

---

# Bootstrap Latent-Predictive Representations for Multitask Reinforcement Learning

---

Zhaohan Daniel Guo<sup>\*1</sup> Bernardo Avila Pires<sup>\*1</sup> Bilal Piot<sup>1</sup> Jean-Bastien Grill<sup>2</sup> Florent Alché<sup>2</sup>  
Rémi Munos<sup>2</sup> Mohammad Gheshlaghi Azar<sup>1</sup>

## Abstract

Learning a good representation is an essential component for deep reinforcement learning (RL). Representation learning is especially important in multitask and partially observable settings where building a representation of the unknown environment is crucial to solve the tasks. Here we introduce Predictions of Bootstrapped Latents (PBL), a simple and flexible self-supervised representation learning algorithm for multitask deep RL. PBL builds on multistep predictive representations of future observations, and focuses on capturing structured information about environment dynamics. Specifically, PBL trains its representation by predicting latent embeddings of future observations. These latent embeddings are themselves trained to be predictive of the aforementioned representations. These predictions form a bootstrapping effect, allowing the agent to learn more about the key aspects of the environment dynamics. In addition, by defining prediction tasks completely in latent space, PBL provides the flexibility of using multimodal observations involving pixel images, language instructions, rewards and more. We show in our experiments that PBL delivers across-the-board improved performance over state of the art deep RL agents in the DMLab-30 and Atari-57 multitask suites.

## 1. Introduction

Deep reinforcement learning (RL) has seen many successes in the recent years (Mnih et al., 2015; Levine et al., 2016; Silver et al., 2017; Vinyals et al., 2019). However, there are still many improvements to be made in complex, mul-

---

<sup>\*</sup>Equal contribution <sup>1</sup>DeepMind, London, UK <sup>2</sup>DeepMind Paris, France. Correspondence to: Mohammad Gheshlaghi Azar <mazar@google.com>.

titask and partially observable environments. Previous work (Mirowski et al., 2017; Jaderberg et al., 2017; Hessel et al., 2019; Gregor et al., 2019) has demonstrated that augmenting deep RL agents with auxiliary tasks for representation learning helps improve RL task performance. Representation learning has also had broader impact in deep RL, as representations have been used to generate intrinsic motivation signal for exploration (Houthoofd et al., 2016; Pathak et al., 2017; Burda et al., 2019; Azar et al., 2019; Puigdomènech Badia et al., 2020) and as the basis for search models in planning agents (Oh et al., 2017; Amos et al., 2018; Hafner et al., 2019; Schrittwieser et al., 2019).

A common approach for representation learning is through training a (recurrent) neural network (RNN) to make predictions on future observations in a self-supervised fashion (Oh et al., 2017; Amos et al., 2018; Moreno et al., 2018; Guo et al., 2018; Hafner et al., 2019; Schrittwieser et al., 2019; Gregor et al., 2019). Predictive models for representation learning have been widely used as auxiliary tasks for deep RL (Jaderberg et al., 2017; Oord et al., 2018; Zhang et al., 2019; Aytar et al., 2018; Gregor et al., 2019; Song et al., 2020). However, to learn a rich and useful representation, these methods may demand a high-degree of accuracy in multistep prediction of future observations (Gregor et al., 2015; 2019). This degree of accuracy can be difficult to achieve in many problems, especially in partially observable and multitask settings, where uncertainty in the agent state and complex, diverse observations make prediction more challenging. In these settings, instead of trying to predict everything accurately, we would ideally want to predict a latent embedding of future observations that focuses on the key, structural information of the tasks.

In this paper, we introduce Predictions of Bootstrapped Latents (PBL, “pebble”), a new representation learning technique for deep RL agents. PBL learns a representation of the history (*agent state*, Sutton & Barto, 2018) by predicting the future *latent* embeddings of observations. The novel contribution is in how we train these latent embeddings, which we do by having the latent observation embedding at each timestep be predictive of the agent state at the same timestep. This leads to a bootstrapping effect, where the

agent states learn to predict future latent observations, and future latent observations learn to predict the corresponding future agent states. As an example, consider a task where an agent must find a key and use it to unlock a door. All histories that lead to an unlocked door must also contain the observation of finding the key. PBL trains the latent embedding of observations of the unlocked door to predict the agent states that capture these histories, which means the latent embedding is encouraged to encode information pertaining to the key. This example illustrates how representations trained with PBL may capture contextual, structured information about the dynamics of the environment. This is especially beneficial in settings with diverse and complex observations. Furthermore, PBL allows for bootstrapping information from the far future, since the predicted agent state is also predictive of future. Finally, by making predictions solely in latent space, PBL makes it easier to incorporate observations with multiple modalities such as images, voice and language instruction.

We evaluate PBL in the challenging partially observable, multitask setting, where the task indices are kept hidden from the agent (Brunskill & Li, 2013). The multitask setting is the standard setting to evaluate deep RL agents as it is important for an AI agent to perform well, across the board, on a variety of tasks (Finn et al., 2017; Rothfuss et al., 2018; Espeholt et al., 2018; Hessel et al., 2019; Song et al., 2020). As the benchmark, we use DMLab-30 (Beattie et al., 2016) task suite as a standard partially observable and multitask RL domain, and compare the performance of various representation learning techniques when they are used as auxiliary representation learning tasks for the PopArt-IMPALA agent (Hessel et al., 2019), i.e., both RL and representation learning techniques train the same RNN. In particular we report the performance of PBL, pixel control (state-of-the-art), DRAW (Gregor et al., 2019) and contrastive predictive coding (CPC) (Oord et al., 2018) in this setting. We also evaluate our approach on multi-task Atari-57 suite (Belle-mare et al., 2013), comparing its performance with the RL baseline (PopArt-IMPALA) with various auxiliary representation learning task. We show, in section 4, that PBL, is able to outperform all these methods both in terms of final performance and sample efficiency in DMLab-30 suite. PBL also slightly improves the performance in comparison with the other methods in Atari-57 suite. This result suggests that PBL can be used as an alternative for the existing representation learning methods in the context of multitask deep RL.

## 2. Background

**Partially Observable Environments.** We model environments as partially observable Markov decision processes (POMDP, Cassandra et al., 1994) where an agent

does not directly perceive the underlying states but only a partial view of the underlying states. More formally, a POMDP is a tuple  $(\mathbb{S}, \mathbb{A}, \mathbb{O}, p_{\mathbb{S}}, p_{\mathbb{O}}, r, \gamma)$ , where  $\mathbb{S}$  is the state space,  $\mathbb{A}$  the action space and  $\mathbb{O}$  the observation space. The transition probability kernel  $p_{\mathbb{S}}$  determines the probability of the next state  $S_{t+1}$  given the current state and action  $(S_t, A_t)$ :  $p_{\mathbb{S}}(S_{t+1} = s' | S_t = s, A_t = a)$ . The observation kernel  $p_{\mathbb{O}}$  dictates the observation process:  $p_{\mathbb{O}}(O_t = o | S_t = s)$ .  $r \in \mathbb{R}^{\mathbb{S} \times \mathbb{A}}$  represents the reward function and  $\gamma$  is the discount factor.<sup>1</sup>

Policies dictate the action process: A policy is a mapping from sequences of observations and actions to a distribution over actions, that is,  $\pi : \{\mathbb{O} \times (\mathbb{A} \times \mathbb{O})^n : n \in \mathbb{N}\} \rightarrow \Delta(\mathbb{A})$ , where  $\Delta(\mathbb{A})$  denotes the space of action distributions. For a fixed policy  $\pi$  and for all  $t \geq 0$ , we define recursively, the states  $S_t^\pi$ , the observations  $O_t^\pi$ , the histories  $H_t^\pi$  and actions  $A_t^\pi$  starting from some initial distribution  $\rho$  and following  $\pi$ :

$$\begin{aligned} \text{Initialization: } & \begin{cases} S_0^\pi \sim \rho, & O_0^\pi \sim p_{\mathbb{O}}(S_0^\pi), \\ H_0^\pi \doteq O_0^\pi, & A_0^\pi \sim \pi(H_0^\pi). \end{cases} \\ \text{Recurrence: } & \begin{cases} S_{t+1}^\pi \sim p_{\mathbb{S}}(S_t^\pi, A_t^\pi), & O_{t+1}^\pi \sim p_{\mathbb{O}}(S_{t+1}^\pi), \\ H_{t+1}^\pi \doteq (H_t^\pi, A_t^\pi, O_{t+1}^\pi), \\ A_{t+1}^\pi \sim \pi(H_{t+1}^\pi). \end{cases} \end{aligned}$$

The *partial history*  $H_{t,k}^\pi \doteq (H_t^\pi, A_t^\pi, \dots, A_{t+k-1}^\pi)$  is formed by the history  $H_t^\pi$  and the  $k$  subsequent actions  $A_t^\pi, \dots, A_{t+k-1}^\pi$ .

The RL (Sutton & Barto, 2018; Szepesvári, 2010) problem in the POMDP setting is finding the policy that maximizes the expected and discounted sum of rewards:

$$\max_{\pi} E \left( \sum_{t=0}^{\infty} \gamma^t r(S_t^\pi, A_t^\pi) \right).$$

A common strategy for RL in partially observable domains is to compress the full history  $H_t^\pi$  as the *agent state*  $B_t^\pi$  (Sutton & Barto, 2018) using a neural network and solve

$$\max_{\pi} E \left( \sum_{t=0}^{\infty} \gamma^t r(B_t^\pi) \right),$$

where  $r(B_t^\pi) \doteq E(r(S_t^\pi, A_t^\pi) | B_t^\pi)$ . Accordingly we consider the policy  $\pi$  as a mapping from the agent state  $B_t^\pi$  to distributions over actions as well. To avoid clutter, we will drop the  $\pi$  superscript from the POMDP variables.

**Predictive Representations.** Deep RL methods use neural networks to produce the agent state, and learn it from

<sup>1</sup>We note that the multitask learning setting (Brunskill & Li, 2013), where the identity of task is kept hidden from the agent, is absorbed into the POMDP formulation, since the unobserved state can also index different per-task dynamics and reward functions.

RL and self-supervised auxiliary tasks (Mnih et al., 2016; Jaderberg et al., 2017). A common approach to learn agent states is to learn predictive representations (Oh et al., 2015; Jaderberg et al., 2017; Oh et al., 2017; Amos et al., 2018; Guo et al., 2018; Ha & Schmidhuber, 2018; Moreno et al., 2018; Oord et al., 2018; Gregor et al., 2019; Hafner et al., 2019; Schrittwieser et al., 2019). One way to do so is to train representations to predict statistics of future observations conditioned on partial histories (Guo et al., 2018; Gregor et al., 2019). This technique uses RNNs to compress full histories  $H_t$  as  $B_t$  and partial histories  $H_{t,k}$  as  $B_{t,k}$ , as inputs to auxiliary prediction tasks.

Figure 1 outlines the RNN architecture for compressing histories. The horizontal direction shows the RNN




Figure 1. Recurrent architecture for compressing partial histories. Networks used for processing observations and actions have been omitted, and dashed lines connect histories and partial histories aligned in time.

$h_f$  compressing full histories, with actions and observations as inputs. The agent state is updated as  $B_{t+1} \doteq h_f(B_t, O_{t+1}, A_t)$  (with  $B_0$  equal zero).

The vertical direction shows the RNN  $h_p$  compressing partial histories, starting from agent states, and continuing with actions as the only inputs. The state of  $h_p$  is updated as

$$\begin{aligned} B_{t,1} &\doteq h_p(B_t, A_t) \\ B_{t,k+1} &\doteq h_p(B_{t,k}, A_{t+k}), \end{aligned}$$

One can use the agent state  $B_t$  as the basis (input) for the decision making of the RL algorithm (Jaderberg et al., 2017; Espeholt et al., 2018; Hessel et al., 2019). The partial history RNN  $h_p$  is only used for predictions about the future.

### 3. Predictions of Bootstrapped Latents (PBL)

Our method, Predictions of Bootstrapped Latents (PBL, ‘‘pebble’’), consists of two auxiliary prediction tasks: 1) a forward, action-conditional prediction from compressed partial histories to future latent observations; and 2) a reverse prediction from latent observations to agent states.

The forward, action-conditional prediction task of PBL consists of predicting latent embeddings,  $Z_{t+k} = f(O_{t+k})$ , of future observations, from a compressed partial history  $B_{t,k}$ :

$$E(Z_{t+k}|B_{t,k}) \approx E(Z_{t+k}|H_{t,k}) \quad (1)$$

$$= E(Z_{t+k}|H_t, A_t, \dots, A_{t+k-1}). \quad (2)$$

To learn rich representations we make many predictions into the future—we take  $k$  to range from one to a maximum index in the future, which we call the *horizon*.

To predict  $Z_{t+k}$ , we solve:

$$\min_{h \in \mathbb{H}, g \in \mathbb{G}} \sum_{t,k} \|g(B_{t,k}) - Z_{t+k}\|_2^2, \quad (3)$$

where  $\mathbb{H}$  and  $\mathbb{G}$  are hypothesis spaces induced by neural networks,  $g$  is a feed-forward neural network, and  $h = (h_f, h_p)$  are the RNN networks that compute  $B_t$  and  $B_{t,k}$ . We call this problem a *forward prediction* (from compressed partial histories to latents), and its schematic is given in fig. 2.

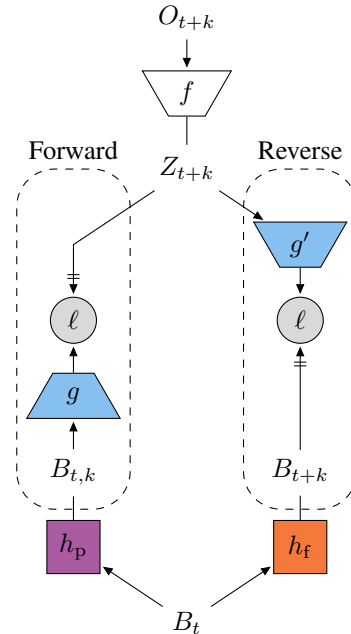


Figure 2. Architecture diagram for PBL, forward prediction (left) and reverse prediction (right). Crossed arrows block gradients and  $\ell(x, y) \doteq \|x - y\|_2^2$ .

The choice of  $Z_t$  is an essential part of our representation

learning. By picking  $Z_t$  to be a *latent* embedding of observations, we are able to learn meaningful observation features, as well as combine multiple modalities of observations (images and natural language sentences, for example) into a single representation. Specifically, PBL trains the latent embedding neural network  $f$  by predicting the learned  $B_t$  from  $Z_t$ . Concretely, we solve

$$\min_{f \in \mathbb{F}, g' \in \mathbb{G}'} \sum_t \|g'(\overbrace{f(O_t)}^{Z_t}) - B_t\|_2^2, \quad (4)$$

where  $\mathbb{F}$  and  $\mathbb{G}'$  are hypothesis spaces induced by neural networks, and  $g'$  is another feed-forward neural network. We call this problem *reverse prediction*, from latents to compressed histories, and its schematic is also given in fig. 2. See algorithm 1 for a high-level overview of the algorithm.

Forward prediction trains agent states to predict future latent observations. Reverse prediction trains latent observations to predict immediate agent states. Together, they can form a cycle, bootstrapping useful information between compressed histories and latent observations.

Because of the bootstrapping effect, the reverse prediction has the potential to train meaningful latent embeddings of observations that encode structural information about the paths that the agent takes to reach these observations. Furthermore, since we are training agent states to be predictive, this bootstrapping effect may also bootstrap useful information from far into the future.

Note that both forward and reverse predictions are one-way, i.e., we do not pass gradients into the targets of the predictions. In practice, we optimize both objectives together, as they optimize disjoint sets of parameters. A natural question that arises is what makes the bootstrapped predictions work, as it seems like there is the possibility that the prediction tasks collapse into trivial solutions. We hypothesize that due to forward and reverse predictions being one-way, the training dynamics actively move away from the trivial solution. Initially, the latent embedding of observations and agent states are the output of randomly initialized neural networks, which can be considered as random, non-linear projections. The forward prediction encourages the agent state to move away from collapsing in order to accurately predict future random projections of observations. Similarly, the reverse prediction encourages the latent observation away from collapsing in order to accurately predict the random projection of a full history. As we continue to train forward and reverse predictions, this seems to result in a virtuous cycle that continuously enriches both representations. In section 4.3, we empirically analyze this issue.

## 4. Experiments

In this section, we present our main experimental results (section 4.1) where we show that PBL, when used as the auxiliary task to shape the representation, outperforms other representation learning auxiliary tasks in DMLab 30 benchmark. Then we empirically investigate the scalability of PBL with the horizon of prediction and the effect of the reverse prediction in section 4.2, the stability of PBL in section 4.3, the neural network architecture choice in section 4.4, and what is captured by the learned representation in section 4.6. We discuss implementation details in section 4.7. Additionally, we performed experiments in the standard Atari-57 domain (Bellemare et al., 2013), which we report in section 4.5.

### 4.1. DMLab 30

We evaluated PBL in the DMLab 30 task set (Beattie et al., 2016). In this benchmark the agent must perform navigation and puzzle-solving tasks in 3D environments. There are a total of 30 tasks, each with a number of procedural variations over episodes. In the multitask setup of DMLab 30 each episode may come from a different task, and the agent *does not* receive the task identifier as an input. The only input the agent receives is an image of the first-person view of the environment, a natural language instruction (in some tasks), and the reward.

We compared PBL with different representation learning techniques, when they are used as auxiliary tasks for the deep RL agent. Pixel control (Jaderberg et al., 2017) is the current state-of-the-art self-supervised representation learning technique used in DMLab 30, having been applied to PopArt-IMPALA (Hessel et al., 2019) and, more recently, V-MPO (Song et al., 2020). CPC has been successful for representation learning in different settings (Oord et al., 2018). One advantage of CPC is its simplicity; it is lightweight and easy to incorporate in different applications. Simcore DRAW (Gregor et al., 2019) has achieved strong performance in some of the DMLab 30 tasks in the single-task regime. It is also a good method to compare with PBL, since it learns the representation with multistep predictions of the future. In contrast to CPC (which uses a contrastive approach for representation learning), Simcore DRAW uses a strong generative model of frames (Gregor et al., 2015).

All the representation learning techniques have been combined with PopArt-IMPALA (Hessel et al., 2019) as the base RL algorithm. PopArt-IMPALA has been shown to perform well in DMLab 30, where the use of PopArt makes the agent suitable for handling different reward scales across tasks. For PBL, Simcore DRAW, and CPC, we predict up to 20 steps into the future.

The most common performance measure for DMLab 30

**Algorithm 1** Training Step Pseudocode for PBL

---

**Require:** Minibatch of trajectories  $B = \{O_t^{(i)}, A_t^{(i)}, R_t^{(i)}\}$ , RNN  $h_p$ , RNN  $h_f$ , MLPs  $g, g', f$ , future prediction horizon  $k$ , RLLoss (reinforcement learning loss)

Encode observation  $Z_t^{(i)} \doteq f(O_t^{(i)})$

Let  $B_0^{(i)} \doteq \mathbf{0}$  and  $B_t^{(i)} \doteq h_f(B_{t-1}^{(i)}, O_t^{(i)}, A_{t-1}^{(i)}) \doteq B_{t,0}^{(i)}$  and  $B_{t,k}^{(i)} \doteq h_p(B_{t,k-1}^{(i)}, A_{t+k-1}^{(i)})$

Forward( $B_t^{(i)}$ )  $\doteq \frac{1}{k} \sum_{j=1}^k \|g(B_{t,j}^{(i)}) - \text{StopGradient}(Z_{t+j}^{(i)})\|_2^2$

Reverse( $Z_t^{(i)}$ )  $\doteq \|g'(Z_t^{(i)}) - \text{StopGradient}(B_t^{(i)})\|_2^2$

Take gradient step of  $\min_{\frac{1}{|B|} \sum_{i,t}} (\text{Forward}(B_t^{(i)}) + \text{Reverse}(Z_t^{(i)}) + \text{RLLoss}(B_t^{(i)}, R_t^{(i)}))$

---

is the mean over tasks of the capped human normalized score. The human normalized score of an agent on a task  $i$  is calculated as  $100 \cdot (a_i - u_i) / (h_i - u_i)$ , where  $a_i$  is the agent’s expected return on that task,  $u_i$  is the expected return of an agent that selects actions uniformly at random, and  $h_i$  is the expected return of a human performing the task. Reference values for the random and human scores can be found in the work of [Espeholt et al. \(2018\)](#). The *capped* human normalized score of an agent on a task is the human normalized score capped at 100. Therefore, the mean capped human normalized emphasizes performance gains toward attaining human performance across tasks and disregards improvements that are superhuman. In this work, we primarily report this score when aggregating results across tasks, but the mean human normalized score can be found in appendix B.

Figure 3 shows the mean capped human normalized score for PBL and our baselines, averaged across 16 independent runs, and with 95% confidence intervals. Appendix A.10 gives additional details on the plotting protocol for our experiments. We see that PBL outperforms all other methods

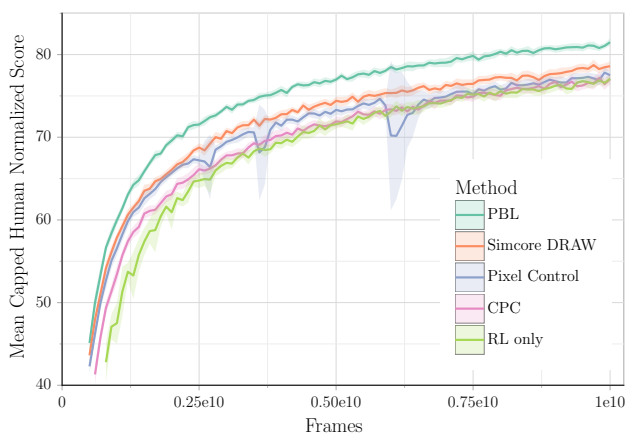


Figure 3. Mean capped human normalized score for compared methods.

throughout the whole training process by a significant margin. Since we are comparing how representation learning

techniques lead to performance improvement, it is worth comparing how each method improves over the RL only baseline. The mean capped human normalized score of this baseline, over the last 5% ( $5 \cdot 10^8$  frames) of training, is 76.70. Pixel control, the previous published state of the art in DMLab 30, improves final performance from to 77.59. PBL achieves 81.80, which is 5.31 times the gain with pixel control. By comparison, Simcore DRAW, the second best in fig. 3, attains 78.62—2.16 times the gain with pixel control, and thus only half of the gains obtained by using PBL.

While Simcore DRAW has been show to deliver strong improvements when training in single tasks, we believe that it may struggle with the large variety of tasks. The large diversity of images in multitask environments greatly increases the challenge of pixel prediction. Also, losses in pixel space can overlook details that are small but important for solving the tasks. Similarly, diverse images produced by different tasks are easy to distinguish from each other, meaning that CPC may not require a rich and informative representation to distinguish between the positive and negative examples, resulting in worse overall performance.

We show a per-task breakdown of PBL final performance relative to pixel control in fig. 4. We see that PBL is able to boost performance across the majority of the tasks, while achieving parity with pixel control on the rest of the tasks. This result shows that the representation learned by PBL encodes meaningful information which is useful to improve the performance of the agent across the board.

## 4.2. Scalability with the Horizon of Prediction

To have a better understanding of the effect of horizon on the performance of PBL we tried different horizon lengths for forward prediction in DMLab-30. Figure 5 shows that indeed, the performance improves monotonically as we predict farther into the future. However, as we increase the horizon, we get diminishing performance gains. Interestingly, predicting only one step into the future matches the performance of pixel control. This shows that multistep prediction enables PBL to encode crucial information for solving the tasks in its representation. To gain a better




Figure 4. Difference in Human Normalized Score per Task between PBL and pixel control. Error bars denote 95% confidence intervals.

understanding on the effect of the prediction horizon, we investigate the case where we disable reverse prediction (see fig. 2), which means the latents are mere random projections of the observations. Interestingly, we found that increasing the horizon makes no difference to performance, so we only report the performance for horizon 20 in fig. 5. Overall, it is slightly worse than pixel control during training, but achieves the same final performance. These results emphasize that it is not enough to only try to predict farther into the future, but that it is crucial to predict meaningful latents about the future. This suggests that reverse prediction in PBL is able to learn a meaningful representation of future observations that does extends far into the future.

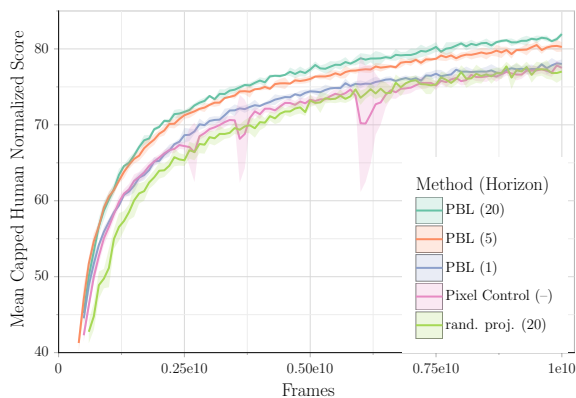


Figure 5. PBL Performance Across Forward Prediction Horizon, compared to pixel control and random projection.

### 4.3. Stability of PBL

From our experimental results, PBL is perfectly stable, and indeed does not collapse into a trivial solution. To further investigate this issue, we have tested a variant of PBL combined with random projection—we predict two separate latent observation embeddings where one is learned as usual and the other is a random projection. This variant grounds the forward prediction, which explicitly prevents collapsing. Interestingly, this variant achieves identical performance to PBL. This suggests that the training dynamics coupled with randomly initialized neural networks provide enough signal at the beginning to move away from collapse. As mentioned in section 3, we suspect that the forward and reverse predictions seem to form a virtuous cycle that continuously enriches both the compressed full history and the latent observations, instead of driving them to a trivial solution.

One subtle note is that for these experiments in DMLab-30, the agent state  $B_t$  is simultaneously trained with the RL loss as well as with PBL. However we have performed alternative experiments that also do not collapse where we have tried training two separate agent states where one is trained solely through PBL, and another is trained solely through RL but with the PBL agent state inputted as an extra observation (without passing gradients back). Also, in the upcoming section 4.6, the experiments have no RL at all, and the representations are purely trained using PBL.

While it might seem from these experiments that it is easy to avoid collapse, we have also tried to see what happens if instead of predicting future latent observations, we directly try predict future agent states using the forward prediction, i.e. predict  $B_{t+k}$  from  $B_t$  (without passing gradients through to the target). Unfortunately, even with the the additional signal from the RL loss, the representation collapses and the agent performance is completely flat.

### 4.4. Architecture Choice

In our experiments, we used a variant of the architecture used by Gregor et al. (2019); Song et al. (2020) (see table 1 in appendix A for full details). We compared the base RL method (PopArt-IMPALA with no auxiliary tasks) with this larger architecture, against the same method with the architecture used by Hessel et al. (2019).

Figure 6 shows the mean capped human normalized score of the base RL method with the two architectures, across four different independent runs. We chose to adopt the larger architecture because it gives a significant boost in overall performance.




Figure 6. Mean capped human normalized score for compared architectures.

#### 4.5. Atari-57

We also ran PBL along with pixel control, CPC and RL-only on the Atari-57 (Bellemare et al., 2013) in a *multitask* setup (Espeholt et al., 2018; Hessel et al., 2019). The main goal of this experiment is to show that PBL is a general representation approach that can be useful in benchmarks other than DMLab 30. We choose Atari-57 since it is a standard benchmark in deep RL.

The result from Simcore DRAW is omitted as it is designed to output RGB images in the DMLab 30 format, and is incompatible with standard stacked, greyscale Atari frames format. We only changed two hyperparameters from DMLab 30: one is the entropy cost, which is now set to 0.01; the second is the unroll length, which is set to 30. In Espeholt et al. (2018), the unroll length was set to 20, but we increased it to 30 which facilitates multistep prediction in the case of PBL and CPC. We also clip the rewards to be between  $-1$  and  $1$  which is standard for Atari-57 benchmark Espeholt et al. (2018).

Figure 7 shows the human normalized score (median across all 57 levels) for the compared approaches. We see that PBL is able to improve the overall performance in this benchmark. Interestingly, it seems that adding auxiliary tasks for representation learning only offers at most minimal improvements in performance. This may be due to Atari being essentially a fully observable domain as opposed to the partially observable nature of DMLab.

Table 10 shows the performance of different methods across the 57 Atari tasks with 95% confidence intervals. PBL outperforms the other methods, including pixel control, for seven tasks. In the other tasks, no method statistically outperforms PBL.

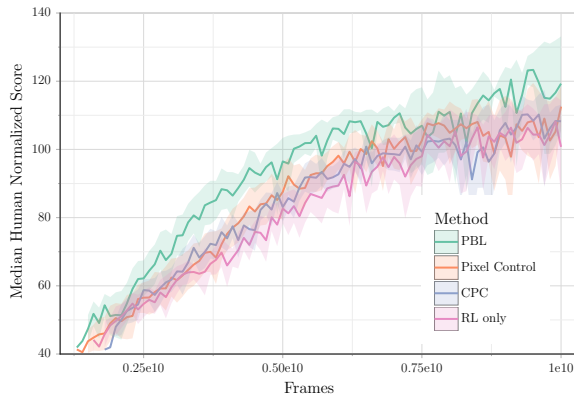


Figure 7. Median human normalized score for compared methods on Atari57.

#### 4.6. Decoding the PBL Representation

In this section, we try to shed some light on what information is being learned by the representation trained with PBL. We designed a simple DMLab-like 3D environment in which the agent is in a small room. There is a single red cube that spawns in a random location in the room. We use a fixed, uniformly random policy and only train the representation without any reinforcement learning.

Our criteria for a good representation is one that is able to encode the position of the red cube once the agent has seen it. To measure this, we train (in tandem with the agent’s representation) a predictor from the agent’s representation to the indicator of the position of the object in a top-down view of the environment. This position predictor does not affect the agent’s representation (that is, no gradients pass from the predictor to the representation)—this type of predictor is known as a probe (Amos et al., 2018) or a glass-box evaluator (Guo et al., 2018). Figure 8 shows a first-person view of the room and a top-down grid-view of the environment indicating the ground-truth position of the object. We




Figure 8. 3D room example: Agent’s first-person view (left) and top-down grid-view indicating the object position (right).

compared PBL against the random projection version where reverse prediction is disabled. Figure 9 shows the loss for the object position predictor between these two methods, and clearly indicates that PBL is learning a representation that is significantly better at encoding the position of the

object. Taking a qualitative look at what happens during

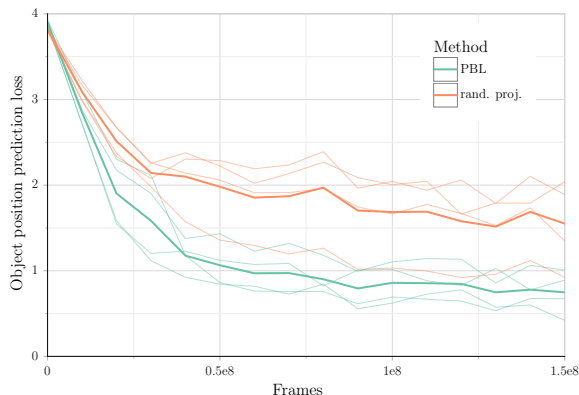


Figure 9. Object position prediction loss for random projection vs PBL. Light lines denote different independent runs.

individual episodes, we see that in videos of the glass-box evaluation, all methods are able to accurately localize the object right after the agent sees the object. However after the agent turns away, PBL’s representation is able to remember the position of the object, while RP’s representation quickly forgets. This shows an intimate connection between predicting the future and effectively compressing the past history—if the agent can predict that at some point far in the future, it will turn around and see the object again, then that means the representation has remembered the position of the object.

#### 4.7. Implementation Overview

In this section, we give an overview of the most relevant implementation details. A full account of our method, with details on architecture choices and parameters, can be found in appendix A.

For the agent architecture (used in all methods), we use the same as Hessel et al. (2019) except for the following changes: 1) We increased the LSTM from one 256 unit hidden layer to two 512 unit hidden layers; 2) We changed the linear policy and value heads to have each an additional 512 unit hidden layer. This is in fact akin to Gregor et al. (2019) but with wider layers. We duplicate the LSTM architecture to use for the partial history unroll. Our larger architecture resulted in strictly better performance for all methods.

For PBL, we duplicate the agent’s observation-processing architecture to use for the latent embedding ( $f$  in fig. 2). For the forward and reverse prediction networks ( $g$  and  $g'$  in fig. 2) we use MLPs with 2 hidden layers of 512 units each.

Most hyperparameters (e.g. PopArt hyperparameters) match the ones in Hessel et al. (2019). However, one significant difference is that we do not use population-based training

(PBT). Instead, we picked the optimizer parameters and entropy penalty that worked best for the base RL agent (without any representation learning). For all representation learning methods, the only hyperparameter we tuned was the scale of the representation loss relative to the RL loss.

Computationally, predicting every future latent from one step ahead to the horizon from every timestep can be very expensive. We mitigate this issue by subsampling from which timesteps we do the forward prediction, and which future latents we predict (Gregor et al., 2019). This does not change the objective, and in practice can greatly speed up computation with minimal loss of performance.

## 5. Related Work

Littman et al. (2001); Singh et al. (2003) introduced predictive state representations (PSRs) and an algorithm for learning such representations. PSRs are akin to action-conditional predictions of the future where the prediction tasks are indicators of events on the finite observation space. The power of PSRs is the ability to act as a compact replacement for the state that can be constructed from observable quantities. That is, PSRs are compact sufficient statistics of the history. Inspired by this line of research, we focus on learning neural representations that are predictive of observations multiple timesteps into the future, conditioned on the agent’s actions.

A number of other works investigated learning predictive representations (with action-conditional predictions of the future) to improve RL performance. Oh et al. (2015) use generated observations in addition to the actions when compressing partial histories. Amos et al. (2018) learned representations by maximizing the likelihood of future proprioceptive observations using a PreCo model, which interleaves an action-dependent (predictor) RNN with an observation dependent (corrector) RNN for compressing full histories, and uses the predictor RNN for the action-only sequences in partial histories (*cf.* fig. 1, where one RNN is used for the full histories, and another for the partial histories). Oh et al. (2017); Schrittwieser et al. (2019) learn to predict the base elements needed for Monte Carlo tree search conditioned on actions (Schrittwieser et al., 2019) or options (Oh et al., 2015): Rewards, values, logits of the search prior policy (in the case of Schrittwieser et al., 2019), and termination/discount (in the case of Oh et al., 2017). Guo et al. (2018); Moreno et al. (2018); Gregor et al. (2019) used an architecture similar to fig. 1 with a particular focus on learning representations of belief state.

Ha & Schmidhuber (2018); Hafner et al. (2019) used variational autoencoders (VAEs, Kingma & Welling, 2014; Gregor et al., 2015) to shape the full history RNN representation, and trained latent models from the partial histories



to maximize likelihood of the respective full histories. The Simcore DRAW technique (Gregor et al., 2019) is a demonstrably strong VAE-based representation learning technique for single task deep RL. DRAW trains the representation by trying to autoencode observations. In contrast PBL tries to learn meaningful features of observations related to important dynamics, rather than a generative model of pixel images.

Jaderberg et al. (2017) introduced pixel control as a representation learning task for reinforcement learning. This technique applies Q-learning for controlling the change in pixels, which can be considered as a technique for learning predictive representations based on future pixel changes.

(CPC Oord et al., 2018) is a powerful and lightweight alternative to generative models in representation learning. Similar to PBL, CPC operates in latent space. In contrast to our approach, which is based on predicting latent embeddings, CPC uses a contrastive approach to representation learning. (For an overview of CPC, see appendix A.5).

Other work that strives to learn a latent agent state representation include (DeepMDP Gelada et al., 2019) and (CRAR François-Lavet et al., 2019) that try to learn a transition model entirely in latent space. However DeepMDP relies on learning the reward function to ground its representation, making it very difficult to learn in sparse reward settings, and CRAR explicitly depends on a form of entropy maximisation to prevent collapse.

The idea of using bootstrapped latent embeddings to train the representations have been used recently in the context self-supervised learning of image representations (Grill et al., 2020). In this case Grill et al. (2020) show that one can improve the quality of the learned representation by using only one network for the latent embedding. The prediction task then consists of an *online* embedding network aiming at predicting the output of a *target* embedding network. The target network in this case is essentially a moving average of the previous copies of the online network. Thus Grill et al. (2020) remove the need to train separate embeddings of the observations.

## 6. Conclusion

We considered the problem of representation learning for deep RL agents in partially observable, multitask settings. We introduced Predictions of Bootstrapped Latents (PBL), a representation learning technique that provides a novel way of learning meaningful future latent observations through bootstrapped forward and reverse predictions. We demonstrated that PBL outperforms the state-of-the-art representation learning technique in DMLab-30.

Our results show that agents in multitask setting significantly

benefit from learning to predict a *meaningful* representation of the future in response to its actions. This is highlighted in the fact that when we simply used random projections of future observations, there was no benefit in trying to predict farther into the future. On the other hand, by learning meaningful latent embeddings of observations, PBL continued to improve as we increased the horizon. Going forward, a promising direction would be to investigate ways to predict more relevant and diverse information, over long horizons.

Latent-space prediction is an easy way to combine multiple observation modalities. It would be worth seeing how far we can push the quality of learned representations by incorporating richer sources of observation. To name a few, proprioception in robotics and control tasks (Levine et al., 2016; Tassa et al., 2018), more sophisticated natural language, and even other sensory inputs like touch, smell, sound and temperature. In fact, the ability to learn representations from multimodal data makes PBL a good candidate for evaluation in other machine learning domains, beyond RL, such as the settings considered by Oord et al. (2018).

A promising extension to PBL may be to generalize the forward and reverse predictions from learning mean predictions (with squared loss) to learning generative models of latents. While this may not necessarily help in DMLab-30 where dynamics are deterministic, it may help in other settings where the dynamics are stochastic. Having a generative model may also be used to generate rollouts for Monte Carlo planning (similar to Schrittwieser et al., 2019).

We think it would be fruitful to evaluate PBL in a transfer learning setting. Since PBL improves performance virtually across all DMLab-30 tasks, it may have learned a representation that better captures the shared, latent structure across tasks, with the potential to generalize to new unseen tasks

Interesting future work also includes a more in-depth study of what predictive representations learn to encode, perhaps using a glass-box approach (Guo et al., 2018; Amos et al., 2018). In particular, it is worth finding out how much information about the POMDP *state* is captured and preserved by these representations.

## Acknowledgements

We thank Hamza Merzic, Hubert Soyer and Satinder Singh Baveja for their support and useful feedback, and the ICML 2020 anonymous reviewers and program committee for their feedback and discussions.

## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard,

- M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- Amos, B., Dinh, L., Cabi, S., Rothörl, T., Colmenarejo, S. G., Muldal, A., Erez, T., Tassa, Y., de Freitas, N., and Denil, M. Learning awareness models. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- Aytar, Y., Pfaff, T., Budden, D., Paine, T. L., Wang, Z., and de Freitas, N. Playing hard exploration games by watching YouTube. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 2935–2945, 2018.
- Azar, M. G., Piot, B., Pires, B. A., Grill, J.-B., Altché, F., and Munos, R. World discovery models. *arXiv preprint arXiv:1902.07685*, 2019.
- Beattie, C., Leibo, J. Z., Teplyaev, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., Schrittwieser, J., Anderson, K., York, S., Cant, M., Cain, A., Bolton, A., Gaffney, S., King, H., Hassabis, D., Legg, S., and Petersen, S. DeepMind lab, 2016.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Brunskill, E. and Li, L. Sample complexity of multi-task reinforcement learning. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence, UAI'13*, pp. 122–131, Arlington, Virginia, USA, 2013. AUAI Press.
- Burda, Y., Edwards, H., Storkey, A. J., and Klimov, O. Exploration by random network distillation. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- Cassandra, A. R., Kaelbling, L. P., and Littman, M. L. Acting optimally in partially observable stochastic domains. In *Proceedings of the 12th National Conference on Artificial Intelligence, Seattle, WA, USA, July 31 - August 4, 1994, Volume 2*, pp. 1023–1028, 1994.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., and Kavukcuoglu, K. IMPALA: scalable distributed deep-RL with importance weighted actor-learner architectures. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, pp. 1406–1415, 2018.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- François-Lavet, V., Bengio, Y., Precup, D., and Pineau, J. Combined reinforcement learning via abstract representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3582–3589, 2019.
- Gelada, C., Kumar, S., Buckman, J., Nachum, O., and Bellemare, M. G. Deepmdp: Learning continuous latent space models for representation learning. *arXiv preprint arXiv:1906.02736*, 2019.
- Gregor, K., Danihelka, I., Graves, A., Rezende, D. J., and Wierstra, D. DRAW: A recurrent neural network for image generation. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pp. 1462–1471, 2015.
- Gregor, K., Jimenez Rezende, D., Besse, F., Wu, Y., Merzic, H., and van den Oord, A. Shaping belief states with generative environment models for RL. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 13475–13487. Curran Associates, Inc., 2019.
- Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- Guo, Z. D., Azar, M. G., Piot, B., Pires, B. A., Pohlen, T., and Munos, R. Neural predictive belief representations. *arXiv preprint arXiv:1811.06407*, 2018.
- Ha, D. and Schmidhuber, J. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 2455–2467, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. In Chaudhuri, K. and Salakhutdinov, R.

- (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2555–2565, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778, 2016a. doi: 10.1109/CVPR.2016.90.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016b.
- Hessel, M., Soyer, H., Espeholt, L., Czarnecki, W., Schmitt, S., and van Hasselt, H. Multi-task deep reinforcement learning with PopArt. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3796–3803, 2019.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735.
- Houthoofd, R., Chen, X., Duan, Y., Schulman, J., Turck, F. D., and Abbeel, P. VIME: variational information maximizing exploration. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 1109–1117, 2016.
- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., and Kavukcuoglu, K. Reinforcement learning with unsupervised auxiliary tasks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(1):1334–1373, January 2016. ISSN 1532-4435.
- Littman, M. L., Sutton, R. S., and Singh, S. P. Predictive representations of state. In *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]*, pp. 1555–1561, 2001.
- Mirowski, P., Pascanu, R., Viola, F., Soyer, H., Ballard, A., Banino, A., Denil, M., Goroshin, R., Sifre, L., Kavukcuoglu, K., Kumaran, D., and Hadsell, R. Learning to navigate in complex environments. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pp. 1928–1937, 2016.
- Moreno, P., Humplik, J., Papamakarios, G., Pires, B. A., Buesing, L., Heess, N., and Weber, T. Neural belief states for partially observed domains. In *NeurIPS 2018 workshop on Reinforcement Learning under Partial Observability*, 2018.
- Nair, V. and Hinton, G. E. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pp. 807–814, 2010.
- Oh, J., Guo, X., Lee, H., Lewis, R. L., and Singh, S. P. Action-conditional video prediction using deep networks in Atari games. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 2863–2871, 2015.
- Oh, J., Singh, S., and Lee, H. Value prediction network. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pp. 6118–6128, 2017.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pp. 2778–2787, 2017.
- Puigdomènech Badia, A., Sprechmann, P., Vitvitskiy, A., Guo, D., Piot, B., Kapturowski, S., Tieleman, O., Arjovsky, M., Pritzel, A., Bolt, A., and Blundell, C. Never

- give up: Learning directed exploration strategies. In *International Conference on Learning Representations*, 2020.
- Rothfuss, J., Lee, D., Clavera, I., Asfour, T., and Abbeel, P. Promp: Proximal meta-policy search. *arXiv preprint arXiv:1810.06784*, 2018.
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., et al. Mastering atari, go, chess and shogi by planning with a learned model. *arXiv preprint arXiv:1911.08265*, 2019.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of Go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Singh, S. P., Littman, M. L., Jong, N. K., Pardoe, D., and Stone, P. Learning predictive state representations. In *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pp. 712–719, 2003.
- Song, H. F., Abdolmaleki, A., Springenberg, J. T., Clark, A., Soyer, H., Rae, J. W., Noury, S., Ahuja, A., Liu, S., Tirumala, D., Heess, N., Belov, D., Riedmiller, M., and Botvinick, M. M. V-MPO: On-policy maximum a posteriori policy optimization for discrete and continuous control. In *(to appear) 8th International Conference on Learning Representations, ICLR 2020*, 2020.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Szepesvári, C. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., et al. DeepMind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M. J., and Levine, S. SOLAR: Deep structured representations for model-based reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pp. 7444–7453, 2019.