

Reinforcement Learning, Fast and Slow

Matthew Botvinick^{1,2}, Sam Ritter^{1,3}, Jane X. Wang¹
Zeb Kurth-Nelson^{1,2}, Charles Blundell¹, Demis Hassabis^{1,2}

¹DeepMind, UK, ²University College London, UK ³Princeton University, USA

March 22, 2019

Abstract

Deep reinforcement learning (RL) methods have driven impressive advances in artificial intelligence in recent years, exceeding human performance in domains ranging from Atari to Go to no-limit poker. This progress has drawn the attention of cognitive scientists interested in understanding human learning. However, the concern has been raised that deep RL may be too sample inefficient – that is, it may simply be too slow – to provide a plausible model of how humans learn. In the present review, we counter this critique by describing recently developed techniques that allow deep RL to operate more nimbly, solving problems much more quickly than previous methods. Although these techniques were developed in an AI context, we propose that they may have rich implications for psychology and neuroscience. A key insight, arising from these AI methods, concerns the fundamental connection between fast RL and slower, more incremental forms of learning.

Powerful but slow: The first wave of Deep RL

Over just the past few years, revolutionary advances have occurred in artificial intelligence research, where a resurgence in neural network or ‘deep learning’ methods [1, 2] has fueled breakthroughs in image understanding [3, 4], natural language processing [5, 6], and many other areas. These developments have attracted growing interest from psychologists, psycholinguists and neuroscientists, curious about whether developments in AI might suggest new hypotheses concerning human cognition and brain function [7, 8, 9, 10, 11].

One area of AI research that appears particularly inviting from this perspective is *deep reinforcement learning* (Box 1). Deep RL marries neural network modeling with reinforcement learning, a set of methods for learning from rewards and punishments rather than from more explicit instruction [12]. After decades as an aspirational rather than practical idea, deep RL has within the last five years exploded into one of the most intense areas of AI research, generating super-human performance in tasks from video games [13] to poker [14] to multiplayer contests [15] to complex board games including go and chess [16, 17, 18, 19].

Beyond its inherent interest as an AI topic, deep RL would appear to hold special interest for psychology and neuroscience. The mechanisms that drive learning in deep RL were originally inspired by animal conditioning research [20], and are believed to relate closely to neural mechanisms for reward-based learning centering on dopamine [21]. At the same time, deep RL leverages neural networks to learn powerful representations that support generalization and transfer, key abilities of biological brains. Given these connections, deep RL would appear to offer a rich source of ideas and hypotheses for researchers interested in human and animal learning, both at the behavioral and neuroscientific levels. And indeed, researchers have started to take notice [7, 8].

At the same time, commentary on the first wave of deep RL research has also sounded a note of caution. On first blush it appears that deep RL systems learn in a fashion quite different from humans. The hallmark of this difference, it has been argued, lies in the sample efficiency of human learning versus deep RL. *Sample efficiency* refers to the amount of data required for a learning system to attain any chosen target level of performance. On this measure, the initial wave of deep RL systems indeed appear drastically different from human learners. To attain expert human-level performance on tasks such as Atari video games or chess, deep RL systems have required many orders of magnitude more training

data than human experts themselves [22]. In short, deep RL, at least in its initial incarnation, appears much too slow to offer a plausible model for human learning. Or so the argument has gone [23, 24].

The critique is indeed applicable to the first wave of deep RL methods, reported beginning around 2013 [e.g. 25]. However, even in the short time since then, important innovations have occurred in deep RL research, which show how the sample efficiency of deep RL can be dramatically increased. These methods mitigate the original demands made by deep RL for huge amounts of training data, effectively allowing deep RL to be fast. The emergence of these computational techniques revives deep RL as a candidate model of human learning and a source of insight for psychology and neuroscience.

In the present review, we consider two key deep RL methods, which mitigate the sample efficiency problem: *episodic deep RL* and *meta-reinforcement learning*. We examine how these techniques enable fast deep RL, and consider their potential implications for psychology and neuroscience.

Sources of slowness of deep RL

A key starting point for considering techniques for fast RL is to examine why initial methods for deep RL were in fact so slow. Here, we describe two of the primary sources of sample inefficiency. At the end of the paper, we will circle back to examine how the constellations of issues described by these two concepts are in fact connected.

The first source of slowness in deep RL is the requirement for incremental parameter adjustment. Initial deep RL methods (which are still very widely used in AI research) employed gradient descent to sculpt the connectivity of a deep neural network mapping from perceptual inputs to action outputs (see Box 1). As has been widely discussed not only in AI but also in psychology [26], the adjustments made during this form of learning must be small, in order to maximize generalization [27] and avoid overwriting the effects of earlier learning (an effect sometimes referred to as catastrophic interference). This demand for small step-sizes in learning is one source of slowness in the methods originally proposed for deep RL.

A second source is weak inductive bias. A basic lesson of learning theory is that any learning procedure necessarily faces a bias-variance trade-off: The stronger the initial assumptions the learning procedure makes about the patterns to be learned — that is, the stronger the initial inductive bias of the learning procedure — the less data will be required for learning to be accomplished (assuming the initial inductive bias matches what’s in the data!). A learning procedure with weak inductive bias will be able to master a wider range of patterns (greater variance), but will in general be less sample efficient [28]. In effect, strong inductive bias is what allows fast learning. A learning system that only considers a narrow range of hypotheses when interpreting incoming data will, perforce, hone in on the correct hypothesis more rapidly than a system with weaker inductive biases (again, assuming the correct hypothesis falls within that narrow initial range). Importantly, generic neural networks are extremely low-bias learning systems; they have many parameters (connection weights) and are capable of using these to fit a wide range of data. As dictated by the bias-variance trade-off, this means that neural networks, in the generic form employed in the first deep RL models (see Box 1) tend to be sample-inefficient, requiring large amounts of data to learn.

Together, these two factors — incremental parameter adjustment and weak inductive bias — explain the slowness of first-generation deep RL models. However, subsequent research has made clear that both of these factors can be mitigated, allowing deep RL to proceed in a much more sample-efficient manner. In what follows, we consider two specific techniques, one of which confronts the incremental parameter adjustment problem, and the other of which confronts the problem of weak inductive bias. In addition to their implications within the AI field, both of these AI techniques bear suggestive links with psychology and neuroscience, as we shall detail.

Episodic deep RL: Fast learning through episodic memory

If incremental parameter adjustment is one source of slowness in deep RL, then one way to learn faster might be to avoid such incremental updating. Naively increasing the learning rate governing gradient descent optimization leads to the problem of catastrophic interference. However, recent research shows that there is another way to accomplish the same goal, which is to keep an explicit record of past events, and use this record directly as a point of reference in making new decisions. This idea, referred to as episodic RL [29, 30], parallels ‘non-parametric’ approaches in machine learning [28] and resembles ‘instance-’ or ‘exemplar-based’ theories of learning in psychology [31, 32]. When a new situation is

encountered and a decision must be made concerning what action to take, the procedure is to compare an internal representation of the current situation with stored representations of past situations. The action chosen is then the one associated with the highest value, based on the outcomes of the past situations that are most similar to the present. When the internal state representation is computed by a multi-layer neural network, we refer to the resulting algorithm as “episodic deep RL”. A more detailed explanation of the mechanics of episodic deep RL is presented in Box 2.

In episodic deep RL, unlike the standard incremental approach, the information gained through each experienced event can be leveraged immediately to guide behavior. However, whereas episodic deep RL is able to go ‘fast’ where earlier methods for deep RL went ‘slow,’ there is a twist to this story: Episodic deep RL’s fast learning depends critically on slow incremental learning. This is the gradual learning of the connection weights that allows the system to form useful internal representations or embeddings of each new observation. The format of these representations is itself learned through experience, using the same kind of incremental parameter updating that forms the backbone of standard deep RL. Ultimately, the speed of episodic deep RL is enabled by this slower form of learning. That is, fast learning is enabled by slow learning.

This dependence of fast learning on slow learning is no coincidence. As we will argue below, it is a fundamental principle, applicable to psychology and neuroscience no less than AI. Before turning to a consideration of this general point, however, we examine its role in the second recently developed AI technique for rapid deep RL: meta-reinforcement learning.

Meta-reinforcement learning: Speeding up deep RL by learning to learn

As discussed earlier, a second key source of slowness in standard deep RL, alongside incremental updating, is weak inductive bias. As formalized in the idea of the bias-variance tradeoff, fast learning requires the learner to go in with a reasonably sized set of hypotheses concerning the structure of the patterns that it will face. The narrower the hypothesis set, the faster learning can be. However, as foreshadowed earlier, there is a catch: A narrow hypothesis set will only speed learning if it contains the correct hypothesis. While strong inductive biases can accelerate learning, they will only do so if the specific biases the learner adopts happen to fit with the material to be learned. As a result of this, a new learning problem arises: How can the learner know what inductive biases to adopt?

One natural answer to this question is to draw on past experience. Of course, this self-evidently occurs all the time in daily life. Consider, for example, the task of learning to use a new smart-phone. In this context, one’s past experience with smart-phones and other related devices will inform one’s hypotheses concerning the way the new phone should work, and will guide exploration of the phone’s operation. These initial hypotheses correspond to the ‘bias’ in the bias-variance tradeoff, and they are responsible for the ability to quickly learn how to use the new phone. A learner without these biases (i.e., with higher ‘variance’) would consider a wider range of hypotheses about the phone’s operation, at the expense of learning speed.

The leveraging of past experience to accelerate new learning is referred to in machine learning as meta-learning [33]. However, not surprisingly, the idea originates from psychology, where it has been called ‘learning to learn.’ In the first paper to use this term, Harlow [34] presented an experiment that captures the principle neatly. Here, monkeys were presented with two unfamiliar objects, and permitted to grab one of them. Beneath lay either a food reward or an empty well. The objects were then placed before the animal again, possibly left-right reversed, and the procedure repeated, and went on doing so for a total of six rounds. Two new and unfamiliar objects were then substituted in, and another six trials ensued with these objects. Then another pair of objects, and so forth. Across many object pairs, the animals were able to figure out that a simple rule always held: One object yielded food and the other did not, regardless of left-right position. When presented with a new pair of objects, Harlow’s monkeys were able to learn in one shot which the preferable object was, a simple but vivid example of learning to learn (see Box 3).

Returning now to AI, recent work has shown how learning to learn can be leveraged to speed up learning in deep RL. This general idea has been implemented in a variety of ways [35, 36]. However, one approach that has particular relevance to neuroscience and psychology was proposed simultaneously by Wang [37] and Duan [38] and their colleagues. Here, a recurrent neural network is trained on a series of interrelated RL tasks. The weights in the network are adjusted very slowly, so they can absorb what is common across tasks, but cannot change fast enough to support the solution of any single task. In




Figure 1: **Schematic of meta-reinforcement learning, illustrating the inner and outer loops of training.** The outer loop trains the parameter weights θ , which determine the inner-loop learner (“Agent”, instantiated by a recurrent neural network) that interacts with an environment for the duration of the episode. Every cycle of the outer loop, a new environment is sampled from a distribution of environments, which share some common structure.

this setting, something rather remarkable occurs. The activity dynamics of the recurrent network come to implement their own separate RL algorithm, which ‘takes responsibility’ for quickly solving each new task, based on knowledge accrued from past tasks (Figure 1). Effectively, one RL algorithm gives birth to another, and hence the moniker ‘meta-reinforcement learning’.

As with episodic deep RL, meta-reinforcement learning again involves an intimate connection between fast and slow learning. The connections in the recurrent network are updated slowly across tasks, allowing general principles that span tasks to be ‘built into’ the dynamics of the recurrent network. The resulting network dynamics implement a new learning algorithm which can solve new problems quickly, because they have been endowed with useful inductive biases by the underlying process of slow learning (see Box 3). Once again, fast learning arises from, and is enabled by, slow learning.

Episodic meta-reinforcement learning

Importantly, the two techniques we have discussed above are not mutually exclusive. Indeed, recent work has explored an approach to integrating meta-learning and episodic control, capitalizing on their complementary benefits [39, 40]. In *episodic meta-RL*, meta-learning occurs within a recurrent neural network, as described in the previous section and Box 3. However, superimposed on this is an episodic memory system, the role of which is to *reinstat*e patterns of activity in the recurrent network. As in episodic deep RL, the episodic memory catalogues a set past events, which can be queried based on the current context. However, rather than linking contexts with value estimates, episodic meta-RL links them with stored activity patterns from the recurrent network’s internal or hidden units. These patterns are important because, through meta-RL, they come to summarize what the agent has learned from interacting with individual tasks (see Box 3 for details). In episodic meta-RL, when the agent encounters a situation that appears similar to one encountered in the past, it reinstates the hidden activations from the previous encounter, allowing previously learned information to immediately influence the current policy. In effect, episodic memory allows the system to recognize previously encountered tasks, retrieving stored solutions.

Through simulation work in bandits and navigation tasks, Ritter et al. [39] showed that episodic meta-RL, just like ‘vanilla’ meta-RL, learns strong inductive biases that enable it to rapidly solve novel tasks. More importantly, when presented with a previously encountered task, episodic meta-RL immediately retrieves and reinstates the solution it previously discovered, avoiding the need to re-explore. On the first encounter with a new task, the system benefits from the rapidity of meta-RL; on the second and later encounters, it benefits from the one-shot learning ability conferred by episodic control.

Implications for neuroscience and psychology

As we discussed at the outset, the problem of sample inefficiency has been adduced as a reason for questioning the relevance of deep RL to learning in humans and other animals [23, 24]. One important implication of episodic deep RL and meta-RL, from the point of view of psychology and neuroscience, is that they undermine this argument, by showing that deep RL in fact need not be slow. This demonstration goes some distance toward defending deep RL as a potential model of human and animal learning. Beyond this general point, however, the specifics of episodic deep RL and meta-RL also point to interesting new hypotheses in psychology and neuroscience.

To begin with episodic deep RL, we have already noted the interesting connection between this and classical instance-based models of human memory, where cognition operates over stored information about specific previous observations [31, 32]. Episodic RL offers a possible account for how instance-based processing might subserve reward-driven learning. Intriguingly, recent work on reinforcement learning in animals and humans has increasingly emphasized the potential contribution of episodic memory, with evidence accruing that estimates of state and action value are based on retrieved memories for specific past action-outcome observations [41, 42, 43, 44]. Episodic deep RL provides a forum for exploring how this general principle might scale to rich, high-dimensional sequential learning problems. Perhaps more fundamentally, it highlights the important role that representation learning and metric learning might play in RL based on episodic memory. Episodic deep RL thus suggests that it may be fruitful to investigate the way that fast episodic RL in humans and animals may interact with and depend upon slower learning processes. While this link between fast and slow learning has been engaged in work on memory per se, including work on multiple memory systems [26, 45, 46, 47], its role in reward-based learning has not been deeply explored [although see 48].

Moving to meta-reinforcement learning, this framework also has interesting potential implications for psychology and neuroscience. Indeed, Wang and colleagues [49] have proposed a very direct mapping from the elements of meta-RL to neural structures and functions. Specifically, they propose that slow, dopamine-driven synaptic change may serve to tune the activity dynamics of prefrontal circuits, in such a way that the latter come to implement an independent set of learning procedures (Box 4). Through a set of computer simulations, Wang and colleagues [49] demonstrated how meta-RL, interpreted in this way, can account for a diverse range of empirical findings from the behavioral and neurophysiology literatures (see Box 4).

Equally direct links connect episodic meta-RL with psychology and neuroscience. Indeed, the reinstatement mechanism involved in episodic meta-RL was directly inspired by neuroscience data indicating that episodic memory circuits can serve to reinstate patterns of activation in cerebral cortex, including areas supporting working memory [see 40]. Ritter and colleagues [39] show how such a function could itself be configured through reinforcement learning, giving rise to a system that can strategically reinstate information about tasks encountered earlier [see also 50, 51, 52]. In addition to the initial inspiration it draws from neuroscience, this work links back to biology by providing a parsimonious explanation for recently reported interactions between episodic and model-based control in human learning [53]. On a broader level, the work reported by Ritter and colleagues [39] provides an illustration of how meta-learning can operate over systems involving multiple memory systems [26], slowly tuning these systems and their interactions in a way that allows them to collectively support rapid learning.

Fast and slow RL: Broader implications

In discussing both episodic RL and meta-RL, we have emphasized the role of ‘slow’ learning in enabling fast, sample-efficient learning. In meta-RL, as we have seen, the role of slow, weight-based learning is to establish inductive biases that can guide inference and thus support rapid adaptation to new tasks. The role of slow, incremental learning in episodic RL can be viewed in related terms. Episodic RL inherently depends on judgments concerning resemblances between situations or states. Slow learning shapes the way that states are internally represented, and thus puts in place a set of inductive biases concerning which states are most closely related.

Looking at episodic RL more closely, one might also say that there is an inductive bias built into the learning architecture. In particular, episodic RL inherently assumes a kind of smoothness principle: Similar states will generally call for similar actions. Rather than being learned, this inductive bias is wired into the structure of the learning system that defines episodic RL. In current AI parlance, this is a case of ‘architectural’ or ‘algorithmic bias,’ in contradistinction to ‘learned bias,’ as seen in meta-RL.

A wide range of current AI research is focused on finding useful inductive biases to speed learning, whether this is accomplished via learning or through the direct hand-design of architectural or algorithmic biases. Indeed, the latter strategy has itself been largely responsible for the current resurgence of neural networks in AI. Convolutional neural networks, which provided the original impulse for this resurgence, build in a very specific architectural bias, connected with translational invariance in image recognition. However, over the past few years, an increasingly large portion of AI research has been focused, either explicitly or implicitly, on the problem of inductive bias [see, e.g., 54, 55, 56].

At a high level, these developments of course parallel some long-standing concerns in psychology. As we have already noted, the idea that inductive biases may be acquired through learning in fact originally derives from psychology [34], and has remained an intermittent topic of psychological research [see, e.g., 23]. However, meta-learning in neural networks may provide a new context to explore the mechanisms and dynamics of such learning-to-learn processes, especially in the reinforcement learning setting.

Psychology, and in particular developmental psychology, has also long considered the possibility that certain inductive biases are ‘built in,’ that is, innately specified [57]. However, the more mechanistic notion of architectural bias, and of biases built into neural network learning algorithms, have been less widely considered [although see 58, 59, 60]. Here again, current methods for deep learning and deep RL provide a toolkit that may be helpful in further exploration, perhaps, for example, in considering the implications of growing information about brain connectivity [61, 62, 63].

It is worth noting that, whereas AI work draws a sharp distinction between inductive biases that are acquired through learning and biases that are ‘wired in’ by hand, in a biological context a more general, unifying perspective is available. Specifically, one can view architectural and algorithmic biases as arising through a different learning process, driven by evolution. Here, evolution is the ‘slow’ learning process, gradually sculpting architectural and algorithmic biases that allow faster lifetime learning. On this account, meta-learning plays out not only within a single lifespan, but also over evolutionary time. Interestingly, this view implies that evolution does not select for a truly ‘general-purpose’ learning algorithm, but for an algorithm which exploits the regularities in the particular environments in which brains have evolved. In this context, recent developments in AI may once again prove handy in exploring implications for neuroscience and psychology. Recent AI work has delved increasingly into methods for building agent architectures, as well as reward functions, through evolutionary algorithms, inspired by natural selection [15, 64, 65, 66, 67].

Whether it focuses on hand-engineering or evolution, AI work on architectural and algorithmic bias furnishes us with a new laboratory for investigating how evolution might sculpt the nervous system to support efficient learning. Possibilities suggested by AI research include constraints on the initial pattern of network connectivity [36]; shaping of synaptic learning rules [35, 68, 69]; and factors encouraging the emergence of disentangled or compositional representations [54, 70, 71] and internal predictive models [51].

Such work, when viewed through lens of psychology, neuroscience, and evolutionary and developmental studies, yields a picture at which learning is operating at many time-scales at once – literally running from millennia to milliseconds – with slower time-scales yielding biases that enable faster learning at levels above, with all of this playing out across evolutionary, developmental and diurnal time-scales, following a trajectory that is strongly influenced by the structure of the environment [see 72]. From this perspective, evolution shapes architecture and algorithm, which embed inductive biases; these then shape lifetime learning, which itself develops further inductive biases based on experience. Within this picture, evolution appears as the ‘slowest’ learning process in a cascade, with each layer supporting a next, ‘faster’ level by serving inductive biases to that next level.

As many readers will recognize, this meta-learning-based perspective is not entirely new to cognitive science. As pointed out earlier, the idea of learning-to-learn has held a place in psychology for decades [34, 73], and hierarchical Bayesian models of learning have long emphasized closely related ideas, with prior probabilities providing the medium for inductive biases [23, 74, 75, 76, 77]. Research on biological evolution has also touched on related ideas [66]. Despite these precedents, the recent surge in AI research investigating the role of slow and fast learning in neural networks, and in reinforcement learning, provides a genuinely new set of perspectives, and a new set of opportunities.

Concluding remarks and future directions

The rapidly developing field of deep reinforcement learning holds great interest for psychology and neuroscience, given its integrated focus on representation learning and goal-directed behavior. In the present review, we have described recently emerging forms of deep RL that overcome the apparent

problem of sample inefficiency, allowing deep RL to work ‘fast.’ These techniques not only reinforce the potential relevance of deep RL to psychology and neuroscience, countering some recent skeptical assessments; they enrich the potential connections by establishing connections to themes such as episodic memory and learning to learn. Furthermore, ideas arising from deep RL research increasingly suggest concrete and specific directions for new research in psychology and neuroscience.

As we have stressed, a key implication of recent work on sample-efficient deep RL is that where fast learning occurs, it necessarily relies on slow learning, which establishes the representations and inductive biases that enable fast learning. This computational dialectic provides a theoretical framework for studying multiple memory systems in the brain, as well as their evolutionary origins. However, human learning likely involves multiple interacting processes, beyond those discussed in this review, and we imagine that any deep RL model will need to integrate all of these in order to approach human-like learning. At a broader level, understanding the relationship between fast and slow in reinforcement learning provides a compelling, organizing challenge for psychology and neuroscience. Indeed, this may be a key area where AI, neuroscience and psychology can synergistically interact, as has always been the ideal in cognitive science.

Outstanding questions

- Can AI methods for sample-efficient deep RL scale to the kinds of rich task environments humans cope with? In particular, can these methods engender rich abstractions of the kind that underlie human intelligence? What kind of training environments might be necessary for this?
- Does flexible, sample-efficient human learning arise from mechanisms related to those currently being explored in AI? If so, what is their neural implementation? Does something like gradient descent learning – so important for current AI techniques – occur in the brain, or is the same role played by some different mechanism?
- What inductive biases are most important for learning in the environment that human learners inhabit? To what extent were these acquired through evolution and genetically or developmentally specified, and to what extent are they acquired through learning?
- One thing that makes human learners so efficient is that we are active, strategic information-seekers. What are the principles that structure and motivate human exploration, and how can we replicate those in AI systems?

Acknowledgements

The authors were funded by DeepMind.

Box 1: Deep reinforcement learning

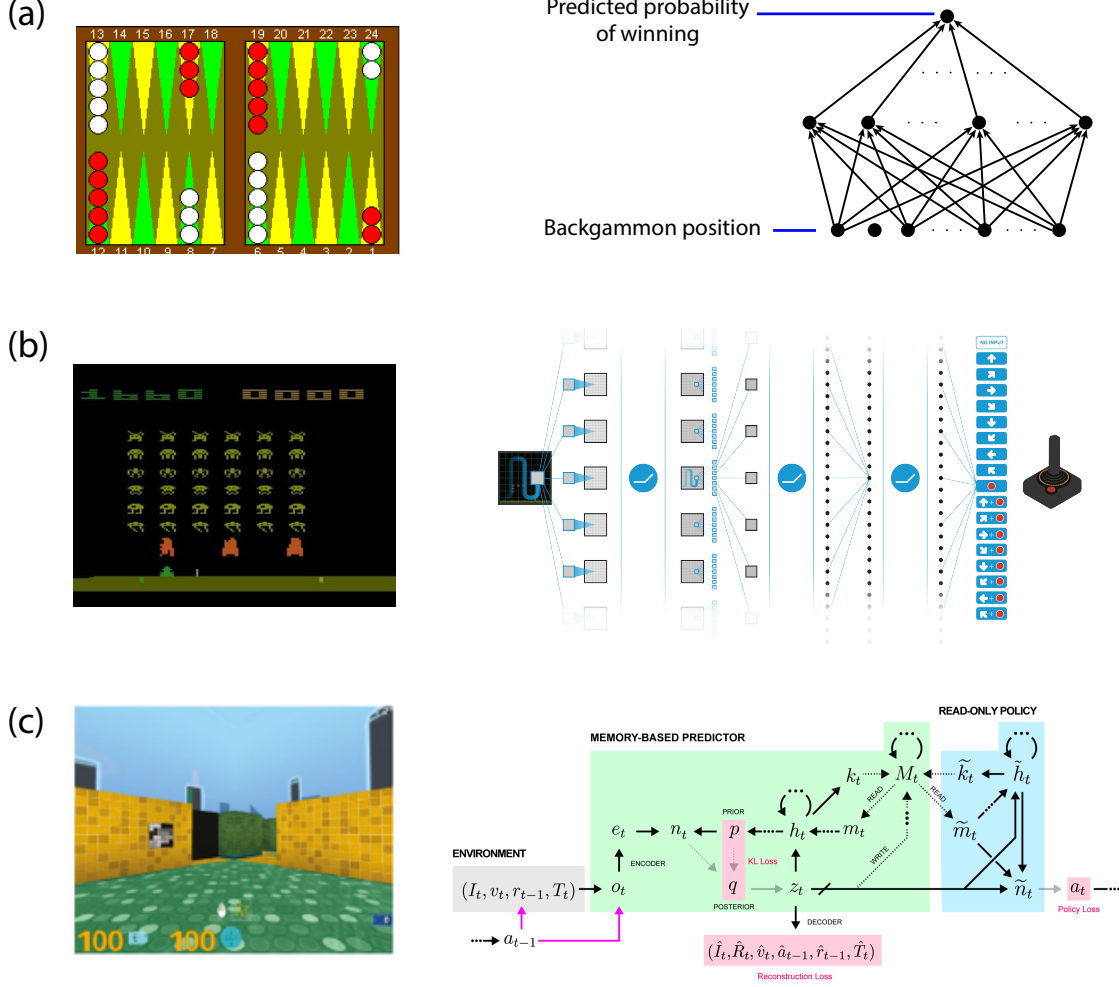


Figure I: **Representative examples of deep RL.**

Reinforcement learning centers on the problem of learning a behavioral *policy* —a mapping from states or situations to actions—which maximizes cumulative long-term reward [12]. In simple settings, the policy can be represented as a look-up table, listing the appropriate action for any state. In richer environments, however, this kind of simple listing is infeasible, and the policy must instead be encoded implicitly as a parameterized function. Pioneering work in the 1990s showed how this function could be approximated using a multi-layer (or deep) neural network [78, 79], allowing gradient-descent learning to discover rich, non-linear mappings from perceptual inputs to actions (see panel a and below). However, technical challenges prevented the integration of deep neural networks with RL until 2015, when breakthrough work demonstrated how deep RL could be made to work in complex domains such as Atari video games [13] (see panel b and below). Since then, rapid progress has been made toward improving and scaling deep RL [80], allowing its application to complex task domains such as Go [16] and capture the flag [81]. In many cases, the later advances have involved integrating deep RL with architectural and algorithmic complements, such as tree search [16] or slot-based, episodic-like memory [52] (see panel c and below). Other developments have focused on the goal of learning speed, allowing deep RL to make progress based on just a few observations, as reviewed in the main text.

The figure illustrates the evolution of deep RL methods, starting in panel a with Tesauro’s ground-breaking backgammon-playing system ‘TD-gammon’ [78]. This centered on a neural network which took as input a representation of the board and learned to output an estimate of the ‘state value,’ defined as expected cumulative future rewards, here equal simply to the estimated probability of eventually winning a game from the current position. Panel b shows the Atari-playing DQN network reported by Mnih and

colleagues [13]. Here, a convolutional neural network (see reference [3]) takes screen pixels as input and learns to output joystick actions. Panel c shows a schematic representation of a state-of-the art deep RL system reported by Wayne and colleagues [51]. A full description of the detailed ‘wiring’ of this RL agent is beyond the scope of the present paper (but can be found in reference [51]). However, as the figure indicates, the architecture comprises multiple modules, including a neural network that leverages an episodic-like memory to predict upcoming events, which ‘speaks’ to a reinforcement-learning module that selects actions based on the predictor module’s current state. The system learns, among other tasks, to perform goal-directed navigation in maze-like environments as shown in the figure.

Box 2: Episodic deep RL

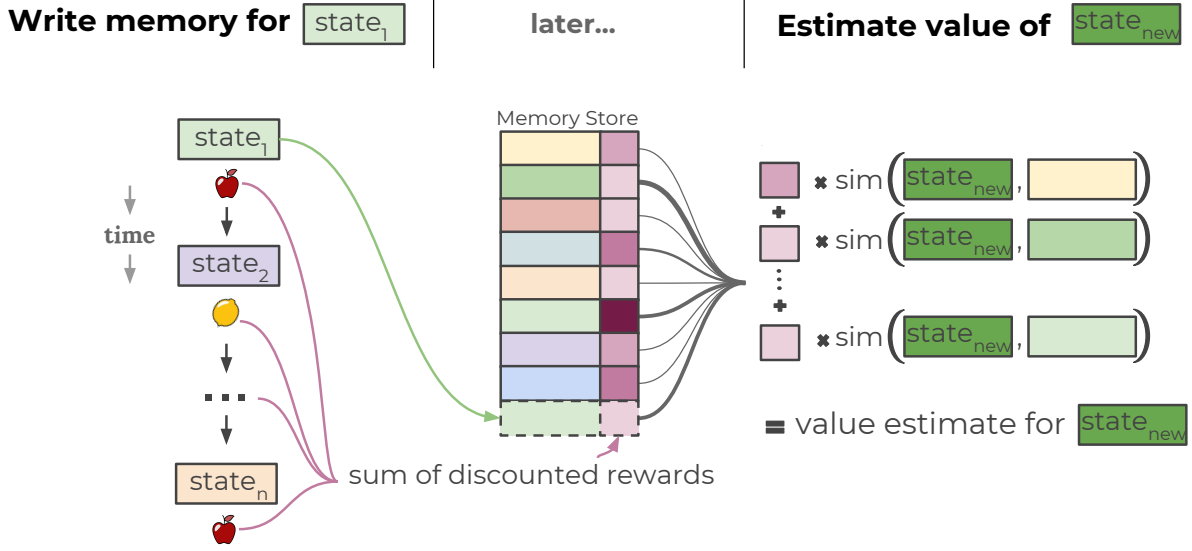


Figure I: Illustration of an example episodic RL algorithm.

Episodic reinforcement learning (RL) algorithms estimate the value of actions and states using episodic memories [30, 43, 44]. Consider for example the episodic valuation algorithm depicted in Figure I, wherein the agent stores each encountered state along with the discounted sum of rewards obtained during the next n time steps. These two stored items comprise an *episodic memory* of the encountered state and the reward that followed. To estimate the value of a new state, the agent computes a sum of the stored discounted rewards, weighted by the similarity (sim.) between stored states and the new state. This algorithm can be extended to estimate action values by recording the actions taken along with the states and reward sums in the memory store, then querying the store to find only memories in which the to-be-evaluated action was taken. In fact, [82] used such an episodic RL algorithm to achieve strong performance on Atari games.

The success of episodic RL depends on the *state representations* used to compute state similarity. In a follow-up to [82], [29] showed that performance could be improved by gradually shaping these state representations using gradient descent learning. These results demonstrated strong performance and state of the art data efficiency on the 57 games in the Atari Learning Environment [29], showcasing the benefits of combining slow (representation) learning and fast (value) learning.

Box 3: Meta-reinforcement learning

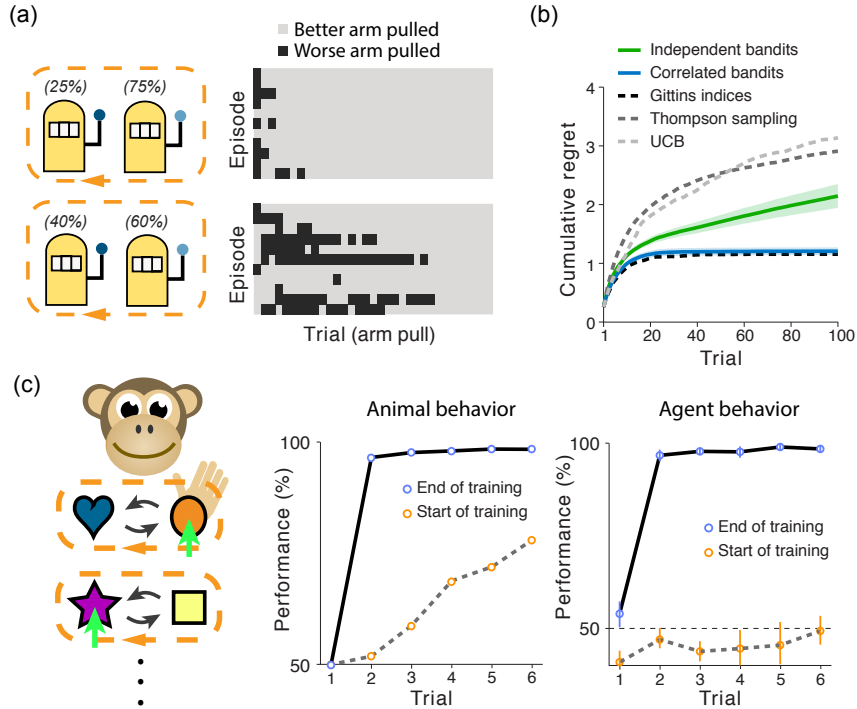


Figure I: **Bandits and the Harlow task.** (a) Example behavior of meta-RL on two-armed bandit problems at two difficulty levels. (b) Performance on correlated (blue) and independent (green) bandit problems, compared to performance of standard machine learning algorithms. (c) The Harlow task, illustration and comparison of animal to agent behavior over the course of training.

Meta-learning occurs when one learning system progressively adjusts the operation of a second learning system, such that the latter operates with increasing speed and efficiency [83, 84, 85, 86, 87]. This scenario is often described in terms of two “loops” of learning, an “outer loop” that uses its experiences over many task contexts to gradually adjust parameters that govern the operation of an “inner loop,” so that the inner loop can adjust rapidly to new tasks (see Figure 1). *Meta-reinforcement learning* refers to the case where both the inner and outer loop implement RL algorithms, learning from reward outcomes and optimizing toward behaviors that yield maximal reward.

As discussed in the main text, one implementation of this idea employs a form of recurrent neural network commonly used in machine learning: long-short term memory units (LSTM) [83, 88]. In this setup, the ‘outer loop’ of learning involves tuning of the network’s connection weights by a standard deep RL algorithm. Over the course of many tasks, this slowly gives rise to an ‘inner loop’ learning algorithm, which is implemented in the activity dynamics of the recurrent network [37, 38, 50, 83, 88]. A simple illustration involves training a recurrent neural network on a ‘two-armed bandit’ task. On every trial the agent must select between two actions (pull the left arm or right arm), and is rewarded based on latent parameters that govern the payoff probability for each action. During training, a series of bandit problems is presented, each with a randomly sampled pair of payoff parameters. The agent interacts with one bandit problem for a fixed number of steps and then moves on to another. The challenge for the agent, in each new bandit problem, is to balance exploration of the two alternatives with exploitation of information gained so far, seeking to maximize cumulative reward. Meta-RL learns to solve this problem, learning to explore more on a difficult bandit problem with arm reward probabilities that are more closely matched ([40%, 60%] vs [25%, 75%]; compare lower panel to upper panel of Figure Ia). After training on a number of problems, the network can – even with its connection weights fixed – explore a new bandit problem and converge on the richer arm using a strategy that compares favorably with hand-crafted machine-learning algorithms (see Figure Ib).

Critically, the learning algorithm that arises in the network’s activity dynamics is not only independent of the ‘outer loop’ RL algorithm that trains the network’s weights. It can also work *better* than that outer loop algorithm, because it learns inductive biases that align with the tasks on which the system is trained. An illustration is provided in Figure Ib (from Wang et al. [49]). The green line depicts performance (cumulative regret; lower is better) after training on a set of bandit tasks across

which the arm reward probabilities are sampled independently ("Independent bandits"), whereas the blue line represents training on a distribution in which the arm probabilities are linked ("Correlated bandits"; specifically, they always add to one). As the figure shows, training on correlated bandits leads the algorithm to hone in on the superior arm more rapidly. This reflects an adaptive inductive bias, implemented in the network's recurrent dynamics.

Figure 1c illustrates a richer application of meta-RL. As described in the main text, Harlow [34] introduced the idea of "learning to learn" using a task where a monkey chose between pairs of unfamiliar objects, eventually showing one-shot learning based on consistent structure in the task. Specifically, for each pair of objects, exactly one is consistently associated with reward, regardless of location. The monkeys eventually learned to randomly select an object on trial 1, and then use the reward outcome to perform perfectly for trials 2-6 (Figure 1c, middle panel). In recent computational modeling work, Wang and colleagues [49] showed that meta-RL, employing a recurrent neural network, gives rise to the same pattern of one-shot learning (Figure 1c, right panel).

Box 4: Meta-RL and rapid learning in the brain

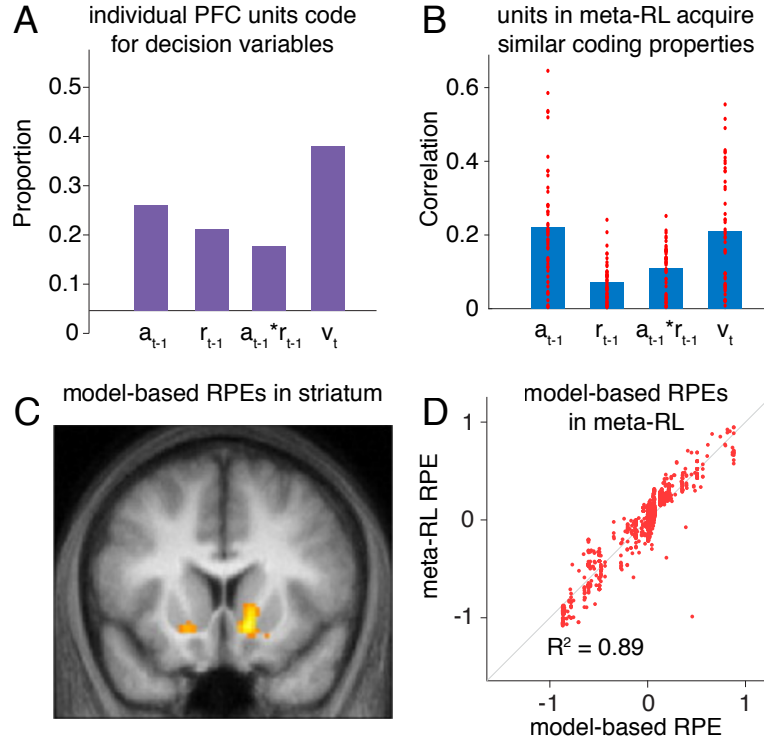


Figure I: (a,b) Proportion of units coding for task-related variables: last action, last reward, interaction of last action and last reward, and current value. (c,d) Like RPEs measured in the brain with fMRI, RPEs in meta-RL contain model-based information. Figures from Wang et al. [49].

Wang and colleagues [49] proposed that meta-RL, as described in Box 3, can model learning in biological brains. They argue that recurrent networks centered on prefrontal cortex (PFC) implement the inner loop of learning, and that this inner loop algorithm is slowly sculpted by an outer loop of dopamine-driven synaptic plasticity.

We first focus on the inner loop. PFC is central to rapid learning [46, 89], and neurons in PFC code for variables that underpin such learning. For example, Tsutsui et al. [90] recorded from primate dorsolateral PFC (dlPFC) during a foraging task in which environmental variables are continually changing. They found that individual neurons coded not only for the value of a current option, but also for the previous action taken, the previous reward received, and the interaction of previous action and previous reward (Figure IA). These are key variables for implementing an effective learning policy in this task.

Wang et al. [49] trained meta-RL (see Box 3) on the same task. They found that, over the course of training, the artificial recurrent network came to implement a fast learning policy whose behavior mimicked animals. Moreover, individual units in this trained network had coding properties remarkably similar to primate dlPFC neurons (Figure IB). Wang et al. [49] proposed that meta-learning over the course of an animal's lifetime shapes the recurrent networks in PFC to have tuning properties that are useful for naturalistic learning tasks.

We now turn to the outer loop. Classically, midbrain dopamine neurons are thought to carry the reward prediction error (RPE) signal of temporal difference learning [91, 92, 93]. In this standard theory, dopamine drives incremental adjustments to cortico-striatal synapses, and these adjustments predispose the animal to repeat reinforced actions. This model-free learning system is often viewed as complementary to a model-based system living in mostly distinct brain areas [94, 95, 96].

However, this standard theory is complicated by the observation that dopamine carries model-based information [97, 98, 99]. For example, in a task designed to isolate model-based and model-free signals, Daw et al. [97] found that, surprisingly, signals in ventral striatum – thought to be the strongest fMRI correlate of dopaminergic prediction errors – contained robust model-based information (Figure IC).

Wang et al. [49] trained meta-RL on a variant of this task [100]. In terms of behavior, they found that the trained recurrent network instantiated in its activity dynamics a robustly model-based learning

procedure – despite the fact that the training procedure of meta-RL was model-free. When they examined meta-RL’s RPEs, they found that the RPEs contained significant model-based information (Figure ID). Much work has been devoted to understanding how the brain learns and uses models [46, 101, 102, 103]. Meta-RL provides a single elegant explanation: that dopamine-dependent model-free learning automatically shapes recurrent networks to become a model-based learning algorithm.

In summary, meta-reinforcement learning may be an important principle of brain organization. In this model, slow, incremental learning processes shape recurrent brain circuits into algorithms that exploit consistent structure in the environment to learn efficiently [for related work, see 87, 104, 105].

Glossary

Neural network: A learnable set of weights and biases arranged in layers which process an input in order to produce an output. To learn more, see McClelland and Rumelhart’s seminal primer [106].

Hidden layer: A layer of a neural network in between the input and output layers.

Deep neural network: A neural network with one or (typically) more hidden layers.

Embedding: A learned representation residing in a layer of a neural network.

Non-parametric: In a non-parametric model, the number of parameters is not fixed, and can grow as more data is provided to the model.

Recurrent neural network: A neural network that runs at each time step in a sequence, passing its hidden layer activations from each step to the next.

References

- [1] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- [2] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.
- [3] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [4] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- [5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [6] Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *SSW*, page 125, 2016.
- [7] Adam H Marblestone, Greg Wayne, and Konrad P Kording. Toward an integration of deep learning and neuroscience. *Frontiers in computational neuroscience*, 10:94, 2016.
- [8] H Francis Song, Guangyu R Yang, and Xiao-Jing Wang. Reward-based training of recurrent neural networks for cognitive and value-based tasks. *Elife*, 6:e21492, 2017.
- [9] Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356, 2016.
- [10] David Sussillo, Mark M Churchland, Matthew T Kaufman, and Krishna V Shenoy. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience*, 18(7):1025, 2015.
- [11] Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11):e1003915, 2014.
- [12] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [14] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017.
- [15] Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. Human-level performance in first-person multiplayer games with population-based deep reinforcement learning. *arXiv preprint arXiv:1807.01281*, 2018.
- [16] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- [17] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashan Kumaran, Thore Graepel, et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*, 2017.

- [18] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [19] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [20] Richard S Sutton and Andrew G Barto. Toward a modern theory of adaptive networks: expectation and prediction. *Psychological review*, 88(2):135, 1981.
- [21] Wolfram Schultz, Peter Dayan, and P Read Montague. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997.
- [22] Pedro A Tsividis, Thomas Pouncy, Jacqueline L Xu, Joshua B Tenenbaum, and Samuel J Gershman. Human learning in atari. 2017.
- [23] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350, 2015.
- [24] Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018.
- [25] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [26] Dhharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7):512–534, 2016.
- [27] Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.
- [28] Christopher M Bishop. Pattern recognition and machine learning (information science and statistics) springer-verlag new york. *Inc. Secaucus, NJ, USA*, 2006.
- [29] Alexander Pritzel, Benigno Uria, Sriram Srinivasan, Adria Puigdomenech, Oriol Vinyals, Demis Hassabis, Daan Wierstra, and Charles Blundell. Neural episodic control. *arXiv preprint arXiv:1703.01988*, 2017.
- [30] Samuel J Gershman and Nathaniel D Daw. Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual review of psychology*, 68, 2017.
- [31] Gordon D Logan. Toward an instance theory of automatization. *Psychological review*, 95(4):492, 1988.
- [32] Edward E Smith and Douglas L Medin. *Categories and concepts*, volume 9. Harvard University Press Cambridge, MA, 1981.
- [33] Tom Schaul and Jürgen Schmidhuber. Metalearning. *Scholarpedia*, 5(6):4650, 2010. doi: 10.4249/scholarpedia.4650. revision #91489.
- [34] Harry F Harlow. The formation of learning sets. *Psychological review*, 56(1):51, 1949.
- [35] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pages 3981–3989, 2016.
- [36] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- [37] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dhharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

- [38] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. R^2 : Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [39] Samuel Ritter, Jane X Wang, Zeb Kurth-Nelson, Siddhant M Jayakumar, Charles Blundell, Razvan Pascanu, and Matthew Botvinick. Been there, done that: Meta-learning with episodic recall. *International Conference on Machine Learning (ICML)*, 2018.
- [40] Samuel Ritter, Jane X Wang, Zeb Kurth-Nelson, and Matthew M Botvinick. Episodic control as meta-reinforcement learning. *bioRxiv*, page 360537, 2018.
- [41] Samuel J Gershman and Nathaniel D Daw. Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annual review of psychology*, 68:101–128, 2017.
- [42] Máté Lengyel and Peter Dayan. Hippocampal contributions to control: the third way. In *Advances in neural information processing systems*, pages 889–896, 2008.
- [43] Aaron M Bornstein, Mel W Khaw, Daphna Shohamy, and Nathaniel D Daw. Reminders of past choices bias decisions for reward in humans. *Nature Communications*, 8:15958, 2017.
- [44] Aaron M Bornstein and Kenneth A Norman. Reinstated episodic context guides sampling-based decisions for reward. *Nature neuroscience*, 20(7):997, 2017.
- [45] Dorothy Tse, Rosamund F Langston, Masaki Takeyama, Ingrid Bethus, Patrick A Spooner, Emma R Wood, Menno P Witter, and Richard GM Morris. Schemas and memory consolidation. *Science*, 316(5821):76–82, 2007.
- [46] Timothy EJ Behrens, Timothy H Muller, James CR Whittington, Shirley Mark, Alon B Baram, Kimberly L Stachenfeld, and Zeb Kurth-Nelson. What is a cognitive map? organizing knowledge for flexible behavior. *Neuron*, 100(2):490–509, 2018.
- [47] James L McClelland, Bruce L McNaughton, and Randall C O’reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995.
- [48] Ian C Ballard, Anthony D Wagner, and Samuel M McClure. Hippocampal pattern separation supports reinforcement learning. *bioRxiv*, page 293332, 2018.
- [49] Jane X Wang, Zeb Kurth-Nelson, Dharsan Kumaran, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Demis Hassabis, and Matthew Botvinick. Prefrontal cortex as a meta-reinforcement learning system. *Nature neuroscience*, 21(6):860, 2018.
- [50] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850, 2016.
- [51] Greg Wayne, Chia-Chun Hung, David Amos, Mehdi Mirza, Arun Ahuja, Agnieszka Grabska-Barwinska, Jack Rae, Piotr Mirowski, Joel Z Leibo, Adam Santoro, et al. Unsupervised predictive memory in a goal-directed agent. *arXiv preprint arXiv:1803.10760*, 2018.
- [52] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471, 2016.
- [53] Oliver Vikbladh, Daphna Shohamy, and Nathaniel Daw. Episodic contributions to model-based reinforcement learning. In *Annual conference on cognitive computational neuroscience, CCN*, 2017.
- [54] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- [55] Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 3675–3683, 2016.

- [56] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pages 3630–3638, 2016.
- [57] Elizabeth S Spelke and Katherine D Kinzler. Core knowledge. *Developmental science*, 10(1):89–96, 2007.
- [58] Matthew M Botvinick. Multilevel structure in behaviour and in the brain: a model of fuster’s hierarchy. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1485):1615–1626, 2007.
- [59] Nicolas P Rougier, David C Noelle, Todd S Braver, Jonathan D Cohen, and Randall C O’Reilly. Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences*, 102(20):7338–7343, 2005.
- [60] David C Plaut. Graded modality-specific specialisation in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*, 19(7):603–639, 2002.
- [61] Olaf Sporns, Dante R Chialvo, Marcus Kaiser, and Claus C Hilgetag. Organization, development and function of complex brain networks. *Trends in cognitive sciences*, 8(9):418–425, 2004.
- [62] Ed Bullmore and Olaf Sporns. The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336, 2012.
- [63] Dharmendra S Modha and Raghavendra Singh. Network architecture of the long-distance pathways in the macaque brain. *Proceedings of the National Academy of Sciences*, 107(30):13485–13490, 2010.
- [64] Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- [65] Chrisantha T Fernando, Jakub Sygnowski, Simon Osindero, Jane Wang, Tom Schaul, Denis Teplyashin, Pablo Sprechmann, Alexander Pritzel, and Andrei A Rusu. Meta learning by the baldwin effect. *arXiv preprint arXiv:1806.07917*, 2018.
- [66] Geoffrey E Hinton and Steven J Nowlan. How learning can guide evolution. *Complex systems*, 1(3):495–502, 1987.
- [67] Jane X Wang, Edward Hughes, Chrisantha Fernando, Wojciech M Czarnecki, Edgar A Duenez-Guzman, and Joel Z Leibo. Evolving intrinsic motivations for altruistic behavior. *arXiv preprint arXiv:1811.05931*, 2018.
- [68] Y Bengio, S Bengio, and J Cloutier. Learning a synaptic learning rule. In *Neural Networks, 1991., IJCNN-91-Seattle International Joint Conference on*, volume 2, pages 969–vol. IEEE, 1991.
- [69] Max Jaderberg, Wojciech Marian Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, and Koray Kavukcuoglu. Decoupled neural interfaces using synthetic gradients. *arXiv preprint arXiv:1608.05343*, 2016.
- [70] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- [71] Irina Higgins, Arka Pal, Andrei A Rusu, Loic Matthey, Christopher P Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. Darla: Improving zero-shot transfer in reinforcement learning. *arXiv preprint arXiv:1707.08475*, 2017.
- [72] Matthew Botvinick, Ari Weinstein, Alec Solway, and Andrew Barto. Reinforcement learning, efficient coding, and the statistics of natural tasks. *Current Opinion in Behavioral Sciences*, 5: 71–77, 2015.
- [73] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- [74] Thomas L. Griffiths, Charles Kemp, and Joshua B. Tenenbaum. *Bayesian Models of Cognition*, page 59–100. Cambridge Handbooks in Psychology. Cambridge University Press, 2008.

- [75] Thomas L Griffiths, Nick Chater, Charles Kemp, Amy Perfors, and Joshua B Tenenbaum. Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in cognitive sciences*, 14(8):357–364, 2010.
- [76] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022):1279–1285, 2011.
- [77] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. *arXiv preprint arXiv:1801.08930*, 2018.
- [78] Gerald Tesauro. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3):58–68, 1995.
- [79] Long-Ji Lin. Reinforcement learning for robots using neural networks. Technical report, Carnegie-Mellon Univ Pittsburgh PA School of Computer Science, 1993.
- [80] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*, 2017.
- [81] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.
- [82] Charles Blundell, Benigno Uria, Alexander Pritzel, Yazhe Li, Avraham Ruderman, Joel Z Leibo, Jack Rae, Daan Wierstra, and Demis Hassabis. Model-free episodic control. *arXiv preprint arXiv:1606.04460*, 2016.
- [83] Sepp Hochreiter, A Steven Younger, and Peter R Conwell. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, pages 87–94. Springer, 2001.
- [84] Sebastian Thrun and Lorien Pratt. Learning to learn: Introduction and overview. In *Learning to learn*, pages 3–17. Springer, 1998.
- [85] Jonathan Baxter. Theoretical models of learning to learn. In *Learning to learn*, pages 71–94. Springer, 1998.
- [86] Jurgen Schmidhuber, Jieyu Zhao, and Marco Wiering. Simple principles of metalearning. Technical report, SEE, 1996.
- [87] Nicolas Schweighofer and Kenji Doya. Meta-learning in reinforcement learning. *Neural Networks*, 16(1):5–9, 2003.
- [88] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [89] Matthew FS Rushworth, MaryAnn P Noonan, Erie D Boorman, Mark E Walton, and Timothy E Behrens. Frontal cortex and reward-guided learning and decision-making. *Neuron*, 70(6):1054–1069, 2011.
- [90] Ken-Ichiro Tsutsui, Fabian Grabenhorst, Shunsuke Kobayashi, and Wolfram Schultz. A dynamic code for economic object valuation in prefrontal cortex neurons. *Nature communications*, 7:12554, 2016.
- [91] P Read Montague, Peter Dayan, and Terrence J Sejnowski. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of neuroscience*, 16(5):1936–1947, 1996.
- [92] Paul W Glimcher. Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Supplement 3):15647–15654, 2011.
- [93] Mitsuko Watabe-Uchida, Neir Eshel, and Naoshige Uchida. Neural circuitry of reward prediction error. *Annual review of neuroscience*, 40:373–394, 2017.

- [94] Ray J Dolan and Peter Dayan. Goals and habits in the brain. *Neuron*, 80(2):312–325, 2013.
- [95] Nathaniel D Daw, Yael Niv, and Peter Dayan. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12):1704, 2005.
- [96] Bernard W Balleine and Anthony Dickinson. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4-5):407–419, 1998.
- [97] Nathaniel D Daw, Samuel J Gershman, Ben Seymour, Peter Dayan, and Raymond J Dolan. Model-based influences on humans’ choices and striatal prediction errors. *Neuron*, 69(6):1204–1215, 2011.
- [98] Brian F Sadacca, Joshua L Jones, and Geoffrey Schoenbaum. Midbrain dopamine neurons compute inferred and cached value prediction errors in a common framework. *Elife*, 5:e13665, 2016.
- [99] Melissa J Sharpe, Chun Yun Chang, Melissa A Liu, Hannah M Batchelor, Lauren E Mueller, Joshua L Jones, Yael Niv, and Geoffrey Schoenbaum. Dopamine transients are sufficient and necessary for acquisition of model-based associations. *Nature neuroscience*, 20(5):735, 2017.
- [100] Kevin J Miller, Matthew M Botvinick, and Carlos D Brody. Dorsal hippocampus contributes to model-based planning. *Nature neuroscience*, 20(9):1269, 2017.
- [101] Jan Gläscher, Nathaniel Daw, Peter Dayan, and John P O’Doherty. States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, 66(4):585–595, 2010.
- [102] Nils Kolling, Marco K Wittmann, Tim EJ Behrens, Erie D Boorman, Rogier B Mars, and Matthew FS Rushworth. Value, search, persistence and model updating in anterior cingulate cortex. *Nature neuroscience*, 19(10):1280, 2016.
- [103] Randall C O’Reilly and Michael J Frank. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural computation*, 18(2):283–328, 2006.
- [104] Shin Ishii, Wako Yoshida, and Junichiro Yoshimoto. Control of exploitation–exploration meta-parameter in reinforcement learning. *Neural networks*, 15(4-6):665–687, 2002.
- [105] Mehdi Khamassi, Pierre Enel, Peter Ford Dominey, and Emmanuel Procyk. Medial prefrontal cortex and the adaptive regulation of reinforcement learning parameters. In *Progress in brain research*, volume 202, pages 441–464. Elsevier, 2013.
- [106] James L McClelland and David E Rumelhart. *Explorations in parallel distributed processing: A handbook of models, programs, and exercises*. MIT press, 1989.