

ADAPT: Zero-Shot Adaptive Policy Transfer for Stochastic Dynamical Systems

James Harrison¹, Animesh Garg², Boris Ivanovic², Yuke Zhu², Silvio Savarese², Li Fei-Fei², Marco Pavone³

Abstract Model-free policy learning has enabled good performance on complex tasks that were previously intractable with traditional control techniques. However, this comes at the cost of requiring a perfectly accurate model for training. This is infeasible due to the very high sample complexity of model-free methods preventing training on the target system. This renders such methods unsuitable for physical systems. Model mismatch due to dynamics parameter differences and unmodeled dynamics error may cause suboptimal or unsafe behavior upon direct transfer. We introduce the Adaptive Policy Transfer for Stochastic Dynamics (ADAPT) algorithm that achieves provably safe and robust, dynamically-feasible zero-shot transfer of RL-policies to new domains with dynamics error. ADAPT combines the strengths of offline policy learning in a black-box source simulator with online tube-based MPC to attenuate bounded dynamics mismatch between the source and target dynamics. ADAPT allows online transfer of policies, trained solely in a simulation offline, to a family of unknown targets without fine-tuning. We also formally show that (i) ADAPT guarantees bounded state and control deviation through state-action tubes under relatively weak technical assumptions and, (ii) ADAPT results in a bounded loss of reward accumulation in case of direct transfer with ADAPT as compared to the policy trained and evaluated in the source environment. We evaluate ADAPT on 2 continuous, non-holonomic simulated dynamical systems with 4 different disturbance models, and find that ADAPT performs between 50%-300% better on mean reward accrual than direct policy transfer.

James Harrison
Department of Mechanical Engineering, Stanford University, Stanford, CA 94305
e-mail: jharrison@stanford.edu

Animesh Garg, Boris Ivanovic, Yuke Zhu, Li Fei-Fei, Silvio Savarese
Department of Computer Science, Stanford University, Stanford, CA 94305
e-mail: [{garg, borisi, yukez, feifeili, ssilvio}@cs.stanford.edu"> {garg, borisi, yukez, feifeili, ssilvio}@cs.stanford.edu](mailto)

Marco Pavone
Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305
e-mail: pavone@stanford.edu

1 Introduction

Deep reinforcement learning (RL) has achieved remarkable advances in sequential decision making in recent years, often outperforming humans on tasks such as Atari games [17]. However, model-free variants of deep RL are not directly applicable to physical systems because they exhibit poor sample complexity, often requiring millions of training examples on an accurate model of the environment. One approach to using model-free RL methods on robotic systems is thus to train in a relatively accurate simulator (a source domain), and transfer the policy to the physical robot (a target domain). This naive transfer may, in practice, perform arbitrarily badly and so online fine-tuning may be performed [1]. During this fine-tuning, the robot may behave unsafely however, and so it is desirable for a system to be able to train in a simulator with slight model inaccuracies but still be able to perform well on the target system on the first iteration. We refer to this as the *zero-shot policy transfer problem*.

The zero-shot transfer problem involves training a policy on a system possessing different dynamics than the target system, and evaluating performance as the average initial return in target domain *without* training in the target domain. This problem is challenging for robotic systems since simplified simulated models may not always accurately capture all relevant dynamics phenomena, such as friction, structural compliance, turbulence and so on, as well as parametric uncertainty in the model. In spite of the renewed focus on this problem, few studies in deep policy adaptation offer insightful analysis or guarantees regarding feasibility, safety, and robustness in policy transfer.

In this paper, we introduce a new algorithm which we refer to as ADAPT, that achieves provably safe and robust, dynamically-feasible zero-shot direct transfer of RL policies to new domains with dynamics mismatch. The key insight here is to leverage the global optimality of learned policy with local stabilization from MPC based methods to enable dynamic feasibility, thereby building on strengths of two different methods. In the offline stage, ADAPT first computes a nominal trajectory (without disturbance) by executing the learned policy on the simulator dynamics. Then in the online stage, ADAPT adapts the nominal trajectory to the target dynamics with an auxiliary MPC controller.

Statement of Contributions

1. We develop the ADAPT algorithm, which allows online transfer of policy trained solely in a simulation offline, to a family of unknown targets without fine-tuning.
2. We also formally show that (i) ADAPT guarantees state and control safety through state-action tubes under the assumption of Lipschitz continuity of the divergence in dynamics and, (ii) ADAPT results in a bounded loss of reward accumulation in case of direct transfer with ADAPT as compared to a policy trained only on target.
3. We evaluate ADAPT on two continuous, non-holonomic simulated dynamical systems with four different disturbance models, and find that ADAPT performs between 50%-300% better on mean reward accrual than direct policy transfer as compared to mean reward.

Organization This paper is structured as follows. In Section 2 we review related work in robust control, robust reinforcement learning, and transfer learning. In Section 3 we formally state the policy transfer problem. In Section 4 we present ADAPT and discuss algorithmic design features. In Section 5 we prove the accrued reward for ADAPT is lower bounded. In Section 6 we present experimental results on a simulated car environment and a two-link robotic manipulator, as well as present results for ADAPT with robust policy learning methods. Finally, in Section 7 we draw conclusions and discuss future directions.

2 Related Work and Background

A plethora of work in both learning and control theory has addressed the problem of varying system dynamics, especially in the context of safe policy transfer and robust control.

Transfer in reinforcement learning The problem of high sample complexity in reinforcement learning has generated considerable interest in policy transfer. Taylor et al. provide an excellent review of approaches to the transfer learning problem [28]. A series of approaches focused on reducing the number of rollouts performed on a physical robot, by alternating between policy improvement in simulation and physical rollouts [1], [13]. In those works, a time-dependent term is added to the dynamics after each physical rollout to account for unmodeled error. This approach, however, does not address robustness in the initial transfer, and the system could sustain or cause damage before the online learning model converges.

The EPOPT algorithm [23] randomly samples dynamics parameters from a Gaussian distribution prior to each training run, and optimizes the reward for the worst-performing ϵ -fraction of dynamics parameters. However, it is not clear how robust it is against disturbances not explicitly experienced in training. This approach is conceptually similar to that in [19], in which more traditional trajectory optimization methods are used with an ensemble of models to increase robustness. Similarly, [14] and [22] use adversarial disturbances instead of random dynamics parameters for robust policy training. Tobin et al. [29] and Peng et al. [21] randomize visual inputs and dynamics parameters respectively. Bousmalis et al. [2] meanwhile adapt rendered visual inputs to reality using a framework based on generative adversarial networks, as opposed to strictly randomizing them. While this may improve adaptation to a target environment in which these parameters are varied, this may not improve performance on dynamics changes outside of those varied; in effect, it does not mitigate errors due to the “unknown unknowns”.

Christiano et al. [4] approach the transfer problem by training an inverse dynamics model on the target system and generating a nominal trajectory of states. The inverse dynamics model then generates actions to connect these states. However, there are no guarantees that an action exists in the target dynamics to connect two learned adjacent states. Moreover, this requires training on the target environment; in this work we

consider zero-shot learning where this is not possible. Recently, the problem of transfer has been addressed in part by rapid test adaptation [6], [24]. These approaches have focused on training modular networks that have both “task-specific” and “robot-specific” modules. This then allows the task-specific module to be efficiently swapped out and retrained. However, it is unclear how error in the learned model affects these methods.

In this work we aim to perform zero-shot policy transfer, and thus efficient model-based approaches are not directly applicable. However, our approach uses an auxiliary control scheme that leverages model learning for an approximate dynamics model. When online learning is possible, sample-efficient model-based reinforcement learning approaches can dramatically improve sample complexity, largely by leveraging tools from planning and optimal control [11]. However, these models require an accurate estimate of the true system dynamics in order to learn an effective policy. A variety of model classes have been used to represent system dynamics, such as neural networks [9], Gaussian processes [5], and local linear models [8], [13].

Robust control Trajectory optimization methods have been widely used for robotic control [27]. Among these optimization methods, model predictive control (MPC) is a class of online methods that perform trajectory optimization in a receding-horizon fashion [20]. This receding-horizon approach, in which a finite-horizon, open-loop trajectory optimization problem is continuously re-solved, results in an online control algorithm that is robust to disturbances. Several works have attempted to combine trajectory optimization methods with dynamics learning [16] and policy learning [10]. In this work, we develop an auxiliary robust MPC-based controller to guarantee robustness and performance for learned policies. Our method combines the strengths of deep policy networks [25] and tube-based MPC [15] to offer a controller with good performance as well as robustness guarantees.

3 Problem Setup and Preliminaries

Consider a finite-horizon Markov Decision Process (\mathcal{M}) defined as a tuple $\mathcal{M} : \langle \mathcal{S}, \mathcal{A}, p, r, T \rangle$. Here \mathcal{S} and \mathcal{A} represent continuous, bounded state and action spaces for the agent, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function that maps a state-action tuple to a scalar, and T is the problem horizon. Finally, $p : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the transition distribution that captures the state transition dynamics in the environment and is a distribution over states conditioned on the previous state and action. The goal is to find a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the expected cumulative reward over the choice of policy:

$$\pi^*(s) = \operatorname{argmax}_{\pi(s)} \mathbb{E} \left[\sum_{t=0}^T r(s_t, a_t) \right]. \quad (1)$$

The above reflects a standard setup for policy optimization in continuous state and action spaces. In this work, we are interested in the case in which we only have an approximately correct environment, which we refer to as the source environment

(e.g. a physics simulator). We may sample this simulator an unlimited number of times, but we wish to maximize performance on the first execution in a target environment. Without any assumptions on the correctness of the simulator, this problem is of course intractable as the two sets of dynamics may be arbitrarily different. However, relatively loose assumptions about the correctness of the simulator are very reasonable, based on the modeling fidelity of the simulator. We assume the simulator (denoted \mathcal{M}_S) has deterministic, twice continuously-differentiable dynamics $s_{t+1} = f(s_t, a_t)$. Then, let the dynamics of the target environment (denoted \mathcal{M}_T) be denoted $s_{t+1} = f(s_t, a_t) + w_t$, for iid additive noise w_t with compact, convex support \mathcal{W} that contains the origin. Generally, the noise distribution may be state and action dependent, so this formulation reduces to standard formulations in both robust and stochastic control [32]. We assume all other components of the MDPs defining the source and target environments are the same (e.g. reward function). Finally, we assume the reward function r is Lipschitz continuous, an assumption that we discuss in more detail in section 5. Based on the above definitions, we can now state the problem we aim to solve.

Problem Statement Given the simulator dynamics and the problem defined by the MDP \mathcal{M}_S , we wish to learn a policy to maximize the reward accrued during operation in the target system, \mathcal{M}_T . Formally, if we write the realization of the disturbance at time t as \tilde{w}_t , we wish to solve the problem:

$$\begin{aligned} \max_{\{a_t\}_{t=0}^T} \quad & \mathbb{E} \left[\sum_{t=0}^T r(s_t, a_t) \right] \\ \text{s.t.} \quad & s_{t+1} = f(s_t, a_t) + \tilde{w}_t, \text{ and } s_t \in \mathcal{S}, a_t \in \mathcal{A} \quad \forall t \in [0, T], \end{aligned} \quad (2)$$

while only having access to the simulator, \mathcal{M}_S , for training.

4 ADAPT: Adaptive Policy Transfer for Stochastic Dynamics

In this section we present the ADAPT algorithm for zero-shot transfer. A high level view of the algorithm is presented in Algorithm 1. First, we assume that a policy is trained in simulation. Our approach is to first compute a nominal trajectory (without disturbance) by continuously executing the learned policy on the simulator dynamics. Then, when transferred to the target environment, we use an auxiliary model predictive control-based (MPC) controller to stabilize around this nominal trajectory. In this work, we use a reward formulation for operation in the primary environment (i.e., the aim is to maximize reward), and a cost formulation for the auxiliary controller (i.e., the aim is to minimize cost to thus minimize deviation from the nominal trajectory). This is in part to disambiguate the distinction between the primary and auxiliary optimization problems.

Policy Training We use model-free policy optimization on the black-box simulated model. Our theoretical guarantees rely on the auxiliary controller avoiding saturation. Therefore, if a policy operates near the limits of its control authority and thus the

auxiliary controller saturates when used on the target environment, this policy is trained using restricted state and action spaces $\mathcal{S}' \subseteq \mathcal{S}$, $\mathcal{A}' \subseteq \mathcal{A}$. We let \mathcal{M}' denote an MDP with restricted state and action spaces. This follows the approach of [15], where it is used to prevent auxiliary controller saturation. Intuitively, restricting the state and action space ensures any nominal trajectory in those spaces can be stabilized by the auxiliary controller. Therefore, if saturation is rare, restricting these sets is unnecessary.

ADAPT is invariant to the choice of policy optimization method. During online operation, a nominal trajectory $\tau = \{(\bar{s}_t, \bar{a}_t)\}_{t=0}^T$ is generated by rolling out the policy on the simulator dynamics, $\mathcal{M}_{\mathcal{S}}$. The auxiliary controller then tracks this trajectory in the target environment.

Approximate Dynamics Model Because the model of the simulator is treated as a black-box, it is impractical to use for the auxiliary controller in an optimal control framework. As such, we rely on an approximate model of the dynamics, separate from the simulator dynamics f , which we refer to as \hat{f} . The specific representation of the model (e.g. linear model, feedforward neural network, etc.) depends on both the accuracy required as well as the method used to solve the auxiliary control problem. This model may be either learned from the simulator, or based on prior knowledge. A substantial body of literature exists on dynamics model learning from black-box systems [18]. Alternatively, this model may be based on external knowledge, either from learning a dynamics model in advance from the target system or from, for example, a physical model of the system.

Auxiliary MPC Controller Our auxiliary nonlinear MPC controller is based on that of [15]. Specifically, we write the auxiliary control problem:

$$\begin{aligned} \min_{\{a_k\}_{k=t}^{t+N}} & \sum_{k=t}^{t+N} (s_k - \bar{s}_k)^T Q_k (s_k - \bar{s}_k) + (a_k - \bar{a}_k)^T R_k (a_k - \bar{a}_k) \\ \text{s.t.} & \quad s_{k+1} = \hat{f}(s_k, a_k), \text{ and } s_k \in \mathcal{S}, a_k \in \mathcal{A} \quad \forall k \in [t, t+N], \end{aligned} \quad (3)$$

where N is the MPC horizon, Q_k and R_k are positive definite cost matrices for the state deviation and control deviation respectively, and \hat{f} is the approximate dynamics model. In some cases, this problem is convex, but generally it may not be. In our experiments, this optimization problem is solved with iterative relinearization based on [30]. However, whereas they iteratively linearize the nonlinear optimal control problem and solve an LQR problem over the full horizon of the problem, we explicitly solve the problem over the MPC horizon. We do not consider terminal state costs or constraints. This formulation of the auxiliary controller by [15] allows us to guarantee, under our assumptions, that our true state stays in a tube around the nominal trajectory, where the tube is defined by level sets of the value function (the details of this are addressed in Section 5).

The solution to the MPC problem is iterative. First, we linearize around the nominal trajectory τ . We introduce the notation $\{(\hat{s}_k, \hat{a}_k)\}_{k=t}^{k=t+N}$, which is the solution for the last iteration. These are initialized as $\hat{s}_t \leftarrow \bar{s}_t$ and $\hat{a}_t \leftarrow \bar{a}_t$. Then, we introduce the deviations from this solution as

$$\delta s_t = s_t - \hat{s}_t, \quad \delta a_t = a_t - \hat{a}_t. \quad (4)$$

Algorithm 1 Adaptive Policy Transfer for Stochastic Dynamics (ADAPT)

Input: Source Env: \mathcal{M}_S , Target Env: \mathcal{M}_T , Initial State: s_0

Offline:

- 1: $\mathcal{A}', \mathcal{S}' \leftarrow \text{bound_set}(\mathcal{A}, \mathcal{S})$ // Calculate constrained state & action space
- 2: $\pi \leftarrow \text{policy_opt}(\mathcal{M}'_S)$ // Train a policy for \mathcal{M}'_S using constrained $\mathcal{S}', \mathcal{A}'$
- 3: $\hat{f} \leftarrow \text{fit_dynamics}(\mathcal{M}_S)$ // Fit Dynamics for \mathcal{M}_S

Online:

- 4: $\tau \leftarrow \text{rollout}(s_0, \pi, \mathcal{M}_S, T)$ // Roll out π on \mathcal{M}_S to get nominal trajectory
- 5: $s \leftarrow s_0$
- 6: **for** $t \in [0, T]$ **do**
- 7: $a \leftarrow \text{aux_MPC}(s, \tau, \hat{f}, \tau, N)$ // NMPC with iterative linearization
- 8: $s \leftarrow f(s, a) + w$ // Rollout the first step of action seq. on \mathcal{M}_T
- 9: **end for**

Then, taking the linearization of our dynamics

$$A_t = \left. \frac{\partial \hat{f}}{\partial s_t} \right|_{s_t = \hat{s}_t, a_t = \hat{a}_t} \quad B_t = \left. \frac{\partial \hat{f}}{\partial a_t} \right|_{s_t = \hat{s}_t, a_t = \hat{a}_t}, \quad (5)$$

we can rewrite the MPC problem as:

$$\begin{aligned} & \min_{\{\delta a_k\}_{k=t}^{t+N}} \sum_{k=t}^{t+N} (\delta s_k + \hat{s}_k - \bar{s}_k)^T Q_k (\delta s_k + \hat{s}_k - \bar{s}_k) + (\delta a_k + \hat{a}_k - \bar{a}_k)^T R_k (\delta a_k + \hat{a}_k - \bar{a}_k) \\ & \text{s.t.} \quad \delta s_{k+1} = A_k \delta s_k + B_k \delta a_k, \text{ and } \delta s_k + \hat{s}_k \in \mathcal{S}, \delta a_k + \hat{a}_k \in \mathcal{A}, \quad \forall k \in [t, t+N]. \end{aligned} \quad (6)$$

Note that the optimization is over the action deviations $\{\delta a_k\}_{k=t}^{t+N}$. Once this problem is solved, we use the update rule $\hat{s}_t \leftarrow \hat{s}_t + \delta s_t$, $\hat{a}_t \leftarrow \hat{a}_t + \delta a_t$. The dynamics are then relinearized, and this is iterated until convergence. Because we use iterative linearization to solve the nonlinear program, it is necessary to choose a dynamics representation \hat{f} that is efficiently linearizable. In our experiments, we use an analytical nonlinear dynamics representation for which the linearization can be computed analytically (see [31] for details), as well as fit a time-varying linear model. Choices such as, e.g., a Gaussian process representation, may be expensive to linearize.

5 ADAPT: Analysis

The following section develops the main theoretical analysis of this study. We will first show that ADAPT results in bounded deviation from the nominal trajectory τ under a set of technical assumptions. This result is then used to show that the deviation between cumulative reward of the realized rollout on the target system and the cumulative reward of the nominal trajectory on the source environment, is upper bounded. This is to say, the decrease in performance below the ideal case is bounded.

5.1 Safety Analysis in ADAPT

Using the notation from Eq (3), let us denote the solution at time k as $C_N^*(s_k, k)$ for MPC horizon N . This is the minimum cost associated with the finite horizon problem that is solved iteratively in the MPC framework. Note that this problem is solved with the approximate dynamics model; in the case where the approximate dynamics model exactly matches the target environment model, the solution to this problem would have value zero as the trajectory would be tracked exactly. We denote by $\kappa_N(s_k)$ the action at time k from the solution to the MPC problem. Then, let $L_d(k) \triangleq \{s \mid C_N^*(s, k) \leq d\}$ denote the level set of the cost function for some value $d \in (0, c)$ (for some constant c ; see [15]) at time k .

We assume the error between approximate dynamics representation \hat{f} and the simulator dynamics f is outer approximated by a compact, convex set \mathcal{D} that contains the origin. Therefore, for all state, action pairs $(s, a) \in \mathcal{S} \times \mathcal{A}$, $f(s, a) - \hat{f}(s, a) \in \mathcal{D}$. In the case where the state and action spaces are bounded, there always exists an outer approximation which satisfies this assumption. However, in practice, it is likely considerably smaller than this worst case.

Let $\mathcal{T}_s(s_0) \triangleq \{L_d(k) \mid k \in \mathbb{Z}_{\geq 0}\}$ denote a state tube defined by the time-dependent level sets of the auxiliary cost function. We may now state our first result, noting that the auxiliary stabilizing policy κ_N is the result of the MPC optimization problem relying solely on the approximate dynamics \hat{f} .

Theorem 1. *Every state trajectory $\{s_t\}_{t=0}^T$ generated by the target dynamics $s_{t+1} = f(s_t, \kappa_N(s_t)) + w_t$ with initial state s_0 , lies in the state tube $\mathcal{T}_s(s_0)$.*

Proof. Note that $\mathcal{W} + \mathcal{D}$, where the addition denotes a Minkowski sum, is compact, convex, and contains the origin. Then, the result follows from Theorem 1 of [15] by replacing the set of disturbances (which the authors refer to as \mathbb{W}) with $\mathcal{W} + \mathcal{D}$. \square

The above result combined with Proposition 2i of [15], which shows that for some constant c_1 , $C_N^*(s_k, k) \geq c_1 \|s_k - \bar{s}_k\|^2$, gives insight into the safety of ADAPT. In particular, note that for an arbitrarily long trajectory, the realized trajectory stays in a region around the nominal trajectory despite using an inaccurate dynamics representation in the MPC optimization problem. While this result shows that the deviation from the nominal trajectory is bounded, it does not allow construction of explicit tubes in the state space, and thus can not be used directly for guarantees on obstacle avoidance. Recent work by Singh et al. [26] establishes tubes of this form, and this is thus a promising extension of the ADAPT framework.

5.2 Robustness Analysis in ADAPT

We will now show that due to the boundedness of state deviation, the deviation in the total accrued reward over a rollout on the target system is bounded. Let $V_S^\pi(s)$ and $V_T^\kappa(s)$ denote the value functions associated with some state s and the primary

policy executed on the source environment, and the ADAPT policy on the secondary environment respectively.

Theorem 2. *Under the technical assumptions made in Section 3 and 5.1, $|V_T^K(s_0) - V_S^\pi(s_0)| \leq c_2 \sum_{t=0}^T \sqrt{C_N^*(s_t, t)}$, where c_2 is some constant and $s_{t+1} = f(s_t, \kappa_N(s_t)) + w_t$.*

Proof. First, note $|V_\kappa(s_0) - V_\pi(s_0)| \leq \sum_{t=0}^T |r(s_t, \kappa_N(s_t)) - r(\bar{s}_t, \pi(\bar{s}_t))|$, where $s_{t+1} = f(s_t, \kappa_N(s_t)) + w_t$ and $\bar{s}_{t+1} = f(\bar{s}_t, \pi(\bar{s}_t))$. Additionally, letting $a = \kappa_N(s)$ and $\bar{a} = \pi(\bar{s})$, note that similarly to Proposition 2i of [15], we can establish a bound on the action deviation from the nominal trajectory in terms of the auxiliary cost function, $C^*(s_t, t) \geq c_3 \|a_t - \bar{a}_t\|^2$ for all t (where the norm is in the Euclidean sense), by taking c_3 as the minimum eigenvalue of R_t . By the Lipschitz continuity of the reward function, and writing the Lipschitz constant of the reward function L_r , we have

$$|r(s, a) - r(\bar{s}, \bar{a})| \leq L_r (\|s - \bar{s}\| + \|a - \bar{a}\|). \quad (7)$$

Then, noting that the quadratic auxiliary cost function C_N^* is always positive, the result is proved by applying Proposition 2i of [15] and the bound on action deviation from the nominal to the right hand side of Equation 7. \square

This result may then be restated in terms of the disturbance sets. Let $\|\mathcal{W} + \mathcal{D}\| \triangleq \max_{w \in \mathcal{W}, d \in \mathcal{D}} \|w + d\|$.

Theorem 3. *Under the same technical assumptions as Theorem 2, the following inequality holds for some constant $c_4 > 0$:*

$$|V_T^K(s_0) - V_S^\pi(s_0)| \leq c_4 T \sqrt{\|\mathcal{W} + \mathcal{D}\|} \quad (8)$$

Proof. The result follows from combining Theorem 2 with Proposition 4ii of [15]. \square

These results shows that along with guarantees on spatial deviation from the nominal trajectory, we may also establish bounds on the accrued reward relative to what is received with the nominal policy in the source environment, in effect demonstrating that zero-shot transfer is possible. The Lipschitz continuity of the reward function is essential to this result, and this illustrates several aspects of the policy transfer problem.

The ADAPT algorithm is based on tracking a nominal rollout in simulation. Critical in the success of this approach is gradual variation of the reward function. Sparse reward structures are likely to fail with this approach to transfer, as tracking the nominal trajectory, even relatively closely, may result in poor reward. On the other hand, a slowly varying reward function, even if tracked relatively roughly may result in accrued reward close to the nominal rollout on the source environment.

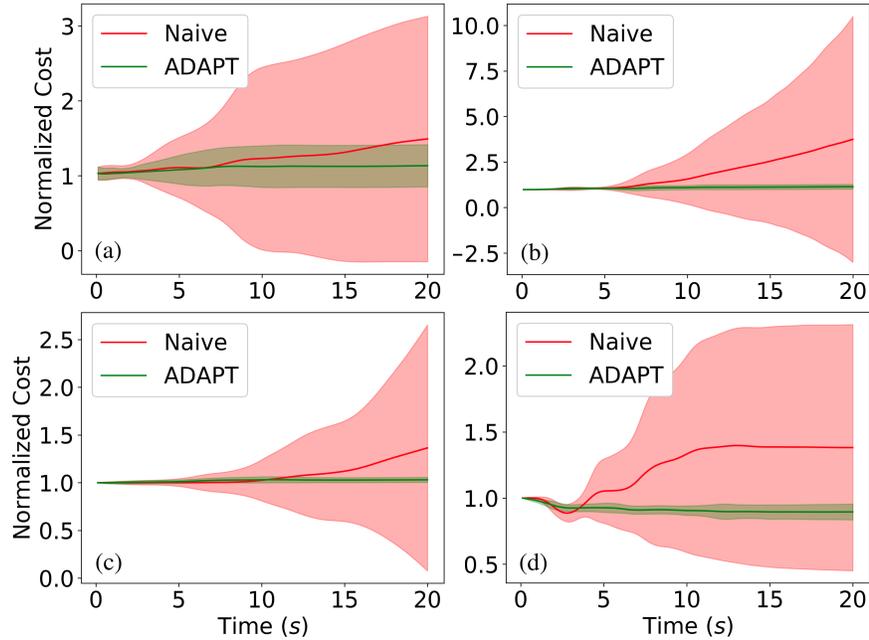


Fig. 1: Mean cumulative cost over the length of an episode for 50 episodes on the kinematic car environment. The confidence intervals are standard error. The costs are normalized to the cost of the naive policy being rolled out on the simulated environment from the same initial state, to allow more direct comparison across episodes. The *naive* rollout is the nominal policy executed on the target environment. The disturbances tested are **a)** a hill landscape, **b)** additive control error, **c)** process noise, and **d)** dynamics parameter error.

6 Experimental Evaluation

We implemented ADAPT on a nonlinear, non-holonomic kinematic car model with a 5-dimensional state space as well as on the `Reacher` environment in OpenAI’s Gym [3]. We train policies using Trust Region Policy Optimization (TRPO) [25]. The policy is parameterized as a neural network with two hidden layers, each with 64 units and ReLU nonlinearities. In all of our experiments, we report *normalized cost*. This is the cost (negative reward) realized by a trial in the target environment, divided by the cost of the nominal policy rolled out on the simulated environment from the same initial state. This allows more direct comparison between episodes for environments with stochastic initial states. We generally compare the *naive* trial, which is the nominal policy rolled out on the target environment (e.g., standard transfer with no adaptation) to ADAPT.

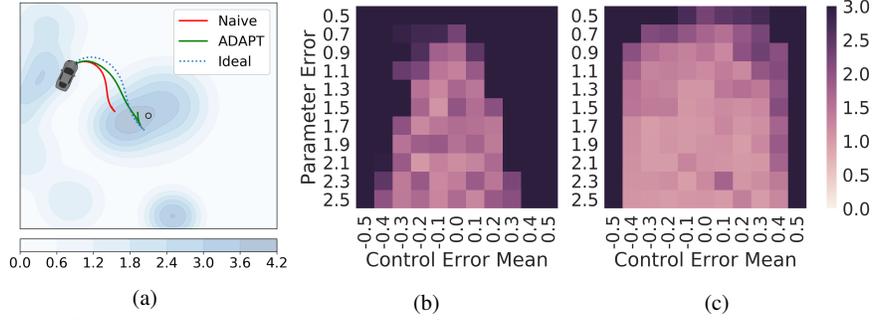


Fig. 2: **a)** The car environment with the paths for the ideal case (nominal policy on simulated environment), the naive case (nominal policy on the target environment), and the ADAPT case (ADAPT on the target environment). The contour plot shows the height of the added hills. Figures **(b)** and **(c)** show the normalized cost for varying disturbances due to additive control error and dynamics parameter error for **b)** the naive case and **c)** ADAPT (lower is better). In addition to the listed disturbances, disturbances due to hills are also added for all trials. Each grid cell is the mean of 50 trials.

6.1 Environment I: 5-D Car

We implemented ADAPT on a nonlinear, nonholonomic 5-dimensional kinematic car model that has been used previously in the motion planning literature [31]. Specifically, the car has state $s = [x, y, \theta, v, \kappa]^T$, where x and y denote coordinates in the plane, θ denotes heading angle, v denotes speed, and κ denotes trajectory curvature. The system has dynamics $\dot{s} = [v \cos \theta, v \sin \theta, v \kappa, a_v, a_\kappa]$, where $a_v \in [-2, 2]$ and $a_\kappa \in [-0.5, 0.5]$ are the controlled acceleration and curvature derivative. The policy is trained to minimize the quadratic cost $L(s, a) = \sum_{t=0}^T \ell(s_t, a_t)$, where $\ell(s_t, a_t) = x_t^2 + y_t^2 + a_{v,t}^2 + a_{\kappa,t}^2$, which results in policies that drive to the origin. In each trial, the vehicle is initialized in a random state, with position $x, y \in [-5, 5]$, with random heading and zero velocity and curvature.

Our auxiliary controller used an MPC horizon of 2 seconds (20 timesteps). Our state deviation penalty matrix, Q , has value 1 along the diagonal for the position terms, and zero elsewhere. Thus, the MPC controller penalizes only deviation in position. The matrix R had small terms (10^{-3}) along the diagonal to slightly penalize control deviations. In practice, this mostly acts as a small regularizing term to prevent large oscillatory control inputs by the auxiliary controller. The behavior of the auxiliary controller is dependent on the matrices Q and R , but in practice good performance may be achieved across environments with fixed values. Because of the relatively high quadratic penalty on control in policy training, the nominal policy rarely approaches the control limits. Thus, we can set $\mathcal{A}' = \mathcal{A}$, and we set $\mathcal{S}' = \mathcal{S}$. For our dynamics model, we use the linearization reported in [31].

6.2 Disturbance Models

We investigate four disturbance types:

1. **Environmental Uncertainty:** We add randomly-generated hills to the target environment such that the car experiences accelerations due to gravity. This noise is therefore state-dependent. Figure 2a shows a randomly generated landscape. We randomly sample 20 hills in the workspace, each of which is circular and has varying radius and height. The vehicle experiences an additive longitudinal acceleration proportional to the landscape slope at its current location, and no lateral acceleration.
2. **Control noise:** Nonzero-mean additive control error drawn from a uniform distribution.
3. **Process noise:** Additive, zero-mean noise added to the state. Disturbances are drawn from a uniform distribution.
4. **Dynamics parameter error:** We add a scaling factor γ to the control of $\dot{\kappa}$, such that $\dot{\kappa} = \gamma a_{\kappa}$.

For the last three, the noise terms were drawn i.i.d. from a uniform distribution at each time t . These disturbances were investigated both independently (Figure 1) and simultaneously (Figure 2). Figure 1 shows the normalized cost of the naive transfer and ADAPT for each of the four disturbances individually.

In our experiments, ADAPT substantially outperforms naive transfer, achieving normalized costs 1.5-5x smaller. Additionally, the variance of the naive transfer is considerably higher, whereas the realized cost for ADAPT is clustered relatively tightly around one (e.g., approximately equal cost to the ideal case). In Figure 1d, the normalized cost of ADAPT is actually below one, implying that the transferred policy performs better than the ideal policy. In fact, this is because the dynamics parameter error in this trial results in oversteer, and so the agent accumulates less cost to turn to face the goal than in the nominal environment. Thus, pointing toward the goal is more “cost-efficient” in the target environment. The performance of direct transfer and ADAPT with varying parameter error may be seen in Figure 2b and Figure 2c. In Figure 2a, a case is presented where the direct policy transfer fails to make it up a hill, whereas the ADAPT policy tracks the nominal trajectory well.

6.3 ADAPT with Robust Offline Policy

Whereas ADAPT’s approach to policy transfer relies primarily on stabilization in the target environment, recent work has focused on training robust policies in the source domain, and then performing direct transfer. In the EPOPT policy training framework [23], an agent is trained over a family of MDPs in which model parameters are drawn from distributions before each training rollout. Then, a Conditional Value-at-Risk (CVaR) objective function is optimized as opposed to an expectation over all training runs. We apply ADAPT on top of an EPOPT-1 policy (equivalent to optimizing expected reward, with model parameters varying), and find that for disturbances explicitly varied during training, the performance of EPOPT-only transfer and ADAPT are comparable. We add parameters γ_i to the state derivative as follows: $\dot{s} = [\gamma_1 v \cos \gamma_2 \theta, \gamma_1 v \sin \gamma_2 \theta, \gamma_1 v \gamma_3 \kappa, \gamma_4 a_v, \gamma_5 a_{\kappa}]$. Each of these γ_i are drawn

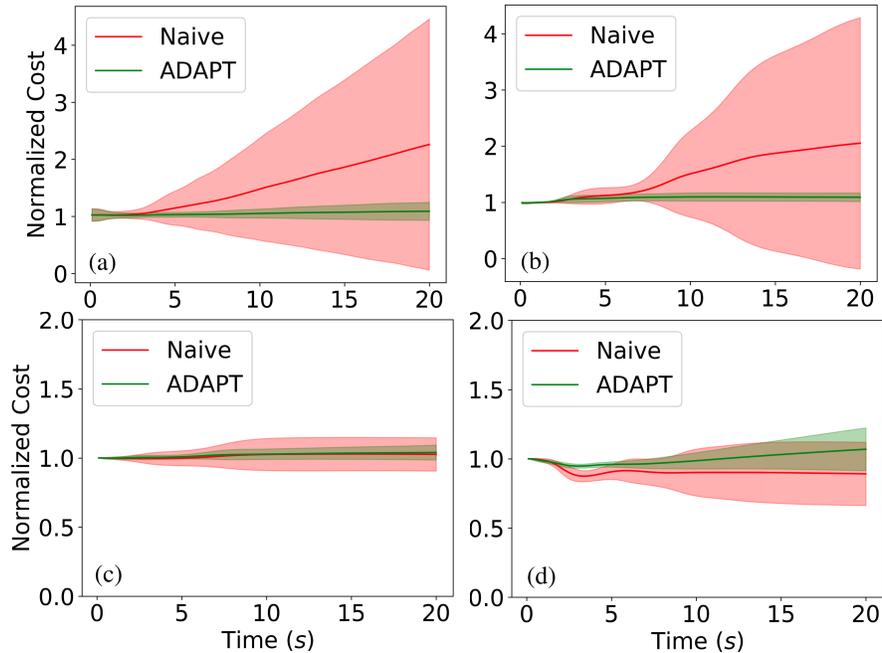


Fig. 3: Mean cumulative cost over the length of an episode for 50 episodes on the 5-D car environment, using an EPOPT-1 robust policy. The confidence intervals are standard error. The disturbances tested are **a)** a hill landscape, **b)** additive control error, **c)** process noise, and **d)** dynamics parameter error. The details of each noise source is presented in the supplementary materials.

from Gaussian distributions before each training run, and are fixed during the training run. Although some of these parameters do not have a physical interpretation, the resulting policies are still robust to both parametric error, as well as process noise. In these experiments, an MPC horizon of 1 second was used (10 timesteps). The matrices Q and R were set as in Section 6.1.

In Figure 3, the comparison between the direct transfer of EPOPT policies and ADAPT policies is presented. We can see that, for disturbances that are explicitly considered in training (specifically, model parameter error), naive transfer performs slightly better, albeit with higher variance. For other disturbances, like the addition of hills or control noise, ADAPT significantly outperforms the directly-transferred policy. Indeed, while the performance of the ADAPT policy is comparable to direct transfer for disturbances directly considered in training, unmodelled disturbances are handled substantially better by ADAPT. Thus, to extract the best performance, we recommend applying the two approaches in tandem.

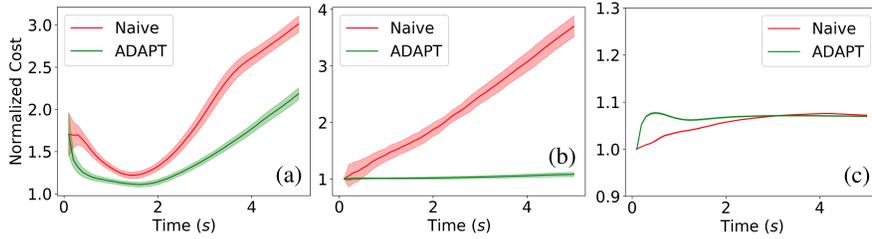


Fig. 4: Mean cumulative cost over the length of an episode for 50 episodes on the `reacher` environment. The confidence intervals are standard error. The costs are normalized to the cost of the naive policy being rolled out on the simulated environment from the same initial state, to allow more direct comparison across episodes. The *naive* rollout is the nominal policy executed on the target environment. The disturbances tested are **a)** additive control error, **b)** process noise, and **c)** dynamics parameter error.

6.4 Environment II: 2-Link Planar Robot Arm

We next evaluate the performance of ADAPT on the `Reacher` environment of Gym [3]. This environment is a two link robotic arm that receives reward for proximity to a goal in the workspace, and is penalized for control effort. The state is a vector of the \sin and \cos of the joint angles, as well as joint angular velocities, the goal position, and the distance from the arm end-effector to the goal. In our tests, we fix one goal location and one starting state for all tests to more directly compare between trials. As such, the variance in normalized cost in experiments is much smaller than in the car experiments. For these experiments, the same noise models were used as in the previous section, with the exception of the “hills” disturbance.

As an approximate dynamics model used for the auxiliary controller, we use the time-varying linear dynamics from [12]. This model is fit from rollouts in simulation. Since this model is linear, the MPC problem is convex, and the iterative MPC converges in one iteration. These dynamics are only valid in a local region, and thus must be fit for each desired policy rollout in the target environment. However, since the model is fit from simulation data, it is generated quickly and inexpensively.

The results for normalized cost comparisons between naive transfer and ADAPT are presented in Figure 4. We note that ADAPT achieves significantly lower cost for additive control error and process noise, but achieves comparable cost for parameter error. The parameter varied in these experiments was the mass of the links of the arm. The effect of this change is to increase the inertia of the manipulator as a whole. In fact, this can be seen in the Figure 4c. While the cost of the naive transfer increases slowly, the cost of the ADAPT trials spikes at approximately time $t = 0.25$. As ADAPT is tracking the nominal trajectory, it increases the torque applied, thus suffering a penalty for the increased control action, but resulting in better tracking of the nominal trajectory.

A similar effect can be observed in Figure 4a. The added control error actually drives the manipulator toward the goal, resulting in the dip in the normalized cost for both trajectories. However, the naive policy overshoots the goal substantially, and thus accrues substantially higher normalized cost than the ADAPT experiments.

7 Conclusion and Outlook

We have presented the ADAPT algorithm for robust transfer of learned policies to target environments with unmodeled disturbances or model parameters. We have also provided guarantees on the lower bounds of the accrued reward in the target environment for a policy transferred with ADAPT. Our results were demonstrated on two different environments with four disturbance models investigated. We additionally discuss usage of robust policies with ADAPT. The results presented demonstrate that this method improves performance on unmodeled disturbances by 50-300%.

In this work, we construct our analysis on the Lipschitz continuity of the dynamics. Indeed, the smoothness of the deviation in dynamics is fundamental to the guarantees we establish. An immediate avenue of future investigation is, therefore, expanding the work presented here to environments with discrete and discontinuous dynamics such as contact. Recently, Farshidian *et al.* [7] have extended an iteratively linearized nonlinear MPC, similar to ours, to switching linear systems, which may have potential as a foundation on which to develop a capable contact formulation of ADAPT. Additionally, recent work has developed robust, receding horizon tube controllers that allow the establishment of explicit tubes in the state space [26]. This approach has the potential to establish explicit safety constraints for operation in cluttered environments. Finally, these methods will also be evaluated on a physical systems.

References

- [1] P. Abbeel, M. Quigley, and A. Y. Ng, "Using inaccurate models in reinforcement learning", in *Proceedings of the 23rd international conference on Machine learning*, ACM, 2006.
- [2] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, *et al.*, "Using simulation and domain adaptation to improve efficiency of deep robotic grasping", *ArXiv preprint arXiv:1709.07857*, 2017.
- [3] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym", *ArXiv preprint arXiv:1606.01540*, 2016.
- [4] P. Christiano, Z. Shah, I. Mordatch, J. Schneider, T. Blackwell, J. Tobin, P. Abbeel, and W. Zaremba, "Transfer from simulation to real world through learning deep inverse dynamics model", *ArXiv preprint arXiv:1610.03518*, 2016.
- [5] M. Deisenroth and C. E. Rasmussen, "Pilco: a model-based and data-efficient approach to policy search", in *Proc. of the 28th Int'l Conf. on Machine Learning (ICML-11)*, 2011.
- [6] C. Devin, A. Gupta, T. Darrell, P. Abbeel, and S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer", *ArXiv preprint arXiv:1609.07088*, 2016.
- [7] F. Farshidian, D. Pardo, and J. Buchli, "Sequential linear quadratic optimal control for nonlinear switched systems", *ArXiv preprint arXiv:1609.02198*, 2016.
- [8] S. Gu, T. Lillicrap, I. Sutskever, and S. Levine, "Continuous deep q-learning with model-based acceleration", *ICML*, 2016.
- [9] N. Heess, G. Wayne, D. Silver, T. Lillicrap, T. Erez, and Y. Tassa, "Learning continuous control policies by stochastic value gradients", in *NIPS*, 2015.
- [10] G. Kahn, T. Zhang, S. Levine, and P. Abbeel, "Plato: Policy learning using adaptive trajectory optimization", *ArXiv preprint arXiv:1603.00622*, 2016.
- [11] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey", *The International Journal of Robotics Research*, p. 0 278 364 913 495 721, 2013.

- [12] S. Levine and P. Abbeel, “Learning neural network policies with guided policy search under unknown dynamics”, in *Advances in Neural Information Processing Systems*, 2014.
- [13] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies”, *Journal of Machine Learning Research*, vol. 17, no. 39, pp. 1–40, 2016.
- [14] A. Mandlekar*, Y. Zhu*, A. Garg*, L. Fei-Fei, and S. Savarese (* equal contribution), “Adversarially robust policy learning through active construction of physically-plausible perturbations”, in *IEEE Int’l Conf. on Intelligent Robots and Systems (IROS)*, 2017.
- [15] D. Q. Mayne, E. C. Kerrigan, E. Van Wyk, and P. Falugi, “Tube-based robust nonlinear model predictive control”, *International Journal of Robust and Nonlinear Control*, 2011.
- [16] D. Mitrovic, S. Klanke, and S. Vijayakumar, “Adaptive optimal feedback control with learned internal dynamics models”, in *From Motor Learning to Interaction Learning in Robots*, Springer, 2010, pp. 65–84.
- [17] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, “Human-level control through deep reinforcement learning”, *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [18] T. Moerland, J. Broekens, and C. Jonker, “Learning multimodal transition dynamics for model-based reinforcement learning”, *ArXiv preprint arXiv:1705.00470*, 2017.
- [19] I. Mordatch, K. Lowrey, and E. Todorov, “Ensemble-cio: full-body dynamic motion planning that transfers to physical humanoids”, in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, IEEE, 2015, pp. 5307–5314.
- [20] M. Neunert, C. de Crousaz, F. Furrer, M. Kamel, F. Farshidian, R. Siegwart, and J. Buchli, “Fast nonlinear model predictive control for unified trajectory optimization and tracking”, in *IEEE Int’l Conf. on Robotics and Automation (ICRA)*, 2016.
- [21] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Sim-to-real transfer of robotic control with dynamics randomization”, *ArXiv preprint arXiv:1710.06537*, 2017.
- [22] L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, “Robust adversarial reinforcement learning”, *ArXiv preprint arXiv:1703.02702*, 2017.
- [23] A. Rajeswaran, S. Ghotra, S. Levine, and B. Ravindran, “EPOpt: learning robust neural network policies using model ensembles”, *ArXiv preprint arXiv:1610.01283*, 2016.
- [24] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, “Progressive neural networks”, *ArXiv preprint arXiv:1606.04671*, 2016.
- [25] J. Schulman, S. Levine, P. Moritz, M. Jordan, and P. Abbeel, “Trust region policy optimization”, *ICML*, 2015.
- [26] S. Singh, A. Majumdar, J.-J. Slotine, and M. Pavone, “Robust online motion planning via contraction theory and convex optimization”, in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*.
- [27] Y. Tassa, T. Erez, and E. Todorov, “Synthesis and stabilization of complex behaviors through online trajectory optimization”, in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, IEEE, 2012, pp. 4906–4913.
- [28] M. E. Taylor and P. Stone, “Transfer learning for reinforcement learning domains: A survey”, *Journal of Machine Learning Research*, vol. 10, 2009.
- [29] 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”, *ArXiv preprint arXiv:1703.06907*, 2017.
- [30] E. Todorov and W. Li, “A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems”, in *American Control Conference, 2005. Proceedings of the 2005*, IEEE, 2005, pp. 300–306.
- [31] D. J. Webb and J. van den Berg, “Kinodynamic rrt*: asymptotically optimal motion planning for robots with linear dynamics”, in *IEEE Int’l Conf. on Robotics and Automation (ICRA)*, 2013.
- [32] K. Zhou, J. C. Doyle, K. Glover, *et al.*, *Robust and optimal control*. 1996, vol. 40.