

A Generalized Projected Bellman Error for Off-policy Value Estimation in Reinforcement Learning

Andrew Patterson

AP3@UALBERTA.CA

Adam White

AMW8@UALBERTA.CA

Martha White

WHITEM@UALBERTA.CA

Department of Computing Science and the Alberta Machine Intelligence Institute (Amii)

University of Alberta

Edmonton, Alberta, Canada

Editor: George Konidaris

Abstract

Many reinforcement learning algorithms rely on value estimation, however, the most widely used algorithms—namely temporal difference algorithms—can diverge under both off-policy sampling and nonlinear function approximation. Many algorithms have been developed for off-policy value estimation based on the linear mean squared projected Bellman error ($\overline{\text{PBE}}$) and are sound under linear function approximation. Extending these methods to the nonlinear case has been largely unsuccessful. Recently, several methods have been introduced that approximate a different objective—the mean-squared Bellman error ($\overline{\text{BE}}$)—which naturally facilitate nonlinear approximation. In this work, we build on these insights and introduce a new generalized $\overline{\text{PBE}}$ that extends the linear $\overline{\text{PBE}}$ to the nonlinear setting. We show how this generalized objective unifies previous work and obtain new bounds for the value error of the solutions of the generalized objective. We derive an easy-to-use, but sound, algorithm to minimize the generalized objective, and show that it is more stable across runs, is less sensitive to hyperparameters, and performs favorably across four control domains with neural network function approximation.

Keywords: Off-policy learning, Temporal difference learning, Reinforcement learning

1. Introduction

Value functions play a central role in reinforcement learning: value-based methods act greedily towards a learned action value function; policy gradient methods often employ a critic to estimate the value of the policy in order to reduce variance; model-based methods often employ learned values for tree search, model-based value prediction, or simulation based planning to update value estimates from hypothetical experience. Fundamental improvements in algorithms for learning value functions can have a significant impact in a variety of problem settings; as such, many algorithms have been developed to improve learning values functions. This includes a variety of variance reduction improvements for off-policy temporal difference algorithms (Precup et al., 2000; Munos et al., 2016; Mahmood et al., 2017); gradient TD methods with linear function approximation (Sutton et al., 2009; Mahadevan et al., 2014; Liu et al., 2016; Ghiassian et al., 2020) and nonlinear function

approximation (Maei et al., 2009); and algorithms using approximations to the mean squared Bellman error ($\overline{\text{BE}}$) (Dai et al., 2017, 2018; Feng et al., 2019).

The most widely used value function learning algorithm is also one of the oldest: temporal difference (TD) learning. The continued popularity of TD stems from both (a) the algorithms simplicity and empirical performance, and (b) the lack of technical tools required to improve it. TD, however, does not follow the gradient of any known objective function (Baird, 1995; Antos et al., 2008), and without a clear objective for TD, it was difficult to extend the algorithm. Related TD-like Residual gradient algorithms directly optimize the $\overline{\text{BE}}$, but suffers from the double sampling problem (Baird, 1995; Scherrer, 2010). Without a strategy to optimize the $\overline{\text{BE}}$, in the absence of a simulator, it was difficult to pursue the $\overline{\text{BE}}$ as an alternative.

Value function learning algorithms that converge under general conditions fall into two camps defined by the objectives they optimize. The first is the formalization of the objective underlying TD—the mean squared projected Bellman error ($\overline{\text{PBE}}$) (Antos et al., 2008; Sutton et al., 2009)—which projects the Bellman error into the space spanned by the function approximator. Several algorithms were introduced to minimize the $\overline{\text{PBE}}$, most of which built on the originally introduced variants: GTD2 and TD with gradient corrections (TDC) (Sutton et al., 2009). Most of these algorithm, however, are limited to linear function approximation because the $\overline{\text{PBE}}$ is defined only for the linear case with the exception of nonlinear GTD (Maei et al., 2009), which used a locally linear approximation for the projection. The gradient computation, however, requires a Hessian-vector product and has not been widely adopted. As such, the extension of the linear $\overline{\text{PBE}}$ for nonlinear function approximation remains open.

The second camp of methods pursue the $\overline{\text{BE}}$, which is naturally defined for nonlinear function approximation. The idea is to use a conjugate form for the $\overline{\text{BE}}$ (Dai et al., 2017), to reformulate it into a saddlepoint problem with an auxiliary variable, removing the double sampling problem because the resulting saddlepoint uses an estimator in place of one of the samples. The SBEED algorithm (Dai et al., 2018) later extended the conjugate $\overline{\text{BE}}$ to the control case by using a smoothed Bellman optimality operator and parameterizing both the policy and value function estimates.

In this work, we bring together these threads of research and provide a novel objective for nonlinear value function estimation in both prediction and control. In particular, we introduce a generalized $\overline{\text{PBE}}$, by using the conjugate reformulation of the $\overline{\text{BE}}$. To understand the role of objectives in off-policy value estimation, it is useful to first understand the algorithmic development in off-policy learning that led to several of these objectives. We first lay out this development, and then describe the contributions in this work.

1.1 A Short History of Off-policy Temporal Difference Learning

The story of off-policy learning begins with Q-learning (Watkins, 1989). Off-policy learning, in some sense, allows side-stepping the exploration-exploitation tradeoff: the agent makes use of an independent exploration policy to select actions while learning the value function for the optimal policy. The exploration policy need not maximize reward, but can instead select actions in order to generate data that improves the optimal policy through learning. Ultimately, the full potential of Q-learning—and this ability to learn about one policy from

a data generated by a totally different exploration—proved limited. Baird’s well-known counterexample (Baird, 1995) provided a clear illustration of how, under function approximation, the weights learned by Q-learning can become unstable.¹ Baird’s counterexample highlights that divergence can occur when updating off-policy with function approximation and with bootstrapping (as in temporal difference (TD) learning); even when learning the value function of a fixed target policy from a fixed data-generating policy.

The instability of TD methods is caused by how we correct the updates to the value function to account for the potential mismatch between the target and exploration policies. Off-policy training involves estimating the total expected future reward (the value function) that would be observed while selecting actions according to the target policy with training data (states, actions, and rewards) generated while selecting actions according to an exploration policy. One approach to account for the differences between the data produced by these two policies is based on using importance sampling corrections: scaling the update to the value function based on the agreement between the target and exploration policy at the current state. If the target and exploration policy would always select the same action in a state, then they completely agree. Alternatively, if they never take the same action in a state they completely disagree. More generally, there can be degrees of agreement. We call this approach *posterior corrections* because the corrections account for the mismatch between policies ignoring the history of interaction up to the current time step—it does not matter what the exploration policy has done in the past.

Another approach, called *prior corrections*, uses the history of agreement between the exploration and target policy in the update. The likelihood that the trajectory could have occurred under the target policy is used to scale the update. The most extreme version of prior corrections uses the trajectory of experience from the beginning of time, corresponding to what has sometimes been referred to as the *alternative life* framework. Prior and posterior corrections can be combined to achieve stable off-policy TD updates (Precup et al., 2000), though finite variance of the updates cannot be guaranteed (Precup et al., 2001).

Learning about many different policies in parallel has long been a primary motivation for off-policy learning, and this usage suggested that perhaps prior corrections are not essential. Several approaches require learning many value functions or policies in parallel, including approaches based on option models (Sutton et al., 1999b), predictive representations of state (Littman and Sutton, 2002; Tanner and Sutton, 2005; Sutton et al., 2011), and auxiliary tasks (Jaderberg et al., 2016). In a parallel learning setting, it is natural to estimate the future reward achieved by following each target policy until termination from the states encountered during training—the value of taking *excursions* from the behavior policy.

This excursion view of off-policy learning lead to the development of an new objective function for off-policy TD learning called the mean squared projected Bellman error ($\overline{\text{PBE}}$). The resultant family of *Gradient*-TD methods that optimize the $\overline{\text{PBE}}$ use posterior corrections via importance sampling and are guaranteed to be stable under function approximation (Sutton et al., 2009; Maei, 2011). This new excursion objective has the same fixed point as TD, thus Gradient-TD methods converge to the same solution in the cases for which TD converges. The major critiques of these methods are (1) the additional complexity due

1. The action-value star MDP can be found in the errata of Baird’s paper (Baird, 1995).

to a second set of learned parameters, and (2) the variance due to importance sampling corrections.

A wide array of algorithms arose to address these limitations and improve sample efficiency. The strategies include (1) using a saddlepoint reformulation that facilitates the use of optimization accelerations (Liu et al., 2016; Du et al., 2017; Liu et al., 2015) and a finite sample analysis (Touati et al., 2018), (2) hybrid TD methods that behave like TD when sampling is on-policy (Hackman, 2013; White and White, 2016), and (3) variance reduction methods for posterior corrections, using different eligibility trace parameters (Precup et al., 2000; Munos et al., 2016; Wang et al., 2016; Mahmood et al., 2017).

The second major family of off-policy methods revisits the idea of using prior corrections. The idea is to incorporate prior corrections only from the beginning of the excursion. In this way, the values of states that are more often visited under the target policy are emphasized, but the high variance of full prior corrections—to the beginning of the episode—is avoided. The *Emphatic* TD(λ) algorithm is based on this idea (Sutton et al., 2016) and was later extended to reduce the variance of the emphasis weights (Hallak et al., 2016). These Emphatic TD methods are guaranteed stable under both on-policy and off-policy sampling with linear function approximation (Sutton et al., 2016; Yu, 2015; Hallak et al., 2016).

Practitioners often bemoan both the complexity and poor performance of these sound off-policy algorithms, favouring instead to combine off-policy TD with heuristics to reduce variance. Many of the algorithms discussed above have not been developed for the case of nonlinear function approximation. On the other hand, off-policy updating is used pervasively in deep learning architectures for learning from a replay buffer, demonstrations, and trajectories generated by other agents (i.e., asynchronous architectures). In order to get the benefits of off-policy updates and mitigate variance we can truncate the corrections: the Retrace algorithm does exactly this (Munos et al., 2016; Wang et al., 2016). Retrace and the much older Tree Backup algorithm can be viewed as adapting the eligibility trace to reduce variance, similar to the ABQ algorithm (Mahmood et al., 2017). These three methods can be extended to utilize gradient corrections to achieve stability, but this would only yield another batch of theoretically sound but practically ignored methods. In this paper, we explore an alternative path: extending gradient TD methods to both control and nonlinear function approximation through the development of a new objective function.

1.2 Revisiting the Objective for Value Estimation

Central to all this development has been the definition of a precise objective: the linear PBE. The proliferation of algorithms, however, has resulted in a focus on comparing and discussing specific algorithms (Dann et al., 2014; White and White, 2016; Touati et al., 2018). This breeds confusion in exactly how the algorithms are related and what objective they are attempting to optimize. By separating the objective and the optimization strategies, we can more systematically understand the differences in solutions under objectives and how effective the optimization strategies are at reaching those solutions.

Two key questions emerge about the specification of the objective: the form of the Bellman error and the weighting across states. The first question is really a revisitation of the long-standing question about whether we should use the $\overline{\text{BE}}$ or the $\overline{\text{PBE}}$. In terms of solution quality, it is mixed: the $\overline{\text{BE}}$ avoids some counterexamples that exist for the

$\overline{\text{PBE}}$ (Scherrer, 2010), however the $\overline{\text{PBE}}$ often produces a better solution (Dann et al., 2014; Scherrer, 2010; Sutton and Barto, 2018). Further, the $\overline{\text{BE}}$ has been shown to have an identifiability problem (Sutton and Barto, 2018). Though the evidence comparing the $\overline{\text{BE}}$ and $\overline{\text{PBE}}$ is inconclusive, the $\overline{\text{PBE}}$ has been the default choice because many algorithms optimize it. The $\overline{\text{BE}}$, on the other hand, is typically avoided due to the double sampling problem, where it is unclear how to obtain an unbiased sample of the gradient without a simulator.

Recently, however, this technical challenge has been overcome with the introduction of a saddlepoint form for the $\overline{\text{BE}}$ (Dai et al., 2017, 2018; Feng et al., 2019). The resulting algorithms are similar to the saddlepoint algorithms for the linear $\overline{\text{PBE}}$: a second estimator is used to estimate a part of the gradient. The $\overline{\text{BE}}$ is particularly alluring, as it equally applies to the linear and nonlinear value estimation settings. The $\overline{\text{PBE}}$, on the other hand, was defined for the linear setting, due to the difficulty in computing the projection operator for the nonlinear setting. These potential advantages, as well as a viable strategy for optimizing the $\overline{\text{BE}}$, motivates the utility of answering which of these objectives might be preferable.

The second question is about the weighting on states in the objective, which must be specified for either the $\overline{\text{BE}}$ or $\overline{\text{PBE}}$. The importance of the weighting on states has been well-recognized for many years, and is in fact the reason TD diverges on Baird’s counterexample: using the stationary distribution of the behavior policy, rather than the target policy, results in an iterative update that is no longer a contraction. The emphatic algorithms were introduced to adjust this weighting to ensure convergence.

This reweighting, however, is not only about convergence; it also changes the fixed point and the quality of the solution. There has been some work investigating the impact of the state weighting on the optimal solution, not just on the behavior of the updates themselves. The most stark result is a simple example where using the solution to the $\overline{\text{PBE}}$ can result in an arbitrarily poor mean squared value error ($\overline{\text{VE}}$) (Kolter, 2011). Several later results extended on-policy bounds on the $\overline{\text{VE}}$ to the off-policy setting, showing that $\overline{\text{VE}}$ could be bounded in the off-policy setting using careful choices on the state weighting—namely using state weightings given by Emphatic TD (Hallak et al., 2016; White, 2017). Despite these insights, the role of the weighting on the quality of the solution in practice remains open. A natural question is: how do we choose the state weighting in the objective, and how much does it matter?

1.3 Contributions

In this work, we focus on the question: what objective should we use for off-policy value estimation? We first summarize many existing off-policy algorithms as optimizing the linear $\overline{\text{PBE}}$ in different ways and in some cases with different state weightings. This summary separates the optimization strategy from the definition of the objective, allowing us to move away from specific algorithms to understanding the differences in fixed points obtained under the different objectives. We then propose a generalized $\overline{\text{PBE}}$ that uses a generalized projection operator that both extends the $\overline{\text{PBE}}$ to the nonlinear setting and unifies the $\overline{\text{BE}}$ and linear $\overline{\text{PBE}}$ under one objective. Using these insights, we provide the following specific contributions:

1. We show how the generalized $\overline{\text{PBE}}$ helps resolve the non-identifiability of the $\overline{\text{BE}}$, where a particular projection in the generalized $\overline{\text{PBE}}$ provides an Identifiable $\overline{\text{BE}}$.
2. We highlight the role of the state weighting in this generalized objective, both extending theoretical results bounding the $\overline{\text{VE}}$ and empirically showing that the emphatic weighting can significantly improve the quality of the solution.
3. We show that these insights also extend to control by defining an objective for learning action values with (soft) maximal operators. We use this objective to derive a sound gradient variant of Q-learning.
4. We exploit the connection to the linear $\overline{\text{PBE}}$ to develop a more effective algorithm for the generalized $\overline{\text{PBE}}$ using gradient corrections rather than the saddlepoint update.
5. Finally, we demonstrate the utility of these prediction and control algorithms in several small benchmark domains.

2. Problem Definition and Background

We consider the problem of learning the value function for a given policy under the Markov Decision Process (MDP) formalism. The agent interacts with the environment over a sequence of discrete time steps, $t = 1, 2, 3, \dots$. On each time step, the agent observes a partial summary of the state $S_t \in \mathcal{S}$ and selects an action $A_t \in \mathcal{A}$. In response, the environment transitions to a new state S_{t+1} according to transition function $P(S_{t+1}|S_t, A_t)$ and emits a scalar reward $R_{t+1} \in \mathcal{R}$. The agent selects actions according to a stochastic, stationary *target policy* $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$.

We study the problem of *value estimation*: the computation or estimation of the expected discounted sum of future rewards for policy π from every state. The *return* at time t , denoted $G_t \in \mathbb{R}$, is defined as the discounted sum of future rewards. The discount factor can be variable, dependent on the entire transition: $\gamma : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, with $\gamma_{t+1} \stackrel{\text{def}}{=} \gamma(S_t, A_t, S_{t+1})$. The return is defined as

$$\begin{aligned} G_t &\stackrel{\text{def}}{=} R_{t+1} + \gamma_{t+1}R_{t+2} + \gamma_{t+1}\gamma_{t+2}R_{t+3} + \gamma_{t+1}\gamma_{t+2}\gamma_{t+3}R_{t+4} + \dots \\ &= R_{t+1} + \gamma_{t+1}G_{t+1}. \end{aligned}$$

When γ_t is constant, γ_c , we get the familiar return $G_t = R_{t+1} + \gamma_c R_{t+2} + \gamma_c^2 R_{t+3} + \dots$. Otherwise, variable γ_t can discount per transition, including encoding termination when it is set to zero. This generalization ensures we can discuss episodic problems without introducing absorbing states (White, 2017). It also enables the derivations and theory to apply to both the continuing and episodic settings. The *value function* $v : \mathcal{S} \rightarrow \mathbb{R}$ maps each state to the expected return under policy π starting from that state

$$v_\pi(s) \stackrel{\text{def}}{=} \mathbb{E}_\pi[G_t | S_t = s] \text{ , for all } s \in \mathcal{S} \tag{1}$$

where the expectation operator $\mathbb{E}_\pi[\cdot]$ reflects that the distribution over future actions is given by π , to distinguish from a potentially different behavior policy.

In this paper, we are interested in problems where the value of each state cannot be stored in a table; instead we must approximate the value with a parameterized function. The approximate value function $\hat{v}_w(s_t)$ can have arbitrary form as long as it is everywhere

differentiable with respect to the weights $\mathbf{w} \in \mathbb{R}^d$. Typically, the number of components in \mathbf{w} is much fewer than the number of possible states ($d \ll |\mathcal{S}|$), thus \hat{v} will generalize values across many states in \mathcal{S} . An important special case is when the approximate value function is linear in the parameters and in features of the state. In particular, the current state S_t is converted into feature vector $\mathbf{x}_t \in \mathbb{R}^d$ by some fixed mapping $\mathbf{x} : \mathcal{S} \rightarrow \mathbb{R}^d$. The value of the state can then be approximated with an inner product: $\hat{v}_{\mathbf{w}}(s_t) = \mathbf{w}^\top \mathbf{x}_t \approx v_\pi(s_t)$. Another typical parameterization for $\hat{v}_{\mathbf{w}}(s_t)$ is a neural network where \mathbf{w} consists of all the weights in the network. We refer to \mathbf{w} exclusively as the *weights*, or weight vector, and reserve the word *parameter* for variables like the discount-rate and stepsize parameters.

We first describe how to learn this value function for the on-policy setting, where the behavior policy equals the target policy. Temporal difference learning (Sutton, 1988) is perhaps the best known and most successful approach for estimating \hat{v} directly from samples generated while interacting with the environment. Instead of waiting until the end of a trajectory to update the value of each state, the TD(λ) algorithm adjusts its current estimate of the weights toward the difference between the discounted estimate of the value in the next state and the estimated value of the current state plus the reward along the way:

$$\delta_t \stackrel{\text{def}}{=} \delta(S_t, A_t, S_{t+1}) \stackrel{\text{def}}{=} R_{t+1} + \gamma_{t+1} \hat{v}_{\mathbf{w}}(S_{t+1}) - \hat{v}_{\mathbf{w}}(S_t). \quad (2)$$

We use the value function’s own estimate of future reward as a placeholder for the future rewards defining G_t that are not available on time-step $t + 1$. In addition, the TD(λ) algorithm also maintains an eligibility trace vector $\mathbf{z}_t \in \mathbb{R}^d$ that stores a fading trace of recent feature activations. The components of \mathbf{w}_t are updated on each step proportional to the magnitude of the trace vector. This simple scheme allows update information to propagate impacting the value estimates for previously encountered states.

The update equations for TD(λ) are straightforward:

$$\begin{aligned} \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t \\ \mathbf{z}_t &\leftarrow \gamma_t \lambda \mathbf{z}_{t-1} + \nabla \hat{v}_{\mathbf{w}}(S_t), \end{aligned}$$

where $\alpha \in \mathbb{R}$ is the scalar stepsize parameter that controls the speed of learning and $\lambda \in \mathbb{R}$ controls the length of the eligibility trace. Under linear function approximation, intermediate values of λ between zero and one often perform best. TD(λ) is only sound for the linear function approximation setting, though TD(0) is often used outside the linear setting and often obtains good performance.

These updates need to be modified for the *off-policy case*, where the agent selects actions according to a *behavior policy* $b : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ that is different from the target policy. The value function for target policy π is updated using experience generated from a behavior policy that is *off*, away, or distant from the target policy. For example, consider the most well-known off-policy algorithm, Q-learning. The target policy might be the one that maximizes future discounted reward, while the behavior is nearly identical to the target policy, but instead selects an exploratory action on each time step with some small probability. More generally, the target and behavior policies need not be so closely coupled. The main requirement linking these two policies is that the behavior policy *covers* the actions selected by the target policy in each state visited by b , that is: $b(a|s) > 0$ for all states and actions in which $\pi(a|s) > 0$.

3. Off-policy Corrections and the Connection to State Weightings

The key problem in off-policy learning is to estimate the value function for the target policy conditioned on samples produced by actions selected according to the behavior policy. This is an instance of the problem of estimating an expected value under some target distribution from samples generated by some other behavior distribution. In statistics, this problem can be addressed with importance sampling, and most methods of off-policy reinforcement learning use such corrections.

There are two distributions that we could consider correcting: the distribution over actions, given the state, and the distribution over states. When observing a transition (S, A, S', R) generated by taking the action according to $b(\cdot|S)$, we can correct the update for that transition so that, in expectation, we update as if actions were taken according to $\pi(\cdot|S)$. However, these updates are still different than if we evaluated π on-policy, because the frequency of visiting state S under b will be different than under π . All methods correct for the distribution over actions (posterior corrections), given the state, but several methods also correct for the distribution over states (prior corrections).

In this section, we discuss the difference in the underlying objective for these updates. We give a brief introduction to prior and posterior corrections, and provide a more in-depth example of the differences in their updates in Appendix A for readers interested in a more intuitive explanation. We show that the underlying objective differs in terms of the state weighting and discuss how emphatic weightings further modify this state weighting in the objective.

3.1 Posterior Corrections

The most common approach to developing sound off-policy TD algorithms makes use of posterior corrections based on importance sampling. One of the simplest examples of this approach is *Off-policy TD*(λ). The procedure is easy to implement and requires constant computation per time step, given knowledge of both the target and behavior policies. On the transition from S_t to S_{t+1} via action A_t , we compute the ratio between π and b :

$$\rho_t \stackrel{\text{def}}{=} \rho(A_t|S_t) \stackrel{\text{def}}{=} \frac{\pi(A_t|S_t)}{b(A_t|S_t)}. \quad (3)$$

These importance sampling corrections simply weight the eligibility trace update:

$$\begin{aligned} \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t^\rho \\ \mathbf{z}_t^\rho &\leftarrow \rho_t (\gamma_t \lambda \mathbf{z}_{t-1}^\rho + \mathbf{x}_t), \end{aligned} \quad (4)$$

where δ_t is defined in Equation (2). This way of correcting the sample updates ensures that the approximate value function \hat{v} estimates the expected value of the return as if the actions were selected according to π . Posterior correction methods use the target policy probabilities for the selected action to correct the update to the value of state S_t using only the data from time step t onward. Values of π from time steps prior to t have no impact on the correction.

Many modern off-policy prediction methods use some form of posterior corrections; including the Gradient-TD methods, Tree Backup(λ), V-trace(λ), and Emphatic TD(λ). In fact, all off-policy prediction methods with stability guarantees make use of posterior

corrections via importance sampling. Only correcting the action distribution, however, does not necessarily provide stable updates, and Off-policy TD(λ) is not guaranteed to converge (Baird, 1995). To obtain stable Off-policy TD(λ) updates, we need to consider corrections to the state distribution as well, as we discuss next.

3.2 Prior Corrections

We can also consider correcting for the differences between the target and behavior policy by using the agreement between the two over a trajectory of experience. *Prior correction* methods keep track of the product of either $\prod_{k=1}^t \pi(A_k|S_k)$ or $\prod_{k=1}^t \frac{\pi(A_k|S_k)}{b(A_k|S_k)}$, and correct the update to the value of S_t using the current value of the product. Therefore, the value of S_t is only updated if the product is not zero, meaning that the behavior policy never selected an action for which $\pi(A_k|S_k)$ was zero—the behavior never completely deviated from the target policy.

To appreciate the consequences of incorporating these prior corrections into the TD update consider a state-value variant of Precup et al. (2000) Off-policy TD(λ) algorithm:

$$\begin{aligned} \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t^\rho \\ \mathbf{z}_t^\rho &\leftarrow \rho_t \left(\gamma_t \lambda \mathbf{z}_{t-1} + \prod_{k=1}^{t-1} \rho_k \mathbf{x}_k \right) \end{aligned} \tag{5}$$

where $\mathbf{z}_0^\rho = \mathbf{0}$. We refer to the above algorithm as *Alternative-life TD(λ)*, as the product of importance sampling ratios in the trace simulate an agent experiencing an entirely different trajectory from the beginning of the episode—an alternative life. Note that the trace is always reinitialized at the start of the episode.

The Alternative-life TD(λ) algorithm has been shown to converge under linear function approximation, but in practice exhibits unacceptable variance (Precup et al., 2001). The Emphatic TD(λ) algorithm, on the other hand, provides an alternative form for the prior corrections that is lower variance but still guarantees convergence. To more clearly explain why, next we will discuss how different prior corrections account for different weightings in optimizing the mean squared Projected Bellman Error ($\overline{\text{PBE}}$).

3.3 The Linear $\overline{\text{PBE}}$ under Posterior and Prior Corrections

In this section, we describe how different forms of prior corrections correspond to optimizing similar objectives, but with different weightings over the state. This section introduces the notation required to explain the many algorithms that optimize the linear $\overline{\text{PBE}}$ and clarifies convergence properties of algorithms, including which algorithms converge and to which fixed point. We start with the linear $\overline{\text{PBE}}$ as most algorithms have been designed to optimize it. In the next section, we extend beyond the linear setting to discuss the generalized $\overline{\text{PBE}}$.

We begin by considering a simplified setting with $\lambda = 0$ and a simplified variant of the linear $\overline{\text{PBE}}$, called the NEU (norm of the expected TD update (Sutton et al., 2009))

$$\text{NEU}(\mathbf{w}) = \left\| \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi [\delta(S, A, S') \mathbf{x}(S) \mid S = s] \right\|_2^2, \tag{6}$$

where $d : \mathcal{S} \rightarrow [0, \infty)$ is a positive weighting on the states. Note the transition (S, A, S') is random as is the TD-error. Equation 6 does not commit to a particular sampling strategy. If the data is sampled on-policy, then $d = d_\pi$, where $d_\pi : \mathcal{S} \rightarrow [0, 1]$ is the stationary distribution for π which represents the state visitation frequency under behavior π in the MDP. If the data is sampled off-policy, then the objective might instead be weighted by the state visitation frequency under b , i.e., $d = d_b$.

We first consider how to sample the NEU for a given a state. The behavior selects actions in each state s , so the update $\delta_t \mathbf{x}_t$ must be corrected for the action selection probabilities of π in state s , namely a posterior correction:

$$\begin{aligned}
 \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = s] &= \sum_{a \in \mathcal{A}} \pi(a|s) \sum_{s' \in \mathcal{S}} P(s'|s, a) \delta(s, a, s') \mathbf{x}(s) \\
 &= \sum_{a \in \mathcal{A}} \frac{b(a|s)}{b(a|s)} \pi(a|s) \sum_{s' \in \mathcal{S}} P(s'|s, a) \delta(s, a, s') \mathbf{x}(s) \\
 &= \sum_{a \in \mathcal{A}} b(a|s) \sum_{s' \in \mathcal{S}} P(s'|s, a) \frac{\pi(a|s)}{b(a|s)} \delta(s, a, s') \mathbf{x}(s) \\
 &= \mathbb{E}_b[\rho(A_t|S_t) \delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = s]. \tag{7}
 \end{aligned}$$

Therefore, the update $\rho_t \delta_t \mathbf{x}_t$ provides an unbiased sample of the desired expected update $\mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = s]$. Most off-policy methods (except Q-learning and Tree-backup) use these posterior corrections.

We can also adjust the state probabilities from d_b to d_π , using prior corrections. Consider the expected update using prior corrections starting in s_0 and taking two steps following b :

$$\begin{aligned}
 &\mathbb{E}_b[\rho_0 \rho_1 \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = S_2] \mid S_0 = s_0] \\
 &= \mathbb{E}_b \left[\rho_0 \sum_{a_1 \in \mathcal{A}} b(a_1|S_1) \sum_{s_2 \in \mathcal{S}} P(s_2|S_1, a_1) \rho(a_1|S_1) \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = s_2] \mid S_0 = s_0 \right] \\
 &= \mathbb{E}_b \left[\rho_0 \sum_{a_1 \in \mathcal{A}} \pi(a_1|S_1) P(s_2|S_1, a_1) \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = s_2] \mid S_0 = s_0 \right] \\
 &= \mathbb{E}_b[\rho_0 \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_{t-1} = S_1] \mid S_0 = s_0] \\
 &= \sum_{a_0 \in \mathcal{A}} \pi(s_0, a_0) \sum_{s_1 \in \mathcal{S}} P(s_1|s_0, a_0) \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_{t-1} = s_1] \\
 &= \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_0 = s_0].
 \end{aligned}$$

More generally, we get

$$\mathbb{E}_b \left[\rho_1 \dots \rho_{t-1} \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_t = s] \mid S_0 = s_0 \right] = \mathbb{E}_\pi[\delta(S_t, A_t, S_{t+1})\mathbf{x}(S_t) \mid S_0 = s_0].$$

These corrections adjust the probability of the sequence from the beginning of the episode as if policy π had taken actions A_1, \dots, A_{t-1} to get to S_t , from which we do the TD(λ) update.

A natural question is which objective should be preferred: the alternative-life ($d \propto d_\pi$) or the excursions objective ($d \propto d_b$). As with all choices for objectives, there is no crisp

answer. The alternative-life objective is difficult to optimize because prior corrections can become very large or zero—causing data to be discarded—and is thus high variance. There is some work that directly estimates $d_\pi(s)/d_b(s)$ and pre-multiplies to correct the update (Hallak and Mannor, 2017; Liu et al., 2020); this approach, however, requires estimating the visitation distributions. On the other hand, the fixed-point solution to the excursion objective can be arbitrarily poor compared with the best value function in the function approximation class if there is a significant mismatch between the behavior and target policy (Kolter, 2011). Better solution accuracy can be achieved using an excursion’s weighting that includes d_b , but additionally reweights to make the states distribution closer to d_π as is done with Emphatic TD(λ). We discuss this alternative weighting in the next section.

The above discussion focused on a simplified variant of the $\overline{\text{PBE}}$ with $\lambda = 0$, but the intuition is the same for the $\overline{\text{PBE}}$ with $\lambda > 0$. To simplify notation we introduce a conditional expectation operator:

$$\mathbb{E}_d[Y] = \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi[Y \mid S = s].$$

We can now define

$$\begin{aligned} \mathbf{C} &\stackrel{\text{def}}{=} \mathbb{E}_d[\mathbf{x}(S)\mathbf{x}(S)^\top] \\ \mathbf{A} &\stackrel{\text{def}}{=} -\mathbb{E}_d[(\gamma(S, A, S')\mathbf{x}(S') - \mathbf{x}(S))\mathbf{z}(S)^\top] \\ \mathbf{b} &\stackrel{\text{def}}{=} \mathbb{E}_d[R(S, A, S')\mathbf{z}(S)^\top] \end{aligned}$$

where the expected eligibility trace $\mathbf{z}(S) \in \mathbb{R}^k$ is defined recursively $\mathbf{z}(s) \stackrel{\text{def}}{=} \mathbf{x}(s) + \lambda \mathbb{E}_\pi[\gamma(S_{t-1}, A_{t-1}, S_t)\mathbf{z}(S_{t-1}) \mid S_t = s]$. We can write the TD(λ) fixed point residual as:

$$\mathbb{E}_d[\delta(S, A, S')\mathbf{z}(S)] = -\mathbf{A}\mathbf{w} + \mathbf{b} \tag{8}$$

where $\mathbb{E}_d[\delta(S, A, S')\mathbf{z}(S)] = \mathbf{0}$ at the fixed point solution for on-policy TD(λ). The linear $\overline{\text{PBE}}$ can then be defined given the definition above:

$$\text{linear } \overline{\text{PBE}}(\mathbf{w}) \stackrel{\text{def}}{=} (-\mathbf{A}\mathbf{w} + \mathbf{b})^\top \mathbf{C}^{-1} (-\mathbf{A}\mathbf{w} + \mathbf{b}) = \|-\mathbf{A}\mathbf{w} + \mathbf{b}\|_{\mathbf{C}^{-1}}^2. \tag{9}$$

The only difference compared with the NEU is the weighted ℓ_2 norm, weighted by \mathbf{C}^{-1} , instead of simply $\|-\mathbf{A}\mathbf{w} + \mathbf{b}\|_2^2$.

Notation Remark: From the rest of the paper, we will not explicitly write the random variables as functions of (S, A, S') . For example, we will use $\mathbb{E}_\pi[\delta \mid S = s]$ to mean $\mathbb{E}_\pi[\delta(S, A, S') \mid S = s]$ and $\mathbb{E}_\pi[\gamma \mid S = s]$ to mean $\mathbb{E}_\pi[\gamma(S, A, S') \mid S = s]$.

3.4 Emphatic Weightings as Prior Corrections

Emphatic Temporal Difference learning, ETD(λ), provides an alternative strategy for obtaining stability under off-policy sampling without computing the gradient of the linear $\overline{\text{PBE}}$. The key idea is to incorporate some prior corrections so that the weighting d results in a positive definite matrix \mathbf{A} . Given such an \mathbf{A} , the TD(λ) update—a semi-gradient algorithm—can be shown to converge. Importantly, this allows for a stable off-policy algorithm with only a single set of learned weights. Gradient-TD methods, on the other hand, use two stepsize parameters and two weight vectors to achieve stability.

Emphatic TD(λ) or ETD(λ) minimizes a variant of the linear $\overline{\text{PBE}}$ defined in Equation 9, where the weighting d is defined based on the *followon* weighting. The followon reflects (discounted) state visitation under the target policy when doing excursions from the behavior: starting from states sampled according to d_b . The followon is defined as

$$f(s_t) \stackrel{\text{def}}{=} d_b(s_t) + \sum_{s_{t-1}, a_{t-1}} d_b(s_{t-1}) \pi(a_{t-1}|s_{t-1}) P(s_t|s_{t-1}, a_{t-1}) \gamma(s_{t-1}, a_{t-1}, s_t) + \dots \quad (10)$$

The emphatic weighting is $m(s_t) = d_b(s_t)\lambda + (1 - \lambda)f(s_t)$. This is the state weighting ETD(λ) uses in the linear $\overline{\text{PBE}}$ in Equation 9, setting $d(s) = m(s)$. ETD(λ) uses updates

$$\begin{aligned} F_t &\leftarrow \rho_{t-1} \gamma_t F_{t-1} + 1 && \text{with } M_t = \lambda_t + (1 - \lambda_t) F_t \\ \mathbf{z}_t^\rho &\leftarrow \rho_t (\gamma_t \lambda \mathbf{z}_{t-1}^\rho + M_t \mathbf{x}_t) \\ \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t^\rho, \end{aligned}$$

with $F_0 = 1$ and $\mathbf{z}_0^\rho = \mathbf{0}$.

To gain some intuition for this weighting, consider the trace update with alternative-life TD(0) and ETD(0) with constant γ . For ETD(0), $M_t = F_t = \sum_{j=0}^t \gamma^j \prod_{i=1}^j \rho_{t-i}$, giving $\mathbf{z}_t^\rho \leftarrow \rho_t \left(\gamma \lambda \mathbf{z}_{t-1}^\rho + \sum_{j=0}^t \gamma^j \prod_{i=1}^j \rho_{t-i} \mathbf{x}_t \right)$. The trace for Alternative-life TD(0) (see Equation 5) is $\mathbf{z}_t^\rho \leftarrow \rho_t (\gamma_t \lambda \mathbf{z}_{t-1}^\rho + \prod_{i=1}^t \rho_i \mathbf{x}_t)$. Both adjust the weighting on \mathbf{x}_t to correct the state distributions. Alternative-Life TD more aggressively downweights states that would not have been visited under the target policy, because it only has a product, whereas ETD uses a sum over all t up to that point.

Emphatic TD(λ) has strong convergence guarantees in the case of linear function approximation. The ETD(λ) under off-policy training has been shown to converge in expectation using the same expected update analysis used to show that TD(λ) converges under on-policy training. Later, Yu (2015) extended this result to show that ETD(λ) converges with probability one.

This weighting also resolves the issues raised by Kolter’s example (Kolter, 2011). Kolter’s example demonstrated that for a particular choice of π and b , the solution to the linear $\overline{\text{PBE}}$ could result in arbitrarily bad value error compared with the best possible approximation in the function class. In other words, even if the true value function can be well approximated by the function class, the off-policy fixed point from the linear $\overline{\text{PBE}}$ with weighting $d = d_b$ can result in an arbitrarily poor approximation to the values. In Section 6, we explain why the fixed points of the linear $\overline{\text{PBE}}$ with the emphatic weighting do not suffer from this problem, expanding on (Hallak et al., 2016, Corollary 1) and (White, 2017, Theorem 1).

3.5 Broadening the Scope of Weightings

To determine which weightings to consider, we need to understand the role of the weighting. There are actually two possible roles. The first is to specify states of interest: determine the relative importance of a state for the accuracy of our value estimates compared to the true values. This provides the weighting in the value error: $\overline{\text{VE}}(\mathbf{w}) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} d(s) (\hat{v}_{\mathbf{w}}(s) - v_\pi(s))^2$. The second is the choice of weighting in our objective, such as the linear $\overline{\text{PBE}}$, which is a surrogate for the value error.

For the first question, we need to determine the relative importance of states. The choice depends on the purpose for the value estimation. For example, if the policy is being evaluated for deployment in an episodic problem, a common choice is to put all weight on the set of start states (Sutton et al., 1999a; Bottou et al., 2013) because it reflects the expected return in each episode. On the other hand, if many value functions are learned in parallel—as in predictive representations (Sutton et al., 2011; White, 2015) or auxiliary tasks (Jaderberg et al., 2016)—it may be better if the predictions are accurate in many states. It is also possible that values from some states might be queried much more often, or that states might correspond to important catastrophic events from which it is important to have accurate predictions for accurate decision-making.

Overall, the choice of d is subjective. Once we make this choice for our evaluation objective, we can ask the second question: which optimization objectives and what weightings are most effective for minimizing our evaluation objective? It is not obvious that the minimum of the linear $\overline{\text{PBE}}$ with weighting d_b provides the best solution to the $\overline{\text{VE}}$ with weighting d_b , for example. In fact, we know that the linear $\overline{\text{PBE}}$ with $d = d_b$ suffers from a counterexample (Kolter, 2011), whereas using the emphatic weighting in the linear $\overline{\text{PBE}}$ provides an upper bound on the $\overline{\text{VE}}$ under weighting d_b . We discuss the potential utility of using a different weighting for the objective than the desired weighting in the $\overline{\text{VE}}$ in Section 6.

4. Broadening the Scope of Objectives

In this section, we discuss how to generalize the linear $\overline{\text{PBE}}$, to obtain the generalized $\overline{\text{PBE}}$. This objective allows for nonlinear value estimation and unifies the $\overline{\text{BE}}$ and linear $\overline{\text{PBE}}$ under one objective.

4.1 An Overview of Existing Objectives

Let us start by discussing the standard evaluation objective used for policy evaluation: the mean squared value error ($\overline{\text{VE}}$)

$$\overline{\text{VE}}(\mathbf{w}) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} d(s) (\hat{v}_{\mathbf{w}}(s) - v_{\pi}(s))^2. \quad (11)$$

The approximation $\hat{v}_{\mathbf{w}}(s)$ is penalized more heavily for inaccurate value estimates in highly weighted states s . This objective cannot be directly optimized because it requires access to $v_{\pi}(s)$. One way to indirectly optimize the $\overline{\text{VE}}$ is to use the mean squared return error ($\overline{\text{RE}}$):

$$\overline{\text{RE}}(\mathbf{w}) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_{\pi} \left[(\hat{v}_{\mathbf{w}}(s) - G_t)^2 \mid S_t = s \right]. \quad (12)$$

The minima of the $\overline{\text{RE}}(\mathbf{w})$ and $\overline{\text{VE}}(\mathbf{w})$ are the same because their gradients are equal

$$\begin{aligned} \nabla \overline{\text{RE}}(\mathbf{w}) &= \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_{\pi} [(\hat{v}_{\mathbf{w}}(s) - G_t) \nabla \hat{v}_{\mathbf{w}}(s) \mid S_t = s] \\ &= 2 \sum_{s \in \mathcal{S}} d(s) (\hat{v}_{\mathbf{w}}(s) - \mathbb{E}_{\pi}[G_t \mid S_t = s]) \nabla \hat{v}_{\mathbf{w}}(s) \\ &= 2 \sum_{s \in \mathcal{S}} d(s) (\hat{v}_{\mathbf{w}}(s) - v_{\pi}(s)) \nabla \hat{v}_{\mathbf{w}}(s) = \nabla \overline{\text{VE}}(\mathbf{w}). \end{aligned}$$




Figure 1: The visualization above characterizes the true v_π , $\overline{\text{PBE}}$ solution and how projections operate on successive approximations. Assume the estimate of v_π starts from \mathbf{v} in red. The Bellman operator pushes the value estimate out of the space of representable functions represented by the plane (Note this corresponds to $\mathcal{F} = \mathcal{H}$ introduced in Section 4.2). The projection brings the approximation back down to the nearest representable function on the plane. This process is repeated over and over until the value estimates converge to the blue dot at the base of the black line. Subsequent updates push the approximation to v_π out of the space of representable functions and the projection back onto the plane. The true value in this case is outside of \mathcal{F} , with the $\overline{\text{VE}}$ being the distance between the \mathbf{v} at $\overline{\text{PBE}} = 0$ and v_π . Note the projection of v_π onto \mathcal{F} need not be equal to $\overline{\text{PBE}}$ solution.

In practice, the $\overline{\text{RE}}$ is rarely used because it requires obtaining samples of entire returns (only updating at the end of episodes). Instead, bootstrapping is used and so forms of the Bellman error are used, as in the $\overline{\text{BE}}$ and $\overline{\text{PBE}}$. The $\overline{\text{BE}}$ reflects the goal of approximating the fixed-point formula given by the Bellman operator \mathcal{T} , defined as

$$\mathcal{T}\hat{v}_w(\cdot)(s) \stackrel{\text{def}}{=} \mathbb{E}_\pi[R + \gamma\hat{v}_w(S') \mid S = s] \text{ for all } s. \quad (13)$$

When equality is not possible, the difference is minimized as in the $\overline{\text{BE}}$

$$\overline{\text{BE}}(\mathbf{w}) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} d(s) (\mathcal{T}\hat{v}_w(\cdot)(s) - \hat{v}_w(s))^2 = \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]^2 \quad (14)$$

where we write $\delta(\mathbf{w})$ to be explicit that this is the TD error for the parameters \mathbf{w} .

There has been much discussion, formal and informal, about using the $\overline{\text{BE}}$ versus the $\overline{\text{PBE}}$. The $\overline{\text{BE}}$ can be decomposed into the $\overline{\text{PBE}}$ and a projection penalty term (Scherrer, 2010). To understand why, recall the definition of a projection operator. For a vector space \mathcal{F} , the projection of a vector v onto \mathcal{F} is the closest point under a given (weighted) norm $\|\cdot\|_d$: $\min_{u \in \mathcal{F}} \|u - v\|_d$. This definition also applies to function spaces. Let $\Pi_{\mathcal{F},d}$ be the weighted projection on the space of value functions, defined as $\Pi_{\mathcal{F},d}u \stackrel{\text{def}}{=} \arg \min_{u \in \mathcal{F}} \|u - v\|_d$. For a given vector $v \in \mathbb{R}^{|\mathcal{S}|}$, composed of value estimates for each state, we get

$$\|v - \mathcal{T}v\|_d^2 = \underbrace{\|v - \Pi_{\mathcal{F},d}\mathcal{T}v\|_d^2}_{\overline{\text{PBE}}} + \underbrace{\|\mathcal{T}v - \Pi_{\mathcal{F},d}\mathcal{T}v\|_d^2}_{\text{Projection Penalty}} \quad (15)$$

This penalty causes the $\overline{\text{BE}}$ to prefer value estimates for which the projection does not have a large impact near the solution. The $\overline{\text{PBE}}$ can find a fixed point where applying the




Figure 2: A comparison of the $\overline{\text{BE}}$ and $\overline{\text{PBE}}$ solutions when the true value function is not representable. As before, we visualize how the approximation that minimizes the $\overline{\text{PBE}}$ at convergence can be far from v_π with a large projection penalty. The approximate value function that minimizes the $\overline{\text{BE}}$ on the other hand is closer to v_π and typically has a smaller projection penalty (note the Bellman operator would indeed push v_{BE} outside \mathcal{F}).

Bellman operator $\mathcal{T}v$ moves far outside the space of representable functions, as long as the projection back into the space stays at v . The projection penalty is sensible, and in fact prevents some of the counterexamples on the solution quality for the $\overline{\text{PBE}}$ discussed in Section 6. We visualize the projection penalty and a potential solution under the $\overline{\text{PBE}}$ in Figure 1, and contrast to a potential solution under the $\overline{\text{BE}}$ in Figure 2.

Despite the potential utility of the $\overline{\text{BE}}$, it has not been widely used due to difficulties in optimizing this objective without a model. The $\overline{\text{BE}}$ is difficult to optimize because of the well-known double sampling problem for the gradient. To see why, consider the gradient

$$\begin{aligned} \nabla_{\mathbf{w}} \overline{\text{BE}}(\mathbf{w}) &= \sum_{s \in \mathcal{S}} d(s) \nabla_{\mathbf{w}} \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]^2 \\ &= 2 \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] \mathbb{E}_\pi[\nabla_{\mathbf{w}} \delta(\mathbf{w}) \mid S = s] \\ &= 2 \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] \mathbb{E}_\pi[\gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S') - \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) \mid S = s] \end{aligned}$$

To estimate this gradient for a given $S = s$, we need two independent samples of the next state and reward. We use the first to get a sample $\delta(\mathbf{w})$ and the second to get a sample of $\gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S') - \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s)$. The product of these two samples gives an unbiased sample of the product of the expectations. If we instead only used one sample, we would erroneously obtain a sample of $\mathbb{E}_\pi[\delta(\mathbf{w})(\gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S') - \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s)) \mid S = s]$.

One promising attempt to approximate the $\overline{\text{BE}}$ used a non-parametric approach (Feng et al., 2019). The objective, called the $\overline{\text{KBE}}$, takes pairs of samples from a buffer to overcome the double sampling problem. Unfortunately, this cannot overcome the issue of non-identifiability in the $\overline{\text{BE}}$. There is a simple example where the same data is generated by two different MDPs, with different optima for the corresponding $\overline{\text{BE}}$ (Sutton and Barto, 2018, Chapter 11.6). The agent cannot hope to use the data to identify which of the two parameters is the optimal solution; that work used the term that the $\overline{\text{BE}}$ is not *learnable*.

The linear $\overline{\text{PBE}}$, on the other hand, is practical to optimize under linear function approximation, as discussed above: the whole family of (gradient) TD algorithms is designed to optimize the linear $\overline{\text{PBE}}$. Unfortunately, the $\overline{\text{PBE}}$ is hard to optimize for the general nonlinear setting, because the projection is hard to compute. Prior attempts to extend GTD

to the nonlinear $\overline{\text{PBE}}$ (Maei et al., 2009) resulted in an algorithm that requires computing Hessian-vector products. In the next section, we discuss how to overcome these issues for the $\overline{\text{BE}}$ and $\overline{\text{PBE}}$, with a unified objective, that is a generalization of the $\overline{\text{PBE}}$.

Finally, for completeness, we conclude with a description of the Mean-Squared TD error ($\overline{\text{TDE}}$), even though it is rarely used. The $\overline{\text{TDE}}$ was introduced to characterize the TD solution as a semi-gradient method. For the objective

$$\overline{\text{TDE}}(\mathbf{w}) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi \left[(R + \gamma \hat{v}_{\mathbf{w}}(S') - \hat{v}_{\mathbf{w}}(s))^2 \mid S = s \right] \quad (16)$$

the gradient includes the gradient of $\hat{v}_{\mathbf{w}}(S')$. TD omits this term, and so is called a semi-gradient method. It is not typical to do gradient descent on the $\overline{\text{TDE}}$, due to commonly held views of poor quality and a counterexample for the residual gradient algorithm which uses the $\overline{\text{TDE}}$ (Sutton and Barto, 2018). We highlight the significant bias when using the $\overline{\text{TDE}}$, in Appendix B, providing further evidence that it is likely not a useful direction.

4.2 An Identifiable $\overline{\text{BE}}$

Before discussing the generalized $\overline{\text{PBE}}$, we start by showing a conjugate form for the $\overline{\text{BE}}$. This reformulation uses the strategy introduced by Dai et al. (2017), which more generally introduces this conjugate form for several objectives that use conditional expectations. They show how to use it for the $\overline{\text{BE}}$ as an example, but defined it slightly differently because they condition on states and actions. For this reason, and because we will build further, we provide the explicit steps to derive the conjugate form for the $\overline{\text{BE}}$.

Let \mathcal{F} be the space of parameterized value functions and \mathcal{F}_{all} the space of all functions. The reformulation uses the fact that the biconjugate of the square function is $y^2 = \max_{h \in \mathbb{R}} 2yh - h^2$ and the fact that the maximum can be brought outside the sum (interchangeability), as long as a different scalar h can be chosen for each state s , as it can be for \mathcal{F}_{all} the space of all functions.

$$\begin{aligned} \overline{\text{BE}}(\mathbf{w}) &= \sum_{s \in \mathcal{S}} d(s) \mathbb{E}_\pi [\delta(\mathbf{w}) \mid S = s]^2 \\ &= \sum_{s \in \mathcal{S}} d(s) \max_{h \in \mathbb{R}} (2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] h - h^2) &> \text{using the biconjugate function} \\ &= \max_{h \in \mathcal{F}_{\text{all}}} \sum_{s \in \mathcal{S}} d(s) (2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] h(s) - h(s)^2) &> \text{using interchangeability.} \end{aligned}$$

The optimal $h^*(s) = \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]$, because

$$\begin{aligned} &\arg \max_{h \in \mathcal{F}_{\text{all}}} \sum_{s \in \mathcal{S}} d(s) (2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] h(s) - h(s)^2) \\ &= \arg \max_{h \in \mathcal{F}_{\text{all}}} \sum_{s \in \mathcal{S}} d(s) \left(2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] h(s) - h(s)^2 - \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]^2 \right) \\ &= \arg \max_{h \in \mathcal{F}_{\text{all}}} - \sum_{s \in \mathcal{S}} d(s) (\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] - h(s))^2 \\ &= \arg \min_{h \in \mathcal{F}_{\text{all}}} \sum_{s \in \mathcal{S}} d(s) (\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] - h(s))^2. \end{aligned}$$

The function $h^*(s) = \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]$ provides the minimal error of zero. This optimal solution also makes it clear why the above is simply a rewriting of the $\overline{\text{BE}}$ because

$$2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] h^*(s) - h^*(s)^2 = 2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]^2 - \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]^2 = \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]^2.$$

More generally, for the continuous state case, interchangeability also holds, as long as the function $h(s) = \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]$ satisfies $h \in \mathcal{F}_{\text{all}}$. We can more generically expressed the $\overline{\text{BE}}$ using expectations over states: $\overline{\text{BE}}(\mathbf{w}) = \mathbb{E}[\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S]^2]$, where the outer expectation is over S with distribution d . For $g(h, s) = \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s] h - h^2$, the $\overline{\text{BE}}$ is

$$\mathbb{E} \left[\max_{h \in \mathbb{R}} g(h, S) \right] = \int_{\mathcal{S}} d(s) \max_{h \in \mathbb{R}} g(h, s) ds = \max_{h \in \mathcal{F}_{\text{all}}} \int_{\mathcal{S}} d(s) g(h(s), s) ds. \quad (17)$$

Because $h(s) = \mathbb{E}_\pi[\delta \mid S = s]$ satisfies $h \in \mathcal{F}_{\text{all}}$, we know that a maximizer exists, as $h^* = h \in \mathcal{F}_{\text{all}}$. Then we can show that $\mathbb{E}[\max_{h \in \mathbb{R}} g(h, S)] = \mathbb{E}[g(h^*(S), S)] = \max_{h \in \mathcal{F}_{\text{all}}} \mathbb{E}[g(h(S), S)]$.²

As highlighted in (Sutton and Barto, 2018, Chapter 11.6), the $\overline{\text{BE}}$ is not identifiable. In that example, however, the inputs given to the value function learner are partially observable. In terms of the above formulation, this would mean the agent can only observe a part of the state for learning \mathbf{w} but the whole state to learn h . Naturally, however, the input-space for h should be similarly restricted to only observable information. This leads us to a new set for h , which includes all functions on the same inputs $\phi(s)$ as given to v , rather than on state:

$$\mathcal{H}_{\text{all}} \stackrel{\text{def}}{=} \{h = f \circ \phi \mid \text{where } f \text{ is any function on the space produced by } \phi\}.$$

The resulting h is restricted to functions of the form $h(s) = f(\phi(s))$. We call the resulting $\overline{\text{BE}}$ an *Identifiable* $\overline{\text{BE}}$, written as:

$$\text{Identifiable } \overline{\text{BE}}(\mathbf{w}) \stackrel{\text{def}}{=} \max_{h \in \mathcal{H}_{\text{all}}} \mathbb{E} [2\mathbb{E}_\pi[\delta(\mathbf{w}) \mid S] h(S) - h(S)^2].$$

Notice that $\mathcal{H}_{\text{all}} \subseteq \mathcal{F}_{\text{all}}$, and so the solution to the Identifiable $\overline{\text{BE}}$ may be different from the solution to the $\overline{\text{BE}}$. In particular, we know the Identifiable $\overline{\text{BE}}(\mathbf{w}) \leq \overline{\text{BE}}(\mathbf{w})$, because the inner maximization is more constrained. In fact, in many cases restricting h can be seen as a projection on the errors in the objective, as we discuss next, making the Identifiable $\overline{\text{BE}}$ an instance of the generalized $\overline{\text{PBE}}$.

4.3 From the Identifiable Bellman Error back to a Projected Bellman Error

The previous section discussed a conjugate form for the $\overline{\text{BE}}$, which led to an identifiable $\overline{\text{BE}}$. Even this Identifiable $\overline{\text{BE}}$, however, can be difficult to optimize, as we will not be able to perfectly represent any h in \mathcal{H}_{all} . In this section, we discuss further approximations, with $h \in \mathcal{H} \subseteq \mathcal{H}_{\text{all}}$, leading to a new set of objectives based on projecting the Bellman error.

To practically use the minimax formulation for the $\overline{\text{BE}}$, we need to approximate h as an auxiliary estimator. This means h must also be a parameterized function, and we will instead only obtain an approximation to the Identifiable $\overline{\text{BE}}$. Let \mathcal{H} be a convex space of

2. This argument is similar to (Dai et al., 2017, Lemma 1), but we do not need to assume g is upper semi-continuous and concave. They use this condition to ensure the existence of a maximum for $g(\cdot, s)$. We know the form of our g and can directly determine the existence of a maximizer (expected TD error).

parameterized functions for this auxiliary function h . As we show below, this \mathcal{H} defines the projection in the generalized $\overline{\text{PBE}}$.

In order to define the $\overline{\text{PBE}}$, we first define a projection operator $\Pi_{\mathcal{H},d}$ which projects any vector $u \in \mathbb{R}^{|\mathcal{S}|}$ onto convex subspace $\mathcal{H} \subseteq \mathbb{R}^{|\mathcal{S}|}$ under state weighting d :

$$\Pi_{\mathcal{H},d}u \stackrel{\text{def}}{=} \arg \min_{h \in \mathcal{H}} \|u - h\|_d. \quad (18)$$

We define the generalized $\overline{\text{PBE}}$ as

$$\overline{\text{PBE}}(\mathbf{w}) \stackrel{\text{def}}{=} \|\Pi_{\mathcal{H},d}(\mathcal{T}\hat{v}_{\mathbf{w}} - \hat{v}_{\mathbf{w}})\|_d^2 \quad (19)$$

where each choice of \mathcal{H} results in different projection operators. This view provides some intuition about the role of approximating h . Depending on how errors are projected, the value function approximation will focus more or less on the Bellman errors in particular states. If the Bellman error is high in a state, but those errors are projected to zero, then no further approximation resources will be used for that state. Under no projection—the set for h being the set of all functions—no errors are projected and the values are learned to minimize the Bellman error. If $\mathcal{H} = \mathcal{F}$, the same space is used to represent h and v , then we obtain the projection originally used for the $\overline{\text{PBE}}$.

We now show the connection between this $\overline{\text{PBE}}$ and the $\overline{\text{BE}}$. In the finite state setting, we have a vector $u \in \mathbb{R}^{|\mathcal{S}|}$ composed of entries $\mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s]$: $u = \mathcal{T}\hat{v}_{\mathbf{w}} - \hat{v}_{\mathbf{w}}$. The projection onto a convex set, with a diagonal weighting matrix $D \stackrel{\text{def}}{=} \text{diag}(d)$ is guaranteed to be an orthogonal projection. This property means $u = \Pi_{\mathcal{H},d}u + \tilde{u} = h + \tilde{u}$, where $h = \Pi_{\mathcal{H},d}u$ and \tilde{u} is the component in u that is orthogonal in the weighted space: $h^\top D \tilde{u} = 0$ for $D \stackrel{\text{def}}{=} \text{diag}(d)$. Then we can write the conjugate form for the $\overline{\text{BE}}$, now with restricted $\mathcal{H} \subset \mathcal{H}_{\text{all}}$

$$\begin{aligned} & \max_{h \in \mathcal{H}} \sum_{s \in \mathcal{S}} d(s) (2\mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s] h(s) - h(s)^2) \\ &= \max_{h \in \mathcal{H}} \sum_{s \in \mathcal{S}} d(s) (2u(s)h(s) - h(s)^2) && \triangleright \text{rewriting } u(s) = \mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s] \\ &= \sum_{s \in \mathcal{S}} d(s) (2u(s)h(s) - h(s)^2) && \triangleright \text{where } h = \Pi_{\mathcal{H},d}u \\ &= \sum_{s \in \mathcal{S}} d(s) (2[h(s) + \tilde{u}(s)]h(s) - h(s)^2) && \triangleright \text{because } u(s) = h(s) + \tilde{u}(s) \\ &= \sum_{s \in \mathcal{S}} d(s) (2h(s)^2 - h(s)^2) + 2 \sum_{s \in \mathcal{S}} d(s) \tilde{u}(s)h(s) \\ &= \sum_{s \in \mathcal{S}} d(s) h(s)^2 + 2 \sum_{s \in \mathcal{S}} d(s) \tilde{u}(s)h(s) \\ &= \sum_{s \in \mathcal{S}} d(s) h(s)^2 && \triangleright \text{where } \sum_{s \in \mathcal{S}} d(s) \tilde{u}(s)h(s) = 0 \text{ because} \\ &= \|\Pi_{\mathcal{H},d}(\mathcal{T}\hat{v}_{\mathbf{w}} - \hat{v}_{\mathbf{w}})\|_d^2 && h \text{ is orthogonal to } \tilde{u}, \text{ under weighting } d \\ &= \overline{\text{PBE}}(\mathbf{w}) \end{aligned}$$

The key assumption above is that \mathcal{H} is a convex set. Convexity is required to ensure that the projection operator $\Pi_{\mathcal{H},d}$ is an orthogonal projection (Dudek and Holly, 1994).




Figure 3: A visual interpretation of how the Bellman operator can push the value estimates outside the space of representable functions and the role of the projection operator. The set \mathcal{F} corresponds to the (parameterized) space of value functions and \mathcal{H} is the set of functions that approximate (project) the Bellman error $\mathcal{T}v - v$. Potential settings include $\mathcal{F} = \mathcal{H}$ (visualized in Figure 1), $\mathcal{F} \subset \mathcal{H}$ visualized in (a) and $\mathcal{F} \neq \mathcal{H}$ visualized in (b). In (a), we highlight two cases: $\mathcal{T}v$ is not representable by any function in \mathcal{F} or \mathcal{H} , or $\mathcal{T}v$ is representable by functions in \mathcal{H} but not \mathcal{F} . In (b) we see examples of projections when \mathcal{F} intersects \mathcal{H} .

This assumption is easily satisfied by linear functions with a fixed basis $\phi(s)$, with bounded weights. For any two functions with weights \mathbf{w}_1 and \mathbf{w}_2 , the function defined by weights $(1 - c)\mathbf{w}_1 + c\mathbf{w}_2$ is also in the set. More generally, typical convex nonlinear function approximation sets used in machine learning are reproducing kernel Hilbert spaces. Many classes of neural networks have been shown to be expressible as RKHSs (see Bietti and Mairal (2019) for a nice overview), including neural networks with ReLU activations as are commonly used in RL. Therefore, this is not an overly restrictive assumption.

4.4 The Connection to Previous $\overline{\text{PBE}}$ Objectives

In this section we show how the generalized $\overline{\text{PBE}}$ lets us express the linear $\overline{\text{PBE}}$, and even the nonlinear $\overline{\text{PBE}}$, by selecting different sets \mathcal{H} . First let us consider the linear $\overline{\text{PBE}}$. The easiest way to see this is to use the saddlepoint formulations developed for the linear $\overline{\text{PBE}}$ (Mahadevan et al., 2014; Liu et al., 2016; Touati et al., 2018). The goal there was to re-express the existing linear $\overline{\text{PBE}}$ using a saddlepoint form, rather than to re-express the $\overline{\text{BE}}$ or find connections between them. The linear $\overline{\text{PBE}} = \|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2$ can be rewritten using the conjugate for the two norm. The conjugate for the two-norm is $\frac{1}{2}\|\mathbf{y}\|_{\mathbf{C}^{-1}} = \max_{\mathbf{h}} \mathbf{y}^\top \mathbf{h} - \frac{1}{2}\|\mathbf{h}\|_{\mathbf{C}}^2$, with optimal $\mathbf{h} = \mathbf{C}^{-1}\mathbf{y}$. Correspondingly, we get

$$\frac{1}{2}\|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2 = \max_{\mathbf{h} \in \mathbb{R}^d} (\mathbf{b} - \mathbf{A}\mathbf{w})^\top \mathbf{h} - \frac{1}{2}\|\mathbf{h}\|_{\mathbf{C}}^2$$

where the solution for $\mathbf{h} = \mathbf{C}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{w})$. This solution makes the first term equal to $\|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2$ and the second term equal to $-\frac{1}{2}\|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2$; adding them gives $\frac{1}{2}\|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2$.

We can obtain the same formulation under the $\overline{\text{PBE}}$, by restricting \mathcal{F} and \mathcal{H} to be the same set of linear functions. Let $\mathcal{L} = \{f : \mathcal{S} \rightarrow \mathbb{R} : f(s) = \mathbf{x}(s)^\top \mathbf{w}, \mathbf{w} \in \mathbb{R}^d\}$. For $\mathcal{F} = \mathcal{H} = \mathcal{L}$, we have that $h^* = \arg \min_{h \in \mathcal{L}} \sum_{s \in \mathcal{S}} d(s) (\mathbb{E}_\pi[\delta(\mathbf{w}) | S = s] - h(s))^2$ satisfies $h^*(s) = \mathbf{x}(s)^\top \mathbf{h}^*$. This h^* is the linear regression solution for targets $\delta(\mathbf{w})$, so $\mathbf{h}^* = \mathbb{E}[\mathbf{x}\mathbf{x}^\top]^{-1} \mathbb{E}[\mathbf{x}\delta(\mathbf{w})]$ which equals $\mathbf{C}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{w})$. We can further verify that the resulting $\overline{\text{PBE}}$ matches the linear $\overline{\text{PBE}}$

(see Appendix C). This result is alluded to in the connection between the NEU and the $\overline{\text{KBE}}$, in (Feng et al., 2019, Corollary 3.5), but not explicitly shown.

This connection also exists with the nonlinear $\overline{\text{PBE}}$, but with a surprising choice for the parameterization of h : using the gradient of the value estimate as the features. The nonlinear $\overline{\text{PBE}}$ is defined as (Maei et al., 2009)

$$\text{nonlinear } \overline{\text{PBE}}(\mathbf{w}) = \mathbb{E}[\delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)]^{\top} \mathbb{E}\left[\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)^{\top}\right]^{-1} \mathbb{E}[\delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)].$$

This corresponds to the linear $\overline{\text{PBE}}$ when $\mathcal{F} = \mathcal{L}$, because $\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) = \mathbf{x}(s)$. Define set $\mathcal{G}_{\mathbf{w}} = \{f : \mathcal{S} \rightarrow \mathbb{R} : f(s) = \mathbf{y}(s)^{\top}\mathbf{h}, \mathbf{h} \in \mathbb{R}^d \text{ and } \mathbf{y}(s) = \nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)\}$. Notice that this function set for h changes as \mathbf{w} changes. Then we get that

$$h_{\text{nl}}^* = \arg \min_{h \in \mathcal{G}_{\mathbf{w}}} \sum_{s \in \mathcal{S}} d(s) (\mathbb{E}_{\pi}[\delta(\mathbf{w}) | S = s] - h(s))^2$$

satisfies $h_{\text{nl}}^*(s) = \nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)^{\top} \mathbf{h}_{\text{nl}}^*$ where $\mathbf{h}_{\text{nl}}^* = \mathbb{E}\left[\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)^{\top}\right]^{-1} \mathbb{E}[\delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)]$.

Plugging this optimal h back into the formula, we get that

$$\begin{aligned} & \max_{h \in \mathcal{G}_{\mathbf{w}}} \sum_{s \in \mathcal{S}} d(s) (2\mathbb{E}_{\pi}[\delta(\mathbf{w}) | S = s] h(s) - h(s)^2) = \sum_{s \in \mathcal{S}} d(s) (2\mathbb{E}_{\pi}[\delta(\mathbf{w}) | S = s] h_{\text{nl}}^*(s) - h_{\text{nl}}^*(s)^2) \\ & = \left(\sum_{s \in \mathcal{S}} d(s) 2\mathbb{E}_{\pi}[\delta(\mathbf{w}) | S = s] \nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)^{\top} \right) \mathbf{h}_{\text{nl}}^* - \sum_{s \in \mathcal{S}} d(s) (\mathbf{h}_{\text{nl}}^*)^{\top} \nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) \nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)^{\top} \mathbf{h}_{\text{nl}}^* \\ & = 2\mathbb{E}[\delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)]^{\top} \mathbf{h}_{\text{nl}}^* - (\mathbf{h}_{\text{nl}}^*)^{\top} \mathbb{E}[\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s)^{\top}] \mathbf{h}_{\text{nl}}^* \\ & = 2\text{nonlinear } \overline{\text{PBE}}(\mathbf{w}) - \text{nonlinear } \overline{\text{PBE}}(\mathbf{w}) \\ & = \text{nonlinear } \overline{\text{PBE}}(\mathbf{w}) \end{aligned}$$

This nonlinear $\overline{\text{PBE}}$ is not an instance of the generalized $\overline{\text{PBE}}$, as we have currently defined it, because the \mathcal{H} changes with \mathbf{w} . It is possible that such a generalization is worthwhile, as using the gradient of the values as features is intuitively useful. Further, interchangeability should still hold, as the exchange of the maximum was done for a fixed \mathbf{w} . Therefore, it is appropriate to explore an \mathcal{H} that changes with \mathbf{w} , and in our experiments we test $\mathcal{H} = \mathcal{G}_{\mathbf{w}}$.

In summary, in this section we introduced the generalized $\overline{\text{PBE}}$ and highlighted connections to the linear $\overline{\text{PBE}}$ and $\overline{\text{BE}}$. The generalized $\overline{\text{PBE}}$ provides a clear path to develop value estimation under nonlinear function approximation, providing a strict generalization of the linear $\overline{\text{PBE}}$. Two secondary benefits are that the generalized $\overline{\text{PBE}}$ provides a clear connection between the $\overline{\text{BE}}$ and $\overline{\text{PBE}}$, based on a difference in the choice of projection (\mathcal{H}), and resolves the identifiability issue in the $\overline{\text{BE}}$.

5. Understanding the Impact of Choices in the Generalized $\overline{\text{PBE}}$

The two key choices in the Generalized $\overline{\text{PBE}}$ is the state-weighting and the (projection) set \mathcal{H} . There are at least three clear criteria for selecting \mathcal{H} and the state-weighting: (1) the quality of the solution, (2) the feasibility of implementation and (3) the estimation error during learning and the impact on learning the primary weights. In this section, we

provide some conceptual and empirical insight into how to choose \mathcal{H} , and empirically show that the choice of weighting can significantly change the quality of the solution. After first showing that these choices clearly matter, and some intuition for why, we then provide theory characterizing the quality of the solution in the following section.

5.1 The Projection Set and the Quality of the Solution

The first criteria parallels the long-standing question about the quality of the solution under the linear $\overline{\text{PBE}}$ versus the $\overline{\text{BE}}$. The examples developed for that comparison provide insights on \mathcal{H} . In this section, we revisit these examples, now in context of the generalized $\overline{\text{PBE}}$.

Objectives based on Bellman errors perform *backwards bootstrapping*, where the value estimates in a state s are adjusted both toward the value of the next state and the value of the previous state. In the case of the $\overline{\text{BE}}$, backwards bootstrapping can become an issue when two or more states are heavily aliased and these aliased states lead to successor states with highly different values. Because the aliased states look no different to the function approximator, they must be assigned the same estimated value. For each of these aliased states, backwards bootstrapping forces the function approximator to balance between accurately predicting the successor values for all aliased states, as well as adjusting the successor values to be similar to those of the aliased states.

The $\overline{\text{PBE}}$, on the other hand, projects the error for the aliased states, ignoring the portion of the Bellman error that forces the function approximator to balance the similarity between the aliased state value and the successor state value. This allows the function approximator the freedom to accurately estimate the values of the successor states without trading off error in states which it cannot distinguish.

To make this concrete, consider the following 4-state MDP from Sutton et al. (2009). State A_1 and A_2 are aliased under the features for \mathcal{F} . For the linear $\overline{\text{PBE}}$, $\mathcal{H} = \mathcal{F}$, and so the states are also aliased when approximating h . For the $\overline{\text{BE}}$, they are not aliased for h . A_1 transitions to B and terminates with reward 1. A_2 transition to C and terminates with reward 0. The linear $\overline{\text{PBE}}$ results in the correct values for B and C —1 and 0 respectively—because it does not suffer from backwards bootstrapping. The $\overline{\text{BE}}$, on the other hand, assigns them values $\frac{3}{4}$ and $\frac{1}{4}$, to reduce Bellman errors at A_1 and A_2 . A generalized $\overline{\text{PBE}}$ with other $\mathcal{H} \neq \mathcal{F}$ would suffer the same issue as the $\overline{\text{BE}}$ in this example, unless the projection $\Pi_{\mathcal{H}}$ mapped errors in the aliased states to zero.

On the other hand, the linear $\overline{\text{PBE}}$ can find solutions where the Bellman error is very high, even though the projected Bellman error is zero. Consider the plane of value functions that can be represented with a linear function approximator. The Bellman operator can take the values far off of this surface, only to be projected back to this surface through the projection operator. At the fixed-point, this projection brings the value estimate back to the original values and the distance that the value estimate moved on the plane is zero, thus the $\overline{\text{PBE}}$ is zero. The $\overline{\text{PBE}}$ can be zero even when the $\overline{\text{BE}}$ is large. Kolter (2011) provides an example where the solution under the $\overline{\text{PBE}}$ can be made arbitrarily far from the true value function. We expand on this example in Figure 4, and show that the solution under the linear $\overline{\text{PBE}}$ can be arbitrarily poor, even though the features allow for an ϵ accurate value estimate and the solution under the $\overline{\text{BE}}$ is very good.




Figure 4: The visualization above shows how the $\overline{\text{PBE}}$ solution can result in arbitrarily bad value error under some behaviours. The blue line above is the same as the visualization used in prior work to demonstrate issues with minimizing $\overline{\text{PBE}}$ (see Kolter (2011) for a description of the counterexample). The vertical axis measures $\overline{\text{VE}}$ and the horizontal axis different behavior policies. This figure differs from Kolter (2011); we show that the $\overline{\text{BE}}$ solution exhibits low error and highlight the impact of changing \mathcal{H} . The size of the set \mathcal{H} increases from the left subplot to the right. More behavior policies result in low generalized $\overline{\text{PBE}}$ as the set \mathcal{H} increases.

If we use an emphasis weighting for the states, rather than behavior visitation, then the solution under the $\overline{\text{PBE}}$ becomes reasonable. Further, even just a small change to \mathcal{H} so that $\mathcal{H} \neq \mathcal{F}$ resolves this counterexample. Ideally, we would use an \mathcal{H} similar to \mathcal{F} , to avoid backwards bootstrapping. Conceptually, a potentially reasonable choice for \mathcal{H} is therefore either (1) $\mathcal{H} = \mathcal{F}$ with some consideration on adjusting the state weighting and (2) an \mathcal{H} that is only slightly bigger than \mathcal{F} , potentially with the inclusion of an additional feature.

A Simple Experiment for Solution Quality Under Different Weightings and \mathcal{H}

We empirically investigate the quality of the solution under the $\overline{\text{PBE}}$ and $\overline{\text{BE}}$ with three different weightings: d_b , d_π and m . The solution quality is measured by the $\overline{\text{VE}}$ under d_b and d_π . We compute the fixed-point of each objective on a 19-state random walk with randomly chosen target and behavior policies. To isolate the impact of representation on the fixed-points, we investigate several forms of state representation where v_π is outside the representable function class. We include the *Dependent* features from Sutton et al. (2009), randomly initialized sparse ReLU networks, tile-coded features, and state aggregation.

The random-walk has 19 states with the left-most and right-most state being terminal. The reward function is zero everywhere except on transitioning into the right-most terminal state where the agent receives +1 reward, and on the left-most terminal state where the agent receives -1 reward. The discount factor is set to $\gamma = 0.99$.

We run each experimental setting one million times with a different randomly initialized neural network, random offset between tilings in the tile-coder, and randomly sampled target and behavior policy. The policies are chosen uniformly randomly on the standard simplex. The neural network is initialized with a Xavier initialization (Glorot and Bengio, 2010), using 76 nodes in the first hidden layer and 9 nodes in the final feature layer. Then 25% of the neural network weights are randomly set to zero to encourage sparsity between




Figure 5: Investigating the \overline{VE} of the fixed-points of \overline{PBE} and \overline{BE} under d_b , d_π , and m on a 19-state random walk. All errors are computed closed form given access to the reward and transition dynamics. The fixed-point of the \overline{PBE} with emphatic weighting consistently has the lowest error across several different state representations (light color); while the fixed-point of the \overline{PBE} under d_b has the highest error (dark blue). Results are averaged over one million randomly generated policies and state representations.

connections and to increase variance between different randomly generated representations. The tile-coder uses 4 tilings each offset randomly and each containing 4 tiles. The state aggregator aggressively groups the left-most states into one bin and the right-most states into another, creating only two features.

Figure 5 shows the normalized log-error of the fixed-points of the \overline{PBE} and \overline{BE} under each weighting. A normalized error between $[0, 1]$, for each representation, is obtained by (1) computing the best value function representable by those features, $\min_{v \in \mathcal{F}} \overline{VE}(v)$ under d_b or d_π and (2) subtracting this minimal \overline{VE} , and normalizing by the maximum \overline{VE} for each column (across objectives and weightings for a fixed representation). The fixed-points are computed using their least-squares closed form solutions given knowledge of the MDP dynamics. Plotted is the mean error across the one million randomly initialized experimental settings. The standard error between settings is negligibly small.

Interestingly, the fixed-points corresponding to weighting d_b consistently have the highest error across feature representations, even on the excursion \overline{VE} error metric with weighting d_b . The \overline{PBE} under emphatic weighting, m , consistently has the lowest error across all feature representations, though is slightly outperformed by \overline{PBE} with weighting d_π for the \overline{VE} with weighting d_π . In these experiments, the \overline{BE} appears to have no advantages over the \overline{PBE} , meaning that the more restricted \mathcal{H} for \overline{PBE} produces sufficiently high quality solutions.

5.2 Feasibility of the Implementation

There are many feasible choices for estimating h . Likely the simplest is to use the same approximator for h as for the values. For example, this might mean that h and v use the same features, or that we have two heads on a shared neural network. However, we could

feasibly consider a much bigger class for h , because h is only used during training, not prediction. For example, we might want v to be efficient to query, and so use a compact parametric function approximator. But h could use a more computationally costly function approximator, updated with replay between agent-environment interaction steps.

To expand the space \mathcal{H} , one feasible approach is to use a separate set of features for \mathcal{H} or learn a separate neural network. The separate neural network implicitly can learn a different set of features, and so allows h to use different features than v . If we allow this second neural network to be much bigger, then we expand the space \mathcal{H} and make the generalized $\overline{\text{PBE}}$ closer to the $\overline{\text{BE}}$.

We can take this expansion further by using nonparametric function approximators for h . For example, a reservoir of transitions can be stored, where $\mathbb{E}_\pi[\delta_t \mid S_t = s]$ is approximated using a weighted average over δ_t in the buffer, where the weighting is proportional to similarity between that state and s . This is the strategy taken by the Kernel $\overline{\text{BE}}$ (Feng et al., 2019), precisely to reduce bias in h and so better approximate the $\overline{\text{BE}}$.

When learning online, this non-parametric approach is less practical. Either a large buffer needs to be maintained, or a sufficient set of representative transitions identified and stored. Further, it is not clear that estimating the $\overline{\text{BE}}$ more closely is actually desirable, as discussed in the previous section. In this work, where we learn online, we advocate for the simplest approach: a shared network, with two heads (see Section 8.2). We also show that learning two separate neural networks performs comparably, in Section 9.5.

5.3 Estimation Error and the Impact on Primary Weight Updates

The generalized $\overline{\text{PBE}}$ presents an additional trade-off between approximation error and estimation error in h . A rich \mathcal{H} may reduce the approximation error (and projection penalty) at the expense of higher error in estimating δ_t via h . A more restricted \mathcal{H} may yield lower estimation error, because less data is needed to estimate h . Note that this trade-off is for approximating/estimating the objective itself. It is different from—and secondary to—the approximation-estimation trade-off for the value function with set \mathcal{F} .

One strategy to restrict \mathcal{H} is to add regularization on h . For example, an ℓ_2 regularizer constrains h to be closer to zero—reducing variance—and improves convergence rates. This strategy was introduced in an algorithm called TD with Regularized Corrections (TDRC) and the control variant, Q-learning with Regularized Corrections (QRC) (Ghiassian et al., 2020). Empirically, these algorithms performed comparably to their TD counterparts, in some cases performing significantly better. This particular constraint on \mathcal{H} was particularly appropriate, because the bias from regularization asymptotically disappears: at the TD fixed-point, the true parameters θ are zero and regularization biases θ towards zero.

More generally, the criteria for selecting \mathcal{H} is about improving the primary update, rather than necessarily reducing approximation error or estimation error for \mathcal{H} . Characterizing how \mathcal{H} improves updates for the primary weights remains an open question. One could imagine algorithmic strategies to identify such an \mathcal{H} using meta-learning, with the objective to optimize features for \mathcal{H} to make the primary weights learn more quickly. This question is particularly difficult to answer, as h can be used in two ways: within the standard saddlepoint update or for gradient corrections, as we discuss further in Section 7. With gradient corrections, the interim bias in h is less problematic than in the saddlepoint update.

In this work, we again advocate for a simple choice: reducing estimation error for \mathcal{H} using ℓ_2 regularization, within a gradient corrections update as in TDRC (Ghiassian et al., 2020). We find this choice to be generally effective, and easy to use. Nonetheless, there is clearly much more investigation that can be done to better understand the choice of \mathcal{H} .

6. Bounding Value Error & the Impact of Weighting on Solution Quality

The desired objective to minimize is the value error with weighting d_{eval} . We, however, optimize a surrogate objective, like the $\overline{\text{PBE}}$, with a potentially different weighting d . In fact, in the last section in Figure 5, we saw that it can be better to pick $d \neq d_{\text{eval}}$, where optimizing $\overline{\text{PBE}}$ with $d = m$ produce better solutions in terms of $\overline{\text{VE}}$ with $d_{\text{eval}} = d_b$ than $\overline{\text{PBE}}$ with $d = d_b$. In this section, we characterize the solution quality under the generalized $\overline{\text{PBE}}$, which depends both on \mathcal{H} and d .

Let $v_{\mathbf{w}}$ be the vector consisting of value function estimates $v(s, \mathbf{w})$. Further, let $v_{\mathbf{w}_{\mathcal{H},d}}$ be the solution to the generalized $\overline{\text{PBE}}$. Similarly to prior theoretical work (Yu and Bertsekas, 2010, Equation 5), our goal is to find bounds of the form

$$\underbrace{\|v_{\mathbf{w}_{\mathcal{H},d}} - v_{\pi}\|_{d_{\text{eval}}}}_{\text{Value error}} \leq C(d_{\text{eval}}, d, \mathcal{H}) \underbrace{\|\Pi_{\mathcal{F},d}v_{\pi} - v_{\pi}\|_d}_{\text{Approximation error}} \quad (20)$$

where the constant $C(d_{\text{eval}}, d, \mathcal{H})$ in the bound depends on the two weightings and the projection set \mathcal{H} . The term $\|\Pi_{\mathcal{F},d}v_{\pi} - v_{\pi}\|_d = \min_{v \in \mathcal{F}} \|v - v_{\pi}\|_d$ represents the approximation error: the lowest error under function class \mathcal{F} if we could directly minimize $\overline{\text{VE}}$ under our weighting d . Compared to prior theoretical work (Yu and Bertsekas, 2010, Equation 5), here we generalize to the nonlinear setting and where d may not equal d_{eval} . We start in the case where $\mathcal{H} = \mathcal{F}$, and then generalize to $\mathcal{H} \supset \mathcal{F}$ —where \mathcal{H} is a superset of \mathcal{F} —in the following subsection. Many of these results build on existing work, which we reference throughout; we also provide a summary table of existing results in the appendix, in Table 1.

6.1 Upper Bound on $\overline{\text{VE}}$ when $\mathcal{H} = \mathcal{F}$

Throughout this section we will assume that $\mathcal{H} = \mathcal{F}$, so that the projection operator for both the objective and value function space is the same. This matches the setting analyzed for the linear $\overline{\text{PBE}}$, though here we allow for nonlinear functions.

Our goal is to characterize the solution to the $\overline{\text{PBE}}$, the fixed point $v_{\mathbf{w}_{\mathcal{F},d}} = \Pi_{\mathcal{F},d}\mathcal{T}v_{\mathbf{w}_{\mathcal{F},d}}$. The typical approach is to understand the properties of \mathcal{T} under norm $\|\cdot\|_d$, as in Bertsekas and Tsitsiklis (1996, Lemma 6.9) or White (2017, Theorem 1). However, we can actually obtain more general results, by directly characterizing $\Pi_{\mathcal{F},d}\mathcal{T}$ and making assumptions about the norm only for a subset of value functions. This approach builds on the strategy taken by Kolter (2011, Theorem 2), where the choice of d was constrained to ensure a contraction, and on the strategy taken by Ghosh and Bellemare (2020, Theorem 4.3), where the set of value functions is constrained to ensure a contraction. We combine the two ideas and get a more general condition, as well as an extension to nonlinear function approximation.

Assumption 1 (Convex Function Space) *The set \mathcal{F} is convex.*

This convexity assumption is needed to ensure the projection operator has the typical properties, particularly that $\|\Pi_{\mathcal{F},d}(v_1 - v_2)\|_d \leq \|v_1 - v_2\|_d$ for all $v_1, v_2 \in \mathcal{F}$.

To characterize the Bellman operator, it will be useful to directly define the discounted transition operator (matrix) under π , $\mathcal{P}_{\pi,\gamma} \in \mathbb{R}^{|S| \times |S|}$, where

$$\mathcal{P}_{\pi,\gamma}(s, s') \stackrel{\text{def}}{=} \sum_a \pi(a|s) P(s'|s, a) \gamma(s, a, s') \quad (21)$$

For a constant discount of $\gamma_c < 1$ in the continuing setting, this simplifies to $\mathcal{P}_{\pi,\gamma}(s, s') = \gamma_c P_\pi$ for $P_\pi(s, s') \stackrel{\text{def}}{=} \sum_a \pi(a|s) P(s'|s, a)$. We can characterize when the projected Bellman operator is a contraction, by using either the norm of this discounted transition operator or the norm of the projected discounted transition operator for a restricted set of value functions.

Definition 2 (Discounted Transition Constant) *Define the discounted transition constant $c_d \stackrel{\text{def}}{=} \|\mathcal{P}_{\pi,\gamma}\|_d$, the weighted spectral norm of the discounted transition operator $\mathcal{P}_{\pi,\gamma}$.*

Definition 3 (Operator Constant) *Define the projected Bellman operator constant $c_{\mathcal{F},d} > 0$ for the set of value functions $\mathcal{F}_{\text{sub}} \subseteq \mathcal{F}$ as the constant that satisfies*

$$\|\Pi_{\mathcal{F},d} \mathcal{P}_{\pi,\gamma}(v_1 - v_2)\|_d \leq c_{\mathcal{F},d} \|v_1 - v_2\|_d \quad \text{for any } v_1, v_2 \in \mathcal{F}_{\text{sub}}. \quad (22)$$

Notice that $c_{\mathcal{F},d} \leq c_d$ because $\|\Pi_{\mathcal{F},d} \mathcal{P}_{\pi,\gamma} v\|_d \leq \|\mathcal{P}_{\pi,\gamma} v\|_d$ for any v .

We show next that the $\overline{\text{VE}}$ of the solution $v_{\mathbf{w}_{\mathcal{F},d}}$ to the generalized $\overline{\text{PBE}}$ is upper bounded by the approximation error times a constant. This constant depends only on the discounted transition constant, if it is less than 1, and otherwise depends also on the operator constant.

Theorem 4 *Assume $c_{\mathcal{F},d} < 1$. Let*

$$C(d, \mathcal{F}) \stackrel{\text{def}}{=} \begin{cases} \frac{1+c_d}{1-c_{\mathcal{F},d}} & \text{if } c_d \geq 1; \\ \frac{1}{1-c_d} & \text{if } c_d < 1. \end{cases} \quad (23)$$

Then

$$\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d \leq C(d, \mathcal{F}) \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d. \quad (24)$$

Proof Case 1: $c_d < 1$. This follows using the standard strategy in Bertsekas and Tsitsiklis (1996, Lemma 6.9) or White (2017, Theorem 1). Notice first that $v_{\mathbf{w}_{\mathcal{F},d}} = \Pi_{\mathcal{F},d} \mathcal{T} v_{\mathbf{w}_{\mathcal{F},d}}$, since it is a solution to the $\overline{\text{PBE}}$. Also note that $v_\pi = \mathcal{T} v_\pi$.

$$\begin{aligned} \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d &\leq \|v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d} v_\pi\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ &= \|\Pi_{\mathcal{F},d} \mathcal{T} v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d} \mathcal{T} v_\pi\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ &= \|\Pi_{\mathcal{F},d} \mathcal{T}(v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi)\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ &\leq \|\mathcal{T}(v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi)\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ &= \|\mathcal{P}_{\pi,\gamma}(v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi)\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ &\leq \|\mathcal{P}_{\pi,\gamma}\|_d \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ &\leq c_d \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d + \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d \\ \implies (1 - c_d) \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d &\leq \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d. \end{aligned}$$

Note that in the above $\mathcal{T}(v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi) = \mathcal{T} v_{\mathbf{w}_{\mathcal{F},d}} - \mathcal{T} v_\pi = \mathcal{P}_{\pi,\gamma} v_{\mathbf{w}_{\mathcal{F},d}} - \mathcal{P}_{\pi,\gamma} v_\pi$ because the reward term in the Bellman operator cancels in the subtraction.

Case 2: $c_d \geq 1$. For this case, we use the approach in Kolter (2011, Theorem 2). We cannot use the above approach, since $(1 - c_d)$ is negative. We start again by adding and subtracting $\Pi_{\mathcal{F},d}v_\pi$, but bound the first term differently.

$$\begin{aligned} \|v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}v_\pi\|_d &= \|\Pi_{\mathcal{F},d}\mathcal{T}v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}\mathcal{T}v_\pi\|_d \\ &= \|\Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}v_\pi\|_d \\ &\leq \|\Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}\Pi_{\mathcal{F},d}v_\pi\|_d + \|\Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}\Pi_{\mathcal{F},d}v_\pi - \Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}v_\pi\|_d. \end{aligned}$$

By assumption $c_{\mathcal{F},d} < 1$ and both $v_{\mathbf{w}_{\mathcal{F},d}} \in \mathcal{F}$ and $\Pi_{\mathcal{F},d}v_\pi \in \mathcal{F}$. Therefore, for the first term

$$\|\Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}\Pi_{\mathcal{F},d}v_\pi\|_d \leq c_{\mathcal{F},d}\|v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}v_\pi\|_d \leq c_{\mathcal{F},d}\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d.$$

For the second term, we have that

$$\|\Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}\Pi_{\mathcal{F},d}v_\pi - \Pi_{\mathcal{F},d}\mathcal{P}_{\pi,\gamma}v_\pi\|_d \leq \|\mathcal{P}_{\pi,\gamma}\Pi_{\mathcal{F},d}v_\pi - \mathcal{P}_{\pi,\gamma}v_\pi\|_d \leq c_d\|\Pi_{\mathcal{F},d}v_\pi - v_\pi\|_d.$$

Putting this all together, we have

$$\begin{aligned} \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d &\leq \|v_{\mathbf{w}_{\mathcal{F},d}} - \Pi_{\mathcal{F},d}v_\pi\|_d + \|\Pi_{\mathcal{F},d}v_\pi - v_\pi\|_d \\ &\leq c_{\mathcal{F},d}\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d + (1 + c_d)\|\Pi_{\mathcal{F},d}v_\pi - v_\pi\|_d \\ \implies (1 - c_{\mathcal{F},d})\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d &\leq (1 + c_d)\|\Pi_{\mathcal{F},d}v_\pi - v_\pi\|_d. \quad \blacksquare \end{aligned}$$

Finally, we can get the desired result in Equation (20), by considering a different state weighting d_{eval} for the evaluation versus the state weighting d we use in the $\overline{\text{PBE}}$. For example, in the off-policy setting, d_{eval} could correspond to d_π but we learn under $d = d_b$ or the emphatic weighting $d = m$. It is appropriate to separate the evaluation and solution weightings, as it may not be feasible to use $d = d_{\text{eval}}$ and further it is even possible we can improve our solution by carefully selecting d different from d_{eval} . Some insights about when this choice can be beneficial has been provided in seminal work on policy gradient methods (Kakade and Langford, 2002). We do not explore when we can obtain such improvements in this work. Our bounds assume worst case differences, as per the next definition, and so the bound on solution quality is bigger when $d \neq d_{\text{eval}}$.

Definition 5 (State Weighting Mismatch) *Define the state weighting mismatch between the desired weighting d_{eval} and the weighting used in the solution d as*

$$\kappa(d_{\text{eval}}, d) \stackrel{\text{def}}{=} \max_{s \in \mathcal{S}} \frac{d_{\text{eval}}(s)}{d(s)}. \quad (25)$$

Corollary 6 *Again assuming $c_{\mathcal{F},d} < 1$, we can further bound the error under a different weighting d_{eval}*

$$\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_{d_{\text{eval}}} \leq \sqrt{\kappa(d_{\text{eval}}, d)} C(d, \mathcal{F}) \|\Pi_{\mathcal{F},d}v_\pi - v_\pi\|_d.$$

Proof We use the state mismatch $\kappa(d_{\text{eval}}, d) = \max_{s \in \mathcal{S}} \frac{d_{\text{eval}}(s)}{d(s)}$ and get

$$\begin{aligned} \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_{d_{\text{eval}}}^2 &= \sum_{s \in \mathcal{S}} d_{\text{eval}}(s) (v_{\mathbf{w}_{\mathcal{F},d}}(s) - v_\pi(s))^2 = \sum_{s \in \mathcal{S}} d_{\text{eval}}(s) \frac{d(s)}{d(s)} (v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi)^2 \\ &\leq \kappa(d_{\text{eval}}, d) \sum_{s \in \mathcal{S}} d(s) (v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi)^2 = \kappa(d_{\text{eval}}, d) \|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d^2. \end{aligned}$$

Combined with the inequality in Equation (24), we obtain

$$\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_{d_{\text{eval}}} \leq \sqrt{\kappa(d_{\text{eval}}, d)} C(d, \mathcal{F}) \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d$$

■

This corollary gives us the desired result, specifically with $C(d_{\text{eval}}, d, \mathcal{F}) = \sqrt{\kappa(d_{\text{eval}}, d)} C(d, \mathcal{F})$.

The next question is when we can expect $c_{\mathcal{F},d} < 1$ or $c_d < 1$. More is known about when $c_d < 1$, because this is the condition used to prove that on-policy TD and emphatic TD are convergent: under $d = d_\pi$ and $d = m$, we know that $s_d < 1$ (White, 2017, Theorem 1). It is clear that other weights d will give this result, but we also know certain off-policy settings where $s_d > 1$, such as in Baird’s counterexample. Even if $s_d \geq 1$, we can still ensure the projected Bellman operator is a contraction by either limiting the set of weightings d or \mathcal{F} so that $c_{\mathcal{F},d} < 1$. Of course, if $d = d_\pi$ or m , then automatically $c_{\mathcal{F},d} < 1$ because for all d and \mathcal{F} , $c_{\mathcal{F},d} \leq c_d$. We provide some characterization of these constants in the next proposition, including both known results and a few novel ones.

Proposition 7 *The following statements are true about c_d .*

1. If $d = d_\pi$ or $d = m$, then $c_d < 1$.
2. If $d = d_\pi$ and the discount is a constant $\gamma_c < 1$ (a continuing problem), then $c_d = \gamma_c$.
3. If $d = d_\pi$ and for some constant γ_c , $\gamma(s, a, s') \leq \gamma_c$ for all (s, a, s') , then $c_d \leq \gamma_c$.
4. If $d_\pi(s) > 0$ iff $d(s) > 0$, then $c_d \leq s_{d_\pi} \sqrt{\kappa(d, d_\pi) \kappa(d_\pi, d)}$.

The proof for this result, as well as the remaining proofs, are in Appendix D. The previous proofs were included in the main body, as they provide useful intuition.

There is an important relationship between quality of the solution and convergence of TD, given by $c_{\mathcal{F},d}$. If $c_{\mathcal{F},d} < 1$, then the projected Bellman operator is a contraction. Correspondingly, the update underlying TD converges. Off-policy TD, therefore, can converge, as long as d or \mathcal{F} are chosen such that $c_{\mathcal{F},d} < 1$. Additionally, notice that we only have bounds on the solution quality under this same criteria. This suggests the following pessimistic conclusion: gradient TD methods, even if they ensure convergence, may not ensure convergence to a quality solution. Therefore, we should eschew them altogether and instead should focus on controlling the weighting d or the function space \mathcal{F} .

However, this conclusion only arises due to the looseness of the bound. The theorem statement requires $c_{\mathcal{F},d} < 1$ for all of \mathcal{F} . The proof itself only requires this contraction property for $v_{\mathbf{w}_{\mathcal{F},d}}$ and $\Pi_{\mathcal{F},d} v_\pi$. The projected operator could be an expansion during the optimization and the final quality of the solution under gradient descent will still be good as long as $c_{\mathcal{F}_{\text{opt}},d} < 1$ for $\mathcal{F}_{\text{opt}} \subset \mathcal{F}$ a small subset of \mathcal{F} that contains $v_{\mathbf{w}_{\mathcal{F},d}}$ and $\Pi_{\mathcal{F},d} v_\pi$.

Corollary 8 *Assume there is a convex subset $\mathcal{F}_{\text{opt}} \subset \mathcal{F}$ where $v_{\mathbf{w}_{\mathcal{F},d}}, \Pi_{\mathcal{F},d} v_\pi \in \mathcal{F}_{\text{opt}}$ and $c_{\mathcal{F}_{\text{opt}},d} < 1$. Then*

$$\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d \leq C(d, \mathcal{F}_{\text{opt}}) \|\Pi_{\mathcal{F},d} v_\pi - v_\pi\|_d.$$

An example of a problem that satisfies $c_{\mathcal{F}_{\text{opt}},d} < 1$ but not $c_{\mathcal{F},d} < 1$ is Baird’s counterexample (Baird, 1995). In fact, $v_\pi \in \mathcal{F}$, and $\|v_{\mathbf{w}_{\mathcal{F},d}} - v_\pi\|_d = 0$. Rather, the initialization is

started in a part of the space where the projected Bellman operator is an expansion, causing off-policy TD to diverge.

On the other hand, Kolter’s counterexample is one where no such convex subset exists. The distribution d can be chosen to make \mathbf{A} more and more singular, correspondingly causing the TD solution $\mathbf{w} = \mathbf{A}^{-1}\mathbf{b}$ to grow. The distance between the solution of the $\overline{\text{PBE}}$ and the optimal value function can be arbitrarily large and no such convex set \mathcal{F}_{opt} can exist.

Overall, these results suggest that we should combine all our of techniques to improve convergence and solution quality. The choice of \mathcal{F} and d should be considered, particularly to improve solution quality. If our function class contains the true value function v_π , then the approximation error $\|\Pi_d v_\pi - v_\pi\|_d = 0$ and the $\overline{\text{VE}}$ is zero regardless of the operator constant. Even in this ideal scenario, TD methods can diverge. Gradient methods should be used to avoid convergence issues, because the projected Bellman operator may not be a contraction in all parts of the space. Gradient methods can help us reach a part of the space \mathcal{F}_{opt} where $c_{\mathcal{F}_{\text{opt}},d} < 1$, and so provide an upper bound on the quality of the solution. An exciting open question remains as how to select \mathcal{F} and d , to obtain $c_{\mathcal{F}_{\text{opt}},d} < 1$.

6.2 Upper Bound on $\overline{\text{VE}}$ when $\mathcal{H} \supseteq \mathcal{F}$

We next consider more general projections, i.e., for any function space $\mathcal{H} \supseteq \mathcal{F}$, that includes the $\overline{\text{BE}}$ and $\overline{\text{PBE}}$ as special cases. We show two different approaches to upper bounding the $\overline{\text{VE}}$: re-expressing the objective using oblique projections and characterizing the difference to the $\overline{\text{BE}}$, for which there is a straightforward bound on the $\overline{\text{VE}}$.

6.2.1 THE $\overline{\text{PBE}}$ AS AN OBLIQUE PROJECTION

The goal of this section is to rewrite the $\overline{\text{PBE}}$ as an oblique projection, which then gives a generic bound on the value error. The bound relies on the norm of the oblique projection, which is not easily computable. But, it provides a potential direction for special cases where the norm of this projection might be able to be simplified.

We start by re-expressing the generalized $\overline{\text{PBE}}$ as a weighted $\overline{\text{VE}}$, using the same approach as Schoknecht (2003) and Scherrer (2010). Notice first that $v_\pi = (I - \mathcal{P}_{\pi,\gamma})^{-1}r_\pi$. Then the generalized $\overline{\text{PBE}}$ for any $v_w \in \mathcal{F}$, written in projection form, is

$$\begin{aligned} \|\Pi_{\mathcal{H},d}(Tv_w - v_w)\|_d^2 &= \|\Pi_{\mathcal{H},d}(r_\pi + \mathcal{P}_{\pi,\gamma}v_w - v_w)\|_d^2 \\ &= \|\Pi_{\mathcal{H},d}(r_\pi - (I - \mathcal{P}_{\pi,\gamma})v_w)\|_d^2 \\ &= \|\Pi_{\mathcal{H},d}[(I - \mathcal{P}_{\pi,\gamma})v_\pi - (I - \mathcal{P}_{\pi,\gamma})v_w]\|_d^2 &> r_\pi = (I - \mathcal{P}_{\pi,\gamma})v_\pi \\ &= \|\Pi_{\mathcal{H},d}(I - \mathcal{P}_{\pi,\gamma})(v_\pi - v_w)\|_d^2 \\ &= \|v_\pi - v_w\|_H^2 &> H \stackrel{\text{def}}{=} (I - \mathcal{P}_{\pi,\gamma})^\top \Pi_{\mathcal{H},d}^\top D \Pi_{\mathcal{H},d} (I - \mathcal{P}_{\pi,\gamma}) \end{aligned}$$

Minimizing the generalized $\overline{\text{PBE}}$ therefore corresponds to minimizing the $\overline{\text{VE}}$ with a reweighting over states that may no longer be diagonal, as H is not a diagonal matrix. In fact, we can see that the solution to the generalized $\overline{\text{PBE}}$ is a projection of v_π onto set \mathcal{F} under weighting H , namely $v = \Pi_{\mathcal{F},H}v_\pi$. A projection under such a non-diagonal weighting is called an oblique projection. Using this form, we can obtain an upper bound using a similar approach to (Scherrer, 2010, Proposition 3), with proof in Appendix D.

Theorem 9 *If $\mathcal{H} \supseteq \mathcal{F}$, then the solution $v_{\mathbf{w}_{\mathcal{H},d}}$ to the generalized \overline{PBE} satisfies*

$$\|v_\pi - v_{\mathbf{w}_{\mathcal{H},d}}\|_d \leq \|\Pi_{\mathcal{F},H}\|_d \|v_\pi - \Pi_{\mathcal{F},d}v_\pi\|_d. \quad (26)$$

We can next extend a previous result, which expressed this projection under linear function approximation for both the TD fixed point and the solution to the \overline{BE} . We can now more generally express the oblique projection for interim \mathcal{H} , with proof in Appendix D.

Corollary 10 *Assume \mathcal{F} is the space of linear functions with features \mathbf{x} and \mathcal{H} the space of linear functions with features ϕ , with $\mathcal{F} \subseteq \mathcal{H}$. Then the solution to the \overline{PBE} is*

$$v_{\mathbf{w}_{\mathcal{H},d}} = \mathbf{X}\mathbf{w}_{\mathcal{H},d} \quad \text{for } \mathbf{w}_{\mathcal{H},d} = (M^\top(I - \mathcal{P}_{\pi,\gamma})\mathbf{X})^{-1}M^\top r_\pi \quad (27)$$

for $M \stackrel{\text{def}}{=} \Pi_{\mathcal{H},d}^\top D(I - \mathcal{P}_{\pi,\gamma})\mathbf{X}$. Further, $v_{\mathbf{w}_{\mathcal{H},d}} = \Pi_{\mathcal{F},H}v_\pi$ for

$$\Pi_{\mathcal{F},H} = \mathbf{X}(M^\top(I - \mathcal{P}_{\pi,\gamma})\mathbf{X})^{-1}M^\top(I - \mathcal{P}_{\pi,\gamma}). \quad (28)$$

6.2.2 THE DISTANCE BETWEEN THE \overline{PBE} AND THE \overline{BE}

There is a well known upper bound on the \overline{VE} , in terms of the \overline{BE} . If we can characterize the distance of the \overline{PBE} to the \overline{BE} , we could potentially exploit this known result to upper bound the \overline{VE} in terms of the \overline{PBE} . Such a characterization should be possible because, for a sufficiently large \mathcal{H} , the \overline{PBE} is equivalent to the \overline{BE} . For smaller \mathcal{H} , we can view the \overline{PBE} as an approximation to the \overline{BE} , and so can characterize that approximation error.

For a given $v \in \mathcal{F}$, let $h_v^*(s) = \mathbb{E}[\delta(v)|S = s]$ and define

$$\text{ApproxError}(\mathcal{H}, v) \stackrel{\text{def}}{=} \min_{h \in \mathcal{H}} \|h_v^* - h\|_d$$

with worst-case approximation error

$$\text{ApproxError}(\mathcal{H}) \stackrel{\text{def}}{=} \max_{v \in \mathcal{F}} \text{ApproxError}(\mathcal{H}, v) \quad (29)$$

where we overload notation because the distinction is clear from the arguments given. Then we can obtain the following result, with proof in Appendix D.

Theorem 11 *For any $v \in \mathcal{F}$,*

$$\begin{aligned} \|v_\pi - v\|_d &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \underbrace{\|\mathcal{T}v - v\|_d}_{\overline{BE}} \\ &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \left(\underbrace{\|\Pi_{\mathcal{H},d}\mathcal{T}v - v\|_d}_{\overline{PBE}} + \text{ApproxError}(\mathcal{H}, v) \right) \\ &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \left(\|\Pi_{\mathcal{H},d}\mathcal{T}v - v\|_d + \text{ApproxError}(\mathcal{H}) \right). \end{aligned}$$

Note that if $c_d < 1$, then $\|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \leq (1 - c_d)^{-1}$.

The upper bound depends on the value of the $\overline{\text{PBE}}$ and the approximation error of \mathcal{H} . For $\mathcal{H} = \mathcal{F}$, the $\overline{\text{PBE}}$ is zero but the approximation error is likely higher. As \mathcal{H} gets bigger, the approximation error gets smaller, but the $\overline{\text{PBE}}$ is also larger. Once \mathcal{H} is big enough to include h^* for a given v , then the approximation error is zero but the $\overline{\text{PBE}}$ is at its maximum, which is to say it is equal to the $\overline{\text{BE}}$. This upper bound is likely minimized for an interim value of \mathcal{H} , that balances between the error from the $\overline{\text{PBE}}$ and the approximation error.

It is important to recognize that the approximation error for \mathcal{F} and \mathcal{H} can be quite different. For example, if $v_\pi \in \mathcal{F}$, then $\text{ApproxError}(\mathcal{H}, v_\pi) = 0$ if $0 \in \mathcal{H}$. The function that returns zero for every state should be in \mathcal{H} to ensure zero approximation error, but need not be in \mathcal{F} . For many convex function spaces we consider, all convex combinations of functions are in \mathcal{F} , and it is likely that we have 0 in \mathcal{F} . But nonetheless, there are likely instances where \mathcal{F} contains near optimal value functions, but that same set produces high approximation error for \mathcal{H} . An important future direction is to better understand the differences in function spaces needed for \mathcal{F} and \mathcal{H} .

7. Algorithms for the Generalized $\overline{\text{PBE}}$

The linear $\overline{\text{PBE}}$ is often optimized using gradient correction algorithms as opposed to saddlepoint methods. The canonical methods are TDC (gradient corrections) and GTD2 (saddlepoint), where TDC has been consistently shown to perform better than GTD2 (White and White, 2016; Ghiassian et al., 2020). We show that similar gradient correction algorithms arise for the generalized $\overline{\text{PBE}}$, and discuss such an algorithm called TDRC.

7.1 Estimating the Gradient of the Generalized $\overline{\text{PBE}}$

To see why (at least) two classes of algorithms arise, consider the gradient for the generalized $\overline{\text{PBE}}$, for a given $h(s) \approx \mathbb{E}_\pi[\delta(\mathbf{w}) \mid S = s]$ with a stochastic sample $\delta(\mathbf{w})$ from $S = s$:

$$-\nabla_{\mathbf{w}}\delta(\mathbf{w})h(s) = h(s)[\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) - \gamma\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S')].$$

This is the standard *saddlepoint update*. The key issue with this form is that any inaccuracy in h has a big impact on the update of \mathbf{w} , and h can be highly inaccurate during learning. Typically, it is initialized to zero, and so it multiplies the update to the primary weights by a number near zero, making learning slow.

The *gradient correction update* is preferable because it relies less on the accuracy of h . The first term uses only the sampled TD-error.

$$\Delta\mathbf{w} \leftarrow \delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) - h_\theta(s)\gamma\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S')$$

where $\Delta\theta \leftarrow (\delta(\mathbf{w}) - h_\theta(s))\nabla_\theta h_\theta(s)$ just like the saddlepoint update. But, the update is biased because it assumes it has optimal $h^* \in \mathcal{H}$ for part of the gradient. To see why, we extend the derivation for the linear setting.

$$\begin{aligned} -\nabla_{\mathbf{w}}\delta(\mathbf{w})h(s) &= h(s)[\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) - \gamma\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S')] \\ &= h(s)\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) - h(s)\gamma\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S') \\ &= (h(s) - \delta(\mathbf{w}) + \delta(\mathbf{w}))\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) - h(s)\gamma\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S') \\ &= \delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) + (h(s) - \delta(\mathbf{w}))\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(s) - h(s)\gamma\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S') \end{aligned}$$

This resembles the gradient correction update, except it has extra term $(h(s) - \delta(\mathbf{w})) \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s)$. In the linear setting, if we have the linear regression solution for h^* with parameters θ , then this second term is zero in expectation. This is because $\nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) = \mathbf{x}(s)$, giving

$$\mathbb{E}_{\pi}[(h(s) - \delta(\mathbf{w})) \nabla_{\mathbf{w}} v(s, \mathbf{w}) \mid S = s] = \mathbf{x}(s) \mathbf{x}(s)^{\top} \theta - \mathbf{x}(s) \mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s]$$

and so in expectation across all states, because $\theta = \mathbb{E}[\mathbf{x}(S) \mathbf{x}(S)^{\top}]^{-1} \mathbb{E}[\mathbf{x}(S) \delta]$, we get that

$$\begin{aligned} \mathbb{E}[(h(S) - \delta(\mathbf{w})) \nabla_{\mathbf{w}} v(S, \mathbf{w})] &= \mathbb{E}[\mathbf{x}(S) \mathbf{x}(S)^{\top}] \mathbb{E}[\mathbf{x}(S) \mathbf{x}(S)^{\top}]^{-1} \mathbb{E}[\mathbf{x}(S) \delta(\mathbf{w})] - \mathbb{E}[\mathbf{x}(S) \delta(\mathbf{w})] \\ &= \mathbb{E}[\mathbf{x}(S) \delta(\mathbf{w})] - \mathbb{E}[\mathbf{x}(S) \delta(\mathbf{w})] = 0. \end{aligned}$$

Therefore, given the optimal $h \in \mathcal{H}$ for \mathcal{H} the set of linear functions, this term can be omitted in the stochastic gradient and still be an unbiased estimate of the full gradient.

More generally, the same reasoning applies if $h(s)$ can be re-expressed as a linear function of $\nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S)$. This provides further motivation for using features produced by the gradient of the values, as in the nonlinear $\overline{\text{PBE}}$, to estimate h . Another choice is to use the features in the last layer of the neural network used for $\hat{v}_{\mathbf{w}}(S)$. Because the output is a linear weighting of features from the last layer, $\nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S)$ includes this last layer as one part of the larger vector. A head for h can be added to the neural network, where h is learned as a linear function of this layer. Its updates do not influence the neural network itself—gradients are not passed backwards through the network—to ensure it is a linear function of the last layer.

Unlike the saddlepoint update, however, the gradient correction update is no longer a straightforward gradient update, complicating analysis. It is possible, however, to analyze the dynamical system underlying these updates.

The asymptotic solution does not require the omitted term, under certain conditions on h , as discussed above. If the dynamical system moves towards this stable asymptotic solution, then convergence can be shown. The TDC update relies on just such a strategy: the joint update is rewritten as a linear system, that is then shown to be a contraction that iterates towards a stable solution (Maei, 2011). The extension to the nonlinear setting is an important open problem.

Theoretical work for these updates under nonlinear function approximation has been completed, for an algorithm called SBEED (Dai et al., 2018). SBEED uses a gradient correction update, but the theory is for the saddlepoint version. This work investigates a slightly different saddlepoint update, learning $h(s) \approx \mathbb{E}[R + \gamma \hat{v}_{\mathbf{w}}(S') \mid S = s]$, instead of estimating the entire TD error. A TD-error estimate can be obtained using $h(s) - \hat{v}_{\mathbf{w}}(s)$. Even under this alternate route, it was reconfirmed that the gradient correction update was preferable—since the final proposed algorithm used this form—but the theory non-trivial.

Remark: Another way to interpret gradient correction algorithms is as approximations of the gradient of the $\overline{\text{BE}}$. We can consider two forms for the negative gradient of the $\overline{\text{BE}}$:

$$\begin{aligned} &\mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s] \mathbb{E}_{\pi}[\nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) - \gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S') \mid S = s] \\ \text{or} \quad &\mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s] \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) - \mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s] \mathbb{E}_{\pi}[\gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S') \mid S = s] \end{aligned}$$

because $\hat{v}_{\mathbf{w}}(s)$ is not random. We can estimate the first form of the gradient using an estimate $h(s) \approx \mathbb{E}_{\pi}[\delta(\mathbf{w}) \mid S = s]$: $h(s)(\nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) - \gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S'))$. This corresponds to a saddlepoint update. To estimate the second form of the gradient, notice that we do not

have a double sampling problem for the first term. This means we can use δ to compute an unbiased sample for the first term: $\delta \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) - h(s) \gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S')$. This strategy corresponds to the gradient correction update.

7.2 TDRC: A Practical Gradient-Based Prediction Algorithm

In this section we incorporate the above insights into an easy-to-use gradient-based algorithm called TD with Regularized Corrections (TDRC). Empirical work with TDC suggests that it performs well because the secondary stepsize is set small, and thus it is effectively using conventional off-policy TD (Ghiassian et al., 2020). TDRC builds on this insight, providing a method that performs similarly to TD when TD performs well, while maintaining the theoretical soundness of TDC. The key idea is to regularize h to be near zero, using an ℓ_2 regularizer, where if $h(s) = 0$ then the update corresponds to the TD update.

The motivation for this approach is that it can help learn h more quickly, without incurring much bias. In the linear setting with $\mathcal{F} = \mathcal{H}$, the parameters $\theta = 0$ for h^* at the TD fixed point. Applying ℓ_2 regularization to the secondary weights, therefore, incurs little bias asymptotically. But, adding a strongly convex regularizer can improve the convergence rate of the secondary variable, and reduce variance of the estimator.

In fact, this perspective on $h(s)$ helps to explain the unreasonably good performance of TD in many practical settings. TD is equivalent to TDC with $h(s) = 0$, which is a zero variance but highly biased estimate of the expected TD error. Asymptotically as our estimates for the primary variable improve, the bias of this heuristic choice of h decreases until we converge to the same fixed-point as TDC (in the cases where TD converges).

The TDRC update equations with importance sampling ratios $\rho(s, a) = \frac{\pi(a|s)}{b(a|s)}$ are

$$\begin{aligned} \Delta \mathbf{w} &\leftarrow \rho(s, a) \delta(\mathbf{w}) \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(s) - \rho(s, a) h_{\theta}(s) \gamma \nabla_{\mathbf{w}} \hat{v}_{\mathbf{w}}(S') \\ \Delta \theta &\leftarrow \rho(s, a) (\delta(\mathbf{w}) - h_{\theta}(s)) \nabla_{\theta} h_{\theta}(s) - \beta \theta \end{aligned}$$

where \mathbf{w} and θ are parameters of $v : \mathcal{S} \rightarrow \mathbb{R}$ and $h : \mathcal{S} \rightarrow \mathbb{R}$ respectively. Notice that the only difference between TDRC and TDC is the inclusion of the $-\beta\theta$ term in the update for θ . It can be shown in the linear setting with $\mathcal{F} = \mathcal{H}$ that the fixed-points of the TDRC dynamical system are the same as the TDC system and that, under some conditions on the learning rates, TDRC converges to its fixed-point on any MDP (Ghiassian et al., 2020).

By adding an ℓ_2 regularizer with regularization parameter β to TDC, we can interpolate between TD and TDC. With β very large the solution set of $h(s)$ becomes diminishingly small until it contains only the point $h(s) = 0$ as $\beta \rightarrow \infty$. On the other hand, as β approaches zero, then the solution set of $h(s)$ becomes that of the $\overline{\text{PBE}}$ and TDRC approaches TDC. Because β scales between two known good algorithms, this suggests that the impetus to highly tune β is small: many choices of β should yield a reasonably performing algorithm.

There are many options to approximate both $v(s)$ and $h(s)$ using neural networks. In our experiments, we chose to use a single network with two sets of outputs, one to predict $v(s)$ and the other to predict $h(s)$. To avoid difficulties with balancing between two loss functions, we only use the gradients from the primary update to adjust the weights of the network and use the gradients from the secondary update only to adjust the weights for that head. The resulting architecture for estimating action values is shown in Figure 6.

8. Extending to Control

The previous development was strictly for policy evaluation. The formulation of a sensible generalized $\overline{\text{PBE}}$ for control, however, can be obtained using a similar route. The conjugate form has already been used to develop a novel control algorithm for nonlinear function approximation, called SBEED (Dai et al., 2018). The SBEED algorithm explicitly maintains a value function and policy, to incorporate entropy regularization, and uses the gradient correction update. We develop an alternative control algorithm that learns only action-values and uses the gradient correction update.

8.1 The Control Objective

Our goal is to approximate $q^* : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, the action-values for the optimal (soft) policy. Instead of the Bellman operator, we use the Bellman optimality operator or generalizations that use other forms of the max but are still guaranteed to be contractions, like the mellow-max operator (Asadi and Littman, 2017). Let m be the given max operator, that takes action-values and returns a (soft) greedy value. In Q-learning, we use a hard-max $m(q(s, \cdot)) = \max_a q(s, a)$ and in mellow-max, $m(q(s, \cdot)) = \tau^{-1} \log \left(\frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \exp(\tau q(s, a)) \right)$. As $\tau \rightarrow \infty$, the mellow-max operator approaches the hard-max operator. The Bellman optimality operator \mathcal{T}_m corresponds to

$$(\mathcal{T}_m q)(s, a) \stackrel{\text{def}}{=} \mathbb{E} [R + \gamma(S, A, S') m(q(S', \cdot)) \mid S = s, A = a] \quad (30)$$

where the expectation is over next state and reward. The fixed point of this operator is q^* .

To learn $\hat{q}_w(s, a)$ approximating $q^*(s, a)$, we define the $\overline{\text{BE}}$ for control

$$\overline{\text{BE}}(\mathbf{w}) \stackrel{\text{def}}{=} \sum_{s, a} d(s, a) \mathbb{E} [\delta(\mathbf{w}) \mid S = s, A = a]^2 \quad (31)$$

where $\delta(\mathbf{w}) \stackrel{\text{def}}{=} R + \gamma(S, A, S') m(\hat{q}_w(S', \cdot)) - \hat{q}_w(S, A)$ and $d : \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$ is some weighting. We overload the notation for the weighting d , to make the connection to the previous objectives clear. We can rewrite this objective using conjugates, to get

$$\begin{aligned} \overline{\text{BE}}(\mathbf{w}) &= \sum_{s \in \mathcal{S}, a \in \mathcal{A}} d(s, a) \max_{h \in \mathbb{R}} (2\mathbb{E}[\delta(\mathbf{w}) \mid S = s, A = a] h - h^2) \\ &= \max_{h \in \mathcal{F}_{\text{all}}} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} d(s, a) (2\mathbb{E}[\delta(\mathbf{w}) \mid S = s, A = a] h(s, a) - h(s, a)^2). \end{aligned}$$

As before, this maximization can be rewritten as a minimization,

where the optimal $h^*(s, a) = \mathbb{E}[\delta(\mathbf{w}) \mid S = s, A = a]$. This equivalence is true for the hard-max operator or the mellow-max, even though the operator is no longer smooth. Finally, in practice, we will learn an approximate h , from the set \mathcal{H} , resulting in a $\overline{\text{PBE}}$ for control:

$$\overline{\text{PBE}}(\mathbf{w}) \stackrel{\text{def}}{=} \max_{h \in \mathcal{H}} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} d(s, a) (2\mathbb{E}[\delta(\mathbf{w}) \mid S = s, A = a] h(s, a) - h(s, a)^2).$$

This objective is the first generalized $\overline{\text{PBE}}$ for learning action-values for control.




Figure 6: Visualization of the neural network architecture used to approximate $\hat{q}(s, a)$ and $h(s, a)$ for the QRC algorithm. TDRC uses the same network architecture, except with only one output for each head of the network predicting $\hat{v}(s)$ and $h(s)$ respectively. The green shaded regions represent gradient information from the update to $\hat{q}(s, a)$ and the red shaded region represents the gradient information from the update $h(s, a)$. Notice only $\hat{q}(s, a)$ modifies the weights of the earlier layers of the network.

The algorithm is a simple modification of the policy evaluation algorithms above, described briefly here and expanded upon more in the next section. The update to $h(s, a)$ is still a gradient of a squared error to the TD error. The saddlepoint gradient update, for a given $h(s, a)$ with a stochastic sample $\delta(\mathbf{w})$ from $S = s, A = a$:

$$-\nabla_{\mathbf{w}}\delta(\mathbf{w})h(s, a) = h(s, a)[\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a) - \gamma\nabla_{\mathbf{w}}m(\hat{q}_{\mathbf{w}}(S', \cdot))]$$

with gradient correction form

$$\delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a) - \gamma h(s, a)\nabla_{\mathbf{w}}m(\hat{q}_{\mathbf{w}}(S', \cdot)).$$

Both involve taking a gradient through the max operator m . For the hard-max operator, this results in a subgradient. The mellow-max operator, on the other hand, is differentiable with derivative

$$\frac{\partial}{\partial w_i}m(\hat{q}_{\mathbf{w}}(s, \cdot)) = \frac{1}{\sum_{a \in \mathcal{A}} \exp(\tau \hat{q}_{\mathbf{w}}(s, a))} \sum_{a \in \mathcal{A}} \exp(\tau \hat{q}_{\mathbf{w}}(s, a)) \frac{\partial}{\partial w_i} \hat{q}_{\mathbf{w}}(s, a)$$

We can reason similarly about the validity of the gradient correction update.

$$\begin{aligned} -\nabla_{\mathbf{w}}\delta(\mathbf{w})h(s, a) &= (h(s, a) - \delta(\mathbf{w}) + \delta(\mathbf{w}))\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a) - h(s, a)\gamma\nabla_{\mathbf{w}}m(\hat{q}_{\mathbf{w}}(S', \cdot)) \\ &= \delta(\mathbf{w})\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a) + (h(s, a) - \delta(\mathbf{w}))\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a) - h(s, a)\gamma\nabla_{\mathbf{w}}m(\hat{q}_{\mathbf{w}}(S', \cdot)). \end{aligned}$$

As before, we can conclude that we can drop this second term, as long as the optimal $h \in \mathcal{H}$ is representable as a linear function of $\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a)$. The fixed point for the gradient correction updates that drop the term $(h(s, a) - \delta(\mathbf{w}))\nabla_{\mathbf{w}}\hat{q}_{\mathbf{w}}(s, a)$ will still converge to the same fixed point, if they converge. The key question that remains is, if the dynamical system produced by these equations does in fact converge.

8.2 QRC: A Practical Gradient-Based Control Algorithm

In this section, we modify the TDRC algorithm specified in Section 7.2 for control. Let θ_{t, A_t} be the weights for only the secondary head predicting $h(S_t, A_t)$, \mathbf{x}_t the last layer of the

network—giving the features for h —and \mathbf{w}_t all of the remaining weights in the network. We use the mellowmax operator, to get the following updates for the QRC algorithm

$$\begin{aligned}\delta_t &= R_{t+1} + \gamma m(\hat{q}_{\mathbf{w}}(S_{t+1}, \cdot)) - \hat{q}_{\mathbf{w}}(S_t, A_t), \text{ and } h_t = \theta_{t,A_t}^\top \mathbf{x}_t \\ \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \nabla_{\mathbf{w}} \hat{q}_{\mathbf{w}}(S, A) - \alpha \gamma h_t \nabla_{\mathbf{w}} m(\hat{q}_{\mathbf{w}}(S', \cdot)) \\ \theta_{t+1,A_t} &\leftarrow \theta_{t,A_t} + \alpha [\delta_t - h_t] \mathbf{x}_t - \alpha \beta \theta_{t,A_t}\end{aligned}$$

QRC is similar to SBEED (Dai et al., 2018), but has two key differences. The first is that SBEED learns a state-value function and an explicit policy. We learn action-values and use a direct mellowmax on the action-values to compute the policy. The other key difference is that the SBEED update uses a different form of gradient correction, which estimates only the Bellman step instead of the TD error in order to correct the gradient. Recall that QRC interpolates between a gradient correction update at one extreme of its hyperparameter ($\beta = 0$) and Q-learning at the other extreme (large β). SBEED, on the other hand, interpolates between a gradient correction update and a residual-gradient method minimizing the mean squared TD error. When QRC converges, it converges to the same fixed-point for all values of its hyperparameter. SBEED, however, interpolates between the fixed-point for the $\overline{\text{PBE}}$ and the mean squared TD error.

9. Empirical Investigation of the Generalized $\overline{\text{PBE}}$ for Control

In this section, we empirically investigate QRC. We first provide a comparison of QRC with Q-learning and SBEED across four benchmark domains. Then we delve more deeply into the design choices in QRC. We compare using gradient corrections within QRC to the saddlepoint form in Section 9.4. We then investigate how QRC performs with separate bases for h , including a more powerful basis for h , in Section 9.5.

9.1 Benchmark Environments and Experimental Design

We used four simulation domains with neural network function approximation. We chose simulation domains with a sufficiently small state dimension to efficiently compare algorithms across many different random initializations of neural network, while still computing the required number of experiment repetitions required for statistical significance. On the other hand, to tease out differences between algorithms, we require domains with sufficiently complex learning dynamics. We chose three classic control domains known to be challenging when approximation resources and agent-environment interactions are limited: Acrobot (Sutton, 1996), Cart Pole (Barto et al., 1983), and Mountain Car (Moore, 1990). We also used Lunar Lander (Brockman et al., 2016) to investigate performance in a domain with a dense reward function and moderately higher-dimensional state.

The network architectures were as follows. For Acrobot and Mountain Car, we used two layer fully-connected neural networks with 32 units in each layer and a ReLU transfer. The output layer has an output for each action-value and uses a linear transfer. For the Cart Pole and Lunar Lander domains, we used the same architecture except with 64 units in each hidden layer. We use a shared network with multiple heads for all algorithms unless

. Code for all experiments is available at <https://github.com/rlai-lab/Generalized-Projected-Bellman-Errors>

otherwise specified. In experiments with policy-gradient-based methods the parameterized policy uses an independent neural network. We do not use target networks.

We swept consistent values of the hyperparameters for every experiment. We swept the stepsize parameter over a wide range $\alpha \in \{2^{-12}, 2^{-11}, \dots, 2^{-7}\}$ for every algorithm. For algorithms which chose a stepsize on the boundary of this range—for instance, GQ often chose the smallest stepsize—we performed a one-off test to ensure that the range was still representative of the algorithm’s performance. All algorithms used mellowmax, with τ swept in the range $\tau \in \{0, 10^{-4}, 10^{-3}, \dots, 10^0\}$, including 0 to allow algorithms to choose to use a hard-max. Algorithms based on the SBEED update have an additional hyperparameter η which interpolates between the gradient correction update and a residual gradient update. For all experiments we swept values of $\eta \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0\}$ and the ratio between the actor and critic stepsizes $\nu \in \{2^{-4}, 2^{-3}, \dots, 2^1\}$, often giving SBEED algorithms twenty-four times as many parameter permutations to optimize over compared to other algorithms. Likewise, we allowed saddlepoint methods (GQ) to optimize over the regularization parameter $\beta \in \{0, 0.5, 1, 1.5\}$, to give them an opportunity to perform well.

The remaining hyper-parameters were not swept, but instead set to reasonable defaults. We used a replay buffer to store the last 4000 transitions, then sampled 32 independent transitions without replacement to compute mini-batch averaged updates. We used the ADAM optimizer (Kingma and Ba, 2015) for all experiments with the default hyperparameters, a momentum term of $\beta_1 = 0.9$ and a squared-gradient term of $\beta_2 = 0.999$. We additionally tested Stochastic Gradient Descent and RMSProp and found that most conclusions remain the same, so choose not to include these results to focus the presentation of results. For each of the four domains we use a discount factor of $\gamma = 0.99$ and cutoff long-running episodes at 500 steps for Acrobot and Cart Pole and 1000 steps for Mountain Car. On episode cutoff events, we do not make an update to the algorithm weights to avoid bootstrapping over this imaginary transition and on true episode termination steps we update with $\gamma = 0$.

We use a non-conventional performance measure to more fairly report algorithm performance. A common performance metric is to report the cumulative reward at the end of each episode, running each algorithm for a consistent number of episodes. This choice causes algorithms to have different amounts of experience and updates. Some algorithms use more learning steps in the first several episodes and achieve higher asymptotic performance because they effectively learned for more steps. We instead report the cumulative reward from the current episode on each step of the current episode. For example in Mountain Car, if the k th episode takes 120 steps, then we would record -120 for each step of the episode. We then run each algorithm for a fixed number of steps instead of a fixed number of episodes, so that each algorithm gets the same number of learning steps and a consistent amount of data from the environment. We record performance over 100,000 steps, recorded every 500 steps—rather than every step—to reduce storage costs.

To avoid tuning the hyperparameters for each algorithm for every problem, we start by investigating a single set of hyperparameters for each algorithm across all four benchmark domains. We evaluate the hyperparameters according to mean performance over runs, for each domain. We then use a Condorcet voting procedure to find the single hyperparameter setting that performs best across all domains.




Figure 7: *Learning curves using the best performing hyperparameters across 4 domain.* The learning curves above are averaged over 100 independent runs and shaded regions correspond to one standard error.




Figure 8: *Distribution of average returns over hyperparameter settings for each benchmark domain.* The vertical axis represents the average performance of each hyperparameter setting (higher is better) and the width of each curve represents the proportion of hyperparameters which achieve that performance level, using a fitted kernel density estimator. The solid horizontal bars show the maximum, mean, and minimum performance respectively and the dashed horizontal bar represents the median performance over hyperparameters. QRC in blue generally performs best and exhibits less variability across hyperparameter settings.

9.2 Overall Results in the Benchmark Environments

Figure 7 shows the learning curves for each algorithm with the single best performing hyperparameter setting across domains. QRC was the only algorithm to consistently be among the best performing algorithms on every domain and was the only algorithm with a single hyperparameter setting that could solve all four domains. Although SBEED was given twenty-four times as many hyperparameter combinations to optimize over, its performance was consistently worse than all other benchmark algorithms. This suggests that the voting procedure was unable to identify a single hyperparameter setting that was consistently good across domains. We additionally include the nonlinear control variant of TDC with a locally linear projection which we call QC-LL (Maei et al., 2009). QC-LL performed well on the two simpler domains, Acrobot and Cart Pole, but exhibited poor performance in the two more challenging domains. We report other voting procedures as well as the performance of the best hyperparameters tuned for each domain independently in Appendix H.

To understand how hyperparameter selection impacts the performance of each algorithm, we visualize the performance distribution over hyperparameters in Figure 8. Ideally, we prefer the variability in performance across hyperparameters to be small and the distribution to be concentrated around a single performance level near the top of the plot. Plots with an hourglass shape represent bimodal distributions where several hyperparameter values perform well and several perform poorly. Plots where the mean and median horizontal markers are quite separated indicate highly skewed distributions.

All algorithms exhibit different hyperparameter performance distributions across the four domains. The SBEED algorithm often has one (or a few) hyperparameter setting(s) which perform well, especially in the Cart Pole domain. This suggests that SBEED is highly sensitive to its hyperparameters. and even small deviations from the ideal hyperparameters can lead to very low performance. Q-Learning generally exhibits wide spread of performance over hyperparameters, with a high skew in two of the problem settings. QRC has much lower spread of performance over hyperparameters often with the bulk of the distribution located near the highest performing hyperparameter setting.

9.3 Investigating Variability Across Runs

In the previous section we investigated performance averaged across runs; in this section, we investigate how much the algorithms vary across runs. Two algorithms could have similar average returns across runs, even if one has many poor runs and many high performing runs and the other has most runs that have reasonable performance.

To observe these differences, we visualize the distribution of returns across 100 runs, as in Figure 9. The horizontal axis shows the return achieved for a given run, averaged over the last 25% of steps. The vertical axis is the proportion of independent runs that achieve that return, using both a histogram with 30 bins and a fitted Gaussian kernel density estimator. In this example plot, we can see that QRC and Q-learning have similar distributions over returns across runs, QC-LL is peaked at a lower return and SBEED exhibits some bimodality, with several reasonable runs and a few runs with low return.




Figure 9: *The performance distribution over runs* for the best performing hyperparameter settings for each algorithm on Lunar Lander. The horizontal axis represents the average episodic return over the last 25% of steps. The vertical axis for each subplot represents the proportion of trials that obtained a given level of performance. The plot shows the empirical histogram and kernel density estimator for the performance distribution over 100 independent trials. Mass concentrated to the right indicates better performance.




Figure 10: *Sensitivity to stepsize parameter*. Distribution of the return per episode for the last 25% of episodes across choice of stepsize. Each row of this figure corresponds to the performance on each algorithm across domains for one value of the stepsize parameter. Each subplot is exactly like Figure 9: the distribution of performance for all four algorithms using a particular stepsize parameter value on a single domain. The highlighted plots in each column represent the best performing stepsize parameter value. QRC consistently exhibits a narrow distribution of performance where the bulk of the distribution is on the upper end of the performance metric (towards the right is better). Q-learning and Nonlinear QC both have wide performance distributions on all domains and exhibit bimodal distributions on Mountain Car. SBEED tends to exhibit bimodal performance often, with a non-trivial proportion of runs which fail to learn beyond random performance.

Because each algorithm prefers different stepsizes, we provide these distribution plots across stepsizes, optimizing the remaining hyperparameters for each stepsize. Figure 10 shows the distribution for every algorithm, across every domain and swept stepsize value. For every algorithm and domain, the distribution with the highest mean is highlighted with a bold color; distributions for all other stepsizes are shown with a faded color. QRC consistently has narrow performance distributions over runs. The worse performing hyperparameter settings for QRC often exhibit narrow distributions as well. In contrast, the other three algorithms demonstrate skewed performance distributions for their best hyperparameters, pulling the mean towards lower performance levels. SBEED often exhibits bimodality in its performance distributions, especially on the Mountain Car domain.³

9.4 Gradient Correction Updates versus Saddlepoint Updates

In this section we compare the saddlepoint and gradient correction forms of the gradient. In Section 7.1, we discussed these two strategies for estimating the gradient of the generalize PBE, and advocated for using the gradient correction form in QRC. To ablate this choice of

3. Our results with SBEED may appear pessimistic compared to results in the literature. It is possible previously published results with SBEED were achieved through domain specific hyper-parameter tuning. Further, our results were averaged over 100 runs, whereas prior work used five. We do not have access to the code used in prior work, and so we can only speculate.




Figure 11: *Comparing gradient correction-based updates (QRC) and saddlepoint methods (GQ, GQ-Grad).* GQ-Grad utilizes the gradient of v as features for the secondary variable, h . Allowing the saddlepoint methods to estimate $h(s)$ by using a linear function of the gradients of the primary variable yields slightly higher performance. Nonetheless, saddlepoint methods suffer from wide performance distributions with the bulk of the distribution being further left than the gradient correction-based updates.

update strategy, we fix all other design decisions. Figure 11 shows the saddlepoint method compared to the gradient correction method across our four domains.

The distribution plots suggest that the saddlepoint method *can* learn reasonable policies in some runs, however, the performance distributions tend to be skewed right with a bulk of the distribution around poor performance. QRC is restricted to regularization parameter $\beta = 1$; to give the saddlepoint method a greater chance of success, we additionally sweep β for it, giving it four times as many hyperparameter permutations. Even with the increased hyperparameter search space, it is clear that the saddlepoint method under-performs the gradient correction-based methods.

9.5 Using a Shared or Separate Basis for h

In this section we investigate the impact of using a shared network in QRC. By sharing a network, we effectively restrict $\mathcal{H} = \mathcal{F}$; by using a different network, \mathcal{H} is less restricted. A shared network is simpler and, in section 7.1, we show that defining h to be a function of the gradients for \hat{v}_w reduces the bias of gradient correction methods, such as QRC. However, in Section 6.2 we show a less restricted \mathcal{H} can improve performance. So, which is better? We compare QRC with a shared network, QRC with a separate network for h (QRC-Sep), and QRC with gradient features for h (QRC-Grad) to understand the impact of these choices.

To help remove the confounding variable of number of learnable parameters, we swept over network architectures for the separate network version of QRC. We kept the same general structure of the neural network by restricting the number of hidden units in both hidden layers to be the same, then swept over the number of units per layer: $n \in \{8, 16, 32, 64\}$ for the network for h in Acrobot and Mountain Car and $n \in \{16, 32, 64, 128\}$ for Cart Pole and Lunar Lander. At the smallest end of the range, QRC with separate networks has fewer

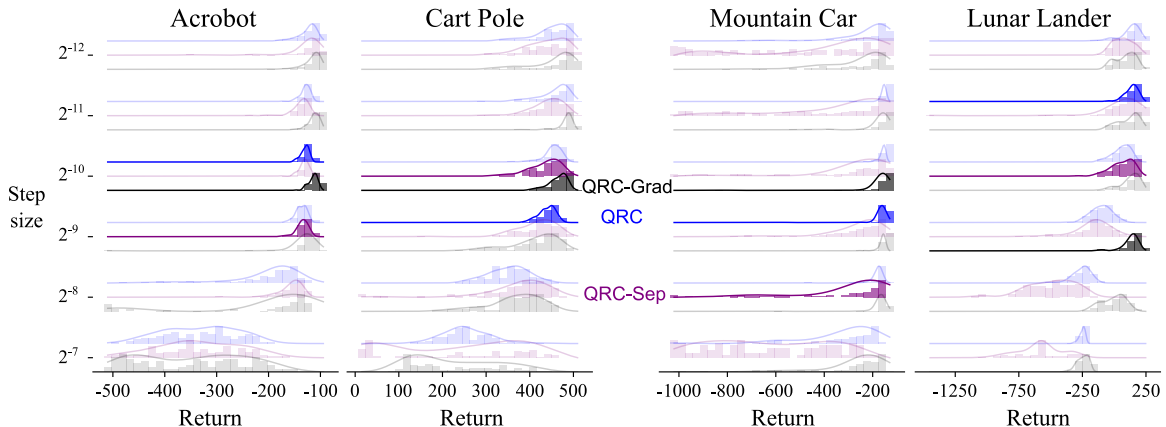


Figure 12: *How to represent h* : ablating the choice of basis function for the secondary variable, by comparing a shared network with two heads (QRC), two separate networks (QRC-Sep), and one network for the primary variable and a linear function of the primary variable’s gradients for the secondary (QRC-Grad).

learnable parameters than when using shared heads, and at the upper end of the range, it has considerably more learnable parameters.

Figure 12 summarizes the results using separate networks for the dual. Performance is comparable for each of the three methods, with QRC-Sep performing marginally worse overall and QRC-Grad performing marginally better. Notably the performance ordering appears to follow the bias of each method, with QRC-Sep having more bias due to needing to learn two separate networks, and QRC-Grad having slightly less bias due to h being a function of the gradients of \hat{v}_w . In general, the simpler choice of using a shared network appears to be as effective in these benchmark environments.

9.6 Using Parameterized Policies

In this section we attempt to better understand the poor performance of SBEED by investigating both action-value and actor-critic variants of the algorithms. Parameterized policies add a significant source of complexity both to the agent design and to the learning dynamics of the problem, making it challenging to directly compare SBEED to QRC. In this section we complete the square by introducing a parameterized policy version of TDRC, which can be viewed as an Actor-Critic method with a TDRC critic, and introduce an action-value version of SBEED. We show that the action-value version of SBEED significantly outperforms its parameterized policy counterpart, but still under-performs QRC in all domains and TDRC-PG in most domains.

The actor-critic version of TDRC uses the following update

$$\begin{aligned}
 \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \nabla_{\mathbf{w}} v(S_t, \mathbf{w}_t) - \alpha \gamma \theta_t^\top \mathbf{x}_t \nabla_{\mathbf{w}} v(S_{t+1}, \mathbf{w}_t) \\
 \theta_{t+1} &\leftarrow \theta_t + \alpha \left[\delta_t - \theta_t^\top \mathbf{x}_t \right] \mathbf{x}_t - \alpha \beta \theta_t \\
 \theta_{\pi, t+1} &\leftarrow \theta_{\pi, t} + \alpha \delta_t \nabla_{\theta_\pi} \ln \pi_{\theta_\pi}(A_t | S_t),
 \end{aligned}$$




Figure 13: *Impact of policy parameterization: action values versus parameterized policies.* The plots above compare the QRC and SBEED update functions both for learning state-value functions with parameterized policies and action-value functions with mellowmax.

which can be viewed as a standard actor-critic update with TDRC as the critic. Our design choices for TDRC-PG mirror those of SBEED, for instance using a two-headed neural network for the value function estimator and a separate neural network for the policy gradient estimator to avoid the need to balance between loss functions.

In Figure 13 we investigate the distribution of performance for each choice of update cross-producted with each choice of action-value or parameterized policy. The width of each violin plot represents the proportion of parameter settings whose area under the learning curve achieved the given performance level indicated on the vertical axis. In all cases, QRC had the most narrow performance distribution centered around the highest or near-highest reward. The parameterized-policy variants generally exhibit wider distributions and lower average AUC than their action-value counterparts, but occasionally achieve higher maximum performance for a single hyperparameter setting. The SBEED update generally performed worse than the QRC update across all four domains.

The empirical superiority of QRC over SBEED in our experiments is partially, but not fully, explained by learning parameterized policies. Because SBEED-Q outperforms SBEED overall, some of SBEED’s poor performance is due to the parameterized policy. SBEED-Q and QRC still perform differently in every domain, with SBEED-Q generally having much higher sensitivity to hyperparameters and occasionally achieving lower maximum performance than QRC, even after tuning over many more hyperparameter settings.

10. Conclusion

In this paper, we introduced the generalized $\overline{\text{PBE}}$ for off-policy value estimation, bringing together insights developed for both the linear $\overline{\text{PBE}}$ and the $\overline{\text{BE}}$. We investigated the quality of the $\overline{\text{PBE}}$ solution, both theoretically by bounding the $\overline{\text{VE}}$ and empirically under different state weightings and for different projection operators. We introduced an algorithm for minimizing the $\overline{\text{PBE}}$ both in prediction and control, incorporating empirical insights from optimizing the linear $\overline{\text{PBE}}$. Finally, we showed empirically that this control algorithm is less sensitive to its hyperparameters and exhibits more stable performance compared to several baselines on four benchmark domains.

This work highlights several open questions for off-policy value estimation. We discuss three key questions that follow-up work can focus on.

Can the empirical benefits of gradient correction methods be combined with the theoretical convenience of saddlepoint methods? Our proposed algorithm uses gradient corrections because these methods typically outperform their saddlepoint counterparts in practice (White and White, 2016; Ghiassian et al., 2018, 2020). In Section 7.1 we discuss how saddlepoint methods use a standard gradient update for min-max problems, and so import the resulting theoretical guarantees. The gradient correction update, on the other hand, is only asymptotically unbiased under certain conditions on the auxiliary variable h . Convergence is only guaranteed under linear function approximation; more needs to be understood about the theoretical properties of this update in the nonlinear setting.

How do we use emphatic weightings for control? In policy evaluation our experiments show that emphatic weights consistently produce better solutions than either the alternative-life solution d_π or the excursion solution d_μ (see Figure 5). In practice, incorporating emphatic weightings can yield high-variance (Sutton and Barto, 2018). An interesting direction is to investigate the use of emphatic weightings in QRC.

What function class should we use for \mathcal{H} ? In this work, we concluded that using a shared network for the primary and auxiliary variables was reasonable effective for optimizing the $\overline{\text{PBE}}$. This effectively restricts $\mathcal{H} = \mathcal{F}$, which mimics the choice for the linear $\overline{\text{PBE}}$, even when we use neural network function approximation. However, restricting \mathcal{H} can result in poor quality solutions, namely those where solutions under the $\overline{\text{BE}}$ are high quality but not under the linear $\overline{\text{PBE}}$. Much more work needs to be done to understand the impacts of different choices for \mathcal{H} .

ACKNOWLEDGMENTS

The authors gratefully acknowledge funding from the Natural Sciences and Engineering Research Council of Canada, CIFAR and Google DeepMind. Thanks to Sina Ghiassian and Richard Sutton for helpful discussions on an earlier version of this work.

References

- Andras Antos, Csaba Szepesvari, and Remi Munos. Learning near-optimal policies with Bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008. doi: 10.1007/s10994-007-5038-2.
- Kavosh Asadi and Michael L. Littman. An Alternative Softmax Operator for Reinforcement Learning. *International Conference on Machine Learning*, 2017.
- Leemon Baird. Residual Algorithms: Reinforcement Learning with Function Approximation. *Machine Learning Proceedings*, pages 30–37, 1995. doi: 10.1016/B978-1-55860-377-6.50013-X.
- Andrew G. Barto, Richard S. Sutton, and Charles W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5):834–846, 1983. doi: 10.1109/TSMC.1983.6313077.

- Dimitri P Bertsekas and J.N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- Alberto Bietti and Julien Mairal. Group invariance, stability to deformations, and complexity of deep convolutional representations. *The Journal of Machine Learning Research*, 20(1): 876–924, 2019.
- Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis Charles, Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 2013.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Bo Dai, Niao He, Yunpeng Pan, Byron Boots, and Le Song. Learning from Conditional Distributions via Dual Embeddings. *International Conference on Artificial Intelligence and Statistics*, page 10, 2017.
- Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. SBEED: Convergent Reinforcement Learning with Nonlinear Function Approximation. *International Conference on Machine Learning*, page 10, 2018.
- Christoph Dann, Gerhard Neumann, and Jan Peters. Policy Evaluation with Temporal Differences: A Survey and Comparison. *Journal of Machine Learning Research*, page 75, 2014.
- Simon S. Du, Jianshu Chen, Lihong Li, Lin Xiao, and Dengyong Zhou. Stochastic Variance Reduction Methods for Policy Evaluation. *International Conference on Machine Learning*, 2017.
- Ewa Dudek and Konstanty Holly. Nonlinear orthogonal projection. *Annales Polonici Mathematici*, 59:1–31, 1994.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. *International Conference on Machine Learning*, 2018.
- Yihao Feng, Lihong Li, and Qiang Liu. A Kernel Loss for Solving the Bellman Equation. *Advances in Neural Information Processing Systems 32*, pages 15456–15467, 2019.
- Sina Ghiassian, Andrew Patterson, Martha White, Richard S. Sutton, and Adam White. Online Off-policy Prediction. *arXiv:1811.02597*, 2018.
- Sina Ghiassian, Andrew Patterson, Shivam Garg, Dhawal Gupta, Adam White, and Martha White. Gradient Temporal-Difference Learning with Regularized Corrections. *International Conference on Machine Learning*, 2020.

- Dibya Ghosh and Marc G. Bellemare. Representations for Stable Off-Policy Reinforcement Learning. *International Conference on Machine Learning*, 2020.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. *International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- Leah Hackman. *Faster Gradient-TD Algorithms*. PhD thesis, University of Alberta, 2013.
- Assaf Hallak and Shie Mannor. Consistent on-line off-policy evaluation. *International Conference on Machine Learning*, pages 1372–1383, 2017.
- Assaf Hallak, Aviv Tamar, Remi Munos, and Shie Mannor. Generalized Emphatic Temporal Difference Learning: Bias-Variance Analysis. *AAAI*, page 7, 2016.
- Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z. Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement Learning with Unsupervised Auxiliary Tasks. *International Conference on Learning Representations*, 2016.
- Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. *International Conference on Machine Learning*, 2002.
- Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*, page 15, 2015.
- J Z Kolter. The Fixed Points of Off-Policy TD. *Advances in Neural Information Processing Systems*, page 9, 2011.
- Michael L. Littman and Richard S Sutton. Predictive Representations of State. *Advances in Neural Information Processing Systems*, pages 1555–1561, 2002.
- Bo Liu, Ji Liu, Mohammad Ghavamzadeh, Sridhar Mahadevan, and Marek Petrik. Finite-Sample Analysis of Proximal Gradient TD Algorithms. *International Conference on Uncertainty in Artificial Intelligence*, 2015.
- Bo Liu, Ji Liu, Mohammad Ghavamzadeh, Sridhar Mahadevan, and Marek Petrik. Proximal Gradient Temporal Difference Learning Algorithms. *International Joint Conference on Artificial Intelligence*, page 5, 2016.
- Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Off-Policy Policy Gradient with Stationary Distribution Correction. *Uncertainty in Artificial Intelligence*, pages 1180–1190, 2020.
- Hamid Reza Maei. *Gradient Temporal-Difference Learning Algorithms*. PhD thesis, University of Alberta, Edmonton, 2011.
- Hamid Reza Maei and Richard S. Sutton. GQ (λ): A general gradient algorithm for temporal-difference prediction learning with eligibility traces. *Conference on Artificial General Intelligence*, 2010.

- Hamid Reza Maei, Csaba Szepesvári, Shalabh Bhatnagar, Doina Precup, David Silver, and Richard S Sutton. Convergent Temporal-Difference Learning with Arbitrary Smooth Function Approximation. *Advances in Neural Information Processing Systems*, pages 1204–1212, 2009.
- Sridhar Mahadevan, Bo Liu, Philip Thomas, Will Dabney, Steve Giguere, Nicholas Jacek, Ian Gemp, and Ji Liu. Proximal Reinforcement Learning: A New Theory of Sequential Decision Making in Primal-Dual Spaces. *arXiv:1405.6757*, 2014.
- Ashique Rupam Mahmood, Huizhen Yu, and Richard S. Sutton. Multi-step Off-policy Learning Without Importance Sampling Ratios. *arXiv:1702.03006*, 2017.
- Odalric-Ambrym Maillard, Remi Munos, Alessandro Lazaric, and Mohammad Ghavamzadeh. Finite-Sample Analysis of Bellman Residual Minimization. *Asian Conference on Machine Learning*, page 16, 2010.
- Andrew William Moore. *Efficient Memory-Based Learning for Robot Control*. PhD thesis, University of Cambridge, 1990.
- Remi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. Safe and Efficient Off-Policy Reinforcement Learning. *Advances in Neural Information Processing Systems*, page 9, 2016.
- Doina Precup, Richard S Sutton, and Satinder Singh. Eligibility Traces for Off-Policy Policy Evaluation. *International Conference on Machine Learning*, page 9, 2000.
- Doina Precup, Richard S Sutton, and Sanjoy Dasgupta. Off-Policy Temporal-Difference Learning with Function Approximation. *International Conference on Machine Learning*, page 9, 2001.
- Bruno Scherrer. Should one compute the Temporal Difference fix point or minimize the Bellman Residual? The unified oblique projection view. *International Conference on Machine Learning*, 2010.
- Ralf Schoknecht. Optimality of Reinforcement Learning Algorithms with Linear Function Approximation. *Advances in Neural Information Processing Systems*, page 8, 2003.
- Richard S. Sutton. Generalization in reinforcement learning: Successful examples using sparse coarse coding. *Advances in Neural Information Processing Systems*, pages 1038–1044, 1996.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, November 2018. ISBN 978-0-262-35270-3.
- Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 99:1057–1063, 1999a.
- Richard S. Sutton, Doina Precup, and Satinder Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1):181–211, 1999b. doi: 10.1016/S0004-3702(99)00052-1.

- Richard S. Sutton, Hamid Reza Maei, Doina Precup, Shalabh Bhatnagar, David Silver, Csaba Szepesvári, and Eric Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. *International Conference on Machine Learning*, pages 1–8, 2009. doi: 10.1145/1553374.1553501.
- Richard S Sutton, Joseph Modayil, Michael Delp, Thomas Degris, Patrick M Pilarski, Adam White, and Doina Precup. Horde: A Scalable Real-time Architecture for Learning Knowledge from Unsupervised Sensorimotor Interaction. *International Conference on Autonomous Agents and Multi-Agent Systems*, page 8, 2011.
- Richard S Sutton, A Rupam Mahmood, and Martha White. An Emphatic Approach to the Problem of Off-policy Temporal-Difference Learning. *Journal of Machine Learning Research*, page 29, 2016.
- Daniel B. Szyld. The many proofs of an identity on the norm of oblique projections. *Numerical Algorithms*, 42(3-4):309–323, 2006.
- Brian Tanner and Richard S. Sutton. TD(λ) networks: Temporal-difference networks with eligibility traces. *International Conference on Machine Learning*, pages 888–895, 2005. doi: 10.1145/1102351.1102463.
- Ahmed Touati, Pierre-Luc Bacon, Doina Precup, and Pascal Vincent. Convergent Tree Backup and Retrace with Function Approximation. *International Conference on Machine Learning*, page 10, 2018.
- J.N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690, 1997. doi: 10.1109/9.580874.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. Dueling Network Architectures for Deep Reinforcement Learning. *International Conference on Machine Learning*, page 9, 2016.
- Christopher John Cornish Hellaby Watkins. *Learning from Delayed Rewards*. PhD thesis, King’s College, Cambridge, 1989.
- Adam White. *Developing a Predictive Approach to Knowledge*. PhD thesis, University of Alberta, 2015.
- Adam White and Martha White. Investigating practical linear temporal difference learning. *International Conference on Autonomous Agents and Multi-Agent Systems*, 2016.
- Martha White. Unifying Task Specification in Reinforcement Learning. *International Conference on Machine Learning*, page 9, 2017.
- Ronald J. Williams and Leemon C. Baird. Tight performance bounds on greedy policies based on imperfect value functions. Technical report, Citeseer, 1993.
- Huizhen Yu. On Convergence of Emphatic Temporal-Difference Learning. *Conference on Learning Theory*, page 28, 2015.

Huizhen Yu and Dimitri P. Bertsekas. Error bounds for approximations from projected linear equations. *Mathematics of Operations Research*, 35(2):306–329, 2010.

Appendix A. An Example of Posterior and Prior Corrections

In this section we go through an explicit example, to highlight the role of prior and posterior corrections. For ease of reference, recall that Off-policy TD(λ) that only has posterior corrections uses update

$$\begin{aligned} \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t^\rho \\ \mathbf{z}_t^\rho &\leftarrow \rho_t (\gamma_t \lambda \mathbf{z}_{t-1}^\rho + \mathbf{x}_t), \end{aligned} \quad (4)$$

and Alternative-life TD(λ) incorporates both prior and posterior corrections

$$\begin{aligned} \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t^\rho \\ \mathbf{z}_t^\rho &\leftarrow \rho_t \left(\gamma_t \lambda \mathbf{z}_{t-1} + \prod_{k=1}^{t-1} \rho_k \mathbf{x}_t \right). \end{aligned} \quad (5)$$




Figure 14: A simple MDP to understand the differences between prior corrections and posterior corrections in off-policy TD algorithms with importance sampling.

Consider the MDP depicted in Figure 14. Each episode starts in the leftmost state denoted ‘x’ and terminates on transition into the terminal state denoted with ‘T’, and each state is represented with a unique tabular state encoding: $x: [1, 0]$, $y: [0, 1]$. In each state there are two possible actions and the behavior policy chooses each action in each state with 0.5 probability. The target policy chooses action a_1 in all states. A posterior correction method like Off-policy TD(λ), will always update the value of a state if action a_1 is taken. For example if the agent experiences the transition $y \rightarrow T$, Off-policy TD(λ) will update the value of state y ; no matter the history of interaction before entering state y .

Although the importance sampling corrections product in the eligibility trace update, Off-policy TD(λ) does not use importance sampling corrections computed from prior time-steps to update the value of the current state. This is easy to see with an example. For simplicity we assume γ_t is 1 everywhere, except upon termination where it is zero. Let’s examine the updates and trace contents for a trajectory where b ’s action choices perfectly agree with π :

$$x \rightarrow y \rightarrow T.$$

After the transition from $x \rightarrow y$, Off-policy TD(λ) will update the value estimate corresponding to x :

$$\begin{bmatrix} \hat{v}_1(x) \\ \hat{v}_1(y) \end{bmatrix} \leftarrow \begin{bmatrix} 0 \\ 0 \end{bmatrix} + \alpha \delta_1 \mathbf{z}_1^\rho = \alpha \delta_1 \begin{bmatrix} \frac{\pi(a_1|x)}{b(a_1|x)} \lambda \\ 0 \end{bmatrix},$$

where $\hat{v}_1(x)$ denotes the estimated value of state x on time step $t = 1$ (after the first transition), and as usual \mathbf{z}_0^ρ and \hat{v} are initialized to zero. After the second transition, $y \rightarrow T$,

the importance sampling corrections will product in the trace, and the value estimates corresponding to both x and y are updated:

$$\begin{bmatrix} \hat{v}_2(x) \\ \hat{v}_2(y) \end{bmatrix} \leftarrow \begin{bmatrix} \hat{v}_1(x) \\ \hat{v}_1(y) \end{bmatrix} + \alpha \delta_2 \begin{bmatrix} \frac{\pi(a_1|y)}{b(a_1|y)} \frac{\pi(a_1|x)}{b(a_1|x)} \lambda^2 \\ \frac{\pi(a_1|y)}{b(a_1|y)} \lambda \end{bmatrix}.$$

The estimated value of state y is only updated with importance sampling corrections computed from state transitions that occur after the visit to y : using $\frac{\pi(a_1|y)}{b(a_1|y)}$, but not $\frac{\pi(a_1|x)}{b(a_1|x)}$.

Finally, consider another trajectory that deviates from the target policy’s choice on the second step of the trajectory:

$$x \rightarrow y \rightarrow y \rightarrow T.$$

On the first transition the value of x is updated as expected, and no update occurs as a result of the second transition. On the third, transition the estimated value of state x is *not* updated; which is easy to see from inspecting the eligibility trace on each time-step:

$$\mathbf{z}_1^\rho = \begin{bmatrix} \frac{\pi(a_1|x)}{b(a_1|x)} \lambda \\ 0 \end{bmatrix}; \quad \mathbf{z}_2^\rho = \mathbf{0}; \quad \mathbf{z}_3^\rho = \begin{bmatrix} 0 \\ \frac{\pi(a_1|y)}{b(a_1|y)} \lambda \end{bmatrix}.$$

The eligibility trace is set to zero on time step two, because the target policy never chooses action a_2 in state y and thus $\frac{\pi(a_2|y)}{b(a_2|y)} = 0$. The value of state S_t is never updated using importance sampling corrections computed on time steps prior to t .

Now consider the updates performed by Alternative-life TD(λ) using different trajectories from our simple MDP (Figure 14). If the agent ever selects action a_2 , then none of the following transitions will result in further updates to the value function. For example, the trajectory $x \rightarrow y \rightarrow y \rightarrow y \cdots y \rightarrow T$ will update $\hat{v}_w(s)$ corresponding to the first $x \rightarrow y$ transition, but $\hat{v}_w(y)$ would never be updated due to the product in Equation 5. In contrast, the Off-policy TD(λ) algorithm described in Equation 4 would update $\hat{v}_w(s)$ on the first transition, and also update $\hat{v}_w(y)$ on the last transition of the trajectory.

Appendix B. The Bias of the $\overline{\text{TDE}}$

In this section, we highlight that the $\overline{\text{TDE}}$ can impose significant bias on the value function solution, to reduce variance of the targets in the TD-error. This could have utility for reducing variance of updates in practice. But, when considering the optimal solution for the $\overline{\text{TDE}}$ —as opposed to the optimization path to get there—it suggests that the $\overline{\text{TDE}}$ is a poor choice of the objective.

We can similarly write the $\overline{\text{TDE}}$ as a decomposition, in terms of both the $\overline{\text{BE}}$ and the $\overline{\text{PBE}}$. The $\overline{\text{TDE}}$ decomposes into the $\overline{\text{BE}}$ and a bias term due to correlation between samples

$$\begin{aligned} & \mathbb{E}_\pi \left[(R + \gamma v(S') - V(s))^2 \mid S = s \right] \\ &= \mathbb{E}_\pi \left[(R + \gamma v(S') - \mathbb{E}_\pi[R + \gamma v(S') \mid S = s] + \mathbb{E}_\pi[R + \gamma v(S') \mid S = s] - V(s))^2 \mid S = s \right] \\ &= \mathbb{E}_\pi \left[(R + \gamma v(S') - \mathbb{E}_\pi[R + \gamma v(S') \mid S = s])^2 \mid S = s \right] + (\mathbb{E}_\pi[R + \gamma v(S') \mid S = s] - V(s))^2 \end{aligned}$$

giving

$$\overline{\text{TDE}} = \|v - \mathcal{T}v\|_D^2 + \mathbb{E}[\mathbf{Var}[R + \gamma v(S') \mid S = s]] \quad (32)$$

This further yields an equality between the $\overline{\text{PBE}}$ and the $\overline{\text{TDE}}$

$$\overline{\text{TDE}} = \|v - \Pi_D \mathcal{T}v\|_D^2 + \|\mathcal{T}v - \Pi_D \mathcal{T}v\|_D^2 + \mathbb{E}[\mathbf{Var}[R + \gamma v(S')|S = s]] \quad (33)$$

This variance penalty encourages finding value functions that minimize the variance of the target. Notice that this connection exists in supervised learning as well, by simply considering the case where $\gamma = 0$. The $\overline{\text{BE}}$ include terms $(\mathbb{E}_\pi[R | S = s] - v(s))^2$ and the $\overline{\text{TDE}}$ is the standard squared error $\mathbb{E}_\pi[(R - v(s))^2 | S = s]$. In regression, this additional variance penalty has no impact on the optimal solution, because it does not include v : $\mathbb{E}[\mathbf{Var}[R + 0 \cdot v(S')|S = s]] = \mathbb{E}[\mathbf{Var}[R|S = s]]$.

Now the question is if it is useful to learn v that results in a lower-variance target. One benefit is that the estimator itself could have lower variance, and so for a smaller number of samples this could warrant the additional bias in v . When comparing these objectives in their ideal forms, over all states under function approximation, it is an undesirable penalty on the value estimates. For instance, even in a 5-state random walk with tabular features, the fixed-point of $\overline{\text{TDE}}$ has non-negligible value error ($\overline{\text{VE}} \approx 0.26$) while the fixed-point of any of the other objectives considered in this work would have zero error.

Appendix C. The Connection to the Saddlepoint Form of the Linear $\overline{\text{PBE}}$

In this section we explicitly show the steps that the generalized $\overline{\text{PBE}}$ with $\mathcal{H} = \mathcal{F}$ under linear function approximation produces the linear $\overline{\text{PBE}}$.

$$\begin{aligned} & \max_{h \in \mathcal{L}} \sum_{s \in \mathcal{S}} d(s) (2\mathbb{E}_\pi[\delta(\mathbf{w}) | S = s] h(s) - h(s)^2) \\ &= \sum_{s \in \mathcal{S}} d(s) 2\mathbb{E}_\pi[\delta(\mathbf{w}) | S = s] \mathbf{x}(s)^\top \theta^* - \sum_{s \in \mathcal{S}} d(s) (\mathbf{x}(s)^\top \theta^*)^2 \\ &= \left(\sum_{s \in \mathcal{S}} d(s) 2\mathbb{E}_\pi[\delta(\mathbf{w}) | S = s] \mathbf{x}(s)^\top \right) \theta^* - \sum_{s \in \mathcal{S}} d(s) (\theta^*)^\top \mathbf{x}(s) \mathbf{x}(s)^\top \theta^* \\ &= 2\mathbb{E}[\delta(\mathbf{w}) \mathbf{x}(s)]^\top \theta^* - (\theta^*)^\top \left(\sum_{s \in \mathcal{S}} d(s) \mathbf{x}(s) \mathbf{x}(s)^\top \right) \theta^* \\ &= 2(\mathbf{b} - \mathbf{A}\mathbf{w})\theta^* - (\theta^*)^\top \mathbf{C}\theta^* \\ &= 2\|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2 - (\theta^*)^\top (\mathbf{b} - \mathbf{A}\mathbf{w}) \\ &= 2\|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2 - \|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2 \\ &= \|\mathbf{b} - \mathbf{A}\mathbf{w}\|_{\mathbf{C}^{-1}}^2 \end{aligned}$$

Appendix D. Proofs for Section 6

This section includes proofs for the theorems in the main body.

Proposition 7 The following statements are true about c_d .

1. If $d = d_\pi$ or $d = m$, then $c_d < 1$.
2. If $d = d_\pi$ and the discount is a constant $\gamma_c < 1$ (a continuing problem), then $c_d = \gamma_c$.

3. If $d = d_\pi$ and for some constant γ_c , $\gamma(s, a, s') \leq \gamma_c$ for all (s, a, s') , then $c_d \leq \gamma_c$.
4. If $d_\pi(s) > 0$ iff $d(s) > 0$, then $c_d \leq s_{d_\pi} \sqrt{\kappa(d, d_\pi) \kappa(d_\pi, d)}$.

Proof **Case 1** is proven in (White, 2017, Theorem 1).

Case 2 is standard result, because $\mathcal{P}_{\pi, \gamma} = \gamma_c P_\pi$. Therefore $\|\mathcal{P}_{\pi, \gamma}\|_{d_\pi} = \gamma_c \|P_\pi\|_{d_\pi} = \gamma_c$ because $\|P_\pi\|_{d_\pi} = 1$.

Case 3 is a simple extension of the above.

$$\mathcal{P}_{\pi, \gamma}(s, s') = \sum_a \pi(a|s) \gamma(s, a, s') P(s'|s, a) = \gamma_c \sum_a \pi(a|s) \gamma(s, a, s') / \gamma_c P(s'|s, a) = \gamma_c \tilde{P}_{\pi, \gamma}$$

with $\tilde{P}_{\pi, \gamma} \stackrel{\text{def}}{=} \sum_a \pi(a|s) \gamma(s, a, s') / \gamma_c P(s'|s, a)$. Because $\gamma(s, a, s') / \gamma_c \leq 1$, we know that $\tilde{P}_{\pi, \gamma}$ is a substochastic matrix and so that $\|\tilde{P}_{\pi, \gamma}\|_{d_\pi} \leq 1$. Therefore, $\|\mathcal{P}_{\pi, \gamma}\|_{d_\pi} = \|\gamma_c \tilde{P}_{\pi, \gamma}\|_{d_\pi} \leq \gamma_c$.

Case 4 is an extension of the strategy used in Kolter (2011, Theorem 2).

$$\begin{aligned} \|\mathcal{P}_{\pi, \gamma}\|_d &= \|D^{1/2} \mathcal{P}_{\pi, \gamma} D^{-1/2}\|_2 \\ &= \|D^{1/2} D_\pi^{-1/2} D_\pi^{1/2} \mathcal{P}_{\pi, \gamma} D_\pi^{-1/2} D_\pi^{1/2} D^{-1/2}\|_2 \\ &\leq \|D^{1/2} D_\pi^{-1/2}\|_2 \|D_\pi^{1/2} \mathcal{P}_{\pi, \gamma} D_\pi^{-1/2}\|_2 \|D_\pi^{1/2} D^{-1/2}\|_2 \\ &= \max_{s \in \mathcal{S}} \sqrt{d(s)/d_\pi(s)} \|\mathcal{P}_{\pi, \gamma}\|_{d_\pi} \max_{s \in \mathcal{S}} \sqrt{d_\pi(s)/d(s)} \\ &\leq \|\mathcal{P}_{\pi, \gamma}\|_{d_\pi} \sqrt{\max_{s \in \mathcal{S}} d(s)/d_\pi(s)} \sqrt{\max_{s \in \mathcal{S}} d_\pi(s)/d(s)} \\ &\leq s_{d_\pi} \sqrt{\kappa(d, d_\pi) \kappa(d_\pi, d)}. \end{aligned} \quad \blacksquare$$

Theorem 9 If $\mathcal{H} \supseteq \mathcal{F}$, then the solution $v_{\mathbf{w}_{\mathcal{H}, d}}$ to the generalized $\overline{\text{PBE}}$ satisfies

$$\|v_\pi - v_{\mathbf{w}_{\mathcal{H}, d}}\|_d \leq \|\Pi_{\mathcal{F}, H}\|_d \|v_\pi - \Pi_{\mathcal{F}, d} v_\pi\|_d. \quad (34)$$

Proof If $\Pi_{\mathcal{F}, H}$ is the identity—a trivial projection—then the result immediately follows. This is because this implies $v_\pi \in \mathcal{F}$, and so both sides of the equation are zero.

Otherwise, assume $\Pi_{\mathcal{F}, H}$ is a non-trivial projection. Notice that $\Pi_{\mathcal{F}, H} \Pi_{\mathcal{F}, d} = \Pi_{\mathcal{F}, d}$, because the first projection already projects to the set \mathcal{F} and so applying $\Pi_{\mathcal{F}, H}$ has no effect. Now for any v ,

$$\begin{aligned} (I - \Pi_{\mathcal{F}, H})(I - \Pi_{\mathcal{F}, d})v &= (I - \Pi_{\mathcal{F}, H} - \Pi_{\mathcal{F}, d} + \Pi_{\mathcal{F}, H} \Pi_{\mathcal{F}, d})v \\ &= (I - \Pi_{\mathcal{F}, H} - \Pi_{\mathcal{F}, d} + \Pi_{\mathcal{F}, d})v \\ &= (I - \Pi_{\mathcal{F}, H})v \end{aligned}$$

because by assumption $\Pi_{\mathcal{F}, H} \Pi_{\mathcal{F}, d} = \Pi_{\mathcal{F}, d}$. Therefore we get that

$$\begin{aligned} \|v_\pi - v\|_d &= \|v_\pi - \Pi_{\mathcal{F}, H} v_\pi\|_d \\ &= \|(I - \Pi_{\mathcal{F}, H})v_\pi\|_d \\ &= \|(I - \Pi_{\mathcal{F}, H})(I - \Pi_{\mathcal{F}, d})v_\pi\|_d \\ &\leq \|I - \Pi_{\mathcal{F}, H}\|_d \|(I - \Pi_{\mathcal{F}, d})v_\pi\|_d \\ &= \|I - \Pi_{\mathcal{F}, H}\|_d \|v_\pi - \Pi_{\mathcal{F}, d} v_\pi\|_d \\ &= \|\Pi_{\mathcal{F}, H}\|_d \|v_\pi - \Pi_{\mathcal{F}, d} v_\pi\|_d \end{aligned}$$

where the last step follows from the fact that for a non-trivial projection operator, $\|\Pi_{\mathcal{F},H}\|_d = \|I - \Pi_{\mathcal{F},H}\|_d$ (Szyld, 2006, Theorem 2.3). \blacksquare

Corollary 10 Assume \mathcal{F} is the space of linear functions with features \mathbf{x} and \mathcal{H} the space of linear functions with features ϕ , with $\mathcal{F} \subseteq \mathcal{H}$. Then the solution to the $\overline{\text{PBE}}$ is

$$v_{\mathbf{w}_{\mathcal{H},d}} = \mathbf{X}\mathbf{w}_{\mathcal{H},d} \quad \text{for } \mathbf{w}_{\mathcal{H},d} = (M^\top(I - \mathcal{P}_{\pi,\gamma})\mathbf{X})^{-1}M^\top r_\pi \quad (35)$$

for $M \stackrel{\text{def}}{=} \Pi_{\mathcal{H},d}^\top D(I - \mathcal{P}_{\pi,\gamma})\mathbf{X}$. Further, $v_{\mathbf{w}_{\mathcal{H},d}} = \Pi_{\mathcal{F},H}v_\pi$ for

$$\Pi_{\mathcal{F},H} = \mathbf{X}(M^\top(I - \mathcal{P}_{\pi,\gamma})\mathbf{X})^{-1}M^\top(I - \mathcal{P}_{\pi,\gamma}). \quad (36)$$

Proof This result follows from two simple facts. First, the $\overline{\text{PBE}}$ can be rewritten as a linear regression problem using $\Psi \stackrel{\text{def}}{=} \Pi_{\mathcal{H},d}(I - \mathcal{P}_{\pi,\gamma})\mathbf{X}$:

$$\begin{aligned} \overline{\text{PBE}}(\mathbf{w}) &= \|\Pi_{\mathcal{H},d}\mathbf{X}\mathbf{w} - \Pi_{\mathcal{H},d}r_\pi - \Pi_{\mathcal{H},d}\mathcal{P}_{\pi,\gamma}\mathbf{X}\mathbf{w} - \|_d \\ &= \|\Pi_{\mathcal{H},d}(I - \mathcal{P}_{\pi,\gamma})\mathbf{X}\mathbf{w} - \Pi_{\mathcal{H},d}r_\pi\|_d = \|\Psi\mathbf{w} - \Pi_{\mathcal{H},d}r_\pi\|_d \end{aligned}$$

with solution

$$v_{\mathbf{w}_{\mathcal{H},d}} = \mathbf{X}\mathbf{w}_{\mathcal{H},d} \quad \text{for } \mathbf{w}_{\mathcal{H},d} = (\Psi^\top D\Psi)^{-1}\Psi^\top D\Pi_{\mathcal{H},d}r_\pi$$

Second, this solution can be simplified because $\Pi_{\mathcal{H},d}^\top D\Pi_{\mathcal{H},d} = D\Pi_{\mathcal{H},d}$. To see why, notice that in the linear case $\Pi_{\mathcal{H},d} = \Phi(\Phi^\top D\Phi)^{-1}\Phi^\top D$ and so

$$\begin{aligned} \Pi_{\mathcal{H},d}^\top D\Pi_{\mathcal{H},d} &= D\Phi(\Phi^\top D\Phi)^{-1}\Phi^\top D\Phi(\Phi^\top D\Phi)^{-1}\Phi^\top D \\ &= D\Phi(\Phi^\top D\Phi)^{-1}\Phi^\top D = D\Pi_{\mathcal{H},d} \end{aligned}$$

Therefore, letting $L \stackrel{\text{def}}{=} I - \mathcal{P}_{\pi,\gamma}$, we get that

$$\begin{aligned} \Psi^\top D\Psi &= \mathbf{X}^\top L^\top \Pi_{\mathcal{H},d}^\top D\Pi_{\mathcal{H},d}L\mathbf{X} = \mathbf{X}^\top L^\top D\Pi_{\mathcal{H},d}L\mathbf{X} \\ \Psi^\top D\Pi_{\mathcal{H},d} &= \mathbf{X}^\top L^\top \Pi_{\mathcal{H},d}^\top D\Pi_{\mathcal{H},d} = \mathbf{X}^\top L^\top D\Pi_{\mathcal{H},d} \end{aligned}$$

and so with $M^\top \stackrel{\text{def}}{=} \mathbf{X}^\top L^\top D\Pi_{\mathcal{H},d}$ we have

$$\mathbf{w}_{\mathcal{H},d} = (M^\top L\mathbf{X})^{-1}M^\top r_\pi$$

$\mathbf{X}(M^\top L\mathbf{X})^{-1}M^\top L$ is the definition of an oblique projection onto the span of \mathbf{X} , orthogonal to the span of $L^\top M$. Since $v_\pi = L^{-1}r_\pi$, we get that

$$v_{\mathbf{w}_{\mathcal{H},d}} = \mathbf{X}\mathbf{w}_{\mathcal{H},d} = \mathbf{X}(M^\top L\mathbf{X})^{-1}M^\top LL^{-1}r_\pi = \mathbf{X}(M^\top L\mathbf{X})^{-1}M^\top Lv_\pi$$

with $\Pi_{\mathcal{F},H} = \mathbf{X}(M^\top L\mathbf{X})^{-1}M^\top L$, completing the proof. \blacksquare

Theorem 11 For any $v \in \mathcal{F}$,

$$\begin{aligned} \|v_\pi - v\|_d &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \underbrace{\|\mathcal{T}v - v\|_d}_{\overline{\text{BE}}} \\ &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \left(\underbrace{\|\Pi_{\mathcal{H},d}\mathcal{T}v - v\|_d}_{\overline{\text{PBE}}} + \text{ApproxError}(\mathcal{H}, v) \right) \\ &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \left(\|\Pi_{\mathcal{H},d}\mathcal{T}v - v\|_d + \text{ApproxError}(\mathcal{H}) \right). \end{aligned}$$

Note that if $c_d < 1$, then $\|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \leq (1 - c_d)^{-1}$.

Proof For the first part, we use the same argument as in Maillard et al. (2010, Equation 8), but for any v rather than the v that minimizes the $\overline{\text{BE}}$. Notice first that $v_\pi = (I - \mathcal{P}_{\pi,\gamma})^{-1}r_\pi$ and so $r_\pi = (I - \mathcal{P}_{\pi,\gamma})v_\pi$. This means that

$$\begin{aligned} \mathcal{T}v - v &= r_\pi + \mathcal{P}_{\pi,\gamma}v - v = (I - \mathcal{P}_{\pi,\gamma})v_\pi + (\mathcal{P}_{\pi,\gamma} - I)v \\ &= (I - \mathcal{P}_{\pi,\gamma})(v_\pi - v) \\ \implies v_\pi - v &= (I - \mathcal{P}_{\pi,\gamma})^{-1}(\mathcal{T}v - v) \\ \implies \|v_\pi - v\|_d &\leq \|(I - \mathcal{P}_{\pi,\gamma})^{-1}\|_d \|\mathcal{T}v - v\|_d \end{aligned}$$

To bound the approximation error, let us define the conjugate form as $f(v, h) \stackrel{\text{def}}{=} \sum_s d(s)(2\mathbb{E}[\delta(v)|S = s]h(s) - h(s)^2)$. Notice first that $\overline{\text{BE}}(v) \geq \overline{\text{PBE}}(v)$ for any v because

$$\overline{\text{BE}}(v) - \overline{\text{PBE}}(v) = \max_{h \in \mathcal{H}_{\text{all}}} f(v, h) - \max_{h \in \mathcal{H}} f(v, h)$$

and the first maximization is over a broader set, with $\mathcal{H} \subset \mathcal{H}_{\text{all}}$. Let us additionally define $g(v, h) \stackrel{\text{def}}{=} \sum_s d(s)(\mathbb{E}[\delta(v)|S = s] - h(s))^2$, the minimization form for h we used previously. Then $-g(v, h) = f(v, h) - \sum_s d(s)\mathbb{E}[\delta(v)|S = s]^2$. We can therefore rewrite this difference using a minimization

$$\begin{aligned} \overline{\text{BE}}(v) - \overline{\text{PBE}}(v) &= \max_{h \in \mathcal{H}_{\text{all}}} f(v, h) - \sum_s d(s)\mathbb{E}[\delta(v)|S = s]^2 - \max_{h \in \mathcal{H}} f(v, h) + \sum_s d(s)\mathbb{E}[\delta(v)|S = s]^2 \\ &= \max_{h \in \mathcal{H}_{\text{all}}} -g(v, h) - \max_{h \in \mathcal{H}} -g(v, h) \\ &= \min_{h \in \mathcal{H}} g(v, h) - \min_{h \in \mathcal{H}_{\text{all}}} g(v, h) \end{aligned}$$

where on the last step we used the fact that $\max_{h \in \mathcal{H}} -g(v, h) = -\min_{h \in \mathcal{H}} g(v, h)$. Notice that $\min_{h \in \mathcal{H}} g(v, h) = \|h_v^* - h\|_d = \text{ApproxError}(\mathcal{H}, v)$, because $h_v^*(s) = \mathbb{E}[\delta(v)|S = s]$. Further, $\min_{h \in \mathcal{H}_{\text{all}}} g(v, h) = 0$, since $h_v^* \in \mathcal{H}_{\text{all}}$. Therefore, we get that

$$\overline{\text{BE}}(v) - \overline{\text{PBE}}(v) = \min_{h \in \mathcal{H}} g(v, h) = \text{ApproxError}(\mathcal{H}, v) \leq \text{ApproxError}(\mathcal{H}) \quad \blacksquare$$

Appendix E. Summary of Known Bounds for Value Estimation

In the main paper, we derived bounds for the generalized $\overline{\text{PBE}}$. These results built on existing bounds, in some cases for more specific settings and in others for slightly different settings. We summarize those known results in Table 1.

Bound	Additional Comments
Williams and Baird (1993, Prop 3.1) $\ v - v_\pi\ _\infty \leq \frac{1}{1-\gamma} \ \mathcal{T}v - v\ _\infty$, any v	A seminal result showing the relationship between the worst case value error across states and the worst case Bellman error across states.
Tsitsiklis and Van Roy (1997, Lemma 6) $\ v - v_\pi\ _{d_\pi} \leq \frac{1-\lambda\gamma}{1-\gamma} \ \Pi_{\mathcal{F}, d_\pi} v_\pi - v_\pi\ _{d_\pi}$	A seminal result for the on-policy setting. $v = \arg \min_{v \in \mathcal{F}} \ \Pi_{\mathcal{F}, d_\pi} \mathcal{T}v - v\ _{d_\pi}^2$ is obtained with on-policy TD(λ) using linear function approximation. For $\lambda = 1$, the lhs and rhs are equal.
Antos et al. (2008, Lemma 10) $\ \mathcal{T}v - v\ _d \leq \min_{v \in \mathcal{F}} \ \mathcal{T}v - v\ _d + \max_{v \in \mathcal{F}} \min_{v' \in \mathcal{F}} \ \mathcal{T}v - v'\ _d$ The first term is called the inherent Bellman error of \mathcal{F} and the second is the inherent one-step Bellman error of \mathcal{F} .	Upper bound on the true Bellman error using the empirical loss that resembles the conjugate Bellman error. This work assumes $\mathcal{H} = \mathcal{F}$, and v is the solution to the generalized $\overline{\text{PBE}}$. The bound in the paper also accounts for using an empirical loss; the result here is obtained from within the proof. See Page 124, and note that $\bar{L}(v) = \min_{v \in \mathcal{F}} \bar{L}(v)$.
Maillard et al. (2010, Equation 8) $\ v - v_\pi\ _d \leq \ (I - \mathcal{P}_{\pi, \gamma})^{-1}\ _d \ \mathcal{T}v - v\ _d$ for $v = \arg \min_{v \in \mathcal{F}} \ \mathcal{T}v - v\ _d$	The original result is given for a constant γ . We include the more generic result here, since it is a straightforward extension.
Scherrer (2010, Proposition 3) $\ v_M - v_\pi\ _d \leq \sqrt{\text{spectral-radius}(ABC B^\top)} \ \Pi_{\mathcal{F}, d} v_\pi - v_\pi\ _d$ $v_M = \mathbf{X}(M^\top \mathbf{X})^{-1} M^\top v_\pi$ is the oblique projection of v_π onto the span of \mathbf{X} —linear functions with features $\mathbf{x}(s)$ —orthogonally to span of matrix $L^\top M \in \mathbb{R}^{n \times d}$ for $L = (I - \mathcal{P}_{\pi, \gamma})$.	This result assumes linear function approximation. If $M = D\mathbf{X}$, then v_M is the TD fixed point. If $M = DL\mathbf{X}$, then v_M is a solution to the $\overline{\text{BE}}$. The spectral radius of a matrix is the maximum of the absolute value of the eigenvalues. This result is for matrices $A = \mathbf{X}^\top D\mathbf{X}$, $B = (M^\top L\mathbf{X})^{-1}$ and $C = MLD^{-1}L^\top M$. This result extends and simplifies the earlier results in (Schoknecht, 2003) and in (Yu and Bertsekas, 2010).
Kolter (2011, Theorem 2) $\ v - v_\pi\ _{d_\mu} \leq \frac{1+\gamma \text{cond}(\bar{D})}{1-\gamma} \ \Pi_{\mathcal{F}, d_\mu} v_\pi - v_\pi\ _{d_\mu}$ v obtained with Off-policy LSTD using linear function approximation with conditions on d_μ .	$\text{cond}(\bar{D})$ is the condition number of the matrix $\bar{D} \doteq D_\mu^{-1/2} D_\pi^{1/2}$. For these diagonal matrices, $\text{cond}(\bar{D}) = \kappa(d_\mu, d_\pi) / \kappa(d_\pi, d_\mu)$. This bound assumes d_μ satisfies a certain relaxed contraction property, that states that applying the transition to the values and then projecting is a non-expansion under d_μ (see (Kolter, 2011, Theorem 2)).

<p>Hallak et al. (2016, Corollary 1)</p> $\ v - v_\pi\ _{d_\mu} \leq C(d_\mu, m)^{-1/2} \ \Pi_{\mathcal{F}, m} v_\pi - v_\pi\ _m$ $\ v - v_\pi\ _m \leq (\gamma C(d_\mu, m))^{-1/2} \ \Pi_{\mathcal{F}, m} v_\pi - v_\pi\ _m$ <p>for $C(d_\mu, m) \stackrel{\text{def}}{=} 1 - \frac{\gamma^2}{\beta} (1 - \kappa(d_\mu, m))$.</p>	<p>Bounds provided for ETD($\lambda = 0, \beta$) under linear function approximation to obtain v. Note that $\beta \leq \gamma$ and $\beta = \gamma$ corresponds to ETD(0). Here, γ is a constant less than 1. The result is also extended to $\lambda > 0$ in Hallak et al. (2016, Corollary 1), by computing s_m for the Bellman operator.</p>
<p>White (2017, Theorem 1)</p> $\ v - v_\pi\ _d \leq (1 - s_d)^{-1} \ \Pi_{\mathcal{F}, d} v_\pi - v_\pi\ _d$	<p>Generalized results in Hallak et al. (2016), showing the Bellman operator is a contraction for transition-based discounting. This result was shown for $d = d_\pi$ and $d = m$, under linear function approximation with v the solution to the linear $\overline{\text{PBE}}$ with weighting d.</p>

Table 1: A summary of known bounds on the $\overline{\text{VE}}$ of the solution to the $\overline{\text{PBE}}$ and $\overline{\text{BE}}$.

Appendix F. Extensions to n -step returns and λ -returns

In this section, we provide the TDRC algorithm under n -step returns and λ -returns. We start by discussing how the generalized $\overline{\text{PBE}}$ extends, and then the resulting algorithm.

There are further n -step variants of these objectives, where bootstrapping occurs only after n steps: $(G_{t,n} - \hat{v}_w(S_t))^2$ where $G_{t,n} = R_{t+1} + \gamma_{t+1}R_{t+2} + \dots + \gamma_{t+1:t+n}R_{t+n} + \gamma_{t+1:t+n+1}\hat{v}_w(S_{t+n+1})$ where $\gamma_{t+1:t+n} = \gamma_{t+1}\gamma_{t+2} \dots \gamma_{t+n}$. The extreme of n -step returns is to use the full return with no bootstrapping, as in Monte Carlo methods, with the objective becoming the $\overline{\text{RE}}$. The conjugate form and derivations above extend to n -step returns, simply by considering the n -step Bellman operator and corresponding n -step TD error:

$$\delta^{(n)}(\mathbf{w}) \stackrel{\text{def}}{=} R_{t+1} + \gamma_{t+1}R_{t+2} + \dots + \gamma_{t+1:t+n}R_{t+n} + \gamma_{t+1:t+n+1}\hat{v}_w(S_{t+n+1}) - \hat{v}_w(S_t)$$

with importance sampling ratios included, in the off-policy setting. The n -step generalized $\overline{\text{PBE}}$ is $\max_{h \in \mathcal{H}} \mathbb{E}_d [2\mathbb{E}_\pi[\delta^{(n)}(\mathbf{w}) | S = s] h(s) - h(s)^2]$. The function h is trying to estimate the expected n -step return from s : $\mathbb{E}_d [2\mathbb{E}_\pi[\delta^{(n)}(\mathbf{w}) | S = s]]$. The saddlepoint update for \mathbf{w} is

$$\Delta \mathbf{w} \leftarrow h(S_t) (\nabla_{\mathbf{w}} \hat{v}_w(S_t) - h_\theta(S_t) \gamma_{t+1:t+n+1} \nabla_{\mathbf{w}} \hat{v}_w(S_{t+n+1}))$$

and the gradient correction update is

$$\Delta \mathbf{w} \leftarrow \delta^{(n)}(\mathbf{w}) \nabla_{\mathbf{w}} \hat{v}_w(S_t) - h_\theta(S_t) \gamma_{t+1:t+n+1} \nabla_{\mathbf{w}} \hat{v}_w(S_{t+n+1})$$

where both use the same update for h :

$$\Delta \mathbf{h} \leftarrow -(\delta^{(n)}(\mathbf{w}) - h_\theta(S_t)) \nabla_{\mathbf{h}} h_\theta(S_t)$$

The primary difference when considering n -step returns is that, for large n , it is less necessary to estimate h . For large n , the correlation between $\delta^{(n)}(\mathbf{w})$ and $\hat{v}_w(S_{t+n+1})$ becomes smaller. Consequently, it would not be unreasonable to use $\delta^{(n)}(\mathbf{w}) \nabla_{\mathbf{w}} \hat{v}_w(S_t) -$

$\delta^{(n)}(\mathbf{w})\gamma_{t+1:t+n}\nabla_{\mathbf{w}}\hat{v}_{\mathbf{w}}(S_{t+n+1})$, as the incurred bias is likely small. Further, if the discount per step is less than 1, then the gradient correction term also diminishes in importance, because it is pre-multiplied by $\gamma_{t+1:t+n+1}$. For example, for a constant $\gamma < 1$, we get $\gamma_{t+1:t+n} = \gamma^n$. One might expect that the gradient correction update might have an even greater advantage here over the saddlepoint update. It remains an open question as to the relationship between n and some of these choices.

The other generalization, with λ -returns, involves mixing n -step returns with a potentially state-dependent parameter λ . For a constant $\lambda \in [0, 1]$, the λ -return is

$$G_{t,\lambda} = (1 - \lambda) \sum_{i=1}^{\infty} \lambda^{i-1} G_{t,i} = R_{t+1} + \gamma_{t+1}[(1 - \lambda)\hat{v}_{\mathbf{w}}(S_{t+1}) + \lambda G_{t+1,\lambda}].$$

More generally, for a state-dependent $\lambda : \mathcal{S} \rightarrow [0, 1]$ (Maei and Sutton, 2010), with $\lambda_{t+1} \stackrel{\text{def}}{=} \lambda(S_{t+1})$, we have $G_{t,\lambda} = R_{t+1} + \gamma_{t+1}[(1 - \lambda_{t+1})\hat{v}_{\mathbf{w}}(S_{t+1}) + \lambda_{t+1}G_{t+1,\lambda}]$. This target is used in place of $G_{t,n}$ to define the $\overline{\text{PBE}}$ for λ -returns.

Notice, however, that the λ -return involves an infinite sum into the future, or until the end of an episode. For episodic problems, these targets can be computed, like Monte carlo methods. However, more typically to optimize the $\overline{\text{PBE}}$ for λ -returns, the gradient is estimated using a backwards view, rather than this forward view (Sutton and Barto, 2018). This reformulation only holds under linear function approximation. It remains an open question how effective the approach is in the nonlinear setting.

Appendix G. Summary of Algorithms for the Linear $\overline{\text{PBE}}$

In this section we summarize the current algorithms that optimize the linear $\overline{\text{PBE}}$ by the implicit weighting they use. The two axes indicate whether prior or posterior corrections are used, and whether the weighting includes d_{π} or d_b . The emphatic weighting adjusts the state weighting with prior corrections, but still includes d_b as part of the weighting in the objective. Alternative-life, on the other hand, removes d_b all together.

The primary differences between the algorithms is in how they optimize their respective objectives. They can use gradient updates or hybrid updates, and can incorporate other additions like action-dependent bootstrapping. For example, we can easily obtain a gradient version of ETD, by incorporating emphatic weights into the GTD algorithm. The GTD(λ) algorithm is

$$\begin{aligned} \mathbf{z}_t^{\rho} &\leftarrow \rho_t (\gamma_t \lambda \mathbf{z}_{t-1}^{\rho} + \mathbf{x}_t) \\ \theta_{t+1} &\leftarrow \theta_t + \alpha_h [\delta_t \mathbf{z}_t^{\rho} - (\theta_t^{\top} \mathbf{x}) \mathbf{x}_t] \\ \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_t^{\rho} - \underbrace{\alpha \gamma_{t+1} (1 - \lambda) (\theta_t^{\top} \mathbf{z}_t^{\rho}) \mathbf{x}_{t+1}}_{\text{correction term}} \end{aligned}$$

The Emphatic GTD(λ) algorithm uses the same updates to the two sets of weights, but with emphatic weights in the trace: $\mathbf{z}_t^{\rho} \leftarrow \rho_t (\gamma_t \lambda \mathbf{z}_{t-1}^{\rho} + M_t \mathbf{x}_t)$. More simply, when considering

Emphatic GTD(0), we get

$$\begin{aligned}
 F_t &\leftarrow \rho_{t-1}\gamma_t F_{t-1} + 1 \\
 \theta_{t+1} &\leftarrow \theta_t + \alpha_h \rho_t F_t [\delta_t - (\theta_t^\top \mathbf{x})] \mathbf{x}_t \\
 \mathbf{w}_{t+1} &\leftarrow \mathbf{w}_t + \alpha \rho_t F_t [\delta_t \mathbf{x}_t - \gamma_{t+1}(1 - \lambda)(\theta_t^\top \mathbf{x}_t) \mathbf{x}_{t+1}]
 \end{aligned}$$

where F_t is omitted for standard GTD(0).

Many of these algorithms can incorporate different weightings by pre-multiplying the update, but might use different approaches to approximate the gradient. It is important to separate the weighting and algorithmic approach, as otherwise, comparisons between algorithms (e.g., ETD and GTD) could reflect either the difference in objective (weighting by m or d_b) or in the algorithmic approach (TD-style updates or gradient updates). The algorithms, categorized according to weightings, are summarized in Table 2. For a more complete list of algorithms and their updates, see (White and White, 2016, Appendix A).

	d_π	d_b
Posterior corrections	<ul style="list-style-type: none"> ▷ N/A. Alternative life algorithms cannot only do posterior corrections. 	<ul style="list-style-type: none"> ▷ Off-Policy TD(λ) ▷ GTD(λ) (Sutton et al., 2009), ▷ HTD(λ) (Maei, 2011; White and White, 2016) ▷ Action-dependent bootstrapping <ul style="list-style-type: none"> – Tree Backup(λ) (Precup et al., 2000) – V-trace(λ) (Espeholt et al., 2018) – AB-Trace(λ) (Mahmood et al., 2017) ▷ Saddlepoint methods for GTD2(λ) <ul style="list-style-type: none"> – GTD2-MP(λ) (Liu et al., 2016) – SVRG and SAGA (Du et al., 2017) – Gradient Tree Backup(λ) (Touati et al., 2018)
Prior & Posterior corrections	<ul style="list-style-type: none"> ▷ Alternative-life TD(λ) ▷ Alternative-life GTD(λ), HTD(λ) & Saddlepoint methods 	<ul style="list-style-type: none"> ▷ ETD(λ) (Sutton et al., 2016) ▷ ETD(λ, β) (Hallak et al., 2016) ▷ Emphatic GTD(λ), HTD(λ) & Saddlepoint methods

Table 2: A summary of off-policy value estimation methods for the linear $\overline{\text{PBE}}$, based on weightings in the objective and whether they incorporate prior and posterior corrections. The algorithms in grey are hypothetical algorithms that can easily be derived by applying the same derivations as in their original works, but with alternative weightings.




Figure 15: Average return over number of learning steps for best hyperparameters chosen for each algorithm on each domain. When tuning hyperparameters on a per domain basis, the gaps between algorithm performance begin to decrease. QRC still consistently is among the top performing algorithms in each benchmark domain. The performance difference between SBEED and the other competitors on Cart Pole can be explained by the forced stochastic ϵ -greedy policy used by action value methods compared to the increasingly deterministic policy learned by SBEED.

Appendix H. Impact of Voting Procedure

In Section 9.2 we investigate the performance of the benchmark algorithms across four simulation environments with a single hyperparameter setting chosen for each algorithm via a voting procedure. In this section, we investigate the impact of the voting procedure used on the conclusions drawn from the results in Figure 7. We start by investigating each algorithm given that we have the ability to tune hyperparameters individually for every domain. We then investigate the impact of excluding Lunar Lander from the voting procedure, allowing all algorithms (especially SBEED) to choose parameters without needing to balance performance on the most challenging domain. Finally, we demonstrate that a more simple voting procedure leads towards less favorable results for algorithms which have higher sensitivity to hyperparameters across domains.

Figure 15 shows the performance for each algorithm for their best performing respective hyperparameter on each domain. Because every algorithm is able to maximize their individual performance for each domain, and because many algorithms successfully solve each domain with their best parameters, performance gaps between algorithms shrink considerably. Still only QRC is able to successfully solve the Mountain Car domain consistently across runs and both QC-LL and SBEED find a sub-optimal policy in Lunar Lander.

Figure 15 is an idealized setting where every domain gets to pick their favorite parameter setting through extensive tuning. A more realistic setting would involve picking a single “default” hyperparameter setting and perhaps a small exploration of parameters along a small number of axes (for example choosing default parameters in the ADAM optimizer, but sweeping a few values of the stepsize). Because it is often computationally infeasible to tune hyperparameters extensively for each domain, this results in a compromise where some algorithms get to use parameters that perform well on only a subset of the domains. Figure 7 illustrates this tradeoff by requiring each algorithm to choose a single hyperparameter setting that balances performance across all domains. For algorithms that are highly insensitive to their hyperparameters, this selection process is easy and the algorithm need not tradeoff performance significantly in each domain. For algorithms that are highly sensitive to their

hyperparameters, this selection process is very challenging and can result in hyperparameter settings that are catastrophic on some domains but highly performing on others.

In Section 9.2, we used a pair-wise voting strategy, Condorcet voting, to give each algorithm maximum opportunity to balance hyperparameters across all possible pairs of domains. Because Lunar Lander is a much more challenging domain than the rest, the inclusion of Lunar Lander significantly impacted the outcome of the voting procedure. To understand the impact of this, in Figure 16 we use the same voting mechanism as in Figure 7 except we exclude Lunar Lander from influencing the outcome of the vote. In this case, each algorithm chooses the best hyperparameter setting with a pair-wise vote across Cart Pole, Mountain Car, and Acrobot, then we show performance for that hyperparameter setting across all four domains. This can additionally be viewed as a form of hold-out set (i.e. leave-one-out cross validation), where we test how the hyperparameter choosing strategy generalizes to a hand-picked challenging sample. In this setting, all algorithms suffer lower performance in Lunar Lander, particularly Q-learning and SBEED.

Finally, in Figure 17 we investigate the impact of a different voting strategy across all four domains. We use an instant runoff vote which is generally known to provide an inferior between outlying “candidates” than pair-wise voting procedures. Our results show that QC-LL and SBEED are both strongly negatively impacted by this choice of voting procedure where Q-learning and QRC achieve similar performance as in Figure 7.




Figure 16: Average return over number of learning steps for hyperparameters chosen via Condorcet voting. The performance of each algorithm on Lunar Lander was excluded from impacting the result of the vote and can be treated as a hold-out set.




Figure 17: Average return over number of learning steps for hyperparameters chosen via instant runoff voting. For algorithms other than QRC, the instant runoff voting procedure could not find a suitable parameter setting which balanced performance across all domains.