Opinion TRENDS in Neurosciences Vol.27 No.8 August 2004

Computational models of the basal ganglia: from robots to membranes

Kevin Gurney[1], Tony J. Prescott[1], Jeffery R. Wickens[2] and Peter Redgrave[1]

[1]Adaptive Behaviour Research Group, Department of Psychology, University of Sheffield, Sheffield S10 2TP, UK
[2]Department of Anatomy and Structural Biology and The Neuroscience Centre, University Otago, Dunedin, New Zealand

With the rapid accumulation of neuroscientific data comes a pressing need to develop models that can explain the computational processes performed by the basal ganglia. Relevant biological information spans a range of structural levels, from the activity of neuronal membranes to the role of the basal ganglia in overt behavioural control. This viewpoint presents a framework for understanding the aims, limitations and methods for testing of computational models across all structural levels. We identify distinct modelling strategies that can deliver important and complementary insights into the nature of problems the basal ganglia have evolved to solve, and describe methods that are used to solve them.

In recent years, an increasing number of computational models have addressed various aspects of basal ganglia function. The motivation for constructing such models derives from a pressing need to interpret the growing mountain of complex biological data associated with the basal ganglia. In the past, the qualitative information-flow ('box-and-arrow') models of microcircuitry [1], of internal connectivity between basal ganglia nuclei [2–4] and of their interactions with external structures [5–7] have been useful for interpreting a wide range of experimental data and have guided much of the recent basal ganglia research. However, the rapid accumulation of anatomical, biochemical, physiological, pharmacological and behavioural information is exposing the inadequacy of qualitative models to explain current data and predict future experimental outcomes (Figure 1). To proceed further in our understanding of the functional dynamics of information processing within the basal ganglia, and its interactions with the rest of the brain, quantitative models of all aspects of basal ganglia biology will be needed. With the expected proliferation of computational models, all claiming various degrees of biological plausibility, it will be important for experimentalists and modellers alike to appreciate the different kinds and levels of model, their underlying assumptions and limitations, how they relate to each other, and how best to validate them. The principal aim of the present viewpoint is, therefore, to offer an organising framework within which a wide spectrum of computational models of the basal ganglia can be placed.

Our proposed framework rests on two basic ideas. The first was originally articulated by David Marr [8] when he proposed that brain functions address the solution of computational problems and that these decompose into three levels of analysis: (i) 'what' is being computed and why – the computational task; (ii) 'how' the computation is carried out – the algorithm; and (iii) 'where' the computation is carried out – the implementation. The second idea is that this tri-level analysis of Marr can be applied at each of several structural levels of description [9] (Box 1). Thus, computational problems might be solved in neural components from the level of membranes to entire brain systems – there is no preferred structural level of modelling because each can deliver important computational insights.

The general applicability of this scheme will be demonstrated by discussing specific examples of recent models that deal with the analysis of computational issues at different structural levels of the basal ganglia. We will start with the highest-level systems models, where it is apparent that two different but potentially complementary modelling strategies have developed. We will then show, with examples, how the proposed framework can also help evaluate lower level microcircuit and membrane models of neural function.

System-level models

The nuclei that constitute the basal ganglia are acknowledged to form a functional sub-system within the wider brain architecture (Figure 1a). Models that have sought to understand the computational role (or roles) of the basal

Box 1. Hierarchy of structural levels for biological descriptions of the basal ganglia and their position within the brain [9]

(i) Central nervous system (whole brain)
(ii) Brain modules (e.g. basal ganglia, cerebellum, cortex and hippocampus)
(iii) Nuclei within modules (e.g. striatum, globus pallidus and substantia nigra)
(iv) Small circuits and microanatomy (e.g. mutual inhibition, convergence and divergence)
(v) Neurons and signal codes (e.g. medium spiny neurons and interneurons)
(vi) Synapses and membranes (e.g. spine and shaft membranes, presynaptic and postsynaptic membranes)
(vii) Intracellular signals (e.g. second messenger systems)

Corresponding author: Peter Redgrave (p.redgrave@sheffield.ac.uk).

 0166-2236/$ - see front matter doi:10.1016/j.tins.2004.06.003

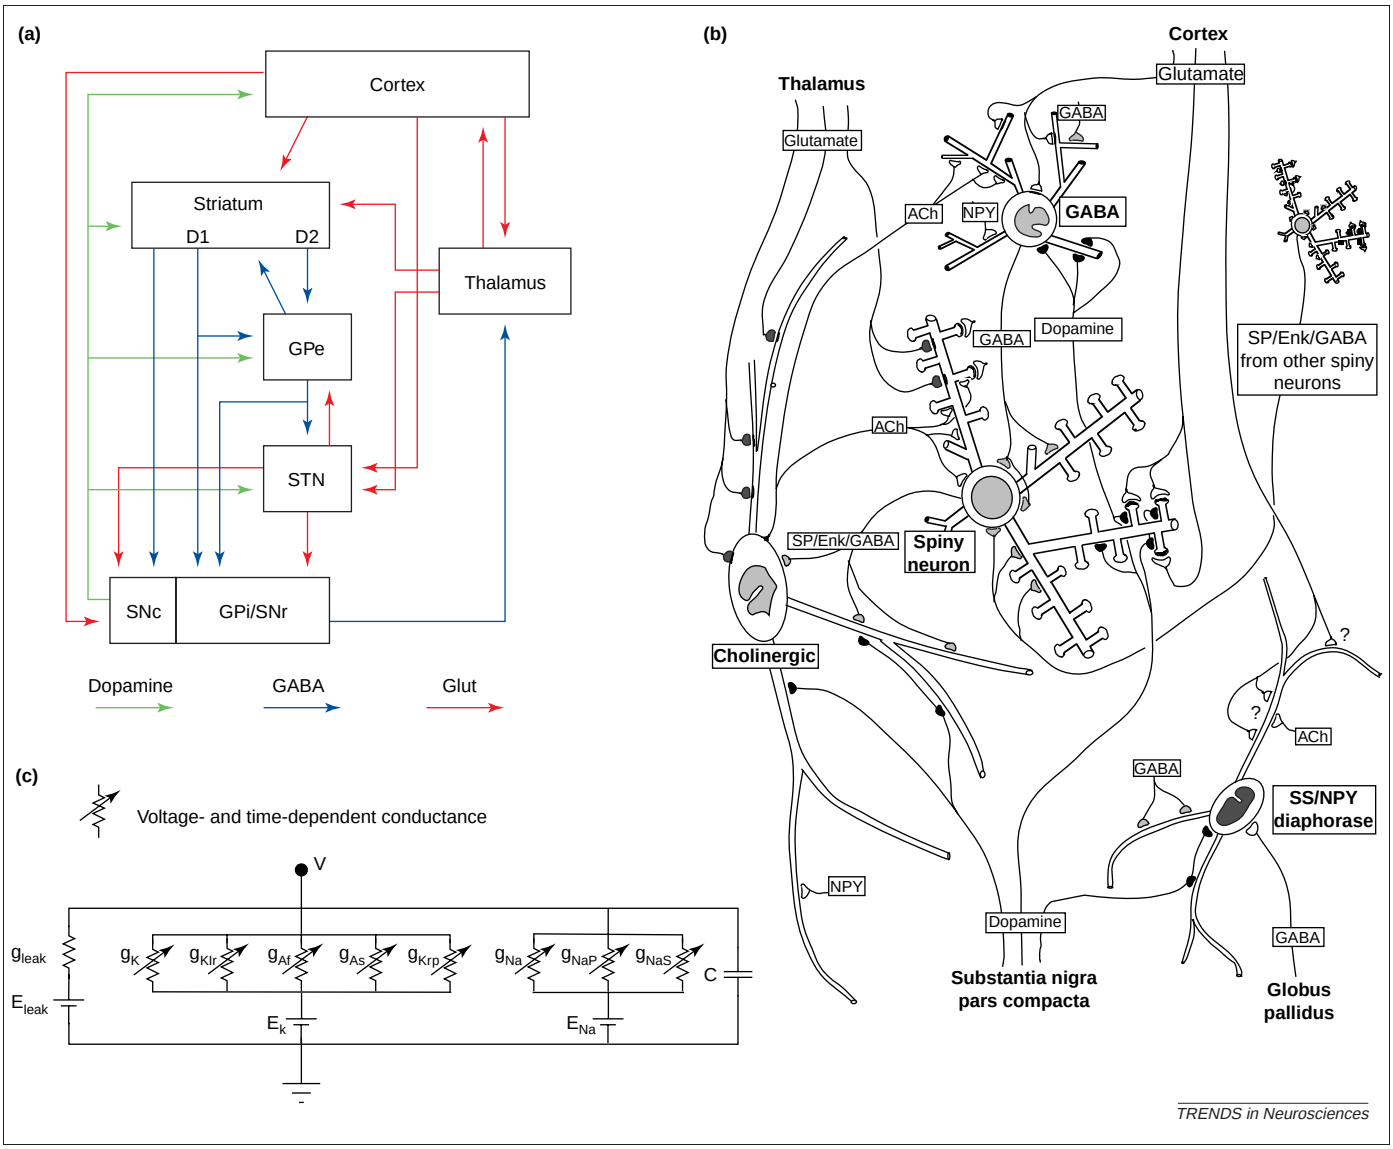

Figure 1. Biological complexity at different structural levels of the basal ganglia. **(a)** A summary box-and-arrow diagram representing currently known connections between sub-nuclei of the basal ganglia, the cerebral cortex and the thalamus. Abbreviations: D, dopamine receptors; Glut, glutamatergic connections; GPe, external globus pallidus; GPi, internal globus pallidus; SNc, substantia nigra pars compacta; SNr substantia nigra pars reticulata; STN, subthalamic nucleus. **(b)** A schematic representation of the connectivity between neuronal elements in the striatum. Reproduced, with permission, from Ref. [59] © (1995) R.G. Landes Co. **(c)** The equivalent electrical circuit of a single isopotential patch or 'compartment' of a neural membrane in which the conductances of multiple ion species are represented [53]. The complexity of all panels illustrates that there is such detail at each level that it is no longer possible to 'think' a way through the dynamic information processing of either neurons or circuits.

ganglia, taken as a whole, have generally been concerned with high-level aspects of action selection [10,11], motor sequence processing [12–15] and/or dimensionality reduction [16,17]. A common feature of many models is an adaptive capability involving some form of instrumental learning, in which nigrostriatal dopamine signals are held to play a key role. This development reflects a striking convergence between empirical studies of the short-latency dopamine response [18] and reinforcement learning methods [19], based on reward prediction error [20], that were originally developed by machine learning theorists [21]. This work has led to several prominent and well-established system-level models that have been reviewed extensively elsewhere [16,17,19,22,23]. These models are prime examples of a 'top-down' strategy of system-level modelling that we will contrast with a mainly 'bottom-up' approach adopted by us and others to investigate potential functions of the basal ganglia. To explain the differences between the two approaches, we will introduce the concepts of 'mechanism mapping' and 'mechanism mining' in the context of two contrasting examples from the recent literature – a model of reinforcement-driven dimensionality reduction presented by Bar Gad *et al.* [17] and a selection model from one of our own laboratories [10,24]

The computational hypothesis

Both top-down and bottom-up strategies usually begin with some form of computational hypothesis that can constrain the search for possible mechanisms. Hence, Bar-Gad *et al.* [17] proposed that the primary computational task of the basal ganglia is to conduct a reinforcement-driven dimensionality reduction that takes input from multiple sensory, motor, affective and cognitive sources, and relays a compressed encoding of this information to areas of the brain involved in executive planning and

selection of action (e.g. frontal cortex). This process is viewed as meeting the need for managed inflow of data into executive systems, and the requirement to encode agent–environment state information in a manner conducive to fast and effective learning of state–action mappings. By contrast, the selection model of Gurney *et al.* [10,24] is based on the hypothesis that the basal ganglia constitute a generic resource-selection mechanism [25]. Here the basal ganglia are seen as a biological solution to the problem of regulating access, by multiple functionally independent action systems, to the limited and largely shared motor resources of the 'final common motor path'. The difference in approach between these two models emerges at the algorithmic level of description.

Mechanism mapping

The Bar Gad model invokes a high-level 'procedurally transparent' algorithm – principal component analysis (PCA) – to perform the required dimension reduction. By procedurally transparent we mean here an algorithm that can be defined in general mathematical terms without any reference to an underlying neural architecture. In this case, PCA is a classical statistical method that compresses data with minimal information loss [26]. However, this model is also typical of its kind in requiring two mappings between the transparent algorithm and its (neural) implementation [27]. First, the high-level procedure (PCA) is mapped into an artificial neural network (ANN) instantiated using mechanisms such as weighted sums of inputs and Hebbian learning rules. A second mapping is then required, whereby parallels are drawn between the activity of elements in the ANN model and details of basal ganglia biology.

Validating top-down models

Use of a procedurally transparent algorithm brings the benefit of analytic clarity to understanding the input–output transformations generated by the ANN. However, the price to pay for this clarity is that the second-stage mapping – from the ANN to its hypothesized implementation in the biological substrate – is more problematic [16,28]. Neural mechanisms have to be found that fit an exact prescription defined by a largely top-down analysis. Thus, although procedurally based models can provide important insights into the kinds of operation needed to perform a high-level task, the key aspect of their validation is a successful 'mechanism mapping' from an ANN onto biological neural circuits. The extent to which this process is successful reveals whether the brain is likely to use analogous operations.

Mechanism mining

In contrast to approaches that are inspired by known mathematical methods, the model of Gurney *et al.* [10,24] has no algorithm currently defined that is procedurally transparent. Instead, it comprises a cluster of biologically constrained neuronal mechanisms that act in concert to achieve a computational objective. The methodology underlying this kind of model is based on 'mining' for potential mechanisms capable of achieving the overall goal. In the case of Gurney *et al.* [10,24], simulations and quantitative analyses demonstrated the disinhibition of appropriate output targets on the basis of relative saliences in competing input channels, thus confirming that the looped channels of the basal ganglia can act as a plausible substrate for action selection.

Sub-component functions

In addition to providing a platform on which to test a high-level computational hypothesis, system-level models with architectures tightly constrained by biological data are also in a position to provide insights into the functional properties of network sub-components. For example, a new hypothesis for the role of the feedback loop between the subthalamic nucleus and globus pallidus (Figure 1a) was an unanticipated result of the Gurney *et al.* model [10,24]. In analysis and simulation, this loop ensured that signals from the output nuclei (substantia nigra pars reticulata and entopeduncular nucleus) remained within strict operating limits, independent of the number of actively competing channels. When the loop is absent ('lesioned'), excitatory drive from the subthalamic nucleus increased with the number of active channels, ultimately overwhelming the ability of the striatum to impose selective disinhibition. This emergent property of 'capacity scaling' was entirely unpredicted and would have been difficult to discern from current biological data.

Validating bottom-up models

It is the case with all models that their output should resemble that of the corresponding biological systems. However, in contrast to the additional methods used to validate procedurally inspired models that have been already discussed here, different strategies are required to evaluate system-level models in which the model architecture is constructed directly from biological data. They can be enumerated as follows:

(i) In all current models, the biological features represented are highly selected and simplified. Thus, a powerful evaluative strategy is to add further biologically constrained detail, and re-test the ability of the model to perform the nominated computational function. In cases where functionality remains intact, or actually improves, strong support will accrue to the original computational hypothesis. Where performance deteriorates, the validity of original conjecture will be questioned. Control procedures, which comprise the addition of 'non-biological' features to the model, would be expected to impair performance. It is with these considerations in mind that the selection model of Gurney *et al.* has been subjected to further tests of 'added biological realism': the addition of basal ganglia–thalamocortical loops [29], the inclusion of collaterals from the 'direct pathway' to globus pallidus [30], the addition of lateral inhibitory connections between elements in the globus pallidus and output nuclei [31] and, more recently, the replacement of leaky integrating elements with spiking neurons. In each case, the overall selection capability of the model was found to improve. In examples where we tested the effects of adding 'non-biological' versions of connectivity, overall selection performance of the model was impaired.

(ii) A further, more stringent test of a system-level model is to embed it within the control architecture of an autonomous agent. Under these circumstances the model is exposed to more realistic sequences of sensory input via the tight dynamic coupling between the robot and its environment. Embedding within a wider system also forces consideration of how the model will contribute to overt behaviour, and how it should be interfaced with sensory and motor systems external to the basal ganglia. Again, it is with these considerations in mind that the biologically constrained model of Gurney *et al.* has been shown successfully to select the actions of a mobile robot engaged in a simulated foraging task [11,32] (Figure 2).

Summary

Two different and complementary strategies of system-level modelling of the basal ganglia have developed: a top-down strategy, where the performance of procedurally transparent algorithms are 'mapped' onto biological systems, and a more bottom-up approach where functional capabilities are 'mined' from biologically constrained architectures. We turn now to show how the analysis proposed by Marr [8] can apply equally well to computational models at lower levels of description in the structural hierarchy [9] (Box 1). These models generally deploy the bottom-up approach with the intention of discovering ('mining') potential mechanisms and computational hypotheses.

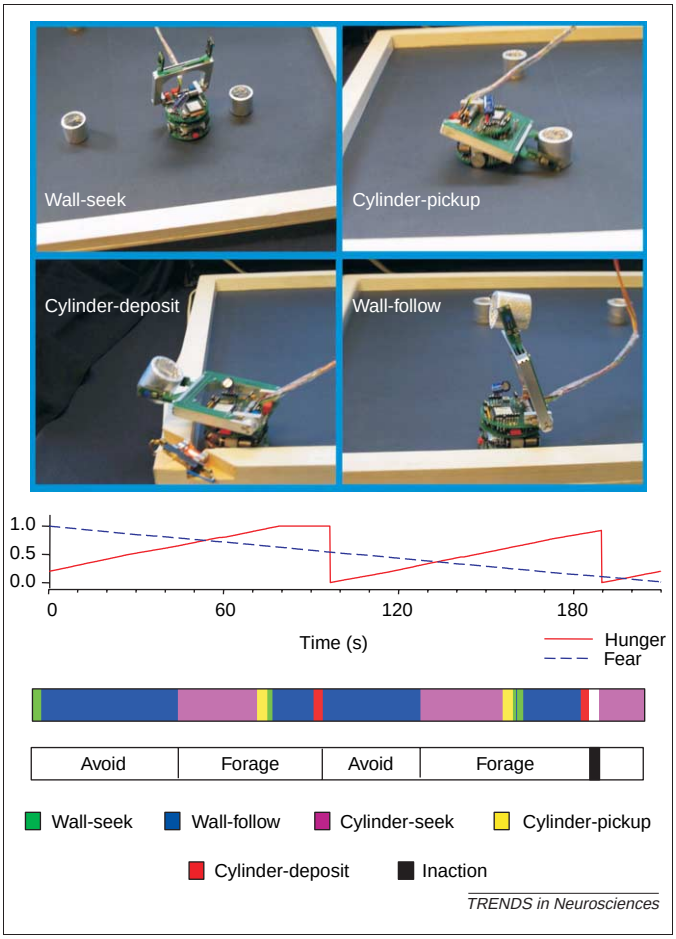

Figure 2. Action selection in a robot controlled by a biologically constrained computational model of the basal ganglia [24] including basal ganglia–thalamocortical loops [60]. The selection competition between action sub-systems (photographs) is resolved at a rate of ~8 cycles s^{-1}. Input to the model encodes the instantaneous salience of the different actions based on their perceptual and motivational affordances. The outcome of each cycle is the removal of 'nigral inhibition' from the motor command generated by the winning sub-system. The behaviour of the robot shows emergent organization both into bouts of specific behaviours (shown by the colour-coded time-bar) and into higher-order sequences of integrated behaviour (avoidance and foraging), interspersed with inactivity (coded black). The temporal organization of behaviour reflects the motivational priorities of the robot encoded by the levels of simulated 'fear' and 'hunger' – the relative intensities of which are represented in arbitrary values over time in the central graph.

Microcircuit models

There is currently little formal analytical knowledge that allows us to deduce the operations of a neural circuit from the details of its component neurons and their interconnections. Even with a complete schematic of the connectivity within a circuit, together with a full description of the physiological properties of the component neurons, we are usually unable to infer the operational properties of the circuit in question [33] (consider Figure 1b). However, by modelling the same circuit it becomes possible to examine input–output relationships distributed over whole networks, and insights about what is being computed can result. This approach is a further instance of 'mechanism mining'. In turn, the resulting mechanisms can be used as clues to the overall computation(s) being performed. This bottom-up strategy for generating circuit-level hypotheses can be seen as a powerful complementary method for validating hypotheses derived from lesion evidence or from theoretical analyses. Consequently, a range of computational models has been used to investigate the properties of microcircuits within and between basal ganglia nuclei including the striatum, globus pallidus and subthalamus [4,34–36]. To illustrate the applicability of the proposed framework to this substantial body of work we will consider the single example of intrastriatal circuitry.

Models of the striatum

Models of striatal microcircuitry were initially constrained by the morphology of spiny projection neurons and their collaterals. In general, each spiny neuron axon makes an extensive arbour of collateral branches in the region of its own dendritic tree, providing inhibitory synaptic connections with other spiny neurons [37,38]. Initial models assumed that this anatomical arrangement would produce a mechanism consisting of a lateral-inhibition-type network with strong, reciprocal connections between neurons forming the basic computational unit. Models based on this assumption exhibited a competitive dynamic such that, under high-gain conditions, the most active cells suppressed activity in their weaker neighbours [39]. In cases of perfect symmetry, such networks show stable peaks of high activity surrounded by valleys of inactivity [40]. The computational task of striatal circuits suggested by these models was consistent with the selection hypothesis already described here [25]; namely, that they perform the selection of a single action from among competing, mutually exclusive alternatives.

Subsequent experimental biology has shown, however, that inhibitory interactions among spiny projection

neurons are not as strong as assumed in these earlier models. In fact, the connections are generally sparse in nature – with probability of contact between adjacent cells much less than one, and primarily asymmetric [35,41–44]. In addition, it has become clear that inhibitory interneurons, although numerically $<1\%$ of the cells in the striatum, contribute close to 10% of the inhibitory synapses, and could thus play a significant role in feedforward inhibition [44,45]. Modelling studies based on these new data have shown that disruption of symmetry can lead to cycles of activity and travelling waves [35,46] that could play a more sophisticated role in the selection and serial organization of behaviour in normal and pathological conditions. For example, other results suggest that feedforward interneurons could have a significant part to play in initializing networks to perform extended sequences [40].

This example demonstrates that studying the dynamics of models based on the best available biological data can produce functional hypotheses that would be difficult to intuit from biological data alone. The history of microcircuit modelling of the neostriatum also shows that the usefulness of such models depends crucially on the accuracy of the biological observations from which they are derived. Detailed quantitative anatomy [47–50] is therefore invaluable in constraining possible connectivity schemes.

Conductance-based models of single neurons

To capture functionality at even lower levels of description, it is necessary to model the dynamics of individual membrane currents in compartments that represent different parts of a single neuron. Ionic current dynamics are usually captured using the Hodgkin–Huxley formalism and expressed via a set of ordinary differential equations [51]. The complexity of interactions between these currents makes it almost impossible to predict outcomes without a model (Figure 1c). A range of biophysical conductance-based models of prominent cell types in the basal ganglia have therefore been presented, including the medium spiny neurons of the striatum [34,52–54], the globus pallidus [55], and the dopamine neurons of substantia nigra pars compacta [56,57]. To illustrate how Marr's scheme can apply at the neuron level we will use as an example a short-term facilitation (STF) effect observed in striatal medium spiny neurons [58].

Short- term facilitation in striatal neurons

When two supra-threshold current pulses are applied to a medium spiny striatal cell, the second pulse is associated with a reduced time to first spike and an increase number of elicited spikes (Figure 3). To try and explain this STF effect, Mahon *et al.* [58] constructed a single-compartment conductance-based model of the medium spiny neuron incorporating several K^+ and Na^+ currents. By adopting a modelling approach, the contributions of different membrane currents to a phenomenon such as STF can be investigated using simulations that deliberately omit specific currents. Using this strategy, Mahon *et al.* [58] showed that the slowly inactivating A-current (I_{As}) was likely to be the most important mechanism for engendering this facilitation (Figure 3b). The consequent suggestion by Mahon *et al.* [58] that the function of STF in the medium spiny neuron could be to prolong the window for detecting temporally distributed inputs is a good example of how mechanism mining can be used to generate a computational hypothesis.

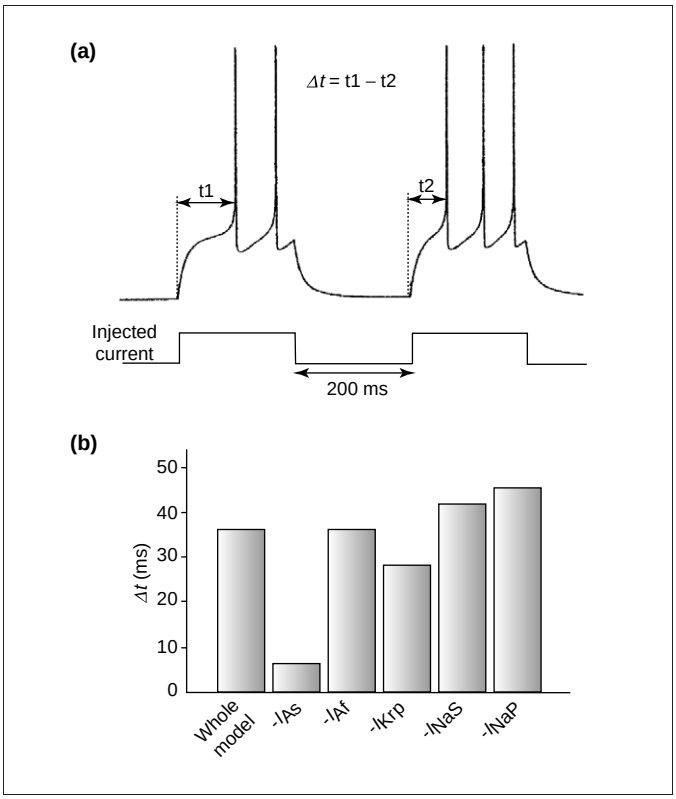

Figure 3. Simulation results from a conductance-based model demonstrating short-term facilitation (STF) in neostriatal medium spiny projection neurons. Adapted, with permission, from Ref. [53]. **(a)** Membrane behaviour of the model under current clamp comprised of two identical pulses. STF is demonstrated by a change in the time to first spike (Δt) and the number of spikes elicited per pulse. **(b)** Δt for the intact model and a series of simulations in which single currents were omitted. The slowly inactivating A-current (I_{As}) produced the most dramatic reduction in Δt, indicating that this is the most likely mechanism for inducing STF. The K^+ currents used included a slowly inactivating A-current (I_{As}), a fast A-current (I_{Af}), and a persistent current (I_{KrP}). The Na^+ currents included a slowly inactivating current (I_{NaS}) and a persistent current (I_{NaP}).

Evaluation of conductance-based models

It is a characteristic of biophysical neuron models that they are rarely inspired by any initial 'top-down' computational requirement; rather, they try to replicate the dynamic behaviour of membrane phenomena in a manner that is constrained by observed biology. In pursuing these constraints, a major problem is the requirement to find the particular configuration of multiple and interacting model parameters that best fit the biological data. This issue has been addressed recently by Wood *et al.* [54] who have developed a deterministic parameter search technique to find the maximal conductances in biophysical models and have applied this technique to modelling the membrane properties of medium spiny neurons [54].

The accurate replication of biological data can be seen as a prior step, to enable subsequent selective modulation of individual features within the model to determine relative contributions to the overall phenomenon

(see the STF case already described). Conductance-based models of this type are therefore particularly useful for: (i) providing mechanistic explanations for known membrane behaviour; (ii) through the use of virtual manipulations, positing further mechanisms that, so far, have not been observed *in vivo* or *in vitro* but that can be tested experimentally; and (iii) using the mechanisms discovered in (i) and (ii) to develop computational hypotheses about neural function.

Inter-relationships between models

Models at all levels should be evaluated, primarily, by their ability to explain observed physiological and behavioural properties of the basal ganglia. As the field develops, however, models should be expected to show increased consistency across levels. The components of higher-level models, for example, should be sufficiently realistic abstractions of their lower level counterparts. Thus, a further important use for biophysical models of basal ganglia neurons will be to inform and optimise the operating characteristics of 'simplified' neuronal elements used in higher microcircuit and system-level models. Although such analyses will lead to increased realism at higher levels, the computational hypotheses that emerge at higher levels of abstraction will provide useful starting points in the search for plausible computational mechanisms further down. Finally, it is important to note that the styles of information representation available at different levels could permit mechanisms at one level that do not translate effectively to another; for example, model neurons using firing rates cannot be expected to support mechanisms that rely on spike codes. Ultimately, this means that accurate models at the highest level might need explicitly to incorporate relatively low-level elements.

Concluding remarks

In this article we have attempted to show how the distinctions introduced by Marr [8,27] provide a useful basis for understanding the aims, limitations and strategies for testing of computational models across all structural levels of the basal ganglia (Box 1). In addition, at the system level, an important distinction has been made between models based on two-stage mapping from transparent procedures to the neural substrate and those that base their development on mining algorithms more directly from the anatomical and physiological data. The process of mapping algorithms onto the biological substrate can illuminate our understanding of the basal ganglia with powerful, tried-and-tested insights from the mathematical sciences. At the same time, this method could lead to the rethinking or refining of these procedures that generates benefits for engineering as well as neuroscience. The mechanism-mining approach, however, appears to be more readily applicable across many structural levels and, by producing models closely fitted to neurobiological constraints, will help unravel the functional mechanisms embodied in neural tissue. This approach also offers the intriguing possibility of the discovery of hitherto unknown procedures that could eventually lead to advances across all the sciences of intelligent systems. Finally, it is to be hoped that a better understanding of the various modelling approaches will encourage often-sceptical experimental neuroscientists to engage with their computational counterparts in a collaborative effort to explain better the increasingly complex datasets that describe the biological features of the basal ganglia.

Acknowledgements

This work was supported by grants from the EPSRC (GR/R95722/01 and GR/S19639/01, T.J.P. and K.G.), Wellcome Trust (068012, P.R.), Health Research Council of New Zealand and the Marsden Fund (J.R.W.)

References

1 Gerfen, C.R. and Wilson, C.J. (1996) The basal ganglia. In *Handbook of Chemical Neuroanatomy Vol 12: Integrated Systems of the CNS, Part III*. (Swanson, L.W. *et al.*, eds), pp. 371–468, Elsevier
2 Albin, R.L. (2001) End of lines and boxes. *Mov. Disord.* 16, 405–406
3 Mink, J.W. (1996) The basal ganglia: Focused selection and inhibition of competing motor programs. *Prog. Neurobiol.* 50, 381–425
4 Gillies, A. *et al.* (2002) Subthalamic–pallidal interactions are critical in determining normal and abnormal functioning of the basal ganglia. *Proc. R. Soc. Lond. B. Biol. Sci.* 269, 545–551
5 Prescott, T.J. *et al.* (1999) Layered control architectures in robots and vertebrates. *Adapt. Behav.* 7, 99–127
6 Alexander, G.E. *et al.* (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu. Rev. Neurosci.* 9, 357–381
7 Joel, D. and Weiner, I. (2000) Striatal contention scheduling and the split circuit scheme of basal ganglia–thalamocortical circuitry: from anatomy to behaviour. In *Brain Dynamics and the Striatal Complex* (Miller, R. and Wickens, J.R., eds), pp. 209–236, Harwood Academic Publishers
8 Marr, D. (1982) *Vision: A Computational Approach*, Freeman & Co
9 Churchland, P.S. and Sejnowski, T.J. (1992) *The Computational Brain*, Bradford Books
10 Gurney, K. *et al.* (2001) A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biol. Cybern.* 84, 411–423
11 Prescott, T.J. *et al.* (2002) The robot basal ganglia: action selection by an embedded model of the basal ganglia. In *Basal Ganglia VII* (Nicholson, L. and Faull, R., eds), pp. 349–356, Plenum Press
12 Beiser, D.G. and Houk, J.C. (1998) Model of cortical–basal ganglionic processing: encoding the serial order of sensory events. *J. Neurophysiol.* 79, 3168–3188
13 Berns, G.S. and Sejnowski, T.J. (1998) A computational model of how the basal ganglia produce sequences. *J. Cogn. Neurosci.* 10, 108–121
14 Suri, R.E. and Schultz, W. (1998) Learning of sequential movements by neural network model with dopamine-like reinforcement signal. *Exp. Brain Res.* 121, 350–354
15 Fukai, T. (1999) Modeling the interplay of short-term memory and the basal ganglia in sequence processing. *Neurocomputing* 26-7, 687–692
16 Joel, D. *et al.* (2002) Actor–critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Netw.* 15, 535–547
17 Bargad, I. *et al.* (2003) Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Prog. Neurobiol.* 7, 439–473
18 Schultz, W. (1998) Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27
19 Schultz, W. (2002) Getting formal with dopamine and reward. *Neuron* 36, 241–263
20 Schultz, W. and Dickinson, A. (2000) Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* 23, 473–500
21 Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning – an Introduction*, MIT Press
22 Gillies, A. and Arbuthnott, G. (2000) Computational models of the basal ganglia. *Mov. Disord.* 15, 762–770
23 Suri, R.E. (2002) TD models of reward predictive responses in dopamine neurons. *Neural Netw.* 15, 523–533
24 Gurney, K. *et al.* (2001) A computational model of action selection in

the basal ganglia. I. A new functional anatomy. *Biol. Cybern.* 84, 401–410
25 Redgrave, P. *et al.* (1999) The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 89, 1009–1023
26 Oja, E. (1982) A simplified neuron model as a principal component analyzer. *J. Math. Biol.* 15, 267–273
27 Marr, D. and Poggio, T. (1977) From understanding computation to understanding neural circuitry. *Neurosci. Res. Prog. Bull.* 15, 470–488
28 Pennartz, C.M.A. (1996) The ascending neuromodulatory systems in learning by reinforcement: comparing computational conjectures with experimental findings. *Brain Res. Rev.* 21, 219–245
29 Humphries, M.D. and Gurney, K. (2001) A pulsed neural network model of bursting in the basal ganglia. *Neural Netw.* 14, 845–863
30 Gurney, K. *et al.* (2001) 'Direct pathway' connections to globus pallidus in a computational model of the basal ganglia. *Society for Neuroscience Abstracts* (Vol. 27), program no. 514.2, Society for Neuroscience
31 Humphries, M.D. *et al.* (2003) 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)* (Kaynak, O. *et al.*, eds), pp. 797–804, Springer-Verlag
32 Girard, B. *et al.* (2003) A basal ganglia inspired model of action selection evaluated in a robotic survival task. *J. Integr. Neurosci.* 2, 179–200
33 Hopfield, J.J. and Tank, D.W. (1986) Computing with neural circuits: a model. *Science* 233, 625–633
34 Kotter, R. and Wickens, J.R. (1995) Interactions of glutamate and dopamine in a computational model of the striatum. *J. Comput. Neurosci.* 2, 195–214
35 Plenz, D. (2003) When inhibition goes incognito: feedback interaction between spiny projection neurons in striatal function. *Trends Neurosci.* 26, 436–443
36 Bevan, M.D. *et al.* (2002) Move to the rhythm: oscillations in the subthalamic nucleus–external globus pallidus network. *Trends Neurosci.* 25, 525–531
37 Wilson, C.J. and Groves, P.M. (1980) Fine structure and synaptic connections of the common spiny neuron of the rat neostriatum: a study employing intracellular inject of horseradish peroxidase. *J. Comp. Neurol.* 194, 599–615
38 Somogyi, P. *et al.* (1981) Monosynaptic cortical input and local axon collaterals of identified striatonigral neurons. a light and electron microscopic study using the golgi-peroxidase transport degeneration procedure. *J. Comp. Neurol.* 195, 567–584
39 Wickens, J.R. *et al.* (1991) Two dynamic modes of striatal function under dopaminergic–cholinergic control: Simulation and analysis of a model. *Synapse* 8, 1–12
40 Alexander, M.E. and Wickens, J.R. (1993) Analysis of striatal dynamics: the existence of two modes of behaviour. *J. Theor. Biol.* 163, 413–438
41 Czubayko, U. and Plenz, D. (2002) Fast synaptic transmission between striatal spiny projection neurons. *Proc. Natl. Acad. Sci. U. S. A.* 99, 15764–15769
42 Tunstall, M.J. *et al.* (2002) Inhibitory interactions between spiny projection neurons in the rat striatum. *J. Neurophysiol.* 88, 1263–1269
43 Taverna, S. *et al.* (2004) Direct physiological evidence for synaptic connectivity between medium-sized spiny neurons in rat nucleus accumbens *in situ*. *J. Neurophysiol.* 91, 1111–1121
44 Koos, T. and Tepper, J.M. (1999) Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nat. Neurosci.* 2, 467–472
45 Gurney, K. and Overton, P.G. A model of short and long range selective processing in neostriatum. *Neurocomputing* 58–60, 555–562
46 Wickens, J.R. *et al.* (1996) Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex *in vitro*. *Neuroscience* 70, 1–5
47 Oorschot, D.E. (1996) Total number of neurons in the neostriatal, pallidal, subthalamic, and substantia nigral nuclei of the rat basal ganglia: a stereological study using the cavalieri and optical dissector methods. *J. Comp. Neurol.* 366, 580–599
48 Bolam, J.P. *et al.* (2000) Synaptic organisation of the basal ganglia. *J. Anat.* 196, 527–542
49 Ingham, C.A. *et al.* (1998) Plasticity of synapses in the rat neostriatum after unilateral lesion of the nigrostriatal dopaminergic pathway. *J. Neurosci.* 18, 4732–4743
50 Zheng, T. and Wilson, C.J. (2002) Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. *J. Neurophysiol.* 87, 1007–1017
51 Meunier, C. and Segev, I. (2002) Playing the devil's advocate: is the Hodgkin–Huxley model useful? *Trends Neurosci.* 25, 558–568
52 Wilson, C.J. (1992) Dendritic morphology, inward rectification and the functional properties of neostriatal neurons. In *Single Neuron Computation* (McKenna, T. *et al.*, eds), pp. 141–171, Academic Press
53 Mahon, S. *et al.* (2000) Role of a striatal slowly inactivating potassium current in short-term facilitation of corticostriatal inputs: a computer simulation study. *Learn. Mem.* 7, 357–362
54 Wood, R. *et al.* A novel parameter optimisation technique for compartmental models applied to a model of a striatal medium spiny neuron. *Neurocomputing* 58–60, 1109–1116
55 Hanson, J.E. *et al.* (2004) Sodium channels and dendritic spike initiation at excitatory synapses in globus pallidus neurons. *J. Neurosci.* 24, 329–340
56 Canavier, C.C. (1999) Sodium dynamics underlying burst firing and putative mechanisms for the regulation of the firing pattern in midbrain dopamine neurons: A computational approach. *J. Comput. Neurosci.* 6, 49–69
57 Wilson, C.J. and Callaway, J.C. (2000) Coupled oscillator model of the dopaminergic neuron of the substantia nigra. *J. Neurophysiol.* 83, 3084–3100
58 Mahon, S. *et al.* (2000) Intrinsic properties of rat striatal output neurones and time-dependent facilitation of cortical inputs *in vivo*. *J. Physiol.* 527, 345–354
59 Bolam, J.P. and Bennett, B.D. (1995) Microcircuitry of the neostriatum. In *Molecular and Cellular Mechanisms of Neostriatal Function* (Ariano, M.A. and Surmeier, D.J., eds), pp. 1–19, R.G. Landes Co
60 Humphries, M.D. and Gurney, K.N. (2002) The role of intra-thalamic and thalamocortical circuits in action selection. *Network* 13, 131–156