Revisiting the Static Model in Robust Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Designing control policies whose performance level is guaranteed to remain above a given threshold in a span of environments is a critical feature for the adoption of reinforcement learning (RL) in real-world applications. The search for such robust policies is a notoriously difficult problem, often cast as a two-player game, whose formalization dates back to the 1970's. This two-player game is strongly related to the so-called dynamic model of transition function uncertainty, where the environment dynamics are allowed to change at each time step. But in practical applications, one is rather interested in robustness to a span of static transition models throughout interaction episodes. The static model is known to be harder to solve than the dynamic one, and seminal algorithms, such as robust value iteration, as well as most recent works on deep robust RL, build upon the dynamic model. In this work, we propose to revisit the static model. We suggest an analysis of why solving the static model under some mild hypotheses is a reasonable endeavor, and formalize the general intuition that robust MDPs can be solved by tackling a series of static problems. We introduce a generic meta-algorithm called IWOCS, which incrementally identifies worst-case transition models so as to guide the search for a robust policy. Discussion on IWOCS sheds light on new ways to decouple policy optimization and adversarial transition functions and opens new perspectives for analysis. We derive a deep RL version of IWOCS and demonstrate it is competitive with state-of-the-art algorithms on classical benchmarks.

1 Introduction

One major obstacle in the way of real-life deployment of reinforcement learning (RL) algorithms is their inability to produce policies that retain, without further training, a guaranteed level of efficiency when controlling a system that somehow differs from the one they were trained upon. This property is referred to as *robustness*, by opposition to *resilience*, which is the ability to recover, through continued learning, from environmental changes. For example, when learning control policies for aircraft stabilization using a simulator, it is crucial that the learned controller be able to control a span of aircraft configurations with different geometries, or masses, or in various atmospheric conditions. Depending on the criticality of the considered application, one will prefer to optimize the expected performance over a set of environments —thus weighting in the probability of occurrence of a given configuration— or, at the extreme, optimize for the worst case configuration. In the present work, we consider such worst case guarantees and revisit the framework of robust Markov Decision Processes (MDPs) [Iyengar, 2005].

Departing from the common perspective which views robust MDPs as two-player games, we investigate whether it is possible to solve them through a series of non-robust problems. The two-player game formulation is called the dynamic model of transition function uncertainty, as an adversarial

Submitted to 37th Conference on Neural Information Processing Systems (NeurIPS 2023). Do not distribute.

environment is allowed to change the transition dynamics at each time step. The solution to this game can be shown to be equivalent, under mild hypotheses, to that of the static model, where the environment retains the same transition function throughout the time steps. Our first contribution is a series of arguments which cast the search for a robust policy as a resolution of the static model (Section 2). We put this formulation in perspective of recent related works in robust RL (Section 3). Then, we introduce a generic meta-algorithm which we call IWOCS for Incremental Worst-Case Search (Section 4). IWOCS builds upon the idea of incrementally identifying worst case transition functions and expanding a discrete uncertainty set, for which a robust policy can be approximated through a finite set of non-robust value functions. We instantiate two IWOCS algorithms, one on a toy illustrative problem with a discrete state space, then another on popular, continuous states and actions, robust RL benchmarks where it is shown to be competitive with state-of-the art robust deep RL algorithms (Section 5).

2 Problem statement

Reinforcement Learning (RL) [Sutton and Barto, 2018] considers the problem of learning a decision making policy for an agent interacting over multiple time steps with a dynamic environment. At each time step, the agent and environment are described through a state $s \in \mathcal{S}$, and an action $a \in \mathcal{A}$ is performed; then the system transitions to a new state s' according to probability $T(s'|s,a)$, while receiving reward $r(s,a,s')$. The tuple $M_T = (\mathcal{S}, \mathcal{A}, T, r)$ forms a Markov Decision Process (MDP) [Puterman, 2014], which is often complemented with the knowledge of an initial state distribution $p_0(s)$. Without loss of generality and for the sake of readability, we will consider a unique starting state s_0 in this paper, but our results extend straightforwardly to a distribution $p_0(s)$. A stationary decision making policy is a function $\pi(a|s)$ mapping states to distributions over actions (writing $\pi(s)$ the action for the special case of deterministic policies). Training a reinforcement learning agent in MDP M_T consists in finding a policy that maximizes the expected discounted return: $J_T^\pi = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, s_{t+1}) | s_0, a_t \sim \pi, s_{t+1} \sim T] = V_T^\pi(s_0)$, where V_T^π is the value function of π in MDP M_T. An optimal policy in M_T will be noted π_T^* and its value function V_T^*. A convenient notation is the state-action value function $Q_T^\pi(s,a) = \mathbb{E}_{s' \sim T}[r(s,a,s') + \gamma V_T^\pi(s')]$ of policy π in MDP M_T, and the corresponding optimal Q_T^*.

Robust MDPs, as introduced by Iyengar [2005] or Nilim and El Ghaoui [2005], introduce an additional challenge. The transition functions T are picked from an uncertainty set $\mathcal{T}$ and are allowed to change at each time step. A common assumption, called *sa-rectangularity*, states that $\mathcal{T}$ is a Cartesian product of independent marginal sets of distributions on $\mathcal{S}$, for each state-action pair. The value of a stationary policy π in the sequence of MDPs induced by $\mathbf{T} = \{T_t\}_{t \in \mathbb{N}}$ is noted $V_{\mathbf{T}}^\pi$. The pessimistic value function for π is $V^\pi(s) = \min_{\mathbf{T}} V_{\mathbf{T}}^\pi(s)$, where the agent plays a sequence of actions $a_t \in \mathcal{A}$ drawn from π, against the environment, which in turn picks transition models $T_t \in \mathcal{T}$ so as to minimize the overall return. The robust value function is the largest such pessimistic value function and hence the solution to $V^*(s) = \max_\pi \min_{\mathbf{T}} V_{\mathbf{T}}^\pi(s)$. The robust MDP problem can be cast as the zero-sum two-player game, where $\hat{\pi}$ denote the decision making policy of the adversarial environment, deciding $T_t \in \mathcal{T}$ based on previous observations. Then, the problem becomes $\max_\pi \min_{\hat{\pi}} V_{\hat{\pi}}^\pi(s)$, where $V_{\hat{\pi}}^\pi$ is the expected value of a trajectory where policies π and $\hat{\pi}$ play against each other. Hence, the optimal policy becomes the minimax policy, which makes it robust to all possible future evolutions of the environment's properties.

Robust Value Iteration. Following Iyengar [2005, Theorem 3.2], the optimal robust value function $V^*(s) = \max_\pi \min_{\mathbf{T}} V_{\mathbf{T}}^\pi(s)$ is the unique solution to the robust Bellman equation $V(s) = \max_a \min_T \mathbb{E}_{s' \sim T}[r(s,a,s') + \gamma V(s')] = \mathcal{L}V(s)$. This directly translates into a robust value iteration algorithm which constructs the $V_{n+1} = \mathcal{L}V_n$ sequence of value functions [Satia and Lave Jr, 1973, Iyengar, 2005]. Such robust policies are, by design, very conservative, in particular when the uncertainty set is large and under the rectangularity assumption. Several attempts at mitigating this over-conservativeness have been made from various perspectives. For instance, Lim et al. [2013] propose to learn and tighten the uncertainty set, echoing other works that incorporate knowledge about this set into the minimax resolution [Xu and Mannor, 2010, Mannor et al., 2012]. Other approaches [Wiesemann et al., 2013, Lecarpentier and Rachelson, 2019, Goyal and Grand-Clement, 2022] propose to lift the rectangularity assumption and capture correlations in uncertainties across states or time steps, yielding significantly less conservative policies. Ho et al. [2018] and

Grand-Clément and Kroer [2021] retain the rectangularity assumption and propose algorithmic schemes to tackle large but discrete state and action spaces.

The static model. In many applications, one does not wish to consider non-stationary transition functions, but rather to be robust to any transition function from $\mathcal{T}$ which remains stationary throughout a trajectory. This is called the *static* model of transition function uncertainty, by opposition to the *dynamic* model where transition functions can change at each time step. Hence, the static model's minimax game boils down to $\max_\pi \min_T V_T^\pi(s)$. If the agent is restricted to stationary policies $\pi(a|s)$, then $\max_\pi \min_{\mathbf{T}} V_{\mathbf{T}}^\pi(s) = \max_\pi \min_T V_T^\pi(s)$ [Iyengar, 2005, Lemma 3.3], that is the static and dynamic problems are equivalent, and the solution to the dynamic problem is found for a static adversary.[1] In this paper, we will only consider stationary policies.

No-duality gap. Wiesemann et al. [2013, Equation 4 and Proposition 9] introduce an important saddle point condition stating that $\max_\pi \min_T V_T^\pi(s) = \min_T \max_\pi V_T^\pi(s)$.

Incrementally solving the static model. Combining the static and dynamic models equivalence and the no-duality gap condition, we obtain that, for rectangular uncertainty sets and stationary policies, the optimal robust value function $V^*(s) = \max_\pi \min_{\mathbf{T}} V_{\mathbf{T}}^\pi(s) = \max_\pi \min_T V_T^\pi(s) = \min_T \max_\pi V_T^\pi(s) = \min_T V_T^*(s)$. The key idea we develop in this paper stems from this formulation. Suppose we are presented with M_{T_0} and solve it to optimality, finding $V_0^*(s) = V_{T_0}^*(s)$. Then, suppose we identify M_{T_1} as a possible better estimate of a worst case MDP in $\mathcal{T}$ than T_1. We can solve for $V_{T_1}^*$ and $V_1^*(s) = \min\{V_{T_0}^*(s), V_{T_1}^*(s)\}$ is the robust value function for the discrete uncertainty set $\mathcal{T}_1 = \{T_0, T_1\}$. The intuition we attempt to capture is that by incrementally identifying candidate worst case MDPs, one should be able to define a sequence of discrete uncertainty sets $\mathcal{T}_i = \{T_j\}_{j\in[0,i]}$ whose robust value function V_i^* may converge to V^*. In other words, it should be possible to incrementally robustify a policy by identifying the appropriate sequence of transition models and solving individually for them, trading the complexity of the dynamic model's resolution for a sequence of classical MDP problems. The algorithm we propose in Section 4 follows this idea and searches for robust stationary policies in the static model, by incrementally growing a finite uncertainty set.

3 Related work

Robust RL as two-player games. A common approach to solving robust RL problems is to cast the dynamic formulation as a zero-sum two player game, as formalized by Morimoto and Doya [2005]. In this framework, an adversary, denoted by $\hat{\pi} : \mathcal{S} \to \mathcal{T}$, is introduced, and the game is formulated as $\max_\pi \min_{\hat{\pi}} \mathbb{E}[\sum_{t=0}^\infty \gamma^t r(s_t, a_t, s_{t+1}) | s_0, a_t \sim \pi(\cdot|s_t), T_t = \hat{\pi}(s_t), s_{t+1} \sim T_t(\cdot|s_t, a_t)]$. Most methods differ in how they constrain $\hat{\pi}$'s action space within the uncertainty set. A first family of methods define $\hat{\pi}(s_t) = T_{ref} + \Delta(s_t)$, where T_{ref} denotes the reference transition function. Among this family, Robust Adversarial Reinforcement Learning (RARL) [Pinto et al., 2017] applies external forces at each time step t to disturb the reference dynamics. For instance, the agent controls a planar monopod robot, while the adversary applies a 2D force on the foot. In noisy action robust MDPs (NR-MDP) [Tessler et al., 2019] the adversary shares the same action space as the agent and disturbs the agent's action $\pi(s)$. Such gradient-based approaches incur the risk of finding stationary points for π and $\hat{\pi}$ which do not correspond to saddle points of the robust MDP problem. To prevent this, Mixed-NE [Kamalaruban et al., 2020] defines mixed strategies and uses stochastic gradient Langevin dynamics. Similarly, Robustness via Adversary Populations (RAP) [Vinitsky et al., 2020] introduces a population of adversaries, compelling the agent to exhibit robustness against a diverse range of potential perturbations rather than a single one, which also helps prevent finding stationary points that are not saddle points. Aside from this first family, State Adversarial MDPs [Zhang et al., 2020, 2021, Stanton et al., 2021] involve adversarial attacks on state observations, which implicitly define a partially observable MDP. The goal in this case is not to address robustness to the worst-case transition function but rather against noisy, adversarial observations. A third family of methods considers the general case of $\hat{\pi}(s_t) = T_t$ where $T_t \in \mathcal{T}$. Minimax Multi-Agent Deep Deterministic Policy Gradient (M3DDPG) [Li et al., 2019] is designed to enhance robustness in multi-agent reinforcement learning settings, but boils down to standard robust RL in the two-agents case. Max-min TD3 (M2TD3) [Tanabe et al., 2022] considers a policy π, defines a value function

[1] This does not imply the solution to the static model is the same as that of the dynamic model in the general case: the optimal static π may be non-stationary and solving the static model is known to be NP-hard.

$Q(s, a, T)$ which approximates $Q^\pi_T(s, a) = \mathbb{E}_{s' \sim T}[r(s, a, s') + \gamma V^\pi_T(s')]$, updates an adversary $\hat{\pi}$ so as to minimize $Q(s, \pi(s), \hat{\pi}(s))$ by taking a gradient step with respect to $\hat{\pi}$'s parameters, and updates the policy π using a TD3 gradient update in the direction maximizing $Q(s, \pi(s), \hat{\pi}(s))$. As such, M2TD3 remains a robust value iteration method which solves the dynamic problem by alternating updates on π and $\hat{\pi}$, but since it approximates Q^π_T, it is also closely related to the method we introduce in the next section.

Regularization. Derman et al. [2021], Eysenbach and Levine [2022] also highlighted the strong link between robust MDPs and regularized MDPs, showing that a regularized policy learned during interaction with a given MDP was actually robust to an uncertainty set around this MDP. Kumar et al. [2023] propose a promising approach in which they derive the adversarial transition function in a closed form and demonstrate that it is a rank-one perturbation of the reference transition function. This simplification results in more streamlined computation for the robust policy gradient.

Domain randomization (DR) [Tobin et al., 2017] learns a value function $V(s) = \max_\pi \mathbb{E}_{T \sim \mathcal{U}(\mathcal{T})} V^\pi_T(s)$ which maximizes the expected return *on average* across a fixed (generally uniform) distribution on $\mathcal{T}$. As such, DR approaches do not optimize the worst-case performance. Nonetheless, DR has been used convincingly in applications [Mehta et al., 2020, OpenAI et al., 2019]. Similar approaches also aim to refine a base DR policy for application to a sequence of real-world cases [Lin et al., 2020, Dennis et al., 2020, Yu et al., 2018].

For a more complete survey of recent works in robust RL, we refer the reader to the work of Moos et al. [2022]. To the best of our knowledge, the approach sketched in the previous section and developed in the next one is the only one that directly addresses the static model. For that purpose, it exploits the equivalence with the dynamic model for stationary policies and solves the dual of the minimax problem, owing to the no-duality gap property.

4 Incremental Worst-case Search

In order to search for robust policies, we consider the no-duality gap property: the best performance one can expect in the face of transition function uncertainty $\max_\pi \min_T V^\pi_T(s_0)$, is also the worst performance the environment can induce for each transition function's optimal policy $\min_T V^*_T(s_0)$. If the value $V^\pi_T(s_0)$ was strictly concave / convex with respect to π / T respectively, we could hope to solve for the robust policy through a (sub)gradient ascent / descent method. Unfortunately, it seems $V^\pi_T(s_0)$ easily admits more convoluted optimization landscapes, involving stationary points, and local minima and maxima. The $\max_\pi$ problem often benefits from regularization [Geist et al., 2019]. Although one could study regularization for the $\min_T$ problem [Grand-Clément and Petrik, 2022] or the equivalence with a regularized objective [Derman et al., 2021], we turn towards a simpler process conceptually. We consider a (small) discrete set of MDPs $\mathcal{T}_i = \{T_j\}_{j \in [0,i]}$, for which we derive the corresponding optimal value functions $Q^*_{T_j}$. Notably, $\mathcal{T}_i$ is a subset of a (supposed) sa-rectangular uncertainty set, but is not sa-rectangular itself, so there is no guarantee that the static-dynamic equivalence holds in $\mathcal{T}_i$. However, one can consider the sa-rectangular set $\tilde{\mathcal{T}}_i = \times_{s,a} \{T_j(\cdot|s, a)\}_{j \in [0,i]}$ composed of the cartesian product of all local $\{T_j(\cdot|s, a)\}_{j \in [0,i]}$ sets for each s, a pair. Then we implicitly define Q_i as the function that maps any pair s, a to the smallest expected optimal outcome $Q_i(s, a) = \min_{j \in [0,i]} \{Q^*_{T_j}(s, a)\}$. Consequently, Q_i is an upper-bound on the robust Q-function with respect to the $\tilde{\mathcal{T}}_i$ uncertainty set. We abusively call Q_i the $\mathcal{T}_i$-robust value function. In s_0, Q_i coincides with the robust value function for the static model of uncertainty with respect to the $\mathcal{T}_i$ uncertainty set. The corresponding greedy policy is $\pi_i(a) = \arg\max_a Q_i(s, a)$ and is a candidate for the robust policy. Let us write $T_{i+1} = \arg\min_{T \in \mathcal{T}} V^{\pi_i}_T(s_0)$. Then, if $V^{\pi_i}_{T_{i+1}}(s_0) = Q_i(s_0, \pi_i(s_0))$, we have found a robust policy for all transition models in $\mathcal{T}$. Otherwise, we can solve for $Q^*_{T_{i+1}}$, append T_{i+1} to $\mathcal{T}_i$ to form $\mathcal{T}_{i+1}$ and repeat. Of course, one bottleneck difficulty lies in solving the $\min_T$ problem accurately enough, but at least this approach decouples this difficulty from the policy optimization process, which is only concerned with static MDPs. Consequently, the idea we develop is to incrementally expand $\mathcal{T}_i$ by solving $\min_{T \in \mathcal{T}} V^{\pi_i}_T(s_0)$ using optimization methods that can cope with ill-conditioned optimization landscapes. We call Incremental Worst Case Search (IWOCS) this general meta-algorithm, which we summarize in Algorithm 1. We illustrate IWOCS in the current section on a toy example, while Section 4 extends it to the case of high-dimensional state and action spaces, which involve a need for function approximation, and permit only collecting samples in a sub-part of the state space.

Algorithm 1: Incremental Worst-Case Search meta-algorithm

Input: $\mathcal{T}$, T_0, max number of iterations M, tolerance on robust value ϵ

for $i = 0$ *to* M **do**

  Find $Q^*_{T_i} = \max_\pi Q^\pi_{T_i}$ `/* Non-robust policy optimization */`

  Define $\mathcal{T}_i = \{T_j\}_{j\in[0,i]}$

  Define $Q_i : s, a \mapsto \min_{j\in[0,i]}\{Q^*_{T_j}(s,a)\}$ `/*` $\mathcal{T}_i$`-robust value function */`

  Define $\pi_i(s) = \arg\max_a(Q_i(s,a))$ `/* Candidate policy */`

  Find $T_{i+1} = \arg\min_{T\in\mathcal{T}} V^{\pi_i}_T(s_0)$ `/* Identify worst` T `*/`

  if $|V^{\pi_i}_{T_{i+1}}(s_0) - Q_i(s_0, \pi_i(s_0)| \leqslant \epsilon$ **then**

    return π_i, T_{i+1}, $V^{\pi_i}_{T_{i+1}}(s_0)$ `/* Early termination condition */`

return π_M, T_{M+1}, $V^{\pi_M}_{T_{M+1}}(s_0)$

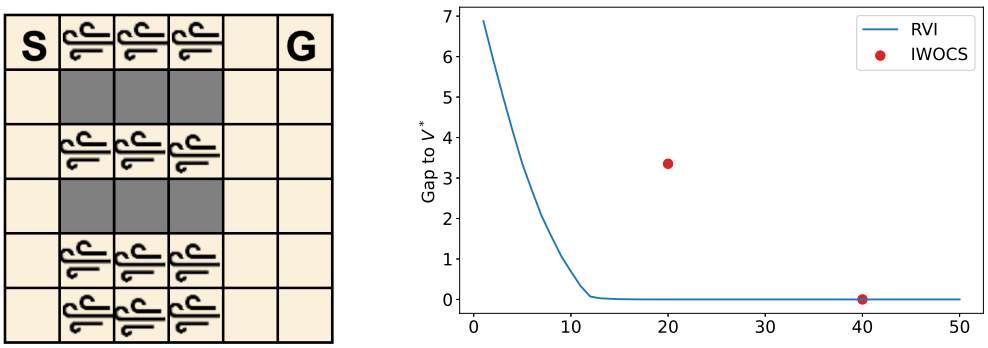

(a) Windy walk grid-world.

(b) Convergence to V^* versus Bellman iterates.

Figure 1: IWOCS and RVI.

We implement an IWOCS algorithm for robust MDPs, using value iteration (VI) as the policy optimization algorithm and a brute force search across transition functions to identify worst-case MDPs ($V^{\pi_i}_T(s_0)$ is evaluated through Monte-Carlo rollouts). Detailed pseudo-code is provided in Appendix C. Note that any policy optimization method, or optimization method for T, is a suitable fit here. The goal here is to illustrate the behavior of IWOCS, compare it to the seminal robust value iteration (RVI) algorithm, and validate empirically that IWOCS is able to find worst-case static MDPs and robust policies.

This vanilla IWOCS is evaluated on the "windy walk" grid-world MDP illustrated on Figure 1a, where an agent wishes to navigate from a starting position S to a goal G. Actions belong to the discrete $\{N, S, E, W\}$ set and transitions are deterministic, except in the "corridors", where wind can knock back the agent to the left. In the topmost corridor, the probability of being knocked left is α, in the middle corridor it is α^3 and it is α^6 in the bottom corridor (details in Appendix C). Hence, the uncertainty set is fully parameterized by α, which takes 25 discrete values, uniformly distributed in $[0, 0.5]$. Rewards are -1 at each time step and the goal is an absorbing state yielding zero reward.

Figure 1b illustrates how IWOCS converges to the robust value function V^*. RVI builds the sequence $V_{n+1} = \mathcal{L}V_n$ and we plot $|V_n(s_0) - V^*(s_0)|$ versus the number of robust Bellman iterates n. On the other hand, IWOCS requires its first policy optimization to terminate before it can report its first $Q_i(s_0, \pi_i(s_0))$. Thus, we plot $|Q_i(s_0, \pi_i(s_0)) - V^*(s_0)|$ after a fixed number of 100 standard Bellman backups for VI. It is important to note that one iterate of the standard Bellman operator requires solving a $\max_a$ in each state, while an iterate of the robust Bellman operator requires solving a more costly $\max_a \min_T$ problem in each state. Therefore, the x-axis does not account for computational time. IWOCS finds the worst-case static model after two iterations and converges to the same value as RVI.

5 Deep IWOCS

We now turn towards challenging robust control problems and introduce an instance of IWOCS meant to accommodate large and continuous state and action spaces, using function approximators such as neural networks. This instance of IWOCS uses Soft Actor Critic (SAC) Haarnoja et al. [2018] as the policy optimization method. One motivation for using SAC is that it has been proven to yield a locally robust policy around the MDP it is trained upon [Eysenbach and Levine, 2022].

5.1 Method

Accounting for regularization terms. Since SAC learns a regularized Q-function which accounts for the policy's entropy, and lets the importance of this term vary along the optimization, orders of magnitude may change between Q_{T_i} and Q_{T_j}. To avoid the influence of this regularization term when defining the $\mathcal{T}_i$-robust Q-function, we train an additional unregularized Q-network which only takes rewards into account. We call π_T the policy network which approximates the optimal policy of the regularized MDP based on T. This policy's (regularized) value function is approximated by the Q'_T network (our implementation uses double Q-networks as per the common practice — all details in Appendix A.2), while an additional Q_T network (double Q-network also) tracks the unregularized value function of π_T. The $\mathcal{T}_i$-robust Q-function is defined with respect to this unregularized value function as $Q_i(s,a) = \min_{j\in[0,i]}\{Q_{T_j}(s,a)\}$.

Partial state space coverage. In large state space MDPs, it is likely that interactions will not explore the full state space. Consequently, the different Q_{T_j} functions are trained on replay buffers whose empirical distribution's support my vary greatly. Evaluating neural networks outside of their training distribution is prone to generalization errors. This begs for indicator functions specifying on which (s,a) pair each Q_T is relevant. We chose to implement such an indicator function using predictive coding [Rao and Ballard, 1999] on the dynamical model T_j. Note that other choices can be equally good (or better), such as variance networks [Neklyudov et al., 2019], ensembles of neural networks [Lakshminarayanan et al., 2016] or 1-class classification [Béthune et al., 2023]. Our predictive coding model for T_j predicts $\hat{T}_j(s,a) = s'$ for deterministic dynamics, by minimizing the expected value of the loss $\ell(\hat{T}_j(s,a); s') = \|\hat{T}_j(s,a) - s'\|_1$. At inference time, along a trajectory, we consider Q_{T_j} has been trained on sufficient data in s_t, a_t, if $\ell(\hat{T}_j(s_{t-1}, a_{t-1}); s_t) \leqslant \rho_j$, ie. if the prediction error for s_t is below the threshold ρ_j (details about tuning ρ_j in Appendix A.4). We redefine $Q^*_{T_j}$ to be $+\infty$ in all states where $\ell(\hat{T}_j(s_{t-1}, a_{t-1}); s_t) > \rho_j$, so that it does not participate in the definition of Q_i.

Worst case identification. When $V^{\pi_i}_T(s_0)$ is non differentiable with respect to T (or T's parameters), one needs to fall back on black-box optimization to find $T_{i+1} = \arg\min_{T\in\mathcal{T}} V^{\pi_i}_T(s_0)$. We turn to evolutionary strategies, and in particular CMA-ES [Hansen and Ostermeier, 2001] for that purpose, for its ability to escape local minima and efficiently explore the uncertainty set $\mathcal{T}$ (hyperparameters in Appendix A). Note that making $V^{\pi}_T(s_0)$ differentiable with respect to T is feasible by making the critic network explicitly depend on T's parameters, as in the work of Tanabe et al. [2022]. We do not resort to such a model, as it induces the risk for generalization errors, but it constitutes a promising alternative for research. To evaluate $V^{\pi_i}_T(s_0)$ for a given T, we run a roll-out from s_0 by applying $\pi_i(s)$ in each encountered state s. Since we consider continuous action spaces and keep track of the critics Q_{T_j}, Q'_{T_j} and the actor π_{T_j} for all $T_j \in \mathcal{T}_i$, we can make direct use of π_{T_j} which is designed to mimic an optimal policy in M_{T_j}. Specifically, in s, we evaluate $j^* = \arg\min_{j\leqslant i} Q^*_{T_j}(s, \pi_{T_j}(s))$, and apply $\pi_i(s) = \pi_{j*}(s)$. If no $Q^*_{T_j}$ is valid in s, we fall back to a default policy trained with domain randomization.

5.2 Empirical evaluation

Experimental framework. This section assesses the proposed algorithm's worst-case performance and generalization capabilities. Experimental validation was performed on optimal control problems using the MuJoCo simulation environments [Todorov et al., 2012]. The algorithm was benchmarked against state-of-the-art robust reinforcement learning methods, including M2TD3 [Tanabe et al., 2022], M3DDPG [Li et al., 2019], and RARL [Pinto et al., 2017]. We also compare with Domain Randomization (DR) [Tobin et al., 2017] for completeness. For each environment, two versions of the uncertainty set are considered, following the benchmarks reported by Tanabe et al. [2022]. In

the first one, T is parameterized by a global friction coefficient and the agent's mass, making $\mathcal{T}$ isomorphic to $\mathbb{R}^2$. In the second one, a third, environment-dependent parameter is included (details in Appendix A.5). To ensure a fair comparison, performance metrics were collected after 4 million steps for environments with a 2D uncertainty set and after 5 million steps for those with a 3D uncertainty set. All reference methods optimize a single policy along these 4 or 5 million steps, but IWOCS optimizes a sequence of non-robust policies, for which we divided this sample budget: we constrained IWOCS to train its default policy and each subsequent SAC agent for a fixed number of interaction steps, so that the sum is 4 or 5 million steps. All results reported below were averaged over 10 distinct random seeds.

Worst-case performance. Table 1 reports the normalized worst-case scores comparing IWOCS, M2TD3, SoftM2TD3, M3DDPG, RARL, and DR using TD3. The worst-case score for the policy returned by IWOCS is a direct by-product of the algorithm's last steps, since they search for worst-case transition functions that minimize the value of Monte-Carlo rollouts. The worst-case scores for all other methods are drawn from the M2TD3 paper [Tanabe et al., 2022] where they are evaluated by defining a uniform grid over the transition function's parameter space and performing roll-outs for each transition model[2]. To obtain comparable metrics across environments, we normalized each method's score v using the vanilla TD3 (trained on the environment with default transition function parameters) reference score v_{TD3} as a minimal baseline and the M2TD3 score v_{M2TD3} as target score: $(v - v_{TD3})/|v_{M2TD3} - v_{TD3}|$. Hence this metric reports how much a method improves upon TD3, compared to how much M2TD3 improved upon TD3. The non-normalized scores are reported in Table 7 in the Appendix. This instance of IWOCS demonstrates competitive performance, outperforming all other methods in 7 out of the 11 environments (note that we did not report results on environments from the literature that feature a simpler 1D uncertainty set). IWOCS permits a 2.36-fold improvement on average across environments, over the state-of-the-art M2TD3. It seems that Ant is an environment where IWOCS struggles to reach convincing worst case scores (and Humanoid to a lesser extent). We conjecture this is due to the wide range of possible mass and friction parameters, which make the optimization process very noisy (almost zero mass and friction is a worst-case T making the ant's movement rather chaotic and hence induces a possibly misleading replay buffer) and may prevent the policy optimization algorithm to yield good non-robust policies and value functions given its sample budget. However, IWOCS provides a major (up to 10-fold) improvement on Hopper.

Table 1: Avg. of normalized worst-case performance over 10 seeds for each method.

Environment	IWOCS	M2TD3	SoftM2TD3	M3DDPG	RARL	DR (TD3)
Ant 2	-0.19 ± 0.87	$\mathbf{1.00 \pm 0.04}$	0.92 ± 0.06	-0.72 ± 0.05	-1.32 ± 0.04	0.02 ± 0.05
Ant 3	0.45 ± 0.59	$\mathbf{1.00 \pm 0.09}$	0.97 ± 0.18	-0.36 ± 0.20	-1.28 ± 0.06	0.61 ± 0.03
HalfCheetah 2	$\mathbf{1.35 \pm 0.28}$	1.00 ± 0.05	1.07 ± 0.05	-0.02 ± 0.02	-0.05 ± 0.02	0.84 ± 0.04
HalfCheetah 3	$\mathbf{1.84 \pm 0.56}$	1.00 ± 0.14	1.39 ± 0.15	-0.03 ± 0.05	-0.13 ± 0.05	1.10 ± 0.04
Hopper 2	$\mathbf{5.56 \pm 0.74}$	1.00 ± 0.05	1.09 ± 0.06	0.46 ± 0.06	0.61 ± 0.17	0.87 ± 0.03
Hopper 3	$\mathbf{9.32 \pm 0.99}$	1.00 ± 0.09	0.68 ± 0.08	0.22 ± 0.04	0.56 ± 0.17	0.73 ± 0.13
HumanoidStandup 2	0.82 ± 0.18	1.00 ± 0.12	$\mathbf{1.25 \pm 0.16}$	0.98 ± 0.12	0.88 ± 0.13	1.14 ± 0.14
HumanoidStandup 3	0.90 ± 0.63	$\mathbf{1.00 \pm 0.11}$	0.96 ± 0.07	0.97 ± 0.07	0.88 ± 0.13	0.86 ± 0.06
InvertedPendulum 2	$\mathbf{2.82 \pm 0.00}$	1.00 ± 0.37	0.38 ± 0.08	-0.00 ± 0.00	-0.00 ± 0.00	0.15 ± 0.01
Walker 2	$\mathbf{1.24 \pm 0.34}$	1.00 ± 0.14	0.83 ± 0.15	0.04 ± 0.04	-0.08 ± 0.01	0.71 ± 0.17
Walker 3	$\mathbf{1.87 \pm 0.50}$	1.00 ± 0.23	1.03 ± 0.20	0.06 ± 0.05	-0.10 ± 0.01	0.65 ± 0.19
Aggregated	$\mathbf{2.36 \pm 0.51}$	1.00 ± 0.13	0.96 ± 0.11	0.14 ± 0.06	0.0 ± 0.07	0.70 ± 0.08

Average performance. While our primary aim is to maximize the worst-case performance, we also appreciate the significance of average performance in real-world scenarios. Table 2 reports the normalized average score (non-normalized scores in Appendix 8) obtained by the resulting policy over a uniform grid of 100 transition functions in 2D uncertainty sets (1000 in 3D ones). Interestingly, M3DDPG and RARL feature negative normalized scores and perform worse on average than vanilla TD3 on most environments (as M2TD3 on 3 environments). DR and IWOCS display the highest average performance. Although this outcome was anticipated for DR, it may initially seem surprising for IWOCS, which was not explicitly designed to optimize mean performance. We posit this might be attributed to two factors. First, in MDPs which have not been identified as worst-cases, encountered

[2]The same grid was used as a sanity check for IWOCS to verify that CMA-ES indeed found the worst-case transition function.

Table 2: Avg. of normalized average performance over 10 seeds for each method.

Environment	IWOCS	M2TD3	SoftM2TD3	M3DDPG	RARL	DR (TD3)
Ant 2	0.45 ± 0.51	1.00 ± 0.02	1.04 ± 0.00	-0.13 ± 0.12	-1.04 ± 0.02	$\mathbf{1.28 \pm 0.03}$
Ant 3	-3.56 ± 2.96	-1.00 ± 0.44	-0.36 ± 0.46	-6.98 ± 0.44	-8.94 ± 0.18	$\mathbf{0.92 \pm 0.22}$
HalfCheetah 2	1.74 ± 0.26	1.00 ± 0.03	1.10 ± 0.04	-1.08 ± 0.07	-1.94 ± 0.03	$\mathbf{1.84 \pm 0.09}$
HalfCheetah 3	1.90 ± 0.32	1.00 ± 0.07	1.17 ± 0.03	-1.43 ± 0.14	-2.48 ± 0.05	$\mathbf{2.33 \pm 0.12}$
Hopper 2	$\mathbf{1.90 \pm 0.80}$	1.00 ± 0.07	0.74 ± 0.12	-0.40 ± 0.11	1.86 ± 0.92	0.36 ± 0.08
Hopper 3	$\mathbf{6.17 \pm 3.30}$	-1.00 ± 0.23	-1.20 ± 0.13	-2.20 ± 0.37	1.63 ± 1.53	1.17 ± 0.23
HumanoidStandup 2	$\mathbf{3.00 \pm 3.00}$	-1.00 ± 0.67	0.67 ± 0.83	-1.83 ± 0.67	-3.00 ± 1.33	0.50 ± 0.67
HumanoidStandup 3	$\mathbf{1.87 \pm 4.12}$	1.00 ± 0.75	0.38 ± 0.37	0.00 ± 0.50	-1.75 ± 0.87	0.38 ± 0.87
InvertedPendulum 2	$\mathbf{2.86 \pm 0.00}$	1.00 ± 0.68	1.06 ± 0.46	-1.10 ± 0.25	-0.47 ± 0.32	2.47 ± 0.03
Walker 2	$\mathbf{1.03} \pm 0.45$	1.00 ± 0.06	0.83 ± 0.16	-0.53 ± 0.11	-1.21 ± 0.02	0.91 ± 0.15
Walker 3	$\mathbf{1.19 \pm 0.50}$	1.00 ± 0.13	0.96 ± 0.18	-0.57 ± 0.09	-1.43 ± 0.04	1.13 ± 0.10
Aggregated	$\mathbf{1.68 \pm 1.47}$	0.45 ± 0.28	0.58 ± 0.25	-1.47 ± 0.25	-1.70 ± 0.48	1.20 ± 0.23

states are likely to have no valid Q_{T_j} value function. In these MDPs, if we were to apply any of the π_{T_j}, its score could be as low as the worst cast value (but not lower, otherwise the MDP should have been identified as a worst case earlier). But since IWOCS' indicator functions identify these states as unvisited, the applied policy falls back to the DR policy, possibly providing a slightly better score above the worst case value for these MDPs. Second, the usage of indicator functions permits defining the IWOCS policy as an aggregate of locally optimized policies, possibly avoiding averaging issues. As for the worst-case scores, IWOCS does not perform well on Ant environments. However, it provides better average scores than both DR and M2TD3 on the Humanoid benchmarks. Overall, it performs better than all other methods on 7 out of the 11 environments and is only slightly outperformed by DR on 2 more.

Tracking the convergence of IWOCS. IWOCS aims at solving iteratively the robust optimization problem by covering the worst possible case at each iteration. Consequently, we could expect the value of the candidate robust policy π_i to increase throughout iterations. Figure 2 reports the score of π_i across iterations, averaged over 10 optimization runs, along with the observed standard deviation, on all 3D uncertainty sets. In the three environments where IWOCS dominates over other reference methods, we observe a nice steady increase in value, with limited standard deviations across runs. Notably, in the Hopper environment, the search for a worst-case environment seems very beneficial to IWOCS, enabling a large leap in score over the first iteration (causing the 10-fold increase reported in Table 1). Results on the Ant benchmark seem to confirm the conjecture made to explain the poor behavior of IWOCS: the large variance in π_i scores tend to indicate that optimizing for non-robust policies is somewhat noisy and might prevent finding robust policies overall.

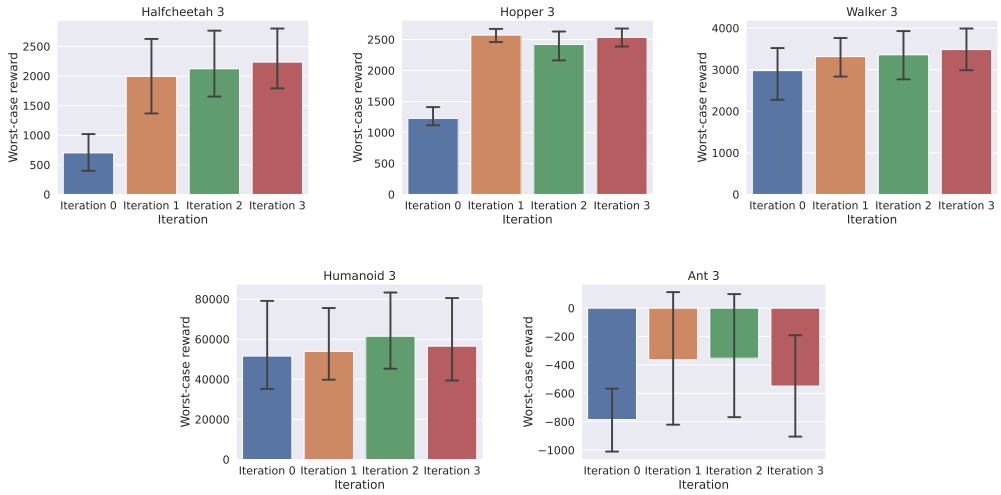

Figure 2: Worst-case performance for each iteration of IWOCS on 3D uncertainty sets

Table 3: Halfcheetah 2, worst parameters search for each iteration over 10 seeds.

	$J_{T_0}^{\pi_{T_0}}$	ψ_1^1	ψ_2^1	$J_{T_1}^{\pi_0}$	$J_{T_1}^{\pi_{T_1}}$	ψ_1^2	ψ_2^2	$J_{T_2}^{\pi_1}$	$J_{T_2}^{\pi_{T_2}}$	ψ_1^3	ψ_2^3	$J_{T_3}^{\pi_2}$
0	5753.2	4.0	7.0	2128.3	20559.5	4.0	0.1	3626.3	9724.7	4.0	7.0	3745.6
1	109.755	4.0	7.0	2027.0	3732.0	4.0	7.0	3827.6	6183.4	4.0	7.0	3827.6
2	5897.4	0.1	7.0	2354.0	9930.0	4.0	7.0	2497.4	3670.5	4.0	7.0	3874.4
3	6396.9	4.0	7.0	2473.0	3815.1	4.0	0.6	3053.9	6022.4	4.0	7.0	3825.5
4	2143.26	0.1	7.0	1633.3	3807.9	4.0	7.0	1978.4	6282.5	4.0	7.0	1978.4
5	4476.0	4.0	7.0	1983.2	13040.2	4.0	0.1	2597.4	3694.9	4.0	0.1	2822.3
6	6275.2	4.0	7.0	-91.1	3639.7	4.0	0.1	3448.4	5347.6	4.0	7.0	3986.0
7	6213.8	0.1	7.0	1587.1	2046.31	4.0	7.0	2411.1	3757.7	4.0	7.0	3769.5
8	6409.8	4.0	7.0	2053.8	3709.9	4.0	0.1	3182.3	3712.3	4.0	7.0	3872.9
9	6801.4	0.1	0.1	2341.4	3620.1	0.5	4.4	3619.3	5737.3	0.6	6.8	5431.3
avg	5047.6			1849.0	6790.1			3024.2	5413.3			3713.4
std	2210.6			740.5	5940.3			623.4	1886.8			876.3

Worst case paths in the uncertainty set. IWOCS seems to reliably find worst case MDPs and policies in a number of cases. Table 3 permits tracking the worst case MDPs and policies along iterations for the HalfCheetah 2D environment (all other results in Appendix D). Specifically, it reports the score of the first optimized policy $J_{T_0}^{\pi_{T_0}}$, then the parameters ψ_1^1 and ψ_2^1 (friction and mass) of the worst case transition function T_1, and the value $J_{T_1}^{\pi_0}$ of the candidate robust policy π_0 in T_1. The next columns repeats these values for later iterations. Each line corresponds to a different random seed. At iteration 0, the first trained policy π_{T_0} displays a large variance in score across runs, as is often the case with deep RL methods (SAC included). In turn, it is no surprise that the identified worst case MDPs differ from one line to the other and that their scores features some variability either. However, all instances of T_1 tend to lie on some corner of the uncertainty set. As IWOCS move on to iteration 1, then 2, worst-case MDPs tend to concentrate on the same corner (the true worst-case MDP) and robust scores tend to increase (as Figure 2 showed in the 3D case).

6 Conclusion

The search for robust policies in uncertain MDPs is a long-standing challenge. In this work, we proposed to revisit the static model of transition function uncertainty, which is equivalent to the dynamic model in the case of sa-rectangular uncertainty sets and stationary policies. We proposed to exploit this equivalence and the no-duality-gap property to design an algorithm that trades the resolution of a two-player game, for a sequence of one-player MDP resolutions. This led to the IWOCS (Incremental Worst-Case Search) meta-algorithm, which incrementally builds a discrete, non-sa-rectangular uncertainty set and a sequence of candidate robust policies. An instance of IWOCS, using SAC for policy optimization, and CMA-ES for worst-case search, appeared as a relevant method on popular robust RL benchmarks, and outperformed the state-of-the-art algorithms on a number of environments. IWOCS proposes a new perspective on the resolution of robust MDPs and robust RL problems, which appears as a competitive formulation with respect to traditional methods. It also poses new questions, like the tradeoffs between policy optimization precision and overall robustness, gradient-based methods for worst-case search, bounds due to approximate value functions, or validity of using Q_i as a surrogate of the robust value function for the $\mathcal{T}_i$ uncertainty set. All these open interesting avenues for future research.

References

Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2): 257–280, 2005.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

Jay K Satia and Roy E Lave Jr. Markovian decision processes with uncertain transition probabilities. *Operations Research*, 21(3):728–740, 1973.

Shiau Hong Lim, Huan Xu, and Shie Mannor. Reinforcement learning in robust markov decision processes. *Advances in Neural Information Processing Systems*, 26, 2013.

Huan Xu and Shie Mannor. Distributionally robust markov decision processes. *Advances in Neural Information Processing Systems*, 23, 2010.

Shie Mannor, Ofir Mebel, and Huan Xu. Lightning does not strike twice: robust mdps with coupled uncertainty. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, pages 451–458, 2012.

Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.

Erwan Lecarpentier and Emmanuel Rachelson. Non-stationary markov decision processes, a worst-case approach using model-based reinforcement learning. *Advances in neural information processing systems*, 32, 2019.

Vineet Goyal and Julien Grand-Clement. Robust markov decision processes: Beyond rectangularity. *Mathematics of Operations Research*, 2022.

Chin Pang Ho, Marek Petrik, and Wolfram Wiesemann. Fast bellman updates for robust mdps. In *International Conference on Machine Learning*, pages 1979–1988. PMLR, 2018.

Julien Grand-Clément and Christian Kroer. Scalable first-order methods for robust mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12086–12094, 2021.

Jun Morimoto and Kenji Doya. Robust reinforcement learning. *Neural computation*, 17(2):335–359, 2005.

Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pages 2817–2826. PMLR, 2017.

Chen Tessler, Yonathan Efroni, and Shie Mannor. Action robust reinforcement learning and applications in continuous control. In *International Conference on Machine Learning*, pages 6215–6224. PMLR, 2019.

Parameswaran Kamalaruban, Yu ting Huang, Ya-Ping Hsieh, Paul Rolland, C. Shi, and V. Cevher. Robust reinforcement learning via adversarial training with langevin dynamics. *Neural Information Processing Systems*, 2020.

Eugene Vinitsky, Yuqing Du, Kanaad Parvate, Kathy Jang, Pieter Abbeel, and Alexandre Bayen. Robust reinforcement learning using adversarial populations. *arXiv preprint arXiv:2008.01825*, 2020.

Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Mingyan Liu, Duane Boning, and Cho-Jui Hsieh. Robust deep reinforcement learning against adversarial perturbations on state observations. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21024–21037, 2020.

Huan Zhang, Hongge Chen, Duane S Boning, and Cho-Jui Hsieh. Robust reinforcement learning on state observations with learned optimal adversary. In *International Conference on Learning Representations*, 2021.

Samuel Stanton, Rasool Fakoor, Jonas Mueller, Andrew Gordon Wilson, and Alex Smola. Robust reinforcement learning for shifting dynamics during deployment. In *Workshop on Safe and Robust Control of Uncertain Systems at NeurIPS*, 2021.

Shihui Li, Yi Wu, Xinyue Cui, Honghua Dong, Fei Fang, and Stuart Russell. Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4213–4220, 2019.

Takumi Tanabe, Rei Sato, Kazuto Fukuchi, Jun Sakuma, and Youhei Akimoto. Max-min off-policy actor-critic method focusing on worst-case robustness to model misspecification. In *Advances in Neural Information Processing Systems*, 2022.

Esther Derman, Matthieu Geist, and Shie Mannor. Twice regularized mdps and the equivalence between robustness and regularization. *Advances in Neural Information Processing Systems*, 34, 2021.

Benjamin Eysenbach and Sergey Levine. Maximum entropy rl (provably) solves some robust rl problems. In *International Conference on Learning Representations*, 2022.

Navdeep Kumar, Esther Derman, Matthieu Geist, Kfir Levy, and Shie Mannor. Policy gradient for s-rectangular robust markov decision processes. *arXiv preprint arXiv:2301.13589*, 2023.

Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017.

Bhairav Mehta, Manfred Diaz, Florian Golemo, Christopher J. Pal, and Liam Paull. Active domain randomization. In *Proceedings of the Conference on Robot Learning*, volume 100, pages 1162–1176, 2020.

OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik's cube with a robot hand. *arXiv preprint arXiv: Arxiv-1910.07113*, 2019.

Zichuan Lin, Garrett Thomas, Guangwen Yang, and Tengyu Ma. Model-based adversarial meta-reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 10161–10173, 2020.

Michael Dennis, Natasha Jaques, Eugene Vinitsky, A. Bayen, Stuart J. Russell, Andrew Critch, and S. Levine. Emergent complexity and zero-shot transfer via unsupervised environment design. *Neural Information Processing Systems*, 2020.

Wenhao Yu, C. K. Liu, and Greg Turk. Policy transfer with strategy optimization. *International Conference On Learning Representations*, 2018.

Janosch Moos, Kay Hansel, Hany Abdulsamad, Svenja Stark, Debora Clever, and Jan Peters. Robust reinforcement learning: A review of foundations and recent advances. *Machine Learning and Knowledge Extraction*, 4(1):276–315, 2022.

Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR, 2019.

Julien Grand-Clément and Marek Petrik. On the convex formulations of robust markov decision processes. *arXiv preprint arXiv:2209.10187*, 2022.

Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. *arXiv preprint arXiv: Arxiv-1812.05905*, 2018.

Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.

Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variance networks: When expectation does not meet your expectations. In *International Conference on Learning Representations*, 2019.

B Lakshminarayanan, A Pritzel, and C Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. arxiv. *arXiv preprint arXiv:1612.01474*, 2016.

Louis Béthune, Paul Novello, Thibaut Boissin, Guillaume Coiffier, Mathieu Serrurier, Quentin Vincenot, and Andres Troya-Galvis. Robust one-class classification with signed distance function using 1-lipschitz neural networks. *arXiv preprint arXiv:2303.01978*, 2023.

Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001. doi: 10.1162/106365601750190398.

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.

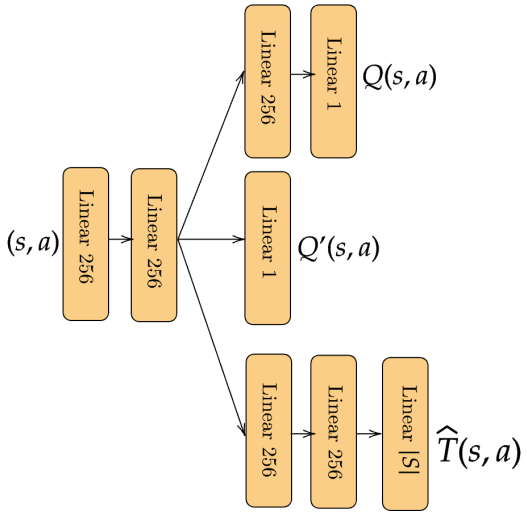

Figure 3: Critic and dynamic model architecture with shared encode

A Hyperparameters

A.1 Soft actor critic

Table 4: Hyperparameter Settings for SAC Experiment

Hyperparameter	**Value**
Learning rate actor	3e-4
Learning rate critic	1e-3
Batch size	256
Memory size	1e6
Gamma	0.99
Polyak update	0.995
Number of steps before training	1e4

A.2 Neural Network

A.3 Computer settings

A.4 Ours algorithm

Table 5: Hyperparameter Settings for 2D Experiment

Hyperparameter	**Value**
Expert train step	800000
Domain randomization train step	1600000
Number of iterations	3
Generations	10
Population size	100
CMA ES mean	0.5
CMA ES std	0.5

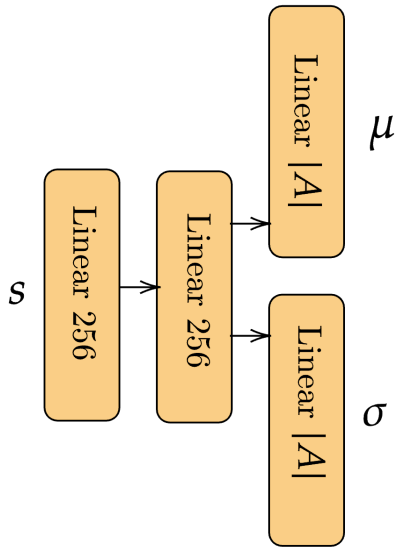

Figure 4: Actor architecture

Table 6: Hyperparameter Settings for 3D Experiment

Hyperparameter	**Value**
Expert train step	800000
Domain randomization train step	1800000
Number of iterations	4
Generations	10
Population size	100
CMA ES mean	0.5
CMA ES std	0.5

A.5 Complementary results

B Uncertainty set for MuJoCo

C Tabular environment

D Worst case path

D.1 Median score

Table 7: Avg. $\pm$ std. error of worst-case performance over 10 seeds for each method

Environment	Ours	M2TD3	SoftM2TD3	M2-DDPG	M3DDPG	RARL (DDPG)	DR (TD3)	TD3 (ref)
Ant 2 $(\times 10^3)$	1.1 ± 2.2	$\mathbf{4.13} \pm 0.11$	3.92 ± 0.14	0.95 ± 0.20	-0.25 ± 0.13	-1.77 ± 0.09	1.64 ± 0.13	1.59 ± 0.08
Ant 3 $(\times 10^3)$	-0.5 ± 0.64	$\mathbf{0.10} \pm 0.10$	0.07 ± 0.20	-1.13 ± 0.28	-1.38 ± 0.22	-2.38 ± 0.07	-0.32 ± 0.03	-0.99 ± 1.13
HalfCheetah 2 $(\times 10^3)$	$\mathbf{3.71} \pm 0.87$	2.61 ± 0.16	2.82 ± 0.16	2.54 ± 0.23	-0.58 ± 0.06	-0.70 ± 0.05	2.12 ± 0.13	-0.53 ± 0.06
HalfCheetah 3 $(\times 10^3)$	$\mathbf{2.23} \pm 0.87$	0.93 ± 0.21	1.53 ± 0.23	1.20 ± 0.22	-0.66 ± 0.08	-0.81 ± 0.07	1.09 ± 0.06	-0.61 ± 0.08
Hopper 2 $(\times 10^2)$	$\mathbf{28.7} \pm 3.78$	5.33 ± 0.28	5.79 ± 0.29	4.30 ± 0.57	2.58 ± 0.29	3.34 ± 0.89	4.68 ± 0.15	0.21 ± 0.04
Hopper 3 $(\times 10^2)$	$\mathbf{25.3} \pm 2.67$	2.84 ± 0.25	1.98 ± 0.22	2.25 ± 0.29	0.73 ± 0.11	1.64 ± 0.46	2.10 ± 0.35	0.14 ± 0.03
HumanoidStandup 2 $(\times 10^4)$	5.46 ± 1.06	6.50 ± 0.70	$\mathbf{7.94} \pm 0.90$	6.24 ± 0.54	6.37 ± 0.72	5.78 ± 0.73	7.31 ± 0.78	0.73 ± 0.07
HumanoidStandup 3 $(\times 10^4)$	5.63 ± 3.55	$\mathbf{6.20} \pm 0.64$	5.99 ± 0.37	5.96 ± 0.58	6.01 ± 0.38	5.54 ± 0.76	5.41 ± 0.34	0.57 ± 0.04
InvertedPendulum 2 $(\times 10^2)$	$\mathbf{10} \pm 0$	3.56 ± 1.32	1.36 ± 0.30	1.10 ± 0.62	0.02 ± 0.00	0.02 ± 0.00	0.57 ± 0.02	0.03 ± 0.00
Walker 2 $(\times 10^3)$	$\mathbf{3.82} \pm 0.97$	3.14 ± 0.39	2.64 ± 0.43	0.85 ± 0.12	0.39 ± 0.11	0.06 ± 0.04	2.31 ± 0.50	0.28 ± 0.07
Walker 3 $(\times 10^3)$	$\mathbf{3.48} \pm 0.88$	1.94 ± 0.40	2.00 ± 0.35	0.82 ± 0.13	0.28 ± 0.09	0.00 ± 0.02	1.32 ± 0.34	$0.17 + 0.06$

Table 8: Avg. $\pm$ std. deviation of average performance over 10 seeds for each method

Environment	IWOCS	M2TD3	SoftM2TD3	M3DDPG	RARL	DR (TD3)	TD3 (ref)
Ant 2 $(\times 10^3)$	3.69 ± 1.60	5.44 ± 0.05	5.56 ± 0.01	1.86 ± 0.38	-1.00 ± 0.06	$\mathbf{6.32} \pm \mathbf{0.09}$	2.28 ± 0.09
Ant 3 $(\times 10^3)$	1.38 ± 1.48	2.66 ± 0.22	2.98 ± 0.23	-0.33 ± 0.22	-1.31 ± 0.09	$\mathbf{3.62} \pm \mathbf{0.11}$	3.16 ± 1.00
HalfCheetah 2 $(\times 10^3)$	5.63 ± 0.44	4.35 ± 0.05	4.52 ± 0.07	0.77 ± 0.12	-0.70 ± 0.05	$\mathbf{5.79} \pm 0.15$	2.63 ± 0.20
HalfCheetah 3 $(\times 10^3)$	4.98 ± 0.42	3.79 ± 0.09	4.02 ± 0.04	0.58 ± 0.18	-0.81 ± 0.07	$\mathbf{5.54} \pm 0.16$	2.47 ± 0.18
Hopper 2 $(\times 10^3)$	$\mathbf{3.38} \pm 0.78$	2.51 ± 0.07	2.26 ± 0.12	1.15 ± 0.11	3.34 ± 0.89	1.89 ± 0.08	1.54 ± 0.17
Hopper 3 $(\times 10^3)$	$\mathbf{3.0} \pm 0.99$	0.85 ± 0.07	0.79 ± 0.04	0.49 ± 0.11	1.64 ± 0.46	1.50 ± 0.07	1.15 ± 0.14
HumanoidStandup 2 $(\times 10^5)$	$\mathbf{1.21} \pm 0.18$	0.97 ± 0.04	1.07 ± 0.05	0.92 ± 0.04	0.85 ± 0.08	1.06 ± 0.04	1.03 ± 0.03
HumanoidStandup 3 $(\times 10^5)$	$\mathbf{1.16} \pm 0.33$	1.09 ± 0.06	1.04 ± 0.03	1.01 ± 0.04	0.87 ± 0.07	1.04 ± 0.07	1.01 ± 0.03
InvertedPendulum 2 $(\times 10^2)$	$\mathbf{10} \pm 0$	6.13 ± 1.42	6.26 ± 0.95	1.76 ± 0.51	3.07 ± 0.66	9.18 ± 0.07	4.05 ± 0.52
Walker 2 $(\times 10^3)$	$\mathbf{4.78} \pm 0.90$	4.72 ± 0.12	4.37 ± 0.32	1.63 ± 0.22	0.26 ± 0.05	4.54 ± 0.31	2.70 ± 0.20
Walker 3 $(\times 10^3)$	$\mathbf{4.58} \pm 0.83$	4.27 ± 0.21	4.21 ± 0.30	1.65 ± 0.15	0.21 ± 0.07	4.48 ± 0.16	2.60 ± 0.18

Table 9: List of environment and parameters for the experiements

Environment	Uncertainty Set $\mathcal{T}$	Reference Parameter	Uncertainty Parameter Name
Ant 2	$[0.1, 3.0] \times [0.01, 3.0]$	$(0.33, 0.04)$	torso mass $\times$ front left leg mass
Ant 3	$[0.1, 3.0] \times [0.01, 3.0] \times [0.01, 3.0]$	$(0.33, 0.04, 0.06)$	torso mass $\times$ front left leg mass $\times$ front right leg mass
HalfCheetah 2	$[0.1, 4.0] \times [0.1, 7.0]$	$(0.4, 6.36)$	world friction $\times$ torso mass
HalfCheetah 3	$[0.1, 4.0] \times [0.1, 7.0] \times [0.1, 3.0]$	$(0.4, 6.36, 1.53)$	world friction $\times$ torso mass $\times$ back thigh mass
Hopper 2	$[0.1, 3.0] \times [0.1, 3.0]$	$(1.00, 3.53)$	world friction $\times$ torso mass
Hopper 3	$[0.1, 3.0] \times [0.1, 3.0] \times [0.1, 4.0]$	$(1.00, 3.53, 3.93)$	world friction $\times$ torso mass $\times$ thigh mass
HumanoidStandup 2	$[0.1, 16.0] \times [0.1, 8.0]$	$(8.32, 1.77)$	torso mass $\times$ right foot mass
HumanoidStandup 3	$[0.1, 16.0] \times [0.1, 5.0] \times [0.1, 8.0]$	$(8.32, 1.77, 4.53)$	torso mass $\times$ right foot mass $\times$ left thigh mass
InvertedPendulum 2	$[1.0, 31.0] \times [1.0, 11.0]$	$(4.90, 9.42)$	pole mass $\times$ cart mass
Walker 2	$[0.1, 4.0] \times [0.1, 5.0]$	$(0.7, 3.53)$	world friction $\times$ torso mass
Walker 3	$[0.1, 4.0] \times [0.1, 5.0] \times [0.1, 6.0]$	$(0.7, 3.53, 3.93)$	world friction $\times$ torso mass $\times$ thigh mass

Table 10: Humanoid 2 worst parameters for each iteration over 10 seeds

Iteration number	1		2		3	
Seed number	torsomass	rightfootmass	torsomass	rightfootmass	torsomass	rightfootmass
0	13.30	2.66	12.06	4.79	11.71	6.67
1	13.59	5.78	12.32	4.02	2.25	0.32
2	12.55	0.14	7.55	4.71	7.55	4.71
3	15.72	5.17	14.10	3.70	14.10	3.70
4	13.16	2.28	1.43	0.10	1.43	0.10
5	13.73	0.88	11.44	4.56	13.05	2.37
6	3.20	0.28	9.98	5.72	10.41	0.25
7	8.64	1.57	8.68	0.24	8.68	0.24
8	14.65	5.38	14.03	0.12	14.03	0.12
9	13.94	0.28	9.53	5.67	13.94	0.28

Table 11: Humanoid 3 worst parameters for each iteration over 10 seeds

Iteration number	1			2			3			4		
Seed number	torsomass	leftthighmass	rightfootmass	torsomass	leftthighmass	rightfootmass	torsomass	leftthighmass	rightfootmass	torsomass	leftthighmass	rightfootmass
0	11.55	2.20	3.72	12.54	2.19	3.95	14.70	1.02	2.62	5.56	1.70	2.32
1	13.33	4.39	2.21	6.82	2.41	0.10	6.82	2.41	0.10	6.82	2.41	0.10
2	14.17	3.93	0.13	13.31	3.83	2.63	13.31	3.83	2.63	13.31	3.83	2.63
3	3.36	3.58	7.10	15.16	3.57	0.40	2.39	0.17	6.29	2.39	0.17	6.29
4	14.67	4.16	3.77	3.75	4.49	4.74	3.75	4.49	4.74	14.29	0.38	0.44
5	13.06	1.88	0.83	9.60	2.62	7.73	16.00	5.00	8.00	4.91	3.05	3.34
6	14.85	3.80	0.32	13.39	4.93	5.96	10.06	3.41	1.83	10.06	3.41	1.83
7	14.30	1.83	1.33	8.57	2.98	1.43	6.06	1.22	7.14	7.72	0.53	0.47
8	15.60	2.67	3.00	12.04	0.48	4.05	12.04	0.48	4.05	13.34	0.47	0.68
9	14.61	4.15	3.85	10.56	4.62	0.65	10.56	4.62	0.65	10.56	4.62	0.65

Table 12: InvertedPendulum 2 worst parameters for each iteration over 10 seeds

Iteration number	1		2		3	
Seed number	polemass	cartmass	polemass	cartmass	polemass	cartmass
0	16.76	8.50	16.76	8.50	16.76	8.50
1	11.84	4.23	11.84	4.23	11.84	4.23
2	25.32	2.40	25.32	2.40	25.32	2.40
3	16.02	4.55	16.02	4.55	16.02	4.55
4	27.92	10.88	27.92	10.88	27.92	10.88
5	16.49	8.04	16.49	8.04	16.49	8.04
6	11.32	9.65	11.32	9.65	11.32	9.65
7	22.62	4.35	22.62	4.35	22.62	4.35
8	16.76	8.50	16.76	8.50	16.76	8.50
9	11.84	4.23	11.84	4.23	11.84	4.23

Table 13: Halfcheetah 2 worst parameters for each iteration over 10 seeds

Iteration number	1		2		3	
Seed number	worldfriction	torsomass	worldfriction	torsomass	worldfriction	torsomass
0	4.00	7.00	4.00	0.10	4.00	7.00
1	4.00	7.00	4.00	7.00	4.00	7.00
2	0.10	7.00	4.00	7.00	4.00	7.00
3	4.00	7.00	4.00	0.55	4.00	7.00
4	0.10	7.00	4.00	7.00	4.00	7.00
5	4.00	7.00	4.00	0.14	4.00	0.10
6	4.00	7.00	4.00	0.10	4.00	7.00
7	0.10	7.00	4.00	7.00	4.00	7.00
8	4.00	7.00	4.00	0.10	4.00	7.00
9	0.10	0.10	0.51	4.38	0.61	6.85

Table 14: Halfcheetah 3 worst parameters for each iteration over 10 seeds

Iteration number			1			2			3			4
Seed number	ψ_{11}	ψ_{12}	ψ_{13}	worldfriction	torsomass	backthighmass	worldfriction	torsomass	backthighmass	worldfriction	torsomass	backthighmass
0	4.00	0.10	0.10	4.00	0.10	0.10	4.00	7.00	3.00	4.00	7.00	0.10
1	4.00	7.00	0.10	4.00	7.00	3.00	0.10	0.10	3.00	4.00	7.00	3.00
2	0.10	0.10	3.00	4.00	7.00	3.00	4.00	7.00	0.25	0.10	0.10	0.10
3	4.00	7.00	3.00	4.00	0.10	3.00	4.00	1.21	3.00	4.00	0.10	3.00
4	4.00	0.10	0.39	0.10	7.00	3.00	4.00	7.00	3.00	4.00	0.10	0.10
5	0.10	0.10	3.00	0.32	0.22	0.28	0.32	0.22	0.28	0.32	0.22	0.28
6	0.10	7.00	3.00	0.35	0.34	0.51	0.35	0.34	0.51	0.35	0.34	0.51
7	4.00	7.00	0.10	4.00	7.00	3.00	4.00	7.00	3.00	4.00	2.28	3.00
8	4.00	7.00	3.00	4.00	7.00	0.10	4.00	0.10	3.00	4.00	0.10	3.00
9	4.00	7.00	0.25	4.00	7.00	3.00	4.00	0.90	3.00	4.00	0.90	3.00

Table 16: Ant 3 worst parameters for each iteration over 10 seeds

Iteration number			1			2			3			4
Seed number	torsomass	frontleftlegmass	frontrightlegmass	torsomass	frontleftlegmass	frontrightlegmass	torsomass	frontleftlegmass	frontrightlegmass	torsomass	frontleftlegmass	frontrightlegmass
0	0.10	0.01	1.25	3.00	0.01	1.80	3.00	2.20	0.01	1.65	1.03	0.26
1	2.12	1.15	0.51	0.34	0.18	2.80	0.34	0.18	2.80	2.11	3.00	3.00
2	1.55	2.14	0.01	3.00	2.15	3.00	3.00	2.15	3.00	3.00	2.15	3.00
3	2.49	2.09	1.89	1.94	0.01	3.00	3.00	3.00	0.01	3.00	3.00	0.01
4	0.96	1.66	0.32	0.10	3.00	3.00	0.10	3.00	3.00	0.10	3.00	3.00
5	0.63	0.36	2.91	0.17	2.55	2.23	0.10	3.00	3.00	1.91	1.07	1.00
6	1.84	0.88	0.18	1.07	1.61	0.91	2.65	2.22	0.78	2.65	2.22	0.78
7	0.53	0.82	1.51	1.09	0.01	0.01	3.00	3.00	0.01	3.00	3.00	2.98
8	0.10	0.01	3.00	3.00	3.00	0.01	3.00	3.00	0.01	3.00	3.00	2.54
9	3.00	0.01	0.01	3.00	3.00	0.15	0.76	2.31	2.53	0.14	2.16	0.06

Table 17: Walker 2 worst parameters for each iteration over 10 seeds

Iteration number	1		2		3	
Seed number	worldfriction	torsomass	worldfriction	torsomass	worldfriction	torsomass
0	4.00	5.00	4.00	5.00	4.00	5.00
1	4.00	5.00	4.00	1.92	4.00	5.00
2	4.00	0.10	4.00	5.00	4.00	5.00
3	4.00	5.00	4.00	3.51	4.00	3.84
4	4.00	0.10	4.00	0.10	4.00	0.10
5	4.00	5.00	4.00	5.00	4.00	5.00
6	4.00	4.48	4.00	5.00	4.00	5.00
7	4.00	0.10	4.00	0.13	4.00	0.10
8	4.00	5.00	4.00	0.10	4.00	0.10
9	4.00	5.00	4.00	0.10	4.00	5.00

Table 18: Walker 3 worst parameters for each iteration over 10 seeds

Iteration number	1			2			3			4		
Seed number	worldfriction	torsomass	thighmass	worldfriction	torsomass	thighmass	worldfriction	torsomass	thighmass	worldfriction	torsomass	thighmass
0	4.00	5.00	6.00	3.00	0.30	0.28	4.00	0.30	0.24	4.00	2.90	6.00
1	4.00	0.10	0.10	4.00	5.00	6.00	0.26	5.00	0.10	4.00	5.00	5.82
2	4.00	5.00	6.00	4.00	5.00	6.00	4.00	5.00	0.10	4.00	1.63	6.00
3	4.00	5.00	6.00	4.00	5.00	6.00	4.00	5.00	4.26	4.00	5.00	6.00
4	4.00	0.10	0.10	4.00	0.10	0.10	4.00	5.00	6.00	4.00	5.00	6.00
5	4.00	5.00	6.00	4.00	0.10	0.10	4.00	5.00	6.00	4.00	2.24	6.00
6	1.47	5.00	6.00	4.00	5.00	6.00	0.10	0.10	0.10	2.41	5.00	6.00
7	4.00	0.64	6.00	4.00	0.10	6.00	4.00	0.10	6.00	4.00	0.10	6.00
8	1.95	5.00	6.00	4.00	5.00	6.00	4.00	5.00	6.00	4.00	5.00	6.00
9	1.35	0.10	0.10	0.35	1.74	2.35	0.10	0.10	0.10	0.10	0.10	0.10

Table 19: Hopper 2 worst parameters for each iteration over 10 seeds

Iteration number	1		2		3	
Seed number	worldfriction	torsomass	worldfriction	torsomass	worldfriction	torsomass
0	3.00	0.10	3.00	2.87	2.64	1.46
1	3.00	0.10	3.00	3.00	2.92	0.23
2	3.00	3.00	2.99	0.55	1.73	0.47
3	3.00	0.10	1.36	2.64	2.98	0.59
4	3.00	0.10	3.00	3.00	2.93	2.34
5	3.00	0.10	1.58	2.95	2.60	0.24
6	1.81	0.17	0.89	3.00	2.94	2.88
7	3.00	0.10	1.49	3.00	0.10	3.00
8	3.00	0.92	1.45	0.25	2.31	0.12
9	2.95	0.12	1.76	2.94	2.97	0.13

Table 20: Hopper 3 worst parameters for each iteration over 10 seeds

Iteration number	1			2			3			4		
Seed number	worldfriction	torsomass	thighmass	worldfriction	torsomass	thighmass	worldfriction	torsomass	thighmass	worldfriction	torsomass	thighmass
0	3.00	0.10	0.10	3.00	3.00	4.00	3.00	3.00	4.00	2.82	1.81	2.35
1	3.00	0.10	0.10	2.83	0.50	0.70	2.10	0.55	0.78	0.78	0.60	0.86
2	3.00	0.10	0.10	3.00	3.00	4.00	3.00	3.00	4.00	3.00	3.00	4.00
3	3.00	0.10	0.10	3.00	3.00	4.00	0.41	0.56	0.56	0.37	0.15	0.40
4	3.00	0.10	0.10	2.15	0.59	0.47	1.67	2.05	0.87	3.00	0.10	0.10
5	0.65	0.42	0.40	1.36	1.80	2.97	1.60	0.18	0.12	2.96	0.10	0.10
6	3.00	0.10	0.10	2.92	0.60	0.84	0.10	0.10	0.10	2.99	1.87	1.16
7	3.00	0.10	0.10	2.31	3.00	4.00	2.15	3.00	4.00	2.05	3.00	0.10
8	3.00	0.11	0.10	3.00	3.00	4.00	1.38	0.53	4.00	3.00	3.00	4.00
9	3.00	0.10	0.10	2.00	2.00	3.41	3.00	3.00	4.00	1.10	0.21	0.16

Table 21: Halfcheetah 2 worst parameters for each iteration over 10 seeds

	$J^{\pi_0}_{T_0}$	ψ^1_1	ψ^1_2	$J^{\pi_0}_{T_1}$	$J^{\pi_1}_{T_1}$	ψ^2_1	ψ^2_2	$J^{\pi_1}_{T_2}$	$J^{\pi_2}_{T_2}$	ψ^3_1	ψ^3_2	$J^{\pi_2}_{T_3}$
0	5753.2	4.0	7.0	2128.3	20559.5	4.0	0.1	3626.3	9724.7	4.0	7.0	3745.6
1	109.755	4.0	7.0	2027.0	3732.0	4.0	7.0	3827.6	6183.4	4.0	7.0	3827.6
2	5897.4	0.1	7.0	2354.0	9930.0	4.0	7.0	2497.4	3670.5	4.0	7.0	3874.4
3	6396.9	4.0	7.0	2473.0	3815.1	4.0	0.6	3053.9	6022.4	4.0	7.0	3825.5
4	2143.26	0.1	7.0	1633.3	3807.9	4.0	7.0	1978.4	6282.5	4.0	7.0	1978.4
5	4476.0	4.0	7.0	1983.2	13040.2	4.0	0.1	2597.4	3694.9	4.0	0.1	2822.3
6	6275.2	4.0	7.0	-91.1	3639.7	4.0	0.1	3448.4	5347.6	4.0	7.0	3986.0
7	6213.8	0.1	7.0	1587.1	2046.31	4.0	7.0	2411.1	3757.7	4.0	7.0	3769.5
8	6409.8	4.0	7.0	2053.8	3709.9	4.0	0.1	3182.3	3712.3	4.0	7.0	3872.9
9	6801.4	0.1	0.1	2341.4	3620.1	0.5	4.4	3619.3	5737.3	0.6	6.8	5431.3

Table 22: Halfcheetah THREE worst parameters for each iteration over 10 seeds

M				0				1				2				3
Parameters	wf	tm	bm	wr	wf	tm	bm	wr	wf	tm	bm	wr	wf	tm	bm	wr
0	4.0	0.1	0.1	572.7	4.0	0.1	0.1	1329.2	4.0	7.0	3.0	1221.3	4.0	7.0	0.1	1781.1
1	4.0	7.0	0.1	650.8	4.0	7.0	3.0	1932.2	0.1	0.1	3.0	1835.7	4.0	7.0	3.0	1863.6
2	0.1	0.1	3.0	338.1	4.0	7.0	3.0	1116.4	4.0	7.0	0.2	1133.7	0.1	0.1	0.1	1659.6
3	4.0	7.0	3.0	824.7	4.0	0.1	3.0	2016.6	4.0	1.2	3.0	2380.5	4.0	0.1	3.0	2402.4
4	4.0	0.1	0.4	1019.4	0.1	7.0	3.0	85.4	4.0	7.0	3.0	1414.6	4.0	0.1	0.1	1016.5
5	0.1	0.1	3.0	812.0	0.3	0.2	0.3	4363.3	0.3	0.2	0.3	4363.3	0.3	0.2	0.3	4363.3
6	0.1	7.0	3.0	-109.3	0.3	0.3	0.5	2414.9	0.3	0.3	0.5	2414.9	0.3	0.3	0.5	2414.9
7	4.0	7.0	0.1	158.4	4.0	7.0	3.0	2315.3	4.0	7.0	3.0	1996.3	4.0	2.3	3.0	2307.8
8	4.0	7.0	3.0	1077.4	4.0	7.0	0.1	2415.3	4.0	0.1	3.0	2428.4	4.0	0.1	3.0	2503.1
9	4.0	7.0	0.2	1676.7	4.0	7.0	3.0	1955.0	4.0	0.9	3.0	2047.2	4.0	0.9	3.0	2047.2