Policy Gradient with Tree Expansion

Anonymous Author(s)
Affiliation
Address
email

Abstract

Policy gradient methods are notorious for having a large variance and high sample complexity. To mitigate this, we introduce SoftTreeMax—a generalization of softmax that employs planning. In SoftTreeMax, we extend the traditional logits with the multi-step discounted cumulative reward, topped with the logits of future states. We analyze SoftTreeMax and explain how tree expansion helps to reduce its gradient variance. We prove that the variance decays exponentially with the planning horizon as a function of the chosen tree-expansion policy. Specifically, we show that the closer the induced transitions are to being state-independent, the stronger the decay. With approximate forward models, we prove that the resulting gradient bias diminishes with the approximation error while retaining the same variance reduction. Ours is the first result to bound the gradient bias for an approximate model. In a practical implementation of SoftTreeMax, we utilize a parallel GPU-based simulator for fast and efficient tree expansion. Using this implementation in Atari, we show that SoftTreeMax reduces the gradient variance by three orders of magnitude. This leads to better sample complexity and improved performance compared to distributed PPO.

1 Introduction

Policy Gradient (PG) methods [Sutton et al., 1999] for Reinforcement Learning (RL) are often the first choice for environments that allow numerous interactions at a fast pace [Schulman et al., 2017]. Their success is attributed to several factors: they are easy to distribute to multiple workers, require no assumptions on the underlying value function, and have both on-policy and off-policy variants.

Despite these positive features, PG algorithms are also notoriously unstable due to the high variance of the gradients computed over entire trajectories [Liu et al., 2020, Xu et al., 2020]. As a result, PG algorithms tend to be highly inefficient in terms of sample complexity. Several solutions have been proposed to mitigate the high variance issue, including baseline subtraction [Greensmith et al., 2004, Thomas and Brunskill, 2017, Wu et al., 2018], anchor-point averaging [Papini et al., 2018], and other variance reduction techniques [Zhang et al., 2021, Shen et al., 2019, Pham et al., 2020].

A second family of algorithms that achieved state-of-the-art results in several domains is based on planning. Planning is exercised primarily in the context of value-based RL and is usually implemented using a Tree Search (TS) [Silver et al., 2016, Schrittwieser et al., 2020]. In this work, we combine PG with TS by introducing a parameterized differentiable policy that incorporates tree expansion. Namely, our SoftTreeMax policy replaces the standard policy logits of a state and action, with the expected value of trajectories that originate from these state and action. We consider two variants of SoftTreeMax, one for cumulative reward and one for exponentiated reward.

Combining TS and PG should be done with care given the biggest downside of PG—its high gradient variance. This raises questions that were ignored until this work: (i) How to design a PG method based on tree-expansion that is stable and performs well in practice? and (ii) How does the tree-expansion

Submitted to 38th Conference on Neural Information Processing Systems (NeurIPS 2024). Do not distribute.

policy affect the PG variance? Here, we analyze SoftTreeMax, and provide a practical methodology to choose the expansion policy to minimize the resulting variance. Our main result shows that a desirable expansion policy is one, under which the induced transition probabilities are similar for each starting state. More generally, we show that the gradient variance of SoftTreeMax decays at a rate of $|\lambda_2|^d$, where d is the depth of the tree and λ_2 is the second eigenvalue of the transition matrix induced by the tree expansion policy. This work is the first to prove such a relation between PG variance and tree expansion policy. In addition, we prove that the with an approximate forward model, the bias of the gradient is bounded proportionally to the approximation error of the model.

To verify our results, we implemented a practical version of SoftTreeMax that exhaustively searches the entire tree and applies a neural network on its leaves. We test our algorithm on a parallelized Atari GPU simulator [Dalton et al., 2020]. To enable a tractable deep search, up to depth eight, we also introduce a pruning technique that limits the width of the tree. We do so by sampling only the most promising nodes at each level. We integrate our SoftTreeMax GPU implementation into the popular PPO [Schulman et al., 2017] and compare it to the flat distributed variant of PPO. This allows us to demonstrate the potential benefit of utilizing learned models while isolating the fundamental properties of TS without added noise. In all tested Atari games, our results outperform the baseline and obtain up to 5x more reward. We further show in Section 6 that the associated gradient variance is smaller by three orders of magnitude in all games, demonstrating the relation between low gradient variance and high reward.

We summarize our key contributions. (i) We show how to combine two families of SoTA approaches: PG and TS by **introducing SoftTreeMax:** a novel parametric policy that generalizes softmax to planning. Specifically, we propose two variants based on cumulative and exponentiated rewards. (ii) We **prove that the gradient variance of SoftTreeMax in its two variants decays exponentially** with its tree depth. Our analysis sheds new light on the choice of tree expansion policy. It raises the question of optimality in terms of variance versus the traditional regret; e.g., in UCT [Kocsis and Szepesvári, 2006]. (iii) We prove that with an approximate forward model, the **gradient bias is proportional to the approximation error**, while retaining the variance decay. This quantifies the accuracy required from a learned forward model. (iv) We **implement a differentiable deep version of SoftTreeMax** that employs a parallelized GPU tree expansion. We demonstrate how its gradient variance is reduced by three orders of magnitude over PPO while obtaining up to 5x reward.

2 Preliminaries

Let Δ_U denote simplex over the set U. Throughout, we consider a discounted Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma, \nu)$, where $\mathcal{S}$ is a finite state space of size S, $\mathcal{A}$ is a finite action space of size A, $r : \mathcal{S} \times \mathcal{A} \to [0, 1]$ is the reward function, $P : \mathcal{S} \times \mathcal{A} \to \Delta_{\mathcal{S}}$ is the transition function, $\gamma \in (0, 1)$ is the discount factor, and $\nu \in \mathbb{R}^S$ is the initial state distribution. We denote the transition matrix starting from state s by $P_s \in [0, 1]^{A \times S}$, i.e., $[P_s]_{a,s'} = P(s'|a, s)$. Similarly, let $R_s = r(s, \cdot) \in \mathbb{R}^A$ denote the corresponding reward vector. Separately, let $\pi : \mathcal{S} \to \Delta_{\mathcal{A}}$ be a stationary policy. Let P^π and R_π be the induced transition matrix and reward function, respectively, i.e., $P^\pi(s'|s) = \sum_a \pi(a|s) \Pr(s'|s, a)$ and $R_\pi(s) = \sum_a \pi(a|s) r(s, a)$. Denote the stationary distribution of P^π by $\mu_\pi \in \mathbb{R}^S$ s.t. $\mu_\pi^\top P^\pi = P^\pi$, and the discounted state visitation frequency by d_π so that $d_\pi^\top = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \nu^\top (P^\pi)^t$. Also, let $V^\pi \in \mathbb{R}^S$ be the value function of π defined by $V^\pi(s) = \mathbb{E}^\pi \left[\sum_{t=0}^{\infty} \gamma^t r\left(s_t, \pi(s_t)\right) \mid s_0 = s\right]$, and let $Q^\pi \in \mathbb{R}^{S \times A}$ be the Q-function such that $Q^\pi(s, a) = \mathbb{E}^\pi \left[r(s, a) + \gamma V^\pi(s')\right]$. Our goal is to find an optimal policy $\pi^\star$ such that $V^\star(s) \equiv V^{\pi^\star}(s) = \max_\pi V^\pi(s),\ \forall s \in \mathcal{S}$.

For the analysis in Section 4, we introduce the following notation. Denote by $\Theta \in \mathbb{R}^S$ the vector representation of $\theta(s)\ \forall s \in \mathcal{S}$. For a vector u, denote by $\exp(u)$ the coordinate-wise exponent of u and by $D(u)$ the diagonal square matrix with u in its diagonal. For a matrix A, denote its i-th eigenvalue by $\lambda_i(A)$. Denote the k-dimensional identity matrix and all-ones vector by I_k and $\mathbf{1}_k$, respectively. Also, denote the trace operator by Tr . Finally, we treat all vectors as column vectors.

2.1 Policy Gradient

PG schemes seek to maximize the cumulative reward as a function of the policy $\pi_\theta(a|s)$ by performing gradient steps on θ. The celebrated Policy Gradient Theorem [Sutton et al., 1999] states that

$$\frac{\partial}{\partial\theta}\nu^\top V^{\pi_\theta} = \mathbb{E}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left[\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s,a)\right].$$

The variance of the gradient is thus

$$\mathrm{Var}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left(\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s,a)\right). \tag{1}$$

In the notation above, we denote the variance of a vector random variable X by

$$\mathrm{Var}_x\left(X\right) = \mathrm{Tr}\left[\mathbb{E}_x\left[\left(X - \mathbb{E}_x X\right)^\top \left(X - \mathbb{E}_x X\right)\right]\right],$$

similarly as in [Greensmith et al., 2004]. From now on, we drop the subscript from Var in (1) for brevity. When the action space is discrete, a commonly used parameterized policy is softmax: $\pi_\theta(a|s) \propto \exp\left(\theta(s,a)\right)$, where $\theta : \mathcal{S}\times\mathcal{A}\to\mathbb{R}$ is a state-action parameterization.

3 SoftTreeMax: Exponent of trajectories

We introduce a new family of policies called SoftTreeMax, which are a model-based generalization of the popular softmax. We propose two variants: Cumulative (C-SoftTreeMax) and Exponentiated (E-SoftTreeMax). In both variants, we replace the generic softmax logits $\theta(s,a)$ with the score of a trajectory of horizon d starting from (s,a), generated by applying a behavior policy π_b. In C-SoftTreeMax, we exponentiate the expectation of the logits. In E-SoftTreeMax, we first exponentiate the logits and then only compute their expectation.

Logits. We define the SoftTreeMax logit $\ell_{s,a}(d;\theta)$ to be the random variable depicting the score of a trajectory of horizon d starting from (s,a) and following the policy π_b:

$$\ell_{s,a}(d;\theta) = \gamma^{-d}\left[\sum_{t=0}^{d-1}\gamma^t r_t + \gamma^d\theta(s_d)\right]. \tag{2}$$

In the above expression, note that $s_0 = s,\ a_0 = a,\ a_t \sim \pi_b(\cdot|s_t)\ \forall t \geq 1$, and $r_t \equiv r\left(s_t, a_t\right)$. For brevity of the analysis, we let the parametric score θ in (2) be state-based, similarly to a value function. Instead, one could use a state-action input analogous to a Q-function. Thus, SoftTreeMax can be integrated into the two types of implementation of RL algorithms in standard packages. Lastly, the preceding γ^{-d} scales the θ parametrization to correspond to its softmax counerpart.

C-SoftTreeMax. Given an inverse temperature parameter β, we let C-SoftTreeMax be

$$\pi^{\mathrm{C}}_{d,\theta}(a|s) \propto \exp\left[\beta\mathbb{E}^{\pi_b}\ell_{s,a}(d;\theta)\right]. \tag{3}$$

C-SoftTreeMax gives higher weight to actions that result in higher expected returns. While standard softmax relies entirely on parametrization θ, C-SoftTreeMax also interpolates a Monte-Carlo portion of the reward.

E-SoftTreeMax. The second operator we propose is E-SoftTreeMax:

$$\pi^{\mathrm{E}}_{d,\theta}(a|s) \propto \mathbb{E}^{\pi_b}\exp\left[\left(\beta\ell_{s,a}(d;\theta)\right)\right]; \tag{4}$$

here, the expectation is taken outside the exponent. This objective corresponds to the exponentiated reward objective which is often used for risk-sensitive RL [Howard and Matheson, 1972, Fei et al., 2021, Noorani and Baras, 2021]. The common risk-sensitive objective is of the form $\log\mathbb{E}[\exp(\delta R)]$, where δ is the risk parameter and R is the cumulative reward. Similarly to that literature, the exponent in (4) emphasizes the most promising trajectories.

SoftTreeMax properties. SoftTreeMax is a natural model-based generalization of softmax. For $d = 0$, both variants above coincide since (2) becomes deterministic. In that case, for a state-action parametrization, they reduce to standard softmax. When $\beta\to 0$, both variants again coincide and sample actions uniformly (exploration). When $\beta\to\infty$, the policies become deterministic and

greedily optimize for the best trajectory (exploitation). For C-SoftTreeMax, the best trajectory is defined in expectation, while for E-SoftTreeMax it is defined in terms of the best sample path.

SoftTreeMax convergence. Under regularity conditions, for any parametric policy, PG converges to local optima [Bhatnagar et al., 2009], and thus also SoftTreeMax. For softmax PG, asymptotic [Agarwal et al., 2021] and rate results [Mei et al., 2020b] were recently obtained, by showing that the gradient is strictly positive everywhere [Mei et al., 2020b, Lemmas 8-9]. We conjecture that SoftTreeMax satisfies the same property, being a generalization of softmax, but formally proving it is subject to future work.

SoftTreeMax gradient. The two variants of SoftTreeMax involve an expectation taken over S^d many trajectories from the root state s and weighted according to their probability. Thus, during the PG training process, the gradient $\nabla_\theta \log \pi_\theta$ is calculated using a weighted sum of gradients over all reachable states starting from s. Our method exploits the exponential number of trajectories to reduce the variance while improving performance. Indeed, in the next section we prove that the gradient variance of SoftTreeMax decays exponentially fast as a function of the behavior policy π_b and trajectory length d. In the experiments in Section 6, we also show how the practical version of SoftTreeMax achieves a significant reduction in the noise of the PG process and leads to faster convergence and higher reward.

4 Theoretical Analysis

In this section, we first bound the variance of PG when using the SoftTreeMax policy. Later, we discuss how the gradient bias resulting due to approximate forward models diminishes as a function of the approximation error, while retaining the same variance decay.

We show that the variance decreases exponentially with the tree depth, and the rate is determined by the second eigenvalue of the transition kernel induced by π_b. Specifically, we bound the same expression for variance as appears in [Greensmith et al., 2004, Sec. 3.5] and [Wu et al., 2018, Sec. A, Eq. (21)]. Other types of analysis could instead have focused on the estimation aspect in the context of sampling [Zhang et al., 2021, Shen et al., 2019, Pham et al., 2020]. Indeed, in our implementation in Section 5, we manage to avoid sampling and directly compute the expectations in Eqs. (3) and (4). As we show later, we do so by leveraging efficient parallel simulation on the GPU in feasible run-time. In our application, due to the nature of the finite action space and quasi-deterministic Atari dynamics [Bellemare et al., 2013], our expectation estimator is noiseless. We encourage future work to account for the finite-sample variance component. We defer all the proofs to Appendix A.

We begin with a general variance bound that holds for any parametric policy.

Lemma 4.1 (Bound on the policy gradient variance)**.** *Let $\nabla_\theta \log \pi_\theta(\cdot|s) \in \mathbb{R}^{A\times \dim(\theta)}$ be a matrix whose a-th row is $\nabla_\theta \log \pi_\theta(a|s)^\top$. For any parametric policy π_θ and function $Q^{\pi_\theta} : \mathcal{S}\times\mathcal{A}\to\mathbb{R}$,*

$$\mathrm{Var}\left(\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s,a)\right) \le \max_{s,a}\left[Q^{\pi_\theta}(s,a)\right]^2 \max_s \|\nabla_\theta \log \pi_\theta(\cdot|s)\|_F^2.$$

Hence, to bound (1), it is sufficient to bound the Frobenius norm $\|\nabla_\theta \log \pi_\theta(\cdot|s)\|_F$ for any s.

Note that SoftTreeMax does not reduce the gradient uniformly, which would have been equivalent to a trivial change in the learning rate. While the gradient norm shrinks, the gradient itself scales differently along the different coordinates. This scaling occurs along different eigenvectors, as a function of problem parameters (P, θ) and our choice of behavior policy (π_b), as can be seen in the proof of the upcoming Theorem 4.4. This allows SoftTreeMax to learn a good "shrinkage" that, while reducing the overall gradient, still updates the policy quickly enough. This reduction in norm and variance resembles the idea of gradient clipping Zhang et al. [2019], where the gradient is scaled to reduce its variance, thus increasing stability and improving overall performance.

A common assumption in the RL literature [Szepesvári, 2010] that we adopt for the remainder of the section is that the transition matrix P^{π_b}, induced by the behavior policy π_b, is irreducible and aperiodic. Consequently, its second highest eigenvalue satisfies $|\lambda_2(P^{\pi_b})| < 1$.

From now on, we divide the variance results for the two variants of SoftTreeMax into two subsections. For C-SoftTreeMax, the analysis is simpler and we provide an exact bound. The case of E-SoftTreeMax is more involved and we provide for it a more general result. In both cases, we show that the variance decays exponentially with the planning horizon.

4.1 Variance of C-SoftTreeMax

We express C-SoftTreeMax in vector form as follows.

Lemma 4.2 (Vector form of C-SoftTreeMax)**.** *For $d \geq 1$, (3) is given by*

$$\pi^{C}_{d,\theta}(\cdot|s) = \frac{\exp\left[\beta\left(C_{s,d} + P_s\left(P^{\pi_b}\right)^{d-1}\Theta\right)\right]}{\boldsymbol{1}_A^\top \exp\left[\beta\left(C_{s,d} + P_s\left(P^{\pi_b}\right)^{d-1}\Theta\right)\right]}, \tag{5}$$

where

$$C_{s,d} = \gamma^{-d} R_s + P_s\left[\sum_{h=1}^{d-1}\gamma^{h-d}\left(P^{\pi_b}\right)^{h-1}\right] R_{\pi_b}.$$

The vector $C_{s,d} \in \mathbb{R}^A$ represents the cumulative discounted reward in expectation along the trajectory of horizon d. This trajectory starts at state s, involves an initial reward dictated by R_s and an initial transition as per P_s. Thereafter, it involves rewards and transitions specified by R_{π_b} and P^{π_b}, respectively. Once the trajectory reaches depth d, the score function $\theta(s_d)$ is applied,.

Lemma 4.3 (Gradient of C-SoftTreeMax)**.** *The C-SoftTreeMax gradient is given by*

$$\nabla_\theta \log \pi^{C}_{d,\theta} = \beta\left[I_A - \boldsymbol{1}_A(\pi^{C}_{d,\theta})^\top\right] P_s\left(P^{\pi_b}\right)^{d-1},$$

in $\mathbb{R}^{A\times S}$, where for brevity, we drop the s index in the policy above, i.e., $\pi^{C}_{d,\theta} \equiv \pi^{C}_{d,\theta}(\cdot|s)$.

We are now ready to present our first main result:

Theorem 4.4 (Variance decay of C-SoftTreeMax)**.** *For every $Q : \mathcal{S}\times\mathcal{A} \to \mathbb{R}$, the C-SoftTreeMax policy gradient variance is bounded by*

$$\mathrm{Var}\left(\nabla_\theta \log \pi^{C}_{d,\theta}(a|s)Q(s,a)\right) \leq 2\frac{A^2S^2\beta^2}{(1-\gamma)^2}|\lambda_2(P^{\pi_b})|^{2(d-1)}.$$

We provide the full proof in Appendix A.4, and briefly outline its essence here.

Proof outline. Lemma 4.1 allows us to bound the variance using a direct bound on the gradient norm. The gradient is given in Lemma 4.3 as a product of three matrices, which we now study from right to left. The matrix P^{π_b} is a row-stochastic matrix. Because the associated Markov chain is irreducible and aperiodic, it has a unique stationary distribution. This implies that P^{π_b} has one and only one eigenvalue equal to 1; all others have magnitude strictly less than 1. Let us suppose that all these other eigenvalues have multiplicity 1 (the general case with repeated eigenvalues can be handled via Jordan decompositions as in [Pelletier, 1998, Lemma1]). Then, P^{π_b} has the spectral decomposition $P^{\pi_b} = \mathbf{1}_S\mu_{\pi_b}^\top + \sum_{i=2}^{S}\lambda_i v_i u_i^\top$, where λ_i is the i-th eigenvalue of P^{π_b} (ordered in descending order according to their magnitude) and u_i and v_i are the corresponding left and right eigenvectors, respectively, and therefore $(P^{\pi_b})^{d-1} = \mathbf{1}_S\mu_{\pi_b}^\top + \sum_{i=2}^{S}\lambda_i^{d-1}v_i u_i^\top$.

The second matrix in the gradient relation in Lemma 4.3, P_s, is a rectangular transition matrix that translates the vector of all ones from dimension S to A : $P_s\mathbf{1}_S = \mathbf{1}_A$. Lastly, the first matrix $\left[I_A - \mathbf{1}_A(\pi^{\mathrm{C}}_{d,\theta})^\top\right]$ is a projection whose null-space includes the vector $\mathbf{1}_A$, i.e., $\left[I_A - \mathbf{1}_A(\pi^{\mathrm{C}}_{d,\theta})^\top\right]\mathbf{1}_A = 0$. Combining the three properties above when multiplying the three matrices of the gradient, it is easy to see that the first term in the expression for $(P^{\pi_b})^{d-1}$ gets canceled, and we are left with bounded summands scaled by $\lambda_i(P^{\pi_b})^{d-1}$. Recalling that $|\lambda_i(P^{\pi_b})| < 1$ and that $|\lambda_2| \geq |\lambda_3| \geq \ldots$ for $i = 2,\ldots,S$, we obtain the desired result. □

Theorem 4.4 guarantees that the variance of the gradient decays exponentially with d. It also provides a novel insight for choosing the behavior policy π_b as the policy that minimizes the absolute second eigenvalue of the P^{π_b}. Indeed, the second eigenvalue of a Markov chain relates to its connectivity and its rate of convergence to the stationary distribution [Levin and Peres, 2017].

Optimal variance decay. For the strongest reduction in variance, the behavior policy π_b should be chosen to achieve an induced Markov chain whose transitions are state-independent. In that case, P^{π_b}

is a rank one matrix of the form $\mathbf{1}_S \mu_{\pi_b}^\top$, and $\lambda_2(P^{\pi_b}) = 0$. Then, $\mathrm{Var}\left(\nabla_\theta \log \pi_\theta(a|s) Q(s,a)\right) = 0$. Naturally, this can only be done for pathological MDPs; see Appendix C.1 for a more detailed discussion. Nevertheless, as we show in Section 5, we choose our tree expansion policy to reduce the variance as best as possible.

Worst-case variance decay. In contrast, and somewhat surprisingly, when π_b is chosen so that the dynamics is deterministic, there is no guarantee that it will decay exponentially fast. For example, if P^{π_b} is a permutation matrix, then $\lambda_2(P^{\pi_b}) = 1$, and advancing the tree amounts to only updating the gradient of one state for every action, as in the basic softmax.

4.2 Variance of E-SoftTreeMax

The proof of the variance bound for E-SoftTreeMax is similar to that of C-SoftTreeMax, but more involved. It also requires the assumption that the reward depends only on the state, i.e. $r(s,a) \equiv r(s)$. This is indeed the case in most standard RL environments such as Atari and Mujoco.

Lemma 4.5 (Vector form of E-SoftTreeMax)**.** *For $d \geq 1$, (4) is given by*

$$\pi_{d,\theta}^E(\cdot|s) = \frac{E_{s,d} \exp(\beta\Theta)}{1_A^\top E_{s,d} \exp(\beta\Theta)}, \tag{6}$$

where

$$E_{s,d} = P_s \prod_{h=1}^{d-1} \left(D\left(\exp(\beta\gamma^{h-d} R)\right) P^{\pi_b} \right).$$

The vector R above is the S-dimensional vector whose s-th coordinate is $r(s)$.

The matrix $E_{s,d} \in \mathbb{R}^{A \times S}$ has a similar role to $C_{s,d}$ from (5), but it represents the exponentiated cumulative discounted reward. Accordingly, it is a product of d matrices as opposed to a sum. It captures the expected reward sequence starting from s and then iteratively following P^{π_b}. After d steps, we apply the score function on the last state as in (6).

Lemma 4.6 (Gradient of E-SoftTreeMax)**.** *The E-SoftTreeMax gradient is given by*

$$\nabla_\theta \log \pi_{d,\theta}^E = \beta \left[I_A - \mathbf{1}_A (\pi_{d,\theta}^E)^\top \right] \times \frac{D\left(\pi_{d,\theta}^E\right)^{-1} E_{s,d} D(\exp(\beta\Theta))}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \quad \in \quad \mathbb{R}^{A \times S},$$

where for brevity, we drop the s index in the policy above, i.e., $\pi_{d,\theta}^E \equiv \pi_{d,\theta}^E(\cdot|s)$.

This gradient structure is harder to handle than that of C-SoftTreeMax in Lemma 4.3, but here we also can bound the decay of the variance nonetheless.

Theorem 4.7 (Variance decay of E-SoftTreeMax)**.** *There exists $\alpha \in (0,1)$ such that,*

$$\mathrm{Var}\left(\nabla_\theta \log \pi_{d,\theta}^E(a|s) Q(s,a)\right) \in \mathcal{O}\left(\beta^2 \alpha^{2d}\right),$$

for every Q. Further, if P^{π_b} is reversible or if the reward is constant, then $\alpha = |\lambda_2(P^{\pi_b})|$.

Theory versus Practice. We demonstrate the above result in simulation. We draw a random finite MDP, parameter vector $\Theta \in \mathbb{R}_+^S$, and behavior policy π_b. We then empirically compute the PG variance of E-SoftTreeMax as given in (1) and compare it to $|\lambda_2(P^{\pi_b})|^d$. We repeat this experiment three times for different P^{π_b} : (i) close to uniform, (ii) drawn randomly, and (iii) close to a permutation matrix. As seen in Figure 1, the empirical variance and our bound match almost identically. This also suggests that $\alpha = |\lambda_2(P^{\pi_b})|$ in the general case and not only when P^{π_b} is reversible or when the reward is constant.

4.3 Bias with an Approximate Forward Model

The definition of the two SoftTreeMax variants involves the knowledge of the underlying environment, in particular the value of P and r. However, in practice, we often can only learn approximations of the dynamics from interactions, e.g., using NNs [Ha and Schmidhuber, 2018, Schrittwieser et al., 2020]. Let $\hat{P}$ and $\hat{r}$ denote the approximate kernel and reward functions, respectively. In this section, we study the consequences of the approximation error on the C-SoftTreeMax gradient.

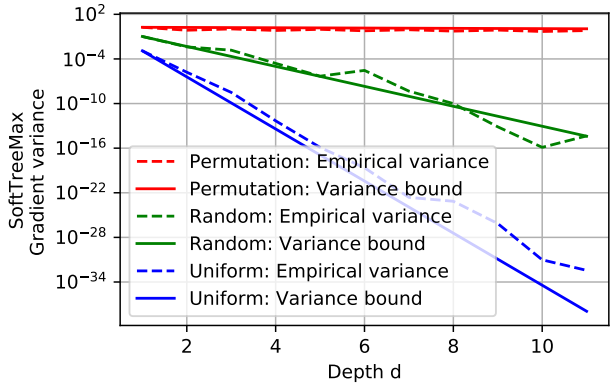

Figure 1: A comparison of the empirical PG variance and our bound for E-SoftTreeMax on randomly drawn MDPs. We present three cases for P^{π_b} : (i) close to uniform, (ii) drawn randomly, and (iii) close to a permutation matrix. This experiment verifies the optimal and worse-case rate decay cases. The variance bounds here are taken from Theorem 4.7 where we substitute $\alpha = |\lambda_2(P^{\pi_b})|$. To account for the constants, we match the values for the first point in $d = 1$.

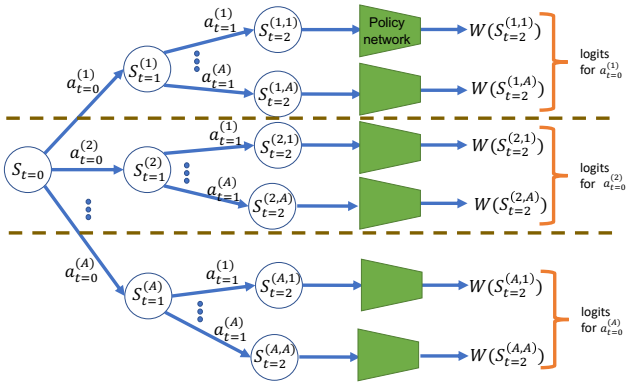

Figure 2: **SoftTreeMax policy**. Our exhaustive parallel tree expansion iterates on all actions at each state up to depth d ($= 2$ here). The leaf state of every trajectory is used as input to the policy network. The output is then added to the trajectory's cumulative reward as described in (2). I.e., instead of the standard softmax logits, we add the cumulative discounted reward to the policy network output. This policy is differentiable and can be easily integrated into any PG algorithm. In this work, we build on PPO and use its loss function to train the policy network.

Let $\hat{\pi}^{\mathrm{C}}_{d,\theta}$ be the C-SoftTreeMax policy defined given the approximate forward model introduced above. That is, let $\hat{\pi}^{\mathrm{C}}_{d,\theta}$ be defined exactly as in (5), but using $\hat{R}_s, \hat{P}_s, \hat{R}_{\pi_b}$ and $\hat{P}^{\pi_b}$, instead of their unperturbed counterparts from Section 2. Then, the variance of the corresponding gradient again decays exponentially with a decay rate of $\lambda_2(\hat{P}^{\pi_b})$. However, a gradient bias is introduced. In the following, we bound this bias in terms of the approximation error and other problem parameters. The proof is provided in Appendix A.9.

Theorem 4.8. *Let ϵ be the maximal model mis-specification, i.e., let $\max\{\|P - \hat{P}\|, \|r - \hat{r}\|\} = \epsilon$. Then the policy gradient bias due to $\hat{\pi}^{C}_{d,\theta}$ satisfies*

$$\left\| \frac{\partial}{\partial \theta} \left(\nu^\top V^{\pi^{C}_{d,\theta}} \right) - \frac{\partial}{\partial \theta} \left(\nu^\top V^{\hat{\pi}^{C}_{d,\theta}} \right) \right\| = \quad \mathcal{O}\left(\frac{1}{(1-\gamma)^2} S \beta^2 d \epsilon \right). \tag{7}$$

To the best of our knowledge, Theorem 4.8 is the first result that bounds the bias of the gradient of a parametric policy due to an approximate model. It states that if the learned model is accurate enough, we expect similar convergence properties for C-SoftTreeMax as we would have obtained with the true dynamics. It also suggests that higher temperature (lower β) reduces the bias. In this case, the logits get less weight, with the extreme of $\beta = 0$ corresponding to a uniform policy that has no bias. Lastly, the error scales linearly with d : the policy suffers from cumulative error as it relies on further-looking states in the approximate model.

5 SoftTreeMax: Deep Parallel Implementation

Following impressive successes of deep RL [Mnih et al., 2015, Silver et al., 2016], using deep NNs in RL is standard practice. Depending on the RL algorithm, a loss function is defined and gradients on the network weights can be calculated. In PG methods, the scoring function used in the softmax is commonly replaced by a neural network W_θ: $\pi_\theta(a|s) \propto \exp\left(W_\theta(s,a)\right)$. Similarly, we implement SoftTreeMax by replacing $\theta(s)$ in (2) with a neural network $W_\theta(s)$. Although both variants of

SoftTreeMax from Section 3 involve computing an expectation, this can be hard in general. One approach to handle it is with sampling, though these introduce estimation variance into the process. We leave the question of sample-based theory and algorithmic implementations for future work.

Instead, in finite action space environments such as Atari, we compute the exact expectation in SoftTreeMax with an exhaustive TS of depth d. Despite the exponential computational cost of spanning the entire tree, recent advancements in parallel GPU-based simulation allow efficient expansion of all nodes at the same depth simultaneously [Dalal et al., 2021, Rosenberg et al., 2022]. This is possible when a simulator is implemented on GPU [Dalton et al., 2020, Makoviychuk et al., 2021, Freeman et al., 2021], or when a forward model is learned [Kim et al., 2020, Ha and Schmidhuber, 2018]. To reduce the complexity to be linear in depth, we apply tree pruning to a limited width in all levels. We do so by sub-sampling only the most promising branches at each level. Limiting the width drastically improves runtime, and enables respecting GPU memory limits, with only a small sacrifice in performance.

To summarize, in the practical SoftTreeMax algorithm we perform an exhaustive tree expansion with pruning to obtain trajectories up to depth d. We expand the tree with equal weight to all actions, which corresponds to a uniform tree expansion policy π_b. We apply a neural network on the leaf states, and accumulate the result with the rewards along each trajectory to obtain the logits in (2). Finally, we aggregate the results using C-SoftTreeMax. We leave experiments E-SoftTreeMax for future work on risk-averse RL. During training, the gradient propagates to the NN weights of W_θ. When the gradient $\nabla_\theta \log \pi_{d,\theta}$ is calculated at each time step, it updates W_θ for all leaf states, similarly to Siamese networks [Bertinetto et al., 2016]. An illustration of the policy is given in Figure 2.

6 Experiments

We conduct our experiments on multiple games from the Atari simulation suite [Bellemare et al., 2013]. As a baseline, we train a PPO [Schulman et al., 2017] agent with 256 GPU workers in parallel [Dalton et al., 2020]. For the tree expansion, we employ a GPU breadth-first as in [Dalal et al., 2021]. We then train C-SoftTreeMax [1] for depths $d = 1 \dots 8$, with a single worker. For depths $d \geq 3$, we limited the tree to a maximum width of 1024 nodes and pruned trajectories with low estimated weights. Since the distributed PPO baseline advances significantly faster in terms of environment steps, for a fair comparison, we ran all experiments for one week on the same machine. For more details see Appendix B.

In Figure 3, we plot the reward and variance of SoftTreeMax for each game, as a function of depth. The dashed lines are the results for PPO. Each value is taken after convergence, i.e., the average over the last 20% of the run. The numbers represent the average over five seeds per game. The plot conveys three intriguing conclusions. First, in all games, SoftTreeMax achieves significantly higher reward than PPO. Its gradient variance is also orders of magnitude lower than that of PPO. Second, the reward and variance are negatively correlated and mirror each other in almost all games. This phenomenon demonstrates the necessity of reducing the variance of PG for improving performance. Lastly, each game has a different sweet spot in terms of optimal tree depth. Recall that we limit the run-time in all experiments to one week The deeper the tree, the slower each step and the run consists of less steps. This explains the non-monotone behavior as a function of depth. For a more thorough discussion on the sweet spot of different games, see Appendix B.3.

7 Related Work

Softmax Operator. The softmax policy became a canonical part of PG to the point where theoretical results of PG focus specifically on it [Zhang et al., 2021, Mei et al., 2020b, Li et al., 2021, Ding et al., 2022]. Even though we focus on a tree extension to the softmax policy, our methodology is general and can be easily applied to other discrete or continuous parameterized policies as in [Mei et al., 2020a, Miahi et al., 2021, Silva et al., 2019]. **Tree Search.** One famous TS algorithm is Monte-Carlo TS (MCTS; [Browne et al., 2012]) used in AlphaGo [Silver et al., 2016] and MuZero [Schrittwieser et al., 2020]. Other algorithms such as Value Iteration, Policy Iteration and DQN were also shown to

[1] We also experimented with E-SoftTreeMax and the results were almost identical. This is due to the quasi-deterministic nature of Atari, which causes the trajectory logits (2) to have almost no variability. We encourage future work on E-SoftTreeMax using probabilistic environments that are risk-sensitive.

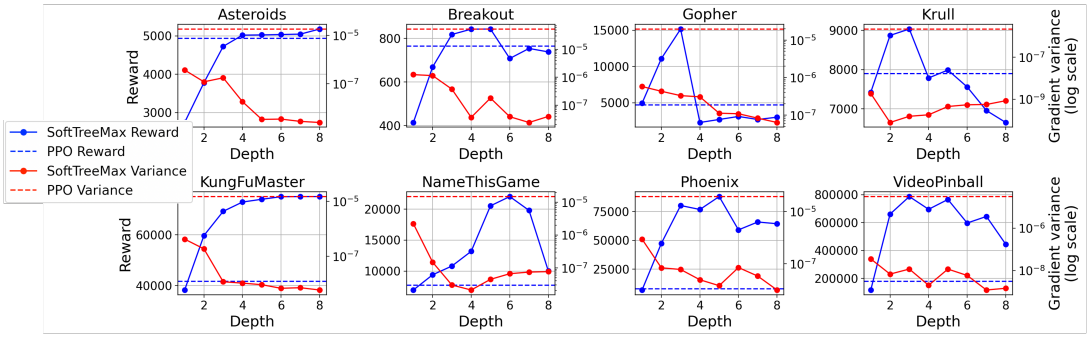

Figure 3: **Reward and Gradient variance: GPU SoftTreeMax (single worker) vs PPO (256 GPU workers).** The blue reward plots show the average of 50 evaluation episodes. The red variance plots show the average gradient variance of the corresponding training runs, averaged over five seeds. The dashed lines represent the same for PPO. Note that the variance y-axis is in log-scale.

give an improved performance with a tree search extensions [Efroni et al., 2019, Dalal et al., 2021]. **Parallel Environments.** In this work we used accurate parallel models that are becoming more common with the increasing popularity of GPU-based simulation [Makoviychuk et al., 2021, Dalton et al., 2020, Freeman et al., 2021]. Alternatively, in relation to Theorem 4.8, one can rely on recent works that learn the underlying model [Ha and Schmidhuber, 2018, Schrittwieser et al., 2020] and use an approximation of the true dynamics. **Risk Aversion.** Previous work considered exponential utility functions for risk aversion [Chen et al., 2007, Garcıa and Fernández, 2015, Fei et al., 2021]. This utility function is the same as E-SoftTreeMax formulation from (4), but we have it directly in the policy instead of the objective. **Reward-free RL.** We showed that the gradient variance is minimized when the transitions induced by the behavior policy π_b are uniform. This is expressed by the second eigenvalue of the transition matrix P^{π_b}. This notion of uniform exploration is common to the reward-free RL setup [Jin et al., 2020]. Several such works also considered the second eigenvalue in their analysis [Liu and Brunskill, 2018, Tarbouriech and Lazaric, 2019].

8 Discussion

In this work, we introduced for the first time a differentiable parametric policy that combines TS with PG. We proved that SoftTreeMax is essentially a variance reduction technique and explained how to choose the expansion policy to minimize the gradient variance. It is an open question whether optimal variance reduction corresponds to the appealing regret properties the were put forward by UCT [Kocsis and Szepesvári, 2006]. We believe that this can be answered by analyzing the convergence rate of SoftTreeMax, relying on the bias and variance results we obtained here.

As the learning process continues, the norm of the gradient and the variance *both* become smaller. On the face of it, one can ask if the gradient becomes small as fast as the variance or even faster can there be any meaningful learning? As we showed in the experiments, learning happens because the variance reduces fast enough (a variance of 0 represents deterministic learning, which is fastest).

Finally, our work can be extended to infinite action spaces. The analysis can be extended to infinite-dimension kernels that retain the same key properties used in our proofs. In the implementation, the tree of continuous actions can be expanded by maintaining a parametric distribution over actions that depend on θ. This approach can be seen as a tree adaptation of MPPI [Williams et al., 2017].

9 Reproducibility and Limitations

In this submission, we include the code as part of the supplementary material. We also include a docker file for setting up the environment and a README file with instructions on how to run both training and evaluation. The environment engine is an extension of Atari-CuLE [Dalton et al., 2020], a CUDA-based Atari emulator that runs on GPU. Our usage of a GPU environment is both a novelty and a current limitation of our work.

References

A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *J. Mach. Learn. Res.*, 22(98):1–76, 2021.

M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.

L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016.

S. Bhatnagar, R. S. Sutton, M. Ghavamzadeh, and M. Lee. Natural actor–critic algorithms. *Automatica*, 45(11):2471–2482, 2009.

C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.

S. Chatterjee and E. Seneta. Towards consensus: Some convergence theorems on repeated averaging. *Journal of Applied Probability*, 14(1):89–97, 1977.

X. Chen, M. Sim, D. Simchi-Levi, and P. Sun. Risk aversion in inventory management. *Operations Research*, 55(5):828–842, 2007.

G. Dalal, A. Hallak, S. Dalton, S. Mannor, G. Chechik, et al. Improve agents without retraining: Parallel tree search with off-policy correction. *Advances in Neural Information Processing Systems*, 34:5518–5530, 2021.

S. Dalton et al. Accelerating reinforcement learning through gpu atari emulation. *Advances in Neural Information Processing Systems*, 33:19773–19782, 2020.

Y. Ding, J. Zhang, and J. Lavaei. On the global optimum convergence of momentum-based policy gradient. In *International Conference on Artificial Intelligence and Statistics*, pages 1910–1934. PMLR, 2022.

Y. Efroni, G. Dalal, B. Scherrer, and S. Mannor. How to combine tree-search methods in reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3494–3501, 2019.

Y. Fei, Z. Yang, Y. Chen, and Z. Wang. Exponential bellman equation and improved regret bounds for risk-sensitive reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 20436–20446, 2021.

C. D. Freeman, E. Frey, A. Raichuk, S. Girgin, I. Mordatch, and O. Bachem. Brax-a differentiable physics engine for large scale rigid body simulation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.

J. Garcıa and F. Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.

E. Greensmith, P. L. Bartlett, and J. Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(9), 2004.

D. Ha and J. Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018.

R. A. Howard and J. E. Matheson. Risk-sensitive markov decision processes. *Management science*, 18(7):356–369, 1972.

C. Jin, A. Krishnamurthy, M. Simchowitz, and T. Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning*, pages 4870–4879. PMLR, 2020.

S. W. Kim, Y. Zhou, J. Philion, A. Torralba, and S. Fidler. Learning to simulate dynamic environments with gamegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1231–1240, 2020.

L. Kocsis and C. Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pages 282–293. Springer, 2006.

D. A. Levin and Y. Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

G. Li, Y. Wei, Y. Chi, Y. Gu, and Y. Chen. Softmax policy gradient methods can take exponential time to converge. In *Conference on Learning Theory*, pages 3107–3110. PMLR, 2021.

Y. Liu and E. Brunskill. When simple exploration is sample efficient: Identifying sufficient conditions for random exploration to yield pac rl algorithms. *arXiv preprint arXiv:1805.09045*, 2018.

Y. Liu, K. Zhang, T. Basar, and W. Yin. An improved analysis of (variance-reduced) policy gradient and natural policy gradient methods. *Advances in Neural Information Processing Systems*, 33: 7624–7636, 2020.

V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.

A. S. Mathkar and V. S. Borkar. Nonlinear gossip. *SIAM Journal on Control and Optimization*, 54 (3):1535–1557, 2016.

J. Mei, C. Xiao, B. Dai, L. Li, C. Szepesvári, and D. Schuurmans. Escaping the gravitational pull of softmax. *Advances in Neural Information Processing Systems*, 33:21130–21140, 2020a.

J. Mei, C. Xiao, C. Szepesvari, and D. Schuurmans. On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR, 2020b.

E. Miahi, R. MacQueen, A. Ayoub, A. Masoumzadeh, and M. White. Resmax: An alternative soft-greedy operator for reinforcement learning. 2021.

V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

E. Noorani and J. S. Baras. Risk-sensitive reinforce: A monte carlo policy gradient algorithm for exponential performance criteria. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 1522–1527. IEEE, 2021.

M. Papini, D. Binaghi, G. Canonaco, M. Pirotta, and M. Restelli. Stochastic variance-reduced policy gradient. In *International conference on machine learning*, pages 4026–4035. PMLR, 2018.

M. Pelletier. On the almost sure asymptotic behaviour of stochastic algorithms. *Stochastic processes and their applications*, 78(2):217–244, 1998.

N. Pham, L. Nguyen, D. Phan, P. H. Nguyen, M. Dijk, and Q. Tran-Dinh. A hybrid stochastic policy gradient algorithm for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 374–385. PMLR, 2020.

A. Raffin, A. Hill, M. Ernestus, A. Gleave, A. Kanervisto, and N. Dormann. Stable baselines3, 2019.

A. Rosenberg, A. Hallak, S. Mannor, G. Chechik, and G. Dalal. Planning and learning with adaptive lookahead. *arXiv preprint arXiv:2201.12403*, 2022.

J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Z. Shen, A. Ribeiro, H. Hassani, H. Qian, and C. Mi. Hessian aided policy gradient. In *International conference on machine learning*, pages 5729–5738. PMLR, 2019.

A. Silva, T. Killian, I. D. J. Rodriguez, S.-H. Son, and M. Gombolay. Optimization methods for interpretable differentiable decision trees in reinforcement learning. *arXiv preprint arXiv:1903.09338*, 2019.

D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

C. Szepesvári. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.

J. Tarbouriech and A. Lazaric. Active exploration in markov decision processes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 974–982. PMLR, 2019.

P. S. Thomas and E. Brunskill. Policy gradient methods for reinforcement learning with function approximation and action-dependent baselines. *arXiv preprint arXiv:1706.06643*, 2017.

G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou. Information theoretic mpc for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721. IEEE, 2017.

C. Wu, A. Rajeswaran, Y. Duan, V. Kumar, A. M. Bayen, S. Kakade, I. Mordatch, and P. Abbeel. Variance reduction for policy gradient with action-dependent factorized baselines. In *International Conference on Learning Representations*, 2018.

P. Xu, F. Gao, and Q. Gu. An improved convergence analysis of stochastic variance-reduced policy gradient. In *Uncertainty in Artificial Intelligence*, pages 541–551. PMLR, 2020.

J. Zhang, T. He, S. Sra, and A. Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. *arXiv preprint arXiv:1905.11881*, 2019.

J. Zhang, C. Ni, C. Szepesvari, M. Wang, et al. On the convergence and sample efficiency of variance-reduced policy gradient method. *Advances in Neural Information Processing Systems*, 34: 2228–2240, 2021.

Appendix

A Proofs

A.1 Proof of Lemma 4.1 – Bound on the policy gradient variance

For any parametric policy π_θ and function $Q : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$,

$$\mathrm{Var}\left(\nabla_\theta \log \pi_\theta(a|s) Q(s,a)\right) \leq \max_{s,a} \left[Q(s,a)\right]^2 \max_s \|\nabla_\theta \log \pi_\theta(\cdot|s)\|_F^2,$$

where $\nabla_\theta \log \pi_\theta(\cdot|s) \in \mathbb{R}^{A \times \dim(\theta)}$ is a matrix whose a-th row is $\nabla_\theta \log \pi_\theta(a|s)^\top$.

Proof. The variance for a parametric policy π_θ is given as follows:

$$\begin{aligned}\mathrm{Var}\left(\nabla_\theta \log \pi_\theta(a|s) Q(a,s)\right) =& \mathbb{E}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left[\nabla_\theta \log \pi_\theta(a|s)^\top \nabla_\theta \log \pi_\theta(a|s) Q(s,a)^2\right] - \\ & \mathbb{E}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left[\nabla_\theta \log \pi_\theta(a|s) Q(s,a)\right]^\top \mathbb{E}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left[\nabla_\theta \log \pi_\theta(a|s) Q(s,a)\right],\end{aligned}$$

where $Q(s,a)$ is the currently estimated Q-function and d_{π_θ} is the discounted state visitation frequency induced by the policy π_θ. Since the second term we subtract is always positive (it is of quadratic form $v^\top v$) we can bound the variance by the first term:

$$
\begin{aligned}
\mathrm{Var}\left(\nabla_\theta \log \pi_\theta(a|s) Q(a,s)\right) \leq & \mathbb{E}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left[\nabla_\theta \log \pi_\theta(a|s)^\top \nabla_\theta \log \pi_\theta(a|s) Q(s,a)^2\right] \\
&= \sum_s d_{\pi_\theta}(s) \sum_a \pi_\theta(a|s) \nabla_\theta \log \pi_\theta(a|s)^\top \nabla_\theta \log \pi_\theta(a|s) Q(s,a)^2 \\
&\leq \max_{s,a}\left[\left[Q(s,a)\right]^2 \pi_\theta(a|s)\right] \sum_s d_{\pi_\theta}(s) \sum_a \nabla_\theta \log \pi_\theta(a|s)^\top \nabla_\theta \log \pi_\theta(a|s) \\
&\leq \max_{s,a}\left[Q(s,a)\right]^2 \max_s \sum_a \nabla_\theta \log \pi_\theta(a|s)^\top \nabla_\theta \log \pi_\theta(a|s) \\
&= \max_{s,a}\left[Q(s,a)\right]^2 \max_s \|\nabla_\theta \log \pi_\theta(\cdot|s)\|_F^2.
\end{aligned}
$$

□

A.2 Proof of Lemma 4.2 – Vector form of C-SoftTreeMax

In vector form, (3) is given by

$$
\pi^{\mathrm{C}}_{d,\theta}(\cdot|s) = \frac{\exp\left[\beta\left(C_{s,d} + P_s\left(P^{\pi_b}\right)^{d-1}\Theta\right)\right]}{\mathbf{1}_A^\top \exp\left[\beta\left(C_{s,d} + P_s\left(P^{\pi_b}\right)^{d-1}\Theta\right)\right]}, \tag{8}
$$

where

$$
C_{s,d} = \gamma^{-d} R_s + P_s\left[\sum_{h=1}^{d-1} \gamma^{h-d}\left(P^{\pi_b}\right)^{h-1}\right] R_{\pi_b}. \tag{9}
$$

Proof. Consider the vector $\ell_{s,\cdot} \in \mathbb{R}^{|\mathcal{A}|}$. Its expectation satisfies

$$
\begin{aligned}
\mathbb{E}^{\pi_b} \ell_{s,\cdot}(d;\theta) &= \mathbb{E}^{\pi_b}\left[\sum_{t=0}^{d-1} \gamma^{t-d} r_t + \theta(s_d)\right] \\
&= \gamma^{-d} R_s + \sum_{t=1}^{d-1} \gamma^{t-d} P_s (P^{\pi_b})^{t-1} R_{\pi_b} + P_s (P^{\pi_b})^{d-1}\Theta.
\end{aligned}
$$

As required. □

A.3 Proof of Lemma 4.3 – Gradient of C-SoftTreeMax

The C-SoftTreeMax gradient of dimension $A \times S$ is given by

$$
\nabla_\theta \log \pi^{\mathrm{C}}_{d,\theta} = \beta\left[I_A - \mathbf{1}_A (\pi^{\mathrm{C}}_{d,\theta})^\top\right] P_s \left(P^{\pi_b}\right)^{d-1},
$$

where for brevity, we drop the s index in the policy above, i.e., $\pi^{\mathrm{C}}_{d,\theta} \equiv \pi^{\mathrm{C}}_{d,\theta}(\cdot|s)$.

Proof. The (j,k)-th entry of $\nabla_\theta \log \pi^{\mathrm{C}}_{d,\theta}$ satisifes

$$
\begin{aligned}
\left[\nabla_\theta \log \pi^{\mathrm{C}}_{d,\theta}\right]_{j,k} &= \frac{\partial \log(\pi^{\mathrm{C}}_{d,\theta}(a^j|s))}{\partial \theta(s^k)} \\
&= \beta[P_s(P^{\pi_b})^{d-1}]_{j,k} - \frac{\sum_a\left[\exp\left[\beta\left(C_{s,d} + P_s\left(P^{\pi_b}\right)^{d-1}\Theta\right)\right]\right]_a \beta\left[P_s(P^{\pi_b})^{d-1}\right]_{a,k}}{\mathbf{1}_A^\top \exp\left[\beta\left(C_{s,d} + P_s\left(P^{\pi_b}\right)^{d-1}\Theta\right)\right]} \\
&= \beta[P_s(P^{\pi_b})^{d-1}]_{j,k} - \beta \sum_a \pi^{\mathrm{C}}_{d,\theta}(a|s)\left[P_s(P^{\pi_b})^{d-1}\right]_{a,k} \\
&= \beta[P_s(P^{\pi_b})^{d-1}]_{j,k} - \beta\left[(\pi^{\mathrm{C}}_{d,\theta})^\top P_s(P^{\pi_b})^{d-1}\right]_k \\
&= \beta[P_s(P^{\pi_b})^{d-1}]_{j,k} - \beta\left[\mathbf{1}_A(\pi^{\mathrm{C}}_{d,\theta})^\top P_s(P^{\pi_b})^{d-1}\right]_{j,k}.
\end{aligned}
$$

Moving back to matrix form, we obtain the stated result. □

A.4 Proof of Theorem 4.4 – Exponential variance decay of C-SoftTreeMax

The C-SoftTreeMax policy gradient is bounded by

$$\mathrm{Var}\left(\nabla_\theta \log \pi^{\mathrm{C}}_{d,\theta}(a|s)Q(s,a)\right) \leq 2\frac{A^2S^2\beta^2}{(1-\gamma)^2}|\lambda_2(P^{\pi_b})|^{2(d-1)}.$$

Proof. We use Lemma 4.1 directly. First of all, it is know that when the reward is bounded in $[0,1]$, the maximal value of the Q-function is $\frac{1}{1-\gamma}$ as the sum as infinite discounted rewards. Next, we bound the Frobenius norm of the term achieved in Lemma 4.3, by applying the eigen-decomposition on P^{π_b}:

$$P^{\pi_b} = \mathbf{1}_S\mu^\top + \sum_{i=2}^{S}\lambda_i u_i v_i^\top, \tag{10}$$

where μ is the stationary distribution of P^{π_b}, and u_i and v_i are left and right eigenvectors correspondingly.

$$\begin{aligned}
\|\beta\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)P_s(P^{\pi_b})^{d-1}\|_F &= \beta\|\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)P_s\left(\mathbf{1}_S\mu^\top+\sum_{i=2}^{S}\lambda_i^{d-1}u_iv_i^\top\right)\|_F \\
\textit{(P_s is stochastic)} \quad &= \beta\|\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)\left(\mathbf{1}_A\mu^\top+\sum_{i=2}^{S}\lambda_i^{d-1}P_su_iv_i^\top\right)\|_F \\
\textit{(projection nullifies $\mathbf{1}_A\mu^\top$)} \quad &= \beta\|\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)\left(\sum_{i=2}^{S}\lambda_i^{d-1}P_su_iv_i^\top\right)\|_F \\
\textit{(triangle inequality)} \quad &\leq \beta\sum_{i=2}^{S}\|\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)\left(\lambda_i^{d-1}P_su_iv_i^\top\right)\|_F \\
\textit{(matrix norm sub-multiplicativity)} \quad &\leq \beta|\lambda_2^{d-1}|\sum_{i=2}^{S}\|I_{A,A}-\mathbf{1}_A\pi^\top\|_F\|P_s\|_F\|u_iv_i^\top\|_F \\
&= \beta|\lambda_2^{d-1}|(S-1)\|I_{A,A}-\mathbf{1}_A\pi^\top\|_F\|P_s\|_F.
\end{aligned}$$

Now, we can bound the norm $\|I_{A,A}-\mathbf{1}_A\pi^\top\|_F$ by direct calculation:

$$\begin{aligned}
\|I_{A,A}-\mathbf{1}_A\pi^\top\|_F^2 &= \mathrm{Tr}\left[\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)^\top\right] &(11)\\
&= \mathrm{Tr}\left[I_{A,A}-\mathbf{1}_A\pi^\top-\pi\mathbf{1}_A^\top+\pi^\top\pi\mathbf{1}_A\mathbf{1}_A^\top\right] &(12)\\
&= A-1-1+A\pi^\top\pi &(13)\\
&\leq 2A. &(14)
\end{aligned}$$

From the Cauchy-Schwartz inequality,

$$\|P_s\|_F^2 = \sum_a\sum_s\left[[P_s]_{a,s}\right]^2 = \sum_a\|[P_s]_{a,\cdot}\|_2^2 \leq \sum_a\|[P_s]_{a,\cdot}\|_1\|[P_s]_{a,\cdot}\|_\infty \leq A.$$

So,

$$\begin{aligned}
\mathrm{Var}\left(\nabla_\theta \log \pi^{\mathrm{C}}_{d,\theta}(a|s)Q(s,a)\right) &\leq \max_{s,a}\left[Q(s,a)\right]^2\max_s\|\nabla_\theta\log\pi^{\mathrm{C}}_{d,\theta}(\cdot|s)\|_F^2 \\
&\leq \frac{1}{(1-\gamma)^2}\|\beta\left(I_{A,A}-\mathbf{1}_A\pi^\top\right)P_s(P^{\pi_b})^{d-1}\|_F^2 \\
&\leq \frac{1}{(1-\gamma)^2}\beta^2|\lambda_2(P^{\pi_b})|^{2(d-1)}S^2(2A^2),
\end{aligned}$$

which obtains the desired bound. □

A.5 A lower bound on C-SoftTreeMax gradient (result not in the paper)

For completeness we also supply a lower bound on the Frobenius norm of the gradient. Note that this result does not translate to the a lower bound on the variance since we have no lower bound equivalence of Lemma 4.1.

Lemma A.1. *The Frobenius norm on the gradient of the policy is lower-bounded by:*

$$\|\nabla_\theta \log \pi^{C}_{d,\theta}(\cdot|s)\|_F \geq C \cdot \beta |\lambda_2(P^{\pi_b})|^{(d-1)}. \tag{15}$$

Proof. We begin by moving to the induced l_2 norm by norm-equivalence:

$$\|\beta \left(I_{A,A} - \mathbf{1}_A \pi^\top\right) P_s (P^{\pi_b})^{d-1}\|_F \geq \|\beta \left(I_{A,A} - \mathbf{1}_A \pi^\top\right) P_s (P^{\pi_b})^{d-1}\|_2.$$

Now, taking the vector u to be the eigenvector of the second eigenvalue of P^{π_b}:

$$\begin{aligned}
\|\beta \left(I_{A,A} - \mathbf{1}_A \pi^\top\right) P_s (P^{\pi_b})^{d-1}\|_2 &\geq \|\beta \left(I_{A,A} - \mathbf{1}_A \pi^\top\right) P_s (P^{\pi_b})^{d-1} u\|_2 \\
&= \beta \| \left(I_{A,A} - \mathbf{1}_A \pi^\top\right) P_s u\|_2 \\
&= \beta |\lambda_2(P^{\pi_b})|^{(d-1)} \| \left(I_{A,A} - \mathbf{1}_A \pi^\top\right) P_s u\|_2.
\end{aligned}$$

Note that even though $P_s u$ can be 0, that is not the common case since we can freely change π_b (and therefore the eigenvectors of P^{π_b}). □

A.6 Proof of Lemma 4.5 – Vector form of E-SoftTreeMax

For $d \geq 1$, (4) is given by

$$\pi^{\mathrm{E}}_{d,\theta}(\cdot|s) = \frac{E_{s,d} \exp(\beta\Theta)}{1_A^\top E_{s,d} \exp(\beta\Theta)}, \tag{16}$$

where

$$E_{s,d} = P_s \prod_{h=1}^{d-1} \left(D\left(\exp[\beta\gamma^{h-d} R]\right) P^{\pi_b}\right) \tag{17}$$

with R being the $|S|$-dimensional vector whose s-th coordinate is $r(s)$.

Proof. Recall that

$$\ell_{s,a}(d;\theta) = \gamma^{-d} \left[r(s) + \sum_{t=1}^{d-1} \gamma^t r(s_t) + \gamma^d \theta(s_d) \right]. \tag{18}$$

and, hence,

$$\exp[\beta \ell_{s,a}(d;\theta)] = \exp\left[\beta\gamma^{-d} \left(r(s) + \sum_{t=1}^{d-1} \gamma^t r(s_t) + \gamma^d \theta(s_d) \right) \right]. \tag{19}$$

Therefore,

$$\mathbb{E}[\exp \beta \ell_{s,a}(d;\theta)] = \mathbb{E}\left[\exp\left[\beta\gamma^{-d} \left(r(s) + \sum_{t=1}^{d-1} \gamma^t r(s_t) \right) \right] \mathbb{E}\left[\exp\left[\beta\left(\theta(s_d)\right)\right] | s_1, \ldots, s_{d-1} \right] \right] \tag{20}$$

$$= \mathbb{E}\left[\exp\left[\beta\gamma^{-d} \left(r(s) + \sum_{t=1}^{d-1} \gamma^t r(s_t) \right) \right] P^{\pi_b}(\cdot|s_{d-1}) \right] \exp(\beta\Theta) \tag{21}$$

$$= \mathbb{E}\left[\exp\left[\beta\gamma^{-d} \left(r(s) + \sum_{t=1}^{d-2} \gamma^t r(s_t) \right) \right] \exp[\beta\gamma^{-1} r(s_{d-1})] P^{\pi_b}(\cdot|s_{d-1}) \right] \exp(\beta\Theta). \tag{22}$$

By repeatedly using iterative conditioning as above, the desired result follows. Note that $\exp(\beta\gamma^{-d} r(s))$ does not depend on the action and is therefore cancelled out with the denominator. □

A.7 Proof of Lemma 4.6 – Gradient of E-SoftTreeMax

The E-SoftTreeMax gradient of dimension $A \times S$ is given by

$$\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} = \beta \left[I_A - \mathbf{1}_A (\pi^{\mathrm{E}}_{d,\theta})^\top \right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} E_{s,d} D(\exp(\beta\Theta))}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)},$$

where for brevity, we drop the s index in the policy above, i.e., $\pi^{\mathrm{E}}_{d,\theta} \equiv \pi^{\mathrm{E}}_{d,\theta}(\cdot|s)$.

Proof. The (j,k)-th entry of $\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}$ satisfies

$$\begin{aligned}
[\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}]_{j,k} &= \frac{\partial \log(\pi^{\mathrm{E}}_{d,\theta}(a^j|s))}{\partial \theta(s^k)} \\
&= \frac{\partial}{\partial \theta(s^k)} \left(\log[(E_{s,d})_j^\top \exp(\beta\Theta)] - \log[\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)] \right) \\
&= \frac{\beta (E_{s,d})_{j,k} \exp(\beta\theta(s^k))}{(E_{s,d})_j^\top \exp(\beta\Theta)} - \frac{\beta \mathbf{1}_A^\top E_{s,d} e_k \exp(\beta\theta(s^k))}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \\
&= \frac{\beta (E_{s,d} e_k \exp(\beta\theta(s^k)))_j}{(E_{s,d})_j^\top \exp(\beta\Theta)} - \frac{\beta \mathbf{1}_A^\top E_{s,d} e_k \exp(\beta\theta(s^k))}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \\
&= \beta \left[\frac{e_j^\top}{e_j^\top E_{s,d} \exp(\beta\Theta)} - \frac{\mathbf{1}_A^\top}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \right] E_{s,d} e_k \exp(\beta\theta(s^k)).
\end{aligned}$$

Hence,

$$[\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}]_{\cdot,k} = \beta \left[D(E_{s,d} \exp(\beta\Theta))^{-1} - (\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta))^{-1} \mathbf{1}_A \mathbf{1}_A^\top \right] E_{s,d} e_k \exp(\beta\theta(s^k))$$

From this, it follows that

$$\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} = \beta \left[D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} - \mathbf{1}_A \mathbf{1}_A^\top \right] \frac{E_{s,d} D(\exp(\beta\Theta))}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)}. \tag{23}$$

The desired result is now easy to see. □

A.8 Proof of Theorem 4.7 — Exponential variance decay of E-SoftTreeMax

There exists $\alpha \in (0,1)$ such that, for any function $Q : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$,

$$\mathrm{Var}\left(\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}(a|s) Q(s,a) \right) \in \mathcal{O}\left(\beta^2 \alpha^{2d}\right).$$

If all rewards are equal ($r \equiv$ const), then $\alpha = |\lambda_2(P^{\pi_b})|$.

Proof outline. Recall that thanks to Lemma 4.1, we can bound the PG variance using a direct bound on the gradient norm. The definition of the induced norm is

$$\|\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}\| = \max_{z:\|z\|=1} \|\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} z\|,$$

with $\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}$ given in Lemma 4.6. Let $z \in \mathbb{R}^S$ be an arbitrary vector such that $\|z\| = 1$. Then, $z = \sum_{i=1}^S c_i z_i$, where c_i are scalar coefficients and z_i are vectors spanning the S-dimensional space. In the full proof, we show our specific choice of z_i and prove they are linearly independent given that choice. We do note that $z_1 = \mathbf{1}_S$.

The first part of the proof relies on the fact that $(\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}) z_1 = 0$. This is easy to verify using Lemma 4.6 together with (6), and because $\left[I_A - \mathbf{1}_A (\pi^{\mathrm{E}}_{d,\theta})^\top \right]$ is a projection matrix whose null-space is spanned by $\mathbf{1}_S$. Thus,

$$\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} z = \nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} \sum_{i=2}^S c_i z_i.$$

In the second part of the proof, we focus on $E_{s,d}$ from (6), which appears within $\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta}$. Notice that $E_{s,d}$ consists of the product $\prod_{h=1}^{d-1} \left(D\left(\exp(\beta\gamma^{h-d}R\right) P^{\pi_b}\right)$. Even though the elements in this product are not stochastic matrices, in the full proof we show how to normalize each of them to a stochastic matrix B_h. We thus obtain that

$$E_{s,d} = P_s D(M_1) \prod_{h=1}^{d-1} B_h,$$

where $M_1 \in \mathbb{R}^S$ is some strictly positive vector. Then, we can apply a result by Mathkar and Borkar [2016], which itself builds on [Chatterjee and Seneta, 1977]. The result states that the product of stochastic matrices $\prod_{h=1}^{d-1} B_h$ of our particular form converges exponentially fast to a matrix of the form $\mathbf{1}_S \mu^\top$ s.t. $\|\mathbf{1}_S \mu^\top - \prod_{h=1}^{d-1} B_h\| \leq C\alpha^d$ for some constant C.

Lastly, $\mathbf{1}_S \mu_{\pi_b}^\top$ gets canceled due to our choice of $z_i,\ i = 2, \ldots, S$. This observation along with the above fact that the remainder decays then shows that $\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} \sum_{i=2}^{S} z_i = \mathcal{O}(\alpha^d)$, which gives the desired result. □

Full technical proof. Let $d \geq 2$. Recall that

$$E_{s,d} = P_s \prod_{h=1}^{d-1} \left(D\left(\exp[\beta\gamma^{h-d}R]\right) P^{\pi_b}\right), \tag{24}$$

and that R refers to the S-dimensional vector whose s-th coordinate is $r(s)$. Define

$$B_i = \begin{cases} P^{\pi_b} & \text{if } i = d-1, \\ D^{-1}(P^{\pi_b} M_{i+1}) P^{\pi_b} D(M_{i+1}) & \text{if } i = 1, \ldots, d-2, \end{cases} \tag{25}$$

and the vector

$$M_i = \begin{cases} \exp(\beta\gamma^{-1}R) & \text{if } i = d-1, \\ \exp(\beta\gamma^{i-d}R) \circ P^{\pi_b} M_{i+1} & \text{if } i = 1, \ldots, d-2, \end{cases} \tag{26}$$

where $\circ$ denotes the element-wise product. Then,

$$E_{s,d} = P_s D(M_1) \prod_{i=1}^{d-1} B_i. \tag{27}$$

It is easy to see that each B_i is a row-stochastic matrix, i.e., all entries are non-negative and $B_i \mathbf{1}_S = \mathbf{1}_S$.

Next, we prove that all non-zeros entries of B_i are bounded away from 0 by a constant. This is necessary to apply the next result from Chatterjee and Seneta [1977]. The j-th coordinate of M_i satisfies

$$(M_i)_j = \exp[\beta\gamma^{i-d}R_j] \sum_k [P^{\pi_b}]_{j,k} (M_{i+1})_k \leq \|\exp[\beta\gamma^{i-d}R]\|_\infty \|M_{i+1}\|_\infty. \tag{28}$$

Separately, observe that $\|M_{d-1}\|_\infty \leq \|\exp(\beta\gamma^{-1}R)\|_\infty$. Plugging these relations in (26) gives

$$\|M_1\|_\infty \leq \prod_{h=1}^{d-1} \|\exp[\beta\gamma^{h-d}R]\|_\infty = \prod_{h=1}^{d-1} \|\exp[\beta\gamma^{-d}R]\|_\infty^{\gamma^h} = \|\exp[\beta\gamma^{-d}R]\|_\infty^{\sum_{h=1}^{d-1}\gamma^h} \leq \|\exp[\beta\gamma^{-d}R]\|_\infty^{\frac{1}{1-\gamma}}. \tag{29}$$

Similarly, for every $1 \leq i \leq d-1$, we have that

$$\|M_i\|_\infty \leq \prod_{h=i}^{d-1} \|\exp[\beta\gamma^{-d}R]\|_\infty^{\gamma^h} \leq \|\exp[\beta\gamma^{-d}R]\|_\infty^{\frac{1}{1-\gamma}}. \tag{30}$$

The jk-th entry of $B_i = D^{-1}(P^{\pi_b} M_{i+1}) P^{\pi_b} D(M_{i+1})$ is

$$(B_i)_{jk} = \frac{P^{\pi_b}_{jk} [M_{i+1}]_k}{\sum_{\ell=1}^{|S|} P^{\pi_b}_{j\ell} [M_{i+1}]_\ell} \geq \frac{P^{\pi_b}_{jk}}{\sum_{\ell=1}^{|S|} P^{\pi_b}_{j\ell} [M_{i+1}]_\ell} \geq \frac{P^{\pi_b}_{jk}}{\|\exp[\beta\gamma^{-d}R]\|_\infty^{\frac{1}{1-\gamma}}}. \tag{31}$$

Hence, for non-zero $P^{\pi_b}_{jk}$, the entries are bounded away from zero by the same. We can now proceed with applying the following result.

Now, by [Chatterjee and Seneta, 1977, Theorem 5] (see also (14) in [Mathkar and Borkar, 2016]), $\lim_{d\to\infty}\prod_{i=1}^{d-1} B_i$ exists and is of the form $\mathbf{1}_S\mu^\top$ for some probability vector μ. Furthermore, there is some $\alpha \in (0,1)$ such that $\varepsilon(d) := \left(\prod_{i=1}^{d-1} B_i\right) - \mathbf{1}_S\,\mu^\top$ satisfies

$$\|\varepsilon(d)\| = O(\alpha^d). \tag{32}$$

Pick linearly independent vectors $w_2, \ldots, w_S$ such that

$$\mu^\top w_i = 0 \text{ for } i = 2, \ldots, d. \tag{33}$$

Since $\sum_{i=2}^S \alpha_i w_i$ is perpendicular to μ for any $\alpha_2, \ldots \alpha_S$ and because $\mu^\top \exp(\beta\Theta) > 0$, there exists no choice of $\alpha_2, \ldots, \alpha_S$ such that $\sum_{i=2}^S \alpha_i w_i = \exp(\beta\Theta)$. Hence, if we let $z_1 = \mathbf{1}_S$ and $z_i = D(\exp(\beta\Theta))^{-1} w_i$ for $i = 2, \ldots, S$, then it follows that $\{z_1, \ldots, z_S\}$ is linearly independent. In particular, it implies that $\{z_1, \ldots, z_S\}$ spans $\mathbb{R}^S$.

Now consider an arbitrary unit norm vector $z := \sum_{i=1}^S c_i z_i \in \mathbb{R}^S$ s.t. $\|z\|_2 = 1$. Then,

$$\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} z = \nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} \sum_{i=2}^S c_i z_i \tag{34}$$

$$= \beta\left[I_A - \mathbf{1}_A(\pi^{\mathrm{E}}_{d,\theta})^\top\right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} E_{s,d} D(\exp(\beta\Theta))}{\mathbf{1}_A^\top E_{s,d}\exp(\beta\Theta)} \sum_{i=2}^S c_i z_i \tag{35}$$

$$= \beta\left[I_A - \mathbf{1}_A(\pi^{\mathrm{E}}_{d,\theta})^\top\right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} E_{s,d}}{\mathbf{1}_A^\top E_{s,d}\exp(\beta\Theta)} \sum_{i=2}^S c_i w_i \tag{36}$$

$$= \beta\left[I_A - \mathbf{1}_A(\pi^{\mathrm{E}}_{d,\theta})^\top\right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} \left[\mathbf{1}_S\mu^\top + \varepsilon(d)\right]}{\mathbf{1}_A^\top E_{s,d}\exp(\beta\Theta)} \sum_{i=2}^S c_i w_i \tag{37}$$

$$= \beta\left[I_A - \mathbf{1}_A(\pi^{\mathrm{E}}_{d,\theta})^\top\right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} \varepsilon(d)}{\mathbf{1}_A^\top E_{s,d}\exp(\beta\Theta)} \sum_{i=2}^S c_i w_i \tag{38}$$

$$= \beta\left[I_A - \mathbf{1}_A(\pi^{\mathrm{E}}_{d,\theta})^\top\right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1} \varepsilon(d) D(\exp(\beta\Theta))}{\mathbf{1}_A^\top E_{s,d}\exp(\beta\Theta)} (z - c_1\mathbf{1}_S), \tag{39}$$

where (34) follows from the fact that $\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} z_1 = \nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} \mathbf{1}_S = 0$, (35) follows from Lemma 4.6, (36) holds since $z_i = D(\exp(\beta\Theta))^{-1} w_i$, (38) because μ is perpendicular w_i for each i, while (39) follows by reusing $z_i = D(\exp(\beta\Theta))^{-1} w_i$ relation along with the fact that $z_1 = \mathbf{1}_S$.

From (39), it follows that

$$\|\nabla_\theta \log \pi^{\mathrm{E}}_{d,\theta} z\| \leq \beta \|\varepsilon(d)\| \left\| \left[I_A - \mathbf{1}_A (\pi^{\mathrm{E}}_{d,\theta})^\top \right] \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1}}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \right\| \|D(\exp(\beta\Theta))\| \, \|z - c_1 \mathbf{1}_S\| \tag{40}$$

$$\leq \beta \alpha^d (\|I_A\| + \|\mathbf{1}_A (\pi^{\mathrm{E}}_{d,\theta})^\top\|) \left\| \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1}}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \right\| \exp(\beta \max_s \theta(s)) \|z - c_1 \mathbf{1}_S\| \tag{41}$$

$$\leq \beta \alpha^d (1 + \sqrt{A}) \left\| \frac{D\left(\pi^{\mathrm{E}}_{d,\theta}\right)^{-1}}{\mathbf{1}_A^\top E_{s,d} \exp(\beta\Theta)} \right\| \exp(\beta \max_s \theta(s)) \|z - c_1 \mathbf{1}_S\| \tag{42}$$

$$\leq \beta \alpha^d (1 + \sqrt{A}) \left\| D^{-1}(E_{s,d} \exp(\beta\Theta)) \right\| \exp(\beta \max_s \theta(s)) \|z - c_1 \mathbf{1}_S\| \tag{43}$$

$$\leq \beta \alpha^d (1 + \sqrt{A}) \frac{1}{\min_s [E_{s,d} \exp(\beta\Theta]_s} \exp(\beta \max_s \theta(s)) \|z - c_1 \mathbf{1}_S\| \tag{44}$$

$$\leq \beta \alpha^d (1 + \sqrt{A}) \frac{\exp(\beta \max_s \theta(s))}{\exp(\beta \min_s \theta(s)) \min_s |M_1|} \|z - c_1 \mathbf{1}_S\| \tag{45}$$

$$\leq \beta \alpha^d (1 + \sqrt{A}) \frac{\exp(\beta \max_s \theta(s))}{\exp(\beta \min_s \theta(s)) \exp(\beta \min_s r(s))} \|z - c_1 \mathbf{1}_S\| \tag{46}$$

$$\leq \beta \alpha^d (1 + \sqrt{A}) \exp(\beta [\max_s \theta(s) - \min_s \theta(s) - \min_s r(s)]) \|z - c_1 \mathbf{1}_S\|. \tag{47}$$

Lastly, we prove that $\|z - c_1 \mathbf{1}_S\|$ is bounded independently of d. First, denote by $c = (c_1, \ldots, c_S)^\top$ and $\tilde{c} = (0, c_2, \ldots, c_S)^\top$. Also, denote by Z the matrix with z_i as its i-th column. Now,

$$\|z - c_1 \mathbf{1}_S\| = \|\sum_{i=2}^{S} c_i z_i\| \tag{48}$$

$$= \|Z\tilde{c}\| \tag{49}$$

$$\leq \|Z\| \|\tilde{c}\| \tag{50}$$

$$\leq \|Z\| \|c\| \tag{51}$$

$$= \|Z\| \|Z^{-1} z\| \tag{52}$$

$$\leq \|Z\| \|Z^{-1}\|, \tag{53}$$

where the last relation is due to z being a unit vector. All matrix norms here are l_2-induced norms.

Next, denote by W the matrix with w_i in its i-th column. Recall that in (33) we only defined $w_2, \ldots, w_S$. We now set $w_1 = \exp(\beta\Theta)$. Note that w_1 is linearly independent of $\{w_2, \ldots, w_S\}$ because of (33) together with the fact that $\mu^\top w_1 > 0$. We can now express the relation between Z and W by $Z = D^{-1}(\exp(\beta\Theta))W$. Substituting this in (53), we have

$$\|z - c_1 \mathbf{1}_S\| \leq \|D^{-1}(\exp(\beta\Theta))W\| \|W^{-1} D(\exp(\beta\Theta))\| \tag{54}$$

$$\leq \|W\| \|W^{-1}\| \|D(\exp(\beta\Theta))\| \|D^{-1}(\exp(\beta\Theta))\|. \tag{55}$$

It further holds that

$$\|D(\exp(\beta\Theta))\| \leq \max_s \exp(\beta\theta(s)) \leq \max\{1, \exp[\beta \max_s \theta(s)])\}, \tag{56}$$

where the last relation equals 1 if $\theta(s) < 0$ for all s. Similarly,

$$\|D^{-1}(\exp(\beta\Theta))\| \leq \frac{1}{\min_s \exp(\beta\theta(s))} \leq \frac{1}{\min\{1, \exp[\beta \min_s \theta(s)])\}}. \tag{57}$$

Furthermore, by the properties of the l_2-induced norm,

$$\|W\|_2 \leq \sqrt{S}\|W\|_1 \tag{58}$$
$$= \sqrt{S} \max_{1\leq i\leq S} \|w_i\|_1 \tag{59}$$
$$= \sqrt{S} \max\{\exp(\beta\Theta), \max_{2\leq i\leq S} \|w_i\|_1\} \tag{60}$$
$$\leq \sqrt{S} \max\{1, \exp[\beta \max_s \theta(s)], \max_{2\leq i\leq S} \|w_i\|_1)\}. \tag{61}$$

Lastly,

$$\|W^{-1}\| = \frac{1}{\sigma_{\min}(W)} \tag{62}$$
$$\leq \left(\prod_{i=1}^{S-1} \frac{\sigma_{\max}(W)}{\sigma_i(W)}\right) \frac{1}{\sigma_{\min}(W)} \tag{63}$$
$$= \frac{(\sigma_{\max}(W))^{S-1}}{\prod_{i=1}^{S} \sigma_i(W)} \tag{64}$$
$$= \frac{\|W\|^{S-1}}{|\det(W)|}. \tag{65}$$

The determinant of W is a sum of products involving its entries. To upper bound (65) independently of d, we lower bound its denominator by upper and lower bounds on the entries $[W]_{i,1}$ that are independent of d, depending on their sign:

$$\min\{1, \exp[\beta \min_s \theta(s)])\} \leq [W]_{i,1} \leq \max\{1, \exp[\beta \max_s \theta(s)])\}. \tag{66}$$

Using this, together with (53), (55), (56), (57), and (61), we showed that $\|z - c_1\mathbf{1}_S\|$ is upper bounded by a constant independent of d. This concludes the proof. □

A.9 Bias Estimates

Lemma A.2. *For any matrix A and $\hat{A}$,*

$$\hat{A}^k - A^k = \sum_{h=1}^{k} \hat{A}^{h-1}(\hat{A} - A)A^{k-h}.$$

Proof. The proof follows from first principles:

$$\sum_{h=1}^{k} \hat{A}^{h-1}(\hat{A} - A)A^{k-h} = \sum_{h=1}^{k} \hat{A}^{h-1}\hat{A}A^{k-h} - \sum_{h=1}^{k} \hat{A}^{h-1}AA^{k-h} \tag{67}$$
$$= \sum_{h=1}^{k} \hat{A}^{h}A^{k-h} - \sum_{h=1}^{k} \hat{A}^{h-1}A^{k-h+1} \tag{68}$$
$$= \hat{A}^k - A^k + \sum_{h=1}^{k-1} \hat{A}^{h}A^{k-h} - \sum_{h=2}^{k} \hat{A}^{h-1}A^{k-h+1} \tag{69}$$
$$= \hat{A}^k - A^k. \tag{70}$$

□

Henceforth, $\|\cdot\|$ will refer to $\|\cdot\|_\infty$, i.e. the induced infinity norm. Also, for brevity, we denote $\pi^{\mathrm{C}}_{d,\theta}$ and $\hat{\pi}^{\mathrm{C}}_{d,\theta}$ by π_θ and $\hat{\pi}_\theta$, respectively. Similarly, we use d_{π_θ} and $d_{\hat{\pi}_\theta}$ to denote $d_{\pi^{\mathrm{C}}_{d,\theta}}$ and $d_{\hat{\pi}^{\mathrm{C}}_{d,\theta}}$. As for the induced norm of the matrix P and its perturbed counterpart $\hat{P}$, which are of size $S \times A \times S$, we slightly abuse notation and denote $\|P - \hat{P}\| = \max_s\{\|P_s - \hat{P}_s\|\}$, where P_s is as defined in Section 2.

Definition A.3. Let ϵ be the maximal model mis-specification, i.e., $\max\{\|P-\hat{P}\|, \|r-\hat{r}\|\} = \epsilon$.

Lemma A.4. *Recall the definitions of R_s, P_s, R_{π_b} and P^{π_b} from Section 2, and respectively denote their perturbed counterparts by $\hat{R}_s, \hat{P}_s, \hat{R}_{\pi_b}$ and $\hat{P}^{\pi_b}$. Then, for ϵ defined in Definition A.3,*

$$\max\{\|R_s - \hat{R}_s\|, \|P_s - \hat{P}_s\|, \|R_{\pi_b} - \hat{R}_{\pi_b}\|, \|P^{\pi_b} - \hat{P}^{\pi_b}\|\} = O(\epsilon). \tag{71}$$

Proof. The proof follows easily from the fact that the differences above are convex combinations of $P - \hat{P}$ and $r - \hat{r}$. □

Lemma A.5. *Let π_θ be as in* (5)*, and let $\hat{\pi}_\theta$ also be defined as in* (5)*, but with R_s, P_s, P^{π_b} replaced by their perturbed counterparts $\hat{R}_s, \hat{P}_s, \hat{P}^{\pi_b}$ throughout. Then,*

$$\|\pi^{\mathrm{C}}_{d,\theta} - \hat{\pi}^{\mathrm{C}}_{d,\theta}\| = O(\beta d \epsilon). \tag{72}$$

Proof. To prove the desired result, we work with (5) to bound the error between $R_s, P_s, P^{\pi_b}, R_{\pi_b}$ and their perturbed versions.

First, we apply Lemma A.2 together with Lemma A.4 to obtain that $\|(P^{\pi_b})^k - (\hat{P}^{\pi_b})^k\| = O(k\epsilon)$. Next, denote by M the argument in the exponent in (5), i.e.

$$M := \beta[C_{s,d} + P_s(P^{\pi_b})^{d-1}\Theta].$$

Similarly, let $\hat{M}$ be the corresponding perturbed sum that relies on $\hat{P}$ and $\hat{r}$. Combining the bounds from Lemma A.4, and using the triangle inequality, we have that $\|\hat{M} - M\| = O(\beta d \epsilon)$.

Eq. (5) states that the C-SoftTreeMax policy in the true environment is $\pi_\theta = \exp(M)/(1^\top \exp(M))$. Similarly define $\hat{\pi}_\theta$ using $\hat{M}$ for the approximate model. Then,

$$\hat{\pi}_\theta = (\pi_\theta \circ \exp(M - \hat{M}))1^\top \exp(M)/(1^\top \exp(\hat{M})),$$

where $\circ$ denotes element-wise multiplication. Using the above relation, we have that $\|\hat{\pi}_\theta - \pi_\theta\| = \|\pi_\theta\|\|\frac{\exp(M-\hat{M})1^\top \exp(M)}{1^\top \exp(\hat{M})} - 1\|$. Using the relation $|e^x - 1| = O(x)$ as $x \to 0$, the desired result follows.

□

Theorem A.6. *Let ϵ be as in Definition A.3. Further let $\hat{\pi}^{\mathrm{C}}_{d,\theta}$ being the corresponding approximate policy as given in Lemma 4.2. Then, the policy gradient bias is bounded by*

$$\left\| \frac{\partial}{\partial \theta}\left(\nu^\top V^{\pi_\theta}\right) - \frac{\partial}{\partial \theta}\left(\nu^\top V^{\hat{\pi}_\theta}\right) \right\| = \mathcal{O}\left(\frac{1}{(1-\gamma)^2} S \beta^2 d \epsilon\right). \tag{73}$$

We first provide a proof outline for conciseness, and only after it the complete proof.

Proof outline. First, we prove that $\max\{\|R_s - \hat{R}_s\|, \|P_s - \hat{P}_s\|, \|R_{\pi_b} - \hat{R}_{\pi_b}\|, \|P^{\pi_b} - \hat{P}^{\pi_b}\|\} = \mathcal{O}(\epsilon)$. This follows from the fact that the differences above are suitable convex combinations of either the rows of $P - \hat{P}$ or $r - \hat{r}$. We use the above observation along with the definitions of $\pi^{\mathrm{C}}_{d,\theta}$ and $\hat{\pi}^{\mathrm{C}}_{d,\theta}$ given in (5) to show that $\|\pi^{\mathrm{C}}_{d,\theta} - \hat{\pi}^{\mathrm{C}}_{d,\theta}\| = O(\beta d \epsilon)$. The proof for the latter builds upon two key facts: (a) $\|(P^{\pi_b})^k - (\hat{P}^{\pi_b})^k\| \le \sum_{h=1}^k \|\hat{P}^{\pi_b}\|^{h-1}\|\hat{P}^{\pi_b} - P^{\pi_b}\|\|p^{\pi_b}\|^{k-h} = O(k\epsilon)$ for any $k \ge 0$, and (b) $|e^x - 1| = O(x)$ as $x \to 0$. Next, we decompose the LHS of (7) to get

$$\sum_s \left(\prod_{i=1}^4 X_i(s) - \prod_{i=1}^4 \hat{X}_i(s) \right) = \sum_s \sum_{i=1}^4 \hat{X}_1(s) \cdots \hat{X}_{i-1}(s) \left(X_i(s) - \hat{X}_i(s) \right) \times X_{i+1}(s) \cdots X_4(s),$$

where $X_1(s) = d_{\pi^{\mathrm{C}}_{d,\theta}}(s) \in \mathbb{R}$, $X_2(s) = (\nabla_\theta \log \pi^{\mathrm{C}}_{d,\theta}(\cdot|s))^\top \in \mathbb{R}^{S \times A}$, $X_3(s) = D(\pi^{\mathrm{C}}_{d,\theta}(\cdot|s)) \in \mathbb{R}^{A \times A}$, $X_4(s) = Q^{\pi^{\mathrm{C}}_{d,\theta}}(s,\cdot) \in \mathbb{R}^{A \times A}$, and $\hat{X}_1(s), \ldots, \hat{X}_4(s)$ are similarly defined with $\pi^{\mathrm{C}}_{d,\theta}$ replaced by $\hat{\pi}^{\mathrm{C}}_{d,\theta}$. Then, we show that, for $i = 1, \ldots, 4$, (i) $\|X_i(s) - \hat{X}_i(s)\| = O(\epsilon)$ and (ii) $\max\{\|X_i\|, \|\hat{X}_i\|\}$ is bounded by problem parameters. From this, the desired result follows. □

Proof. We have

$$\frac{\partial}{\partial\theta}\left(\nu^\top V^{\pi_\theta}\right) - \frac{\partial}{\partial\theta}\left(\nu^\top V^{\pi'_\theta}\right) \tag{74}$$

$$= \mathbb{E}_{s\sim d_{\pi_\theta}, a\sim\pi_\theta(\cdot|s)}\left[\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s,a)\right] - \mathbb{E}_{s\sim d_{\hat{\pi}_\theta}, a\sim\hat{\pi}_\theta(\cdot|s)}\left[\nabla_\theta \log \hat{\pi}_\theta(a|s) Q^{\hat{\pi}_\theta}(s,a)\right] \tag{75}$$

$$= \sum_{s,a}\left(d_{\pi_\theta}(s)\pi_\theta(a|s)\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s,a) - d_{\hat{\pi}_\theta}(s)\hat{\pi}_\theta(a|s)\nabla_\theta \log \hat{\pi}_\theta(a|s) Q^{\hat{\pi}_\theta}(s,a)\right) \tag{76}$$

$$= \sum_{s}\Big(d_{\pi_\theta}(s)(\nabla_\theta \log \pi_\theta(\cdot|s))^\top D(\pi_\theta(\cdot|s)) Q^{\pi_\theta}(s,\cdot) \tag{77}$$

$$- d_{\hat{\pi}_\theta}(s)(\nabla_\theta \log \hat{\pi}_\theta(\cdot|s))^\top D(\hat{\pi}_\theta(\cdot|s)) Q^{\hat{\pi}_\theta}(s,\cdot)\Big) \tag{78}$$

$$= \sum_{s}\left(\prod_{i=1}^{4} X_i(s) - \prod_{i=1}^{4}\hat{X}_i(s)\right) \tag{79}$$

$$= \sum_{s}\sum_{i=1}^{4}\hat{X}_1(s)\cdots\hat{X}_{i-1}(s)\left(X_i(s) - \hat{X}_i(s)\right)X_{i+1}(s)\cdots X_4(s), \tag{80}$$

where $X_1(s) = d_{\pi_\theta}(s) \in \mathbb{R}$, $X_2(s) = (\nabla_\theta \log \pi_\theta(\cdot|s))^\top \in \mathbb{R}^{S\times A}$, $X_3(s) = D(\pi_\theta(\cdot|s)) \in \mathbb{R}^{A\times A}$, $X_4(s) = Q^{\pi_\theta}(s,\cdot) \in \mathbb{R}^{A\times A}$, and $\hat{X}_1(s), \ldots, \hat{X}_4(s)$ are similarly defined with π_θ replaced by $\hat{\pi}_\theta$.

Therefore,

$$\left\|\frac{\partial}{\partial\theta}\left(\nu^\top V^{\pi_\theta}\right) - \frac{\partial}{\partial\theta}\left(\nu^\top V^{\pi'_\theta}\right)\right\| \le \left(\max_s \Gamma(s)\right) S, \tag{81}$$

where

$$\Gamma(s) = \|\sum_{s}\sum_{i=1}^{4}\hat{X}_1(s)\cdots\hat{X}_{i-1}(s)\left(X_i(s) - \hat{X}_i(s)\right)X_{i+1}(s)\cdots X_4(s)\|. \tag{82}$$

Next, since $d_{\pi_\theta}, d_{\hat{\pi}_\theta}, \pi_\theta$, and $\hat{\pi}_\theta$ are all distributions, we have

$$\max\{|X_1(s)|, |\hat{X}_1(s)|, |X_3(s,a)|, |\hat{X}_3(s,a)|\} \le 1. \tag{83}$$

Separately, using Lemma 4.3, we have

$$\|X_2\| = \|\nabla_\theta \log \pi_\theta(a|s)\| \le \beta(\|I_A\| + \|\mathbf{1}_A \pi_\theta^\top\|)\|P_s\|\|(P^{\pi_b})^{d-1}\|. \tag{84}$$

Since all rows of the above matrices have non-negative entries that add up to 1, we get

$$\|Y\| \le 2\beta. \tag{85}$$

In the rest of the proof, we bound each of $\|X_1 - \hat{X}_1\|, \ldots, \|X_4 - \hat{X}_4\|$.

Finally,

$$\|X_4\| \le \frac{1}{1-\gamma}. \tag{86}$$

Similarly, the same bounds hold for $\hat{X}_1, \hat{X}_2, \hat{X}_3$ and $\hat{X}_4$.

From, we have

$$\|X_1 - \hat{X}_1\| \le (1-\gamma)\sum_{t=0}^{\infty}\gamma^t\|\nu^\top(P^{\pi_\theta})^t - \nu^\top(P^{\hat{\pi}_\theta})^t\| \tag{87}$$

$$\le (1-\gamma)\|\nu\|\sum_{t=0}^{\infty}\gamma^t t d\epsilon \tag{88}$$

$$\le (1-\gamma)d\epsilon\sum_{t=0}^{\infty}\gamma^t t \tag{89}$$

$$= \frac{\gamma d\epsilon}{1-\gamma}. \tag{90}$$

The last relation follows from the fact that $(1-\gamma)^{-1} = \sum_{t=0}^{\infty} \gamma^t$, which in turn implies

$$\gamma \frac{\partial}{\partial \gamma} \left(\frac{1}{1-\gamma} \right) = \sum_{t=0}^{\infty} t\gamma^t. \tag{91}$$

From Lemma A.5, it follows that

$$\|X_3 - \hat{X}_3\| = O(\beta d \epsilon). \tag{92}$$

Next, recall that from Lemma 4.3 that

$$X_2(s,\cdot) = \beta \left[I_A - \mathbf{1}_A(\pi_\theta)^\top \right] P_s \left(P^{\pi_b} \right)^{d-1}.$$

Then,

$$\|X_2(s,\cdot) - \hat{X}_2(s,\cdot)\| \leq \|\beta \left[I_A - \mathbf{1}_A(\pi_\theta)^\top \right] P_s\| \| \left(P^{\pi_b} \right)^{d-1} - \left(\hat{P}^{\pi_b} \right)^{d-1} \| \tag{93}$$

$$+ \|\beta \left[I_A - \mathbf{1}_A(\pi_\theta)^\top \right] \| \|P_s - \hat{P}_s\| \| \left(\hat{P}^{\pi_b} \right)^{d-1} \| \tag{94}$$

$$+ \beta \|\mathbf{1}_A(\pi_\theta)^\top - \mathbf{1}_A(\hat{\pi}_\theta)^\top \| \|\hat{P}_s \left(\hat{P}^{\pi_b} \right)^{d-1} \|. \tag{95}$$

Following the same argument as in (85) and applying Lemma A.2, we have that (93) is $O(\beta d \epsilon)$. Similarly, from the argument of (85), Eq. (94) is $O(\beta \epsilon)$. Lastly, (95) is $O(\beta d \epsilon)$ due to Lemma A.5. Putting the above three terms together, we have that

$$\|X_2(s,\cdot) - \hat{X}_2(s,\cdot)\| = O(\beta d \epsilon). \tag{96}$$

Since the state-action value function satisfies the Bellman equation, we have

$$Q^{\pi_\theta} = r + \gamma P Q^{\pi_\theta} \tag{97}$$

and

$$Q^{\hat{\pi}_\theta} = \hat{r} + \gamma \hat{P} Q^{\hat{\pi}_\theta}. \tag{98}$$

Consequently,

$$\|Q^{\pi_\theta} - Q^{\hat{\pi}_\theta}\| \leq \|r - \hat{r}\| + \gamma \|PQ^{\pi_\theta} - PQ^{\hat{\pi}_\theta}\| + \gamma \|PQ^{\hat{\pi}_\theta} - \hat{P}Q^{\hat{\pi}_\theta}\| \tag{99}$$

$$\leq \epsilon + \gamma \|P\| \|Q^{\pi_\theta} - Q^{\hat{\pi}_\theta}\| + \gamma \|P - \hat{P}\| \|Q^{\hat{\pi}_\theta}\| \tag{100}$$

$$\leq \epsilon + \gamma \|Q^{\pi_\theta} - Q^{\hat{\pi}_\theta}\| + \frac{\gamma}{1-\gamma} \epsilon, \tag{101}$$

which finally shows that

$$\|X_4 - \hat{X}_4\| = \|Q^{\pi_\theta} - Q^{\hat{\pi}_\theta}\| \leq \frac{\epsilon}{(1-\gamma)^2}. \tag{102}$$

□

B Experiments

B.1 Implementation Details

The environment engine is the highly efficient Atari-CuLE [Dalton et al., 2020], a CUDA-based version of Atari that runs on GPU. Similarly, we use Atari-CuLE for the GPU-based breadth-first TS as done in Dalal et al. [2021]: In every tree expansion, the state S_t is duplicated and concatenated with all possible actions. The resulting tensor is fed into the GPU forward model to generate the tensor of next states $(S_{t+1}^0, \ldots, S_{t+1}^{A-1})$. The next-state tensor is then duplicated and concatenated again with all possible actions, fed into the forward model, etc. This procedure is repeated until the final depth is reached, for which $W_\theta(s)$ is applied per state.

We train SoftTreeMax for depths $d = 1 \ldots 8$, with a single worker. We use five seeds for each experiment.

For the implementation, we extend Stable-Baselines3 [Raffin et al., 2019] with all parameters taken as default from the original PPO paper [Schulman et al., 2017]. For depths $d \geq 3$, we limited the tree to a maximum width of 1024 nodes and pruned non-promising trajectories in terms of estimated weights. Since the distributed PPO baseline advances significantly faster in terms of environment steps, for a fair comparison, we ran all experiments for one week on the same machine and use the wall-clock time as the x-axis. We use Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz equipped with one NVIDIA Tesla V100 32GB.

B.2 Time-Based Training Curves

We provide the training curves in Figure 4. For brevity, we exclude a few of the depths from the plots. As seen, there is a clear benefit for SoftTreeMax over distributed PPO with the standard softmax policy. In most games, PPO with the SoftTreeMax policy shows very high sample efficiency: it achieves higher episodic reward although it observes much less episodes, for the same running time.

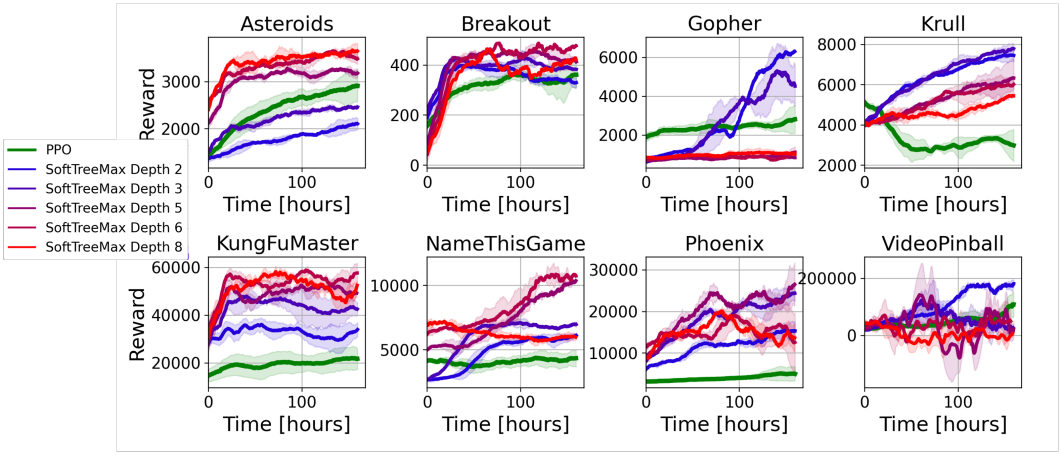

Figure 4: **Training curves: GPU SoftTreeMax (single worker) vs PPO (256 GPU workers).** The plots show average reward and standard deviation over 5 seeds. The x-axis is the wall-clock time. The runs ended after one week with varying number of time-steps. The training curves correspond to the evaluation runs in Figure 3.

B.3 Step-Based Training Curves

In Figure 5 we also provide the same convergence plots where the x-axis is now the number of online interactions with the environment, thus excluding the tree expansion complexity. As seen, due to the complexity of the tree expansion, less steps are conducted during training (limited to one week) as the depth increases. In this plot, the monotone improvement of the reward with increasing tree depth is noticeable in most games.

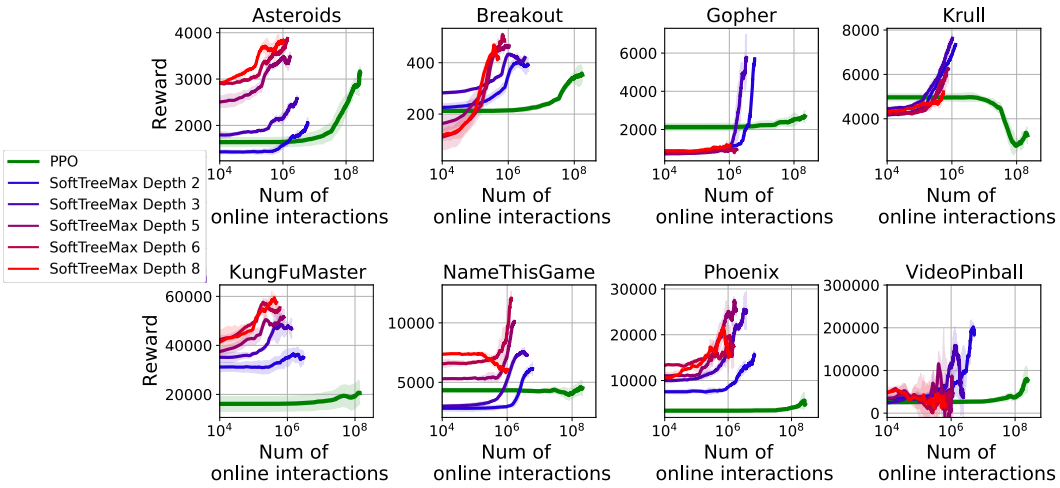

Figure 5: **Training curves: GPU SoftTreeMax (single worker) vs PPO (256 GPU workers).** The plots show average reward and standard deviation over 5 seeds. The x-axis is the number of online interactions with the environment. The runs ended after one week with varying number of time-steps. The training curves correspond to the evaluation runs in Figure 3.

We note that not for all games we see monotonicity. Our explanation for this phenomenon relates to how immediate reward contributes to performance compared to the value. Different games benefit differently from long-term as opposed to short-term planning. Games that require longer-term planning need a better value estimate. A good value estimate takes longer to obtain with larger depths, in which we apply the network to states that are very different from the ones observed so far in the buffer (recall that as in any deep RL algorithm, we train the model only on states in the buffer). If the model hasn't learned a good enough value function yet, and there is no guiding dense reward along the trajectory, the policy becomes noisier, and can take more steps to converge – even more than those we run in our week-long experiment.

For a concrete example, let us compare Breakout to Gopher. Inspecting Fig. 5, we observe that Breakout quickly (and monotonically) gains from large depths since it relies on the short term goal of simply keeping the paddle below the moving ball. In Gopher, however, for large depths (>=5), learning barely started even by the end of the training run. Presumably, this is because the task in Gopher involves multiple considerations and steps: the agent needs to move to the right spot and then hit the mallet the right amount of times, while balancing different locations. This task requires long-term planning and thus depends more strongly on the accuracy of the value function estimate. In that case, for depth 5 or more, we would require more train steps for the value to "kick in" and become beneficial beyond the gain from the reward in the tree.

The figures above convey two key observations that occur for at least some non-zero depth: (1) The final performance with the tree is better than PPO (Fig. 3); and (2) the intermediate step-based results with the tree are better than PPO (Fig. 5). This leads to our main takeaway from this work — there is no reason to believe that the vanilla policy gradient algorithm should be better than a multi-step variant. Indeed, we show that this is not the case.

C Further discussion

C.1 The case of $\lambda_2(P^{\pi_b}) = 0$

When P^{π_b} is rank one, it is not only its variance that becomes 0, but also the norm of the gradient itself (similarly to the case of $d \to \infty$). Note that such a situation will happen rarely, in degenerate MDPs. This is a local minimum for SoftTreeMax and it would cause the PG iteration to get stuck, and to the optimum in the (desired but impractical) case where π_b is the optimal policy. However, a similar phenomenon was also discovered in the standard softmax with deterministic policies:

$\theta(s,a) \to \infty$ for one a per s. PG with softmax would suffer very slow convergence near these local equilibria, as observed in Mei et al. [2020a]. To see this, note that the softmax gradient is $\nabla_\theta \log \pi_\theta(a|s) = e_a - \pi_\theta(\cdot|s)$, where $e_a \in [0,1]^A$ is the vector with 0 everywhere except for the a-th coordinate. I.e., it will be zero for a deterministic policy. SoftTreeMax avoids these local optima by integrating the reward into the policy itself (but may get stuck in another, as discussed above).

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: [NA]

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We included a relevant section at the end of the paper.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: All proofs can be found in the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Yes. We also attached the repository, together with a docker environment, as supplementary material.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We attached the repository, together with a docker environment, as supplementary material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: All relevant information can be found in the paper and the appendix.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: [NA]

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: The paper has no scietal impact.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: [NA]

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: The paper does not release new assets.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.