whisker deflection (deflecting the whiskers of an anaesthetised rat) are therefore beginning to give way to studies of more natural ‘active’ deflection where moving whiskers encounter stationary obstacles. These studies show that whisker motion plays a key role in signal formation. There is no evidence of proprioception in the whisker musculature but angle cells may provide equivalent information (Szwed et al., 2003). Either these cells, or the temporal relationship between whisker motion and whisker deflection, are thought to provide the information necessary to transform deflections of moving whiskers into an appropriate head-centered reference frame (Szwed et al., 2003; Ahissar and Arieli, 2001).

Whisking motor patterns vary substantially with behavioural circumstance, but discernable ‘bouts’ of more-or-less periodic whisking at 6-

10Hz interspersed by periods of inactivity, are typical. Whisk frequency tends to be relatively constant within a bout (Hill et al., 2008) but other kinematic parameters can vary substantially driven, apparently, by both internal and external variables. The strongest observed external influence is whisker-environment contact, which rarely fails to modulate whisking (Grant et al., 2009; Mitchinson et al., 2007). For instance, a unilateral unexpected whisker-environment contact generally leads to suppression of protraction ipsilaterally (i.e. on the side the contact was made) and to increased protraction amplitude contralaterally (see figure 0.2). We hypothesise that this is the outcome of a control policy we term ‘Minimal Impingement, Maximal Contact’ (MIMC), which tends to maximise the count of whisker-environment contacts, whilst keeping the depth of those contacts within a managed range to maintain signal quality. A further observation (Grant et al., 2009) that spread between whisker columns is reduced during environmental contact is consistent with this policy, with rearward, non-contacting, whiskers brought forward to meet an ipsilateral obstruction. Another, internal, modulatory influence is head-rotation, whereby the animal appears to preempt upcoming head rotations by moving its whiskers backward (forward) on the side to which (away from which) the head will turn (Towal and Hartmann, 2006). In addition to these asymmetries, a temporary loss of bilateral synchrony in whisker movements is often observed following a unilateral contact (Mitchinson et al., 2007), whilst repeated contacts with the environment can lead to longer periods of desynchronization (unpublished results from our laboratory).

Psychophysical and behavioural experiments (see (Prescott et al., In press) for review) show that, using only the data gathered by their macrovibrissae, rats can locate objects accurately in space (Knutsen et al., 2006), perform fine textural discriminations (Carvell and Simons, 1990), and judge gap widths (Krupa et al., 2001), and that both macro- and micro- vibrissae are required for effective prey capture (Anjum et al., 2006). However, a reasonable hypothesis is that macrovibrissae are primarily used for locating objects, and then microvibrissae are brought to bear for close investigation (Brecht et al., 1997). For instance, the microvibrissae seem to be used preferentially in a shape-discrimination task (Brecht et al., 1997; Fox et al., 2009), and in our own laboratory, where we do not constrain the animal’s behaviour, we consistently see the microvibrissae used for close investigation of surfaces and objects.

As a consequence of these findings, and from inspecting many in-house video recordings of rats exploring environments and objects, we




Figure 0.2 Whisking asymmetry induced by contact with a surface. Frames from two example high-speed video sequences recorded in our laboratory, showing exploring rats with whiskers at the maximum protraction phase of the whisk cycle, but with the whiskers ipsilateral to an object of interest held back towards the cheek, whilst the contralateral whisker field pushes forward towards the surface. Electromyograph recordings from the whisking muscles confirm that the contralateral whiskers are driven significantly harder than the ipsilateral ones (Mitchinson et al., 2007). We interpret behaviour such as this as evidence for a ‘Minimal Impingement, Maximal Contact’ (MIMC) active sensing control strategy.

consider the ‘orient’ behaviour, in which a rat positions its head such that the front of its snout is brought to bear on its apparent focus of attention (Figure 0.1, inset top-right), to be a key component of active sensing. Indeed, orienting should perhaps be considered as the primary active sensing strategy employed by the animal, with repetitive whisker motion (whisking) adding a second component that provides wider coverage of space, contact-detach cycles without head motion, and more precise control over the nature of contacts. Observing that the body must also be moved if the rat is to orient its snout to locations a little distance away, then we could consider that locomotion of a rat in a novel environment may be well described as a stream of orients of the snout to one location after another. That is, the rat shifts its focus of attention and the head, whiskers, and body follow. Thus, we might consider orienting to constitute the foundation of exploratory behaviour in general, and therefore to be a prerequisite for effective active sensing in any whiskered entity, animal or robot. Note that, in a familiar environment, episodes of locomotion with a specific destination in mind, as opposed

to as a series of orients to immediately-sensed features, are also seen. In such conditions, where locomotion is not motivated by sensing, we might expect different whisking behaviour, attuned more to supporting locomotion (e.g. to ensure a sound footing and avoid collisions) rather than to maximising the gain of new sensory information. Experiments to establish whether rat whisking behaviour is noticeably different in these circumstances are currently in progress in our laboratory.

Orients are generally observed to occur on the timescale of one or two whisking periods (Prescott et al., 2009). Contact usually occurs during whisker protraction, and repositioning of the snout may complete quickly enough such that the battery of contacts due to the subsequent protraction sample the neighbourhood of the attended object. Supplementary video 1 shows an example orient that completes in about one whisking period. In this clip, the orient has begun by 40ms following contact, and completes around 160ms after contact (with the peak of the subsequent protraction occurring about 120ms after contact). Some orients may take two (perhaps, more) whisks to complete—for instance supplementary video 2 shows an orient completed in the space of two whisks.

After orienting, the animal will often keep its snout near to an attended object for a few whisks in order to investigate it more closely using the sensory equipment around the snout. This activity can be complex, and is thus less easy to describe, but we often see an investigative behaviour we refer to as ‘dabbing’, whereby the microvibrissae are lightly touched or brushed against the object in synchrony with macrovibrissal protractions (Hartmann, 2001; Prescott et al., 2005). The result is that tactile information is obtained at high spatial density towards the center of the dab, through the microvibrissal array, whilst, within the same narrow time window, surrounding surfaces are sampled in a sparser fashion by the macrovibrissae. Supplementary video 1 shows the animal, immediately following the orient, performing five ‘dabs’ at the attended feature (the corner of a block) before appearing to move on, whisking and dabbing across the wider extent of the object. The whole operation, from contact through orient and dabbing to moving on, is completed within three quarters of a second.

0.2.3 Neurobiology of the rat vibrissal system

Anatomical loops at multiple levels are present in the rat whisker system (Kleinfeld et al., 1999) (see Figure 0.3). Within this complex circuit, the

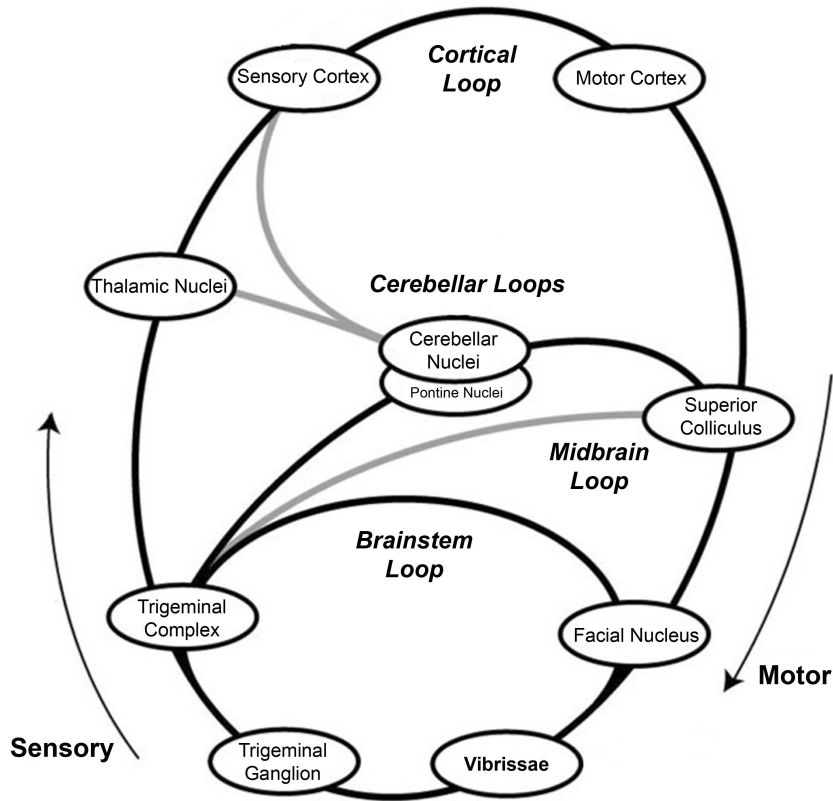


Figure 0.3 Neuroanatomical loops in the whisker system of the rat (modified from Kleinfeld et al. (2006)). Loops through non-whisker musculature (neck, shoulders, etc.) are not shown. Whisker sensory signals pass up from the vibrissae via the trigeminal nerves to the trigeminal sensory complex. From here, they pass along multiple loop paths, including: directly to the facial nucleus and back to the muscles driving the vibrissae; through superior colliculus; through cerebellar nuclei; through sensory and motor cortex via the thalamic nuclei.

most studied pathway is that carrying whisker signals upstream from the follicle, through the trigeminal complex, and ventro-posteromedial and posteromedial thalamic nuclei (known as 'VPM' and 'POm', respectively), to the primary ('barrel') somatosensory cortex (S1, Waite and Tracey (1995)). This, and related, cortical pathways are likely to be involved in extracting behaviourally-relevant features from vibrissal signals, and have been shown to be required for some whisker-driven tasks

(Krupa et al., 2001). However, it is important not to over-emphasise the role of this feedforward pathway to the cortex. As is clear from the diagram, the more general character of the system is that it consists of a set of nested closed loops which likely have different, though overlapping, functional roles. Each of these loops connects sensation to actuation at relatively short latencies, particularly at the lower levels of the neuraxis.

For instance, the pathway via the trigeminal complex to the facial nucleus, which contains the motoneurons that drive the intrinsic and extrinsic whisking musculature, provides a fast and direct, brainstem-only pathway through which contact information could affect the ongoing movement of the whiskers (Nguyen and Kleinfeld, 2005; Mitchinson et al., 2007). A midbrain loop through the superior colliculus (SC) (Drager and Hubel, 1976) is likely to underlie whisker-initiated orienting and avoidance responses (Sahibzada et al., 1986). The latency for whisker deflection signals to reach the SC may be as short as 5 milliseconds (Cohen et al., 2008) thus providing the capacity for very rapid orients as might be required, for instance, when the source of the stimulus is a moving prey animal. From the S1 cortex signals pass to the motor cortex (MCx) which has a substantial area devoted to the whisker system (Haiss and Schwarz, 2005). Although it is known that MCx sends a large projection to the brain areas involved in whisking pattern generation (Kleinfeld et al., 1999), its precise role in modulating or initiating whisking is poorly understood. Areas such as the trigeminal complex, SC, and S1 also project to two brain sub-systems that perform more general purpose roles within the brain architecture, the cerebellum and basal ganglia (BG). The cerebellum is thought to function as an adaptive filter (Dean et al., 2010) that could, for instance, use current state to predict future sensory signals. In contrast, the role of BG seems to be to provide action selection (Redgrave et al., 1999), that is, to decide which of the many possible behaviours available to the animal are engaged at any one time.

In the accounts of our whiskered robots that follow we will describe implementations that constitute embodied tests of functional hypotheses concerning several of the neural circuits and brain sub-systems mentioned here.

0.3 Whiskerbot

Our first foray into whiskered robotics was Whiskerbot (2003-2007) (Pearson et al., 2007b), which ran simultaneously with an effort to model neural components of the rat whisker sensory system. Computational models were developed of transduction in the whisker follicle (Mitchinson et al., 2004, 2008), processing in the trigeminal nucleus (Pearson et al., 2007a), transformation of whisker signals into a head-centered reference frame, orienting to whisker contacts (Mitchinson et al., 2006), and whisking pattern generation and its modulation by contact signals (Mitchinson et al., 2007). A number of these models used arrays of artificial ‘leaky-integrate-and-fire’ (LIF) neurons and were designed run on embedded onboard digital hardware in real-time (Pearson et al., 2006b, 2007a). The robot, including deployment of these models, is discussed in this section.

0.3.1 Robot

Whiskers and transducers

Since our control models are based on the neural systems of the rat, we expect them to work with signals of the type used by the rat. Where those systems are not fully understood, we cannot assume that they will work with signals with different characteristics. Our approach, therefore, is to engineer our whiskers to match those of the animal as closely as possible. Biological whiskers display an enviable combination of stiffness and damping, and even light contacts on the tip can elicit orienting, indicating that they effectively propagate mechanical deformation. Our hope was to design artificial whiskers for our robot with similar physical characteristics to their natural counterparts so that our robotic sensors would share these beneficial properties. Inevitably, however, since we were designing an engineered system composed of non-biological materials, the desire to accurately mimic the properties of natural whiskers and of the tissues that support them was balanced against engineering practicalities, necessitating some compromises, described below.

In the past, potentiometers and springs have been used to measure torque at the whisker base (Russell, 1992; Ueno and Kaneko, 1994). More recently, real rat whiskers were bonded to the diaphragms of electret microphones to measure high frequency deflections as the whiskers were moved across textured surfaces (Yokoi et al., 2005). In Whiskerbot, we chose resistive strain gauges bonded directly to the base of the arti-

ficial whiskers and configured to measure strain in two axes, which we denote x (AP) and y (DV) (Pearson et al., 2007b). The advantages of this approach were two-dimensional measurement, high sensitivity, and large bandwidth (limited only by sampling speed). The disadvantage of strain gauges is that high strains can reduce bond integrity, leading to calibration drift. Similar gauges have also been used recently configured, as in Whiskerbot, to measure two-dimensional strain at the whisker base (Quist and Hartmann, 2008).

Given this gauge technology, manufacturing constraints (particularly, attaching the gauges securely) forced us to build our whiskers physically scaled-up by a factor of four. Moreover, a complex robot of similar size to a rat would generate significant engineering challenges that we are not currently equipped to face. To mitigate the effect of this change on whisker dynamics, we adopted a temporal scaling-down by a factor of two (that is, our neural models ran at half wall-clock speed when deployed on the robot). This temporal scaling means that we match the animal’s behaviour by whisking at a more leisurely pace of around 3-4Hz.

We tested a variety of materials for the whisker shaft. A major finding was that, to generate a strain at the base of the whisker that was well clear of the noise floor, we were forced to build from materials much stiffer than those of real whiskers. This compromise did not prevent useful modelling—indeed, experiments described later in this chapter demonstrate effective contact detection and texture discrimination using these whiskers. What remains to be discovered is whether additional, or higher quality, sensory information might be recovered from whiskers that are a closer physical match to biology. An active line of research in our own and other laboratories is mechanical modelling of whiskers, in an attempt to gain insight into this question (Birdwell et al., 2007; Fox et al., 2008).

Whisking

As we have seen, actuating the whiskers, i.e. ‘whisking’, seems to be key to sensing in the animal. Several groups have shown that moving artificial whiskers across surfaces can provide useful information about surface features such as texture or shape (Russell, 1992; Ueno and Kaneko, 1994; Wilson and Chen, 1995; Gopal and Hartmann, 2007). Some mobile platforms have used non-actuated whiskers for obstacle avoidance and perception (Jung and Zelinsky, 1996; Seth et al., 2004b), and some have used actuated whisker arrays where all the whiskers, on both sides, are

moved together (Fend et al., 2004; Kim and Möller, 2007). Physical whisking mechanics as complex as that found in the biology (Hill et al., 2008) would be very challenging to implement. However, we noted above that AP angular motion is sufficient to reproduce a substantial part of the whisker motion observed in the animal (Grant et al., 2009). Our current robot platforms therefore reproduce just this degree of freedom for whisker motion, although unlike previous whisking robots we do allow for independent AP movement of individual whiskers (Whiskerbot) or whisker columns (SCRATCHbot). Shaft encoders provide sensory data encoding the angle of each whisker carrier (denoted θ), in analogy to biological angle cells.

The Whiskerbot platform was designed to have nine whiskers on each side. Shape-memory alloy wire was used to independently actuate each whisker to minimise weight and power consumption. Passing current through this material generates heat, causing a linear muscle-like contraction, which generated whisker protraction. Springs played the role of tissue elasticity, providing passive retraction (Dörfl, 1982). This system was able to whisk at up to 5Hz when fans were used to cool the actuating wires; however, operating in this region limited the lifetime of the wires significantly. Most experiments, therefore, were performed in the 1-2Hz range.

Platform

The basic platform layout is a two-wheel differential drive unit, with the head fixed in relation to the body, and a ‘snout’ area at the front of the head (no microvibrissae were included, but the area at the front of the head between the macrovibrissal fields was designated as the snout to act as the target area for orienting). The whiskers were mounted in rows on either side of the head. In common with most robotic platforms, the actuators were driven by local feedback controllers, in response to set-point signals (position or velocity). This represents a significant departure from biological actuation, which is open-loop at short timescales. We contend, however, that this is an advantageous approach during early development, since it decouples actuation mechanics from motor control. Not only does this avoid the need to develop mechanics that are functionally similar to those of the biology, as well as non-trivial plant-tailored open-loop controllers, it also frees us to update our model of the mechanical plant (in the light of new results) without having to change the mechanics themselves. Thus, we can simulate any biological plant in software, and charge the servos with following its outputs. The disadvantage of




Figure 0.4 Photographs of whisking actuation technology used for **a)** Whiskerbot, shape memory wire (protraction) and springs (passive retraction), and **b)** SCRATCHbot, motor and gearbox driving each column

this approach is that interactions between the plant mechanics and the mechanics of the environment are only indirectly modelled. Since these interactions could be a contributor to whisker transduction, we may seek to move away from artificial closed-loop control as our physical models of whiskers and whisking mature.

Computing

The ever-present constraint of any mobile robotics platform is power, which limits the amount of processing that can be deployed. This constraint is hardened by the fact that closed-loop controllers need to receive update signals on a regular period to avoid erratic behaviour. Thus, software components must operate in strict real-time, constraining model complexity (size, resolution, scope etc.). To mitigate this on Whiskerbot, we offloaded some of the large (but simple) neural models onto re-programmable hardware devices (Field Programmable Gate Arrays, FPGAs). Using parallel computing and function-specific hardware, in

this way, we were able to run detailed spiking neural network models yet maintain strict adherence to the real-time constraint (Pearson et al., 2007a, 2006a). Coordination of this processing architecture, incorporating both hardware and software components, was left to BRAHMS a framework for integrated heterogeneous computing that we have developed in-house (Mitchinson et al., 2010). The main control resource of the platform was a re-configurable computing platform called a ‘Bennuey’ PC-104+ motherboard (Nallatech, 2007). This consists of a PC-104 Single Board Computer (SBC) and a number of expansion slots for FPGA modules. One FPGA was used as a bridge between hardware communications systems and the PCI bus of the PC. Other FPGAs were configured for hardware acceleration of spiking neural models in a ‘neural coprocessor’ and a ‘follicle coprocessor’ (Pearson et al., 2005, 2006b) (see Section 0.3.2).

0.3.2 Control Architecture

An overview of the robot control architecture is illustrated in Figure 0.5, which we map loosely onto parts of the biological architecture of Figure 0.3; the figure shows the SCRATCHbot configuration. At the bottom left is the interface to the robot platform; this consists of the neural and follicle coprocessors (Whiskerbot) as well as the sensors (x, y, θ) and actuators (whiskers, wheels, neck). The remainder of the architecture can be described as an inner loop (small circular arrow) mediating whisking pattern modulation, and a middle loop (large circular arrow) mediating behaviour. Higher loops (curved arrow to left) modelling cortical and hippocampal systems, for such competences as object discrimination and spatial mapping, are the subject of current work in our laboratory (Fox et al., 2009) (see also Discussion); the current architecture exhibits only immediate responses, and has no long-term memory. The Whiskerbot configuration is similar, but uses the spiking output of the follicle coprocessor to drive the coordinate transform (and does not perform reafferent noise removal) and implements the Whisker Pattern Generator (WPG) as a spiking neuron model in the neural coprocessor, with modulation from the follicle coprocessor output.

Whisker pattern generation

The Whisker Pattern Generator (WPG) is a model of the central pattern generator present (Gao et al., 2001) (though not yet located (Cramer et al., 2007)) in the rat brain and whose activity underlies the rhythmic

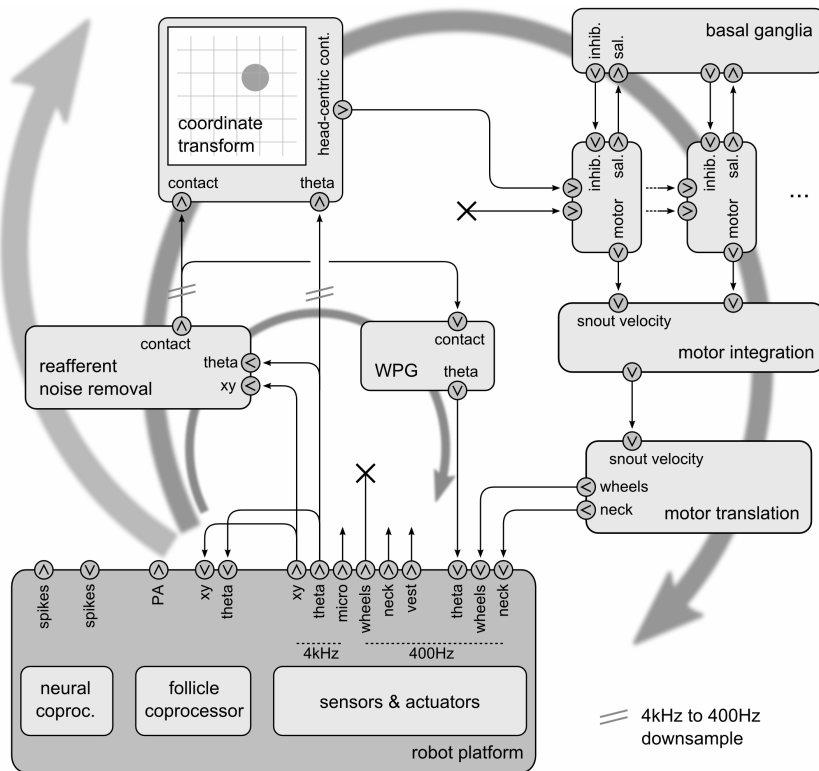


Figure 0.5 Overview diagram of control architecture of SCRATCHbot/Whiskerbot, SCRATCHbot configuration shown (see text for details). Straight arrowed lines represent numeric data streams at 400Hz or 4kHz. **xy** and **theta** are raw sensor data. Reafferent noise is removed from **xy** to generate **contact**, the clean per-whisker contact strength signal, which modulates the WPG. **contact** is transformed into head-centric space using the instantaneous values of **theta**, and is passed to the actions. Competition between actions for the motors is mediated by the Basal Ganglia. Finally, the selected action sends snout velocity to the motor translation layer, which generates platform-specific control signals for the actuators.

whisker motions observed in the behaving animal. Although bilaterally asynchronous whisker movements are sometimes observed in animals, particularly following, or during, interactions with surfaces, movements of the two whisker fields generally appear to be tightly coupled such that perturbations are usually corrected within a small number of whisker cycles (often one). Likewise, within each whisker field some asynchronous

movements have been observed, but this has yet to be quantified or explained. Rather than address the issue of multiple coupled WPGs at the outset of our study, for Whiskerbot we chose to use a single WPG to generate the base whisking signal, and derive movement patterns for each whisker from this signal.

Typical whisking, as described above, can be broadly described as modulated periodic oscillations. Therefore, our model is based around modulation of periodic oscillators. The simplest possible model is a single oscillator generating the angular position of each individual whisker through a gain. In this model, the whiskers are constrained to move synchronously (all whiskers in phase), symmetrically (whisking on the two sides having the same profile), and periodically (each whisking cycle is identical). Each of these constraints is relaxed in a long-enough series of rat whisking, as described above; in the robot, we can relax these constraints to test the impact of different modulation strategies. Although not enough is known about the rat WPG to constrain a detailed model we nevertheless chose to implement our model WPG in a neurally-inspired manner. For this purpose a base oscillator was formed from two, alternately active, spiking neuron populations, one population acting as the integrator and the other as the reset signal. An additional pair of populations, one on each side, relayed the oscillator activity to the actuator control board (and could be modulated, see below), where each of the signals is used to drive all whiskers on that side of the face. The WPG model was implemented in the neural coprocessor.

Sensory transduction through a model of the whisker follicle and trigeminal ganglion

In order to use the Whiskerbot platform to investigate embedded neural processing we needed to recode the x, y whisker deflections recorded using the strain gauges into spike trains in simulated sensory nerves. For this purpose we used a model of transduction that we developed in order to accurately reproduce observations made under both passive and active deflection conditions (Mitchinson et al., 2004). This model was in two parts. First, we developed a model of the mechanical contributions to transduction consisting of six anatomical masses, related by springs and dampers that represented the response of tissues and fluids in the ‘whisker-follicle assembly’. The parameters of this part of the model were largely drawn from anatomical studies, making direct connection with the anatomy, and the model confirms that the mechanics are substantially involved in signal formation. The second part of the

model was a conventional noisy LIF cell model that was used to simulate the sensory cells within the trigeminal ganglion (these cells do not merely transmit the response of the mechanoreceptors, but play a key role in signal formation). The parameters of this second model, along with remaining parameters of the mechanical part, were chosen to balance plausibility (e.g. smoothness) with reproduction of the results of biological experiments (Lichtenstein et al., 1990; Shoykhet et al., 2000; Szwed et al., 2003). Two distinct classes of deflection cell were modelled, ‘slowly adapting’ (SA) and ‘rapidly adapting’ (RA), reflecting a dichotomy widely described in the literature. These classes are defined by their responses to prolonged deflections: SA cells fire throughout these deflections, whilst RA cells fire only during their onset and offset. Both classes of cell are directionally sensitive, so information on the direction of whisker deflection is encoded. This model was computed on the Whiskerbot ‘follicle coprocessor’, in FPGA (see Figure 0.5).

Our control architecture also included a population of model whisker angle cells (Mitchinson et al., 2006) similar to those described in the animal by (Szwed et al., 2003). For this purpose, angular position, as measured by each shaft encoder, was used to drive a bank of cells associated with each whisker. Each cell responded strongly only when the measured angle was near to its preferred angle, with some overlap between cells. Thus, the identity and response of the active cells implies the whisker angle, consistent with biological data.

Coincidence detection

Whisker deflection signals, in the animal and in the robot control architecture, are generated whenever the whiskers are obstructed during protraction. The contact signals from each whisker are provided in a ‘whisker-centred’ frame-of-reference. In order to integrate information from multiple whiskers (and with information from other sensory modalities) signals therefore need to be transformed into a single reference frame (this is the ‘coordinate transform’ of Figure 0.5). In the rat, it has been hypothesised that this is performed by neural mechanisms that detect coincidences between firing in deflection cells and angle cells (Szwed et al., 2003). In the Whiskerbot control architecture a sheet of LIF cells representing the head frame was driven by deflection cells and angle cells together, with innervation patterns calculated based on the known geometry. Thus, coincident firing generated activity in the correct location on the sheet (Mitchinson et al., 2006), and indicated contact at the encoded location. Since the architecture does not yet include a model

to determine the distance along the whisker at which contact occurred (although extensions to allow this are possible e.g. (Birdwell et al., 2007; Evans et al., 2008)), it is instead assumed that contact occurred close to the tip of the whisker. In practice, this assumption has proved adequate for the robotic experiments performed using the Whiskerbot platform, allowing us to defer the problem of computing radial distance to contact for later.

Action selection

A fundamental problem faced by all but the simplest organisms is deciding which action, of those possible, to take at one moment. Failure to efficiently select a unique action to perform can lead to confused use of the actuators (or muscles), and flawed behaviour (Prescott et al., 2006, 2007). Appropriate action selection should take into account not just exteroceptive sensory signals, such as those derived from whisker-environment contacts, but also proprioceptive signals (odometry, for example) and internal indicators of homeostatic and motivational state. One proposal is that these different signals are integrated in a collection of brain nuclei called the Basal Ganglia (BG) whose intrinsic circuitry appear to be optimised to perform efficient and robust selection between competing behavioural alternatives (Redgrave et al., 1999). This theoretical proposal has been developed into a number of computational neuroscience models that have been evaluated both in simulation and on a robot platform (Gurney et al., 2004; Prescott et al., 2006). We use a version of this BG model to perform action selection within our robot control architecture. Briefly, each action that the robot can express ‘bids’ for use of the motors by indicating its current ‘salience’—salience is higher when the action is strongly indicated by the current sensory input. The model BG chooses a winner, and allows that action to have use of the platform, avoiding mix-ups. Actions that Whiskerbot exhibited include: **dead reckon** (use odometry to reach a pre-programmed location, analogous to path integration in rodents (Etienne et al., 1996)); various forms of **explore** (random walk, whisking from side to side to detect obstacles); and **orient** (orient snout towards focus of attention). Since we are interested in viewing the majority of motor output (at least during exploratory behaviour) as directly consequent to a desired positioning and orientation of the snout, all of these actions are designed to generate a desired snout velocity vector. Motor Integration across actions, then, means simply summing the (snout) velocity vector from all active actions, relying on the BG to suppress motor output from

non-selected actions. Note that we were not concerned with learning, here—the BG model has fixed weights and mediates between bids by actions, the salience of each of which is a pre-programmed function of current sensory inputs.

0.3.3 Whiskerbot Experiments

The Whiskerbot platform was used to validate the embedded computational neuroscience models (i.e. to demonstrate that they could adequately perform their intended role), to evaluate active control strategies for vibrissal sensing, and to develop and test classification methods for texture discrimination using whisker signals.

An example of the type of experiment performed, was our investigation of the likely consequences of a Minimal Impingement (MI) control strategy on the whisker deflection signals processed in the brain. As noted previously, our own behavioural observations in animals had indicated that whiskers rapidly cease to protract following contact with an object during exploration. We hypothesized that this result implied a control strategy that sought to minimize the extent to which whiskers were allowed to bend against surfaces. To implement MI in our robot control architecture the total activity across all whisker deflection cells on one side of the face was fed back to suppress activity in the ipsilateral WPG relay. This has the desired effect that protraction ceased rapidly after contact, as seen in the animal (Mitchinson et al., 2007). Figure 0.6 shows how MI affects WPG output, the signals consequently generated in the strain gauges, and the response of the simulated deflection cells, during whisking against a stationary obstacle (examples of whisker movement and deflection with MI off and on are shown in supplementary videos 3 and 4, respectively). With MI enabled, the signals are cleaner and more closely match those observed in the animal (Mitchinson et al., 2006; Pearson et al., 2007b). This result is in line with our predictions for the effect of this control strategy on signal quality in the animal (Mitchinson et al., 2007), and suggests that the whisker signals being relayed to the sensory cortex (and elsewhere) in awake, exploring animals will be quite different from those generated in the same animals in the absence of feedback control (as for instance, in whisker deflection experiments performed under anaesthesia).

We also tested Whiskerbot’s ability to orient to obstacles following contact. In these experiments the robot was allowed to proceed across a smooth floor and interact with one or more point obstacles (narrow

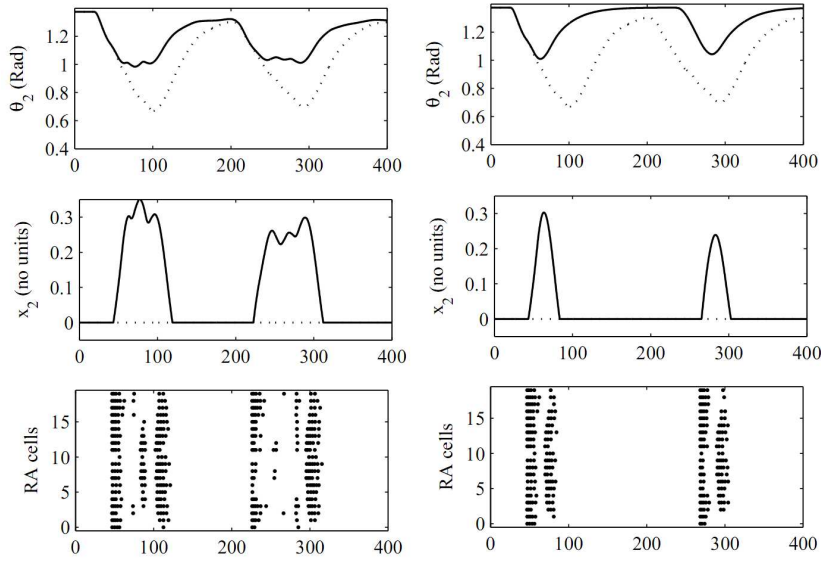


Figure 0.6 Signals recorded during two whisks where the rostral whiskers contact a stationary obstacle; left/right column of panels is without/with MI control. Top panels show rostral whisker column angles (solid line); dotted line is the same in both panels, and shows the output of the pattern generator (i.e. the requested column angle) for the unmodulated case. In the unmodulated case (top left), the whisker presses hard against the obstacle, and the column is physically prevented from moving further forward. In the modulated case (top right), the whisker briefly touches the obstacle, and the column is retracted by the MI policy. Other panels show strain gauge output (middle) and simulated rapidly adapting (RA) deflection cells (bottom, with each dot representing a spike). All panels reproduced from figures in (Mitchinson et al., 2006).

cylinders). Initially, the robot control architecture selects the default **explore** behaviour and the robot proceeds in a random walk, whisking as it goes. On detecting one of the obstacles, the robot switches to **orient**, and turns its snout to the obstacle before pausing—that is, it expresses an ‘orient’. With no ability to ‘dab’ on this platform, **orient** concludes with the robot backing off and heading in another direction. An example, using a version of the platform equipped with only two whiskers, is shown in supplementary video 5.

On occasion an orient was performed when no obstacle had been contacted. Under highspeed video observation, we see complex whisker dynamics excited by the motion of the whisker carrier—that is, the

whiskers ‘whip’ around when they are moved. This generates a relatively high power noise signal in the x data stream, leading to spurious spiking in contact cells. These ‘ghost orients’ were traced to this noise source (Pearson et al., 2007b) and are considered further in the section below describing SCRATCHbot.

In addition to providing an embodied test-bed for neural models an important goal for Whiskerbot was to investigate tactile feature extraction from artificial whiskers signals. To that end, we have shown that simple Gaussian classifiers, together with either hand-picked or biomimetic features derived from whisker contact signals, can be used as the basis for effective robotic texture classification (Fox et al., 2009). These results serve to further demonstrate the potential of vibrissal tactile sensing for applications in autonomous robotics such as navigation and spatial-mapping in darkness.

0.4 SCRATCHbot

SCRATCHBOT constituted a fundamental redesign of many of the physical and mechanical aspects of our whiskered robot, and a more high-level approach to the development of the embedded brain-based control system. The following sub-sections summarise and explain the main changes and extensions.

0.4.1 Changes to the Robot Design

Signal transduction

One limitation of Whiskerbot was that the strain gauges used to measure whisker deflection were apt to come loose such that repairing a snapped whisker, or re-attaching strain gauges, was a laborious process. Therefore, an important practical improvement for SCRATCHbot was to move to a more robust (though less sensitive) method of transduction (Figure 0.7). In the new whisker assembly the whisker shaft is supported in a rigid hollow sleeve by a thick layer of relatively soft polyurethane around the base; thus, it returns elastically to its rest position when it is released. A small, axially-magnetised, disc magnet is bonded to the very base of the whisker, and a miniature tri-axis Hall-effect sensor (HS) is positioned underneath. The HS generates two voltages linearly related to the displacement of the magnet in the two axes. This design took its inspiration from the work of Kim *et al.* (Kim and Möller, 2007). Each




Figure 0.7 Exploded view of the Hall effect based macro-vibrissae sensor assembly and an example output voltage plot taken from the sensor during 3 ramp-hold-release displacement profiles of varying magnitude. The voltage and displacement data has been normalised to the the maximum output voltage (5V), calibrated for a 40mm deflection of the whisker applied at a point 60% along the length from the base. Note that the magnitude of the third displacement profile exceeds the calibrated range resulting in the output voltage saturating.

whisker assembly is calibrated after construction by programming the rest position and deflection limits in each axis into the HS. This effectively sets the sensitivity of the whisker to deflections, since the voltages generated by the HS are linearly scaled to the programmed deflection range. Figure 0.7 shows some typical voltage trace outputs from a calibrated whisker during a range of controlled displacements in one axis

(deflection outside the programmed range results in saturation). Importantly, these sensory modules remain undamaged under large deflections. Where such deflections cause a whisker to break, a new whisker is simply inserted into the old module, and the HS recalibrated, to return to normal operation.

A secondary consequence of this change was the freedom to explore alternative, less stiff, materials for the whisker shaft. Nonetheless, it has still proved difficult to match the simultaneous high sensitivity and high damping displayed by rat whiskers. Partly, this is an issue of materials: biological whiskers are highly composited, and achieve good energy absorption simultaneously with good signal transmission; reproducing such a material will require substantial further development. Partly, it is an issue of scale: the smaller the whisker, the less prone to prolonged oscillation; a substantially smaller (factor four) whisker/whisking assembly is currently under development in our laboratory—this will be a close match to the size of a rat’s larger whiskers, and we expect to achieve a better sensitivity/damping trade-off as a result.

Microvibrissae

An important enhancement over Whiskerbot is the addition of an array of microvibrissae, mounted at the center of the snout between the two actuated macrovibrissal arrays. The central array consists of nine short (80mm) plastic whiskers, mounted at their base into a common polyurethane sheet, and again instrumented with magnets and Hall-effect sensors to measure deflections in 2-dimensions. These microvibrissae are able to initiate orienting in the same manner as the macrovibrissae if an obstacle is encountered snout-first. In addition, they are able to sample the object at a relatively high spatial resolution once the robot has oriented. The classifier systems originally developed for discriminating texture using Whiskerbot macrovibrissae are currently being adapted and extended for use with both the macro- and micro- vibrissal arrays (Evans et al., 2009), and to distinguish other object properties such orientation, distance, and velocity relative to the snout.

Whisker array geometry and actuation

A further issue that emerged with Whiskerbot was that the shape-memory alloy actuators chosen for whisker actuation, although lightweight and low power, were fiddly to work with and tended to deteriorate rapidly with use. Therefore, we moved to more conventional whisker actuation by DC motor for SCRATCHbot. To reduce the number of




Figure 0.8 Photograph of the SCRATCHbot platform with detailed view of the front right motor drive unit (top right panel) and the three degree of freedom neck assembly (lower right panel).

motors required, three whisker carriers were mounted on either side of a light-weight plastic ‘head’, with each carrier carrying three whiskers in a column. The geometry of the head was such that all the whiskers would point directly ahead of the robot when fully protracted. Each column can rotate (AP) through 120° , which is similar to the maximum angular range available to the rat. A second actuated axis of rotation (DV) was implemented on the SCRATCHbot platform, though this was limited to a single actuator for each side (three columns) and constrained to $\pm 15^{\circ}$ of rotation about the vertical.

Head and body design and actuation

We argued earlier that active sensing in rats is as much a matter of positioning the head and body as of positioning the whiskers. The Whiskerbot platform, with its head fixed relative to its body and nonholonomic constraints on the body, was very restricted in this respect (Prescott et al., 2009). With a view to opening up our investigations of the role of head and body positioning in sensing, SCRATCHbot was endowed with several additional degrees of freedom (DOF)—see Figure 0.8. The head is fixed to the top stage of a three DOF ‘neck’ assembly (Elumotion, 2009) with its base fixed to the body. This neck allows the head to be moved through axes referred to as elevation, pitch and yaw, as well as providing a certain amount of ‘reach’. In addition, the platform’s

three wheels can be turned through 180^0 , approximating a holonomic platform.

Computing

Finally, whilst demonstrating embedded onboard processing for neural models was a key goal in the development of Whiskerbot, this was less of a core objective in the design of SCRATCHbot. Onboard FPGA co-processors have been retained to allow for the possibility of computationally-intensive onboard processing, however we have also added a wireless signal link allowing the inclusion into the control architecture of asynchronous heavy-duty off-board processing. Three levels of computation are thus available on the new platform: short-latency, relatively inflexible (onboard FPGA); medium-latency, limited computational power, synchronous (in the PC-104); high latency, high computational power, asynchronous (off-board PC). This gives a range of computing options that are well suited to implement aspects of the lower, middle and upper loops of the neural architecture.

0.4.2 Changes and Additions to the Control Architecture

On SCRATCHbot, we have so far chosen not to model the encoding of x, y (whisker deflection signals) and θ (whisking angle) in simulated spiking neurons. Rather, these signals are propagated as continuously-valued numeric variables to other parts of the control architecture. Consequently, the brain-based models investigated in SCRATCHbot should be thought of as more abstract approximations to the computations performed by the rat vibrissal system than some of the spike-based models previously investigated in Whiskerbot. Modelling at this higher level, we believe, can allow us to make more rapid progress in identifying the computations that are performed by whisker-related circuitry in the rat brain. We remain very interested, however, in understanding how these more algorithmic models might be implemented in neural tissue. Indeed, for problems such as whisking pattern generation, we are currently conducting investigations at multiple levels of abstraction. We are using more detailed neurally-based models to explore, in simulation, the role of particular cell populations (such as the whisking motoneurons in the facial nucleus) whilst, at the same time, employing more abstract algorithmic WPG models to generate whisker movement patterns on the

robot platform that are more directly comparable with rat whisking behaviour.

Naturally, we have an extensive list of additions that we hope to make to the SCRATCHbot control architecture. The following sub-sections summarise some changes that have already been implemented, and also provide a brief outline of where we expect this work to go in the near future.

Using MIMC to control whisker spread

Predictable variation in whisker spread (the angular separation between the whiskers) was noted previously as a characteristic of animals that are exploring surfaces (Grant et al., 2009). To investigate the possible causes of this variability we extended the modulation options of SCRATCHbot’s WPG by implementing a separate relay for each column (rather than having just one for each side of the head, as in Whiskerbot). Whisker-environment contact excites all of these relays, whilst suppressing only those relays driving the whiskers that contacted the environment. The result is that, in addition to the per-side MIMC elicited in Whiskerbot, more rearward whiskers move more rapidly than they would otherwise, and are thus brought forward to meet a contacted obstacle. The net result is a reduction in inter-column spread following contact, as seen in the animal. Another way of putting this, is that, by implementing MIMC at the per-column level, ‘control’ of whisker spread appears as an automatic consequence of this general active sensing strategy - the whiskers are brought forward to meet the environment wherever possible, whilst being restrained from bending too far against it.

Head and body movement

A key task for the motor system is to generate control signals for the wheels and neck that achieve the desired snout movement; this takes place in the Motor Translation Layer (MTL) of our control architecture. Conventional robotic approaches to controlling multi-DOF systems (e.g. potential-field or sampling-based) can be expensive to solve, may suffer from local minima, may not be robust, and are not generally bio-plausible. We use, instead, an algorithm we call ‘Snake’, which takes a bio-inspired approach, causing free (uncontrolled) nodes of the mechanics to follow adjacent, controlled, nodes according only to mechanical constraints (that is, there is no explicit motion planning). Thus, actuators are ‘recruited’ to contribute to the movement in a distal-first pattern, as has been seen in the animal during fictive orienting

(Sahibzada et al., 1986), and more massy central nodes tend to be moved less than lightweight peripheral nodes. This algorithm results in motion that appears quite natural to the human observer. Furthermore, nodes are moved no more than necessary to achieve the target velocity of controlled nodes, and the computations for each node are local to that node and cheap, all of which are bio-plausible characteristics. The algorithm can be inferior to explicit motion planning under some conditions; for instance, the trajectory of the overall plant may pass through illegal configurations, or uncontrolled nodes may intersect obstacles. We hypothesise that, in biology, failures of this class are easier to deal with (reactively) than is computationally-heavy proactive global motion planning. In SCRATCHbot, usually only the snout node is controlled (snout location being the hypothesised goal), and the joints and base of the neck follow as if the robot was being led by a ring through its nose.

Predicting and cancelling sensory signals due to self-movement

Other than the increased complexity of the motor translation layer (due to the increased degrees of freedom in the head and neck), and the abstraction away from the spiking neuron level, SCRATCHbot initially employed similar algorithms for generating orienting responses to those used in Whiskerbot. The robot therefore displayed ‘ghost orients’, just as Whiskerbot did, for the same reason that movements of the robot body, head, or whiskers, can induce large transients in the transduced whisker signals.

To mitigate this problem, we took inspiration from a hypothesis of motor-reafferent noise removal in cerebellum (Anderson et al., 2009; Dean et al., 2010), and used a bank of linear adaptive filters, one for each whisker, to attempt to predict this noise signal based on the whisker angle signal measured from the shaft encoder attached to each column. After learning the parameters of the filters independently for each whisker from example data, this noise removal proved to be very effective. The nature of the motor-reafferent noise generated at the transducers depends both on the whisker material and length, so it would appear that both the animal and the robot would benefit from mechanisms, such as efficient damping, that can act to minimise this noise at source.

Work in progress

We are currently working on a range of extensions to the robot control architecture. These include systems for tactile spatial-mapping (mod-

elled on the rat hippocampus), for feature detection (modelled on S1 cortex), and for improved decision-making by BG. Longer-term we also plan to explore the role whisker motor cortex (MCx) in initiating and modulating whisking bouts, and of further cortical areas involved in somatosensory processing such as the S2 cortex. The cerebellum is known to be involved in a number of distinct loops with the vibrissal system that may each implement a similar computation (adaptive filtering) but in a context that fulfills a quite different role with respect to the behaviour of the animal. For instance, one hypothesis that we wish to explore is that the cerebellum may be involved in predictive tracking of a moving target that would allow a whiskered predator to track rapidly moving prey using only their vibrissae. Another is that cerebellum may be involved in the tuning of open-loop motor control operations, such as the orient action implemented on the robot.

0.4.3 Scratchbot Experiments

Observing general SCRATCHbot behaviour (locate and orient), we typically use 30-60 second runs in a featureless flat arena, with one or more obstacles, or an experimenter’s hand. On startup, a behaviour called **unpark** bids strongly to the BG, and is given control of the robot, whilst other actions (and whisking) are suppressed. The neck axes are driven such that the head is moved from its ‘park’ position to an unstable point we call the ‘unpark’ position. On arrival, **unpark** stops bidding, the default **explore** behaviour is given control, and whisking begins. At the preset end time, software control ceases, and the low-level controllers automatically return the neck to the park position.

Figure 0.9, and supplementary video 6, illustrate what happens when contact is made by one or more of the macrovibrissae. The deflection of the whisker due to the whisker coming into contact with the experimenter’s hand (Figure 0.9 Frame 1) causes the salience of **orient** to increase. If the salience of **orient** is held for long enough (a few tens of milliseconds), the BG switches, selecting the new action by inhibiting the output of **explore** and disinhibiting the output of **orient**. The orient action pattern consists of two phases: first, the snout is oriented to the point of contact; second, poise is maintained for a period of time suitable for fine-scale exploratory behaviour with the microvibrissae (currently being implemented). Note that the salience of the second phase of **orient** is lower than that of the first; thus, the robot can be more easily interrupted whilst exploring the object than whilst completing

the orient itself. When **orient** completes, it stops bidding for the plant, **explore** is again selected, and the robot straightens up as it resumes its exploration. The removal of motor reafferent noise is very effective, and SCRATCHbot does not express ghost orients.

Currently, as described (and shown in supplementary video 5 for WhiskerBot, supplementary video 6 for SCRATCHbot), locomotion during robot exploration results from switching between **explore** and **orient**. As described in Section 0.2.2, we hypothesise that rat locomotion in similar circumstances might be viewed as a series of orients, with the focus of attention being constantly shifted, often ahead of the animal. In future, we will test this hypothesised approach to locomotion using SCRATCHbot, removing **explore** and generalising **orient**, such that the robot can ‘attend to’ (thus, orient to) a more general target, which may not be something immediately detected as interesting, but rather a location about which it intends to gather sensory information.

We have also tested the effect of the addition of control of whisker spread on the nature of signals collected by the whiskers. In Whiskerbot, we showed that MI implemented on each side of the face effectively cleaned up contacts between a single whisker and the environment. In SCRATCHbot, we were able to demonstrate that per-column MIMC was effective in (a) cleaning up contacts on multiple whiskers and (b) generating more whisker-environment contacts than would otherwise have occurred. An example of this new version of MIMC is illustrated in Figure 0.10 and supplementary video 7. For this experiment we fixed the robot head in a position facing a stationary ‘wall’ (similar to that typically recorded using the experimental set-up described by Grant *et al.* (Grant et al., 2009)). In each trial, the first whisk against the wall is unmodulated (left upper panel), and shows that the more rostral column whiskers are heavily deflected, whilst the most caudal column whiskers do not touch the wall at all. On the second whisk, the MIMC modulation has taken effect (right upper panel), the rostral whiskers are less protracted and thus are deflected less heavily, and the caudal whiskers are brought forward and generate contacts with the wall (i.e. whisker ‘spread’ is reduced). The effect of this modulation on the contact signals collected can be seen in the lower panel of the figure—across the three columns, there is a tendency towards normalisation of contact depth.

0.5 Discussion

0.5.1 Why build robot models of animals and their nervous systems?

When neuroscientists think of model systems they usually think, not of robots, but of using one animal as a model of another (e.g. the rat as a model of a human); of an animal in an altered condition (e.g. under anaesthesia, or awake but immobilized) as a model of the awake, behaving one; or of an isolated part, such as a brain slice or a muscle, as a model of that component within the intact, functioning system. However, consider the task that our research group is attempting which is to understand the whisker system of the rat. Some of the biological models that are available to us, the properties that they share with our target system, and their amenability to experimental investigation of their internal processes are summarized in Table 0.1.

System		Properties				Amenable to experimental investigation of internal processes
		Similar substrates	Intact & integrated	Closed loop	Capable of active control	
Awake, free-moving rat (Target organism)		*****	*****	*****	*****	*
Animal models	Awake, restrained rat	*****	*****	*****	***	**
	Anaesthetized rat	*****	*****	***	*	***
	Rat brain slice	*****	*	**	*	****
<i>Whiskered robot</i>		*	*****	*****	*****	*****

Table 0.1 *Comparison of material model of the awake, free moving rat.*

The right-hand column of table 0.1 qualitatively illustrates the difficulty of studying the internal processes of the mammalian nervous system in each of the available models. This is hardest to do in the awake, freely-moving animal (the behaving rat as a model of itself) and becomes progressively easier as we move to the restrained preparation, then anaesthetized, then to an isolated brain slice. Despite recent advances in embedded, wireless chronic recording systems, that are making more naturalistic experiments possible, access to the neural processes of unrestrained behaving animals will always be very limited. The problem, however, as illustrated in the other columns of the table, is that as we move from the target animal to the more amenable biological models, we progressively lose many of the interesting properties that we wish to

understand—how the free-moving animal controls its head and body so as to optimize sensing; how sensorimotor loops are closed through the environment; and how the component parts operate within an integrated and intact system. The more restricted animal preparations can still be useful, of course, but insight concerning these properties will need to be inferred where they cannot be obtained directly. As a consequence we may then ask how the compromises made to create a particular model (e.g. treatment with anaesthesia) have impacted on the results and on the inferences drawn.

Now, examine the bottom line of table 0.1. All of the interesting properties just listed are, or could be, exhibited in a robot model of the behaving animal in a system that is highly amenable to investigation of its internal processes, indeed, far more so than any of the available animal models. The snag, and of course it's a big one, is that the robot's physical and computational substrates—its body and its brain—may not approximate the target animal in an adequate way. Thus the results obtained from experiments with the robot may not be valid when translated back to the animal. But note that this is not too different from the situation with the animal models—in both cases, animal and robot, we have had to compromise and allow some aspects of the model to differ significantly from our target. What we have with robotics, at least, is the possibility that we can continue to refine the model, so that if we are worried that some aspect is substantially wrong we can fix it and re-run the experiment to see if the result changes. As the biomimicry improves, we can expect richer and more convincing insights into the properties of the biological system we are trying to understand.

But how good does the biomimicry need to be to make the model useful? According to the pioneer of cybernetics Norbert Wiener 'the best material model of a cat is another, or preferably the same, cat' ((Rosenblueth and Wiener, 1945) p. 320). But this is 'best' in the sense of most accurate, not most useful. That is, if you want to be really fussy about the fidelity of the model compared to the target, then you will end up coming back to the original system (the same cat, or, in our case, the awake, free moving rat) as the only model that is good enough. In practice, and this is exactly the point that Wiener was making, all systems that are useful as models will only approximate their target. Indeed, Rosenblueth and Wiener (Rosenblueth and Wiener, 1945) suggested that, for a useful model, sufficient and necessary conditions are, not that it should be accurate in every detail (some degree of abstraction is decidedly a good thing), but that it should translate the problem

into a domain we understand better, and it should allow us to conduct experiments with relative ease. We contend that—for exploring the relationships between brain, body, and behaviour—robotics, which translates problems in biology into problems in computing and engineering, meets these requirements well.

0.5.2 Insights from whiskered robots into the biology of the rat vibrissal system

We would like to conclude by reviewing some of the contributions that we believe robotics can make to neuroethology, illustrating each with an example from our research on whiskered robots.

Discovering important questions

We have discovered, as the consequence of our collaboration, that the engineer's mind-set is really very different from that of the experimentalist, and that both can gain from the interchange of ideas and experience. When an engineer is asked to design a robot that mimics some aspect of animal behaviour, his or her first questions are likely to include 'What is your specification for this robot?' and 'What would you like it to do?'. The biologist is then likely to reply that they can specify some aspects of the design, and of the desired behaviour, but not all, since much of the biology is still unknown at this stage. Further, of those things they can specify, much of what they can tell you is approximate or qualitative and some of it is disputed. At this point the engineer may wonder if the problem is adequately posed! If both sides are still keen to pursue the collaboration the experimentalist might then ask what is it that the engineer needs to know in order to create a sufficient specification and the ensuing list of questions can then be integrated into the ongoing programme of empirical work. Often the questions that the engineer will ask are very different from those that have hitherto been addressed by any experiment. This is because the task of synthesis differs radically from that of analysis (Braitenberg, 1986). Synthesis often imposes an order on design decisions which makes the answers to some questions more important, or at least more urgent, than others. For instance, when we began to design our first whisking robot one of the first questions from the engineers was 'What happens to the whiskers once they touch a surface - do they keep moving and bend against the object, or do they stop?'. A simple enough question, but when we looked at the experimental literature that existed at that time there was no clear answer.

Thus began a programme of experimental work that showed that, during exploratory whisking, protraction ceases rapidly following contact, leading to our hypothesis of the ‘minimal impingement’ whisker control strategy (Mitchinson et al., 2007).

Testing the sufficiency of theories

Many good reasons for building models, in simulation or in robots, have been nicely summarized by Epstein (Epstein, 2008). One of the most important of these is that a model allows us determine whether our theories are adequate to account for the behaviour we are trying to explain. If the robot can replicate the behaviour of the animal then there is no longer any question of whether the theory is sufficient (although we can still ask whether it does the task as well as the animal or in the same way). As it turns out, some of the tasks that look easy are harder to replicate than you would expect, and some apparently hard tasks turn out to be relatively easy. For instance, orienting to a whisker contact sounds easy, but when we implemented it we found that the robot was easily distracted by ghost contacts generated by its own movement. Resolving this problem using an adaptive filter was non-trivial and suggested a new theory about the role of vibrissal loops through the cerebellum. A theory of vibrissal orienting without noise cancellation was an insufficient theory, but this was not apparent until we built the robot model and tried it out.

Suggesting new hypotheses

The task of devising an effective solution to whisker-guided orienting has also suggested a new hypothesis concerning the representations of vibrissal stimuli in the midbrain superior colliculus (SC). In mammals, the SC is known to be a seat of spatially-organized sensory information, and to drive rapid orienting and aversion responses (Dean et al., 1989). It is also well evidenced that SC is an important locus for cross-sensory integration of spatial information (Sparks, 1986). Recent histological and electrophysiological measurements indicate that, in the rat, there are strong neural projections from the brainstem trigeminal nuclei to the intermediate layers of SC (Kleinfeld et al., 1999), providing a substrate for whisker-visual integration. In visual animals, SC uses a retino-centric coordinate system; in the rat, retino-centric and head-centric are very closely related, since eye movement is minimal. If spatial information from the whiskers is also integrated in SC, the coordinate transform modelled in Whiskerbot (from whisker to head-centered reference frame)

must take place either in SC itself, or on the way to it (i.e. in the trigeminal nuclei). Whisker representations in SC in the anaesthetized animal are very broad (Drager and Hubel, 1976), particularly in the AP direction, with individual whiskers having representations almost as large as the whisker rows to which they belong. From our modeling work we contend that these broad, AP-biased, fields are to be expected. Specifically, if accurate integration is to be performed between visual and vibrissal stimuli, then since whiskers move relative to the head as the result of whisking, their representations must also move around with respect to the head-centered map in SC, leading to broad overall fields. Recording from multiple cells in an animal whilst the whiskers move, it should be possible to show that the strongest response to individual whiskers moves around in SC in a manner consistent with this hypothesis. Whilst we can only be confident of this prediction in the awake animal, it might also be observed in a lightly anaesthetised animal, with whisking invoked electrically.

Investigating the role of embodiment

It is increasingly recognized that behaviour is the consequence of the interaction between the brain, the body, and the environment (Chiel and Beer, 1997; Chiel et al., 2009) and that robotics may be one of the most cost-effective ways of studying this interaction. Current simulations of real-world physics are not up to the task of adequately capturing all of the dynamical properties (e.g. collision elasticity, surface friction) of the interaction between two objects such as a whisker shaft and a moving or irregular surface (Fox et al., 2008), thus the world remains its own best model for this aspect of our work. As noted above, our recent efforts to build artificial vibrissae fall short of producing artefacts with all the desirable mechanical properties of the rat whisker-shaft/follicle. Therefore, our current approach combines simulation, to better understand the biomechanics of natural whiskers, with robotic experiments to determine what properties of surfaces can be effectively discerned through artificial vibrissal sensing. Through this combined approach we hope to be able to better identify the contribution of the morphology to the task of tactile sensing for this system.

Investigating ‘What if?’ scenarios

Robots do not have to mimic the biology, and one way to find out why the biology is-as-it-is might be to build a system that works differently and compare. For instance, we can directly contrast the sensing properties

of whiskers that do not taper or do not curve with those that do, in order to better understand why natural whiskers have evolved to do both. Likewise, we can investigate ways of controlling the movement of the whiskers that are either natural or unnatural and observe the consequences of these different movement strategies for the type, quality and quantity of tactile information obtained by the robot (animal).

Doing the experiments that cannot be done

One of the most important uses of a physical model is that it can allow experiments to be performed that could not be done on the animal, or at least would be very difficult to do *in vivo* given our current knowledge and available methods. For example, what impact does minimal impingement (MI) control of vibrissal movement (Mitchinson et al., 2007) have on the sensory signals ascending from the whiskers to higher processing centers such as the S1 cortex? This question is important because much of the research conducted on cortex is performed in anaesthetized animals, or even in brain slices, using input signals that are intended to mimic the effects of natural contacts. Unfortunately we cannot perform a straightforward experiment in the animal where we turn-off MI control and see the difference that this makes to signals (in the anaesthetized preparation we can replace natural whisker motion with a form of fictive whisking (Szwed et al., 2003), however, there are differences here from the awake, behaving animal that could make interpretation of the findings problematic). In contrast, in the robot model, the required experiment is trivial to perform (see 0.6), and we can exactly compare the contacts resulting from movement, with and without MI, and the spike trains that these contacts generate in the simulated trigeminal nerve (which, in the animal, is just two synapses away from the signals arriving at the S1 cortex). These robot experiments can therefore help neuroscientists to select more plausible signal trains with which to stimulate animal models during *in vivo* or *in vitro* experiments.

Evaluating the usefulness of biological solutions to problems in robotics

Finally, one of the benefits for engineers of engaging in this kind of interdisciplinary collaboration is to determine whether the biological system has a solution to an existing practical problem that they might usefully copy. The rat is a successful and versatile mammal that uses its whiskers in a range of tasks from object detection and recognition, through guidance of locomotion across all kinds of terrain, to prey capture (Prescott

et al., In press). In mammals more generally, tactile hairs are widely deployed across the body (e.g. on the legs, paws, or back) for detecting unexpected contacts and for tactile discrimination of surface properties. A sensory capacity that has proved so effective for animals could lead to increased flexibility and performance when deployed on robots. Indeed an artificial whisker system could prove particularly useful for robots that must operate in environments where vision systems can provide only degraded or ambiguous input, such as in smoke- or dust-filled buildings or for covert operations in darkness. Towards this end, the next stage in the development of SCRATCHbot will be to devise tactile-based strategies for environment exploration and local and global navigation. For instance, we plan to incorporate a spatial memory based on current understanding of the rat hippocampal formation. By associating odometry and head-centered contact information with tactile sensory features, such as texture and object shape, a spatial map will be constructed as the robot explores. For the biologists, this will allow investigation into how tactile sensory information is presented to long-term memory systems. For the roboticists, this will provide valuable insight into how a touch-based platform could be effectively deployed for robot guidance in the absence of vision.

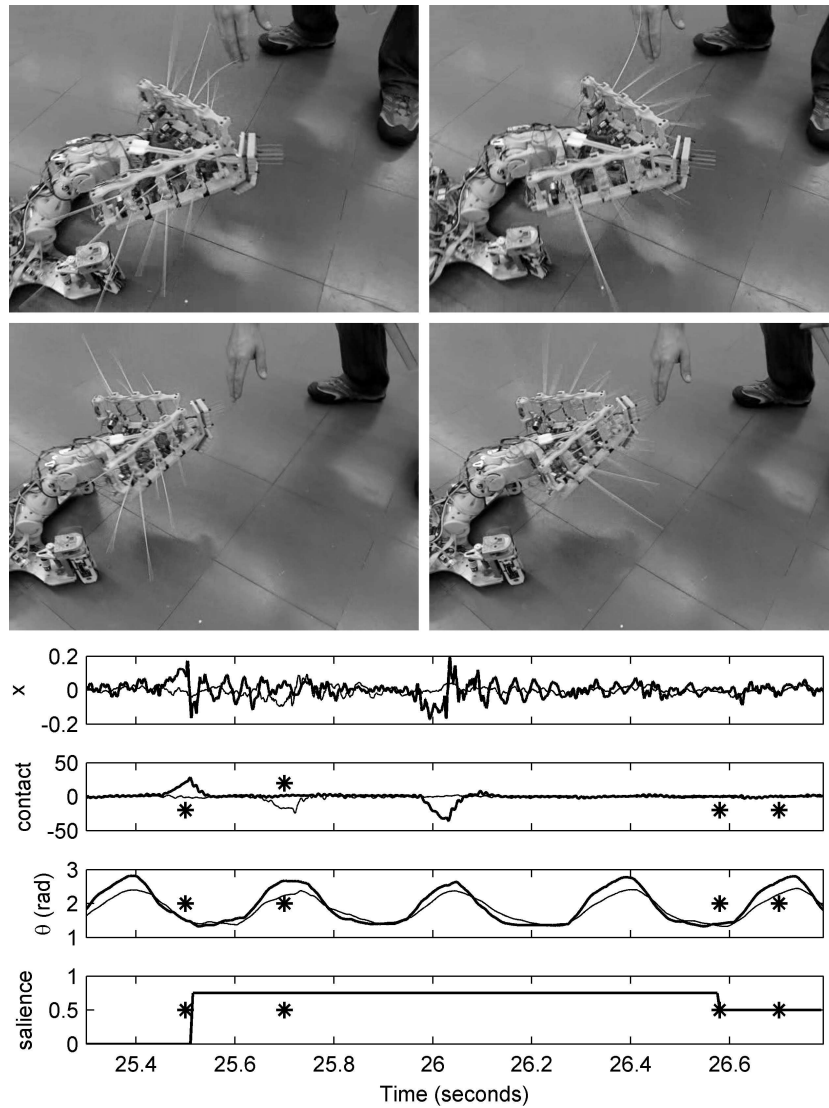


Figure 0.9 Data from SCRATCHbot during an orient to a whisker contact. Upper three plots show data from 5th whisker (middle row/column, thick line) and 7th whisker (top row, rear column, thin line) on the left. Upper: raw x-component of whisker sensory signal; Second: re-afferent noise removed, greatly improving the signal-to-noise ratio; Third: angular position (θ) of middle/rear columns (π radians indicates straight ahead); Lower: saliency of **orient**. Frame timing is indicated in plots by star symbol. Contact on whisker 5 during retraction (Frame 1) is followed by increase in **orient** saliency and action selection. Additional contacts on whisker 7 (Frame 2) and 5 during the orient are ignored. Snout arrives at the point of initial whisker contact (Frame 3), completing the orient; saliency is reduced. The second phase of **orient** (Frame 4) is a placeholder, during which the micro-vibrissae will be used for fine inspection of the contacted feature, in future work.

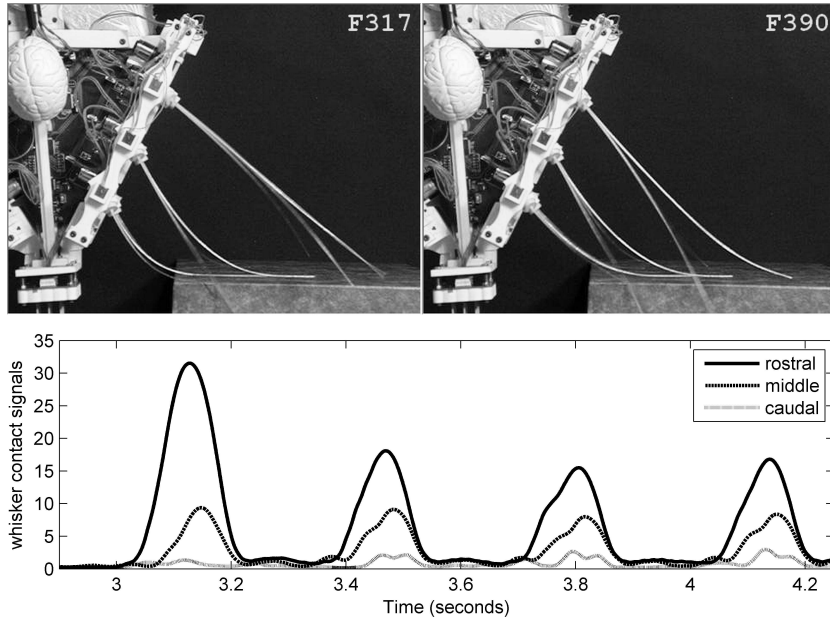


Figure 0.10 Investigating the impact of per-column MIMC on whisking patterns and data collection (see text). Upper panel shows two frames taken from a highspeed video of a trial of whisking against a wall. The frames are taken from the moment of peak whisker protraction in the first whisk (left) and in the second (right); in the second whisk, modulation has taken effect, and the contacts of each column with the wall are normalised. Lower panel shows the contact signals retrieved from the centre whisker in each of the three columns, during the first four whisks of the same trial—signals are normalised in second, and later, whisks.

References

- Ahissar, E., and Arieli, A. 2001. Figuring Space by Time. *Neuron*, **32**(2), 185–201.
- Anderson, S. R., Porrill, J., Pearson, M. J., Pipe, A. G., Prescott, T. J., and Dean, P. 2009. *Cerebellar-inspired forward model of whisking enhances contact detection by vibrissae of robot rat*. Society For Neuroscience, Chicago IL, Poster 77.2.
- Anjum, F., Turni, H., Mulder, P. G., van der Burg, J., and Brecht, M. 2006. Tactile guidance of prey capture in Etruscan shrews. Pages 16544–16549 of: *Proc Natl Acad Sci USA*, vol. 103.
- Berg, R. W., and Kleinfeld, D. 2003. Rhythmic Whisking by Rat: Retraction as Well as Protraction of the Vibrissae Is Under Active Muscular Control. *J Neurophysiol*, **89**(1), 104–117.
- Bermejo, R., Vyas, A., and Zeigler, H. P. 2002. Topography of Rodent Whisking—I. Two-dimensional Monitoring of Whisker Movements. *Somatosens Mot Res*, **19**(4), 341–346.
- Birdwell, J. A., Solomon, J. H., Thajchayapong, M., Taylor, M. A., Cheely, M., Towal, R. B., Conradt, J., and Hartmann, M. J. Z. 2007. Biomechanical Models for Radial Distance Determination by the Rat Vibrissal System. *J Neurophysiol*, **98**, 2439–2455.
- Braitenberg, Valentino. 1986. *Vehicles: Experiments in synthetic psychology*. The MIT Press.
- Brecht, M., Preilowski, B., and Merzenich, M. M. 1997. Functional Architecture of the Mystacial Vibrissae. *Behavioural Brain Research*, **84**, 81–97.
- Carvell, G. E., and Simons, D. J. 1990. Biometric Analyses of Vibrissal Tactile Discrimination in the Rat. *J Neurosci*, **10**(8), 2638–2648.
- Carvell, G. E., Simons, D. J., Lichtenstein, S. H., and Bryant, P. 1991. Electromyographic Activity of Mystacial Pad Musculature During Whisking Behavior in the Rat. *Somatosens Mot Res*, **8**(2), 159–164.
- Chiel, H J, and Beer, R D. 1997. The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, **20**(12), 553–7.
- Chiel, Hillel J, Ting, Lena H, Ekeberg, Orjan, and Hartmann, Mitra J Z.

2009. The brain in its body: motor control and sensing in a biomechanical context. *J Neurosci*, **29**(41), 12807–14.
- Cohen, J. D., Hirata, A., and Castro-Alamancos, M. A. 2008. Vibrissa Sensation in Superior Colliculus: Wide-Field Sensitivity and State-Dependent Cortical Feedback. *J Neurosci*, **28**(44), 11205–11220.
- Cramer, N. P., Li, Y., and Keller, A. 2007. The Whisking Rhythm Generator: A Novel Mammalian Network for the Generation of Movement. *J Neurophys*, **97**, 2148–2158.
- Dean, P., Redgrave, P., and Westby, G W. 1989. Event or emergency? Two response systems in the mammalian superior colliculus. *Trends in Neurosciences*, **12**(4), 137–47.
- Dean, P., Porrill, J., Ekerot, C. F., and Jörntell, H. 2010. The cerebellar micro-circuit as an adaptive filter: experimental and computational evidence. *Nat Rev Neurosci*, **11**(1), 30–43.
- Dörfl, J. 1982. The Musculature of the Mystacial Vibrissae of the White Mouse. *J Anat*, **135**(Pt 1), 147–154.
- Drager, U. C., and Hubel, D. H. 1976. Topography of visual and somatosensory projections to mouse superior colliculus. *J Neurophys*, **39**(1), 91–101.
- Ebara, S., Kumamoto, K., Matsuura, T., Mazurkiewicz, J. E., and Rice, F. L. 2002. Similarities and Differences in the Innervation of Mystacial Vibrissal Follicle-sinus Complexes in the Rat and Cat: A Confocal Microscopic Study. *J Comp Neurol*, **449**(2), 103–119.
- Elumotion. 2009. *Biologically inspired robotic devices (website)*. <http://www.elumotion.com>.
- Epstein, JM. 2008. Why Model? *Journal of Artificial Societies and Social Simulation*, **11**(4), 12.
- Etienne, A. S., Maurer, R., and Séguinot, V. 1996. Path integration in mammals and its interaction with visual landmarks. *J Exp Biol*, **199**(1), 201–9.
- Evans, M., Fox, C. W., Pearson, M., and Prescott, T. J. 2008. *Radial distance to contact estimation from dynamic robot whisker information*. Barrels XX (Society For Neuroscience Satellite Meeting).
- Evans, M., Fox, C. W., Pearson, M. J., and Prescott, T. J. 2009. Spectral template based classification of robotic whisker sensor signals in a floor texture discrimination task. Pages 19–24 of: *Proceedings of Towards Autonomous Robotic Systems (TAROS)*.
- Fend, M., Bovet, S., and Hafner, V. 2004. The Artificial Mouse - A Robot with Whiskers and Vision. In: *Proceedings of the 35th International Symposium on Robotics*.
- Fox, C., Evans, M., Stone, J., and Prescott, T. 2008. Towards temporal inference for shape recognition from whiskers. In: *Proceedings of Towards Autonomous Robotic Systems (TAROS)*.
- Fox, C. W., Mitchinson, B., Pearson, M. J., Pipe, A. G., and Prescott, T. J. 2009. Contact type dependency of texture classification in a whiskered mobile robot. *Autonomous Robots*, **26**(4), 223–239.

- Gao, P., Bermejo, R., and Zeigler, H. P. 2001. Whisker Deafferentation and Rodent Whisking Patterns: Behavioral Evidence for a Central Pattern Generator. *J Neurosci*, **21**(14), 5374–5380.
- Gopal, V., and Hartmann, M. J. Z. 2007. Using hardware models to quantify sensory data acquisition across the rat vibrissal array. *Bioinsp Biomim*, **2**, S135–145.
- Grant, Robyn A., Mitchinson, Ben, Fox, Charles W., and Prescott, Tony J. 2009. Active touch sensing in the rat: Anticipatory and regulatory control of whisker movements during surface exploration. *Journal of Neurophysiology*, **101**, 862–874.
- Gurney, K. N., Humphries, M., Wood, R., Prescott, T. J., and Redgrave, P. 2004. Testing computational hypotheses of brain systems function: a case study with the basal ganglia. *Network*, **15**(4), 263–90.
- Haiss, F., and Schwarz, C. 2005. Spatial Segregation of Different Modes of Movement Control in the Whisker Representation of Rat Primary Motor Cortex. *J Neurosci*, **25**(6), 1579–1587.
- Hartmann, M. J. 2001. Active Sensing Capabilities of the Rat Whisker System. *Autonomous Robots*, **11**, 249–254.
- Hartmann, M. J., Johnson, N. J., Towal, R. B., and Assad, C. 2003. Mechanical Characteristics of Rat Vibrissae: Resonant Frequencies and Damping in Isolated Whiskers and in the Awake Behaving Animal. *J Neurosci*, **23**(16), 6510–6519.
- Hill, D. N., Bermejo, R., Zeigler, H. P., and Kleinfeld, D. 2008. Biomechanics of the Vibrissa Motor Plant in Rat: Rhythmic Whisking Consists of Triphasic Neuromuscular Activity. *J Neurosci*, **28**(13), 3438–3455.
- Jung, D., and Zelinsky, A. 1996. Whisker-based mobile robot navigation. Pages 497–504 of: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 2.
- Kim, D., and Möller, R. 2004. A Biomimetic Whisker for Texture Discrimination and Distance Estimation. In: *Proc of the Int Conf on the Sim of Adap Behav (SAB)*.
- Kim, D., and Möller, R. 2007. Biomimetic whiskers for shape recognition. *Robot Autonom Syst*, **55**(3), 229–243.
- Kleinfeld, D., Berg, R. W., and O’Connor, S. M. 1999. Anatomical Loops and Their Electrical Dynamics in Relation to Whisking by Rat. *Somatosens Mot Res*, **16**(2), 69–88.
- Kleinfeld, D., Ahissar, E., and Diamond, M. E. 2006. Active Sensation: Insights from the Rodent Vibrissa Sensorimotor System. *Curr Opin Neurobiol*, **16**(4), 435–44.
- Knutsen, P., Biess, A., and Ahissar, E. 2008. Vibrissal Kinematics in 3D: Tight Coupling of Azimuth, Elevation, and Torsion. *Neuron*, **59**(1), 35–42.
- Knutsen, P. M., Pietr, M., and Ahissar, E. 2006. Haptic Object Localization in the Vibrissal System: Behavior and Performance. *J Neurosci*, **26**(33), 8451–8464.
- Krupa, D. J., Matell, M. S., Brisben, A. J., Oliveira, L. M., and Nicolelis, A. L. 2001. Behavioral Properties of the Trigeminal Somatosensory System in

- Rats Performing Whisker-dependent Tactile Discriminations. *J Neurosci*, **21**(15), 5752–5763.
- Lichtenstein, S. H., Carvell, G. E., and Simons, D. J. 1990. Responses of Rat Trigeminal Ganglion Neurons to Movements of Vibrissae in Different Directions. *Somatosens Mot Res*, **7**(1), 47–65.
- Lottem, E., and Azouz, R. 2009. Mechanisms of Tactile Information Transmission through Whisker Vibrations. *J Neurosci*, **29**(37), 11686–11697.
- Mitchinson, B., Gurney, K. N., Redgrave, P., Melhuish, C., Pipe, A. G., Pearson, M., Gilhespy, I., and Prescott, T. J. 2004. Empirically Inspired Simulated Electro-mechanical Model of the Rat Mystacial Follicle-sinus Complex. *Proc R Soc Lond B Biol Sci*, **271**(1556), 2509–2516.
- Mitchinson, Ben, Pearson, Martin, Melhuish, Chris, and Prescott, Tony J. 2006. A model of sensorimotor co-ordination in the rat whisker system. Pages 77–88 of: *From Animals to Animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour*.
- Mitchinson, Ben, Martin, Chris J., Grant, Robyn A., and Prescott, Tony J. 2007. Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact. *Royal Society Proceedings B*, **274**(1613), 1035–1041.
- Mitchinson, Ben, Arabzadeh, Ehsan, Diamond, Mathew E., and Prescott, Tony J. 2008. Spike-timing in primary sensory neurons: A model of somatosensory transduction in the rat. *Biological Cybernetics*, **98**(3), 185–194.
- Mitchinson, Ben, Chan, Tak-Shing, Chambers, Jon, Pearson, Martin, Humphries, Mark, Fox, Charles, Gurney, Kevin, and Prescott, Tony J. 2010. BRAHMS: Novel middleware for integrated systems computation. *Advanced Engineering Informatics*, **24**(1), 49–61.
- Nallatech. 2007. *PCI-104 Series: COTS PCI-104 FPGA computing stacks product brief (NT190-0331 Version 1.0)*. available on-line: <http://www.nallatech.com/mediaLibrary/images/english/6406.pdf>.
- Neimark, M. A., Andermann, M. L., Hopfield, J. J., and Moore, C. I. 2003. Vibrissa Resonance as a Transduction Mechanism for Tactile Encoding. *J Neurosci*, **23**(16), 6499–6509.
- Nguyen, Q. T., and Kleinfeld, D. 2005. Positive Feedback in a Brainstem Tactile Sensorimotor Loop. *Neuron*, **45**(3), 447–457.
- Pearson, M., Nibouche, M., Pipe, A. G., Melhuish, C., Gilhespy, I., Mitchinson, B., Gurney, K., Prescott, T., and Redgrave, P. 2006a. A Biologically Inspired FPGA Based Implementation of a Tactile Sensory System for Object Recognition and Texture Discrimination. Pages 1–4 of: *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL)*.
- Pearson, M., Nibouche, M., Gilhespy, I., Gurney, K., Melhuish, C., Mitchinson, B., and Pipe, A. G. 2006b. A Hardware Based Implementation of a Tactile Sensory System For Neuromorphic Signal Processing Applications. Pages IV–IV of: *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 4.

- Pearson, M. J., Melhuish, C., Pipe, A. G., Nibouche, M., Gilhespy, I., Gurney, K., and Mitchinson, B. 2005. Design and FPGA implementation of an embedded real-time biologically plausible spiking neural network processor. Pages 582–585 of: *International Conference on Field Programmable Logic and Applications*.
- Pearson, M. J., Pipe, A. G., Mitchinson, B., Gurney, K., Melhuish, C., Gilhespy, I., and Nibouche, M. 2007a. Implementing spiking neural networks for real-time signal-processing and control applications: A model-validated FPGA approach. *IEEE Trans on Neural Networks*, **18**(5), 1472–1487.
- Pearson, M. J., Pipe, A. G., Melhuish, C., Mitchinson, B., and Prescott, T. J. 2007b. Whiskerbot: A Robotic Active Touch System Modeled on the Rat Whisker Sensory System. *Adaptive Behaviour*, **15**(3), 223–240.
- Petersen, C. C. H. 2007. The Functional Organization of the Barrel Cortex. *Neuron*, **56**, 339–355.
- Prescott, T. J., Mitchinson, B., Redgrave, P., Melhuish, C., and Dean, P. 2005. *Three-dimensional reconstruction of whisking patterns in freely moving rats*. Society For Neuroscience, November 12–16, Poster 625.3.
- Prescott, T. J., Gonzalez, F. M. Montes, Gurney, K., Humphries, M. D., and Redgrave, P. 2006. A Robot Model of the Basal Ganglia: Behavior and Intrinsic Processing. *Neural Netw*, **19**(1), 31–61.
- Prescott, T. J., Bryson, J. J., and Seth, A. K. 2007. Introduction. Modelling natural action selection. *Philos Trans R Soc Lond B Biol Sci*, **362**(1485), 1521–9.
- Prescott, T. J., Pearson, M. J., Mitchinson, B., Sullivan, J. C. W., and Pipe, A. G. 2009. Whisking with robots: from rat vibrissae to biomimetic technology for active touch. *IEEE Robotics and Automation Magazine*, **16**(3), 42–50.
- Prescott, T. J., Grant, R. A., and Mitchinson, B. In press. Vibrissal function and behavior. *Scholarpedia*.
- Quist, B. W., and Hartmann, M. J. 2008. A two-dimensional force sensor in the millinewton range for measuring vibrissal contacts. *J Neurosci Methods*, **172**(2), 158–167.
- Redgrave, P., Prescott, T. J., and Gurney, K. 1999. The Basal Ganglia: a Vertebrate Solution To the Selection Problem?. *Neuroscience*, **89**(4), 1009–1023.
- Rice, F. L., Mance, A., and Munger, B. L. 1986. A Comparative Light Microscopic Analysis of the Sensory Innervation of the Mystacial Pad. I. Innervation of Vibrissal Follicle-sinus Complexes. *J Comp Neurol*, **252**(2), 154–174.
- Rosenblueth, A., and Wiener, N. 1945. The role of models in science. *Philosophy of Science*, **12**(4).
- Russell, R. A. 1992. Using tactile whiskers to measure surface contours. Pages 1295–1299 of: *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2.
- Russell, R. A., and Wijaya, J. A. 2005. Recognising and manipulating objects using data from a whisker sensor array. *Robotica*, **23**(5), 653–664.

- Sachdev, R. N., Berg, R. W., Champney, G., Kleinfeld, D., and Ebner, F. F. 2003. Unilateral Vibrissa Contact: Changes in Amplitude but Not Timing of Rhythmic Whisking. *Somatosens Mot Res*, **20**(2), 163–169.
- Sahibzada, N., Dean, P., and Redgrave, P. 1986. Movements resembling orientation or avoidance elicited by electrical stimulation of the superior colliculus in rats. *J Neurosci*, **6**(3), 723–733.
- Seth, A. K., McKinstry, J. L., Edelman, G. M., and Krichmar, J. L. 2004a. Active sensing of visual and tactile stimuli by brain-based devices. *J Robot Automat*, **19**(4), 222–238.
- Seth, A. K., McKinstry, J. L., Edelman, G. M., and Krichmar, J. L. 2004b. Texture Discrimination by an Autonomous Mobile Brain-based Device with Whiskers. In: *Proc of the IEEE Int Conf on Robotics and Automation (ICRA)*.
- Shoykhet, M., Doherty, D., and Simons, D. J. 2000. Coding of Deflection Velocity and Amplitude by Whisker Primary Afferent Neurons: Implications for Higher Level Processing. *Somatosens Mot Res*, **17**(2), 171–180.
- Solomon, J. H., and Hartmann, M. J. 2006. Biomechanics: Robotic whiskers used to sense features. *Nature*, **443**(7111), 525.
- Sparks, D L. 1986. Translation of sensory signals into commands for control of saccadic eye movements: role of primate superior colliculus. *Physiological Reviews*, **66**(1), 118–71.
- Stüttgen, M. C., Kullmann, S., and Schwarz, C. 2008. Responses of Rat Trigeminal Ganglion Neurons to Longitudinal Whisker Stimulation. *J Neurophysiol*, **100**, 1879–1884.
- Szwed, M., Bagdasarian, K., and Ahissar, E. 2003. Encoding of Vibrissal Active Touch. *Neuron*, **40**(3), 621–630.
- Towal, R. B., and Hartmann, M. J. 2006. Right-left Asymmetries in the Whisking Behavior of Rats Anticipate Head Movements. *J Neurosci*, **26**(34), 8838–46.
- Ueno, N., and Kaneko, M. 1994. Dynamic active antenna - a principle of dynamic sensing. Pages 1784–18790 of: *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2.
- Waite, P. M. E., and Tracey, D. J. 1995. Trigeminal Sensory System. Pages 705–724 of: Paxinos, G. (ed), *The Rat Nervous System*. Academic Press.
- Welker, W. I. 1964. Analysis of Sniffing of the Albino Rat. *Behaviour*, **22**, 223–244.
- Wilson, J. F., and Chen, Z. 1995. A Whisker Probe System for Shape Perception of Solids. *J Dyn Sys Meas Control*, **117**(1), 104–108.
- Wineski, L. E. 1985. Facial morphology and vibrissal movement in the golden hamster. *J Morphol*, **183**, 199–217.
- Yokoi, H., Lungarella, M., Fend, M., and Pfeifer, R. 2005. Artificial whiskers: Structural characterization and implications for adaptive robots. *Journal of Robotics and Mechatronics*, **7**(5), 584–595.
- Zucker, E., and Welker, W. I. 1969. Coding of Somatic Sensory Input by Vibrissae Neurons in the Rat’s Trigeminal Ganglion. *Brain Res*, **12**(1), 138–156.