1 General introduction

Anil K. Seth, Joanna J. Bryson, and Tony J. Prescott

Action selection is the task of deciding 'what to do next'. As a general problem facing all autonomous entities – animals and artificial agents – it has exercised both the sciences concerned with understanding the biological bases of behaviour (e.g., ethology, neurobiology, psychology) and those concerned with building artefacts (e.g., artificial intelligence, artificial life, and robotics). The problem has two parts: what constitutes an action, and how are actions selected?

This volume is dedicated to advancing our understanding of the behavioural patterns and neural substrates that support action selection in animals, including humans. Its chapters investigate a wide range of issues, including (1) whether biological action selection is optimal, and if so what is optimised; (2) the neural substrates for action selection in the vertebrate brain; (3) the role of perceptual selection in decision making, and (4) the interaction between group and individual decision making. The mechanisms of action selection considered in these contexts include abstract neural circuits (e.g., Bogacz *et al*., this volume) through specific brain systems (e.g., Stafford and Gurney, this volume) to policy choices exercised by political parties (Laver *et al*., this volume.) Taken together, this research has broad implications across the natural, social, medical, and computing sciences.

The second aim of the volume is to advance methodological practice, in particular, the practice of computational modelling. Although models cannot generate data about nature, they can generate data about theories. Complex theories can therefore be tested by comparing the outcome of simulation models against other theories in their ability to account for data drawn from nature. Models of 'natural action selection' attempt to account for transitions among different behavioural options, and a wide range of modelling methodologies are currently in use. Formal, mathematical models have been complemented with larger-scale simulations that allow investigation of systems for which analytical solutions are intractable or unknown. These include models of artificial animals (simulated agents or robots) embedded in simulated worlds (individual-based or agent-based models), as well as models of underlying neural control systems (computational neuroscience and connectionist approaches). Over recent years, work in a variety of disciplines has leveraged these and other modelling techniques, as well as analytical

Modelling Natural Action Selection, eds. Anil K. Seth, Tony J. Prescott and Joanna J. Bryson.

mathematics, to shed new light on the natural action selection. A comprehensive selection of this work is showcased in the present volume, work which explicitly addresses and integrates both parts of the action selection problem: the 'what' and the 'how'.

Because computational modelling is still a relatively new constituent of the scientific method, it is important to pinpoint exactly what the scientific contributions of computational models can be. For models of natural action selection, several general challenges can be identified: is the model sufficiently constrained by biological data that it captures interesting properties of the target natural system? Do manipulations of the model result in similar outcomes to those seen in nature? Does the model make predictions? Is there a simpler model that could account for the data equally well? Or is the model so abstract that its connection to data seems trivial? A potential pitfall of more detailed models is that they may trade biological fidelity with comprehensibility. The scientist is then left with two systems, one natural and the other artificial, neither of which is well understood. Hence, the best models are generally those that hit upon a good trade-off between accurately mimicking key properties of a target biological system, while at the same time remaining understandable to the extent that new insights into the natural world are generated.

This volume gathers together a remarkable selection of contributions from leading researchers that define the current and future landscape of modelling natural action selection. It has had a substantial gestation period. Back in 2005, we (the editors) convened a two-day meeting entitled 'Modelling Natural Action Selection' which was held at Edinburgh University and supported by the UK Biotechnology and Biological Sciences Research Council (BBSRC). The meeting attracted almost 100 participants, as well as the interest of the Royal Society, who subsequently commissioned a theme issue of their *Philosophical Transactions Series B* dedicated to the same topic, which appeared in 2007 (362:1485). The present volume updates and extends the contents of this theme issue, with the aim of providing a landmark reference in the field. Some of the chapters are revised versions of the corresponding journal papers, incorporating new results and enhanced cross-referencing. Other chapters are entirely new contributions, included in order to better reflect the changing research landscape.

The volume is divided into three sections. The first investigates the question of rational and/or optimal decision making. When an animal does one thing rather than another, it is natural to ask 'why?' A common explanation is that the action is optimal with respect to some goal. Assessing behaviour from a normative perspective has particular value when observations deviate from predictions, because we are forced to consider the origin of the apparently suboptimal behaviour. The seven chapters within this section address a wide variety of action selection problems from the perspective of optimality, and in doing so integrate a number of distinct modelling techniques including analytical optimisation, agent-based modelling combined with numerical optimisation, and Bayesian approaches.

The second section surveys a range of computational neuroscience models examining potential neural mechanisms underlying action selection. An important open question in this area is whether there are specialised neural mechanisms for action selection, or whether selective behaviour 'emerges' from the interactions among brain systems, bodies, and environments. If a dedicated neural mechanism does exist, then arguably it should exhibit the following properties: (1) it should receive inputs that signal internal and external cues relevant to decision making, (2) it should perform some calculation of urgency of salience appropriate to each action, (3) it should possess mechanisms enabling conflict resolution based on salience, and (4) it should generate outputs that allow expression of 'winning' actions while inhibiting losers. A focus of several chapters in this section is the notion that the vertebrate basal ganglia may implement a selection mechanism of this kind. Other chapters discuss alternative possibilities and related issues, including subcortical and brainstem mechanisms, hierarchical architectures, the integration of perception and action via 'affordance competition', feedback in the oculomotor system, and neural disorders of action selection.

The third and final section addresses the important topic of social action selection. In nature, action selection is rarely purely an individual matter; rather, adaptive action selection usually involves a social context. This context supplies a highly dynamic environment, in which the actions of others are generally tightly coupled to the action of the individual. Social action selection is prevalent in many species at many levels of organismal sophistication. Examples include the troop structure of primate species, nest selection by ant colonies, and patterns of voting in democratic societies. Models in this section address each of these topics and others, with a detailed focus on rigorous methodology.

In summary, the study of action selection integrates a broad range of topics including, but not limited to, neuroscience, psychology, ecology, ethology, and even political science. These domains have in common a complexity that benefits from advanced modelling techniques, exemplifying the notion of 'understanding by building'. These techniques can help answer many important questions such as: why animals, including humans, sometimes act irrationally; how damage to neural selection substrates can lead to debilitating neurological disorders; and how action selection by individuals impacts on the organisation of societies. The present volume is dedicated to presenting, integrating, and advancing research at these frontiers.

Acknowledgements

We wish to thank our editors at Cambridge University Press: Martin Griffiths and Lynette Talbot, as well as our authors and our reviewers. We are also very grateful to Dr Gillian Hayes for assistance with indexing and copyediting.

Part I

Rational and optimal decision making

Cambridge University Press
978-1-107-00049-0 - Modelling Natural Action Selection
Edited by Anil K. Seth, Tony J. Prescott and Joanna J. Bryson
Excerpt
More information

2 Introduction to Part I: rational and optimal decision making

Anil K. Seth, Tony J. Prescott, and Joanna J. Bryson

When an animal does one thing rather than another, it is natural to ask 'why?' A common explanation is that the action is optimal with respect to some goal. For example, when observing the foraging behaviour of a shorebird, one may ask whether the intake of food is being maximised. This 'normative' view, a direct extension of Darwinian principles, has its more recent roots in behavioural ecology (Krebs and Davies, 1997) and optimal foraging theory (Stephens and Krebs, 1986). Adopting a normative perspective on action selection can be very useful in placing constraints on possible underlying mechanisms, for developing and comparing theoretical frameworks relating behaviour to mechanism, and for explaining instances of apparently irrational or suboptimal behaviour. The seven chapters within this section present new insights and modelling results relevant to each of these issues. They also connect with the other parts of this book in important ways. The constraints on underlying mechanisms are thoroughly explored by the computational neuroscience models described in Part II, and patterns of both rational and irrational social behaviour are encountered in a variety of forms in Part III.

2.1 Suboptimality and 'matching'

Assessing animal behaviour from a normative perspective has particular value when observations deviate from predictions, because the scientist is now forced to consider the origin of the apparently suboptimal – or 'irrational' – behaviour. In one important example, many animals behave according to Herrnstein's (1961) 'matching law', in which responses are allocated in proportion to the reward obtained from each response. However, as both Houston *et al*. (this volume) and Seth (this volume) note, matching is not always optimal. One response to this observation is to propose that suboptimal matching arises as a side-effect of some underlying principle of behaviour, such as Thorndike's 'law of effect', which proposes that behaviour shifts towards alternatives that have higher immediate value (Thorndike, 1911). Another is given by the notion of *ecological rationality* – that cognitive mechanisms fit the demands of particular ecological niches and may deliver predictably suboptimal behaviour when operating outside these niches (Gigerenzer *et al*., 1999). In line with ecological rationality, Seth (this

Modelling Natural Action Selection, eds. Anil K. Seth, Tony J. Prescott and Joanna J. Bryson.
Published by Cambridge University Press.

volume) shows that simple decision rules that lead to optimal foraging in competitive environments with multiple foragers also lead individual foragers to obey the matching law. The remainder of Seth's chapter describes a novel methodological approach that combines agent-based modelling and optimal/normative approaches via the use of genetic algorithms. The resulting approach, which is given the label 'optimal agent-based modelling', achieves the important task of integrating function and mechanism in explanations of action selection (McNamara and Houston, 2009).

As Houston *et al*. (this volume) discuss, a further possible explanation of apparent irrationality is that we were wrong about what is being optimised. As an example, they consider *violations of transitivity*, showing that such violations can in fact be optimal when decisions are state-dependent and when choice options persist into the future. Another axiom of standard rationality is *independence from irrelevant alternatives*, the notion that the relative preference of one option over another is unaffected by the inclusion of further options into the choice set. However, as Houston *et al*. show, adding a suboptimal option can affect future expectations because, assuming that decision-making errors happen, the suboptimal option is likely to be wrongly chosen in the future. Apparent violations of rationality in human decision making are also discussed by Bogacz *et al*. (this volume), whose contribution we discuss further below.

2.2 Compromise behaviour

A related way we can be mistaken about what is optimised is when behaviour reflects compromises among multiple goals. As Crabbe (this volume) explains, a *compromise action* is defined as an action that, while not necessarily best suited for satisfying any particular goal, may be the best when all (or several) goals are taken into account. For example, a predator stalking two birds might not move directly towards either one of the birds, but in-between the two, hedging its bets in case one of the birds elects to fly away. Crabbe provides a review of approaches to compromise behaviour in relevant literatures finding that, although the intuition that compromise is useful is widespread, data from empirical studies supporting this assertion is rather thin. Crabbe then presents some detailed models directly addressing the question of when compromise behaviour is optimal. The models are restricted to simple situations involving both prescriptive (e.g., food) and proscriptive (e.g., danger areas) goals, but both are spatially explicit and open to analytical characterisation. Perhaps surprisingly, Crabbe finds that compromise behaviour is usually of little benefit in the scenarios analysed. He suggests the interesting proposal that while this may be true for so-called 'low-level' compromise behaviour (e.g., motor actions), it may be less true for 'high-level' compromise, in which the compromise is among competing behaviours (e.g., get food, find shelter) that are less easily 'blended'.

2.3 Optimal perceptual selection

Action selection can be mediated not only by motor control systems but also by perceptual systems. For example, mechanisms of selective attention can guide action selection by linking a specific motor output to one stimulus among a range of stimuli (Posner, 1980). The issue of perceptual selection is raised by a number of chapters in this collection, including Bogacz *et al*. and Stafford and Gurney (see Part II). As an example, take the problem of detecting coherent motion in a cloud of otherwise randomly moving dots. A popular solution to this problem is provided by the *leaky competing accumulator* (LCA) model (Usher and McClelland, 2001), which proposes that during decision making, noisy evidence supporting each of a range of alternatives is accumulated. Importantly, under certain conditions the LCA model can be shown to be optimal (Usher and McClelland, 2001). Bogacz *et al*. (this volume) review various neural implementations of the LCA model from the perspective of optimality, and describe extensions to this work which show that nonlinear neuronal mechanisms can perform better than linear mechanisms in terms of speed of decision making between multiple alternatives. This result raises the interesting hypothesis that nonlinearities in neuronal response functions may have evolved at least in part as a result of selective pressures favouring rapid decision making.

Bogacz *et al*. next extend the LCA framework beyond perceptual selection to account for so-called ‘value-based’ decisions in which alternatives are compared on the basis of their match to a set of internal motivations as well as to sensory signals. They discuss two examples of apparent irrationality: *risk-aversion*, where humans and other animals prefer the less risky of two alternatives that are equated for expected value, and *preference reversal*, where the preference order between two alternatives can be reversed by the introduction of a third, irrelevant, choice option. They show that the LCA model can account for these phenomena given a nonlinear utility function which is applied to the difference in value between each alternative and a ‘referent’ which may correspond to the present (or expected) state of the decision maker.

An excellent example of the cross-fertilisation of modelling strategies is evident here. In Part III of this volume, Marshall and colleagues show how the LCA model can be applied to collective decision making by colonies of social insects, via direct competition between evidence accumulating populations. And in Part II, Stafford and Gurney use a similar principle (their ‘diffusion’ model) to account for behavioural responses in a Stroop paradigm, in which responses are either congruent or incongruent with inducing stimuli.

2.4 Bayesian approaches to action selection

Bayesian theory has become extremely prominent in theoretical biology and especially neuroscience, as a means of combining new sources of information (called ‘likelihoods’) with expectations or information about the past (‘priors’). In the context of action selection, Bayesian decision theory describes how prior beliefs representing action outcomes should be combined with new knowledge (e.g., provided by a visual system) to best approach optimal outcomes as defined by a normative ‘utility’ function (Körding, 2007). Alternatively, by assuming optimal behaviour, one can use *inverse* Bayesian decision theory to identify the priors and likelihoods used by a decision-making agent. As these observations show, Bayesian decision theory is tightly integrated with normative perspectives on action selection.

In their chapter, Berniker *et al.* provide a comprehensive review of Bayesian approaches to modelling action selection, covering the fundamental issues of combining priors and likelihoods and combining multiple pieces of information. They then address the important problem of Bayesian estimation over time, in which the posterior probabilities from the past are taken to define priors for the future via Kalman filtering, with many applications in motor adaptation (Berniker and Körding, 2008). The problem of estimating environmental structure is then explored, and it is argued that such a process is an essential aspect of any action selection mechanism. Following a discussion of inverse decision theory, the chapter concludes with a discussion of Bayesian inference as a means for expressing optimal control strategies.

2.5 Sequential action selection

The final chapter in this section takes as a starting point the ubiquity of serial ordering in our daily lives. Davelaar notes that action selection mechanisms must convert simultaneously activated representations into a sequence of actions, in order to satisfy constraints of optimality. Instead of focusing on action sequences per se, Davelaar explores sequential memory retrieval, on the grounds that the wealth of modelling material in this domain can be usefully leveraged, and that there are likely to be domain-general mechanisms underlying sequential selection tasks. Davelaar compares two mechanisms: *competitive queuing* and *resampling*, finding that the latter provides a better overall match to empirical data. While both mechanisms create sequences from simultaneously activated representations, they differ in how inter-representation inhibition is applied. The success of the resampling model suggests general constraints on how multiple representations are organised for the expression of adaptive sequences. Moreover, the model is shown to account for deficits in memory selection performance in patients with Alzheimer's and Huntington's diseases.

References

Berniker, M. and K. Körding (2008). Estimating the sources of motor errors for adaptation and generalization. *Nat. Neurosci.* **11**(12): 1454–61.

Gigerenzer, G., P. Todd, and the ABC Research Group (1999). *Simple Heuristics that Make us Smart*. Cambridge: Cambridge University Press.

Herrnstein R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav*. **4**: 267–272.

Körding, K. (2007). Decision theory: what 'should' the nervous system do? *Science* **318**(5850): 606–10.

Krebs, J. and N. Davies (1997). *Behavioral Ecology: An Evolutionary Approach*. Oxford: Blackwell Publishers.