

Adversarially Regularized Policy Learning Guided by Trajectory Optimization

Zhigen Zhao

Simiao Zuo

Tuo Zhao*

Ye Zhao*

ZHIGEN.ZHAO@GATECH.EDU

SIMIAOZUO@GATECH.EDU

TOURZHAO@GATECH.EDU

YE.ZHAO@ME.GATECH.EDU

Georgia Institute of Technology, North Ave NW, Atlanta, GA 30332

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Recent advancement in combining trajectory optimization with function approximation (especially neural networks) shows promise in learning complex control policies for diverse tasks in robot systems. Despite their great flexibility, the large neural networks for parameterizing control policies impose significant challenges. The learned neural control policies are often overcomplex and non-smooth, which can easily cause unexpected or diverging robot motions. To address this issue, we propose **adversarially regularized policy learning** guided by trajectory optimization (VERONICA) for learning smooth control policies. Specifically, our proposed approach controls the smoothness (local Lipschitz continuity) of the neural control policies by stabilizing the output control with respect to the worst-case perturbation to the input state. Our experiments on robot manipulation show that our proposed approach not only improves the sample efficiency of neural policy learning but also enhances the robustness of the policy against various types of disturbances, including sensor noise, environmental uncertainty, and model mismatch.

Keywords: Adversarial Regularization, Policy Learning, Trajectory Optimization

1. Introduction

Robust and generalizable motion planning enables robotic systems to handle various uncertainties and accomplishes diverse tasks. However, learning a dynamically consistent neural control policy (i.e., a neural-network control policy) and executing it reliably remain challenging. First, the function approximators used to model the policy can be highly complex and non-smooth, causing poor generalization performance. Second, the dynamics models involved are often mismatched from the physical robot, leading to the need of learning a robust policy.

Trajectory optimization (TO) (Betts, 1998; Kuindersma et al., 2016; Tassa et al., 2014; Posa et al., 2014) is a powerful model-based approach to generate optimal control sequences for complex robotic systems. However, existing methods for solving TO problems with full robot dynamics require solving large nonlinear programs, resulting in high computational cost. This difficulty prevents the use of TO methods in real-time robot control settings. As such, to alleviate the computational burden at run-time, it is preferable to have a parametric representation of a robot control policy. In comparison, model-free policy search, as in (Deisenroth et al., 2013), aims to automatically learn the controller through random exploration. However, a majority of these methods fail to utilize the prior knowledge on the robot dynamics encoded in the physical model, which causes sample inefficiency.

* Co-corresponding Author

To take advantage of both TO and policy search, [Mordatch and Todorov \(2014\)](#) and [Levine and Koltun \(2013a\)](#) train a robot control policy supervised by optimized trajectory samples, and meanwhile adapting TO to the learned policy. The work in [Mordatch and Todorov \(2014\)](#) observes that the derivatives of a neural control policy can behave irregularly even when the policy matches the optimal trajectory baseline. This is because neural networks have high complexity and flexibility, which makes them highly non-smooth — a small change in the networks’ input can cause a large variation in the output. To mitigate this limitation, existing works attempt to impose some smoothness constraints on the policy. For instance, [Mordatch and Todorov \(2014\)](#) matches the gradient for policy and trajectory samples via tangent propagation. However, tangent propagation requires Jacobian computation on each trajectory point, which does not scale well to large datasets.

To alleviate these issues, we propose a new approach: **adversarially regularized policy learning** guided by trajectory optimization (VERONICA). Specifically, our approach improves the local Lipschitz continuity of the neural control policy via adversarial regularization, which improves generalization performance for inputs not seen during training. We focus on promoting smoothness in policy for non-hybrid robotics tasks that are often governed by differential equations with high-order continuity. For hybrid systems where non-smooth dynamics might occur during physical contact, several works in TO ([Brubaker et al., 2009](#); [Todorov, 2011](#); [Mordatch et al., 2012](#)) and physical simulation MuJoCo ([Todorov et al., 2012](#)) propose to model contact with a smoothed model, where contact forces diminish gradually with contact distance. The work of [Drnach and Zhao \(2021\)](#) proposes a risk-sensitive cost function to represent a stochastic, smoothed variant of the original complementarity contact problem ([Chen and Mangasarian, 1996](#)). In this work, we show that the VERONICA framework also provides robustness benefits for a hybrid locomotion system with physical contacts.

The VERONICA framework is related to existing works ([Miyato et al., 2018](#); [Zhang et al., 2019](#); [Hendrycks et al., 2019](#); [Xie et al., 2019](#); [Jiang et al., 2019](#); [Shen et al., 2020](#); [Jiang et al., 2019](#); [Zuo et al., 2021a](#); [Li et al., 2020](#)). These works consider similar regularization techniques, but target at other applications with different motivations, e.g., semi-supervised learning, unsupervised domain adaptation, harnessing adversarial examples, fine-tuning pre-trained models and reinforcement learning. [Morimoto and Doya \(2000\)](#) and [Pinto et al. \(2017\)](#) solve similar min-max problems to improve the robustness of reinforcement learning.

We further observe that besides promoting policy smoothness, adversarial regularization improves the robustness of the policy against modeling errors and perturbations in the environment. We verify that the VERONICA framework produces stable robot behaviors under sensor noise, environmental uncertainty, and model mismatch.

Conventionally, adversarial regularization involves a min-max game, which is solved by alternating gradient descent-ascent. During training, neither of the players can be advantageous, such that the generated perturbations can be over-strong and hinder model generalization. To resolve this issue, we employ Stackelberg adversarial regularization (SAR), as proposed in [Zuo et al. \(2021b\)](#), which formulates adversarial regularization as a Stackelberg game ([Von Stackelberg, 2010](#)). In SAR, the policy (i.e., the leader) has a higher priority than the perturbation (i.e., the follower). The leader procures its advantage by considering how the follower will respond after observing the leader’s decision, such that the leader anticipates the predicted move of the follower when optimizing its strategy. We note that prioritizing the policy optimization is reasonable and beneficial because we target the performance of the learned policy, instead of the adversary.

Our contributions are: I) We propose VERONICA, an adversarial regularization method for learning smooth neural control policies guided by TO. This improves the generalization perfor-

mance of the learned policy; II) We show that the learned policy achieves better robustness under disturbances such as sensor noise, environmental uncertainty, and model mismatch; III) We reformulate adversarial regularization as a Stackelberg game, which further improves generalization and robustness of the policy compared with the conventional formulation.

2. Related Works

Adversarial Training in Robot Learning: Adversarial training has previously been used to improve safety in robot visuomotor control scenarios (Chen et al., 2020) to improve generalization. The work in Lechner et al. (2021) argues that adversarial training induces unexplored error profiles in *vision-based* robot learning, which studies classification tasks that are not Lipschitz continuous. In contrast, our work focuses on adversarial regularization for neural control policy in *dynamics-based* robot learning, which are intrinsically smooth. Therefore, *vision-based* adversarial training studies fundamentally different problems than ours.

Imitation Learning: Behavioral cloning (BC) uses supervised learning to directly imitate expert trajectories without interacting with the environment (Schaal et al., 1997). However, BC is particularly vulnerable to error compounding (Ross et al., 2011). In our work, we solve a BC problem for policy learning in each iteration of the Alternating Direction Method of Multipliers (ADMM) method, while the ADMM framework offers a coupling mechanism to allow the trajectory optimizer (i.e., the teacher) to not only guide the learned policy (i.e., the student) towards better solutions but also adapt to the student. More importantly, we incorporate an adversarial regularizer to improve policy smoothness, which significantly eases the effect of error compounding.

Along another line of research, generative adversarial imitation learning (Ho and Ermon, 2016; Zolna et al., 2019) uses generative adversarial networks (GAN) to directly generate policies that imitate expert demonstrations. In contrast, the adversaries in our work are the direct perturbations on the input (i.e., the robot state), rather than the discriminator network.

Trajectory-Optimization-Guided Policy Learning: Trajectory optimization has been used to aid and stabilize value function learning in the reinforcement learning (RL) context (Lowrey et al., 2018), while Landry et al. (2021) use a bilevel optimization to learn the value function with adversarial samples. In this work, we focus on supervised learning approaches that train neural control policies from TO.

Guided policy search (GPS) (Levine and Koltun, 2013a,b, 2014) iteratively updates guiding sample using differential dynamic programming (DDP) and trains policies on the distribution over the guiding samples. In contrast, the work of Mordatch and Todorov (2014) seeks consensus between neural network policy and trajectory optimization using ADMM (Boyd et al., 2011). Duburcq et al. (2020) similarly solve for ADMM consensus, but aim to learn a trajectory sequence rather than policy. The ADMM formulation in our work is closely related to Mordatch and Todorov (2014), but we focus on adversarial regularization for policy learning.

3. Method

We introduce VERONICA, our proposed adversarially regularized approach which combines the strength of policy learning and trajectory optimization. First, we define an adversarial regularizer and explain how it improves smoothness and robustness of neural control policies; Second, we describe an ADMM-based algorithm that solves the full joint optimization problem; Third, we develop an extension to our proposed adversarial regularization approach — Stackelberg adversarial regularization. We consider the neural control policy learning process guided by N optimal trajectories

$\{\mathbf{X}, \mathbf{U}\} = \{\mathbf{X}_i, \mathbf{U}_i \mid i = 1, \dots, N\}$, and each optimal trajectory $\{\mathbf{X}_i, \mathbf{U}_i\}$ consists of T state-control pairs $\{\mathbf{x}_i^t \in \mathbb{R}^{d_x}, \mathbf{u}_i^t \in \mathbb{R}^{d_u} \mid t = 1, \dots, T\}$, where \mathbf{x}_i^t and \mathbf{u}_i^t denote the robot state and the control, respectively. In this study, the robot state corresponds to the joint positions, velocities and task parameters such as goal configurations, while the control corresponds to the joint torque. Moreover, let $\pi(\cdot|\mathbf{W})$ denotes the neural control policy, where \mathbf{W} denotes the associated parameters.

3.1. Adversarial Regularization for Neural Control Policy

To promote smoothness of the neural control policy, we consider the following adversarial discrepancy measure:

$$r_{\text{adv}}(\mathbf{x}, \mathbf{W}) = \max_{\|\delta\| \leq \epsilon} r(\mathbf{x}, \mathbf{W}, \delta) = \max_{\|\delta\| \leq \epsilon} \|\pi(\mathbf{x}|\mathbf{W}) - \pi(\mathbf{x} + \delta|\mathbf{W})\|^2,$$

where $\|\cdot\|$ denotes the ℓ_2 norm, $\delta \in \mathbb{R}^{d_x}$ is the adversarial perturbation injected to the state vector \mathbf{x} , and $\epsilon > 0$ is the perturbation strength. Such an adversarial discrepancy measure $r_{\text{adv}}(\mathbf{x}, \mathbf{W})$ essentially computes the maximal deviation of the neural control policy output at state \mathbf{x} given an input perturbation δ whose ℓ_2 norm is bounded by ϵ .

We then apply the adversarial discrepancy measure to control the smoothness of the neural control policy. Specifically, we solve the following joint optimization problem:

$$\begin{aligned} \min_{\mathbf{X}, \mathbf{U}, \mathbf{W}} \quad & \sum_{i=1}^N \mathcal{L}(\mathbf{X}_i, \mathbf{U}_i) + \mathcal{Q}_{\text{BC}}(\mathbf{X}, \mathbf{U}, \mathbf{W}) + \alpha \mathcal{R}_{\text{adv}}(\mathbf{X}, \mathbf{W}), \\ \text{s.t.} \quad & \mathbf{x}^{t+1} = f(\mathbf{x}^t, \mathbf{u}^t), \mathbf{x}^0 = \mathbf{x}_{\text{init}}, \mathbf{X} \in \mathcal{X}, \mathbf{U} \in \mathcal{U}, \end{aligned} \quad (1)$$

where $\mathcal{L}(\mathbf{X}_i, \mathbf{U}_i)$ denotes the loss function of the trajectory optimization (TO) for the i^{th} trajectory, $\mathcal{Q}_{\text{BC}}(\mathbf{X}, \mathbf{U}, \mathbf{W})$ denotes the loss function for policy learning:

$$\mathcal{Q}_{\text{BC}}(\mathbf{X}, \mathbf{U}, \mathbf{W}) = \frac{1}{N} \sum_{i,t} \|\pi(\mathbf{x}_i^t|\mathbf{W}) - \mathbf{u}_i^t\|^2,$$

$\mathcal{R}_{\text{adv}}(\mathbf{X}, \mathbf{W})$ is the adversarial regularizer for controlling the smoothness of the policy:

$$\mathcal{R}_{\text{adv}}(\mathbf{X}, \mathbf{W}) = \frac{1}{N} \sum_{i,t} r_{\text{adv}}(\mathbf{x}_i^t, \mathbf{W}) = \frac{1}{N} \sum_{i,t} \max_{\|\delta_i^t\| \leq \epsilon} \|\pi(\mathbf{x}_i^t|\mathbf{W}) - \pi(\mathbf{x}_i^t + \delta_i^t|\mathbf{W})\|^2,$$

and α is the regularization coefficient weighting between the $\mathcal{Q}_{\text{BC}}(\mathbf{X}, \mathbf{U}, \mathbf{W})$ and \mathcal{R}_{adv} .

Solving the optimization problem in Eq. (1) learns a neural control policy that not only minimizes the TO loss and the behavior cloning loss, but also encourages the adversarial discrepancy measure of the policy to be small at every state of the optimal trajectories.

(I) Adversarial Regularization Improves Generalization: Existing methods usually train neural control policies by only minimizing the trajectory optimization loss and behavior cloning loss. Due to the high capacity of deep neural networks, the learned neural control policies are often over-complex and highly non-smooth. This is inconsistent with observations that many optimal control policies for robots are smooth. Here we exclude the problem involving physical contact dynamics, which exhibits discontinuous and non-smooth phenomenon. Smoothness requires a small perturbation to the state vector \mathbf{x} to only yield a small change to the policy output (Figure 1). Such a property is desirable in robotics tasks, since they often involve differential equations with high-order continuity properties. Therefore, improving smoothness of the learned policy can improve its ability to generalize to states unseen during training.

VERONICA naturally promotes the desired smoothness by imposing a high penalty when the adversarial perturbation δ yields a large deviation to the policy output. More precisely, $r_{\text{adv}}(\mathbf{x}, \mathbf{W})$ essentially upper bounds the deviation of the policy output due to the adversarial perturbation δ with respect to the state \mathbf{x} , and therefore can be viewed as a measure of the local Lipschitz constant within a small neighborhood of \mathbf{x} , i.e., $C_{\mathbf{x}} = \sup_{\|\delta\| \leq \epsilon} \frac{\|\pi(\mathbf{x}|\mathbf{W}) - \pi(\mathbf{x} + \delta|\mathbf{W})\|}{\|\delta\|}$. Accordingly, our proposed adversarial regularizer penalizes the average discrepancy measures of the neural control policy at all trajectory points, which enforces its local Lipschitz continuity.

(II) Adversarial Regularization Gains Robustness: Robot systems measure their states from sensors, which are prone to stochastic or systematic sensor errors. VERONICA naturally gains robustness against such disturbances. Specifically, the adversarial perturbation in VERONICA can be viewed as a proxy to the errors. Therefore, our approach does not require prior knowledge of them. In comparison, existing methods for handling such errors usually assume specific forms, e.g., independent Gaussian noise, which can be restrictive in practice.

Moreover, as suggested in [Asadi et al. \(2018\)](#), the Lipschitz continuity is essential to robustness, especially for control and reinforcement learning problems. This is because for policies without the Lipschitz continuity property, a small error in sensor measurement or state transition potentially leads to a drastic change to the policy output. Due to the dynamic nature of the control problem, it will further yield significant error compounding during policy roll-out. Moreover, when the models used to describe robot dynamics mismatch the real robot, such compounding system errors can be catastrophic. Quantitatively, the upper bound for policy robustness under state disturbance, measured by compounding value function discrepancy, is proportional to the Lipschitz constant of the neural control policy. As the VERONICA approach can effectively control the local Lipschitz continuity of the neural control policy, such an issue can be mitigated. A theoretical analysis can be found in Appendix G of the supplementary material ([Zhao et al., 2021](#)).

3.2. Combined Trajectory Optimization and Adversarially Regularized Policy Learning

We apply ADMM ([Zhao et al., 2020](#); [Zhou and Zhao, 2020](#)) to solve the optimization problem in Eq. (1). Specifically, we reparameterize Eq. (1) into a decomposable form by introducing two auxiliary sets of state and control variables: $(\mathbf{X}^{\text{TO}}, \mathbf{U}^{\text{TO}})$ represents the trajectory samples generated by trajectory optimization (TO), and $(\mathbf{X}^{\text{PL}}, \mathbf{U}^{\text{PL}})$ are copies of $(\mathbf{X}^{\text{TO}}, \mathbf{U}^{\text{TO}})$ for policy learning. Accordingly, the optimization problem in Eq. (1) is reformulated as:

$$\begin{aligned} \min_{\mathbf{X}^{\text{TO,PL}}, \mathbf{U}^{\text{TO,PL}}, \mathbf{W}} \quad & \sum_{i=1}^N \mathcal{L}(\mathbf{X}_i^{\text{TO}}, \mathbf{U}_i^{\text{TO}}) + \mathcal{Q}_{\text{BC}}(\mathbf{X}^{\text{PL}}, \mathbf{U}^{\text{PL}}, \mathbf{W}) + \alpha \mathcal{R}_{\text{adv}}(\mathbf{X}^{\text{PL}}, \mathbf{W}) \\ \text{s.t.} \quad & \mathbf{X}^{\text{TO}} = \mathbf{X}^{\text{PL}}, \mathbf{U}^{\text{TO}} = \mathbf{U}^{\text{PL}}. \end{aligned} \quad (2)$$

ADMM splits the above optimization problem into N individual TO problems and a policy learning problem to be solved in an iterative manner. Let $\lambda_{\mathbf{X}_i}^p, \lambda_{\mathbf{U}_i}^p$ denote the dual variables at the p^{th} iteration and $\rho_x, \rho_u > 0$ denote the penalty parameters. The ADMM primal and policy updates are:




Figure 1: Illustration of policy smoothness at state \mathbf{x} and control \mathbf{u} . If the policy $\pi(\cdot|W)$ is smooth around \mathbf{x} , the perturbed state $\tilde{\mathbf{x}}$ will produce a control \mathbf{u}' similar to \mathbf{u} . If the policy $\pi(\cdot|W)$ is non-smooth around \mathbf{x} , the output control \mathbf{u}'' would deviate significantly from \mathbf{u} .

$$\begin{aligned} \mathbf{X}_i^{\text{TO},p+1}, \mathbf{U}_i^{\text{TO},p+1} = \arg \min_{\mathbf{X}_i, \mathbf{U}_i} & \mathcal{L}(\mathbf{X}_i, \mathbf{U}_i) + \frac{\rho_x}{2} \|\mathbf{X}_i - \mathbf{X}_i^{\text{PL},p} + \lambda_{\mathbf{X}_i}^p\|^2 \\ & + \frac{\rho_u}{2} \|\mathbf{U}_i - \mathbf{U}_i^{\text{PL},p} + \lambda_{\mathbf{U}_i}^p\|^2, \quad (\text{primal TO update}) \end{aligned} \quad (3)$$

$$\mathbf{W}^{p+1} = \arg \min_{\mathbf{W}} \mathcal{Q}_{\text{BC}}(\mathbf{X}^{\text{PL},p}, \mathbf{U}^{\text{PL},p}, \mathbf{W}) + \mathcal{R}_{\text{adv}}(\mathbf{X}^{\text{PL},p}, \mathbf{W}), \quad (\text{policy update}) \quad (4)$$

$$\begin{aligned} \mathbf{X}_i^{\text{PL},p+1}, \mathbf{U}_i^{\text{PL},p+1} = \arg \min_{\mathbf{X}_i, \mathbf{U}_i} & \mathcal{Q}_{\text{BC}}(\mathbf{X}_i^{\text{PL},p}, \mathbf{U}_i^{\text{PL},p}, \mathbf{W}^{p+1}) + \frac{\rho_x}{2} \|\mathbf{X}_i^{\text{TO},p+1} - \mathbf{X}_i + \lambda_{\mathbf{X}_i}^p\|^2 \\ & + \frac{\rho_u}{2} \|\mathbf{U}_i^{\text{TO},p+1} - \mathbf{U}_i + \lambda_{\mathbf{U}_i}^p\|^2. \quad (\text{primal PL update}) \end{aligned} \quad (5)$$

Primal TO update: The update in Eq. (3) involves TO and is solved by either direct optimization methods [Drnach and Zhao \(2021\)](#) or indirect methods such as differential dynamic programming (DDP), as described in [Tassa et al. \(2014\)](#) and [Jacobson and Mayne \(1970\)](#). We defer details of the DDP algorithm to Appendix A in the supplementary material ([Zhao et al., 2021](#)).

Policy update: Note that Eq. (4) is a min-max optimization problem. For notation simplicity, we omit the iteration index p , and we rewrite it as

$$\mathbf{W} = \arg \min_{\mathbf{W}} \mathcal{Q}_{\text{BC}}(\mathbf{X}^{\text{PL}}, \mathbf{U}^{\text{PL}}, \mathbf{W}) + \frac{\alpha}{N} \sum_{i,t} \max_{\|\delta_i^t\| \leq \epsilon} r(\mathbf{x}_i^{\text{PL},t}, \mathbf{W}, \delta_i^t). \quad (6)$$

To solve Eq. (6), we apply an alternating gradient descent/ascent algorithm. Specifically, at the s^{th} iteration, we first apply the projected gradient ascent algorithm to update δ_i^t for K steps,

$$\delta_i^{t,s} = \delta_i^{t,s,K}, \quad \text{where } \delta_i^{t,s,k} = \Pi \left[\delta_i^{t,s,k-1} + \eta_\delta \nabla_{\delta} r(\mathbf{x}_i^{\text{PL},t}, \mathbf{W}^s, \delta_i^{t,s,k-1}) \right] \text{ for } k = 2, \dots, K.$$

Here, $\delta_i^{t,s,1}$ is randomly sampled from $\mathcal{N}(0, \sigma^2 \mathbb{I})$, Π denotes projection to the ℓ_2 ball with a radius ϵ , and $\eta_\delta > 0$ denotes the step size. Then we apply a gradient descent (or stochastic gradient descent) step to \mathbf{W} ,

$$\mathbf{W}^s = \mathbf{W}^{s-1} - \eta_W [\nabla_{\mathbf{W}} \mathcal{Q}_{\text{BC}}(\mathbf{X}^{\text{PL}}, \mathbf{U}^{\text{PL}}, \mathbf{W}^s) + \frac{\alpha}{N} \sum_{i,t} \nabla_{\mathbf{W}} r(\mathbf{x}_i^{\text{PL},t}, \mathbf{W}^s, \delta_i^{t,s})]. \quad (7)$$

Primal PL update: The update in Eq. (5) solves an unconstrained differentiable optimization sub-problem, which can be efficiently solved for each trajectory using stochastic gradient descent.

Dual update: After the above three updates, we perform the dual update as follows:

$$\lambda_{\mathbf{X}_i}^{p+1} = \lambda_{\mathbf{X}_i}^p + \mathbf{X}_i^{\text{TO},p+1} - \mathbf{X}_i^{\text{PL},p+1}, \quad \lambda_{\mathbf{U}_i}^{p+1} = \lambda_{\mathbf{U}_i}^p + \mathbf{U}_i^{\text{TO},p+1} - \mathbf{U}_i^{\text{PL},p+1}. \quad (8)$$

After a certain number of iterations of the above primal-dual policy updates, the joint optimization in Eq. (2) achieves a consensus and the primal and dual residuals meet the ADMM stopping criteria. The overall algorithm is summarized in Algorithm 1 in Appendix B ([Zhao et al., 2021](#)).

3.3. Stackelberg Adversarial Regularization

One major limitation of the adversarial regularizer in Eq. (6) is that it solves a min-max-game-based optimization, where neither of the players can be advantageous. This is problematic because the adversarial player may generate over-strong perturbations that hinder generalization. To mitigate this issue, we employ Stackelberg adversarial regularization ([Zuo et al., 2021b](#)) to solve the policy update in Eq. (6) through a Stackelberg game formulation. In a Stackelberg game, there are two

players, a leader (the policy) and a follower (the perturbations). The leader acknowledges the strategy of the follower, such that it is always in an advantageous position. This effectively eliminates the over-strong perturbations.

To simplify the notation, we omit the indices on the trajectory sample points \mathbf{x} . We solve

$$\begin{aligned} \min_{\mathbf{W}} \mathcal{Q}_{\text{SAR}}(\mathbf{W}) &= \mathcal{Q}_{\text{BC}}(\mathbf{X}, \mathbf{U}, \mathbf{W}) + \frac{\alpha}{N} \sum r(\mathbf{x}, \mathbf{W}, \delta^K), \\ \text{s.t. } \delta^K(\mathbf{W}) &= U^K \circ U^{K-1} \circ \dots \circ U^1(\delta^0). \end{aligned} \quad (9)$$

The policy parameter \mathbf{W} in Eq. (9) is the leader, and the perturbation $\delta(\mathbf{W})$ is the follower. Here, \circ denotes operator composition, i.e., $f(\cdot) \circ g(\cdot) = f(g(\cdot))$. Each U^k for $k = 1, \dots, K$ represents the k^{th} step update operator for the follower’s strategy. The operators are defined by pre-selected optimization algorithms such as stochastic gradient descent (SGD) or Adam (Kingma and Ba, 2014).

In Stackelberg adversarial training, the leader acknowledges the strategy of the follower by treating the perturbations (the follower) as a function of the policy parameters (the leader). Correspondingly, we solve for the policy parameters using gradient descent, where the Stackelberg gradient is

$$\frac{d\mathcal{Q}_{\text{SAR}}(\mathbf{W})}{d\mathbf{W}} = \underbrace{\frac{d\mathcal{Q}_{\text{BC}}(\mathbf{X}, \mathbf{U}, \mathbf{W})}{d\mathbf{W}} + \alpha \frac{\partial r(\mathbf{x}, \mathbf{W}, \delta^K)}{\partial \mathbf{W}}}_{\text{leader}} + \underbrace{\alpha \frac{\partial r(\mathbf{x}, \mathbf{W}, \delta^K)}{\partial \delta^K} \frac{d\delta^K}{d\mathbf{W}}}_{\text{leader-follower interaction}}. \quad (10)$$

In comparison, the conventional adversarial regularization in Eq. (6) uses only the leader term and does not consider the leader-follower interaction.

The most expensive term to compute in Eq. (10) is $d\delta^K/d\mathbf{W}$. Recall that we have $\delta^k = U^k(\delta^{k-1})$, where U^k is an update operator, e.g., a one-step gradient ascent. As a short-hand, we write

$$\delta^k(\mathbf{W}) = \delta^{k-1}(\mathbf{W}) + \Delta(\mathbf{x}, \delta^{k-1}(\mathbf{W}), \mathbf{W}),$$

where $\Delta(\mathbf{x}, \delta^{k-1}(\mathbf{W}), \mathbf{W})$ signifies the update from δ^{k-1} to δ^k . Then we have

$$\frac{d\delta^k}{d\mathbf{W}} = \frac{d\delta^{k-1}}{d\mathbf{W}} + \frac{\partial \Delta(\mathbf{x}, \delta^{k-1}, \mathbf{W})}{\partial \mathbf{W}} + \frac{\partial \Delta(\mathbf{x}, \delta^{k-1}, \mathbf{W})}{\partial \delta^{k-1}} \frac{d\delta^{k-1}}{d\mathbf{W}}.$$

This recursive differentiation can be efficiently computed using deep learning libraries, such as *PyTorch* (Paszke et al., 2019). Please refer to Zuo et al. (2021b) for more details. The overall adversarial regularization algorithm is shown in Algorithm 2 in Appendix C (Zhao et al., 2021).

4. Experiments

We evaluate VERONICA on cart-pole swing-up, Kuka arm manipulation, and hopper locomotion tasks. The experiments are shown in the video¹. We compare smoothness, generalization, and robustness of policies trained with Gaussian perturbations, conventional adversarial regularization (VERONICA-AR), and SAR (VERONICA-SAR). We do not include tangent propagation due to the excessive computational requirements to compute the Jacobian. We also demonstrate that the neural control policy is able to handle simple multi-modal dynamics for the pick and place task.

For Kuka manipulation tasks, the simulation environment is implemented in *PyBullet* (Coumans and Bai, 2016–2021). TO is solved using DDP implemented in *Crocodyl* (Mastalli et al., 2020). For hopper locomotion tasks, we implement both the simulation environment and a direct TO algorithm in *Drake* (Tedrake and the Drake Development Team, 2019). The adversarially regularized policy learning algorithm is implemented in *PyTorch* (Paszke et al., 2019) and *Higher* (Grefenstette et al., 2019). The implementation details can be found in Appendix D (Zhao et al., 2021).

1. The link to the video is <https://youtu.be/2z1AC9Xs8Bg>.

Policy Smoothness: We qualitatively examine the smoothness of our neural control policy by inspecting a typical policy roll-out for cart-pole swing-up and Kuka arm reaching tasks, as shown in Figure 2. Figure 2(a) shows the smoothness comparison during a cart-pole swing-up. VERONICA produced visually smoother force sequences comparing to Gaussian perturbation. Figure 2(b) displays the torque sequence of Kuka joint 2 during a reaching task. The policy trained by Gaussian perturbation generates a non-smooth torque profile around the initial position of the task, indicating that the Gaussian perturbation is not sufficient to prevent overfitting at the initial phase of the trajectory, where the torque changes relatively quickly with respect to state. In comparison, the VERONICA-AR and VERONICA-SAR policies produce smoother control sequences that track the baseline closely. To inspect the smoothness of the neural control policies, we plot the torque output on Kuka joint 2 against the joint angle in Figure 2(c). VERONICA successfully penalize against the non-smooth peak that appeared in the torque profile of the Gaussian perturbed policy.




Figure 2: Comparison of control output smoothness for cart-pole and Kuka arm reaching tasks. Trajectory optimization baseline is marked as a dashed line. (a) Time sequence of forces applied onto the cart during swing-up. (b-c) Torque output for Kuka joint 2 with respect to time and joint 2 angle.




Figure 3: Cost percentile plot for 3-DOF arm reaching task with 100 different initializations and under different disturbances on sensor measurement. Disturbances are drawn from a uniform distribution bounded by ζ . Policies trained with no perturbation, Gaussian perturbation, VERONICA-AR, and VERONICA-SAR are compared against an undisturbed TO baseline. The plot is capped at 2 times the maximum baseline cost. A cost curve that exceeds the plotting cap indicates that a percentage of policy roll-outs lead to unstable robot motion.




Figure 4: Example of an undisturbed policy roll-out for a 3-DOF manipulator reaching task where Gaussian perturbation fails. (a) Cumulative cost for policy roll-out (b-c) Torque outputs on joints 2 and 4.

Generalization Performance: To evaluate the generalization performances of VERONICA, we perform policy roll-outs with 100 different initializations in an undisturbed environment (see

Figure 3(a)). The adversarially regularized policies produce lower costs because the policies trained with no perturbation or Gaussian perturbation are unable to generate stable robot motions under some initializations. Figure 4 displays an example of an arm reaching task that Gaussian perturbation cannot handle. Although a vast majority of roll-outs with the VERONICA-AR policy are stable, a small percentage (2%) produces unstable robot motions that fail to achieve the task. In contrast, the VERONICA-SAR policy leads to stable and near-optimal robot motions across all attempts, demonstrating that VERONICA-SAR enhances numerical stability comparing to VERONICA-AR.

Policy Robustness: We evaluate our policies’ robustness against three different kinds of disturbances. For sensor noise and environmental uncertainty, we add a uniform noise bounded by an ℓ_∞ -norm ball with radius ζ onto the sensor measurement and state transition, respectively. As for model mismatch, we modify the URDF file used in policy roll-out by decreasing the mass of each robot link by 0.25 kg.

We first compare the policies’ robustness against different magnitudes of sensor noise, as shown in Figure 3(b-c). While Gaussian perturbation does provide some robustness comparing to the unregularized policy, VERONICA-AR and VERONICA-SAR consistently outperforms the Gaussian perturbation. Furthermore, VERONICA-AR deviates significantly from the undisturbed TO baseline under a strong sensor noise ($\zeta = 0.05$), while VERONICA-SAR remains able to produce stable robot motion and closely track the TO baseline.

Table 1: Task Error for 3-DOF Manipulator Reaching Task ($\zeta = 0.01$, *Unit: m*)

	Gaussian	VERONICA-AR	VERONICA-SAR
Undisturbed	$1.62e-1 \pm 5.05e-2$	$6.26e-2 \pm 3.96e-2$	$6.39e-2 \pm 2.70e-2$
Sensor Error	$1.75e-1 \pm 7.85e-2$	$7.11e-2 \pm 7.57e-2$	$6.61e-2 \pm 2.42e-2$
Environment Uncertainty	$1.73e-1 \pm 7.59e-2$	$8.66e-2 \pm 1.03e-1$	$7.75e-2 \pm 3.99e-2$
Model Mismatch	$2.14e-1 \pm 8.26e-2$	$5.24e-2 \pm 3.27e-2$	$1.23e-1 \pm 2.16e-2$

Table 1 shows the average task errors - the distance between the goal and the actual final positions for the robot arm’s end-effector - and their standard deviation for 100 manipulator reaching tasks under different types of disturbances. VERONICA provides significantly lower task errors across all clean and disturbed experiments. Furthermore, VERONICA-SAR leads to a lower standard deviation than VERONICA-AR, indicating that the policy learned by VERONICA-SAR is less prone to outliers comparing to VERONICA-AR.

Application to Higher-DOF Manipulators: We investigate how the performance of VERONICA-SAR scales to higher state and control dimensions by evaluating the task errors of manipulator reaching tasks for 3, 5, and 7-DOF Kuka arms (Table 2). The task error increases with the dimensionality of the problem, but not significantly. Note that the 5 and 7-DOF experiments involve manipulation in the 3-D space, which lead to much higher problem complexity than the planar 3-DOF configuration, and require larger neural control policies. Figure 5 indicates that similar to the 3-DOF cases, the proposed Stackelberg adversarial regularization benefits both generalization and robustness performance compared to Gaussian regularization in the 7-DOF Kuka arm reaching tasks.

Table 2: Median Task Errors for M -DOF Manipulator (*Unit: m*)

$M = 3$	$M = 5$	$M = 7$
$6.39e-2$	$1.23e-1$	$1.32e-1$

Preliminary Study of Learning Multimodal Dynamics: In the pick and place task, we train a policy to handle the control of the Kuka arm for both free-moving or object-holding scenarios. In order to train the policy applicable for both cases simultaneously, we include a discrete variable in the network input to signify the grasping state of the object. Figure 6 shows the arm’s torque




Figure 5: Cost percentile plot for 7-DOF arm reaching task with 100 different initializations and under different disturbances on sensor measurements. The plot is capped at 3 times the maximum baseline cost.

output for the same initialization, with or without an object. For simplicity, this experiment assumes that only one object with a known mass, and the object is fixed to a pre-specified position in the gripper when grasped by the arm. In the future, the adaptability of the policy can be improved by augmenting the input with more information such as the weight of the object and the relative position between the object and the gripper.

Application to Hybrid Locomotion Systems: We apply VERONICA in hopper locomotion tasks to evaluate the performance of VERONICA in a single leg 5-DOF hopper system, where the hybrid locomotion trajectories involve intermittent contacts with the terrain. We compare the cost percentile plot between the TO baseline and VERONICA-SAR, as displayed in Figure 7. Note that the open-loop rollout of trajectories generated by TO baseline performs poorly in simulation due to the model mismatch between TO and simulation environments. In contrast, the policy trained with VERONICA-SAR generates a lower cost hopper motions due to the robustness against model mismatch provided by adversarial perturbation. A visual comparison can be found in the video.

5. Conclusion

We present VERONICA, an adversarial regularization framework for combined trajectory optimization and policy learning. We show that the proposed regularizer improves generalization and robustness by enforcing Lipschitz continuity of the policy. Additionally, we propose to further stabilize training by formulating the adversarial regularization as a Stackelberg game. The experiment results in robot manipulation scenarios show that our approach helps to improve the smoothness of the learned policy, which results in a more stable robot motions and lower policy execution costs. Additionally, we demonstrate that policies trained with VERONICA are able to robustly handle various types of disturbances.

Our future work will (i) evaluate the performance of VERONICA in the presence of more types of perturbations and uncertainties, such as varying link moment of inertia and kinematic parameters; (ii) extend VERONICA to solve more complex manipulation and locomotion problems involving physical contact and enhance robustness to contact uncertainties. Adaptive adversarial training, where perturbations are generated by an additional network, can be incorporated to generate variable perturbation radius around contact points.




Figure 6: Comparison for the control policy outputs with or without grasping a 5kg object.




Figure 7: Cost percentile plot for hopper locomotion task with 100 different initializations.

Acknowledgments

This project was partially supported by *Ford Motor Company* under the Ford-GT Alliance Program.

References

- Kavosh Asadi, Dipendra Misra, and Michael Littman. Lipschitz continuity in model-based reinforcement learning. In *International Conference on Machine Learning*, pages 264–273. PMLR, 2018.
- John T Betts. Survey of numerical methods for trajectory optimization. *Journal of guidance, control, and dynamics*, 21(2):193–207, 1998.
- Stephen Boyd, Neal Parikh, and Eric Chu. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- Marcus A Brubaker, Leonid Sigal, and David J Fleet. Estimating contact dynamics. In *2009 IEEE 12th International Conference on Computer Vision*, pages 2389–2396. IEEE, 2009.
- Chunhui Chen and Olvi L Mangasarian. A class of smoothing functions for nonlinear and mixed complementarity problems. *Computational Optimization and Applications*, 5(2):97–138, 1996.
- Xi Chen, Ali Ghadirzadeh, Mårten Björkman, and Patric Jensfelt. Adversarial feature training for generalizable robotic visuomotor control. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1142–1148. IEEE, 2020.
- Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.
- Marc Peter Deisenroth, Gerhard Neumann, Jan Peters, et al. A survey on policy search for robotics. *Foundations and trends in Robotics*, 2(1-2):388–403, 2013.
- Luke Drnach and Ye Zhao. Robust trajectory optimization over uncertain terrain with stochastic complementarity. *IEEE Robotics and Automation Letters*, 6(2):1168–1175, 2021.
- Alexis Duburcq, Yann Chevaleyre, Nicolas Bredeche, and Guilhem Bo eris. Online trajectory planning through combined trajectory optimization and function approximation: Application to the exoskeleton atalante. In *IEEE International Conference on Robotics and Automation*, 2020.
- Edward Grefenstette, Brandon Amos, Denis Yarats, Phu Mon Htut, Artem Molchanov, Franziska Meier, Douwe Kiela, Kyunghyun Cho, and Soumith Chintala. Generalized inner loop meta-learning. *arXiv preprint arXiv:1910.01727*, 2019.
- Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. *arXiv preprint arXiv:1906.12340*, 2019.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 4572–4580, 2016.
- David H Jacobson and David Q Mayne. Differential dynamic programming. 1970.

- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. Smart: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. *arXiv preprint arXiv:1911.03437*, 2019.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Scott Kuindersma, Robin Deits, Maurice Fallon, Andrés Valenzuela, Hongkai Dai, Frank Permenter, Twan Koolen, Pat Marion, and Russ Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous robots*, 40(3):429–455, 2016.
- Benoit Landry, Hongkai Dai, and Marco Pavone. Seagul: Sample efficient adversarially guided learning of value functions. In *Learning for Dynamics and Control*, pages 1–13. PMLR, 2021.
- Mathias Lechner, Ramin Hasani, Radu Grosu, Daniela Rus, and Thomas A Henzinger. Adversarial training is not ready for robot learning. *arXiv preprint arXiv:2103.08187*, 2021.
- Sergey Levine and Vladlen Koltun. Guided policy search. In *International conference on machine learning*, pages 1–9. PMLR, 2013a.
- Sergey Levine and Vladlen Koltun. Variational policy search via trajectory optimization. *Advances in neural information processing systems*, 26:207–215, 2013b.
- Sergey Levine and Vladlen Koltun. Learning complex neural network policies with trajectory optimization. In *International Conference on Machine Learning*, pages 829–837. PMLR, 2014.
- Yan Li, Ethan X Fang, Huan Xu, and Tuo Zhao. Implicit bias of gradient descent based adversarial training on separable data. 2020.
- Kendall Lowrey, Aravind Rajeswaran, Sham Kakade, Emanuel Todorov, and Igor Mordatch. Plan online, learn offline: Efficient learning and exploration via model-based control. *arXiv preprint arXiv:1811.01848*, 2018.
- Carlos Mastalli, Rohan Budhiraja, Wolfgang Merkt, Guilhem Saurel, Bilal Hammoud, Maximilien Naveau, Justin Carpentier, Ludovic Righetti, Sethu Vijayakumar, and Nicolas Mansard. Crocodyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- Igor Mordatch and Emo Todorov. Combining the benefits of function approximation and trajectory optimization. In *Robotics: Science and Systems*, volume 4, 2014.
- Igor Mordatch, Emanuel Todorov, and Zoran Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)*, 31(4):1–8, 2012.
- Jun Morimoto and Kenji Doya. Robust reinforcement learning. In *NIPS*, pages 1061–1067. Citeseer, 2000.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2817–2826. JMLR. org, 2017.
- Michael Posa, Cecilia Cantu, and Russ Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 33(1):69–81, 2014.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Stefan Schaal et al. Learning from demonstration. *Advances in neural information processing systems*, pages 1040–1046, 1997.
- Qianli Shen, Yan Li, Haoming Jiang, Zhaoran Wang, and Tuo Zhao. Deep reinforcement learning with robust and smooth policy. In *International Conference on Machine Learning*, pages 8707–8718. PMLR, 2020.
- Yuval Tassa, Nicolas Mansard, and Emo Todorov. Control-limited differential dynamic programming. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1168–1175. IEEE, 2014.
- Russ Tedrake and the Drake Development Team. Drake: Model-based design and verification for robotics, 2019. URL <https://drake.mit.edu>.
- Emanuel Todorov. A convex, smooth and invertible contact model for trajectory optimization. In *2011 IEEE International Conference on Robotics and Automation*, pages 1071–1076. IEEE, 2011.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.
- Heinrich Von Stackelberg. *Market structure and equilibrium*. Springer Science & Business Media, 2010.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation. *arXiv preprint arXiv:1904.12848*, 2019.

- Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *arXiv preprint arXiv:1901.08573*, 2019.
- Zhigen Zhao, Ziyi Zhou, Michael Park, and Ye Zhao. Sydebo: Symbolic-decision-embedded bilevel optimization for long-horizon manipulation in dynamic environments. *arXiv preprint arXiv:2010.11078*, 2020.
- Zhigen Zhao, Simiao Zuo, Tuo Zhao, and Ye Zhao. Supplementary material for adversarially regularized policy learning guided by trajectory optimization, 2021. URL <https://tinyurl.com/3cnyubv9>.
- Ziyi Zhou and Ye Zhao. Accelerated admm based trajectory optimization for legged locomotion with coupled rigid body dynamics. In *American Control Conference*, pages 5082–5089, 2020.
- Konrad Zolna, Scott Reed, Alexander Novikov, Sergio Gomez Colmenarej, David Budden, Serkan Cabi, Misha Denil, Nando de Freitas, and Ziyu Wang. Task-relevant adversarial imitation learning. *arXiv preprint arXiv:1910.01077*, 2019.
- Simiao Zuo, Chen Liang, Haoming Jiang, Pengcheng He, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Tuo Zhao. Arch: Efficient adversarial regularized training with caching. *arXiv preprint arXiv:2109.07048*, 2021a.
- Simiao Zuo, Chen Liang, Haoming Jiang, Xiaodong Liu, Pengcheng He, Jianfeng Gao, Weizhu Chen, and Tuo Zhao. Adversarial training as stackelberg game: An unrolled optimization approach. *arXiv preprint arXiv:2104.04886*, 2021b.