

Opinion piece



**Cite this article:** Yoo SBM, Hayden BY, Pearson JM. 2021 Continuous decisions. *Phil. Trans. R. Soc. B* **376**: 20190664. <https://doi.org/10.1098/rstb.2019.0664>

Accepted: 30 July 2020

One contribution of 17 to a theme issue ‘Existence and prevalence of economic behaviours among non-human primates’.

**Subject Areas:**  
behaviour

**Keywords:**  
open-world, continuous, embodied

**Author for correspondence:**  
Seng Bum Michael Yoo  
e-mail: sbyoo.ur.bcs@gmail.com

<sup>†</sup>Present address: Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, MA, 02139, USA

Continuous decisions

Seng Bum Michael Yoo<sup>1,†</sup>, Benjamin Yost Hayden<sup>2</sup> and John M. Pearson<sup>3</sup>

<sup>1</sup>Center for Neuroscience Imaging Research, Institute for Basic Science, Suwon, Republic of Korea, 16419

<sup>2</sup>Department of Neuroscience, Center for Neuroengineering, Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN 55455, USA

<sup>3</sup>Department of Biostatistics and Bioinformatics, Center for Cognitive Neuroscience, Department of Neurobiology, Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

BYH, 0000-0002-7678-4281

Humans and other animals evolved to make decisions that extend over time with continuous and ever-changing options. Nonetheless, the academic study of decision-making is mostly limited to the simple case of choice between two options. Here, we advocate that the study of choice should expand to include *continuous decisions*. Continuous decisions, by our definition, involve a continuum of possible responses and take place over an extended period of time during which the response is continuously subject to modification. In most continuous decisions, the range of options can fluctuate and is affected by recent responses, making consideration of reciprocal feedback between choices and the environment essential. The study of continuous decisions raises new questions, such as how abstract processes of valuation and comparison are co-implemented with action planning and execution, how we simulate the large number of possible futures our choices lead to, and how our brains employ hierarchical structure to make choices more efficiently. While microeconomic theory has proven invaluable for discrete decisions, we propose that engineering control theory may serve as a better foundation for continuous ones. And while the concept of value has proven foundational for discrete decisions, goal states and policies may prove more useful for continuous ones.

This article is part of the theme issue ‘Existence and prevalence of economic behaviours among non-human primates’.

1. Introduction

Decisions as studied in the laboratory are typically *discrete*. Consider a paradigmatic example, the Allais Paradox ([1], figure 1). The Allais Paradox involves a single choice between a pair of categorically distinct fully specified options. It is also prototypical—the vast majority of decisions studied in microeconomics, behavioural economics and neuroeconomics have the same properties [2–4]. In discrete decisions, the decision-maker has, in principle, as long as they want to decide. The result of the choice does not affect what options are made available in the future. The choice itself takes place either instantaneously or irrespective of time, and, once made, is irrevocable. The decision-maker can express their preference verbally, with a button press, or with any other response modality. Finally, the decision-maker’s choice is free of context; it is disembodied [5,6].

Many of the decisions we make in our daily lives are discrete. Examples include selecting an item from a menu in a restaurant and selecting a life insurance plan. However, many other decisions we face in our daily lives are continuous. Consider how a diner selects food bite-by-bite in a large meal, or the daily adjustments to the portfolio of an active stock market investor. Or consider the pursuit behaviour of the window-skimmer dragonfly (*Libellula luctuosa*). As it pursues its mosquito prey, the dragonfly must continuously monitor the mosquito’s position while both prey and predator move rapidly through three-dimensional space [7]. In this situation, the dragonfly is not faced with any obvious discrete decision among fixed options: rather, it must carefully regulate its wings to stabilize its body, track and extrapolate the future position of the prey and plot an intercept course. Actions at each moment come from a large



**Figure 1.** The Allais Paradox (left) is a canonical example of what we call a discrete decision. It involves a single choice between two options; the choice is independent of action, context and time. It has no influence on what options are available in the future. A dog moving along a bank preparing to jump and grab a fish (right) is a canonical example of a continuous decision. (Online version in colour.)

continuum of possibilities, each flowing into the next. Moreover, the mosquito will change its behaviour based on what the dragonfly does, and the dragonfly can exploit its knowledge of that fact to improve its foraging.

While such complex behaviour may be conceptualized as a series of discrete choices—one choice at each moment—this conceptualization is both facile and reductive. It misses some of the most important aspects of continuous decisions, such as the fact that each choice opens many new doors and closes many others. Alternatively, one could view the choice as the selection of an overarching strategy, but this just ignores the important moment-by-moment nature of continuous decisions. It is best to view the *decision* as encompassing the entire process, from the continuum of motor adjustments to the selection of flight control strategies. Of course, discrete and continuous decisions are not entirely separate categories of decisions. Instead, they represent two extremes of a spectrum of continuousness. Our central message in this manuscript is that the properties of continuous decisions (including the properties of partially continuous ones that are not shared with purely discrete decisions) are deserving of interest.

## 2. Benefits of studying continuous decisions

The greatest benefit to studying continuous decisions is that they are a type of decision our brains were evolved to make [8–11]. Focusing our research on evolutionarily valid contexts is important because it maximizes our chances of making valid discoveries. If artificial decisions are contrived, then they may, for example, produce preference patterns, including anomalies, that are irrelevant in the natural world.

For example, we and others have argued that the intertemporal choice paradigm, as typically implemented in animals, is quite unnatural and that results from animal intertemporal choice tasks may not paint an accurate picture of natural behaviour [12–17]. Another example comes from the fact that laboratory risk paradigms (including our own) tend to draw possible outcomes from a limited range of possibilities rather than from a continuous probability distribution function [18,19]. Such distributions are mathematically well-behaved, but may be less common than ones with a more continuous and contoured shape. Interestingly, this question of ethological validity applies not just to natural situations but also to markets, which have many properties in common with nature but not with the laboratory (for a compelling discussion of this topic, see [20], figure 2).

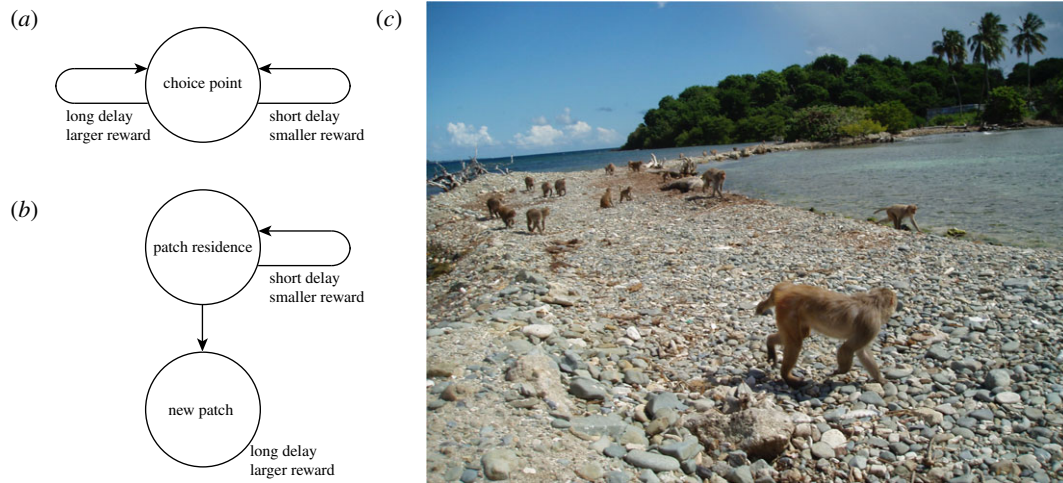
## 3. Continuous decisions are open-world

We take the term ‘open-world’ from the domain of video games [22]. Open-world games like *Grand Theft Auto* and *Minecraft* present players with virtually limitless territory to explore, and terrain filled with a multitude of possible actions. Like the situations in which many real-world decisions occur, they contain more unknowns than knowns, which reward curiosity and search, and a large number of other agents with their own goals [23,24]. Open-world environments pose a significant computational challenge, elided in nearly all laboratory decisions but central to real decisions, of narrowing down a plenitude of possible actions to a small set of ‘live’ options. To meet these challenges, decision-makers must employ cognitive abilities that make these computationally intensive problems tractable, abilities like mental maps, curiosity and heuristics. Note that ‘live’ options have much similarity with *afforded* options, which have begun to have influence in the neuroscience of decision-making, although the mechanisms by which options activate affordance representations remain poorly understood [18,25–27].

In order to delve into the implications of open-world tasks, we will adopt the language of reinforcement learning (RL) [28,29]. Notably, RL does not need to assume either a model of the world, a fixed set of choices, or even a static set of values. A learner that follows RL can learn online (it does not need to be pre-trained). Moreover, while RL has traditionally been most effective in the tabular setting—the case in which options and contexts are discrete and finite, so that the value of each action can be stored in a lookup table—its formalism is generic enough to extend to continuous decisions. Finally, recent progress in deep RL has opened a new dialogue between machine learning and cognitive neuroscience, producing artificial agents capable of solving challenging continuous tasks like video games while underperforming humans in others like imitation learning (e.g. [30–34]).

### (a) States and maps

The first distinction we want to draw between discrete and continuous tasks is in the dimension of what is called the state space. For a given problem, RL always assumes a *state*, the collection of variables needed to fully define the current context [28,35–38]. For discrete choices, the state space (the set of all possible states) usually consists of a small number of variables—perhaps attribute information for each of the choice options available and some information about current preferences.



**Figure 2.** (a) A traditional discrete choice task used to study time preferences is the intertemporal choice task (reviewed in [13]). (b) A patch-leaving task has an ostensibly isomorphic structure, except that the specific options available on the next trial are determined entirely by the choices made on the present trial [12]. This feature has major effects on choice patterns and, although the element of cross-trial dependency is ostensibly more difficult, renders behaviour nearly optimal [21]. (c) The patch-leaving task, but not the intertemporal choice task, is ethologically valid. Here, monkeys at the Cayo Santiago field station travel between their morning foraging site on Large Cay to their afternoon foraging site on Small Cay.

For the dragonfly pursuit described above, the current state might include position and velocity information for all joints, estimated position and velocity of both the dragonfly and its prey, and the relative locations of relevant obstacles. The state might also include internal variables of the organisms that might be relevant to predicting behaviour (e.g. hunger, fatigue, etc.). Consequently, the state space in open-world decisions tends to be enormous. In part, the proliferation of these states is related to the mathematical requirements of RL, which assume that the state contains all information on which behaviour at the current time depends. But this explosion of states is also inherent in the complexity of the behaviour itself and the granularity at which it is modelled.

Moreover, it is not just the number of state variables that are increased in continuous decisions, but their ranges. In discrete RL settings, for  $S$  state variables with  $N$  values each, the size of the state space is  $N^S$ . That is, for a decision problem with five yes/no contextual variables (Am I hungry? For sweet or salty snacks? Is there food close by?), there are a total of  $2^5 = 32$  possible states. In the case of continuous states, where  $N$  becomes infinite, state spaces become infinitely large and tabular approaches must fail. Moreover, in many cases of interest—and not just continuous cases—states are only partially observed, requiring some form of state inference. (For example, in games like poker, opponents' cards are key missing pieces of game state information.)

Fortunately, while practical and algorithmic challenges remain in these cases, they are addressed by existing concepts like models, maps, trajectories and policies that bridge the discrete and continuous cases. The common insight underlying these extensions of tabular RL is that, when states or actions are continuous, we can often make use of a concept of distance, such that 'nearby' states are similar and quantities like value can be assumed to change smoothly [28,39,40]. That is, continuous spaces can often, paradoxically, have fewer variables than discrete ones. For instance, if one is given a table of 100 numbers, this data structure has, apparently, 100 free parameters. But if these numbers are samples from a quadratic function at regularly spaced intervals along the  $x$ -axis, they are encapsulated by only three parameters. These sorts of

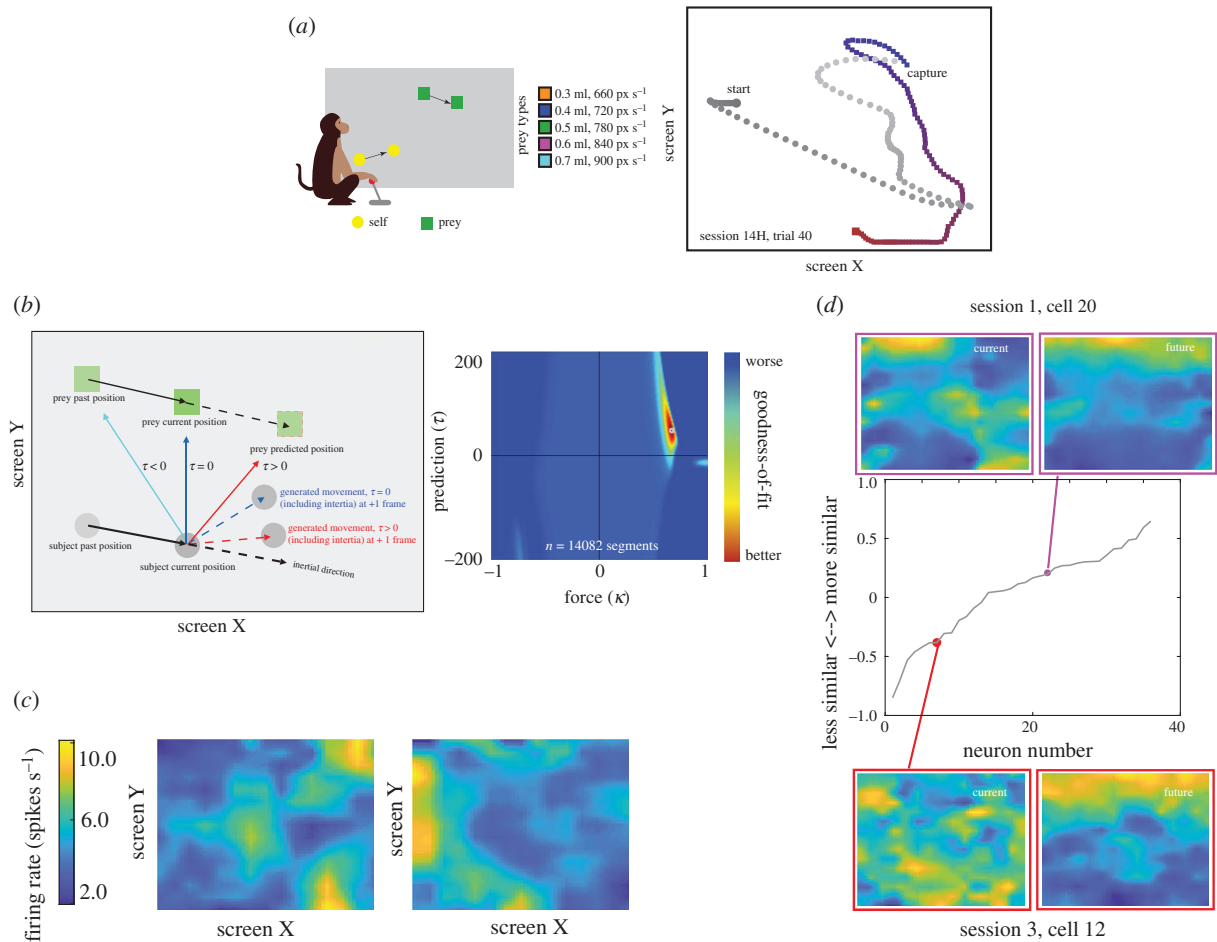
geometric and smoothness assumptions often allow one to replace tables with functions and an infinite number of individualized values with a finite number of function parameters, rendering learning feasible once again.

Along related lines, much recent work on RL has focused on the existence and utility of mental models or maps [36,38,41–43]. That work strongly supports the idea that mental maps—originally proposed by Tolman—have a distinct neural reality [44]. To give a specific example, in a simulated pursuit task, we found evidence for such maps in the firing rates of single neurons in the dorsal anterior cingulate cortex (dACC) [45]. Specifically, dACC neurons have complex, non-regular multimodal firing rate place fields, reminiscent of non-grid cells in the medial entorhinal cortex [45]. These maps in dACC (and potentially in other regions as well) likely serve to make possible the kinds of computations necessary for continuous pursuit.

What is important in these cases is that not that such maps are necessary for learning—RL agents who lack them can still eventually learn—but that these maps capture either regularities in state sequences (like random walks on an unseen graph) or structure within the states (like proportions of bird beaks and legs) that facilitate faster learning. In the language of learning theory, they constitute *inductive biases* that allow learners to make the most of each observation. More importantly, such inductive biases, by making strong assumptions about the possible types of knowledge to be acquired, combat the curse of dimensionality. In the studies above, though numbers of states were large, their dimensionalities remained small. When dimensionality likewise increases, learning becomes an intractable search problem without some form of organizing assumption.

### (b) Exploration

Of course, our formulation of continuous, open-world decisions naturally intersects with the question of how open worlds can be mapped. That is, continuous decisions face a much more extreme version of the explore/exploit dilemma than discrete choice paradigms. Thus, curiosity has recently experienced a surge of interest not only in neuroscience, but



**Figure 3.** Task and data from [45]. (a) Left: cartoon of *virtual pursuit task*. Subject uses a joystick to control an avatar (circle) and pursue prey (square) on a computer screen. Right: avatar and prey trajectories on example trials. Grey: path of avatar; red/blue: path of prey. Colour gradient indicates the time progression through the trial. (b) Left: a generative model explains the trajectory of non-human primates. Right: the model result exhibiting that non-human primates are predicting the future position of the prey. (c) Example filters from neurons that are tuned for the position of the prey within the monitor. (d) Example filters of neurons that are tuned for the future position of the prey.

in RL [23,24,46–49]. This also links to theoretical work on free-energy-inspired models [50,51] that incorporate information seeking, inference and reward into a single framework. However, this rapidly growing body of work is beyond the scope of this review. For our purposes, we are interested in the challenges and opportunities posed by continuous tasks, and focus for the remainder on the necessary integration of action and decision, assuming important quantities like exploratory policy and reward functions as given. More specifically, open-world decisions necessitate exploration, since the world is too big to be known, and that exploration requires strategy because large state spaces cannot be exhaustively visited [48,49]. In short, continuous decision-makers are almost always highly information-starved [23].

### (c) Need for heuristics

Finally, it is important to note that the goal of action is not necessarily to find the optimal course of action, but instead to find a course of action that is sufficiently good [52–55]. Heuristic approaches are, of course, valuable even in simple non-continuous decisions. The additional complexity of natural decisions means that satisficing approaches, such as heuristics, are likely to be used in real-world decision-making. In continuous decisions, where the costs of computation are much greater and the time pressures more unrelenting, heuristics are likely to be even more important. Understanding the use of heuristics in

continuous decisions, then, represents an important future direction for study (figure 3).

## 4. Continuous decisions involve dynamics and feedback

Two more important and closely related features of many continuous decisions are that they are dynamic and involve feedback. Indeed, the importance of both dynamics (sometimes high-dimensional, often nonlinear) and feedback (from both other agents and the environment) is the most obvious distinguishing feature between continuous decisions and their discrete counterparts. The problems imposed by dynamics encompass questions not only of control (e.g. how to move an arm to snatch a passing prey), but also of timing (precisely *when* to make this move), adding additional complexities often excluded in laboratory decision paradigms. Once again, we believe that RL, particularly its intersection with control theory, offers a conceptual framework for addressing these challenges, even in cases where learning is not the primary dynamic [56,57].

Before stepping further, it is helpful to draw the distinction between modelling a set of actions as sequential decisions and the cognitive experience of the sequential decision-making. For instance, while it is intuitive that the iterated Prisoner's Dilemma involves a sequence of

cooperation and defection decisions by each player, it is less intuitive that a short walk down a straight hallway also comprises a series of decisions, one for each step. Yet both can be modelled mathematically as probabilistic sequences of 'decisions,' perhaps, in the case of walking, with rewards for staying upright and forward progress. Yet our subjective experience is one of the long periods of automatic behaviour punctuated by changes of mind, phenomena that are only imperfectly captured by the framework of sequential decisions. Instead, we might prefer to mirror mathematically the dissociation between rare goal changes and continuous execution of unconscious behaviours in time.

Help in thinking about these problems comes from control theory, which has a long and fruitful history in motor neuroscience (e.g. [56–59]). In the control framework, a system is endowed with a set of possible controls to apply (forces, torques, electronic commands), a system to which these controls may be applied (typically modelled as a set of differential equations), a goal or set point and a cost for applying control (energy, price, regret). The aim of control theory is to solve for a *policy*, an algorithm that sets control as a function of the system's state at a given time. The optimal policy is one that best balances control error (how close the system remains to its goal) against control costs. Such a framework automatically addresses two key problems identified above. First, control systems operate in time. Control policies are *functions*, not discrete actions, and specify how the controller co-evolves with the system. They are thus fundamentally dynamic quantities. Second, control policies do not require the semantics of decisions at all. They are conceptualized as automatic processes that operate adaptively, often deterministically, once a goal has been set. Thus, they have intuitive appeal for a framework that combines sparse, conscious decisions with automatic ongoing processes.

Again, this approach is complementary to the RL perspective reviewed above [28], which likewise attempts to learn a policy that maximizes some set of rewards. In fact, many classic RL problems like cart-pole or learning to walk are control problems [56,60,61]. Such links have often been overlooked in decision neuroscience, where RL actions are most frequently associated with decisions [62], but there is nothing in the formalism to force this view. Before the advent of deep RL, RL methods were often limited to low-dimensional, discrete state and action spaces, but modern applications now routinely feature complex, continuous policies parameterized by deep neural networks [30,63]. Thus, RL likewise offers up the notion of a full policy in place of a discrete action as a building block of continuous decisions.

Just as importantly, while control systems are often assumed to be autonomous, with fixed, pre-specified set points, RL can be extended to the case of *hierarchical* learning and control [64–70]. That is, the RL framework can easily accommodate the idea of *switching between policies* or changing goals within the same policy [71,72]. These high-level changes are most often slower (in the case of goals) or sparser (for policy switches) and map neatly onto the experience of rarer deliberative decisions setting in motion automatic behaviours. Moreover, hierarchical formulations of RL allow changes in state to drive high-level actions, facilitating 'bottom-up' feedback that renders policies more flexible than purely 'fire and forget' processes. And while training such hierarchical RL systems has continued to prove challenging [65–67,73], there is ample evidence that brains have evolved by solving exactly this problem [38].

### (a) Example: penalty kick task

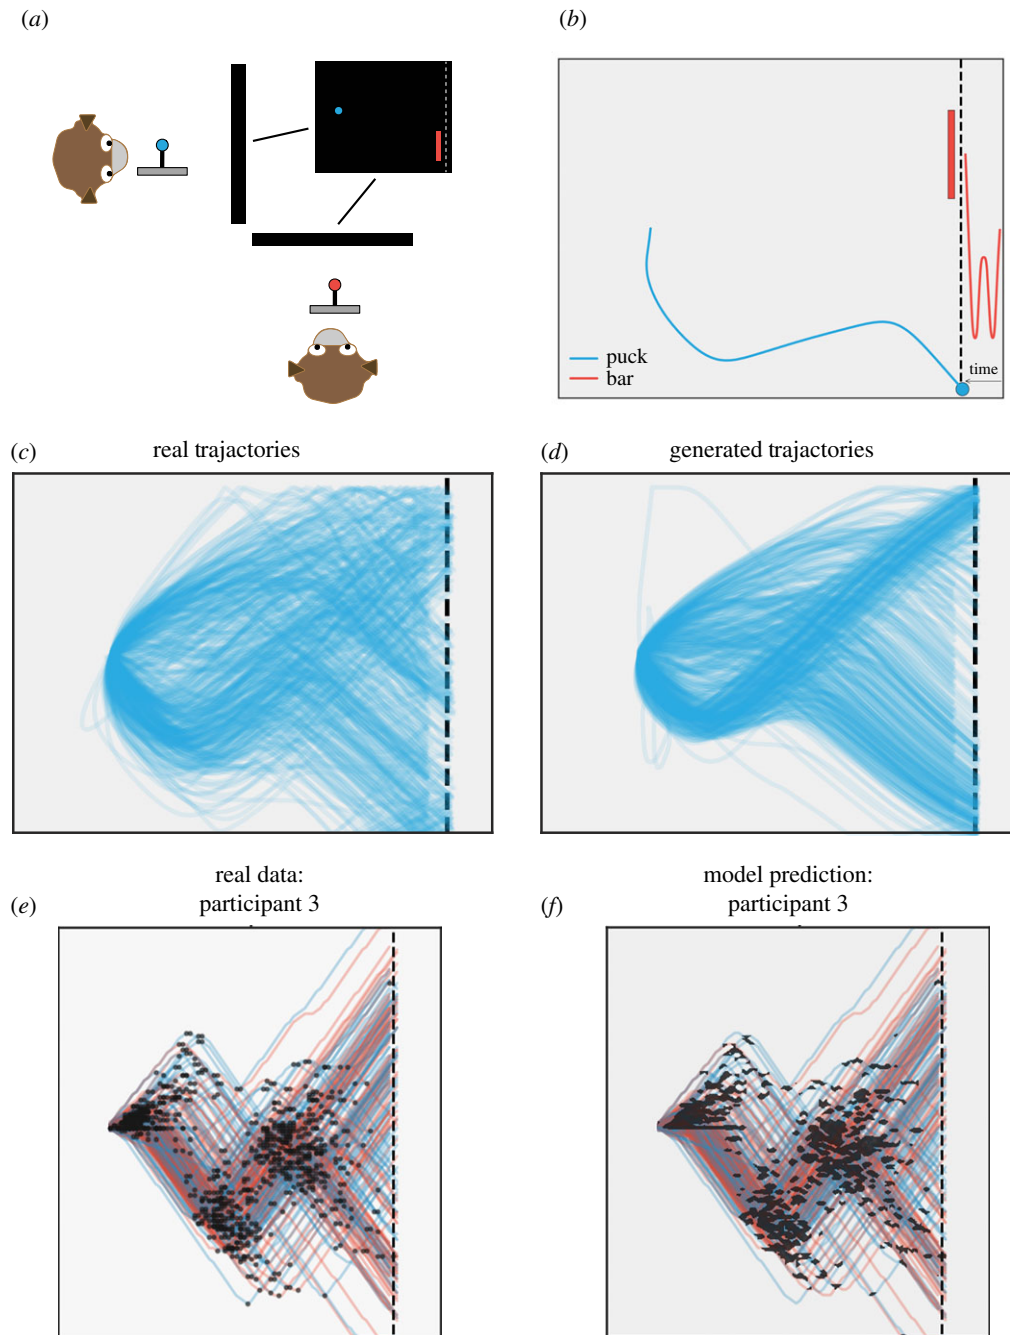
A recent study from our group illustrates this sort of integration of continuous control with sparse strategy changes. In the *penalty kick task*, monkeys and humans played a competitive video game against conspecifics that required continuous joystick input ([71,72], figure 4). The game is based on the idea of a penalty shot in hockey. One player (the shooter) controlled the motion of a small circle (the puck), while the other player (the goalie) controlled the vertical motion of a bar placed at the opposite side of the screen. Goalies were rewarded for blocking the puck with the bar, while shooters received rewards for reaching a goal line behind the bar. Though the game was highly constrained, players' trajectories proved highly variable, revealing a variety of both individual strategies and inter-player dynamics. Such paradigms as this bear some relation to continuous games such as duels and are closely related to pursuit games in differential game theory, but in the latter, the focus is typically on computing optimal solutions rather than modelling players' real behaviour [74,75].

These phenomena can be modelled as the result of a control model applied to an evolving value function incorporating each player's trajectory [72,76], while directly modelling players' policies allows for the calculation of a measure of instantaneous coupling between the players [71]. However, in both human and monkey participants, we found that the key strategic variable was not so much shooters' ability to adjust to their opponents as their ability to advantageously time their final movements on each trial [71]. This work, in other words, provides important information about the specific strategies that human and macaque decision-makers are able to use in such tasks; that information is unavailable in discrete decision-making paradigms.

## 5. Implications for neuroscience: a preliminary sketch

Traditional 'box-and-line' approaches to cognitive neuroscience presume the existence of discrete cognitive functions with intuitive easy-to-name roles. These functions are assumed to be reified in neuroanatomy. For example, if choice consists of evaluation, comparison and selection, then these three conceptually discrete functions ought to correspond to discrete anatomical substrates [77,78]. An alternative viewpoint is distributed; it imagines that choice reflects an emergent process arising from multiple brain regions whose functions may not correspond to nameable processes, and/or that may largely overlap [18,25,79–81].

The greatest debates about modularity concern the prefrontal cortex, which can be readily divided into a few dozen areas based on cytoarchitecture. However, the functions of these areas remain difficult to identify definitively. In fact, this architectural diversity raises the question of why we have so many prefrontal regions to begin with [82]. One possibility is that each region integrates new information from new sources and participates in transforming that information via an untangling process [79]. In this view, each area in sequence provides a partial transformation from one representation format to another [83,84]. Indeed, it is this view, and not that of a modular system, that best accords with recent results in artificial neural networks (e.g. [85–87]). Of course, some cognitive functions are localized at the coarsest levels. Our claim is simply that, once one embraces the framework of continuous



**Figure 4.** Task from [72]. (a) Pairs of monkeys competed in a zero-sum, real-time game with continuous joystick inputs. One player controlled the position of a blue ball, the other the vertical position of a red bar. The goal of the blue player was to guide the ball past the dotted line at the right of the screen, while the red player attempted to block the ball with the bar. (b) Example of a single pair of player trajectories on a single trial. For display purposes, the trajectory of the bar has been extended in time along the  $x$ -axis, though the bar could only move vertically. (c) Continuous control resulted in highly varied trajectories, but these could be captured by a model of players' coevolving goal states that generated realistic trials (d). (e) In a version of the same task in human subjects, players' trajectories were characterized by sparse changes in trajectory (black dots), and these changepoints could be predicted by modelling players' policies as a function of continuous game state (f). Black patches represent regions of increased probability of direction change. Blue traces indicate trials played against a human bar opponent, red traces those played against a computer. Adapted from [71,72].

decisions, with its tight coupling of action and decision, dynamics and feedback, the tight conceptual link between modularity in space and serial processing in time is no longer as natural as it has traditionally appeared.

To elaborate on this last point, consider choosing between two options. In this simple situation, there is no particular reason (aside from reaction time minimization) for the motor action to begin until comparison and choice of option are complete [18]. In other words, a serial design suffices. And indeed, there is plentiful evidence consistent with the serial model of

choice (e.g. [78]). But in continuous decisions, the action is fully interwoven with choice. It is difficult to even think about how choice may occur without consideration of action [18,25].

We explored the idea of how the brain predicts the future in a recent study [45]. In our virtual pursuit project (see above), subjects could gain a strategic advantage over their prey by predicting the prey's future position and moving to it, and a generative model of behaviour allowed us to confirm that subjects indeed moved towards the predicted position of the prey. We then used the mapping algorithm to map the

representation of predicted future positions of prey (specifically, one developed in [88]). That is, we computed the average firing rate of each neuron as a function of the position in space of the future position of the avatar. These maps are akin to place cell maps in the hippocampus, but they are for virtual positions (on a computer screen) and for future, rather than present positions. We found that these maps are present in single dACC neurons and that they come from the same cells—multiplexed—as self-position maps.

When decisions are not discrete selections but processes of continuous adjustment made on the fly, action selection can no longer easily be separated from choice. This suggests a link to embodied decision-making arguments forwarded by Shadlen *et al.* [89], who argue that the responses of single neurons in the lateral intraparietal area (LIP) to visual motion are actually better thought of as an embodied decision variable—the weight of evidence in favour of shifting gaze to a particular location. This theory neatly explains why such signals are found in neurons with response fields, and neatly solves the classic debate about the intentional role of LIP. It is a small move from this theory of perceptual decision-making to economic decisions.

### (a) Modularity and continuous decisions

We explored the idea of how the brain predicts the future in a recent study [45]. Predicting the consequences of our actions and using those predictions to guide choice are common to discrete decisions [90,91]. But it is much more difficult in continuous tasks. In our virtual pursuit project (see above), subjects could gain a strategic advantage over their prey by predicting the prey's future position and moving to it—much like the dragonfly discussed above. We then designed a generative model of behaviour and were able to explicitly test the hypothesis that subjects moved towards the predicted position of the prey. We then used the mapping algorithm to map the representation of predicted future positions of prey. We found that these maps are present in single dACC neurons and that they come from the same cells—multiplexed—as self-position maps.

A consequence of this idea is that decision computations ought to be constrained, at least in part, by the motor system. There is plentiful evidence for this idea. For example, we tend to gaze longer at options that we are about to choose. There is evidence that exogenous manipulations that bias our looking times towards one option increase our chances of choosing that option [92]. Another example is the idea from robotics and cognitive science called 'morphological computation' [93–95]. The idea is that the shape and material properties of the body can be exploited to make central control processes simpler. It implies that the local computation that is based on certain properties of the body would influence the control process [96]. In this case, if a certain effector is used to implement choice, their biomechanical properties would influence the outcome of the choice. This raises an intriguing possibility—continuous decisions expressed in two modalities may differ owing to different physical properties or biomechanical constraints. Because continuous decisions are more tightly integrated with effector systems than discrete ones are, their underlying computations will be more influenced by the affordances of these systems. That is, we may make the same choice whether the output modality is our eyes, our fingers or our feet, but we may make different

decisions when the modality is the same. We may, for example, predict further into the future when using our eyes than our hands.

Further evidence comes from the fact that subjects appear to initiate their action before they have fully decided on an option. That is, they hedge their bets and choose an intermediate path between the two options [97,98]. Then, as their decision process unfolds in their heads, their arm changes its pattern and begins to divert its path towards the chosen option. Critically, the initial hedging path lasts longer if the choice takes more time to unfold. In other words, the beginning of selection takes place before the choice is complete. Taking this idea literally, the choice process occurs simultaneously with the motor action. They are not discrete and the action would provide a feedback signal to the choice, which is opposed to the idea that movement does not influence the choice at all. Ultimately, then, the resolution of the question of modularity requires testing under more naturalistic contexts. That is, natural tasks provide a valuable 'stress test' for modular models.

## 6. Conclusion

Here, we offer a definition for continuous decisions, and, by contrast, a complementary definition for discrete decisions. We argue that the neuroscientific and psychological study of decision-making is impoverished by its near-exclusive focus on discrete decisions, and would be enriched by expanding its scope to include continuous ones.

One practical obstacle is the lack of standardized tasks that can serve as the focus of such research. We propose that simple computerized continuous motion tasks, such as the penalty kick task and pursuit task, can serve that role [45,72]. These two tasks involve a continuum of actions, take place over an extended time, and, because they involve multiple players, exhibit the complex intra-trial feedback relationship for decisions that is characteristic of many naturalistic continuous decisions.

Another practical obstacle to the adoption of continuous decisions is the additional complexity involved in studying continuous tasks. As a result, we propose that the development of new analysis techniques ought to be a major focus moving forward. For example, most traditional analyses in systems neuroscience have relied on averages, both in time and across trials, to reduce the complexity in data and achieve statistical power. But when trial-to-trial behavioural variability is large, these repetitions may no longer be comparable to one another. As a result, single-trial analyses will likely be crucial (e.g. [99,100], also [71,72,76]). And these methods often benefit from denser recording techniques such as calcium imaging [101] and multi-contact multi-electrode recording ([102,103] and [104,105]).

Along a different axis, there is much promise in recent studies that eschew single brain areas in favour of integrative analyses of simple, natural behaviours. Recent technical advances are making such studies possible, and this trend is likely to continue (e.g. [106]). Corresponding analytical tools are being developed too [107,108]. This also requires novel theories about information transformations between areas (e.g. [79,86,109]).

To the degree that some complexity in behaviour is irreducible, the study of continuous decisions is both necessary and inevitable. Not only are these decisions closer to natural

behaviour (and so easier to train in model species), they begin to tap the complexity we want to study, offering new puzzles that force us to integrate our knowledge across cognitive levels, brain regions and time scales. If the behaviours we are interested in—our own included—take place in a changing, dynamic world, then our models of that world must grow to embrace these essential features.

## References

- Allais M. 1953 Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica* **21**, 503–546. (doi:10.2307/1907921)
- Kahneman D, Tversky A. 1979 Prospect theory: an analysis of decision under risk. *Econometrica* **47**, 263–292. (doi:10.2307/1914185)
- Glimcher PW, Fehr E. 2013 *Neuroeconomics: decision making and the brain*. London, UK: Academic Press.
- Platt ML, Huettel SA. 2008 Risky business: the neuroeconomics of decision making under uncertainty. *Nat. Neurosci.* **11**, 398–403. (doi:10.1038/nn2062)
- Lepora NF, Pezzulo G. 2015 Embodied choice: how action influences perceptual decision making. *PLoS Comput. Biol.* **11**, e1004110. (doi:10.1371/journal.pcbi.1004110)
- Pezzulo G, Cisek P. 2016 Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends Cogn. Sci.* **20**, 414–424. (doi:10.1016/j.tics.2016.03.013)
- Gonzalez-Bellido PT, Fabian ST, Nordström K. 2016 Target detection in insects: optical, neural and behavioral optimizations. *Curr. Opin. Neurobiol.* **41**, 122–128. (doi:10.1016/j.conb.2016.09.001)
- Brown AE, De Bivort B. 2018 Ethology as a physical science. *Nat. Phys.* **14**, 653–657. (doi:10.1038/s41567-018-0093-0)
- Kacelnik A, Vasconcelos M, Monteiro T, Aw J. 2011 Darwin's 'tug-of-war' vs. starlings' 'horse-racing': how adaptations for sequential encounters drive simultaneous choice. *Behav. Ecol. Sociobiol.* **65**, 547–558. (doi:10.1007/s00265-010-1101-2)
- Krakauer JW, Ghazanfar AA, Gomez-Marín A, Machner MA, Poeppel D. 2017 Neuroscience needs behavior: correcting a reductionist bias. *Neuron* **93**, 480–490. (doi:10.1016/j.neuron.2016.12.041)
- Pearson JM, Watson KK, Platt ML. 2014 Decision making: the neuroethological turn. *Neuron* **82**, 950–965. (doi:10.1016/j.neuron.2014.04.037)
- Stephens DW, Anderson D. 2001 The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behav. Ecol.* **12**, 330–339. (doi:10.1093/beheco/12.3.330)
- Hayden BY. 2016 Time discounting and time preference in animals: a critical review. *Psychon. Bull. Rev.* **23**, 39–53. (doi:10.3758/s13423-015-0879-3)
- Blanchard TC, Hayden BY. 2014 Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *J. Neurosci.* **34**, 646–655. (doi:10.1523/JNEUROSCI.3151-13.2014)
- Blanchard TC, Pearson JM, Hayden BY. 2013 Postreward delays and systematic biases in measures of animal temporal discounting. *Proc. Natl Acad. Sci. USA* **110**, 15 491–15 496. (doi:10.1073/pnas.1310446110)
- Heilbronner SR, Hayden BY. 2016 The description-experience gap in risky choice in nonhuman primates. *Psychon. Bull. Rev.* **23**, 593–600. (doi:10.3758/s13423-015-0924-2)
- Eisenreich BR, Hayden BY, Zimmermann J. 2019 Macaques are risk-averse in a freely moving foraging task. *Sci. Rep.* **9**, 1–12. (doi:10.1038/s41598-018-37186-2)
- Cisek P. 2012 Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* **22**, 927–936. (doi:10.1016/j.conb.2012.05.007)
- Heilbronner SR. 2017 Modeling risky decision-making in nonhuman animals: shared core features. *Curr. Opin. Behav. Sci.* **16**, 23–29. (doi:10.1016/j.cobeha.2017.03.001)
- d'Acremont M, Bossaerts P. 2008 Neurobiological studies of risk assessment: a comparison of expected utility and mean–variance approaches. *Cogn. Affect. Behav. Neurosci.* **8**, 363–374. (doi:10.3758/cabn.8.4.363)
- Blanchard TC, Hayden BY. 2015 Monkeys are more patient in a foraging task than in a standard intertemporal choice task. *PLoS ONE* **10**, e0117057. (doi:10.1371/journal.pone.0117057)
- Muncy J. 2015 Open-world games are changing the way we play. *Wired*, 12/03/2015.
- Wang MZ, Hayden BY. 2020 Latent learning, cognitive maps, and curiosity. *Curr. Opin. Behav. Sci.* **38**, 1–7. (doi:10.1016/j.cobeha.2020.06.003)
- Cervera RL, Wang MZ, Hayden BY. 2020 Systems neuroscience of curiosity. *Curr. Opin. Behav. Sci.* **35**, 48–55. (doi:10.1016/j.cobeha.2020.06.011)
- Cisek P, Kalaska JF. 2010 Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298. (doi:10.1146/annurev.neuro.051508.135409)
- Cisek P, Kalaska JF. 2005 Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814. (doi:10.1016/j.neuron.2005.01.027)
- Hayden BY, Moreno-Bote R. 2018 A neuronal theory of sequential economic choice. *Brain Neurosci. Adv.* **2**, 2398212818766675. (doi:10.1177/2398212818766675)
- Sutton RS, Barto AG. 2018 *Reinforcement learning: an introduction*. Cambridge, MA: MIT press.
- Dayan P, Abbott LF. 2001 *Theoretical neuroscience: computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.
- Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Petersen S. 2015 Human-level control through deep reinforcement learning. *Nature* **518**, 529–533. (doi:10.1038/nature14236)
- Vinyals O, Babuschkin I, Czarnecki WM, Mathieu M, Dudzik A, Chung J, Oh J. 2019 Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**, 350–354. (doi:10.1038/s41586-019-1724-z)
- Finn C, Yu T, Zhang T, Abbeel P, Levine S. 2017 One-shot visual imitation learning via meta-learning. arXiv preprint arXiv:1709.04905.
- Duan Y, Andrychowicz M, Stadie B, Ho OJ, Schneider J, Sutskever I, Abbeel P, Zaremba W. 2017 One-shot imitation learning. In *Advances in neural information processing systems*, vol. 30 (eds U von Luxburg, S Bengio, R Fergus, R Garnett, I Guyon, H Wallach, SVN Vishwanathan) pp. 1087–1098. Red Hook, NY: Curran Associates.
- Ho J, Ermon S. 2016 Generative adversarial imitation learning. In *Advances in neural information processing systems*, vol. 29 (eds DD Lee, M Sugiyama, UV Luxburg, I Guyon, R Garnett) pp. 4565–4573. Red Hook, NY: Curran Associates.
- Wikenheiser AM, Schoenbaum G. 2016 Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nat. Rev. Neurosci.* **17**, 513–523. (doi:10.1038/nrn.2016.56)
- Schuck NW, Cai MB, Wilson RC, Niv Y. 2016 Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* **91**, 1402–1412. (doi:10.1016/j.neuron.2016.08.019)
- Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. 2014 Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267–279. (doi:10.1016/j.neuron.2013.11.005)
- Niv Y. 2019 Learning task-state representations. *Nat. Neurosci.* **22**, 1544–1553. (doi:10.1038/s41593-019-0470-8)
- Silver D, Lever G, Heess N, Degris T, Wierstra D, Riedmiller M. 2014 Deterministic policy gradient algorithms. In *Proceedings of the 31st Int. Conf. on Machine Learning*, vol. 32 (ICML'14), pp. 1-387–1-395.



40. Lillicrap TP *et al.* 2015 Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971. (<https://arxiv.org/abs/1509.02971>)
41. Behrens TE, Muller TH, Whittington JC, Mark S, Baram AB, Stachenfeld KL, Kurth-Nelson Z. 2018 What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* **100**, 490–509. (doi:10.1016/j.neuron.2018.10.002)
42. Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, Gershman SJ. 2017 The successor representation in human reinforcement learning. *Nat. Hum. Behav.* **1**, 680–692. (doi:10.1038/s41562-017-0180-8)
43. Stachenfeld KL, Botvinick MM, Gershman SJ. 2017 The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643. (doi:10.1038/nn.4650)
44. Tolman EC. 1948 Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189. (doi:10.1037/h0061626)
45. Yoo SBM, Tu JC, Piantadosi ST, Hayden BY. 2020 The neural basis of predictive pursuit. *Nat. Neurosci.* **23**, 252–259. (doi:10.1038/s41593-019-0561-6)
46. Kidd C, Hayden BY. 2015 The psychology and neuroscience of curiosity. *Neuron* **88**, 449–460. (doi:10.1016/j.neuron.2015.09.010)
47. Watters N, Matthey L, Bosnjak M, Burgess CP, Lerchner A. 2019 Cobra: Data-efficient model-based rl through unsupervised object discovery and curiosity-driven exploration. arXiv preprint arXiv:1905.09275. (<https://arxiv.org/abs/1905.09275>)
48. Ecoffet A, Huizinga J, Lehman J, Stanley KO, Clune J. 2019 Go-explore: a new approach for hard-exploration problems. arXiv preprint arXiv:1901.10995.
49. Gottlieb J, Oudeyer PY, Lopes M, Baranes A. 2013 Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends Cogn. Sci.* **17**, 585–593. (doi:10.1016/j.tics.2013.09.001)
50. Friston KJ, Lin M, Frith CD, Pezzulo G, Hobson JA, Oudobaka S. 2017 Active inference, curiosity and insight. *Neural Comput.* **29**, 2633–2683. (doi:10.1162/neco\_a\_00999)
51. Bossaerts P. 2018 Formalizing the function of anterior insula in rapid adaptation. *Front. Integr. Neurosci.* **12**, 61. (doi:10.3389/fnint.2018.00061)
52. Griffiths TL, Lieder F, Goodman ND. 2015 Rational use of cognitive resources: levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* **7**, 217–229. (doi:10.1111/tops.12142)
53. Schwartz B, Ward A, Monterosso J, Lyubomirsky S, White K, Lehman DR. 2002 Maximizing versus satisficing: happiness is a matter of choice. *J. Pers. Soc. Psychol.* **83**, 1178. (doi:10.1037/0022-3514.83.5.1178)
54. Gigerenzer G, Gaissmaier W. 2011 Heuristic decision making. *Annu. Rev. Psychol.* **62**, 451–482. (doi:10.1146/annurev-psych-120709-145346)
55. Lieder F, Griffiths TL. 2020 Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* **43**, e1. (doi:10.1017/s0140525x1900061x)
56. Todorov E, Jordan MI. 2002 Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235. (doi:10.1038/nn963)
57. Wolpert DM, Diedrichsen J, Flanagan JR. 2011 Principles of sensorimotor learning. *Nat. Rev. Neurosci.* **12**, 739–751. (doi:10.1038/nrn3112)
58. Gollivan JP, Chapman CS, Wolpert DM, Flanagan JR. 2018 Decision-making in sensorimotor control. *Nat. Rev. Neurosci.* **19**, 519–534. (doi:10.1038/s41583-018-0045-9)
59. Diedrichsen J, Shadmehr R, Ivry RB. 2010 The coordination of movement: optimal feedback control and beyond. *Trends Cogn. Sci.* **14**, 31–39. (doi:10.1016/j.tics.2009.11.004)
60. Todorov E. 2004 Optimality principles in sensorimotor control. *Nat. Neurosci.* **7**, 907–915. (doi:10.1038/nn1309)
61. Todorov E. 2009 Efficient computation of optimal actions. *Proc. Natl Acad. Sci. USA* **106**, 11 478–11 483. (doi:10.1073/pnas.0710743106)
62. Daw ND, Tobler PN. 2014 Value learning through reinforcement: the basics of dopamine and reinforcement learning. In *Neuroeconomics: decision making and the brain* (ed. PW Glimcher, E Fehr), pp. 283–298, 2nd edn. London, UK: Academic Press.
63. Van Hasselt H, Guez A, Silver D. 2016 Deep reinforcement learning with double q-learning. (AAAI'16) In *Proc. Thirtieth AAAI conference on artificial intelligence*, pp. 2094–2100. Menlo Park, CA: AAAI Press.
64. Doya K, Samejima K, Katagiri KI, Kawato M. 2002 Multiple model-based reinforcement learning. *Neural Comput.* **14**, 1347–1369. (doi:10.1162/089976602753712972)
65. Precup D, Sutton RS. 1998 Multi-time models for temporally abstract planning. In *Advances in neural information processing systems*, vol. 10 (eds MI Jordan, MJ Kearns, SA Solla), pp. 1050–1056. Cambridge, MA: MIT.
66. Sutton RS, Precup D, Singh S. 1999 Between MDPs and semi-MDPs: a framework for temporal abstraction in reinforcement learning. *Artif. Intell.* **112**, 181–211. (doi:10.1016/S0004-3702(99)00052-1)
67. Bacon PL, Harb J, Precup D. 2017 Option-critic architecture. arXiv:1609.05140v2. (<https://arxiv.org/abs/1609.05140>)
68. Pearson JM, Platt ML. 2013 Change detection, multiple controllers, and dynamic environments: insights from the brain. *J. Exp. Anal. Behav.* **99**, 74–84. (doi:10.1002/jeab.5)
69. Merel J, Botvinick M, Wayne G. 2019 Hierarchical motor control in mammals and machines. *Nat. Commun.* **10**, 1–12. (doi:10.1038/s41467-019-13239-6)
70. Botvinick MM, Niv Y, Barto AC. 2009 Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition* **113**, 262–280. (doi:10.1016/j.cognition.2008.08.011)
71. McDonald KR, Broderick WF, Huettel SA, Pearson JM. 2019 Bayesian nonparametric models characterize instantaneous strategies in a competitive dynamic game. *Nat. Commun.* **10**, 1–12. (doi:10.1038/s41467-019-09789-4)
72. Iqbal SN, Yin L, Drucker CB, Kuang Q, Gariépy JF, Platt ML, Pearson JM. 2019 Latent goal models for dynamic strategic interaction. *PLoS Comput. Biol.* **15**, e1006895. (doi:10.1371/journal.pcbi.1006895)
73. Nachum O, Tang H, Lu X, Gu S, Lee H, Levine S. 2019 Why does hierarchy (sometimes) work so well in reinforcement learning? arXiv preprint arXiv:1909.10618. (<https://arxiv.org/abs/1909.10618>)
74. Basar T, Olsder GJ. 1999 *Dynamic noncooperative game theory*, vol. 23. Philadelphia, PA: Society for Industrial and Applied Mathematics.
75. Braun DA, Aertsen A, Wolpert DM, Mehring C. 2009 Learning optimal adaptation strategies in unpredictable motor tasks. *J. Neurosci.* **29**, 6472–6478. (doi:10.1523/jneurosci.3075-08.2009)
76. McDonald KR, Pearson JM, Huettel SA. 2020 Dorsolateral and dorsomedial prefrontal cortex track distinct properties of dynamic social behavior. *Soc. Cogn. Affect. Neurosci.* **15**, 383–393. (doi:10.1093/scan/nsaa053)
77. Rangel A, Camerer C, Montague PR. 2008 A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* **9**, 545–556. (doi:10.1038/nrn2357)
78. Padoa-Schioppa C. 2011 Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* **34**, 333–359. (doi:10.1146/annurev-neuro-061010-113648)
79. Yoo SBM, Hayden BY. 2018 Economic choice as an untangling of options into actions. *Neuron* **99**, 434–447. (doi:10.1016/j.neuron.2018.06.038)
80. Hunt LT, Hayden BY. 2017 A distributed, hierarchical and recurrent framework for reward-based choice. *Nat. Rev. Neurosci.* **18**, 172–182. (doi:10.1038/nrn.2017.7)
81. Pirrone A, Azab H, Hayden BY, Stafford T, Marshall JA. 2018 Evidence for the speed–value trade-off: human and monkey decision making is magnitude sensitive. *Decision* **5**, 129. (doi:10.1037/dec0000075)
82. Passingham RE, Wise SP. 2012 *The neurobiology of the prefrontal cortex: anatomy, evolution, and the origin of insight (No. 50)*. Oxford, UK: Oxford University Press.
83. Azab H, Hayden BY. 2017 Correlates of decisional dynamics in the dorsal anterior cingulate cortex. *PLoS Biol.* **15**, e2003091. (doi:10.1371/journal.pbio.2003091)
84. Azab H, Hayden BY. 2018 Correlates of economic decisions in the dorsal and subgenual anterior cingulate cortices. *Eur. J. Neurosci.* **47**, 979–993. (doi:10.1111/ejn.13865)
85. Yamins DLK, Hong H, Cadieu CF, Solomon EA, Seibert D, DiCarlo JJ. 2014 Performance-optimized hierarchical models predict neural responses in

- higher visual cortex. *Proc. Natl Acad. Sci. USA* **111**, 8619–8624. (doi:10.1073/pnas.1403112111)
86. DiCarlo JJ, Zoccolan D, Rust NC. 2012 How does the brain solve visual object recognition? *Neuron* **73**, 415–434. (doi:10.1016/j.neuron.2012.01.010)
87. Mante V, Sussillo D, Shenoy KV, Newsome WT. 2013 Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84. (doi:10.1038/nature12742)
88. Shadlen MN, Kiani R, Hanks TD, Churchland AK. 2008 Neurobiology of decision making. An intentional framework. In *Better than conscious? Decision making, the human mind, and implications for institutions* (eds C Engel, W Singer), 71–101. Cambridge, MA: MIT. (doi:10.7551/mitpress/9780262195805.003.0004)
89. Shadlen MN, Kiani R, Hanks TD, Churchland AK. 2008 Neurobiology of decision making. An intentional framework. In *Better than conscious? Decision making, the human mind, and implications for institutions* (eds C Engel, W Singer), pp. 71–101. Cambridge, MA: MIT. (doi:10.7551/mitpress/9780262195805.003.0004)
90. Wang MZ, Hayden BY. 2017 Reactivation of associative structure specific outcome responses during prospective evaluation in reward-based choices. *Nat. Commun.* **8**, 1–13. (doi:10.1038/s41467-016-0009-6)
91. Kahnt T, Heinzle J, Park SQ, Haynes JD. 2010 The neural code of reward anticipation in human orbitofrontal cortex. *Proc. Natl Acad. Sci. USA* **107**, 6010–6015. (doi:10.1073/pnas.0912838107)
92. Shimojo S, Simion C, Shimojo E, Scheier C. 2003 Gaze bias both reflects and influences preference. *Nat. Neurosci.* **6**, 1317–1322. (doi:10.1038/nn1150)
93. Paul C. 2006 Morphological computation: a basis for the analysis of morphology and control requirements. *Robot. Auton. Syst.* **54**, 619–630. (doi:10.1016/j.robot.2006.03.003)
94. Brooks RA. 1990 Elephants don't play chess. *Robot. Auton. Syst.* **6**, 3–15. (doi:10.1016/S0921-8890(05)80025-9)
95. Zhang YS, Ghazanfar AA. 2018 Vocal development through morphological computation. *PLoS Biol.* **16**, e2003933. (doi:10.1371/journal.pbio.2003933)
96. Ghazi-Zahedi K, Langer C, Ay N. 2017 Morphological computation: synergy of body and brain. *Entropy* **19**, 456. (doi:10.3390/e19090456)
97. Haith AM, Pakpoor J, Krakauer JW. 2016 Independence of movement preparation and movement initiation. *J. Neurosci.* **36**, 3007–3015. (doi:10.1523/JNEUROSCI.3245-15.2016)
98. Christopoulos V, Schrater PR. 2015 Dynamic integration of value information into a common probability currency as a theory for flexible decision making. *PLoS Comput. Biol.* **11**, 1–26. (doi:10.1371/journal.pcbi.1004402)
99. Sauerbrei BA, Guo JZ, Cohen JD, Mischiati M, Guo W, Kabra M, Hantman AW. 2020 Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**, 386–391. (doi:10.1038/s41586-019-1869-9)
100. Pandarinath C, O'Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, Henderson JM. 2018 Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* **15**, 805–815. (doi:10.1038/s41592-018-0109-9)
101. Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. 2019 Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, eaav7893. (doi:10.1126/science.aav7893)
102. Egger SW, Remington ED, Chang CJ, Jazayeri M. 2019 Internal models of sensorimotor integration regulate cortical dynamics. *Nat. Neurosci.* **22**, 1871–1882. (doi:10.1038/s41593-019-0500-6)
103. Gallego JA, Perich MG, Miller LE, Solla SA. 2017 Neural manifolds for the control of movement. *Neuron* **94**, 978–984. (doi:10.1016/j.neuron.2017.05.025)
104. Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE. 2020 Long-term stability of cortical population dynamics underlying consistent behavior. *Nat. Neurosci.* **23**, 260–270. (doi:10.1038/s41593-019-0555-4)
105. Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Chase SM, Yu BM. 2018 Learning by neural reassociation. *Nat. Neurosci.* **21**, 607–616. (doi:10.1038/s41593-018-0095-3)
106. Steinmetz NA, Koch C, Harris KD, Carandini M. 2018 Challenges and opportunities for large-scale electrophysiology with Neuropixels probes. *Curr. Opin. Neurobiol.* **50**, 92–100. (doi:10.1016/j.conb.2018.01.009)
107. Pinto L, Rajan K, DePasquale B, Thiberge SY, Tank DW, Brody CD. 2019 Task-dependent changes in the large-scale dynamics and necessity of cortical regions. *Neuron* **104**, 810–824. (doi:10.1016/j.neuron.2019.08.025)
108. Semedo JD, Zandvakili A, Machens CK, Byron MY, Kohn A. 2019 Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259. (doi:10.1016/j.neuron.2019.01.026)
109. Bernardi S, Benna MK, Rigotti M, Munuera J, Fusi S, Salzman CD. 2019 The geometry of abstraction in hippocampus and prefrontal cortex. *Biorxiv*, 408633. (doi:10.1101/408633)