On The Effect of Auxiliary Tasks on Representation Dynamics

Clare Lyle* University of Oxford

Mark Rowland* DeepMind

Georg Ostrovski DeepMind

Will Dabney DeepMind

Abstract

While auxiliary tasks play a key role in shaping the representations learnt by reinforcement learning agents, much is still unknown about the mechanisms through which this is achieved. This work develops our understanding of the relationship between auxiliary tasks, environment structure, and representations by analysing the dynamics of temporal difference algorithms. Through this approach, we establish a connection between the spectral decomposition of the transition operator and the representations induced by a variety of auxiliary tasks. We then leverage insights from these theoretical results to inform the selection of auxiliary tasks for deep reinforcement learning agents in sparse-reward environments.

1 Introduction

Auxiliary tasks have provided robust benefits to deep reinforcement learning agents (Jaderberg et al., 2017a; Mirowski et al., 2017; Lin et al., 2019). A commonly-held belief is that these benefits are mediated through improved representation learning. This hypothesis naturally raises a number of questions that, broadly speaking, remain open. What makes a good auxiliary task? Can we predict how an auxiliary task will affect an agent's representation? When should one auxiliary task be used instead of another? More generally, how should this hypothesis about the mechanism of auxiliary tasks itself be tested? The complex interacting components of large-scale deep reinforcement learning agents make it difficult to extract general insights. In this work we aim to shed light on the answers to these questions by distilling the benefits of auxiliary tasks down to the effects on the dynamics of the representations of reinforcement learning agents.

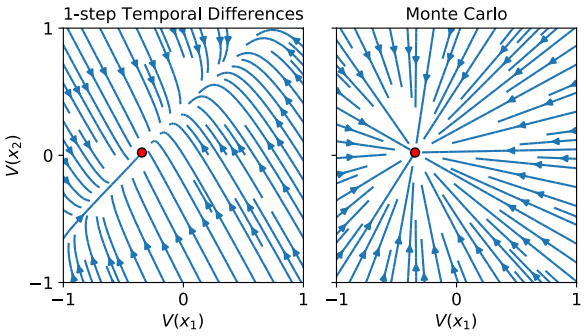

Figure 1: An example of qualitatively different value function dynamics for a two-state MDP for 1-step temporal difference learning and Monte Carlo learning, with fixed point V^π in red.

We begin by considering a *learning dynamics* framework for studying the effects of auxiliary tasks; see Figure 1 for a toy illustration, with full details given in Section 3. The central idea behind this framework is that it is not just *what* an agent learns that dictates how its representation is shaped, but *how* it learns.

This framework provides a *model* for representation learning in RL. Under this model, even in the case of value-based algorithms, it is shown that agents automatically incorporate the transition structure of the environment into their representations. We characterize the dynamics induced by a number of auxiliary tasks, with particular focus on ensemble predictions and random cumulant functions, and prove convergence of the induced representations to subspaces defined by certain decompositions of the environment's transition operator. We then consider the effectiveness of auxiliary tasks in sparse-reward environments, and via the use of the learning dynamics framework, construct a hypothesis as to which auxiliary tasks should be particularly well suited to such environments; we then test these developments in the Arcade Learning Environment (Bellemare et al., 2013), demonstrating strong performance with random cumulant auxiliary tasks.

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. *Equal contribution.

2 Background

We consider a Markov decision process $(\mathcal{X}, \mathcal{A}, P, R)$ comprising a finite state space $\mathcal{X}$, finite action space $\mathcal{A}$, transition kernel $P : \mathcal{X} \times \mathcal{A} \to \mathscr{P}(\mathcal{X})$, and reward distribution function $\mathcal{R} : \mathcal{X} \times \mathcal{A} \to \mathscr{P}(\mathbb{R})$.

2.1 Value-based reinforcement learning

Two key tasks in reinforcement learning are *policy evaluation* and *policy optimisation*. The former is specified by a policy $\pi : \mathcal{X} \to \mathscr{P}(\mathcal{A})$. An agent using the policy π to interact with the environment generates a sequence of states, actions and rewards $(X_t, A_t, R_t)_{t\geq 0}$. The performance of the agent is summarised by the return $\sum_{t\geq 0} \gamma^t R_t$, for a discount factor $\gamma \in [0, 1)$. The goal of policy evaluation is to (approximately) compute the value function

$$V^\pi(x) = \mathbb{E}_\pi[\textstyle\sum_{t\geq 0} \gamma^t R_t \mid X_0 = x] \,,$$

for all $x \in \mathcal{X}$. Policy optimisation consists of finding a policy $\pi^* \in \mathscr{P}(\mathcal{A})^{\mathcal{X}}$ that maximises the expected return from all possible initial states. The value function associated with π^* is denoted V^*.

Crucial to the value-based approach to reinforcement learning are the *Bellman operators*. The one-step evaluation operator associated with a policy π is the function $T^\pi : \mathbb{R}^{\mathcal{X}} \to \mathbb{R}^{\mathcal{X}}$ defined by

$$(T^\pi V)(x, a) = \mathbb{E}_\pi[R_0 + \gamma V(X_1) \mid X_0 = x] \,.$$

Introducing the transition operator $P^\pi \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ defined by $P^\pi(x'|x) = \sum_{a\in\mathcal{A}} \pi(a|x) P(x'|x, a)$, and the expected reward vector $R^\pi \in \mathbb{R}^{\mathcal{X}}$ defined by $R^\pi(x) = \mathbb{E}_\pi[R_0|X_0 = x]$, this can be expressed even more succinctly in operator notation as

$$T^\pi V = R^\pi + \gamma P^\pi V \,.$$

The Bellman optimality operator is the function $T^* : \mathbb{R}^{\mathcal{X}\times\mathcal{A}} \to \mathbb{R}^{\mathcal{X}\times\mathcal{A}}$ defined by

$$(T^* V)(x) = \max_{a\in\mathcal{A}} \mathbb{E}[R_0 + \gamma V(X_1) | X_0 = x, A_0 = a] \,.$$

Repeated application of T^π (resp., T^*) to any initial value function converges to V^π (resp., V^*) (Bertsekas and Tsitsiklis, 1996). Popular algorithms such as Q-learning (Watkins and Dayan, 1992), which form the basis of many deep RL agents (Mnih et al., 2015), can be viewed as approximating the iterative application of T^* and related operators (Tsitsiklis, 1994; Jaakkola et al., 1994; Bertsekas and Tsitsiklis, 1996).

2.2 Features and representations

In many environments, it is impractical to store a value function as a table indexed by states, and further this does not permit generalisation in the course of learning. Instead, it is typical to parametrise $V \in \mathbb{R}^{\mathcal{X}}$ through a *feature map* $\phi : \mathcal{X} \to \mathbb{R}^K$ and *weight vector* $\mathbf{w} \in \mathbb{R}^K$, leading to a factorisation of the form

$$V(x) = \langle \phi(x), \mathbf{w} \rangle \,.$$

Such a parametrisation may be amenable to more efficient learning, for example if ϕ abstracts away unimportant information, allowing for generalisation between similar states. Even more concisely, writing $\Phi \in \mathbb{R}^{\mathcal{X}\times K}$ for the matrix with rows $\phi(x)$ yields

$$V = \Phi\mathbf{w} \,. \tag{1}$$

The quantity Φ is often referred to as the agent's *representation* of the environment (Boyan, 1999; Levine et al., 2017; Bertsekas, 2018; Chung et al., 2018; Bellemare et al., 2019; Dabney et al., 2020). In many small- and medium-scale applications, the representation is fixed ahead of time, and only $\mathbf{w}$ is updated during learning; this is the linear function approximation regime. Many common choices of features relate to various decompositions of operators associated with the transition operators P^π. In deep reinforcement learning, however, Φ and $\mathbf{w}$ are learnt simultaneously.

2.3 Representation learning and auxiliary tasks

To perfectly express a value function V^π in the form of Equation (1), the following condition is necessary:

$$\langle \Phi \rangle \supseteq \langle V^\pi \rangle \,, \tag{2}$$

where $\langle \Phi \rangle$ denotes the column span of Φ, and $\langle V^\pi \rangle$ denotes the one-dimensional subspace of $\mathbb{R}^{\mathcal{X}}$ spanned by V^π. However, this condition is not sufficient for efficient sample-based learning (Du et al., 2019). There are several reasons for this; for some intuition, consider that since value functions are typically learnt through bootstrapping algorithms, the agent is required to accurately express a *sequence* of value functions as its estimates are updated, and thus a good representation should also allow such intermediate value functions to be expressed in the course of learning (Dabney et al., 2020).

Despite the importance of the representation Φ, it remains unclear how exactly the notion of a good representation in this sense should be formalised. In spite of this, representation learning is a hugely important aspect of deep reinforcement learning. A consistent finding in empirical deep RL research is that requiring the agent to use its representation to predict other functions of state, referred to as *auxiliary tasks*, in addition to its primary task of learning an optimal policy,

can lead to considerable boosts in performance. Examples of commonly-used auxiliary tasks include predicting the expected return associated with other reward functions (Sutton et al., 2011), other discount factors (Fedus et al., 2019), and other policies (Dabney et al., 2020), as well as other properties of the return distribution (Bellemare et al., 2017) and other aspects of the environment observations (Jaderberg et al., 2017a), amongst others. We discuss prior work on auxiliary tasks in greater detail in Section 6. A popular hypothesis is that auxiliary tasks add further constraints to Expression (2), requiring the representation Φ to contain more functions of interest than just V^π in its column span (Bellemare et al., 2019; Dabney et al., 2020).

3 Learning dynamics

Our aim in the remainder of the paper is to develop an understanding of the ways in which auxiliary tasks shape representations in RL. Our central results establish connections between decompositions of transition operators, commonly used in static feature selection, and certain classes of auxiliary tasks used in deep reinforcement learning. To build up to these results, in this section we examine learning algorithms in the absence of auxiliary tasks, first considering tabular learning algorithms, and then moving to the case where representations and feature weights are learnt simultaneously.

3.1 Warm-up: Tabular value function dynamics

We consider the following one-step temporal difference (TD) continuous-time learning dynamics:

$$\partial_t V_t(x) = \mathbb{E}_\pi[R_0 + \gamma V_t(X_1) | X_0 = x] - V_t(x)\,,$$

for each $x \in \mathcal{X}$, which may also be written

$$\partial_t V_t(x) = R^\pi(x) + \gamma(P^\pi V_t)(x) - V_t(x)\,,$$

or in full matrix notation,

$$\partial_t V_t = -(I - \gamma P^\pi)V_t + R^\pi\,. \tag{3}$$

The differential equation in (3) is an affine autonomous system, and is straightforwardly solvable.

Lemma 3.1. If $(V_t)_{t\geq 0}$ satisfies Equation (3) with initial condition V_0 at time $t = 0$, then we have

$$V_t = \exp(-t(I - \gamma P^\pi))(V_0 - V^\pi) + V^\pi\,. \tag{4}$$

We recover as a straightforward corollary the well-known result that $V_t \to V^\pi$ as $t \to \infty$, since all eigenvalues of $(I - \gamma P^\pi)$ have strictly positive real part.

However, the solution in Equation (4) also describes the *trajectory* by which V_t reaches this limiting value. Figure 1 provides an illustration of this in a small MDP; the value functions accumulate along a particular affine subspace of $\mathbb{R}^\mathcal{X}$ prior to convergence.

This phenomenon can in fact be formalised. To do so, we need a notion of distance between subspaces of $\mathbb{R}^\mathcal{X}$. The following definition follows Ye and Lim (2016). Intuitively, it can be thought of as generalizing the notion of an angle between vectors to subspaces.

Definition 3.2. For two K-dimensional subspaces $Y_1, Y_2 \leq \mathbb{R}^\mathcal{X}$, the *principal angles* $\theta_1, \ldots, \theta_K \in [0, \pi/2]$ between the subspaces are defined by taking orthonormal matrices $\mathbf{Y}_1 \in \mathbb{R}^{\mathcal{X}\times K}$ and $\mathbf{Y}_2 \in \mathbb{R}^{\mathcal{X}\times K}$ the columns of which span Y_1 and Y_2 respectively, and defining $\theta_k = \cos^{-1}(\sigma_k(\mathbf{Y}_1^\top \mathbf{Y}_2))$, where $\sigma_k(\mathbf{A})$ is the k^{th} singular value of the matrix $\mathbf{A}$. One can check that this definition is independent of the matrices $\mathbf{Y}_1$ and $\mathbf{Y}_2$, depending only on the subspaces Y_1, Y_2 themselves. The *Grassmann distance* $d(Y_1, Y_2)$ between Y_1 and Y_2 is then defined as $\|\theta\|_2 = (\sum_{k=1}^K \theta_k^2)^{1/2}$.

With these definitions in hand, we now give a precise version of the statement alluded to in the discussion and figure above. We make some simplifying assumptions to avoid focusing on technicalities here, and give a discussion of the more general case in Appendix F.

Assumption 3.3. P^π is real-diagonalisable, with strictly decreasing eigenvalue sequence $1 = |\lambda_1| > |\lambda_2| > \cdots > |\lambda_{|\mathcal{X}|}|$, and corresponding right-eigenvectors $U_1, \ldots, U_{|\mathcal{X}|}$.

Proposition 3.4. Under Assumption 3.3, and $(V_t)_{t\geq 0}$ the solution to Equation (3), for almost every[1] initial condition V_0, we have

$$d(\langle V_t - V^\pi \rangle, \langle U_1 \rangle) \to 0\,.$$

The behaviour described by Proposition 3.4 is exhibited in Figure 1, as the value function V_t approaches the affine subspace in direction $(1, 1)$ prior to converging to V^π. A more general version of this statement can also be given with an ensemble of K value functions, which indicates that yet more information about the environment is contained in the learnt collection. The proofs of these results relate to the classical *power method* in linear algebra.

Proposition 3.5. Under Assumption 3.3, and $(V_t^{(k)})_{t\geq 0}$ the solution to Equation (3) for each $k = 1, \ldots, K$, for almost every initial condition $(V_0^{(k)})_{k=1}^K$, we have

$$d(\langle V_t^{(k)} - V^\pi \mid k \in [K] \rangle, \langle U_{1:K} \rangle) \to 0\,.$$

[1] In the measure-theoretic sense that the set of excluded initial conditions V_0 has Lebesgue measure 0.

Key insight.
Even in an environment with no reward signal at all (in which case $V^\pi = 0$), an agent performing TD learning still picks up information about the transition structure of the environment within its value function.

Due to the importance of the vectors $U_{1:K}$ in this analysis, we introduce the term *eigen-basis functions* (EBFs) to describe them.

We observe that a similar analysis, indicating similar behaviour, is possible for related learning algorithms such as n-step temporal difference learning and TD(λ); see Appendix E for further details. In contrast, Monte Carlo learning dynamics correspond to the differential equation

$$\partial_t V_t = (I - \gamma P^\pi)^{-1} R^\pi - V_t \,,$$

which has the solution

$$V_t = e^{-t}(V_0 - (I - \gamma P^\pi)^{-1} R^\pi) + (I - \gamma P^\pi)^{-1} R^\pi \,.$$

The trajectory associated with this solution simply linearly interpolates between V_0 and V^π, as illustrated in Figure 1, and does not pick up any additional information about the environment in the value function as learning proceeds. See Appendix E for further details. This example serves to illustrate that it is not just *what* an agent learns (V^π), but *how* the agent learns that plays a key, measurable role in what environment information is picked up in its value function. We now apply this perspective to representation learning.

3.2 Representation dynamics

Recall the parametrisation of $V \in \mathbb{R}^{\mathcal{X}}$ from Section 2.2, taking the form

$$V = \Phi \mathbf{w} \,,$$

for $\Phi \in \mathbb{R}^{\mathcal{X} \times K}$, $\mathbf{w} \in \mathbb{R}^K$. Central to deep reinforcement learning agents is the idea that Φ and $\mathbf{w}$ are simultaneously learnt from a single RL loss. As in the value function case, we will focus on the dynamics with single-step temporal difference learning; remarks on other learning algorithms are given in Appendix E. The dynamics associated with single-step TD learning are given by

$$\partial_t \Phi_t = -\alpha \frac{1}{2} \nabla_{\Phi_t} \| R^\pi + \mathtt{SG}[\gamma P^\pi \Phi_t \mathbf{w}_t] - \Phi_t \mathbf{w}_t \|_2^2 \,, \quad (5)$$

$$\partial_t \mathbf{w}_t = -\beta \frac{1}{2} \nabla_{\mathbf{w}_t} \| R^\pi + \mathtt{SG}[\gamma P^\pi \Phi_t \mathbf{w}_t] - \Phi_t \mathbf{w}_t \|_2^2 \,, \quad (6)$$

where $\alpha, \beta \in [0, \infty)$ are learning rates, implying that features and weights may be learnt at different rates. Further, $\mathtt{SG}[]$ denotes a *stop-gradient*, indicating that we treat the instances of Φ_t and $\mathbf{w}_t$ within the expression as constants with regard to computing derivatives; this reflects the fact that temporal difference learning is a *semi-gradient* method.

The use of a single loss to learn both the representation and weights corresponds to the approach taken in deep RL, and we will use these dynamics as an idealized model of the deep RL setting. While this model ignores some practicalities of deep RL (such as visitation distributions, implicit bias from the function approximation architecture, and stochasticity introduced by mini-batch training), it allows us to obtain valuable insights into representation dynamics which, as we shall see in Section 5, accurately predict the behaviour of deep RL agents.

Lemma 3.6. Let Φ_t and $\mathbf{w}_t$ parameterize a value function approximator as defined above. Then

$$\partial_t \Phi_t = \alpha (R^\pi + \gamma P^\pi \Phi_t \mathbf{w}_t - \Phi_t \mathbf{w}_t) \mathbf{w}_t^\top \,, \quad (7)$$

$$\partial_t \mathbf{w}_t = \beta \Phi_t^\top (R^\pi + \gamma P^\pi \Phi_t \mathbf{w}_t - \Phi_t \mathbf{w}_t) \,. \quad (8)$$

This joint flow on Φ_t and $\mathbf{w}_t$ leads to much richer behaviour than the flow considered on value functions in the previous section. Without further assumptions, the evolution of the representation Φ_t may be complex, and will not necessarily incorporate environment information as described for the case of value functions in Proposition 3.5. In particular, in sparse reward environments, the agent may learn to predict a near-zero value function by setting the weights $\mathbf{w}_t$ close to zero, which would effectively prevent any further updating of the features Φ_t, ruling out the possibility of a result analogous to Proposition 3.5.

4 Auxiliary task dynamics

Having studied the temporal difference learning dynamics in Equation (7) & (8), we now examine how auxiliary value-prediction tasks influence the behaviour of the agent's representation during the learning process.

As described above, developing a granular description of the joint learning dynamics of the representation and weights of the learner is a complex task, and so we focus on the limiting case in which the number of auxiliary tasks is large relative to the dimensionality of the representation. We conclude that under certain conditions, representations learnt in the many-task limit bear a close connection with the *eigen-basis functions* described in Section 3.1, and also *resolvent singular basis functions*, a new decomposition introduced in Section 4.2. The reader may find it useful to

refer to Appendix G for a more detailed discussion of these decompositions.

4.1 Ensemble value prediction

We begin by considering the auxiliary task of *ensemble value prediction* (Osband et al., 2016; Anschel et al., 2017; Agarwal et al., 2020). Rather than making a single prediction of the value function Q^π, the learner makes $M \in \mathbb{N}$ separate predictions as linear functions of a common representation Φ^M, using M independently initialized weights matrices $\mathbf{w}^m \in \mathbb{R}^K$ $(m = 1, \ldots, M)$. We note that while at initialization $\Phi_0^M \in \mathbb{R}^{\mathcal{X} \times d}$ is independent of M, its dynamics do depend on M through the contribution of the weights. Simultaneous temporal difference learning on all predictions leads to the following dynamics:

$$\partial_t \Phi_t^M = \alpha \sum_{m=1}^{M} (R^\pi + \gamma P^\pi \Phi_t^M \mathbf{w}_t^m - \Phi_t^M \mathbf{w}_t^m)(\mathbf{w}_t^m)^\top , \tag{9}$$

$$\partial_t \mathbf{w}_t^m = \beta (\Phi_t^M)^\top (R^\pi + \gamma P^\pi \Phi_t^M \mathbf{w}_t^m - \Phi_t \mathbf{w}_t^m) . \tag{10}$$

The following result characterises the representation learnt by the agent in the many-tasks limit, again establishing a connection to EBFs; we follow the approach described by Arora et al. (2019b) in fixing the linear weights associated with the value function; this dramatically simplifies our analysis, while still describing practical settings in which the features and weights are trained separately as in Chung et al. (2018).

Theorem 4.1. For $M \in \mathbb{N}$, let $(\Phi_t^M)_{t \geq 0}$ be the solution to Equation (9), with each $\mathbf{w}_t^m$ for $m = 1, \ldots, M$ initialised independently from $N(0, \sigma_M^2)$, and fixed throughout training ($\beta = 0$). We consider two settings: first, where the learning rate α is scaled as $\frac{1}{M}$ and $\sigma_M^2 = 1$ for all M, and second where $\sigma_M^2 = \frac{1}{M}$ and the learning rate α is equal to 1. These two settings yield the following dynamics, respectively:

$$\lim_{M \to \infty} \partial_t \Phi_t^M \overset{P}{=} -(I - \gamma P^\pi)\Phi_t^M \quad , \text{and} \tag{11}$$

$$\lim_{M \to \infty} \partial_t \Phi_t^M \overset{D}{=} -(I - \gamma P^\pi)\Phi_t^M + R^\pi \epsilon^\top \ , \ \epsilon \sim \mathcal{N}(0, I) . \tag{12}$$

The corresponding limiting trajectories for a fixed initialisation $\Phi_0 \in \mathbb{R}^{\mathcal{X} \times K}$, are therefore given respectively by

$$\lim_{M \to \infty} \Phi_t^M \overset{P}{=} \exp(-t(I - \gamma P^\pi))\Phi_0 \quad , \text{and} \tag{13}$$

$$\lim_{M \to \infty} \Phi_t^M \overset{D}{=} \exp(-t(I - \gamma P^\pi))(\Phi_0 - (I - \gamma P^\pi)^{-1} R^\pi \varepsilon^\top) + (I - \gamma P^\pi)^{-1} R^\pi \varepsilon^\top \ , \ \epsilon \sim \mathcal{N}(0, I) . \tag{14}$$

In contrast to the case described in Section 3.2, this result indicates that the introduction of auxiliary tasks leads to useful environment information being incorporated into the representation. Indeed, the dynamics described above imply the following convergence result, analogous to Proposition 3.5.

Corollary 4.2. Under the feature flow (9) with $\mathbf{w}_t^m$ fixed at initialization for each $i = 1, \ldots, M$ and Assumption 3.3, for almost all initialisations Φ_0, we have when $R^\pi = 0$

$$d(\langle \Phi_t \rangle, \langle U_{1:K} \rangle) \to 0 , \quad \text{as } t \to \infty.$$

Key insight.
Under the conditions of Theorem 4.1 and Corollary 4.2, the ensemble auxiliary tasks cause the agent's representation Φ to align with EBFs.

We show in Appendix I that this behaviour is observed in practice when $M \gg K$ and the value of $\mathbf{w}_t^m$ is fixed at initialization. We additionally compare the representations learned when $\mathbf{w}_t^m$ is allowed to vary over training. Here we find empirically that allowing the weights to vary during training induces dynamics that differ from those predicted by Theorem 4.1 for the fixed-weights setting. To illustrate this, we follow the evolution of a single column of Φ_t, i.e. a single feature vector ϕ_t, trained with the ensemble prediction dynamics of Equations (9) & (10) on a simple four-rooms gridworld environment in Figure 2.

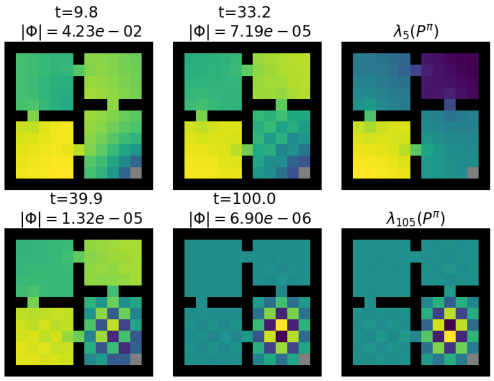

Figure 2: Visualization of a single column of Φ_t (i.e. feature vector) after application of the ODE in Equation (7) for $t \in [0, 100]$ in the four rooms environment, with $K = 10$ and $M = 20$. Early in its trajectory, ϕ_t exhibits similarity to smooth eigenfunctions (e.g. the eigenfunction corresponding to the 5^{th} greatest eigenvalue λ_5 which we plot in the top right) of P^π, but later converges to non-smooth eigenfunctions (e.g. the eigenfunction corresponding to eigenvalue λ_{105}, the most negative eigenvalue, plotted in the bottom right).

Auxiliary task	Dynamics ($r=0$)	Φ_∞ ($r=0$)	Φ_∞ ($r\neq 0$)	Limit of $\langle \Phi_t - \Phi_\infty \rangle$
Ensemble	$-(I-\gamma P^\pi)\Phi_t$	0	$(I-\gamma P^\pi)^{-1} r\epsilon^\top$	EBFs of P^π
Random cumulants	$-(I-\gamma P^\pi)\Phi_t + Z_\Sigma$	$(I-\gamma P^\pi)^{-1} Z_\Sigma$	$(I-\gamma P^\pi)^{-1} Z_\Sigma$	EBFs of P^π
Additional policies	$-(I-\gamma P^{\overline{\pi}})\Phi_t$	0	$(I-\gamma P^{\overline{\pi}})^{-1} R^{\overline{\pi}} \epsilon^\top$	EBFs of $P^{\overline{\pi}}$
Multiple γs	$-(I-\overline{\gamma} P^\pi)\Phi_t$	0	$(I-\overline{\gamma} P^\pi)^{-1} R^\pi \epsilon^\top$	EBFs of P^π

Table 1: Summary of dynamics and limiting solutions under some common auxiliary tasks in the limit of infinitely-many prediction outputs. For additional policies, $\overline{\pi}$ denotes the average of the finite set of policies $\pi_1, \ldots, \pi_L$ under consideration (L fixed and independent of M), and for multiple discount factors, $\overline{\gamma}$ denotes the average of the discount factors $\gamma_1, \ldots, \gamma_L$ under consideration.

We visualize ϕ_t along with two illustrative eigenfunctions of the transition matrix P^π, corresponding to one positive and one negative eigenvalue. We observe that while the feature ϕ_t quickly evolves to resemble the smooth eigenfunction corresponding to the positive eigenvalue for small values of t, it later converges to the non-smooth eigenfunction corresponding to the most negative eigenvalue of the transition matrix P^π. While we leave further analysis to future work, this example hints at an intriguing relationship between the EBFs and the joint representation dynamics.

4.2 Random cumulants

In the case of zero rewards, our previous results show that whilst from the perspective of subspaces the representation approaches the EBF subspace in Grassmann distance, in Euclidean distance the representation is approaching the zero matrix pointwise. This has important implications for the scenario of large-scale sparse-reward environments, in which the agent may not encounter rewards for long periods of time, and indicates that the agent's representation is at risk of collapsing in such cases.

Motivated by this analysis, we consider a means of alleviating this representation collapse, by learning value functions for *randomly generated cumulants* (Osband et al., 2018; Dabney et al., 2020). Mathematically, the agent again makes many predictions from a common representation, with each prediction indexed by $m = 1, \ldots, M$ attempting to learn the value function associated with a randomly drawn reward function $r^m \in \mathbb{R}^{\mathcal{X}}$ under the policy π. Thus, the agent's parameters are the representation Φ and a set of weights $\mathbf{w}^m$ for each prediction. The learning dynamics are then given by:

$$\partial_t \Phi_t^M = \alpha \sum_{m=1}^{M} (r^m + \gamma P^\pi \Phi_t^M \mathbf{w}_t^m - \Phi_t^M \mathbf{w}_t^m)(\mathbf{w}_t^m)^\top, \tag{15}$$

$$\partial_t \mathbf{w}_t^m = \beta (\Phi_t^M)^\top (r^m + \gamma P^\pi \Phi_t^M \mathbf{w}_t^m - \Phi_t \mathbf{w}_t^m). \tag{16}$$

The main result of this section is to show that, even in the absence of reward, the limiting distribution induced by random cumulant auxiliary tasks dynamics described in Equation (15) is not the zero subspace.

Theorem 4.3. For fixed $M \in \mathbb{N}$, let the random rewards $(r^m)_{m=1}^M$ and weights $(\mathbf{w}^m)_{m=1}^M$ be as defined above, let $\alpha = 1$, and consider the representation dynamics in Equation (15), with weights fixed throughout training ($\beta = 0$). Let Σ denote the covariance matrix of the random cumulant distribution. Then

$$\lim_{M\to\infty} \sum_{m=1}^{M} r^m (\mathbf{w}^m)^\top \overset{D}{=} Z_\Sigma \sim \mathcal{N}(0, \Sigma), \text{ and}$$

$$\lim_{M\to\infty} \Phi_t^M \overset{D}{=} \exp(-t(I-\gamma P^\pi))(\Phi_0 - (I-\gamma P^\pi)^{-1} Z_\Sigma) + (I-\gamma P^\pi)^{-1} Z_\Sigma .$$

As the columns of Z_Σ are mean-zero, uncorrelated, with covariance matrices Σ, the limiting distribution of each column of $\Phi_\infty = \lim_{t\to\infty} \lim_{M\to\infty} \Phi_t^M$ has covariance $\Psi\Sigma\Psi^\top$, where Ψ is the resolvent $(I-\gamma P^\pi)^{-1}$.

Corollary 4.4. Under the feature flow (15) with $\mathbf{w}_t^m$ fixed at initialization for each $i = 1, \ldots, M$ and Assumption 3.3, for almost all initialisations Φ_0, we have when $R^\pi = 0$

$$d(\lim_{M\to\infty} \langle \Phi_t^M - \Phi_\infty \rangle, \langle U_{1:K} \rangle) \to 0, \quad \text{as } t \to \infty.$$

Theorem 4.3 indicates that the left-singular vectors of $\Sigma^{1/2}\Psi$ (or equivalently, the right-eigenvectors of $\Psi\Sigma\Psi^\top$) are key to understanding the effects on random cumulants on representations; we introduce the term *resolvent singular basis functions* (RSBFs) to refer to these vectors in the canonical case $\Sigma = I$.

Key insight.
With random cumulant auxiliary tasks, under the assumptions of Theorem 4.3 and Corollary 4.4, the distribution of the limiting representation does not collapse, and is characterized by the RSBFs of P^π, while the trajectory it follows to reach this subspace is determined by the EBFs of P^π.

These decompositions of P^π bear deep connections to prior work on feature learning. EBFs correspond

to the eigendecomposition of the successor representation, which can be explicitly related to the proto-value functions described by Mahadevan (2009) when the transition matrix P^π corresponds to that of a random walk policy (Machado et al., 2018b). For symmetric P^π we obtain an additional correspondence between EBFs and RSBFs, though we note that when P^π is not symmetric the RSBFs may differ from both the EBFs and the singular value decomposition of the transition matrix P^π. We provide further discussion of RSBFs and comparisons against existing concepts in feature selection in Appendix G.

In Appendix H we show that RSBFs can be viewed as Bayes-optimal features in the sense that they minimize the expected value function approximation error given an isotropic Gaussian prior on an unknown reward function.

4.3 Analysis of additional auxiliary tasks

The infinite-task limit simplifies the analysis of a broad range of auxiliary tasks, and analogous results to Theorem 4.1 can be easily derived for families of auxiliary tasks which predict returns associated with additional policies and multiple discounts factors. We provide a summary of these results in Table 1, including their full statements and derivations in Appendix B.

We consider two additional classes of auxiliary task: predicting the values of multiple policies (Dabney et al., 2020), and predicting value under multiple discount factors (Fedus et al., 2019). Under the multiple policies auxiliary task, the agent's objective is to learn a set of value functions $V^1, \dots, V^M$ such that $V^i(x) = \mathbb{E}_{\pi_i}[R^{\pi_i}(x) + \gamma P^{\pi_i} V^i(x)]$. Under the multiple discount factors auxiliary task, the agent's objective is analogously to find $V^i(x) = \mathbb{E}_{\pi_i}[R^\pi(x) + \gamma_i P^\pi V^i(x)]$ for $\gamma_i \in \gamma_1, \dots, \gamma_k$. We consider an ensemble prediction variant of these objectives, where given a fixed set of k policies, we train an ensemble of M predictors $V^{1,1}, \dots, V^{m,1}, \dots, V^{1,k}, \dots, V^{m,k}$, where $m = \frac{M}{k}$ and the value function $V^{i,j}$ is trained on policy (respectively discount factor) π_j (respectively γ_j).

In both cases, under the conditions of the previous theorems the dynamics of the ensemble converge to the dynamics induced by the mean of the set of auxiliary tasks, implying the counter-intuitive result that training with multiple auxiliary tasks doesn't provide additional utility over the single task setting. This apparent shortcoming can be addressed by ensuring that the weights corresponding to each auxiliary task π_i or γ_i are initialized in *orthogonal subspaces*, so that the vector space V in which the representation evolves can be decomposed as $V = \oplus_{i \in [1,k]} V_i$. In this case, we obtain an analogous decomposition of the representation Φ and its corresponding dynamics, obtaining convergence to a direct sum of the limiting representation of each task. This suggests that the benefits of auxiliary tasks might be maximized by appropriate initialization schemes which encourage the representations learned for each task to be independent.

5 Experiments

In this section, we complement the theoretical results above with empirical investigations in both tabular and deep reinforcement learning settings.

5.1 Feature generalisation across the value-improvement path

Having established connections between the representations induced by auxiliary tasks and several decompositions of the environment transition operator, we now turn to the question of how useful these representations are to a reinforcement learning agent. In particular, we address how well representations learnt under one policy generalize under the policy improvement step to approximate future value functions, with particular attention paid to EBFs and RSBFs, the decompositions that feature in our earlier analysis.

To address this question empirically we run tabular policy iteration on a stochastic chain MDP, yielding a sequence of policies $(\pi_j)_{j=1}^J$ and associated value functions $(V_j)_{j=1}^J$. We then compute EBFs and RSBFs associated with P^π, and compute the acute angle between V_j and the subspace spanned by these features, for each $j \in [J]$; this is in fact equal to a generalisation of the Grassmann distance for subspaces of unequal dimension (Ye and Lim, 2016). We also compare against a baseline of isotropic randomly-generated features; full experimental details are provided in Appendix D.

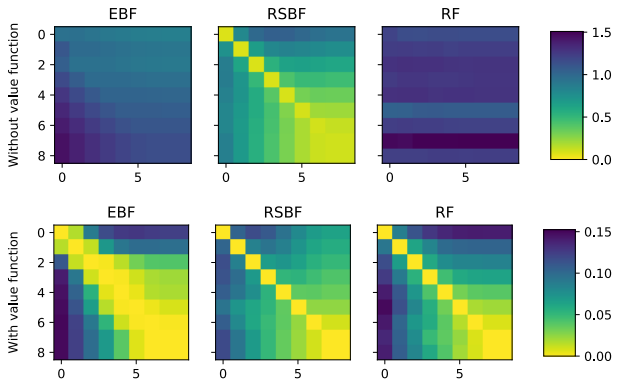

Figure 3: Transfer of EBFs, RSBFs, and RFs across the value-improvement path of a chain MDP, with and without the value function as an additional feature.

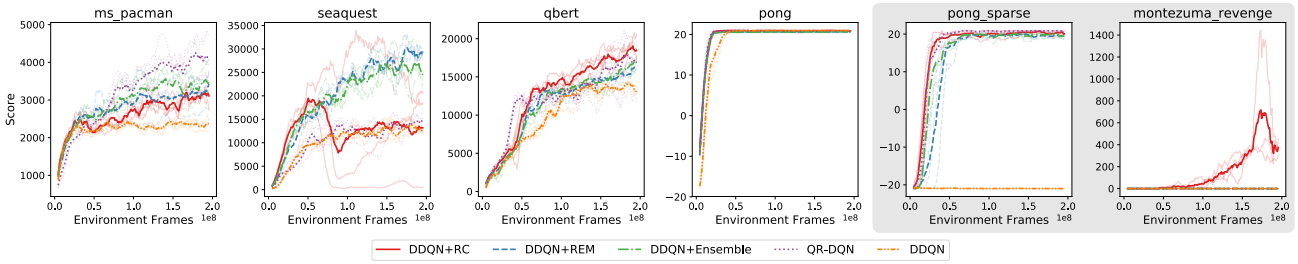

Figure 4: Learning curves for DDQN, DDQN+RC, DDQN+Ensemble, DDQN+REM, and QR-DQN agents on several dense reward ALE environments. Two games with sparse rewards are shown in the shaded box.

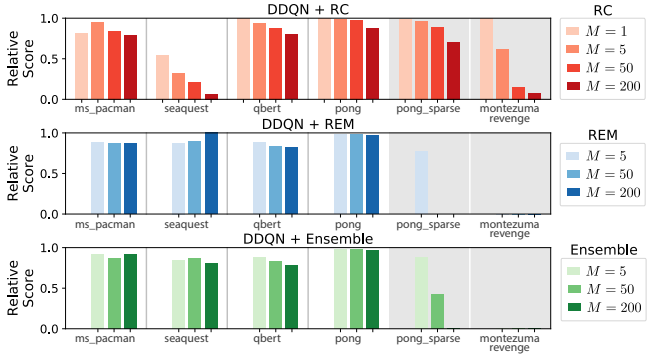

Figure 5: Sweep over number of auxiliary heads for RC, REM and Ensemble. Each bar corresponds to a single game, presented in the same order as Figure 4. Relative score is the per-game score divided by the maximum over all algorithms and hyperparameters.

Results are given in the top row of Figure 3 for the case of four features; each individual heatmap plots Grassmann distances, with rows indexing the policy that generated the features, and columns indexing the policy yielding the target value function. In general, the RSBFs provide better transfer across policies in the improvement path relative to random features and EBFs. For times $j, j' \in [J]$, we observe that the Grassmann distance between the RSBFs of P^{π_j} and the value function of j', $V^{\pi_{j'}}$, increases as $|j - j'|$ does.

We also evaluate transfer when the vector V^{π_j} is added to the set of features, in the bottom row of Figure 3. This contains the subspace to which the value functions described in Proposition 3.5 converge, as the limiting solutions can be described as being of the form $V^{\pi_j} + u$ for $u \in \langle U_{1:K} \rangle$. Surprisingly, we find that in this setting the EBFs for π_j outperform RSBFs specifically in predicting $V^{\pi_{j+1}}$. This can be observed in the upper off-diagonal the EBF plot in Figure 3. We conclude that the dynamics induced by TD updates may be particularly beneficial to transfer between policies in the value-improvement path, and further study of this phenomenon is a promising avenue for future work.

5.2 Auxiliary tasks for large-scale environments with sparse rewards

We now consider the problem of deep RL agents interacting with environments with sparse reward structure. Motivated by the theoretical results obtained in earlier sections, we study the effects of a variety of auxiliary tasks in this setting; our analysis indicates that random cumulants may be particularly effective in preventing representation collapse in such environments.

We modify a Double DQN agent (Van Hasselt et al., 2016) with a variety of auxiliary tasks, including random cumulants (RC) (Dabney et al., 2020), random ensemble mixtures (REM) (Agarwal et al., 2020), an ensembling approach (Anschel et al., 2017), and also compare with QR-DQN, a distributional agent (Dabney et al., 2018). Full details of these agents, including specific implementation details for deep RL versions of these auxiliary tasks, are given in Appendix D.

We evaluate these agents on a series of Atari games from the Arcade Learning Environment (Bellemare et al., 2013; Machado et al., 2018a), comprising Montezuma's Revenge, Pong, MsPacman, Seaquest, and Q*bert. In addition, we evaluate on a more challenging, sparse reward, version of Pong in which the agent does not observe negative rewards.[2]

Figure 4 shows the main results from these experiments. Recall from Section 4 that the random cumulant auxiliary task causes the agent's representation to converge to the RSBFs of P^{π} in the idealized setting. We hypothesize that this auxiliary task will therefore improve agent performance over ensemble-based auxiliary tasks in sparse-reward environments. Our empirical results support our hypothesis, with the random cumulant agent (DDQN+RC) generally per-

[2]We attempted a similar modification of the other three dense reward games, but found no agent or configuration that was able to successfully learn on them. Full details, along with hyperparameters and results on these unsuccessful modifications, are given in Appendix D.

forming well in the sparse-reward environments. Of particular note is the strong performance in Montezuma's Revenge. We expected reduced performance for DDQN+RC in the dense-reward games, but were surprised to observe improved performance here as well. However, we do note the instability seen in Seaquest. Finally, Figure 5 shows the result of a hyperparameter sweep over the number of auxiliary task heads, revealing relevant differences in the three methods considered. Overall, we find that random cumulants are a promising auxiliary task specifically in sparse-reward environments, and believe that this motivates further theoretical development to close the gap between the dynamics of representations in deep RL agents, and the settings studied above.

6 Related work

As described previously, a wide variety of auxiliary tasks have been demonstrated to improve performance in deep reinforcement learning (Sutton et al., 2011; Anschel et al., 2017; Jaderberg et al., 2017b; Bellemare et al., 2017; Barreto et al., 2017; Mirowski et al., 2017; Du et al., 2018; Riedmiller et al., 2018; van den Oord et al., 2018; Dabney et al., 2018; Gelada et al., 2019; Fedus et al., 2019; Kartal et al., 2019; Lin et al., 2019; Stooke et al., 2020; Dabney et al., 2020; Guo et al., 2020; Laskin et al., 2020). These works principally focus on demonstrating the empirical benefits of these tasks on agent performance, leaving an analysis as to why these effects occur to future work. Follow-up work on distributional reinforcement learning, for example, has begun to close the theory-practice gap (Lyle et al., 2019; Rowland et al., 2018). There is also a growing body of work on understanding the impact of representations on the sample efficiency of reinforcement learning; see for example Du et al. (2019); Van Roy and Dong (2019); Lattimore et al. (2020).

Further analysis of auxiliary tasks in deep reinforcement learning focuses on their effect on the representation learned by the agent (Bellemare et al., 2019; Dabney et al., 2020) and its ability to approximate the value functions of several policies. Additionally, Ghosh and Bellemare (2020) propose an auxiliary task based on its effect on the *stability* of learned representations. Kumar et al. (2021) also study representation collapse in deep reinforcement learning, in the absence of auxiliary tasks. Aside from reinforcement learning, there are also related empirical approaches using bootstrapping to shape representations in self-supervised learning (Grill et al., 2020), and theoretical work in characterising the regularising effect of self-distillation (Mobahi et al., 2020) and over-parametrisation (Arora et al., 2019a; Li et al., 2018) in supervised learning.

Recent work in representation learning has its roots in the broader feature selection problem in reinforcement learning. This problem has been extensively studied (Parr et al., 2008, 2007; Mahadevan and Maggioni, 2007; Petrik, 2007; Mahadevan, 2009; Kroon and Whiteson, 2009; Fard et al., 2013; Jiang et al., 2015), particularly in the linear value function approximation setting. Parr et al. (2008) analyze power-iteration-style feature learning methods, to which our analysis of the convergence of features presents notable similarity. Also closely related is the notion of feature adaptation (Menache et al., 2005; Yu and Bertsekas, 2009; Di Castro and Mannor, 2010; Bhatnagar et al., 2013; Prabuchandran et al., 2014, 2016; Barker and Ras, 2019), in which features are adaptively updated simultaneously with the weights used for value function approximation.

7 Conclusion

We have introduced a framework based on learning dynamics to analyse representations in reinforcement learning. This led to a variety of theoretical results concerning learning with and without the presence of auxiliary tasks, as well as several straightforward models for studying representation learning empirically. With this, we were able to thoroughly test a new hypothesis on the effectiveness of particular auxiliary tasks in sparse reward environments, which led to improved understanding of representation learning in RL, as well as practical modifications to deep RL algorithms.

There are many natural follow-up directions to this work. One direction is to further develop the theory associated with the learning dynamics perspective, in order to (i) understand how additional types of auxiliary tasks, in particular auxiliary tasks that don't correspond to value functions, affect the representations in the learning models developed in this paper, (ii) extend the learning models themselves to incorporate further aspects of large-scale learning scenarios, such as sample-based learning and state-visitation distribution corrections, and (iii) investigate other common learning dynamics, such as gradient TD methods (Sutton et al., 2008). There is also scope for further empirical work to develop an understanding of which auxiliary tasks are useful in certain types of environments, extending the observations relating to sparse-reward environments and random cumulants in this paper. We hope that the community will find this framework useful for testing many more hypotheses in a wide range of scenarios, ultimately leading to a better understanding of how reinforcement learning and representation learning fit together.

Acknowledgements

We thank Diana Borsa for detailed feedback on a preliminary version of this paper, and Kris Cao, Marc Bellemare, and the wider DeepMind team for valuable discussions. We also thank the anonymous reviewers for useful comments during the review process. CL is funded by an Open Phil AI Fellowship.

References

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

Oron Anschel, Nir Baram, and Nahum Shimkin. Averaged-DQN: Variance reduction and stabilization for deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.

Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Neural Information Processing Systems (NeurIPS)*, 2019a.

Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning (ICML)*, 2019b.

Edward Barker and Charl Ras. Unsupervised basis function adaptation for reinforcement learning. *Journal of Machine Learning Research*, 20(128):1–73, 2019.

André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado van Hasselt, and David Silver. Successor features for transfer in reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2017.

Bahram Behzadian and Marek Petrik. Feature selection by singular value decomposition for reinforcement learning. In *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*, 47:253–279, 2013.

Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.

Marc G Bellemare, Will Dabney, Robert Dadashi, Adrien Ali Taïga, 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.

Dimitri P Bertsekas. Feature-based aggregation and deep reinforcement learning: A survey and some new implementations. *IEEE/CAA Journal of Automatica Sinica*, 6(1):1–31, 2018.

Dimitri P Bertsekas and John N Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1st edition, 1996.

Shalabh Bhatnagar, Vivek S Borkar, and KJ Prabuchandran. Feature search in the Grassmanian in online reinforcement learning. *IEEE Journal of Selected Topics in Signal Processing*, 7 (5):746–758, 2013.

Justin A Boyan. Least-squares temporal difference learning. In *ICML*, pages 49–56. Citeseer, 1999.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

Wesley Chung, Somjit Nath, Ajin Joseph, and Martha White. Two-timescale networks for nonlinear value function approximation. In *International Conference on Learning Representations (ICLR)*, 2018.

Will Dabney, Mark Rowland, Marc G Bellemare, and Rémi Munos. Distributional reinforcement learning with quantile regression. In *AAAI Conference on Artificial Intelligence*, 2018.

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. *arXiv*, 2020.

Dotan Di Castro and Shie Mannor. Adaptive bases for reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 312–327. Springer, 2010.

Simon S Du, Sham M Kakade, Ruosong Wang, and Lin F Yang. Is a good representation sufficient for sample efficient reinforcement learning? In *International Conference on Learning Representations (ICLR)*, 2019.

Yunshu Du, Wojciech M Czarnecki, Siddhant M Jayakumar, Razvan Pascanu, and Balaji Lakshminarayanan. Adapting auxiliary losses using gradient similarity. *arXiv*, 2018.

Mahdi Milani Fard, Yuri Grinberg, Amir-massoud Farahmand, Joelle Pineau, and Doina Precup. Bell-

man error based feature generation using random projections on sparse spaces. In *Neural Information Processing Systems (NIPS)*, 2013.

William Fedus, Carles Gelada, Yoshua Bengio, Marc G Bellemare, and Hugo Larochelle. Hyperbolic discounting and learning over multiple horizons. In *Reinforcement Learning and Decision Making (RLDM)*, 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.

Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv*, 2020.

Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jeanbastien Grill, Florent Altché, Rémi Munos, and Mohammad Gheshlaghi Azar. Bootstrap latent-predictive representations for multitask reinforcement learning. *arXiv*, 2020.

Tommi Jaakkola, Michael I. Jordan, and Satinder P Singh. On the convergence of stochastic iterative dynamic programming algorithms. *Neural Computation*, 6(6), 1994.

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)*, 2017a.

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)*, 2017b.

Nan Jiang, Alex Kulesza, and Satinder Singh. Abstraction selection in model-based reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2015.

Bilal Kartal, Pablo Hernandez-Leal, and Matthew E Taylor. Terminal prediction as an auxiliary task for deep reinforcement learning. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2019.

Mark Kroon and Shimon Whiteson. Automatic feature selection for model-based reinforcement learning in factored MDPs. In *2009 International Conference on Machine Learning and Applications*, pages 324–330. IEEE, 2009.

Aviral Kumar, Rishabh Agarwal, Dibya Ghosh, and Sergey Levine. Implicit under-parameterization inhibits data-efficient deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2021.

Michael Laskin, Aravind Srinivas, and Pieter Abbeel. CURL: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2020.

Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in RL with a generative model. In *International Conference on Machine Learning (ICML)*, 2020.

Nir Levine, Tom Zahavy, Daniel J Mankowitz, Aviv Tamar, and Shie Mannor. Shallow updates for deep reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2017.

Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory (COLT)*, 2018.

Xingyu Lin, Harjatin Baweja, George Kantor, and David Held. Adaptive auxiliary task weighting for reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2019.

Clare Lyle, Marc G Bellemare, and Pablo Samuel Castro. A comparative analysis of expected and distributional reinforcement learning. In *AAAI Conference on Artificial Intelligence*, 2019.

Marios C Machado, Marc G Bellemare, and Michael Bowling. A Laplacian framework for option discovery in reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.

Marlos C Machado, Marc G Bellemare, Erik Talvitie, Joel Veness, Matthew Hausknecht, and Michael Bowling. Revisiting the Arcade Learning Environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018a.

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)*, 2018b.

Sridhar Mahadevan. Learning representation and control in Markov decision processes: New frontiers. *Foundations and Trends® in Machine Learning*, 1 (4):403–565, 2009.

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*, 8 (Oct):2169–2231, 2007.

Ishai Menache, Shie Mannor, and Nahum Shimkin. Basis function adaptation in temporal difference reinforcement learning. *Annals of Operations Research*, 134(1):215–238, 2005.

Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, Dharshan Kumaran, and Raia Hadsell. Learning to navigate in complex environments. In *International Conference on Learning Representations (ICLR)*, 2017.

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*, 518(7540):529–533, 2015.

Hossein Mobahi, Mehrdad Farajtabar, and Peter L Bartlett. Self-distillation amplifies regularization in Hilbert space. In *Neural Information Processing Systems (NeurIPS)*, 2020.

Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped DQN. In *Neural Information Processing Systems (NIPS)*, 2016.

Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2018.

Ronald Parr, Christopher Painter-Wakefield, Lihong Li, and Michael Littman. Analyzing feature generation for value-function approximation. In *International Conference on Machine Learning (ICML)*, 2007.

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 *International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.

KJ Prabuchandran, Shalabh Bhatnagar, and Vivek S Borkar. An actor critic algorithm based on Grassmanian search. In *IEEE Conference on Decision and Control*, 2014.

KJ Prabuchandran, Shalabh Bhatnagar, and Vivek S Borkar. Actor-critic algorithms with online feature adaptation. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 26(4):1–26, 2016.

John Quan and Georg Ostrovski. DQN Zoo: Reference implementations of DQN-based agents, 2020.

Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degrave, Tom Wiele, Vlad Mnih, Nicolas Heess, and Jost Tobias Springenberg. Learning by playing - solving sparse reward tasks from scratch. In *International Conference on Machine Learning (ICML)*, 2018.

Mark Rowland, Marc G Bellemare, Will Dabney, Rémi Munos, and Yee Whye Teh. An analysis of categorical distributional reinforcement learning. In *Artificial Intelligence and Statistics (AISTATS)*, 2018.

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

Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. *arXiv*, 2020.

Richard S Sutton, Csaba Szepesvári, and Hamid Reza Maei. A convergent $O(n)$ algorithm for off-policy temporal-difference learning with linear function approximation. *Neural Information Processing Systems (NIPS)*, 2008.

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.

John N Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Machine learning*, 16(3): 185–202, 1994.

Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv*, 2018.

Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double Q-

learning. In *AAAI Conference on Artificial Intelligence*, 2016.

Benjamin Van Roy and Shi Dong. Comments on the Du-Kakade-Wang-Yang lower bounds. *arXiv*, 2019.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.

Ke Ye and Lek-Heng Lim. Schubert varieties and distances between subspaces of different dimensions. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1176–1197, 2016.

Huizhen Yu and Dimitri P Bertsekas. Basis function adaptation methods for cost approximation in MDP. In *2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, pages 74–81. IEEE, 2009.