Book Section:

Prescott, A.J., Mitchinson, B., Lepora, N.F. et al. (8 more authors) (2015) The robot vibrissal system: Understanding mammalian sensorimotor co-ordination through biomimetics. In: Krieger, P. and Groh, A., (eds.) Sensorimotor Integration in the Whisker System. Springer , New York , pp. 213-240. ISBN 9781493929757

The final publication is available at Springer via
http://dx.doi.org/10.1007/978-1-4939-2975-7

The robot vibrissal system: Understanding mammalian sensorimotor co-ordination through biomimetics

Tony J Prescott, Ben Mitchinson, Nathan F. Lepora, Stuart P. Wilson, Sean R. Anderson, John Porrill, Paul Dean, Charles W. Fox
University of Sheffield

Martin J Pearson, J. Charles Sullivan, Anthony G. Pipe
Bristol Robotics Laboratory

Correspondence: t.j.prescott@sheffield.ac.uk

Summary

We consider the problem of sensorimotor co-ordination in mammals through the lens of vibrissal touch, and via the methodology of embodied computational neuroscience—using biomimetic robots to synthesize and investigate models of mammalian brain architecture. The chapter focuses on five major brain sub-systems and their likely role in vibrissal system function—superior colliculus, basal ganglia, somatosensory cortex, cerebellum, and hippocampus. With respect to each of these we demonstrate how embodied modelling has helped elucidate their likely function in the brain of awake behaving animals. We also demonstrate how the appropriate co-ordination of these sub-systems, with a model of brain architecture, can give rise to integrated behaviour in a life-like whiskered robot.

Keywords: sensorimotor co-ordination, biomimetic robot, embodied computational neuroscience, layered architecture

Citation: In Krieger, P. and Groh, A. (2015). *Sensorimotor Integration in the Whisker System.* Springer.

Research on active vibrissal touch has the potential to help us understand, perhaps even rethink, many of the key computations underlying sensorimotor co-ordination in the mammalian brain.

Consider, for instance, the task of visually-guided reach and grasp which is widely studied in both humans and primates. Work in humanoid robotics might decompose this task as the following steps: (i) identify a potential target in peripheral vision based on a rapid analysis of superficial salient features (colour, shape, movement); (ii) orient to and fixate on the object using foveal vision to form an internal 3-dimensional model of the object and of its key properties (shape, size, texture, and so forth); (iii) in parallel, form a second set of representations of the position and orientation of the object in space relative to those of the body, arm, and hand; (iv) match the first, "what?", model with a variety of stored "templates" in order to determine whether this particular item is, indeed, a suitable target for reaching; (iv) apply algorithms to the computed "where?" representations of the object and body, and make use of knowledge of the kinematic and dynamic properties of the arm, hand, and digits, to determine appropriate movement trajectories; (v) execute the planned movements largely ballistically but using some sensory feedback in the final approach, to locate, enclose, and lift the object in an effective way.

Now consider the capacity of an animal such as the Etruscan shrew, the smallest living terrestrial mammal—and known to be a remarkably efficient predator—to localise, identify, and entrap an agile prey insect using only its whiskers (Brecht, Naumann et al., 2011). The problem is similar in many ways to that of human (or humanoid) sensory-guided reaching. The visual periphery compares to the macrovibrissae (the longer actuated facial whiskers on either side of the snout), and the visual fovea to the microvibrissae (the shorter non-moving whiskers on the upper lip and chin) and other tactile sensory surfaces around the mouth. The orienting system, as in primate vision, is likely to have the superior colliculus at its core, and will be driven by a very rapid but coarse analysis of features in the whisker signals via a midbrain loop that co-ordinates movements of the whiskers, head and trunk. Further analyses of the whisker signals, from both the macro- and micro-vibrissae will involve the somatosensory cortices, and pathways through to the temporal lobes. These will likely involve the decomposition of sensory signals into components (self-motion, object properties such as shape, texture, etc.), but may also require the reintegration of decomposed features into more complete representations of the target. Alongside determination of object properties, information about the prey animal's spatial position and orientation will also have to be computed from the same set of vibrissal deflection signals. The decision of whether to make an attack will then depend on a comparison of computed features with remembered patterns corresponding to previously successful (and unsuccessful) attacks. Whilst this process will likely involve cortical systems (including hippocampus) it will ultimately involve decision-making mechanisms in basal ganglia to decide if the template fits, and, if so, whether the attack option is appropriate right now (as opposed, say, to further approach behaviour or avoidance). Planning and execution of the strike will also involve the motor cortex, and midbrain and brainstem motor systems which, together with the cerebellum, will co-ordinate precision orienting with biting, and may use additional sensory information from the vibrissae to accurately adjust the final phases of the strike.

Despite the above similarities, however, a number of features of the shrew vibrissal system might lead us to think rather differently about this problem from the way we initially conceived our example task of human visually-guided reach.

First, rather than being able to fixate and examine the target at leisure, the animal must make do with signals from a few fleeting contacts between the vibrissal tips and a small number of unknown locations on the target (Munz, Brecht et al., 2010). Further, both the sensors themselves and target are moving rapidly, the latter with unknown direction and speed. In other words, this is a task, where information about the target is relatively sparse, and where timing and dynamics are crucial. The urgency of the required response means that the preparation of attack behaviour will likely occur alongside the processing of vibrissal signals to determine object properties—so that the former can be put into effect as soon as the weight of evidence lies in its favour. In other words, this task is perhaps more similar to the challenge faced by a batter trying to hit a moving ball in fading light—the target object suddenly and rapidly appears out of nowhere, and a successful response must be executed within a critical and narrow time window.

Second, the shrew's brain is tiny (Roth-Alpermann, Anjum et al., 2010). Not only must its predation behaviour be accomplished with 20,000 times fewer neurons that a human might utilise for reach and grasp, we also know from the speed of the attack (which can be as little as 80 milliseconds (Anjum, Turni et al., 2006)) that the shrew achieves its goal in a far smaller number of processing steps. Whatever phases are necessary for decomposition of sensory signals and their reconstruction as object representations, these will necessarily involve a small number of processing sites each made up of relatively few neurons. The construction of complex internal models, for comparison against rich templates, looks decidedly improbable in this system. More likely, key features are rapidly extracted and mapped, across a small number of synapses, into representations of their potential for action. Indeed, the step of "representing" the prey insect itself may even be missed out entirely. This animal may encode information about objects in its tactile world only in terms of their potential as affordances (Gibson, 1979) to guide different forms of approach, avoidance, or consummatory behaviour. Thus this is a system in which we can explore what is the minimum amount of internal transformation and representation needed in order to support complex, sensory-guided behaviour; and in which we can discover how active sensing systems (Prescott, Diamond et al., 2011) merge perception into action, via closed loop control (Ahissar and Kleinfeld, 2003), without the two ever being truly separate.

Overall, then, while understanding this system will not directly answer the question of how the human brain performs reach-and-grasp, the study of vibrissal-guided behaviour could help us understand many aspects of mammalian sensorimotor control and perhaps rethink a number of assumptions based on more primate-centric analyses of brain processing.

In this chapter we consider five major brain sub-systems and their likely role in vibrissal system function—superior colliculus (SC), basal ganglia (BG), primary somatosensory cortex (S1), cerebellum, and hippocampus—bearing in mind the behavioural domain of whisker-guided predation in animals such as the shrew or rat. One of these, the somatosensory cortex, is specialised for tactile processing, but shares many aspects of its computational architecture with other areas of mammalian cortex. The remaining four (basal ganglia, cerebellum, superior colliculus and hippocampus) are more "general purpose" in the sense that they appear to have some characteristic function that operates in a similar way across different sensory modalities or motor functions but that is also tuned, in some appropriate manner, to the particular requirements of processing and control for vibrissal touch. We are therefore hopeful that the insights obtained by studying the role of each of these sub-systems in the vibrissal processing of rodent-like mammals will generalise to understanding their functional capacities in other domains too. Each of these systems is the subject of a vast neuroscientific literature that we cannot even begin to summarise here. We therefore restrict our focus to providing a brief outline of the hypothesized functional role of each system in vibrissal touch and then describe how we have investigated this from an *embodied computational neuroscience* perspective that seeks to develop and test systems-level computational models of neural circuits embedded within the control system of biomimetic robots.

One might ask why we bother to build robot models of animals and their nervous systems? One answer, suggested the neurobiologist Valentino Braitenberg (Braitenberg, 1986), is that synthesis (engineering a model of a biological system) is quite different from analysis (reverse-engineering an existing biological system); thus, in building a robot model of our target animal, that mimics sufficiently some aspects of its body, brain and behaviour, we can expect to learn a good deal about the original creature. Another answer is that a robot model should allow us to conduct experiments that will help us better understand the biological system, and moreover would be impossible (or at least much more difficult) to perform in the original animal (Mitchinson, Pearson et al., 2011; Rosenblueth and Wiener, 1945). Finally, neurobiological studies have shown 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, more generally, co-ordinates sensing with action. Thus, a further reason for building biomimetic robot models is to provide improved insight into brain architecture as a whole. Indeed, by building robotic whisker systems—see examples of our whiskered robots in Figure 1—we consider that we are taking significant steps towards building an integrated robotic model of the mammalian brain.

FIGURE 1 HERE

A control architecture for behavioural integration in vibrissal touch

We begin our consideration of the sensorimotor co-ordination for vibrissal touch by addressing the overall problem of behavioural integration, or behavioural coherence, that is central to the task of building life-like systems (Prescott, 2007). Living, behaving systems display patterns of behaviour that are integrated over space and time such that the animal controls its effector systems in a co-ordinated way, generating sequences of actions that maintain homeostatic equilibrium, satisfy drives, or meet goals. How animals achieve behavioural integration is, in general, an unsolved problem in anything other than some of the simplest invertebrates. What is clear from the perspective of behaviour is that the problem is under-constrained since similar sequences of overt behaviour can be generated by quite different underlying control architectures (Hallam and Malcolm, 1994). This implies that to understand the solution to the integration problem in any given organism is going to require investigation of mechanism in addition to observations of behaviour. In this regard, physical models—such as robots—can prove useful as a means of embodying hypotheses concerning alternative control architectures whose behavioural consequences can then be measured observationally (Mitchinson, Pearson et al., 2011). Research with robots has repeatedly demonstrated forms of emergent behaviour (Mitchinson, Pearson et al., 2012)—the appearance of integrated behavioural sequences that are not explicitly programmed—demonstrating the value of this embodied testing for suggesting and testing candidate mechanisms.

The biological literature provides a range of different hypotheses concerning the mechanisms that can give rise to behavioural integration; here, we highlight two—behavioural and salience map competition.

The neuroethology literature suggests a decomposition of control into behavioural sub-systems that then compete to control the animal (see (Prescott, Redgrave et al., 1999; Redgrave, Prescott et al., 1999) for a review). This approach has been enthusiastically adopted by researchers in behaviour-based robotics (see, e.g. (Brooks, 1991)) as a means of generating integrated patterns of behaviour in autonomous robots that can be robust to sensory noise, or even to damage to the controller.

An alternative hypothesis emerges from the literature on spatial attention, particularly that on visual attention in primates including humans (Gandhi and Katnani, 2011). This approach suggests that actions, such as eye movements and reaches towards targets, are generated by first computing a ‘salience map’ that integrates information about the relevance (salience) to the animal of particular locations in space into a single topographic representation. Some maximisation algorithm is then used to select the most salient position in space towards which action is then directed. It is usual in this literature to distinguish between the computation of the salience map, the selection of the target within the map, and orienting actions that move the animal, or its effector systems, towards the target. In the mammalian brain these different functions may be supported by distinct (though overlapping) neural mechanisms (Posner, 1980).

Of course, the approaches of behavioural competition and salience map competition are not mutually exclusive and it is possible to imagine various hierarchical schemes, whereby, for instance, a behaviour is selected first and then a point in space to which the behaviour will be directed. Alternatively, the target location might be selected and then the action to be directed at it. Finally, parallel, interacting sub-systems may simultaneously converge on both a target and suitable action (Cisek, 2007). We recently investigated the hypothesis that a salience map model can be used to generate action sequences for a biomimetic whiskered robot snout mounted on a mobile robot platform, and compared this with an earlier control model based on behaviour selection (Mitchinson, Pearson et al., 2012). Both control systems generated life-like sequences which alternate between exploration and orienting behaviour, but in the salience map version these higher-level behavioural `bouts' were an emergent consequence of actions determined by following a shifting focus of spatial attention (determined by a salience map) rather than resulting from the alternation of distinct behavioural primitives.

In the mammalian brain, sensorimotor loops involving the cortex, superior colliculus, basal ganglia, cerebellum and hippocampus may interact to implement a control system similar to this hypothesised salience map model. Figure 2 summarises the multi-level loop architecture used in our recent biomimetic robot Shrewbot, which is derived from our general understanding of the control architecture of the rat vibrissal system. We cannot represent the whole brain in our model from the outset, and there is no general agreement on the function of some of the neural centres. Since the robot must generate behaviour if we are to experiment with it, our breakdown of the control system into modules is by function, but the particular breakdown chosen is deeply inspired by our understanding of brain anatomy. This places us in a strong position to hypothesise relationships between structure and function in the neural system, and these hypotheses are a major outcome of our robot work. Here, the component ‘selection mechanism', modelled on the mammalian basal ganglia and superior colliculus, is responsible for selecting and driving the majority of movements of the robot's body (neck and wheels). Below this system, motor systems implement control commands, and low-level reflex loops support some rapid responses to current conditions (for instance, whisker protraction is inhibited by contact with the environment (Pearson, Mitchinson et al., 2010)). Above this system, we are beginning to add more cognitive components that modulate selection. The component labelled ‘abstraction' (Sullivan, Mitchinson et al., 2012) gleans additional information about what has been contacted by the vibrissal sensors in a manner analogous to processing centres such as the somatosensory cortex. Elsewhere, the component labelled ‘allocentric memory' retains a memory of the robot's past spatial experience and thus models some of the functionality of the mammalian hippocampal system (Fox, Evans et al., 2012a; O'Keefe and Nadel, 1978). ‘Signal conditioning’ indicates the importance of early processing of sensory signals to, for instance, distinguish components of the signal that may be due to the organism’s (or robot’s) own movement rather than to contact with the external world. Some neurobiological evidence, and our own modelling work, suggest an important role for the cerebellum in this regard (Anderson, Pearson et al., 2010; Anderson, Porrill et al., 2012).

FIGURE 2 HERE

Orienting the tactile fovea with the superior colliculus

To demonstrate the capacity of this architecture for generating integrated behaviour we have focused on the problem of orienting to interesting or novel stimuli detected by the robot vibrissae. To develop our model of orienting we first assume a 'tactile fovea' (Brecht, Preilowski et al., 1997), as the region of the snout with the highest density of microvibrissae, and focus on the key component of orienting behaviour in rodent-like mammals of bringing the fovea to a target. For instance, when faced with a task of discriminating between multiple objects, rat behaviour can be described as foveation (targeting the sensory fovea) to each discriminandum in sequence (Brecht, Preilowski et al., 1997). In our control architecture, then, the selection mechanism thus drives movements of the fovea with its output being the desired instantaneous velocity of the foveal position. In this model the movement of the remaining nodes of the animal/robot are unconstrained at the level of the selection mechanism and, instead, are determined at the level of the body (i.e. the motor system). Specifically, nodes such as the neck joints, and body are 'enslaved' to the fovea, and move so as to carry the fovea towards its target as smoothly and directly as possible. One could say that our robots are 'led by the nose'. This is, of course, a simplification of biological behaviour, though we have been surprised by how life-like (and practical) the resulting behaviour can be.

In primates, foveation is well studied with respect to the visual system and is known to be mediated by the Superior Colliculus (SC) (Gandhi and Katnani, 2011). In rats, stimulation of SC can evoke not only eye movements (McHaffie and Stein, 1982), but also orienting-like movements of the snout, circling, and even locomotion (Sahibzada, Dean et al., 1986). Recent neurobehavioural evidence also directly implicates SC as having a major role in rodent prey capture (Favaro, Gouvea et al., 2011). More generally, the SC appears to be a very plausible location to integrate tactile signals for spatial attention. It has the right inputs with signals arriving from the vibrissae via the from trigeminal sensory complex (Cohen, Hirata et al., 2008), and with further inputs converging from several relevant areas of cortex including S1 (Cohen, Hirata et al., 2008; Hemelt and Keller, 2007; Miyashita and Hamada, 1996). The organization of the SC is topographic in both its sensory and motor aspects, with a sensory topography appropriate to encoding a salience map centred on the foveal region of the snout (Benedetti, 1991; Drager and Hubel, 1976) and motor maps suitably configured to generate orienting head movements (Sahibzada, Dean et al., 1986). Inspired by these facts, we have developed a model of foveal velocity vector generation that mirrors the features of SC—that is, we employ a topographic saliency map driven by sensory input and modulated by information from mid- and upper-brain, with a simple motor output transform that drives foveation to the most salient region of local space (Mitchinson, Pearson et al., 2012). In the case of our robots, salience is excited by whisker contact and endogenous noise and suppressed by a top-down 'inhibition-of-return' signal from an allocentric memory component that lowers the salience of regions that have recently been foveated. The selection task, then, is to choose between foveation targets in nearby space.

This saliency map model of tactile attention has recently been extended to incorporate the regulation of vibrissal movement (Mitchinson and Prescott, 2013). To evaluate the model we tested it within a simulated two-dimensional environment containing configurable ‘obstacles’, under conditions analogous to those used in behavioural experiments (Grant, Mitchinson et al., 2009; Mitchinson, Martin et al., 2007; Mitchinson, Grant et al., 2011; Towal and Hartmann, 2006), and showed that it exhibits many of the modulations of whisker movement previously reported and summarised in Figure 3 (Mitchinson and Prescott, 2013). The model can also account for anticipatory aspects of active vibrissal control (see e.g. (Arkley, Grant et al., 2014; Grant, Mitchinson et al., 2009)) that cannot be the outcome of purely reflexive mechanisms. Here again the SC is implicated as a key sub-system in the rodent brain. Stimulation of SC can generate modulatory (non-periodic) whisker movements (Hemelt and Keller, 2008) suggesting a role in determining the protraction amplitude of the whiskers. Accordingly, SC outputs directly target the facial nucleus which is the motor nucleus that drives the whisker musculature (Miyashita and Mori, 1995). The receptive fields of SC neurons that are sensitive to deflection of single macrovibrissae are large and overlapping under anaesthesia (Drager and Hubel, 1976). Since the whiskers sweep back and forth during exploration this raises the possibility that, in the awake behaving animal, vibrissal receptive fields in SC are actually sharply tuned, but encode target locations in a head-centred spatial map (that might be contacted by moving whiskers) rather than contacts on distinct macrovibrissae *per se*.

FIGURE 3 HERE

The role of the cortex and basal ganglia in decision-making

Whilst the SC provides a mechanism that can control the orienting movements of the head and sensory systems, it is only one of many structures involved in identifying and *selecting* targets for foveation. Studies in primates implicate sensory processing in cortical areas (Gold and Shadlen, 2001) coupled with action selection in the basal ganglia (BG) (Chambers, Humphries et al., 2011; Hikosaka, Takikawa et al., 2000; Redgrave, Prescott et al., 1999) as critical substrates for the decision-making aspects of target selection. Our research with whiskered robots is helping us to analyse the contributions of these different neural systems to perceptual decision making in the mammalian brain.

The last two decades have seen major advances in our understanding of decision making as statistically optimal inference from noisy and ambiguous sensations using Bayesian probability theory (Knill and Pouget, 2004). A Bayesian approach to the task of classification involves recording the likelihoods of measurements from example sensory data. Given new test data, these likelihoods can be used with Bayes’ rule to calculate the posterior probability of the test data being drawn from each trained class, that is, the likelihood of the data belonging to any given category given the history of past data. Within the broader class of Bayesian classifiers the approach of *sequential analysis* (Wald, 1947) operates by applying Bayes’ rule repeatedly to accumulating evidence for competing perceptual hypotheses, derived from time series of sensory data, until a preset threshold is reached. This method can be likened to a process that has been observed in parietal cortex when monkeys are required to make perceptual judgements about visual motion direction and where individual neurons have been recorded that noisily ramp-up their firing rates until reaching a decision threshold (Bogacz, Brown et al., 2006; Gold and Shadlen, 2007; Platt and Glimcher, 1999).

Using our whiskered robots, we have explored the possibility that sequential Bayes can provide an effective general classifier for object properties detected using vibrissal sensors. Examining characteristics such as texture, radial distance to contact, speed of object movement and novelty, and using a range of robot platforms deploying different strategies for the control of whisker movement and position, we have shown that sequential Bayes is reliable, accurate (hit rates of >90% on several tasks), and out-performs a number of competing classification methods such as spectral templates, maximum likelihood, and multi-layer neural networks (Lepora, Pearson et al., 2010; Lepora, Fox et al., 2012; Lepora, Evans et al., 2010; Lepora, Sullivan et al., 2012; Sullivan, Mitchinson et al., 2012). That a classification method that matches with primate data can operate effectively with signals from artificial whiskers gives hope that a single theory of perceptual decision-making can be developed that will apply equally to primate vision and rodent vibrissal touch. The further implication of these studies is that a common memory format (log likelihoods) could be used to encode tactile memories for object properties. These models also have the potential to address questions about the nature of tactile memory in animals such as the Etruscan shrew since the ability to classify objects (for instance, as prey items that can be attached) requires efficient and compact memory traces, and the ability to make timely and appropriate decisions based on accumulating evidence.

Alongside evidence that cortex accumulates evidence for competing hypotheses, converging evidence from neurobiology and computational modelling, is showing that the BG anatomy maps onto a network implementation of an optimal statistical method for hypothesis testing that provides for timely and efficient selection of an appropriate response (Bogacz and Gurney, 2007; Gurney, Prescott et al., 2004; Lepora and Gurney, 2012). As noted above, in the sensory component of this process, evidence for the alternative interpretations of a stimulus accumulates in neuronal “evidence bins” (e.g., that a visual stimulus is moving right rather than left) and this accumulated evidence competes within the BG to elicit an action (e.g., press right lever or left lever). In the vibrissal system, the substantial projections from layer 5a of S1 cortex to the striatum—the major input structure of the BG—could provide a neural substrate for decision-making in relation to tactile object properties. We have used insights from experimental data, and from recent recordings in cortical areas during decision-making tasks (e.g. (Diamond, von Heimendahl et al., 2008)), to revise and extend existing primate-based computational models of the decision-making process (Lepora, Fox et al., 2012; Lepora and Gurney, 2012) and are in the process of exploring the implications of these revisions for decision-making in embodied robotic models.

The vibrissal somatosensory cortices—feature maps for detecting behavioural affordances

Within the parietal lobes of all mammals there are localised cortical areas that are more specialised towards particular sensory modalities—such as somatosensation, vision, audition and vestibular sensing—and other areas that are more multisensory in nature. In vibrissal specialists, such as rats, mice and shrews, the somatosensory areas devoted to the region of the snout are massively expanded compared to those of most mammals. In these species, multiple somatotopic maps of the body have been identified, the principal ones being labelled primary and secondary somatosensory cortex (S1 and S2). Whilst both of these areas have large domains devoted to the vibrissae, barrel-like aggregates of neurons (“barrels”) have been identified only in S1. The size of the area of cortex devoted to the large macrovibrisae appears to reflect the high innervation levels of the whisker follicles (see also discussion of cortical area size in (Catania and Catania, This Volume)). In the mouse, S1 cortex represents approximately 13% of the cortical surface area in total and 69% of the somatosensory cortex (Lee and Erzurumlu, 2005).

In rats and mice, S1 barrels exist for both the large and motile macrovibrissae in the posteriomedial barrel field (PMBF) and for the smaller non-actuated microvibrissae in the anterolateral barrel field (ALBF) (Woolsey, Welker et al., 1975), however, research has almost entirely focused on the larger barrels found in PMBF because of their ease of stimulation via the macrovibrissae. S1 and S2 are reciprocally connected with each other and also, via the corpus collosum, with their contralateral other halves (Keller, 1995). S1 is also connected with a number of other cortical regions including the motor and perirhinal cortices. Other major S1 projection areas include the thalamic areas from which it receives input (VPM and POm), the reticular nucleus of the thalamus, and, of particular interest here, sub-cortical targets in the basal ganglia, pontine nuclei (cerebellum), and superior colliculus (Keller, 1995).

For a vibrissal-specialist like the Etruscan shrew, successful prey capture is critically dependent on accurate and rapid detection of tactile stimulus velocity. This leads to the general question of how the brain might extract movement direction and speed from patterns of vibrissal deflection. Since it was first proposed, Jeffress' place theory (Jeffress, 1948) has been a dominant model for understanding how sensory motion is encoded in the brain (Joris, Smith et al., 1998). The idea is that coincidence detector neurons receive input from sensors after delays governed by the distance of each neuron from the corresponding signal sources. The inter-sensor time difference is then encoded by the location of neurons that are active because their connection delays exactly compensate the inter-sensor stimulation interval. The place theory therefore suggests an important role for neural geometry in computing the motion of sensory stimuli. Despite being a general theory of neural computation, most of the evidence for the place theory is provided by studies of the auditory system of auditory specialists such as the barn owl. The evidence from studies of mammalian auditory systems is inconclusive, for example, rabbit auditory cortex neurons are tuned to much longer inter-aural delays than can be accounted for by known axonal connection velocities (Fitzpatrick, Kuwada et al., 2000), and evidence from other sensory modalities is sparse.

In order to provide a further test of the generality of the place theory, we sought to apply it to a model of tactile stimulus processing in rodent barrel cortex (Wilson, Bednar et al., 2011). We asked whether model cortical neurons receiving synaptic inputs via delays governed by realistic connection geometry and plausible axonal propagation speeds would match the range of real responses to paired stimulation of adjacent whiskers. Validating this hypothesis we recreated, in simulation, the broad range of spiking patterns displayed by layer 2/3 barrel cortex neurons when adjacent whiskers are deflected through the range of inter-stimulus intervals, as measured electrophysiologically by Shimegi et al. (Shimegi, Akasaki et al., 2000). These biological experiments have shown that, when two adjacent whiskers are stimulated in a sequence with a few milliseconds interval, the responses of cortical neurons depends strongly on their positions (whether closer to the barrel of the first or second whisker), and are typically stronger than the sum of the responses to independent whisker deflections for a specific time interval (figure 4 left). Our modelling results (figure 4 right) showed—consistent with a place theory interpretation—that this broad range of recorded response profiles emerges naturally from the connection geometry as a function of the anatomical location of the neuron. In practical terms, the result that stimulus-evoked responses can be predicted by neuron location is important because it suggests that neural geometry needs special consideration as we construct theories of cortical processing.

Further consideration of the role of neural geometry may lead to predictions about sensory processing in species that maintain map-like representations compared with those that do not. For example, whilst individual neurons in rodent primary visual cortex respond selectively for the orientation of visual edges, they are arranged randomly in the cortex with respect to their orientation preference (Ohki, Chung et al., 2005). Presumably rodents would therefore be poor at using map-dependent mechanisms to extract stimulus velocity in the orientation domain (i.e., when extracting information about image rotation), compared to primates that have smooth topological maps for orientation preference and so might use place-coding mechanisms. In more general terms, evidence supporting the place theory from a tactile mammalian sensory system provides new insight into understanding how the brain represents moving sensory input. Finally, in the context of the vibrissal system, and the barrel cortex, this study provided a novel account of the fusion of information at the multi-whisker level that both explained existing data, by casting it within a general and powerful framework (place theory), and made testable predictions that could be investigated experimentally.

FIGURE 4 HERE

A key feature of the mammalian sensory cortices is the presence of self-organising topological maps. Cortical maps for features of each sensory modality can be highly plastic and shaped by a combination of physical and environmental constraints (Feldman and Brecht, 2005; Fox and Wong, 2005). We recently conducted a series of experiments driving map self-organisation with activity patterns representing tactile stimulation of an array of artificial whiskers, in order to predict the organisation of object representations in the somatosensory cortex (Wilson, 2011). Inputs to the model were patterns of activity in simulated layer 4, encoding the spatial location and direction of whisker deflections caused by tactile stimuli that varied in shape, direction and speed. Layer 4 activity patterns were then remapped as layer 2/3 activity patterns using distance-dependent signalling delays in the layer 4 to layer 2/3 projection, to additionally encode the relative timing of whisker deflections (Wilson, Bednar et al., 2011). This model represents a biologically grounded method by which to map the full spatial-temporal pattern of multi-whisker inputs to an essentially spatial representation of the stimulus across layer 2/3. Layer 2/3 activity patterns representing the range of multi-whisker stimulus patterns could thus be presented to a self-organising map model of postnatal development in layer 5, using an approach that we have shown previously to recreate known topological feature maps in layer 2/3 (Wilson, Law et al., 2010).

Our model of layer 5 map self-organisation, like our previous model for layer 2/3 map self-organisation, is based on the LISSOM (Laterally Interconnected Synergetically Self-organising Map) algorithm originally developed to capture the self-organising properties of primate visual cortex (Miikkulainen, Bednar et al., 2005). In our barrel cortex model, responses across the cortical sheet became organised into coextensive topological maps, wherein iso-feature contours for tactile stimulus shape, direction, and speed preferences intersected at right angles (see Figure 5). The model therefore makes the critical prediction that orthogonal tactile feature spaces are represented in the somatosensory cortex by orthogonal feature maps (and hence by orthogonal spatial codes). A series of controlled simulation experiments suggested further that i) speed and shape selective neurons align to regions of low selectivity in and between direction pinwheels, ii) direction, shape, and speed maps are acquired in developmental sequence, iii) stimulus direction is resolved by afferent projections to layer 5, whereas shape and speed are resolved by subsequent recurrent interactions in layer 5. These findings constitute specific, testable predictions about the development of functional maps and object representations in somatosensory cortex, i.e., that maps for the tactile motion direction implied by multi-whisker deflection sequences emerge earlier and more robustly than lower-order feature maps representing e.g., stimulus shape and speed. These modelling predictions were validated in experiments that connected self-organising networks to an artificial sensor array stimulated by a table-top positioning robot (Wilson, 2011).

FIGURE 5 HERE

The cerebellum viewed as an adaptive filter and forward model

In experiments with the Scratchbot robot platform we occasionally observed that the robot would orient (foveate) as if to a target when no object is in fact present. On further investigation it appeared that on these occasions the sensory signal generated by active whisking is wrongly interpreted as contact with a target. This empirical observation in our biomimetic robot contrasts with the lack of reports of such 'phantom' orienting in normal rats. On the other hand, it has been shown that in rats sensory signals *are* generated by whisking movements. Specifically, a study by Leiser and Moxon reported that trigeminal ganglion cells of the rat fired during active whisking in air with no object contacts but were silent when the whiskers were at rest (Leiser and Moxon, 2007). The implication of this result is that whisker sensory signals may include self-generated artefacts during whisking. The fact that the robot does show phantom orienting and rats appear not to suggests that, unlike the robot, rats can discriminate between the component of a sensory signal that originates from an external source and the component that is self-generated by its own whisking movements. We can ask the question: how is this discrimination achieved? This leads us to suggest that the rat may actively cancel self-generated sensory signals —what we might call 'self-induced' or 'self-generated' noise.

Interestingly, noise cancellation in biological systems has a well-investigated precedent—interference cancellation in passive electro-sensing in electric fish (for review see Bell et al. (Bell, Han et al., 2008)). Of particular interest here is the evidence that suggests a cerebellar-like structure performs the function of noise cancellation in these animals (Bell, Han et al., 2008; von der Emde and Bell, 1996). Additionally, this cerebellar-like structure is thought to act analogously to an adaptive filter (Montgomery and Bodznick, 1994; Sawtell and Williams, 2008), linking biological noise cancellation to both the signal processing literature (Widrow and Stearns, 1985) as well as the adaptive filter model of the cerebellum (Fujita, 1982). In humans, a similar capacity to predict or cancel self-induced sensory signals is indicated by our inability to entertain ourselves by self-tickling (as opposed to be tickled by someone else). In this case, functional MRI data (Blakemore, Wolpert et al., 1998) also indicates a role for the cerebellum in predicting sensory signals due to self-movement thereby making them seem less amusing!

The above considerations led us to the hypothesis that rats may use their cerebellum to generate a signal which cancels the effects of self-generated whisking noise on incoming sensory signals from the whiskers (see (Anderson, Pearson et al., 2010; Dean, Porrill et al., 2010) for a review of the wider literature on the role of the cerebellum in sensory noise cancellation). Our first step in investigating this hypothesis was to determine whether our proposed mechanism would work *in principle*, by using it to achieve noise cancellation in a whisking robot (Anderson, Pearson et al., 2010). The step had two goals, first to solve the practical problem of noise cancellation in the robot, and secondly to provide a theoretical basis for studying noise cancellation in whisking animals. Our approach was to use inspiration from the signal processing literature to form a prototype hypothesis of a whisking noise cancellation scheme.

The subject of noise cancellation has been studied in the signal processing literature since the 1960s (for early references see Widrow et al. (Widrow and Stearns, 1985)). Much of the formative work was conducted by Bernard Widrow and was linked as an application problem to the least-mean-squares (LMS) adaptive filtering algorithm. The generic noise cancellation scheme is illustrated in figure 6, for detailed explanation see (Anderson, Pearson et al., 2010). The weights of the adaptive filter in the noise cancellation scheme (figure 6a) are adjusted by removing the correlations in the clean signal from the reference noise, implemented via the LMS rule. In the context of whisking, the self-generated noise is thought to be caused by the movement of the whisker, either by inertia of the whisker base in the follicle or the whisker musculature pressing and activating the mechanoreceptors (Leiser and Moxon, 2007). Ultimately, this activation is caused by the motor command to the whisker plant. Hence, we regard the motor command (either high- or low-level) as the reference noise, which is correlated with the noise signal but uncorrelated with signals related to object contacts. In our proposed whisking noise cancellation scheme (illustrated in figure 6b) the cerebellum learns to predict self-generated noise from the motor commands that cause the whisker movements. Hence, the cerebellum learns an internal forward model of the whisker dynamics that transform motor commands into sensory signals.

FIGURE 6 HERE

Adaptive filter model of the cerebellum

In the above whisking noise cancellation scheme we use the adaptive-filter to computationally model the cerebellum, as originally proposed by Fujita (Fujita, 1982). The mapping of this scheme onto the cerebellar microcircuit is illustrated in figure 7. We have previously investigated the computational properties of this model for adaptive motor control (Dean, Porrill et al., 2002; Dean and Porrill, 2008; Porrill, Dean et al., 2004; Porrill and Dean, 2007), and others have proposed that it could be used in principle to learn forward models (see (Dean, Porrill et al., 2010) for review). However, our vibrissal noise study was, to our knowledge, the first instance of the adaptive-filter model of the cerebellum being applied to learning a specific forward model (i.e. of whisker dynamics) for the purposes of noise cancellation.

FIGURE 7 HERE

To develop and validate our proposed whisking noise cancellation scheme we recorded experimental data from Scratchbot during ‘free’ whisking (i.e. with no object contacts). Note that free-whisking is an ideal scenario to test the noise cancellation scheme because during free whisking the whisker sensory signal should be zero. Hence, whilst the whisker dynamics are unknown and therefore the optimal cerebellar filter is also unknown, the output of the cancellation scheme is known: it should be zero. Therefore it is straight-forward to evaluate the performance of the noise cancellation scheme during free-whisking. Figure 8 shows example results of the application of the cerebellar noise cancellation algorithm to free-whisking data.

FIGURE 8 HERE

In a further extension of our sensory noise cancellation model (Anderson, Porrill et al., 2012) we have shown that the addition of sensory information from the whiskers allows the adaptive filter to learn a more complex internal model that performs more robustly than a forward model based on efference copy signals alone, particularly when the whisking-induced interference has a periodic structure. More generally, our analysis of the whisking noise cancellation scheme reveals that the functional role of the cerebellum may be to learn a forward model of the whisker/follicle dynamics. This links to separate speculation over the functional role of the cerebellum in motor control and sensory processing, where it has been suggested that the cerebellum can learn a variety of forward and inverse models in control and state estimation tasks, see for instance (Dean and Porrill, 2014; Wolpert, Miall et al., 1998). Our development of the whisking noise cancellation scheme from a theoretical basis has led to a number of experimental predictions relating to the functionality of different components of the cerebellar micro-circuit: (i) that the mossy fibres transmit a copy of motor command, (ii) that the Purkinje cell output is an estimate of the self-induced noise signal, (iii) that the climbing fibre teaching signal is an estimate of the 'clean' whisker sensory signal, and (iv) that the superior colliculus is the target of the cerebellar output and acts to compare predicted and actual sensory signals (Anderson, Porrill et al., 2012).

Cerebellar/collicular algorithms for orienting and predictive pursuit

Cerebellar circuits are likely to be important for fast predictive control of ballistic movements needed for tasks such as prey tracking and capture since cerebellar damage is known to impair predictive aspects of motor behaviour (Bastian, 2006). An important role might lie in the calibration of sensory maps used to generate fast orienting movements. We have hypothesised (Porrill, Anderson et al., 2010) that the known extensive cerebellar-collicular connectivity (see (Anderson, Porrill et al., 2012; Teune, van der Burg et al., 2000)), together with the adaptive filter cerebellar architecture described above, could play a role in calibrating predictive topographic maps in the colliculus. We are currently investigating how this model can be employed in a predictive architecture in which features appear in the salience map at their predicted rather than their current positions.

The cerebellum may also calibrate sensory information that provides input to the predictive system. For example, as described above, we have developed a cortical algorithm for estimating contact timing for a target moving through the robotic whisker array illustrated in figure 9 left (Wilson, 2011). Figure 9 (centre) shows target velocities recovered from these timings, whilst incorporating a cerebellar learning element produced a more accurate calibration as shown in Figure 9 (right).

FIGURE 9 HERE

Tactile self-localisation and mapping in the hippocampus

The lifestyle of any small mammal, even one as tiny as the Etruscan shrew, requires the capacity to know where you are at all times with respect to key locations such as the nest, important feeding sites, and significant danger zones. Indeed, as we have seen above, simply to explore space efficiently using vibrissal touch requires some long-term memory of locations you have visited in the recent past, and the capacity to update an estimate of your own position as you move around. In the mammalian brain the hippocampal system is known to be important in building and maintaining representations of the environment (the 'place cell' system (O'Keefe and Nadel, 1978)) and in maintaining estimates of changes in position determined through path integration (the 'grid cell' system (Moser, Kropff et al., 2008)). Recent data also demonstrates that the hippocampus also encodes tactile information that describe the environmental context obtained through vibrissal touch (Itskov, Vinnik et al., 2011).

The principal input structures of the hippocampus are the superficial layers of Entorhinal Cortex (EC). EC projects to Dentate Gyrus (DG) which is believed to increase the sparcity of the encoding generated by the EC. Both EC and DG project to CA3, which also receives strong recurrent connections that are disabled (Hasselmo, Schnell et al., 1995) by septal acetylcholine (ACh). CA3 and EC project to CA1, which in turn projects to the deep layers of Entorhinal cortex, there is also a back-projection from CA3 to DG (Scharfman, 2007). Although the classical view of hippocampus is as a single loop, there is also a second loop—EC and CA1 project to Subiculum (Sub), which projects to the midbrain Septum (Sep) via fornix. Septal ACh and GABA fibres then project back to all parts of hippocampus. Figure 10 summarises many of these connections.

There have been two broad schools of hippocampal modelling one based acquiring spatial sequences, and the other on the notion auto-associative memory including pattern reconstruction based on partial or noisy input (see (Fox and Prescott, 2010a) for review). However, the objectives of both auto-associative and spatial sequence memories can be combined by a general Bayesian filter with noisy observations, which seeks infers the (hidden) state of the world. Such a filter that maintains just a single estimate of the current state-of-the-world (e.g. of location in spatial map) is known as a 'unitary particle filter'. We have developed a model of spatial learning in the rodent hippocampus (Fox and Prescott, 2010a), viewed as a unitary particle filter, by mapping key structures in the hippocampal system onto the components of a Temporal Restricted Boltzmann Machine—a probabilistic algorithm for learning sequence data developed by researchers in machine learning (see, e.g. (Taylor, Hinton et al., 2007)). The algorithm approximates Bayesian filtering to infer both auto-associative de-noised percepts and temporal sequences, that is, it can clean-up and fill-out incoming sensory patterns and can use these to recall or forecast sequences of places visited during navigation. The mapping to the hippocampal system (see Figure 10) proposes a novel role for the subiculum, and for ACh from the septum, in detecting when the animal has become lost (by detecting a mismatch between predicted and actual sensory signals). A follow-up paper (Fox and Prescott, 2010b) extended this model to include online learning of connections to and from the simulated hippocampal CA3 region.

FIGURE 10 HERE

Building on our computational models of hippocampus, we have developed tactile Self-Localisation and Mapping (tSLAM) for whiskered mobile robot platforms. tSLAM provides a robot with a means of mapping and navigating a novel environment by touch information alone, something which has never previously been developed in robotics. A critical step, was the development of a hierarchical Bayesian 'blackboard' architecture (Fox, Evans et al., 2012b) to investigate how to fuse information from multiple local tactile feature reports to recognise objects in the world. This work also involved developing techniques for online head-centric spatial localisation of whisker contacts, and their subsequent world-centric transformation. To achieve tSLAM we have developed a particle-filter based mapping and localisation algorithm, taking odometry (path integration) and tactile information from the robot in real-time. This information is then integrated into an occupancy grid map, and a current position estimate.

The tSLAM system has now been piloted on two whiskered robot systems—Crunchbot, a modified Roomba vacuum-cleaner robot with a small array of static whiskers (Fox, Evans et al., 2012b), and Shrewbot a robot with multiple actively controlled whiskers and a 3 degrees-of-freedom neck (see Figure 1 and (Pearson, Fox et al., 2013)). In the case of the Shrewbot platform, odometry derived from the robot base and neck were passed at regular intervals (in phase with the whisking) to a population of particles each maintaining an estimate of head pose and location within a 2-dimensional occupancy grid. The importance of each particle was calculated by fusing the likelihood that each whisker in the array is at the estimated location in the map based on tactile information sampled throughout the previous whisk. The screen shot shown in Figure 11 shows a one hour experimental run of Shrewbot in a 3m diameter arena. The pink regions representing areas of the map that have a high probability of occupancy, the dashed white line representing ground truths taken from the over head video camera. The dashed white representation of Shrewbot is its ground truth location, whilst the solid representation is the current best estimate of pose and location taken from the particle with highest importance (cloud visible as red dots near the head).

FIGURE 11 HERE

We have also used the Shrewbot platform to model the active touch based hunting behaviour of the Etruscan Shrew (Mitchinson, Pearson et al., 2014). A study of vibrissal-guided predation of insects by the shrew (Munz, Brecht et al., 2010) identified three distinct phases of hunting behaviour: search, contact and attack. The search phase was reproduced on Shrewbot using the tactile attention based model of action selection described above, whereby the locus of attention drives the orienting behaviour of the robot between subsequent whisks. Upon making contact, an internal geometric model of two classes of object was compared to the sparse 3-dimensional tactile information derived from the whisker array. The two classes of object were vertical "walls" and the dome shaped covering of a mobile robot referred to as "preybot" (see Figure 12). Shrewbot's reflexive whisker control strategies (Pearson, Mitchinson et al., 2011) caused an increase in the whisking set angle similar to that reported in (Munz, Brecht et al., 2010) as well as an increase in the number and frequency of whiskers making controlled contacts with the object. This information was collated into a "prey belief" metric that influenced the decision to either attack the object (preybot) or to ignore it (walls). In parallel to the attack decision process, the centre of mass of the preybot was also estimated. To accommodate the relatively sparse information from whisker contacts, some of the known characteristics of the preybot were used to better infer its location and orientation and hence its affordances as a potential "prey" object. The velocity of the preybot was derived from this information and thence a prediction of where a particular point on that robot (in this case the tail) should be in the near future. This location in space was then set as the target for an attack.

FIGURE 12 HERE

Conclusion

In this chapter we have briefly summarised an extensive programme of work aimed at describing and simulating, in biomimetic (robotic) models, the control architecture for sensorimotor co-ordination in the vibrissal touch system of small mammals. We have shown how the evolution of our robot models has progressively captured more-and-more of the important features of the biological target system including morphology, sensory transduction, motor control, and internal processing. Focusing initially on the problem of orienting to vibrissal contacts we have shown that models of the superior colliculus and basal ganglia can be combined to generate sequences of exploratory and orienting movements that allow the robot to explore an environment, and orient to unexpected contacts, in a life-like way. Our robot experiments also revealed a need to pre-process sensory signals in order to distinguish real physical contacts from 'ghost' contact signals induced by the whisking movements of the artificial vibrissae. This led to a novel hypothesis about the role of the cerebellum in vibrissal processing and the demonstration that adaptive filter algorithms modelled on cerebellar microcircuitry can be effective in predicting/cancelling self-induced sensory noise. The task of developing integrated sequences of movements in whiskered robots also revealed the need for spatial memory systems that could effectively encode and remember the location of object contacts and allow the robot/animal to maintain a good estimate of its position in space. To make these systems more effective in identifying, and responding appropriately, to tactile environmental affordances, we are also developing models of cortical systems (particular of primary somatosensory cortex), and have shown that model basal ganglia circuits can make timely decisions between alternatives based on cortical encodings of vibrissal touch signals. Whilst we have yet to realise the full architecture shown in Figure 2 in a single robot, our latest robotic models show a capacity for integrated behaviour that has surprised and impressed exhibition and conference audiences into thinking that they are observing something like a 'robot animal'. From the perspective of understanding brain architecture, we also consider that we have made important steps towards understanding and demystifying the neural-basis for sensorimotor co-ordination in mammalian brains including our own.

Acknowledgement

The authors would like to acknowledge the support of the European Union FP7 BIOTACT Project 'Biomimetic Technology for Vibrissal Active Touch' (ICT-215910) and the Weizmann Institute of Science project "Development of motor-sensory strategies for vibrissal active touch". We would also like to thank various collaborators whose experimental and theoretical research has inspired our biomimetic models including Mathew Diamond, Michael Brecht, Ehud Ahissar, David Golomb, Mitra Hartmann, Chris Yeo, Peter Redgrave and Kevin Gurney.

Figure Legends

Figure 1. Biomimetic whiskered robots. From left to right: Whiskerbot, Scratchbot, Shrewbot, Generation 2 Biotact Sensor. Each robot has a snout configured with an array of moveable artificial whiskers. Different mechanisms have been explored for whisker actuation and for sensory transduction in the different devices. We have also gradually evolved the overall design of the whisker morphology and of the neuromimetic control architecture. The most recent model systems (Shrewbot, G2 Biotact Sensor) feature arrays of individually actuated whiskers with intrinsic transduction systems based on Hall effect sensors that can measure whisker deflection in three dimensions. For further details of the 'evolution' of our whiskered robots see (Mitchinson, Pearson et al., 2011; Mitchinson, Sullivan et al., 2013; Prescott, Pearson et al., 2009). Photos by Martin Pearson (Whiskerbot, Scratcbot, Biotact Sensor) and Tony Prescott (Shrewbot).

Figure 2. Model of brain architecture for control of a whiskered mobile robot. The abstract components of the model can be mapped to key sub-systems in the mammalian brain (see text). Figure reproduced from Figure 2 of (Mitchinson, Pearson et al., 2012) with permission from Springer.

Figure 3. Model of the regulation of whisking behaviour by spatial attention. *Top-left.* A mix of exogenous (tactile) and endogenous (other) influences affect the locus/region of spatial attention. This locus drives head movements and is responsible for the modulation of whisker movements from whisk to whisk. *Bottom-left.* Three frames from an overhead video of a rat executing an orient to an unexpected contact. *Centre.* The implementation used to test the model—additional components are a simple oscillator to generate periodic whisking (OSC), an implementation of inhibition-of-return (IOR) to generate sufficiently rich orienting behaviour for testing, and a physical model of whisker/environment interactions. *Right.* Comparison of results from current model (A) and recordings of rat behaviour (B) under three analyses, from top: Contact-induced asymmetry (see, e.g. (Mitchinson, Martin et al., 2007))—if an animal approaches a surface at an oblique angle then protraction of whiskers ipsilateral to the surface is reduced (red color / dark shading inside boundary), whilst protraction of whiskers on the contralateral side is increased (blue color / light shading); Head-turning asymmetry (see, e.g. (Towal and Hartmann, 2006))—as an animal turns the whiskers typically move asymmetrically as if to anticipate obstacles in the direction of the turn; Spread reduction during contact (see. e.g. (Grant, Mitchinson et al., 2009))—whilst exploring a surface the whiskers are brought closer together with the effect of increasing the number of whisker-surface contacts. Adapted from figures 2, 3, 7 and 8 of (Mitchinson and Prescott, 2013) which should be consulted for further explanation of the model and results.

Figure 4. Spatiotemporal interactions of cortical responses to paired whisker stimulations. The left panel shows experimental data replotted from Figure 8e of (Shimegi, Akasaki et al., 2000) with permission from Society for Neuroscience. The right panel shows the behaviour of our model (Wilson, Bednar et al., 2011). In both panels, the response facilitation index – computed as the ratio between the response to a paired stimulation of adjacent whiskers A and B and the linear sum of the responses to either A or B separately – is shown as a function of the time interval that separated the two whisker deflections. Figure reproduced from (Wilson, Bednar et al., 2011).

Figure 5. The emergence of orthogonal coding for tactile stimuli of different shapes (concavity/convexity), directions, and speeds, in a self-organising model of map development in the barrel cortex. See (Wilson, 2011) for further details.

Figure 6. Noise cancellation schemes. (a) A generic adaptive noise cancellation scheme, see for instance Widrow and Stearns (Widrow and Stearns, 1985). (b) A proposed biological whisking noise cancellation scheme. See (Anderson, Pearson et al., 2010) for further details.

Figure 7. Schematic diagram of the organization of the cerebellar microcircuit and its interpretation as an adaptive linear filter. (a) Simplified architecture of cerebellar cortex. (b) Adaptive filter model of the cerebellum. Adapted from Figure 1a,b of (Dean, Porrill et al., 2010) with permission from Nature Publishing Group.

Figure 8. Results from applying the noise cancellation algorithm to the free-whisking sensory signal. (a) Low frequency linear noise cancellation in the range 0-5 Hz. (b) High frequency nonlinear noise cancellation (up to 100 Hz). See (Anderson, Pearson et al., 2010) for further details.

Figure 9. Cortical algorithm for contact timing. a) A planar target is moved through a robotic whisker array on an xy-plotter through eight speeds and eight directions to generate a set of multi-whisker deflection patterns b) plot of target velocities recovered from the relative timing of whisker responses, as computed using a cortical velocity-encoding algorithm based on (Wilson, 2011); plotted against the x and y components of the stimulus motion velocity, the distortion of the cortical estimates compared to a regular grid of true velocities is clear c) estimation of the motion velocities are improved by cerebellar correction after randomised representation of the data set.

Figure 10. Model of spatial memory in the hippocampal system viewed as a particle filter. Structures are labelled with UML notation indicating many-to-many fully connected links (* → *), one-to-one links (1 → 1) and many-to-one links (* → 1). Thick lines are ACh projections, thin lines are glutamate. The model implements a spatial memory system for location based on multisensory

signals from tactile, visual, and path integration signals. See text for abbreviations showing the proposed mapping to regions of the rodent hippocampal system. Reproduced from Figure 1 of (Fox and Prescott, 2010a) with permission from IEEE.

Figure 11. Tactile Self-localisation and Mapping (tSLAM) in the Shrewbot whiskered robot platform. Screen shot taken from a combined video of overhead camera view (right) with an appropriately scaled and rotated 2d occupancy grid representation of the arena in a typical particle after approximately 1 hour of whisker based tactile exploration (left). Figure reproduced from multimedia supplement to (Pearson, Fox et al., 2013) with permission from IEEE.

Figure 12. Tactile identification and tracking of a target in a whiskered robot. Snapshots taken with an overhead camera of Shrewbot approaching the "preybot" during hunting behaviour. The images indicate: *search* (frame 1), *contact* (frame 2 and 3), *attack* (frame 4) and a return to *search* (frame 5). Figure reproduced from Figure 3 of (Mitchinson, Pearson et al., 2014) with permission from Elsevier.

Bibliography

Ahissar E, Kleinfeld D. Closed-loop neuronal computations: focus on vibrissa somatosensation in rat. Cereb Cortex. 2003;13(1):53-62.

Anderson SR, Pearson MJ, Pipe A, Prescott T, Dean P, Porrill J. Adaptive cancelation of self-generated sensory signals in a whisking robot. Robotics, IEEE Transactions on. 2010;26(6):1065-76.

Anderson SR, Porrill J, Prescott TJ, Dean P. An internal model architecture for novelty detection: Implications for cerebellar and collicular roles in sensory processing. PLoS ONE. 2012;7(9):e44560.

Anjum F, Turni H, Mulder PG, van der Burg J, Brecht M. Tactile guidance of prey capture in Etruscan shrews. Proc Natl Acad Sci U S A. 2006;103(44):16544-9.

Arkley K, Grant RA, Mitchinson B, Prescott TJ. Strategy Change in Vibrissal Active Sensing during Rat Locomotion. Current biology : CB. 2014;24(13):1507-12.

Bastian AJ. Learning to predict the future: the cerebellum adapts feedforward movement control. Curr Opin Neurobiol. 2006;16(6):645-9.

Bell CC, Han V, Sawtell NB. Cerebellum-like structures and their implications for cerebellar function. Annu Rev Neurosci. 2008;31:1-24.

Benedetti F. The postnatal emergence of a functional somatosensory representation in the superior colliculus of the mouse. Developmental Brain Research. 1991;60(1):51-7.

Blakemore S-J, Wolpert DM, Frith CD. Central cancellation of self-produced tickle sensation. Nature Neuroscience. 1998;1(7):635-40.

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol Rev. 2006;113(4):700-65.

Bogacz R, Gurney K. The basal ganglia and cortex implement optimal decision making between alternative actions. Neural Comput. 2007;19(2):442-77.

Braitenberg V. Vehicles: Experiments in Synthetic Psychology. Cambridge, MA: MIT Press; 1986.

Brecht M, Preilowski B, Merzenich MM. Functional architecture of the mystacial vibrissae. Behavioural Brain Research. 1997;84(1-2):81-97.

Brecht M, Naumann R, Anjum F, Wolfe J, Munz M, Mende C, et al. The neurobiology of Etruscan shrew active touch. Philos Trans R Soc Lond B Biol Sci. 2011;366(1581):3026-36.

Brooks RA. New approaches to robotics. Science. 1991;253:1227-32.

Catania KC, Catania EH. Comparative Studies of Somatosensory Systems and Active Sensing. In: Krieger P, Groh A, editors. Sensorimotor Integration in the Whisker System. New York: Springer; This Volume.

Chambers J, Humphries M, Gurney KN, Prescott TJ. Mechanisms of choice in the primate brain: A quick look at positive feedback. In: Seth A, Bryson JJ, Prescott TJ, editors. Modelling Natural Action Selection. Cambridge: CUP; 2011. p. 390-418.

Cisek P. Cortical mechanisms of action selection: The affordance competition hypothesis. Phil Trans Roy Soc B. 2007;362(1485):1585-99.

Cohen JD, Hirata A, Castro-Alamancos MA. Vibrissa sensation in superior colliculus: wide-field sensitivity and state-dependent cortical feedback. J Neurosci. 2008;28(44):11205-20.

Dean P, Porrill J, Stone JV. Decorrelation control by the cerebellum achieves oculomotor plant compensation in simulated vestibulo-ocular reflex. Proc Biol Sci. 2002;269(1503):1895-904.

Dean P, Porrill J. Adaptive-filter models of the cerebellum: computational analysis. Cerebellum. 2008;7(4):567-71.

Dean P, Porrill J, Ekerot CF, Jorntell H. The cerebellar microcircuit as an adaptive filter: experimental and computational evidence. Nat Rev Neurosci. 2010;11(1):30-43.

Dean P, Porrill J. Decorrelation learning in the cerebellum: computational analysis and experimental questions. Prog Brain Res. 2014;210:157-92.

Diamond ME, von Heimendahl M, Arabzadeh E. Whisker-mediated texture discrimination. PLoS Biol. 2008;6(8):e220.

Drager UC, Hubel DH. Topography of visual and somatosensory projections to mouse superior colliculus. J Neurophysiol. 1976;39(1):91-101.

Favaro PD, Gouvea TS, de Oliveira SR, Vautrelle N, Redgrave P, Comoli E. The influence of vibrissal somatosensory processing in rat superior colliculus on prey capture. Neuroscience. 2011;176:318-27.

Feldman DE, Brecht M. Map plasticity in somatosensory cortex. Science. 2005;310(5749):810-5.

Fitzpatrick DC, Kuwada S, Batra R. Neural sensitivity to interaural time differences: beyond the Jeffress model. J Neurosci. 2000;20(4):1605-15.

Fox CW, Prescott TJ. The hippocampus as a unitary coherent particle filter. International Joint Conference on Neural Networks (IJCNN); Barcelona, Spain2010a.

Fox CW, Prescott TJ. Learning in the unitary coherent hippocampus. 20th International Conference on Artificial Neural Networks (ICANN); , ThessalonIka, Greece2010b.

Fox CW, Evans M, Pearson M, Prescott TJ. Tactile SLAM with a biomimetic whiskered robot. IEEE International Conference on Robotics and Automation (ICRA)2012a.

Fox CW, Evans MH, Pearson MJ, Prescott TJ. Towards hierarchical blackboard mapping on a whiskered robot. Robot Auton Syst. 2012b;60(11):1356-66.

Fox K, Wong RO. A comparison of experience-dependent plasticity in the visual and somatosensory systems. Neuron. 2005;48(3):465-77.

Fujita M. Adaptive filter model of the cerebellum. Biological Cybernetics. 1982;45(3):195-206.

Gandhi NJ, Katnani HA. Motor functions of the superior colliculus. Annu Rev Neurosci. 2011; 34:205-31.

Gibson JJ. The Ecological Approach to Visual Perception. Boston: Houghton Mifflin; 1979.

Gold JI, Shadlen MN. Neural computations that underlie decisions about sensory stimuli. Trends Cogn Sci. 2001;5(1):10-6.

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

Grant RA, Mitchinson B, Fox C, Prescott TJ. Active touch sensing in the rat: Anticipatory and regulatory control of whisker movements during surface exploration. Journal of Neurophysiology. 2009;101(2):862-74.

Gurney K, Prescott TJ, Wickens JR, Redgrave P. Computational models of the basal ganglia: from robots to membranes. Trends Neurosci. 2004;27(8):453-9.

Hallam JCT, Malcolm CA. Behaviour—perception, action, and intelligence—the view from situated robotics. Philosophical Transactions of the Royal Society of London Series A. 1994;349(1689):29-42.

Hasselmo ME, Schnell E, Barkai E. Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. J Neurosci. 1995;15(7 Pt 2):5249-62.

Hemelt ME, Keller A. Superior sensation: superior colliculus participation in rat vibrissa system. BMC Neurosci. 2007;8:12.

Hemelt ME, Keller A. Superior colliculus control of vibrissa movements. J Neurophysiol. 2008.

Hikosaka O, Takikawa Y, Kawagoe R. Role of the basal ganglia in the control of purposive saccadic eye movements. 2000;80(3):953-78.

Itskov PM, Vinnik E, Diamond ME. Hippocampal representation of touch-guided behavior in rats: persistent and independent traces of stimulus and reward location. PLoS One. 2011;6(1):e16462.

Jeffress LA. A place theory of sound localization. J Comp Physiol Psychol. 1948;41(1):35-9.

Joris PX, Smith PH, Yin TC. Coincidence detection in the auditory system: 50 years after Jeffress. Neuron. 1998;21(6):1235-8.

Keller A. Synaptic organization of the barrel cortex. In: Jones EG, Diamond IT, editors. Cerebral Cortex Volume II: The Barrel Cortex of Rodents. New York: Plenum Press; 1995.

Knill DC, Pouget A. The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci. 2004;27(12):712-9.

Lee L-J, Erzurumlu RS. Altered parcellation of neocortical somatosensory maps in N-methyl-D-aspartate receptor-deficient mice. J Comp Neurol. 2005;485:57-63.

Leiser SC, Moxon KA. Responses of trigeminal ganglion neurons during natural whisking behaviors in the awake rat. Neuron. 2007;53(1):117-33.

Lepora N, Pearson MJ, Mitchinson B, Evans M, Fox CW, Pipe T, et al., editors. Naive Bayes novelty detection for a moving robot with whiskers. IEEE International Conference on Robotics and Biomimetics (ROBIO); 2010; Tainjin, China, December 14-18.

Lepora N, Fox CW, Evans MH, Diamond ME, Gurney K, Prescott TJ. Optimal decision-making in mammals: insights from a robot study of rodent texture discrimination. J R Soc Interface. 2012;9(72):1517-28.

Lepora N, Gurney K. The basal ganglia optimize decision making over general perceptual hypotheses. Neural Computation. 2012;24(11):2924–45.

Lepora NF, Evans M, Fox CW, Diamond ME, Gurney K, Prescott TJ, editors. Naive Bayes texture classification applied to whisker data from a moving robot. Neural Networks (IJCNN), The 2010 International Joint Conference on; 2010 18-23 July 2010.

Lepora NF, Sullivan JC, Mitchinson B, Pearson M, Gurney K, Prescott TJ, editors. Brain-inspired Bayesian perception for biomimetic robot touch. Robotics and Automation (ICRA), 2012 IEEE International Conference on; 2012 14-18 May 2012.

McHaffie JG, Stein BE. Eye movements evoked by electrical stimulation in the superior colliculus of rats and hamsters. Brain Research. 1982;247:243--53.

Miikkulainen R, Bednar JA, Choe Y, oseph Sirosh J. Computational Maps in the Visual Cortex. Berlin: Springer; 2005.

Mitchinson B, Martin CJ, Grant RA, Prescott TJ. Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact. Proc Biol Sci. 2007;274(1613):1035-41.

Mitchinson B, Grant RA, Arkley K, Rankov V, Perkon I, Prescott TJ. Active vibrissal sensing in rodents and marsupials. Philos Trans R Soc Lond B Biol Sci. 2011;366(1581):3037-48.

Mitchinson B, Pearson M, Pipe T, Prescott TJ. Biomimetic robots as scientific models: A view from the whisker tip. In: Krichmar J, editor. Neuromorphic and Brain-based Robots. Boston, MA: MIT Press; 2011.

Mitchinson B, Pearson MJ, Pipe AG, Prescott TJ. The emergence of action sequences from spatial attention: insight from mammal-like robots. Living Machines: Biomimetic and Biohybrid Systems; Barcelona, Spain2012.

Mitchinson B, Prescott TJ. Whisker movements reveal spatial attention: a unified computational model of active sensing control in the rat. PLoS Comput Biol. 2013;9(9):e1003236.

Mitchinson B, Sullivan JC, Pearson MJ, Pipe AG, Prescott TJ, editors. Perception of simple stimuli using sparse data from a tactile whisker array. Biomimetic abd Biohybrid Systems; 2013; London: Springer.

Mitchinson B, Pearson MJ, Pipe AG, Prescott TJ. Biomimetic tactile target acquisition, tracking and capture. Robot Auton Syst. 2014;62(3):366-75.

Miyashita E, Mori S. The superior colliculus relays signals descending from the vibrissal motor cortex to the facial nerve nucleus in the rat. Neurosci Lett. 1995;195(1):69-71.

Miyashita E, Hamada Y. The 'functional connection' of neurones in relation to behavioural states in rats. Neuroreport. 1996;7(14):2407-11.

Montgomery JC, Bodznick D. An adaptive filter that cancels self-induced noise in the electrosensory and lateral line mechanosensory systems of fish. Neurosci Lett. 1994;174(2):145-8.

Moser EI, Kropff E, Moser MB. Place cells, grid cells, and the brain's spatial representation system. Annu Rev Neurosci. 2008;31:69-89.

Munz M, Brecht M, Wolfe J. Active touch during shrew prey capture. Frontiers in Behavioral Neuroscience. 2010;5:12.

O'Keefe JA, Nadel L. The Hippocampus as a Cognitive Map: Oxford University Press; 1978.

Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC. Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature. 2005;433(7026):597-603.

Pearson M, Fox C, Sullivan JCP, T. J., Pipe AG, Mitchinson B, editors. Simultaneous localisation and mapping on a multi-degree of freedom biomimetic whiskered robot. IEEE International Conference on Robotics and Automation (ICRA); 2013; Karlsruhe.

Pearson MJ, Mitchinson B, Welsby J, Pipe T, Prescott TJ. SCRATCHbot: Active tactile sensing in a whiskered mobile robot. In: Meyer J-A, Guillot A, editors. The 11th International Conference on Simulation of Adaptive Behavior; Paris, France2010.

Pearson MJ, Mitchinson B, Sullivan JC, Pipe AG, Prescott TJ. Biomimetic vibrissal sensing for robots. Philos Trans R Soc Lond B Biol Sci. 2011;366(1581):3085-96.

Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. Nature. 1999;400(6741):233-8.

Porrill J, Dean P, Stone JV. Recurrent cerebellar architecture solves the motor-error problem. Proc Biol Sci. 2004;271(1541):789-96.

Porrill J, Dean P. Recurrent cerebellar loops simplify adaptive control of redundant and nonlinear motor systems. Neural Comput. 2007;19(1):170-93.

Porrill J, Anderson S, Dean P. Can cerebellar input calibrate collicular topographic maps? . BMC Neuroscience. 2010;11 (Suppl 1):120.

Posner MI. Orienting of attention. Quarterly Journal of Experimental Psychology. 1980;32(3-25).

Prescott TJ, Redgrave P, Gurney KN. Layered control architectures in robots and vertebrates. Adaptive Behavior. 1999;7(1):99-127.

Prescott TJ. Forced moves or good tricks in design space? Landmarks in the evolution of neural mechanisms for action selection. Adaptive Behavior. 2007;15(1):9-31.

Prescott TJ, Pearson MJ, Mitchinson B, Sullivan JCW, Pipe AG. Whisking with robots: From rat vibrissae to biomimetic technology for active touch. Ieee Robotics & Automation Magazine. 2009;16(3):42-50.

Prescott TJ, Diamond ME, Wing AM. Active touch sensing. Philos Trans R Soc Lond B Biol Sci. 2011;366(1581):2989-95.

Redgrave P, Prescott TJ, Gurney KN. The basal ganglia: A vertebrate solution to the selection problem? Neuroscience. 1999;89:1009-23.

Rosenblueth A, Wiener N. The role of models in science. Philosophy of Science. 1945;12(4):316-21.

Roth-Alpermann C, Anjum F, Naumann R, Brecht M. Cortical organization in the Etruscan shrew (Suncus etruscus). J Neurophysiol. 2010;104(5):2389-406.

Sahibzada N, Dean P, Redgrave P. Movements resembling orientation or avoidance elicited by electrical stimulation of the superior colliculus in rats. J Neurosci. 1986;6(3):723-33.

Sawtell NB, Williams A. Transformations of electrosensory encoding associated with an adaptive filter. Journal of Neuroscience. 2008;28(7):1598-612.

Scharfman HE. The CA3 "backprojection" to the dentate gyrus. Prog Brain Res. 2007;163:627-37.

Shimegi S, Akasaki T, Ichikawa T, Sato H. Physiological and anatomical organization of multiwhisker response interactions in the barrel cortex of rats. J Neurosci. 2000;20(16):6241-8.

Sullivan JC, Mitchinson B, Pearson MJ, Evans M, Lepora NF, Fox CW, et al. Tactile discrimination using active whisker sensors. Ieee Sensors Journal. 2012;12(2):350-62.

Taylor GW, Hinton GE, Roweis ST. Modeling human motion using binary latent variables Advances in neural information processing systems (NIPS). 192007. p. 1345.

Teune TM, van der Burg J, van der Moer J, Voogd J, Ruigrok TJ. Topography of cerebellar nuclear projections to the brain stem in the rat. Prog Brain Res. 2000;124:141-72.

Towal RB, Hartmann MJ. Right-left asymmetries in the whisking behavior of rats anticipate head movements. J Neurosci. 2006;26(34):8838-46.

von der Emde G, Bell CC. Nucleus preeminentialis of mormyrid fish, a center for recurrent electrosensory feedback. I. Electrosensory and corollary discharge responses. J Neurophysiol. 1996;76(3):1581-96.

Wald A. Sequential Analysis. New York: John Wiley and Sons; 1947.

Widrow B, Stearns SD. Adaptive signal processing. Englewood Cliffs, NJ: Prentice-Hall; 1985.

Wilson SP, Law JS, Mitchinson B, Prescott TJ, Bednar JA. Modeling the emergence of whisker direction maps in rat barrel cortex. PLoS ONE. 2010;5(1):e8778.

Wilson SP. Figuring Time by Space: Representing Sensory Motion in Cortical Maps [PhD]: Sheffield; 2011.

Wilson SP, Bednar JA, Prescott TJ, Mitchinson B. Neural computation via neural geometry: a place code for inter-whisker timing in the barrel cortex? PLoS Comput Biol. 2011;7(10):e1002188.

Wolpert DM, Miall RC, Kawato M. Internal models in the cerebellum. Trends Cogn Sci. 1998;2(9):338-47.

Woolsey TA, Welker C, Schwartz RH. Comparative anatomical studies of the SmL face cortex with special reference to the occurrence of "barrels" in layer IV. J Comp Neurol. 1975;164(1):79-94.

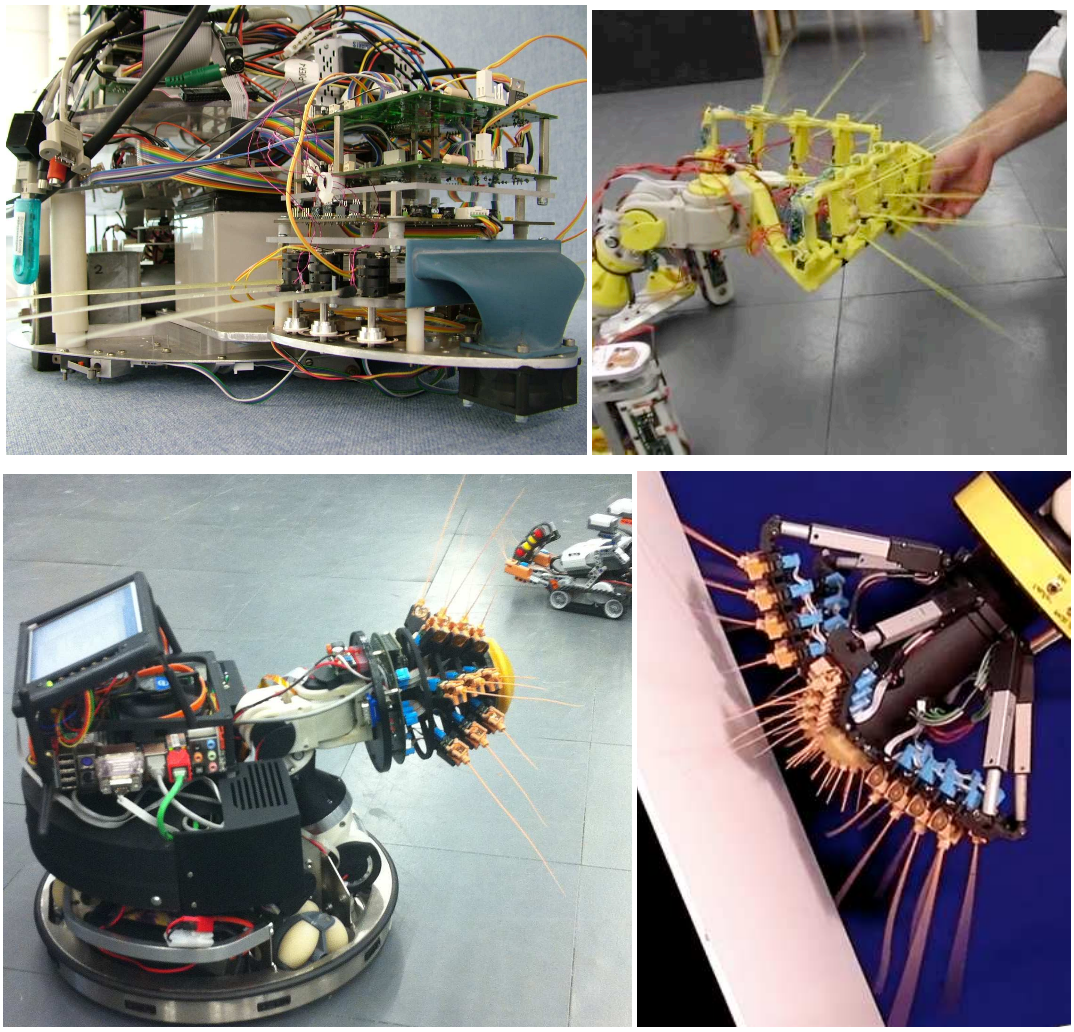

Figure 1

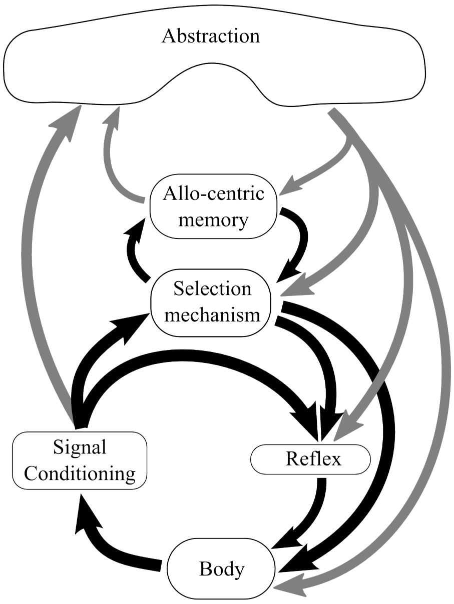

Figure 2

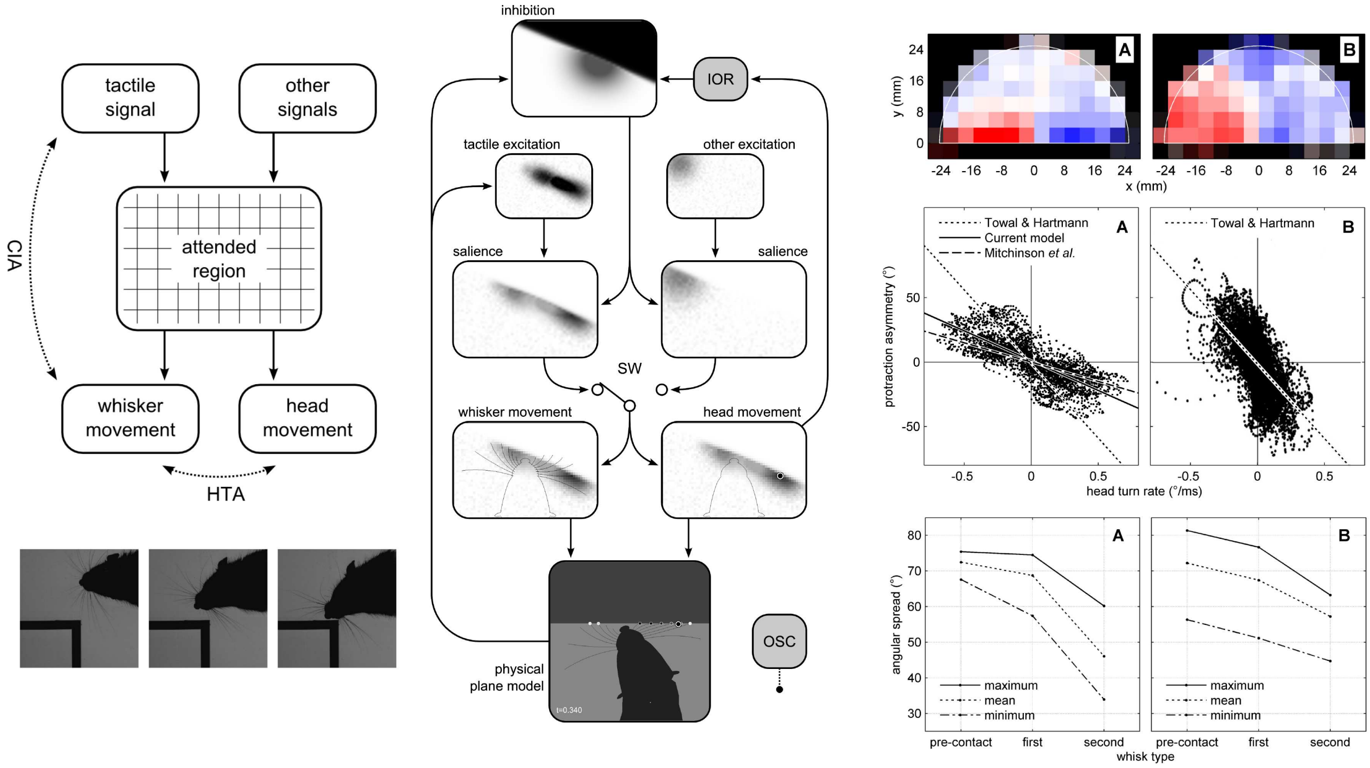

Figure 3

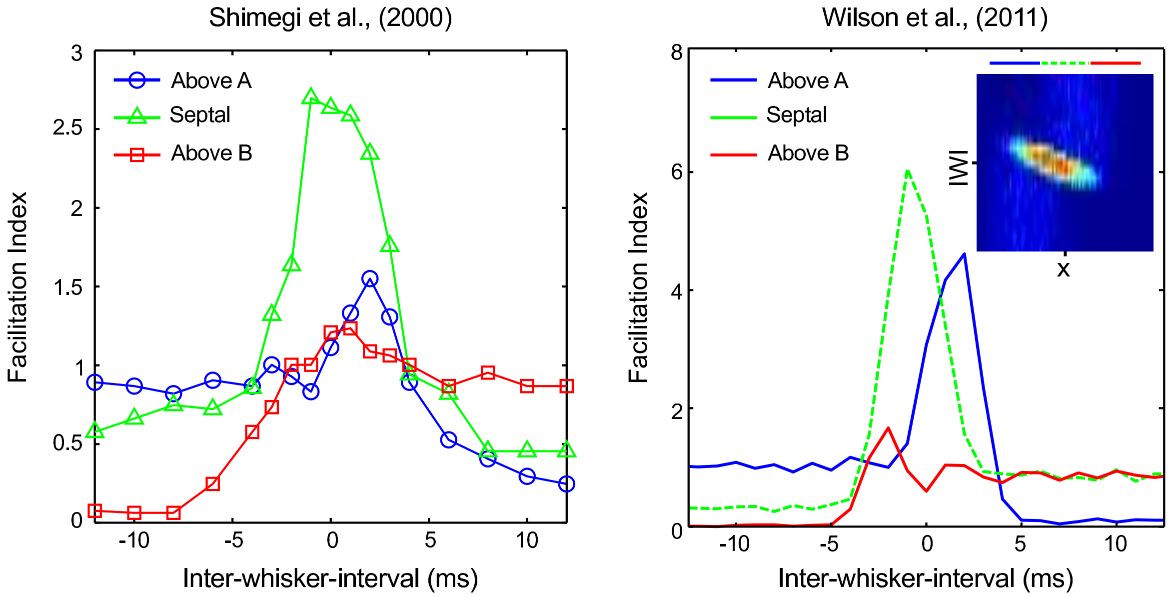

Figure 4

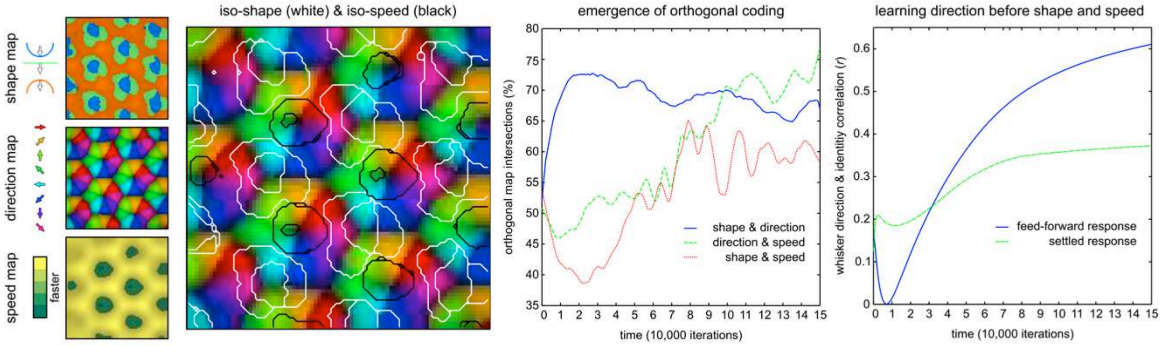

Figure 5

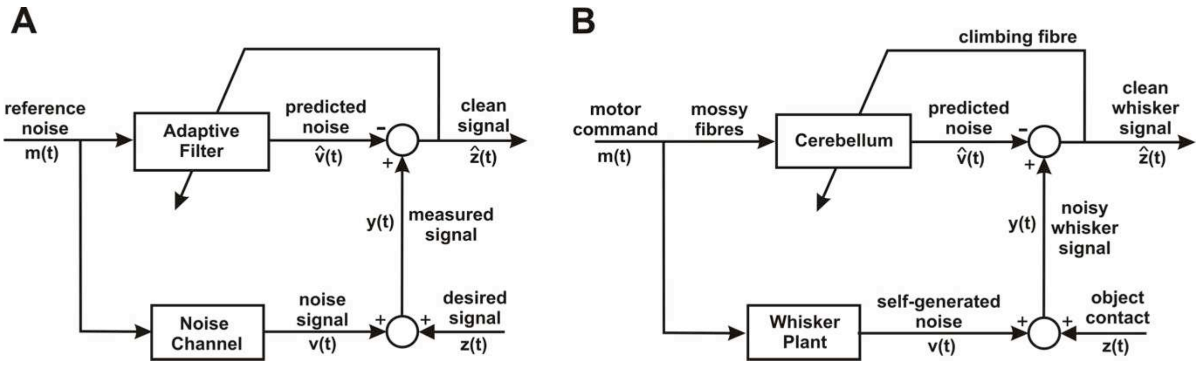

Figure 6

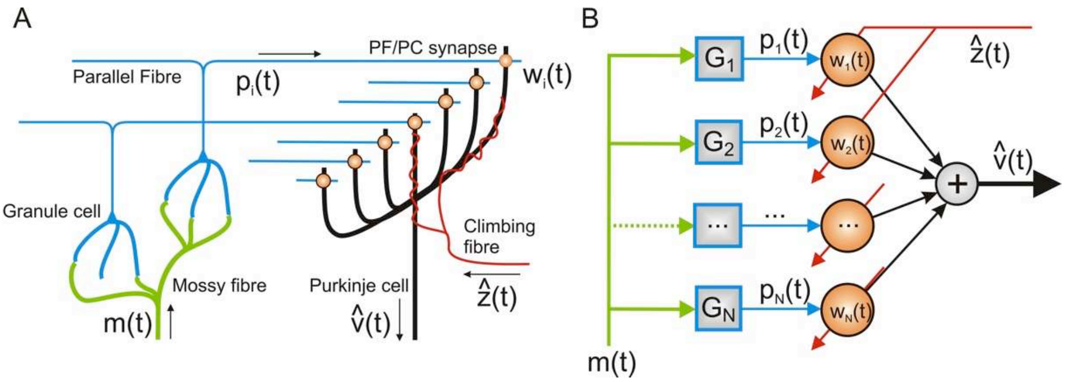

Figure 7

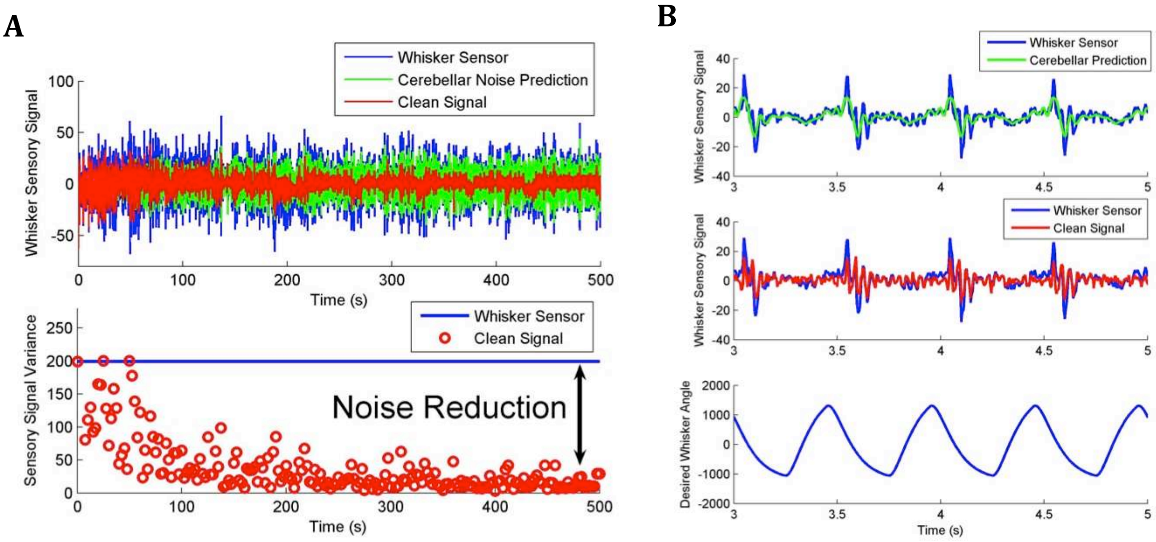

Figure 8

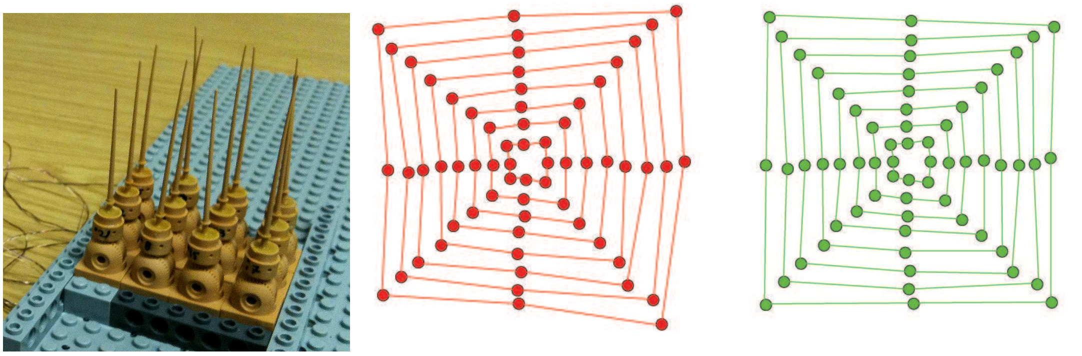

Figure 9

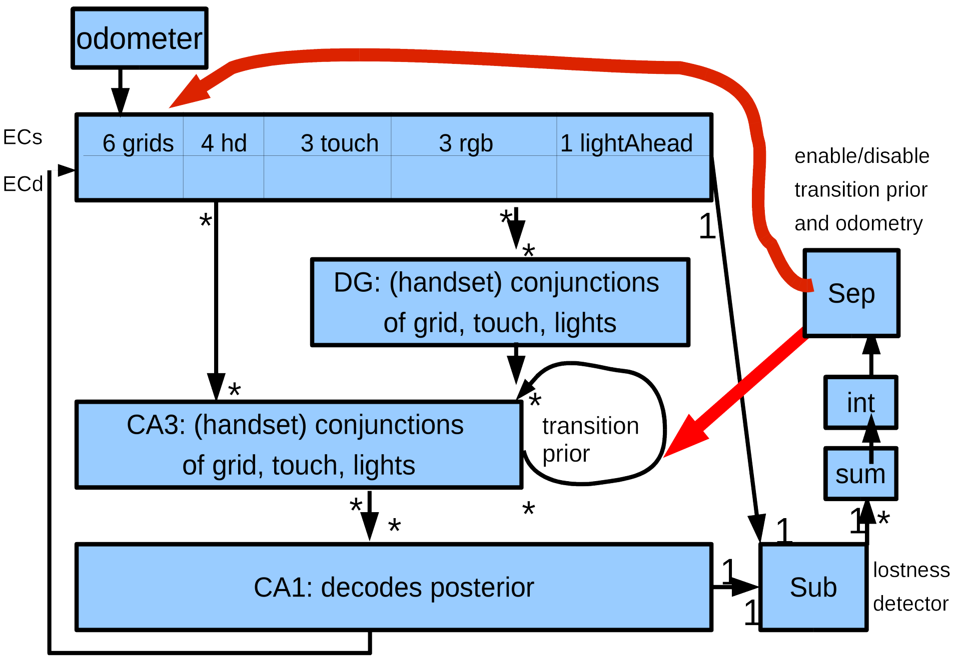

Figure 10

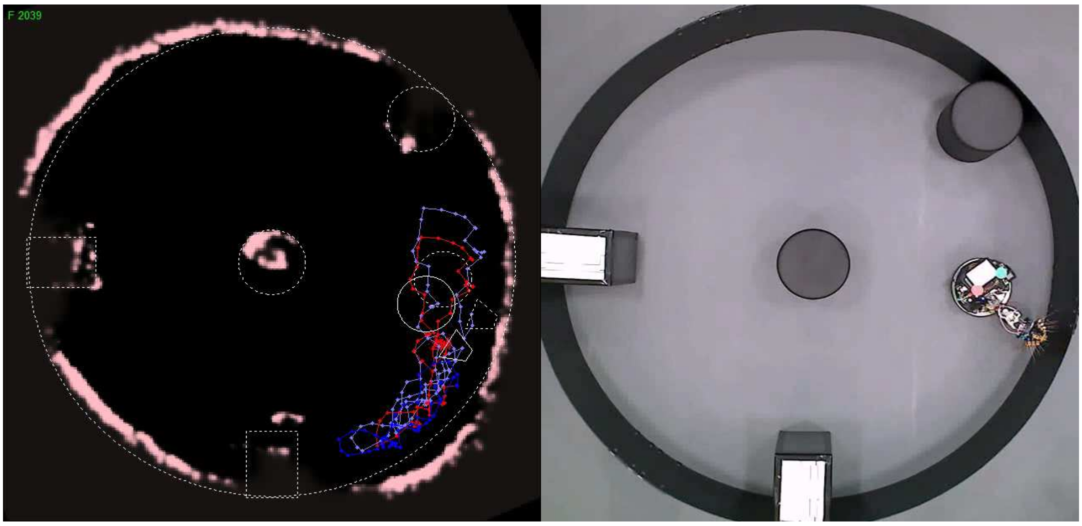

Figure 11

Figure 12