PROTO-VALUE NETWORKS: SCALING REPRESENTATION LEARNING WITH AUXILIARY TASKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Auxiliary tasks improve the representations learned by deep reinforcement learning agents. Analytically, their effect is reasonably well-understood; in practice, however, their primary use remains in support of a main learning objective, rather than as a method for learning representations. This is perhaps surprising given that many auxiliary tasks are defined procedurally, and hence can be treated as an essentially infinite source of information about the environment. Based on this observation, we study the effectiveness of auxiliary tasks for learning rich representations, focusing on the setting where the number of tasks and the size of the agent's network are simultaneously increased. For this purpose, we derive a new family of auxiliary tasks based on the successor measure. These tasks are easy to implement and have appealing theoretical properties. Combined with a suitable off-policy learning rule, the result is a representation learning algorithm that can be understood as extending Mahadevan & Maggioni (2007)'s proto-value functions to deep reinforcement learning – accordingly, we call the resulting object *proto-value networks*. Through a series of experiments on the Arcade Learning Environment, we demonstrate that proto-value networks produce rich features that may be used to obtain performance comparable to established algorithms, using only linear approximation and a small number (~4M) of interactions with the environment's reward function.

1 INTRODUCTION

In deep reinforcement learning (RL), an agent maps observations to a policy or return prediction by means of a neural network. The role of this network is to transform observations into a series of successively refined features, which are linearly combined by the final layer into the desired prediction. A common perspective treats this transformation and the intermediate features it produces as the agent's *representation* of its current state. Under this lens, the learning agent performs two tasks simultaneously: representation learning, the discovery of useful state features; and credit assignment, the mapping from these features to accurate predictions.

Although end-to-end reinforcement learning has been shown to obtain good performance in a wide variety of problems (Mnih et al., 2015; Levine et al., 2016; Bellemare et al., 2020), modern RL methods typically incorporate additional machinery that incentivizes the learning of good state representations: for example, predicting immediate rewards (Jaderberg et al., 2017), future states (Schwarzer et al., 2021a), or observations (Gelada et al., 2019); encoding a similarity metric (Castro, 2020; Agarwal et al., 2021a; Zhang et al., 2021); and data augmentation (Laskin et al., 2020). In fact, it is often possible, and desirable, to first learn a sufficiently rich representation with which credit assignment can then be efficiently performed; in that sense, representation learning has been a core aspect of RL from its early days (Sutton & Whitehead, 1993; Sutton, 1996; Ratitch & Precup, 2004; Mahadevan & Maggioni, 2007; Diuk et al., 2008; Konidaris et al., 2011; Sutton et al., 2011).

An effective method for learning state representations is to have the network predict a collection of auxiliary tasks associated with each state (Caruana, 1997; Jaderberg et al., 2017; Chung et al., 2019). In an idealized setting, auxiliary tasks can be shown to induce a set of features that correspond to the principal components of what is called the auxiliary task matrix (Bellemare et al., 2019; Lyle et al., 2021). This makes it possible to analyze the theoretical approximation error (Petrik, 2007; Parr et al., 2008), generalization (Le Lan et al., 2022), and stability (Ghosh & Bellemare, 2020) of the learned representation. Perhaps surprisingly, there is comparatively little that is known about their empirical

behaviour on larger-scale environments. In particular, the scaling properties of representation learning from auxiliary tasks – i.e., the effect of using more tasks, or increasing network capacity – remain poorly understood. This paper aims to fill this knowledge gap.

Our approach is to construct a family of auxiliary rewards that can be sampled and subsequently predicted by the network. Specifically, we implement the successor measure (Blier et al., 2021; Touati & Ollivier, 2021), which extends the successor representation (Dayan, 1993) by replacing state-equality with set-inclusion. In our case, these sets are defined implicitly by a family of binary functions over states, for example universal hash functions (Carter & Wegman, 1979). We conduct most of our studies on binary functions derived from randomly-initialized networks, whose effectiveness as random cumulants has already been demonstrated (Dabney et al., 2021).

Although our results may hold for other types of auxiliary rewards, our method has a number of benefits: it can be trivially scaled by sampling more random networks to serve as auxiliary tasks, it directly relates to the binary reward functions common of deep RL benchmarks, and can to some extent be theoretically understood. The actual auxiliary tasks consist in predicting the expected return of the random policy for their corresponding auxiliary rewards; in the tabular setting, this corresponds to proto-value functions (Mahadevan & Maggioni, 2007; Stachenfeld et al., 2014; Machado et al., 2018). Consequently, we call our method *proto-value networks* (PVN).

We study the effectiveness of this method on the Arcade Learning Environment (ALE) (Bellemare et al., 2013). Overall, we find that PVN produces state features that are rich enough to support linear value approximations that are comparable to those of DQN (Mnih et al., 2015) on a number of games, while only requiring a fraction of interactions with the environment reward function. We explore the features learned by PVN and show that they capture the temporal structure of the environment, which we hypothesize contributes to their utility when used with linear function approximation.

In an ablation study, we find that increasing the value network's capacity improves the performance of our linear agents substantially, and that larger networks can accommodate more tasks. Perhaps surprisingly, we also find that our method performs best with what might seem like small number of auxiliary tasks: the smallest networks we study produce their best representations from 10 or fewer tasks, and the largest, from 50 to 100 tasks. In a sense, this finding corroborates the result of Lyle et al. (2021, Fig. 5), where optimal performance (on a small set of Atari 2600 games and with the standard DQN network) was obtained with a single auxiliary task. From this finding we hypothesize that individual tasks may produce much richer representations than expected, and the effect of any particular task on fixed-size networks (rather than the idealized, infinite-capacity setting studied in the literature) remains incompletely understood.

2 Related work

Deep RL algorithms have employed auxiliary prediction tasks to learn representations with various emergent properties (Schaul et al., 2015; Jaderberg et al., 2017; Machado et al., 2018; Bellemare et al., 2019; Gelada et al., 2019; Fedus et al., 2019; Dabney et al., 2021; Lyle et al., 2022). While most of these papers optimize auxiliary tasks in support of reward maximization from online interactions, our work investigates learning representations solely from auxiliary tasks on offline datasets. Closely related to our work is the study of random cumulants (Dabney et al., 2021; Lyle et al., 2021), both of which identify random cumulant auxiliary tasks as being especially useful in sparse-reward environments. Our work differs from these prior works in both motivation and implementation. Notably absent in prior work on random cumulants is the study of representational capacity as a function of the number of tasks.

Another body of related work on decoupling representation learning from RL primarily revolves around the use of contrastive learning (Anand et al., 2019; Wu et al., 2019; Stooke et al., 2021; Schwarzer et al., 2021b; Erraqabi et al., 2022). Anand et al. (2019) proposed ST-DIM, a collection of temporal contrastive losses operating on image patches from environmental observations. Although the representations learned by ST-DIM are able to predict annotated state-variables in Atari 2600 games, their pretraining method was never evaluated for control. Stooke et al. (2021) uses contrastive learning for learning the temporal dynamics, resulting in minor improvements in online control from a fixed representation. Additionally, Schwarzer et al. (2021b) augments next-state prediction with goal-conditioned RL and inverse dynamics modelling, enabling strong performance on Atari 100k benchmark (Kaiser et al., 2020). Our work is complementary to these prior works and investigates

the utility of scaling auxiliary tasks for learning good representations, which in principle can be easily combined with existing approaches. Additionally, recent work on using state-similarity metrics tackles the representation learning problem through the lens of behavioral similarity (Castro et al., 2021; Zhang et al., 2021; Agarwal et al., 2021a). We note that, in contrast to our method, the behavioral metrics used in these works are heavily based on the reward structure of the environment.

Related to our method, Touati & Ollivier (2021) consider representation learning with the successor measure (see also Touati, 2021, Algorithm 7). Algorithmically, their approach differs from ours in a number of ways, including the use of a learned state density function in lieu of indicator functions, the decomposition of the successor measure into its so-called forward and backward representations, and a bespoke sampling procedure to generate sample trajectories from which the representation is learned. By comparison, our approach directly constructs a relevant set of auxiliary tasks, which results in a significantly simpler algorithm that is more easily scaled according to available computational resources and to the full gamut of Atari 2600 games, as we will demonstrate.

3 Background

The RL problem can be modeled as a Markov Decision Process (MDP) defined by the 5-tuple $\mathcal{M} = \langle \mathcal{X}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma \rangle$, in which $\mathcal{X}$ is a set of states, $\mathcal{A}$ is a set of actions, $\mathcal{R} : \mathcal{X} \times \mathcal{A} \mapsto \mathbb{R}$ is a scalar reward function, $\mathcal{P} : \mathcal{X} \times \mathcal{A} \mapsto \mathscr{P}(\mathcal{X})$ is a transition function that maps state-action pairs to a distribution over next states, and $\gamma \in [0, 1)$ is a discount factor. A policy $\pi : \mathcal{X} \mapsto \mathscr{P}(\mathcal{A})$ is a function that maps states to a distribution over actions.

The goal of an RL agent is to learn a policy that maximizes the cumulative discounted rewards from the environment, also known as the discounted return. The state-action value function is defined as the expected discounted return when starting in a state and following the policy π:

$$Q^\pi(x, a) := \mathop{\mathbb{E}}_{A_t \sim \pi, \mathcal{P}} \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(X_t, A_t) \,|\, X_0 = x, A_0 = a \right].$$

In this paper, we consider approximating the value function Q^π using a linear combination of features. We call the map $\phi : \mathcal{X} \to \mathbb{R}^k$ a *k-dimensional state representation*; $\phi(x)$ is the feature vector for a state $x \in \mathcal{X}$. The value function approximant at (x, a) is

$$\hat{Q}(x, a) = \phi(x)^\top w_a,$$

where $w_a \in \mathbb{R}^k$ is a weight vector associated with action a. In deep reinforcement learning, the state representation is parameterized by a neural network. Often, the representation is learned end-to-end by optimizing the parameters to make more accurate predictions about the value function. Additional predictions that further shape the state representation are called *auxiliary tasks* (Jaderberg et al., 2017). In this work, we write $\mathcal{T}$ for the set of auxiliary tasks.

The *successor representation* (SR; Dayan, 1993) encodes the temporal structure of the MDP in terms of which states can be reached from any other state. It is given by

$$\psi_{\text{SR}}(x, a, \tilde{x}) = \sum_{t=0}^{\infty} \gamma^t \mathbb{P}\{X_t = \tilde{x} \,|\, X_0 = x, A_0 = a\}.$$

A convenient, recursive form expresses the SR in terms of an indicator function:

$$\psi_{\text{SR}}(x, a, \tilde{x}) = \mathbf{1}\{x = \tilde{x}\} + \gamma \, \mathbb{E}_\pi \left[\psi_{\text{SR}}(X', A', \tilde{x}) \,|\, X = x, A = a \right].$$

This form highlights the fact that for each $\tilde{x}$, the successor representation is the value function associated with the reward function $\mathcal{R}(x, a) = \mathbf{1}\{x = \tilde{x}\}$.

4 Proto-value networks

In this section, we derive our *proto-value networks* algorithm. At a high level, this algorithm learns a state representation that approximates the singular vectors associated with the successor measure, the extension of the SR to continuous state spaces. We do this in order to derive an algorithm that is more suitably tailored to the large state spaces of deep RL domains, where many states are encountered once or never at all, and some notion of distance between states must be accounted for.

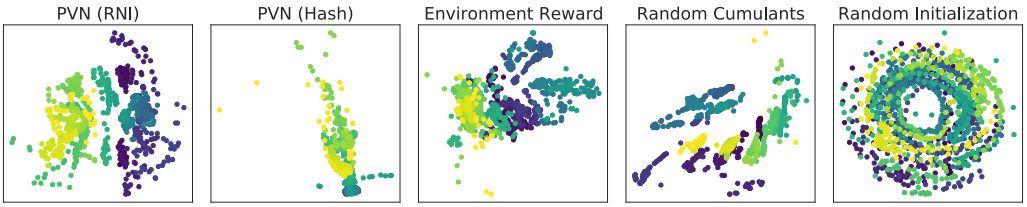

Figure 1: MDS plots of the features learned by PVN and other baseline methods. Each plotted point corresponds to a state in a single trajectory on the game CHOPPER COMMAND. Dark purple points correspond to states near the beginning of the episode, and light yellow points correspond to states at the end of the episode. More plots for different games can be found in Appendix C. Environment Reward corresponds to features learned by optimizing the environment reward directly.

To gain some understanding into this process, let us consider how the method of auxiliary tasks (Jaderberg et al., 2017) can be used to obtain a state representation that approximates the successor representation. In the tabular setting, where $\mathcal{X}$ and $\mathcal{T}$ are of finite sizes n and m respectively, we write the feature matrix $\Phi \in \mathbb{R}^{n\times d}$, so that each state x is associated with a feature vector $\phi(x) \in \mathbb{R}^d$. Given an auxiliary task matrix $\Psi \in \mathbb{R}^{n\times m}$, the method of auxiliary tasks can be shown to be equivalent to minimizing the loss function

$$\mathcal{L}(\Phi, W) = \|\Phi W - \Psi\|_F^2 = \sum_{x\in\mathcal{X}, i\in\mathcal{T}} \left(\phi(x)^\top w_i - \psi_i(x)\right)^2$$

jointly with respect to Φ and W. Here, $W \in \mathbb{R}^{d\times m}$ is a weight matrix with columns $(w_i)_{i=1}^m$ and $\psi_i(x)$ is the entry of Ψ corresponding to state x and task i. In the sequel, we will assume that a near-optimal W can be obtained easily and simply consider the loss

$$\mathcal{L}(\Phi) = \min_W \mathcal{L}(\Phi, W),$$

to be minimized over Φ. It is known (e.g., Bellemare et al., 2019) that any feature matrix that minimizes this loss function must have columns that lie in the subspace spanned by the top d left singular vectors of Ψ. In particular, when Ψ is square and symmetric the auxiliary task method recovers the subspace spanned by its top d eigenvectors.

Here, we are interested in the setting in which Ψ is the successor representation matrix for the *uniformly random policy*. In the symmetric case, the eigenvectors of Ψ form what is called the *proto-value functions* of the MDP (Mahadevan & Maggioni, 2007).[1] These eigenvectors are of special importance because they encode the spatial structure of the MDP in terms of a diffusion process, and have been shown to correlate with neural encodings of spatial location in mammals (Stachenfeld et al., 2014). They are thus promising candidates for learning rich state representations for deep reinforcement learning agents.

4.1 EXTENSION TO THE RANDOM SUCCESSOR MEASURE

Let π be a policy and Σ the power set of $\mathcal{X}$. The successor measure $\psi^\pi : \mathcal{X}\times\mathcal{A}\times\Sigma \to \mathbb{R}$ extends the successor representation to quantify the discounted visitation frequency of an agent, in expectation over the random trajectory and for *various subsets of the state space* (Blier et al., 2021). Given a set $S \subset \mathcal{X}$, we write

$$\psi^\pi(x, a, S) = \sum_{t=0}^{\infty} \gamma^t \mathbb{P}\{X_t \in S \,|\, X_0 = x, A_0 = a\}\,.$$

[1] Behzadian & Petrik (2018) gives the singular-vector extension for the asymmetric case. Because this extension is straightforward and symmetry rarely holds, in this paper we use the term proto-value networks to describe state representations learned in both the symmetric and asymmetric settings.

As with the successor representation, this can be expressed in terms of an expectation over an indicator function, and further decomposed in a Bellman equation:

$$\begin{aligned}\psi^\pi(x,a,S) &= \sum_{t=0}^{\infty} \mathbb{E}_\pi \left[\gamma^t \mathbf{1}\{X_t \in S\} \,|\, X_0 = x, A_0 = a\right] \\ &= \mathbf{1}\{x \in S\} + \gamma\, \mathbb{E}_\pi \left[\psi^\pi(X', A', S) \,|\, X = x, A = a\right].\end{aligned}$$

The passage from state equality to set inclusion is particularly appealing in deep RL: first, because states rarely repeat along a trajectory or between episodes, the indicator $\mathbf{1}\{x = y\}$ is almost always zero. Second, set inclusion allows us to incorporate a notion of closeness to ψ, for example by focusing on subsets S that include semantically similar states. We will return to this point later in the section.

By analogy with the tabular setting, let us now define a loss function which, if suitably minimized, should produce a useful state representation. For ease of exposition, we continue to assume that $\mathcal{X}$ is finite, although perhaps very large; the reader interested in a proper mathematical treatment of the full continuous-state setting is invited to consult Blier et al. (2021) and Pfau et al. (2019).

Let ξ be a distribution over subsets of states and $\Xi \in \mathbb{R}^{n \times n}$ the diagonal matrix with the elements ξ on its diagonal. The *Monte Carlo successor measure loss* is

$$\mathcal{L}_{MCSM}(\Phi) = \min_{w_{S,a} \in \mathbb{R}^d} \mathop{\mathbb{E}}_{S \sim \xi} \Big[\Big(\sum_{x \in \mathcal{X}, a \in \mathcal{A}} (\phi(x)^\top w_{S,a} - \psi(x,a,S)\Big)^2\Big].$$

Theorem 1. *If Φ^* is a feature matrix minimizing $\mathcal{L}_{MCSM}(\Phi)$, then its column space spans the top d left singular vectors of the (infinite-dimensional) successor measure matrix Ψ with respect to the inner product $(x,y)_\Xi = y^\top \Xi x$, for all $x, y \in \mathbb{R}^n$. Here, $\Xi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is the diagonal matrix whose entries are given by ξ.*

In practice, samples of $\psi(x,a,S)$ (which must be estimated from complete trajectories) are not available; instead, it is preferable to learn an approximation by bootstrapping (Sutton & Barto, 2018). The corresponding temporal-difference successor measure loss is

$$\min_{w_{S,a} \in \mathbb{R}^d} \mathop{\mathbb{E}}_{S \sim \xi} \Big[\Big(\sum_{x \in \mathcal{X}, a \in \mathcal{A}} (\mathbf{1}\{x \in S\} + \gamma\, \mathbb{E}\left[\phi(X')^\top w_{S,A'} \,|\, X = x, A = a\right] - \phi(x)^\top w_{S,a}\Big)^2\Big]; \quad (1)$$

we will use this form in the derivations that follow.

4.2 A PRACTICAL IMPLEMENTATION

Our algorithm aims to approximate the loss in Equation 1 using tools from deep RL. We first approximate the expectation over ξ by sampling a collection of sets $(S_i)_{i=1}^m$ from ξ. These sets are kept fixed throughout learning. With this in mind, each set corresponds to an indicator function that we treat as a binary reward function $r_i(x) = \mathbf{1}\{x \in S_i\}$. The actual auxiliary task is then the value function of the random policy associated with this reward.

Denote by $\hat{\psi}_i(x,a)$ the prediction made by our neural network for state x, action a, and the set S_i. Given a sample transition (x,a,x'), we define the sample target

$$r_i(x) + \gamma \frac{1}{|\mathcal{A}|} \sum_{a' \in \mathcal{A}} \hat{\psi}_i(x', a').$$

Notice that the average over the next-action a' arises as a consequence of taking the policy π to be uniformly random. We then train the neural network by performing stochastic gradient descent on the loss derived from this sample target:

$$\Big(r_i(x) + \gamma \frac{1}{|\mathcal{A}|} \sum_{a' \in \mathcal{A}} \hat{\psi}_i(x', a') - \hat{\psi}_i(x,a)\Big)^2.$$

Following common usage, the actual gradient estimate is obtained by aggregating multiple transitions into a minibatch and applying the Adam optimizer (Kingma & Ba, 2015).

Before explaining how the sets S_i are defined, let us remark on a number of appealing properties of these auxiliary tasks, when viewed from a deep RL perspective. First, the use of a random policy means that learning usually proceeds in an off-policy manner. However, we expect this to be a relatively mild form of off-policy learning, one that is in general much more stable than one derived by maximization, as in a Bellman optimality equation. Although one could also learn the value function associated with the current policy (as in SARSA (Rummery & Niranjan, 1994)), this precludes the use of offline datasets for learning the representation, or at least makes the learned representation strongly dependent on the behaviour policy. By contrast, the representation learned by proto-value networks only depends on the availability of data. In effect, these auxiliary tasks *depend only on the structure of the environment, and not on the agent's behaviour*.

We also expect binary reward functions to be easier to tune than, say, those derived from a distance function (dependent on getting the scale parameter correct) or real-valued random rewards (dependent on the underlying distribution). Binary rewards are particularly appealing in domains where the reward function is itself binary or ternary (i.e., Atari 2600 video games), in which case they can be adjusted to have similar statistics to the true reward function. We will demonstrate how to do this in the following section.

4.3 Generating indicator functions

Thus far we have described our algorithm as sampling sets of states $(S_i)_{i=1}^m$ which are then converted into a reward function by means of an indicator. In deep reinforcement learning, this is inconvenient for two reasons: first, because it is not clear from what distribution of states should be sampled (how should one generate arbitrary video-game states?); second, because testing for set inclusion may also be brittle, effectively reducing to repeated equality tests. Instead, we opt here for an *implied set*, defined directly by its indicator function.

Let $\mathcal{F}$ be a family of functions mapping $\mathcal{X}$ to $\{0, 1\}$. Then, for any function $f \in \mathcal{F}$, its implied set is

$$S_f = \{x \in \mathcal{X} : f(x) = 1\}.$$

Of course, this is equivalent to

$$f(x) = \mathbf{1}\{x \in S_f\}.$$

Sampling functions from $\mathcal{F}$ according to some distribution ξ_f and using them in lieu of the indicator is therefore equivalent to sampling sets of states for some distribution ξ implied by ξ_f. The advantage is that testing for inclusion in S_f only requires the evaluation of f at x, which for carefully-chosen functions can be done at little computational cost.

The simplest scenario occurs when the family $\mathcal{F}$ is parametrized by some weight vector θ, so that the random function f_θ corresponds to a random set of states. In this paper we consider two such families of functions: universal hash functions and random network indicators. Both families are *tunable*, in the sense that they are parametrized so that the implied sets S_f each cover a desired fraction of the overall state space. In probabilistic terms, tunable means that we can with minimal or no computation find parameters such that for any given state x,

$$\mathbb{P}\{x \in S_f\} = p.$$

Here, the probabilistic statement is with respect to the draw of f from $\mathcal{F}$. For universal hash functions, the tuning is immediate from the algorithm, and so we describe it first.

A *Carter-Wegman* family of hash functions $\mathcal{F}_{\text{CW}}$ (Carter & Wegman, 1979) consists of functions mapping each integer $x \in \mathbb{N}$ to the set $\{0, \ldots, k-1\}$, with the property that

$$\mathbb{P}\{h(x) = i\} = \tfrac{1}{k} \text{ for } i = 0, \ldots, k-1,$$

where the probabilistic statement is over the random draw of h from $\mathcal{F}_{\text{CW}}$. One may think of a CW family as deterministically assigning labels to integers x (in the sense that f is deterministic), but randomly (in the sense that f is random). See Appendix B.1 for full implementation details.

We construct our tunable indicator function as

$$f(x) = \mathbf{1}\{h(x) = 0\}.$$

By construction, choosing $k = \frac{1}{p}$ yields the desired tuning (up to integer rounding). In our setting, x is a high-dimensional observation (for example, an image) rather than an integer; yet we will

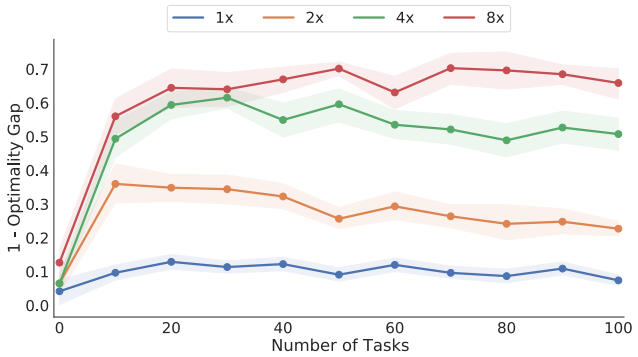

Figure 2: Performance of agents performing linear value approximation on top of a learned PVN. Agents are trained for 4 million environment frames; the legend indicates the network size as a capacity multiplier. Scaling network capacity with the number of auxiliary tasks leads to better performance. Larger networks support a larger number of auxiliary tasks.

see that, perhaps surprisingly, encoding each image as a unique integer is sufficient to produce better-than-random state representations.

One drawback of using universal hash functions to define sets of interest is that they may assign different values to perceptually near-identical states (a single pixel difference suffices). Following common usage (Burda et al., 2019; Dabney et al., 2021), we may use randomly initialized neural networks to map similar states to similar values. Specifically, let us view a randomly initialized, single-output DQN network as a function $g : \mathcal{X} \to \mathbb{R}$. We further decompose this function into a map $g_1 : \mathcal{X} \to \mathbb{R}^l$ and a linear map from $\mathbb{R}^l \to \mathbb{R}$:

$$g(x) = g_1(x)^\top \omega + b,$$

where ω is a parameter vector and $b \in \mathbb{R}$ is a bias term. With this in mind, we may simply construct the indicator function[2]

$$f(x) = \mathbf{1}\{g(x) \geqslant 0\}.$$

The result, however, is not yet tunable: it is hard to choose the right distribution of network weights so that a desired fraction of states satisfy $f(x) = 1$. However, for any $p \in [0, 1]$ and any non-zero fixed ω, g_1, and distribution of states μ, there exists a bias term b' such that

$$\mathbb{P}_{x\sim\mu}\{g_1(x)^\top \omega + b' \geqslant 0\} = p\,.$$

Such a bias term can accurately be determined from a small number of online interactions using the method of quantile regression (Koencker, 2005); the exact update rule is given in Appendix B.2. With this method, we obtain network-derived indicator functions that are tunable and are likely to assign similar values to perceptually similar states. We refer to this class of indicator functions as *random network indicators* (RNIs), which we empirically evaluate in the following section.

5 Empirical Analysis

To disentangle the contributions of the primary and auxiliary tasks on the expressiveness of the learned features, we split our learning procedure in two parts: a representation pre-training phase, and an online RL phase. During the representation pre-training phase, we use transition data from offline Atari datasets in RL Unplugged (Agarwal et al., 2020; Gulcehre et al., 2020) and the procedure described in Section 4 to train an encoder which acts as a feature extractor (see Appendix B for complete implementation details). Note that while this dataset contains environment rewards, none of the methods make use of the environment rewards unless explicitly stated. Following the pre-training phase, we fix the weights of the learned encoder and train an RL agent online directly from this "frozen" representation. Notably, we train for *only 3.75 million* agent steps, compared to the 50 million agent steps (200M Atari 2600 frames) that is standard in most Atari setups. Our agents are implemented using the Acme library (Hoffman et al., 2020). Our hyperparameter choices for both phases of training can be found in Appendix B.4.

[2]This corresponds to using a SIGN nonlinearity at the end of the network.

5.1 SCALING CAPACITY WITH AUXILIARY TASKS

Prior work indicates that the optimal number of auxiliary tasks for representation learning is unexpectedly small, and that scaling up the number of auxiliary tasks can hurt performance (Lyle et al., 2021, Fig. 5). We expect that the representational capacity of the neural network has a strong effect on the number of auxiliary tasks we are able to learn with. To study this effect, we use the Impala-CNN network (Espeholt et al., 2018) and vary its effective width; that is, we multiply the number of convolutional filters and the number of features in the penultimate layer. We select a width multiplier in the set $\{1, 2, 4, 8\}$ and sweep the number of tasks from $\{0, 10, \ldots, 100\}$. For this experiment, we use 5 games (ASTERIX, BEAM RIDER, PONG, QBERT, and SPACE INVADERS) with 3 seeds for pre-training, resulting in 15 encoders for each combination of width multiplier and number of auxiliary tasks. During the online phase, we train with 3 seeds per encoder, resulting in a total of 45 runs per sweep configuration. We evaluate for 100 episodes after 1M agent steps.

We summarize our results using Rliable (Agarwal et al., 2021b). Figure 2 depicts the optimality gap (distance from human-level performance). We find that increasing the representational capacity of the network increases performance, even for a very small number of tasks. This is perhaps surprising, since it indicates that we only need a handful of tasks to train large-scale representations, corroborating results by Lyle et al. (2021). Though a small part of this performance gain might be obtained just by virtue of having more output features (following the lottery ticket hypothesis (Frankle & Carbin, 2019)), we can see that there is a marked improvement when we increase the number of auxiliary tasks from 0 for all network sizes.

We further find that as network capacity is increased, the algorithm can use more auxiliary tasks to improve its representation. For example, while the $2\times$ network achieves maximal performance with 10 tasks, the $8\times$ network performs best in the range of $[50, 100]$ tasks. This gives evidence for the scalability of PVN as an approach for learning rich state representations.

5.2 EVALUATING THE LEARNED REPRESENTATION

Next, we use the insights gained from our scaling experiment to evaluate a model with a large number of auxiliary tasks on a broader suite of Atari games. We use the $8\times$ network that achieved the best performance in our scaling experiment, and fix the number of auxiliary tasks to 100, which empirically performed well. We use the same training setup described in the previous section, though we use all the 46 games available in RL Unplugged. We use one seed during offline pre-training due to the high computational requirements for running on many games, but keep 3 seeds per encoder in the online stage. We train for 3.75M agent steps, and then evaluate for 100 episodes. We compare against the following pre-training baselines:

Random Initialization: Randomly initialized features using the same network architecture. This simple baseline should confirm that the efficacy of our representations come from our pre-training procedure, and not merely because we use a large encoder network.

Random Cumulants (RCs): Random reward functions introduced by Dabney et al. (2021), and later expanded upon by Lyle et al. (2021). This method is similar to ours, but uses a random reward $r_i(x, x') = s \cdot (f(x') - f(x))$ instead of the random indicator function, and replaces the average over next-state actions by a maximization (off-policy learning of the optimal policy for each cumulant). Here, f is also given by a random network.

For each of these methods, we freeze the encoder after the pre-training stage and use the previously-described online training scheme. Figure 3 illustrates that PVN outperforms random initialization and RCs in all aggregate metrics. We also note that PVN using linear function approximation (3.75M agent interactions) is competitive with DQN (50M agent interactions) in many games, as illustrated in the per-game results found in Appendix D.

To gain some insight into why PVN performs well compared to the baseline methods, we visualize the learned representations using multidimensional scaling (MDS) plots in Figure 1. These plots hint at the structure of the learned representations, which clearly differ by method. Notably, the representations learned by PVN (RNI) place temporally-successive states close together, and appears to capture information about the dynamics of the environment without needing access to the environment reward. More plots for different games may be found in Appendix C.

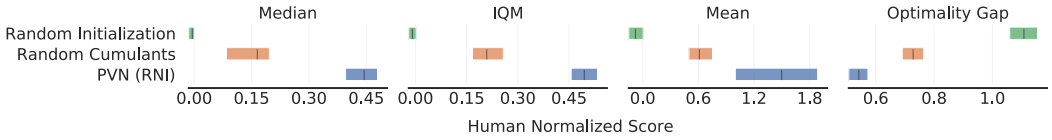

Figure 3: Performance of PVN RNI vs other methods described in Section 5.2, Environment Reward, and Random Cumulants, computed using 125 seeds, aggregated across 46 Atari games.

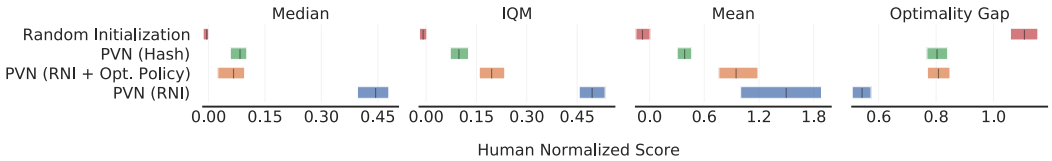

Figure 4: **PVN (Hash)**: PVN with hash indicator functions. **PVN (RNI)**: PVN with random network indicators. A randomly initialized network with $(8\times)$ capacity is plotted for comparison.

5.3 Ablations

We perform ablative experiments to verify the importance of the different PVN components. First, we validate our choice of indicator function by replacing RNIs with the hash indicator functions described in Section 4. We compare their performance in Figure 4, which shows that hash indicator functions perform poorly compared to RNIs; this indicates that the choice of indicator function is an important design decision. We expect that the inductive biases in random convolutional networks allow RNIs to include a notion of state similarity in the tasks they induce.

Next, we hypothesize that using the random policy as the target policy is a key contributor to PVN's performance. To verify this hypothesis, we ablate the TD-target of our learning update to maximize over the next-state action-values, as per the Bellman optimality equation. When the mean function is replaced with the max function in the TD backup, PVN attempts to learn the optimal value function for each indicator function, rather than the value function of the random policy. The result of this experiment can be seen in Figure 4. Using the mean formulation has a much higher median human normalized score than the max formulation. This is likely due to instability that arises from max bias and state coverage due to the off-policy learning required for the optimal value function. Learning the value function of the random policy also requires off-policy learning; however, we predict that it doesn't have such a large effect, as we previously described in Section 4.2.

6 Discussion

While our experiments have shed some light on the scalability of auxiliary tasks, there are a number of remaining open questions that represent exciting opportunities for further exploration. For example, having shown that random network indicators vastly outperform hash function indicators, we expect that there might be other classes of indicator functions that have strong inductive biases that could lead to better representations. There is also the question of how far we can push the scaling results that we found in our experiments. We scaled PVN to networks that fit on a single (albeit large) accelerator, but it is likely that PVN can scale to larger architectures and more tasks. An exciting future direction is to use insights from the literature on scaling models effectively (Tan & Le, 2019) to further scale the auxiliary tasks we introduced here. Orthogonally, it may still be possible to train with more tasks without increasing the capacity of the network. It is surprising that with even a relatively large network with tens of millions of parameters, such as Impala $(8\times)$, the network only supports a handful of tasks. It is not clear why training with more tasks leads to worse performance, especially for smaller Impala architectures. A deeper analysis into the learning dynamics of TD-based auxiliary tasks might yield an avenue to scaling the number of auxiliary tasks without needing such large networks. Finally, in line with Agarwal et al. (2022), we will open-source our pre-trained representations, which we hope would enable researchers to tackle credit assignment on ALE, without the excessive computational cost of re-learning such representations.

REFERENCES

Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2020.

Rishabh Agarwal, Marlos C Machado, Pablo Samuel Castro, and Marc G Bellemare. Contrastive behavioral similarity embeddings for generalization in reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2021a.

Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Neural Information Processing Systems (NeurIPS)*, 2021b.

Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron Courville, and Marc G Bellemare. Beyond tabula rasa: Reincarnating reinforcement learning. *Neural Information Processing Systems (NeurIPS)*, 2022.

Ankesh Anand, Evan Racah, Sherjil Ozair, Yoshua Bengio, Marc-Alexandre Côté, and R Devon Hjelm. Unsupervised state representation learning in atari. In *Neural Information Processing Systems (NeurIPS)*, 2019.

Bahram Behzadian and Marek Petrik. Feature selection by singular value decomposition for reinforcement learning. In *Proceedings of the ICML prediction and generative modeling workshop*, 2018.

Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The Arcade Learning Environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research (JMLR)*, 47:253–279, June 2013.

Marc G. Bellemare, Will Dabney, Robert Dadashi, Adrien Ali Taiga, Pablo Samuel Castro, Nicolas Le Roux, Dale Schuurmans, Tor Lattimore, and Clare Lyle. A geometric perspective on optimal representations for reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2019.

Marc G Bellemare, Salvatore Candido, Pablo Samuel Castro, Jun Gong, Marlos C Machado, Subhodeep Moitra, Sameera S Ponda, and Ziyu Wang. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.

Léonard Blier, Corentin Tallec, and Yann Ollivier. Learning successor states and goal-dependent values: A mathematical viewpoint. *arXiv preprint arXiv:2101.07123*, 2021.

Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.

J. Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 1979.

Rich Caruana. Multitask learning. *Machine Learning*, 28(1), 1997.

Pablo Samuel Castro. Scalable methods for computing state similarity in deterministic markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.

Pablo Samuel Castro, Tyler Kastner, Prakash Panangaden, and Mark Rowland. Mico: Improved representations via sampling-based state similarity for markov decision processes. In *Neural Information Processing Systems (NeurIPS)*, 2021.

Wesley Chung, Somjit Nath, Ajin George Joseph, and Martha White. Two-timescale networks for nonlinear value function approximation. In *International Conference on Learning Representations (ICLR)*, 2019.

Will Dabney, André Barreto, Mark Rowland, Robert Dadashi, John Quan, Marc G. Bellemare, and David Silver. The value-improvement path: Towards better representations for reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2021.

Peter Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 1993.

Carlos Diuk, Andre Cohen, and Michael L Littman. An object-oriented representation for efficient reinforcement learning. In *International Conference on Machine learning (ICML)*, 2008.

Akram Erraqabi, Marlos C. Machado, Mingde Zhao, Sainbayar Sukhbaatar, Alessandro Lazaric, Ludovic Denoyer, and Yoshua Bengio. Temporal abstractions-augmented temporally contrastive learning: An alternative to the laplacian in RL. *arXiv*, 2203.11369, 2022.

Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning (ICML)*, 2018.

William Fedus, Carles Gelada, Yoshua Bengio, Marc G. Bellemare, and Hugo Larochelle. Hyperbolic discounting and learning over multiple horizons. *arXiv*, 1902.06865, 2019.

Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations (ICLR)*, 2019.

Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G. Bellemare. DeepMDP: Learning continuous latent space models for representation learning. In *International Conference on Machine Learning (ICML)*, 2019.

Dibya Ghosh and Marc G Bellemare. Representations for stable off-policy reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2020.

Caglar Gulcehre, Ziyu Wang, Alexander Novikov, Thomas Paine, Sergio Gómez, Konrad Zolna, Rishabh Agarwal, Josh S Merel, Daniel J Mankowitz, Cosmin Paduraru, et al. Rl unplugged: A suite of benchmarks for offline reinforcement learning. *Neural Information Processing Systems (NeurIPS)*, 2020.

Matthew W. Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Maron, Nikola Momchev, Danila Sinopalnikov, Piotr Stańczyk, Sabela Ramos, Anton Raichuk, Damien Vincent, Léonard Hussenot, Robert Dadashi, Gabriel Dulac-Arnold, Manu Orsini, Alexis Jacq, Johan Ferret, Nino Vieillard, Seyed Kamyar Seyed Ghasemipour, Sertan Girgin, Olivier Pietquin, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, Sarah Henderson, Abe Friesen, Ruba Haroun, Alex Novikov, Sergio Gómez Colmenarejo, Serkan Cabi, Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Andrew Cowie, Ziyu Wang, Bilal Piot, and Nando de Freitas. Acme: A research framework for distributed reinforcement learning. *arXiv preprint arXiv:2006.00979*, 2020. URL https://arxiv.org/abs/2006.00979.

Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z. Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations (ICLR)*, 2017.

Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Miłos, Błażej Osiński, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, Afroz Mohiuddin, Ryan Sepassi, George Tucker, and Henryk Michalewski. Model based reinforcement learning for atari. In *International Conference on Learning Representations (ICLR)*, 2020.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015.

Roger Koencker. *Quantile Regression*. Cambridge University Press, 2005.

George D. Konidaris, Sarah Osentoski, and Philip S. Thomas. Value function approximation in reinforcement learning using the fourier basis. In *Proceedings of the 25th Conference on Artificial Intelligence*, 2011.

Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *Neural Information Processing Systems (NeurIPS)*, 2020.

Charline Le Lan, Stephen Tu, Adam Oberman, Rishabh Agarwal, and Marc G Bellemare. On the generalization of representations in reinforcement learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.

Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research (JMLR)*, 17(1), 2016.

Clare Lyle, Mark Rowland, Georg Ostrovski, and Will Dabney. On the effect of auxiliary tasks on representation dynamics. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

Clare Lyle, Mark Rowland, and Will Dabney. Understanding and preventing capacity loss in reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2022.

Marlos C. Machado, Clemens Rosenbaum, Xiaoxiao Guo, Miao Liu, Gerald Tesauro, and Murray Campbell. Eigenoption discovery through the deep successor representation. In *International Conference on Learning Representations (ICLR)*, 2018.

Sridhar Mahadevan and Mauro Maggioni. Proto-value functions: A laplacian framework for learning representation and control in markov decision processes. *Journal of Machine Learning Research (JMLR)*, 2007.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015.

Ronald Parr, Lihong Li, Gavin Taylor, Christopher Painter-Wakefield, and Michael L Littman. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2008.

Marek Petrik. An analysis of Laplacian methods for value function approximation in MDPs. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2007.

David Pfau, Stig Petersen, Ashish Agarwal, David G. T. Barrett, and Kimberly L. Stachenfeld. Spectral inference networks: Unifying deep and spectral learning. In *International Conference on Learning Representations (ICLR)*, 2019.

Bohdana Ratitch and Doina Precup. Sparse distributed memories for on-line value-based reinforcement learning. In *Proceedings of the 15th European Conference on Machine Learning*, 2004.

Gavin A. Rummery and Mahesan Niranjan. On-line Q-learning using connectionist systems. Technical report, Cambridge University Engineering Department, 1994.

Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International Conference on Machine Learning (ICML)*, 2015.

Max Schwarzer, Ankesh Anand, Rishab Goel, R Devon Hjelm, Aaron Courville, and Philip Bachman. Data-efficient reinforcement learning with self-predictive representations. In *International Conference on Learning Representations (ICLR)*, 2021a.

Max Schwarzer, Nitarshan Rajkumar, Michael Noukhovitch, Ankesh Anand, Laurent Charlin, R Devon Hjelm, Philip Bachman, and Aaron C Courville. Pretraining representations for data-efficient reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021b.

Kimberly L. Stachenfeld, Matthew Botvinick, and Samuel J. Gershman. Design principles of the hippocampal cognitive map. In *Neural Information Processing Systems (NeurIPS)*, 2014.

Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2021.

Richard S. Sutton. Generalization in reinforcement learning: Successful examples using sparse coarse coding. In *Neural Information Processing Systems (NeurIPS)*, 1996.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Richard S. Sutton and Steven D. Whitehead. Online learning with random representations. In *International Conference on Machine Learning (ICML)*, 1993.

Richard S. Sutton, Joseph Modayil, Michael Delp, Thomas Degris, Patrick M. Pilarski, Adam White, and Doina Precup. Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2011.

Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, 2019.

Ahmed Touati. *Large State Spaces and Self-supervision in Reinforcement Learning.* PhD thesis, Université de Montréal, 2021.

Ahmed Touati and Yann Ollivier. Learning one representation to optimize all rewards. In *Neural Information Processing Systems (NeurIPS)*, 2021.

Yifan Wu, George Tucker, and Ofir Nachum. The laplacian in RL: learning representations with efficient approximations. In *International Conference on Learning Representations (ICLR)*, 2019.

Amy Zhang, Rowan Thomas McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. In *International Conference on Learning Representations (ICLR)*, 2021.