

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Structure Learning in Motor Control:A Deep Reinforcement Learning Model

Permalink

<https://escholarship.org/uc/item/35v0x009>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors

Weinstein, Ari
Botvinick, Matthew M.

Publication Date

2017

Peer reviewed

Structure Learning in Motor Control: A Deep Reinforcement Learning Model

Ari Weinstein (ariweinstein@google.com)

Deepmind, London, UK

Matthew M. Botvinick (botvinick@google.com)

Deepmind, London, UK

Gatsby Computational Neuroscience Unit, University College London

Abstract

Motor adaptation displays a structure-learning effect: adaptation to a new perturbation occurs more quickly when the subject has prior exposure to perturbations with related structure. Although this ‘learning-to-learn’ effect is well documented, its underlying computational mechanisms are poorly understood. We present a new model of motor structure learning, approaching it from the point of view of deep reinforcement learning. Previous work outside of motor control has shown how recurrent neural networks can account for learning-to-learn effects. We leverage this insight to address motor learning, by importing it into the setting of model-based reinforcement learning. We apply the resulting processing architecture to empirical findings from a landmark study of structure learning in target-directed reaching (Braun et al., 2009), and discuss its implications for a wider range of learning-to-learn phenomena.

Keywords: motor adaptation; reinforcement learning; learning to learn; structure learning; system identification

Introduction

Learning can be defined as a process that improves performance as exposure to a task increases. However, research on human and animal learning makes clear that this simple definition is not quite enough to explain the observed relationship between experience and performance. The full picture must also include ‘learning-to-learn,’ a process whereby growing experience causes learning itself to become more efficient (Harlow, 1949). More specifically, learning-to-learn (also referred to as meta-learning and structure learning) occurs in settings where the learner encounters a series of tasks that share some underlying structure, and gains from these an ability to quickly adapt to a new task that displays the same general form (Thrun & Pratt, 1998).

A vivid example of learning-to-learn, which provides a concrete focus for the present research, comes from research on motor adaptation. Many studies have documented the ability of human subjects to adapt to perturbations of motor dynamics or kinematics, as for example in prism adaptation (Harris, 1963). However, a series of studies by Braun and colleagues (Braun et al., 2009, 2010; Braun & Wolpert, 2012) went beyond this to show that adaptation can occur faster when the subject has prior exposure to perturbations that share structure with the final test conditions. In one specific experiment, upon which we will continue to focus, Braun and colleagues (2009) studied reaching under visuo-motor rotation. They examined the speed with which target-directed reaching adapted to a 60-degree rotation, manipulating between subjects the content of a preceding set of training

trials. In one condition, which we will refer to as *Rot*, subjects dealt with a series of rotations (though never the one presented at test). In a comparison condition *Rot+*, subjects dealt with a more diverse set of transformations, each made up of a rotation along with shear and scale components. Results showed that subjects in the *Rot* group adapted faster to the probe rotation problem (Figure 1). Braun et al. (2009) interpreted this as learning-to-learn effect, which they referred to as “motor structure learning”: Subjects in the *Rot* group evidently learned that the transformations being presented were restricted to a particular structurally coherent set (rotations), and this allowed them to infer and adapt rapidly to the probe transformation. This structure learning was less feasible in the *Rot+* condition because the structure underlying the training set was more complex, thus offering weaker constraint on inference when facing a new transformation.

In the present study, we consider the computational mechanisms underlying motor structure learning, treating it as a case study in learning-to-learn. Despite widespread agreement that learning-to-learn effects are both real and important, the precise computational processes underlying such effects are poorly understood. The most widely proposed idea comes from a Bayesian perspective, and proposes that learning-to-learn involves refining the structure and hyperparameters of a generative model of the relevant task domain (Lake et al., 2015). Braun and colleagues initially proposed, and later investigated (Genewein et al., 2015) a model of this sort to account for their structure learning results.

A different computational proposal, which has been less widely considered in cognitive science, comes from neural network or deep learning research. In classic work, Hochreiter and colleagues (2001) showed how a recurrent neural network (RNN) can learn to learn, by integrating information about past outcomes into predictions concerning new observations. Recent applications of this idea (Wang et al., 2016; Duan et al., 2016) have treated the RNN in Hochreiter’s scheme as a mechanism for directly selecting actions. In the present work, we leverage Hochreiter’s (2001) insight in a different way, using an RNN as an adaptive model of the task domain, which is leveraged by a separate action-selection mechanism. In this sense, the aim of our work is to bridge Bayesian and deep learning perspectives on learning-to-learn. On a more immediate level, we show how the resulting approach can be used to account for the findings of Braun




Figure 1: Results redrawn from Braun et al.(2009), showing mean cumulative error on a series of five reaches under a 60-degree visuomotor rotation.

and colleagues (2009) in motor adaptation.

The motor control literature suggests that actions such as reaching are based, at least in part, on an internal predictive or forward model of reaching dynamics (Wolpert et al., 1995; Miall & Wolpert, 1996), and analyses of motor adaptation have portrayed adaptation as reflecting a progressive adjustment of this internal model to fit with current observations (Berniker & Kording, 2008; Haith & Krakauer, 2013). Following this idea, and a great deal of previous work in computational motor control, we construe action selection as a form of model-based reinforcement learning (RL).

Formulating the problem in this way begins by casting it as a finite-time Markov decision process (MDP) M , which is made of a set of states S , a set of possible actions A , a transition function \mathcal{T} , and reward function R (in many settings a discount factor γ is included, but since we formulate the task as a finite-time problem this is unnecessary). The goal is to select actions that maximize the cumulative reward up to some time T : $\sum_{t=0}^T r_{t+1}$, where t indexes discrete time steps up to some maximum T , and r_t is the reward received on each step. Focusing on target-directed reaching, the task studied by Braun and colleagues (2009), the problem is defined as follows: M is the entire reaching task; S are possible arm configurations; A are possible motor inputs; \mathcal{T} defines the dynamics of the arm based on motor inputs; R is the negative distance from the cursor to the target. In order to be consistent with the literature on structure learning in motor control, we will use the terms reward maximization and penalty (or error) minimization interchangeably.

In model-based RL, a model \hat{M} of the environment M is built, and then used by a planner P in order to construct an action-selection policy. The general form of a model-based learning architecture is diagrammed in Figure 2, left. Here a planner P is informed of a current state s by the true MDP M . Based on the particular policy of P , the planner queries the model \hat{M} with a series of state-action pairs (s_t, a_t) , and in turn receives an estimated next state s_{t+1} and reward r_{t+1} . After the planner completes querying \hat{M} , either because it has




Figure 2: Model-based reinforcement learning with a fixed model (left) and an adaptive model (right).

taken as much data as it needs or due to some outside pressure such as a time limit, it returns an action a which is executed in M , which results in a new state and reward and the process repeats.

In the architecture shown in Figure 2 (left), one way of implementing the forward model M is as a feed-forward neural network. This approach has been explored in a number of previous studies (Jordan & Rumelhart, 1992; Hamrick et al., 2016). However, a feed-forward neural network will not suffice to address the learning-to-learn phenomena we are concerned with here. Indeed, the overall architecture must be fundamentally changed in order to address the learning-to-learn problem.

As introduced earlier, learning-to-learn arises in a setting where the learner encounters a series of interrelated problems or tasks, and must adapt to each one in turn. Using our terminology, each task $M_n, n = 1 \dots N$ becomes a sample from a task distribution \mathbb{M} . As such, the properties of each M_n must be inferred based on observed action-outcome pairs (a process referred to in the engineering literature as *system identification*). On a formal level, this demand changes the MDP we have been considering into a partially observable Markov decision process (POMDP). By definition, a POMDP is an MDP which additionally has an observation space O and observation function Ω which takes its internal state and outputs an observation o to the agent. Instead of the true state, an agent only has access to observations, which unlike state, is generally insufficient to act optimally when considered in isolation. In order for \hat{M} to adjust to each M_n , it must have some form of memory α to keep a relevant summary of interactions with the environment, allowing for integration over previous timesteps in order to accurately estimate problem dynamics.

These requirements yield the interaction and planning structure diagrammed in Figure 2 (right). Instead of states s , information presented is in terms of observations o . The model \hat{M} must now directly consume o and r from M_n at every time step, which causes it to update its memory α . P now only passes a sequence of actions along trajectories to the model. This is because it does not have access to the

true state, and because observations alone are not sufficient for planning or modelling¹. Additionally, during the planning trajectories, the P must signal to \hat{M} when its evaluation of a simulated trajectory is complete so that α can be reset (omitted from figure for clarity).

Note that, unlike in the simpler MDP, in the POMDP setting the internal model \hat{M} cannot be accurately implemented as a feed-forward neural network, because such networks do not have memory or persistent internal state. The key move in the present work is to substitute for the feed-forward network a RNN, whose recurrent connectivity endows it with the memory needed to support system identification and, as we will show, learning-to-learn.

Simulation Study

We predicted that the proposed architecture, if trained in an appropriate multi-task setting, would display learning-to-learn, leveraging experience with past tasks to adapt rapidly to a new task sharing in the same structure. In order to test this idea, we applied the architecture to the task paradigm employed in their study of motor structure learning by Braun and colleagues (2009).

Model implementation and task design

We implement the architecture shown in Figure 2 (right), instantiating the forward model in the form of a recurrent neural network (which is naturally deep as it is unrolled over time). More specifically, this involves one LSTM layer (Hochreiter & Schmidhuber, 1997) followed by two more fully connected layers containing rectified linear units (Nair & Hinton, 2010), where each layer contains 100 units. The planner is an open-loop planner based on cross-entropy optimization, as described in Weinstein & Littman (2013), with the addition of “warm starting.” In warm starting, planning is done from scratch on the first step of a trajectory, but all subsequent steps in the actual domain initiate planning with the result from the previous step. At each time step only the first action in the current plan is executed in the true domain before partial replanning in this manner occurs. For simplicity, we assume (without loss of generality) that \hat{M} has access to the reach-target coordinates and can compute the reward function.

In order to model target-directed reaching, we implemented a simple arm model. While not intended to offer a detailed model of biomechanics, this was intended to capture the most important aspects in terms of possible arm geometry, velocity, and acceleration (Nagasaki, 1989). As simulated, the underlying state space of the problem has four dimensions: horizontal shoulder angle, elbow angle, and corresponding angular velocities. Observations emitted are the Euclidian position of the cursor controlled by the arm’s tip as seen in the experiment (meaning that \hat{M} must also learn to

estimate velocities), and the goal. The two dimensional action space sets the angular accelerations of the joints, and the reward is the negative Euclidean distance of the cursor from the center of the goal region.

In the reaching task, the cursor is always initialized at the origin and is controlled by the transformation of the underlying position of the simulated hand. Before each trial a goal location is selected which is set to be 8 cm from the origin at a uniformly distributed angle.

Training and testing procedure

The simulation study, like the experiment by Braun et al. (2009) was divided into training and testing phases. During training, the RNN model was trained to predict each sequential outcome observation exactly along a trajectory consisting of observations and randomly selected actions, that is, following a random walk. Again as in the empirical study, two versions of the model were trained in different environments. One model, which we label (in a minor abuse of terminology) *Rot*, was trained on a series of visuomotor rotations, simulated by appropriately transforming the observed cursor coordinates. The second model instance, *Rot+*, was trained on a combination of rotations, shears, and scales (following the design described in Braun et al., 2009, Supplemental Data). Following the design imposed by Braun and colleagues, when the rotation to be presented to *Rot+* fell between $\pm 50^\circ$ and 70° , a rotation of $\pm 60^\circ$ was substituted and no linear transform was applied. As a result, both *Rot* and *Rot+* had roughly equal exposure to the transformation used during test trials. In both conditions, the model was trained by backpropagation through time on 2,000 trajectories of random-walk data, with each trajectory containing three seconds of simulation time, and training starting from the initial observation of each trajectory.

In the testing phase the RNN weight parameters were frozen and reaches were elicited only under only pure rotations of $\pm 60^\circ$, as in the testing phase of the experiment by Braun and colleagues. Goal locations were placed at a randomly selected angle 8 cm from the start location of the cursor. The radius of the goal region is 1.6 cm. In order to simulate a series of reaches, the angle of the imposed visuomotor rotation was held constant while the position of the goal varied between reaches. Test reach trajectories ran for a maximum of two seconds, terminating early if the cursor was brought within the goal region for 500 ms.

Results

Training of models for both the *Rot* and *Rot+* conditions were successful, but the model trained on *Rot* was able to achieve an average error of about 0.002 cm per time step for trajectories in the training set, while *Rot+* an error of about 0.03 cm by the same metric. In the *Rot* condition, the RNN model learned to act as an adaptive forward model, adjusting its predictions to fit with accumulating action-outcome observations. Figure 3 shows the average observation-prediction errors of both *Rot* and *Rot+* models during an initial random-

¹Belief states (Kaelbling et al., 1998) or predictive state representations (Littman et al., 2001) are sufficient for planning, but can be computed internally in each module and do not need to be communicated.



Figure 3: Model errors by step in initial trial.

walk trajectory which was not part of the training set. The initial data-point is the error of the model prior to any experience in the test MDP. In interpreting the values on the y-axis of the plot, it should be taken into account that in our simulation two seconds of time takes 28 discrete time steps, and error compounds over these steps. In contrast to the *Rot* model, the *Rot+* model adapts much less successfully, despite having been trained on an identical amount of data.

Figure 4 shows the mean cumulative penalty when the model is coupled with a planner, for each reach at test for both models. This is intended for comparison with the empirical data from humans in preceeding work shown in Figure 1. As predicted, *Rot* is better able to conduct structure learning, by adapting more rapidly and completely to the test rotation (the manipulation both models were exposed to during training) than *Rot+*. This qualitatively replicates the experimental findings from Braun et al. (2009).

Figure 5a shows average trajectories for five successive reaches (normalized by rotation and goal angles), for both *Rot* and *Rot+* models. Both models adapted across reaches (starting with smaller initial angular errors after the initial reach), but the effects were stronger in the *Rot* model. Quite striking is the standard deviation of the final position of the first trial of *Rot*, and *Rot+* in cyan and magenta, respectively. Although on average *Rot+* tracked toward target, there is a tremendous amount of variability in its trajectories, and was not able to consistently reach the goal region, whereas *Rot* usually terminated within the target.

We also consider other indirect metrics of performance which are presented in the human studies such as initial angular error, velocity, and minimum distance to goal region, which are presented in Figures 5b through 5d, respectively. In general the results with these metrics are similar to the previous plots, with *Rot* improving quickly and performing better than *Rot+*. We also note the higher variance of *Rot+*, which manifests itself in wider confidence intervals across all Figures, especially Figure 5d. These results are qualitatively aligned with those reported in the experimental study.

In fact, of these metrics, the only one which shows improvement by *Rot+* is the initial angular error. Even with




Figure 4: Average cumulative penalty by trial.

this improvement, the agent frequently falls short of reaching goal region (which would allow for an early termination of distance penalties). This is most likely due to the fact that on average testing data in *Rot+* has a scaling amount of roughly 1.3 (this design is part of the original human study), and indeed *Rot+* almost uniformly tells the planner that actions will result in greater changes in location than actually occur.

Although *Rot+* was less effective at structure learning than *Rot*, it is not the case that it failed entirely. The average penalty of a trajectory for agent using a uniform random policy is approximately 220 units which is significantly poorer than what *Rot+* was able to achieve.

We note that our goal was not to fit the results from the human data quantitatively, but rather to demonstrate the same phenomenon which is that structure learning becomes more difficult as the the amount of variability in the problem increases. And although *Rot+* was not able to perform well, the overall architecture does have the capacity to do effective structure learning; expanding the data corpus size by a factor of five produces models that have statistically equal, high quality performance on test tasks for both *Rot* and *Rot+*.

Discussion

Learning-to-learn is a fundamental aspect of human behavior, but its computational basis is not yet well understood. We have presented a new model of learning-to-learn in the setting of motor adaptation. This task, defined by Braun and colleagues (2009) involves learning to learn in the sense that the subject must gather data on a current situation in order to infer the hidden parameters of the dynamics, and indeed Braun and colleagues state that learning to learn can be recast as structure learning. On the other hand, a stronger definition of learning to learn could require learning to adapt to a situation it has not experienced in the past, perhaps in terms of new objects to interact with that follow some prelearned rules (Harlow, 1949). This has been considered in a different simulated setting in RL where the agent learns policies (Wang et al., 2016), as opposed to models of the environment as is done here.

Adopting the standard approach, we assume that motor

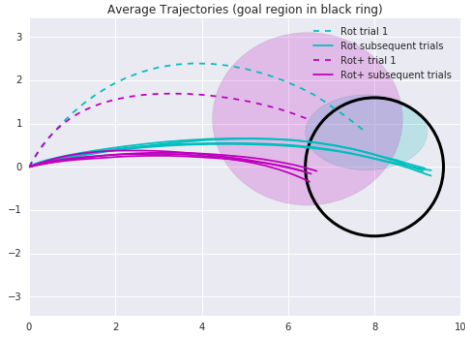


Figure 5a: Average trajectories by trial. Standard deviation of final position of first trial in shaded region. Goal region in black.

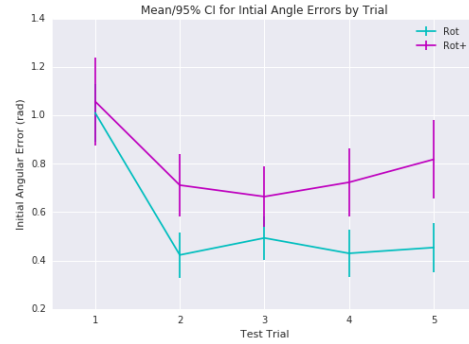


Figure 5b: Average angle error from goal after 200 ms of simulated time.

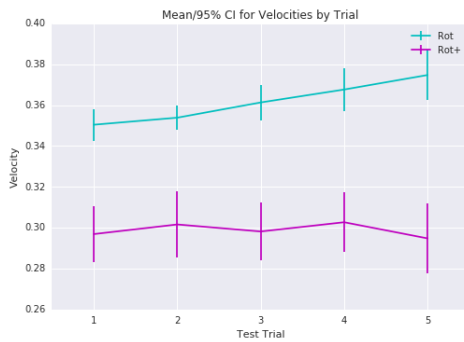


Figure 5c: Average velocities by trial

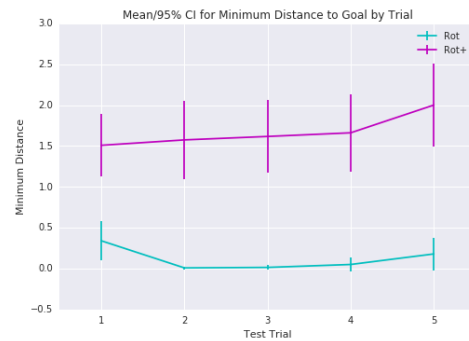


Figure 5d: Average minimum distance to goal by trial

adaptation involves updating an internal forward model of reaching dynamics. Our novel contribution is to instantiate this internal model as a recurrent neural network. Through simulations of a key experimental study, we have shown that the resulting system not only learns to adapt to changing perturbations, but also that its adaptation becomes more effective when there is prior exposure to structurally related conditions, as seen empirically in motor structure learning. Importantly, no special measures were required in order to secure this learning-to-learn effect. Through error-correcting learning, the parameters of the RNN are, perforce, fit to the structure of the pre-training data. That same structure is thus naturally – indeed inevitably – expressed in its later inferences at test.

We consider learning to learn as refining a (potentially implicit) hypothesis set based on experience. If the problem has a large underlying dimension, then the hypothesis set learned by the model must be of corresponding size. This is in turn fundamentally linked to the amount of data required to both train the model, as well as do inference, accurately. For these reasons, it is to be expected that when comparing the data requirements of doing both in *Rot* versus *Rot+*, *Rot* leads to lower data requirements. Just as is the case with Braun and colleagues (2009), we do not attempt to disentangle these issues, although a more detailed investigation warrants future

attention.

As noted earlier, our use of RNN dynamics to capture learning-to-learn effects builds directly on pioneering work by Hochreiter and colleagues (2001), in which an RNN model was applied to the problem of function induction (see also Wang et al., 2016; Santoro et al., 2016). In contrast to that work, we deployed our RNN as a forward model situated within a larger model-based RL system. In this sense, our implementation bridges between Hochreiter’s original proposal and models of motor adaptation that have embedded an adaptive Bayesian model of limb dynamics (e.g. Berniker & Kording, 2008; Genewein et al., 2015). The approach we have introduced also relates to other work in which RNNs have been used as forward models in support of motor adaptation, but where multiple fixed models are assumed (Haruno et al., 2001; Pitti et al., 2013), rather than a single adaptive model used here. These fixed models lack memory, meaning that reweighing fixed models aside, adaptation is only possible by retraining the system. Implicitly, our work implements a sort of Kalman filter which has also been considered previously in recurrent networks (Wolpert et al., 1995). Undertaking a careful comparison between these related approaches and the one we have introduced here offers an important objective for next-step research.

Our implementation of the reaching task was deliberately minimal, simplifying both the underlying biomechanics and the motor planning process, in order to foreground our central computational proposal. Naturally, a more detailed evaluation of the approach, incorporating a higher degree of empirical constraint, will be desirable in further evaluating the viability of our approach as a theory of motor adaptation. A related opportunity is to consider the potential parallel between the recurrent connectivity underlying the function of our adaptive model and the recurrent connectivity inherent in biological neural circuits underlying motor control and adaptation, including circuits running through the basal ganglia and cerebellum.

At the same time, however, we feel it may also be fruitful to apply the model-based framework we have introduced here in domains beyond motor control, in particular other domains that display the characteristics of a POMDP and where learning-to-learn effects have been observed. Such tasks are indeed ubiquitous, ranging from structured bandit tasks to video-game play (Wang et al., 2016; Lake et al., 2015). To the extent that the framework we have presented here can be adapted and (more challenging) effectively scaled to these other settings, it offers to provide a more general new perspective on the problem of learning-to-learn.

Acknowledgements

We would like to thank Daniel Braun and Konrad Kording for valuable feedback.

References

- Berniker, M., & Kording, K. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nature Neuroscience*, *11*, 1454–1461.
- Braun, D. A., Aersten, A., & Wolpert, D. M. (2009). Motor task variation induces structural learning. *Current Biology*, *19*, 352–357.
- Braun, D. A., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behavioural Brain Research*, *206*, 157–165.
- Braun, D. A., & Wolpert, D. M. (2012). Structural learning in sensorimotor control. In *Encyclopedia of the sciences of learning* (pp. 3208–3211).
- Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., & Abbeel, P. (2016). RL²: Fast reinforcement learning via slow reinforcement learning. *CoRR*, *abs/1611.02779*.
- Genewein, T., Hez, E., Razzaghpahan, Z., & Braun, D. A. (2015). Structure learning in bayesian sensorimotor integration. *PLoS Computational Biology*, *11*(8).
- Haith, A. M., & Krakauer, J. W. (2013). Model-based and model-free mechanisms of human motor learning. In *Progress in motor control: Neural, computational and dynamic approaches* (pp. 1–21).
- Hamrick, J. B., Pascanu, R., Vinyals, O., Ballard, A., Heess, N., & Battaglia, P. W. (2016). Imagination-based decision making with physical models in deep neural networks. In *NIPS 2016 workshop on intuitive physics*.
- Harlow, H. F. (1949). The formation of learning sets. *Psychological Review*, *56*(1), 51–65.
- Harris, C. S. (1963). Adaptation to displaced vision: Visual, motor, or proprioceptive change? *Science*, *140*(3568), 812–813.
- Haruno, M., Wolpert, D. M., & Kawato, M. (2001). Mosaic model for sensorimotor learning and control. *Neural Computation*, *13*, 2201–2220.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8).
- Hochreiter, S., Younger, A. S., & Conwell, P. R. (2001). Learning to learn using gradient descent. In *Artificial neural networks ICANN* (Vol. 2130, pp. 87–94).
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, *3*, 307–354.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, *1–2*.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338.
- Littman, M. L., Sutton, R. S., & Singh, S. (2001). Predictive representations of state. In *Neural information processing systems* (Vol. 14).
- Miall, R., & Wolpert, D. (1996). Forward models for physiological motor control. *Neural Networks*, *9*(8), 1265 - 1279.
- Nagasaki, H. (1989). Asymmetric velocity and acceleration profiles of human arm movements. *Experimental Brain Research*, *74*(2), 319–326.
- Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *International conference on machine learning* (pp. 807–814).
- Pitti, A., Braud, R., Mah, S., Quoy, M., & Gaussier, P. (2013). Neural model for learning-to-learn of novel task sets in the motor domain. *Frontiers in Psychology*, *22*.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., & Lillicrap, T. (2016). One-shot learning with memory-augmented neural networks. *arXiv preprint arXiv:1605.06065*.
- Thrun, S., & Pratt, L. (Eds.). (1998). *Learning to learn*. Kluwer Academic Publishers.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., ... Botvinick, M. (2016). Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763v2*.
- Weinstein, A., & Littman, M. L. (2013). Open-loop planning in large-scale stochastic domains. In *AAAI conference on artificial intelligence* (Vol. 27, pp. 1436–1442).
- Wolpert, D., Ghahramani, Z., & Jordan, M. (1995). An internal model for sensorimotor integration. *Science*, *269*(5232), 1880–1882.