



This is a repository copy of *Embodied Models and Neurorobotics*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/130868/>

Version: Accepted Version

Book Section:

Prescott, T.J. orcid.org/0000-0003-4927-5390, Ayers, J., Grasso, F.W. et al. (1 more author) (2016) *Embodied Models and Neurorobotics*. In: Arbib, M.A. and Bonaiuto, J., (eds.) *From Neuron to Cognition via Computational Neuroscience*. MIT Press , pp. 483-512. ISBN 9780262335270

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

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



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

Chapter 17. Embodied Models and Neurorobotics

Tony J. Prescott, Department of Psychology, University of Sheffield, UK
Joseph Ayers, Dept. of Biology & Marine Science Centre, Northeastern University, USA.
Frank W. Grasso, Dept. of Psychology, Brooklyn College, USA,
Paul F. M. J. Verschure, SPECS Lab, Universitat Pompeu Fabra & Catalan Institute of
Advanced Research (ICREA), Spain.

Corresponding author:

Tony J Prescott, Dept. of Psychology, University of Sheffield, Western Bank, Sheffield,
S10 2TN, UK. Email: t.j.prescott@sheffield.ac.uk, tel: +44 114 2226547

The total number of actual words in all of the manuscript: 18,185 (text 12,604, figure legends 1,500, bibliography 3,919).

The total number of figures in the manuscript: 18.

The total word equivalents in the manuscript (= total actual words plus 300 times the number of figures): 23,585.

Acknowledgement

The contributions of Tony Prescott and of Paul Verschure were partly supported by the EU Framework 7 Projects CSNII (601167), BIOTACT (215910), EFAA (270490), and WYSIWYD (612139).

Whereas most other chapters in this book study computer or mathematical models of various aspects of brain function, here we focus on the value of physically-embodied models, in which simulated brain circuits are embedded within robotic bodies. The general goal in doing this is to better understand the interaction of the brain, the body and the environment in generating the adaptive behavior we see in animals. We call these *embodied computational neuroscience* models (Prescott, Montes Gonzalez et al. 2006), and the emerging domain of research in such systems *neurorobotics* (Ayers 2002, Dario, Carrozza et al. 2005, Krichmar 2008). Although the main emphasis of this chapter will be on the value of such models to neuroscience, it is worth noting that the construction of these artifacts to address scientific aims, can also lead to the availability of a new class of technologies that can address outstanding challenges in engineering (Prescott, Lepora et al. 2014).

The structure of our chapter is as follows. First, we introduce the idea of developing embodied models of animals and their nervous systems comparing and contrasting these with other types of models studied in neuroscience. Second, we discuss the relationship between brains and bodies explaining why the notion of embodiment is critical if we are to fully understand the role of brains in determining or influencing behavior. Third, we explore, in some detail, several example robot models of invertebrate sensorimotor co-ordination, showing how useful models can be developed at different levels of description—for instance, in terms of information, or in terms of circuits, and how the best models can combine multiple levels of description. Fourth, we investigate how the layered architecture of the mammalian brain can be approximated in robot models that incorporate key aspects of animal behavior and morphology. Finally, we consider how such models could be scaled-up to address the challenge of understanding the whole mammalian brain.

1. Physical models in neurobiology

As explained in the introductory chapter, the goal of modeling is not to replicate as accurately as possible what we see in biology. Rather it is to help us to understand of the underlying principles that are exploited by biological systems to maintain life.

Today, modeling often takes the form of computer simulation, however, physical models have also played an important role in biological science. For instance, the discovery by Watson and Crick of the structure of our genetic material was based on the construction of a physical model comprised of little metal sheets. This three-dimensional structure allowed Watson and Crick to conceive of a shape—the double helix—that could account for the x-ray diffraction data of DNA gathered by Rosalind Franklin (Crick 1990). The Watson and Crick model facilitated the scientific discovery process without looking to accurately mimic biological reality; rather, it employed a suitably defined abstraction.

In 1945, eight years before Watson and Crick's breakthrough, Rosenblueth and Wiener proposed that the criteria for a good model was not how accurately it captured all the details of the biological phenomenon of interest, but how useful it was in casting new light on that phenomenon. The Watson and Crick model succinctly demonstrates this principle—to intuit the structure of the DNA molecule became much easier once they had a physical model they could manipulate.

A further strength of all good models is that they capture the major common features or properties of the system under study and summarize the individual variation. This is true not only of physical models, robots and computer-based simulations, but also of biological models. Neurobiologists, who can sometimes be dismissive of simulation as “only a model”, should if they are being logically consistent, also dismiss the model systems they study themselves. For instance, someone might use one animal as a model of another (e.g. the rat as a model of a human); or an animal in an altered condition (e.g. under anesthesia, or awake but immobilized) as a model of the same animal but awake and behaving; or an isolated part, such as a brain slice or a muscle, as a model of that component within an intact, functioning system. Hence the neurobiologist, like the modeler, takes significant steps away from the actual system of interest—species difference, anesthesia/immobility, dissection—for the practical benefits of having a system that is easier to study.

For similar reasons of accessibility and ease of study, Rosenbleuth and Wiener argued that a good model should translate a question about the natural world into a domain that we understand better, and it should allow us to conduct experiments with relative ease. In this chapter we contend that—for exploring a range of relationships between brain, body, and behavior—robotics, which translates problems in biology into problems in computing and engineering, meets the requirements of being easier to study and to understand particularly well.

Let’s begin from the standpoint of the neurobiologist looking at potential model systems for understanding the role of the brain in generating natural (ethologically-relevant) behavior. Table 1 compares and contrasts neuroscientific models and robotic models in order to show some of their advantages and disadvantages. Our table illustrates that robotic models can nicely capture interesting properties of awake behaving animals that are lost as we move to more accessible neurobiological models. In particular, robot models are more amenable to investigation of their internal processes, such as the current state of any calculations being performed, than any existing animal model. The main issue for robot models, though, is how to capture enough of the relevant properties of the neural and physical substrates of the target system to create a useful model.

System		Properties				Amenable to experimental investigation of internal processes
		Similar neural & physical substrates	Intact system	Embedded in the environment	Expressing ethologically-relevant behavior	
Awake, free-moving organism (target)		*****	*****	*****	*****	*
Animal models	Awake, restrained	*****	*****	*****	***	**
	Under anesthetic	*****	*****	***	*	***
	Brain slice or similar	*****	*	**	*	****
<i>Robotic model</i>		*	*****	*****	*****	*****

Table 3.1. Advantages and disadvantages of different neuroscientific models compared to robotic (physical) models on an ordinal scale (more *s = better) (Adapted from Mitchinson, Pearson et al. (2011)).

Note that for each of the different neuroscientific model systems considered in Table 1 we might also build a computer simulation to meet our requirement for accessibility and ease of experimentation. Indeed, in this age of high-power computer systems with real-time simulated physics engines, is it ever really worthwhile physically copying aspects of the animal? Our answer is “not always”, or perhaps even, “generally not”. Depending on your question (and one issue, as we will see, is whether you are asking the best question) you may be better off staying in the world of simulation as many of the excellent computational models in this book show. In this chapter we will argue, nevertheless, that there are circumstances where a physical model is going to be a better model, in the sense of being easier to build, study and understand, than a simulation.

As will become clear, tackling the problem of creating a physical model may sometimes also drive the researcher to ask a different set of questions that result in a new way of thinking about that system that is closer to biological reality. Building a physical model can also take us a step closer to creating a novel behaving system that may be of some real benefit in the world (and one of the goals of science must be to create this kind of useful artifact). In general, however, it is rarely a question of either/or. In developing robotic models we will often use simulation as a tool for exploring the design space for biomimetic artifacts before we build them, and we will look to use our completed physical model to validate and extend the insights obtained through simulation.

2. Embodiment matters

The existence of directed swimming in single-celled eukaryotes such as *Paramecia* (Jensen 1959) reminds us that nervous conduction is not a prerequisite for movement or action. Indeed, the first nervous systems evolved some hundreds of millions of years after the appearance of the first active mobile animals (Bengston 1994). What the nervous system certainly does, however, is to enable movement to be more rapid, more efficient, and better adapted to the environment.

The importance of the body in generating coordinated movement is beautifully illustrated in machines that exploit the natural dynamics of their parts for periodic motion. For example, McGeer (1990) showed that a bipedal walking machine, with no control system whatsoever, can generate a stable walking gait on a suitably sloped surface by relying on the passive dynamics of suitably configured mechanical parts. This synergy between form and function has also been called morphological computation (Pfeifer and Bongard 2006). Animals provide control through their nervous systems in a manner that complements their natural body dynamics. In walking or running, for example, many legged animals exploit the pendulum-like natural motion of jointed limbs to help generate a suitable cyclic pattern. In a similar way, a parent pushing a child on a swing controls the back-and-forth motion by periodically interjecting energy into the natural dynamics of the system. By relinquishing some aspects of control to the body animals also benefit from the energy recycling capability of elastic tissues. Muscles and tendons, for instance, convert kinetic energy to potential energy as the foot hits the ground, providing a store of energy to be released in the next step cycle. Designing controllers modeled on animal locomotion pattern generators, that exploit these principles, provides a very promising path for building efficient legged robots (see Chapter 9, Motor Pattern, Generation, Ijspeert (2014), and Section 3 below). Such controllers can be simpler—for instance, having fewer control parameters—than more

traditional forms of continuous robot control, and will entrain themselves to the dynamics of the body making them highly adaptable. A broader lesson is that rather than thinking of the nervous system as the coordinator of the body, in the way that a conductor directs an orchestra, we might think of it more as “one of a group of players engaged in jazz improvisation” (Chiel and Beer 1997, p. 555) influencing, but also being influenced by, the body and environment in which it is embedded.

Exploiting the physical properties of body parts leads to a simplification of the problem of locomotion for walking and running animals and machines. Similarly, exploiting the physical properties of the world allows animals and robots to generate apparently complex behavior through simple principles. This has been recognized many times but some examples are worth noting.

The neurologist Grey Walter built one of the first brain-inspired robots in the 1950s when he created his robot ‘tortoises’ Elmer and Elsie (Walter 1953). Each tortoise (figure 1 left) had a single light sensitive ‘eye’, a shell that was sensitive to touch, and a pilot light that was configured to extinguish when a second light source was detected. Control was provided by simple analogue electrical circuits, termed relays, designed to mimic some of the properties of biological neural networks. As shown in Figure 1 (centre), these relays implemented a layered control hierarchy. The lowest level behavior provided forward movement (*explore*), if light was detected this was suppressed and replaced by a movement towards the light (*phototaxis*), and if the shell touched against an obstacle, the robot would stop anything else it was doing and exhibit a simple escape maneuver (*wriggle*). By photographing his robot in darkness, with a lengthy shutter time, Grey Walter was able to record its trajectory when moving in environments such as that illustrated in Figure 1 (right panel). Here, after interacting with various obstacles the robot eventually approaches a light source. In other photographs the robot can be seen homing in on its ‘hutch’ in order to recharge its battery, or traversing alongside a mirror with an oscillating locomotion movement pattern (alternately approaching and then moving away from its own reflection) that Grey Walter termed a ‘dance’.




Figure 1. Grey Walter's tortoise "Elsie". The robot (left) generates apparently complex behavior as the result of a simple layered control system (centre) interacting with an environment of obstacles and light sources. The photograph on the right illustrates the robot's trajectory recorded using a long shutter speed, with a light source placed on the robot's outer shell. (Walter 1953).

This idea of configuring a control system as a set of reflexes with a fixed priority scheme was adopted with enthusiasm in the 1980s by researchers interested in generating robust behavior for fast-acting mobile robots (e.g. Brooks (1986)) that can engage in simple but reliable activities such as wall-following (see also Mataric (1997)). Similar

principles may also have underpinned the behavior of bilaterian creatures who left spiraling or meandering foraging trails that then became trace fossils in Precambrian rocks, and that are our earliest evidence of the behavior of multi-celled animals (Prescott and Ibbotson 1997). These fossil meanders belong to the class of *stigmergic* processes in which the behavior of the organism restructures its environment which in turn effects the activity of that organism (Theraulaz and Bonabeau 1995). Stigmergy can be thought of as a form of “information offloading” (Hutchins 1995, Beer and Williams 2014), whereby some aspects of the processing that are required to implement a behavior are externalized by altering the environment’s structure. In other words, we should think of the system formed by the *organism in its environment* as the substrate for the computation that results in the observed behavioral pattern, not just the organism and its nervous system.




Figure 2. Robot modeling of stigmergic processes. Left: Fossilized animal from the Precambrian era likely to have been left by an early bilaterian animal, from Crimes and Anderson (1985). Right: A custom-built robot configured to leave a meandering paper trail using a fixed control hierarchy of reactive mechanisms. For details see Prescott and Ibbotson (1997).

Grey Walter’s tortoises showed that the interaction of a reactive, but appropriately configured, control system with the environment, via suitable, but often primitive, sensors and actuators can give rise to behavior that is surprisingly complex. Thirty years later, in a series of thought experiments, the neurobiologist Valentino Braitenberg (1986) imagined a series of robot “vehicles” of increasing sophistication, beginning with machines with only sensory and motor elements, and finishing with variants that had simple artificial nervous systems capable of properties such as logic, memory and abstraction. These internal processing elements hint at some likely functions of animal brains that go beyond reactive control. Indeed, in many animal nervous systems, theoretical and computational analyses point to the presence of “hybrid” control architectures that combine elements of reactive control with integrative mechanisms that operate both in space, co-ordinating different parts of the body, and in time, organizing behavior over multiple time-scales (for discussion see, Verschure, Krose et al. (1992), Arbib and Liaw (1995), Prescott, Redgrave et al. (1999) and Sections 3 to 5 below).

Braitenberg relied on his imagination to conceive of how his vehicles might behave. The ability to do this degrades rapidly as internal or external complexity is added to the scenario envisaged, imagination is also susceptible bias. Embodied computational neuroscience on the other hand, via the medium of robotics, allows us to explore theories of brain architecture that reflect more of the true complexity of animal nervous

systems. By embedded hypothesized control mechanisms into the dynamics of a real-world interaction they are able to take seriously the role of body and environment in generating behavior. We can even do this in the same environment as that inhabited by the organism whose behavior we wish to understand. We present some specific examples of this possibility in section 2 below.

One of the first examples of neurorobotic modeling is the work by Edelman and co-workers on the NOMAD series of real-world artifacts (Edelman, Reeke et al. 1992). These robot models, also known as *brain-based devices* (see also Seth, McKinstry et al. (2004), Krichmar (2008), Krichmar and Wagatsuma (2011)), aimed at validating the theory of neuronal group selection, or “neural Darwinism”, that Edelman had proposed. In particular, they tested the putatively fundamental principle that differential selection of neuronal connections can lead to repertoires of neurons that are tuned to specific states of the world. Reentry or recurrence across these populations can allow the network to achieve higher-level function, and their behavior can be shaped through value-based learning such as that derived from reward and punishment.

Edelman’s brain-based devices helped to demonstrate that the interaction with the environment is important not just for understanding how we control behavior, but also for how we *acquire* new forms of control. From birth, action shapes the animal’s experience which in turn has a profound influence on the coding systems used by the brain (Verschure, Voegtlin et al. 2003, Lungarella, Pegors et al. 2005). These self-organize so as to efficiently represent the kinds of signals that occur most frequently, and also, perhaps, those signals that will be most effective in predicting how the sensory world will unfold in the near future (see, also Clark (2013)), or that are most behaviorally-relevant to the activity in which the animal is currently engaged (Prescott, Diamond et al. 2011). In the brain both of these forms of learning occur and we will later describe *Distributed Adaptive Control* theory (Verschure, Krose et al. 1992, Verschure, Voegtlin et al. 2003), which explores how different learning, and memory systems, in the mammalian brain, work together to generate perception, cognition and action. In chapter 14 (Neural Maps), Bednar and Williams also detail how self-organization shapes the function and development of neural maps that underlie the major mammalian perceptual systems.

The role of morphology in computation

Over time, evolution selects physical morphologies that create and enhance behaviorally relevant signals prior to any processing by the nervous system.

For instance, the harbor seal has perfected the skill of tracking the hydrodynamic trails left by fish prey as they move through the water allowing them to detect and follow an animal whose movement has disturbed the water up to several minutes previously (Dehnhardt, Mauck et al. 2001). The sensor mechanism by which seals achieve this task is a specialization of the facial vibrissae, which in seals have richly innervated follicles that allow the detection of minute changes in water currents. The capacity to detect these changes is enhanced by an adaptation in the morphology of the whisker whereby the smoothly tapered vibrissal shaft of land mammals has been replaced by a non-tapered and undulating shape (see Figure 3, left). Hydrodynamic tests of this whisker shape show that it has a remarkable capacity to eliminate the sensory noise due to the animal’s own movement (Hanke, Witte et al. 2010). Whereas a smooth and tapered

whisker would bend as the animal swims through the water due to the effects of drag, the undulating shape of the seal whiskers creates micro-vortices that cancel and largely eliminate drag giving a much improved signal-to-noise ratio for detecting disturbances in the water created by prey animals. These changes to the sensory system of the seal are matched by adaptations of the body, such as streamlining for efficient swimming, that demonstrate how mechanisms for perception co-evolve with those for action.

A second example of morphology simplifying the problem faced by the nervous system is provided by the mating behavior of the female cricket which requires it to identify and track-down a male of its species by attending to its distinctive chirping sound. The task of distinguishing and orienting towards a male “song” is facilitated by the physical structure of the cricket’s tracheal tube whose four openings (see Figure 3 right) differentially amplify sounds arriving at the left and right ears depending on the sound-source direction (Michelsen, Popov et al. 1994). The physics of this system is such that directional selectivity is greatest for auditory vibrations at the specific wavelength produced by the male animal, thus providing a non-neural mechanism that assists the female in orienting towards ethologically-appropriate stimuli. Using robot models, Webb and her co-workers have explored the relative contributions of physical morphology and of identified neural circuits in the cricket brain to generating integrate phonotaxis behavior (see Webb (2002) for review). Further examples of the importance of morphology for simplifying perception and control in animals and robots are discussed by Chiel and Beer (1997), and Pfeifer and Bongard (2006).




Figure 3. Examples of morphological computation in animal sensory systems. Left: The whisker of harbor seal (*Phoca vitulina*) in dorsal (A) and frontal view (B) showing the flattening, in the dorso-ventral direction, and presence of undulations that serve to reduce the drag caused by the whisker (from Hanke et al., 2010). Right: The tracheal tube of the cricket showing the four openings (two on each side) that provide enhanced directional-sensitivity to sounds at the wavelength produced by the male cricket (from Webb, 2002).

When is the real world better than simulation?

ENIAC, the first general purpose programmable electronic computer was used by the US military to calculate artillery firing tables based on measures such as the azimuth of the gun barrel, weight and shape of the missile, and the strength and direction of the wind (Reed 1952). Thus, since their origins, one of the important uses of computers has been to simulate the physics of the real world. Today, a range of simulated physics systems exist from the quick-and-dirty, but real-time, engines that drive gaming environments, to computationally intensive programs that provide accurate, but domain-limited simulation for specific problem spaces such as fluid dynamics, weather forecasting or the dynamics of global warming. Simulators are available that can be configured to replicate the physical design and sensory and actuating systems of standard robot platforms (see, e.g. Michel and Heudin (1998), Koenig and Howard (2004)). Special purpose simulators have also been developed that capture aspects of the physical

morphology of target animals such as the lamprey (Beauregard, Kennedy et al. 2006), quadrupeds such as the salamander (See Chapter 9 and Ijspeert (2000)), and, of course, humans (Fenner, Brook et al. 2008, Tikhonoff, Cangelosi et al. 2008). However, current simulations of real-world physics struggle to capture all of its mechanical and dynamical properties such as collision elasticity, surface friction, absorption, etc. Interactions between objects that are deformable, possess irregular surfaces, or are subject to environmental disturbances are a particular challenge (Ijspeert 2014). Validation of these properties, which may be essential for trusting the scientific conclusions obtained with simulations, can be an immense undertaking.

The use of simulation for exploring nervous system, body, and environment interactions is therefore not always a ‘no-brainer’. Simulators for games platforms are designed to give a believable impression of real-world physics but cut corners on the true physics in their pursuit of perceptual acceptability; special-purpose simulators are likewise generally not developed with the application of testing embodied brain models as their focus. Adapting any simulator to provide an embedded testing environment for a computational neuroscience model is therefore likely to require some work. Consider, for instance, the problem of how animals are able to track odor plumes in turbulent liquid flows. Experimental studies of animal behavior indicate that animals are sampling, and altering their behavior, at time scales (sub-second) and spatial scales (sub-millimeter) that are outside the limits of currently available simulators of real flows (see below). Below we explore how research in embodied robotics has allowed scientists to side-step this problem by embedding suitably-configured robots in real environments.

Although the world comes “for free” with a robot model, something that replicates key properties of the body has to be provided as the interface between model nervous system and the environment. Ideally the properties of the physical robot model should match relevant biomechanical constraints, identified in the animal, at least as far as these are thought to be relevant for the behaviors targeted. One of the trade-offs in modeling is then to decide when the benefits of having an embodied model, in terms of access to real-world physics, justifies the cost of building and running a bespoke physical system. As we will see in the examples discussed below, valuable results can often be obtained by attending to a key set of morphological constraints, rather than mimicking an entire creature. Indeed, holding to our view that a physical model should be useful rather than accurate, one usually abstracts away many of the physical details in order to get a good understanding of both the problem that a biological system is solving and the way in which this is accomplished. However, the challenge is also to be cogent and clear about these assumptions and how they affect the results and their interpretation (see Prescott, Montes Gonzalez et al. (2006) for further discussion of this issue).

3. Embodied models as means to understand animal sensorimotor and information processing capabilities

Often when we look at the animal world we see abilities that neither we, nor our technologies, possess. Hydrodynamic trail-tracking in seals is one example we have already mentioned, another is the memory abilities of Clarke’s nutcrackers, a bird who

can recall the locations of up to 30,000 seeds stored in the summer the following winter over a 25 km area (Balda and Kamil 1992). Such feats humble those of us who lose our car keys between the time we return home from work or school in the evening and the time we leave the next morning! The homing ability of salmon is another often cited, but under-appreciated, example. These fish are able to swim thousands of miles across the (to us) featureless expanse of the Pacific ocean to an exact location to which they had just a brief exposure years before, using learned magnetic, olfactory and visual cues (Crossin, Hinch et al. 2007). The secrets possessed by these birds and fish are not yet understood to a degree that permits their instantiation in designs for implementable technology. Therefore their study in streamlined simpler devices that are not encumbered by biological processes that are essential for the animal but not required for the task (waste elimination, reproduction, etc.), are an attractive approach to understanding how they do it.

Plume-tracking in a robot lobster

Nor are vertebrates the sole keepers of the secrets to information processing abilities not realized by human technology. Invertebrates also possess abilities that we have yet to understand and that we would like to have in our technologies. American lobsters, as well as several other crustacean species, are able to locate the source of an attractive smell by walking to it, following the odor over distances of many meters (Grasso and Basil 2002). This does not seem so startling until one realizes the complicated and unpredictable path that each whiff of aroma travels between the aromatic object and the lobster's "nose". The odor does not just make its leisurely slow way through molecular diffusion to the animal (this would take an inordinate amount of time) but it is carried by water flow. Thus the lobster's behavior, also known as *rheotaxis*, is "go with (or against) the flow." In the real ocean, where the lobster makes its living, real water flows are also chaotic or, in the language of fluid mechanics, they are *turbulent*.

Turbulence is not completely random but consists of eddies and vortices that appear and disappear on many scales and that interact with each other. The structure of turbulent flow is related to the Reynolds number, Re , where:

$$Re = \frac{\rho U l}{\mu}.$$

Here U is the mean velocity of the object relative to the fluid; l is a characteristic length used to specify spatial scale; μ is the dynamic viscosity of the fluid (it's "stickiness"); and ρ is the density of the fluid.

The value of Re gives an indication of the expected degree of turbulence. The ratio here is one of inertial forces (the numerator) to viscous forces (the denominator). If the value is less than one, then the viscous forces dominate and diffusion determines the rate and pattern of aroma transport from the source to the lobster's sensor. If it is greater than one, however, inertial forces dominate, and the flow patterns determine the rate and pattern of movement of aromas through space. The higher the Reynolds number the more chaotic. For each given range of numbers (say, 10-100, 100-1000 etc.), however, the pattern of flow, and therefore the movement of aromas through space, is "dynamically similar" regardless of the medium. In other words, within each range, characteristic dispersal patterns emerge that are the same if the fluid is air, water, or

liquid mercury, or if the spatial scale is a few centimeters in a fish tank or the atmosphere of the planet Jupiter. A fuller discussion of turbulence and of experimental and simulation approaches to understanding turbulence is provided by Benzi and Frish (2010).

It makes sense, then, that animal brains would evolve to make use of this universal patterning and, indeed, they have (see Vogel (1994) for a fascinating comprehensive tour of the relationship between fluid flow and design of biological systems). In the case of the lobster, it gets *better* at tracking the longer it is engaged in following a given turbulent odor plume; specifically, it moves faster, as if more confident, and also reduces its steering error relative to the target. This suggests that memories of recent encounters with whiffs of odor inform its later decisions (Moore and Atema 1991, Basil and Atema 1994). Unfortunately, the principle of “dynamic similarity” does not allow us to model the exact time series of aromas encountered by the animal using a simulation, rather, it merely tells us that the patterns are “similar”. Further, the more powerful Navier-Stokes equations, which describe all of turbulent flow, provide statistical averages of flows at points in space and time but not the *exact sequence* of inputs a single animal moving through a plume would receive. In the case of chemical plume tracking in turbulent flows, then, standard modeling techniques cannot provide us with an accurate series of concentration values along the path an animal would take, only a series of independent estimates of the averages. A test of the hypothesis that the lobster remembers the aroma it encounters during tracking, if conducted in such a simulator, might miss out on the key information that the lobster uses and therefore be no test at all.

On the other hand, applying a suitable *robotic* model of lobster information processing should be an effective way to adequately test the hypothesis. Simply place the physical model into real plumes and let its behavior provide evidence for the effectiveness of the proposed model in explaining the information processing that occurs in the animal (Figure 4). Indeed, when researchers did this they found a role for memory in a robot lobster (Grasso and Atema 2002) that might not have been obvious if the researchers had studied a simulated plume with a simulated lobster. Specifically, Grasso and Atema (2002) compared three different control schemes, exploiting (i) a single chemo-sensor, (ii) a stereo sensor pair, and (iii) a stereo pair with memory, as illustrated in Figure 5. A measure of overall tracking performance compared with a lobster showed significant benefit of the stereo sensor pair, and of the additional memory component, although the best model fell some way short of lobster performance, suggesting that there may yet be better algorithms to be discovered (perhaps by looking in more detail at the neural substrates).




Figure 4: Robot 'lobster' tracking a chemical plume. The problem of locating the plume source in this robot is decomposed into transforming a series of chemical concentration signals (visualized by dye in the photo) into a series of differential commands to the wheels of the robot which form a path in space. This robot demonstrated that a 'memory' of recent concentration series lead the robot to more efficient paths to the source; revealing the plausibility of the information processing method and memory for the biological lobster. From (Grasso and Atema 2002).

In this study the robot was a relatively simple device with two driven wheels capable of moving on the bottom of a fume or flow tank, or a relatively smooth portion of the ocean floor as shown in figure 4, with sensors for flow detection that broadly replicated the sensitivity of lobster antennae rather than seeking to match specific chemosensory mechanisms. Key features such as the scale and positioning of the sensors relative to the flow were tuned to match the biological target. The results do not rule out alternative models of lobster rheotaxis but, at the same time, they do provide proof-of-principle of the memory-based method as a plausible strategy that can work with real-world flows. More generally, these results suggest that as the physics of the part of the world that is of interest becomes more difficult to model, an effective approach is to use the world as its own model and to test physical models of embodied sensorimotor processing against it.




Figure 5. Illustration of the plume-tracking strategy used by the robot. The sequence of motion is

indicated by the progression of arrows. Chemical samples are collected at each successive location to inform the next move. Upstream motion (U), triggered by detection of the plume, is in a series of upstream steps. Failure to detect the plume at one of the sampling positions leads to an across stream motion called a *cast*. The choice of casting left (Cl) or right (Cr) is one of the parameters that is under algorithmic control. Grasso and Atema (2002) found that a strategy using memory of past patterns of stimulation to control the direction of cast gave results closest to those found with living animals. From Grasso and Atema (2002).

Modeling the neural circuits underlying bilaterian pattern generation

The example of lobster rheotaxis, described above, illustrates how robots can be used to evaluate theories of animal behavior couched at an algorithmic or informational level, however, embodied modeling can also be used to directly test the capacity of model neural circuits to generate adaptive behavior.

The nervous systems of all bilateral animals have been characterized as having an innate neuronal architecture based on *command* neurons, *coordinating* neurons and *central pattern generators* (CPGs) modulated by phase and amplitude modulating reflexes (Kennedy and Davis 1977, Stein 1978, Stein, Grillner et al. 1997, Pearson 2005, Jordan, Liu et al. 2008). These network components appear to underlie innate behaviors such as the control of posture, repetitive movement and inter-segmental coordination (Ayers 2002). In invertebrates, these components have been demonstrated and analyzed at the level of identifiable neurons and underlying molecular processes (Selverston and Ayers 2006) and key underlying mechanisms have been shown to be conserved between invertebrates and vertebrates (Grillner, Hellgren et al. 2005, Grillner, Markram et al. 2005, Pearson 2005). Biological cellular CPGs have been described in model animals where the relevant circuitry can be established by pairwise recordings between neurons (Selverston, Russell et al. 1976, Buchanan and McPherson 1995). An important aspect of these biological model systems is that they accommodate complex integrative phenomena such as neuromodulation (Hasselmo 1995, Dickinson 2006, Harris-Warrick 2011) and that the details of their function can be characterized by anatomical (Weeks 1981) and cellular dissection (Selverston 1980). (Note that, in Chapter 8, Fellous, Hasselmo and Canavier describe a number of rather different functions of neuromodulation in mammals.)

Key to developing a robot model that can capture the functionality of these circuits is the capability to compute in real time and to achieve realistic mechanistic models of phenomena such as neuronal integration (Ayers, Rulkov et al. 2010) and excitation/contraction coupling of artificial muscle (Witting, Ayers et al. 2000) within a physical plant that adequately approximates the biomechanics of the animal. This can be achieved by using phenomenological models of neurons and their synapses, incorporated into a CPG control architecture organized around exteroceptive reflexes (i.e. those concerned with external stimuli), and embedded within a biomimetic robot body plan (Ayers and Rulkov 2007, Ayers and Witting 2007, Westphal, Rulkov et al. 2011).

A critical component of bilaterian innate behavior is the CPG that drives central motor programs that mediate the behavioral contributions of single body segments (Selverston 2010). Although some CPGs produce relatively fixed action patterns (Hume, Getting et al. 1982) most are subject to extensive neuromodulation that allows the underlying networks to reconfigure to generate different behavioral acts (Heinzel 1988). In others,

gating synaptic input can switch between alternative patterns such as walking in different directions (Ayers and Davis 1977, Ayers and Witting 2007).

Using the lobster, lamprey, and honeybee as target animals, and focusing on the neural circuitry underlying locomotion, Ayers and Rulkov (2007) developed embodied models using electronic *Discrete Time Map-based* (DTM) neurons whose membrane dynamics evolve according to the state-space plot illustrated in Figure 6 which is intended to capture some of the dynamics of biological neurons. Here the membrane voltage in a given cycle $n+1$ is specified relative to its value in cycle n according to two difference equations based on synaptic current input and two control parameters α and σ whose values are specific to each neuron type, shape the function of the map, and define the characteristics of individual neurons. A detailed specification of this model is given in Box 1. As shown in the figure, when configured into different regions of α/σ space, neurons are either silent, tonically firing, bursting, or, in the highlighted saddle region between bursting and spiking, exhibit chaotic discharge (Shilnikov and Rulkov 2003).

To create synthetic neural networks that could simulate locomotor CPGs, Ayers and Rulkov configured networks of electronic neurons in bursting mode (Figure 6a) using predominantly inhibitory model synapses. The difference between postsynaptic voltage and the reversal potential of the spiking state of the presynaptic neuron was used to define the model synapses, with the topology of the network maintained in a look-up table of pre and post- synaptic neurons. Because the neuron model is based on difference equations rather than differential equations it is possible to control a large number of neurons and synapses in real time and tune their properties to behavioral context. This allows us to model phenomena such as neuromodulation (see below) using behaving robots (Ayers, Blustein et al. 2012). Moreover the option to vary chaos (figure 6b) allows us to model the adaptations underlying the wiggling and squirming behaviors observed in animals.




Figure 6. State space for Discrete Time Map-based (DTM) neurons. The yellow region (b) indicates where chaotic firing occurs. a & c indicate regions of bursting & spiking discharge Adapted from Rulkov (2002).

Figure 7 illustrates CPGs configured in this way as either reciprocal half-center (7a) or endogenous pacemaker inhibitory networks (7b) to model lamprey-like swimming (Westphal, Rulkov et al. 2011) and lobster-like walking respectively (Ayers and Davis 1977). In both types of networks, *synergies*—pools of neurons operating together—are formed by excitatory synapses between elements. Relative timing between synergies is then adjusted by differentially varying the control parameters α , σ , and the relative strength of the inhibitory synapses between synergies.

The half-center model of lamprey swimming, illustrated for a segment of the spinal cord in figure 7a, relies on reciprocal inhibition between the cross caudal interneurons (CCs) and on post inhibitory rebound (that is, increased neuronal excitability following the cessation of inhibition) in these cells (Perkel and Mulloney 1974, Buchanan and McPherson 1995, Buchanan 1996). This mechanism is sufficient to generate proportional symmetric alternating bursts driving bending of the tail to the left or right (Westphal, Rulkov et al. 2011).

In the lobster walking model (7b), an endogenous pacemaker-configured neuron, *elev* (*elevator*), establishes an oscillation with an antigravity synergy, *dep* (*depressor*), and a *stance* synergy that recovers later than *dep* from *elev* inhibition to create the late swing epoch (figure 7c). The *stance* synergy alternates with a *swing* synergy. *Elev* and *swing* generate the early phase while *swing* and *dep* generate the late swing phase (Figure 7c; see also figure 10 for patterns of model neural activity). The *stance* phase of stepping is mediated by the *stance* and *dep* synergies. Command neurons gate off synapses between protractor and retractor bi-functional synergies during forward and backward walking to mediate forward and backward walking, and between swing and stance and extensor and flexor synergies during lateral walking (Ayers and Davis 1977).




Figure 7. Central pattern generators for lamprey-like swimming (a) and lobster-like omnidirectional walking (b). In these diagrams the circles represent individual DTM neurons. In A, the contralateral synergies alternate due to intrinsic bursting mechanisms and reciprocal inhibition. Abbreviations: CC, cross caudal interneurons; LIN, lateral inhibitory neurons; EIN, excitatory interneuron (adapted from Westphal, Rulkov et al. (2011)). In B, differential inhibitory synaptic strength between the elevator synergy and depressor or swing synergy creates a three-phase rhythm characteristic of walking in the four directions. Descending walking commands gate connectivity between swing and stance phase interneurons and bi-functional motor neurons, at a pattern-generating nexus, to mediate walking in different directions. The three-phase pattern of walking is indicated in C.

Ayers and co-workers have used DTM networks to control three biomimetic robots *RoboLobster*, *RoboLamprey* and *RoboBee* (Ayers, Blustein et al. 2012), as shown in Figure

8. These robotic implementations have verified that the command neuron/coordinating neuron/CPG architecture, when instantiated with phenomenological models of neurons and synapses, can achieve adaptive behavior through swimming, walking and flying (Ayers, Blustein et al. 2012).




Figure 8. Biomimetic Robots (a). RoboLobster. (b) A helicopter proxy RoboBee. (c). RoboBee. (f). RoboLamprey. See Ayers, Blustein et al. (2012) for further details.

To achieve intersegmental coordination in a model system such as RoboLobster, interneurons are incorporated that pass information from one governing oscillator to another causing the phase advances or delays that mediate gait (Ayers and Selverston 1979). The system is also modulated parametrically by inter-segmental commands that shift the elements into a bursting mode and control average frequency (Ayers and Witting 2007). This architecture is illustrated in figure 9a.

Command neurons exist in a heterogeneous population that mediates a variety of behavioral acts (Bowerman and Larimer 1974a, Bowerman and Larimer 1974b). In the lamprey some commands preferentially excite anterior segments to mediate forward swimming while others excite more posterior segments to mediate backward swimming (Matsushima and Grillner 1992).

Exteroceptive reflexes link sensors to taxic and compensatory behavior (Figure 8b-c) and can operate in parallel when triggered by environmental contingencies. For example, parallel exteroceptive reflex circuits for collision and heading control in RoboLamprey are illustrated in figure 8c. In robot implementations, small microprocessors are used to analyze analog sensor reports and generate 'labeled-line' codes (Bullock 1978). For example, to guide heading control, a heading estimate, provided by an analog compass, is compared with a desired heading; a current proportional to this difference is then used to activate a "heading deviation" neuron. Modulatory interneurons can target the neuronal oscillator or the motor neurons directly to modulate amplitude.




Fig. 9. Inter-segmental modulatory systems. A. Command neuron, coordinating neuron central pattern generator organization of *RoboLobster*. B. Pathways for command (CN) and MI modulatory control of the *RoboLobster* CPGs by exteroceptive sensors. C. Layered exteroceptive reflex architecture of *RoboLamprey* to mediate collision and heading control. A hierarchy of lamprey inter-segmental and brain commands project the sensor inputs for heading deviation and rapid deceleration.

The exteroceptive reflex architecture works interchangeably between the three robot platforms demonstrating its adaptability to different physical morphologies (Ayers 2002). For example, a neuronal compass based on heading deviation neurons works equally well during walking in *RoboLobster*, swimming in *RoboLamprey* or flying in *RoboBee* (Ayers, Blustein et al. 2012, Westphal, Blustein et al. 2013). Reflexes have also been implemented for optical flow (optical flow sensors), hydrodynamic flow (antennae), heading (compass), collision (accelerometer), gravity (inclinometer), odometry (optical flow sensors with neuronal integrators) and beacon tracking (short baseline sonar array) and these operate well in parallel (Westphal, Blustein et al. 2013).

By approximating the biomechanics of the animal model, the dynamics of the robot vehicle behavior provides a good match to that of the living animal (Ayers 2004). Furthermore, in underwater robots, the use of shape memory alloy actuators (Mohd Jani, Leary et al. 2014) allows employment of heat to mediate excitation/contraction coupling much as living muscle uses intracellular calcium (Witting, Ayers et al. 2000). Thus the motor programs characteristic of the model neuromuscular system perform in an analogous fashion to that of the animal.

The processes that give rise to the motor rhythm (endogenous bursting, postinhibitory rebound, etc.) are also excellent models of the corresponding processes in the living networks (Selverston and Ayers 2006). A key feature of these models is that because they seek to capture the nonlinear dynamical behavior of neurons, rather than being neuronal conductance models, they are simpler, can operate in real time and are therefore suitable for robot control applications (Ayers and Rulkov 2007). As shown in Figure 10, the use of neuronal network models, instead of finite state systems, allows one to replicate in great detail the real behavior of the neurobiological system (a network) and, thanks to spiking nature of the models, provides a link between the electronic neurons and experimental measurements of neuronal activity from the animal.




Figure 10. Motor programs for forward and backward walking selected from an electronic neuronal network shown in Fig. 7b. The forward command gates off synapses linking the stance phase and protractor neurons as well as those linking the swing phase and retractor neurons.

Chemosensors based on principles of synthetic biology are currently under development that will allow tests of circuit-based models of odor-guided plume-following building on the more algorithmic approach described earlier (Grasso and Atema 2002). By merging the information processing hypotheses developed by Grasso and co-workers, with the more mechanistic understanding developed by Ayers and his associates it should be possible, in the future, to develop embodied models in which we have confidence that we are matching the target biological system at multiple levels of description. By matching multiple constraints in this way we can achieve a form *convergent validation* (Verschure 1996). That is, we not only obtain a more complete understanding, we also increase the likelihood that our model captures how the living system actually works rather than merely how it ‘might’ work.

4. Neurorobotics of the mammalian vibrissal system

The above embodied models have focused largely on invertebrate target animals, but can neurorobotics also help us to gain insight into brains more similar to our own; that is, those of other mammals?

One of the most popular model systems in which to study processing in mammalian brains is the whisker, or *vibrissal*, sensorimotor system owing to its discrete organization from the sensory apparatus (the whisker shaft) all the way to the sensory cortex (see Figure 11 and Diamond, von Heimendahl et al. (2008)), its relative ease of manipulation (for a living biological system), and, not least, its presence in laboratory rats and mice. Whisker signals are processed at multiple brain sites, crossing a minimum of two synapses, in the brainstem and in the thalamus, before reaching the sensory cortex. The whisker cortex, which is known as the *barrel cortex* due to the

presence of cellular aggregates that have a one-to-one mapping with the facial vibrissae, is itself a target for huge research effort (see, e.g. Petersen (2007)). This is due to the ability to be able to tweak a whisker at the periphery, in a known and quantifiable way, and then record in a precise area of cortex to determine the effect of that stimulus. Barrel cortex is therefore widely viewed as a preparation in which neuroscientific research could unlock the secret of the ‘cortical microcircuit’. That is, if we can understand the processing going on in the barrel field, then we may understand something about the generic processing capacities of mammalian six-layered cortex that is replicated, with some variation, across both of the cerebral hemispheres and in all mammalian species. It is no surprise then, that new data on the barrel cortex, is published on a weekly, if not daily, basis.



Figure 11. The rodent vibrissal system. The vibrissal system (left) of rats and mice is a major target of studies aimed at understanding mammalian brain architecture, partly due to the one-to-one correspondence (right) between single vibrissae on the snout and cellular aggregates in the vibrissal “barrel” cortex. Right figure adapted from Diamond, von Heimendahl et al. (2008).

But what of the stimuli with which neuroscientists are probing the rat brain in order to understand what the barrel cortex is doing? Perhaps unsurprisingly, many of the studies of barrel cortical function have been performed in anaesthetized animals or in brain slices. Less frequent historically, but now increasing in number, are studies in immobilized animals (head-fixed) that are awake and able to move their whiskers. Due to the difficulty of stably recording from electrodes implanted in free moving animals (where the recording device is usually connected via a springy umbilical cable) the number of studies that have looked at barrel cortex processing in more natural circumstances remains comparatively small. Moreover, the practical difficulties associated with recording in moving animals mean that such studies generally record extra-cellular activity—i.e. they use electrodes that pick-up activity in multiple nearby cells; the capacity to distinguish what is happening in single neurons is therefore limited. Overall then, the picture is much as we have described it in table 1—as you move from the freely moving rat to more tractable experimental models—head-fixed, anaesthetized, or slice—the ability to interrogate the processing system improves but the capacity of those signals to tell you about the nature of processing in awake naturally behaving animals falls away. This situation indicates there could scope for robotic models that can capture some of the physics of how whiskers interact with surfaces in exploring animals, and thus perhaps shed light on the nature of the processing occurring in this part of the mammalian sensorimotor system.

The closed loop of active vibrissal touch

There are two other strong motivations for investigating physical models of rodent vibrissal processing. First, as we have already encountered with lobster plume tracking, simulation of real sensory transduction can be difficult. In the case of whiskers, you have a tapered flexible shaft one end of which is embedded in a soft deformable body—the whisker follicle—packed with mechanoreceptors. Tactile sensing involves contact of the (usually) moving whisker with surfaces that often have complex microgeometry (e.g. texture), which may also be moving. Various simulation approaches have been developed including a recent dynamic model that includes effects of damping, inertia, and collisions (Quist, Seghete et al. 2014), however, accurate simulation of the encounters of whiskers with interesting surface geometry remains an unsolved problem that may be more easily addressed by building a physical model.

The second motivation concerns the nature of vibrissal sensing itself. Like many natural sensing systems, vibrissae are not deployed passively to detect the consequences of objects brushing against them. Rather, many whiskered animals, and particularly rodents, actively move their whiskers against objects and surfaces of interest (see figure 12). Indeed, rats and mice move their whiskers back and forth many times per second (a rate of about 8hz in rats), in a behavior known as “whisking”, generating multiple touches of the whiskers against salient objects. The control of whiskers happens alongside orienting movements of the head and body (Grant, Mitchinson et al. 2009, Grant, Sperber et al. 2012), thus we can think of whisking as one component of an active sensing system that uses the musculoskeletal system to help isolate and enhance the stimulus features of the environment that are of particular interest (Prescott, Diamond et al. 2011). We can immediately begin to see a major compromise that studies of brain slices and anaesthetized animals have to make when they study barrel cortex. Electrical stimulation of a brain slice, or passive deflection of the whisker of a sleeping mouse or rat, are both poor substitutes for the kinds of signals that are ascending from the brainstem of an awake and active animal. Even in the head-restrained case, the animal is detecting with its whiskers what the researcher has chosen to expose it too rather than deciding for itself, through head and body movements, where to deploy its whiskers (Mitchinson and Prescott 2013). Indeed, it is well-known that in head-fixed animals the normal whisking behavior is not generally expressed and when it is evoked (e.g. by stimulation of the olfactory system) may have somewhat different characteristics from the whisking of the free-moving animal.




Figure 12. Active vibrissal touch is supported by a nested-loop brain architecture. Many animals generate rhythmic whisking movements, shown here is a bout of rat whisking at two temporal scales as measured optoelectronically in a head fixed animal (Gao, Bermejo et al. 2001). The upper trace shows the average left and right whisker movement over the course of an 8.5s bout. The lower trace shows that the movement of the left and right whisker fields are usually strongly coupled. Vibrissal signals are processed in multiple nested sensorimotor loops that show a tight coupling between sensing and actuation (whisker movement) (adapted from Kleinfeld, Ahissar et al. (2006)).

How important is this for understanding the signals that are being processed in barrel cortex? This only becomes clear when we carefully study the nature of the whisker movements expressed by animals, and the architecture of the neural circuits that process vibrissal signals. In the case of the latter, it is now well known (e.g. Kleinfeld, Ahissar et al. (2006)) that the barrel cortex does not simply sit atop a feed-forward circuit relaying whisker deflection systems up from the periphery as somewhat misleadingly implied in figure 11 above. Rather, as indicated in figure 12 (right), it is embedded within a complex architecture of nested sensorimotor loops at the brainstem, midbrain, and cortical levels, each of which implements a relatively short-latency coupling between sensory input and motor output. Thus both the behavior of the animal (movement of the whiskers) and the neural circuits themselves, suggest that barrel cortical activity cannot be properly understood in isolation from the sensorimotor activity—exploratory whisking—that generates it. We return to consideration of this larger architecture after first homing in on the whiskers and their role in generating the sensory signals we find in the brain.

Towards a robot model—how engineers sometimes ask the right questions

What can we find out about the vibrissal system and processing in key brain areas such as the barrel cortex by building a robot model of the rat vibrissae? First off we have to build such a model. How then should the robot move its whiskers? Back-and-forth certainly, and at some speed, but when the whiskers touch an object should they keep moving forward regardless or should they change their behavior? Surprisingly, when researchers began to develop robot vibrissae the answer to this key question was not available in the scientific literature despite almost a century of effort to try to understand rat vibrissal sensing. Therefore, urged on by the robot engineers, Mitchinson, Martin et al. (2007) carefully recorded free-moving rats interacting with an object using their whiskers and studied the trajectories and timing of whisker motion,

before and after object contacts, in high-speed digital video recordings. These studies, and subsequent ones (Voigts, Sakmann et al. 2008, Grant, Mitchinson et al. 2009), have clearly demonstrated that rats modulate the movement of their whiskers on a millisecond basis depending on the nature of the contacts they are making with objects and surfaces. That is, the neural circuits really are implementing some tight feedback control of the sensor apparatus. This stands as an example of a question that was important to ask, but for which an important impetus came from an attempt to build a physical model. A more general lesson, that was also recognised Braitenberg (1986), is that when we approach a complex biological system with the aim of synthesizing an artifact that operates in a similar way, we are prompted to ask different questions than if we simply look at the target system and try to analyse its function.

Advancing neurobiological knowledge by doing experiments that would be impossible in the animal

Knowing how the whiskers should move, the next step was to build an approximation to the whisker processing apparatus found in the animal. We cannot *exactly* replicate the mechanics of rat whiskers, or of their embedding in the follicle, or the complex responses of the mechanoreceptors to whisker deflection, or the responses of the primary afferent cells in the brainstem that process the signals from these receptors (hundreds at a time). However, perhaps we could get close enough to learn something interesting. Mitchinson, Gurney et al. (2004) implemented, in simulation, an electromechanical model of the whisker follicle (see figure 13) and fed it with inputs from an artificial whisker (a tapered, flexible plastic shaft) instrumented with strain gauge transducers. By suitable tuning of the parameters of this model they showed that they could reproduce, using deflections of the artificial whisker, patterns of firing in model primary afferent cells that were qualitatively similar to those recorded when the whiskers of anaesthetized animals are deflected (Mitchinson, Gurney et al. 2004, Mitchinson, Arabzadeh et al. 2008).

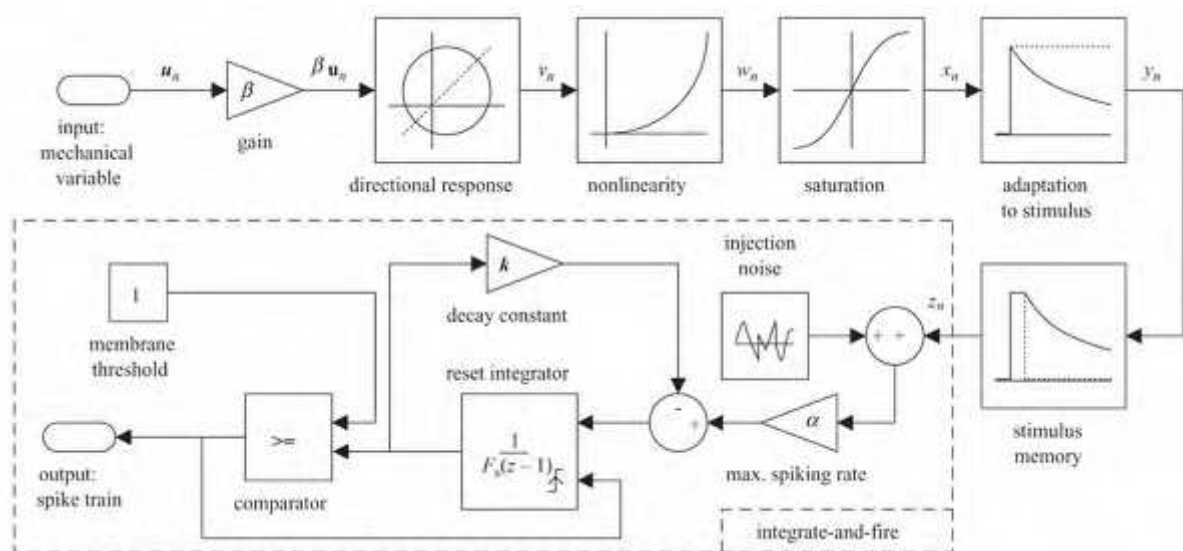


Figure 13. Using strain signals from an artificial whisker to model the response properties of vibrissal primary afferent neurons. The input (mechanical strain, u) is transformed, in turn, by processing units representing directional response, nonlinearity, saturation, adaptation to stimulus and stimulus memory; the resulting signal (z) drives a conventional integrate-and-fire membrane model. See Mitchinson, Gurney et al. (2004) for further details.

Now it becomes possible to perform an experiment that cannot be done in the animal (recall that a key criteria for a good model is that it should allow us to do interesting experiments with relative ease).

In order to understand the signals that are ascending from the vibrissae (as shown in Figure 11) for processing in the barrel cortex of awake behaving animals it is clear that we need to know what effect the feedback control that animal's exert on their whiskers is having on the ascending signals. The video studies showed that rats were modifying their whisker movements when they contacted objects. More precisely, the whiskers stop protracting (moving forward) and start retracting (moving backwards) earlier when they contact surfaces than when whisking in free air (see figure 14). The consequence is that whiskers bend less against surfaces than they would otherwise—you could say that the rat explores surfaces using a “light touch” (Mitchinson, Martin et al. 2007). The neural circuits that implement this feedback circuit in the animal are not well understood. Indeed, even if we knew where they were, experimentally interfering with them (e.g. attempting to block the feedback), would have unknown effects on the animal. Would an animal continue exploring objects as normal, or express anything like its natural behavior, if it were unable to regulate whisker movement in the usual way? However, the roboticists could perform the equivalent experiment in their model in a very simple way.




Figure 14. Effect of whisker-object contact on whisker motion. High-speed video recordings show that the protraction of the whisker often stops rapidly (12-14ms) after contact with a surface. In this sequence of frames the right whisker field touches an object in frame 2 (this is indicated by the white dot on the object just below the contact point), frames 3 and 4 show that the whiskers on that side of the snout then begin to retract whilst those on the opposite side are still moving forward. The overall effect is that whiskers make relatively light contacts with surfaces. [Figure adapted from Mitchinson, Martin et al. (2007) which should be consulted for more details of how this sequence was recorded].

A whisker pattern generation model was implemented that generated a sinusoidal whisking motion similar to that seen in the animal, and modulation was imposed on this to suppress further protraction, and initiate earlier retraction, following contact (Pearson, Pipe et al. 2007). Now this feedback mechanism could be easily turned on or off by toggling a variable in the control program. Figure 15 shows the effect of the presence or absence of feedback control on an important classes of model primary afferent neurons—so-called ‘rapidly adapting’ (RA) cells, as simulated with the Mitchinson, Gurney et al. (2004) electromechanical model. The figure shows that the effect of the feedback is to significantly alter the nature of the signals generated by these model cells (presented here as simulated spike rasters). Specifically, the responses of the cells were both briefer and less noisy when the feedback control was in place than when it was switched off. This outcome occurred, naturally enough, because the effect of the feedback was to produce a much smoother deflection pattern than when the feedback was deactivated. Of course, turning off protraction early should produce a smaller deflection, but what was interesting here was that despite having a smooth

sinusoidal drive and touch against a smooth surface, in the no feedback condition, there was a highly variable deflection pattern for the duration of the contact resulting in rather noisy model primary afferent responses.




Figure 15. Top. Whisker deflection pattern (left) and activity in model rapidly adapting (RA) primary afferent neurons in the condition where feedback control of sensor movement is switched for the robot *Whiskerbot* (pictured bottom left). For further details see Pearson, Pipe et al. (2007). Bottom right. The whiskered robot *Shrewbot* has recently been deployed to perform further neurobotic investigations of vibrissal touch (Pearson, Mitchinson et al. 2011) including tactile simultaneous localization and mapping (Pearson, Fox et al. 2013) and modeling of whisker-based predator-prey pursuit (Mitchinson, Pearson et al. 2014).

Based on these results, and their experimental studies in animals, Pearson et al. (2011) have proposed, and further demonstrated experimentally, in a new whiskered robot *Shrewbot* (figure 15, bottom right), that the effect of feedback in this system could be to provide more uniform signals, within a constrained dynamic range, for relay up to cortical areas. This could be very useful information for experimentalists trying to design neuroethologically-valid deflection patterns or neural stimulus trains with which to investigate their animal models, and for theoreticians trying to understand the nature of vibrissal processing in the barrel cortex microcircuit.

Further work has developed models of other neural structures—superior colliculus, cerebellum, barrel cortex, hippocampus—and evaluated their performance in the context of sensory inputs derived from active artificial vibrissal sensors (Prescott, Pearson et al. 2009, Anderson, Pearson et al. 2010, Pearson, Mitchinson et al. 2010, Mitchinson, Pearson et al. 2011). A review of this broader work which is seeking to render an embodied model of the full sensorimotor loop architecture of the rodent vibrissal system shown in figure 13 is available in Prescott, Mitchinson et al. (In press).

5. Scaling up to the whole brain

We have argued throughout this chapter that an effective strategy for understanding the relationship between brain and behavior is through abstraction to a physical model of reduced but appropriate complexity. In this final section we argue that, at an appropriate level of abstraction, the strategy of neurorobotics can also scale up to understanding the whole brain.

An elaborate robot-based theory that seeks to extend embodied modeling to encapsulate the control architecture of the vertebrate brain is the *Distributed Adaptive Control* (DAC) series of models proposed by Verschure and colleagues (see Verschure, Krose et al. (1992), Verschure (2012) for a review). DAC addresses the fundamental question that the theoretician and modeler faces: what level of abstraction is sufficient to allow a parsimonious description while in the meantime being sufficiently constrained by data to be meaningful? DAC solves this dilemma by advancing a dual agenda: an integrated systems-level architecture to identify the overall behavioral and functional constraints as described above, paired with anatomically and physiological constrained models to test specific assumptions and predictions.

We earlier saw the notion of nested-loop architecture applied to the vibrissal system. The DAC theory aims at conceptualizing a wider organizational scheme for the whole brain from which we can extract general principles to be mapped both to behavior and to its underlying neuronal substrate. The DAC theory starts from the fundamental question of what the function of the brain could be and proposes that it can be seen as optimizing four fundamental objective functions to generate action in the physical world or the '**How**' of survival:

1. **Why:** the motivation for action in terms of needs, drives and goals.
2. **What:** the objects in the world that actions pertain to.
3. **Where:** the location of objects in the world and the self.
4. **When:** the timing of action relative to the dynamics of the world.

We can call this the **H4W** problem for short and each of the Ws designates a large set of sub-objectives of varying complexity. Answering H4W requires the tight coordination of a number of processes ranging from control and sensing of the body through to the cognitive and integrative processes that underlie self-awareness. One can argue that a similar taxonomy could hold for any brain. This might be the case, and indeed the DAC theory has also been mapped to invertebrate systems such as ants and moths (Pyk, Bermudez et al. 2006, Mathews, Lechón M et al. 2009, Mathews, Verschure et al. 2010). However, vertebrate brains provide systems that explicitly process all aspects of H4W as opposed to invertebrates.

The DAC theory proposes that the brain solves H4W as an integrated control system comprising four layers called: *Soma*, *Reactive*, *Adaptive* and *Contextual*. Three functional columns intersect these layers delineating subsystems dealing with: exosensing states of the *world*, endosensing states of the *self* and the interfacing of world and self in *action* (Figure 16). DAC proposes that we can think of the brain, and its organization along the neuroaxis, as a matrix. Further, DAC proposes that as we move from the reactive to the contextual layer we advance from evolutionary ancient and predefined systems of the

brainstem towards the memory-dependent decision making circuits of the frontal lobes of the neo-cortex.

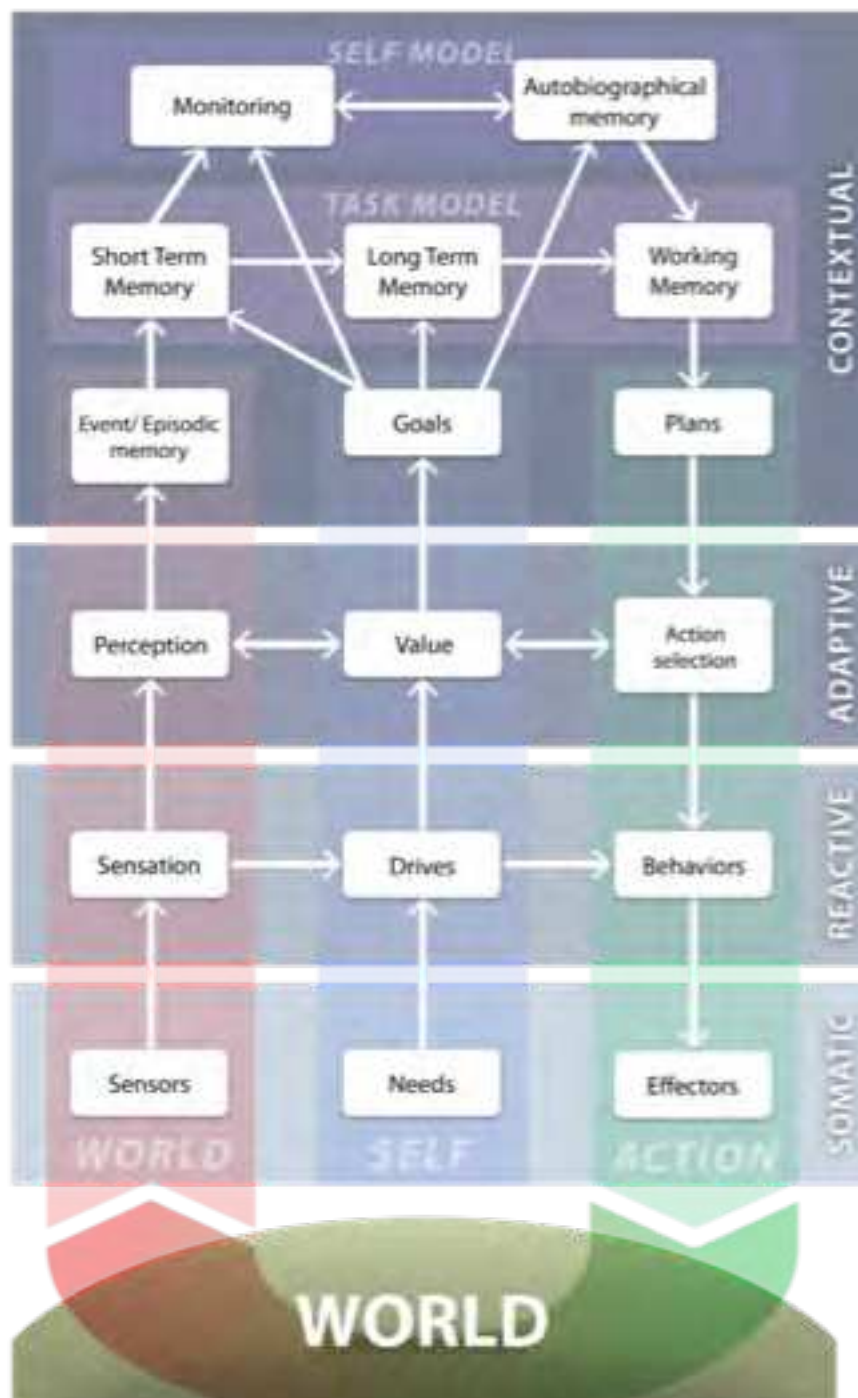


Figure 16. Conceptual scheme of the DAC architecture. DAC comprises four tightly coupled layers: *Soma*, *Reactive*, *Adaptive* and *Contextual*. Across these layers we distinguish three functional columns of organization: exosensing defined as the sensation, perception and retaining states of the *world* (left), endosensing which is detecting, signaling and retaining states of the physically instantiated *self* (centre), and the interface to the world through *action* (right). The arrows show the primary flow of information mapping exo- and endosensing into action. See text for further explanation. Figure from Verschure, Pennartz et al. (2014).

An attentive reader might object that Figure 16 is incomplete, or even incorrect, by raising concerns such as: Why does one go through episodic memory to get to short-term memory. Why does working memory depend on LTM but episodic memory does not? Where is procedural memory? Why is value lower than goals? However, obviously a two-dimensional that must fit on one page cannot express all of the relevant aspects of structural and functional brain organization. Rather, as the theory it reflects, it seeks to provide a suitable abstraction (Verschure, Pennartz et al. 2014). Nevertheless, the DAC theory has generated some very precise predictions concerning the organization of the neuronal substrate underlying adaptive behavior that demonstrate its utility. We will describe a few of these after reviewing the basic organization of the theory. For a more detailed analysis of the mapping of the DAC theory to the brain structures underlying goal-oriented see Verschure, Pennartz et al. (2014).

Within the DAC architecture, the *somatic layer* designates the body itself and defines three fundamental sources of information: *sensation* driven by external and internal sources of stimulation (or exo- and endosensing respectively), *needs* defined by the essential variables that assure survival (e.g. oxygen, glucose, carbohydrates, sleep, sex, etc.) and the skeletal-muscle system for *actuation* and the proprioception it provides. To keep the soma intact is the primary objective of the brain. This ability for self-maintenance was recognized by the 19th century physiologists Claude Bernard and Ivan Pavlov under the notion of *homeostasis* (see Verschure, Pennartz et al. (2014)). Later variations include the idea of *autopoiesis*, where the overall ability to maintain oneself is in turn a homeostatic variable (Varela, Maturana et al. 1974). From a robotics perspective we can operationalize these essential variables as the need to optimize energy and minimize damage. However, we do need to keep in mind that the robots we consider are not literal animals but rather operational approximations that allow our models to generate behavior in the real world. The DAC theory is of interest because it suggests how self-regulation can be realized in the embodied brain considered as an integrated system.

The *reactive layer* (Verschure, Krose et al. 1993) comprises fast predefined sensorimotor loops that support direct behaviors underlying basic functionality such as exploration, feeding, grooming, defense, orienting, mobility, sex, etc. Each of these reflexive behavior systems is triggered by low complexity signals, sensing, and directly coupled to specific need dependent internal states of the agent or drives. Hence, the primitive organizational elements of the reactive layer are sense-affect-act triads that are regulated following homeostatic principles (Sanchez-Fibla, Bernardet et al. 2010). This layer can be seen as a more advanced version of the reflex-driven robot tortoise built by Grey Walter, the “vehicles” proposed by Valentino Braitenberg or Brooks’ subsumption architecture (see section 1 above). The key difference is that the DAC reactive layer is modeled after the primitive brainstem structures that govern these behavior systems in the vertebrate brain (Panksepp and Biven 2012).

In the reactive layer every behavior system is conceived as a homeostatic control system that in turn is regulated by an integrative *allostatic* process. Whereas homeostasis considers a dynamic process as seeking stability, allostasis considers self-regulation, in general, as a continuous dynamic process—seeking stability through change (Sterling and Eyer 1988). In our case this means that each individual behavioral homeostatic system, such as the need for fight or flight, is regulated in real-time relative to the

pertinent needs of the behaving system. The control model proposed here is that behavior systems are individually tuned to specific properties of the environment, such as the home base for security or the open field for exploration, while being continuously tuned with respect to the overall behavioral needs of the agent. Robotic experiments show that these drive-reducing reactive control systems can give rise to complex patterns of environment exploration (Sanchez-Fibla, Bernardet et al. 2010) similar to those measured in animals (see figure 17). Recent research has applied this model to provide allostatic control of a humanoid robot engaged in human-robot social interaction (Vouloutsis, Lallée et al. 2013).




Figure 17. Behavior generated by the reactive layer of DAC controlling a robot in the real world while *exploring* a familiar environment (left) or a *novel* environment (right). The familiar environment triggers exploration behavior that brings the robot out of the home base at the lower-right corner, while the novel environment induces security seeking so that robot moves primarily within the home-base region. Insets show example trajectories from rats under similar conditions. Adapted from Sanchez-Fibla, Bernardet et al. (2010).

The *adaptive layer* of DAC is interfaced to the full sensorium of the agent, its internal needs and effector systems (Verschure and Coolen 1991, Verschure, Krose et al. 1992, Duff and Verschure 2010), extending the sensorimotor loops of the reactive layer with acquired sensor and action states. Through learning, a state space of sensation and action helps the agent deal with the fundamental unpredictability of both the internal and the external environment. This is signified in figure 16 by “Perception” and “Action” which are both dependent on plasticity and memory. Moreover, the adaptive layer of DAC provides a solution to the fundamental *symbol-grounding problem* that real-world systems face (Searle 1980, Harnad 1990, Verschure, Krose et al. 1992, Verschure 1998). The core issue is that a system that needs to learn to act has to associate actions with states of the world. These states must be picked up with its sensors and the resulting internal representations grounded in the experience of the agent itself, as opposed to being predefined by an external entity such as the system programmer.

Learning in the adaptive layer occurs in a restricted temporal window of relatively immediate interaction. In order to go beyond “the now”, other learning systems must be engaged that can expand this time window. The adaptive layer and the state-space it acquires thus provide a foundation for the more advanced learning and memory

systems of the contextual layer. More specifically, the *contextual layer* of DAC develops behavioral plans or policies (Voegtlin and Verschure 1998, Verschure and Voegtlin 1999) the atomic elements of which are formed by the state space of exo- and endosensing built by the adaptive layer, and by its sensorimotor contingencies. The contextual layer forms sequential short- and long-term memory representations (STM and LTM) that reflect states of the environment and actions generated by the agent or its acquired sensory-motor contingencies. The acquisition and retention of these sequences is conditional on the goal achievement of the agent. The behavioral plans can be recalled through sensory matching and internal chaining among the elements of the retained memory sequences. The dynamic states that this process entails can be interpreted as DAC's working memory system.

The development of the DAC architecture has taken place through a close interaction between theoretical investigations of the neurobiology and psychology of living things, and a wide variety of robotics models of different kinds, and at different levels of abstraction compared to the biological target. An example that shows the importance of embodied modeling, in the development of the DAC framework, is research that lead to the understanding of the phenomenon of behavioral feedback in the coupling between perceptual and behavioral learning (Verschure, Voegtlin et al. 2003). Behavioral feedback describes a process through which learning leads to the development of behavioral habits that, in turn, lead to the creation of an effective interaction space that is of lower complexity than the overall environment. This reduction in complexity means that the environment is sub-sampled and the perceptual learning systems further adjust to this reduced set of interactions. Behavioral feedback therefore directly couples behavioral and perceptual learning through behavior itself (Verschure, Voegtlin et al. 2003) (see also "closing the loop" in section 1 above). This direct impact of situatedness on perception was first analyzed through embodied learning using a robotic model.

Answering the "where" question

We conclude this section by looking at the practical impact of the DAC research strategy and by analyzing a specific problem—how to build a map of the world and localize oneself in it, or the problem of "where" of the H4W taxonomy.

Earlier we noted some of the remarkable feats of navigation and spatial knowledge shown by some species of birds and fish. For many animals, an essential aspect of survival is to build up appropriate knowledge of the environment and to know the position within it of the self and of other significant locations and objects. In robotics this is widely referred to as the *Simultaneous Localization and Mapping* problem (SLAM) (Leonard and Durrant-Whyte 1991, Thrun, Burgard et al. 2005). SLAM aims at simultaneously estimating locations of novel landmarks, and of the robot itself, while incrementally building a map of the environment. Solving the SLAM problem is seen as the technological breakthrough required for robots to perform autonomous tasks requiring mobility (Thrun 2010). The most popular methods for SLAM are based on estimation-based approaches (Dissanayake, Newman et al. 2001), grounded in adaptive filter techniques starting with the Kalman filter (Kalman 1960) (the workhorse of many control engineering applications). However, solutions to SLAM face challenges due to various sources of uncertainty such as errors in path integration, sensor noise, and, above all, the dynamics of the world itself. Biological systems display robust solutions to SLAM, it is thus no surprise that attention has been paid to the ways in which animals

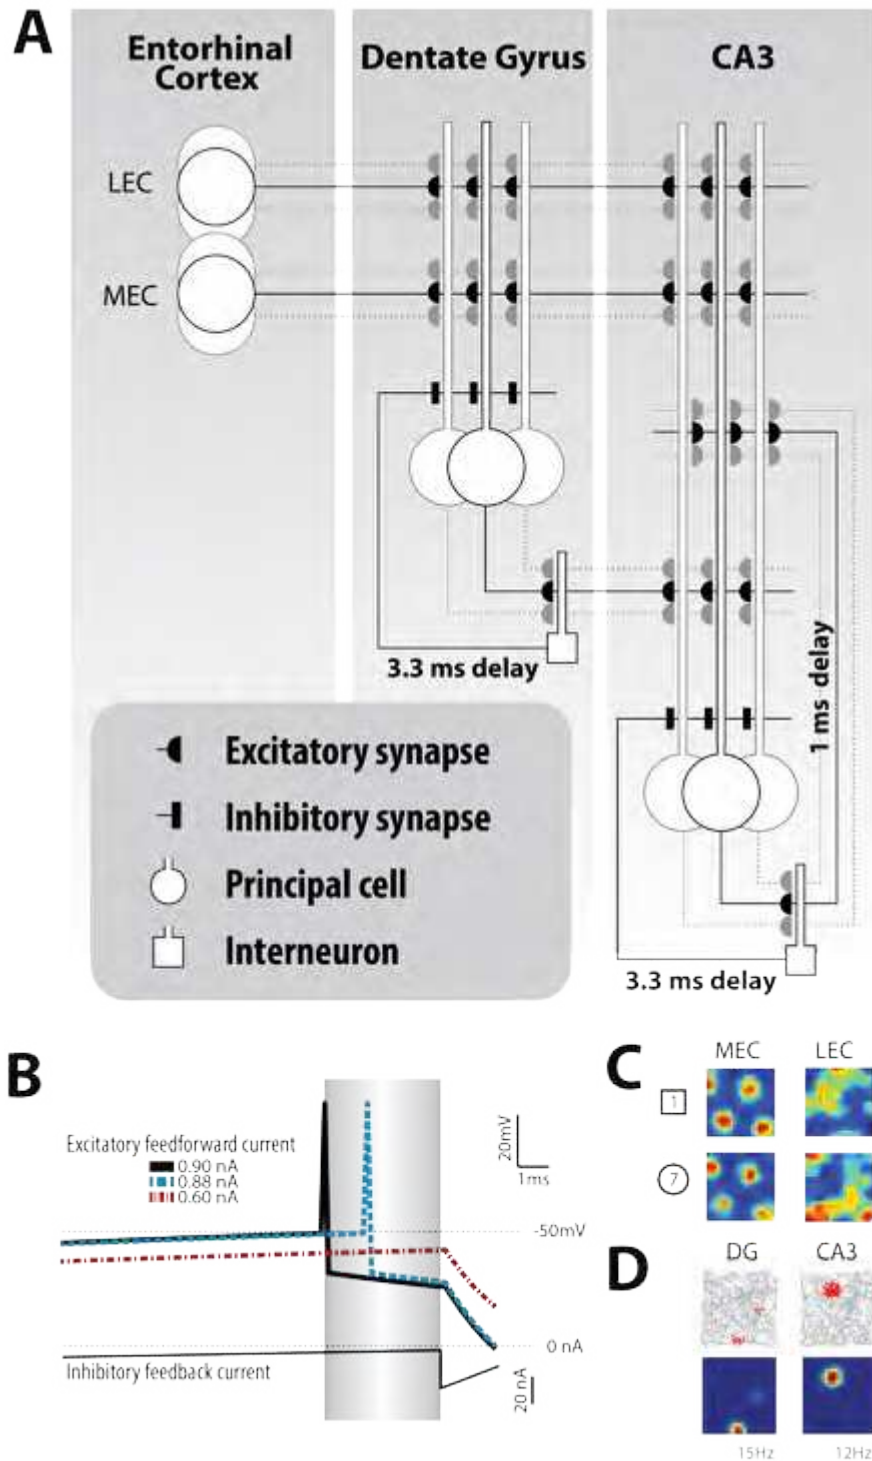
solve spatial tasks (Milford and Wyeth 2010, Sünderhauf and Protzel 2010). Understanding how animals solve the SLAM problem could thus serve a dual purpose, both to provide good theories of animal spatial cognition and to provide a useful biomimetic advance in mobile robotics.

In mammals, spatial cognition is associated with the hippocampus and adjacent areas in the medial temporal lobe. Within this area of the brain a number of cell types have been described that relate to the position in space of the animal such as place cells (O'Keefe and Dostrovsky 1971, O'Keefe and Conway 1978), head-direction cells (Taube, Muller et al. 1990), and grid cells (Hafting, Fyhn et al. 2005). Place cells show a response to a specific location in space, head-direction cells indicate the heading direction of the animal, while grid cells show a response field that includes a grid of precisely-tuned location-specific responses. A common brain-based approach towards solving simple navigation tasks is to combine model place cells with reinforcement learning to associate actions to locations in space (e.g. Burgess, Donnet et al. (1997), Foster, Morris et al. (2000), Arleo and Gerstner (2001)). Recent studies have claimed that such hippocampus-inspired models are capable of competing with state-of-the-art SLAM algorithms in large-scale benchmarks (Milford and Wyeth 2010). However, a limitation of these models is that they make the strong assumption that perception itself does not depend on behavior, for example, they often rely on a predefined discretized model of space, and do not always take into account critical physiological and anatomical properties of the relevant brain structures.

We turn to a specific model developed within the DAC framework (Verschure, Krose et al. 1992, Verschure 2012) that illustrates how predictions formed at the level of the overall neuronal architecture can be mapped to the neuronal substrate. As described earlier, DAC postulates a contextual layer that operates on highly integrated, multi-modal representations of sensorimotor couplets that comprise both sensing and action. This assumption constitutes a testable hypothesis and the question thus becomes whether the brain operates on such integrated sensorimotor representations.

The hippocampus interfaces with the neocortex via the *entorhinal cortex* (EC), receiving sensory information via the lateral EC, and spatial information from the grid cells of the medial EC. Given these inputs, Lisman (2005, 2007) has hypothesized that the *dentate gyrus* (DG), which is the first processing stage of the hippocampus, may form the sense-act couplets predicted by the DAC architecture (Lisman 2007). Using an anatomically and physiologically constrained model, Rennó-Costa, Lisman et al. (2010, 2014) have shown that representations in the lateral and medial EC may be multiplexed in the responses of the neurons in the DG (Figure 18). The model proposes that the first step in the construction of sense-act couples predicted by DAC occurs in the fixed divergent mapping of the EC onto the DG. This model also explains the specific phenomenon of “rate remapping” of the population response in DG when visual cues are smoothly varying due to the “morphing” of the environment (e.g. when a square environment is gradually modified to become circular). Rate remapping describes the finding that the correlation of the population response of the DG in morphed environments smoothly degrades (remaps) as environments become more dissimilar (Leutgeb, Leutgeb et al. 2007). The DAC model of the DG can account for this effect but, more importantly, also shows that the neurons in the DG provide an integrated encoding of action and sensory states, as is assumed for the DAC contextual layer, which drives place cell activity in the

hippocampal CA3 region. Rennó-Costa Rennó-Costa, Luvizotto et al. (2011) have also shown that this detailed neuronal model of mammalian spatial cognition can provide for robust SLAM in a mobile robot, providing additional evidence of the sufficiency of this account as a theory of animal navigation behavior.



***Figure 18. A computational model of the mammalian temporal lobe as a test of the sensorimotor multiplexing hypothesis of DAC and its impact on spatial memory.** (A) Excitatory granule cells of DG receive convergent input from *entorhinal cortex* (EC) (convergence ratio of 2700:1) with delayed inhibitory feedback (delay: 3.3 ± 0.4 msec) from local

interneurons. Excitatory cells in hippocampal CA3 receive convergent input from both EC (ratio 2900:1) and from *dentate gyrus* (DG) (ratio 50:1) together with delayed feedback inhibition from local interneurons (delay: 3.3 ± 0.4 msec) and recurrent excitatory input (delay: 1 msec). (B) Delayed feedback inhibition defines a winner-take-all competition that selects which cell fires in a given cycle. The diagram provides a trace of three sample cells with different strength of excitatory feed-forward input. Time is represented by the horizontal axis, the gray area designates the window between the first spike and the onset of global inhibition. Cell voltage and input currents are shown on the ordinate. (C) Rate maps of sample EC grid cells exposed to two shapes of an environment. (D) Action potentials (red dots) with overlaid trajectory of a simulated agent (gray line) and equivalent rate maps of sample DG and CA3 cells showing a well-defined place cell response in CA3.

Conclusion—advancing brain theory through neurorobotics

The approach of embodied modeling has a broad range of benefits and contributions to bring to brain science. Neurorobotics adds to the computational neuroscience approach the capacity to directly test models in the real world, avoiding the need to simulate in detail the physics of organism-environment interaction, and providing validation that our theoretical models can generate integrated and appropriate behavior in real-time. Perhaps as importantly, the challenge to create neurorobotic models, demands that we also close the loop between perception and action, and this serves to remind us that there is a tight coupling between brains, bodies, and the environment, and that behavior, and therefore any explanation of behavior, must take into account all three.

Box 1. Discrete-time Map based (DTM) Neurons

The DTM architecture is a map-based model that captures the dynamics of spiking and bursting neurons and their synaptic interactions (Rulkov 2002). It allows the construction of complex sensorimotor networks that operate in real time on small digital signal processors (Westphal, Rulkov et al. 2011, Ayers, Blustein et al. 2012). It also allows dynamic tuning of synaptic strengths, time constants (Ayers, Rulkov et al. 2010). Our goal with this model is to program the motor programs with neurons and synapses, hence, there are no algorithms in the controller. Each neuron and synapse is represented as an object in a look-up table. In a run-loop, the membrane voltage of a model cell in cycle $n+1$ is determined as a function of its value in cycle n according to two difference equations based on synaptic current input and two control parameters σ (equation 1) and α (equation 2) and, whose values specify the neuron type:

$$\begin{aligned} x_{n+1} &= f_{\alpha}(x_n, x_{n-1}, y_n + \beta_n), \\ y_{n+1} &= y_n - \mu(x_n + 1) + \mu\sigma + \mu\sigma_n, \end{aligned} \quad (1)$$

where x_n is the fast and y_n is the slow (due to $0 < \mu < 1$) dynamical variables. The nonlinear function is written in the following form

$$x_{n+1} = f_n(x_n, x_{n-1}, u) = \begin{cases} \alpha/(1 - x_n) + u, & x_n \leq 0, \\ \alpha + u, & 0 < x_n < \alpha + u \text{ and } x_{n-1} \leq 0, \\ -1, & x_n \geq \alpha + u \text{ or } x_{n-1} > 0, \end{cases} \quad (2)$$

where the third argument $u = y_n$ or a combination of input variables that depend on the model type.

The parameters α and σ (figure 6) shape the function of the map and define characteristics of individual neurons. Input variables β_n and σ_n incorporate the action of synaptic inputs I_{syn} and can be written as $\beta_n = \beta_E I_{syn}$, $\sigma_n = \sigma_E (I_{syn})$ where β_E and σ_E are constants that control how quickly neurons respond to the input and support dynamical mechanisms for spike frequency adaptation to DC pulses of current (Rulkov 2002). If y is omitted, the 1-D model generates tonic spiking activity.

The simplest map-based model for a synaptic current can be written as:

$$I_{n+1}^{syn} = \gamma I_n^{syn} - \begin{cases} g_{syn}(x_n^{post} - x_{rp}), & \text{spike}_{pre}, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where g_{syn} is the strength of synaptic coupling, indexes *pre* and *post* stand for the presynaptic and postsynaptic variables, respectively. Here γ controls the relaxation rate of the synapse ($0 < \gamma < 1$) and x_{rp} defines the reversal potential and, therefore, the type of synapse: excitatory or inhibitory. The synaptic current from equation 3 is summed with equation 2 resulting in equation 1.

Synapse objects are used to connect neurons. The synapse equations evaluate the difference between postsynaptic voltage and reversal potential and the spiking state of the presynaptic neuron (equation 3). The topology of the network is defined in the look-up table of pre and post synaptic neurons for each synapse. Every neuron can project to a variety of neurons and can receive inputs from multiple neurons.

The model can be executed on a variety of processors including the Lego Mindstorms™ Brick (Westphal, Rulkov et al. 2011, Blustein, Rosenthal et al. 2013, Westphal, Blustein et al. 2013) and can be implemented in LabView™ or procedural C. In procedural C, an initiation function sets the parameters for each neuron and synapse while a run-time loop updates all objects at a discrete time interval. The neurons and synapses otherwise proceed asynchronously.

Bibliography

- Anderson, S. R., M. J. Pearson, A. Pipe, T. Prescott, P. Dean and J. Porrill (2010). "Adaptive cancelation of self-generated sensory signals in a whisking robot." Robotics, IEEE Transactions on **26**(6): 1065-1076.
- Arbib, M. A. and J.-S. Liaw (1995). "Sensorimotor transformations in the worlds of frogs and robots." Artificial Intelligence **72**(1-2): 53-79.
- Arleo, A. and W. Gerstner (2001). "Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity." Biol. Cybern. **83**: 287--299.
- Ayers, J. (2002). A conservative biomimetic control architecture for autonomous underwater robots. Neurotechnology for Biomimetic Robots. J. Ayers, Davis, J. and Rudolph, A. Cambridge, MA, MIT Press: 234-252.
- Ayers, J. (2004). "Underwater walking." Arthropod Structure & Development **33**(3): 347-360.
- Ayers, J., D. Blustein and A. Westphal (2012). "A Conserved Biomimetic Control Architecture for Walking, Swimming and Flying Robots." Biomimetic and Biohybrid Systems **7375**: 1-12.
- Ayers, J. and W. J. Davis (1977). "Neuronal Control of Locomotion in the Lobster, (*Homarus americanus*) I. Motor Programs for forward and Backward walking." Journal of Comparative Physiology **115**: 1-27.
- Ayers, J., Davis, J. and Rudolph, A. (2002). Neurotechnology for Biomimetic Robots. Cambridge, MIT Press.
- Ayers, J. and N. Rulkov (2007). Controlling Biomimetic Underwater Robots with Electronic Nervous Systems. Bio-mechanisms of Animals in Swimming and Flying. N. Kato and S. Kamimura. Tokyo, Springer-Verlag: 295-306.
- Ayers, J., N. Rulkov, D. Knudsen, Y.-B. Kim, A. Volkovskii and A. Selverston (2010). "Controlling Underwater Robots with Electronic Nervous Systems." Applied Bionics and Biomimetics **7**: 57-67.
- Ayers, J. and A. I. Selverston (1979). "Monosynaptic entrainment of an endogenous pacemaker network: a cellular mechanism for von Holst's magnet effect." Journal of comparative physiology A, Neuroethology, sensory, neural, and behavioral physiology **129**: 5-17.
- Ayers, J. and J. Witting (2007). "Biomimetic Approaches to the Control of Underwater Walking Machines." Phil. Trans. R. Soc. Lond. A **365**: 273-295.
- Balda, R. P. and A. C. Kamil (1992). "Long-term spatial memory in clark's nutcracker, *Nucifraga columbiana*." Animal Behaviour **Volume 44**(4): 761-769.
- Basil, J. A. and J. Atema (1994). "Lobster orientation in turbulent odor plumes: simultaneous measurement of tracking behavior and temporal odor patterns." Biol. Bull. **187**: 272-283.
- Beauregard, M., P. J. Kennedy and J. Debenham (2006). "Fast simulation of animal locomotion: lamprey swimming." Professional Practice in Artificial Intelligence **218**: 247-256.
- Beer, R. D. and P. L. Williams (2014). "Information Processing and Dynamics in Minimally Cognitive Agents." Cognitive Science: n/a-n/a.
- Bengston, S., Ed. (1994). Early Life on Earth. New York, Columbia University Press.
- Benzi, R. and U. Frisch (2010). "Turbulence." Scholarpedia **5**(3): 3439.
- Blustein, D., N. Rosenthal and J. Ayers (2013). "Designing and implementing nervous system simulations on LEGO robots." JoVE (Journal of Visualized Experiments)(75): e50519-e50519.

- Bowerman, R. F. and J. L. Larimer (1974). "Command fibres in the circumoesophageal connectives of crayfish I. Tonic fibres." J. Exp. Bio. **60**: 95-117.
- Bowerman, R. F. and J. L. Larimer (1974). "Command fibres in the circumoesophageal connectives of crayfish II. Phasic fibres." J. Exp. Biol. **60**: 119-134.
- Braitenberg, V. (1986). Vehicles: experiments in synthetic psychology. Cambridge, MA, MIT Press.
- Brooks, R. A. (1986). "A robust layered control system for a mobile robot." IEEE Journal on Robotics and Automation **RA-2**: 14-23.
- Buchanan, J. T. (1996). "Lamprey spinal interneurons and their roles in swimming activity." Brain, behavior and evolution **48**(5): 287-296.
- Buchanan, J. T. and D. R. McPherson (1995). "The neuronal network for locomotion in the lamprey spinal cord: evidence for the involvement of commissural interneurons." J Physiol Paris **89**(4-6): 221-233.
- Bullock, T., Orkand, R. Grinnell, A. (1978). Introduction to Nervous Systems. San Francisco, CA, W. H. Freeman
- Burgess, N., J. G. Donnet, K. J. Jeffery and O. K. J (1997). "Robotic and neuronal simulation of the hippocampus and rat navigation." Phil. Trans. R. Soc. Lond. B. Biol. Sci. **352**: 1535--1543.
- Chiel, H. J. and R. D. Beer (1997). "The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment." Trends in Neurosciences **20**: 553-557.
- Clark, A. (2013). "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." Behavioural and Brain Sciences **36**(3): 181-204.
- Crick, F. (1990). What Mad Pursuit: A Personal View of Scientific Discovery, Basic Books (AZ).
- Crimes, T. P. and M. M. Anderson (1985). "Trace fossils from late Precambrian—early Cambrian strata of southeastern Newfoundland (Canada): Temporal and environmental implications." Journal of Paleontology **59**(2): 310–343.
- Crossin, G., S. Hinch, S. Cooke, D. Welch, S. Batten, D. Patterson, G. Van Der Kraak, J. Shrimpton and A. Farrell (2007). "Behaviour and physiology of sockeye salmon homing through coastal waters to a natal river." Marine Biology **152**(4): 905-918.
- Dario, P., M. C. Carrozza, E. Guglielmelli, C. Laschi, A. Menciassi, S. Micera and F. Vecchi (2005). "Robotics as a future and emerging technology: biomimetics, cybernetics, and neuro-robotics in European projects." Robotics & Automation Magazine, IEEE **12**(2): 29-45.
- Dehnhardt, G., B. Mauck, W. Hanke and H. Bleckmann (2001). "Hydrodynamic trail-following in harbor seals (*Phoca vitulina*)." Science **293**(5527): 102-104.
- Diamond, M. E., M. von Heimendahl, P. M. Knutsen, D. Kleinfeld and E. Ahissar (2008). "'Where' and 'what' in the whisker sensorimotor system." Nat Rev Neurosci **9**(8): 601-612.
- Dickinson, P. S. (2006). "Neuromodulation of central pattern generators in invertebrates and vertebrates." Current Opinion in Neurobiology **16**(6): 604-614.
- Dissanayake, M. W. M. G., P. Newman, S. Clark, H. F. Durrant-Whyte and M. Csorba (2001). "A solution to the simultaneous localization and map building (SLAM) problem." Robotics and Automation, IEEE Transactions on **17**(3): 229-241.
- Duff, A. and P. F. M. J. Verschure (2010). "Unifying Perceptual and Behavioral Learning with a Correlative Subspace Learning Rule." Neurocomputing **73**: 1818–1830.

- Edelman, G. M., G. N. Reeke, Jr., W. E. Gall, G. Tononi, D. Williams and O. Sporns (1992). "Synthetic neural modeling applied to a real-world artifact." Proc Natl Acad Sci U S A **89**(15): 7267--7771.
- Fenner, J. W., B. Brook, G. Clapworthy, P. V. Coveney, V. Feipel, H. Gregersen, D. R. Hose, P. Kohl, P. Lawford, K. M. McCormack, D. Pinney, S. R. Thomas, S. Van Sint Jan, S. Waters and M. Viceconti (2008). The EuroPhysiome, STEP and a roadmap for the virtual physiological human.
- Foster, D. J., R. G. M. Morris and P. Dayan (2000). "A model of hippocampally dependent navigation, using the temporal difference learning rule." Hippocampus **10**: 1--16.
- Gao, P., R. Bermejo and H. P. Zeigler (2001). "Whisker deafferentation and rodent whisking patterns: behavioral evidence for a central pattern generator." J Neurosci **21**(14): 5374-5380.
- Grant, R. A., B. Mitchinson, C. Fox and T. J. Prescott (2009). "Active touch sensing in the rat: Anticipatory and regulatory control of whisker movements during surface exploration." Journal of Neurophysiology **101**(2): 862-874.
- Grant, R. A., A. L. Sperber and T. J. Prescott (2012). "The role of orienting in vibrissal touch sensing." Front Behav Neurosci **6**: 39.
- Grasso, F. W. and J. Atema (2002). "Integration of Flow and Chemical Sensing for Guidance of Autonomous Marine Robots in Turbulent Flows." Environmental Fluid Mechanics **2**(1): 95-114.
- Grasso, F. W. and J. A. Basil (2002). "How lobsters, crayfishes, and crabs locate sources of odor: current perspectives and future directions." Current Opinion in Neurobiology **12**(6): 721-727.
- Grillner, S., J. Hellgren, A. Ménard, K. Saitoh and M. A. Wikström (2005). "Mechanisms for selection of basic motor programs – roles for the striatum and pallidum." Trends in Neurosciences **28**(7): 364-370.
- Grillner, S., H. Markram, E. De Schutter, G. Silberberg and F. E. LeBeau (2005). "Microcircuits in action—from CPGs to neocortex." TRENDS in Neurosciences **28**(10): 525-533.
- Hafting, T., M. Fyhn, S. Molden, M. B. Moser and E. I. Moser (2005). "Microstructure of a spatial map in the entorhinal cortex." Nature **436**(7052): 801-806.
- Hanke, W., M. Witte, L. Miersch, M. Brede, J. Oeffner, M. Michael, F. Hanke, A. Leder and G. Dehnhardt (2010). "Harbor seal vibrissa morphology suppresses vortex-induced vibrations." J Exp Biol **213**(Pt 15): 2665-2672.
- Harnad, S. (1990). "The symbol grounding problem." Physica **42**: 335--346.
- Harris-Warrick, R. M. (2011). "Neuromodulation and flexibility in central pattern generator networks." Current opinion in neurobiology **21**(5): 685-692.
- Hasselmo, M. E. (1995). "Neuromodulation and cortical function: modeling the physiological basis of behavior." Behavioural brain research **67**(1): 1-27.
- Heinzel, H. G. (1988). "Gastric mill activity in the lobster. II. Proctolin and octopamine initiate and modulate chewing." J Neurophysiol **59**(2): 551-565.
- Hume, R. I., P. A. Getting and M. A. Del Beccaro (1982). "Motor organization of Tritonia swimming. I. Quantitative analysis of swim behavior and flexion neuron firing patterns." J Neurophysiol **47**(1): 60-74.
- Hutchins, E. (1995). Cognition in the wild. Cambridge, MA, MIT Press.
- Ijspeert, A. J. (2000). "A 3-D biomechanical model of the salamander." Virtual Worlds **1834**: 225-234.
- Ijspeert, A. J. (2014). "Biorobotics: Using robots to emulate and investigate agile locomotion." Science **346**(6206): 196-203.

- Jensen, D. D. (1959). "A theory of the behaviour of *Paramecium aurelia* and the behavioral effects of feeding, fission and ultraviolet microbeam irradiation." Behaviour **15**(1): 82-122.
- Jordan, L., J. Liu, P. Hedlund, T. Akay and K. Pearson, G., (2008). "Descending command systems for the initiation of locomotion in mammals." Brain Research Reviews **57**(1): 183 - 191.
- Kalman, R. E. (1960). "A new approach to linear filtering and prediction problems." Transactions ASME Journal of Basic Engineering **82**: 35--45.
- Kennedy, D. and W. J. Davis (1977). Organization of Invertebrate Motor Systems. Handbook of Physiology. S. R. K. Geiger, E. R., Brookhart, J. M., Montcastle, V. B. Bethesda, MD, American Physiological Society. **1, part 2**: 1023-1087.
- Kleinfeld, D., E. Ahissar and M. E. Diamond (2006). "Active sensation: insights from the rodent vibrissa sensorimotor system." Curr Opin Neurobiol **16**(4): 435-444.
- Koenig, N. and A. Howard (2004). Design and use paradigms for Gazebo, an open-source multi-robot simulator. Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on.
- Krichmar, J. (2008). "Neurorobotics." Scholarpedia **3**(3): 1365.
- Krichmar, J. and H. Wagatsuma, Eds. (2011). Neuromorphic and Brain-based Robots. Cambridge, UK, Cambridge University Press.
- Leonard, J. J. and H. F. Durrant-Whyte (1991). Simultaneous map building and localization for an autonomous mobile robot, Ieee.
- Leutgeb, J. K., S. Leutgeb, M.-B. Moser and E. I. Moser (2007). "Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus." Science **315**: 961-966.
- Lisman, J. E. (2005). "The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme." Hippocampus **15**(12631759488508144345related:2aqyMFQGTa8J): 913-922.
- Lisman, J. E. (2007). "Role of the dual entorhinal inputs to hippocampus: a hypothesis based on cue/action (non-self/self) couplets." Prog Brain Res **163**: 615--625.
- Lungarella, M., T. Pegors, D. Bulwinkle and O. Sporns (2005). "Methods for quantifying the informational structure of sensory and motor data." Neuroinformatics **3**(3): 243-262.
- Mataric, M. J. (1997). "Behavior-based control: examples from navigation, learning, and group behavior." Experimental and Theoretical Artificial Intelligence, special issue on Software Architectures for Physical Agents **9**(2-3): 323-336.
- Mathews, Z., Lechón M, Calvo JB, Dhir A, Duff A, Badia SBi and V. PFMJ (2009). Insect-Like Mapless Navigation Based on Head Direction Cells and Contextual Learning Using Chemo-Visual Sensors. The 2009 IEEE/RSJ International Conference on Intelligent RObots and Systems IROS, IEEE.
- Mathews, Z., P. F. M. J. Verschure and S. Bermudez i Badia (2010). An insect-based method for learning landmark reliability using expectation reinforcement in dynamic environments. Robotics and Automation (ICRA), 2010 IEEE International Conference on.
- Matsushima, T. and S. Grillner (1992). "Neural mechanisms of intersegmental coordination in lamprey: local excitability changes modify the phase coupling along the spinal cord." J Neurophysiol **67**(2): 373-388.
- Mcgeer, T. (1990). "Passive Dynamic Walking." International Journal of Robotics Research **9**(2): 62-82.
- Michel, O. and J.-C. Heudin (1998). Webots: Symbiosis Between Virtual and Real Mobile Robots Virtual Worlds, Springer Berlin / Heidelberg. **1434**: 254-263.

- Michelsen, A., A. V. Popov and B. Lewis (1994). "Physics of directional hearing in the cricket *Gryllus bimaculatus*." Journal of Comparative Physiology A **175**(2): 153-164.
- Milford, M. and G. Wyeth (2010). "Persistent navigation and mapping using a biologically inspired SLAM system." The International Journal of Robotics Research **29**(9): 1131-1153.
- Mitchinson, B., E. Arabzadeh, M. E. Diamond and T. J. Prescott (2008). "Spike-timing in primary sensory neurons: a model of somatosensory transduction in the rat." Biol Cybern **98**(3): 185-194.
- Mitchinson, B., K. N. Gurney, P. Redgrave, C. Melhuish, A. G. Pipe, M. Pearson, I. Gilhespy and T. J. Prescott (2004). "Empirically inspired simulated electro-mechanical model of the rat mystacial follicle-sinus complex." Proc Biol Sci **271**(1556): 2509-2516.
- Mitchinson, B., C. J. Martin, R. A. Grant and T. J. Prescott (2007). "Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact." Proc Biol Sci **274**(1613): 1035-1041.
- Mitchinson, B., M. Pearson, T. Pipe and T. J. Prescott (2011). Biomimetic robots as scientific models: A view from the whisker tip. Neuromorphic and Brain-based Robots. J. Krichmar and H. Wagatsuma. Cambridge, UK, Cambridge University Press.
- Mitchinson, B., M. J. Pearson, A. G. Pipe and T. J. Prescott (2014). "Biomimetic tactile target acquisition, tracking and capture." Robotics and Autonomous Systems **62**(3): 366-375.
- Mitchinson, B. and T. J. Prescott (2013). "Whisker Movements Reveal Spatial Attention: A Unified Computational Model of Active Sensing Control in the Rat." PLoS Computational Biology **9**(9): e1003236.
- Mohd Jani, J., M. Leary, A. Subic and M. A. Gibson (2014). "A review of shape memory alloy research, applications and opportunities." Materials & Design **56**(0): 1078-1113.
- Moore, P. A. and J. Atema (1991). "Spatial Information in the Three-Dimensional Fine Structure of an Aquatic Odor Plume." The Biological Bulletin **181**(3): 408-418.
- O'Keefe, J. and D. H. Conway (1978). "Hippocampal place units in the freely moving rat: why they fire where they fire." Exp Brain Res **31**(4): 573-590.
- O'Keefe, J. and J. Dostrovsky (1971). "The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat." Brain Res **34**(1): 171-175.
- Panksepp, J. and L. Biven (2012). The Archaeology of Mind: Neuroevolutionary Origins of Human Emotions (Norton Series on Interpersonal Neurobiology), WW Norton & Company.
- Pearson, K. G. (2005). Common Principles of Motor Control in Vertebrates and Invertebrates.
- Pearson, M., C. Fox, J. C. P. Sullivan, T. J., A. G. Pipe and B. Mitchinson (2013). Simultaneous localisation and mapping on a multi-degree of freedom biomimetic whiskered robot. IEEE International Conference on Robotics and Automation (ICRA), Karlsruhe.
- Pearson, M. J., B. Mitchinson, J. C. Sullivan, A. G. Pipe and T. J. Prescott (2011). "Biomimetic vibrissal sensing for robots." Philosophical Transactions of the Royal Society B-Biological Sciences **366**(1581): 3085-3096.
- Pearson, M. J., B. Mitchinson, J. Welsby, T. Pipe and T. J. Prescott (2010). SCRATCHbot: Active tactile sensing in a whiskered mobile robot. The 11th International Conference on Simulation of Adaptive Behavior. J.-A. Meyer and A. Guillot. Paris, France.
- Pearson, M. J., A. G. Pipe, C. Melhuish, B. Mitchinson and T. J. Prescott (2007). "Whiskerbot: A robotic active touch system modeled on the rat whisker sensory system." Adaptive Behavior **15**(3): 223-240.

- Perkel, D. H. and B. Mulloney (1974). "Motor Pattern Production in Reciprocally Inhibitory Neurons Exhibiting Post-Inhibitory Rebound." Science **185**: 181-183.
- Petersen, C. C. (2007). "The functional organization of the barrel cortex." Neuron **56**(2): 339-355.
- Pfeifer, R. and J. C. Bongard (2006). How the Body Shapes the Way We Think: A New View of Intelligence (Bradford Books), The MIT Press.
- Prescott, T. J., M. E. Diamond and A. M. Wing (2011). "Active touch sensing." Philos Trans R Soc Lond B Biol Sci **366**(1581): 2989-2995.
- Prescott, T. J. and C. Ibbotson (1997). "A robot trace-maker: modeling the fossil evidence of early invertebrate behavior." Artificial Life **3**: 289-306.
- Prescott, T. J., N. Lepora and P. F. M. J. Vershure (2014). A future of living machines?: International trends and prospects in biomimetic and biohybrid systems.
- Prescott, T. J., B. Mitchinson, N. F. Lepora, S. P. Wilson, S. R. Anderson, J. Porrill, P. Dean, C. W. Fox, M. J. Pearson, J. C. Sullivan and A. G. Pipe (In press). The robot vibrissal system: Understanding mammalian sensorimotor co-ordination through biomimetics. Sensorimotor integration in the whisker system. P. Krieger and A. Groh. USA, Springer.
- Prescott, T. J., F. M. Montes Gonzalez, K. Gurney, M. D. Humphries and P. Redgrave (2006). "A robot model of the basal ganglia: behavior and intrinsic processing." Neural Netw **19**(1): 31-61.
- Prescott, T. J., M. J. Pearson, B. Mitchinson, J. C. W. Sullivan and A. G. Pipe (2009). "Whisking with robots: From rat vibrissae to biomimetic technology for active touch." Ieee Robotics & Automation Magazine **16**(3): 42-50.
- Prescott, T. J., P. Redgrave and K. N. Gurney (1999). "Layered control architectures in robots and vertebrates." Adaptive Behavior **7**(1): 99-127.
- Pyk, P., S. Bermudez, U. Bernardet, P. Knuesel, M. Carlsson, J. Gu, E. Chanie, B. Hansson, T. C. Pearce and P. F. M. J. Verschure (2006). "An artificial moth: Chemical source localization using a robot based neuronal model of moth optomotor anemotactic search." Autonomous Robots: special issue on robot olfaction **20**(3): 197-213.
- Quist, B. W., V. Seghete, L. A. Huet, T. D. Murphey and M. J. Z. Hartmann (2014). "Modeling Forces and Moments at the Base of a Rat Vibrissa during Noncontact Whisking and Whisking against an Object." The Journal of Neuroscience **34**(30): 9828-9844.
- Reed, H. L. (1952). Firing table computations on the Eniac. Proceedings of the 1952 ACM national meeting (Pittsburgh). Pittsburgh, Pennsylvania, ACM: 103-106.
- Rennó-Costa, C., J. E. Lisman and P. F. M. J. Verschure (2010). "The Mechanism of Rate Remapping in the Dentate Gyrus." Neuron **68**(6): 1051-1058.
- Rennó-Costa, C., J. E. Lisman and P. F. M. J. Verschure (2014). "A Signature of Attractor Dynamics in the CA3 Region of the Hippocampus." PLoS Comput Biol **10**(5): e1003641.
- Rennó-Costa, C., A. Luvizotto, E. Marcos, A. Duff, M. Sánchez-Fibla and P. F. M. J. Verschure (2011). Integrating Neuroscience-based Models Towards an Autonomous Biomimetic Synthetic. IEEE International Conference on Robotics and BIOMimetics (IEEE-ROBIO 2011). Phuket Island, Thailand, IEEE.
- Rosenbleuth, A. and N. Wiener (1945). "The role of models in science." Philosophy of Science **12**(4): 316-321.
- Rulkov, N. (2002). "Modeling of spiking-bursting neural behavior using two-dimensional map." Phys Rev E **65**: 041922.
- Sanchez-Fibla, M., U. Bernardet, E. Wasserman, T. Pelc, M. Mintz, J. C. Jackson, C. Lansink, C. Pennartz and P. Verschure (2010). "Allostatic control for robot behavior regulation: a comparative rodent-robot study." Advances In Complex Systems.

- Sanchez-Fibla, M., U. Bernardet, E. Wasserman, T. Pelc, M. Mintz, J. C. Jackson, C. Lansink, C. Pennartz and P. F. M. J. Verschure (2010). "Allostatic control for robot behavior regulation: a comparative rodent-robot study." Advances In Complex Systems **13**(3): 377-403.
- Searle, J. R. (1980). "Minds Brains and Programs." The Behavioral and Brain Sciences **3**: 417--457.
- Selverston, A. I. (1980). "Are central pattern generators understandable?" The Behavioral and Brain Sciences **3**: 535-571.
- Selverston, A. I. (2010). "Invertebrate central pattern generator circuits." Philos Trans R Soc Lond B Biol Sci **365**(1551): 2329-2345.
- Selverston, A. I. and J. Ayers (2006). "Oscillations and oscillatory behavior in small neural circuits." Biological Cybernetics **95**(6): 537-554.
- Selverston, A. I., D. F. Russell, D. G. King and J. P. Miller (1976). "The stomatogastric nervous system: structure and function of a small neural network." Prog Neurobiol **7**(3): 215-290.
- Seth, A. K., J. L. McKinstry, G. M. Edelman and J. L. Krichmar (2004). "Active sensing of visual and tactile stimuli by brain-based devices." International Journal of Robotics & Automation **19**(4): 222-238.
- Shilnikov, A. and N. Rulkov (2003). "Origin Of Chaos In A Two-Dimensional Map Modeling Spiking-Bursting Neural Activity." Int. J. Bifurcation Chaos **13**(11): 3325-3340.
- Stein, P. S. G. (1978). "Motor Systems, with specific reference to the control of locomotion." Ann. Rev. Neurosci. **1**: 61-81.
- Stein, P. S. G., S. Grillner, A. I. Selverston and D. G. Stuart (1997). Neurons, Networks and Motor Behavior. Cambridge, MA., MIT Press.
- Sterling, P. and J. Eyer (1988). Allostasis: A new paradigm to explain arousal pathology. Handbook of life stress, cognition, and health. S. Fisher and J. T. Reason. Chicester, NY, Wiley.
- Sünderhauf, N. and P. Protzel (2010). "Learning from Nature: Biologically Inspired Robot Navigation and SLAM—A Review." KI-Künstliche Intelligenz **24**(3): 215-221.
- Taube, J. S., R. U. Muller and J. B. Ranck, Jr. (1990). "Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations." J Neurosci **10**(2): 436-447.
- Theraulaz, G. and E. Bonabeau (1995). "Modeling the collective building of complex architectures in social insects with lattice swarms." Journal Of Theoretical Biology **177**(4): 381-400.
- Thrun, S. (2010). "Toward robotic cars." Commun. ACM **53**(4): 99-106.
- Thrun, S., W. Burgard and D. Fox (2005). Probabilistic robotics. Cambridge, MA, MIT Press
- Tikhanoff, V., A. Cangelosi, P. Fitzpatrick, G. Metta, L. Natale and F. Nori (2008). An open-source simulator for cognitive robotics research: the prototype of the iCub humanoid robot simulator. Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems. Gaithersburg, Maryland, ACM.
- Varela, F. G., H. R. Maturana and R. Uribe (1974). "Autopoiesis: the organization of living systems, its characterization and a model." Biosystems **5**(4): 187-196.
- Verschure, P. F. M. J. (1996). Connectionist Explanation: Taking positions in the Mind-Brain dilemma. Neural Networks and a New Artificial Intelligence. G. Dorffner, London: Thompson: 133--188.

- Verschure, P. F. M. J. (1998). Synthetic epistemology: the acquisition, retention, and expression of knowledge in natural and synthetic systems. Proceedings of the IEEE World Conference on Computational Intelligence, WCCI'98, Anchorage, IEEE: 147-153.
- Verschure, P. F. M. J. (2012). "The Distributed Adaptive Control Architecture of the Mind, Brain, Body Nexus." Biologically Inspired Cognitive Architecture - BICA **1**(1).
- Verschure, P. F. M. J. and A. C. C. Coolen (1991). "Adaptive Fields: Distributed representations of classically conditioned associations." Network **2**: 189-206.
- Verschure, P. F. M. J., B. Krose and R. Pfeifer (1992). "Distributed Adaptive Control: The self-organization of structured behavior." Robotics and Autonomous Systems **9**: 181--196.
- Verschure, P. F. M. J., B. Krose and R. Pfeifer (1993). "Distributed adaptive control: The self-organization of structured behavior." Robotics and Autonomous Systems **9**: 181-196.
- Verschure, P. F. M. J., C. M. A. Pennartz and G. Pezzulo (2014). "The why, what, where, when and how of goal-directed choice: neuronal and computational principles." Philosophical Transactions of the Royal Society B: Biological Sciences **369**(1655).
- Verschure, P. F. M. J. and T. Voegtlin (1999). "A bottom up approach towards the acquisition and expression of sequential representation applied to a behaving real-world device: Distributed Adaptive Control III." Neural Networks **11**: 1531--1549.
- Verschure, P. F. M. J., T. Voegtlin and R. J. Douglas (2003). "Environmentally mediated synergy between perception and behaviour in mobile robots." Nature **425**: 620--624.
- Voegtlin, T. and P. Verschure (1998). "What Can Robots Tell Us About Brains? A Synthetic Approach Towards the Study of Learning and Problem Solving." Reviews in the Neurosciences **10**(3-4): 291-310.
- Vogel, S. (1994). Life in moving fluids: the physical biology of flow. Princeton, NJ, Princeton university press.
- Voigts, J., B. Sakmann and T. Celikel (2008). "Unsupervised whisker tracking in unrestrained behaving animals." J Neurophysiol.
- Vouloutsis, V., S. Lallée and P. M. J. Verschure (2013). Modulating Behaviors Using Allostatic Control. Biomimetic and Biohybrid Systems. N. Lepora, A. Mura, H. Krapp, P. M. J. Verschure and T. Prescott, Springer Berlin Heidelberg. **8064**: 287-298.
- Walter, W. G. (1953). The Living Brain. London, Penguin Books.
- Webb, B. (2002). "Robots in invertebrate neuroscience." Nature **417**(6886): 359-363.
- Weeks, J. C. (1981). "Neuronal basis of leech swimming: separation of swim initiation, pattern generation, and intersegmental coordination by selective lesions." J Neurophysiol **45**(4): 698-723.
- Westphal, A., D. Blustein and J. Ayers (2013). A Biomimetic Neuronal Network-Based Controller for Guided Helicopter Flight. Biomimetic and Biohybrid Systems, Springer: 299-310.
- Westphal, A., N. Rulkov, J. Ayers, D. Brady and M. Hunt (2011). "Controlling a Lamprey-Based Robot with an Electronic Nervous System." Smart Structures and Systems **8**(1): 37-54.
- Witting, J., J. Ayers and K. Safak (2000). Development of a biomimetic underwater ambulatory robot: advantages of matching biomimetic control architecture with biomimetic actuators. Sensor Fusion and Decentralized Control in Robotic Systems III. G. McKee and P. Schenker. **4196**: 54-61.