Final report for EPSRC grant GR/R95722/01 (abridged version)

Robot Control using a Model of Central Structures in the Vertebrate Brain

Principal Investigator: Tony J. Prescott
Co-investigators: Kevin Gurney, Peter Redgrave
Report Co-author: Mark D. Humphries

Adaptive Behaviour Research Group, Department of Psychology,
University of Sheffield, Western Bank, Sheffield, S10 2TN.
http://www.shef.ac.uk/~abrg

Background

A worthwhile strategy for discovering useful designs for robot control systems is to reverse-engineer biological control systems. Research carried out by the *Adaptive Behaviour Research Group* in Sheffield has therefore focused on investigating how the vertebrate brain operates as a control architecture, in part to determine whether this can provide useful clues for the design of control systems for robots. Neurobiological research on the vertebrate brain suggests that control is distributed across multiple 'vertical' layers (cortex, forebrain, midbrain, brain-stem, spinal cord), a form of organisation that makes the brain extremely robust as lower level systems still operate when high-level ones (e.g. cortex) are damaged [1]. However, it is also possible to identify a second organising principle in the brain architecture which is the presence of a centralised (or *centrencephalic*) 'core'. This is a group of structures, principally the *basal ganglia (BG)* in the forebrain and midbrain, and the *reticular formation (RF)* in the hindbrain, that appear to play a general role in determining which actions or behaviours are selected for expression at any given time (see Figure 1a). In other words, it seems likely that these centralised components of the brain architecture are important for *action selection*, which is a fundamental problem for the control of any autonomous agent, biological or artificial.

In work prior to the current grant, our group established itself as an international leader in computational modelling of the basal ganglia, and performed some preliminary research to evaluate the effectiveness of a model basal ganglia as the action selection component of a robot control architecture (see [2] for review). Building on this earlier work the current project proposed to develop a new model of the brainstem reticular formation, improve our model of the basal ganglia, evaluate both sets of structures as possible controllers for autonomous robots, and explore the interaction between these structures in a more complete model of the brain's centrencephalic core. Our goal was not only to obtain insights into the neural control of behaviour that could lead to better designs of robot control systems, but also to obtain a better understanding of the function of these structures in normal and damaged brains. From this latter perspective we hoped that the research would also impact on the understanding on human brain disorders affecting these central brain structures such as Parkinson's disease.

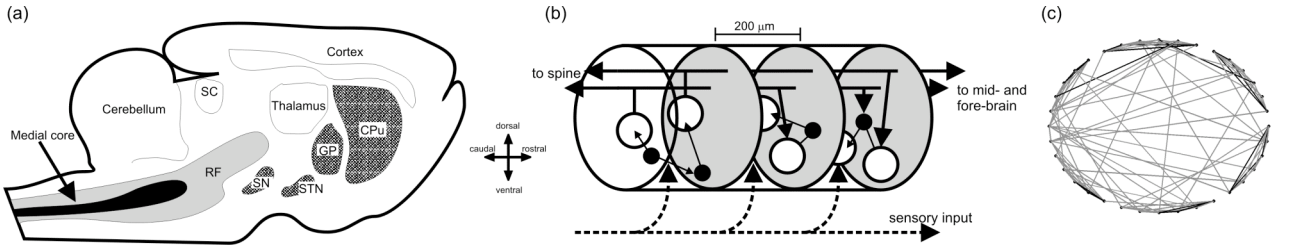

Figure 1. (a) Vertical, front-to-back, brain slice showing relative positions of the basal ganglia nuclei (hashed) and reticular formation (RF), all components of the brain's centralised core for neural action selection. (b) Proposed cluster organisation of the medial RF. (c) The network formed by the cluster organisation is a *small-world*.

Key advances and supporting methodology

Grant Objective I: To develop novel robot control architectures with embedded components based on biologically-accurate models of central structures in the vertebrate brain.

The reticular formation

The neuroscientist and cybernetics pioneer Warren McCulloch proposed that the reticular formation acts as 'command computer' which sets the global behavioural state (or 'mode') of the animal. In 1969, in a landmark paper, he and his colleagues presented a computational model of the reticular formation which demonstrated mode selection in simulation [3]. The existence of this model partly inspired the current project since it already encapsulates the idea of the reticular formation as a selection mechanism. Our research therefore began by implementing the McCulloch model and by replicating some of the original simulation experiments. This work demonstrated, however, that the original model did not have all of the selection properties claimed by authors. In particular, it did not always converge rapidly on a solution,

and could thus become stuck in seemingly endless computation. This would seem intolerable for a robot action selection system as it implies that in some circumstances no action is selected [4].

Inevitably, given its age, the McCulloch model does not capture some important aspects of reticular formation anatomy and physiology that have been reported since. Indeed, in the light of recent research, we are now able to be more specific about the region of interest: it is the fields that make up the *medial* reticular formation (mRF). In addition, the mathematical form of the McCulloch model is very different to current forms of computational neuroscience modelling, and thus it was desirable to produce an updated model in which selection capabilities could be demonstrated using more realistic neural models.

Our programme of work to create a new mRF model began with a review of this structures' internal anatomy and led us to propose a new anatomical model of the mRF as follows. The mRF is composed of two main neuron types: a medium-to-giant *projection* neuron with a bifurcating axon that projects back to the spinal cord and forward towards the thalamus (forebrain), and uses an excitatory neurotransmitter; and a small-to-medium inter-neuron which projects predominantly within the RF, and uses an inhibitory neurotransmitter. These cells are arranged in *clusters* comprising a mixture of projection and inter-neurons, in which the inter-neurons project only within the cluster, and the projection neurons only contact neurons outside the cluster. A cluster's rostral (forward) and caudal (rear) borders are thus defined by the first collateral (axonal branch) in each direction from the projection neurons' axons. The cluster organisation of the mRF is schematically depicted in Figure 1b.

We specified a statistical model which generated this organisation, given numbers of cluster and cells-per-cluster, proportions of neuron types, and probabilities of each synaptic connection type (values for each parameter being limited by ranges sourced from the literature). We then conducted an extensive exploration of the network properties of this model, to see what the structure of the mRF could imply about its function. Specifically, we principally looked at whether the network formed by the model was a *small-world* or *scale-free network*, both the subject of recent intense developments in the literature. A small-world network has two defining properties: its nodes are more clustered—that is, more locally inter-connected—than would be expected if the same number of total links were made at random; its nodes are also linked by shorter paths than would be expected if the same number of total links were made uniformly (see Figure 1c). Both of these properties are common to many real-world networks, including social networks and food webs. A scale-free network is one in which the distribution of links-per-node follows a power law, and thus has the property of scale-invariance (the distribution does not change with a scaling of the network). These networks are also found in many places, such as the World Wide Web and paper co-author networks. The existence of both these network types in such wide ranging areas suggests the existence of some common underlying organisational principles.

We showed that, to the extent our model captures the mRF's organisation, the mRF is a small-world, but not scale-free, network at the individual neuron level [5]. Moreover, we showed that this result is robust: the network properties remained the same across variation in the bounded parameters, two alternative long-range connection schemes, and two alternative algorithms for determining the structure. This has interesting implications for the putative action selection mechanism supported by the mRF. The structural properties of a small-world network imply dynamic properties—rapid cross-network synchronisation, consistent stabilisation, and persistent activity—that may all be critical to the representation of actions, and to their co-ordination. Moreover, just as we had previously argued that the basal ganglia is a central switch that minimises wiring costs [1], the small-world network also minimises wiring costs for a given component (neuron) placement when optimising the distance between components. This result also marks the crossing of a new frontier in research in computational neuroscience since it is the first demonstration of small world properties in an identified neural network in the vertebrate brain.

We next used our anatomical model to define a population-level dynamic model of the mRF, suitable for the initial exploration of its role as an action selection mechanism. Added to the model were definitions of sensory input organisation and internal weights, again both constrained by the literature. Simulations of this model were used to explore the possible levels of action representation within this structure. We concluded that the activity of projection neurons in each cluster is most likely to represent that a *sub-action* (a component of a complete action) is selected, and that multiple compatible sub-actions may be recruited together. In a simple case, we showed that this sub-action configuration resulted in reliable, correct selection across the majority of input patterns [6]. Selection failed when inputs represent signals of approximately equal salience (urgency); it is thus an intriguing possibility that more complex action selection mechanisms evolved to resolve these situations.

The basal ganglia

Further work toward understanding the role of the basal ganglia in action selection was necessary before we could explore its integration with a second action selection system supported by the mRF. Thus, concurrent with the work on the mRF, we continued to improve our existing model of the basal ganglia both by incorporating new neurophysiological data, and by further assessing its selection capabilities. This modelling work followed two lines. First, we added new connections to the model that were originally omitted due to a lack of detailed knowledge available at the time. At the same time we developed a systematic approach to assessing computational neuroscience models based on functional hypotheses. We found that the individual addition of the new connections resulted in improvements to the basal ganglia's

selection performance, and that the benefits accrued when the connections were added in combination [7, 8]. This research was complemented by a further strand of work, partly supported from grant funds, in which we developed biophysical models of single neurons within the basal ganglia, and models of large-scale basal ganglia networks based on 'leaky-integrate and fire' model neurons, with the goal of investigating to what extent more-detailed and physiologically-accurate models of these circuits possess similar selection characteristics to those observed in our more abstract models. Work at these lower levels has, to date, provided additional support to the hypothesis the basal ganglia acts as an action selection mechanism. In 2004 we published a review of our own and related basal ganglia modelling work, including the robot research described below, in the high impact journal *Trends in Neurosciences* [9].

Grant Objective II: To evaluate the effectiveness of these [neurally-inspired] architectures in generating fault-tolerant, integrated behaviour in a mobile robot, in benchmark tasks modelled on animal foraging behaviour, and using measures of appropriate selection, clean switching, and behavioural maintenance.

The reticular formation

Despite the problems identified with the original McCulloch model, the objective of our first robotic experiment was to examine its potential role as a robot controller. Our motivation here was partly to confirm that the problems observed in simulation persisted in an embodied implementation, and partly because, despite it being the authors' original intention, this model has never been evaluated as a robot controller. We used a two-resource benchmark task that we have also used to evaluate our basal ganglia model as a robot controller (see below): The robot roams freely around a flat, rectangular walled arena, whose otherwise grey floor contained two black and two white circles. When on a black circle, the robot is able to charge a potential energy store; when on a white circle, it is able to convert that stored potential energy into usable energy. The robot's energy decreases at a constant rate throughout its lifetime— when the energy reaches zero, the robot expires. Thus, the aim of the task is co-ordinate the finding and using of the two energy resources so as to maximise the life-time of the robot.

The performance of the robot controlled by the McCulloch model was compared against robots controlled using a winner-takes-all (WTA) strategy, or using random selection. Each action the robot could perform had an associated salience value calculated for it, which indicated how urgent and appropriate that action was. The WTA strategy simply picked the most salient action. We found that the McCulloch model, as originally presented, could not control selection any better than random selection, however, by fixing the structure of the model throughout the lifetime of robot, some configurations were found that improved over random selection. A genetic algorithm (GA) was then used to find the best model configurations. The GA's goal was to maximise mean energy over the life-time of robot (which was set at a maximum so that the robot did not run forever). The GA found model configurations which were as good as a WTA strategy. Interestingly, as in our earlier simulation studies, effective model configurations found using the GA did not always converge on a solution, yet these models were successful despite this limitation. Thus, we drew from this result a strong lesson in the usefulness of testing models in both simulation and embodied form: an apparent flaw in simulation may not have any effect in a real-world situation [4].

More recent work has focused on evaluating the effectiveness of the new mRF model as a robot controller. The same two-resource benchmark task was used, however, the robot's action repertoire was changed, in light of the simulation results for this new model, so that *sub-actions* were selected and co-ordinated to achieve the same goal of maximising the life-time of the robot. So, for example, the action of moving in a specific direction required the simultaneous selection of move forward and turn (to the appropriate side). Sampling from the space of possible anatomical models resulted in mRF models which could not sustain the robot's life-time much beyond the minimum possible survival time. Even using a GA to optimise the anatomical structures gave the same result. By exploring the effect of the task parameters, we found that the robot's survival time was critically dependent on the rate of energy recharge, i.e. how much energy it stored or converted per unit time spent on the appropriate coloured circle. By increasing the recharge rate by an order of magnitude (compared to that used to assess the basal ganglia model, the McCulloch model, and the WTA strategy), the GA found a small number of structures that could perform as well as the other tested controllers. However, the comparatively small number of effective structures, and the necessary large increase in energy recharge rate both suggest that the mRF's ability to perform action selection is fragile, and could break-down following changes in its own structure (due to disease or mutation) or in the environment (such as food shortages). Thus, this leaves open the intriguing possibility that more complex action selection systems, such as the basal ganglia, evolved to increase the robustness of the animal. Results of this study are expected to be published in 2006.

Basal ganglia

Our funding provided some travel money to support a collaboration with scientists at the University of Paris 6, led by Dr Jean-Arcady Meyer (funded as a visiting fellow), with whom we had already begun investigating robot implementations of computational neuroscience models. This work led to the development of the two resource problem, described above, as a benchmark test of robot action selection and also provided evidence that an embodied basal ganglia controller can act as a more effective action selection mechanism than WTA (at least in the context of this task). This research resulted in a joint article with the Paris group [10], and the collaboration will be continued as part of a recently funded EU

framework 6 project (see below). At Sheffield, we also developed and tested a revised and more biologically-accurate model of the basal ganglia and investigated its effectiveness as an action selection mechanism, against a range of metrics, as part of the controller for a Khepera™ robot engaged in a simulated foraging task [11]. Our analysis of robot behaviour, and of intrinsic processing by the basal ganglia model, allowed us to draw some interesting comparisons with the outcomes of behavioural experiments with animals, most notably with respect to (i) the role of the basal ganglia in behavioral sequencing, (ii) the activity of neurons in basal ganglia input (striatum and sub-thalamic nucleus) and output (substantia nigra) nuclei during ongoing behavior, and (iii) the behavior of animals in situations of behavioral conflict.

Grant Objective III: To investigate the trade-off between hierarchical and heterarchical control, and the benefits of layered decomposition of control, in both the vertebrate brain and in candidate robot control systems

Having developed a new and more accurate model of the mRF, and an improved basal ganglia model, we turned to the problem of how these two action selection systems might be integrated. The potential interactions cover many possibilities: a hierarchy (one tells the other what to do), a 'subsumption' style architecture (the functions of one are totally assumed by the other, creating a redundant system), a heterarchy (a 'flat' system in which both participate equally), or some combination of these. We hoped to reduce the space of possible interactions by considering the anatomical and functional connections between the two substrates of these selection mechanisms.

To lay the ground for this work, we reviewed the development of action selection systems through evolution [12]. The onset of bilateral organisation (that is, body halves) appears to have been a principle driving force for the development of a central nervous system, and thus the onset of hierarchical organisation with a distinction between peripheral effectors and central processing. Vertebrates show evidence of a hierarchy even at the lowest brain and spinal levels, however, even at the spinal level, feedback to higher brain regions rules out the possibility of a strict hierarchy of control. A comprehensive review of the known interconnections between the mRF and basal ganglia showed us that the direct connections between the two systems were sparse, and not well documented. However, we concluded that there exists a critical structure, the *pedunculopontine* (PPN) nucleus, mediating their communication. The PPN receives projections from components of the basal ganglia (most importantly, it receives projections from the output nuclei), and, in turn, projects back to some of these components, suggesting a tight integration with the basal ganglia. The PPN also projects extensively into the mRF. Thus, the anatomy is suggestive of a partial hierarchy, as shown in Figure 2. This idea of a partial hierarchy is supported by functional data from the literature on animal locomotion: stimulation of the PPN can invoke walking, in turn, driving neurons in the mRF; stimulation of basal ganglia output nuclei prevents walking from occurring, suggesting that it has top-down control of the PPN. A full model of the integrated system illustrated here is currently being assembled. This model proposes that channel-based representations of actions in the basal ganglia—as represented by the basal ganglia output nucleus SNr (substantia nigra pars reticulata)—map onto channels in the PPN. These in turn project to groups of cell clusters (c_i in Figure 2) in the mRF. These groups project to the spinal regions (and cranial nerves) which contain the circuitry necessary for the sub-actions that those clusters represent. A reduction in inhibitory SNr activity signals a selected action, which results in the disinhibition of the corresponding PPN channel. In turn, the target mRF cluster group of that PPN channel are excited, and activate the necessary spinal circuits. An open question, currently under investigation, is whether the cluster group consists of mutually compatible sub-actions (that act to further inhibit the other groups) or of mutually incompatible sub-actions, whose competition is resolved at the level of the mRF on the basis of the available sensory data. Results from this part of the project will be published in an invited article to appear in *The Philosophical Transactions of the Royal Society. B.* in late 2006.

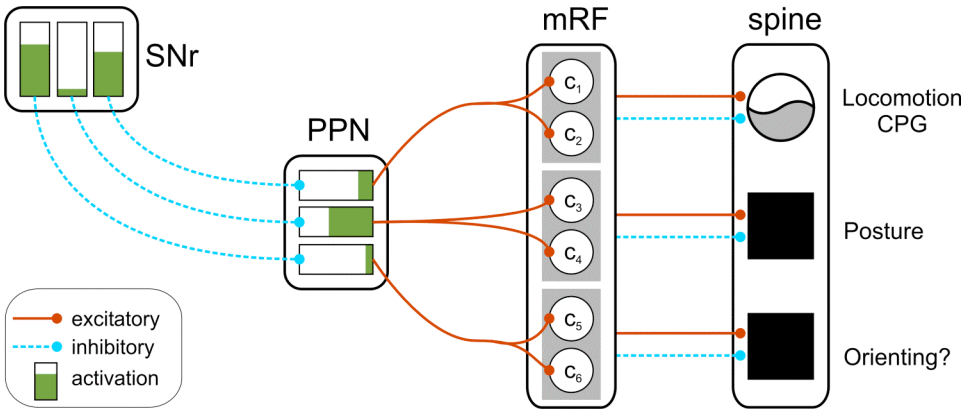

Figure 2. Schematic of a plausible integration of the basal ganglia (BG) and mRF action selection systems. SNr is a basal ganglia output nucleus which appears to mediate mRF activity via a partial hierarchy involving the PPN.

Related to this issue of layered control we also began examining the relationship between cortical and mid-brain systems involved in saccadic eye movements and developed a model that can account for a number of results in studies of primate oculomotor control [13]. Research on layered control in the vertebrate brain is to be continued through two substantial projects funded by the EPSRC and the EU Framework 6 program.

Grant Objective IV: To use embedded (robot) models to advance the understanding of central integrative systems in the vertebrate brain, and of the human brain disorders affecting these structures such as Parkinson's disease.

In order to investigate potential neuropsychological implications, we improved the capacity of our model of basal ganglia to simulate disease-like states. Damage to the basal ganglia which results in a disease state most often occurs in the dopamine system, a modulatory input to the striatum (the principal input nucleus of the basal ganglia). Dopamine depletion is associated with the onset of Parkinson's Disease, whilst other disorders, such as schizophrenia, Tourette's syndrome, Huntington's disease and ADHD are also linked with dopamine dysfunction. We improved our model of dopamine's effects in striatum by deriving a simple yet powerful high-level model from the known low-level, ion-channel based, effects of dopamine on striatal neuron membranes [14]. The new model accounted for a well-known result—that increased dopamine in turn increases the signal-to-noise ratio of a neuron's input-output encoding. The model also showed that a natural extension of this effect is that an over-abundance of dopamine leads to neurons which respond in an all-or-none fashion to an input and thus no longer reliably encode it. The effects of dopamine modulation on basal ganglia control of action selection were also investigated in our robot model engaged in a simulated foraging task [15]. In this situation dopamine depletion resulted in premature termination of behaviour, inability to select an appropriate behaviour, slowed movement and (in case of extreme depletion) akinesia. All of these outcomes are consistent with known symptoms of Parkinson's disease. In the robot model, excessive dopamine resulted in a variety of outcomes including increased 'distractibility' and inability to suppress unwanted behaviours. These results are consistent with symptoms of disorders such Huntington's disease and schizophrenia that can be partly characterised as involving excessive dopamine activity in the basal ganglia. Draft articles on these topics are currently in preparation and will be submitted to journals in 2006.

Research impact and benefits to society

The research described here has resulted in a large number of refereed articles in academic journals and conference proceedings. We have targeted outlets that emphasise research on biologically-inspired artificial systems (e.g. *Neural Networks, Adaptive Behaviour, Network),* those that will reach a primarily biological audience *(Trends in Neuroscience, Journal of Integrative Neuroscience)* and those intended for a general science audience *(Proceedings of the Royal Society)*. The work has also been disseminated by poster and/oral presentations at conferences (*International Conference on Artificial Neural Networks, Society for Neuroscience, International Meeting of the Basal Ganglia Society, British Neuroscience Association),* and focused workshops *(Artificial Rodents, Paris, July 2004; Modelling Natural Action Selection, Edinburgh, July, 2005).* The research has also been reported on by mainstream (*Focus Magazine*, October 2002, p. 74–6., "Spare parts for the brain?" *The Economist: Science & Technology Quarterly*, June 19th 2003) and special-interest media (*BioVenture View. EPSRC Newsline, AISB Quarterly*). In July of this year, Dr Tony Prescott co-organised an international workshop on "Modelling Natural Action Selection" (MNAS), which was the first ever international meeting to focus exclusively on this topic. The two-day workshop attracted ~50 participants, featured three invited speakers, and was supported by top-up funding from the BBSRC and from the EPSRC Biologically-inspired Robotics Network (BIROnet). The results from the current project were presented to that meeting. Interest in output of this work has been expressed by leading research groups, both nationally and internationally, leading to ongoing research collaborations, most notably with robotic groups in Bristol, UK (Prof. C. Melhuish); and Paris, France (Dr J-A. Meyer). We have also established a collaboration with an industrial partner, BAE systems, who is interested investigating the potential of bio-inspired action selection mechanisms in robot control and is involved in both of the research projects described next.

Partly as the result of work performed in this project the research groups has been successful in obtaining follow-up funding for two substantial multi-disciplinary projects. First, we have obtained a 5 year grant from the EPSRC 'Novel Computation' initiative entitled "Integrative computation for autonomous agents: a novel approach based on the vertebrate brain (REVERB)" worth c£1.9M (c£780K to Sheffield) with Dr K Gurney a principle investigator. This multi-centre grant brings together neuroscientists and computational modellers at Sheffield with a range of experts in other fields from seven different UK universities. Second, we have obtained a 4 year grant from European Union Framework 6 Cognitive Systems Programme, entitled "Integrating Cognition, Emotion and Autonomy (ICEA)" worth 5.8M Euros, (c850k Euros to Sheffield) with Dr T J Prescott as principal Sheffield investigator, and with Dr M. Humphries employed as a post-doctoral researcher. This consortium consists of 10 partner institutions based in seven European countries. The ICEA project will develop a cognitive systems architecture based on the architecture and physiology of the mammalian brain. In both cases the contribution of the Sheffield group will primarily be concerned with reverse-engineering of brain circuits in order to develop novel control systems for autonomous robots. Both projects will also continue the focus on layered brain architectures initiated in the current grant.

References (*indicates a refereed research result from the current grant)

1. Prescott, T.J., P. Redgrave, and K.N. Gurney, *Layered control architectures in robots and vertebrates.* Adaptive Behavior, 1999. **7**(1): p. 99-127.
2. Prescott, T.J., K. Gurney, and P. Redgrave, *Basal Ganglia*, in *The Handbook of Brain Theory and Neural Networks*, M.A. Arbib, Editor. 2002, MIT Press: Cambridge, MA. p. 147–151.
3. Kilmer, W.L., W.S. McCulloch, and J. Blum, *A model of the vertebrate central command system*. International Journal of Man-Machine Studies, 1969. **1**: p. 279-309.
4. *Humphries, M., K. Gurney, and T.J. Prescott, *Is there an integrative center in the vertebrate brain-stem? A robotic evaluation of a model of the reticular formation viewed as an action selection device.* Adapt. Behavior, 2005. **13**: p. 97-113.
5. *Humphries, M.D., K. Gurney, and T.J. Prescott, *The brainstem reticular formation is a small world, not scale-free network.* Proceedings of the Royal Society B. Biological Sciences, 2005.
6. *Humphries, M., K. Gurney, and T. Prescott. *Action selection in a macroscopic model of the brainstem reticular formation*. in *Modelling Natural Action Selection: An International Workshop*. July 2005. Edinburgh. AISB Press.
7. *Humphries, M.D., T.J. Prescott, and K.N. Gurney, *The interaction of recurrent axon collateral networks in the basal ganglia*, in *Artificial Neural Networks and Neural Information Processing - ICANN/ICONIP 2003 (Lecture Notes in Computer Science 2714)*, E.A. O. Kaynak, E. Oja & L. Xu, Editor. 2003, Springer-Verlag: Berlin. p. 797-804.
8. *Gurney, K., M.D. Humphries, R. Wood, T.J. Prescott, and P. Redgrave, *Testing computational hypotheses of brain systems function: a case study with the basal ganglia.* Network: Computation in Neural Systems, 2004. **15**: p. 263-290.
9. *Gurney, K., T.J. Prescott, J. Wickens, and P. Redgrave, *Computational models of the basal ganglia: from membranes to robots.* Trends in Neurosciences, 2004. **27**: p. 453-459.
10. *Girard, B., V. Cuzin, A. Guillot, K.N. Gurney, and T.J. Prescott, *A basal ganglia inspired model of action selection evaluated in a robotic survival task.* Journal of Integrative Neuroscience, 2003. **2**(2): p. 179-200.
11. *Prescott, T.J., F.M. Gonzalez, M.D. Humphries, K. Gurney, and P. Redgrave, *A robot model of the basal ganglia: behaviour and intrinsic processing.* Neural Networks, In Press.
12. *Prescott, T.J. *Forced moves or good tricks in design space? Great moments in the evolution of the neural substrate for action selection*. in *Modelling Natural Action Selection: An International Workshop*. July 2005. Edinburgh. AISB Press.
13. Chambers, J., M. Humphries, K.N. Gurney, and T.J. Prescott. *Mechanisms of choice in the primate brain: A quick look at positive feedback*. in *Modelling Natural Action Selection: An International Workshop*. July 2005. Edinburgh. AISB Press.
14. Humphries, M.D., *High level modeling of dopamine mechanisms in striatal neurons*. 2003, Dept. Psychol, Univ. Sheffield.
15. Prescott, T.J., K. Gurney, F.M. Gonzalez, M. Humphries, and P. Redgrave, *Action selection in a robot model of the basal ganglia: The role of simulated dopamine.* 2003: Society for Neuroscience Abstract, Program No. 915.5.