The emergence of action sequences from spatial attention: insight from mammal-like robots.

Ben Mitchinson[1], Martin J Pearson[2], Anthony G Pipe[2], and Tony J Prescott[1]

[1] ATLAS Research Group, The University Of Sheffield, UK
[2] Bristol Robotics Laboratory, Bristol, UK.

Abstract. Animals display patterns of behaviour that are integrated over space and time. One plausible candidate for the decomposition of control is that behavioural sub-systems compete to command effectors. An alternative hypothesis, inspired by research on primate visual attention, is that actions are directed at a sequence of foveation targets selected using a topographic 'salience map'. In small terrestrial mammals, many behaviours are underpinned by foveation, since important effectors (teeth, tongue) are co-located with foveal sensors (microvibrissae, lips, nose), suggesting a central role for foveal selection in generating integrated behaviour. Here, we investigate control architectures for a biomimetic robot equipped with a rodent-like vibrissal tactile sensing system, explicitly comparing a salience map model for action guidance with an earlier model implementing behaviour selection. Both architectures generate life-like action sequences, but in the salience map version higher-level behavioural 'bouts' are an emergent consequence of following a shifting focus of attention.

Keywords: brain-based robotics, action selection, tactile sensing, behavioural integration, saliency map

1 Introduction

The problem of behavioural integration, or behavioural coherence, is central to the task of building life-like systems [?,?]. Living, behaving systems such as animals display patterns of behaviour that are integrated over space and time such that the organism controls its effector systems in a co-ordinated way and generates sequences of actions that serve to maintain its homeostatic equilibrium, satisfy its drives, or meet its goals. How animals achieve behavioural integration is, in general, an unsolved problem in anything other than some of the simplest invertebrates. What is clear from the perspective of behaviour is that the problem is under-constrained since similar sequences of overt behaviour can be generated by quite different underlying control architectures [?]. This implies that to understand the solution to the integration problem in any given organism is going to require investigation of mechanism in addition to observations of behaviour. In this regard, physical models—such as robots—can prove useful as a means of embodying hypotheses concerning alternative control architectures whose behavioural consequences can then be measured observationally [1]. Research with
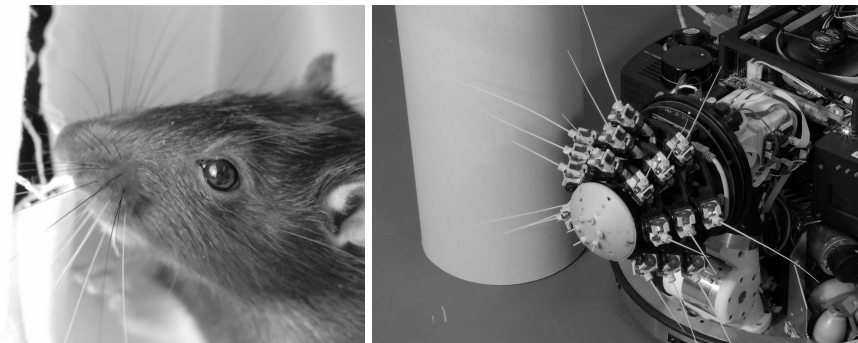
Fig. 1. (Left) A laboratory rat, showing the prominent macrovibrissae arrayed around the snout. (Right) Shrewbot, the robot used in the reported experiments, pictured alongside the cardboard cylinder that is also described. Visible are: eighteen large macrovibrissae surrounding the small microvibrissae (six fitted, here); the head on which they are mounted at the end of a 3 d.o.f. neck; the Robotino mobile platform.

robots has repeatedly demonstrated forms of emergent behaviour—the appearance of integrated behavioural sequences that are not explicitly programmed—demonstrating the value of this embodied testing for suggesting and testing candidate mechanisms.

The biological literature provides for a range of different hypotheses concerning the mechanisms that can give rise to behavioural integration; here, we highlight two, and explore and discuss their behaviour in a robot. The neuroethology literature suggests a decomposition of control into behavioural sub-systems that then compete to control the animal (see [2] for a review). This approach has been enthusiastically adopted by researchers in behaviour-based robotics as a means of generating integrated patterns of behaviour in autonomous robots that can be robust to sensory noise, or even to damage to the controller. An alternative hypothesis emerges from the literature on spatial attention, particularly regarding visual attention in primates, including humans [?]. This approach suggests that actions, such as eye movements and reaches towards targets, are generated by first computing a 'salience map' that integrates information about the relevance (salience) to the animal of particular locations in space into a single topographic representation. Some maximisation algorithm is then used to select the most salient position in space towards which action is then directed. It is usual in this literature to distinguish between the computation of the salience map, the selection of the target within the map, and orienting actions that move the animal, or its effector systems, towards the target. In the mammalian brain these different functions may be supported by distinct (though overlapping) neural mechanisms [?,?]. Of course, the approaches of behavioural competition and salience map competition are not mutually exclusive and it is possible to imagine various hierarchical schemes, whereby, for instance, a behaviour is selected first and then a point in space to which the behaviour will be directed. Alternatively,

the target location might be selected and then the action to be directed at it. Finally, parallel, interacting sub-systems may simultaneously converge on both a target and suitable action [**?**].

In our own research we have been investigating alternative architectures for controlling mammal-like robots equipped with active vibrissal sensing systems modelled on those of small rodent-like terrestrial mammals such as mice, rats, shrews and opossums. Our goal is to generate a physical model of the exploratory and orienting behaviour of these animals as measured in our laboratory using high-speed video recordings. At the same time, we seek to implement and test hypotheses about how the mammalian brain generates this behaviour. Although small mammals can grasp and manipulate objects with their forelimbs, a key effector system for foraging and object investigation is the snout, which carries multiple sensory systems and an important effector, the mouth. In the current study we investigate the hypothesis that a salience map model can be used to generate action sequences for a biomimetic robot snout mounted on a mobile robot platform, and compare this with an earlier control model based on behaviour selection. Both control systems generate life-like sequences which alternate between exploration and orienting behaviour, but in the salience map version these higher-level behavioural 'bouts' are an emergent consequence of actions determined by following a shifting focus of spatial attention (determined by a salience map) rather than being explicit control primitives. In the mammalian brain sensimotor loops involving the cortex, superior colliculus, and basal ganglia may interact to implement a control system similar to this hypothesised salience map model.

We have developed a lineage of rodent-like biomimetic robots—that is, robots which share the essential body plan of rodents, an orientable head mounted on a mobile base, and that use an array of motile whiskers (modelled on the macrovibrissae of whisking animals [3, 4]) as their primary sense [5]. These robots are biomimetic both in their morphology and in the computational and algorithmic aspects of their control architecture. Below, we describe the two models of behavioural integration in the context of the more general control architecture used by our robots. We go on to describe two experiments using the robot 'Shrewbot' (Figure 1); Shrewbot is equipped, in addition to its macrovibrissae, with an array of small, immotile, whiskers at the tip of the snout (modelled on the microvibrissae [3]).

2   Models

Figure 2 summarises the multi-level loop architecture used in our biomimetic robots, which mirrors the structure of the neural substrate [6]. We cannot represent the whole brain in our control architecture from the outset, and there is no general agreement on the function of many neural centres, or even whether description in such terms is possible. Since the robot must 'function' in some sense if we are to experiment with it, our break down of the control system into modules is by function, but the particular break down chosen is driven by the
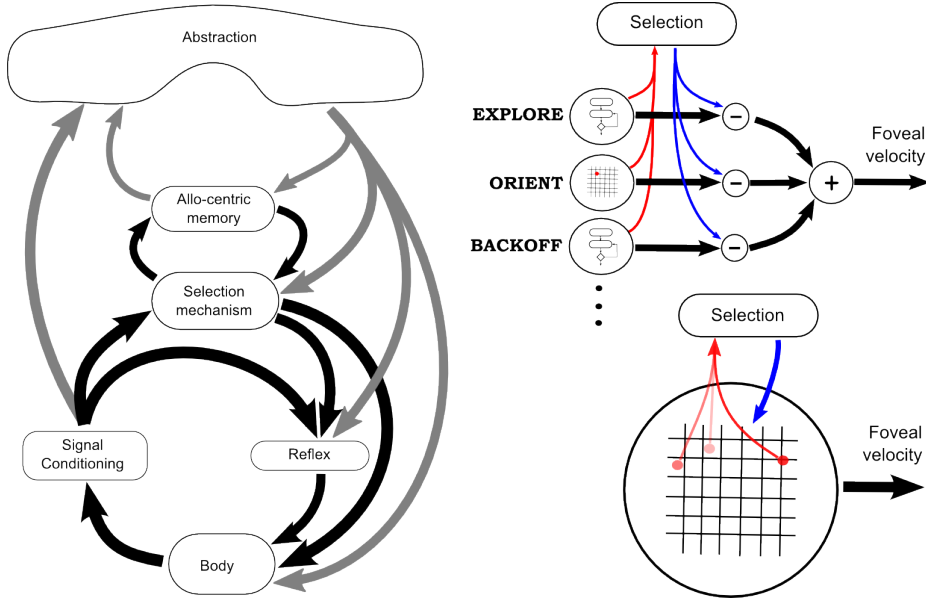
Fig. 2. (Left) Multi-level loop architecture common to the discussed models. (Right) Detail of the component labeled 'Selection Mechanism' in the left-hand panel, for **BehavSel** (above) and **AttenSel** (below).

anatomy. This places us in a strong position to hypothesise relationships between structure and function in the neural system, and these hypotheses are a major outcome of our robot work [1]. Here, we focus on the component 'Selection Mechanism' which is responsible for the majority of movements of the robot's body (neck and wheels). Below this system, low-level reflex loops effect rapid responses to current conditions (for instance, whisker protraction is inhibited by contact with the environment [7]). Above this system, we are beginning to add more cognitive components that modulate selection, either by gleaning detailed information about what has been contacted from the sensory signals (in the component labelled 'Abstraction' [8]), or by retaining a memory of the robot's past spatial experience (in the component labelled 'Allo-centric memory' [9]). The details of the wider system, beyond the selection mechanism, are covered elsewhere [5].

We identify a 'tactile fovea' [3], a short distance in front of the microvibrissae; the selection mechanism drives movements of this fovea, its output being the instantaneous foveal velocity. Our focus is on the key component of behaviour in rodent-like mammals, highlighted above, of bringing the fovea to a target. For instance, when faced with a task of discriminating between multiple objects, rat behaviour can be described as foveation to each discriminandum in sequence [3]. Beyond that it supports this foveation, the movement of the remaining nodes of the animal/robot is unconstrained, so these nodes (neck joints, body) are slaved

to the fovea in our models—that is, our robots are 'led by the nose'. This is, of course, a simplification of biological behaviour, though we have been surprised by how life-like (and practical) the resulting behaviour is. Our robots do not have manipulators, so that foveation is the only behaviour that they express: the remainder of this section compares two approaches we have taken to generating the foveal velocity vector.

The first model [10, 11, 7] of selection is an abstraction of the biology into individual behaviours—that is, it is biomimetic at the ethological level. It consists of a list of pre-defined behaviours which are arbitrated by a model of Basal Ganglia (BG), after a previous model of robot foraging [12], see Figure 2. The BG chooses one behaviour to take control of the motor plant (i.e. provide the foveal velocity) at any one time. The individual behaviours themselves are responsible for 'bidding' for this privilege, with an intensity that reflects the degree to which they are appropriate given the immediate conditions. Note that this model has minimal memory—its response, modulation from higher systems aside, is a function only of its recent inputs (and endogenous noise). Thus, this model fits neatly into the paradigm of 'behaviour-based robotics'; we denote it as **BehavSel**.

Whilst, in principle, a list of any number of behaviours could be included, in practice we have only ever implemented four (EXPLORE, ORIENT, BACKOFF, GOTO). This was based on need—we found that we were able to simulate rat behaviour to a sufficiently good approximation for our experiments with only these (in fact, the heuristic behaviour GOTO was included as an experimental convenience, rather than to correspond to any particular animal behaviour, and is not described here). Each behaviour in **BehavSel** is a 'fixed action pattern' in the sense that, once initiated, it will usually complete [12]. A higher priority behaviour can interrupt an ongoing, lower priority, behaviour. For instance, detection of an obstacle causes EXPLORE to be interrupted by ORIENT. Furthermore, some of the behaviours are parametrised such that immediate sensory information can change their details (owing to this, the term 'modal action pattern' may be preferred by some ethologists [12]). For instance, ORIENT directs the fovea to a recently-detected item in the environment.

In an early form of this model, the selection mechanism was the only source of foveal control. Since arbitration takes time [12], and processing resources on the robot are limited, this system did not prove sufficiently fast to protect the robot from damage when it foveated inaccurately. For this reason, we added a low-level reflex such that strong contact on the microvibrissae leads the neck to immediately and rapidly concertina, retracting the snout and protecting it from damage in the period of a few tens of milliseconds before the selection mechanism is able to respond to the same signals. This approach has been very satisfactory, both because we no longer damage microvibrissae and because it does not interfere at all in our investigations of the broader behaviour of the system (that is, we can investigate selection without thinking about the influence of this reflex).

The long-term behaviour of **BehavSel** can be summarised, as follows. Without stimulus, EXPLORE has the highest salience, and the robot proceeds for-

ward sweeping its fovea from side to side. On contacting a stimulus, the salience of ORIENT is raised, and the robot foveates the point of contact. Meanwhile, the salience of BACKOFF is raised somewhat, so that when the ORIENT action completes, BACKOFF is chosen and the robot moves away from the contacted obstacle, before falling back into EXPLORE. With this arrangement, the robot can be safely left to freely explore a simple environment, so long as the geometry of the environment is fairly smooth (the robot has no reversing cameras).

One question raised by the **BehavSel** model, from a biomimetic standpoint is: where are the EXPLORE and BACKOFF behaviours implemented in the brain? Behaviours, after all, are things we observe from outside the animal, and not necessarily things that are explicit in any algorithm. In the field of bio-inspired robotics, being unable to answer this question is not problematic. In biomimetic robotics, however, communicability between biological data (with the exception of ethological data) and the design of and results from models suffers if they do not share the same encoding and algorithms. As discussed, all of the behaviours in **BehavSel** are expressed through foveation. (Visual) foveation in primates is well studied and is mediated by the Superior Colliculus (SC) [13]. In rats, stimulation of SC can evoke not only eye movements [14], but also orienting-like movements of the snout, circling, and even locomotion, amongst other behaviours [15]. Inspired by these facts, we have developed a second model of foveal velocity vector generation that mirrors the features of SC—that is, a topographic saliency map driven by sensory input and modulated by information from mid- and upper-brain, with a simple motor output transform that drives foveation to the most salient region of local space. Selection, then, is between foveation targets in local space, rather than between behaviours (see Figure 2). In the case of our robots, salience is excited by whisker contact and endogenous noise and suppressed by a top-down 'inhibition-of-return' (IOR) signal from an allo-centric memory component which lowers the salience of regions that have recently been foveated (after a related model of IOR in the primate visual system [16]). This model is denoted **AttenSel**.

The long-term behaviour of **AttenSel** can be summarised, as follows. Initially, salience is driven only by endogenous noise, and the robot foveates stochastically. The noise is spatially-biased, such that foveations in a forward direction are more likely, and the robot tends to proceed forward. On contacting a stimulus, the salience in the corresponding region of the map is raised, and the robot foveates the point of contact. That location then becomes less salient, owing to the IOR signal. The foveation targets selected following this, then, drive the robot away from that location. Some time after, the robot has returned to the stochastic foveation pattern observed initially. That is, the observed behaviour of the robot using **AttenSel** is quite similar to that observed using **BehavSel**—or, to put it another way, quite similar behaviours emerge from the latter model to those that are explicitly designed into the former model.
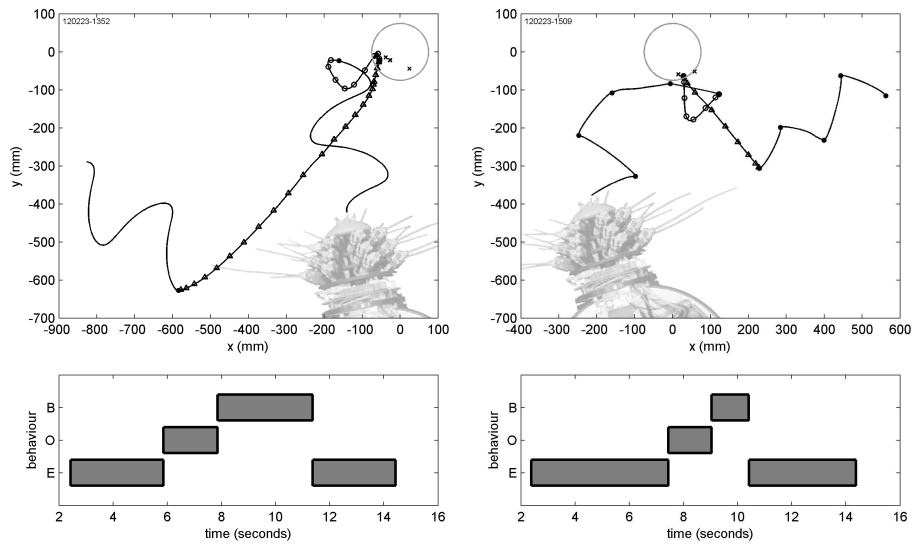
Fig. 3. Results of Experiment 1. (Upper panels) Projected onto the horizontal plane, the time course of the fovea location (black line), the moments of re-selection (black dots), and the ranges of **EXPLORE** (no markers), **ORIENT** (circles), and **BACK-OFF** (triangles) behaviours. The obstacle boundary is also marked. (Lower panels) Ethograms of the switching between behaviours against time (E, O, B, for **EXPLORE**, **ORIENT**, **BACKOFF**).

3   Results

In Experiment 1, we placed Shrewbot nearby and facing a cardboard obstacle (150mm diameter), with each of the two models in control, in turn. In each case, the robot proceeds forwards, exploring, until it contacts the obstacle with its macrovibrissae. It then orients to the obstacle. Finally, it backs away and moves off exploring in another direction. The results of each experiment are shown in Figure 3. In the case of **BehavSel**, the ethogram is recovered from recorded signals; in the case of **AttenSel**, the ethogram was generated by an observer, charged with reviewing the video *post hoc* and judging which of the three behaviours explicit in **BehavSel** was being exhibited over time.
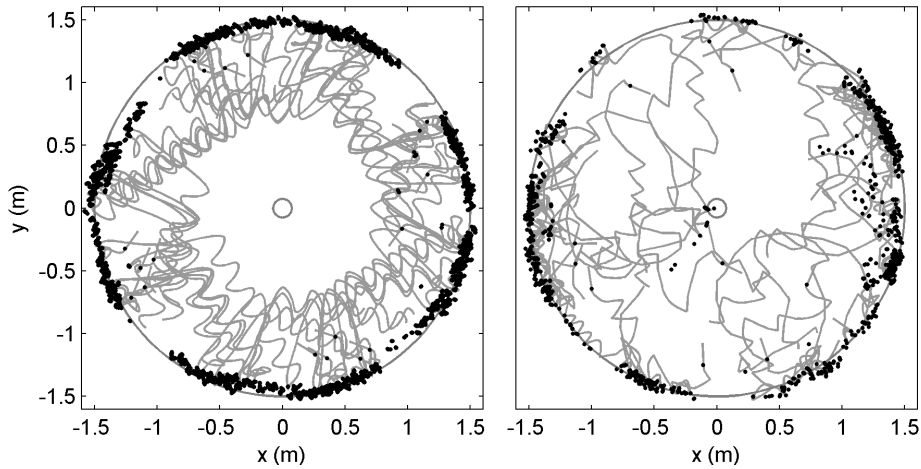

Fig. 4. Results of Experiment 2, for **BehavSel** (left) and **AttenSel** (right). Arena boundary and cardboard tube at the centre are marked as grey circles. The foveal location over time, projected onto the horizontal plane, is shown as a grey trace. Contact locations over the course of the whole experiment are marked as black dots. Some odometry drift is apparent (groups of dots not aligned with obstructions), as is some noise in the contact signal (isolated dots far away from obstructions).

In Experiment 2, we placed Shrewbot inside a circular arena (1500mm radius, bounded by 500mm high smooth vertical walls) at the centre of which was the cardboard tube from Experiment 1. The robot was then allowed to behave freely for ten minutes. The mobile platform suffers from significant odometry drift—to mitigate this for the sake of this illustration, we halted the robot every thirty seconds and re-measured its position and orientation using an overhead camera. To avoid calibration problems that arise if the whiskers are in contact with something when the robot is restarted, the robot was occasionally moved away from a wall before being restarted, which is why the foveal location traces are discontinuous in the figure. The general result of this experiment is that both

models perform a high-level behaviour that might be described as 'exploration' of the arena. The differences in the details of the exploratory behaviour expressed by the two models, visible in the figure, are not relevant to the discussion, here.

4   Discussion

The **BehavSel** and **AttenSel** models can be characterised by what is represented explicitly (or, what is 'encoded') and what is represented implicitly (or, that which 'emerges'). In **BehavSel**, behaviours are represented explicitly (and selected between) whilst 'spatial attention', as measured by the location that is foveated, emerges from the interaction of those behaviours with the physical environment. Conversely, in **AttenSel**, locations of spatial attention are represented explicitly (and selected between) whilst behaviours, as measured by a human observer, emerge from interaction with the environment. Modelling at multiple levels of abstraction is undoubtedly useful; indeed, parts of the wider control system outlined above but not described are represented using a variety of encodings. Here, we discuss the two approaches to modelling that we have explored in the context of biological research, robotic performance, and the controller design process.

Biological research The correspondence in representation between **AttenSel** (and **AttenSel**-like models) and the biology mean that inferences we draw about the relationship between structure and function in the robot control system can be transposed directly to the biological control system. Conversely, hypotheses about structure and function that arise from biological research can be tested very directly in the robot. The wider structure of both models is biomimetic at a high level (see Figure 2), but the **AttenSel** component also reflects the biology at a low level. This choice favours communicability between the model and the biological data. SC is one of the most clearly understood (and well studied) regions of the sub-cortical mammalian brain, and its most prominent feature, common to most accounts in most species, is its domination by encodings of sensory and motor space [13], though it is directly involved in mediating behaviour—or, at least, 'behavioural primitives' [15]. The **AttenSel** model of selection reflects the encoding and algorithm of SC and is directly comparable to the biological system. More generally, communicability is affected by the signal encoding used in any model. Cortex aside (and de-corticate animals can express rich behaviour [17]), the neuroscience literature is filled with descriptions of spatially and temporally encoded signals, not of behaviour generators.

Since behaviours are not specified in advance in the **AttenSel** model, the implicit behavioural space is large and continuous. In contrast, the **BehavSel** model behavioural space—parameterisation aside—is populated by a limited number of singular points. As a result, no time window of observation of the **AttenSel** model can be said to have precisely isolated 'behaviour X'. The human observer can and will pick out particular behaviours, certainly—we describe behaviour

using behavioural terms—and we have chosen to describe the behaviour of **AttenSel** in the language of **BehavSel** in Experiment 1, above. But, the behaviour of the model at any one time is chosen from that continuous implicit behavioural space. Another way of stating this is that the things we might describe as discrete behaviours—ORIENT, EXPLORE, BACKOFF—bleed into one another, so that the specific behaviour of the robot in any time window is an overlapping integration of these. A certain class of animal behavioural experiments can be summarised, methodologically, as long periods of watching the animals freely behave punctuated by brief recording opportunities where clean examples of behaviours of interest, defined in advance, are expressed [18, 19]. This methodology reflects the fact that stereotyped descriptions of the behaviour of the animal are approximations, at best. In this sense, **AttenSel** is a more complete model of foveation behaviour than **BehavSel**.

Robot performance The **AttenSel** model retains a desirable aspect of the behavioural-based robotics paradigm. That is, the immediate information about current conditions remains largely in the outside world, rather than being internalised, and the long-term behaviour represents overlapping contributions from—implicitly represented—simple behaviours. However, its encoding permits, very naturally, interaction between behaviours and the modulation of behaviour by systems that deal, explicitly, in something other than behaviour. For instance, a BACKOFF-like behaviour emerges from the interaction between endogenous noise in the saliency map and the spatial memory that the area right in front of the robot has previously been visited. A higher-level planning system that was explicitly concerned with reaching a particular location (recalling the GOTO behaviour mentioned above) could evoke locomotion from the model by modulating the salience map appropriately. This communicability between components of the system is analogous to that between the model and the biological data—the concern of all of the listed influences is space.

The continuum of possible behaviours in the **AttenSel** model represents behavioural flexibility. Speaking in terms more suited to the **BehavSel** model substrate, the use of a little bit of BACKOFF in a particular ORIENT behaviour, which in a particular situation might be advantageous, would be the result of an automatic mixing of different influences on the robots attention. To achieve the same thing in the **BehavSel** model, the model would require a specification in advance of how behaviours should mix under different conditions, an exponentially-increasing specification as the list of behaviours increases. In **AttenSel**, a specification is required, conversely, for how different influences on attention should mix—whilst this is not a trivial specification to derive either, it may be easier to derive than that for **BehavSel**, owing to the common encoding used by those influences. Thus, and in analogy to the richness of behaviour observed in animals, the **AttenSel** model might favour flexibility to the particular conditions faced by the robot in the environment, a key requirement for autonomous systems.

The design process One significant disadvantage of an **AttenSel**-like control model is that it may be difficult to get the robot to do 'the right thing'. Behaviours cannot be directly designed; rather, they emerge from the interaction between the environment and signal encodings and transforms that are specified by the designer. Thus, we cannot 'program the robot to do X'; rather, we can program how the robot generates and processes encodings, in a condition-dependent way. The design loop is closed by critiquing behaviour and modifying the way the robot computes those encodings, accordingly. Being divorced from behaviour in this way might not suit the goals of industrial robotics, for example, where 'the right thing' is tightly specified in advance. However, one of the major challenges of contemporary robotics is persuading robots to behave sensibly under conditions that are unknown *a priori*, where 'the right thing' may be, accordingly, unknown. One approach to this challenge is to use learning, applicable—in general—regardless of the choice of encoding, in the hope of side-stepping the design process to some degree. But, designing learning algorithms that are able to derive appropriate behaviour from scratch is difficult even when the task is very constrained (see [20] for several practical examples). When the task is much more general, this problem is only exacerbated by the sizes of the sensory and motor spaces that must be related—learning generally doesn't 'scale well' [20]. As Brooks notes: 'Most animals have significant behavioral expertise build in without having to explicitly learn it all from scratch.' [21]. The **AttenSel** model is an example of borrowing some of that expertise to help us to choose the structure of a complex control system. In this report, we have shown that a design process, informed by biological data, is sufficient on its own to produce behaviour as sensible as that which we designed ourselves, explicitly, in an earlier model.

References

1. Mitchinson, B., Pearson, M.J., Pipe, A.G., Prescott, T.J.: Biomimetic robots as scientific models: A view from the whisker tip. In Krichmar, J.L., Wagatsuma, H., eds.: Neuromorphic and brain-based robots. Cambridge University Press (2010) 23–57
2. Redgrave, P., Prescott, T.J., Gurney, K.: The basal ganglia: a vertebrate solution to the selection problem?. Neuroscience **89**(4) (1999) 1009–1023
3. Brecht, M., Preilowski, B., Merzenich, M.M.: Functional architecture of the mystacial vibrissae. Behavioural Brain Research **84** (1997) 81–97
4. Mitchinson, B., Grant, R.A., Arkley, K., Rankov, V., Perkon, I., Prescott, T.J.: Active vibrissal sensing in rodents and marsupials. Phil. Trans. R. Soc. B **366** (2011) 3037–3048
5. Pearson, M.J., Mitchinson, B., Sullivan, J.C., Pipe, A.G., Prescott, T.J.: Biomimetic vibrissal sensing for robots. Phil. Trans. R. Soc. B **366** (2011) 3085–3096

6. Kleinfeld, D., Berg, R.W., O'Connor, S.M.: Anatomical loops and their electrical dynamics in relation to whisking by rat. Somatosens Mot Res **16**(2) (1999) 69–88

7. Pearson, M.J., Mitchinson, B., Welsby, J., Pipe, T., Prescott, T.J.: Scratchbot: Active tactile sensing in a whiskered mobile robot. In: Proceedings of From Animals To Animats 11, Paris, 25–29 August. Volume 6226 of LNCS. (2010) 93–103

8. Sullivan, J., Mitchinson, B., Pearson, M., Evans, M., Lepora, N., Fox, C., Melhuish, C., Prescott, T.: Tactile discrimination using active whisker sensors. IEEE Sensors Journal **12**(2) (2012) 350–362

9. Fox, C., Evans, M., Lepora, N., Hamm, A., Prescott, T.: Crunchbot: A mobile whiskered robot platform. In: Proceedings of Towards Autonomous Robotic Systems (TAROS). Volume 6856 of LNCS. (2011) 102–113

10. Mitchinson, B., Pearson, M., Melhuish, C., Prescott, T.J.: A model of sensorimotor co-ordination in the rat whisker system. In: From Animals to Animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour. (2006) 77–88

11. Pearson, M.J., Pipe, A.G., Melhuish, C., Mitchinson, B., Prescott, T.J.: Whiskerbot: A robotic active touch system modeled on the rat whisker sensory system. Adaptive Behaviour **15**(3) (2007) 223–240

12. Prescott, T.J., Gonzalez, F.M.M., Gurney, K., Humphries, M.D., Redgrave, P.: A robot model of the basal ganglia: Behavior and intrinsic processing. Neural Netw **19**(1) (Jan 2006) 31–61

13. Gandhi, N.J., Katnani, H.A.: Motor functions of the superior colliculus. Annu. Rev. Neurosci. **34** (2011) 205–31

14. McHaffie, J.G., Stein, B.E.: Eye movements evoked by electrical stimulation in the superior colliculus of rats and hamsters. Brain Research **247** (1982) 243–253

15. Sahibzada, N., Dean, P., Redgrave, P.: Movements resembling orientation or avoidance elicited by electrical stimulation of the superior colliculus in rats. J Neurosci **6**(3) (1986) 723–733

16. Itti, L., Koch, C.: Computational modelling of visual attention. Nature Reviews Neuroscience **2** (2001) 194–203

17. Vanderwolf, C.H., Kolb, B., Cooley, R.K.: Behavior of the rat after removal of the neocortex and hippocampal formation. J Comp. and Phys. Psychology **92**(1) (1978) 156–175

18. Mitchinson, B., Martin, C.J., Grant, R.A., Prescott, T.J.: Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact. Royal Society Proceedings B **274**(1613) (2007) 1035–1041

19. Grant, R.A., Mitchinson, B., Fox, C.W., Prescott, T.J.: Active touch sensing in the rat: Anticipatory and regulatory control of whisker movements during surface exploration. Journal of Neurophysiology **101** (2009) 862–874

20. Kaelbling, L.P., Littman, M.L., Moore, A.W.: Reinforcement learning: A survey. J Artificial Intelligence Research **4** (1996) 237–285

21. Brooks, R.A.: A robot that walks; emergent behaviors from a carefully evolved network. Neural Computation **1** (1989) 253–262