

Learning-Based Robust Control: Unifying Exploration and Distributional Robustness for Reliable Robotics via Free Energy

Hozefa Jesawada^{*1}, Giovanni Russo², Abdalla Swikir³, and Fares Abu-Dakka¹

Abstract—A key challenge towards reliable robotic control is devising computational models that can both learn policies and guarantee robustness when deployed in the field. Inspired by the free energy principle in computational neuroscience, to address these challenges, we propose a model for policy computation that jointly learns environment dynamics and rewards, while ensuring robustness to epistemic uncertainties. Expounding a distributionally robust free energy principle, we propose a modification to the maximum diffusion learning framework. After explicitly characterizing robustness of our policies to epistemic uncertainties in both environment and reward, we validate their effectiveness on continuous-control benchmarks, via both simulations and real-world experiments involving manipulation with a Franka Research 3 arm. Across simulation and zero-shot deployment, our approach narrows the sim-to-real gap, and enables repeatable tabletop manipulation without task-specific fine-tuning.

I. INTRODUCTION

A popular paradigm to design autonomous robots has become that of learning control policies from simulations [1]. Yet, even with high-fidelity simulators, these policies can fail when deployed in real-world conditions that deviate, even slightly, from those seen during training [2]. Such mismatches can arise in, e.g., contact-rich tasks, or in the presence of sensing/actuation noise, or even when some dynamics – such as nonlinear frictions, delays, calibration drifts [3] – is not perfectly modeled. In these settings, small discrepancies can cause large failures, with potentially catastrophic consequences for the robot and its environment.

In this context, the minimization of the free energy – and its instantiation into active inference robotics architectures – has emerged as a promising approach bridging machine learning, robotics and neuroscience to design inference, learning and control algorithms for effective robot control [4], [5], [6], [7], [8], [9], [10]. Yet, despite the success of free energy based policies, there is currently no computational model that, simultaneously, learns the policy and provides explicit robustness guarantees against epistemic uncertainties in the environment model learned by the agent and the task reward. Motivated by this, we present a computational model that: (i) without having available a model of the environment and reward function, learns a policy that minimizes the (variational) free energy; (ii) provides a-priori explicit robustness guarantees of the policy that can be used as *certificates* for robot deployment in the field.

^{*} Corresponding: hj3000@nyu.edu. ¹Mechanical Engineering Program, New York University Abu Dhabi, Abu Dhabi, UAE. ²Department of Information and Electrical Engineering & Applied Mathematics, University of Salerno, Fisciano, Italy. ³Mohammed bin Zayed University for Artificial Intelligence, Abu Dhabi, UAE. Work not supported by any organization

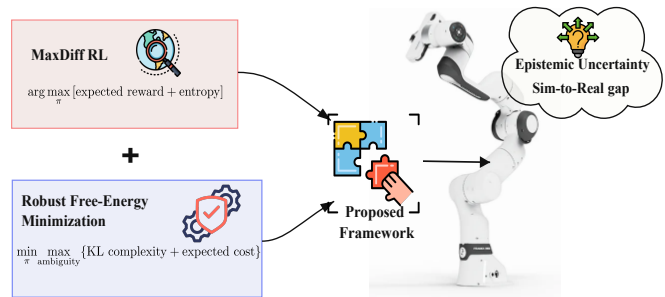


Fig. 1: Modifying Maximum Diffusion RL with Distributionally Robust Free Energy Principle to address epistemic uncertainty, narrowing sim-to-real gaps for reliable robotics.

To achieve this, expounding a model-based distributional free energy principle (DR-FREE) for policy computation [11], we propose a modification to a recent learning framework [12], Maximum Diffusion (MaxDiff). MaxDiff, which generalizes maximum entropy methods – see, e.g., [13], [14], [15] – and exceeds state-of-the-art performance on popular robotics benchmarks, learns policies with maximally diffusive path statistics without having access to the environment model and reward function. However, despite its impressive control performance, MaxDiff robustness is only implicit: it depends on the entropy of the optimal policy, not known a priori [13], and explicit bounds are only available for discrete action spaces. In contrast, DR-FREE provides explicit distributional robustness guarantees in free energy minimization problems, but requires access to both the dynamics model and the reward. Remarkably, as we recall in Section II, the MaxDiff optimization can be recast as a free energy minimization problem, enabling us to integrate DR-FREE into MaxDiff and obtain explicit robustness guarantees without sacrificing MaxDiff’s desirable features. The effectiveness of our method is illustrated via continuous control benchmarks. Numerical and hardware experiments – involving manipulation with a Franka Emika Panda arm – show that our method outperforms baseline MaxDiff [12], narrows the sim-to-real gap and enables zero-shot deployment of the policy.

Related Work: the minimization of the free energy offers a unifying account across machine learning, robotics and neuroscience. See [5], [6], [7], [8], [9], [10] for surveys and results instantiating this framework into robots. Here, we focus on robustness in learning based control for robotics.

Adversarial and ensemble-based training strategies like RARL [16] and EPOpt [17] expose policies to disturbances

or challenging domains during training, enhancing tolerance to parametric spread. Domain randomization has been a key sim2real tactic in robotics, i.e., randomizing inertial and contact parameters for transfer [18], [19], [20]. While effective in broadening training coverage, these methods typically optimize a *nominal* objective at test time; robustness relies on the alignment between training distributions and deployment shifts, with no explicit control over deployment-time worst-case model misspecification. In contrast, we optimize a *distributionally robust* objective at execution through per-state-action KL ambiguity, directly aligning decisions with epistemic risk.

Robust MDPs and their distributionally robust variants optimize the worst-case value over uncertainty sets on transition kernels, yielding robust Bellman operators and bounds under bounded misspecification [21], [22], [23]. Wasserstein and f -divergence sets (incl. KL) offer tractable inner maximizations and interpretable radii [24]. Our formulation instantiates *per-state-action* KL balls shaped by the learned model’s uncertainty and solves the inner step via a free-energy (Donsker–Varadhan) reduction [11] to a scalar convex problem, enabling real-time planning. In contrast to most DR-MDP applications that target tabular or abstract benchmarks, we tailor ambiguity calibration and solver design to continuous-state/action robotics and couple it with a diffusion prior.

Max-entropy RL and control-as-inference (e.g., LMDP, PI^2 , soft-actor methods) temper policies toward high entropy, offering stochastic smoothing and connections to risk sensitivity [25]. These objectives regularize the policy distribution (or passive dynamics) [26], but do not defend against adversarial shifts in the transition model. We borrow the same free-energy machinery but place the KL on next-state distributions (epistemic dynamics), yielding a robust soft policy whose conservatism adapts to model uncertainty.

Information-theoretic MPC/MPPI performs trajectory sampling with exponential cost weighting and has seen extensive robotic deployments (autonomous driving, quadrotors, manipulators), including GPU-accelerated real-time variants [27], [28]. Standard NN-MPPI, however, assumes a single predictive model plus temperature tuning; it does not enforce distributional robustness to transition misspecification. Our controller preserves MPPI’s practicality but embeds a principled inner robustification via per-step KL ambiguity.

Probabilistic ensembles (PETS) [29] and model-based policy optimization (MBPO) [30] achieve strong sample efficiency and competitive control on continuous-control and robotics benchmarks. Model-Ensemble Exploration and Exploitation (MEEE) [31] tackles the problem of exploration and exploitation simultaneously by taking into account the future observation novelty and leveraging a model-ensemble. However, none of these methods provide any formal guarantees of robustness against adversarial uncertainty.

Maximum Diffusion RL [12] enhances state-space coverage by amplifying diffusion and mitigating temporal correlations. MaxDiff outperforms MaxEnt RL [13], NN-MPPI [28], and Soft-Actor Critic (SAC) [15]. We incorporate a

MaxDiff-style diffusion prior into our planner to ensure effective exploration in simulation, while KL ambiguity maintains robustness during execution, ensuring clear separation of concerns.

Robotics-focused RL demonstrations—SAC on real platforms, vision-in-the-loop manipulation, and sim2real locomotion [32], [33]—either rely on large real-world data, on-policy adaptation, or heavy randomization. In contrast, using a simulation-trained dynamics model, (i) we align control with epistemic risk via per-state KL ambiguity, (ii) preserve sample-based planning (MPPI/PETS), and (iii) couple exploration (MaxDiff) with free-energy DR robustness without deployment fine-tuning. We also show that KL-bounded dynamics robustness extends to a class of reward perturbations through an augmented-state free-energy bound.

Contributions: we propose a modification to MaxDiff, combining it with the distributional free energy principle, DR-FREE. For our resulting computational model, we characterize policy robustness bounds and evaluate performance across both simulations, via OpenAI Gym [34] and MuJoCo [35], and real hardware experiments with a Franka Emika Panda arm in a manipulation task. Experiments convincingly show that our framework outperforms standard MaxDiff baselines. Moreover, our model also proves useful for zero-shot deployments. In particular, we train a policy for a *Franka Research 3* from a simulator, with the model available during training different from the real robotic arm. Then, after training, we deploy the trained policy on the real arm and show that, without any fine-tuning, the robot is able to complete a tabletop manipulation task.

To the best of our knowledge, this work provides the first free energy computational model that, simultaneously, learns policies for continuous control tasks and provides explicit, a priori robustness guarantees that can be used to certify the use of the robot in field settings. Indeed: 1) in MaxDiff, and related literature, continuous policies can be learned without having access to the environment model and reward function. However, robustness only emerges a posteriori as a by-product of the entropy of the optimal policy; 2) In contrast, DR-FREE can guarantee robustness a priori. However, it requires a model and reward available to the agent.

Our framework combines the strengths of these approaches, enabling a priori robustness guarantees as in DR-FREE, while simultaneously learning policies without requiring access to models or rewards as in MaxDiff.

II. BACKGROUND

After introducing the key notation and definitions, we formalize the optimization problems relevant to our work. Then, leveraging this problem, we introduce the key ingredients behind our computational model, MaxDiff and DR-FREE, motivating our approach to learning-based robust control.

A. Notation and Definitions

We use calligraphic letters to denote sets e.g., \mathcal{X} and boldface to denote vectors e.g., \mathbf{x} . A random variable is written as \mathbf{V} , with a particular realization denoted by \mathbf{v} .

The *probability density function* (pdf) of \mathbf{v} is denoted by $p(\mathbf{v})$. Whenever integrals involving a pdf are considered, we implicitly assume that they exist. The expectation of a function $\mathbf{h}(\cdot)$ applied to \mathbf{v} is defined as $\mathbb{E}_p[\mathbf{h}(\mathbf{v})] := \int \mathbf{h}(\mathbf{v})p(\mathbf{v})d\mathbf{v}$, where the integral is taken over the support of $p(\mathbf{v})$. The conditional pdf of \mathbf{v}_1 given \mathbf{v}_2 is written as $p(\mathbf{v}_1 | \mathbf{v}_2)$. For countable sets, we use the notation $wk_1 : k_n$, where w_k represents a generic element, $k_1 (k_n)$ denotes the index of the first (last) element, and $k_1 : k_n$ represents the set of consecutive integers from k_1 through k_n . We denote by \mathcal{D} the set of all probability density functions.

Definition 2.1: The KL divergence of $p(\mathbf{v})$ w.r.t $q(\mathbf{v})$ with $\text{supp } p \subseteq \text{supp } q$ is:

$$D_{\text{KL}}(p || q) := \int_{\mathbf{v}} p(\mathbf{v}) \ln \frac{p(\mathbf{v})}{q(\mathbf{v})} d\mathbf{v}.$$

The KL divergence is a measure of the proximity of the pair of pdfs. For discrete pmfs the integral can be substituted with a summation over \mathbf{v} .

Let $\mathcal{X} \subseteq \mathbb{K}^n$ and $\mathcal{U} \subseteq \mathbb{K}^p$ be the state and action spaces (with $\mathbb{K} \in \{\mathbb{R}, \mathbb{Z}\}$), and let the decision horizon be $k = 1, \dots, N$. At each k , we consider:

- A *trained (nominal) dynamics model* $\bar{p}_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$;
- A *generative model* $q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ for states and a *generative prior* $q_k(\mathbf{u}_k | \mathbf{x}_{k-1})$ for actions;
- the trajectory measures

$$p_{0:N} = p_0(\mathbf{x}_0) \prod_{k=1}^N p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \pi_k(\mathbf{u}_k | \mathbf{x}_{k-1}),$$

$$q_{0:N} = p_0(\mathbf{x}_0) \prod_{k=1}^N q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) q_k(\mathbf{u}_k | \mathbf{x}_{k-1}).$$

- stage costs $c_k^{(x)} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ and $c_k^{(u)} : \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$.

B. Free-Energy Minimization Problem

Free-energy principle arises naturally in decision-making across information theory, active inference, learning, and control. In RL, they underpin schemes such as Maximum Entropy RL [13], [12], Bayesian inference [7], and entropy-based transport methods.

In our setting, the free-energy functional is

$$\mathcal{F}(p_{0:N}) = \underbrace{D_{\text{KL}}(p_{0:N} || q_{0:N})}_{\text{complexity}} + \underbrace{\mathbb{E}_{p_{0:N}} \left[\sum_{k=1}^N c_k^{(x)}(\mathbf{x}_k) + c_k^{(u)}(\mathbf{u}_k) \right]}_{\text{expected cost}}, \quad (1)$$

where the first term measures divergence from a reference distribution $q_{0:N}$ and the second is expected cost. Notably, by defining a cost-shaped reference $\tilde{q}_{0:N} \propto q_{0:N} \exp(-\sum_k c_k^{(x)} + c_k^{(u)})$, (1) reduces to $\mathcal{F}(p_{0:N}) = D_{\text{KL}}(p_{0:N} || \tilde{q}_{0:N})$, shows that minimization of (1) is equivalent to KL projection onto a cost-weighted reference.

C. Maximum Diffusion (MaxDiff) RL

MaxDiff derives a path-distribution P_{max} that *maximizes trajectory entropy* under continuity constraints, yielding Markov and ergodic sample paths and diffusion-like exploration. Writing the finite-horizon path entropy via the chain rule,

$$S[P_{\text{max}}[x_{1:N}]] = \sum_{t=1}^N S[p_{\text{max}}(\mathbf{x}_{t+1} | x_t)] \propto \frac{1}{2} \sum_{t=1}^N \log \det C[x_t],$$

where $C[x]$ is the local state-increment covariance; thus maximizing path entropy promotes large $\log \det C[x]$ along the trajectory.

A practical MaxDiff RL objective (SOC form) augments rewards with a diffusion bonus:

$$\pi_{\text{MaxDiff}}^* \in \arg \max_{\pi} \mathbb{E}_{(\mathbf{x}_{1:N}, \mathbf{u}_{1:N}) \sim P_{\pi}} \left[\sum_{k=1}^N r(\mathbf{x}_k, \mathbf{u}_k) + \frac{\alpha}{2} \log \det C_{\pi}[x_k] \right], \quad (2)$$

which is the objective used for the empirical results in [12]. Equivalently, MaxDiff can be cast as minimizing $D_{\text{KL}}(P_{\pi} || P_{\text{max}}^r)$ between the agent's path distribution and a maximally diffusive, reward-shaped reference P_{max}^r , leading to an SOC with modified running reward or cost.

D. Distributionally Robust Free-Energy Principle

DR-FREE [11] defines a sequential min-max problem over policies $\{\pi_k(\mathbf{u}_k | \mathbf{x}_{k-1})\}_{1:N}$ and environment models $\{p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)\}_{1:N}$:

$$\min_{\{\pi_k\}} \max_{\{p_k \in \mathcal{B}_{\eta}(\bar{p}_k)\}} \mathcal{F}(p_{0:N}), \quad (3)$$

subject to $p_k \in \mathcal{B}_{\eta}(\bar{p}_k)$ for all $k = 1, \dots, N$. The ambiguity set is

$$\mathcal{B}_{\eta}(\bar{p}_k) = \left\{ p_k \in \mathcal{D} : D_{\text{KL}}(p_k || \bar{p}_k) \leq \eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) \right\},$$

where $\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) \geq 0$ bounds the statistical complexity relative to the trained model \bar{p}_k , and the support constraint ensures the free-energy objective remains finite.

At each time-step k , the resolution engine solves a stage-wise bi-level problem:

$$\min_{\pi_k} D_{\text{KL}}(\pi_k || q_k^{(u)}) + \mathbb{E}_{\pi_k} \left[c_k^{(u)}(\mathbf{u}_k) + \max_{p_k \in \mathcal{B}_{\eta}(\bar{p}_k)} \left\{ D_{\text{KL}}(p_k || q_k^{(x)}) + \mathbb{E}_{p_k} [\tilde{c}_k(\mathbf{X}_k)] \right\} \right], \quad (4)$$

where the cost-to-go is defined recursively as $\tilde{c}_k(\mathbf{x}_k) = c_k^{(x)}(\mathbf{x}_k) + \hat{c}_{k+1}(\mathbf{x}_k)$, with $\hat{c}_{k+1}(\mathbf{x}_k)$ being the optimal value of (4) evaluated at $k+1$ (initialized at $\hat{c}_{N+1} = 0$). The inner maximization reduces to a scalar convex optimization problem whose non-negative optimal value yields the *cost of ambiguity* $\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) + \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$.

The optimal policy has an explicit Gibbs form:

$$\pi_k^*(u | \mathbf{x}_{k-1}) \propto q_k(u | \mathbf{x}_{k-1}) \exp\left(-c_k^{(u)}(u) - \eta_k(\mathbf{x}_{k-1}, \mathbf{u}) - \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u})\right), \quad (5)$$

i.e., the generative prior $q_k(\mathbf{u}_k | \mathbf{x}_{k-1})$ is modulated by an exponential kernel containing the action cost and the cost of ambiguity. Actions associated with higher ambiguity are assigned lower probability, while as ambiguity vanishes, the policy recovers that of a standard free-energy minimizing agent.

III. MODIFYING MAXDIFF WITH DR-FREE

First, we note that, given the reward $r = -(c^{(x)} + c^{(u)})$, one can define

$$\tilde{q}_{0:N} \propto q_{0:N} \exp\left(-\sum_{k=1}^N (c_k^{(x)}(\mathbf{X}_k) + c_k^{(u)}(\mathbf{U}_k))\right).$$

Then, as noted in [11], when there is no ambiguity the DR-FREE objective can be recast as the minimization of the KL divergence between $p_{0:N}$ and $\tilde{q}_{0:N}$ [36]. In turn, this means that the optimization problem becomes the same problem tackled by MaxDiff when the generative model is chosen as $q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) = p_{\max}(\mathbf{x}_k | \mathbf{x}_{k-1})$ and a fixed action prior $q_k(\mathbf{u}_k | \mathbf{x}_{k-1})$ (e.g., $\bar{p}(\mathbf{u}_k | \mathbf{x}_{k-1})$), so that the reference density becomes $\propto \prod_k p_{\max}(\mathbf{x}_k | \mathbf{x}_{k-1}) \exp(\sum_k r(\mathbf{x}_k, \mathbf{u}_k))$.

A. Modifying MaxDiff: constructing p_{\max} and using it as the state generative model

a) *Principle*: in DR-FREE, the complexity term $D_{\text{KL}}(p_{0:N} || q_{0:N})$ biases the optimal policy toward a *generative model* $q_{0:N}$ whose single-step factors are $q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and $q_k(\mathbf{u}_k | \mathbf{x}_{k-1})$. To inject MaxDiff's exploration into this bias, we *choose the state generative kernel to be maximally diffusive*, i.e.

$$q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \equiv p_{\max}(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k),$$

while keeping any convenient prior $q_k(\mathbf{u}_k | \mathbf{x}_{k-1})$ for actions. This choice plugs MaxDiff's path-entropy objective into DR-FREE via the complexity term, while the ambiguity machinery (inner maximization over p_k) and the soft optimal policy remain as in DR-FREE (scalar, convex inner step; Gibbs-form policy kernel).

b) *What p_{\max} means (MaxDiff)*: MaxDiff defines maximally diffusive Markov transitions by maximizing path entropy; for Gaussian increments this reduces to maximizing the sum of local entropies,

$$S[p_{\max}[x_{1:N}]] = \sum_{t=1}^N S[p_{\max}(\mathbf{x}_{t+1} | x_t)] \propto \frac{1}{2} \sum_{t=1}^N \log \det C[x_t],$$

which yields the implemented MaxDiff-RL objective $\max_{\pi} \mathbb{E}[\sum_t r(\mathbf{x}_t, \mathbf{u}_t) + \frac{\alpha}{2} \log \det C_{\pi}[x_t]]$. In practice, when $C[x]$ can be low-rank/high-dimensional, one may use leading-eigenvalue or log-trace approximations for numerical stability (with the usual caveats).

c) *Computing p_{\max} for DR-FREE: a KL-MaxEnt construction*: for each $(\mathbf{x}_{k-1}, \mathbf{u}_k)$, the trained nominal dynamics are $\bar{p}_k(\cdot | \mathbf{x}_{k-1}, \mathbf{u}_k)$. We compute a *single-step* maximally diffusive kernel by solving the *maximum-entropy trust-region*

problem

$$\begin{aligned} p_{\max}(\cdot | \mathbf{x}_{k-1}, \mathbf{u}_k) &\in \arg \max_{p \in \mathcal{D}} \mathcal{H}(p) \\ \text{s.t. } D_{\text{KL}}(p || \bar{p}_k(\cdot | \mathbf{x}_{k-1}, \mathbf{u}_k)) &\leq \varepsilon_k(\mathbf{x}_{k-1}, \mathbf{u}_k), \\ \text{supp } p &\subseteq \text{supp } \bar{p}_k. \end{aligned} \quad (6)$$

Lagrangian calculus gives $p_{\max}(dx') \propto \bar{p}_k(dx')^{\beta_k}$ with a unique $\beta_k \in (0, 1]$ chosen to meet the KL budget; in exponential families this is equivalent to *flattening* the nominal precision (increasing entropy) while staying within KL radius.

Gaussian closed form: assume $\bar{p}_k(\cdot | \mathbf{x}_{k-1}, \mathbf{u}_k) = \mathcal{N}(\bar{\mu}_k, \bar{\Sigma}_k)$ with $\bar{\mu}_k := \bar{f}_k(\mathbf{x}_{k-1}, \mathbf{u}_k)$. Then

$$p_{\max}(\cdot | \mathbf{x}_{k-1}, \mathbf{u}_k) = \mathcal{N}(\bar{\mu}_k, \Sigma_{\max,k}), \Sigma_{\max,k} = \lambda_k \bar{\Sigma}_k, \lambda_k > 0,$$

with *scalar* λ_k determined by the KL constraint. If we enforce $D_{\text{KL}}(p_{\max} || \bar{p}_k) = \varepsilon_k$, the unique solution satisfies

$$\frac{n}{2} (\lambda_k - 1 - \log \lambda_k) = \varepsilon_k \quad (\lambda_k \geq 1), \quad (7)$$

so $\Sigma_{\max,k}$ is a *uniform inflation* of $\bar{\Sigma}_k$ that maximizes entropy, i.e., $\log \det \Sigma_{\max,k} = \log \det \bar{\Sigma}_k + n \log \lambda_k$ increases monotonically with λ_k —exactly aligned with MaxDiff's $\frac{1}{2} \log \det(\cdot)$ objective. Equation (7) is a one-dimensional root find (bisection/Newton), strictly increasing in λ_k .

d) *Plug-in to DR-FREE*: Once p_{\max} is computed, set $q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \leftarrow p_{\max}(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and proceed with DR-FREE: (i) the *inner* maximization of $D_{\text{KL}}(p_k || q_k) + \mathbb{E}_{p_k}[\tilde{c}_k(\mathbf{X}_k)]$ over $p_k \in \mathcal{B}_{\eta}(\bar{p}_k)$ yields the cost of ambiguity $\eta(\mathbf{x}_{k-1}, \mathbf{u}_k) + \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$, where $\tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k) = \min_{\alpha \geq 0} \tilde{V}_{\alpha}(\mathbf{x}_{k-1}, \mathbf{u}_k)$ is obtained via scalar convex optimization¹; (ii) the *outer* minimization over π_k returns the Gibbs policy (5).

e) *Ambiguity-free limit*: when $\eta_k \rightarrow 0$ the DR-FREE policy reduces to the free-energy minimizer that explicitly contains $D_{\text{KL}}(\bar{p}_k || q_k)$ and $\mathbb{E}_{\bar{p}_k}[\tilde{c}_k]$ in its exponent; with the MaxDiff plug-in $q_k = p_{\max}$ (Gaussian), this term is

$$D_{\text{KL}}(\mathcal{N}(\bar{\mu}_k, \bar{\Sigma}_k) || \mathcal{N}(\bar{\mu}_k, \lambda_k \bar{\Sigma}_k)) = \frac{n}{2} \left(\frac{1}{\lambda_k} - 1 + \log \lambda_k \right),$$

which is a closed-form function of the λ_k used to build p_{\max} .

f) *When estimating through data*: when estimating a state-diffusion matrix $C[x]$ from data as in [13], use the log-det objective or its stable approximations for large n . Any q_k is admissible in DR-FREE (it can be a time-series model or encode goals); the plug-in $q_k = p_{\max}$ is therefore fully compatible with the framework.

The overall decision cycle is illustrated in Fig. 2, which complements Algorithm 1 by showing the interaction between replay buffer, model training, the computation of p_{\max} , the min-max optimization, and action execution.

¹Explicitly, $\tilde{V}_{\alpha} = \alpha \ln \mathbb{E}_{\bar{p}_k}[(\bar{p}_k \exp \tilde{c}_k / q_k)^{1/\alpha}] + \alpha \eta_k$ for $\alpha > 0$; see [11] for the boundary case $\alpha = 0$.

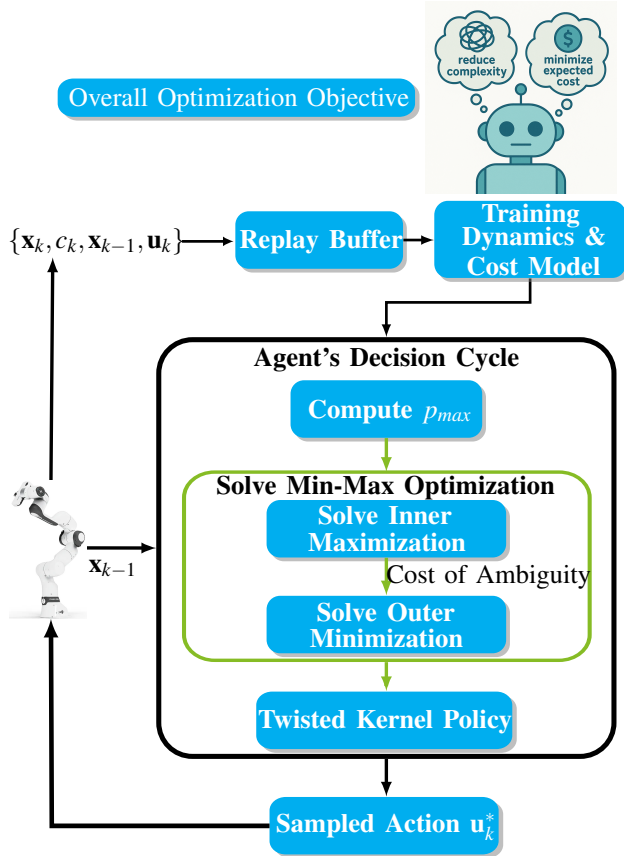


Fig. 2: overview of the proposed control loop. At each step, the agent collects data into a replay buffer, updates the dynamics and cost models, computes the maximally diffusive kernel p_{\max} , and solves the min-max optimization. The resulting twisted kernel produces an action that is applied to the system, and the cycle repeats.

Algorithm 1 Our Computational Model

Require: Horizon N , ambiguity radius $\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k)$, MaxDiff trust ε_q , action prior $q(\mathbf{u}_k | \mathbf{x}_{k-1})$.

- 1: Initialize the replay buffer \mathbb{D} , nominal dynamics model $\bar{p}(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and cost model $c_k(\mathbf{x}_k, \mathbf{u}_k)$ as neural networks.
 - 2: **for** episode $e = 1, 2, \dots, E$ **do**
 - 3: Initialize state \mathbf{x}_0
 - 4: **for** time step $k = 0, 1, \dots, N-1$ **do**
 - 5: **Build** p_{\max} using (6)–(7).
 - 6: Sample random action from \mathcal{U} .
 - 7: **Solve** scalar convex optimization as in (11).
 - 8: **Solve** outer minimization (3) and obtain optimal policy (5).
 - 9: **Execute action:** sample $\mathbf{u}_k \sim \pi_k^*(\cdot | \mathbf{X}_k)$, apply to system, collect $\{\mathbf{x}_k, c_k, \mathbf{x}_{k-1}, \mathbf{u}_k\} \rightarrow \mathbb{D}$.
 - 10: **end for**
 - 11: **Update** $\bar{p}(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and $c_k(\mathbf{x}_k, \mathbf{u}_k)$ using \mathbb{D} and backpropagation.
 - 12: **end for**
-

IV. JOINT ROBUSTNESS TO DYNAMICS AND COST PERTURBATIONS

In this section, we extend our framework to account not only for dynamics misspecification but also for bounded perturbations in the stage cost. The idea of augmenting the system state with a running-cost variable, so that cost perturbations appear as additional uncertainty in the transition kernel, was briefly discussed in [13]. Our formal analysis, however, builds directly on the distributional free-energy results in [11]. Leveraging these results, we show that the policy obtained by solving (3) remains robust under stage cost perturbations. This extension allows the adversary's KL budget to be allocated jointly across dynamics and cost channels, while preserving the tractable min-max structure and Gibbs-form policy of DR-FREE.

A. Augmented State Formulation

Let the nominal (stage) cost be

$$\ell(\mathbf{x}, \mathbf{u}) := c^{(x)}(\mathbf{x}) + c^{(u)}(\mathbf{u}).$$

We reserve $\tilde{c}(\mathbf{x}, \mathbf{u})$ exclusively for the *ambiguity cost* returned by the DR-FREE inner maximization. To account for perturbations $\delta c(\mathbf{x}, \mathbf{u})$ with $|\delta c| \leq \Delta_k(\mathbf{x}, \mathbf{u})$, we augment the observation with the running cost C_k , so the augmented state is (\mathbf{X}_k, C_k) .

The corresponding transition kernels are

$$\begin{aligned} \bar{p}_k^{\text{aug}}(\mathbf{x}', c' | \mathbf{x}, C, \mathbf{u}) &= \bar{p}_k(\mathbf{x}' | \mathbf{x}, \mathbf{u}) \mathcal{N}(c' | C + \ell(\mathbf{x}, \mathbf{u}), \sigma^2), \\ p_k^{\text{aug}, \delta}(\mathbf{x}', c' | \mathbf{x}, C, \mathbf{u}) &= p_k(\mathbf{x}' | \mathbf{x}, \mathbf{u}) \\ &\quad \mathcal{N}(c' | C + \ell(\mathbf{x}, \mathbf{u}) + \delta c(\mathbf{x}, \mathbf{u}), \sigma^2), \end{aligned}$$

where \bar{p}_k is the nominal next-state model and p_k is an admissible (possibly misspecified) next-state model. As in DR-FREE, we impose per-state-action KL balls:

$$D_{\text{KL}}(p_k \| \bar{p}_k) \leq \eta_k^{\text{dyn}}(\mathbf{x}, \mathbf{u}).$$

Theorem 4.1: Suppose the augmented ambiguity radii obey

$$\eta_k^{\text{aug}}(\mathbf{x}, \mathbf{u}) \geq \eta_k^{\text{dyn}}(\mathbf{x}, \mathbf{u}) + \frac{\Delta_k(\mathbf{x}, \mathbf{u})^2}{2\sigma^2} \quad \text{for all } (\mathbf{x}, \mathbf{u}), \quad (8)$$

Then, for each k and every (\mathbf{x}, \mathbf{u}) , the true augmented kernels $p_k^{\text{aug}, \delta}$ lie in the DR-FREE augmented ambiguity sets $\mathcal{B}_{\eta_k^{\text{aug}}}(\bar{p}_k^{\text{aug}})$, and the saddle-point policy solving

$$\min_{\{\pi_k\}} \max_{\{p_k^{\text{aug}} \in \mathcal{B}_{\eta_k^{\text{aug}}}(\bar{p}_k^{\text{aug}})\}} D_{\text{KL}}(p_{0:N}^{\text{aug}} \| \bar{q}_{0:N}^{\text{aug}}) \quad (9)$$

is *simultaneously robust* to both dynamics and cost perturbations, while preserving the DR-FREE Gibbs-form policy.

Proof: **Step 1 — KL decomposition and Gaussian bound.** By the KL chain rule on $(\mathbf{X}_{k+1}, C_{k+1})$,

$$\begin{aligned} D_{\text{KL}}(p_k^{\text{aug}, \delta} \| \bar{p}_k^{\text{aug}}) &= D_{\text{KL}}(p_k(\cdot | \mathbf{x}, \mathbf{u}) \| \bar{p}_k(\cdot | \mathbf{x}, \mathbf{u})) \\ &\quad + \mathbb{E}_{p_k(\mathbf{x}' | \mathbf{x}, \mathbf{u})} \left[D_{\text{KL}}(\mathcal{N}(\cdot | C + \ell + \delta c, \sigma^2) \| \right. \\ &\quad \left. \mathcal{N}(\cdot | C + \ell, \sigma^2)) \right]. \end{aligned} \quad (10)$$

For Gaussians with equal variance, the second term equals $(\delta c)^2/(2\sigma^2) \leq \Delta_k(\mathbf{x}, \mathbf{u})^2/(2\sigma^2)$. Using $D_{\text{KL}}(p_k \|\bar{p}_k) \leq \eta_k^{\text{dyn}}(\mathbf{x}, \mathbf{u})$, we obtain

$$D_{\text{KL}}(p_k^{\text{aug}, \delta} \|\bar{p}_k^{\text{aug}}) \leq \eta_k^{\text{dyn}}(\mathbf{x}, \mathbf{u}) + \frac{\Delta_k(\mathbf{x}, \mathbf{u})^2}{2\sigma^2} = \eta_k^{\text{aug}}(\mathbf{x}, \mathbf{u}),$$

so $p_k^{\text{aug}, \delta} \in \mathcal{B}_{\eta_k^{\text{aug}}}(\bar{p}_k^{\text{aug}})$ pointwise.

Step 2 — Pathwise feasibility. Per-step feasibility composes along the horizon, hence $p_{0:N}^{\text{aug}, \delta} \in \prod_k \mathcal{B}_{\eta_k^{\text{aug}}}(\bar{p}_k^{\text{aug}})$. For any fixed $\{\pi_k\}$,

$$D_{\text{KL}}(p_{0:N}^{\text{aug}, \delta} \|\bar{q}_{0:N}^{\text{aug}}) \leq \sup_{\{p_k^{\text{aug}} \in \mathcal{B}_{\eta_k^{\text{aug}}}\}} D_{\text{KL}}(p_{0:N}^{\text{aug}} \|\bar{q}_{0:N}^{\text{aug}}).$$

Step 3 — Inner maximization, Lagrangian, and ambiguity cost. At a fixed (\mathbf{x}, \mathbf{u}) , introduce the multiplier $\lambda_k(\mathbf{x}, \mathbf{u}) \geq 0$ for the per-step KL constraint. Writing $p_k^{(x)}$ for the next-state component and $p_k^{(r)}$ for the running-cost component, the inner Lagrangian is

$$\begin{aligned} \mathcal{L}_k(p_k^{(x)}, p_k^{(r)}, \lambda_k) &= \mathbb{E}_{p_k^{(x)}, p_k^{(r)}} [R_k + V_{k+1}(\mathbf{X}_{k+1})] \\ &- \lambda_k D_{\text{KL}}(p_k^{(x)} \|\bar{p}_k^{(x)}) - \lambda_k \mathbb{E}_{p_k^{(x)}} D_{\text{KL}}(p_k^{(r)} \|\bar{p}_k^{(r)}) + \lambda_k \eta_k^{\text{aug}}(\mathbf{x}, \mathbf{u}). \end{aligned}$$

Where, $R_k := C_{k+1} - C_k$ and $V_{k+1}(\mathbf{X}_{k+1}) \equiv \hat{c}_{k+1}(\mathbf{X}_{k+1})$. Maximizing over $p_k^{(r)}$ (pointwise) gives the running-cost dual term

$$\psi_k(\mathbf{x}, \mathbf{u}; \lambda_k) = \lambda_k \log \mathbb{E}_{\bar{p}_k^{(r)}} [e^{R_k/\lambda_k}] = \ell(\mathbf{x}, \mathbf{u}) + \frac{\sigma^2}{2\lambda_k}.$$

Then, maximizing over $p_k^{(x)}$ yields the DV dual

$$\tilde{V}_k(\mathbf{x}, \mathbf{u}; \lambda_k) = \lambda_k \log \mathbb{E}_{\bar{p}_k^{(x)}(\cdot|\mathbf{x}, \mathbf{u})} \left[\exp\left(\frac{V_{k+1}(\mathbf{X}_{k+1}) + \psi_k(\mathbf{x}, \mathbf{u}; \lambda_k)}{\lambda_k}\right) \right].$$

Hence, the *ambiguity cost* at (\mathbf{x}, \mathbf{u}) is the scalar convex program

$$\tilde{c}(\mathbf{x}, \mathbf{u}) = \min_{\lambda_k \geq 0} \left\{ \tilde{V}_k(\mathbf{x}, \mathbf{u}; \lambda_k) + \lambda_k \eta_k^{\text{aug}}(\mathbf{x}, \mathbf{u}) \right\}. \quad (11)$$

Step 4 — Policy form. Substituting the inner value into the outer step (with action prior and control cost) yields the DR-FREE Gibbs policy in (5).

Step 5 — Conclusion. Steps 1–4 establish that the augmented radius (8) ensures robustness to joint dynamics and cost perturbations, with the same implementable inner–outer structure as DR-FREE. ■

a) Discussion. The result strictly extends DR-FREE: when $\Delta_k \equiv 0$ (no cost perturbations), $\eta_k^{\text{aug}} = \eta_k^{\text{dyn}}$ and the formulation reduces to the original DR-FREE engine. When $\Delta_k > 0$, the adversary’s KL budget can be allocated across dynamics and the cost channel, yet the ambiguity cost still arises from a *one-dimensional* convex minimization over λ_k , preserving tractability and the Gibbs policy form.

Remark 4.1: The resolution engine still yields the exponential policy kernel $q_k(\mathbf{u}_k | \mathbf{x}_{k-1}) \exp(-c_k^{(u)}(\mathbf{u}_k) - \eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) - \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k))$, whose normalization gives $\pi_k^*(\mathbf{u}_k | \mathbf{x}_{k-1})$. The only change is that the adversary’s KL budget can now flow into the cost channel, which is precisely the claimed robustness.

V. EXPERIMENTS

We evaluate the proposed controller in three settings with two sources of epistemic uncertainty: (i-ii) *model learning error*, where a neural network approximates the true dynamics but inevitably incurs prediction errors (HalfCheetah, Franka simulation), and (iii) *sim-to-real transfer*, where a simulation-trained model is deployed on physical hardware (Franka Research 3). All experiments use identical greedy controllers (horizon $N=1$) for both the proposed method and baselines, isolating the effect of the robustness mechanism. The reference control prior $q_k(\mathbf{u}_k | x_k)$ is uniform over the action space.

We set $\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) = D_{\text{KL}}(\mathcal{N}(\mathbf{x}^*, \Sigma_*) \|\bar{p}_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k))$, linking robustness to model uncertainty. For sensitivity analysis, we scale by coefficient ρ : $\eta_k^{\text{eff}} = \rho \cdot \eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k)$. The baseline $\rho=0$ recovers MaxDiff, $\rho=1$ is the default, $\rho > 1$ increases conservatism. For all the experiments, ambiguity arises from mismatch between the learned neural network dynamics model $\bar{p}_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and true dynamics, and also the Gaussian noise injected into the environment reward.

Given this setup, the resulting greedy policy takes the Gibbs form

$$\pi^*(\mathbf{u}_k | \mathbf{x}_{k-1}) \propto \exp(-\eta(\mathbf{x}_{k-1}, \mathbf{u}_k) - \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)), \quad N = 1. \quad (12)$$

We report (i) episode *return*, and, for manipulation tasks, (ii) the *minimum distance-to-goal* with a 5 cm success threshold. Figures referenced below show mean \pm one standard deviation across seeds. Across all tasks the maximum episode length is set at 1000 steps. For implementation details and code refer to: <https://tinyurl.com/2v9yc9dk>.

A. HalfCheetah-v5 (MuJoCo)

Task and priors. The desired state \mathbf{x}^* encodes target forward velocity and an approximately upright pose. The controller uses the greedy policy (12) at every step.

Results. Figure 3-top compares learning curves for DR-FREE and a MaxDiff baseline. DR-FREE shows steady improvement in return with lower variance early in training. Snapshots in Fig. 3-middle and bottom panels illustrate qualitatively smoother, more stable gaits relative to the higher-variance baseline. For the HalfCheetah benchmark, after training we performed 20 evaluation rollouts: our proposed method successfully reached the goal 18 times, while the MaxDiff baseline succeeded only 6 times.

B. Franka Obstacle Task (simulation)

Environment. A tabletop manipulation scene with a vertical obstacle between the start pose and the goal (grasp-and-place of a green block).

Results. Figure 4-top shows returns increasing as the minimum distance-to-goal falls below the 5 cm threshold. The snapshots in Fig. 4b depict a learned, collision-free path that adapts around the obstacle. The controller’s cautious lateral adjustment emerges from the ambiguity cost, which is larger near contact-uncertain poses and thus biases the Gibbs kernel toward safer actions.

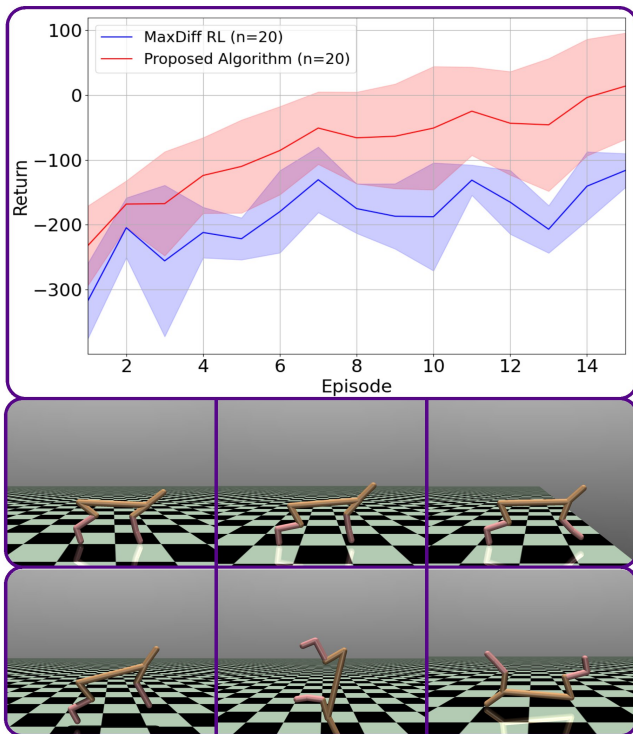


Fig. 3: HalfCheetah results. Top: training performance of the proposed algorithm vs. MaxDiff RL on HalfCheetah-v5 (mean ± 1 STD over $n = 20$ runs). Middle: rollout frames from the proposed algorithm showing a stable stride reaching the target. Bottom: rollout frames from MaxDiff RL illustrating an unstable gait leading to failure.

C. Franka Research 3 (real robot)

Deployment. We deploy the proposed algorithm on the real Franka Research 3 arm for a pick-and-place task in a cluttered tabletop scene, directly transferring the learned dynamics and executing the greedy controller without any task-specific tuning. The robot must reach a first goal position to grasp a green cube and then transport it to a second goal location. Two representative executions are shown in Fig. 5:

- 1) **No-obstacle:** After grasping the cube at the first goal, the policy selects the optimal straight path across the table and successfully places the object at the second goal; see Fig. 5-top.
- 2) **Obstacle-present:** When an obstacle is encountered en route, the policy autonomously plans a collision-free strategy by elevating the gripper above the obstruction before advancing toward and placing the object at the second goal; see Fig. 5-bottom.

Interpretation. Qualitatively distinct, yet collision-free behaviors in these settings arise directly from the optimal policy in (12). With a non-informative action prior, adaptation is governed by the Gibbs exponent: the goal-centered state prior pulls the system toward the target state \mathbf{x}^* , while the obstacle cost increases in regions near obstacles. This modulation naturally leads the policy to take the direct path when unobstructed, or to switch to a lift-over maneuver when obstacles are present, ensuring safe task execution.

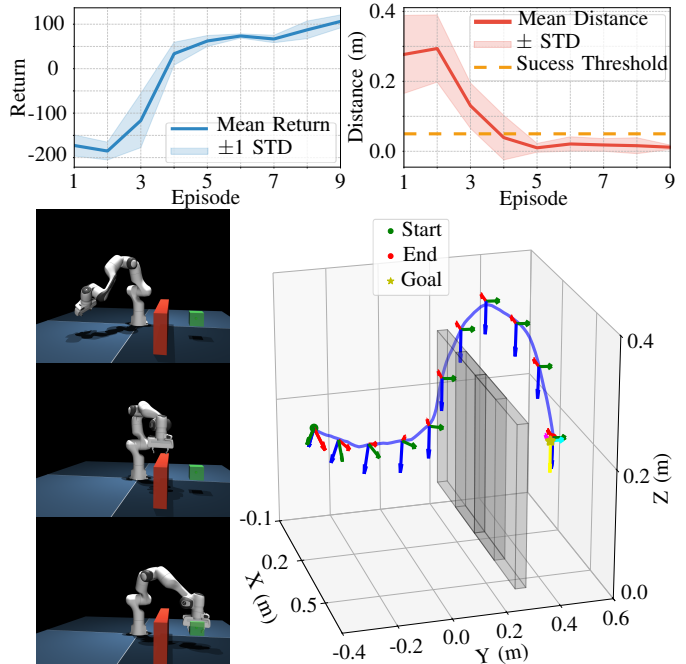


Fig. 4: Franka obstacle-avoidance results with the proposed method. Top: learning curves showing episode return and minimum distance to goal (success threshold 5 cm, mean ± 1 STD over $n = 20$ runs). Bottom: i) a rollout demonstrating collision-free manipulation around the obstacle. ii) trajectory plot with goal orientation highlighted in cyan, magenta, and yellow, illustrating efficient training and task execution.

Metric	$\rho = 0$	0.5	1	5	100	1000	2000
Normalized Avg. Episode Cost (\downarrow)	1.87	1.01	1.00	1.24	9.50	56.8	148
Standard Deviation	0.076	0.014	0.055	0.099	7.96	11.4	28.3
Success Rate (\uparrow)	100%	100%	100%	80%	60%	20%	0%

TABLE S-1: Performance across ρ over 5 trials. Episode cost is normalized to the best case ($\rho=1$). A trial succeeds if the robot reaches the goal while staying ≥ 0.07 m from all obstacle centers throughout the episode.

In Table S-1, we report a sensitivity analysis of the scaling coefficient ρ for obstacle-aware pick-and-place. As ρ increases, the learned policy becomes increasingly goal-seeking and less responsive to obstacles, leading to more frequent collisions.

VI. CONCLUSION

Building on the distributionally robust free energy principle, we introduced a modification to the maximum diffusion framework that simultaneously enables policy learning for continuous robotics control tasks and provides a priori robustness guarantees. The resulting computational model enables robust exploration by combining maximum-diffusion state generation with per-state-action KL ambiguity sets, yielding a tractable min-max objective and a robust optimal policy. The same min-max structure extends to bounded stage-cost perturbations via an augmented-state formulation.



Fig. 5: Deployment on the Franka Research 3 for cluttered tabletop pick-and-place. Top: grasp at the first goal and place at the second. Bottom: with an obstacle, the controller lifts to avoid collision and completes the placement.

Both numerical and hardware experiments illustrate that our policy computation model convincingly outperforms MaxDiff baselines in the HalfCheetah environment and enables zero-shot deployments on real hardware.

REFERENCES

- [1] J. Kober, J. A. Bagnell, and J. Peters, “Reinforcement learning in robotics: A survey,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [2] M. Kejrival, E. Kildebeck, R. Steininger, and A. Shrivastava, “Challenges, evaluation and opportunities for open-world learning,” *Nature Machine Intelligence*, vol. 6, no. 6, p. 580–588, 2024.
- [3] S. Thrun, “Probabilistic robotics,” *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.
- [4] K. Friston, “The free-energy principle: a rough guide to the brain?” *Trends in cognitive sciences*, vol. 13, no. 7, pp. 293–301, 2009.
- [5] T. J. Prescott and S. P. Wilson, “Understanding brain functional architecture through robotics,” *Science Robotics*, vol. 8, no. 78, 2023.
- [6] P. Vijayaraghavan, J. F. Queißer, S. V. Flores, and J. Tani, “Development of compositionality through interactive learning of language and action of robots,” *Science Robotics*, vol. 10, no. 98, 2025.
- [7] P. Lanillos, C. Meo, C. Pezzato, A. A. Meera, M. Baioumy, W. Ohata, A. Tschantz, B. Millidge, M. Wisse, C. L. Buckley, and J. Tani, “Active inference in robotics and artificial agents: Survey and challenges,” 2021.
- [8] M. Baioumy, P. Duckworth, B. Lacerda, and N. Hawes, “Active inference for integrated state-estimation, control, and learning,” in *IEEE International Conference on Robotics and Automation*, 2021, p. 4665–4671.
- [9] S. Wakayama and N. Ahmed, “Active inference for autonomous decision-making with contextual multi-armed bandits,” in *IEEE International Conference on Robotics and Automation*, 2023, pp. 7916–7922.
- [10] C. Pezzato, R. Ferrari, and C. Hernández, “A novel adaptive controller for robot manipulators based on active inference,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, p. 2973–2980, 2020.
- [11] A. Shafiei, H. Jesawada, K. Friston, and G. Russo, “Distributionally robust free energy principle for decision-making,” *Nature Communications*, vol. 17, no. 1, p. 707, 2025.
- [12] T. A. Berrueta, A. Pinosky, and T. D. Murphey, “Maximum diffusion reinforcement learning,” *Nature Machine Intelligence*, vol. 6, no. 5, pp. 504–514, 2024.
- [13] B. Eysenbach and S. Levine, “Maximum entropy rl (provably) solves some robust rl problems,” in *International Conference on Learning Representations*, 2022.
- [14] E. Hazan, S. Kakade, K. Singh, and A. Van Soest, “Provably efficient maximum entropy exploration,” in *International Conference on Machine Learning*, 2019, pp. 2681–2691.
- [15] T. Haamoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International conference on machine learning*, 2018, pp. 1861–1870.
- [16] L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, “Robust adversarial reinforcement learning,” in *International conference on machine learning*, 2017, pp. 2817–2826.
- [17] A. Rajeswaran, S. Ghotra, B. Ravindran, and S. Levine, “Epopt: Learning robust neural network policies using model ensembles,” in *International Conference on Learning Representations*, 2017.
- [18] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Sim-to-real transfer of robotic control with dynamics randomization,” in *IEEE international conference on robotics and automation*, 2018, pp. 3803–3810.
- [19] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *IEEE/RSJ international conference on intelligent robots and systems*, 2017, pp. 23–30.
- [20] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, “Learning dexterous in-hand manipulation,” *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
- [21] G. Iyengar, “Robust dynamic programming,” *Math. Oper. Res.*, vol. 30, pp. 257–280, 05 2005.
- [22] A. Nilim and L. El Ghaoui, “Robust control of markov decision processes with uncertain transition matrices,” *Operations Research*, vol. 53, no. 5, pp. 780–798, 2005.
- [23] E. Derman and S. Mannor, “Distributional robustness and regularization in reinforcement learning,” *arXiv preprint arXiv:2003.02894*, 2020.
- [24] H. Jesawada, A. Acernese, D. Del Vecchio, G. Russo, and C. D. Vecchio, “Dr-pets: Learning-based control with planning in adversarial environments,” *IEEE Control Systems Letters*, vol. 9, pp. 2006–2011, 2025.
- [25] E. Todorov, “Linearly-solvable Markov decision problems,” in *Advances in Neural Information Processing Systems*, vol. 19, 2007, pp. 1369–1376.
- [26] D. Gagliardi and G. Russo, “On a probabilistic approach to synthesize control policies from example datasets,” *Automatica*, vol. 137, p. 110121, 2022.
- [27] G. Williams, A. Aldrich, and E. A. Theodorou, “Model predictive path integral control: From theory to parallel computation,” *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [28] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, “Information-theoretic model predictive control: Theory and applications to autonomous driving,” *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [29] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep reinforcement learning in a handful of trials using probabilistic dynamics models,” *Advances in neural information processing systems*, vol. 31, 2018.
- [30] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to trust your model: Model-based policy optimization,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [31] Y. Yao, L. Xiao, Z. An, W. Zhang, and D. Luo, “Sample efficient reinforcement learning via model-ensemble exploration and exploitation,” in *IEEE International Conference on Robotics and Automation*, 2021, pp. 4202–4208.
- [32] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection,” *The International journal of robotics research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [33] A. Kumar, Z. Li, J. Zeng, D. Pathak, K. Sreenath, and J. Malik, “Adapting rapid motor adaptation for bipedal robots,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 1161–1168.
- [34] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.
- [35] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *IEEE/RSJ international conference on intelligent robots and systems*, 2012, pp. 5026–5033.
- [36] E. Garrabe, H. Jesawada, C. Del Vecchio, and G. Russo, “On convex data-driven inverse optimal control for nonlinear, non-stationary and stochastic systems,” *Automatica*, vol. 173, p. 112015, 2025.