

---

# 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>

### 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, at least 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 avoid 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 not be adequate for representing the sequence of value functions leading to it (McCallum, 1996; Li et al., 2006). We take this argument one step further and note that we should avoid myopically overfitting the representation to *any* value function in the sequence of functions generated by an RL agent, 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* (Jaderberg et al., 2017; Bellemare et al., 2019): additional prediction problems that shape the representation used by the agent. 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

<sup>1</sup>DeepMind <sup>2</sup>Google Brain. Correspondence to: Will Dabney <wdabney@google.com>.

build on this analysis to propose novel auxiliary tasks designed with the value-improvement path in mind.

This paper provides several primary 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 sets of auxiliary tasks. This study shows that representations which accurately approximate the past value-improvement path may also 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.

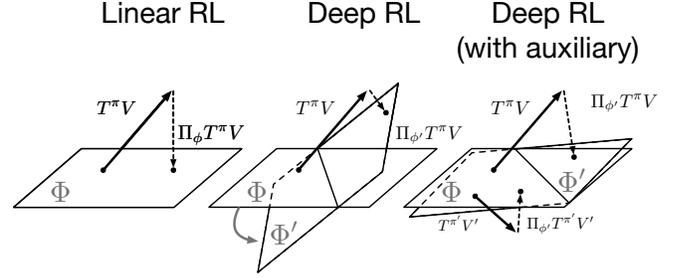


Figure 1. Linear approximation in RL projects the Bellman targets onto the space spanned by features. In *Deep RL* the representation itself moves to better fit the targets. However, this can lead to overfitting to the current value function. With *auxiliary tasks* the representation is regularized,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

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 & 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  $\hat{Q}^\pi(x, a) = \langle \phi(x), \theta_a \rangle$ , where  $\langle \cdot, \cdot \rangle$  denotes the inner product and  $\phi : \mathcal{X} \rightarrow \mathbb{R}^K$  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  $\Phi = \langle \phi \rangle := \text{Span}(\{\phi_1^\pi, \dots, \phi_k^\pi\})$  (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 & Barto, 2018) and the *Fourier basis* (Konidaris et al., 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 a policy  $\pi$ , and are closely related to successor representations in the brain (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 (see Figure 1) (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 & 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 in order 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 & Barto, 2018). This creates a cyclic dependence of the approximation on itself that can lead to instabilities (Baird, 1995; Bertsekas

& Tsitsiklis, 1996). 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 an 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  $i = 0, 1, \dots$ , where  $Q' \succeq Q$  indicates  $Q'(x, a) \geq Q(x, a)$  for all  $(x, a) \in \mathcal{X} \times \mathcal{A}$ .

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.

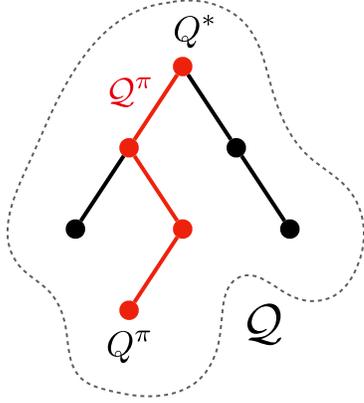


Figure 2. Schematic representation of policy iteration’s value-improvement path  $\mathcal{Q}^\pi$ . Each point in the diagram represents a value function  $Q^\pi \in \mathcal{Q}$ . Note any time two value-improvement paths intersect, they merge.

### 3.1. A 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  $\mathcal{Q}^\pi$  to refer to policy iteration’s value-improvement path starting at  $Q^\pi$ . The value-improvement path  $\mathcal{Q}^\pi$  has several interesting properties, also illustrated in Figure 2:

1. **Order:**  $\mathcal{Q}^\pi$  is a totally ordered set, since for any two  $Q^{\pi'}, Q^{\pi''} \in \mathcal{Q}^\pi$  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  $\mathcal{Q}^\pi$  and  $\mathcal{Q}^{\pi'}$  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

$Q^{\pi_i} \in \mathcal{Q}^\pi$  and  $Q^{\pi_i} \in \mathcal{Q}^{\pi'}$ , given  $Q^{\pi_j} \in \mathcal{Q}^\pi$  such that  $Q^{\pi_j} \succeq Q^{\pi_i}$ , it must be the case that  $Q^{\pi_j} \in \mathcal{Q}^{\pi'}$ .

3. **Size:** If  $Q^\pi \prec Q^{\pi'}$ , then  $Q^\pi \notin \mathcal{Q}^{\pi'}$ . Also, although  $Q^{\pi'} \in \mathcal{Q}^\pi$  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  $\mathcal{Q}^\pi$  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 features  $\phi(x)$  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 have a big effect on the way we approach the representation learning problem.

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 end-point 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 & 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 (Bellemare 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  $U \subseteq \mathbb{R}^{\mathcal{X}}$ , 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  $U^{\mathcal{A}} \subseteq \mathbb{R}^{\mathcal{X} \times \mathcal{A}}$  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 2, Appendix A),

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

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 1 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 & Tsitsiklis, 1996).

##### 4.1. 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 3 gives an illustration of each method.

**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

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

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 a similar perspective to that of Sutton et al. (2011) and construct auxiliary tasks by learning optimal value-functions

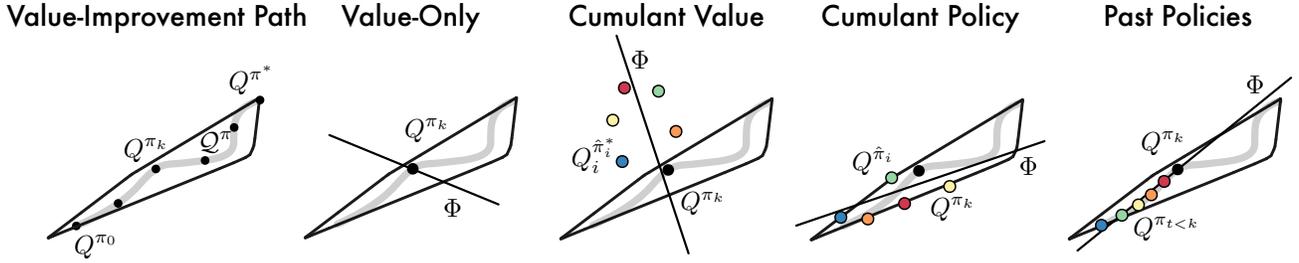


Figure 3. Illustration of the relationship between the value-improvement path,  $\mathcal{Q}^\pi$ , and auxiliary tasks discussed in this work (shown as colored dots). The outlined space shows the hull of the value polytope, while the gray curve denotes the value-improvement path. Potential representations are shown by a solid line denoted  $\Phi$ .

of a diverse collection of cumulants.<sup>1</sup> 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  $Q^{\hat{\pi}_i}$  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:  $\langle (I - \gamma P^{\pi^*})^{-1} c | \forall c \rangle$ . 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, in practice 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  $Q^{\hat{\pi}_i}$  using a separate approximator and use  $Q^{\hat{\pi}_i}$  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 1. 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

<sup>1</sup>Recall that a cumulant is a pseudo-reward function  $c : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \rightarrow \mathcal{R}$  that in general does not coincide with the actual reward defined in the MDP.

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  $Q^{\hat{\pi}_i} \prec Q^\pi$  are not in the path  $\mathcal{Q}^\pi$ . The value functions produced by the above two methods may be entirely unrelated to those in the value-improvement path. Thus, given limited capacity, perhaps we can do better.

**Past policies in the value path.** Indeed, 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 past-policies where each auxiliary task estimates a different mixture over past policies. 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, past-policies and past-mixture, have not been previously proposed, but when viewed through 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 it is 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 & Sridharan, 2013; Kalai & Vempala, 2005). At an abstract level, representation learning in the context of value

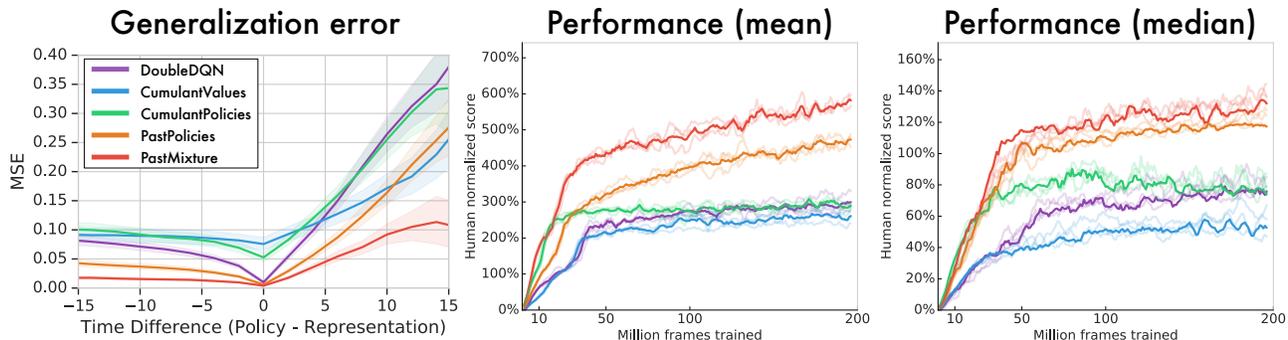


Figure 4. (Left) *Generalization*: Mean-squared error between representations  $\phi_t$  and value functions  $Q_k$  for pairs  $t, k$  drawn throughout training. (Right) *Performance*: Human-normalized performance on Atari-57. Results averaged over three seeds. Full set of results in Figure 7 of the Appendix.

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. This problem differs from many typical online learning problems in an interesting way; the action space is not a convex subset of Euclidean space, but is the *Grassmannian* space of fixed-dimension subspaces. Nevertheless, we believe this connection represents an interesting direction for future work, and a rich source of ideas for further auxiliary tasks in deep RL.

## 4.2. 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. We tested this hypothesis 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 et al., 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 B). 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  $k = t - 15, t - 14, \dots, t + 15$  (see Appendix B for full details).

Figure 4 (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 4 (center, right) shows the human-normalized mean and median scores on Atari-57 (additional results in Appendix D).

These results clearly show two trends. First, the methods’ ability to generalize to future value functions reflects to a large extent 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 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.

As a final observation, note that the CumulantValue agent obtains worse long-term performance than the other auxiliary tasks. This could suggest that auxiliary tasks should be defined in terms of the MDP’s actual reward, which is consistent with the main argument of this paper, but we believe this is missing a broader point. A closer inspection of the results reveals that in some games the cumulant-value auxiliary tasks actually performs *best*—notably on those games where exploration is particularly difficult (see Figure 7 in the Appendix). 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 investigations would be to provide a formal justification for cumulant-based auxiliary tasks analogous to the ones provided here for their reward-based counterparts.

## 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. Learning value functions with different discount factors as auxiliary tasks is an interesting approach; however, the connection between different discounts and the true discount of the MDP is unclear.

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 & 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 E.

## 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 excellent performance on our experiments and can be readily combined with value-based agents in the literature.
4. A novel study investigating the effect of different 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.

## References

- Baird, L. Residual algorithms: Reinforcement learning with function approximation. In *International Conference on Machine Learning (ICML)*, 1995.
- Barreto, A., Dabney, W., Munos, R., Hunt, J., Schaul, T., van Hasselt, H., and Silver, D. Successor features for transfer in reinforcement learning. In *Neural Information Processing Systems (NIPS)*, 2017.
- Barth-Maron, G., Hoffman, M. W., Budden, D., Dabney, W., Horgan, D., TB, D., Muldal, A., Heess, N., and Lillicrap, T. Distributed distributional deterministic policy gradients. In *International Conference on Learning Representations (ICLR)*, 2018.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The Arcade Learning Environment: an evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Bellemare, M. G., Dabney, W., and Munos, R. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.
- Bellemare, M. G., Dabney, W., Dadashi, R., Taïga, A. A., Castro, P. S., Roux, N. L., Schuurmans, D., Lattimore, T., and Lyle, C. A geometric perspective on optimal representations for reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Neuro-dynamic programming*, volume 5. Athena Scientific Belmont, MA, 1996.
- Borsa, D., Barreto, A., Quan, J., Mankowitz, D. J., van Hasselt, H., Munos, R., Silver, D., and Schaul, T. Universal successor features approximators. In *International Conference on Learning Representations (ICLR)*, 2019.
- Dabney, W., Ostrovski, G., Silver, D., and Munos, R. Implicit quantile networks for distributional reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2018a.
- Dabney, W., Rowland, M., Bellemare, M. G., and Munos, R. Distributional reinforcement learning with quantile regression. In *AAAI Conference on Artificial Intelligence*, 2018b.
- Dadashi, R., Taïga, A. A., Roux, N. L., Schuurmans, D., and Bellemare, M. G. The value function polytope in reinforcement learning. *International Conference on Machine Learning (ICML)*, 2019.
- Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5(4):613–624, 1993.
- Fedus, W., Gelada, C., Bengio, Y., Bellemare, M. G., and Larochelle, H. Hyperbolic discounting and learning over multiple horizons. *arXiv*, 2019.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., and Silver, D. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI Conference on Artificial Intelligence*, 2018.
- Howard, R. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA, 1960.
- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., and Kavukcuoglu, K. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Kalai, A. and Vempala, S. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Kolter, J. Z. and Ng, A. Y. Regularization and feature selection in least-squares temporal difference learning. In *International Conference on Machine Learning (ICML)*, 2009.
- Konidaris, G., Osentoski, S., and Thomas, P. Value function approximation in reinforcement learning using the Fourier basis. In *AAAI Conference on Artificial Intelligence*, 2011.
- Li, L., Walsh, T., and Littman, M. Towards a unified theory of state abstraction for MDPs. In *International Symposium on Artificial Intelligence and Mathematics (ISAIM)*, 2006.
- Mahadevan, S. Proto-value functions: Developmental reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2005.
- McCallum, A. K. *Reinforcement Learning with Selective Perception and Hidden State*. PhD thesis, University of Rochester, 1996.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Munos, R. Error bounds for approximate policy iteration. In *International Conference on Machine Learning (ICML)*, 2003.

- Parr, R., Li, L., Taylor, G., Painter-Wakefield, C., and Littman, M. L. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2008.
- Puterman, M. L. *Markov Decision Processes*. John Wiley & Sons, Inc., 1994.
- Rakhlin, S. and Sridharan, K. Optimization, learning, and games with predictable sequences. In *Neural Information Processing Systems (NIPS)*, 2013.
- Rowland, M., Dadashi, R., Kumar, S., Munos, R., Bellemare, M. G., and Dabney, W. Statistics and samples in distributional reinforcement learning. *International Conference on Machine Learning (ICML)*, 2019.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal value function approximators. In *International Conference on Machine Learning (ICML)*, 2015.
- Song, Z., Parr, R. E., Liao, X., and Carin, L. Linear feature encoding for reinforcement learning. In *Neural Information Processing Systems (NIPS)*, 2016.
- Sprekeler, H. On the relation of slow feature analysis and Laplacian eigenmaps. *Neural computation*, 23(12):3287–3302, 2011.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., and Precup, D. Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2011.
- van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double Q-learning. In *AAAI Conference on Artificial Intelligence*, 2016.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine Learning*, 8(3):279–292, 1992.
- Ye, Y. 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, 2011.

## Appendices

### A. Proofs

**Definition 2.** For a distribution  $d_\mu$ , optimal policy  $\pi^*$  and sequence of policies due to policy iteration  $\pi_k$ , define the stochastic matrices (Munos, 2003),

$$Q_k = \frac{(1-\gamma)^2}{2}(I - \gamma P^{\pi^*})^{-1} \times \\ \left[ P^{\pi_{k+1}}(I - \gamma P^{\pi_{k+1}})^{-1} + P^{\pi^*}(I - \gamma P^{\pi_k})^{-1} \right], \\ \tilde{Q}_k = \frac{(1-\gamma)^2}{2}(I - \gamma P^{\pi^*})^{-1} \times \\ \left[ P^{\pi_{k+1}}(I - \gamma P^{\pi_{k+1}})^{-1}(I + \gamma P^{\pi_k}) + P^{\pi^*} \right].$$

Then, the *related distributions* are given by  $d_{\mu_k} := d_\mu Q_k$ .

#### Proof of Theorem 1

*Proof.* Our proof extends the results of Munos (2003) from state-value functions to action-value functions. First, we bound the norm of the value function approximation error at iteration  $k$ ,

$$\|V_k - V^{\pi_k}\|_\mu^2 = \|\mathbb{E}_{\pi_k} Q_k - \mathbb{E}_{\pi_k} Q^{\pi_k}\|_\mu^2 \leq \|Q_k - Q^{\pi_k}\|_{\mu \cdot \pi_k}^2.$$

Thus, bounding the per-iteration approximation error of the action-value function similarly bounds the value function approximation error. Second, we lower-bound the suboptimality of the policy at iteration  $k$ ,

$$\|Q^* - Q^{\pi_k}\|_{\mu \cdot \pi_k}^2 = \\ \mathbb{E}_{(x,a) \sim \mu \cdot \pi_k} \left( \gamma \mathbb{E}_{x' \sim P(\cdot|x,a)} [V^*(x') - V^{\pi_k}(x')] \right)^2, \\ \leq \gamma^2 \|V^* - V^{\pi_k}\|_{\mu \cdot \pi_k P}^2.$$

Combining these two with Munos’s results for value functions yields our result.  $\square$

### B. Experimental details

**Hyper-parameter sweep** As a prelude to our primary empirical study we performed a comparable hyper-parameter sweep on all algorithms on a set of six games (Assault, Asterix, Frostbite, Gopher, Seaquest, MsPacman), and used the best performing hyper-parameters for each algorithm when running on the full Atari-57 suite.

For all methods we swept over learning rates (0.00025, 0.0001, 0.00005, 0.00001) and number of auxiliary tasks (4, 16, 64, 128). For past-policies, due to memory constraints we instead considered the range (2, 4, 8, 16) for number of tasks. For cumulant-values and cumulant-policies we swept (jointly with other parameters) over the type of cumulant functions generated (i.e. using the output of a random network passed through either a tanh or sign function).

**Generalization error analysis** For each agent we performed the generalization analysis on one seed (randomly chosen) and all 57 games. The results in Figure 4 are averages over all games for each method, with mean-squared errors first normalized to be between  $[0, 1]$  using the full range of errors generated by all methods *on that game*. This allows us to show an aggregated plot, despite different games yielding significantly different errors. Note that this means results are averaged over 57 independent runs, each with up to 100 saved networks, per algorithm.

The horizontal axis gives the time difference,  $k - t$ , between the policy whose value function,  $Q^{\pi_k}$ , provides the prediction target, and the representation,  $\phi_t$ , upon which we perform linear regression to predict that target. The procedure is as follows: to generate the data we run the  $\epsilon$ -greedy policy with respect to  $Q^{\pi_k}$  for 50000 frames, with  $\epsilon = 0.005$ , storing frames and value function outputs of the network. This results in the training examples  $\{(s_i, a_i), Q_k(s_i, a_i)\}$  for  $k = t - 15, t - 14, \dots, t + 15$ . We then divide this data set randomly into a training set and testing set with a 90%/10% split, perform linear regression on the training set,  $\min_{w_a} \sum_i (\langle \phi_t(s_i), w_a \rangle - Q_k(s_i, a_i))^2$  for  $k = t - 15, t - 14, \dots, t + 15$  and  $a \in \mathcal{A}$ , and report mean-squared error on the test set.

**Additional details** Other training details match those of van Hasselt et al. (2016). In particular, note that we use clipped rewards, end-of-episode life-loss signal, and the standard  $\gamma = 0.99$ . Experimental results were averages over 3 seeds.

### C. Auxiliary tasks

In the main text we give brief descriptions of each of the auxiliary tasks used in our experiments. Here, we attempt to give a more detailed and reproducible account of how each of these was implemented. For all the auxiliary tasks the additional value functions are parameterized as additional linear heads off of the same shared hidden layer as the primary value function. As these value function heads are part of the main network they use target networks in the same way as is standard with DQN agents. We will use the variable  $n$  to denote the number of such additional action-value heads.

**Cumulant Value** The *Cumulant Value* task involves learning action-value functions of other reward functions (cumulants). There are two components to this task: (1) the algorithm used to estimate the action-value function, and (2) how the cumulants are produced. For the first, we use the same Double Q-learning update as used for the primary value function. For the second, we use the same neural network as for a standard DQN, but with output dimension of  $n$  instead of  $|\mathcal{A}|$ . For an input state, consisting of stacked

frames,  $x$  we denote the output of this randomly initialized network as  $f_i(x) \in \mathbb{R}$ . Given an observed transition from state  $x_t$  to next state  $x_{t+1}$ , we define the corresponding cumulant as

$$c_i(x_t, x_{t+1}) := \tanh(s \times [f_i(x_{t+1}) - f_i(x_t)]),$$

where  $s = 100$  was tuned by hand to produce cumulants across the full range  $(-1, 1)$ . However, we note that this same end could be achieved by changing the initialization of the network weights. Thus, the auxiliary value functions here are best viewed as estimating the optimal action-value function for their corresponding cumulants,  $Q_i^{\hat{\pi}_i}$ . Although these cumulants have the same range of values as the primary (clipped) reward function, they tend to be significantly less sparse. Finally, note that we stop gradients from flowing into the cumulant network, so they are fixed at their randomly initialized values.

**Cumulant Policy** The *CumulantPolicy* task builds upon the CumulantValue task but uses an entirely separate (third) network to estimate the cumulant-specific action-value functions. These action-value functions are used to define a cumulant policy  $\hat{\pi}_i$ , which is the greedy policy on the corresponding cumulant value functions. The auxiliary value heads on the main network are trained to *evaluate* these cumulant policies with respect to the primary reward function. That is, the policy  $\hat{\pi}_i$  being targeted is the same as in the CumulantValue task, but the reward function is the true reward instead of the cumulant for which the policy was optimized. We denote these auxiliary action-value functions with  $Q^{\hat{\pi}_i}$ .

**Past Policies** In the *PastPolicies* task the auxiliary action-value function heads are trained to evaluate different policies under the true reward function. Here, these auxiliary policies,  $\hat{\pi}_i$ , come from the previous target networks of the primary action-value function. Specifically, let  $\theta_m$  and  $\bar{\theta}_m = \theta_{m-1}$  denote the network parameters and those of the target network respectively after  $m - 1$  target network updates. Then, the  $n$  auxiliary policies are defined as  $\hat{\pi}_i(x) := \arg \max_{a \in \mathcal{A}} Q(x, a; \bar{\theta}_{m-i})$ . Thus, the target policy for auxiliary head  $i$  comes from a network that is  $i$  target updates old, and thus forms a sliding window over recent past policies of the main action-value function.

**Past Mixture** Although similarly motivated, the *PastMixture* task takes quite a different approach. In this case, each auxiliary action-value head tracks the primary reward and policy, but with a different learning rate. That is, each auxiliary head takes as its target the Q-learning target with the bootstrap term composed of the average of all auxiliary head predictions, but the loss for each head is weighted differently ( $\alpha_i = (i + 1)/(n + 1)$ ). This causes each auxiliary head to track changes in value faster or slower than each

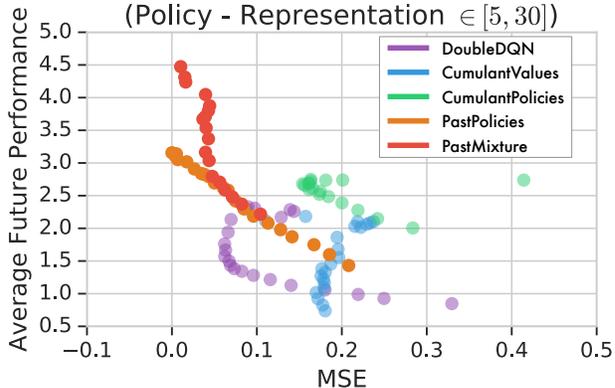


Figure 5. Scatter plot comparing, relative to a feature representation and averaged over all 57 games, an algorithm’s future performance (averaged over policies [5, 30] million frames in the future) against the generalization MSE on the same time period. Colors indicate algorithm consistent with the main text. We emphasize, with opacity, the algorithms which are *consistent* with the value-improvement path arguments we introduced.

Algorithm	Pearson coefficient	p-value
Double DQN	-0.518	0.014
Cumulant Values	0.370	0.091
Cumulant Policies	-0.287	0.248
Past Targets	-0.995	2.38e-21
Past Mixture	-0.895	5.50e-7

Table 1. Correlation between generalization errors for future values in the value-improvement path and future performance.

other, providing different mixtures between the current and past value estimates.

### D. Correlation analysis

We attempted to further evaluate our claims that generalization error on the value-improvement path is predictive of future performance. To this end, we computed, for each representation and averaged over games, the average future performance of policies between 5 and 30 million frames in the future. We also computed the mean-squared error for that representation linearly predicting the future value functions in the same window. In Figure 5 we show a scatter plot of these results and in Table 1 we give the Pearson correlation coefficients and resulting  $p$ -values.

Notice the particularly high degree, and significance, of correlation between generalization error and future performance for the two auxiliary tasks that are conceptually best aligned with the value-improvement path.

There are limitations of this analysis, primarily it only shows correlation, and so cannot establish a clear causal relationship. That is, both are strongly affected by the current quality of the agent’s representation (*i.e.*, better performing agents also have better representations, and stronger future performance). However, we believe it provides additional insight into the nature of the representation learning problem in RL.

### E. Understanding distributional RL through the concept of value-improvement path

In *distributional reinforcement learning* (DRL), the aim is to learn the full distribution of the random returns  $Z^\pi(x, a) = \sum_{t \geq 0} \gamma^t R_t | X_0 = x, A_0 = a$  for each  $(x, a) \in \mathcal{X} \times \mathcal{A}$ , rather than just their expected value. Typically, practical approaches to distributional RL involve learning parametrized approximations to these distributions, including categorical distributions with fixed support (Bellemare et al., 2017) and discrete distributions with varying support (Dabney et al., 2018b). Recently, it has been shown that all such approaches can be interpreted as learning a finite collection of auxiliary statistics of the return distribution (Rowland et al., 2019), and thus these methods fit naturally in the framework of auxiliary tasks. This interpretation may help to explain their significantly improved empirical performance (Hessel et al., 2018; Barth-Maron et al., 2018; Dabney et al., 2018a).

Distributional RL may confer some of the same benefits as the previous auxiliary task (see Proposition 2). However, viewed in the context of the value-improvement path, distributional RL provides additional avenues for analysis and understanding.

An interesting interpretation of some forms of distributional RL is that they are a natural way of decomposing the representation problem into simpler sub-problems, leading to a sort of *divide-and-conquer* strategy. We will use quantiles to illustrate this point, though the intuition may apply to other forms of distributional RL as well. Recall that a quantile function is defined as  $Z_\tau^\pi \doteq F_\pi^{-1}(\tau)$ , where  $\tau \in [0, 1]$  and  $F_\pi$  is the cumulative distribution function over possible returns under policy  $\pi$ . From this definition it should be clear that, given an MDP, all stochastic policies  $\pi$  share the same extreme quantile functions  $Z_0^\pi$  and  $Z_1^\pi$ .<sup>2</sup> This means that, once we have learned  $Z_0^\pi$  and  $Z_1^\pi$  for *any* policy, we can “re-use” them in the representation of *all* policies  $\pi'$  along the value-improvement path. Now, if we increase  $\tau$  from 0 only slightly (or, equivalently, slightly reduce it from  $\tau = 1$ ), we would expect the corresponding  $Z_\tau^\pi$  not to vary much as a function of  $\pi$ , so these quantile functions should still be quite transferable from one policy to the other. As  $\tau \rightarrow 0.5$ , we would expect  $Z_\tau^\pi$  to vary more and more as

a function of  $\pi$ , but even when  $\tau = 0.5$ , it should not vary much more than the mean—that is, the value function itself.

In summary, the “divide” step consists in decomposing the problem of approximating the mean  $Q^\pi$  into  $n$  problems of estimating the quantiles  $Z_\tau^\pi$ . The “conquer” step, in turn, is to average the solutions of the individual problems, that is,  $\hat{Q}^\pi = \frac{1}{n} \sum_{i=0}^n \hat{Z}_{\tau_i}^\pi$ . If each of the sub-problems is indeed simpler than the original, in the sense that it is possible to leverage more information from past estimates of  $Z_\tau^\pi$  to represent the current one, then we should expect the representation learned by distributional RL to be useful along the value-improvement path. In what follows we will make these statements more formal and give some examples of scenarios in which the intuition above holds.

**Proposition 1.** For a policy  $\pi$ , let  $Z_\tau^\pi : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$  be an auxiliary task trained by asymmetric regression (e.g. quantile regression) at threshold  $\tau \in [0, 1]$ . Assume that  $Z_\tau^\pi$  is bounded, monotonically increasing in  $\tau$  and Lipschitz continuous such that  $\|Z_\tau^\pi - Z_{\tau'}^\pi\| \leq \beta|\tau - \tau'|$ . Then,

$$\max_{\pi, \pi'} \|Z_\tau^\pi - Z_{\tau'}^{\pi'}\| \leq 2\beta \min\{\tau, 1 - \tau\}, \forall \tau \in [0, 1].$$

*Proof.* Begin by observing that, for this class of statistics of the return distribution, for all pair of policies  $\pi, \pi'$

$$\begin{aligned} \pi(x, a) > 0, \pi'(x, a) > 0, \forall a \in \mathcal{A} \\ \implies Z_{\tau=0}^\pi = Z_{\tau=0}^{\pi'}, Z_{\tau=1}^\pi = Z_{\tau=1}^{\pi'}. \end{aligned}$$

Let  $\tau' \in \{0, 1\}$ . Then

$$\begin{aligned} \|Z_\tau^\pi - Z_{\tau'}^{\pi'}\| &= \|Z_\tau^\pi - Z_{\tau'}^\pi + Z_{\tau'}^\pi - Z_{\tau'}^{\pi'}\|, \\ &\leq \|Z_\tau^\pi - Z_{\tau'}^\pi\| + \|Z_{\tau'}^\pi - Z_{\tau'}^{\pi'}\|, \\ &\leq 2\beta|\tau - \tau'|. \end{aligned}$$

Optimizing over  $\tau' \in \{0, 1\}$  yields the result.  $\square$

In Figure 6 we show a concrete example of the consequences of Proposition 1. The MDP is a 3-state chain, with  $\gamma = 0.7$  and zero reward except for the terminal state (where  $r = 1$ ). The actions transition left (right) with probability 0.9 and transition in a random direction with probability 0.1. The convex hull of the value polytope is shown by the gray surface. We track policies as they interpolate between two fixed policies, varying the mixture  $\alpha \in [0.1, 0.3, 0.5, 0.7, 0.9]$ . The value functions for these interpolating policies are shown, by color, with the filled points. Finally, we demonstrate the effects of Proposition 1 by plotting, as a color-matched line, the spectrum of quantile functions for each policy. Observe, consistent with our result, that each policy’s quantile function converges to the same point as  $\tau \rightarrow 0$  and  $\tau \rightarrow 1$ , and that they vary maximally as  $\tau \approx 0.5$ . Furthermore, we observe that they reflect the divide-and-conquer property, where the largest variation is no larger than the variability of the value polytope itself.

<sup>2</sup>Stochastic policies with non-zero probability on all actions.

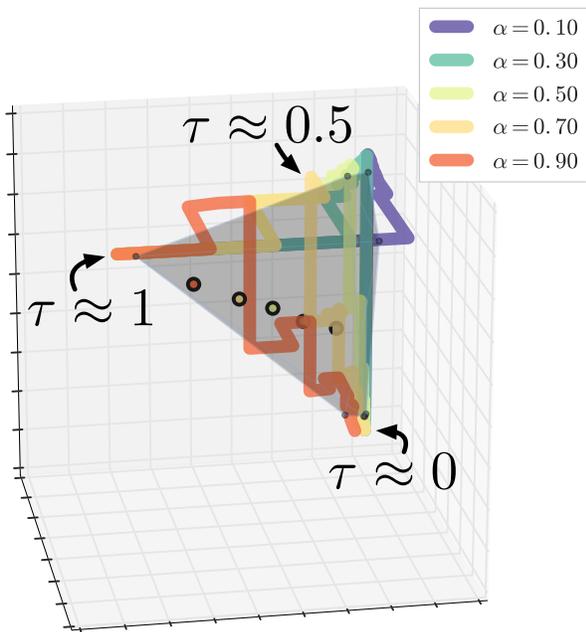


Figure 6. Example of Proposition 1.

**Definition 3.** A **deterministic MDP** is an MDP with deterministic transition dynamics and immediate rewards.

**Proposition 2.** Suppose we have the following distributional RL learning update for the action-value distribution  $\eta_n$ :

$$\eta_{n+1} := (1 - \alpha)\eta_n + \alpha\eta_n^*,$$

where  $\eta_n^*$  is the state-action target distribution at update  $n$  which we assume to be Diracs, and  $\alpha \in (0, 1)$ . We have that for  $\alpha$  small enough, the  $N$  quantiles of  $\eta_n$  match exactly  $N$  past values of the target history  $\{\eta_j^*\}_{j \leq n}$ .

*Proof.* We can express  $\eta_n(x, a)$  as a mixture of past targets:

$$\begin{aligned} \eta_n(x, a) &= (1 - \alpha)\eta_{n-1}(x, a) + \alpha\eta_{n-1}^*(x, a) \\ &= \sum_{i=0}^{n-1} \alpha(1 - \alpha)^{n-1-i} \eta_i^*(x, a) + (1 - \alpha)^n \eta_0(x, a) \end{aligned}$$

Suppose  $\alpha < \frac{1}{N}$ , the distribution  $\eta_n(x, a)$  is as mixture of  $n$  Diracs  $\eta_i^*(x, a)$  with associated weight  $\alpha(1 - \alpha)^{n-1-i} < \frac{1}{N}$ . Therefore the  $N$  quantiles of  $\eta_n(x, a)$  match  $N$  past target distributions.

□

# The Value-Improvement Path

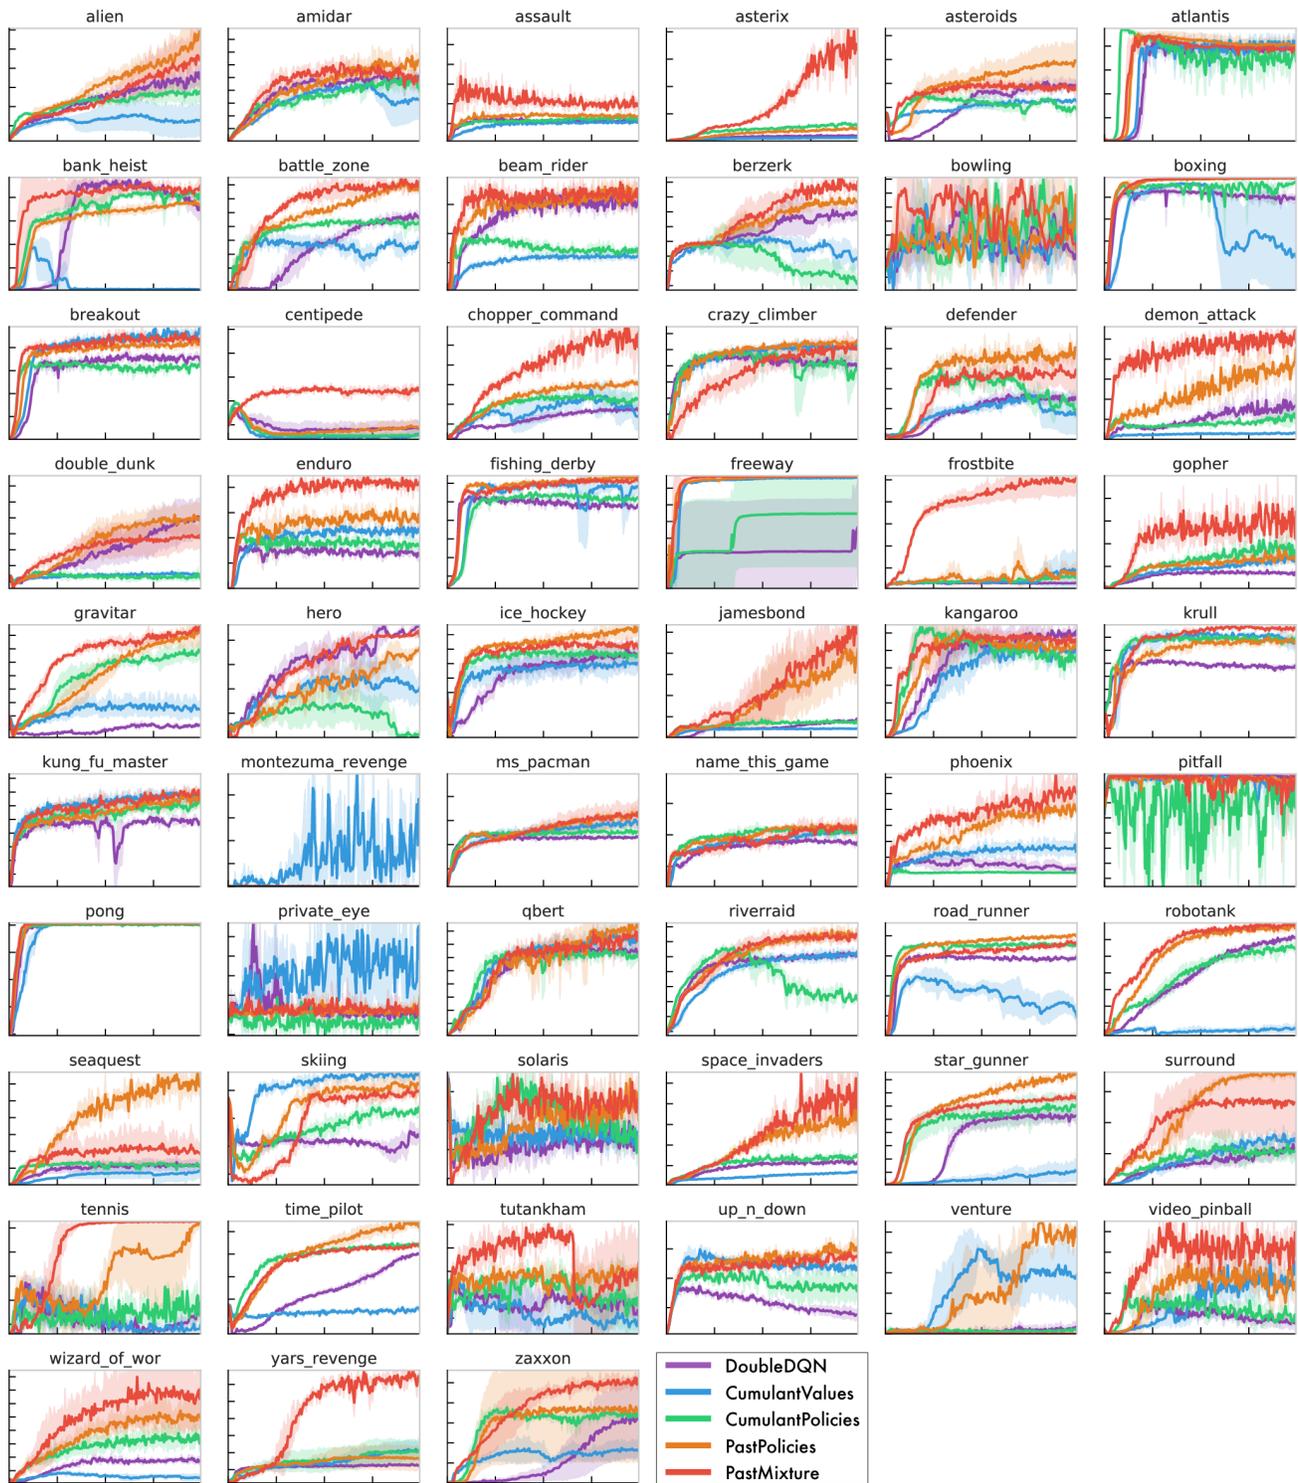


Figure 7. Atari-57 results for all methods, averages over three seeds and error bands show standard deviation.