Value-free reinforcement learning: policy optimization as a minimal model of operant behavior

Daniel Bennett^{a,b}, Yael Niv^{a,c}, Angela J. Langdon^a

^a*Princeton Neuroscience Institute, Princeton University, USA*

^b*Department of Psychiatry, Monash University, Australia*

^c*Department of Psychology, Princeton University, USA*

Abstract

Reinforcement learning is a powerful framework for modelling the cognitive and neural substrates of learning and decision making. Contemporary research in cognitive neuroscience and neuroeconomics typically uses value-based reinforcement-learning models, which assume that decision-makers choose by comparing learned values for different actions. However, another possibility is suggested by a simpler family of models, called *policy-gradient reinforcement learning*. Policy-gradient models learn by optimizing a behavioral policy directly, without the intermediate step of value-learning. Here we review recent behavioral and neural findings that are more parsimoniously explained by policy-gradient models than by value-based models. We conclude that, despite the ubiquity of ‘value’ in reinforcement-learning models of decision making, policy-gradient models provide a lightweight and compelling alternative model of operant behavior.

Keywords: value, decision-making, reinforcement learning, policy gradient, computational modelling

Highlights

- Reinforcement learning models can be divided into value-based and policy-gradient
- Policy-gradient RL gives a more parsimonious account of many operant behaviors
- These behaviors can therefore be explained without invoking the notion of ‘value’

1. Introduction

1 What is value? In spite of the ubiquity of this term in the field of value-
2 based decision making, there are a number of different ways of defining
3 value, and one’s chosen definition has important implications for the kinds of
4 inferences that one is likely to draw about the cognitive and neural processes
5 that subserve operant behavior [1, 2, 3, 4, 5].

6 In typical usage, ‘value’ is an explanatory variable that quantifies the
7 degree to which an individual prefers (or is willing to work for) one good
8 or outcome over others [6, 7]. So, for instance, a child who consistently
9 chooses a chocolate bar over fresh fruit would be said to value chocolate
10 more highly. Value—as typically defined—is therefore a latent construct
11 that scaffolds choice behavior by providing a common currency for compar-
12 ison of different actions (often loosely identified with the similar economic
13 concept of *utility* [8, 9]). In support of the hypothesis that choice behavior
14 is supported by the computation of value, proponents frequently note (e.g.,
15 [10]) that reward-related dopaminergic neural activity is consistent with a
16 class of value-learning algorithms from the field of *reinforcement learning*
17 (RL) [11, 12]. This argument implicitly assumes that different theoretical
18 frameworks each use the term ‘value’ in the same way. However, there are
19 important differences between typical usage of ‘value’ and value as defined
20 in RL.

21 In RL, value is defined as *expected cumulative future reward* [13] (see Box
22 1). As such, an action’s value in RL is not merely a relative quantification
23 of preference, as in the typical usage of ‘value’; instead, it exactly quanti-
24 fies the total amount of future reward¹ that an agent expects to receive if
25 it takes that action. At first glance, this definition of value appears con-
26 sonant with typical usage, since it seems reasonable that one should prefer
27 actions with greater expected cumulative future reward. Crucially, however,
28 although in RL differences in learned value between two actions do indeed
29 entail differences in choice behavior, the converse is not true. That is, we
30 may observe differences in an individual’s preference for different actions
31 despite the agent not having learned different values (in the RL sense) for
32 those actions. This is because there is an entire separate branch of RL
33 algorithms—*policy-gradient* RL—in which agents learn to choose actions by
34 optimizing a behavioral policy directly, without ever taking the intermediate
35 step of learning their values [14].

¹After accounting for the temporal discounting of future rewards.

36 Here we give an introduction to the distinction between value-based and
37 policy-gradient RL (also termed ‘indirect’ and ‘direct’ actors, respectively
38 [15]), and review recent work suggesting that policy-gradient RL provides a
39 good account of neural and behavioral data. We suggest that value (in the
40 RL sense) is not strictly entailed by data from typical laboratory tasks, and
41 propose that value need only be invoked as an explanatory latent construct
42 for phenomena that cannot be accounted for under simpler algorithms like
43 policy-gradient RL. This does not exclude the possibility that the brain does
44 compute value in some circumstances; however, we would argue that policy-
45 gradient RL models are generally favoured by the principle of parsimony,
46 and that ‘value’ should only be added to these models to explain data that
47 cannot be explained by policy-gradient models alone.

48 **2. Value-based RL versus policy-gradient RL**

49 Value-based RL is a prominent computational model of the cognitive
50 processes subserving simple operant behavior and of the neural substrates of
51 these processes (e.g., [16, 17, 18, 19]). Briefly, in a prototypical value-based
52 RL model (Figure 1A), the agent learns the value of each of a set of discrete
53 actions by trial-and-error, and chooses between actions by mapping these
54 estimated values into a behavioral *policy* using a policy mapping function
55 (see Box 1). As such, using a value-based RL model as a model of a subject’s
56 behavior implicitly assumes that, at an algorithmic level of description, the
57 subject learns action-values and makes choices by comparing them.

58 According to the principle of parsimony, however, in modelling subjects’
59 behavior we should attempt to find the model that best explains the data
60 while minimizing model complexity (i.e., invoking as few latent explanatory
61 constructs as possible) [20]. This raises the question: do value-based models
62 provide the simplest RL account of behavior, or can simpler models (such
63 as policy-gradient RL; Figure 1B) account for behavior equally well?

A. Value-based reinforcement learning

Example: *Q-learning with softmax choice rule*

Compute policy from action-values:
$$\pi(a) = \frac{e^{\beta \cdot Q(a)}}{\sum_{\hat{a} \in A} e^{\beta \cdot Q(\hat{a})}}$$

Update action-values:
$$\Delta Q(a) = \alpha (r_t - Q(a))$$

B. Policy-gradient reinforcement learning

Example: *Gradient-bandit algorithm*

Parameterised policy:
$$\pi^\theta(a) = \frac{e^{\theta_a}}{\sum_{\hat{a} \in A} e^{\theta_{\hat{a}}}}$$

Update parameters:
$$\Delta \theta_a = \begin{cases} \alpha \cdot [1 - \pi^\theta(a)] \cdot r_t & \text{if } a \text{ was chosen} \\ -\alpha \cdot \pi^\theta(a) \cdot r_t & \text{if } a \text{ was unchosen} \end{cases}$$

Figure 1: Update schematics for example value-based and policy-gradient RL algorithms. Shaded diamond nodes denote observable variables, unshaded circular nodes denote latent variables that are internal to the RL agent, and arrows denote dependencies. For simplicity, in these algorithms we do not show the environmental state, which would be an additional (potentially partially) observable variable. **A:** in a value-based RL algorithm (such as the *Q-learning* model presented here), actions (a , chosen from a discrete set A) are a product of the agent’s policy π , which in turn is determined (dotted cyan arrow) by the learned action-values (Q). The update rule for action-values (dashed green arrow) depends on the action-values and received reward (r) at the previous timestep, and only indirectly on the policy. This algorithm has two adjustable parameters: the learning rate α and the softmax inverse temperature β . **B:** a policy-gradient algorithm (such as the gradient-bandit algorithm presented here; see [13]) selects actions according to a parameterised policy π^θ , and updates the parameters θ of this policy directly (dashed magenta arrow; in the gradient-bandit algorithm, θ is a vector of action preferences), without the intermediate step of learning action-values. In the policy-gradient algorithm, by contrast with the value-based algorithm, the size of the update to θ depends more directly on the current policy, since the size of the update to each action preference is scaled by the probability of that action under the policy.

64 In a policy-gradient algorithm, the agent chooses actions according to
65 a parameterized policy, observes the outcomes of these actions, and then

66 adjusts the parameters of its policy so as to increase the probability of ac-
 67 tions associated with more reward and decrease the probability of actions
 68 associated with less reward (i.e., it adjusts the parameters of its policy ac-
 69 cording to the gradient² of reward with respect to the parameters of its
 70 policy). Rather than the indirect algorithm of value-based RL (learning the
 71 values of different actions, and acting by mapping these values into a policy),
 72 policy-gradient RL is a conceptually simpler algorithm that directly adjusts
 73 the policy without troubling with the intermediate latent construct of value.
 74 Consequently, policy-gradient RL is also simpler than value-based RL in im-
 75 plementation as a cognitive model: in its simplest form, policy-gradient RL
 76 requires one adjustable parameter per participant (a learning rate for up-
 77 dates to the policy parameters), compared to two adjustable parameters for
 78 the simplest form of value-based RL (a learning rate for value updates, plus
 79 a policy mapping parameter such as the softmax inverse temperature β)³.

80 In the simplest choice setting, where a subject repeatedly chooses be-
 81 tween a fixed set of actions in a single environmental state, policy-gradient
 82 RL algorithms can, in principle, explain behavior as well as value-based RL
 83 algorithms. For instance, the gradient-bandit algorithm described by Sut-
 84 ton and Barto [13] (based on the REINFORCE algorithm of Williams [23])
 85 uses a softmax policy parameterized by a vector of *preferences* for different
 86 actions. These preferences are reminiscent of values, but unlike values they
 87 are not interpretable as expected discounted future reward; instead, as pol-
 88 icy parameters, they directly determine choice probabilities. For instance,
 89 consider a state with three actions associated with reward magnitudes of
 90 1 unit, 9 units, and 10 units, respectively. A gradient-bandit agent that
 91 learns an optimal policy will show a low preference for the second-best (9-
 92 unit) option relative to the best (10-unit) option, in spite of the fact that

²There exist other policy-optimization algorithms that update the policy without using a gradient, such as trust-region policy optimization [21], but these are beyond the scope of this article.

³In fitting value-based RL models with a softmax choice rule, these two parameters are frequently strongly anticorrelated, leading to difficulty in parameter identifiability [22]. We argue that this non-identifiability is a consequence of a deeper conceptual issue: to explain choices using a value-based algorithm, we must posit, in addition to value-learning, an additional cognitive operation by which learned values are mapped into a behavioral policy. However, for any given choice between different actions, there are infinitely many combinations of underlying action-values and policy-mapping functions that will produce identical choice probabilities. This many-to-one correspondence is the root cause of parameter non-identifiability in value-based RL models, and is avoided in policy-gradient RL.

the *values* of these two actions are relatively similar. As such, value-based and policy-gradient RL have different representations of actions in the environment: whereas value-based algorithms maintain a representation the expected future reward of different actions, policy-gradient algorithms can be thought of as representing actions in terms of which should be taken and which should be avoided. This means that policy-gradient algorithms do not represent the reward amounts associated with different actions. In general, however, this sparser representation does not prevent a policy-gradient RL agent from learning to behave in a manner that maximises its expected future reward.

In more complex choice settings, the policy optimized by the RL agent need not be parameterized by preferences. In environments with a continuous one-dimensional action space, for instance, actions might be selected according to a Gaussian policy (e.g., [24]); in such a case, the parameters of the agent’s policy would control the mean and variance of the probability distribution over actions that is produced by the policy.

3. Behavioral and neural evidence for policy-gradient RL

The greater parsimony of policy-gradient RL is one reason to prefer it over value-based RL as a model of simple operant behavior. A second reason is that, in some decision-making tasks, policy-gradient methods provide a good account of behavioral and neural data that are more difficult to explain with value-based RL models.

3.1. Behavioral evidence for policy-gradient RL

Two behavioral phenomena better explained by policy-gradient RL than value-based RL are context-dependent preference learning and continuous-action learning.

Context-dependent preference learning. When given a choice between two options with equivalent reinforcement histories, but where one option was learned in a rich environmental context (high reward availability) and the other was learned in a lean context (low reward availability), humans and other animals display a marked preference for the lean-context option [25, 26, 27, 28]. That is, animals’ learned preferences between different options are a function not only of each option’s reinforcement history, but also of the environmental context within which each option was experienced. Since action-values in value-based RL are defined cardinally as expected cumulative future reward conditional on taking an action (i.e., action-values are

not contextually modulated by the value of their associated environmental state), this phenomenon is difficult to account for using value-based RL models absent additional post-hoc assumptions regarding the reward function (e.g., [28]).

By contrast, context-dependent preference learning emerges straightforwardly from the principles of policy-gradient RL. In a gradient-bandit algorithm, for instance, because the goal of the agent is to update action-preferences to optimize its policy (rather than to learn action-values), learned preferences for each action are solely a function of whether taking an action within its environmental state will lead to maximization of future reward. As such, a gradient-bandit agent will learn positive preferences for actions that are the best available in a state (i.e., actions that should be taken when possible), and learn negative preferences for actions that are the worst available in a state (i.e., actions that should be avoided). If, after learning, the agent is given a choice between the worst action from a rich state versus the best option from a lean state, it will tend to prefer the lean-state action (because of its positive learned preference as the best action in its state) over the rich-state action (because of its negative learned preference as the worst action in its state), even if the reinforcement histories of the two actions are identical.

Continuous-action learning. In natural environments, behavior frequently involves selecting actions from a *continuous* action space (e.g., controlling a mouse cursor, or driving a car; see [29]). Standard value-based RL models like Q -learning typically perform poorly in such environments, because they operate over a tabular representation of actions in which all actions are equally (dis)similar to one another. This tabular representation leads to inefficient learning in continuous action spaces, since the algorithm cannot generalize between similar actions (e.g., steering a car 10 degrees left is similar to steering 11 degrees left, but dissimilar from steering 40 degrees right). Although there exist continuous extensions of value-based RL involving value-function approximation (e.g., [30]), these approaches have historically focused on continuous state spaces rather than continuous action spaces. Applying value-function approximation to continuous action spaces is not straightforward, because actions need not only be evaluated, but also sampled according to their estimated value. This implies estimation of the whole function (we need evaluate only the current state, but to select an action we need to evaluate all possible actions), which, even for simple choice rules such as those considered here, is computationally intractable for most value-function approximators [31].

168 By contrast, since policy-gradient methods optimise an overall policy,
 169 rather than learning the values of a number of distinct actions, policy-
 170 gradient RL algorithms are straightforwardly applicable to continuous ac-
 171 tion spaces, providing that the functional form of the policy can be sampled
 172 from (e.g., a two-parameter gamma distribution for reaction time choices;
 173 [32]). In learning a continuous action space, a policy-gradient RL agent will
 174 respond to the reinforcement of an action by adjusting the policy such as
 175 to increase the choice probability not only of the chosen action, but also
 176 of similar actions. This leads to efficient generalization of learning across
 177 continuous action spaces.

178 In line with this proposal, recent research in the field of motor control
 179 has suggested that online recalibration processes for motor execution are
 180 well-explained by direct updating of an implicit behavioral policy [33, 34].
 181 For instance, a recent study by Hadjiosif et al. [35] used mirror-reversed
 182 visual feedback in a sensorimotor adaptation task, and found that subjects’
 183 patterns of post-reversal errors were well explained by a model in which
 184 an implicit behavioral policy was updated according to a sensory prediction
 185 error. As such, policy-gradient RL also represents a point of contact between
 186 models of simple operant behavior and models of higher-order motor control.
 187 Indeed, even apparently simple operant behaviours such as discrete choice
 188 can also be conceptualized as involving continuous actions if we consider that
 189 the latency and vigour of these choices constitute a continuously-distributed
 190 action space [32].

191 3.2. Neural evidence for policy-gradient RL

192 *Sensory and action components of midbrain dopamine responses.* Phasic
 193 bursts from midbrain dopamine (DA) neurons have been proposed as a
 194 neural correlate of the reward prediction error (RPE) signal central to RL
 195 theories [36, 11, 37, 38]. This hypothesis regarding the role of dopamine
 196 stands for policy-based RL as well. Indeed, the difference between value-
 197 and policy-based RL frameworks lies not in the dependence on an RPE, but
 198 in what kind of representation is updated by this teaching signal. Decades
 199 of evidence in classical conditioning tasks have demonstrated phasic DA re-
 200 sponses to sensory stimuli that predict the delivery of a reward, consistent
 201 with the learning of a ‘subjective value’ for the predictive stimulus. However,
 202 in instrumental settings, DA activity (both phasic and tonic) has also been
 203 associated with movement signaling and control: phasic responses in dorsal-
 204 striatum-projecting DA neurons can trigger locomotion [39], and movement
 205 initiation and response vigor are bidirectionally modulated by transient ac-
 206 tivity in DA neurons in the substantia nigra pars compacta [40].

207 Action dependence of DA activity is not restricted to the substantia ni-
 208 gra; ventral tegmental area DA responses are attenuated unless the correct
 209 movement is initiated in an operant go/no-go task [41]. More recently, close
 210 inspection of DA responses during very early learning in a classical condi-
 211 tioning task show that the timing of the initiation of licking accounts for
 212 phasic DA response to rewards at this stage, before more stereotypical sen-
 213 sory components of the RPE signal have emerged [42]. These results suggest
 214 that representations of action are, at least in these settings, central to the
 215 formation and update of reward predictions. That is, neural reward predic-
 216 tions in dopaminergic circuits appear to be deeply intertwined with neural
 217 representations of actions themselves. This stands in contrast with simple
 218 value-based RL models, where the learned values of actions are represented
 219 separately and only transformed into a policy as needed at the time of choice
 220 (top panel, Figure 1).

221 *Neural signatures of policy learning.* The basal ganglia are a central locus in
 222 the brain for the initiation and control of movement [43]. Accordingly, much
 223 of the evidence for policy learning in the brain hinges on the interpretation of
 224 neural activity in the striatum during various instrumental tasks. Represent-
 225 tations associated with action-values have been reported in striatal regions
 226 [19, 44, 45]; however, unambiguous identification of these activity patterns
 227 as value signals rather than action preferences or other related constructs
 228 is complicated by the presence of temporal correlations in neural activity
 229 recorded from different brain regions [46].

230 Distinguishing value representations from other arbiters of preference,
 231 such as policy, suffers also from the confound that these quantities are cor-
 232 related in almost all learning models. In fact, specific considerations in task
 233 design are necessary to establish the condition in which predictions from
 234 value- and policy-based accounts of learning can be disambiguated: policy-
 235 based methods predict updates to all (within-context) action probabilities
 236 subsequent to the outcome of a single action, leading to a fundamentally
 237 relative representation of action preference. Indeed, in a task in which in-
 238 formation about the outcome of forgone as well as chosen options is pro-
 239 vided, the BOLD signal in the striatum was consistent with a policy-update
 240 signal, rather than an action-value update [47]. Neural correlates of context-
 241 dependent effects in choice implicate the medial prefrontal cortex and ven-
 242 tral striatum in the relative update of preferences [28], while dopaminergic
 243 error-signals related to counterfactual outcomes have been identified in hu-
 244 man striatum [48]. These findings speak to interactions between chosen and
 245 non-chosen options that are the hallmark of policy-based RL, and suggest

such learning recruits the same neural regions and mechanisms that have previously been closely identified with value-based learning in the brain.

4. The intersection of value and policy

In the preceding sections, we have drawn a sharp contrast between value-based and policy-gradient RL. However, this is a false dichotomy, since an entire class of algorithms—*actor-critic* RL—marries value-based and policy-gradient RL. Although we have not discussed these algorithms here, they form the basis of much state-of-the-art work in computer science applications of reinforcement learning [49, 50, 51], alongside other methods that incorporate components of both value-based and policy-gradient RL (e.g., [52, 53]). In the domain of behavioral science, actor-critic RL has shown great promise as a model of diverse phenomena—including conditioned avoidance, matching behavior in response to variable-interval reinforcement, and mood [54, 55, 56]—that are not readily accounted for by either value-based or policy-gradient models alone. In neuroscience, similarly, it has been suggested that dopaminergic neural activity in the basal ganglia is accounted for well by actor-critic models [57, 58]. As such, actor-critic RL represents a potentially fruitful avenue for developing a general model of operant behavior that may resolve the tensions between value-based and policy-gradient RL that we have reviewed here.

More broadly, our advocating for policy-gradient RL as a minimal model of learning in operant settings does not preclude the possibility that more nuanced representations of outcomes—and even value as defined in RL—are recruited in various behavioral tasks. To be clear, there exist a number of phenomena that appear better explained by value-based and model-based RL models than by policy-gradient RL alone. For instance, reward expectations are clearly formed and exploited in classical conditioning [59, 60, 61], outcome devaluation procedures, suggest that a specific expectation of future outcomes is formed and used to guide choice in an adaptive manner [62], and sensory preconditioning studies show that conditioning can result in the formation of associations between different stimuli, and not only between stimuli and responses or stimuli and rewards [63]. Indeed, theories of model-based learning and decision making rely on specifying the precise interaction between states, values and policies to account for flexible behavior in uncertain and dynamic environments [64, 65]. Once again, our argument is not that the brain only ever uses policy-gradient RL; rather, we suggest that, in modelling data, policy-gradient models should typically be favoured in the interests of parsimony, and that the latent construct of value should

284 only be invoked to explain phenomena that are not explicable under this
285 simpler model.

286 Finally, part of the disjunction between value-based and policy-based RL
287 theories of behavior derives from the historical delineation between operant
288 and classical conditioning paradigms, with the former assumed to pertain
289 to policies for action, and the latter construed as a selective window on
290 value learning. In general, however, the construct of policy is not restricted
291 to discrete choice paradigms; indeed, any engagement in a situation that
292 involves motivationally relevant outcomes will require some targeted regu-
293 lation of movement, which implies the existence of a policy. More expansive
294 consideration of policies that include continuous action spaces [29, 66], the
295 withholding as well as execution of movements [67], and the inclusion of
296 ‘internal’ actions such as the control of attentional focus [68, 69], natu-
297 rally complicate the boundary between operant and classical conditioning,
298 merging perspectives from policy- and value-learning into a more integrated
299 whole.

Box 1: Glossary of Reinforcement-Learning Terms

Policy: a function, often denoted π , that specifies a probability distribution over actions given the agent’s current state. Actions are sampled from the policy at the time of choice. A policy is an essential component of all RL algorithms—including value-based RL algorithms, which must specify a *policy-mapping function* (see below) that computes a policy given a set of action-values.

Value: expected discounted cumulative future reward. In value-based RL, values can be defined both for states ($V^\pi(s)$: the expected future reward associated with being in state s and choosing actions according to the policy π) and actions ($Q^\pi(s, a)$: the expected future reward associated with taking action a in state s , and choosing actions according to the policy π thereafter).

Policy-mapping function: a function that specifies choice probabilities for a set of actions given their estimated values, in value-based RL algorithms. Also known as a ‘choice rule’. Examples include arg max, ϵ -greedy, and softmax.

- arg max: a policy-mapping function that deterministically selects the action with the highest estimated value
- ϵ -greedy: a policy mapping function that selects either the action with the highest estimated value (with probability $1 - \epsilon$) or a random action (with probability ϵ)
- softmax: a policy-mapping function that stochastically selects actions with a probability that increases based on their estimated value relative to the values of alternative options:

$$\pi(a) = \frac{e^{\beta \cdot Q^\pi(s, a)}}{\sum_{\hat{a} \in A} e^{\beta \cdot Q^\pi(s, \hat{a})}} \quad (1)$$

The degree of stochasticity in this mapping is controlled by the inverse temperature parameter β , such that all actions are equally likely when $\beta = 0$ and the softmax function becomes the arg max function as $\beta \rightarrow \infty$

Policy gradient: For RL algorithms that use a parameterized policy, the policy gradient is a vector that indicates how much extra reward the agent expects to receive if it made an incremental change to each of the parameters of its policy (technically, the gradient is the vector of partial derivatives of the expected reward function with respect to policy parameters). The

gradient of a policy is the key variable estimated (or approximated) by policy-gradient RL algorithms, and is used to adjust the parameters of the policy in the direction in which expected reward is expected to increase most steeply. For this reason, policy-gradient algorithms must use a policy that is everywhere differentiable with respect to its parameters.

- [1] J. P. O’Doherty, The problem with value, *Neuroscience & Biobehavioral Reviews* 43 (2014) 259–268.
 - [2] K. J. Miller, A. Shenhav, E. A. Ludvig, Habits without values, *Psychological Review* 126 (2019) 292–311.
- * This paper presents evidence that habit formation—a process previously linked with model-free learning of action-values—may be instead produced by a value-free process in which choosing an action directly strengthens its future choice probability. Although this is not strictly speaking a policy-gradient model, it nevertheless points to the feasibility of explaining operant behavior in terms of modulation of a policy, without recourse to the explanatory construct of value.
- [3] K. Juechems, C. Summerfield, Where does value come from?, *Trends in Cognitive Sciences* 23 (2019) 836–850.
 - [4] G. Suri, J. J. Gross, J. L. McClelland, Value-based decision making: An interactive activation perspective, *Psychological Review* 127 (2020) 153.
 - [5] B. Hayden, Y. Niv, The case against economic values in the brain, 2020. Preprint hosted at PsyArXiv.
 - [6] E. T. Rolls, *Emotion Explained*, Oxford University Press, 2005.
 - [7] A. Rangel, C. Camerer, P. R. Montague, A framework for studying the neurobiology of value-based decision making, *Nature Reviews Neuroscience* 9 (2008) 545–556.
 - [8] M. L. Platt, P. W. Glimcher, Neural correlates of decision variables in parietal cortex, *Nature* 400 (1999) 233–238.
 - [9] D. J. Levy, P. W. Glimcher, The root of all value: a neural common currency for choice, *Current Opinion in Neurobiology* 22 (2012) 1027–1038.
 - [10] P. W. Glimcher, Value-based decision making, in: P. W. Glimcher, E. Fehr (Eds.), *Neuroeconomics*, Elsevier, 2 edition, 2014, pp. 373–391.
 - [11] W. Schultz, P. Dayan, P. R. Montague, A neural substrate of prediction and reward, *Science* 275 (1997) 1593–1599.

- [12] J. P. O’Doherty, P. Dayan, K. Friston, H. Critchley, R. J. Dolan, Temporal difference models and reward-related learning in the human brain, *Neuron* 38 (2003) 329–337.
 - [13] R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 2018.
 - [14] G. Mongillo, H. Shteingart, Y. Loewenstein, The misbehavior of reinforcement learning, *Proceedings of the IEEE* 102 (2014) 528–541.
- * This paper provides an accessible overview of the algorithmic differences between value-based and policy-gradient reinforcement learning, as well as reviewing behavioral and neural evidence in favor of each model family.
- [15] P. Dayan, L. F. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, The MIT Press, 2001.
 - [16] K. Wunderlich, A. Rangel, J. P. O’Doherty, Neural computations underlying action-based decision making in the human brain 106 (2009) 17199–17204.
 - [17] M. Ito, K. Doya, Validation of decision-making models and analysis of decision variables in the rat basal ganglia 29 (2009) 9861–9874.
 - [18] N. D. Daw, S. J. Gershman, B. Seymour, P. Dayan, R. J. Dolan, Model-based influences on humans’ choices and striatal prediction errors, *Neuron* 69 (2011) 1204–1215.
 - [19] X. Cai, S. Kim, D. Lee, Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice, *Neuron* 69 (2011) 170–182.
 - [20] J. Vandekerckhove, D. Matzke, E.-J. Wagenmakers, et al., Model comparison and the principle of parsimony, in: J. R. Busemeyer, Z. Wang, J. T. Townsend, A. Eidels (Eds.), *Oxford Handbook of Computational and Mathematical Psychology*, 2015, pp. 300–319.
 - [21] J. Schulman, S. Levine, P. Abbeel, M. Jordan, P. Moritz, Trust region policy optimization, in: F. Bach, D. Blei (Eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, PMLR, Lille, France, 2015, pp. 1889–1897.

- [22] I. C. Ballard, S. M. McClure, Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models, *Journal of Neuroscience Methods* 317 (2019) 37–44.
- [23] R. J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Machine Learning* 8 (1992) 229–256.
- [24] T. Degris, P. M. Pilarski, R. S. Sutton, Model-free reinforcement learning with continuous action in practice, in: *2012 American Control Conference (ACC)*, IEEE, pp. 2177–2182.
- [25] L. Pompilio, A. Kacelnik, State-dependent learning and suboptimal choice: When starlings prefer long over short delays to food, *Animal Behaviour* 70 (2005) 571–578.
- [26] L. Pompilio, A. Kacelnik, S. T. Behmer, State-dependent learned valuation drives choice in an invertebrate, *Science* 311 (2006) 1613–1615.
- [27] J. Aw, R. Holbrook, T. Burt de Perera, A. Kacelnik, State-dependent valuation learning in fish: Banded tetras prefer stimuli associated with greater past deprivation, *Behavioural Processes* 81 (2009) 333–336.
- [28] S. Palminteri, M. Khamassi, M. Joffily, G. Coricelli, Contextual modulation of value signals in reward and punishment learning, *Nature Communications* 6 (2015) 1–14.

* This paper uses a carefully controlled task design to show that human participants display behavior consistent with context-sensitive preference learning. Specifically, participants’ preferences for an action were modulated by the value of the action’s context. One result of this was that avoidance responses accrue positive preference in contexts with a negative expected value.

- [29] S. B. M. Yoo, B. Y. Hayden, J. M. Pearson, Continuous decisions, *Philosophical Transactions of the Royal Society B* 376 (2021) 20190664.

* This paper points out that naturalistic decision-making frequently involves continuous action spaces and prolonged choice windows, and offers a cogent argument that many common laboratory tasks and models are poorly suited for studying these features of decision making. It suggests that policy-gradient reinforcement learning and control theory provide a solid foundation for studying continuous decisions.

- [30] K. Doya, Reinforcement learning in continuous time and space, *Neural Computation* 12 (2000) 219–245.
 - [31] J. C. Santamaria, R. S. Sutton, A. Ram, Experiments with reinforcement learning in problems with continuous state and action spaces, *Adaptive Behavior* 6 (1997) 163–217.
 - [32] Y. Niv, The effects of motivation on habitual instrumental behavior, Ph.D. thesis, The Hebrew University of Jerusalem, 2007.
 - [33] A. M. Haith, J. W. Krakauer, Model-based and model-free mechanisms of human motor learning, in: *Progress in motor control*, Springer, 2013, pp. 1–21.
 - [34] S. D. McDougle, R. B. Ivry, J. A. Taylor, Taking aim at the cognitive side of learning in sensorimotor adaptation tasks, *Trends in Cognitive Sciences* 20 (2016) 535–544.
 - [35] A. M. Hadjiosif, J. W. Krakauer, A. M. Haith, Did we get sensorimotor adaptation wrong? implicit adaptation as direct policy updating rather than forward-model-based learning, *Journal of Neuroscience* 41 (2021) 2747–2761.
- * Using a motor control task with mirror-reversed visual feedback, this paper provides behavioral evidence that implicit motor adaptation is more consistent with the direct learning of a behavioral policy than with the use of a predictive forward model to plan movements.
- [36] M. Watabe-Uchida, N. Eshel, N. Uchida, Neural circuitry of reward prediction error, *Annual Review of Neuroscience* 40 (2017) 373–394.
 - [37] N. Eshel, M. Bukwich, V. Rao, V. Hemmelder, J. Tian, N. Uchida, Arithmetic and local circuitry underlying dopamine prediction errors, *Nature* 525 (2015) 243–246.
 - [38] M. R. Roesch, D. J. Calu, G. Schoenbaum, Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards, *Nature neuroscience* 10 (2007) 1615.
 - [39] M. W. Howe, D. A. Dombeck, Rapid signalling in distinct dopaminergic axons during locomotion and reward, *Nature* 535 (2016) 505–510.

- [40] J. A. da Silva, F. Tecuapetla, V. Paixão, R. M. Costa, Dopamine neuron activity before action initiation gates and invigorates future movements, *Nature* 554 (2018) 244–248.
 - [41] E. C. Syed, L. L. Grima, P. J. Magill, R. Bogacz, P. Brown, M. E. Walton, Action initiation shapes mesolimbic dopamine encoding of future rewards, *Nature Neuroscience* 19 (2016) 34–36.
- * The authors show that RPE-like changes in DA concentration in the nucleus accumbens (measured using fast-scan cyclic voltammetry) are only evident during the anticipation or execution of overt actions in a go/no-go task, suggesting the roles of dopamine in signaling prediction errors and in motivating purposeful movement are interrelated.
- [42] L. T. Coddington, J. T. Dudman, The timing of action determines reward prediction signals in identified midbrain dopamine neurons, *Nature neuroscience* 21 (2018) 1563.
- ** Demonstrates movement initiation produces phasic responses in midbrain dopamine neurons early in learning during a classical conditioning task in mice, suggesting action-related components of neural prediction error signals both precede and are independent of responses related to learning about reward-predictive sensory cues.
- [43] A. Klaus, J. Alves da Silva, R. M. Costa, What, if, and when to move: basal ganglia circuits and self-paced action initiation, *Annual review of neuroscience* 42 (2019) 459–483.
 - [44] K. Samejima, Y. Ueda, K. Doya, M. Kimura, Representation of action-specific reward values in the striatum, *Science* 310 (2005) 1337–1340.
 - [45] T. H. FitzGerald, K. J. Friston, R. J. Dolan, Action-specific value signals in reward-related regions of the human brain, *Journal of Neuroscience* 32 (2012) 16417–16423.
 - [46] L. Elber-Dorozko, Y. Loewenstein, Striatal action-value neurons reconsidered, *eLife* (2018) 32.

* Challenges the popular idea that populations of striatal neurons unambiguously reflect action values, showing that previous analyses that putatively isolate correlates of action value in neural

activity cannot dissociate value representations from alternative representations (e.g., policy).

- [47] J. Li, N. D. Daw, Signals in human striatum are appropriate for policy update rather than value prediction, *Journal of Neuroscience* 31 (2011) 5504–5511.
- [48] K. T. Kishida, I. Saez, T. Lohrenz, M. R. Witcher, A. W. Laxton, S. B. Tatter, J. P. White, T. L. Ellis, P. E. Phillips, P. R. Montague, Subsecond dopamine fluctuations in human striatum encode superposed error signals about actual and counterfactual reward, *Proceedings of the National Academy of Sciences* 113 (2016) 200–205.
- [49] J. Schulman, P. Moritz, S. Levine, M. Jordan, P. Abbeel, High-dimensional continuous control using generalized advantage estimation, *arXiv preprint arXiv:1506.02438* (2015).
- [50] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, K. Kavukcuoglu, Asynchronous methods for deep reinforcement learning, in: *International Conference on Machine Learning*, pp. 1928–1937.
- [51] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, et al., Grandmaster level in StarCraft II using multi-agent reinforcement learning, *Nature* 575 (2019) 350–354.
- [52] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, M. Riedmiller, Deterministic policy gradient algorithms, in: *Proceedings of the 31st International Conference on Machine Learning*, Beijing, China, p. 9.
- [53] O. Nachum, M. Norouzi, K. Xu, D. Schuurmans, Bridging the gap between value and policy based reinforcement learning, in: *Advances in Neural Information Processing Systems*, pp. 2775–2785.
- [54] Y. Sakai, T. Fukai, The actor-critic learning is behind the matching law: matching versus optimal behaviors, *Neural Computation* 20 (2008) 227–251.
- [55] T. V. Maia, Two-factor theory, the actor-critic model, and conditioned avoidance, *Learning & Behavior* 38 (2010) 50–67.
- [56] D. Bennett, G. Davidson, Y. Niv, A model of mood as integrated advantage, *Psychological Review* (in press).

- [57] A. G. Barto, Adaptive critics and the basal ganglia, in: *Models of Information Processing in the Basal Ganglia*, The MIT Press, Cambridge, MA, 1994.
- [58] D. Joel, Y. Niv, E. Ruppín, Actor-critic models of the basal ganglia: New anatomical and computational perspectives, *Neural Networks* 15 (2002) 535–547.
- [59] M. S. Fanselow, K. M. Wassum, The origins and organization of vertebrate Pavlovian conditioning, *Cold Spring Harbor Perspectives in Biology* 8 (2016).
- [60] N. T. Lichtenberg, Z. T. Pennington, S. M. Holley, V. Y. Greenfield, C. Cepeda, M. S. Levine, K. M. Wassum, Basolateral amygdala to orbitofrontal cortex projections enable cue-triggered reward expectations, *Journal of Neuroscience* 37 (2017) 8374–8384.
- [61] R. A. Rescorla, A. R. Wagner, A theory of pavlovian conditioning, *Classical Conditioning II: Current Theory and Research* (1971).
- [62] B. W. Balleine, A. Dickinson, Goal-directed instrumental action: contingency and incentive learning and their cortical substrates, *Neuropharmacology* 37 (1998) 407–419.
- [63] M. J. Sharpe, C. Y. Chang, M. A. Liu, H. M. Batchelor, L. E. Mueller, J. L. Jones, Y. Niv, G. Schoenbaum, Dopamine transients are sufficient and necessary for acquisition of model-based associations, *Nature Neuroscience* 20 (2017) 735–742.
- [64] A. J. Langdon, M. J. Sharpe, G. Schoenbaum, Y. Niv, Model-based predictions for dopamine, *Current Opinion in Neurobiology* 49 (2018) 1–7.
- [65] P. Dayan, K. C. Berridge, Model-based and model-free Pavlovian reward learning: Revaluation, revision, and revelation, *Cognitive, Affective, & Behavioral Neuroscience* 14 (2014) 473–492.
- [66] Y. Niv, N. Daw, P. Dayan, How fast to work: Response vigor, motivation and tonic dopamine, in: Y. Weiss, B. Schölkopf, J. Platt (Eds.), *Advances in Neural Information Processing Systems*, volume 18, MIT Press, 2006.

- [67] A. G. Collins, M. J. Frank, Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive, *Psychological Review* 121 (2014) 337.
- [68] Y. C. Leong, A. Radulescu, R. Daniel, V. DeWoskin, Y. Niv, Dynamic interaction between reinforcement learning and attention in multidimensional environments, *Neuron* 93 (2017) 451–463.
- [69] A. Radulescu, Y. Niv, I. Ballard, Holistic reinforcement learning: the role of structure and attention, *Trends in Cognitive Sciences* 23 (2019) 278–292.