

**Towards a model of the emergence of action space maps in  
the motor cortex**

A thesis submitted for the degree of Master of Philosophy

October 2014

Hiroki Urashima

**Department of Psychology**

**The University of Sheffield**



## Thesis summary

Self-organising maps can recreate many of the essential features of the known functional organisation of primary cortical areas in the mammalian brain. According to such models, cortical maps represent the spatial-temporal structure of sensory and/or motor input patterns registered during the early development of an animal, and this structure is determined by interactions between the neural control architecture, the body morphology, and the environmental context in which the animal develops. We present a minimal model of pseudo-physical interactions between an animat body and its environment, which includes each of these elements, and show how cortical map self-organisation is affected by manipulations to each element in turn. We find that maps robustly self-organise to reveal a homuncular organisation, where nearby body parts tend to be represented by adjacent neurons, but suggest that a homunculus caricature of these maps masks the true organisation as one that remaps from sensory coordinates into ‘action spaces’ for controlling movements of the body to obtain environmental reward. The results motivate a reappraisal of the classic motor cortex homunculus, and demonstrate the utility of an animat modelling approach for investigating the essential constraints that affect cortical map self-organisation.

## **Acknowledgements**

I would like to express my deep gratitude to my supervisor Dr Stuart P Wilson for his patient guidance, enthusiastic encouragement, and useful critiques of this research work, and for help with English in the final write-up. I would also like to thank my friends Siyang Wang and Ryoko Higashi, who patiently supported me. Lastly, I would like to thank my boss, Makoto Hirota, for allowing and encouraging me to undertake research at the University of Sheffield.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	The classic homunculus . . . . .	2
1.2	Activity-dependent self-organisation . . . . .	5
1.3	Self-organising homunculi . . . . .	6
1.4	Body maps versus action maps . . . . .	9
1.5	Motivation behind the thesis . . . . .	11
1.6	Organisation of the thesis . . . . .	13
<b>2</b>	<b>Cortical self-organisation as an interaction of brain, body, and world</b>	<b>14</b>
2.1	Introduction . . . . .	14
2.2	Methods . . . . .	17
2.2.1	A self-organising animat model of brain, body, and world . . . . .	17
2.2.2	Modelling pseudo-physical interactions via shape-matching . . . . .	17
2.2.3	A simple thermotactic controller . . . . .	21
2.2.4	Learning under supervision . . . . .	30
2.2.5	Learning without supervision . . . . .	36
2.3	Results . . . . .	37

2.3.1	Emergent sensorimotor mapping . . . . .	37
2.3.2	Experiment 1 . . . . .	39
2.3.3	Experiment 2 . . . . .	41
2.3.4	Experiment 3 . . . . .	46
2.4	Discussion . . . . .	50
<b>3</b>	<b>Extending the animat modelling framework</b>	<b>55</b>
3.1	Introduction . . . . .	55
3.2	Modelling assumptions . . . . .	57
3.3	Modelling framework . . . . .	59
3.4	Modelling implementation . . . . .	61
<b>4</b>	<b>General discussion</b>	<b>70</b>
4.1	Conclusion: Towards a model of the emergence of action space maps in the motor cortex . . . . .	76

# Chapter 1

## Introduction

### Chapter summary

The primate primary motor cortex has traditionally been considered in terms of the motor homunculus; a functional organisation of responses in different parts of the parietal cortex such that stimulation of adjacent neurons elicits activity in adjacent muscle groups. However, a more recent theory proposes that the functional organisation is better understood as a reflection of ‘action space’; the statistical structure of patterns to the sensory-motor system, where similar body movement primitives incidentally involve adjacent body parts. Models of the emergence of cortical action space maps have looked at the statistical structure of an animals’ movement patterns, and have been able to reproduce aspects of the action space on simulated cortical sheets, which can account for aspects of the motor homunculus as a natural by-product. Modelling cortical development using self-organising maps to a) discover the statistical structure behind the space from which primate motor patterns are drawn, and b) establish predictions about the potential mapping from action space to the two dimensions of the cortical sheet, has led to several novel predictions about motor cortical organisation. An important implication of these results is that whilst the structure of the action space may be relevant to the function of the primary motor cortex areas, the motor homunculus itself need not be. In this chapter we review the evidence motivating

a view of the motor cortex functional organisation as representing a map of action space versus a map of the homunculus body plan. The review suggests that aspects of the brain architecture, the body morphology, and the sensory environment, may contribute to different degrees to establishing the functional organisation of motor cortex. We conclude by proposing a programme of modelling that could help tease these two ideas apart, in which the relative contribution of the brain architecture, the body morphology, and the sensory environment, can be systematically investigated.


## **1.1 The classic homunculus**

Penfield and Boldrey (1937) made electrophysiological recordings from (what later became known as) the somatosensory and motor cortices of human patients undergoing surgery. They found that responses to tactile stimulation varied depending on the location of the stimulus on the body, and on the location of the electrode in the cortical tissue. In the somatosensory areas, responses were functionally organised so that stimulation of adjacent regions of the body evoked selective responses in adjacent regions of the tissue. Likewise, stimulation of adjacent regions of the motor cortex evoked twitches of the muscles in adjacent groups on the body. These observations led to the now classic image of the homunculus maps of the body, laid out on the cortical surface in topological register with the body (see Figure 1.1). Later visualisations of the somatosensory and motor homunculi, which reflected the location and amount of cortical territory dedicated to processing information about each body part, portray an image of the “little man” with areas like that representing the hands and face exaggerated. These images have captured the imagination of neuroscientists for generations since. However, despite a wealth of data obtained that describe the organisation of homunculus maps in almost every mammalian



species tested, relatively little is currently understood about the functional implications (or not) of having maps in the brain.

Since the classic somatosensory and motor homunculus maps were first described, various other cortical maps have been established. For example, in the primary visual cortices, similar map-like organisations describing neuronal selectivity for the location at which visual stimuli appear on the retina have been determined (Hubel and Wiesel, 1962). Embedded within these large-scale maps of the physical layout of the sensory surfaces (e.g., the skin, retinae, cochleae), finer-scale maps appear to reflect a more abstract structure that relates specific features of those stimuli to one-another. For example, embedded within the retinotopic maps in the primate primary visual cortices, organisation of responses appear uniformly and vary smoothly across the cortex with respect to the orientation of elongated bars of light (and dark) presented to the retina: similar orientations evoke specific responses in adjacent regions of the cortex (Swindale, 1998). The pinwheel-dense spatial patterning of orientation maps measured in primate V1 have been described mathematically as representing an optimal trade-off between map continuity, i.e., adjacent features are represented next to each other, and map coverage, i.e., all features of one map are represented within the space occupied by a given feature of another map (see also Swindale et al., 2000), although this description has been very difficult to test directly. However, as discussed by Purves et al. (1992), it is unclear whether spatial patterning in the anatomical organisation of the cortex performs any useful function, or whether such patterns are merely epiphenomena of self-organising developmental algorithms selected for their efficiency in some non-computational regard.



**Figure 1.1. Penfield homunculus (Penfield and Rasmussen, 1950).** *Left:* The sensory homunculus, which is a map of somatosensory areas in the brain. *Right:* The classical motor homunculus of the human brain. The regions of the motor cortex correspond to specific body parts. In this theory, adjacent neurons respond to adjacent parts of the body.

## 1.2 Activity-dependent self-organisation

The dominant theory that explains how these maps arise is known as activity-driven self-organisation (e.g., see Miikkulainen et al., 2005). According to this theory, and variants of it, maps emerge due to a combination of Hebbian ‘cells that fire together wire together’ learning (Hebb, 1949) and local interactions between nearby neurons. Nearby cells form representations of similar features due to unsupervised learning processes that are strongest between adjacent cells due to recurrent interactions between them. This idea is perhaps best understood with reference to computational neuroscience models of the developmental process. The most influential and widely used of these models is the self-organising map algorithm (SOM) developed by Kohonen (1982, 1984); Kohonen and Honkela (2007). The crucial steps of the SOM algorithm, described here in terms relevant to the emergence of body maps, are as follows.

A set of cortical neurons arranged in a two-dimensional grid that corresponds to the cortical sheet, are each associated with a series of synaptic weights that connect them to a similar sheet of input units representing the sensory surface. In such a network, the value of activation of one input unit might be used to represent the intensity of tactile stimulation at a particular body part such as the hand, and activity at another unit might represent the stimulation of a different body part such as the arm. Prior to learning, i.e., prior to simulation of a developmental process in which the cortical neurons interact laterally and apply Hebb’s learning rule, the strength of the synaptic weights from the various input units to each cortical neuron is a random value (corresponding to an absence of functional organisation prior to learning). Simulated development then proceeds by i) setting the activation of the input units so as to reflect a typical combination of body part co-activations that might occur when the body makes contact with an object, ii)

computing the responses of all cortical neurons to the input pattern, as the sum over the input unit activations multiplied by the corresponding synaptic weight between them, iii) locating the cortical neuron that responds maximally to the input pattern, iv) modifying the synaptic weights to the maximally active neuron and those nearby on the cortical sheet by a Hebbian learning rule (which shifts the vector of synaptic weight values towards the vector of input unit activations), and then v) iteratively repeating the learning process for many example input patterns, generated under the assumption that objects contact multiple neighbouring input units together at random locations on the input sheet (e.g., nearby body parts at random location on the body).

### **1.3 Self-organising homunculi**

If the activation of nearby input units on the input sheet is constructed in such a way that nearby input units tend to be correlated, e.g., the ‘arm unit’ is more likely to be activated at the same time as the adjacent ‘hand unit’ than the distal ‘foot unit’, Kohonen’s SOM algorithm will lead to nearby cortical neurons developing preferential responses to inputs from the arm and the hand, and likewise to the strongest responses to the arm and the foot emerging in more distant cortical neurons. The net result of this iterative learning process is an organisation of preferential responses amongst the cortical neurons that reflects the topological organisation of the body parts with respect to each other. Hand and arm representations become co-localised in the cortex, as do arm and shoulder, as do shoulder and neck etc., so as to form a continuous ‘body map’.

Farah (1998) proposed that such mechanisms might explain several distinctive known features of homunculus maps. Whilst primate homunculus maps tend to be continuously organised with hand representations next to arm representations next to shoulder representation etc., they are known also to contain specific discontinuities. For example, maps are discontinuous at the junction between the representation of the hands and the face, which are co-localised in classic images of the somatosensory homunculus despite being distant on the body. Farah proposed that the hands and face representations might be co-localised in somatosensory maps because, as the developing foetus is tucked up in the womb in the foetal position during prenatal development, its hands and face are often touching, and therefore tactile activation of these distal body parts will be highly correlated in time, to the extent that a SOM-like developmental process could lead to those distal body parts becoming represented by adjacent neurons, as any other consistently co-active pair of body part would be.

Farah's hypothesis was tested by Stafford and Wilson (2007), who varied the degree of correlation in the activity between distal input units on the input sheet labelled as the 'hands' and 'face' units. They found that the more highly correlated these units were during training of a Kohonen SOM model, the more likely the two distal input units were to be represented by nearby cortical units as the (developmental) learning process unfolded. This model, and related models representing similar assumptions about how the correlational structure of sensory input is presented to the developing cortex, offers a proof of principle validation of the more general hypothesis that cortical maps are established through activity-dependent self-organising processes. Models making similar assumptions have been successful in explaining almost all known features of somatosensory and visual cortex maps, as well as many of the properties of maps in auditory cortices. Work by Urashima has recently extended this approach to explain aspects of the map organisation in primate vestibular cortical areas (Urashima, 2013). Further-

more, in a series of experiments on primate somatosensory cortex, Merzenich and Kaas (1982) found a dramatic functional reorganisation of the receptive field properties of SI neurons, and dramatic reshaping of the overall somatotopic map representations of intact fingers following finger amputation. This provides evidence for a model in which cortical territory assigned to spared digits invades those no longer receiving input from the missing digits, consistent with simulations of somatotopic reorganisation by self-organisation developed by Ritter et al. (1992). The evidence from such models suggests that the cortex is originally a blank slate, onto which area-specific organisation is assigned by developmental experience.

Strong support for this view is provided by classic experiments in which input to the developing neocortex is redirected; Sur et al. (1999) and Sharma et al. (2000) showed that re-routing subcortical input from the eyes to putative auditory cortex can lead to the emergence of visual map representations (orientation preference maps) in the new target cortex. Similar, albeit circumstantial, evidence from a model of somatosensory map self-organisation by Wilson et al. (2010) showed that the same mechanisms used to model visual orientation map organisation can lead to the emergence of organisations for similar features of somatosensory stimuli when inputs to those mechanisms are constructed to represent associations between tactile stimuli typically experienced during tactile development. Together these studies suggest that activity-driven self-organisation is a general-purpose, modality-independent mechanism for establishing map organisation throughout the developing neocortex.

## 1.4 Body maps versus action maps

Several fundamental questions about maps remain unanswered. Perhaps the most important of these is the question of whether the organisation of cortical maps holds any significance for the nature of the computations that are performed in a given cortical region. A recent review article suggests that whilst variations on the underlying principles of self-organisation can explain the emergence of a variety of map organisation, this is not sufficient to explain the potential benefit to the brain to organising its responses in terms of topological maps (Wilson and Bednar, 2014). These authors claim that whilst map continuity seems to be an intuitively useful aspect of cortical organisation, we are currently lacking a good theory for how map continuity aids cortical computation. Wilson and Bednar explain that despite the intuitive idea that continuous maps minimize structural and metabolic costs associated with connecting representations of similar stimuli over larger distances, current theories such as the influential wiring optimization hypothesis of Koulakov and Chklovskii (2001) do not satisfactorily explain why evolution should have been constrained to represent similar features together. This argument is subtle but important; if we assume that distances between representations of similar features (e.g., similar body parts) should be short, then self-organised cortical maps can be shown to minimize the costs of having long wires, but this does not explain why the distances between the representation of similar features should necessarily need to be short in the first place.

A further challenge to the idea that homunculi represent the statistical structure of developmental input patterns was presented by Parpia (2011), who argued that whilst the assumptions of SOM models can explain the topological structure of cortical maps (such that similar stimuli are represented by adjacent neurons), the consistent ordering of homunculi within a species, e.g.,

with the hands and face represented at the dorsal region of the primary somatosensory cortex, cannot be explained by activity-driven self-organisation alone. Simply put, SOMs can explain why the hands and face are next to each other, but they cannot explain why the face should always be represented at the top. Parpia (2011) proposes that topological discontinuities in the organisation of homunculus maps are better explained in terms of a divergence of cortical input stemming from the thalamus or brainstem. He suggests that discontinuities between the hands and face at the cortical level mirror a segregation of peripheral inputs from these areas of the body well before they reach the cortex, and thus Parpia would accept only a relatively minor influence of self-organising processes in maintaining cortical organisations once they have been established by earlier (presumably genetic) influences on brain development.

Debate about the origin and functional significance of cortical maps motivate us to rethink the implicit assumption underpinning much of the neuroscientific study of maps conducted since Penfield and Boldrey (1937) first discovered the cortical homunculi; the assumption that maps are fundamental units of neural computation. Knudsen popularised the term ‘computational map’ over 30 years ago (Knudsen et al., 1987), but even in these early theoretical arguments there was a hint that a good theory of precisely what computations are aided by map organisation is beyond our understanding - the only concrete example of a computation aided by a map organisation that Knudsen could provide relied on the use of signalling delays between neurons in the map to extract temporal information about stimulus movements, which is a proposal that remains controversial (see Wilson et al., 2011; Wilson and Bednar, 2014).

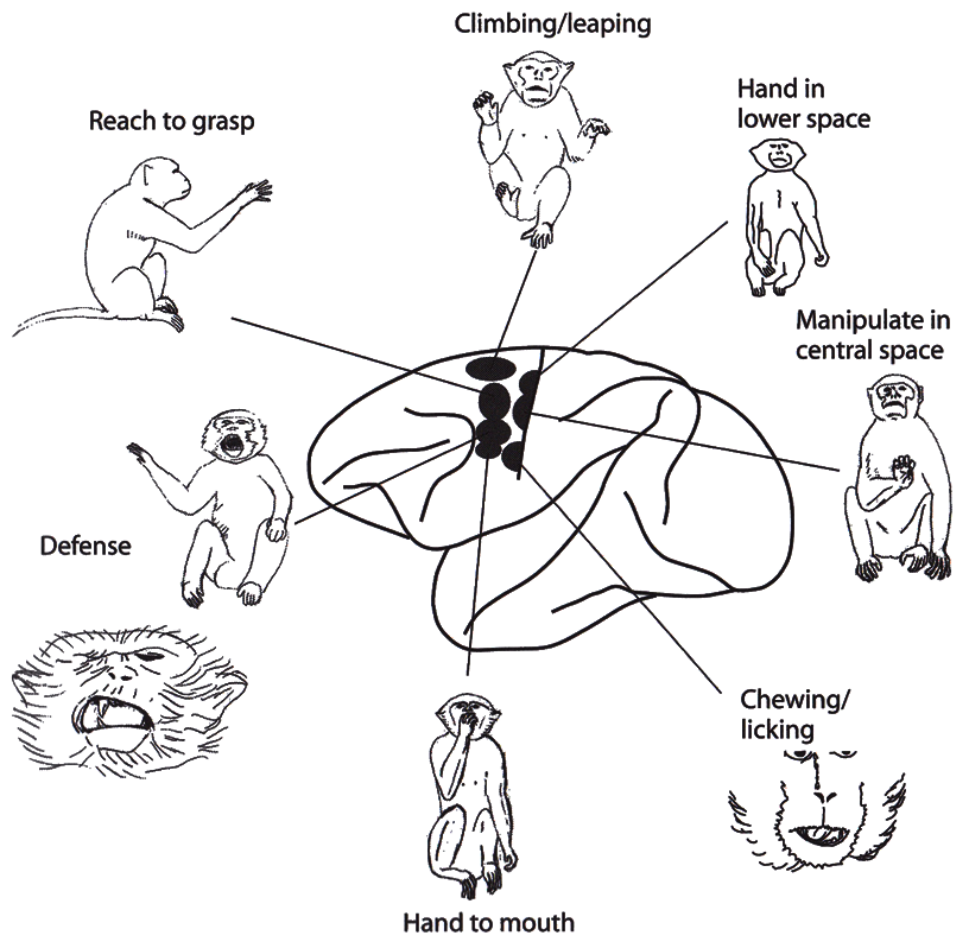
A relatively recent theory proposed by Graziano and colleagues (see Graziano, 2009 for an overview) further challenges the intuition that cortical homunculi represent the layout of tactile and motor units on the body per se. This work has shown that a correlational structure suffi-



cient to lead to the emergence of homunculus maps by activity dependent self-organisation, is present in the organisation of a fundamentally different topological structure, that he attributes instead to ‘action spaces’. That is, the structure that drives map organisation in the motor cortex relates not to the position of body parts in co-ordinates of the body, but rather to a space in which limbs can be positioned to carry out actions and obtain reward from the environment. The idea is that what we refer to as the primary motor cortex contains homunculus maps, not because the cortex is primed to represent adjacent body parts in adjacent cortical regions, but instead because motor repertoires that lead to similar action outcomes incidentally involve similar body parts being moved together. For example, if reaching movements of the arms and grasping movements of the hands tend to result in similar outcomes (i.e., manipulation of an object at a particular external location), such that outcomes associated with both movements are correlated, then the hands and arms will become associated in the cortex by virtue of the movements being correlated. The implication of Graziano’s thesis again suggests a subtle but important departure from the classic interpretation of cortical homunculi as maps of the body - homunculi are perhaps more accurately described as action maps than body maps.

## **1.5 Motivation behind the thesis**

An important lesson from considerations of the works by Koulakov and Chklovskii (2001), Knudsen et al. (1987), Stafford and Wilson (2007), and Graziano (2009), is that whilst activity-dependent self-organisation can explain many of the defining features of cortical maps, a complete understanding of the functional significance of map organisation must be formulated in the context of that organisation being of some use to the animal. To this end, an important first



**Figure 1.2. Motor homunculus representing action space (Graziano, 2009).** The black regions in the figure correspond to specific action groups, which are triggered by electrical stimulation to the motor cortex. Similar actions like climbing and leaping are evoked by a stimulation to nearby neurons.

step is to establish a program of computational modelling that extends the traditional approach to simulating map self-organisation via Kohonen's SOM algorithm, to allow us to study how maps can self-organise in the context of environmental reward. Understanding how maps can be useful requires that we study their emergence as an animal (or simulated agent) is learning to achieve some action outcome, i.e., as it is learning to obtain reward. This should help us to tease apart the subtle distinctions between body maps and action maps, and to start to formulate testable predictions about the functional importance (or otherwise) of topological maps. The work reported in this thesis represents initial steps towards this goal.

## **1.6 Organisation of the thesis**

In chapter 2 we develop an 'animat model' of the interactions between the brain, body, and world in the context of environmental reward, showing how an animated body can be used to disentangle the sources of correlational structure that lead to body maps versus action maps. In chapter 3 we consider the animat model within a broader modelling framework, and explain how this framework can be extended to investigate the emergence of cortical maps for full action spaces. In chapter 4 we discuss the results of the animat model in terms of body maps versus action maps.

# Chapter 2

## Cortical self-organisation as an interaction of brain, body, and world

### 2.1 Introduction

Input driven self-organisation represents the dominant theory for how cortical maps emerge in the mammalian neocortex. According to this theory, three main ingredients are required for topological maps to emerge; Hebbian learning between neurons that are active together in time, local interactions between neurons that are located together in space, and spatial-temporal contingency in the typical patterns of sensory input experienced during development. Accordingly, cortical maps reflect (and may be used to represent) spatial-temporal structure in the world. In addition, for a given organism, three major factors determine the spatial-temporal structure in the sensory world; the morphology of the sensory receptor surfaces and their layout on the body of the organism, the physical environment in which the organism moves and the distribution of rewarding and aversive stimuli within that environment, and the control strategy used by the organism to move the body with respect to its environment. Most of the ideas behind self-organising maps have been developed with respect to representations of sensory features such

as visual edges, tactile stimulation, or auditory tones, as more biological data are available to constrain these models. Some important models have also been developed for mapping sensory inputs to appropriate motor outputs (see in particular Ritter et al., 1992, and also Hoffmann et al., 2010), and we are likewise concerned with modelling the full sensorimotor interaction.

The realisation that sensorimotor experience is an interaction between the brain, the body, and the environment has led some to rethink what the apparent organisation of topological maps in the cortex might truly represent. Graziano and colleagues propose that maps measured in the primate primary motor cortex are organised not in terms of the classic motor homunculus, wherein adjacent neurons correspond with adjacent muscle groups on the body, but rather as a reflection of the structure of ‘action spaces’, where similar actions incidentally involve the coordination of neighbouring muscle groups (see Graziano, 2009, for an overview). A self-organising network model, when trained on patterns that represent typical primate motor repertoires, converged on a functional organisation of motor cortex in terms of the ‘action space’ in which those movement patterns are correlated, whilst also being able to account for somatotopic homunculus maps typically used to characterize motor cortical areas, as a by-product of the network mapping the high-dimensional action space onto the two dimensions of the cortical sheet (Aflalo, 2006; Graziano and Aflalo, 2007). For example, according to this view the muscles of the arms and shoulders that are involved in reaching for and grasping objects are adjacent in motor cortex maps, not because those muscle groups are next to each other on the body *per se*, but because the control of reaching and grasping requires the coordination of overlapping muscle groups.

To explore this idea further we have been developing a novel computational model and simulation platform, in which the movements of simple animat organisms are governed by the

distribution of reward in a physical environment, and in which pseudo-physical interactions between the body and the world drive self-organising map models of cortical development.

In this chapter we report the development of a model via two stages, following a similar development of the model by Ritter et al. (1992). First, we present a supervised learning network, which allows us to state clearly and explicitly what the goal of the animat is (i.e., to orient towards a heat source) and to see how this explicit statement of the goal should, in principle, be reflected in the functional organisation of the network. As well as being a crucial step in model development, this should help reveal how features of the more complicated mappings we expect to emerge through unsupervised learning should be organised. The goal of this modelling effort is to develop an unsupervised algorithm for investigating how interactions between the brain, body, and world shape cortical maps, and to use this model to test how the structure of these maps is influenced by modification to each. We conducted a series of experiments with this model to determine the relative contribution of the body morphology, the distribution of reward in a physical environment, and the neural architecture used to control the movement of the animat body with respect to its environment, to resulting cortical map organisation. Our modelling results demonstrate that homuncular organisation tends to emerge naturally from the interaction between these processes, but that the true organisation reflects a projection from the body surface to an environmentally adaptive space for coordinating motor commands. Our results thus serve as an existence-proof for the plausibility of the hypothesis that homuncular organisation in motor cortex masks the true functional organisation, which may be better understood as a reflection of a remapping from the body topology to an ‘action space’ of environmentally appropriate commands for controlling movements of the body.

## **2.2 Methods**

### **2.2.1 A self-organising animat model of brain, body, and world**

We present a minimal animat model for investigating how interactions between the neural control architecture, the body morphology, and the distribution of reward in the sensory environment affect cortical map self-organisation. The model comprises i) a pseudo-physical simulation for moving a simple animat body in its environment, used to generate patterns of thermo-tactile input for driving simulated cortical development, ii) a simple control architecture for controlling movements of the body to maximise reward, i.e., to move the body towards a heat source, and iii) a self-organising map model with a set of Hebbian-modifiable weighted connections to thermo-tactile sensors on the animat body, and a set of Hebbian-modifiable weighted connections to units that control simple movements of the animat body to optimise thermo-tactile reward.

### **2.2.2 Modelling pseudo-physical interactions via shape-matching**

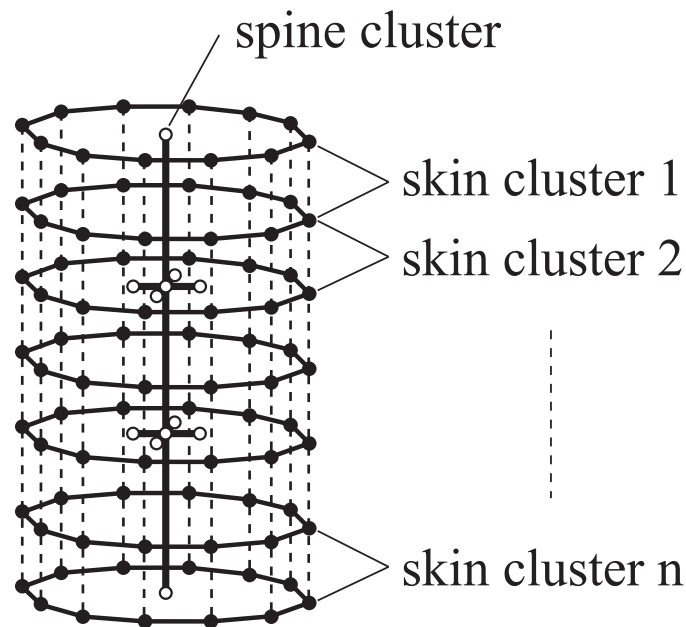
In order to generate input patterns to the self-organising map model we use the influential shape-matching algorithm of Müller et al. (2005), which is used by animators to efficiently simulate realistic-looking deformations of a body surface under gravitational and other environmental forces. We implemented the cluster-based quadratic-deformation variant of the shape-matching algorithm reported by Müller et al. (2005), which allows maximal deformation of the body surface under the influences of gravity and within the constraints imposed by contact between

the animat skin surface and planar surfaces that constitute its physical environment. The basic version of our animat body is a simple point-cloud, arranged so as to form an open-ended cylindrical ‘skin’ surface from a contiguous array of spherical thermotactile elements (taxels) shown in Figure 2.2. Enclosed within the skin are three bone-like structures, which can be rotated relative to each other in three-dimensions to form a rudimentary spinal cord. When the bones are rotated, the shape-matching algorithm is used to reposition the taxels such that in the absence of environmental constraints the skin would (by gradient descent) tend towards the original undeformed point-cloud configuration. The result is an apparently realistic pseudo-physical simulation of skin deformation that we can use to generate patterns of thermotactile stimulation to drive cortical map self-organisation.

The point cloud is divided into two types of clusters: the spine and the skin. The spine cluster, which includes the end-points of the bones, is animated to control the movements of the body. The skin cluster is divided into multiple clusters, each comprising two adjacent rings of taxels, which allows for greater flexibility in the deformation of the body. Bones are rotated as rigid bodies (i.e., by choosing rigid parameters for the shape matching) and the skin is a more highly deformable body (i.e., by cluster-based quadratic parameterisation of the shape-matching algorithm). Figure 2.1 shows how the taxels on the skin surface are related to the bones of the spine.

For full details of the shape-matching algorithm we refer the interested reader to Müller et al. (2005), and we provide here an overview. In order to find the goal positions, from the initial position of points in a pre-defined point cloud,  $x^0$ , and the current position of the point cloud  $x$ ,





**Figure 2.1. Clusters constituting the animat body.** There are two types of clusters: the spine cluster and the skin clusters. Skin cluster has a two-ring shape corresponding to taxels, and contains points of the spine cluster if they are inside the rings. While the skin clusters are deformed based on the quadratic deformation algorithm, the spine cluster is deformed based on the rigid deformation algorithm.

the idea is to find the rotational matrix  $\mathbf{R}$  which minimizes the following equation:

$$\sum_i m_i (\mathbf{R}(\mathbf{x}_i^0 - \mathbf{t}_0) + \mathbf{t} - \mathbf{x}_i)^2 \quad (2.1)$$

where  $\mathbf{x}_i^0$  and  $\mathbf{x}_i$ , defined as 3D vectors, are  $i$ th points in the point clouds,  $\mathbf{x}^0$  and  $\mathbf{x}$ , respectively,  $\mathbf{t}_0$  and  $\mathbf{t}$  are translation vectors, and  $m_i$  is the mass of each point. Equation 2.1 minimises discrepancy between the locations of the points in the point cloud and the original undeformed point cloud by formal gradient descent algorithm.  $\mathbf{R}$  can be solved by the equation,

$$\mathbf{R} = \mathbf{A}(\sqrt{\mathbf{A}^T \mathbf{A}})^{-1} \quad (2.2)$$

$$\mathbf{A} = \sum_i m_i (\mathbf{x}_i - \mathbf{x}_{\text{cm}})(\mathbf{x}_i^0 - \mathbf{x}_{\text{cm}}^0)^T \quad (2.3)$$

where  $\mathbf{x}_{\text{cm}}^0$  and  $\mathbf{x}_{\text{cm}}$ , defined as 3D vectors, are the centre of mass of initial position and current position, respectively, and  $\mathbf{A}$  is a matrix containing rotational part of  $\mathbf{R}$ , which solved by polar decomposition as shown in equation 2.2.

Then, goal positions are solved by

$$\mathbf{g}_i = \mathbf{R}(\mathbf{x}_i^0 - \mathbf{x}_{\text{cm}}^0) + \mathbf{x}_{\text{cm}} \quad (2.4)$$

The next position of point cloud after time  $h$ ,  $\mathbf{x}_i(t+h)$ , can be calculated as

$$\mathbf{x}_i(t+h) = \mathbf{x}_i(t) + h \left( \mathbf{v}_i(t) + \alpha_{\text{sm}} \frac{\mathbf{g}_i(t) - \mathbf{x}_i(t)}{h} + h \frac{\mathbf{f}_{\text{ext}}(t)}{m_i} \right) \quad (2.5)$$

where  $\mathbf{v}(t)$  is the velocity of the point cloud at time  $t$ ,  $\mathbf{f}_{\text{ext}}$  is the external forces applied to the


point cloud, e.g., gravity, and  $\alpha_{sm}$  expresses the stiffness of the animat body. Note that for our implementation of the shape-matching algorithm for simulating the skin, we use a deformation coefficient of 0.1, a rigidity constant of  $\alpha_{sm} = 0.7$ , and a simulation timestep of  $h = 0.01$ .

### 2.2.3 A simple thermotaxic controller

The first task is to get the animat to move according to a simple control strategy. To keep the model simple, we seek to define a mapping from a one-dimensional ‘sensory space’ to a one-dimensional ‘motor space’. Our sensory space is inherently  $n$ -dimensional, for  $n$  taxels on the skin surface, and our three-bones by three-rotation-axes spinal cord yields a motor space with several degrees of freedom. We can simplify matters somewhat by defining a thermotaxic control strategy based on mapping the overall difference in temperature sensed on the left and right halves of the body to a one-dimensional rate of ‘orienting’ behaviour. If we are able to set the animat orienting behaviour with a single function, then we can expect such a mapping to yield a rich, goal-directed, behaviour, mediated by the distribution of reward (heat) in the animat environment (Braitenberg, 1984).

To this end, we must first define some constraints on how the bones of the animat are able to rotate. We choose to rotate the peripheral bones with respect to the central bone by applying sine-waves of equal amplitude and wavelength to the pitch  $x$  and yaw  $y$  axes of the left  $a$  and right  $b$  (see Figure 2.3). Sine-waves take the form:

$$\theta(t) = A \sin\left(\frac{t}{M}\pi + \tau\right), \quad (2.6)$$




**Figure 2.2. Pseudo physical simulation.** The animat comprises a three-bone rigid spinal cord, and a deformable skin of thermotactile sensors. The skin is deformed by the cluster-based quadratic ‘shape matching’ algorithm of Müller et al. (2005). The sensors of the skin detect the temperature of the floor with Gaussian distributed heat.

where  $\theta$  is the angle of the bone,  $\tau$  is the phase offset of the sine wave determining the  $\theta$  value,  $A = \frac{1}{3}\pi$  sets the amplitude and  $M = 64$  timesteps sets the period of the sine wave. With reference to the (arbitrarily chosen)  $x$  axis of bone  $a$ , the phase offset of the remaining three sine-waves can then be used to define a three-dimensional motor control space, parameterised by offset values  $\tau_{ay} \in \{0, 2\pi\}$ ,  $\tau_{bx} \in \{0, 2\pi\}$ , and  $\tau_{by} \in \{0, 2\pi\}$ .


Preliminary simulations of the animat movement revealed a strong tendency for the animat to roll around the arena, rotating about its long axis, which made directing the animat very difficult. To help stabilize the movement we effectively defined a vertical axis through the body by setting the mass of the taxels depending on their height;  $m_i(\zeta) = \frac{100}{n} + 0.2 \sin \zeta$ , where  $\zeta$  is the angle of the taxel from the horizontal (note that the total mass is 100).

With the animat now stabilized we simulated its movements for combinations of  $\tau_{ay}$ ,  $\tau_{bx}$ , and  $\tau_{by}$  that evenly tile the full motor control space, and we quantify for each combination the rate and direction of turning of the animat (see Figure 2.4, and 2.5).

To measure the trajectory of the animat we plot its centre of gravity in the horizontal plane over time. Noting the sensitivity of the centre of gravity to the periodicity inherent in our sine-wave movement strategy, we proceed using an appropriately sub-sampled and low-pass filtered version of the raw trajectory to compensate. We then compute a derivative of the animat trajectory. The sign of the derivative gives us the clockwise/anti-clockwise turning direction, and the average absolute value of the derivative across a trial of 1000 timesteps gives an approximation to the turning rate. We note that for the vast majority of parameter combinations the computed second-derivative of the trajectory remained fairly constant throughout the trial, hence taking the time-average of the first derivative provides a sensible estimate of the turning rate. Finally



**Figure 2.3. Pitch and yaw angles of the peripheral bones.** The peripheral bones of the animat's spine can be angled by applying pitch rotation and yaw rotation. The pitch and yaw angles for joint  $a$  are denoted as  $\theta_{ay}$  and  $\theta_{ax}$ , respectively. Similarly, ones for  $b$  are  $\theta_{by}$  and  $\theta_{bx}$ .



**Figure 2.4. Analysing animat trajectories.** This table shows how the animat trajectory changes when  $\tau_{ay}$  and  $\tau_{by}$  values change from 0 to  $2\pi$  each. The animat movement is categorised into 6 patterns: moving forward, turning clockwise and counter-clockwise, and backward version of them. The blue and red coloured lines indicate actual trajectories and low-pass filtered trajectories of animat, respectively.

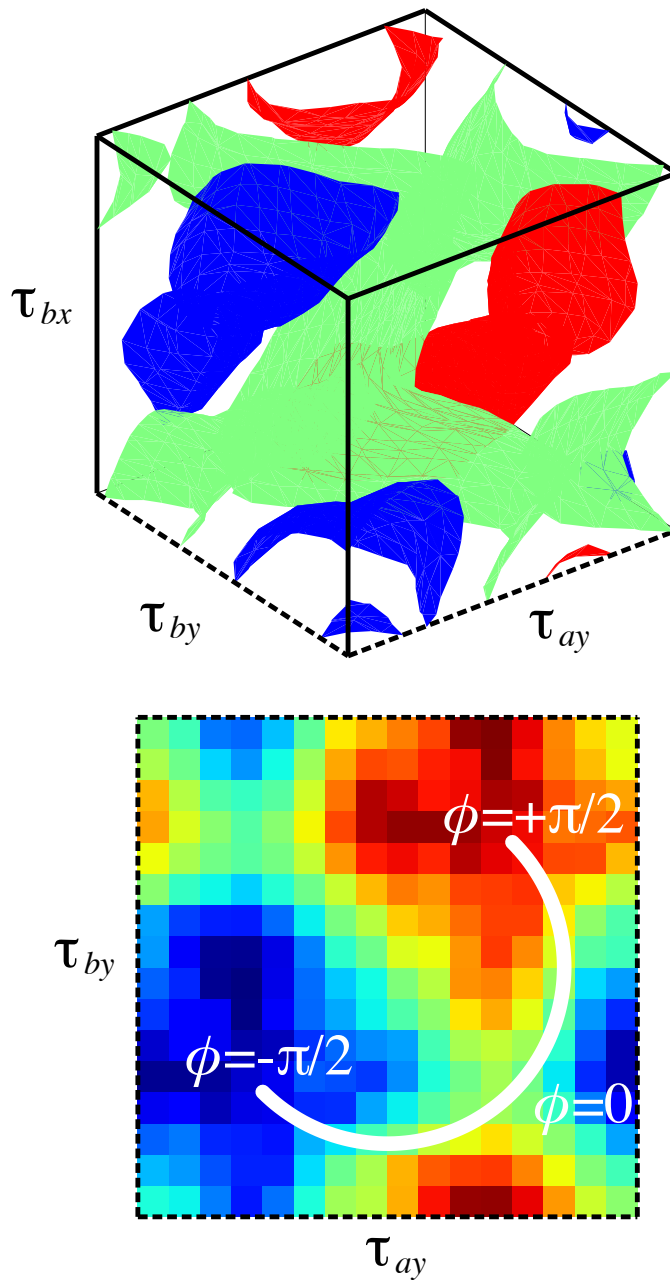
we define the ‘orienting’ behaviour of the animat by weighting the turning rate of its trajectory with the sign of the turning direction, and we plot the orienting at each point in the parameter space. We observed that the mapping from combinations of sine-wave phase offsets to orienting behaviours is locally continuous.

A smooth mapping between sine-wave phase offsets and orienting responses makes it relatively straightforward to define a continuous function curve in the motor control space that generates smooth changes in orienting behaviour. By inspection, we chose for this a simple half-circle curve defined on the plane  $\tau_{bx} = 0$ , which passes smoothly between the maximal observed turning rates in the clockwise and anti-clockwise directions. The resulting function  $f(\phi) : \phi \in \{-\frac{\pi}{2}, +\frac{\pi}{2}\} \rightarrow (\tau_{bx}, \tau_{ay}, \tau_{by})$  that we use to control orienting returns values:

$$\begin{aligned} \tau_{bx} &= 0, \\ \tau_{ay} &= \frac{\pi}{\sqrt{2}} \cos(\phi + \pi/4) + \pi, \\ \tau_{by} &= \frac{\pi}{\sqrt{2}} \sin(\phi + \pi/4) + \pi. \end{aligned} \tag{2.7}$$

Accordingly, animat orienting is completely determined by  $\phi$ . By this one-dimensional scheme, setting  $\phi = 0$  yields forward motion of the animat, setting  $\phi = -\pi/2$  yields maximal clockwise turning, and setting  $\phi = +\pi/2$  yields maximal anti-clockwise turning.





**Figure 2.5.** A one-dimensional mapping of a high-dimensional motor space. The 3D map are divided by the angular velocity value of the animat movement. The green coloured surface indicates the forward/backward movement of the animat. The enclosed areas by the red and blue coloured surface mean the counter-clockwise and clockwise movement with more than specific angular velocity, respectively. The 2D graph is mapped from the 3D one when  $\tau_{bx}$  value is 0. The specified arc is chosen to control the animat movement.

To define our simple thermotactic controller, we simply set:


$$\phi = g(\boldsymbol{\chi}) = \frac{\pi (\sum_{i \in L} \chi_i - \sum_{i \in R} \chi_i)}{2 (\sum_{i \in L} \chi_i + \sum_{i \in R} \chi_i)} \quad (2.8)$$

where  $L$  and  $R$  are the set of left and right taxels, respectively. If temperature is evenly distributed across the two halves of the animat body it will maintain its course, but if temperature is warmer on the left the animat will orient to the right, and vice versa, with the degree of orienting varying monotonically with the absolute difference in temperature.


Each panel of Figure 2.6 shows the trajectory of the animat for given heat source location, defined by a two-dimensional Gaussian profile of decreasing temperature from an arbitrarily chosen point on the arena floor. The temperature at point  $\mathbf{p} = \{p_x, p_y\}$ , relative to a heat source at  $\mathbf{p}_h = \{p_{hx}, p_{hy}\}$  is given by  $T_p = e^{-((p_{hx} - p_x)^2 + (p_{hy} - p_y)^2)/1.4}$ , where the decay parameter 1.4 is in units of the animat body length.

Confirming our intuition, the mapping results in a simple thermotactic behaviour of the animat, leading it to be attracted to the heat source, and resulting in a trajectory like that taken by a moth as it spirals around a lightbulb.


behaviour 1




behaviour 2



behaviour 3



behaviour 4




**Figure 2.6. Trajectories of animat using a simple thermotaxic controller.** In each figure, the black dot designates the start point of animat, and the red circle indicates the heat centre. Animat draws flower-like trajectory in every case.

### **2.2.4 Learning under supervision**

To simulate input-driven cortical map self-organisation, and thus to derive predictions about how the form of cortical maps is affected by the movement of the animat relative to its environment, we adapt a model originally outlined by Ritter et al. (1992). This model itself is an extension to Kohonen's self-organising map model (Kohonen, 1982, 1984; Kohonen and Honkela, 2007) to include an additional set of Hebbian-modifiable weighted 'output' connections from a cortical sheet to a small set of control units, and a mechanism for reward-based learning.

The demonstration of the network learning used by Ritter et al. (1992) was an application to a pole-balancing task, i.e., the network could learn to map deviations of a simulated pole from a target vertical angle, to appropriate forces to be applied at the base to keep the pole vertically balanced. The extension of this modelling approach that we develop for thermotactic control of the animat is conceptually similar, however we end up with a rather different formulation of the model.

The objective for our network will be to discover an appropriate mapping between body surface temperatures and appropriate motor commands  $\phi$  that can generate the thermotactic behaviour that in the previous section was hard-coded into the animat. We can get an idea of how the unsupervised network should configure itself to achieve such a mapping by training the weights of the network first using supervised learning. Detailing this process of supervised learning allows us here to introduce the basic architecture of the network that will later be built upon for unsupervised learning.



**Figure 2.7. A self-organising animat model.** We started from hard-coded thermotaxis, expand it to supervised learning (left), then to unsupervised learning (right) based on heat reward. The model consists of a cylindrical animat body in its environment in the context of heat reward, a self-organising map in hexagonal shape connected from thermotactile sensors on the animat body, and a control architecture driving the animat body based on the output from the map.

Our animat cortex comprises 3169 neurons arranged on a regular hexagonal lattice, such that each neuron  $j$  has a coordinate  $x$  and  $y$ . Following the model of Ritter et al. (1992), each neuron has two sets of weighted connections; one set of afferent weights that connect directly to thermotactile sensors on the surface of the animat body  $w_{\text{in}}$ , and a set of output weights  $w_{\text{out}}$  from each cortical neuron.

On each timestep  $t$ , the response  $\rho$  of cortical neuron  $j$  is determined by the current weighted activity of the taxels  $\chi$  on the animat body:

$$\rho_j(t) = \sum_i w_{\text{in}_i} \chi_i(t) \quad (2.9)$$

According to the Kohonen self-organising map algorithm, the neuron with the maximum response  $\rho$  is crowned the ‘winning’ neuron, and we denote its properties using e.g.,  $\hat{x}$ ,  $\hat{w}_{\text{in}}$ ,  $\hat{w}_{\text{out}}$  etc. Once the winning unit is identified, a ‘neighbourhood’  $h$  is defined around it, with elements:

$$h_j = \begin{cases} e^{-d_j/D}, & d_j < \omega \\ 0 & \text{else} \end{cases} \quad (2.10)$$

where  $d_j = \sqrt{(\hat{x} - x_j)^2 + (\hat{y} - y_j)^2}$ ,  $D$  is a decay function, and  $\omega$  is a cut off function. The two functions are defined as  $D = 0.5^{1+t/t_{\text{max}}}$ , and  $\omega = e^{-0.2t/t_{\text{max}}}$  where  $t_{\text{max}}$  is the total time of the simulation. As the simulation runs the network learns by Hebbian modification of the weights in the neighbourhood, which encourages similar neurons to learn to represent similar input/output patterns:

$$\begin{aligned}\Delta \mathbf{w}_{\text{in}} &= \epsilon \alpha \mathbf{h}(\boldsymbol{\chi} - \mathbf{w}_{\text{in}}), \\ \Delta \mathbf{w}_{\text{out}} &= \epsilon \alpha \mathbf{h}(\hat{w}_{\text{out}} - \mathbf{w}_{\text{out}}),\end{aligned}\tag{2.11}$$

where  $\epsilon$  is the learning rate, and we here set  $\alpha = 1$  to be a placeholder in lieu of its later use during unsupervised learning. In the unsupervised learning, the output weight of the winning neuron will be used to specify the animat orienting behaviour  $\phi$  at each iteration. However, for the supervised learning case, the value of  $\phi$  is chosen to be:

$$\phi(t) = \left(1 - \frac{t}{t_{\text{max}}}\right)g(\boldsymbol{\chi}) + \frac{t}{t_{\text{max}}}\hat{w}_{\text{out}}.\tag{2.12}$$

Hence at the beginning of the simulation the orienting of the animat is entirely determined by our hard-coded thermotactic controller, but over time control of orienting is subsumed by the network weights. Essentially, this process gradually consolidates a look-up table that maps from  $\boldsymbol{\chi}$  to appropriate  $\phi$ , into the network weights (see Ritter et al., 1992), such that the network weights eventually take full control of the animat.

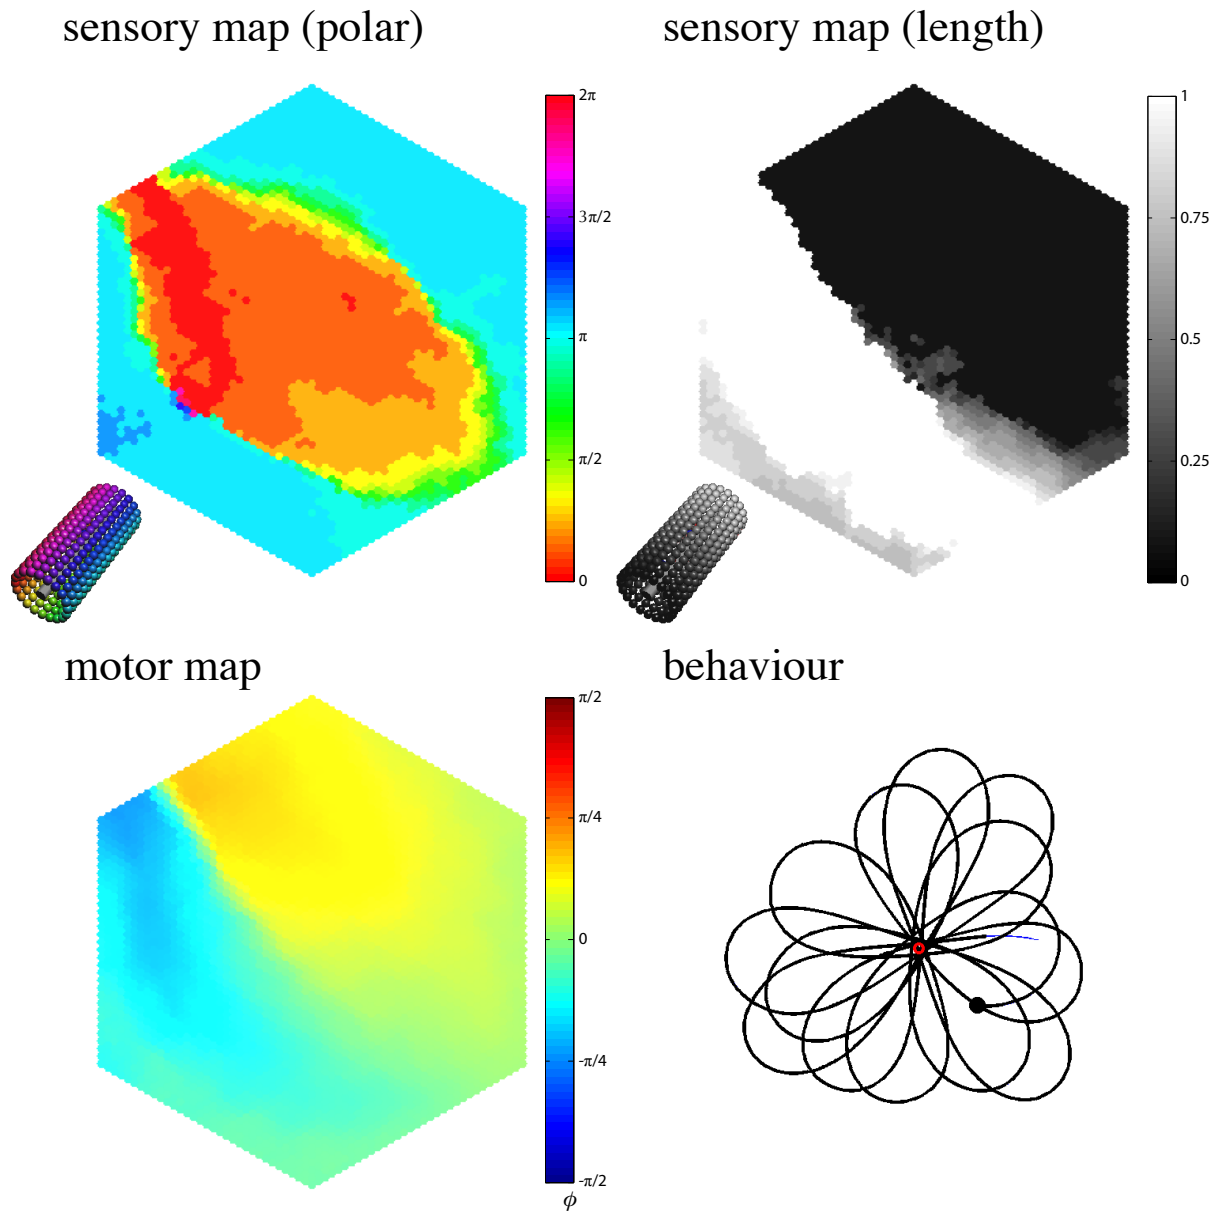
For all animat simulations reported in this thesis, 250 learning trials, each consisting of 10,000 simulation time-steps, were run. On each learning trial the animat was dropped onto the arena floor from a height of around 1 body length, and at a random  $x, y$  location in the environment in the range  $[-0.8, 0.8]$  animat body lengths. Testing trials were always run with learning turned off for duration of 200,000 timesteps unless otherwise stated.

The supervised learning process allows us to derive for our animat cortex a set of network weights that represent thermotaxic control (see Figure 2.8). In these plots the sensory maps show for each neuron the location on the body of the taxel to which that neuron has the largest synaptic weight. For the sensory weights we show two plots, one where neurons are coloured to show the radial location of the preferred taxel on the body, and the other using grey-scale to show its position with respect to the long axis of the animat body. Motor map plots show the value of the  $\phi$  (i.e., the motor output weight) of the neuron against a jet colourmap where blue is low (i.e., if a blue neuron is the winning neuron the animat will turn clockwise), green is zero (i.e., the animat will go forward), and red is positive (i.e., the animat will turn counter-clockwise).

In these plots we are drawn to two important features. First, the representation of the body surface, as revealed in a plot indicating the taxel on the animat body surface for which each neuron develops the strongest synaptic weight  $w_{in}$ , is locally continuous, such that adjacent neurons represent adjacent points on the animat skin. Thus, the network has learnt a homuncular representation of the body. Second, the final output weights, which determine the orienting behaviour of the animat, vary smoothly across the cortical sheet.

The final panel in Figure 2.8 shows the trajectory of a representative animat after supervised learning, for an identical environment and animat starting position. The close correspondence between the two trajectories reassures us that the final network weights have consolidated the thermotaxic control strategy.





**Figure 2.8. Self-organised sensory and motor maps for thermotaxis via supervised learning.** The colour in the sensory maps shows which taxel neurons prefer to in the position of polar coordinate and length coordinate. The motor map shows the preference of  $\phi$  value; if colour is red,  $\phi$  value is high, which means the animat turns anti-clockwise. If colour is blue, value is low, which indicates the animat turns clockwise. In the case of supervised learning, the motor map is relatively smooth. Behaviour shows a trajectory that animat took. Black dot is the start point of the animat and red circle is the heat centre.

### 2.2.5 Learning without supervision

The learning procedure outlined above gives us a means of establishing a cortical mapping that translates patterns of sensory stimuli into appropriate actions based on our thermotactic control strategy, and thus it gives us an idea of the configuration of sensory input and motor output weights that we should expect an unsupervised learning algorithm to learn were it to discover to maximise heat reward. For the network to discover such a mapping without explicit instruction, we first define a reward function, which for our thermotactic animat is simply the average temperature of its skin  $r = \frac{1}{n} \sum_{i=0}^n \chi_i$ , and then we adapt the network weights only for positive changes in reward  $\Delta r = r(t) - r(t - 1)$ :

$$\alpha = \begin{cases} 1 & \Delta r > \kappa \\ 0 & \text{else} \end{cases} \quad (2.13)$$

Here  $\kappa$  sets a dynamic threshold, maintained for each neuron, which is modified only for neurons in the winning neighbourhood by:

$$\Delta \kappa = \gamma(\Delta r - \kappa), \quad (2.14)$$

where  $\gamma = 0.05$  sets a constant for the dynamic threshold. The dynamic threshold essentially tracks the contribution of the winning neuron to the value of the reward function, permitting only weight changes that, relative to the recent history of that neuron, increase the reward. Again following Ritter et al. (1992), the process is made stochastic by setting the orienting behaviour of the animat  $\phi = \hat{w}_{\text{out}} + \hat{\sigma}\eta$ , where  $\eta$  is normally distributed noise, and a variance associated

with each neuron in the winning neighbourhood is gradually reduced after every timestep using to  $\Delta\sigma = \epsilon h(t/t_{\max} - \sigma)$ .

A full summary of the algorithm is reported in Figure 2.9.

## 2.3 Results

We concentrate now only on results from the unsupervised learning case. We refer to this as the standard animat model, to distinguish it from three manipulations that we make to the simulation. We first report observations about the cortical map organisations that emerge from the network, and about the subsequent behaviour of the animat with respect to the goal. Then, in three separate experiments we change the nature of the control task faced by the network and examine the effect on map organisation and behaviour.

### 2.3.1 Emergent sensorimotor mapping

The standard animat model ends up developing a continuous map of the body in the afferent weights, as shown in Figure 2.10. The plot shows the preferred taxel of each neuron (that to which it develops the strongest synaptic weight). Notice that the sensory map is dominated by the representation of a relatively small proportion of the circular axis. The reason why the neurons corresponding to the small part of the body dominates these sensory maps is that the body was weighted, i.e., the bottom is heavier, and the experiment in Figure 2.11 shows that

**1. initialisation**

- 1.1. define point cloud (including skin and spine)
- 1.2. define mass of each point
- 1.3. divide points to clusters
- 1.4. define environment
- 1.5. define self-organising map

**2. learning/testing trials****2.1. trial initialisation**

- 2.1.1. reset the position of the point cloud
- 2.1.2. precompute centre of mass of clusters
- 2.1.3. update the position of heat centre

**2.2. trial body**

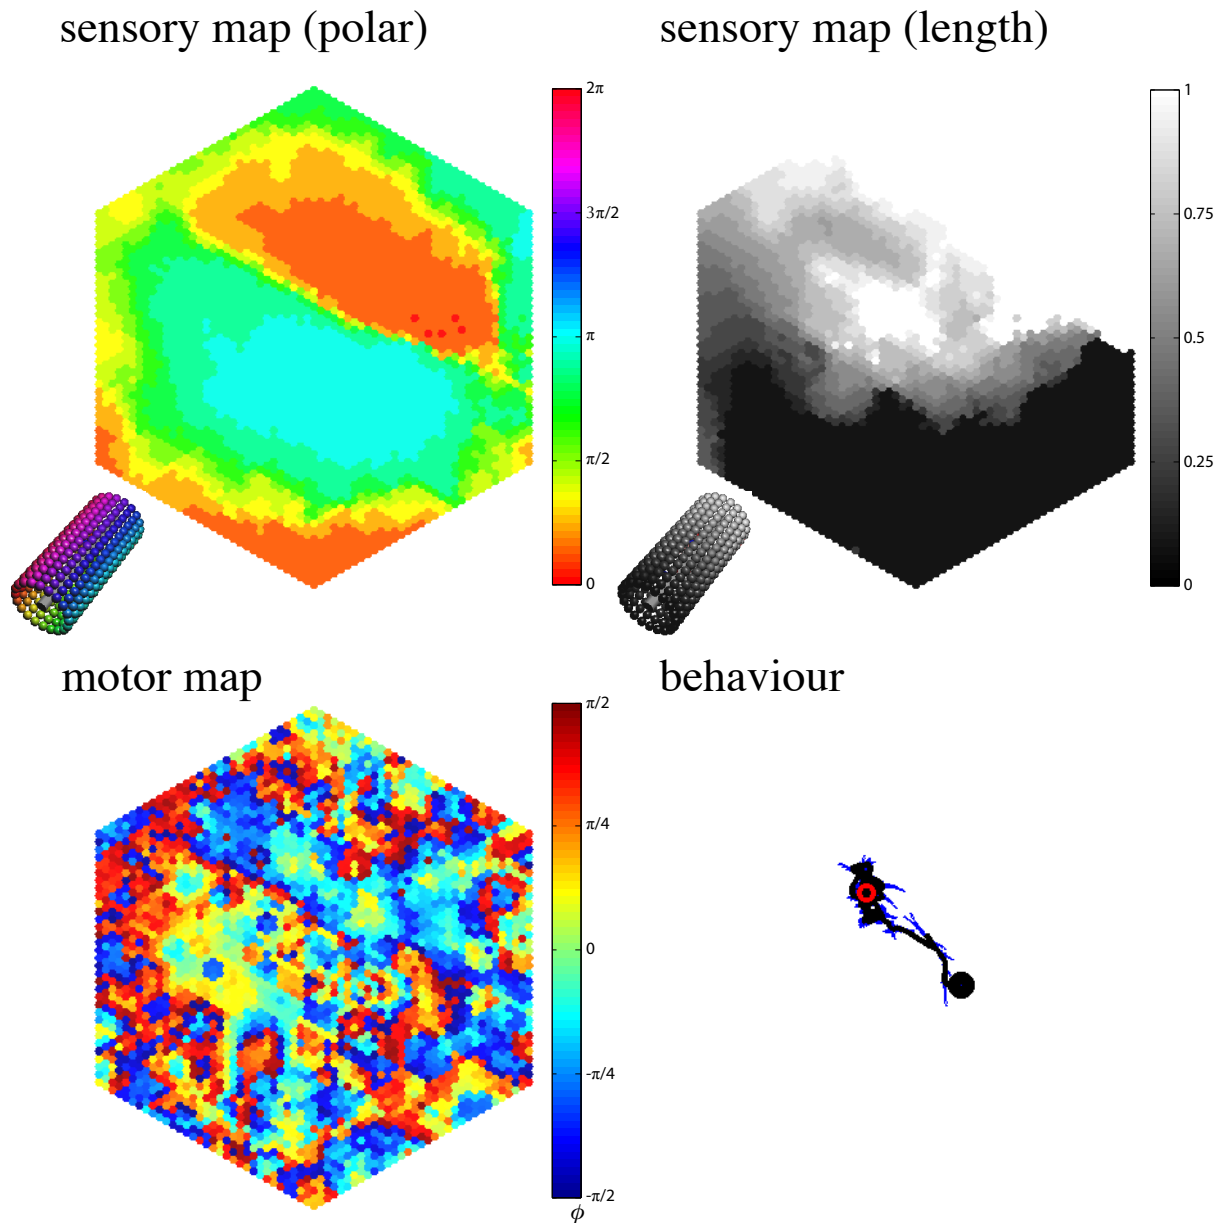
- 2.2.1. rotate peripheral bones of spine
- 2.2.2. update centre of mass of spine cluster
- 2.2.3. calculate external force (gravity)
- 2.2.4. detect/resolve collision with floor
- 2.2.5. calculate goal position of point cloud
- 2.2.6. move point cloud towards goal position
- 2.2.7. sense temperature from skin in contact with floor
- 2.2.8. locate winning neuron
- 2.2.9. calculate output based on the position of winning neuron (with/without noise)
- 2.2.10. update input and output weights

**Figure 2.9. Detailed algorithm of the animat model.** It consists of mainly two parts: initialisation and trials. After the initialisation, multiple learning trials and testing trials are executed.

the bias goes away when the weighting of the taxels is removed. Notice that the contours in the plots of the sensory weights coloured independently for the radial and long axes of the animat body are generally orthogonal, hence the network has learnt a smooth map of the layout of the body. Discontinuities in the length plot tend to be related to the two ends of the animat. Furthermore, the motor map showing the disrupted representation does not correspond to either of the sensory maps, which indicates that the motor map does not represent the body layout like Penfield's motor homunculus in Figure 1.1.

### **2.3.2 Experiment 1**

In the simulations reported thus far the mass of the taxels was unevenly distributed around the cyclic axis of the animat, so as to effectively give it a vertical axis and hence give it a chance to stay upright. In the first experiment we remove this constraint by assigning equal mass to all taxels (in fact these were the very first experiments we ran before deciding to present the stabilized animat as our basic model). The overall mass of the evenly and unevenly weighted animats was held constant. When we run the simulation, we observe that the animat has a tendency to roll around the arena (see Figure 2.11). Typical trajectories around the arena appear far less well directed to the goal than in the standard model. In many cases the animat will never reach the goal, and in situations where it does so it is unclear whether this occurs simply by chance. Once the animat begins to move and rotate, the task may in fact be impossible for the animat to achieve, simply because of the constraint that bones were only allowed to rotate on one axis, and therefore corrections for the rolling by orienting around an orthogonal axis were not possible. Analysis of the trajectory of the animat is therefore less interesting to us than the cortical maps that emerge. Particularly clear in the colouring of the sensory



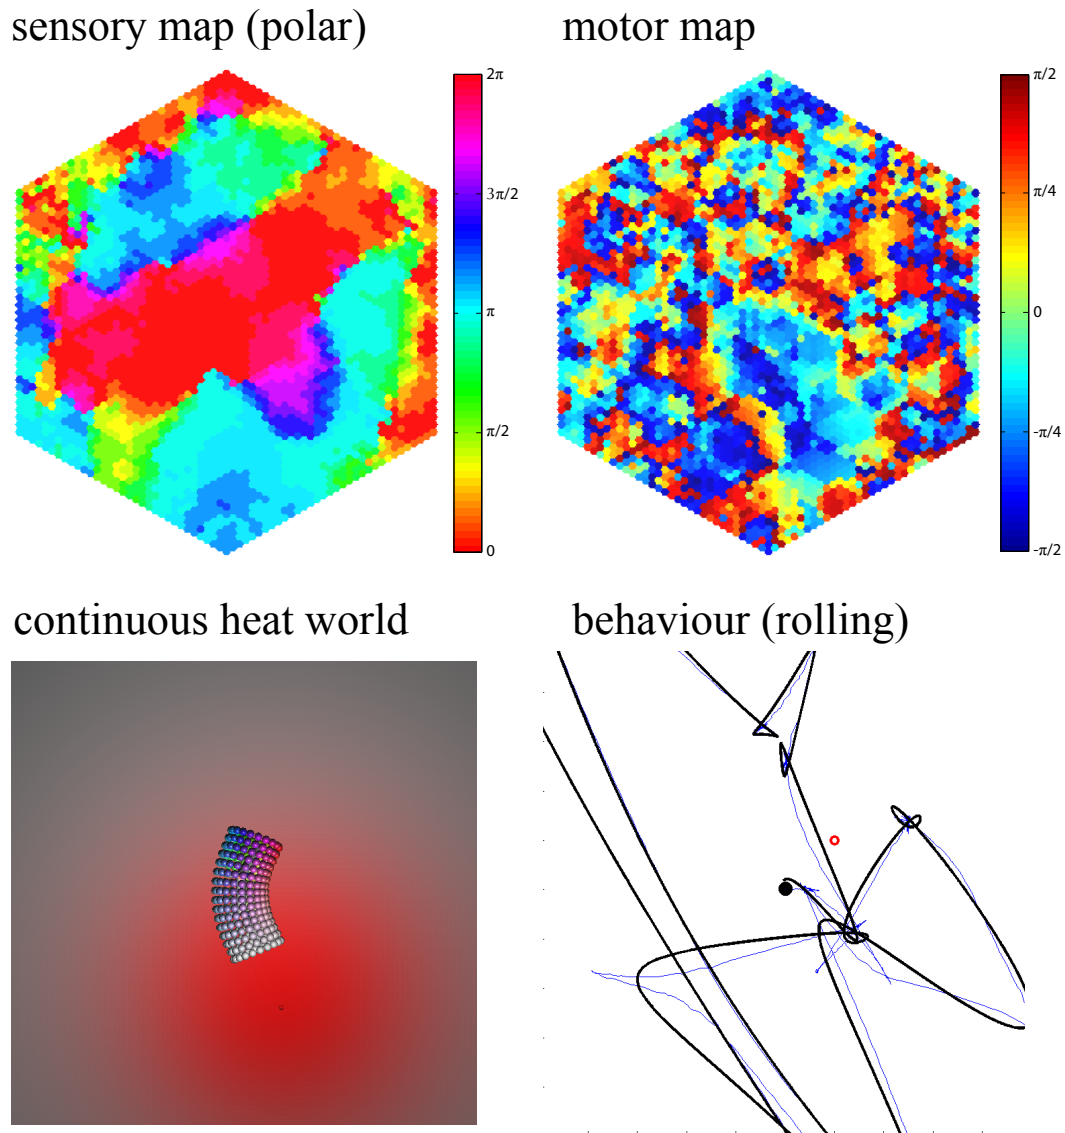
**Figure 2.10. Self-organised sensory and motor maps for optimising heat reward via unsupervised learning.** The colour in the sensory maps shows which taxel neurons prefer to in the position of polar coordinate and length coordinate. The motor map show the preference of  $\phi$  value. In the case of unsupervised learning, the motor map is completely disrupted. Behaviour shows a trajectory that animat took. The animat was almost directly going to the heat source.

connection weights by the cyclic index of the preferred taxel, is that a more complete sensory homunculus map tends to emerge, now with taxels labelled by the full range of colours being represented. The sensory map is again largely continuous with respect to both axes of the body, and approximately even coverage of skin surface is apparent. This can be explained in terms of the nature of the sensory inputs; by rolling around the arena, all taxels end up generating input, and hence all have the opportunity to be candidates for Hebbian association with the winning neuron. This manipulation makes it clear that what we might interpret as an improvement in the continuity and coverage of the mapping of the body has little bearing on the success of the animat in its task of obtaining heat reward.

### **2.3.3 Experiment 2**

In experiment 1 we changed the nature of the physical interaction between the body and the environment with respect to the cyclic axis of the animat, to explore the extent to which this interaction may be reflected in biased distributions of feature preferences in the self-organised maps. In this second experiment we instead investigated the extent to which the shape of the body, i.e., the body morphology, affects the form of the emergent homunculus mapping, this time along the longitudinal axis of the animat. Between conditions we systematically morphed the shape of the body from a cylinder to a cone, whilst keeping all other parameters (distribution of mass, number of taxels, relationship of taxels to bones etc.) constant as per the standard model.

For successive rings of taxels indexed by  $i$  that comprise an animat body of unit length, the radius of each is given by:



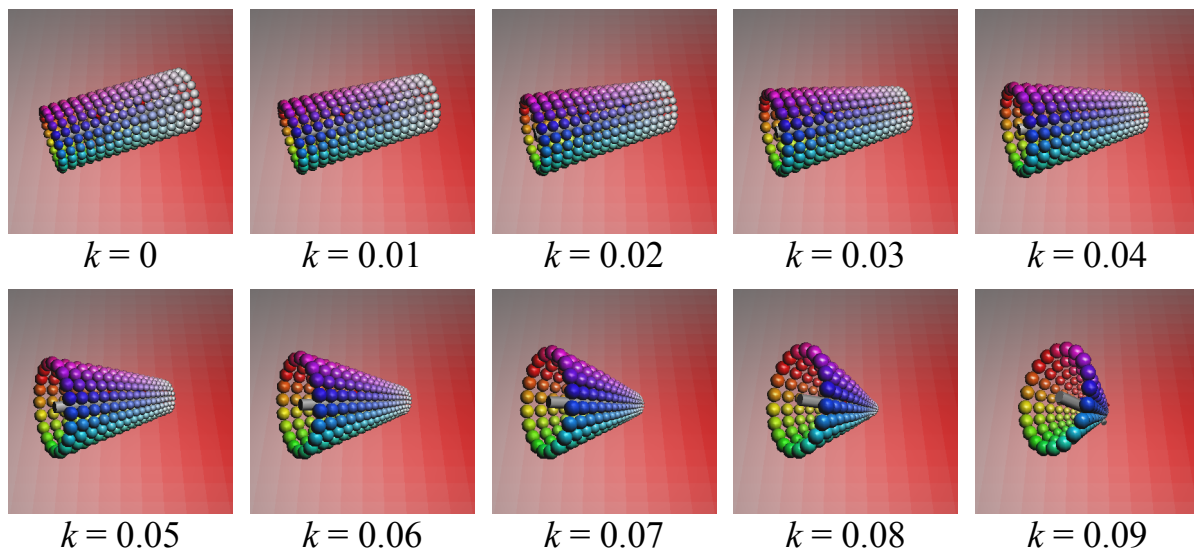
**Figure 2.11. Manipulating the body morphology.** The animat is no longer bottom heavy, which leads to it rolling around in the arena and generating a full sensory homunculus map.




$$z_i = \frac{1 - k + 2\frac{n}{\pi} \sum_{j=0}^{i-1} r_j}{2k + \frac{n}{\pi} \sqrt{(2k)^2 + 1}},$$

where  $n$  is the number of taxels in each ring and  $k > 0$  changes the angle of the cone. A contiguous point cloud of taxels of radius  $r_i = \pi(z_i 2k + 0.5 - k)/n$  is maintained if each is placed at  $(x_{i,\theta} = (z_i 2k + 0.5 - k) \cos \theta, y_{i,\theta} = (z_i 2k + 0.5 - k) \sin \theta, z_i)$ , where  $\theta \in (1, \dots, n) \frac{2\pi}{n}$  gives the position of each taxel in the ring.

Cone-shaped animats of various angles set by  $k$  were then simulated as before. We found that the sensory map coloured by the location of the preferred taxel along the length of the body had showed a large difference as  $k$  increased. Figure 2.13 shows the ratio of units whose preferred location on the body is on the front half, as a function of  $k$ . When  $k$  was zero (i.e., as for the cylindrical animat used for the previous experiments), the ratio was around 50% meaning that maps represented the body evenly. However, as  $k$  was increased from 0 to 0.07, the ratio gradually increased before jumping up to 1 when  $k$  was 0.09, at which point we observed a strong tendency for the animat to get stuck in a position where its slimmer tail was pointing upwards. This serves as a clear demonstration that the morphology of the animat plays an important role in shaping the functional organisation of the cortex; more collisions occurring between the floor and the taxels at the front of the animat, lead to an exaggerated representation of those locations in the sensory map, in line with what we expect from Hebbian learning.



**Figure 2.12. Manipulating the body shape.** As  $k$  increases, one side of the cone getting bigger and the other side smaller.



**Figure 2.13. Manipulating the body shape.** The graph shows how the ratio of units that prefer front side of the body over  $k$  value.

### **2.3.4 Experiment 3**

In experiments 1 and 2, continuous homunculus body maps emerged in the distribution of receptive fields across the afferent weights to the network. These smooth body maps were affected by the nature of the physical interaction between the body and the world, due to manipulations that we made to the body morphology of the animat. Experiment 3 was designed to explore a different question; how does the spatial distribution of stimuli across the environment affect the organisation of the body map? Specifically, experiment 3 was designed to test a) to what extent is smoothness in the sensory world important for establishing smooth body maps?, and b) do smoother body maps lead to increased behavioural performance? Based on our intuitions from the study of self-organising maps, we expect that the answer to the first question will be that smoother environments will lead to smoother body maps, and as we will show this is indeed what we observe. The second question is perhaps more interesting; if it is possible to construct a regime in which body map smoothness does not predict behavioural success, then this would be an important first step towards defining an action map as distinct from a body map.

Experiment 3 was therefore defined as follows. Standard animats were trained (independently) in different environments in which the spatial distribution of thermal reward was systematically varied from noisy to smooth. The animats were then tested (with learning switched off) in the smooth Gaussian environment used in experiments 1 and 2. We refer to environments in which reward was noisily distributed (in which mappings were learnt) as the learning environments, and we refer to the Gaussian environment in which the animats' performances were tested as the testing environment.

To construct learning environments, we defined a 'heating tile' which is virtually laid out over

the environment floor. The tile is a 400 pixel x 400 pixel rectangle and each pixel is a square of width 0.02 animat body lengths. Temperature on the pixel,  $T_{i,j}$ , is determined by the following equation.

$$T_{i,j} = \frac{f^2 \sum_{s=0}^M \sum_{t=0}^N \beta(\sqrt{(i-s)^2 + (j-t)^2}, 1/f) T_{s,t}^0}{2\pi},$$

where  $N$  and  $M$  are 400,  $T_{s,t}^0$  is a uniformly distributed random number over  $[0, 1]$ , and  $f$  is a spatial-frequency of the environment, as defined by application of a filtering function  $\beta$ :

$$\beta(l, u) = \begin{cases} 1 & l < u \\ 0 & \text{else} \end{cases} \quad (2.15)$$

where  $u$  is a filter width. Essentially, changing the width of the filter that we apply to the noise allows us to manipulate the degree of spatial continuity of heat reward in the environment.

In these learning arenas the animats were allowed to learn in the same way as for the previous experiments. After learning the weights were then stored, learning was turned off, and the animat was placed in the smooth testing environment. For each animat eight simulations were conducted in the testing phase. The animat began each of these simulations at exactly the same position, with the heat centre at a fixed distance from the animat centre. However the heat source was located in one of eight directions relative to the animat starting position. Averaging performance across these eight testing simulations is important to control against potential artefacts related to animats having arbitrarily learnt to orient in a specific direction irrespective of the goal location. Our performance measure was the average distance of the centre of the animat from the centre of the heat source across all timesteps (where all test trials were of equal

duration). A baseline performance level was established by a separate set of simulations where the Gaussian test environment was used also as the learning environment, followed by an identical testing procedure. All conditions were repeated using 10 different randomly seeded animats (where the seed is applied both to the randomisation of the initial network weights and to the noisy pattern of reward in the learning environment).

Figure 2.14 shows the main results of experiment 3. The baseline average distance that animats maintained from the heat source was  $1.3 \pm 0.1$  animat body lengths, averaged across simulation timesteps, across eight directions of the target in the testing environment, and across ten randomly seeded animats/learning environments. When trained in the lowest spatial frequency learning environment ( $f = .00333$ ; i.e., white noise filtered at a width of 300 floor units), the animat was on average  $2.8 \pm 0.5$  body lengths away from the heat source. As the spatial frequency was increased the average distance of the animat increased (as did the variance) to  $5.8 \pm 1.0$  body lengths for  $f = 0.2$  (a filter width of 5 floor elements). The average distance then fell sharply to  $1.5 \pm 0.1$  for  $f = 0.3$ , which is indistinguishable from baseline performance, before rising steadily to  $2.3 \pm 0.2$  as the noise filter width was reduced to just 1 floor unit.

It is interesting that smoother training environments resulted in the worst performance, whereas intermediate levels of noise resulted in performance approaching baseline. To try to explain this we constructed plots of the paths taken in the testing environment, by an example animat from each of the training environments. Trajectory plots in the lower panels of Figure 2.14 are colour coded by the orientation of the eight target locations, and all eight testing trajectories for a given animat are shown overlaid. The black dot at the centre of each plot shows the animat starting position (always the same, and with the same initial orientation), and the coloured dots show the location of the heat centre.

The first point to note is that in most cases the trajectories taken by a particular animat vary so as to form a continuous colour wheel in the environment when the plots are overlaid. When a clear colour wheel is apparent it means that the animat was indeed moving in a way that was systematically related to the location of the heat centre, rather than moving randomly irrespective of the heat source. In the lowest spatial frequency training case (bottom left panel), the animat clearly moves towards the target and stays close to it. For the next two spatial frequencies the trajectories appear more random. In the worst case ( $f = 0.2$ ) the animat makes large loops around the arena in a way that seems unrelated to the heat source location. An example from the best case ( $f = 0.3$ ) shows a clear thermotactic behaviour. In the noisiest learning environment ( $f = 1$ ) it is interesting that the animat was clearly still able to learn an appropriate sensor-motor mapping from its environment. The colour wheel produced by the overlay of its test trajectories shows a clear relationship to the target location, despite the turning rate appearing to be comparatively slow, and hence the average distance from the target location being large.

Hence it seems that sensor-motor maps sufficient to direct the animat to a target heat source in an environment that has never seen before can be learnt even in environments that contain very noisy distributions of reward. More continuous learning environments do not necessarily lead to better learnt action outcomes.

It is interesting to examine the form of the cortical maps that were generated during the learning process. Maps constructed from the (post-training) motor output weights of the network are shown in the second row from the bottom in Figure 2.14, and are comparable across training conditions. However, maps constructed from the sensory inputs showing the preferred body part location of each neuron vary in a more interesting way. These plots are constructed as before with hue showing the location on the body of the strongest weight around the radial

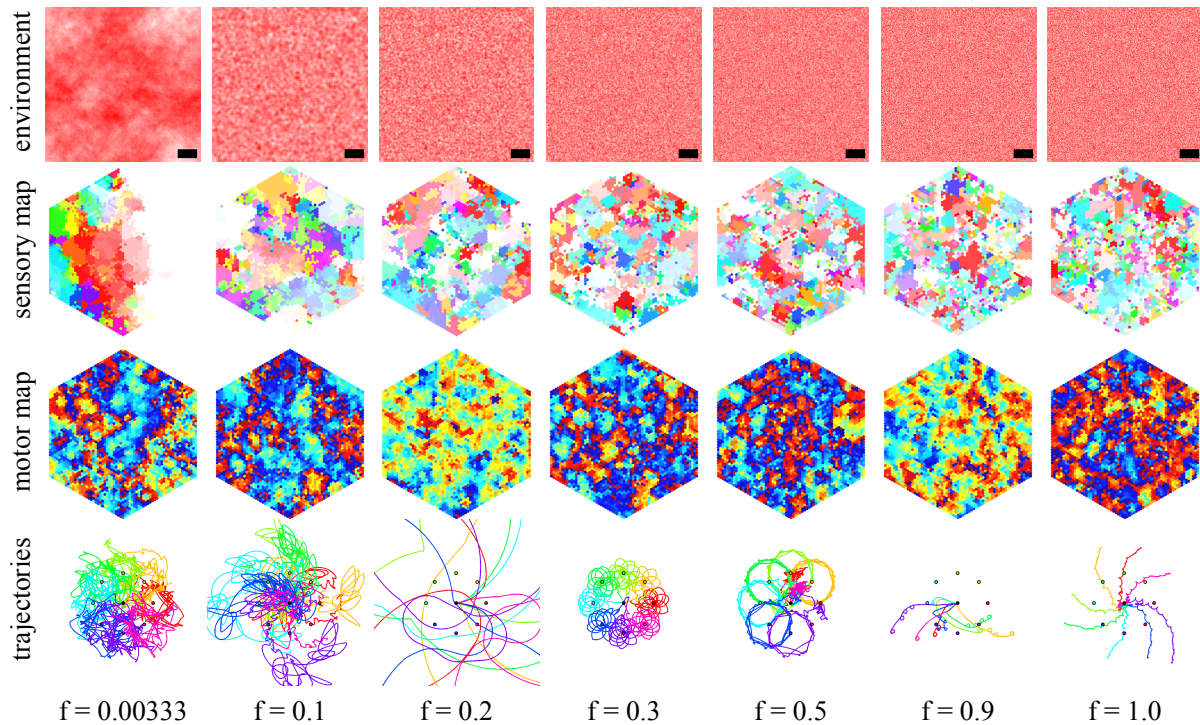
axis of the animat body, and saturation showing the body location along the longer axis. By visual inspection, the continuity of the self-organised maps decreased as the spatial frequency of the noise was increased in the training environment. Sensory maps that emerge in low spatial frequency environments are largely continuous (with respect to the body topology), resulting in a smooth progression of colour and saturation across the map plots, particularly for the lowest spatial frequency learning environment.

Hence, whilst our more continuous learning environments allowed the topology of the body to be better reflected in the organisation of sensory weights in the animats' cortical networks, this does not necessarily mean that those sensory maps were better able to solve the problem of thermotaxis. Continuous learning environments resulted in smooth body maps but poor behaviour, compared to more discontinuous environments which led to disordered body maps but good behaviour.


## **2.4 Discussion**

In chapter 2 we have explored how an artificial body (the animat) can be used to generate tactile inputs to a self-organising model of the development of somatosensory and motor maps in the cortex. A key factor of this model is that the development of the cortical maps occurs in the context of the animat achieving a goal, i.e., increasing the amount of heat that it obtains from its environment. With the increase of heat serving as a reward function during the developmental process, we were able to disentangle the factors that contribute to the emergence of sensory maps that are smooth with respect to the body topology. A key finding was that maps of the





**Figure 2.14. Manipulating spatial frequency of noisy environment.** We prepared seven kinds of noisy environments with different spatial frequency ranged from 0.00333 to 1.0 for investigating how environmental feature affects the development of the cortical map. As for environment figures, the warmer the colour is red, the cooler it is white. The black rectangle for each figure designates the size of animat. Sensory map is relatively smooth when the frequency is 0.00333, and then it is getting disrupted as the frequency rises. The motor maps seem to not smooth in either case. In the bottom raw, trajectories are coloured in 8 colours depending on the position of the heat source.




**Figure 2.15. Result of manipulating spatial frequency of noisy environment.** This shows average distance between animat and the heat source when the spatial frequency of environment changes. The baseline distance coloured in red means that the animat learnt the environment in the context of Gaussian distributed heat reward.

body (i.e., the spatial organisation of sensory receptive fields with respect to the layout of the body), emerge to be smooth only when the distribution of reward in the environment is also smooth. Crucially, the performance of the animat in obtaining reward, when tested in a smooth environment that it had not encountered during map learning, was not critically dependent on the smoothness of the emergent topological relationship between the sensory weights and the body topology.

Although there were some effects of the smoothness of the distribution of reward in the environment in experiment 3, over a wide range of body-map smoothness, the resulting behaviour of the animat was good. The maps that the animat learnt in noisy environments were sufficient to enable it to move towards heat sources when tested. Indeed, when reward was distributed at low spatial frequency across the learning environment, body maps emerged to be largely continuous, but we found that these maps were not always as good at directing the animat to reward as disorganised body maps that emerged in higher spatial frequency (i.e., noisy) learning environments. Body map continuity was therefore not the main predictor of behavioural success. Figure 2.16 which shows maximally active units taken from across the entire learning regime suggests that the decision of the control of the animat is determined by interaction between all neurons on the network and just not individual isolated neurons.

This is an important first step towards capturing the distinctions that theorists like Graziano (2009) have made between body maps and action maps.



**Figure 2.16. Maximally active units during training.** This shows the frequency of maximally active units during the development of the map. Blue coloured units are rarely activated while red are frequently activated.

# Chapter 3

## Extending the animat modelling framework

### Chapter summary

The aim of this chapter is to demonstrate the possibility of extending the animat modelling framework to operate in more complicated sensory-motor environments. The specific test-case we choose is an extension of the brain-body-world framework to include interactions between multiple actors. Through this example we explain how other aspects of the brain-body-world interaction can in principle be explored in future animat experiments. This represents an important next step towards understanding the emergence of a complete action-space map.

### 3.1 Introduction

In chapter 2 we saw that the animat model that we constructed was able to develop body maps when trained by unsupervised learning in environments in which reward was distributed

smoothly. However, we suggest that in the continuous environment smooth body maps emerge only as an incidental by-product of the process of learning an action map when environmental stimuli vary continuously. The three experiments conducted in chapter 2 help us to disentangle body maps from action maps, but they do not capture the full idea of an action space map, because the task for the animat was one-dimensional, i.e. increase thermal reward. Learning to represent a full action space would mean developing a mapping that could allow optimization of reward with respect to multiple action-outcomes.

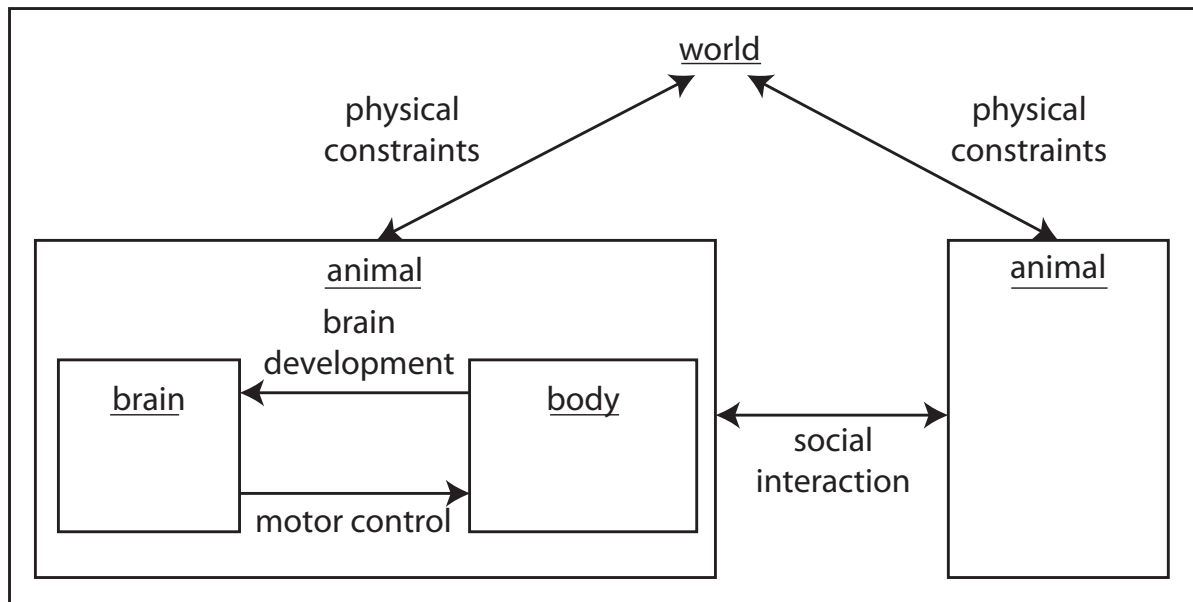
Comparing the ability of the animat to map a multi-dimensional action space, like that in which both reaching and grasping reside for primates, would require a trade-off in obtaining the reward associated with competing action-outcome pairs. To clarify, consider the argument of Graziano (2009) that homunculi reflect a statistical structure (the body topology) inherent in the relationship between actions like reaching and grasping that result in similar outcomes (e.g., manipulation of an object) because such actions incidentally involve nearby body parts like the arm and hand. We have shown how only one dimension of an action space, mapping the distribution of bodily heat onto rewarding turning behaviour, can be learnt.

It is beyond the scope of our modelling effort to map a multi-dimensional action space using the animat model. However it is within the scope of this project to show how the approach could in principle be extended towards this end.

## **3.2 Modelling assumptions**

One of the main ideas explored in chapter 2 was that a complete theory of the emergence of cortical maps, leading to an understanding of the importance (if any) of spatially continuous neural representation, must respect the constraints on the brain (the neural control architecture and information processing constraints), the body (its morphology, and sensor and actuator properties), and the world (including the distribution of rewarding and aversive stimuli, and physics).

In the words of David Marr, understanding the brain at the level of individual neurons is like ‘...trying to understand bird flight by studying only feathers: It just cannot be done’ (Marr, 1982, page 115). Our modelling approach is inspired by this approach; we are attempting to build a framework in which the most important factors affecting cortical map organisation can be studied together rather than in isolation. With this framework in place, we can later systematically replace individual components of it with improved algorithms that capture more relevant details. Figure 3.1 shows how our assumptions about the key ingredients of such a framework fit together in (a multi-animat extension) of our animat model. In the following sections we explore in more detail how our assumptions about the brain-body-world interaction is captured by the software architecture of the animat model, and how these assumptions can be better captured with improved sub-components.



**Figure 3.1. Modelling assumption.** We modeled the natural world as the illustrated system. This system has a world where multiple animals are living. We assume that there exist physical laws in the world. These laws have affect on animal's behaviour like gravity pulling down the body. Animals are living in this world as they socialise with each other. Animal's brain and body influence each other; The brain develops based on the information from sensory organs on the body, and the body behaves by commands from the brain.



### **3.3 Modelling framework**

We consider the animat model in terms of a broader modelling framework for modelling the effect of brain, body, world interactions on the developing brain. Our simulations are realised using C++ to take advantage of the flexibility object-oriented programming, and allowing us to separate the rather large repository of code that we have developed into understandable units or classes.

The world class/object describes structures like the environment terrain, gravitational and frictional forces that are common to all agents, and for the case of multiple animats it is where we resolve collisions between them.

This data is transferred between animat objects through inter-object communication, e.g., with function calls. Animats also communicate with each other through inter-object communication, which could in principle compare their body shapes to avoid intersection, or be used to exchange information relating to social cues like gesturing. Each animat comprises both a brain and a body, and the neural networks of the brain are processed using a separate class/object; this allows the brain and body to exchange information about the state of the body sensors and the requested motor outputs efficiently.

For the simulations reported so far, the code was optimized to run serially on a single processing core. However, these objects can in principle be run on different processors and executed in parallel to make multi-agent simulations run more efficiently. For example, each object could be run on separate processes or threads; in that case, communication between them would be replaced with inter-thread communication or inter-process communication, respectively.

Parallel computing can use other CPU and/or computer resources, speeding up compute time. However, in the case of parallel computing, synchronisation between processes/threads would be required from the point of view of data protection.

As a demonstration of the multiple agent system providing the interaction between brain, body, and world, we have begun work on an example using the animat model introduced in the previous chapter as shown in figure 3.2. 16 cylindrical-shaped animats are placed on the planar environment in a grid-like fashion and are simulated in the usual way. The system of this demonstration is similar to the animat model in the previous chapter but has the implementation of the collision detection between animats. This takes advantage of the contiguous layout of spherical taxels on the animat body because sphere-sphere interactions are very easy to resolve in three dimensions.


However this approach lead to an unanticipated problem; when shape-matching allows the point clouds to be highly deformed (as occurs during multi-agent collisions), the taxels can spread leaving gaps between adjacent taxels on one body that taxels of another body can squeeze between, leading multiple agents to sometimes become irresolvably entangled. Resolving this issue is of course possible by re-introducing a mesh to the point cloud to resolve collisions, and this is interesting avenue for future work.

A future version of this model could enable us to explore how more complicated forms of physical communication between agents (e.g., heat exchange) might lead to the emergence of social interactions at the multi-agent level, and how these group-level interactions might in turn affect the organisation of cortical maps.

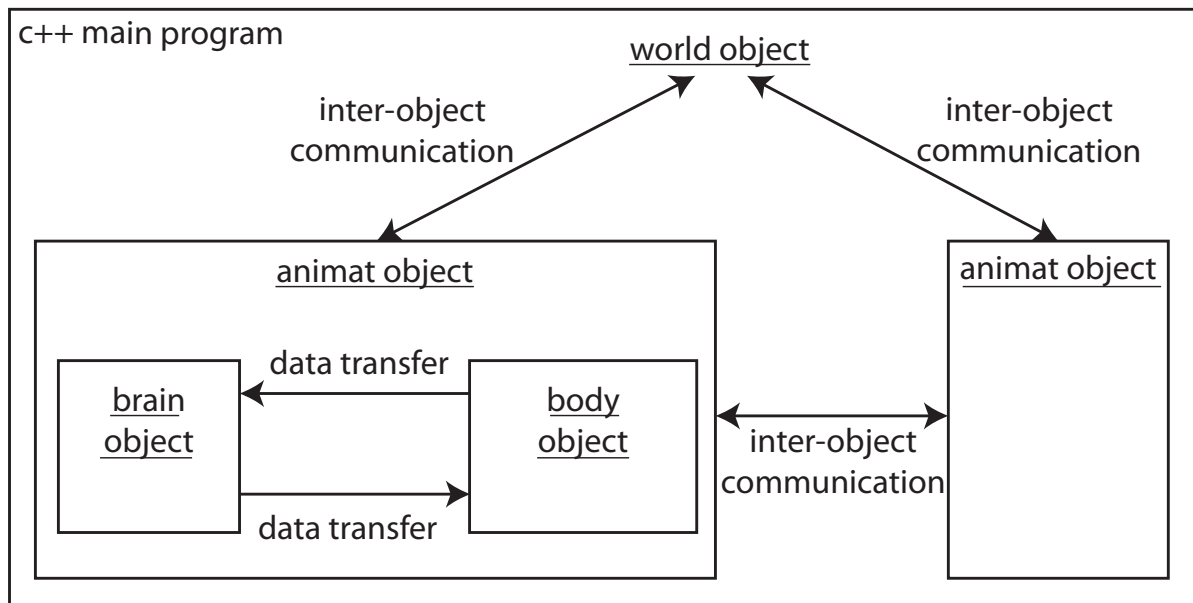
- To what extent does social behaviour have an effect on the development of the functional organisation of the primary cortex?
- How much of a difference is there between the development of the functional organisation of individual brains?
- What kind of influence does the environment where multiple agents are living have on their behaviour?

### **3.4 Modelling implementation**

At the conceptual level we are building a framework for exploring brain, body, world interactions. We have already populated this framework with candidate algorithms (shape-matching, self-organising map, collision detection etc.) which allowed us to simulate the entire system to some level of detail. It is worth considering how these candidate algorithms might be iteratively substituted for sub-components that better capture the physical and biological constraints that operate on the developing brain. This is a rather different approach to that normally taken in computational neuroscience, which typically starts by tackling much more specific problems of information processing, and perhaps later builds up towards tackling the more complex tasks of motor control faced by real biological brains; instead our approach encourages a kind of bootstrapping where the full interaction is up and running in the model from the start, and we can now develop sub-components of it in isolation (see Figure 3.5). We discuss here how aspects of the physical modelling might be improved, i.e., considering alternatives to shape-matching.



**Figure 3.2. Multiple agents in the simulation environment.** *Left:* 16 animats put on the environment in gridlike fashion as an initial state. *Right:* Animats after several time units from the initial state. As each animat having its own brain, it can evolve in the environment individually.



**Figure 3.3. Modelling framework.** We modeled a framework in order to reproduce the natural world assumed in figure 3.1 by code. Signs and texts in this figure correspond to ones in the previous figure. This framework is implemented in one of the object-oriented languages, c++. Taking advantage of the object-oriented language, we made a class for each component. Accordingly, there are world object, animat object, brain object, and body object; they correspond to world, animal, brain and body in the previous figure, respectively. Physical constraints and social interaction are expressed as data transfer between objects. Brain development and motor control are also expressed in the same way.

Shape-matching is limited because it cannot simulate physical quantities directly; rather it is a stable and efficient alternative to the use of a full physics engine. In future models, if accurate simulation of physical quantities because necessary such accuracy is required, a full physics engine could be used instead of the shape-matching algorithm for computing physical constraints. The Open Dynamics Engine might be a good candidate (Smith, 2001).

To capture aspects of social interaction between animats, we implemented a basic collision detection algorithm, which detects the intersections between bodies of multiple animats and if they collide with each other, moves them back into their positions from the previous time-step. In our implementation, only dynamic friction is programmed; static friction cannot be simulated. However, using other software like ODE by Smith (2001) to realise collision detection, we could improve our simulation quality.


The body of animats are implemented as cylindrically arranged spherical taxels. Starting from this simple cylindrical shape, it would be easy to understand the physical properties of the simulation. Yet this shape is far from that of real animals. A 3D modelling tool could be used to sculpt an animal's figure and use it instead of the shape primitive that are easiest for us to analyse. Laser scanning could be used to obtain realistic body morphologies, and the meshes produced by such techniques can be easily incorporated into the shape-matching approach.

To control the movement of the animat we can use alternative central pattern generator models, such as those proposed by Ijspeert (2008), which involve controlling self-organisation of movement dynamics of the spine, using neural network models of the basal ganglia in amphibians, instead of the simple sine-wave generators we use here to define a simple motor space.

The shape-matching-algorithm yields realistic pseudo-physical animation at low computational cost. However, the clustering of taxels is rather complicated; clusters have to be rather carefully constructed and paired to the bones, which is easy for bodies made of simple shape primitives like the cylindrical animat, but more difficult for more complex body morphologies. As an alternative to this algorithm, skin deformation techniques such as linear blending skinning (LBS) and dual quaternion skinning (DQS) offer good candidates (Kavan et al., 2008). These techniques are often used in the gaming industry for the purpose of animation. Figure 3.4 shows our design of an animat model using the process of vertex skinning. The editing software Blender (Blender Online Community, 2013) was used to configure the skin and bones in 3D, and these configurations were exported into the custom software we have been developing, and into which we implemented LSB Figure 3.4.

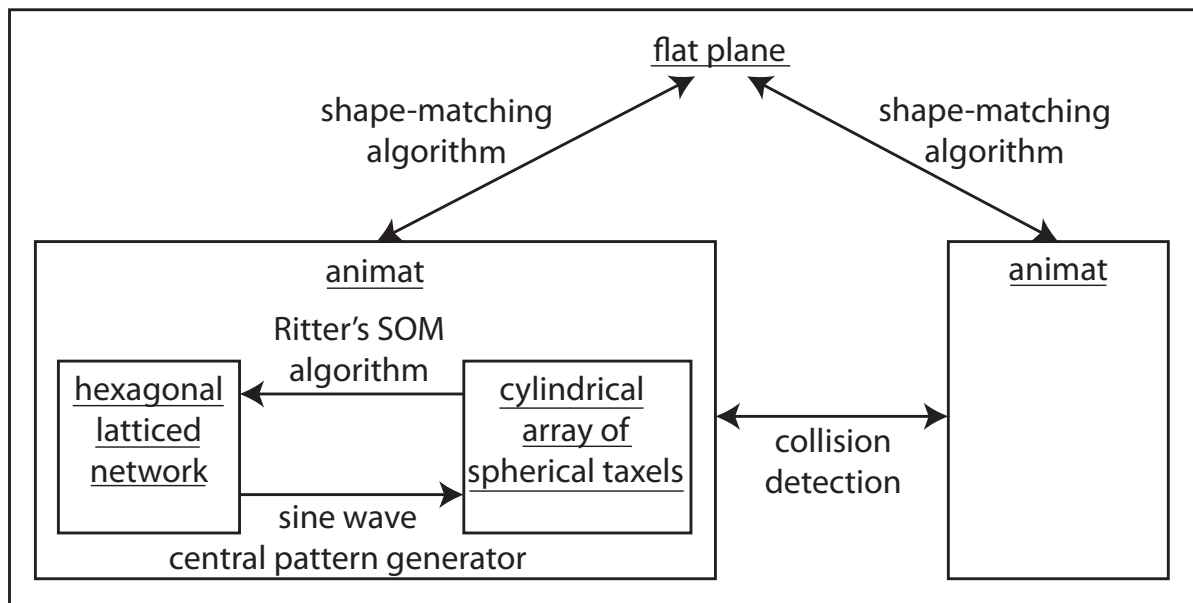
The basic approach of the skinning technique is preparing 3D surface model and skeleton, and associating vertices of the model with bones so that when a bone rotates at a joint, vertices follow the bone movement depending on how much these vertices are associated with bones. This approach allows an object to deform naturally; however, we found that unlike shape-matching this algorithm is unable to express interactive animation such as one object squashing the other object, as we were able to achieve with the shape-matching algorithm. One way to add this feature would be combining the skinning technique and quadratic deformation of the shape-matching algorithm. After the skinning method determines the basic shape of the object based on the bone position, the shape-matching algorithm deforms the object using the basic shape as an initial point cloud. In this case, the shape-matching algorithm has only one cluster, and hence computational cost would remain low.

Keeping the computational cost low is important for facilitating the modelling process. In




**Figure 3.4. Designing process of skinning.** *Left:* At first, design detailed surface model and skeleton. In this case, a cylindrical model and three-bone skeleton are made *Centre:* Then, associate each vertex of the model with one or more bones with weight, which means how much bone movement influences the vertex movement. The colour of the vertices in the figure means how much degree vertices are associated with the lower bone. Red coloured vertex is more associated and blue coloured is less. *Right:* Lastly, move bones of the skeleton, resulting in skin movement depending on how much degree vertices are associated to each bones. The figure shows the cylindrical model are smoothly deformed at the bending positions.





**Figure 3.5. Modelling implementation.** This figure shows how each components in the previous figure are implemented in this system. In our system, brain, body and world are implemented as a hexagonal latticed network, cylindrical array of spherical taxels, and a flat plane in 3D space, respectively. Physical constraints between world and animal and social interaction between animals are implemented as shape-matching algorithm and collision detection, respectively. As for the implementation between brain and body, the link from the body to the brain is implemented as Ritter’s self-organising map algorithm (Ritter et al., 1992), one in the opposite direction is done as sine wave.

our experience, considerable insights into the behaviour of a complex model like the one we developed in chapter 3 can be gained by the ability to interact with the model, in something that approximates real-time. Figure shows a demonstration that we were able to construct that allows the user to interact with the model in real-time. In the top-right of the figure is an interactive window that allows the user to move the center of the heat source and see the animat respond by changing its trajectory in real-time. The demonstration runs on a standard laptop computer, and we were able to present it in real-time at a conference on Living Machines along other biomimetic robotics projects Urashima and Wilson (2014).



**Figure 3.6. Interactive simulation of the animat model.** Left window shows an animat simulation in realtime. Right top window is a control panel that shows the information of the arena (the white rectangle is the heat source, green rectangle is the current position of the animat, and black line is the recent trajectory of the animat) and allows users to change the position of the heat centre by a mouse click. Right bottom windows show the current view of the sensory and motor maps.

# Chapter 4

## General discussion

We have presented a novel approach to investigating the relative contribution of neural plasticity, body morphology, and environmental context, in shaping neural maps by input-driven self-organisation. This approach involved integrating two existing algorithms, shape-matching and self-organising maps, to create an animat model with a rudimentary brain, body, and goal. The shape-matching algorithm of Müller et al. (2005) was used to generate patterns of tactile input to the animat sensorimotor cortex, which was modelled using a self-organising map algorithm extended to include reward-based learning for motor control (Ritter et al., 1992).

In previous models from our group we have used self-organising networks to explain experimentally measured map organisations, e.g., Stafford and Wilson (2007); Wilson et al. (2010), and hence recreating experimentally determined cortical maps has often been our primary goal as modellers. However, the current model was constructed specifically to explore whether our goal as modellers mirrors any that we may ascribe to brain evolution, i.e., does spatial patterning in the distribution of receptive fields across a neuronal sheet subserve a specific function? The goal of our self-organising network was therefore defined not in terms of its ability to recreate a particular spatial organisation, but in terms of its ability to modulate central pattern generators that rotate the animat bones, in order to maximise reward obtained from the environment. In-

tegrating shape-matching and self-organising maps thus allowed the loop between sensing and action to be closed.

In designing a sensory environment wherein neural maps of the body might emerge, we were inspired by the recent study of McGregor et al. (2011), in which the topology of various bodies were recreated using tactile stimulation and an unsupervised algorithm similar in spirit to the self-organising map. In this study, feed-forward networks were driven by passive stimulation of rigid bodies, either by propelling them across an uneven floor (under gravity), or by bombarding them with small moving objects. Either strategy gave rise, in the patterns of correlation between nearby taxels, to information about various body topologies sufficient to yield faithful body reconstructions. Our animat model may be considered an extension to this approach, with tactile stimulation evoked by actively controlled movements of the body, and it confirms the intuition of McGregor et al. (2011) that self-organising maps may likewise be useful for unsupervised learning of body schemata.

Like McGregor et al. (2011), our model could have been embodied in a full physics engine, and several good platforms for doing so exist, e.g., Cofer et al. (2010). Instead we implemented shape-matching for its efficiency in generating realistic-looking movements and deformations of a three-dimensional body, governed by a small parameter set. Animators may be interested in our method of coupling one shape-matching body (the animat bones) with another in which it is enclosed (the animat skin) to enable self-generated motion relative to the environment; custom-written c++/python/openGL libraries are available on request. However we consider these implementational details as a means to an end. Simulated physics here serves simply to embody our neural network in an external environment, and to embed it in a control task resembling that faced by biological brains. It has been argued that brains were evolved primarily

to optimise the relative rotations of body parts in order to obtain reward (and avoid harm), e.g., Tweed (2003); Wolpert and Ghahramani (2000), and in these terms learning to rotate an animat spine to gain heat represents an appropriate simplification and abstraction of the challenges faced by biological neural networks.

The task of guiding a deformable body up a temperature gradient is not a particularly difficult motor control problem. Indeed we were able to solve it using a simple Braitenberg-style sensor-motor mapping (Braitenberg, 1984). But crucially the task allows us to disentangle two causes of correlation in the typical activities of adjacent sensors: Spatial continuity of stimuli in the world external to the body, and spatial contiguity in the sampling of those stimuli at the body surface. In other words, things that are next to each other tend to elicit similar responses in nearby sensors. Self-organising maps exploit this correlational structure. The distinction between continuous stimuli and contiguous sampling may seem trivial to those who have trained self-organising nets as models of primary visual cortex development, for example, by presenting spatially localized patterns of photoreceptor intensity to networks before deriving retinotopic mappings from them. In such an exercise, replacing the spatially localized blobs of activation with random patterns of retinal activation will of course fail to generate retinotopic maps, irrespective of whether we imagine the noise to have originated from an untuned television set or from a scrambled afferent projection from the retina and/or thalamus. The importance of distinguishing between stimulus continuity and sampling contiguity is highlighted by our final experiment (Figure 2.14), showing how disrupted stimulus continuity (i.e., a noisy distribution of heat/reward on the ground) yields discontinuous body maps in the network afferent weights. The results of this experiment reveal that poorly organised body maps emerging in discontinuous environments can be functionally equivalent to somatotopic homunculi emerging in continuous environments. Either is sufficient to control the body, hence somatotopic

organisation in the afferent weights seems to be functionally irrelevant.

Previous work from our group has shown that the principles of self-organisation can explain discontinuities in the mapping of the body surface to the brain. We tested the hypothesis of Farah (1998) that a tendency for the hands and face to touch in the womb, despite being distal on the body, can lead to discontinuous homunculi in which the two are co-local, as in the primate primary somatosensory cortex (Stafford and Wilson, 2007). However, this study also showed that self-organisation alone cannot explain the consistent medial-lateral ordering of the body representation between the maps of conspecifics, without evoking some additional (presumably genetic) constraints. In a later critique of self-organisation as a theory of somatosensory map organisation, Parpia (2011) highlighted the divergence of projections from subcortical hands and face areas. Whether maps self-organise by activity-dependent Hebbian processes, or are governed by cell migration along genetically pre-specified chemical gradients, or reflect some combination of the two, our argument is unperturbed. Homuncular organisation in the cortex, by whatever mechanism it arises, has not yet been demonstrated to be functionally relevant for neural computation. We *believe* that homunculi are useful, perhaps allowing distance-dependent intracortical signal delays to subserve computation of spatial-temporal derivatives of sensory inputs localised to continuous cortical maps (Wilson et al., 2011; Wilson, 2011). However, we are not yet in a position to reject the alternative hypothesis, as expressed succinctly by Deacon (2012): “Although a map is similar in structure to what it represents, it is not intrinsically meaningful”.

In the context of our modelling results it is interesting to reconsider what it means to measure a continuous homunculus map in the somatosensory or motor cortices of an animal. In our model, homunculi emerge as just one component of a network organised to convert patterns of sensory

input to appropriate motor commands. While it is possible from outside the system for us to observe body-related organisation, this need not have any direct relevance to neural computation. For the animat, the *purpose* of its afferent weights is not to maintain a body schema but simply to help increase thermal reward. Graziano (2009) and colleagues contend that what we might be tempted to interpret as maps for the somatotopic layout of mechanoreceptors and muscle fibers, in the organisation of responses in primary somatosensory and motor homunculus maps, in fact reflect the topological/correlational structure of another, albeit related, space in which motor actions involving similar muscle groups are correlated. In this ‘action space’ neuronal activities are thought to be functionally organised, via self-organising mechanisms (Aflalo, 2006; Graziano and Aflalo, 2007), by their capacity to elicit similar action outcomes, such as reaching, grasping, and defensive movements of the arms. Although the action spaces relevant for our animats were limited to single actions, i.e., “turn towards heat” or “get warmer”, and therefore no real action space map can be derived, the implications of our modelling results echo the warning of Graziano (2009), that somatotopic organisation may place only partial constraints on the underlying correlational structure exploited by the local application of competitive Hebbian learning.

In similar terms, a recent review paper argued that neuroscience is currently lacking a good (evolutionary or computational) theory of what continuity in cortical maps might be for (Wilson and Bednar, 2014). Briefly, the argument is that map continuity may be an inevitable consequence of efficient developmental algorithms that were selected primarily for their ability to distribute stimulus representations to cover the entire space from which those stimuli are drawn. Hence, whilst map continuity may have since been selected for, the original driver of map-generating self-organising processes may have been map *coverage* rather than map *continuity*. This view shifts the focus from the map to the developmental algorithm by which the



map emerges. Clearly input-driven self-organisation is a general-purpose mechanism for creating adaptive neural networks, as reflected by our success in applying the algorithm of Ritter et al. (1992), with little modification, to a quite different control problem (Ritter et al., 1992, originally solved a robot pole-balancing problem). It will be interesting in future to explore how additional properties of biologically grounded self-organising models like LISSOM (Mikkulainen et al., 2005; Wilson et al., 2010) and GCAL (Stevens et al., 2013), e.g., Hebbian modification of additional recurrent intracortical weights, might advantage networks extended in a similar way for reward-based motor control.

The model that we present here represents our first steps in establishing a computational framework through which to tease apart the evolutionary and developmental origins of neural maps. A key aspect of this effort is to place neural map self-organisation in a developmental context in which the resulting map organisation may be useful, and hence may later be selected for. The developmental context, here thermoregulation, was chosen because it opens the door to testing our thesis scientifically. Self-organised thermoregulatory control strategies could be evaluated in the context of a wealth of existing experimental data on the role of thermoregulation in governing the collective behaviour of juvenile rodents, as they huddle together to keep warm throughout the first two postnatal weeks. In a recent related paper we show how aggregation patterns formed in groups of two-dimensional pups orienting towards heat sources (which include littermates) match those measured in behavioural experiments with juvenile mice (Glancy et al., 2014; Canals and Bozinovic, 2011; Alberts, 2007). Our next steps will therefore be to attempt to recreate the emergent group-level behaviours predicted by our two-dimensional huddling model, in litters of three-dimensional animats, with maps organised by each optimising a local reward function that we have shown can lead to self-organised group-level huddling behaviour. In such a model we hope to explore the potential for self-organisation to attribute

correlational structure in sensory input patterns to the motion of others versus the motion of the self.

## **4.1 Conclusion: Towards a model of the emergence of action space maps in the motor cortex**

We proposed a computational system for modelling the interaction between brain, body, and world in the context of reward in order to produce action space maps on a cortical sheet. By using this system, we have obtained sensory and motor homunculus maps on the cortical sheet. These maps have a continuous topological relationship to the body, as we expected, and they allow the animat to move towards rewarding places efficiently, as we hoped. This shows that a cortical map can be useful. Furthermore, experimentally manipulating components of the brain, body, and world, suggested that the homunculus pattern on the cortical maps has no functional significance, because animats were controllable even when their body maps were not continuous.

In addition, we explored how the model could be extended in future research, within a general animat modelling framework, by exploring how specific algorithms comprising the model could be replaced with more sophisticated ones.

Although we have extended Kohonen's self-organising map model (in line with Ritter's earlier work) to start to investigate the learning of actions, there is some exciting and interesting further modelling work to be done towards mapping a full multi-dimensional action space. Our

extensible modelling framework provides the first steps towards this end.

# References

- Aflalo, T. N. (2006). Possible Origins of the Complex Topographic Organization of Motor Cortex: Reduction of a Multidimensional Space onto a Two-Dimensional Array. *Journal of Neuroscience*, 26(23):6288–6297.
- Alberts, J. R. (2007). Huddling by rat pups: ontogeny of individual and group behavior. *Developmental psychobiology*, 49(1):22–32.
- Blender Online Community (2013). Blender - a 3D modelling and rendering package. URL: <http://www.blender.org/>.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. MIT Press.
- Canals, M. and Bozinovic, F. (2011). Huddling behavior as critical phase transition triggered by low temperatures. *Complexity*, 17(1):35–43.
- Cofer, D., Cymbalyuk, G., Reid, J., Zhu, Y., Heitler, W. J., and Edwards, D. H. (2010). Animat-Lab: a 3D graphics environment for neuromechanical simulations. *Journal of neuroscience methods*, 187(2):280–8.
- Deacon, T. W. (2012). *Incomplete Nature: How Mind Emerged from Matter*. W. W. Norton.
- Farah, M. J. (1998). Why does the somatosensory homunculus have hands next to face and feet next to Genitals? A hypothesis. *Neural Computation*, 10(8):1983–1985.
- Glancy, J., Groß, R., Stone, J. V., and Wilson, S. P. (2014). A self-organizing model of thermoregulatory huddling (submitted). *PLoS Comput Biol*.
- Graziano, M. S. A. (2009). *The Intelligent Movement Machine : An Ethological Perspective on the Primate Motor System: An Ethological Perspective on the Primate Motor System (Google eBook)*, volume 2008. Oxford University Press.
- Graziano, M. S. A. and Aflalo, T. N. (2007). Mapping behavioral repertoire onto the cortex. *Neuron*, 56(2):239–51.

- Hebb, D. (1949). *The Organization of Behavior*. New York: John Wiley.
- Hoffmann, M., Marques, H., Arieta, A., Sumioka, H., Lungarella, M., and Pfeifer, R. (2010). Body Schema in Robotics: A Review. *IEEE Transactions on Autonomous Mental Development*, 2(4):304–324.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160:106–54.
- Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: a review. *Neural networks : the official journal of the International Neural Network Society*, 21(4):642–53.
- Kavan, L., Collins, S., Žára, J., and O'Sullivan, C. (2008). Geometric skinning with approximate dual quaternion blending. *ACM Transactions on Graphics*, 27(4):1–23.
- Knudsen, E. I., du Lac, S., and Esterly, S. D. (1987). Computational maps in the brain. *Annual review of neuroscience*, 10:41–65.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69.
- Kohonen, T. (1984). *Self-Organization and Associative Memory*. Springer-Verlag, Berlin.
- Kohonen, T. and Honkela, T. (2007). Kohonen network. *Scholarpedia*, 2(1):1568.
- Koulakov, A. A. and Chklovskii, D. B. (2001). Orientation preference patterns in mammalian visual cortex: a wire length minimization approach. *Neuron*, 29(2):519–27.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.
- McGregor, S., Polani, D., and Dautenhahn, K. (2011). Generation of tactile maps for artificial skin. *PloS one*, 6(11):e26561.
- Merzenich, M. and Kaas, J. (1982). Reorganization of mammalian somatosensory cortex following peripheral nerve injury. *Trends in Neurosciences*, 5:434–436.
- Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (2005). *Computational Maps in the Visual Cortex*. Springer Science & Business Media.
- Müller, M., Heidelberger, B., Teschner, M., and Gross, M. (2005). Meshless deformations based on shape matching. *ACM Transactions on Graphics*, 24(3):471–478.
- Parpia, P. (2011). Reappraisal of the somatosensory homunculus and its discontinuities. *Neural computation*, 23(12):3001–15.

- Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60(4):389–443.
- Penfield, W. and Rasmussen, T. (1950). *The cerebral cortex of man: a clinical study of localization of function*. Macmillan.
- Purves, D., Riddle, D. R., and LaMantia, A. S. (1992). Iterated patterns of brain circuitry (or how the cortex gets its spots). *Trends in neurosciences*, 15(10):362–8.
- Ritter, H., Martinetz, T., and Schulten, K. (1992). *Neural Computation and Self-Organizing Maps; An Introduction*. Addison-Wesley Longman Publishing Co., Inc.
- Sharma, J., Angelucci, A., and Sur, M. (2000). Induction of visual orientation modules in auditory cortex. *Nature*, 404(6780):841–7.
- Smith, R. (2001). ODE: Open Dynamics Engine.
- Stafford, T. and Wilson, S. P. (2007). Self-organisation can generate the discontinuities in the somatosensory map. *Neurocomputing*, 70(10-12):1932–1937.
- Stevens, J.-L. R., Law, J. S., Antolík, J., and Bednar, J. A. (2013). Mechanisms for stable, robust, and adaptive development of orientation maps in the primary visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(40):15747–66.
- Sur, M., Angelucci, A., and Sharma, J. (1999). Rewiring cortex: the role of patterned activity in development and plasticity of neocortical circuits. *Journal of neurobiology*, 41(1):33–43.
- Swindale, N. V. (1998). Cortical organization: Modules, Polymaps and mosaics. *Current Biology*, 8(8):270–273.
- Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., and Hübener, M. (2000). Visual cortex maps are optimized for uniform coverage. *Nature neuroscience*, 3(8):822–6.
- Tweed, D. (2003). *Microcosms of the Brain: What Sensorimotor Systems Reveal about the Mind*. Oxford University Press.
- Urashima, H. (2013). *Predicting the organisation of motion maps in the vestibular cortex using a self-organising map*. Master’s thesis, The University of Sheffield.
- Urashima, H. and Wilson, S. P. (2014). A Self-organising Animat Body Map. In *Biomimetic and Biohybrid Systems*, volume 8608 of *Lecture Notes in Computer Science*, pages 439–441. Springer International Publishing.
- Wilson, S. P. (2011). *Figuring time by space: Representing sensory motion by spatial coding*. Phd thesis, The University of Sheffield.

- Wilson, S. P. and Bednar, J. A. (2014). Are visual cortex maps optimized for continuity? (in press). *Developmental Neurobiology*.
- Wilson, S. P., Bednar, J. A., Prescott, T. J., and Mitchinson, B. (2011). Neural computation via neural geometry: a place code for inter-whisker timing in the barrel cortex? *PLoS computational biology*, 7(10):e1002188.
- Wilson, S. P., Law, J. S., Mitchinson, B., Prescott, T. J., and Bednar, J. A. (2010). Modeling the emergence of whisker direction maps in rat barrel cortex. *PloS one*, 5(1):e8778.
- Wolpert, D. M. and Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature neuroscience*, 3 Suppl:1212–7.