
SPO: Sequential Monte Carlo Policy Optimisation

Matthew V Macfarlane *
University of Amsterdam
m.v.macfarlane@uva.nl

Edan Toledo
InstaDeep

Donal Byrne
InstaDeep

Paul Duckworth
InstaDeep

Alexandre Larterre
InstaDeep

Abstract

Leveraging planning during learning and decision-making is central to the long-term development of intelligent agents. Recent works have successfully combined tree-based search methods and self-play learning mechanisms to this end. However, these methods typically face scaling challenges due to the sequential nature of their search. While practical engineering solutions can partly overcome this, they often result in a negative impact on performance. In this paper, we introduce SPO: Sequential Monte Carlo Policy Optimisation, a model-based reinforcement learning algorithm grounded within the Expectation Maximisation (EM) framework. We show that SPO provides robust policy improvement and efficient scaling properties. The sample-based search makes it directly applicable to both discrete and continuous action spaces without modifications. We demonstrate statistically significant improvements in performance relative to model-free and model-based baselines across both continuous and discrete environments. Furthermore, the parallel nature of SPO’s search enables effective utilisation of hardware accelerators, yielding favourable scaling laws.

1 Introduction

The integration of reinforcement learning (RL) and neural-guided planning methods has recently achieved considerable success. Such methods effectively leverage planning during training via iterative imitation learning [7, 68, 18, 8]. Applying additional computation via planning generates improved policies which is then amortised [4] into a neural network policy. Using this policy as part of planning itself creates a powerful self-improvement cycle. They have demonstrated state of the art performance in applications ranging from chess [75] and matrix multiplication [24], to language modelling [61]. However, one of the most commonly-used search-based policy improvement operators, MCTS [16]: i) performs poorly on low budgets [33], ii) is inherently sequential limiting its scalability [51, 73], and iii) requires modifications to adapt to large or continuous action spaces [57, 41]. These limitations underscore the need for more scalable, efficient and generally applicable algorithms.

In this work, we introduce SPO: Sequential Monte Carlo Policy Optimisation, a model-based RL algorithm that utilises scalable sampled-based Sequential Monte Carlo planning for policy improvement. We formalise SPO as an approximate policy iteration algorithm, and show that it is formally grounded within the Expectation Maximisation (EM) framework.

*Work done during internship at InstaDeep

SPO is unique by leveraging both breadth and depth of search in order to provide better estimates of the target distribution, derived via EM optimisation, compared to previous works. When viewed relative to popular algorithms it combines the breadth used in MPO [1] and the depth used in V-MPO [76], while enforcing KL constraints on updates to ensure stable policy improvement. SPO achieves strong performance across a range of benchmarks, outperforming AlphaZero, a leading model-based expert iteration algorithm. We also benchmark against a leading Sequential Monte Carlo (SMC) algorithm from Piché et al. [65] and show that leveraging posterior estimates from SMC for policy improvement is central to strong performance.

We demonstrate that i) SPO outperforms baselines across both high-dimensional continuous control and challenging discrete planning tasks without algorithmic modifications, and ii) the sampling-based approach is inherently parallelisable, enabling the use of hardware accelerators to improve training speed. This provides significant advantages over MCTS-based methods [16], whose sequential process can be particularly inefficient during training. 3) The explicit use of KL targeting leads to strong and stable policy improvement, reducing the need for a grid search over exploration hyperparameters in search. We find that SPO, empirically has strong scaling behaviours, with improved performance given additional search budget both during training and inference.

2 Related Work

There is an extensive history of prior works that frame control and RL as an inference problem [82, 45]. Expectation Maximization [19, 28, 60] is a popular approach to directly optimising the evidence lower bound (ELBO) of the probability of optimality of a given policy. A number of model free methods implement EM, including RWR [63], REPS [64], MPO [1], V-MPO [76], and AWAC [59] each with different approaches to the E-step and M-step (summarised in Furuta et al. [29]). Model based algorithms such as MD-GPS [58] and AlphaZero [74] can also be viewed under the EM framework using LQR [5] and MCTS [9, 16] during the E-step respectively. Grill et al. [33] show MCTS is a form of regularised policy optimisation equivalent to optimising the ELBO. Such model-based algorithms, that iterate between policy improvement using search, and projecting this improvement to the space of parameterisable policies, are also often referred to as Expert iteration (ExIt) [74, 7] or Dual Policy Iteration [77]. Such approaches have also been applied to active inference [27] where MCTS has been used to estimate expected free energy, with this computation then amortised into a neural policy [25]. Our work can also be framed as ExIt, but differs by using Neural SMC for planning with an explicit KL regularisation on the improvement.

Sequential Monte Carlo (SMC) [31, 43, 52], also referred to as particle filters [66, 6], is an inference method often employed to sample from intractable distributions. Gu et al. [34] explore the parameterisation of the proposal in SMC using LSTMs [40], showing that posterior estimates from SMC can be used to train a parameterised proposal in non-control based tasks. Li et al. [49] demonstrate the same concept but using an MCMC sampler [30]. SMC can be applied to sample from distributions over trajectories in RL. For example, SMC-Learning [44] updates SMC weights using a Boltzmann exploration strategy, but unlike SPO it does not control policy improvement with KL regularisation or leverage neural approximators. Piché et al. [65] derive an SMC weight update to directly estimate the posterior over optimal trajectories. However this update requires the optimal value function, which they approximate using a SAC [37] value function. This approach requires high sample counts and does not leverage SMC posterior estimates for policy improvement, instead utilising a Gaussian distribution or SAC to train the proposal. CriticSMC [50] estimates the same posterior but utilises a soft action-value function to score particles [36], reducing the number of steps performed in the environment, enabling more efficient exploration and improved performance on lower SMC budgets compared with Piché et al. [65]. However their method is slower in wall clock time and uses a static proposal, therefore not performing iterative policy improvement. A useful property of SMC search methods is that it is inherently parallelisable. Parallelising MCTS with a virtual loss has been explored, however this often leads to performance degradation [51], increasing exploration and leading to out-of-distribution states that are difficult to evaluate Dalal et al. [17].

3 Background

Sequential decision-making can be formalised under the Markov Decision Process (MDP) framework [67]. An MDP is a tuple $(S; A; T; r; \gamma; \mu)$ where, S is the set of possible states, A is the set of actions, $T : S \times A \rightarrow P(S)$ is the state transition probability function, $r : S \times A \rightarrow \mathbb{R}$ is the reward function, $\gamma \in [0; 1]$ is the discount factor, and μ is the initial state distribution. An agent interacts with the MDP using a policy $\pi : S \rightarrow P(A)$, that associates to every state a distribution over actions. We quantify the quality of a policy by the expected discounted return, that the agent seeks to maximise $V^\pi = \mathbb{E} [\sum_{t=0}^{\infty} \gamma^t r_t]$; where $r_t = r(s_t; a_t)$ is the reward received at time t and π is the set of all realisable policies. The value function $V^\pi(s_t) = \mathbb{E} [\sum_{t=t}^{\infty} \gamma^t r_t | s_t]$ maps a state s_t to the expected discounted sum of future rewards when acting according to π . Similarly, the state-action value function $Q^\pi(s_t; a_t) = \mathbb{E} [\sum_{t=t}^{\infty} \gamma^t r_t | s_t; a_t]$ maps a state s_t to the expected discounted return, when taking the initial action a_t and following π thereafter.

Control as inference formulates the RL objective as an inference problem within a probabilistic graphical model [82, 42, 45]. For a horizon T , the distribution over trajectories $\mathcal{T} = (s_0; a_0; s_1; a_1; \dots; s_T; a_T)$ is given by $p(\mathcal{T}) = \mu(s_0) \prod_{t=0}^{T-1} p(a_t) T(s_{t+1} | s_t; a_t)$, which is a function of the initial state distribution, the transition dynamics, and an action prior. Note that this distribution is insufficient for solving control problems, because it has no notion of rewards. We therefore have to introduce an additional *optimality variable* into this model, which we will denote O_t , that is defined such that $p(O_t = 1j) \propto \exp(r_t)$. Therefore, a high reward at time t means a high probability of having taken the optimal action at that point. We are then concerned with the target distribution $p(O_{1:T})$, which is the distribution of trajectories given optimality at every step. We denote $O_{1:T}$ as O going forward. The RL objective is then formulated as finding a policy to maximise $\log p(O = 1) = \log \int p(O = 1j) d\mathcal{T}$, which intuitively can be thought of as maximising the distribution of optimality at all timesteps, given actions sampled according to π . To optimise this challenging objective, we can derive the evidence lower bound (ELBO) using an auxiliary distribution q [60]:

$$\begin{aligned} \log p(O = 1) &= \log \int p(O = 1j) d\mathcal{T} = \log \int q(\mathcal{T}) \frac{p(O = 1j)}{q(\mathcal{T})} d\mathcal{T} \\ &= \mathbb{E}_q \left[\sum_t \frac{r_t}{\gamma^t} \right] - \text{KL}(q(\mathcal{T}) \parallel p(\mathcal{T})). \end{aligned} \quad (1)$$

Since $p(O = 1j) \propto \exp(\sum_t r_t)$ and Z is a normalising constant, ensuring a valid probability distribution.

Expectation Maximisation (EM) [20] has been widely applied to solve the optimisation problem $\max_{\pi} \log p(O = 1)$ [19, 63, 64, 1, 76]. After deriving the lower bound $\mathcal{J}(q; \pi)$ on the objective in eq. (1) using an auxiliary non-parametric distribution q , EM then performs a coordinate ascent, iterating between optimizing the bound with respect to q (E-step) and with respect to the parametric policy π (M-step). This generates a sequence of policy pairs $(\pi_0; q_0); (\pi_1; q_1); \dots; (\pi_n; q_n)$ such that in each step i in the EM sequence, $\mathcal{J}(q_{i+1}; \pi_{i+1}) \geq \mathcal{J}(q_i; \pi_i)$. Viewed through a traditional policy improvement lens, the E-step corresponds to a policy evaluation phase where we perform rollouts, generating states with their associated estimates for q and value estimates. The M-step corresponds to a policy improvement phase where π and V are updated. This step can also be thought of as amortising the probabilistic inference computation performed in the E-step, into a neural network forward pass operation.²

Sequential Monte Carlo (SMC) methods [31] are designed to sample from an intractable distribution p referred to as the *target distribution* by using a tractable proposal distribution q (we use q to prevent overloading of q). *Importance Sampling* does this by sampling from q and weighting samples by $w(x) = \frac{p(x)}{q(x)}$. The estimate is performed using a set of N particles $\{x^{(n)}\}_{n=1}^N$, where $x^{(n)}$ represents a sample from the support that the target and proposal distributions are defined over, along with the associated importance weights $w^{(n)}$. Each particle uses a sample from the proposal distribution to improve the estimation of the target, increasing N naturally improves the estimation of p [10]. Once importance weights are calculated, the target is estimated as $\sum_{n=1}^N w^{(n)} x^{(n)}$,

²Many related RL algorithms can be understood by the different approaches to either the E-step or M-step, see appendix F.1 for a summary of EM approaches.

where w is the normalised importance sample weight and $\delta_{x^{(n)}}$ is the dirac measure located at $x^{(n)}$. *Sequential Importance Sampling* generalises this for sequential problems, where $x = (x_1; \dots; x_T)$. In this case importance weights can be calculated iteratively according to:

$$w_t(x_{1:t}) = w_{t-1}(x_{1:t-1}) \frac{p(x_t/x_{1:t-1})}{(x_t/x_{1:t-1})}. \quad (2)$$

Particle filtering methods can be affected by *weight degeneracy* [56]. This occurs when a few particles dominate the normalised importance weights, rendering the remaining particles negligible. As the variance of particle weights is guaranteed to increase over sequential updates [21], this phenomenon is unavoidable. When the majority of particles contribute little to the estimation of the target distribution, computational resources are wasted. *Sequential Importance Resampling* (SIR) [53] mitigates this problem by periodically resampling particles according to their current weights, subsequently resetting these weights.

4 SPO Method

In this section, we present a novel method that combines Sequential Monte Carlo (SMC) sampling with the Expectation Maximisation (EM) framework for policy iteration. We begin by formulating the objective function that we aim to optimise. We then outline our iterative approach to maximising this function, alternating between an expectation step (E-step) and a maximisation step (M-step). Within the E-step, we derive the analytical solution for optimising the objective with respect to the auxiliary distribution q . We then demonstrate how SMC can be employed to effectively estimate this *target distribution*. The M-step can be viewed as a projection of the non-parametric policy obtained in the E-step back onto the space of feasible policies. A comprehensive algorithmic outline of our proposed approach is provided in appendix D.2.

4.1 Objective

We add an additional assumption to the lower bound objective defined in eq. (1) and assume that the auxiliary distribution q , defined over trajectories τ , can be decomposed into individual state dependent distributions, i.e. $q(\tau) = p(s_0) \prod_{t \geq 0} q(a_t/s_t) T(s_{t+1}/s_t, a_t)$. We parameterise q using θ which decomposes in the same way. This enables the objective to be written with respect to individual states instead of full trajectories. Multiplying by $p(\theta)$, the objective can then be written as follows:

$$\mathcal{J}(q; \theta) = \mathbb{E}_q \left[\sum_{t=0}^{\infty} \gamma^t [r_t - \text{KL}(q(a_t/s_t) \| k(a_t/s_t; \theta))] \right] + \log p(\theta). \quad (3)$$

Previous works have explored this formulation where $p(\theta)$ is a prior over the parameters of [38, 71, 1].

4.2 E-step

Within the expectation step of EM, we maximise eq. (3) with respect to q . As in previous work, instead of optimising for the rewards, we consider an objective written according to Q-values [1, 59, 54].

$$\max_q \mathbb{E}_{(s)} \left[\mathbb{E}_{q(a|s)} [Q^q(s; a) - \text{KL}(q(a|s) \| k(a|s; \theta))] \right] \quad (4)$$

This aims to maximise the expected Q-value of a policy q , with the constraint that it doesn't move too far away from k . Framing optimisation with respect to Q-values is useful as it enables the use of function approximators to estimate Q .

Maximising this objective is difficult due to the dependence of both the expectation terms and Q-values on q . Following previous works, we fix the Q-values with respect to a fixed policy π , resulting in a partial E-step optimisation [74, 1, 76]. We use the most recent estimate of q as the fixed policy we perform maximisation with respect to, similar to AlphaZero's approach of acting according to the expert policy. The initial state distribution $p(s)$ is likewise fixed to be distributed according to π . In practice, we use a FIFO replay buffer and sample from it to determine (s, a) .

Equation (4) consists of balancing two objectives. To optimize this, we frame this as a constrained maximization problem, to avoid scaling problems, see appendix G.3 for further information. This objective limits the KL between q and π from exceeding a certain threshold:

$$\begin{aligned} \max_q \mathbb{E}_{(s)} [E_{q(a|s)} [Q(s; a)]] \\ \text{s.t. } \mathbb{E}_{(s)} [\text{KL}(q(a|s) \parallel \pi(a|s))] < \epsilon \end{aligned} \quad (5)$$

The following analytic solution to the constrained optimisation problem can then be derived using the Lagrangian multipliers method outlined in appendix G.1. Likewise q^* is obtained by minimising the convex dual function eq. (7), and intuitively can be thought of as enforcing the KL constraint on q , preventing q_i from moving too far from π_i :

$$q_i(a|s) \propto \pi_i(a|s) \exp\left(\frac{Q(s; a) - V(s)}{\eta^*}\right); \quad (6)$$

where $V(s)$ is an action independent baseline. $Q(s; a) - V(s)$ is referred to as the advantage, where optimising for Q-values vs advantages is equivalent [59]. Optimising with advantages however, has demonstrated to practically outperform optimising Q-values directly [62, 76, 59]. This is also a natural update to perform as the *policy improvement theorem* [78] outlines that an update by policy iteration can be performed if at least one state-action pair has positive advantage and a non-zero probability of reaching such a state. Although we have a closed form solution for q , we do not have either the value function or action-value function in practice. In the next section we outline our approach to estimating this distribution.

Regularised Policy Optimisation: In our closed form solution to the optimisation eq. (6), we are required to solve for η^* which ensures that the KL is constrained. This can be calculated by minimising the following objective [84], see appendix G.2 for derivation:

$$g(\eta) = \mathbb{E}_{(s)} \left[\int \pi_i(a|s, \theta_i) \log\left(\int \pi_i(a|s, \theta_i) \exp\left(\frac{A(s; a)}{\eta}\right) da\right) ds \right] \quad (7)$$

Practically, we estimate eq. (7) using a sample-based estimator according to the distribution of states from the replay buffer and stored values for $\pi_i(a|s, \theta_i)$ and $A(s; a)$.

4.2.1 Estimating Target Distribution

Building upon previous work using Sequential Monte Carlo for trajectory distribution estimation [44, 65], we propose a novel approach that integrates SMC within the Expectation Maximisation

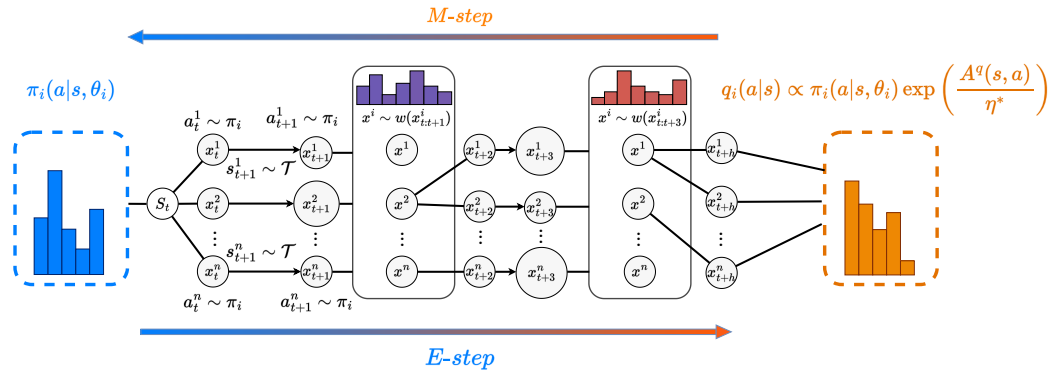


Figure 1: SPO search: n rollouts, represented by particles $x^1; \dots; x^n$, each of which represents an SMC trajectory sample, are performed in parallel according to π_i (left to right). At each environment step, the weights of the particles are adjusted (indicated in the diagram by circle size). We show two resampling regions where particles are resampled, favouring those with higher weights, and their weights are reset. The target distribution is estimated from the initial actions of the surviving particles (rightmost particles). This target estimate, q_i , is then used to update π_i in the M-step.

framework to estimate the target distribution defined in Equation 6. Our SMC-based method enables sampling from q_i over trajectories of length h , incorporating multiple predicted future states and rewards for a more accurate distribution estimation. A key property of estimating the target using SMC is that it uses both breadth (we initialise multiple particles to calculate importance weights) but also depth (leveraging future states and rewards) to estimate the distribution. In contrast, MPO evaluates several actions per state, focusing on breadth without depth, while V-MPO calculates advantages using n -step returns from trajectory sequences, leveraging depth but only for a single action sample. Our method combines both aspects, enhancing the accuracy of target distribution estimation, crucial for both action selection and effective policy improvement updates in the M -step. An outline of the SMC algorithm is provided in algorithm 1, with a visual representation in fig. 1.

SMC estimates the target by maintaining a set of particles $\bar{f}X_{t,i}^{(n)} g_{n=1}^N$ each of which maintains an importance weight for a particular sample. Given calculated importance weights SMC, estimates the target distribution according to $\hat{q}_i(\cdot) = \sum_{n=1}^N w^{(n)} X_{t,i}^{(n)}(\cdot)$ where $w^{(n)}$ is the normalised importance sample weight for a particular sample. We next outline the sequential importance sampling weight update needed to estimate eq. (6). Since we can sample from $(a_t; s_{t+1} | i)$, we leverage this as our proposal distribution. Given the target distribution q can be decomposed into individual state dependent distributions, we can define $p_i(s_{t+1} | s_t; a_t)$ and $(a_t | s_t; i)$ as:

$$p_i(s_{t+1} | s_t; a_t) \propto T(s_{t+1} | s_t; a_t) \exp\left(\frac{A_i(a_t; s_t)}{i}\right) \quad (8)$$

$$(a_t | s_t; i) \propto T_{\text{model}}(s_{t+1} | s_t; a_t) \quad (9)$$

leading to the following convenient SMC weight update according to eq. (2):

$$w(s_{1:t} | a_{1:t-1}) \propto w(s_{1:t-1} | a_{1:t-2}) \left(\frac{T(s_t | s_{t-1}; a_{t-1})}{T_{\text{model}}(s_t | s_{t-1}; a_{t-1})} \right) \frac{\exp(A(a_t; s_t) = *) (a_t | s_t; i)}{(a_t | s_t; i)}, \quad (10)$$

where $Q(a_t; s_t) = V(s_t)$ is simplified to $A(a_t; s_t)$, $s_{1:t}$ is a sequence of state, action pairs from timestep 1 to t , and T_{model} is the environment transition function of the planning model. Note that our work assumes the availability of a model that accurately represents the transition dynamics of the environment T , and therefore simplify the update to $w(s_{1:t} | a_{1:t-1}) \propto w(s_{1:t-1} | a_{1:t-2}) \exp(A(a_t; s_t) = *)$.

Algorithm 1 outlines the method for estimating eq. (6) over h planning steps for N particles (line 3) using the advantage based weight update (line 7) in eq. (10). Once h steps in the environment model have been performed in parallel, we marginalise all but the first actions as a sample based estimate of q_i (line 14).

The advantage function A is typically unknown, so we compute a 1-step estimate at each iteration using the value function and observed environment rewards $\hat{A}(s_t; a) = r_t + V(s_{t+1}) - V(s_t)$. Practically, we parameterise the value function V using a neural network and train it using GAE [70] on true environment rollouts collected. After h steps, the importance weights leverage h -steps of observed states and rewards during planning and corresponding advantage estimates. Compared to tree-based methods such as MCTS, SMC does not require maintaining the full tree in memory. Instead, after each planning step, it only needs to retain the initial action, current state, and current particle weight. It also does not require spending computation sending backward messages to update statistics of previous nodes everytime a new node is added to the tree.

Algorithm 1 SMC q target estimation (timestep t)

- 1: Initialize $\bar{f}S_t^{(n)} = s_t g_{n=1}^N$
 - 2: Set $\bar{f}W_t^{(n)} = 1 g_{n=1}^N$
 - 3: **for** $i \in \{1, \dots, t+h\}$ **do**
 - 4: $\bar{f}a_i^{(n)} = (j s_i^{(n)}) g_{n=1}^N$
 - 5: $\bar{f}S_{i+1}^{(n)} = T_{\text{model}}(s_i^{(n)}; a_i^{(n)}) g_{n=1}^N$
 - 6: $\bar{f}r_i^{(n)} = r_{\text{model}}(s_i^{(n)}; a_i^{(n)}) g_{n=1}^N$
 - 7: $\bar{f}W_i^{(n)} = w_{i-1}^{(n)} \exp(\hat{A}(s_i^{(n)}; r_i^{(n)}; s_{i+1}^{(n)}) = *) g_{n=1}^N$
 - 8: **if** $(i - t) \bmod p = 0$ **then**
 - 9: $\bar{f}X_{t,i}^{(n)} g_{n=1}^N = \text{Mult}(n; w_i^{(1)}; \dots; w_i^{(N)})$
 - 10: $\bar{f}W_i^{(n)} = 1 g_{n=1}^N$
 - 11: **end if**
 - 12: **end for**
 - 13: $\bar{f}a_t^{(n)} g$ as the set of first actions of $\bar{f}X_{t,t+h}^{(n)} g_{n=1}^N$
 - 14: $\hat{q}(a | s_t) = \sum_{n=1}^N w^{(n)} a_t^{(n)}(a)$
-

Resampling Adaptive Search: To mitigate the issue of *weight degeneracy* [56], SPO conducts periodic resampling (lines 8-11). This involves generating a new set of particles from the existing set

by duplicating some trajectories and removing others, based on the current particle weights. This process enhances computational efficiency in estimating the target distribution. By resampling, we avoid updating weights for trajectories with low likelihood under the target, thereby reallocating computational resources to particles with high importance weights [22].

4.3 M-step

After completing the E-step (which generates q_i), we proceed with the M-step, which optimises eq. (3) with respect to θ , parametrised by ϕ . By eliminating terms that are independent of θ , we optimise the following objective, corresponding to a maximum a posteriori estimation with respect to the distribution q_i :

$$\max_{\theta} \mathcal{J}(q_i; \theta) = \max_{\theta} \mathbb{E}_{q_i(s)} \left[\mathbb{E}_{q_i(\cdot|s)} [\log \pi(\theta; a|s)] \right] + \log p(\theta) \quad (11)$$

This optimisation can be viewed as projecting the non-parametric policy q_i back to the space of parametrisable policies $\pi(\theta; \cdot|s)$, as performed in expert iteration style methods such as AlphaZero. $p(\theta)$ represents a prior over the parameter θ . Previous works find that utilising a prior for θ to be close to the estimate from the previous iteration θ_{i-1} leads to stable training [1, 76]. Therefore we assume a gaussian prior over the current policy parameters, see appendix G.5 where we show utilising such a prior leads to the following constrained objective:

$$\begin{aligned} \max_{\theta} \mathbb{E}_{q_i(s)} \left[\mathbb{E}_{q_i(a|s)} [\log \pi(\theta; a|s)] \right] \\ \text{s.t. } \mathbb{E}_{q_i(s)} [\text{KL}(\pi(\theta; \cdot|s); \pi(\theta_{i-1}; \cdot|s))] < \epsilon \end{aligned} \quad (12)$$

4.4 Policy Improvement

Previous approaches to using SMC for RL either do not perform iterative policy improvement using SMC [65], or lack policy improvement constraints needed for iterative improvement [44]. Assuming that SMC provides perfect estimates of the analytic *target distribution* at each step i , Expectation Maximisation algorithm will guarantee that successive iterations will result in monotonic improvements in our lower bound objective, with details outlined in appendix G.4 [54].

Proposition 1. *Given a non-parametric variational distribution q_i and a parametric policy π_{i-1} . Given q_{i+1} , the analytical solution to E-step optimisation eq. (3), and θ_{i+1} , the solution to maximisation problem in the M-step eq. (12) then the ELBO \mathcal{J} is guaranteed to be monotonically increasing: $\mathcal{J}(q_{i+1}; \theta_{i+1}) \geq \mathcal{J}(q_i; \theta_i)$.*

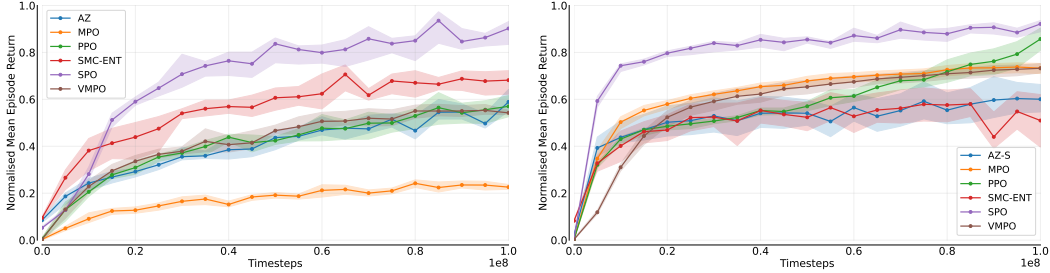
In practice we are unlikely to generate perfect estimates of the target through sample based inference and leave derivations regarding the impact of this estimation on overall convergence for future work. We also draw the connection between our EM optimisation method and Mirror Descent Guided Policy Search (MD-GPS) [58]. Our objective can be viewed as a specific instance of MD-GPS (see appendix G.6). Depending on whether dynamics are linear or not, optimising the EM objective can be viewed either as exact or approximate mirror descent [11]. Monotonic improvement guarantees in MD-GPS follow from those of mirror descent.

5 Experiments

In this section, we focus on three main areas of analysis. First, we demonstrate the improved performance of SPO in terms of episode returns, relative to both model-free and model-based algorithms. We conduct evaluations across a suite of common environments for both continuous control and discrete action spaces. Secondly, we examine the scaling behaviour of SPO during training, showing that asymptotic performance scales with particle count and depth. Finally, we explore the performance-to-speed trade-off at test time by comparing SPO directly to AlphaZero as the search budget increases.

5.1 Experimental Setup

In order to ensure the robustness of our conclusions, we follow the evaluation methodology proposed by Agarwal et al. [3], see appendix H for further details. This evaluation methodology groups performance across tasks within an environment suite, enabling clearer conclusions over the significance of



(a) **Discrete:** Rubik’s Cube 7 and Boxoban Hard. (b) **Continuous:** Ant, HalfCheetah, Humanoid.

Figure 2: Learning curves for discrete and continuous environments. The Y-axis represents the interquartile mean of min-max normalised scores, with shaded regions indicating 95% confidence intervals, across 5 random seeds.

learning algorithms as a whole. We include individual results in appendix C, along with additional analysis measuring the statistical significance of our results.

Environments: For continuous control we evaluate on the Brax [26] benchmark environments of: Ant, HalfCheetah, and Humanoid. For discrete environments, we evaluate on Boxoban [35] (a specific instance of Sokoban), commonly used to assess planning methods, and Rubik’s Cube, a sparse reward environment with a large combinatorial state-action space. See appendix A for further details regarding environments.

Baselines: Our model-free baselines include PPO [72], MPO [1] and V-MPO [76] for both continuous and discrete environments. For model-based algorithms, we compare performance to AlphaZero (including search improvements from MuZero [68]) and an SMC method introduced by Piché et al. [65]³ (which we refer to as SMC-ENT due to its use of maximum entropy RL to train the proposal and value function used within SMC). Our AlphaZero benchmark follows the official open-source implementation⁴. For continuous environments we baseline our results to Sampled MuZero [41], a modern extension to MuZero for large and/or continuous action spaces. We utilise a true environment model, aligning our implementation of Sampled MuZero more closely with AlphaZero. Our core experiments configure SPO with 16 particles and a horizon of 4, and AlphaZero with 64 simulations to equalise search budgets. For remaining parameters, see appendix E.1 and appendix D.3. Each environment and algorithm were evaluated with five random seeds.⁵

5.2 Results

Policy Improvement: In fig. 2 we show empirical evidence that SPO outperforms all baseline methods across both discrete (left) and continuous (right) benchmarks for policy improvement. This conclusion also holds for the per-environment results reported in appendix C. SMC-ENT [65] is a strong SMC based algorithm, however, SPO outperforms it for both discrete and continuous environments providing strong evidence of the direct benefit of using SMC posterior estimates for policy improvement. We also highlight the importance regularising policy optimisation, which is practically achieved by solving for θ^* eq. (7), and is re-estimated at each iteration. In appendix B we ablate the impact of this adaptive method by comparing varying fixed temperature schedules. While a good temperature can be found through expensive grid search for each new environment we find the adaptive temperature ensures stable updates leading to strong learning performance across all environments, with minimal overhead to compute θ^* .

We find that SPO performance is statistically significant compared to AlphaZero across the evaluated environments, see appendix C for detailed results. The substantial variation in AlphaZero performance across these environments underscores the difficulty in tuning it for diverse settings. For instance, while AlphaZero performs well on Boxoban and HalfCheetah, its performance drops significantly on other discrete and continuous problems. This inconsistency poses a major challenge

³We use MDQN [83], and SAC [37] for discrete and continuous environments, respectively. Note that in the discrete case, SMC-ENT using SAC was very unstable, so we chose MDQN as a better alternative.

⁴Open-source implementation available at mctx.

⁵Inference code and checkpoints are available at <https://github.com/instantdeepai/spo>

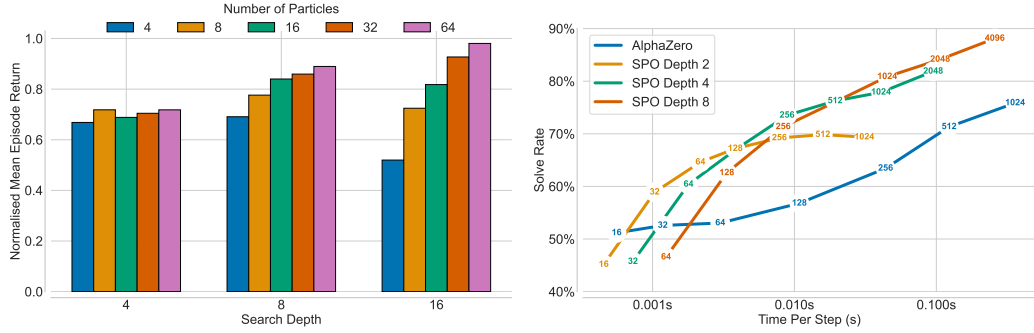


Figure 3: (left) Scaling: Mean normalised performance across all continuous environments on 10^8 environment steps, varying particle numbers N and horizon h for SPO during training. (right) Wall Clock Time Comparison: Performance on Rubik’s cube plotted against wall-clock time for AlphaZero and 3 versions of SPO (varying by SMC search depth), with total search budget labeled at each point.

for its applicability to real-world problems. In contrast, SPO performs consistently well across all environments (both discrete and continuous), highlighting its general applicability and robustness to various problem settings.

5.3 Scaling SPO during training

In fig. 3 (left) we investigate how SPO performance is impacted by scaling particle counts N and horizon length h during training, both of which we find improve the estimation of the *target distribution*. Our results show that scaling both the particle count and the horizon leads to improvements in the asymptotic performance of SPO. It also suggests that both variables should be scaled together for maximum benefit as having a long horizon with low particle counts can actually negatively impact performance. Secondly, we highlight that while for our main results we use a total search budget (particles \times depth) of 64 for policy improvement, our scaling results show that competitive results can be achieved by reducing this budget by a factor of four, which demonstrates competitive performance at low compute budgets.

5.4 Scaling SPO at test time

We can scale particle counts and horizon during training to improve target distribution estimates for the M-step. However, this scaling can also be performed after training when θ_i is fixed, enhancing performance, since actions are sampled in the environment according to the improved policy q . This equates to additional E-step optimization, which can enhance performance due to the representational constraints of the space of parameterised policies. In fig. 3 (right) we demonstrate how the performance of both SPO and AlphaZero search scales with additional search budget at test time using the same high performing checkpoint. We plot the time taken for a single actor step (measured on a TPUv3-8) against solve rate for the Rubik’s Cube problem with time on a logarithmic axis. Specifically, we evaluate on 1280 episodes for cubes ten scrambles away from solved. While we recognise that such analysis can be difficult to perform due to implementation discrepancies, we used the official JAX implementation of AlphaZero (MCTX) within our codebase and compare this to SPO. Additionally, we exclusively measure inference thus no training implementation details affect our measurements.

This provides evidence that AlphaZero has worse scaling when compared to SPO on horizons four and eight, noting that for horizon of two, SPO performance converges early, as low depths can act as a bottleneck for further performance improvements. This highlights the benefits of the SPO parallelism. This chart also shows how for different compute preferences at test time, a different horizon length is preferable. For limited compute, low horizon lengths and relatively higher particle counts provide the best performance, but as compute availability increases, the gains from increasing particle counts decrease and instead horizon length h should be increased.

Lastly, the plot illustrates the increase in available search budget, by using SPO, given a time restriction rather than a compute restriction. For example, SPO can use 4 times more search budget, when compared to AlphaZero, given a requirement of 0.1 second per step.

6 Conclusion

Planning-based policy improvement operators have proven to be powerful methods to enhance the learning of policies for complex environments when compared to model-free methods. Despite the success of these methods, they are still limited in their usefulness, due to the requirement of high planning budgets and the need for algorithmic modifications to adapt to large or continuous action spaces. In this paper, we introduce a modern implementation of Sequential Monte Carlo planning as a policy improvement operator, within the Expectation Maximisation framework. This results in a general training methodology with the ability to scale in both discrete and continuous environments.

Our work provides several key contributions. First, we show that SPO is a powerful policy improvement operator, outperforming our model-based and model-free baselines. Additionally, SPO is shown to be competitive across both continuous and discrete domains, without additional environment-specific alterations. This illustrates the effectiveness of SPO as a generic and robust approach, in contrast to prior methods that require domain-specific enhancements. Furthermore, the parallelisable nature of SPO results in efficient scaling behaviour of search. This allows SPO to achieve a significant boost in performance at inference time by scaling the search budget. Finally, we demonstrate that scaling SPO results in faster wall-clock inference time compared to previous work utilising tree-based search methods. The presented work culminates in a versatile, neural-guided, sample-based planning method that demonstrates superior performance and scalability over baselines.

Limitations & Future Work: This work demonstrates the efficacy of combining probabilistic inference with reinforcement learning and amortisation. While our work considers a relatively standard form of Sequential Monte Carlo, future work could investigate making improvements to this inference method, in order to further improve the estimate of the target distribution, while maintaining the scalability benefits demonstrated. We also draw the connection to Active Inference [27] which has been explored in the context of agent behaviour in complex environments, where Monte Carlo Tree Search has been used to scale previous methods along with deep learning [25]. Leveraging Sequential Monte Carlo in such methods along with amortisation could be a promising direction of research to further scale such methods. Our work exclusively uses exact world models of the environment in order to isolate the effects of various planning methods on performance. Extending our research to include a learned world model would broaden the applicability of SPO to more complex problems that may not have a perfect simulator and are therefore unsuitable for direct planning. Additionally, we apply Sequential Monte Carlo (SMC) only in deterministic settings. Adapting our approach for stochastic environments presents challenges. Specifically, the importance weight update in SMC can lead to the selection of dynamics most advantageous to the agent, potentially fostering risk-seeking behaviour in stochastic settings. However, mitigation strategies exist, as discussed in Levine [45].

Acknowledgments and Disclosure of Funding

The authors would like to thank Teodora Pandeva for detailed discussions on SPO theory and for detailed revisions of the manuscript. Clément Bonnet for brainstorming discussions and feedback and David Kuric for revisions of the final manuscript. We also thank members of the InstaDeep team for their support in the preparation of the final manuscript. Lastly we thank the anonymous reviewers for comments and helpful discussions that helped improve the final version of the paper.

Research was supported with Cloud TPUs from Google’s TPU Research Cloud (TRC). Matthew Macfarlane is supported by the LIFT-project 019.011, which is partly financed by the Dutch Research Council (NWO).

References

- [1] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=S1ANxQW0b>.

- [2] Jacob Adkins, Michael Bowling, and Adam White. A method for evaluating hyperparameter sensitivity in reinforcement learning. In *Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks*.
- [3] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Belle-mare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34:29304–29320, 2021.
- [4] Brandon Amos et al. Tutorial on amortized optimization. *Foundations and Trends® in Machine Learning*, 16(5):592–732, 2023.
- [5] Brian DO Anderson and John B Moore. *Optimal control: linear quadratic methods*. Courier Corporation, 2007.
- [6] Christophe Andrieu, Arnaud Doucet, Sumeetpal S Singh, and Vladislav B Tadic. Particle methods for change detection, system identification, and control. *Proceedings of the IEEE*, 92(3):423–438, 2004.
- [7] Thomas Anthony, Zheng Tian, and David Barber. Thinking fast and slow with deep learning and tree search. *Advances in neural information processing systems*, 30, 2017.
- [8] Ioannis Antonoglou, Julian Schrittwieser, Sherjil Ozair, Thomas K Hubert, and David Silver. Planning in stochastic environments with a learned model. In *International Conference on Learning Representations*, 2021.
- [9] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- [10] Alan Bain and Dan Crisan. *Fundamentals of stochastic filtering*, volume 3. Springer, 2009.
- [11] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [12] Clément Bonnet, Daniel Luo, Donal John Byrne, Shikha Surana, Paul Duckworth, Vincent Coyette, Laurence Illing Midgley, Sasha Abramowitz, Elshadai Tegegn, Tristan Kalloniatis, et al. Jumanji: a diverse suite of scalable reinforcement learning environments in jax. In *The Twelfth International Conference on Learning Representations, 2023*.
- [13] Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin Szeto, Nazanin Mohammadi Sepahvand, Edward Raff, Kanika Madan, Vikram Voleti, et al. Accounting for variance in machine learning benchmarks. *Proceedings of Machine Learning and Systems*, 3:747–769, 2021.
- [14] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [15] Cédric Colas, Olivier Sigaud, and Pierre-Yves Oudeyer. Gep-pg: Decoupling exploration and exploitation in deep reinforcement learning algorithms. In *International conference on machine learning*, pages 1039–1048. PMLR, 2018.
- [16] Rémi Coulom. Efficient selectivity and backup operators in Monte-Carlo tree search. In *International conference on computers and games*, pages 72–83. Springer, 2006.
- [17] Gal Dalal, Assaf Hallak, Steven Dalton, Shie Mannor, Gal Chechik, et al. Improve agents without retraining: Parallel tree search with off-policy correction. *Advances in Neural Information Processing Systems*, 34:5518–5530, 2021.
- [18] Ivo Danihelka, Arthur Guez, Julian Schrittwieser, and David Silver. Policy improvement by planning with gumbel. In *International Conference on Learning Representations*, 2021.
- [19] Peter Dayan and Geoffrey E Hinton. Using expectation-maximization for reinforcement learning. *Neural Computation*, 9(2):271–278, 1997.

- [20] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39 (1):1–22, 1977.
- [21] Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and computing*, 10:197–208, 2000.
- [22] Arnaud Doucet, Nando De Freitas, Neil James Gordon, et al. *Sequential Monte Carlo methods in practice*, volume 1. Springer, 2001.
- [23] Rotem Dror, Segev Shlomov, and Roi Reichart. Deep dominance-how to properly compare deep neural models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2773–2785, 2019.
- [24] Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Francisco J. R. Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, David Silver, Demis Hassabis, and Pushmeet Kohli. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature*, 610(7930):47–53, 2022. doi: 10.1038/s41586-022-05172-4.
- [25] Zafeirios Fountas, Noor Sajid, Pedro Mediano, and Karl Friston. Deep active inference agents using monte-carlo methods. *Advances in neural information processing systems*, 33:11662–11675, 2020.
- [26] C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax - a differentiable physics engine for large scale rigid body simulation, 2021. URL <http://github.com/google/brax>.
- [27] Karl Friston. The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127–138, 2010.
- [28] Thomas Furnstom and David Barber. Variational methods for reinforcement learning. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 241–248. JMLR Workshop and Conference Proceedings, 2010.
- [29] Hiroki Furuta, Tadashi Kozuno, Tatsuya Matsushima, Yutaka Matsuo, and Shixiang Shane Gu. Co-adaptation of algorithmic and implementational innovations in inference-based deep reinforcement learning. *Advances in neural information processing systems*, 34:9828–9842, 2021.
- [30] Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov chain Monte Carlo in practice*. CRC press, 1995.
- [31] Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE proceedings F (radar and signal processing)*, volume 140, pages 107–113. IET, 1993.
- [32] Rihab Gorsane, Omayma Mahjoub, Ruan John de Kock, Roland Dubb, Siddarth Singh, and Arnu Pretorius. Towards a standardised performance evaluation protocol for cooperative marl. *Advances in Neural Information Processing Systems*, 35:5510–5521, 2022.
- [33] Jean-Bastien Grill, Florent Altché, Yunhao Tang, Thomas Hubert, Michal Valko, Ioannis Antonoglou, and Rémi Munos. Monte-carlo tree search as regularized policy optimization. In *International Conference on Machine Learning*, pages 3769–3778. PMLR, 2020.
- [34] Shixiang Shane Gu, Zoubin Ghahramani, and Richard E Turner. Neural adaptive sequential monte carlo. *Advances in neural information processing systems*, 28, 2015.
- [35] Arthur Guez, Mehdi Mirza, Karol Gregor, Rishabh Kabra, Sebastien Racaniere, Theophane Weber, David Raposo, Adam Santoro, Laurent Orseau, Tom Eccles, Greg Wayne, David Silver, Timothy Lillicrap, and Victor Valdes. An investigation of model-free planning: boxoban levels. <https://github.com/deepmind/boxoban-levels/>, 2018.

- [36] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *International conference on machine learning*, pages 1352–1361. PMLR, 2017.
- [37] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- [38] Hirotaka Hachiya, Jan Peters, and Masashi Sugiyama. Efficient sample reuse in em-based policy search. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2009, Bled, Slovenia, September 7-11, 2009, Proceedings, Part I 20*, pages 469–484. Springer, 2009.
- [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [40] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- [41] Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Mohammadamin Barekatin, Simon Schmitt, and David Silver. Learning and planning in complex action spaces. In *International Conference on Machine Learning*, pages 4476–4486. PMLR, 2021.
- [42] Hilbert J Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine learning*, 87:159–182, 2012.
- [43] Genshiro Kitagawa. Monte carlo filter and smoother for non-gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25, 1996.
- [44] Alessandro Lazaric, Marcello Restelli, and Andrea Bonarini. Reinforcement learning in continuous action spaces through sequential monte carlo methods. *Advances in neural information processing systems*, 20, 2007.
- [45] Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- [46] Sergey Levine and Vladlen Koltun. Guided policy search. In *International conference on machine learning*, pages 1–9. PMLR, 2013.
- [47] Haim Levy. Stochastic dominance and expected utility: Survey and analysis. *Management Science*, pages 555–593, 1992.
- [48] Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *First International Conference on Informatics in Control, Automation and Robotics*, volume 2, pages 222–229. SciTePress, 2004.
- [49] Yingzhen Li, Richard E Turner, and Qiang Liu. Approximate inference with amortised mcmc. *arXiv preprint arXiv:1702.08343*, 2017.
- [50] Vasileios Lioutas, Jonathan Wilder Lavington, Justice Sefas, Matthew Niedoba, Yunpeng Liu, Berend Zwartsenberg, Setareh Dabiri, Frank Wood, and Adam Scibior. Critic sequential monte carlo. In *The Eleventh International Conference on Learning Representations*, 2022.
- [51] Anji Liu, Jianshu Chen, Mingze Yu, Yu Zhai, Xuewen Zhou, and Ji Liu. Watch the unobserved: A simple approach to parallelizing monte carlo tree search. In *International Conference on Learning Representations*, 2019.
- [52] Jun S Liu and Rong Chen. Sequential monte carlo methods for dynamic systems. *Journal of the American statistical association*, 93(443):1032–1044, 1998.
- [53] Jun S Liu, Rong Chen, and Tanya Logvinenko. A theoretical framework for sequential importance sampling with resampling. In *Sequential Monte Carlo methods in practice*, pages 225–246. Springer, 2001.

- [54] Zuxin Liu, Zhepeng Cen, Vladislav Isenbaev, Wei Liu, Steven Wu, Bo Li, and Ding Zhao. Constrained variational policy optimization for safe reinforcement learning. In *International Conference on Machine Learning*, pages 13644–13668. PMLR, 2022.
- [55] Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60, 1947.
- [56] Simon Maskell and Neil Gordon. A tutorial on particle filters for on-line nonlinear/non-gaussian bayesian tracking. *IEE Target Tracking: Algorithms and Applications (Ref. No. 2001/174)*, pages 2–1, 2001.
- [57] Thomas M Moerland, Joost Broekens, Aske Plaat, and Catholijn M Jonker. A0c: Alpha zero in continuous action space. *arXiv preprint arXiv:1805.09613*, 2018.
- [58] William Montgomery and Sergey Levine. Guided policy search as approximate mirror descent. *arXiv preprint arXiv:1607.04614*, 2016.
- [59] Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- [60] Gerhard Neumann et al. Variational inference for policy search in changing situations. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, pages 817–824, 2011.
- [61] Jing-Cheng Pang, Pengyuan Wang, Kaiyuan Li, Xiong-Hui Chen, Jiacheng Xu, Zongzhang Zhang, and Yang Yu. Language model self-improvement by reinforcement learning contemplation. *arXiv preprint arXiv:2305.14483*, 2023.
- [62] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.
- [63] Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pages 745–750, 2007.
- [64] Jan Peters, Katharina Mulling, and Yasemin Altun. Relative entropy policy search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, pages 1607–1612, 2010.
- [65] Alexandre Piché, Valentin Thomas, Cyril Ibrahim, Yoshua Bengio, and Chris Pal. Probabilistic planning with sequential monte carlo methods. In *International Conference on Learning Representations*, 2018.
- [66] Michael K Pitt and Neil Shephard. Auxiliary variable based particle filters. *Sequential Monte Carlo methods in practice*, pages 273–293, 2001.
- [67] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [68] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- [69] John Schulman. Trust region policy optimization. *arXiv preprint arXiv:1502.05477*, 2015.
- [70] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1506.02438>.
- [71] John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017.

- [72] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [73] Richard B Segal. On the scalability of parallel uct. In *International Conference on Computers and Games*, pages 36–47. Springer, 2010.
- [74] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [75] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [76] H Francis Song, Abbas Abdolmaleki, Jost Tobias Springenberg, Aidan Clark, Hubert Soyer, Jack W Rae, Seb Noury, Arun Ahuja, Siqi Liu, Dhruva Tirumala, et al. V-mpo: On-policy maximum a posteriori policy optimization for discrete and continuous control. In *International Conference on Learning Representations*, 2019.
- [77] Wen Sun, Geoffrey J Gordon, Byron Boots, and J Bagnell. Dual policy iteration. *Advances in Neural Information Processing Systems*, 31, 2018.
- [78] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [79] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
- [80] Edan Toledo. Stoix: Distributed single-agent reinforcement learning end-to-end in jax, April 2024. URL <https://github.com/EdanToledo/Stoix>.
- [81] Edan Toledo, Laurence Midgley, Donal Byrne, Callum Rhys Tilbury, Matthew Macfarlane, Cyprien Courtot, and Alexandre Laterre. Flashbox: Streamlining experience replay buffers for reinforcement learning with jax, 2023. URL <https://github.com/instadeepai/flashbox/>.
- [82] Marc Toussaint and Amos Storkey. Probabilistic inference for solving discrete and continuous state markov decision processes. In *Proceedings of the 23rd international conference on Machine learning*, pages 945–952, 2006.
- [83] Nino Vieillard, Olivier Pietquin, and Matthieu Geist. Munchausen reinforcement learning. *Advances in Neural Information Processing Systems*, 33:4235–4246, 2020.
- [84] Christian Wirth, Johannes Fürnkranz, and Gerhard Neumann. Model-free preference-based reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

Appendix

Table of Contents

A	Environments	17
A.1	Sokoban	17
A.2	Rubik’s Cube	17
A.3	Brax	18
B	Ablations	19
B.1	E-step KL constraint	19
B.2	E-step Q-value Optimisation	19
B.3	SMC Target Estimation Validation	20
C	Expanded Results	21
C.1	Detailed Summary Results	21
C.2	Hardware	22
C.3	Individual Environment Results	23
D	SPO	24
D.1	Temperature Loss	24
D.2	Approximate Policy Iteration Algorithm	24
D.3	Hyperparameters	25
D.4	Intuition for Temperature and Exploration:	26
E	Baselines	26
E.1	Hyperparameters	26
F	Expectation Maximisation	30
F.1	Overview of methods	30
G	Proofs and Discussions	31
G.1	E-step Analytic solution	31
G.2	E-step KL constraining temperature	31
G.3	E step: Constraint Optimisation	32
G.4	Proof of Proposition 1: Monotone Improvement Guarantee	32
G.5	M-step objective	32
G.6	Connection to Mirror Descent Guided Policy Search	33
H	Statistical Precipice	34
H.1	Overview	34
H.2	Hyperparameters	34

A Environments

A.1 Sokoban

A.1.1 Overview

Table 1: Summary of Boxoban dataset levels

Difficulty Level	Dataset Size
Unfiltered-Train	900k
Unfiltered-Validation	100k
Unfiltered-Test	1k
Medium	450k
Hard	50k

We use the specific instance of Sokoban outlined in Guez et al. [35] illustrated in Figure 4 with the codebase available at <https://github.com/instantdeepai/jumanji>. The datasets employed in this study are publicly accessible at <https://github.com/google-deepmind/boxoban-levels>. These datasets are split into different levels of difficulty, which are categorised in Table 1.

In this research, we always train on the *Unfiltered-Train* dataset. Evaluations are performed on the *Hard* dataset, important for differentiating the strongest algorithms.

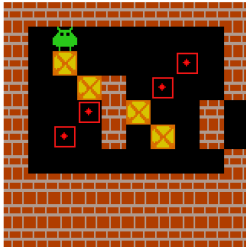


Figure 4: Example of a Boxoban Problem

A.1.2 Network Architecture

Observations are represented as an array of size (10,10) where each entry is an integer representing the state of a cell. We used a deep ResNet [39] architecture for the torso network to embed the observation. We define a block as consisting of the following layers: [CNN, ResNet, ResNet]. Four such blocks are stacked, each characterized by specific parameters:

- Output Channels: (256, 256, 512, 512)
- Kernel Sizes: (3, 3, 3, 3)
- Strides: (1, 1, 1, 1)

Additionally, the architecture includes two output heads (policy and value), each comprising two layers of size 128 with ReLU activations. The output heads share the same torso network. We use the same architecture for all algorithms.

A.2 Rubik’s Cube

A.2.1 Overview

For our Rubik’s Cube experiments, we utilize the implementation in Bonnet et al. [12] available at <https://github.com/instantdeepai/jumanji>. Rubik’s cube problems can be made progressively difficult by scaling the number of random actions performed in a solution state to generate problem instances. We always perform training on a uniform distribution sampled from the range [3,7] of scrambled states, followed by exclusively evaluating on 7 scrambles. The observation is represented using an array of size (6,3,3). The action space is represented using an array of size (6,3) corresponding to each face and the depth at which it can be rotated.

A.2.2 Network Architecture

We utilise an embedding layer to first embed the face representations to size (6,3,3,4). The embedding is flattened and we use a torso layer consisting of a two layer residual network with layer size 512 and ReLU activations. Additionally, the architecture includes two output heads (policy and value), each comprising of a single layer of size 512 with ReLU activations. We use the same architecture for all algorithms.

A.3 Brax

A.3.1 Overview

For experimentation on continuous settings, we make use of Brax [26]. Brax is a library for rigid body simulation much like the popular MuJoCo setting [79] simulator. However, Brax is written entirely in JAX [14] to fully exploit the parallel processing capabilities of accelerators like GPUs and TPUs.

It is important to note, that at the time of writing, there are 4 different physics back-ends that can be used. These back-ends have differing levels of computational complexity and the results between them are not comparable. The results generated in this paper utilise the Spring backend.

Table 2: Observation and action space for Brax environments

Environment	Observation Size	Action Size
Halfcheetah	18	6
Ant	27	8
Humanoid	376	17

Table 2 contains the observation and action specifications of the scenarios used for our experiments. We specify the dimension size of the observation and action vectors.

A.3.2 Network Architecture

In practice, we found the smaller networks used in the original Brax publication to limit overall performance. We instead use a 4 layer feed forward network⁶ to represent both the value and policy network. Aside from the output layers, all layers are of size 256 and use SiLU non-linearities. We use the same architecture for all algorithms. Unlike Rubik’s Cube and Sokoban, we do not use a shared torso to learn embeddings.

⁶except for SMC-ENTR, for which we use 3 to stabilise performance.

B Ablations

B.1 E-step KL constraint

The KL constraint for the E-step within the Expectation Maximisation framework is an important component of ensuring policy improvement, due to the KL term in the ELBO. We investigate the impact of the adaptive temperature for both discrete and continuous environments. We compare SPO with an adaptive temperature updated every iteration to a variety of fixed temperatures.

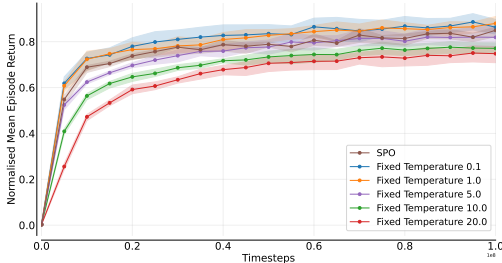


Figure 5: Brax

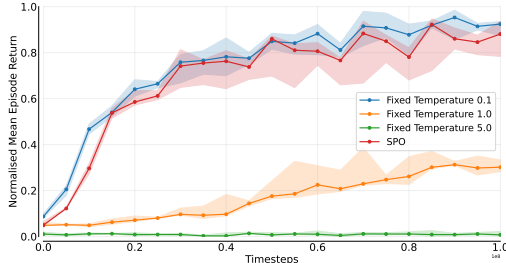


Figure 6: Sokoban and Rubik's Cube

This ablation shows that SPO with an adaptive temperature is among the top performing hyperparameter settings across all environments. However we also note that it is possible to tune a temperature that works well when considering a wide range of temperatures. This is consistent with previous results in Peng et al. [62] that also find practically for specific problems a fixed temperature can be used. Of course in practice having an algorithm that can learn this parameter itself is practically beneficial, removing the need for costly hyperparameter tuning, since the appropriate temperature is likely problem dependant.

Subsequently, we evaluated whether the partial optimisation of the temperature parameter effectively maintained the desired KL divergence constraint and how different values of this constraint affected performance.

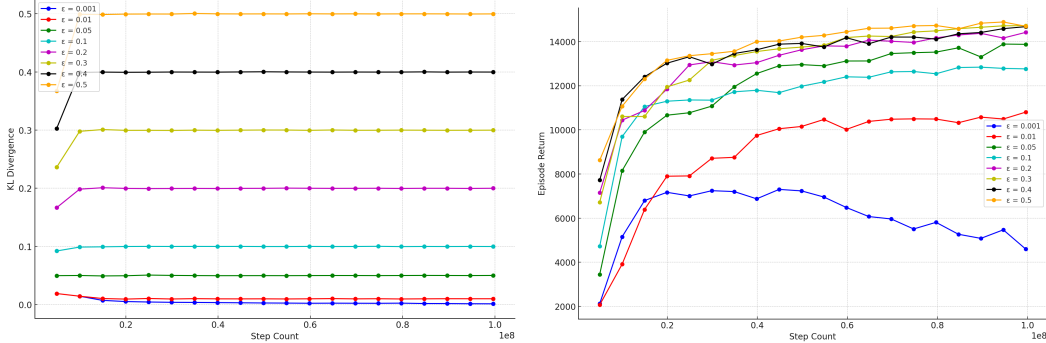


Figure 7: (a) The estimated KL divergence between the prior policy and the target policy q generated by SMC for different values of ϵ during training on the Brax Ant task. (b) Evaluation performance during training for different values of ϵ .

As shown in Figure 7(a), the KL divergence constraint is tightly maintained throughout training for different values of ϵ . Figure 7(b) demonstrates that, for the Ant task, allowing a larger KL divergence between the prior and target policies (q and π , respectively) leads to improved performance.

B.2 E-step Q-value Optimisation

For our analytic solution to the optimisation problem in the E-step we choose to add a value baseline such that our target distribution re-weights with respect to Advantages instead of Q-values. In practice there are no restrictions on utilising Q-values instead for the importance weight update of SMC eq. (10). Below we show results on the continuous environments comparing SPO using advantages to SPO using Q-values.

We see in Figure 8, that the use of Q values produces a reduction in performance that is relatively constant throughout training. It is possible that this reduction is due to different hyperparameters being required to effectively utilise the Q-value distribution but in order to accurately judge the effect, we kept hyperparameters constant

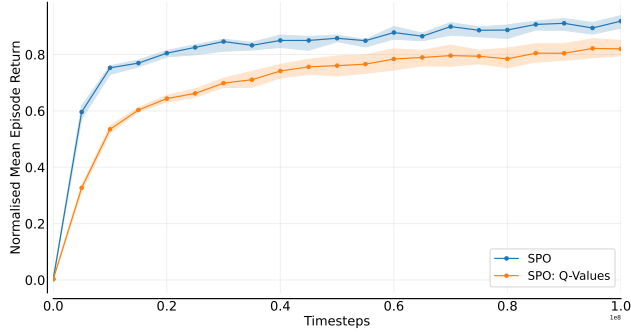


Figure 8: Q-value ablation on Brax tasks

across the runs. Secondly we note this is consistent with existing results for EM algorithms, demonstrating that Advantages outperform Q-values Peng et al. [62], Song et al. [76], Nair et al. [59].

B.3 SMC Target Estimation Validation

Estimating the target distribution accurately is important for policy improvement in SPO. Therefore, we aim to validate the ability of SMC search to better approximate the true target distribution than a simple 1-step function evaluation as is done in algorithms such as MPO [1]. To do this, while the true target distribution is not directly accessible, we approximate it using an unbiased Monte Carlo oracle. We use a large computational budget of 1280 rollouts, to the end of the episode, for every state and every action. We evaluate the impact of varying planning horizons (depth) and particle counts on the KL divergence between the SMC-estimated target policy and the Monte Carlo oracle in the Sokoban environment.

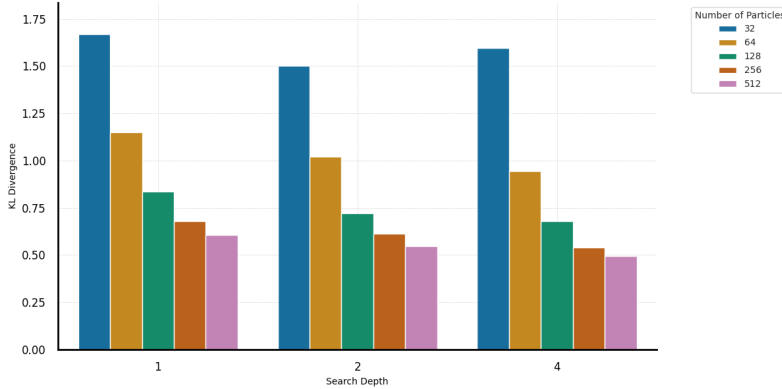


Figure 9: Comparison of the KL divergence to large Monte Carlo simulation of target policy for different planning horizons and particle counts for Sokoban

As shown in Figure 9, increasing the number of particles and the planning depth reduces the KL divergence to the oracle, indicating improved estimation of the target distribution. This aligns with SMC theory [22], which suggests that more particles and deeper planning lead to better approximations. These results highlight the importance of leveraging both breadth (more particles) and depth (longer planning horizons) in SMC for accurate target estimation.

C Expanded Results

We provide more detailed analysis of the results presented in the main body of the paper, including the final point aggregate performance of algorithms, the probability of improvement plots and the performance profiles. All plots are generated according to the methodology outlined by Agarwal et al. [3] and Gorsane et al. [32] by performing the final evaluation using parameters that achieved the best performance throughout intermediate evaluation runs during training.

The point aggregation plots show the rankings of normalised episode returns of each algorithm when different aggregation metrics are used. The probability of improvement plots illustrate the likelihood that algorithm X outperforms algorithm Y on a randomly selected task. It is important to note that a statistically significant result from this metric is a probability of improvement greater than 0.5 where the confidence intervals (CIs) do not contain 0.5. The metric utilises the Mann-Whitney U-statistic [55]. See [3] for further details. The performance profiles illustrate the fraction of the runs over all training environments from each algorithm that achieved scores higher than a specific value (given on the X-axis). Functionally, this serves the same purpose as comparing average episodic returns for each algorithm in a tabular form but in a simpler format. Additionally, if an algorithm’s curve is strictly greater than or equal to another curve, this indicates “stochastic dominance” [47, 23].

C.1 Detailed Summary Results

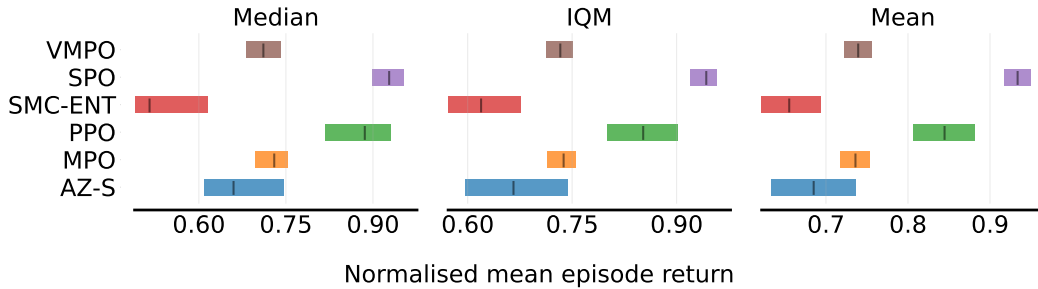


Figure 10: Aggregate point metrics for Brax suite. 95% confidence intervals generated from stratified bootstrapping across tasks and seeds are reported.

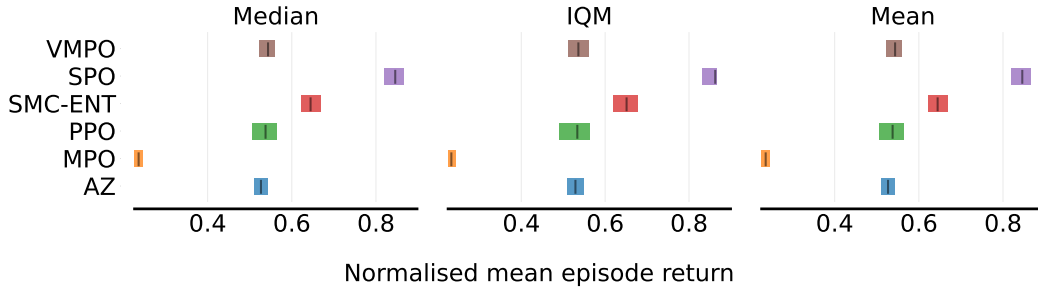


Figure 11: Aggregate point metrics for Sokoban and Rubik’s Cube. 95% confidence intervals generated from stratified bootstrapping across tasks and seeds are reported.

In Figures 10 and 11, we present the absolute metrics following the recommendations by Agarwal et al. [3], Colas et al. [15] and Gorsane et al. [32]. These metrics evaluate the best set of network parameters identified during 20 evenly spaced training evaluation intervals. The absolute evaluations measure performance across 10 times the number of episodes periodically assessed during training, resulting in a total of 1280 episodes evaluated. We report point estimates and their 95% confidence intervals, which were calculated using stratified bootstrapping. Agarwal et al. [3] advocate for the Interquartile Mean (IQM) as a more robust and valid point estimate metric. Our results indicate that SPO surpasses all other algorithms in every point estimate. Notably, the IQM and mean point estimates demonstrate statistical significance.

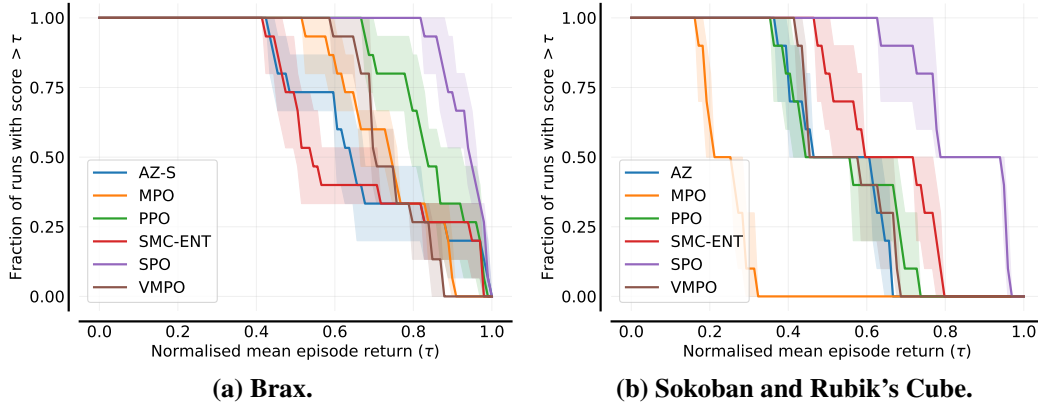


Figure 12: Performance profiles. The Y-axis represents the fraction of runs that achieved greater than a specific normalised score represented on the x-axis, with shaded regions indicating 95% confidence intervals generated from stratified bootstrapping across both tasks and random seeds.

In Figure 12, we present the performance profiles which visually illustrates the full distribution of scores across all tasks and seeds for each algorithm. We see that SPO has a higher lower bound on performance, in addition to upper bound, indicating lower variance across tasks and seeds. Additionally, we see that SPO is strictly above all other algorithms indicating that SPO is stochastically dominant⁷ to all baselines.

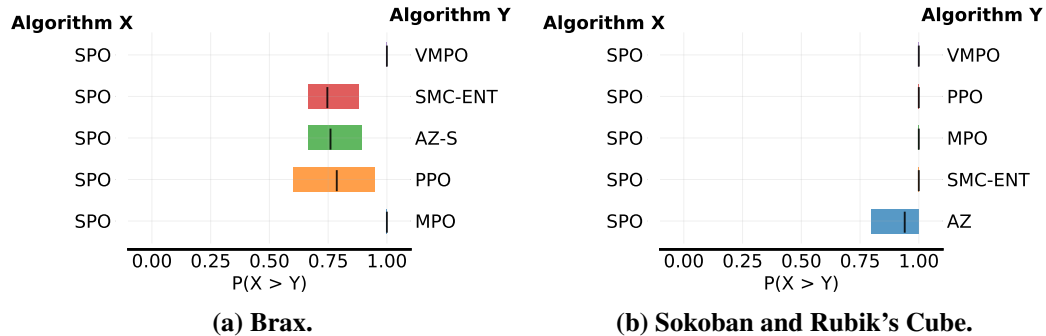


Figure 13: Probability of Improvement.

Lastly, we present the probability of improvement plots in Figure 13. We see that SPO has a high probability of improvement compared to all baselines. Additionally, all probabilities are larger than 0.5 and have their CIs above 0.5 thus indicating statistical significance as specified by Agarwal et al. [3]. Furthermore, we see all CIs have upper bounds greater than 0.75, thereby indicating that these results are statistically meaningful as specified by Bouthillier et al. [13].

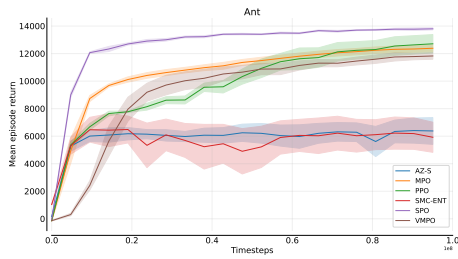
C.2 Hardware

Training was performed using a mixture of Google v4-8 and v3-8 TPUs. Each experiment was run using a single TPU and only v3-8 TPUs were used to compare wall-clock time.

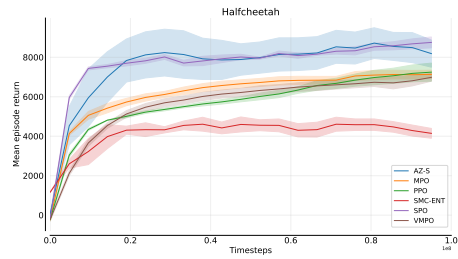
⁷A random variable X is termed stochastically dominant over another random variable Y if $P(X > \tau) \geq P(Y > \tau)$ for all τ , and for some τ , $P(X > \tau) > P(Y > \tau)$.

C.3 Individual Environment Results

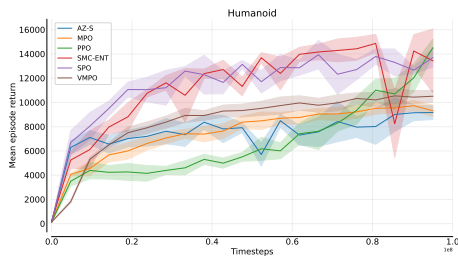
All individual task results are presented in fig. 14. Specifically, we present the IQM of returns achieved during the evaluation intervals throughout training.



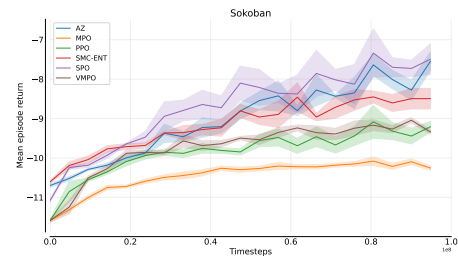
(a) Ant.



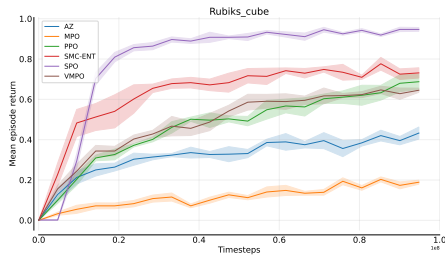
(b) HalfCheetah.



(c) Humanoid.



(d) Sokoban.



(e) Rubik's Cube 7 scrambles

Figure 14: Performance across different environments.

D SPO

Below we outline the practical policy- based losses used within SPO,

$$L(\theta) = \sum_{s; a \sim \mathcal{D}} q_i(a|s) \log q_i(a|s)$$

$$g(\theta) = \mathbb{E}_s \int (s) \log \left(\int (a|s; \theta) \exp \left(\frac{A(s; a)}{\tau} \right) da \right) ds:$$

$$L(\theta; \theta_{old}) = \mathbb{E}_{s \sim p(s)} [\text{sg} [D_{KL}(\theta_{old} \| \theta)]]$$

$$L_{KL}(\theta; \theta_{old}) = \text{sg} [\mathbb{E}_{s \sim p(s)} [D_{KL}(\theta_{old} \| \theta)]]$$

D.1 Temperature Loss

The temperature loss is calculated using advantages collected during SPO search with rollouts performed according to policy θ . If resampling is utilised, after a resample the distribution over trajectories shifts towards the target distribution away from θ . Therefore in order to calculate the loss $g(\theta)$, we only utilise the advantages up to the first resampling period with rollouts performed according to θ .

D.2 Approximate Policy Iteration Algorithm

Below we outline the overall SPO algorithm with the main loop split into the E-step and M-step.

Algorithm 2 SPO Algorithm

```

1:  $B \leftarrow \emptyset$ ;
2: Initialize the policy, value function, temperature, and alpha:
3: Initialize  $\theta_0, \theta_0, \theta_0, \theta_0$ 
4: for iteration  $k = 0; 1; 2; \dots; \text{num\_iterations}$  do
5:   Expectation Step (E-step):
6:   for agent  $m = 1; 2; \dots; M$  in parallel do
7:     Initialize state  $s_0$ 
8:     for timestep  $i = 1; 2; \dots; \text{rollout\_length}$  do
9:        $a = \hat{f}a^{(n)} g_{n=1}^N \hat{q}_{SMC}(S_i; k; k; k)$  . sample-based estimate of  $q_i$ 
10:       $a \sim \hat{f}a^{(n)} g_{n=1}^N$  . Sample an action to execute in the environment
11:       $s_{i+1} \sim T(S_i; a)$  . Execute action  $a$ , observe next state  $s_{i+1}$ 
12:       $r_i = r(s_i; a)$  . Observe reward  $r_i$ 
13:      Store  $(s_i; r_i; a)$  in  $B$ 
14:    end for
15:  end for
16:  Maximization Step (M-step):
17:  for batch  $b = 1; 2; \dots; \text{num\_batches}$  do
18:    Sample batch from replay buffer:
19:    Sample batch of size  $N$  from replay buffer  $(s; r; a) \sim B$ 
20:    Update value function using GAE targets:
21:     $q_{k+1} = \arg \min_{q'} \mathbb{E}_{s \sim B} [(V_{GAE}(s; q') - V(s; q))^2]$ 
22:    Update parameterized policy using sample-based estimate of  $q$ :
23:     $\theta_{k+1}, \tau_{k+1} = \arg \min_{\theta, \tau} \mathbb{E}_{s \sim B} [L(\theta; \tau) + L_{KL}(\theta; \tau)]$ 
24:    Update dual temperature:
25:     $\tau_{k+1} = \arg \min_{\tau} \mathbb{E}_{s \sim B} [g(\tau)]$ 
26:     $\theta_{k+1} = \text{polyak}(\theta; \theta_{k+1})$ 
27:  end for
28: end for

```

D.3 Hyperparameters

Table 3: SPO Hyperparameters for Continuous and Discrete Environments

Parameter	Continuous Environments	Discrete Environments
Actor & Critic Learning Rate	$3e^{-4}$	$3e^{-4}$
Dual Learning Rate	$1e^{-3}$	$3e^{-4}$
Discount Factor	0.99	0.99
GAE Lambda	0.95	0.95
Replay Buffer Size	$6.5e^4$	$6.5e^4$
Batch Size	32	64
Batch Sequence Length	32	17
Max Grad Norm	0.5	0.5
Number of Epochs	128	16
Number of Envs	1024	768
Rollout Length	32	21
(Target Smoothing)	$5e^{-3}$	$5e^{-3}$
Number of Particles	16	16
Search Horizon	4	4
Resample Period	4	4
Initial	10	0.5
Initial	-	0.5
Initial	10	-
Initial	500	-
	0.2	0.5
	-	$1e^{-3}$
	$5e^{-2}$	-
	$5e^{-4}$	-
Dirichlet Alpha	-	0.03
Root Exploration Weight	-	0.25

D.4 Intuition for Temperature and Exploration:

The temperature parameter τ plays a crucial role in balancing exploration and exploitation during the search process in SPO. In this work, it is derived via first choosing the desired KL divergence ϵ between the target policy q_i generated by search and the current policy π_i . If the temperature is too low (i.e., τ is small), the exponential weighting $\exp(A(s; a) / \tau)$ becomes sharply peaked around the highest-advantage actions. This causes the importance weights to concentrate on a few particles, and during resampling, the particles can collapse onto a single root action. Such collapse reduces diversity in the search and renders the rest of the planning process ineffective. Conversely, if the temperature is too high (i.e., τ is large), the weighting becomes flatter, leading to excessive exploration. While exploration is necessary, too much can prevent the search from effectively focusing on promising paths. This work shows that deriving this value via a target KL leads to stable and strong performance across different environments.

E Baselines

For the baseline implementations, we expanded and adapted the existing implementations from Stoix⁸ [80]. All implementations were conducted using JAX [14], with off-policy algorithms leveraging Flashbox [81].

E.1 Hyperparameters

E.1.1 Continuous Control

Table 4: Hyperparameters for PPO and Sampled AlphaZero

Parameter	PPO	Sampled AlphaZero
Actor Learning Rate	0.00069	$3e^{-4}$
Critic Learning Rate	0.00054	$3e^{-4}$
Rollout Length	16	32
Number of Epochs	4	64
Number of Minibatches	16	-
Buffer Size	-	65536
Batch Size	-	32
Sample Sequence Length	-	32
Discount Factor	0.99	0.99
GAE Lambda	0.95	0.95
Clip Epsilon	0.1	-
Entropy Coefficient	0.005	0.005
Value Function Coefficient	0.5	-
Max Grad Norm	0.5	0.5
Decay Learning Rates	True	-
Standardize Advantages	True	-
Number of Simulations	-	64
Dirichlet Alpha	-	0.03
Dirichlet Exploration Fraction	-	0.25
Number of Samples	-	20
Gaussian Noise Exploration Fraction	-	0.001

⁸Available at: <https://github.com/EdanToledo/Stoix>

Table 5: Hyperparameters for MPO and VMPO

Parameter	MPO	VMPO
Rollout Length	8	32
Number of Epochs	72	16
Buffer Size	200000	-
Batch Size	256	-
Sample Sequence Length	16	-
Period	1	-
Actor Learning Rate	$1e^{-4}$	$3e^{-4}$
Critic Learning Rate	-	$3e^{-4}$
Dual Learning Rate (Target Smoothing)	$1e^{-3}$	$1e^{-2}$
Discount Factor	0.005	0.005
Max Grad Norm	0.99	0.99
Decay Learning Rates	0.5	0.5
Number of Samples	True	-
	128	-
	0.05	0.05
	0.05	0.05
	0.0005	0.0005
Initial	10.0	10.0
Initial	10.0	10.0
Initial	500	500
GAE Lambda	0.95	0.95
Actor Target Period	-	25

Table 6: Hyperparameters for SMC-Ent

Parameter	SMC-Ent
Rollout Length	16
Number of Epochs	512
Buffer Size	500000
Batch Size	512
Actor Learning Rate	$3e^{-4}$
Q Learning Rate (Target Smoothing)	$3e^{-4}$
Discount Factor	0.005
Reward Scale	0.99
Number of Particles	10.0
Max Depth	16
Resample Temperature	4
Number of Samples for Value Function	1.0
	64

E.1.2 Discrete Control

Table 7: Hyperparameters for AlphaZero and PPO

Parameter	AlphaZero	PPO
Learning Rate	$3e^{-4}$	$3e^{-4}$
Rollout Length	21	85
Number of Epochs	16	4
Buffer Size	65536	-
Batch Size	64	-
Sample Sequence Length	17	-
Discount Factor	0.99	0.99
GAE Lambda	0.95	0.95
Max Grad Norm	0.5	0.5
Number of Simulations	64	-
Number of Minibatches	-	64
Clip Epsilon	-	0.1
Entropy Coefficient	-	$1e^{-2}$
Standardize Advantages	-	True

Table 8: Hyperparameters for MPO and VMPO

Parameter	MPO	VMPO
Actor Learning Rate	$2e^{-4}$	$3e^{-4}$
Decay Learning Rates	True	-
Dual Learning Rate	0.02	0.01
Number of Epochs	72	4
	0.07	0.5
	0.00015	0.001
GAE Lambda	0.95	0.95
Discount Factor	0.95	0.99
Initial	3.0	3.0
Initial	3.0	3.0
Max Grad Norm	0.5	0.5
Number of Samples	128	-
Q Learning Rate	0.001	-
Rollout Length	16	86
Sample Sequence Length	17	-
(Target Smoothing)	0.005	-
Batch Size	256	-
Buffer Size	500000	-
Actor Target Period	-	64

Table 9: Hyperparameters for SMC-Ent

Parameter	SMC-Ent
Rollout Length	32
Number of Epochs	64
Buffer Size	1000000
Batch Size	2048
Q Learning Rate	$3e^{-4}$
(Target Smoothing)	0.005
Discount Factor	0.97
Max Grad Norm	0.5
Huber Loss Parameter	1.0
Entropy Temperature	0.03
Munchausen Coefficient	0.9
Clip Value Min	$-1e^3$
Number of Particles	16
Search Horizon	4
Resampling Period	4
Resample Temperature	0.1
Reward Scaling	10.0 (Rubik's Cube) / 1.0 (Sokoban)

F Expectation Maximisation

F.1 Overview of methods

EM algorithms differ on a small number of dimensions through which the algorithm can be understood. The first dimension, G , is the optimization objective that is maximized in the E-step. This includes whether advantages, Q-values, or rewards are used, and with respect to which policy. We then consider whether a trust region constraint is used both in the E-step and M-step.

With a defined optimization objective, various methods can be used to estimate it. Most methods derive an analytic solution to the optimization problem and then estimate this distribution using techniques such as TD(0) or function approximation. For example, AlphaZero does not explicitly derive the analytic solution but estimates the solution to this optimization objective through Monte Carlo Tree Search (MCTS).

Table 10 provides a summary of some common EM methods, highlighting their core differences.

Table 10: Summary of EM based Algorithms

Method	G	E-step TR	M-step TR	G estimate	Depth	Breadth
MPO	Q^p	Yes	Yes	Analytic + Function Approximation	0	M
V-MPO	A^p	Yes	Yes	Analytic + n-step TD + top-k Adv	T	1
AWR	A^p	No	No	Analytic + n step TD	T	1
AWAC	A^p	No	No	Analytic + Function Approximation	0	1
PoWER	$\log Q^p$	Yes	No	Analytic + n step TD	T	1
RWR	$\log r$	No	No	-	1	1
REPS	A^p	Yes	No	Analytic + TD(0)	1	1
AlphaZero	Q^q	Yes	No	MCTS	> 0	> 0
SPO	A^q	Yes	Yes	Analytic + SMC	D	M

Note that T refers to an episode length. In addition we add two important dimensions of the G estimate (breadth and depth). Methods like MPO estimate the analytic solution directly using a Q-function. This can be leveraged to estimate the target distribution for a selection of M actions, but without leveraging depth, so rewards and future states are not used to improve the estimates. In contrast, V-MPO leverages depth to form an n-step TD estimate, but only for one of the actions, leading to a relatively poor estimate of the analytic solution.

G Proofs and Discussions

G.1 E-step Analytic solution

We outline the analytic solution to eq. (5), by writing a Lagrangian equation and solving for q . We can first represent our constrained optimisation exactly with the following optimisation problem and associated 2 constraints. We add a state dependent baseline $V(s)$ to the optimisation objective and notate $Q^q(s; a) = V^q(s)$ as $A^q(s; a)$.

$$\begin{aligned} \max_q \int q(s) \left[\int q(a|s) [Q^q(s; a) - V^q(s)] da \right] ds \\ \text{s.t. } \int q(s) [\text{KL}(q(a|s) \| p(a|s))] ds < \epsilon; \\ \int q(s) \left[\int q(a|s) da \right] ds = 1; \end{aligned} \quad (13)$$

The following Lagrangian L can be constructed,

$$\begin{aligned} L(q; \lambda; \mu) = \int q(s) \left[\int q(a|s) A^q(s; a) da \right] ds \\ + \left(\int q(s) \left[\int q(a|s) \log \left(\frac{q(a|s)}{p(a|s)} \right) da \right] ds \right) \\ + \left(1 - \int q(s) \left[\int q(a|s) da \right] ds \right) \lambda; \end{aligned} \quad (14)$$

We can then solve for the value of q that maximises this expression by taking the derivative with respect to q and setting it to 0.

$$\frac{\partial L(q; \lambda; \mu)}{\partial q} = A^q(s; a) - \log q(a|s) - \log p(a|s) - \lambda = 0; \quad (15)$$

The analytic form for the optimal distribution q can then be calculated as:

$$q(a|s) = p(a|s) \exp \left(\frac{A^q(s; a)}{p(a|s)} \right) \exp \left(-\frac{\lambda}{p(a|s)} \right); \quad (16)$$

G.2 E-step KL constraining temperature

We now derive the dual function to be minimised in order to obtain λ^* within the analytic solution for q , which is crucial for enforcing the KL constraint.

The final term $-\frac{\lambda}{p(a|s)}$ acts as a normalising constant since it is independent of q and therefore we can construct the following equality.

$$\exp \left(-\frac{\lambda}{p(a|s)} \right) = \int p(a|s) \exp \left(\frac{A^q(a; s)}{p(a|s)} \right) da; \quad (17)$$

Now that we have expressions for $-\frac{\lambda}{p(a|s)}$ and q we can form the dual function $g(\lambda)$ by substituting these terms back into the original Lagrangian. After simplifying we recover the following expression.

$$g(\lambda) = \lambda + \int q(s) \log \left(\int p(a|s) \exp \left(\frac{A^q(a; s)}{p(a|s)} \right) da \right) ds \quad (18)$$

The optimal dual variable can be calculated as follows

$$\lambda^* = \arg \min g(\lambda); \quad (19)$$

G.3 E step: Constraint Optimisation

In optimising eq. (4), the first term $\mathbb{E}_{q(a|s)} [Q^q(s; a)]$ is dependent on the scale of reward in the environment. This can make it hard to choose β , as the first and second terms are on arbitrary scales. Practically this can mean that for each new environment a new β should be used requiring costly hyperparameter tuning to find. As explored in previous work [1, 76] we choose to enforce a hard constraint, as opposed to a soft constraint. Instead of choosing β we choose δ which is the maximum KL divergence between q and p which is practically less sensitive to reward scale. We found that choosing an δ to be far easier, resulting in a value that generalised across all environments explored. This is important as there has been a trend in recent years to add hyperparameters to Reinforcement Learning algorithms [2] resulting in the need for costly hyperparameter sweeps for algorithms to work on each new problem they are applied to.

G.4 Proof of Proposition 1: Monotone Improvement Guarantee

The optimality of policy π^* , $\log p^*(O = 1)$, is lower bounded by the following ELBO objective:

$$\mathcal{J}(q; \pi^*) = \mathbb{E}_{\sim q} \left[\sum_{t=0}^{\infty} (\gamma^t r_t - D_{KL}(q(jS_t)k(jS_t; \pi^*))) \right] + \log p^*(\pi^*)$$

We improve \mathcal{J} by optimizing the ELBO alternatively via EM. We will outline the proof of the monotonic improvement guarantee of the ELBO using the EM procedure, at the i -th iteration of training. We note that this holds under the assumption that we calculate the true value of the closed form solution q_i in the E-step, which is in practice unlikely to be the case.

E-step: By the definition of E-step, we improve the ELBO with respect to q . We show the E-step update will increase ELBO:

$$\begin{aligned} q_{i+1} &= \arg \max_q \mathbb{E}_q [\mathbb{E}_{a \sim q(\cdot|s)} [A^{q_i}(s; a)] - D_{KL}[q(jS)k(jS; i)]] \\ &= \arg \max_q \mathbb{E}_{\sim q} \left[\sum_{t=0}^{\infty} (\gamma^t r_t - D_{KL}(q(jS_t)k(jS_t; i))) \right] \\ &= \arg \max_q \mathcal{J}(q; i) \\ &\geq \mathcal{J}(q_{i+1}; i) \geq \mathcal{J}(q_i; i): \end{aligned}$$

M-step: We update π by

$$\begin{aligned} \pi_{i+1} &= \arg \max_{\pi} \mathbb{E}_{q_{i+1}} [\mathbb{E}_{a \sim q_{i+1}(\cdot|s)} [\log \pi(a|s)] + \log p^*(\pi)] \\ &= \arg \max_{\pi} \mathbb{E}_{q_{i+1}} [-D_{KL}[q_{i+1}(jS_t)k(jS_t; i)] + \log p^*(\pi)] \\ &= \arg \max_{\pi} \mathbb{E}_{q_{i+1}} [\mathbb{E}_{a \sim q_{i+1}(\cdot|s)} [A^{q_i}(s; a)] - D_{KL}[q_{i+1}(jS_t)k(jS_t; i)] + \log p^*(\pi)] \\ &= \arg \max_{\pi} \mathcal{J}(q_{i+1}; \pi): \end{aligned}$$

Therefore, we have: $\mathcal{J}(q_{i+1}; \pi_{i+1}) \geq \mathcal{J}(q_{i+1}; i)$. Combining these two results we have that after successive applications of the E-step and M-step,

$$\mathcal{J}(q_{i+1}; \pi_{i+1}) \geq \mathcal{J}(q_i; i):$$

G.5 M-step objective

In Reinforcement Learning it can be beneficial to constrain policies from moving too far from the current policy, often leading to increased stability or performance of an algorithm [69, 72]. In this work we utilise a Gaussian prior around π_i , our current value of π , optimised from the previous iteration.

Therefore, $\pi \sim \mathcal{N}(\pi; \Sigma)$ where $\Sigma = \Sigma_i$ and $\Sigma_i^{-1} = F(\pi_i)$, where F is the Fisher information matrix. The prior term in eq. (11) can then be written as $\log p(\pi) = -(\pi - \pi_i)^T F(\pi_i)^{-1} (\pi - \pi_i) + c$. The first term in this expression is the quadratic approximation of the KL divergence [69] while the second term c is a term that does not depend on π and therefore when optimising eq. (11) with respect to π , can be dropped. Using this approximation we rewrite eq. (11) as:

$$\max_{\pi} \mathbb{E}_{s \sim q(s)} [\mathbb{E}_{a \sim q(a|s)} [\log \pi(a|s; i)] - \text{KL}(\pi(a|s; i)k(a|s; i))]$$

Like the E-step, choosing π can be non-trivial so we convert it into a hard constraint optimisation eq. (12). Note that π in eq. (12) is different from eq. (5). We note that considering a Gaussian prior in the M-step is not strictly required for SPO performance, however practically it adds stability to training.

G.6 Connection to Mirror Descent Guided Policy Search

We highlight the connection between the use of Expectation Maximisation to improve the evidence lower bound and Mirror Descent Guided Policy Search (MD-GPS). Mirror Descent Guided Policy Search builds upon previous guided policy search work [46]. Rather than only enforcing the constraint between q and π at convergence (referred to as the local policy: p_i and global policy: q respectively), they constrain q_i against the current policy π_i at every iteration. This results in a very similar algorithm to EM optimisation. Below we provide the outline of MD-GPS algorithm, clearly showing the equivalence to EM [58]:

Algorithm 3 Mirror descent guided policy search (MDGPS): convex linear variant

```

1: for iteration  $k \in \{1, \dots, K\}$  do
2:   C-step:  $p_i = \arg \min_{p_i} E_{p_i(\cdot)} \left[ \sum_{t=1}^T \gamma^t \langle \mathbf{x}_t, \mathbf{u}_t \rangle \right]$  such that  $D_{\text{KL}}(p_i(\cdot) \parallel q(\cdot)) \leq \epsilon$ 
3:   S-step:  $q_i = \arg \min_{q_i} \sum_j D_{\text{KL}}(p_i(\cdot) \parallel q_i(\cdot))$ 
4: end for

```

Comparing MD-GPS to EM, the C-step is equivalent to the E-step and S-step equivalent to the M-step, with the loss ℓ being the negative of the expected discounted sum of returns over trajectories. MD-GPS assumes that the S-step can perfectly minimise the objective in the case of linear dynamics and quadratic costs [48], resulting in exact mirror descent. In practice, most applications of MD-GPS will not satisfy such constraints and so are akin to approximate mirror descent. However, as long as a constraint between the local and global policy can be enforced in terms of KL, various bounds on cost of the global policy can be constructed, see Montgomery and Levine [58] for further details.

H Statistical Precipice

H.1 Overview

Advancements in computational power and algorithmic capabilities have led to a shift in the evaluation of reinforcement learning algorithms, now typically assessed through extensive suites of tasks. Performance metrics such as the mean or median score per task are commonly used, but these metrics can fail to account for the statistical uncertainties arising from a limited number of runs and varying random seeds. The trend towards computationally intensive benchmarks further complicates the issue, as each run can span from hours to weeks, making it impractical to conduct numerous runs per task and thereby increasing the uncertainty in the reported metrics. To address these challenges, we adopt the evaluation methodologies proposed by Agarwal et al. [3] and Gorsane et al. [32].

Each algorithm is evaluated across M tasks within a specified environment suite, with N independent runs per task $m \geq M$. During each run $n \geq N$, performance is measured over E episodes, each consisting of T timesteps. At each interval i , the mean return $G_{m,n}^i$ is computed, and the model is checkpointed. The model with the highest mean return across all intervals is selected for final evaluation.

Following the completion of a training run $n \geq N$, the best model is further evaluated over $10 \cdot E$ episodes. We normalise scores $x_{m,n}$ for each task $m = 1; \dots; M$ and run $n = 1; \dots; N$, scaling them based on the minimum and maximum scores observed across all runs. This normalization produces a set of normalised scores $x_{1:M;1:N}$ per algorithm. These scores are then aggregated into a single scalar estimate, x .

To ensure robust statistical confidence, we employ a 95% confidence interval derived from stratified bootstrapping over the $M \cdot N$ experiments, treating these as random samples. This method integrates the performance across all tasks and runs, simulating the statistical reliability of multiple runs on a single task while considering task diversity. Results are reported for the entire suite.

Metrics:

We utilize normalised scores to evaluate algorithm performance, employing metrics that go beyond simple median and mean calculations:

- **Interquartile Mean (IQM):** This metric calculates the mean of the central 50% of runs, excluding the lower and upper 25%. It is more robust to outliers than the mean and less biased than the median, offering higher statistical efficiency and detecting improvements with fewer runs [3].
- **Probability of Improvement:** This measures the likelihood that algorithm X will outperform algorithm Y on a random task m , using the Mann-Whitney U-statistic. It is defined as:

$$Pr(X > Y) = \frac{1}{M} \sum_{m=1}^M Pr(X_m > Y_m) \tag{20}$$

$$Pr(X_m > Y_m) = \frac{1}{NK} \sum_{i=1}^N \sum_{j=1}^K S(x_{m;i}; y_{m;j}) \tag{21}$$

$$S(x; y) = \begin{cases} 1; & \text{if } y < x \\ \frac{1}{2}; & \text{if } y = x \\ 0; & \text{if } y > x \end{cases} \tag{22}$$

Statistical significance is determined using the Neyman-Pearson criterion, based on the confidence interval bounds [13].

Additionally, performance profiles can be employed to visually compare methods, illustrating the fraction of runs that exceed a given threshold. These profiles aid in identifying stochastic dominance and empirical performance bounds. We also plot the interquartile mean score against environment steps to evaluate sample efficiency.

H.2 Hyperparameters

Table 11: Statistical Precipice Evaluation Hyperparameters

Parameter	Value
Number of Games for In-Training Evaluation - E	128
Number of Games for End-of-Training Evaluation - $10 \cdot E$	1280
Seeds per environment - N	5
Tasks per suite - M	Continuous=3, Discrete=2

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: [NA]

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: [NA]

Justification: We do not derive any new theoretical results in this paper

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: [NA]

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: [No]

Justification: Inference code to be released at publication

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: We include error bars for the main results of the paper in addition to statistical significance tests in the appendix. For additional scaling plots we do not include error bars.

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]

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: [No]

Justification: Our work advances the theoretical understanding and practical implementation of model-based RL methods, with no explicit societal impact, while indirectly contributing to future research and technological innovation

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: [NA]

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: [NA]

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: [NA]

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: [NA]

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.