

# The Value-Improvement Path: Towards Better Representations for Reinforcement Learning

Will Dabney,<sup>1</sup> André Barreto,<sup>1</sup> Mark Rowland,<sup>1</sup>  
Robert Dadashi,<sup>2</sup> John Quan,<sup>1</sup> Marc G. Bellemare,<sup>2</sup> David Silver<sup>1</sup>

<sup>1</sup> DeepMind <sup>2</sup> Google Research  
wdabney@google.com

## Abstract

In value-based reinforcement learning (RL), unlike in supervised learning, the agent faces not a single, stationary, approximation problem, but a sequence of value prediction problems. Each time the policy improves, the nature of the problem changes, shifting both the distribution of states and their values. In this paper we take a novel perspective, arguing that the value prediction problems faced by an RL agent should not be addressed in isolation, but rather as a single, holistic, prediction problem. An RL algorithm generates a sequence of policies that, approximately, improve towards the optimal policy. We explicitly characterize the associated sequence of value functions and call it the *value-improvement path*. Our main idea is to approximate the value-improvement path holistically, rather than to solely track the value function of the current policy. Specifically, we discuss the impact that this holistic view of RL has on representation learning. We demonstrate that a representation that spans the *past* value-improvement path will also provide an accurate value approximation for *future* policy improvements. We use this insight to better understand existing approaches to auxiliary tasks and to propose new ones. To test our hypothesis empirically, we augmented a standard deep RL agent with an auxiliary task of learning the value-improvement path. In a study of Atari 2600 games, the augmented agent achieved approximately double the mean and median performance of the baseline agent.

## 1 Introduction

Whether receiving prescriptive feedback (supervised learning), or delayed evaluative feedback (reinforcement learning), machine learning requires generalization from a finite collection of examples to an unseen population. In supervised learning, generalization is sometimes framed as avoiding *overfitting* to a finite data set, and we might use methods such as regularization or cross-validation to accomplish this. In reinforcement learning (RL), generalization is no less important, and we have the same risk of overfitting to limited samples. But, additionally, the agent faces a *sequence* of learning problems, as the policy is incrementally improved towards optimality. Ideally the agent should also generalize across these problems. Although other types of generalization play a vital role in all areas of machine learning, we focus on this particular form of generalization due to its unique

role in reinforcement learning. We argue that the RL problem should be addressed with this peculiarity in mind.

The key problem underlying the ability to approximate any function is *representation learning*. Although we are ultimately interested in the optimal value function, it has been shown that a representation specialized to this function may be inadequate for representing the sequence of functions leading to it (McCallum 1996; Li, Walsh, and Littman 2006). We take this argument one step further and note that we should avoid myopically overfitting the representation to *any* value function in this sequence of functions, as each intermediate value function serves as a mere “stepping stone” along the path towards the optimal value function.

In this paper we explicitly characterize the sequence of value functions produced by RL’s policy improvement process, which we call the *value-improvement path*. We prove that the efficacy of representation learning depends upon its ability to represent this path. We use this observation both to construct new algorithms and to understand existing algorithms for representation learning.

One common and successful way to approach the representation learning problem is through the use of *auxiliary tasks*: additional prediction problems that shape the representation used by the agent (Jaderberg et al. 2017; Bellemare et al. 2019). We suggest that predictions based upon the value-improvement path provide a natural basis for representation learning. Furthermore, we analyze how well existing auxiliary tasks actually span the value-improvement path. We build on this analysis to propose novel auxiliary tasks designed with the value-improvement path in mind.

This paper provides several contributions. First, we characterize the nature of the value-improvement path. Second, we analyze and discuss both existing and novel auxiliary tasks in relation to their effect on representation generalization. Third, we provide theoretical insights that begin to explain the role of the value-improvement path in representation learning for RL. Finally, we present results from an extensive experimental study of different auxiliary tasks. This study gives preliminary evidence that representations which accurately approximate the past value-improvement path may better approximate future functions on this path.

## 2 Background

We consider a Markov decision process, or MDP,  $M \doteq (\mathcal{X}, \mathcal{A}, P, r, \gamma)$ , with finite state space  $\mathcal{X}$  and action space  $\mathcal{A}$ , transition kernel  $P : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{X})$ , reward function  $r : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \rightarrow \mathbb{R}$ , and discount factor  $\gamma \in [0, 1)$ . Given a policy  $\pi : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{A})$ , the action-value function associated with  $\pi$  gives the expected return conditioned on each possible starting state-action pair:  $Q^\pi(x, a) \doteq \mathbb{E}_\pi[\sum_{t \geq 0} \gamma^t R_t | X_0 = x, A_0 = a]$ , where  $R_t \doteq r(X_t, A_t, X_{t+1})$  are rewards. The task of *evaluation* of a policy  $\pi$  consists of computing  $Q^\pi$ . Given a policy  $\pi$ , we define the associated *Bellman operator* as

$$\mathcal{T}^\pi Q(x, a) \doteq \mathbb{E}_{\pi, P} [r(x, a, x') + \gamma Q(x', a')]; \quad (1)$$

it is well known that  $Q^\pi$  is the fixed point of  $\mathcal{T}^\pi$  (Puterman 1994). The task of *control* consists of finding a policy  $\pi^*$  maximizing the associated action-value function  $Q^* \doteq Q^{\pi^*}$ . RL algorithms based on dynamic programming approach the control problem by alternating policy evaluation (1) with *policy improvement*, in which  $Q^\pi$  is used to compute an improved policy

$$\pi'(x) \doteq \arg \max_a Q^\pi(x, a). \quad (2)$$

It can be shown that  $Q^{\pi'} \succeq Q^\pi$ , that is,  $Q^{\pi'}(x, a) \geq Q^\pi(x, a)$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ . The alternation between policy evaluation (1) and policy improvement (2) can happen at many levels of granularity. For example, if (2) is followed by one application of (1) we have the well known *value iteration* algorithm. If instead we compute  $Q^\pi$ , which corresponds to applying  $\mathcal{T}^\pi$  an infinite number of times, we recover *policy iteration* (Howard 1960). Under some mild assumptions the alternation between (1) and (2) at any level of granularity converges to  $Q^*$  (Puterman 1994). We will generically refer to algorithms obtained by alternating one application of (2) with  $n$  applications of (1), where  $n$  is possibly infinite, as *value-based* algorithms.

In RL it is assumed that the agent does not have access to the dynamics of the MDP, and thus the expectation in (1) is replaced by samples from  $P(\cdot | x, a)$ . Many RL algorithms can thus be understood as stochastic approximations of their dynamic programming counterparts; for example, the stochastic version of value iteration is the well known  $Q$ -learning algorithm (Watkins and Dayan 1992). For the sake of exposition, we will refer to both the state-value function  $V^\pi$  and the action-value function  $Q^\pi$  simply as *value functions*, using their respective symbols to clarify when needed.

**Representation learning.** Usually the state-action space  $\mathcal{X} \times \mathcal{A}$  is too big to allow for an explicit representation of the functions  $Q^\pi$ , so one must resort to an approximation  $\hat{Q}^\pi$ . A common approach is to parametrize  $\hat{Q}^\pi(x, a) = \phi(x)^\top \theta_a$ , where  $\phi : \mathcal{X} \rightarrow \mathbb{R}^K$  are features,  $\theta_a \in \mathbb{R}^K$  are modifiable parameters associated with each  $a \in \mathcal{A}$ , and  $K \in \mathbb{N}$ . Given a set of (possibly non-linear) features  $\phi : \mathcal{X} \rightarrow \mathbb{R}^K$ , policy evaluation comes down to computing the linear weights  $(\theta_a | a \in \mathcal{A})$  such that  $\hat{Q}^\pi \approx Q^\pi$ . This involves the projection of the Bellman operator onto the space spanned by  $\phi$ ,  $\Pi_\phi \mathcal{T}^\pi \hat{Q}^\pi$  for

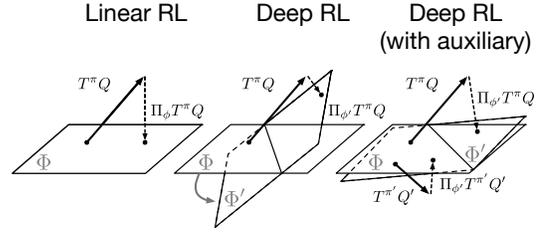


Figure 1: *Linear* approximation in RL projects the Bellman targets onto the space spanned by features. In *Deep RL* the representation itself moves to better fit the targets. However, this can lead to overfitting to the current value function. With *auxiliary tasks* the representation is regularized, reducing representation overfitting.

$\Phi = \langle \phi \rangle := \text{Span}(\{\phi_1^\pi, \dots, \phi_k^\pi\})$  (Figure 1, left). The problem of constructing the basis functions  $\phi$  is known as *representation learning*. Generally  $K \ll |\mathcal{X}|$ , so, for a fixed  $\phi$ , the space of expressible value functions is much smaller than the space of all possible value functions. Intuitively, then, a well-trained map  $\phi$  should be such that  $\phi(x)$  captures the salient information of  $x \in \mathcal{X}$  for return prediction, and ignores irrelevant details.

Representation learning has been extensively studied as a separate step in which one learns features  $\phi$  to be later used with linear function approximation. Fixed basis methods such as *tile-coding* (Sutton and Barto 2018) and the *Fourier basis* (Konidaris, Osentoski, and Thomas 2011) often work well for small input dimensions. However, the field has continually searched for methods that would adapt the representation to the MDP structure. For example, proto-value functions (Mahadevan 2005) attempt to capture the eigenvectors of the matrix  $(I - \gamma P^\pi)^{-1}$ , where  $P^\pi(x' | x) := \mathbb{E}_\pi P(x' | x, a)$  for policy  $\pi$ , and are closely related to successor representations (Dayan 1993), slow-feature analysis (Sprekeler 2011), and successor features (Barreto et al. 2017).

In *deep reinforcement learning*  $\phi$  and  $\theta_a$  are jointly learned as a (deep) neural network (Mnih et al. 2015). As long as the network has sufficient representational capacity, given enough training experience the learned representation  $\phi(x)$  will be able to approximate a policy's value function arbitrarily well. Thus, some may ask if representation learning in RL is still an open problem. We argue for the affirmative, largely due to the need for generalization to future value functions. When training in deep RL the representation itself changes to better fit the value function (Figure 1, center). This is desirable, but can also lead to highly temporally correlated features (Kolter and Ng 2009), and a representation that is degenerate, in terms of having limited span. That is, in common machine learning terms, deep RL can, and does, overfit to the current value function.

**Auxiliary tasks.** It has been argued in the literature that one way to capture the relevant information for a good representation is to learn about many aspects of the world in addition to learning a value function (Parr et al. 2008; Song et al. 2016). One idea in this direction is to define pseudo-rewards, or *cumulants*,  $c : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \rightarrow \mathbb{R}$ , and treat them

as actual rewards, either learning the value of a fixed policy or solving the induced control problem in parallel with the solution of the problem of interest (Sutton et al. 2011; Jaderberg et al. 2017). The intuition here is that these additional tasks, called *auxiliary tasks*, help shape the representation  $\phi$  (Figure 1, right), and thus limit overfitting.

### 3 The Value-Improvement Path

The fact that in value-based RL policies are computed through (2) allows us to think about the problem strictly in terms of value functions. Ultimately, we are interested in the optimal value function  $Q^*$ , from which an optimal policy can be readily computed. However, unlike in supervised learning, in RL we do not have access to samples of  $Q^*$ , and to estimate this function the agent must traverse a path across the space of value functions  $\mathcal{Q} = \{Q^\pi | \pi \in \mathcal{P}(\mathcal{A})^{\mathcal{X}}\}$ .

This special structure of the RL problem creates a number of challenges. Since we do not have direct access to the target function we are trying to approximate, we generally use the approximation itself to build the targets—a strategy sometimes referred to as “bootstrapping” (Sutton and Barto 2018). This creates a cyclic dependence of the approximation on itself that can lead to instabilities (Baird 1995; Bertsekas and Tsitsiklis 1996; Van Hasselt et al. 2018). In fact, many of the techniques currently adopted in deep RL, like target networks and replay buffers, can be interpreted as strategies to ameliorate this instability (Mnih et al. 2015).

However, in this paper we focus on another challenging aspect of the RL problem that has perhaps been overlooked so far. We argue that, when learning a representation  $\phi(x)$ , we should keep in mind that we are traversing the space of value functions, and thus over-specializing  $\phi(x)$  to a particular value function is analogous to overfitting to a finite dataset in supervised learning. In the same way that we take measures to prevent overfitting in supervised learning, we should adopt strategies to avoid over-specialization of  $\phi(x)$ .

But how can we tailor  $\phi(x)$  to a set of value functions that is not known in advance? One possible approach is to characterize the *entire* space of value functions  $\mathcal{Q}$  and try to shape  $\phi(x)$  in order to represent this space as well as possible. Dadashi et al. (2019) showed that the space of state value functions  $\mathcal{V} = \{V^\pi | \pi \in \mathcal{P}(\mathcal{A})^{\mathcal{X}}\}$  forms a polytope. Based on the theory developed by Dadashi et al. (2019), it is straightforward to show that, as an affine image of  $\mathcal{V}$ ,  $\mathcal{Q}$  is also a polytope, so we will use “value polytope” to generically refer to both  $\mathcal{V}$  and  $\mathcal{Q}$ .

Building on Dadashi et al.’s insight, Bellemare et al. (2019) proposed to shape  $\phi(x)$  by learning a set of auxiliary tasks corresponding to value functions that cover the value polytope as well as possible. Although Bellemare et al.’s approach is a clear step forward towards recognizing the nature of the approximation problem in RL, we argue that shaping the representation taking the entire value polytope  $\mathcal{Q}$  (or  $\mathcal{V}$ ) into account may be a stringent requirement in practice. This is based on the observation that the value functions of interest form a set that is generally much smaller than  $\mathcal{Q}$ .

As discussed in Section 2, any value-based RL algorithm computes a sequence of functions that, under some assumptions, end in the optimal value function  $Q^*$ . We will call this

trajectory in function space the *value-improvement path*, and formally define it as follows:

**Definition 1.** A sequence  $\{Q_0, Q_1, \dots, Q^*\}$  is called a *value-improvement path* if  $Q_{i+1} \succeq Q_i$  for  $i = 0, 1, \dots$

Value-improvement paths are worth investigating because they tend to (approximately) reflect the behavior of algorithms of practical interest. In addition, it might be possible to exploit the structure in this type of sequence to improve the generalization ability of the associated algorithm. In the next section we illustrate these points with a specific example of value-improvement path.

#### Prototypical Example of Value-Improvement Path

In order to provide intuition on the concept of value-improvement path, it might be instructive to consider for a moment the scenario studied by dynamic programming, where it is assumed that the dynamics of the MDP  $P(\cdot | x, a)$  are known (Puterman 1994). This allows for the definition of algorithms whose value-improvement paths can be easily analyzed. Perhaps the dynamic programming algorithm whose value-improvement path is easiest to visualize is policy iteration (Howard 1960). Policy iteration has very simple dynamics: starting from a policy  $\pi_0$ , compute its value function,  $Q^{\pi_0}$ , derive an improved policy  $\pi_1$  based on (2), and so on, until  $Q^{\pi_i} = Q^{\pi_{i+1}}$ .

The initial value function  $Q^{\pi_0}$  is sufficient to fully define policy iteration’s value path  $\{Q^{\pi_0}, Q^{\pi_1}, \dots, Q^*\}$ ; we will thus use  $Q^\pi$  to refer to policy iteration’s value-improvement path starting at  $Q^\pi$ . The value-improvement path  $Q^\pi$  has several interesting properties (see Appendix Figure 4):

1. **Order:**  $Q^\pi$  is a totally ordered set, since for any two  $Q^{\pi'}, Q^{\pi''} \in Q^\pi$  it must be the case that either  $Q^{\pi'} \succeq Q^{\pi''}$  or  $Q^{\pi'} \preceq Q^{\pi''}$ . This is in contrast with  $\mathcal{Q}$ , which is a partially ordered set.
2. **Structure:** As long as there is a deterministic way to break ties in (2), we can think of the space composed of all policy iteration’s value-improvement paths as a tree-like structure in which the optimal value function  $Q^*$  is the root, the first level has all the value functions that lead to  $\pi^*$  in one application of (2), and so on. Seen this way, it is clear that two value-improvement paths  $Q^\pi$  and  $Q^{\pi'}$  can intersect at arbitrary levels of the tree, and if they meet in  $Q^{\pi_i}$  they overlap from that point up, all the way to the root of the tree. More formally, if  $Q^{\pi_i} \in Q^\pi$  and  $Q^{\pi_i} \in Q^{\pi'}$ , given  $Q^{\pi_j} \in Q^\pi$  such that  $Q^{\pi_j} \succeq Q^{\pi_i}$ , then  $Q^{\pi_j} \in Q^{\pi'}$ .
3. **Size:** If  $Q^\pi \prec Q^{\pi'}$ , then  $Q^\pi \notin Q^{\pi'}$ . Also, although  $Q^{\pi'} \in Q^\pi$  implies that  $Q^{\pi'} \succeq Q^\pi$ , the converse is not necessarily true. Importantly, for a fixed discount factor  $\gamma$ , the size of any value-improvement path  $Q^\pi$  is polynomial in  $|\mathcal{X}|$  and  $|\mathcal{A}|$ , even though the number of improving policies can be exponentially large in  $|\mathcal{X}|$  (Ye 2011).

The properties above shed some light on the RL representation learning problem. Property 2 indicates that the fea-

tures  $\phi(x)$  should always be able to provide a good approximation of  $Q^*$ —a fact that is not very surprising. Perhaps more insightful is the fact that, although all value-improvement paths end at the same point  $Q^*$ , the trajectory they define in the value-function space  $\mathcal{Q}$  can be quite distinct. This suggests that the representation learning problem is context-dependent, in the sense that it can change considerably depending on the value function used as a starting point for the policy iteration process. Another interesting fact, implied by Property 3, is that, once we know  $Q^\pi$ , we should only care about the value functions  $Q^{\pi'} \in \mathcal{Q}^\pi$ . Since this set of value functions is in general much smaller than the entire polytope  $\mathcal{Q}$  (see Property 3), focusing our attention to  $\mathcal{Q}^\pi$  can significantly influence how we approach representation learning.

The structure of the value-improvement path may change depending on how exactly policy evaluation and policy improvement are applied. For example, if policy improvement (2) is applied to a subset of the state space  $\mathcal{X}$  only, there might be many paths from a given function  $Q^{\pi_i}$  to the endpoint of the path,  $Q^*$ . This means that the tree structure described in Property 2 would be replaced by a directed acyclic graph. Similarly, one should expect the structure of the value-improvement path to change if policy evaluation is not carried out to completion. For example, the value iteration algorithm alternates between a single application of the Bellman operator  $\mathcal{T}^\pi$  defined in (1) and one application of the policy improvement operator (2); in this case the resulting value-improvement path will also be quite distinct from the one induced by policy iteration—in fact, it has been shown that the intermediate functions obtained by value iteration do not belong to the value polytope  $\mathcal{Q}$  (Dadashi et al. 2019). Another example of modified policy evaluation is when policy iteration is performed using approximations  $\hat{Q}^\pi \approx Q^\pi$ . In this case the resulting value path  $\hat{\mathcal{Q}}^\pi \doteq \{\hat{Q}^{\pi_1}, \hat{Q}^{\pi_2}, \dots\}$  may no longer be a value-improvement path (Bertsekas and Tsitsiklis 1996). The value-improvement path also changes when we move from dynamic programming to RL, in which it is assumed that the agent does not have access to the dynamics of the MDP. Since in this case policy evaluation (1) is applied based on samples from  $P(\cdot|x, a)$ , one has a *distribution* over possible value-improvement paths.

In this paper we will repeatedly refer to policy iteration’s value-improvement path  $\mathcal{Q}^\pi$  as a prototypical example of this type of trajectory in function space. In the same way that knowledge of the structure underlying the true space  $\mathcal{Q}$  helps to shape the representation  $\phi(x)$ , we argue that the properties of  $\mathcal{Q}^\pi$  as defined above can help us determine a suitable  $\phi(x)$  regardless of the specific way policy evaluation and policy improvement are carried out. We elaborate on this point next.

## 4 Representation Learning Through the Lens of the Value-Improvement Path

We now revisit the formulation of the representation learning problem in light of the concept of value-improvement path. Currently representation learning in deep RL is tackled in two ways that are, in some sense, the extremes of a

spectrum of possibilities. On one extreme of the spectrum we have the common practice of ignoring the special structure of the RL problem and shaping the representation  $\phi(x)$  looking only at the current value function. As discussed, this can lead to overfitting. On the other extreme we have the recently-proposed approach of shaping  $\phi(x)$  considering the entire value polytope  $\mathcal{Q}$ , which may not be scalable (Belle-mare et al. 2019). Here we propose an intermediate formulation: representation learning in deep RL should be seen as the search for  $\phi(x)$  that allows for good approximations of all value functions in an algorithm’s value-improvement path. Using again policy iteration as a prototypical reference point, we now formally motivate this objective by restating a result by Munos (2003) in terms of policy iteration’s approximate value path  $\hat{\mathcal{Q}}^\pi$ .

Let  $\|Q\|_{d_\mu}^2 = \sum_{(x,a) \in \mathcal{X} \times \mathcal{A}} d_\mu(x, a) Q(x, a)^2$  be the (squared) Euclidean norm on  $\mathbb{R}^{\mathcal{X} \times \mathcal{A}}$  weighted by the training state-action distribution  $d_\mu$ , and, for any subspace  $U \subseteq \mathbb{R}^{\mathcal{X} \times \mathcal{A}}$ , let  $\Pi_U^{d_\mu} : \mathbb{R}^{\mathcal{X} \times \mathcal{A}} \rightarrow \mathbb{R}^{\mathcal{X} \times \mathcal{A}}$  denote the orthogonal projection into the subspace  $U \subseteq \mathbb{R}^{\mathcal{X} \times \mathcal{A}}$  with respect to  $\|\cdot\|_{d_\mu}^2$  (when it is clear from context, we will drop notational dependence on  $d_\mu$ ).

**Theorem 1.** Consider a policy  $\pi$ , a distribution over state-action pairs,  $d_\mu$ , and a representation  $\phi^\pi : \mathcal{X} \rightarrow \mathbb{R}^K$ . Suppose that policy iteration’s approximate value path  $\hat{\mathcal{Q}}^\pi$  is well-approximated in the sense that for some  $\epsilon \geq 0$  and related distributions  $d_{\mu_k}$  (Definition A, Appendix A),

$$\|\Pi_{(\phi^\pi)} Q - Q\|_{d_{\mu_k}} \leq \epsilon, \quad \forall Q \in \hat{\mathcal{Q}}^\pi, k \in [K]. \quad (3)$$

Then, for  $\pi_k$  representing the policy at the  $k^{\text{th}}$  iteration of approximate policy iteration starting with  $\pi_0 = \pi$ , we have

$$\limsup_{k \rightarrow \infty} \|Q^* - Q^{\pi_k}\|_{d_\mu} \leq \frac{2\gamma\epsilon}{(1-\gamma)^2}.$$

The proofs of our theoretical results can be found in the Appendix. Theorem 4 extends Munos’ result to action-value functions, and indicates how long-term performance of approximate policy iteration is affected by the representation’s ability to approximate functions in the value-improvement path. Similar results can be derived for value iteration (Bertsekas and Tsitsiklis 1996).

### Auxiliary Tasks and the Value-Improvement Path

The representation learning problem can potentially be approached in different ways. However, in this section we focus on a method that is commonly used in practice which can be interpreted as a way of addressing the problem as formulated above: auxiliary tasks. Specifically, we will analyze how well the representation induced by different auxiliary tasks span the value-improvement path. We begin with the standard no-auxiliary-task setting and consider methods progressively more aligned with approximating the value-improvement path. As we proceed it will be useful to consider, for each auxiliary task, what subspace the induced representation attempts to capture, and how this compares with that of the value-improvement path. To this end, Figure 2 gives an illustration of each method.

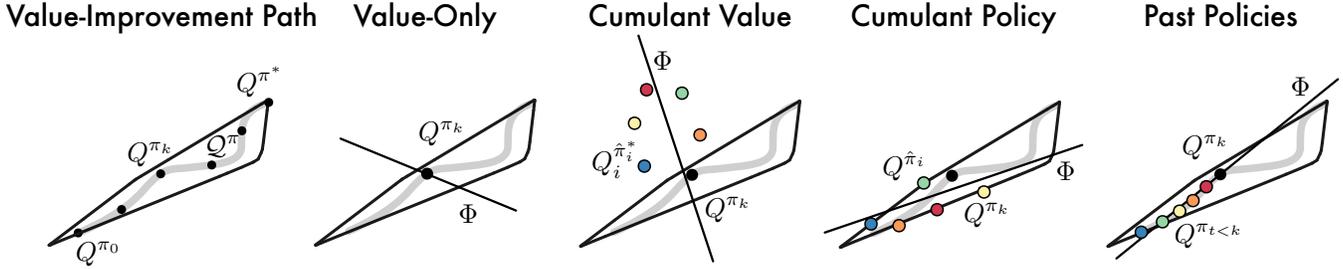


Figure 2: Illustration of the relationship between the value-improvement path,  $Q^\pi$ , and auxiliary tasks discussed in this work (shown as colored dots). The outlined space shows the hull of the value polytope, while the gray curve denotes the value-improvement path. Potential representations are shown by a solid line denoted  $\Phi$ .

**No auxiliary tasks (value-only).** When the training objective is the accurate evaluation of a fixed policy  $\pi$ , generalization to other policies, improvements or otherwise, can be very poor. Mathematically, under this objective, the aim is to find a representation  $\phi^\pi$  and weights  $(\theta_a^\pi | a \in \mathcal{A})$  which obtain a low value for the objective

$$\sum_{(x,a) \sim \mathcal{X} \times \mathcal{A}} d_\mu(x,a) (\phi^\pi(x)^\top \theta_a^\pi - Q^\pi(x,a))^2, \quad (4)$$

where  $d_\mu$  is the distribution over training state-action pairs. If  $K \geq |\mathcal{A}|$ , it is possible to achieve *zero* error on this objective, by ensuring that the subspace spanned by the coordinates of  $\phi$ , written  $\langle \phi \rangle$ , contains each of the action-value functions  $Q^\pi(\cdot, a)$  for  $a \in \mathcal{A}$ . To understand how such a representation generalizes to the approximation of other value functions, let  $Q$  be a new value function—corresponding to an improved policy, for example. The optimal approximation to  $Q$  using the representation  $\phi^\pi$  is given by  $\Pi_{\langle \phi^\pi \rangle} Q$ . Clearly, any representation  $\phi^\pi$  which achieves zero error on the objective (4) will achieve at least as good an approximation performance of  $Q$  as the representation comprising the features  $\{Q^\pi(\cdot, a) | a \in \mathcal{A}\}$ , but no further guarantees can be given. Thus, the following inequality is tight:

$$\|\Pi_{\langle \phi^\pi \rangle} Q - Q\| \leq \|\Pi_{\langle Q^\pi(\cdot, a) | a \in \mathcal{A} \rangle} Q - Q\|.$$

**Cumulant value functions.** Alternatively, we may take on the perspective of Sutton et al. (2011) and construct auxiliary tasks by learning optimal value-functions of a diverse collection of cumulants. Let  $\{c_1, \dots, c_n\}$  be a set of  $n$  cumulant functions, and for  $i = 1, \dots, n$ , let  $\hat{\pi}_i$  be a policy maximizing the value for cumulant  $c_i$ , and  $Q_i^{\hat{\pi}_i}$  its action-value function under  $c_i$ . We can think of these auxiliary tasks as capturing a subspace of the space of optimal value functions induced by all possible cumulants:  $\langle (I - \gamma P^{c_i})^{-1} c | \forall c \rangle$ . Given infinite representational capacity, adding additional auxiliary tasks would monotonically improve the generalization error. However, when  $|\phi|$  is finite, adding auxiliary tasks necessitates a trade-off between approximation errors. In this case, there is no reason to expect the space of all optimal value functions to be well-aligned with the value-improvement path, except for the final point at the optimal policy. For this reason we cannot provide much in the way of generalization guarantees for this auxiliary task. Nonetheless, these can still provide regularization and decorrelation

benefits. For example, the UNREAL agent’s *pixel control* auxiliary loss is of this type (Jaderberg et al. 2017).

**Cumulant policies.** Suppose instead we learn  $Q_i^{\hat{\pi}_i}$  using a separate approximator and use  $Q_i^{\hat{\pi}_i}$  as our auxiliary task (*i.e.* the evaluation of the auxiliary policy  $\hat{\pi}_i$  on the *true* reward function). This would make sure that the value functions used as auxiliary tasks belong to  $\mathcal{Q}$ . Clearly, the resulting method will be sensitive to the distribution of policies generated by the cumulant functions. However, assuming the cumulant functions are sufficiently expressive, the resulting set of auxiliary policies will eventually cover all deterministic policies. Thus, the representation captures the principal components of the value polytope itself (Dadashi et al. 2019). Because the value-improvement path is a subset of the value polytope, this immediately allows us to apply Theorem 4. The *adversarial value functions* method provides a principled approach to solving the problem of generating a set of such policies that span the value polytope, but exact solutions can be intractable (Bellemare et al. 2019).

The value-improvement path is in general much smaller than the value polytope itself (see Property 3). Consider for example that for any policy  $\pi$  all policies  $\hat{\pi}_i$  with  $Q_i^{\hat{\pi}_i} \prec Q^\pi$  are not in the path  $Q^\pi$ . The value functions produced by the above two methods may be entirely unrelated to those in the value-improvement path.

**Past policies in the value path.** Perhaps we can improve generalization error by explicitly restricting the auxiliary value functions to the elements of the value-improvement path. The future policies, and their value functions, are unknown, but we have already passed through some sequence of value-policy pairs during training. Rather than using arbitrary cumulants to generate policies, we can take advantage of the trajectory of improving policies itself to source these auxiliary values. This method involves taking as the auxiliary tasks the value function for the past  $k$  policies in our trajectory of policy improvement,  $(Q^{\pi_{t-k}}, \dots, Q^{\pi_{t-1}})$ . We can also consider a softened version of the PastPolicies where each auxiliary task estimates a different mixture over past value-functions. This can be implemented by simply using a different learning rate for each auxiliary task and using their average as the bootstrap target for all tasks. These two auxiliary tasks, PastPolicies and PastMixture, have not been previously proposed, but viewed

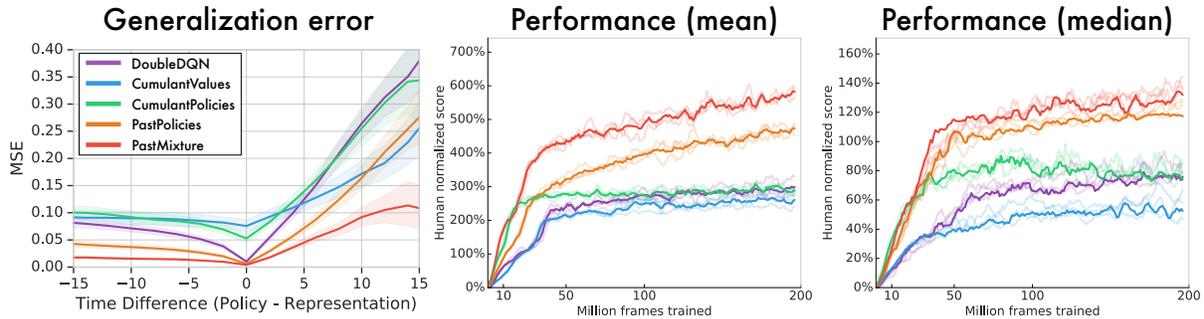


Figure 3: (Left) *Generalization*: Mean-squared error between representations  $\phi_t$  and value functions  $Q_k$  for pairs  $t, k$  drawn throughout training. (Right) *Performance*: Human-normalized performance on Atari-57. Results averaged over three seeds.

in the context of the value-improvement path we may expect them to perform well.

The PastPolicies and PastMixtures approaches should be expected to generalize well backwards toward previous policies in the trajectory, as these are what they are being trained on; however, generalization to future values will depend on the exact shape of the value-improvement path. Interestingly, there is a close connection that can be drawn to optimistic approaches to regret minimization with predictable sequences (Rakhlin and Sridharan 2013; Kalai and Vempala 2005). At an abstract level, representation learning in the context of value paths can be cast as an online learning problem: at each time step, we must select a collection of features, or equivalently, a subspace of  $\mathbb{R}^{\mathcal{X} \times \mathcal{A}}$ , and a Q-function is revealed to us, with our loss depending on how well the Q-function is approximable with our chosen subspace. The hypothesis surrounding PastPolicies and PastMixtures is that this problem is not entirely adversarial, and some useful information for future predictions is contained within previous losses; in fact, the PastPolicies approach precisely embodies the *follow-the-leader* approach to this online learning problem.

### Empirical Analysis

Our goal in this section is to empirically study the effect of the previously discussed auxiliary tasks on the quality of the learned representation. Our hypothesis is that methods with representations that generalize well along the value-improvement path lead to better long-term control performance. Whereas Theorem 4 proves a more rigorous version of this hypothesis, in this section we test our hypothesis empirically, using a novel evaluation of an agent’s representations in terms of the generalization error to past and future value functions in a value path.

For these experiments, we use the Atari-57 benchmark from the Arcade Learning Environment (Bellemare et al. 2013, ALE). We use Double DQN (van Hasselt, Guez, and Silver 2016) as our baseline non-auxiliary algorithm, and compare with each of the auxiliary tasks: CumulantValues, CumulantPolicies, PastPolicies, and PastMixtures. We generated the cumulants for CumulantValues and CumulantPolicies using a random network (details in Appendix C). Each auxiliary task is trained as a linear function of the

last hidden layer of the neural network used by Double DQN, thereby shaping the representation in different ways. To test our hypothesis that tasks aligned with the value-improvement path will lead to improved long-term performance, we first explicitly evaluate this alignment (in hindsight) by measuring approximation error between a representation at one point in time and the value functions along the value path.

Specifically, we carried out the experiment as follows. While training each agent, for 200 million environment frames, we saved the current network every 2 million frames. After training, we evaluated how well, in terms of mean-squared error, each *representation* was able to linearly fit each *value function*. Specifically, we assess how well the representation at time  $t$ ,  $\phi_t(s)$ , can linearly approximate the value functions  $\hat{Q}_k$  for  $k = t - 15, t - 14, \dots, t + 15$  (Appendix C for details).

Figure 3 (left) shows a comparison of the generalization errors for each agent on a held-out set of transitions. Each curve can be interpreted as showing the generalization of a representation to other value functions in the value path: negative values correspond to past value functions and positive values to future value functions. Note that in RL we are generally interested in minimizing the latter. Figure 3 (center, right) shows the human-normalized mean and median scores on Atari-57 (additional results in Appendix E).

These results clearly show two trends. First, the methods’ ability to generalize to future value functions largely reflects what our analysis based on the value-improvement path would predict. Second, and perhaps more important, the generalization error for future value functions is remarkably, although not perfectly, predictive of long-term performance, corroborating the main argument of this paper that the value-improvement path is the space an RL agent should generalize over. Note that the best performing sets of auxiliary sets, PastPolicies and PastMixtures, are actually novel. It should be straightforward to combine these auxiliary tasks with most value-based algorithms in the literature.

Note that CumulantValue obtains worse long-term performance than other auxiliary tasks. This could (incorrectly) lead us to conclude that auxiliary tasks should only be defined in terms of the MDP’s actual reward. Although this is

consistent with the main argument of this paper, we believe there are situations in which having other forms of auxiliary tasks may be beneficial. A closer inspection of the results reveals that in some games the CumulantValue task actually performs *best*—notably on those games where exploration is particularly difficult (Appendix Figure 8). We speculate that auxiliary tasks based on the value of cumulants may be most useful when the actual reward is sparse. An interesting direction for future work would be to provide a formal justification for cumulant-based auxiliary tasks analogous to the ones provided here for their reward-based counterparts.

As a final demonstration, we combine PastMixtures with the state-of-the-art Rainbow agent (Hessel et al. 2018). Figure 5 (Appendix) shows that, despite an existing auxiliary task effect from distributional RL, PastMixtures leads to improved performance of this state-of-the-art agent.

## 5 Discussion and Related Work

Auxiliary tasks were introduced with Sutton et al.’s (2011) Horde architecture, though at the time they were not explicitly aimed at improving an agent’s representation. Later, the UNREAL agent introduced an auxiliary task, *pixel control*, to an A3C-like agent, and showed that these significantly improved performance (Jaderberg et al. 2017). However, the trade-offs between the auxiliary losses and primary loss required close tuning of hyper-parameters. Fedus et al. (2019) proposed learning the distribution of returns for multiple discount factors. Analysis of the relationship between these functions and the value-improvement path is an interesting direction for future work.

Bellemare et al. (2019) recently argued for a geometric approach to the representation learning problem based upon the insights surrounding the *value polytope* (Dadashi et al. 2019). Their proposed representational loss takes the form of a minimization of the maximum projection error against a finite set of *adversarial value functions*.

In this paper we focused on how the concept of value-improvement path can be leveraged for representation learning through the use of auxiliary tasks. One can take the ideas presented one step further and treat the entire value-improvement path as a stationary object that can be approximated as a single function (Schaul et al. 2015; Borsa et al. 2019). This opens up interesting possibilities in terms of how to adjust such a function: since in general every sample transition can be linked to a specific policy  $\pi$  in the value-improvement path, one could think of a training regime that is exclusively “on-policy” (Sutton and Barto 2018).

There is also an intriguing connection between the value-improvement path and distributional RL. As discussed, one way to implement the PastMixtures auxiliary tasks is to adopt a different learning rate for each auxiliary task and use their average as the bootstrap target for all tasks. Interestingly, this can be seen as performing an update similar to quantile regression distributional RL, but in which the asymmetric weights are replaced with symmetric weights (Dabney et al. 2018a,b). This connection suggests an explanation as to why distributional RL is so effective in shaping the representation—a question still open in the literature. We discuss this subject in more detail in Appendix F.

## 6 Conclusions

In this paper we discussed the problem of representation learning in the context of RL. We argued that one should address this problem keeping in mind that learning a representation for RL is considerably different from the corresponding problem in supervised learning. As a consequence, the common practice of treating each policy evaluation as a conventional supervised learning problem may lead to an over-specialization to intermediate target value functions whose interest for RL is only transient. This is analogous to the problem of overfitting in supervised learning. Under this premise, we presented the following contributions:

1. A new formulation of representation learning in RL in terms of an algorithm’s value-improvement path.
2. An interpretation of the commonly-used practice of using auxiliary tasks as a way of addressing the representation learning problem under our new formulation. Specifically, we analyzed how well auxiliary tasks used in the literature, and also new ones, span the value-improvement path.
3. Two novel auxiliary tasks inspired by the concept of value-improvement path, PastPolicies and PastMixtures, that showed strong performance and can be readily combined with value-based agents in the literature.
4. A novel study investigating the effect of auxiliary tasks on the quality of the representation learned. This study is based on an original way of assessing the generalization ability of an approximator that estimates how much it spans a value-improvement path by looking at how well it can represent past and future value functions.

We believe the insights above shed light on the representation learning problem in the context of RL, allowing a better understanding of practices already used and potentially serving as an inspiration for the design of new methods.

## Acknowledgments

The authors wish to thank colleagues at DeepMind and Google for their encouragement and feedback throughout the process of working on this paper. In particular, to Georg Ostrovski and Doina Precup for discussions and thoughts on a previous draft. As well, we thank the anonymous reviewers for their constructive feedback.

## References

- Baird, L. 1995. Residual algorithms: Reinforcement learning with function approximation. In *International Conference on Machine Learning (ICML)*.
- Barreto, A.; Dabney, W.; Munos, R.; Hunt, J.; Schaul, T.; van Hasselt, H.; and Silver, D. 2017. Successor Features for Transfer in Reinforcement Learning. In *Neural Information Processing Systems (NIPS)*.
- Barth-Maron, G.; Hoffman, M. W.; Budden, D.; Dabney, W.; Horgan, D.; TB, D.; Muldal, A.; Heess, N.; and Lillicrap, T. 2018. Distributed Distributional Deterministic Policy Gradients. In *International Conference on Learning Representations (ICLR)*.

- Bellemare, M. G.; Dabney, W.; Dadashi, R.; Taiga, A. A.; Castro, P. S.; Roux, N. L.; Schuurmans, D.; Lattimore, T.; and Lyle, C. 2019. A Geometric Perspective on Optimal Representations for Reinforcement Learning. In *Neural Information Processing Systems (NeurIPS)*.
- Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The Arcade Learning Environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research* 47: 253–279.
- Bertsekas, D. P.; and Tsitsiklis, J. N. 1996. *Neuro-dynamic programming*, volume 5. Athena Scientific Belmont, MA.
- Borsa, D.; Barreto, A.; Quan, J.; Mankowitz, D. J.; van Hasselt, H.; Munos, R.; Silver, D.; and Schaul, T. 2019. Universal Successor Features Approximators. In *International Conference on Learning Representations (ICLR)*.
- Dabney, W.; Ostrovski, G.; Silver, D.; and Munos, R. 2018a. Implicit Quantile Networks for Distributional Reinforcement Learning. In *International Conference on Machine Learning (ICML)*.
- Dabney, W.; Rowland, M.; Bellemare, M. G.; and Munos, R. 2018b. Distributional reinforcement learning with quantile regression. In *AAAI Conference on Artificial Intelligence*.
- Dadashi, R.; Taïga, A. A.; Roux, N. L.; Schuurmans, D.; and Bellemare, M. G. 2019. The value function polytope in reinforcement learning. *International Conference on Machine Learning (ICML)*.
- Dayan, P. 1993. Improving generalization for temporal difference learning: The successor representation. *Neural Computation* 5(4): 613–624.
- Fedus, W.; Gelada, C.; Bengio, Y.; Bellemare, M. G.; and Larochelle, H. 2019. Hyperbolic discounting and learning over multiple horizons. *arXiv preprint arXiv:1902.06865*.
- Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI Conference on Artificial Intelligence*.
- Howard, R. 1960. *Dynamic Programming and Markov Processes*. Cambridge, MA: MIT Press.
- Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2017. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations (ICLR)*.
- Kalai, A.; and Vempala, S. 2005. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences* 71(3): 291–307.
- Kolter, J. Z.; and Ng, A. Y. 2009. Regularization and feature selection in least-squares temporal difference learning. In *International Conference on Machine Learning (ICML)*.
- Konidaris, G.; Osentoski, S.; and Thomas, P. 2011. Value function approximation in reinforcement learning using the Fourier basis. In *AAAI Conference on Artificial Intelligence*.
- Li, L.; Walsh, T.; and Littman, M. 2006. Towards a unified theory of state abstraction for MDPs. In *International Symposium on Artificial Intelligence and Mathematics (ISAIA)*.
- Mahadevan, S. 2005. Proto-value functions: Developmental reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- McCallum, A. K. 1996. *Reinforcement Learning with Selective Perception and Hidden State*. Ph.D. thesis, University of Rochester.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540): 529.
- Munos, R. 2003. Error bounds for approximate policy iteration. In *International Conference on Machine Learning (ICML)*.
- Parr, R.; Li, L.; Taylor, G.; Painter-Wakefield, C.; and Littman, M. L. 2008. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- Puterman, M. L. 1994. *Markov Decision Processes*. John Wiley & Sons, Inc.
- Rakhlin, S.; and Sridharan, K. 2013. Optimization, learning, and games with predictable sequences. In *Neural Information Processing Systems (NIPS)*.
- Rowland, M.; Dadashi, R.; Kumar, S.; Munos, R.; Bellemare, M. G.; and Dabney, W. 2019. Statistics and samples in distributional reinforcement learning. *International Conference on Machine Learning (ICML)*.
- Schaul, T.; Horgan, D.; Gregor, K.; and Silver, D. 2015. Universal value function approximators. In *International Conference on Machine Learning (ICML)*.
- Song, Z.; Parr, R. E.; Liao, X.; and Carin, L. 2016. Linear feature encoding for reinforcement learning. In *Neural Information Processing Systems (NIPS)*.
- Sprekeler, H. 2011. On the relation of slow feature analysis and Laplacian eigenmaps. *Neural computation* 23(12): 3287–3302.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. MIT Press.
- Sutton, R. S.; Modayil, J.; Delp, M.; Degris, T.; Pilarski, P. M.; White, A.; and Precup, D. 2011. Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Van Hasselt, H.; Doron, Y.; Strub, F.; Hessel, M.; Sonnerat, N.; and Modayil, J. 2018. Deep reinforcement learning and the deadly triad. *arXiv preprint arXiv:1812.02648*.

van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double Q-learning. In *AAAI Conference on Artificial Intelligence*.

Watkins, C. J.; and Dayan, P. 1992. Q-learning. *Machine Learning* 8(3): 279–292.

Ye, Y. 2011. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Mathematics of Operations Research* 36(4): 593–603.