

DiffOP: Reinforcement Learning of Optimization-Based Control Policies via Implicit Policy Gradients

Yuexin Bian, Jie Feng, Yuanyuan Shi

University of California, San Diego
yubian@ucsd.edu, jif005@ucsd.edu, yyshi@ucsd.edu

Abstract

Real-world control systems require policies that are not only high-performing but also interpretable and robust. A promising direction toward this goal is model-based control, which learns system dynamics and cost functions from historical data and then uses these models to inform decision-making. Building on this paradigm, we introduce DiffOP, a novel framework for learning optimization-based control policies defined implicitly through optimization control problems. Without relying on value function approximation, DiffOP jointly learns the cost and dynamics models and directly optimizes the actual control costs using policy gradients. To enable this, we derive analytical policy gradients by applying implicit differentiation to the underlying optimization problem and integrating it with the standard policy gradient framework. Under standard regularity conditions, we establish that DiffOP converges to an ϵ -stationary point within $\mathcal{O}(\epsilon^{-1})$ iterations. We demonstrate the effectiveness of DiffOP through experiments on nonlinear control tasks and power system voltage control with constraints.

Code — <https://github.com/alwaysbyx/DiffOP>

Extended version — <https://arxiv.org/abs/2411.07484>

Introduction

Operating and controlling complex physical systems in an effective manner is of critical importance to society. Real-world physical systems, such as power grids (Machowski et al. 2020), commercial and industrial infrastructures (Jang et al. 2024), transportation networks (Negenborn, De Schutter, and Hellendoorn 2008), and robotic systems (Spong, Hutchinson, and Vidyasagar 2020), require control policies that are not only high-performing but also interpretable with reliability and safety guarantees. To this end, optimization-based policies such as model predictive control (MPC) have been explored with known (Morari and Lee 1999; Grüne et al. 2017) or learned system dynamics (Chen, Shi, and Zhang 2019), offering a principled framework for performance optimization, constraint satisfaction, and transparent decision-making.

Optimization-based policies formulate control as a mathematical optimization problem, where the goal is to minimize

a system cost subject to dynamic equations and state-action constraints. Prior work in this area often focuses on learning system dynamics and cost functions separately from historical data to improve predictive accuracy (Chen, Shi, and Zhang 2019; Salzmänn et al. 2023; Zhang, Sun, and Shi 2021). However, this decoupled learning procedure can lead to the objective mismatch problem (Donti, Amos, and Kolter 2017; Lambert et al. 2020), where the learned models perform well in isolation (e.g., accurate prediction), but fail to guide optimal control decisions—since the training objectives do not directly reflect control performance.

To enable end-to-end learning of optimization-based policies, recent works (Srinivas et al. 2018; Amos et al. 2018; Jin et al. 2020; Xu, Molloy, and Gould 2024) have made the optimization layers differentiable, allowing both the dynamics and cost function to be jointly learned. These approaches primarily focus on supervised imitation learning settings, where the optimization policy is trained to mimic expert demonstrations and achieve effective control.

More recently, research has moved beyond imitation learning to explore reinforcement learning (RL)-based training of optimization-based policies, aiming to improve the control performance through direct interaction with the environment. In particular, existing works have investigated the integration of MPC with RL (Esfahani, Kordabad, and Gros 2021; Lin et al. 2023; Reiter et al. 2025; Lawrence et al. 2025), leveraging the fact that MPC naturally defines a state-action value function Q , which can be updated via Q-learning to enhance closed-loop control performance.

Building on this line of work, our approach departs from value function approximation and instead treats the optimization-based policy as a fully differentiable control module. This work introduces DiffOP, a novel optimal control framework based on a **Differentiable Optimization-based Policy**. DiffOP computes analytical policy gradients by integrating trajectory derivatives derived from Pontryagin’s Maximum Principle (PMP) (Jin et al. 2020; Xu, Molloy, and Gould 2024) with the standard policy gradient algorithm (Sutton and Barto 2018), enabling end-to-end learning in model-free RL settings.

Our main contributions are as follows:

- We propose DiffOP, which learns optimization-based control policies via policy gradients. A key technical innovation is the joint learning of both the cost function

and system dynamics with environmental rewards, using implicit differentiation and policy gradient updates.

- We provide, to the best of our knowledge, the first non-asymptotic convergence guarantee for learning optimization-based policies via policy gradients. Specifically, we show that DiffOP converges to an ϵ -stationary point within $\mathcal{O}(\epsilon^{-1})$ iterations.
- We validate DiffOP on a set of nonlinear control tasks and a real-world voltage control problem with safety constraints, demonstrating its ability to learn cost and dynamics models with superior control performance compared to existing RL-based MPC methods and differentiable optimization approaches in supervised settings.

Related Work

Differentiable Optimization for Control. Recent advances have enabled end-to-end learning of model-based control policies (e.g., MPC) by making the underlying optimization differentiable. This allows computing gradients w.r.t. optimization parameters, enabling learning via gradient-based methods. One approach is to unroll the optimization steps within the computational graph and apply automatic differentiation, as proposed in (Srinivas et al. 2018; Liu et al. 2018). However, this method is memory-intensive and may yield inaccurate gradients due to truncation. To address the limitations of unrolling, (Amos et al. 2018) proposed differentiable MPC by approximating the original problem with iterative LQR and differentiating through its KKT conditions, though this approach does not support state constraints. Subsequently, (Jin et al. 2020) derived exact analytical gradients by differentiating through Pontryagin’s Maximum Principle (PMP), enabling support for state-action constraints and improving computational efficiency. More recently, (Xu, Molloy, and Gould 2024) introduced IDOC, which applies variable elimination to the KKT system to directly compute trajectory derivatives. While these methods have primarily been applied to imitation learning, our work builds on (Xu, Molloy, and Gould 2024) to enable gradient computation for direct policy optimization.

Synthesis of MPC and RL. Supervised learning of MPC policies often relies on minimizing proxy losses (e.g., prediction error), which may not align with the true control objectives. To overcome this, there is increasing interest in integrating RL with MPC to enhance closed-loop performance. For example, (Chen, Cai, and Bergés 2019) leverages Proximal Policy Optimization (PPO) to adapt linear dynamics models within the differentiable-MPC framework (Amos et al. 2018), relying on LQR approximations. Given that MPC naturally defines a state-action value function, (Gros and Zanon 2019, 2021; Lawrence et al. 2025) employ Q-learning with temporal-difference updates to learn improved MPC policies. Additionally, (Seel, Gros, and Gravidahl 2023) combines Q-learning with deterministic policy gradients to improve convergence. Our work builds on this direction but differs in that, instead of relying on value function approximation or LQR-based structures, we formulate the control policy as the solution to a parameterized nonlinear optimization problem and derive exact ana-

lytical gradients via implicit differentiation. Beyond algorithmic design, our work contributes theoretical convergence results for the proposed framework, building on policy gradient theory (Papini et al. 2018; Yang, Zheng, and Pan 2021; Zhang et al. 2021) and recent work in bi-level optimization (Ji, Yang, and Liang 2021; Kwon et al. 2023).

DiffOP: A Differentiable Optimization Policy

In this section, we introduce DiffOP, the optimization-based control policy, and formalize its learning process as a bi-level optimization problem. Throughout this paper, we use i to denote the planning step within the optimization horizon, and t to denote the actual time step in the control system.

Optimization-based Control Policy

The optimization-based control policy is defined as:

$$u_{0:H-1}^*(x_{\text{init}}; \theta) = \arg \min_u \sum_{i=0}^{H-1} c(x_i, u_i; \theta_c) + c_H(x_H; \theta_H)$$

subject to $x_0 = x_{\text{init}},$
 $x_{i+1} = f(x_i, u_i; \theta_f),$
 $g(x_i, u_i) \leq 0, i = 0, \dots, H - 1$

(1)

where H is the planning horizon, $x_{\text{init}} \in \mathbb{R}^n$ is the initial state. We denote $x_i \in \mathbb{R}^n, u_i \in \mathbb{R}^m$ as the system state and control action at the i -th planning step, respectively. $c(x_i, u_i; \theta_c), c_H(x_H; \theta_H)$ model the instantaneous and terminal costs with learnable parameter θ_c and θ_H , $f(x_i, u_i; \theta_f)$ models system dynamics with learnable parameters θ_f , and $g(\cdot)$ represents the state and control constraints, which are known to the control agent. In summary, our policy (1) is represented by some parametric representation of the *unknown* cost and dynamics functions with parameters $\theta = (\theta_c, \theta_H, \theta_f) \in \mathbb{R}^d$.

Given a state x , DiffOP generates a sequence of control actions by solving (1). We define the policy output as:

$$\pi_\theta(x) = \{u_0^*(x; \theta), \dots, u_{h-1}^*(x; \theta)\}, \quad (2)$$

where u_i^* denotes the i -th action from the optimized sequence, and $1 \leq h \leq H$ specifies how many actions are executed before receiving the next observation.

Remark 1. Although our policy optimizes over a finite horizon, we do not name it as “MPC” since it is not limited to receding-horizon execution. DiffOP allows execution of multiple actions per optimization step ($h \geq 1$). As shown in both the algorithm and experiments, DiffOP naturally accommodates arbitrary h through its policy gradient formulation, supporting both step-wise (when $h = 1$ as in standard MPC) and trajectory-level control (when $h > 1$). Moreover, any standard optimization solver (Andersson et al. 2019; Wächter and Biegler 2006) can be used to solve the policy.

Policy Learning as a Bi-Level Optimization

We are interested in learning the parameters in the policy $\theta = (\theta_c, \theta_H, \theta_f) \in \mathbb{R}^d$ to minimize the overall control costs.

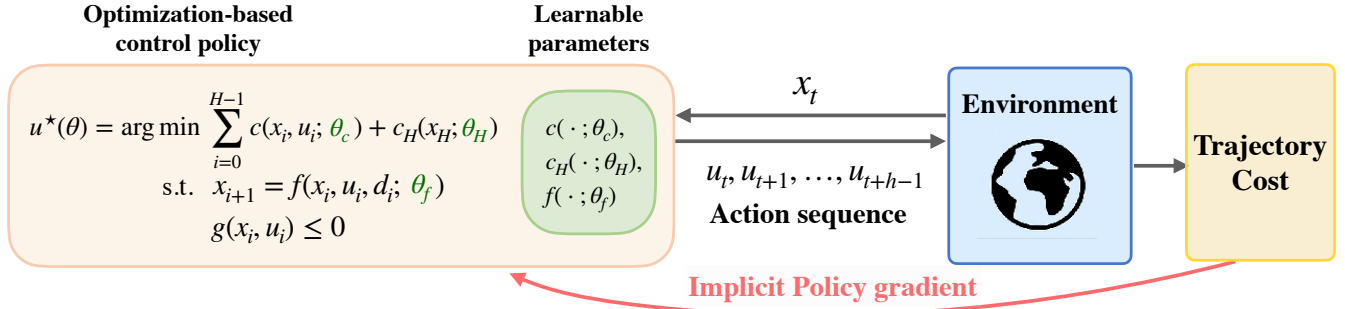


Figure 1: Overview of the DiffOP framework. An optimization-based control policy generates an action sequence by solving a parameterized optimal control problem with learnable dynamics and cost models. The environment executes the action sequence and returns cost feedback, which is used to compute policy gradients for updating the parameters.

The policy optimization problem is formulated as:

$$\min_{\theta} C(\theta) := \mathbb{E} \left[\sum_{t=0}^T c(x_t, u_t; \phi_c) \right] \quad (3a)$$

$$\text{subject to } x_{t+1} = f(x_t, u_t; \phi_f), \quad \forall t, \quad (3b)$$

$$u_t, \dots, u_{t+h-1} \sim \pi_{\theta}(x_t). \quad (3c)$$

Upper-level: Unknown system (3a)(3b): Equation (3a) describes the ground-truth system cost model with parameters ϕ_c and equation (3b) describes ground-truth system dynamics parameterized by ϕ_f , both are *unknown* to the agent. The decision maker seeks to learn an optimization-based policy in the form of (1) with parameters $\theta = (\theta_c, \theta_H, \theta_f)$ that minimizes the expected accumulated costs in the real system.

Lower-level: Stochastic optimization-based policy (3c): The lower-level optimization involves solving the optimization policy to determine the control action at each time step. To enable effective exploration during policy learning, we introduce stochasticity by reparameterizing the policy $u_t, \dots, u_{t+h-1} \sim \pi_{\theta}(x_t)$ as a truncated Gaussian distribution. Specifically, at each decision point t , the policy generates a sequence of h actions by solving the optimization, and each action is sampled as:

$$u_i \sim \mathcal{N}(u_i^*, \sigma^2 I), \text{ clipped to } [u_i^* \pm \beta \sigma^2 I], \quad (4)$$

for $i = t, \dots, t + h - 1$. Here u_i^* denotes the deterministic action at time step obtained by solving the optimization-based policy (2). The term $\sigma^2 I$ is the covariance matrix, with $\sigma^2 > 0$ controlling the variance of the exploration noise, and $I \in \mathbb{R}^{m \times m}$ is the identity matrix. The truncation bounds $u_i^* \pm \beta \sigma^2 I$ ensure the sampled actions remain within a controlled range, with $\beta > 0$ serving as a hyperparameter that regulates the extent of exploration. This independent Gaussian sampling is used for simplicity and analytical tractability in convergence analysis. In practice, DiffOP can be extended to sample the action sequence jointly from a multivariate Gaussian, allowing for correlated noise across the planning horizon. The expectation \mathbb{E} is taken over the initial state distribution $x_0 \sim \mathcal{D}$ and the trajectory $(x_0, u_0, \dots, x_T, u_T)$ induced by the stochastic policy (3c). By choosing a truncated Gaussian distribution over a stan-

dard Gaussian distribution, we explicitly bound the deviation of sampled actions from their optimal counterparts. This bounded exploration helps control the variance of the policy gradients, which is used for establishing gradient *boundedness* and *convergence* analysis of our algorithm.

Crucially, we view problem (3) as a *policy optimization* task, where the policy is implicitly defined through an optimization problem. This formulation enables the adoption of policy gradient methods for end-to-end training.

Proposed Algorithm

The proposed algorithm is summarized in Algorithm 1. Our algorithm builds on the classic policy gradient framework (Sutton and Barto 2018), and incorporates trajectory derivatives derived from Pontryagin’s Maximum Principle (PMP) (Jin et al. 2020; Xu, Molloy, and Gould 2024) to compute gradients of the actual cost with respect to the policy parameters $\theta := \{\theta_c, \theta_H, \theta_f\}$, which include the stage cost parameters θ_c , terminal cost parameters θ_H , and dynamics model parameters θ_f .

Algorithm 1: DiffOP

Input: $\theta^{(0)} = (\theta_c^{(0)}, \theta_H^{(0)}, \theta_f^{(0)})$, learning rate η

for $k = 0, 1, \dots, K - 1$ **do**

for $n = 1, \dots, N$ **do**

 Sample the action as in eq (2)(4), $\forall t$

 Collect the trajectory data $(x_t^{(n)}, u_t^{(n)}, u_t^{*(n)})$, $\forall t$

 Calculate the gradient $\nabla_{\theta} u_t^{*(n)}$ as in Proposition 1

end for

 Estimate policy gradient $\widehat{\nabla}_{\theta} C(\theta^{(k)})$ as in Proposition 2

 Update policy $\theta^{(k+1)} \leftarrow \theta^{(k)} - \eta \widehat{\nabla}_{\theta} C(\theta^{(k)})$

end for

At each training iteration k , the algorithm collects a batch of N trajectories by sampling actions from the current optimization-based policy. For each trajectory, it computes the gradients $\nabla_{\theta} u^*$ of the optimal actions with respect to the policy parameters. These gradients are then used to estimate the overall policy gradient $\widehat{\nabla}_{\theta} C(\theta)$, which guides

the parameter update via standard gradient descent. To derive these gradients, we first present Proposition 1, which characterizes the differentiability of the optimization-based policy w.r.t. its parameters using the first-order optimality conditions and the implicit function theorem.

We begin by defining x_i^*, u_i^* as the optimal state and action obtained by solving the optimization policy (1), where $i = 0, \dots, H - 1$ denotes the i -th planning step. We further define $\zeta^* = (x_0^*, u_0^*, x_1^*, u_1^*, \dots, u_{H-1}^*, x_H^*)$ as the optimal solution trajectory over the horizon H . We stack all constraints into a single vector-valued constraint $\kappa := (\kappa_{-1}, \kappa_0, \kappa_1, \dots, \kappa_H) \in \mathbb{R}^{n_\kappa}$, $n_\kappa = (H + 1)n + q$. Specifically, $\kappa_{-1} = x_0^*$ is the input to the policy, and

$$\kappa_i = \begin{cases} (\tilde{g}(x_i^*, u_i^*), x_{i+1}^* - f(x_i^*, u_i^*; \theta_f)), & i = 0, \dots, H - 1, \\ \tilde{g}(x_H^*), & i = H, \end{cases} \quad (5)$$

where $\tilde{g}(x_i^*, u_i^*)$ are the subset of active inequality constraints at i -th step, and q is the total number of active inequality constraints. The optimization problem in (1) can be solved using general-purpose solvers (Gill, Murray, and Saunders 2005; Diamond and Boyd 2016) to determine both the optimal solution ζ^* and its active constraints κ .

Proposition 1 (Gradient of the optimization-based policy). *Suppose u^* is the solution of the optimization-based policy (1) and denote ζ^* as the resulting trajectory. Assume $c(\cdot)$, $c_H(\cdot)$, $f(\cdot)$, $g(\cdot)$ are twice differentiable in a neighborhood of (θ, ζ^*) . Let*

$$\begin{aligned} A &= \nabla_\zeta \kappa(\zeta^*; \theta), \\ B &= \nabla_{\theta\zeta}^2 J(\zeta^*; \theta) - \sum_{i=-1}^H \sum_{j=1}^{|\kappa_i|} \lambda_{i,j} \nabla_{\theta\zeta}^2 [\kappa_i(\zeta^*; \theta)]_j, \\ C &= \nabla_\theta \kappa(\zeta^*; \theta), \\ D &= \nabla_{\zeta\zeta}^2 J(\zeta^*; \theta) - \sum_{i=-1}^H \sum_{j=1}^{|\kappa_i|} \lambda_{i,j} \nabla_{\zeta\zeta}^2 [\kappa_i(\zeta^*; \theta)]_j, \end{aligned}$$

If $\text{rank}(A) = n_\kappa$ and D is non-singular, then the gradient $\nabla_{\theta} u_i^*$ takes the following form,

$$\nabla_{\theta} u_i^* = [\nabla_{\theta} \zeta^*]_{(n+m) \cdot i + n : (n+m) \cdot (i+1)}, \quad (6)$$

with

$$\nabla_{\theta} \zeta^* = D^{-1} A^\top (A D^{-1} A^\top)^{-1} (A D^{-1} B - C) - D^{-1} B,$$

where n, m are the state and action dimensions, and the Lagrange multiplier $\lambda \in \mathbb{R}^{n_\kappa}$ satisfies $\lambda^\top A = \nabla_\zeta J(\zeta^*; \theta)$.

We note that $\lambda_{i,j}$ is the dual variable corresponding to the j -th element of $\kappa_i(\zeta^*; \theta)$. Proposition 1 is a direct application of the constrained optimization differentiation (Xu, Molloy, and Gould 2024) and its proof is shown in our report (Bian, Feng, and Shi 2025) for completeness.

We now turn to the estimation of the policy gradient. Let $\tau = (x_0, u_0, x_1, u_1, \dots, u_{T-1}, x_T)$ be the trajectory induced by the policy. The analytical policy gradient update rule would be the following:

Proposition 2 (Policy gradient update). *Consider the policy learning problem (3), the policy gradient takes the analytical form of:*

$$\nabla_{\theta} C(\theta) = \mathbb{E} \left[L(\tau) \left(\sum_{t=0}^T \frac{1}{\sigma^2} [\nabla_{\theta} u_t^*]^\top (u_t - u_t^*) \right) \right], \quad (7)$$

where u_t, u_t^* are the actual control action and corresponding optimal solution derived by the policy at time t .

Proof. The result follows from the REINFORCE gradient estimator (Sutton and Barto 2018), applied to a stochastic policy modeled as a Gaussian centered at the optimization-based solution u_t^* . For a detailed derivation, please refer to our report (Bian, Feng, and Shi 2025). \square

By establishing differentiability of u_t^* with respect to θ via the implicit function theorem as described in Proposition 1, we estimate the policy gradient in practice by Monte Carlo sampling over N trajectories:

$$\widehat{\nabla}_{\theta} C(\theta^{(k)}) = \frac{1}{N} \sum_{n=1}^N \left[L(\tau^{(n)}) \sum_{t=0}^T \frac{1}{\sigma^2} \nabla_{\theta} u_t^{*(n)\top} (u_t^{(n)} - u_t^{*(n)}) \right] \quad (8)$$

It is worth noting that although we employ the REINFORCE algorithm (Sutton and Barto 2018) in Algorithm 1 to estimate the policy gradients, there are several techniques available to reduce the variance of the gradient estimator (Zhao et al. 2011; Grathwohl et al. 2017) which could be incorporated into the framework.

Non-Asymptotic Convergence Analysis

We present our main theoretical result, which establishes the convergence properties of DiffOP when trained using Algorithm 1. Since the policy is implicitly defined through an optimization problem, the sensitivity of its solution with respect to the policy parameters is central to the analysis.

Convergence Under Bounded Sensitivity

We begin by stating a sufficient condition for convergence: bounded first- and second-order sensitivity of the optimization-defined policy with respect to its parameters.

Assumption 1 (Bounded First/Second-Order Sensitivity). *There exists a constant $M_{u1} > 0, M_{u2} > 0$ such that for all planning step i , the solution $u_i^*(x; \theta)$ to the optimization-defined policy satisfies:*

$$\|\nabla_{\theta} u_i^*(x; \theta)\| \leq M_{u1}, \quad \|\nabla_{\theta}^2 u_i^*(x; \theta)\| \leq M_{u2}.$$

This assumption ensures that the policy responds smoothly to changes in the parameters θ . Next, we assume standard conditions within the policy optimization (Papini et al. 2018), regarding the initial state distribution and the trajectory cost.

Assumption 2. *The initial state distribution \mathcal{D} is supported in a region with a finite radius D_0 . For any initial state $x_0 \sim \mathcal{D}$, the trajectory cost $L(\tau)$ is bounded, i.e., there exists a positive constant $M_L > 0$ such that $L(\tau) \leq M_L$.*

Now we can characterize the convergence of the proposed algorithm.

Theorem 1 (Convergence of DiffOP with Policy Gradient). *Suppose Assumptions 1 and 2 hold. For any $\epsilon > 0$, and $\nu \in (0, 1)$, define a smoothness constant*

$$L_C = M_L T (\sqrt{m} \beta M_{u_2} + \frac{1}{\sigma^2} M_{u_1}^2 + T m \beta^2 M_{u_1}^2),$$

a stepsize $\eta = \frac{1}{4L_C}$, the number of policy iterations K , and the number of sampled trajectories for each policy gradient step $N = \frac{2m\beta^2 M_{u_1}^2 T^2 M_{u_1}^2}{\epsilon^2} \log \frac{2Kd}{\nu}$. Then, with probability at least $1 - \nu$, we have

$$\min_{k=0, \dots, K-1} \|\nabla_{\theta} C(\theta^{(k)})\|^2 \leq \frac{16L_C (C(\theta^{(0)}) - C(\bar{\theta}))}{K} + 3\epsilon. \quad (9)$$

where $\bar{\theta}$ is the global optimum of (3).

Proof. The detailed proof is provided in our report (Bian, Feng, and Shi 2025), building on standard stochastic approximation techniques (Papini et al. 2018) and adapted to the implicit policy structure induced by optimization-based control. \square

Sufficient Conditions for Bounded Sensitivity

We now present a common class of optimization-based policies that satisfy the sensitivity bounds in Assumption 1. Specifically, we consider the unconstrained strongly convex policy (1), for which concrete sufficient conditions can be established.

Let $x_0 = x_{\text{init}}$ denote the initial state, and let $u = [u_0, \dots, u_{H-1}] \in \mathbb{R}^{mH}$ be the control sequence generated by the optimization-based policy with parameters θ . The unrolled cost function can be written as:

$$J(x_{\text{init}}, u, \theta) = \sum_{i=0}^{H-1} c(x_i, u_i; \theta_c) + c_H(x_H; \theta_H), \quad (10)$$

subject to the dynamics $x_0 = x_{\text{init}}$, $x_{i+1} = f(x_i, u_i; \theta_f)$ for $i = 0, \dots, H-1$. We next introduce technical assumptions that characterize the optimization landscape and smoothness properties of this class of optimization-defined policies.

Assumption 3. *The function $J(x_{\text{init}}, u, \theta)$ is μ -strongly convex with respect to u .*

Strong convexity ensures the uniqueness of the optimal solution and the stability of the mapping from θ to the control sequence $u^*(x; \theta)$. This injectivity yields well-defined gradients and the convergence guarantees in our policy optimization framework.

Assumption 4. *Let $z = (x_{\text{init}}, u, \theta)$ and $\mathcal{Z} \subset \mathbb{R}^{n_z}$ be a compact set. The unrolled cost function $J(z)$ satisfies the following:*

- $\nabla_z J(z)$ is L_1 -Lipschitz: for any $z, z' \in \mathcal{Z}$,

$$\|\nabla_z J(z) - \nabla_z J(z')\| \leq L_1 \|z - z'\|.$$

- $\nabla_{\theta} \nabla_u J(z)$ is L_2 -Lipschitz: for any $z, z' \in \mathcal{Z}$,

$$\|\nabla_{\theta} \nabla_u J(z) - \nabla_{\theta} \nabla_u J(z')\| \leq L_2 \|z - z'\|.$$

- $\nabla_u^2 J(z)$ is L_3 -Lipschitz: for any $z, z' \in \mathcal{Z}$,

$$\|\nabla_u^2 J(z) - \nabla_u^2 J(z')\| \leq L_3 \|z - z'\|.$$

Assumption 4 captures the smoothness of the unrolled cost function and its partial derivatives with respect to the optimization variables and policy parameters. These conditions ensure the gradients and Hessians involved in the policy sensitivity analysis are well-behaved. Given Assumptions 3 and 4, we can show that Assumption 2 is satisfied.

Proposition 3. *Under Assumptions 3 and 4, the optimization-defined policy π_{θ} satisfies the bounded sensitivity condition stated in Assumption 1. In particular, there exist constants*

$$M_{u_1} = \frac{L_1}{\mu}, \quad M_{u_2} = \frac{L_2}{\mu} + \frac{L_1 L_2 + L_1 L_3}{\mu^2} + \frac{L_1^2 L_3}{\mu^3},$$

such that for all planning steps i ,

$$\|\nabla_{\theta} u_i^*(x; \theta)\| \leq M_{u_1}, \quad \|\nabla_{\theta}^2 u_i^*(x; \theta)\| \leq M_{u_2}.$$

Proof. The result follows from applying the implicit function theorem to the first-order optimality conditions of the unconstrained problem. The complete proof is provided in our report (Bian, Feng, and Shi 2025). \square

Our analysis makes a key theoretical contribution by establishing the first non-asymptotic convergence guarantee for learning optimization-based policies via policy gradients. We show that DiffOP converges to an ϵ -stationary point within $\mathcal{O}(1/\epsilon)$ iterations under mild smoothness and strong convexity assumptions. Notably, this matches the standard convergence rate of first-order policy gradient methods, despite the added complexity of differentiating through an implicit optimization-based policy.

The central technical challenge lies in bounding the sensitivity of the optimization solution w.r.t. policy parameters. We address this by introducing structural assumptions on the inner optimization problem and leveraging analytical tools from bilevel optimization theory (Ji, Yang, and Liang 2021; Kwon et al. 2023). This analysis not only underpins our convergence guarantee but also provides general insights into when optimization-based policies admit stable gradient-based learning. In more general settings, such as non-convex or constrained optimization problems, the solution mapping can become non-smooth or even discontinuous, especially near constraint boundaries or saddle points. This can lead to unbounded gradients and unstable updates during training, making theoretical convergence guarantees significantly harder to establish. For instance, in the presence of inequality constraints, small changes in parameters can trigger abrupt changes in the active set, breaking the smoothness assumptions required for standard implicit differentiation. Extending our convergence analysis to cover such cases remains an important and challenging direction for future work.

However, in practice, certain smooth non-convex constrained problems may exhibit locally bounded sensitivities (Scieur et al. 2022) and satisfy Assumption 1. To demonstrate the practical performance of DiffOP, we include experiments on constrained control tasks that highlight its applicability to realistic settings.

Systems	Objective parameter θ_c, θ_H	Dynamic parameter θ_f
Cartpole		length and mass for each link
Two-link robot arm	$c(x_i, u_i) = \ \theta_c^T(x_i - x_{\text{nom}})\ ^2 + \ u_i\ ^2$	length and mass for each link
6-DoF quadrotor maneuvering	$c_H(x_H) = \ \theta_H^T(x_H - x_{\text{nom}})\ ^2$	mass, wing length, inertia matrix

Table 1: Nonlinear experimental settings (Jin et al. 2020).




Figure 2: Control cost versus training iteration on nonlinear control tasks. Solid lines indicate the mean cost over 5 runs, and shaded regions denote the 20th to 80th percentile range.

Experiments

In this section, we evaluate the effectiveness of DiffOP on nonlinear control tasks and a constrained voltage regulation problem. Detailed experimental settings are provided in our report (Bian, Feng, and Shi 2025).

Nonlinear Control Tasks

In this section, we evaluate the proposed DiffOP framework on a set of nonlinear control tasks. We follow the experimental setup introduced in (Jin et al. 2020), and summarize the environment details in Table 1.

We compare against the following optimization-based control baselines. In all methods, the forward pass is solved using CasADi optimization solver (Andersson et al. 2019):

- **PDP (Offline)** (Jin et al. 2020): Learns optimization parameters via Pontryagin Differentiable Programming from expert trajectories. As a purely offline method, it does not adapt online. Reported performance is taken directly from the original paper.
- **RLMPC-TD** (Gros and Zanon 2019): A model-free reinforcement learning approach that updates a short-horizon MPC policy using temporal-difference (TD) Q-learning.
- **RLMPC-DPG** (Seel, Gros, and Gravdahl 2023): A variant of RLMPC-TD that also employs deterministic policy gradient (DPG) for updating the MPC parameters, with the same short-horizon control structure.

We evaluate **DiffOP** under two deployment modes to highlight its flexibility:

- **DiffOP (Step)**: At each time step, apply only the first action of the optimized trajectory, i.e., a receding-horizon scheme mirrors RLMPC-TD and RLMPC-DPG.
- **DiffOP (Traj)**: Generate the entire optimized control sequence once at the start of the episode and execute

it open-loop. This trajectory-optimization setting boosts temporal consistency but sacrifices some robustness due to its open-loop nature.

For fair comparison, the planning horizon is set to 6 for Cartpole and 12 for RobotArm and Quadrotor across RLMPC-TD, RLMPC-DPG, and DiffOP (Step). For all DiffOP experiments, we set $\beta = 0.01$, $\eta = 0.1$ and $N = 10$. RL-based MPC baselines are implemented using the MPC4RL framework (Reinhardt et al. 2024). Results are presented in Figure 2. Compared to RL-based MPC baselines, DiffOP (Step) achieves faster convergence and lower final control cost across all nonlinear control tasks. In particular, it outperforms both RLMPC-TD and RLMPC-DPG which often converge to suboptimal solutions and may fail to consistently reduce control cost.

Compared to PDP (Offline), DiffOP (Traj) consistently achieves lower control cost in Robot Arm and Quadrotor tasks, showing the benefit of policy refinement through online interaction. While PDP relies solely on supervised expert data, DiffOP (Traj) fine-tunes the optimization-based policy using reinforcement learning, enabling it to surpass offline imitation performance and adapt better to the true system dynamics. These results highlight the effectiveness of DiffOP under both receding-horizon and open-loop settings, and its robustness across control tasks with different levels of nonlinearity and complexity.

Voltage Control with Power Injection Constraints

We further evaluate DiffOP on a practical voltage control scenario involving reactive power constraints, using the IEEE 13-bus radial distribution benchmark. We adopt the experimental setup of (Feng et al. 2023), using the Pandapower toolbox (Turitsyn et al. 2011) to simulate the power network under two voltage disturbance scenarios: (1) high

	Transient cost	Steady state cost
DiffOP (Step)	-6.81	-0.11
RLMPC-DPG	-6.11	-0.10
RLMPC-TD	-4.62	-0.07
DiffOP (Traj)	-6.80	-0.11
PDP (Offline)	-5.86	-0.09
Stable-DDPG	-5.61	-0.09
TASRL	-6.76	-0.11

Table 2: Performance on 500 scenarios of the IEEE 13-bus system. Lower cost indicates better performance. DiffOP (Traj) and PDP use horizon $H = 30$ with open-loop execution; DiffOP (Step), RLMPC-DPG, and RLMPC-TD use $H = 6$ with first-action execution.

voltage due to excess PV generation under strong sunlight, and (2) low voltage caused by peak loads with insufficient generation. For each scenario, we vary load profiles to induce deviations of 5%–15% from the nominal voltage v^{nom} . The goal of voltage control is to achieve fast voltage restoration to an acceptable range (5%) with minimal control.

Let $v_i \in \mathbb{R}^3$ denote the squared voltage magnitudes, $q_i \in \mathbb{R}^3$ the reactive power injection, and $u_i \in \mathbb{R}^3$ the control inputs at step i . We define the optimization-based control policy for voltage regulation:

$$\arg \min_u \sum_{i=0}^{H-1} \left(u_i^\top C_u u_i + (v_i - v^{\text{nom}})^\top C_v (v_i - v^{\text{nom}}) \right) + (v_H - v^{\text{nom}})^\top C_v (v_H - v^{\text{nom}}) \quad (11a)$$

$$\text{s.t. } v_{i+1} = A q_i + v^{\text{env}}, q_{i+1} = q_i + u_i, \quad (11b)$$

$$q \leq q_i \leq \bar{q}. \quad (11c)$$

The objective (11a) penalizes both voltage deviation from the nominal value v^{nom} and control effort, where C_u is a fixed parameter reflecting actuation cost. Equation (11b) models the voltage and reactive power dynamics, where v^{env} is the initial voltage determined by load profiles. In general, v^{env} is considered static due to time-scale separation between load change and the voltage control process. Constraint (11c) ensures that the reactive power injection q_t remains within operational limits. The policy parameters include $\theta = \{C_v, A\}$, which are learned to optimize both transient and steady-state performance—ensuring fast convergence, low operational cost, and stable voltage profiles.

Our proposed DiffOP is evaluated in comparison with all prior baselines, as well as two state-of-the-art algorithms designed for voltage control:

- **Stable-DDPG** (Shi et al. 2022): An RL-trained monotone neural policy for voltage control with theoretical voltage restoration guarantees.
- **TASRL** (Feng et al. 2023): An RL-policy that incorporates Stable-DDPG and safe gradient flow to achieve fast voltage restoration while guaranteeing steady-state optimality with limited reactive-power capacity.

All optimization-based policies are trained on 10,000 trajectories, each of length $T = 30$ seconds. The planning horizon

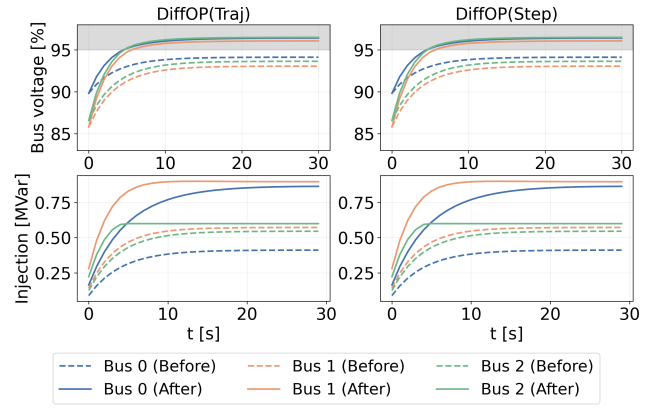


Figure 3: Voltage and reactive power trajectories. Solid lines show post-training responses; dashed lines are pre-training performance. Gray shaded areas denote the safe operating bounds for system stability.

is set to $H = 6$ across RLMPC-TD, RLMPC-DPG, and DiffOP (Step). The overall performance is summarized in Table 2. For comparison, we use results reported in (Feng et al. 2023) for Stable-DDPG and TASRL. DiffOP consistently achieves the lowest transient and steady-state costs across all baselines. Compared to MPC-based approaches, DiffOP effectively learns both cost and dynamics parameters, achieving the best performance in both transient and steady-state costs. Notably, it also outperforms RL-based neural network policies in transient cost and matches TASRL’s steady-state cost.

Further, we visualize the voltage and reactive power injection trajectories in Figure 3. After policy optimization, both variants exhibit significantly improved voltage regulation and faster convergence, with all voltage trajectories stabilized within the safe operating bounds (gray regions), in contrast to the untrained policy. Notably, our framework naturally supports applying a sequence of actions to the system, rather than being limited to receding-horizon execution. Despite the different planning settings used in DiffOP (Step) and DiffOP (Traj), both variants consistently learn effective control policies.

Conclusion

In this paper, we proposed DiffOP, a novel framework for learning optimization-based control policies via analytical implicit policy gradients. Experiments on nonlinear control and constrained voltage regulation tasks show that DiffOP consistently improves performance and outperforms both model-free and MPC-based baselines. This work opens the door to several interesting directions for future research. First, we aim to explore improved sample complexity for learning the optimization-based policy. Secondly, it would be of significant interest to further investigate the robustness of DiffOP in comparison to model-free RL approaches. Finally, we envision extending the experimental results of DiffOP to a wide range of real-world deployments.

Acknowledgments

This work is supported by a Schmidt Sciences AI2050 Early Career Fellowship, NSF grant ECCS-2442689, and DOE grant DE-SC0025495. The work of J. Feng is also supported by the UC-National Laboratory In Residence Graduate Fellowship L24GF7923.

References

- Amos, B.; Jimenez, I.; Sacks, J.; Boots, B.; and Kolter, J. Z. 2018. Differentiable mpc for end-to-end planning and control. *Advances in neural information processing systems*, 31.
- Andersson, J. A.; Gillis, J.; Horn, G.; Rawlings, J. B.; and Diehl, M. 2019. CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1): 1–36.
- Bian, Y.; Feng, J.; and Shi, Y. 2025. DiffOP: Reinforcement Learning of Optimization-Based Control Policies via Implicit Policy Gradients. [arXiv:2411.07484](https://arxiv.org/abs/2411.07484).
- Chen, B.; Cai, Z.; and Bergés, M. 2019. Gnu-rl: A pre-cial reinforcement learning solution for building hvac control using a differentiable mpc policy. In *Proceedings of the 6th ACM international conference on systems for energy-efficient buildings, cities, and transportation*, 316–325.
- Chen, Y.; Shi, Y.; and Zhang, B. 2019. Optimal Control Via Neural Networks: A Convex Approach. In *International Conference on Learning Representations (ICLR)*.
- Diamond, S.; and Boyd, S. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83): 1–5.
- Donti, P.; Amos, B.; and Kolter, J. Z. 2017. Task-based end-to-end model learning in stochastic optimization. *Advances in neural information processing systems*, 30.
- Esfahani, H. N.; Kordabad, A. B.; and Gros, S. 2021. Approximate robust NMPC using reinforcement learning. In *2021 European Control Conference*, 132–137. IEEE.
- Feng, J.; Cui, W.; Cortés, J.; and Shi, Y. 2023. Bridging transient and steady-state performance in voltage control: A reinforcement learning approach with safe gradient flow. *IEEE Control Systems Letters*, 7: 2845–2850.
- Gill, P. E.; Murray, W.; and Saunders, M. A. 2005. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM review*, 47(1): 99–131.
- Grathwohl, W.; Choi, D.; Wu, Y.; Roeder, G.; and Duvenaud, D. 2017. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. *International Conference on Learning Representations (ICLR)*.
- Gros, S.; and Zanon, M. 2019. Data-driven economic NMPC using reinforcement learning. *IEEE Transactions on Automatic Control*, 65(2): 636–648.
- Gros, S.; and Zanon, M. 2021. Reinforcement learning based on mpc and the stochastic policy gradient method. In *2021 American Control Conference*, 1947–1952. IEEE.
- Grüne, L.; Pannek, J.; Grüne, L.; and Pannek, J. 2017. *Non-linear model predictive control*. Springer.
- Jang, D.; Yan, L.; Spangher, L.; and Spanos, C. J. 2024. Active reinforcement learning for robust building control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22150–22158.
- Ji, K.; Yang, J.; and Liang, Y. 2021. Bilevel optimization: Convergence analysis and enhanced design. In *International Conference on Machine Learning (ICML)*, 4882–4892. PMLR.
- Jin, W.; Wang, Z.; Yang, Z.; and Mou, S. 2020. Pontryagin differentiable programming: An end-to-end learning and control framework. *Advances in Neural Information Processing Systems*, 33: 7979–7992.
- Kwon, J.; Kwon, D.; Wright, S.; and Nowak, R. D. 2023. A fully first-order method for stochastic bilevel optimization. In *International Conference on Machine Learning*, 18083–18113. PMLR.
- Lambert, N.; Amos, B.; Yadan, O.; and Calandra, R. 2020. Objective mismatch in model-based reinforcement learning. *arXiv preprint arXiv:2002.04523*.
- Lawrence, N. P.; Loewen, P. D.; Forbes, M. G.; Gopaluni, R. B.; and Mesbah, A. 2025. A view on learning robust goal-conditioned value functions: Interplay between RL and MPC. *arXiv preprint arXiv:2502.06996*.
- Lin, M.; Sun, Z.; Xia, Y.; and Zhang, J. 2023. Reinforcement learning-based model predictive control for discrete-time systems. *IEEE Transactions on Neural Networks and Learning Systems*, 35(3): 3312–3324.
- Liu, R.; Fan, X.; Cheng, S.; Wang, X.; and Luo, Z. 2018. Proximal alternating direction network: A globally converged deep unrolling framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Machowski, J.; Lubosny, Z.; Bialek, J. W.; and Bumby, J. R. 2020. *Power system dynamics: stability and control*. John Wiley & Sons.
- Morari, M.; and Lee, J. H. 1999. Model predictive control: past, present and future. *Computers & chemical engineering*, 23(4-5): 667–682.
- Negenborn, R. R.; De Schutter, B.; and Hellendoorn, J. 2008. Multi-agent model predictive control for transportation networks: Serial versus parallel schemes. *Engineering Applications of Artificial Intelligence*, 21(3): 353–366.
- Papini, M.; Binaghi, D.; Canonaco, G.; Pirota, M.; and Restelli, M. 2018. Stochastic variance-reduced policy gradient. In *International Conference on Machine Learning (ICML)*, 4026–4035. PMLR.
- Reinhardt, D.; Baumgärtner, K.; Frey, J.; Diehl, M.; and Gros, S. 2024. MPC4RL-A Software Package for Reinforcement Learning based on Model Predictive Control. In *2024 IEEE 63rd Conference on Decision and Control*, 1787–1794. IEEE.
- Reiter, R.; Hoffmann, J.; Reinhardt, D.; Messerer, F.; Baumgärtner, K.; Sawant, S.; Boedecker, J.; Diehl, M.; and Gros, S. 2025. Synthesis of model predictive control and reinforcement learning: Survey and classification. *arXiv preprint arXiv:2502.02133*.

Salzmann, T.; Kaufmann, E.; Arrizabalaga, J.; Pavone, M.; Scaramuzza, D.; and Ryll, M. 2023. Real-time neural MPC: Deep learning model predictive control for quadrotors and agile robotic platforms. *IEEE Robotics and Automation Letters*, 8(4): 2397–2404.

Scieur, D.; Gidel, G.; Bertrand, Q.; and Pedregosa, F. 2022. The curse of unrolling: Rate of differentiating through optimization. *Advances in Neural Information Processing Systems*, 35: 17133–17145.

Seel, K.; Gros, S.; and Gravdahl, J. T. 2023. Combining Q-learning and deterministic policy gradient for learning-based MPC. In *2023 62nd IEEE Conference on Decision and Control*, 610–617. IEEE.

Shi, Y.; Qu, G.; Low, S.; Anandkumar, A.; and Wierman, A. 2022. Stability constrained reinforcement learning for real-time voltage control. In *2022 American Control Conference*, 2715–2721. IEEE.

Spong, M. W.; Hutchinson, S.; and Vidyasagar, M. 2020. *Robot modeling and control*. John Wiley & Sons.

Srinivas, A.; Jabri, A.; Abbeel, P.; Levine, S.; and Finn, C. 2018. Universal planning networks: Learning generalizable representations for visuomotor control. In *International Conference on Machine Learning (ICML)*, 4732–4741. PMLR.

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.

Turitsyn, K.; Sulc, P.; Backhaus, S.; and Chertkov, M. 2011. Options for control of reactive power by distributed photovoltaic generators. *Proceedings of the IEEE*, 99(6): 1063–1073.

Wächter, A.; and Biegler, L. T. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106: 25–57.

Xu, M.; Molloy, T. L.; and Gould, S. 2024. Revisiting implicit differentiation for learning problems in optimal control. *Advances in Neural Information Processing Systems*, 36.

Yang, L.; Zheng, Q.; and Pan, G. 2021. Sample complexity of policy gradient finding second-order stationary points. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 10630–10638.

Zhang, J.; Kim, J.; O’Donoghue, B.; and Boyd, S. 2021. Sample efficient reinforcement learning with REINFORCE. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 10887–10895.

Zhang, K.; Sun, Q.; and Shi, Y. 2021. Trajectory tracking control of autonomous ground vehicles using adaptive learning MPC. *IEEE Transactions on Neural Networks and Learning Systems*, 32(12): 5554–5564.

Zhao, T.; Hachiya, H.; Niu, G.; and Sugiyama, M. 2011. Analysis and improvement of policy gradient estimation. *Advances in Neural Information Processing Systems*, 24.