

Learning From Human Directional Corrections

Wanxin Jin¹, Todd D. Murphey¹, Zehui Lu¹, and Shaoshuai Mou¹, *Member, IEEE*

Abstract—This article proposes a novel approach that enables a robot to learn an objective function incrementally from human directional corrections. Existing methods learn from human magnitude corrections; since a human needs to carefully choose the magnitude of each correction, those methods can easily lead to over-corrections and learning inefficiency. The proposed method only requires human directional corrections—corrections that only indicate the direction of an input change without indicating its magnitude. We only assume that each correction, regardless of its magnitude, points in a direction that improves the robot’s current motion relative to an unknown objective function. The allowable corrections satisfying this assumption account for half of the input space, as opposed to the magnitude corrections that have to lie in a shrinking level set. For each directional correction, the proposed method updates the estimate of the objective function based on a cutting plane method, which has a geometric interpretation. We have established theoretical results to show the convergence of the learning process. The proposed method has been tested in numerical examples, a user study on two human–robot games, and a real-world quadrotor experiment. The results confirm the convergence of the proposed method and further show that the method is significantly more effective (higher success rate), efficient/effortless (less human corrections needed), and potentially more accessible (fewer early wasted trials) than the state-of-the-art robot learning frameworks.

Index Terms—Cutting-plane method, human–robot physical interaction, inverse reinforcement learning, learning from corrections, learning from demonstrations (LfD), motion planning.

I. INTRODUCTION

FOR tasks where robots work in proximity to human users, a robot is required to not only guarantee the accomplishment

Manuscript received 29 November 2021; revised 29 April 2022; accepted 13 June 2022. This work was supported in part by the National Science Foundation under Award 1837515. The work of Dr. Shaoshuai Mou’s was supported in part by the Applications for Learning Machines (REALM) Consortium of Northrop Grumman Corporation, in part by the Rolls-Royce Corporation, and in part by the NASA University Leadership Initiative (ULI). This paper was recommended for publication by Associate Editor P. M. Wensing and Editor W. Burgard upon evaluation of the reviewers’ comments. (*Corresponding author: Wanxin Jin.*)

Wanxin Jin is with the General Robotics, Automation, Sensing and Perception (GRASP) Laboratory, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: wanxinjin@gmail.com).

Todd D. Murphey is with the Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208 USA (e-mail: t-murphey@northwestern.edu).

Zehui Lu and Shaoshuai Mou are with the School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN 47906 USA (e-mail: zehuilu789@gmail.com; shaoshuai.mou@aya.yale.edu).

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by Purdue’s Institutional Review Board under Application No. IRB-2021-1539, and performed in the line with Title 45 of the Code of Federal Regulations, Part 46 (45 CFR 46).

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/TRO.2022.3190221>.

Digital Object Identifier 10.1109/TRO.2022.3190221




Fig. 1. Illustration of learning from directional corrections proposed in this article. The robot’s current motion ξ_{θ_k} optimizes an objective function parameterized by θ_k . A human user applies a directional correction a_k to ξ_{θ_k} . Note that the magnitude $\|a_k\|$ of this correction does not matter. After receiving a_k , the robot updates its objective function parameter to θ_{k+1} . This motion correction update process repeats until the convergence of the objective function parameter.

of the task but also complete it in a way that a human user prefers. Different users may have different preferences about how the robot should perform the task. Such customized requirements usually lead to a considerable workload of robot programming, which requires human users to have high expertise and skills.

To circumvent the above limitations of robot programming, learning from demonstrations (LfD) empowers nonexpert users to program robots by providing demonstrations. In LfD [1], a user first provides a robot with demonstrations in a *one-time* manner, and then the robot learns a control policy or objective function *offline* from the demonstrations. While achieving the notable success in various applications [2]–[4], the *one-time* and *offline* nature of LfD could lead to some challenges. For example, when the demonstration data is insufficient to infer an objective function due to the low informativeness [5] or deviation from optima [6], the demonstrations have to be recollected and the robot has to be retrained. This is particularly inconvenient for robots with high degrees of freedom (DoFs).

A recent line of work [6]–[9] addresses the above challenges of LfD by developing a scheme that enables a robot to learn an objective function from *user’s feedback or corrections*. Fig. 1 is an example of those learning schemes. At each iteration, a human user does not need to provide optimal demonstrations to the robot, but merely a correction which is an *incremental improvement* to the robot’s current motion. The robot then leverages the correction to update its objective function. Compared to LfD, this incremental learning scheme reduces the workload of a user, especially for those (such as novices) who cannot provide the optimal demonstrations in a one-time manner [10]. Despite its promise, the state-of-the-art methods [6]–[9] still face some challenges as follows.

First, providing a valid correction—the one that improves the robot motion—can be nontrivial. To obtain a valid correction, the human user must carefully choose the *magnitude* of a correction.

As a robot gets closer to the “desired motion” (i.e., the motion in the human’s perspective), the allowable range of the magnitude of valid corrections gets smaller (see Section II for the detailed explanation). Thus, the human can easily overcorrect a robot near the desired motion, i.e., giving too large corrections that adversely drive the robot away from the desired motion. Such difficulty makes robot learning inefficient (we will experimentally show this in Sections V-C and VI).

Second, most of the existing methods [7]–[9] lack *theoretical guarantee* for the learning performance. While [6] shows that learning from human magnitude corrections can attain a bounded regret [11], this regret bound still cannot explicitly tell whether or not the learned objective function converges to the true one induced by all corrections.

To address the above challenges, this article proposes a new learning method that enables a robot to learn an objective function from human *directional* corrections with *theoretical convergence guarantee*. We highlight the following new features of the proposed method, as shown by Fig. 1.

1) The proposed learning scheme only requires the user’s *directional corrections*. A directional correction is a correction that only concerns directional information and does not necessarily need to be magnitude-specific. For instance, in teaching a mobile robot, a directional correction can be simply as “go left” or “go right” without dictating how far the robot should go. Furthermore, unlike existing work [6]–[9], the proposed approach can *directly* handle the sparse directional corrections without a preprocessing step of creating a human “intended” trajectory, which could introduce artifacts and unfavorably affect the learning performance.

2) We emphasize the theoretical foundations of the proposed method. First, the learning process is based on cutting-plane methods and thus has an intuitive geometric interpretation; second, we establish the theoretical guarantees to show the convergence of our method toward finding the true objective function induced by human directional corrections.

We have conducted extensive experiments to test the proposed method, including various numerical examples, a user study on two human–robot games, and a real-world quadrotor experiment. The results show that the proposed method is significantly more effective (higher success rate), efficient/effortless (less human corrections needed), and potentially more accessible (fewer early wasted trials) than the state-of-the-art robot learning frameworks.

A. Related Work

1) *One-Time Learning From Demonstrations*: To learn an objective function from demonstrations, routine methods are inverse optimal control [12]–[14] or inverse reinforcement learning [15]–[17], where an objective function is learned from optimal demonstrations and subsequently used for control or planning. Successful LfD applications include autonomous driving [2], robot manipulation [3], and motion planning [4]. Despite the notable success [18]–[20], LfD could be inconvenient in practice. First, demonstrations in LfD are usually collected in a *one-time* manner and the robot learning is *offline*. In the case

where demonstration data is insufficient to recover an objective function, such as low data informativeness as discussed in [5], or significantly deviating from the optima, the data has to be recollected and the training has to be rerun. Second, existing LfD [13]–[17] normally requires the optimality of demonstrations, which are challenging to collect for robots with high DoFs. For example, when providing demonstrations for a humanoid robot, a user has to account for the robot’s motion of all DoFs in a spatially and temporally consistent manner [6].

2) *Incrementally Learning From Corrections/Feedback*: Compared to one-time LfD, learning from corrections or feedback enables a user to *incrementally* improve robot motion, making it more suitable for nonexpert users who cannot provide optimal demonstrations in a one-time fashion [10]. The key assumption is that robot motion after correction achieves a higher reward (or a lower cost) than before correction. Under this assumption, [6] develops a coactive learning method to update the robot objective function using user feedback. The user feedback includes a human either selecting the top-ranked robot trajectory among all candidates or physically demonstrating a preferred trajectory to the robot. By defining a regret, which quantifies the average misalignment between the value of robot motion and that of the desired motion under the true objective function, [6] shows the convergence of the regret. But since regret only considers values of an objective function, it cannot directly tell and guarantee that the learned objective function itself is converging toward the true one. As we will show in Section V-C, the objective function can converge to a local solution instead of a true one.

Recently, [7]–[9] handle learning from corrections through the perspective of the partially observable Markov decision process (POMDP). Here, human corrections are viewed as observations about the unknown objective function parameter. By approximating the observation model and applying the maximum *a posteriori* estimation, they obtain a learning formula similar to the coactive learning [6]. Along with this perspective, some variants have been recently developed to particularly account for the uncertainties of human corrections. For example, [21] simultaneously estimates a rationality coefficient to characterize the rationality confidence of human corrections. The work [9] fits the inverse reinforcement learning into a Kalman filter framework to quantitatively capture the uncertainty of the estimation.

Both the above coactive learning and POMDP-based learning require a human to carefully choose the *magnitude* correction that improves robot motion. As we will detail in Section II, choosing a valid magnitude correction is difficult, especially when a robot gets closer to the desired motion, because the allowable magnitude of a correction has to lie in a shrinking level set. Near the desired motion, a human could easily overcorrect the robot, i.e., applying too large corrections that adversely drive the robot away from the desired motion, making the algorithm diverge or even fail. We will experimentally show this difficulty in Section V-C. Moreover, due to the sparsity of human corrections, all the above methods require a dedicated step to obtain a *human “intended” trajectory* for each correction. In the coactive learning, a robot needs to switch to a screening mode to obtain a human preferred trajectory, and in the POMDP-based method,




Fig. 2. Overview of the proposed algorithm.

a human intended trajectory is created by trajectory deformation [22]. These preprocessing steps may not only introduce more hyperparameters but also add some artifacts that could undermine the learning performance [8].

Very recently, [23] directly learns a desired trajectory from human physical interactions, then a robot tracks the learned trajectory using linear–quadratic regulator (LQR) controllers. They minimize a trajectory distance loss, and the learning update is based on the assumption that the direction of a human correction is exactly aligned with the deepest gradient descent of the distance loss. While that work is related to our work in the sense that they both use the direction of a human correction, our work, however, focuses on learning an objective function instead of learning the desired trajectory directly. Importantly, we do not restrict the direction of a correction to be exactly aligned with the deepest gradient descent, but any correction as long as it has a direction that decreases the cost of robot motion. The above two distinctions lead to a new method of our work, which has a strong theoretical guarantee for learning convergence.

Finally, we want to mention that most of the existing methods [7]–[9], [23] lack theoretical guarantee about the learning performance. The only work that attempts to do so is [6], where a regret bound is shown. As discussed previously, such regret is an indicator of the misalignment of the values of the objective function and cannot directly tell or guarantee the convergence of the objective function itself. In this article, we will directly characterize the convergence of the learned objective function toward the true function that is induced by human corrections.

B. Contributions

To give readers a high-level picture of the proposed method, we show an overview in Fig. 2. We claim the following contributions of the proposed method.

- 1) The proposed method learns an objective function from human *directional corrections*. It only requires that a correction, regardless of its magnitude, points in a direction of improving robot motion. As we will show in Sections II and V-C, the allowable corrections that satisfy this requirement always account for half of the input space, making it more flexible for users to choose corrections from. Also, unlike existing methods, the proposed method directly learns from sparse directional corrections without a *preprocessing* step of creating an “intended” trajectory.
- 2) The proposed learning algorithm has strong theoretical foundations. First, the proposed learning method is based on the cutting plane technique, which has straightforward geometric interpretations, as shown in Fig. 2. Second, we

have established the theoretical results to show the convergence of the method toward finding the true objective function induced by all corrections.

Numerical examples, a user study on two human–robot games, and a real-world quadrotor experiment confirm the efficacy and convergence of the proposed algorithm. The user study and real-world experiment further show that the proposed method is significantly more effective (higher success rate), efficient/effortless (fewer corrections needed), and potentially more accessible (fewer early wasted trials) than the state-of-the-art methods.

In what follows, Section II describes the problem formulation. Section III outlines the main algorithm and presents its geometric interpretation. Section IV provides the theoretical results of the algorithm and its detailed implementation. Various numerical examples and comparisons are given in Section V. Section VI presents a user study, and Section VII provides a real-world experiment. Conclusions are drawn in Section VIII. The appendix includes the proof, discussion, and possible extension of this work.

II. PROBLEM FORMULATION

Consider a robot with the following dynamics and the initial condition:

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t, \mathbf{u}_t), \quad \text{with } \mathbf{x}_0 \quad (1)$$

where $\mathbf{x}_t \in \mathbb{R}^n$ is the robot state, $\mathbf{u}_t \in \mathbb{R}^m$ is the control input, $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^n$ is differentiable, and $t = 1, 2, \dots$ is the time step. As commonly used in objective learning methods [5]–[10], [14]–[17], suppose that the robot cost function has the following parameterized form:

$$J(\mathbf{u}_{0:T}, \boldsymbol{\theta}) = \sum_{t=0}^T \boldsymbol{\theta}^\top \boldsymbol{\phi}(\mathbf{x}_t, \mathbf{u}_t) + h(\mathbf{x}_{T+1}) \quad (2)$$

where $\boldsymbol{\phi} : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^r$ is a vector of the *predefined* features (basis functions); $\boldsymbol{\theta} \in \mathbb{R}^r$ is a vector of weights, which are *tunable*; and $h(\mathbf{x}_{T+1})$ is the final cost on the robot final state \mathbf{x}_{T+1} , such as penalizing the distance between \mathbf{x}_{T+1} and a given goal state \mathbf{x}^{goal} , $h(\mathbf{x}_{T+1}) = \|\mathbf{x}^{\text{goal}} - \mathbf{x}_{T+1}\|^2$. Note that depending on specific applications, the final cost term $h(\mathbf{x}_{T+1})$ can be either absent, i.e., $h(\mathbf{x}_{T+1}) = 0$, or also parameterized in a weighted-feature form with the weights to be learned jointly with $\boldsymbol{\theta}$. The method developed in this article can sufficiently handle either cases with little modifications. For a fixed choice of $\boldsymbol{\theta}$, the robot plans a sequence of inputs $\mathbf{u}_{0:T}$ over time horizon T by (locally) optimizing the cost function (2) subject to (1), producing a trajectory

$$\boldsymbol{\xi}_\theta = \{\mathbf{x}_{0:T+1}^\theta, \mathbf{u}_{0:T}^\theta\}. \quad (3)$$

In the following text, we occasionally write the cost function (2) as $J(\boldsymbol{\theta})$ for ease of readability.

For a specific task, suppose that a human user expects the robot to minimize an *implicit* cost function $J(\boldsymbol{\theta}^*)$ in the same form of (2) but with *unknown* $\boldsymbol{\theta}^*$. We call $\boldsymbol{\theta}^*$ the *true weight vector*. In general, a human user may neither explicitly write down the value of $\boldsymbol{\theta}^*$ nor demonstrate the corresponding *desired*

trajectory ξ_{θ^*} , but the human user can tell whether the robot motion trajectory is *desired* or not. A robot trajectory is desired if it minimizes the implicit $J(\theta^*)$; otherwise, it is not desired. In order for the robot to learn $J(\theta^*)$ (and thus generate the desired trajectory ξ_{θ^*}), the human user is only able to make corrections to the robot motion. After receiving each human correction, the robot updates its guess θ toward the true θ^* . This procedure has been shown in Fig. 1.

Specifically, the process of a robot learning from human directional corrections is iterative, as shown in Fig. 2. Each iteration basically includes three steps: robot planning (and execution), human correction, and robot update. Let $k = 1, 2, 3, \dots$ be the iteration index and let θ_k denote the robot's guess of the weight vector at iteration k . At $k = 1$, the robot is initialized with an arbitrary guess θ_1 . At iteration k , the robot first performs *motion planning*, i.e., solving ξ_{θ_k} by minimizing the cost function $J(\theta_k)$ in (2) subject to its dynamics in (1). During the robot executing ξ_{θ_k} , the human user observes ξ_{θ_k} and applies a *correction*, denoted by $\mathbf{a}_{t_k} \in \mathbb{R}^m$, to the robot in its input space. Here, $t_k \in \{0, 1, \dots, T\}$ is called *correction time*, indicating at which time step the correction \mathbf{a}_{t_k} is made. After receiving \mathbf{a}_{t_k} , the robot *updates* its guess θ_k to θ_{k+1} based on an update rule developed later.

One distinguishing feature of our method is that we assume that each human correction \mathbf{a}_{t_k} satisfies

$$\left\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \mathbf{a}_k \right\rangle > 0, \quad k = 1, 2, 3, \dots \quad (4)$$

Here

$$\mathbf{a}_k = \left[\mathbf{0}^T \quad \dots \quad \mathbf{a}_{t_k}^T \quad \dots \quad \mathbf{0}^T \right]^T \in \mathbb{R}^{m(T+1)} \quad (5)$$

with \mathbf{a}_{t_k} filled at the t_k th entry and $\mathbf{0} \in \mathbb{R}^m$ elsewhere. Here, $\langle \cdot, \cdot \rangle$ is the dot product, and $-\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$ is the gradient-descent of $J(\theta^*)$ with respect to $\mathbf{u}_{0:T}$ evaluated at the robot current trajectory $\xi_{\theta_k} = \{\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k}\}$. Note that (4) does not require a specific value to the magnitude of \mathbf{a}_{t_k} but requires its direction roughly around the gradient-descent direction of $J(\theta^*)$. It means that the correction \mathbf{a}_k aims to guide the robot motion ξ_{θ_k} toward a lower cost under $J(\theta^*)$ unless the trajectory is desired. Thus, we call \mathbf{a}_{t_k} a *directional correction*.

In practice, one always needs to account for human's noisy or imperfect corrections. Thus, we modify (4) as follows:

$$\begin{aligned} & \mathbb{E}_{\mathbf{a}_k} \left\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \mathbf{a}_k \right\rangle \\ &= \left\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \mathbb{E}_{\mathbf{a}_k}(\mathbf{a}_k) \right\rangle \\ &= \left\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \bar{\mathbf{a}}_k \right\rangle > 0, \quad k = 1, 2, \dots \end{aligned} \quad (6a)$$

where

$$\bar{\mathbf{a}}_k = \mathbb{E}_{\mathbf{a}_k}(\mathbf{a}_k). \quad (6b)$$

Here, $\mathbb{E}(\cdot)$ is the expectation with respect to the human correction data \mathbf{a}_k . Equation (6) says that the original assumption (4) can be met in expectation; in other words, the robot can receive and average human's *multiple* directional corrections before one update of the cost function is made. There are two




Fig. 3. Magnitude corrections vs. directional corrections. Contours (circles) and desired trajectory $\mathbf{u}_{0:T}^{\theta^*}$ (black dot) of the true cost function $J(\theta^*)$ are plotted. (a) Allowable region of magnitude corrections; green region (a sub-level set) indicates all allowable magnitude corrections $\bar{\mathbf{a}}_k$ that satisfy $J(\mathbf{u}_{0:T}^{\theta_k} + \bar{\mathbf{a}}_k, \theta^*) < J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$. (b) Allowable region of directional corrections; green region (half of the input space) indicates all allowable directional corrections $\bar{\mathbf{a}}_k$ that satisfy $\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \bar{\mathbf{a}}_k \rangle > 0$.

ways to implement multiple corrections. The most convenient way is that a human user applies multiple directional corrections at different time steps but all within the same execution of the robot motion ξ_{θ_k} , i.e., the robot executes the current plan ξ_{θ_k} just once. The second way is that the multiple corrections are applied in the multiple repetitions of the execution of ξ_{θ_k} , i.e., the robot executes the same plan ξ_{θ_k} for multiple times. Whichever implementation, by taking the expectation of human's multiple directional corrections $\bar{\mathbf{a}}_k = \mathbb{E}_{\mathbf{a}_k}(\mathbf{a}_k)$, (6) permits noisy and imperfect human directional corrections. The similar strategies have also been adopted in [6].

The **problem of interest** in this article is that given human (averaged) directional corrections $\bar{\mathbf{a}}_k$ satisfying the assumption (6), we aim to develop a rule to update the robot weight vector guess θ_k such that θ_k converges to θ^* as $k = 1, 2, 3, \dots$

Remark 1: We assume that human corrections $\mathbf{a}_{t_k} \in \mathbb{R}^m$ are in the robot input space, which means that \mathbf{a}_{t_k} can be directly added to robot input \mathbf{u}_{t_k} . This can be readily fulfilled in some applications such as autonomous vehicles, where a user directly changes the steering angle of a vehicle. For other cases where the corrections are not readily in the robot input space, it could be met via some specific interfaces (or computation), which map correction signals to the robot input space. For example, when a human interacts with the end effector of a robot manipulator, one can convert the task-space contact force to the joint torques via Jacobian. Then, \mathbf{a}_{t_k} is the mapped correction. The reason why we do not consider the corrections in the state space is that 1) corrections in input spaces may be easier to implement, as shown in our later experiments, and 2) corrections in state spaces can be infeasible for the underactuated robot systems [24].

Remark 2: Assumption (4) on directional correction \mathbf{a}_{t_k} is less restrictive than ones in [6]–[9] using magnitude corrections, which assume the cost of the corrected robot trajectory $\mathbf{u}_{0:T}^{\theta_k} + \bar{\mathbf{a}}_k$ is lower than that of the robot original motion $\mathbf{u}_{0:T}^{\theta_k}$, i.e., $J(\mathbf{u}_{0:T}^{\theta_k} + \bar{\mathbf{a}}_k, \theta^*) < J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$. As shown in Fig. 3, their assumptions usually lead to the restriction of correction magnitude. Specifically, to satisfy $J(\mathbf{u}_{0:T}^{\theta_k} + \bar{\mathbf{a}}_k, \theta^*) < J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$, $\|\bar{\mathbf{a}}_k\|$ has to be chosen from the $J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$ -sublevel set of $J(\theta^*)$, as shown by the green region in Fig. 3(a). Furthermore, this region will shrink as $\mathbf{u}_{0:T}^{\theta_k}$ gets close to the desired trajectory $\mathbf{u}_{0:T}^{\theta^*}$ (black dot), thus making $\|\bar{\mathbf{a}}_k\|$ more difficult to choose.

In contrast, the directional corrections satisfying (4) always account for half of the input space, as shown by the green region in Fig. 3(b). A human user can choose any correction as long as its direction lies in the half space of the gradient descent of $J(\theta^*)$. Thus, (4) would be more likely to be satisfied for nonexpert users. We will experimentally show this advantage later in Sections V-C and VI.

III. MAIN ALGORITHM OUTLINE AND ITS GEOMETRIC INTERPRETATION

In this section, we will outline the proposed main algorithm and present its geometric interpretation.

A. Hyperplane for Each Directional Correction

Before developing the main algorithm, we first show that condition (6) equals a linear inequality imposed on the true θ^* . This has been formally stated by the following lemma.

Lemma 1: Suppose that the current guess of the weight vector is θ_k , and the robot trajectory $\xi_{\theta_k} = \{x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}\}$ is a result of (locally) minimizing $J(\theta_k)$ in (2) subject to dynamics (1). For ξ_{θ_k} , given a human (averaged) directional correction \bar{a}_{t_k} satisfying (6), one has

$$\langle h_k, \theta^* \rangle + b_k < 0, \quad k = 1, 2, 3 \dots \quad (7)$$

with

$$h_k = \mathbf{H}_1^\top(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}) \bar{a}_k \in \mathbb{R}^r, \quad (8a)$$

$$b_k = \bar{a}_k^\top \mathbf{H}_2(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}) \nabla h(x_{T+1}^{\theta_k}) \in \mathbb{R}. \quad (8b)$$

Here, \bar{a}_k is defined in (6b); $\nabla h(x_{T+1}^{\theta_k})$ is the gradient of the final cost $h(x_{T+1})$ in (2) evaluated at $x_{T+1}^{\theta_k}$; $\mathbf{H}_1(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k})$ and $\mathbf{H}_2(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k})$ are computed as

$$\mathbf{H}_1(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}) = \begin{bmatrix} \mathbf{F}_u \mathbf{F}_x^{-1} \Phi_x + \Phi_u \\ \frac{\partial \phi^\top}{\partial u_{T-1}^{\theta_k}} \end{bmatrix} \in \mathbb{R}^{m(T+1) \times r}, \quad (9a)$$

$$\mathbf{H}_2(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}) = \begin{bmatrix} \mathbf{F}_u \mathbf{F}_x^{-1} \mathbf{V} \\ \frac{\partial f^\top}{\partial u_T^{\theta_k}} \end{bmatrix} \in \mathbb{R}^{m(T+1) \times n} \quad (9b)$$

with

$$\mathbf{F}_x = \begin{bmatrix} I & -\frac{\partial f^\top}{\partial x_1^{\theta_k}} & \dots & 0 & 0 \\ 0 & I & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I & -\frac{\partial f^\top}{\partial x_{T-1}^{\theta_k}} \\ 0 & 0 & \dots & & I \end{bmatrix}, \quad \Phi_x = \begin{bmatrix} \frac{\partial \phi^\top}{\partial x_1^{\theta_k}} \\ \frac{\partial \phi^\top}{\partial x_2^{\theta_k}} \\ \vdots \\ \frac{\partial \phi^\top}{\partial x_T^{\theta_k}} \end{bmatrix}, \quad (10a)$$

$$\mathbf{F}_u = \begin{bmatrix} \frac{\partial f^\top}{\partial u_0^{\theta_k}} & 0 & \dots & 0 \\ 0 & \frac{\partial f^\top}{\partial u_1^{\theta_k}} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{\partial f^\top}{\partial u_{T-1}^{\theta_k}} \end{bmatrix}, \quad \Phi_u = \begin{bmatrix} \frac{\partial \phi^\top}{\partial u_0^{\theta_k}} \\ \frac{\partial \phi^\top}{\partial u_1^{\theta_k}} \\ \vdots \\ \frac{\partial \phi^\top}{\partial u_{T-1}^{\theta_k}} \end{bmatrix}, \quad (10b)$$

$$\mathbf{V} = \begin{bmatrix} 0 & 0 & \dots & 0 & \frac{\partial f^\top}{\partial x_T^{\theta_k}} \end{bmatrix}^\top. \quad (10c)$$

In above, the dimensions are $\mathbf{F}_x \in \mathbb{R}^{nT \times nT}$, $\mathbf{F}_u \in \mathbb{R}^{mT \times nT}$, $\Phi_x \in \mathbb{R}^{nT \times r}$, $\Phi_u \in \mathbb{R}^{mT \times r}$, and $\mathbf{V} \in \mathbb{R}^{nT \times n}$. I is $n \times n$ identity. For a general differentiable function $g(x)$ and a fixed value x^* , we denote $\frac{\partial g}{\partial x^*}$ as the Jacobian of $g(x)$ evaluated at x^* .

Proof: We first derive the explicit form of $\nabla J(u_{0:T}^{\theta_k}, \theta^*)$ in (6), and then show that (6) can be rewritten as (7).

Consider the robot current trajectory $\xi_{\theta_k} = \{x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}\}$, which satisfies the robot dynamics constraint (1). For any $t = 0, 1, \dots, T$, define the infinitesimal perturbation pair $(\delta x_t, \delta u_t)$ at the state-input pair $(x_t^{\theta_k}, u_t^{\theta_k})$. By linearizing the dynamics (1) around $(x_t^{\theta_k}, u_t^{\theta_k})$, we have

$$\delta x_{t+1} = \frac{\partial f}{\partial x_t^{\theta_k}} \delta x_t + \frac{\partial f}{\partial u_t^{\theta_k}} \delta u_t \quad (11)$$

where $\frac{\partial f}{\partial x_t^{\theta_k}}$ and $\frac{\partial f}{\partial u_t^{\theta_k}}$ are the Jacobians of f with respect to x_t and u_t , respectively, evaluated at $(x_t^{\theta_k}, u_t^{\theta_k})$. By stacking (11) for all $t = 0, 1, \dots, T$ and also noting $\delta x_0 = 0$ (because $x_0^{\theta_k} = x_0$ is given), we have the following compact form:

$$-\mathbf{A}^\top \delta x_{1:T+1} + \mathbf{B}^\top \delta u_{0:T} = 0 \quad (12)$$

with $\delta x_{1:T+1} = [\delta x_1^\top, \dots, \delta x_{T+1}^\top]^\top$ and $\delta u_{1:T} = [\delta u_1^\top, \dots, \delta u_T^\top]^\top$

$$\mathbf{A} = \begin{bmatrix} \mathbf{F}_x & -\mathbf{V} \\ 0 & I \end{bmatrix}, \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} \mathbf{F}_u & 0 \\ 0 & \frac{\partial f^\top}{\partial u_T^{\theta_k}} \end{bmatrix} \quad (13)$$

with \mathbf{F}_x and \mathbf{F}_u defined in (10). Given $\delta x_{1:T+1}$ and $\delta u_{1:T}$, the increment of $J(u_{0:T}^{\theta_k}, \theta^*)$, denoted as $\delta J(\theta^*)$, is

$$\delta J(\theta^*) = \mathbf{C}^\top \delta x_{1:T+1} + \mathbf{D}^\top \delta u_{0:T} \quad (14)$$

with

$$\mathbf{C} = \begin{bmatrix} \Phi_x \theta^* \\ \frac{\partial h^\top}{\partial x_{T+1}^{\theta_k}} \end{bmatrix} \quad \text{and} \quad \mathbf{D} = \begin{bmatrix} \Phi_u \theta^* \\ \frac{\partial \phi^\top}{\partial u_T^{\theta_k}} \theta^* \end{bmatrix} \quad (15)$$

with Φ_x and Φ_u defined in (10). Since \mathbf{A} is always invertible (due to full rank), we solve for $\delta x_{1:T+1}$ from (12) and submit it to (14), yielding

$$\begin{aligned} \delta J(\theta^*) &= \mathbf{C}^\top \delta x_{1:T+1} + \mathbf{D}^\top \delta u_{0:T}, \\ &= (\mathbf{C}^\top (\mathbf{A}^{-1})^\top \mathbf{B}^\top + \mathbf{D}^\top) \delta u_{0:T}. \end{aligned} \quad (16)$$

Thus, we have

$$\nabla J(u_{0:T}^{\theta_k}, \theta^*) = \mathbf{B} \mathbf{A}^{-1} \mathbf{C} + \mathbf{D}. \quad (17)$$

Equation (17) can be further written as

$$\begin{aligned} \nabla J(\mathbf{u}_{0:T}^{\theta_k}, \boldsymbol{\theta}^*) &= \mathbf{B}\mathbf{A}^{-1}\mathbf{C} + \mathbf{D} \\ &= \begin{bmatrix} \mathbf{F}_u & \mathbf{0} \\ \mathbf{0} & \frac{\partial \mathbf{f}^\top}{\partial \mathbf{u}_{T^*}^{\theta_k}} \end{bmatrix} \begin{bmatrix} \mathbf{F}_x & -\mathbf{V} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}^{-1} \begin{bmatrix} \Phi_x \boldsymbol{\theta}^* \\ \frac{\partial h^\top}{\partial \mathbf{x}_{T+1}^{\theta_k}} \end{bmatrix} + \begin{bmatrix} \Phi_x \boldsymbol{\theta}^* \\ \frac{\partial \phi^\top}{\partial \mathbf{u}_{T^*}^{\theta_k}} \boldsymbol{\theta}^* \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{F}_u \mathbf{F}_x^{-1} \Phi_x \boldsymbol{\theta}^* + \Phi_x \boldsymbol{\theta}^* + \mathbf{F}_u \mathbf{F}_x^{-1} \mathbf{V} \frac{\partial h^\top}{\partial \mathbf{x}_{T+1}^{\theta_k}} \\ \frac{\partial \phi^\top}{\partial \mathbf{u}_{T^*}^{\theta_k}} \boldsymbol{\theta}^* + \frac{\partial \mathbf{f}^\top}{\partial \mathbf{u}_{T^*}^{\theta_k}} \frac{\partial h^\top}{\partial \mathbf{x}_{T+1}^{\theta_k}} \end{bmatrix} \quad (18) \end{aligned}$$

where we have used Schur complement to compute the inverse of the block matrix \mathbf{A} . By the definitions in (9), (18) becomes

$$\begin{aligned} \nabla J(\mathbf{u}_{0:T}^{\theta_k}, \boldsymbol{\theta}^*) &= \mathbf{H}_1(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k}) \boldsymbol{\theta}^* \\ &\quad + \mathbf{H}_2(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k}) \nabla h(\mathbf{x}_{T+1}^{\theta_k}). \quad (19) \end{aligned}$$

Substituting (19) into (6) and also considering the definitions in (8), we obtain

$$\left\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \boldsymbol{\theta}^*), \bar{\mathbf{a}}_k \right\rangle = -\langle \mathbf{h}_k, \boldsymbol{\theta}^* \rangle - b_k > 0 \quad (20)$$

which leads to (7). This completes the proof. \blacksquare

We have the following remarks on the intuition of Lemma 1.

Remark 3: The key result of Lemma 1 is the linear inequality (7), which is associated with a hyperplane $\langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k = 0$. The parameters \mathbf{h}_k and b_k of this hyperplane, given in (8), are both known and determined by human directional correction $\bar{\mathbf{a}}_k$ and robot current trajectory $\boldsymbol{\xi}_{\theta_k}$. In other words, Lemma 1 states that given robot motion $\boldsymbol{\xi}_{\theta_k}$ and the human directional correction $\bar{\mathbf{a}}_k$, one can write a hyperplane $\langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k = 0$ and a linear inequality (7) for the unknown $\boldsymbol{\theta}^*$.

Remark 4: Obtaining \mathbf{h}_k and b_k requires computing the matrices $\mathbf{H}_1(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$ and $\mathbf{H}_2(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$ in (9). In our previous work [5], $\mathbf{H}_1(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$ and $\mathbf{H}_2(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$ are called the ‘‘recovery matrix.’’ The work [5, Lemma 1] has shown that the ‘‘recovery matrix’’ can be iteratively computed by integrating each point $(\mathbf{x}_t^{\theta_k}, \mathbf{u}_t^{\theta_k})$, $t = 0, 1, \dots, T$. This iterative property of the recovery matrix facilitates the computation of \mathbf{H}_1 and \mathbf{H}_2 by avoiding the inverse of \mathbf{F}_x in (9), significantly reducing the computational burden. Please refer to [5, Lemma 1] for the iteration formula.

B. Outline of the Proposed Algorithm

With Lemma 1, we are ready to develop the main algorithm to learn $\boldsymbol{\theta}^*$. At each iteration k , we maintain a *weight search space*, denoted by $\Omega_k \subset \mathbb{R}^r$, such that $\boldsymbol{\theta}^* \in \Omega_k$ and $\boldsymbol{\theta}_k \in \Omega_k$ for all $k = 1, 2, 3, \dots$. This Ω_k can be thought of as the possible location of $\boldsymbol{\theta}^*$, and $\boldsymbol{\theta}_k$ as the current guess to $\boldsymbol{\theta}^*$. Rather than a direct rule to guide $\boldsymbol{\theta}_k$ toward $\boldsymbol{\theta}^*$, we will develop a rule to update Ω_k to Ω_{k+1} such that a useful measure of the size of Ω_k will converge to 0. With this high-level idea, we outline the proposed main algorithm below.

Main Algorithm (Outline) Initialize the weight search space Ω_0 to be

$$\Omega_0 = \{\boldsymbol{\theta} \in \mathbb{R}^r \mid -\underline{c}_i \leq \boldsymbol{\theta}[i] \leq \bar{c}_i, i = 1, 2, \dots, r\}, \quad (21)$$

where constants $\underline{c}_i \geq 0$ and $\bar{c}_i \geq 0$ denote the lower bound and upper bounds of the i th entry $\boldsymbol{\theta}[i]$ in $\boldsymbol{\theta}$, respectively. Here, \underline{c}_i and \bar{c}_i can be chosen large enough such that $\boldsymbol{\theta}^* \in \Omega_0$. At iteration $k = 1, 2, \dots$, the learning proceeds with the following three steps.

Step 1: Choose a weight vector guess $\boldsymbol{\theta}_k$ from the current weight search space Ω_{k-1} , i.e., $\boldsymbol{\theta}_k \in \Omega_{k-1}$ (we will discuss how to choose $\boldsymbol{\theta}_k \in \Omega_{k-1}$ in Section IV).

Step 2: The robot restarts and plans its motion trajectory $\boldsymbol{\xi}_{\theta_k}$ by solving a trajectory optimization with the cost function $J(\boldsymbol{\theta}_k)$ in (2) and the dynamics in (1). While the robot is executing the plan $\boldsymbol{\xi}_{\theta_k}$, a human applies (averaged) directional correction $\bar{\mathbf{a}}_k$ in (6b). Then, a hyperplane $\langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k = 0$ is computed in (7) and (8).

Step 3: Update the weight search space Ω_{k-1} to Ω_k via

$$\Omega_k = \Omega_{k-1} \cap \{\boldsymbol{\theta} \in \mathbb{R}^r \mid \langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k < 0\}. \quad (22)$$

We provide some remarks to the above outline of the proposed algorithm. For the initialization (21), we allow different entries of $\boldsymbol{\theta}$ to have their own lower and upper bounds, which may come from prior knowledge. Simply but not necessarily, one could also initialize

$$\Omega_0 = \{\boldsymbol{\theta} \in \mathbb{R}^r \mid \|\boldsymbol{\theta}\|_\infty \leq R\} \quad (23)$$

where

$$R = \max\{\underline{c}_i, \bar{c}_i, i = 1, \dots, r\}. \quad (24)$$

In Step 1, a weight vector guess $\boldsymbol{\theta}_k$ is chosen from Ω_{k-1} . We will [in (26) in Lemma 2] show that the true $\boldsymbol{\theta}^*$ *always* lies in Ω_k for any $k = 1, 2, 3, \dots$. Thus, we will expect $\boldsymbol{\theta}_k$ to be closer to $\boldsymbol{\theta}^*$ if a size measure of Ω_k gets smaller. In fact, the weight search space Ω_k is *always nonincreasing* as $\Omega_k \subseteq \Omega_{k-1}$ due to (22) in Step 3. In Section IV, we will further focus on how to pick $\boldsymbol{\theta}_k \in \Omega_{k-1}$ such that the size of Ω_k is *strictly reduced* as k increases. In Step 2, the robot trajectory $\boldsymbol{\xi}_{\theta_k}$ is planned by solving a trajectory optimization with the cost function $J(\boldsymbol{\theta}_k)$ in (2) and the dynamics constraint (1). This can be done by any available trajectory optimization methods, e.g., [25] and existing solvers [26]. After a human applies a directional correction $\bar{\mathbf{a}}_{t_k}$ to the executed $\boldsymbol{\xi}_{\theta_k}$, a hyperplane $\langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k = 0$ can be computed via (7) and (8), which will be used to update the search space by (22) in Step 3. The detailed implementation of the above main algorithm will be described in the next section.

Given the above proposed main algorithm, we next present the following important lemma.

Lemma 2: Under the proposed main algorithm, one has

$$\langle \mathbf{h}_k, \boldsymbol{\theta}_k \rangle + b_k = 0, \quad \forall k = 1, 2, 3, \dots \quad (25)$$

and

$$\boldsymbol{\theta}^* \in \Omega_k, \quad \forall k = 1, 2, 3, \dots \quad (26)$$

A proof of Lemma 2 is given in Appendix A. Lemma 2 has the following intuitive geometric explanations. First, (25) says that the current guess $\boldsymbol{\theta}_k$ is always located on the current hyperplane $\langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k = 0$. Second, (26) says that although the proposed




Fig. 4. Illustration of updating Ω_k . (a) At k th iteration. (b) At $(k+1)$ th iteration.

algorithm directly updates the weight search space Ω_k via (22), the true (but unknown) weight vector θ^* is always contained in the current Ω_k . Thus, intuitively, the smaller the search space Ω_k is, the closer θ^* is to θ_k .

C. Geometric Interpretation

In this section, we will provide an interpretation of the proposed main algorithm through a geometric perspective. For simplicity of illustrations, we assume $\theta \in \mathbb{R}^2$ here.

At the k th iteration shown in Fig. 4(a), by Step 1 of the main algorithm, a weight vector guess θ_k (red dot) is picked from the current search space Ω_{k-1} (light blue region), i.e., $\theta_k \in \Omega_{k-1}$. By Step 2, we obtain a hyperplane $\langle \mathbf{h}_k, \theta \rangle + b_k = 0$ (black dashed line), which cuts through the weight search space Ω_{k-1} into two portions. By (25) of Lemma 2, we know that θ_k always lies on this hyperplane because $\langle \mathbf{h}_k, \theta_k \rangle + b_k = 0$, as shown in Fig. 4(a). By Step 3 of the main algorithm, we only keep one of the two cut portions, which is the intersection between Ω_{k-1} and the half space $\langle \mathbf{h}_k, \theta \rangle + b_k < 0$, and the kept portion will be used as the new search space for the next iteration, that is, $\Omega_k = \Omega_{k-1} \cap \{\theta \in \mathbb{R}^r \mid \langle \mathbf{h}_k, \theta \rangle + b_k < 0\}$, as shown by the blue region in Fig. 4(a). The above procedure repeats at the next iteration $k+1$ in Fig. 4(a) and finally produces a smaller search space Ω_{k+1} , as shown by the darkest blue region in Fig. 4(b). From (22), one has $\Omega_0 \supseteq \dots \supseteq \Omega_{k-1} \supseteq \Omega_k \supseteq \Omega_{k+1} \supseteq \dots$. Moreover, by (26) in Lemma 2, we note that the true θ^* (red point) is always contained in Ω_k whenever k is.

Besides the above geometric illustration, we also have the following observations.

- 1) The key idea of the proposed main algorithm is to cut and remove the weight search space Ω_{k-1} as each human directional correction $\bar{\mathbf{a}}_k$ becomes available. Thus, we always expect that Ω_{k-1} can quickly shrink as k increases because, thereby, we can say that the guess θ_k is close to the true weight vector θ^* . As shown in Fig. 4, the reduction rate of Ω_{k-1} depends on two factors: the human directional correction $\bar{\mathbf{a}}_k$, and how to choose $\theta_k \in \Omega_{k-1}$, both discussed below.
- 2) From (8), we note that human directional correction $\bar{\mathbf{a}}_k$ determines \mathbf{h}_k , which is the normal vector of the hyperplane $\langle \mathbf{h}_k, \theta \rangle + b_k = 0$. When fixing θ_k , we can think of the hyperplane rotates around θ_k with different $\bar{\mathbf{a}}_k$, leading to different removals of Ω_{k-1} . This can be illustrated by comparing Fig. 5(a) and (b).




Fig. 5. Illustration of how different human directional corrections $\bar{\mathbf{a}}_k$ affect the cut of the weight search space Ω_{k-1} . (a) Large cut from Ω_{k-1} . (b) Small cut from Ω_{k-1} .

- 3) The choice of the guess $\theta_k \in \Omega_{k-1}$ defines the specific location of the hyperplane $\langle \mathbf{h}_k, \theta \rangle + b_k = 0$ because the hyperplane is always passing through θ_k by (25) in Lemma 2. Thus, θ_k also affects how Ω_{k-1} is cut and removed. This can be illustrated by comparing Fig. 4(a) with Fig. 5(a).

Based on the above observations, we know that the convergence of the proposed main algorithm is determined by the reduction of the weight search space Ω_{k-1} . This reduction depends on both the human directional correction $\bar{\mathbf{a}}_k$ and the choice of the weight vector guess $\theta_k \in \Omega_{k-1}$. In the next section, we will present a way of choosing θ_k to guarantee the convergence of the proposed main algorithm.

IV. ALGORITHM IMPLEMENTATION AND CONVERGENCE ANALYSIS

In this section, we detail the choice of $\theta_k \in \Omega_{k-1}$, show the convergence of the proposed algorithm, and present a detailed implementation of the algorithm with a termination criterion.

A. Choice of Current Guess $\theta_k \in \Omega_{k-1}$

In the proposed main algorithm, at iteration k , the weight search space Ω_{k-1} is updated using (22), rewritten as follows:

$$\Omega_k = \Omega_{k-1} \cap \{\theta \in \mathbb{R}^r \mid \langle \mathbf{h}_k, \theta \rangle + b_k < 0\}.$$

To show the reduction of a weight search space, it is straightforward to use the volume measure (length or area for 1-D or 2-D cases, respectively) of a (closure) search space Ω_k , denoted as $\text{Vol}(\Omega_k)$. Zero volume implies the convergence of the search space [27]. By (22), we know that $\Omega_k \subseteq \Omega_{k-1}$ and, thus, $\text{Vol}(\Omega_k)$ is nonincreasing. In the following, we need to further find a way such that $\text{Vol}(\Omega_k)$ is strictly decreasing, i.e., there exists a constant $0 \leq \alpha < 1$ such that

$$\text{Vol}(\Omega_k) \leq \alpha \text{Vol}(\Omega_{k-1}). \quad (27)$$

From Fig. 5, we observe that for a fixed choice of $\theta_k \in \Omega_{k-1}$, different human directional corrections $\bar{\mathbf{a}}_k$ can lead to different reduction of Ω_{k-1} . For example, Fig. 5(a) has a large volume reduction from Ω_{k-1} to Ω_k , while Fig. 5(b) leads to a very small reduction. Unfortunately, we cannot assume the specific direction of a human correction $\bar{\mathbf{a}}_k$ (it is the discretion of the user), but we can choose the current guess θ_k “smartly” to avoid the very small reduction regardless of specific human corrections. One intuitive choice of such θ_k is at some *center* of




Fig. 6. Illustration of choosing θ_{k+1} as the center of MVE E_k inscribed in the weight search space Ω_k . (a) Center of MVE in Ω_{k-1} . (b) Center of MVE in Ω_k .

the search space Ω_{k-1} . Thus, we define the following center of the maximum volume ellipsoid (MVE) inscribed in Ω_{k-1} .

Definition 1 (Maximum Volume Inscribed Ellipsoid [28]): Given a compact convex set $\Omega \subset \mathbb{R}^r$, the MVE inscribed in Ω , defined as E , is denoted by

$$E = \{\bar{B}\theta + \bar{d} \mid \|\theta\|_2 \leq 1\}. \quad (28)$$

Here, $\bar{B} \in \mathbb{S}_{++}^r$ (i.e., a $r \times r$ positive definite matrix) and $\bar{d} \in \mathbb{R}^r$ is called the center of E . \bar{B} and \bar{d} solve the optimization

$$\begin{aligned} \max_{\bar{d}, B \in \mathbb{S}_{++}^r} \quad & \log \det B \\ \text{s.t.} \quad & \sup_{\|\theta\|_2 \leq 1} I_\Omega(B\theta + \bar{d}) \leq 0 \end{aligned} \quad (29)$$

where $I_\Omega(\theta) = 0$ for $\theta \in \Omega$ and $I_\Omega(\theta) = \infty$ for $\theta \notin \Omega$.

By Definition 1, let E_k denote the MVE inscribed in Ω_k and d_k denote the center of E_k , $k = 0, 1, \dots$. For the choice of $\theta_{k+1} \in \Omega_k$, we choose

$$\theta_{k+1} = d_k \quad (30)$$

as illustrated in Fig. 6. Other choices of θ_{k+1} using other types of centers of the search space are discussed in Appendix C.

We now give a computational method to solve for d_k , i.e., the center of MVE E_k inscribed in Ω_k . Recall that in the main algorithm, the initial Ω_0 is (21), which is equivalent to

$$\Omega_0 = \left\{ \theta \mid \begin{array}{l} \langle e_i, \theta \rangle - \bar{c}_i \leq 0 \\ -\langle e_i, \theta \rangle - \underline{c}_i \leq 0, \quad i = 1, \dots, r \end{array} \right\} \quad (31)$$

where e_i is a unit vector with the only i th entry as 1. Following the update (22), Ω_k is also a compact polytope, which is

$$\Omega_k = \left\{ \theta \mid \begin{array}{l} \langle e_i, \theta \rangle - \bar{c}_i \leq 0, \quad i = 1, \dots, r; \\ -\langle e_i, \theta \rangle - \underline{c}_i \leq 0, \quad i = 1, \dots, r; \\ \langle h_j, \theta \rangle + b_j < 0, \quad j = 1, \dots, k \end{array} \right\}. \quad (32)$$

As a result, solving (29) for the center d_k becomes a convex program [28], stated in the following lemma.

Lemma 3: For a polytope Ω_k in (32), the center d_k of the MVE E_k inscribed in Ω_k can be solved via the following convex optimization:

$$\begin{aligned} \min_{\bar{d}, B \in \mathbb{S}_{++}^r} \quad & -\log \det B \\ \text{s.t.} \quad & \|B e_i\|_2 + \langle \bar{d}, e_i \rangle \leq \bar{c}_i, \quad i = 1, \dots, r, \\ & \|B e_i\|_2 - \langle \bar{d}, e_i \rangle \leq \underline{c}_i, \quad i = 1, \dots, r, \\ & \|B h_j\|_2 + \langle \bar{d}, h_j \rangle \leq -b_j, \quad j = 1, \dots, k. \end{aligned} \quad (33)$$

The proof of Lemma 3 can be found in [28, pp. 414, Chapter 8.4.2]. Equation (33) can be efficiently solved by available convex program solvers, e.g., [29]. In the implementation, since the number of linear inequalities grows as k increases, the trick of dropping some redundant inequalities in (32) can be adopted [27]. Dropping redundant inequalities does not change Ω_k and its volume reduction (convergence). Please see how to identify the redundant inequalities in [27].

B. Exponential Convergence and Termination Criterion

Now we analyze the convergence of the volume reduction of Ω_k and develop its termination criterion in implementation. First, the convergence of the reduction of $\text{Vol}(\Omega_k)$ is guaranteed by the following lemma.

Lemma 4: Let θ_k be chosen as the center d_{k-1} of the MVE E_{k-1} inscribed in $\Omega_{k-1} \subset \mathbb{R}^r$. Then, the update (22) has

$$\frac{\text{Vol}(\Omega_k)}{\text{Vol}(\Omega_{k-1})} \leq \left(1 - \frac{1}{r}\right). \quad (34)$$

Lemma 4 is a theorem directly from [30]. Lemma 4 indicates

$$\text{Vol}(\Omega_k) \leq \left(1 - \frac{1}{r}\right)^k \text{Vol}(\Omega_0).$$

Thus, the convergence rate of $\text{Vol}(\Omega_k) \rightarrow 0$ is as fast as $(1 - \frac{1}{r})^k \rightarrow 0$, as $k = 0, 1, 2, \dots$

To implement the proposed main algorithm, we will not only need the exponential convergence as established by Lemma 4 but also a termination criterion, which specifies the maximum number of iterations needed for $\text{Vol}(\Omega_k)$ below a certain given threshold. Thus, we have the following theorem.

Theorem 1: In the main algorithm, suppose Ω_0 is given by (21), and at iteration k , θ_k is chosen as the center d_{k-1} of the MVE E_{k-1} inscribed in Ω_{k-1} . Given a termination condition

$$\text{Vol}(\Omega_k) \leq (2\epsilon)^r$$

with ϵ a user-specified threshold, the proposed main algorithm runs for $k \leq K$ iterations, namely, the algorithm terminates in at most K iterations, where

$$K = \frac{r \log(R/\epsilon)}{-\log(1 - 1/r)} \quad (35)$$

with R given in (24).

Proof: Initially, we have $\text{Vol}(\Omega_0) \leq (2R)^r$. From Lemma 4, after k iterations, we have

$$\text{Vol}(\Omega_k) \leq \left(1 - \frac{1}{r}\right)^k \text{Vol}(\Omega_0) \leq \left(1 - \frac{1}{r}\right)^k (2R)^r \quad (36)$$

which yields to

$$\log \text{Vol}(\Omega_k) \leq k \log \left(1 - \frac{1}{r}\right) + \log(2R)^r. \quad (37)$$

When $k = \frac{r \log(R/\epsilon)}{-\log(1 - 1/r)}$

$$\log \text{Vol}(\Omega_k) \leq -r \log(R/\epsilon) + \log(2R)^r. \quad (38)$$

The above equation is simplified to

$$\log \text{Vol}(\Omega_k) \leq \log(2\epsilon)^r \quad (39)$$

Algorithm 1: Learning from directional corrections

Input: Specify a termination threshold ϵ and use it to compute the maximum iteration K by (35).

Initialization: Initial weight search space Ω_0 in (21).

for $k = 1, 2, \dots, K$ **do**

Choose a weight vector guess $\theta_k \in \Omega_{k-1}$ via Lemma 3;

Restart and plan a robot trajectory ξ_{θ_k} by solving a trajectory optimization with the cost function $J(\theta_k)$ in (2) and the dynamics in (1);

Robot executes the planned trajectory ξ_{θ_k} while receiving the human averaged directional correction \bar{a}_k in (6b);

Compute the matrices $H_1(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$ and $H_2(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$, and then compute the hyperplane and half space $\langle \mathbf{h}_k, \theta \rangle + b_k < 0$ via (7) and (8);

Update the weight search space by $\Omega_k = \Omega_{k-1} \cap \{\theta \in \mathbb{R}^r \mid \langle \mathbf{h}_k, \theta \rangle + b_k < 0\}$ via (22);

end

Output: θ_K .

which means that the termination condition $\text{Vol}(\Omega_k) \leq (2\epsilon)^r$ is satisfied. This completes the proof. ■

We have the following comments on Theorem 1.

Remark 5: Since Lemma 2 states that both θ^* and θ_k are in Ω_k for any $k = 1, 2, 3, \dots$, the user-specified threshold ϵ in the termination condition $\text{Vol}(\Omega_k) \leq (2\epsilon)^r$ can also be understood as an indicator for the distance between the true θ^* (usually unknown in practice) and the robot current guess θ_k . ϵ is set based on the desired learning accuracy.¹

C. Detailed Implementation of Main Algorithm

With the termination criterion stated in Theorem 1 and the choice of θ_k in (30) and Lemma 3, we present the detailed implementation of the proposed main method in Algorithm 1. The setting of the initial Ω_0 in (21) and ϵ will be given in Appendix B-D.

V. NUMERICAL EXAMPLES

In this section, we will test the proposed method in numerical simulations and make comparisons with the magnitude-correction methods discussed in the previous related work.

A. Inverted Pendulum

The dynamics of a pendulum is

$$\ddot{\alpha} = \frac{-g}{l} \sin \alpha - \frac{d}{ml^2} \dot{\alpha} + \frac{u}{ml^2} \quad (40)$$

with α being the angle between the pendulum and the direction of gravity, u is the torque applied to the pivot, $l = 1$ m, $m = 1$ kg, and $d = 0.1$ Ns/m are the length, mass, and damping ratio of the pendulum, respectively. We discretize the dynamics

¹In practice, it is usually easy to set ϵ with a small value because the empirical results in Sections VI and VII show that the robot motion trajectory can converge (to the desired motion) more quickly than the objective function itself. This means that one usually sees the convergence of the robot trajectory before the termination condition is reached, as shown in Fig. 10. More discussion about setting ϵ will be given in Appendix B-D.

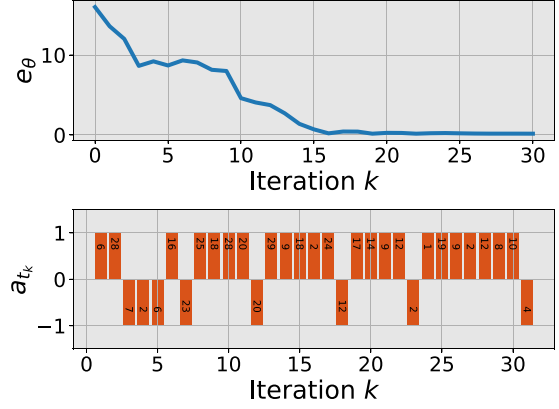


Fig. 7. Learning from simulated directional corrections for the pendulum. The upper panel shows the weight error $e_\theta = \|\theta - \theta^*\|^2$ versus iteration. The bottom shows the simulated directional correction \mathbf{a}_{t_k} (positive or negative sign) at each iteration k ; here, the value in each bar is t_k , which is randomly drawn from a uniform distribution over horizon $[0, 30]$.

using the Euler method with a time interval $\Delta = 0.2$ s. The state and control input of the pendulum are $\mathbf{x} = [\alpha, \dot{\alpha}]^\top$ and $\mathbf{u} = u$, respectively. The initial condition is $\mathbf{x}_0 = [0, 0]^\top$. In the cost function (2), we set the feature and weight vectors as

$$\begin{aligned} \phi &= [\alpha^2, \alpha, \dot{\alpha}^2, u^2]^\top \in \mathbb{R}^4, \\ \theta &= [\theta_1, \theta_2, \theta_3, \theta_4]^\top \in \mathbb{R}^4 \end{aligned} \quad (41)$$

respectively, and the final cost $h(\mathbf{x}_{T+1}) = 10(\alpha - \pi)^2 + 10\dot{\alpha}^2$, as our goal is to swing up the pendulum. The discrete time horizon is $T = 30$.

For numerical analysis, we “simulate” the human directional corrections instead of using real human corrections. The simulated corrections are generated as follows. Suppose that the true weight vector is known: $\theta^* = [0.5, 0.5, 0.5, 0.5]^\top$. At iteration k , a simulated correction \mathbf{a}_{t_k} is generated using the *sign* of the gradient of the true cost function $J(\theta^*)$, i.e.,

$$\mathbf{a}_{t_k} = -\text{sign} \left(\nabla J(\mathbf{u}_{0:T}^{\theta^*}, \theta^*) [t_k] \right) \in \mathbb{R}. \quad (42)$$

Here, $\nabla J(\mathbf{u}_{0:T}^{\theta^*}, \theta^*) [t_k]$ denotes the t_k th entry of ∇J , and the correction time t_k is randomly drawn from a uniform distribution over horizon $[0, T]$. Obviously, the above simulated directional corrections satisfy assumption (4).

The initial weight search space is set as $\Omega_0 = \{\theta \mid 0 \leq \theta[i] \leq 5, i = 1, 2, 3, 4\}$. In Algorithm 1, we set the termination threshold $\epsilon = 10^{-1}$, and the maximum learning iteration solved via (35) is $K = 55$. We apply Algorithm 1 to learn the true θ^* . To illustrate results, we define the weight error $e_\theta = \|\theta - \theta^*\|^2$ and plot e_θ versus iteration k in the top panel of Fig. 7. In the bottom panel of Fig. 7, we plot the simulated directional correction \mathbf{a}_{t_k} at each iteration k , where $+1$ and -1 bar denote the positive and negative signs (i.e., direction) of correction \mathbf{a}_{t_k} in (42), respectively; the number inside the bar denotes the correction time t_k randomly drawn from $[0, T]$.

From the results in Fig. 7, we can see that as the learning iteration k increases, the weight vector guess θ_k converges to




Fig. 8. Learning from simulated directional corrections for the robot arm. The plot is $e_\theta = \|\theta - \theta^*\|^2$ versus iteration. The corrections are given in Fig. 9.

the true $\theta^* = [0.5, 0.5, 0.5, 0.5]^T$. This shows the efficacy of the method, as guaranteed by Theorem 1.

B. Two-Link Robot Arm System

Here, we test the proposed method on a two-link robot arm. The dynamics of the robot arm system (moving horizontally) is $M(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{c}(\mathbf{q}, \dot{\mathbf{q}}) = \boldsymbol{\tau}$, where $M(\mathbf{q})$ is the inertia matrix and $\mathbf{c}(\mathbf{q}, \dot{\mathbf{q}})$ is the Coriolis and centrifugal term; $\mathbf{q} = [q_1, q_2]^T$ is the vector of joint angles, and $\boldsymbol{\tau} = [\tau_1, \tau_2]^T$ is the vector of joint torques. The state and control input are $\mathbf{x} = [\mathbf{q}, \dot{\mathbf{q}}]^T \in \mathbb{R}^4$ and $\mathbf{u} = \boldsymbol{\tau} \in \mathbb{R}^2$, respectively. The initial condition of the robot arm is set as $\mathbf{x}_0 = [0, 0, 0, 0]^T$. All physical parameters in the dynamics are units. We discretize the dynamics using the Euler method with a time interval $\Delta = 0.2$ s. In the cost function (2), we set the feature and weight vectors as

$$\phi = [q_1^2, q_1, q_2^2, q_2, \|\mathbf{u}\|^2]^T \in \mathbb{R}^5, \quad (43a)$$

$$\theta = [\theta_1, \theta_2, \theta_3, \theta_4, \theta_5]^T \in \mathbb{R}^5 \quad (43b)$$

respectively, and the final cost $h(\mathbf{x}_{T+1}) = 100((q_1 - \frac{\pi}{2})^2 + q_2^2 + q_1^2 + q_2^2)$, as the robot arm aims to reach and stop at the pose of $\mathbf{q} = [\frac{\pi}{2}, 0]^T$. The discrete-time horizon is $T = 50$.

We still “simulate” directional corrections. Suppose that we know the true $\theta^* = [1, 1, 1, 1, 1]^T$. At each iteration k , the simulation generates a directional correction \mathbf{a}_{t_k} using the sign of gradient of the true $J(\theta^*)$

$$\mathbf{a}_{t_k} = -\text{sign}\left(\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)[2t_k : 2t_k + 1]\right) \in \mathbb{R}^2 \quad (44)$$

where $\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)[2t_k : 2t_k + 1]$ are the entries at the locations from $2t_k$ to $2t_k + 1$ in ∇J (because $\mathbf{u} \in \mathbb{R}^2$), and the correction time t_k is randomly drawn from (uniform distribution) time horizon $[0, T]$. sign is applied entrywise.

The initial weight search space is $\Omega_0 = \{\theta \mid 0 \leq \theta[i] \leq 4, i = 1, 2, \dots, 5\}$. We set the termination threshold $\epsilon = 10^{-1}$, and the maximum learning iteration is $K = 83$ by (35). We apply Algorithm 1 to learn the true θ^* . We define the weight error $e_\theta = \|\theta - \theta^*\|^2$ and plot e_θ versus iteration in Fig. 8. We also plot $\mathbf{a}_{t_k} = [a_{t_k,1}, a_{t_k,2}]^T$ at each iteration k in Fig. 9, where the value labeled on each bar marks the correction time t_k . The results in Fig. 8 show that as the iteration increases, θ_k converges to the true θ^* , which again confirms Theorem 1.




Fig. 9. Simulated directional correction $\mathbf{a}_{t_k} = [a_{t_k,1}, a_{t_k,2}]^T$ at each iteration k for the robot arm. The number inside the bar is the correction time t_k randomly drawn from range $[0, 50]$.

C. Comparison With Magnitude-Correction Methods

We compare the proposed method with two related works [6], [7], both of which learn an objective function from magnitude corrections. The comparison is conducted on the inverted pendulum used in Section V-A. According to [7], given a magnitude correction, the trajectory deformation [8] is used to obtain a *human intended trajectory*. Specifically, suppose the robot trajectory is $\xi_{\theta_k} = \{\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k}\}$. Given a correction $\bar{\mathbf{a}}_k$, the human intended trajectory, denoted as $\bar{\xi}_{\theta_k} = \{\bar{\mathbf{x}}_{0:T+1}^{\theta_k}, \bar{\mathbf{u}}_{0:T}^{\theta_k}\}$, is solved as

$$\bar{\mathbf{u}}_{0:T}^{\theta_k} = \mathbf{u}_{0:T}^{\theta_k} + M^{-1}\bar{\mathbf{a}}_k \quad (45)$$

where M is a matrix to propagate a single time step correction $\bar{\mathbf{a}}_k$ (5) along the rest of the robot trajectory [22], and $\bar{\mathbf{x}}_{0:T+1}^{\theta_k}$ in $\bar{\xi}_{\theta_k}$ is obtained via rolling out the robot dynamics (1) given $\bar{\mathbf{u}}_{0:T}^{\theta_k}$. The learning update used in both [6] and [7] is

$$\theta_{k+1} = \theta_k + \eta (\phi(\bar{\xi}_{\theta_k}) - \phi(\xi_{\theta_k})) \quad (46)$$

where $\phi(\bar{\xi}_{\theta_k})$ and $\phi(\xi_{\theta_k})$ are the feature vectors of the human intended trajectory $\bar{\xi}_{\theta_k}$ and the uncorrected robot trajectory ξ_{θ_k} , respectively. Here, we set M as the finite differencing matrix [8] and $\eta = 0.0002$ (for best performance).

In comparison, we use the simulated corrections \mathbf{a}_{t_k} , which are generated from the true θ^* , as in (42). We set the magnitude $\|\mathbf{a}_{t_k}\|$ to three levels: $\|\mathbf{a}_{t_k}\| = 0.00125$, $\|\mathbf{a}_{t_k}\| = 0.001$, and $\|\mathbf{a}_{t_k}\| = 0.0008$ because we want to see how sensitive the magnitude-correction methods [6], [7] are to the magnitudes of corrections. Here, the correction time t_k for different magnitude levels is different random draws. To illustrate results, we measure the following three aspects for all methods.

- 1) *Weight error*, defined as $\|\theta_k - \theta^*\|^2$. This is to measure the error between the learned θ_k and the true θ^* .
- 2) *Cost regret* [6], defined as $(\theta^*)^T(\phi(\xi_{\theta_k}) - \phi(\xi_{\theta^*}))$. This is to measure the misalignment of the costs between robot current trajectory ξ_{θ_k} and the desired trajectory ξ_{θ^*} , evaluated under the true cost function $J(\theta^*) = (\theta^*)^T\phi(\xi)$.
- 3) *Trajectory error*, defined as $\|\xi_{\theta_k} - \xi_{\theta^*}\|^2$. It measures the distance between the robot current trajectory ξ_{θ_k} and the desired trajectory ξ_{θ^*} .

Given the simulated corrections \mathbf{a}_{t_k} , we run Algorithm 1 and the method [6], [7] [i.e., updating the weight vector via (46)]. We

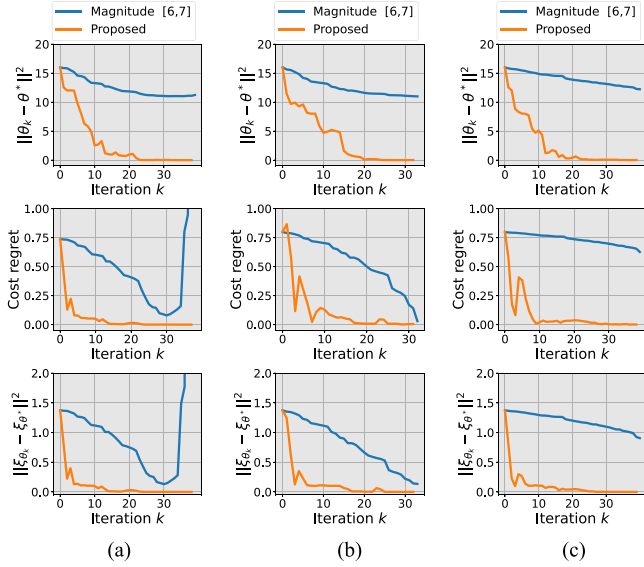


Fig. 10. Comparison between the proposed method (in orange lines) and the magnitude-correction methods [6], [7] (in blue lines). Each column corresponds to a different level of correction magnitude: (a) $\|\mathbf{a}_{t_k}\| = 0.00125$, (b) $\|\mathbf{a}_{t_k}\| = 0.001$, (c) $\|\mathbf{a}_{t_k}\| = 0.0008$, and correction time t_k in different magnitude levels are different random draws. For each magnitude level (each column), the first row shows the weight error versus iteration; the second row is the cost regret versus iteration; and the third row is the trajectory error versus iteration. Detailed analysis is given in the text.

show the results in Fig. 10, where the results of our method are in orange and the results of the magnitude-correction method [6], [7] are in blue. Each column in Fig. 10 corresponds to one level of correction magnitudes (recall that we want to see how robust a method is against different magnitudes of corrections). In each magnitude level (each column), the first row shows the weight error versus iteration, the second row shows the cost regret versus iteration, and the third row shows the trajectory error versus iteration. Based on the results in Fig. 10, we make the following comments.

For the magnitude-correction methods [6], [7] (in blue lines) in Fig. 10, comparing different columns, we note that larger magnitude leads to faster convergence. However, as shown by the left and middle panels in the first row, the magnitude-correction method converges to a local weight vector but not θ^* . This is not a surprise because the magnitude-correction method can only guarantee the convergence of the cost regret, as shown by [6], and not necessarily the convergence of the cost function itself. One conjecture to the above results could be that the local weight vector (different from the true θ^*) results in the same robot trajectory and the same optimal cost as θ^* does. This conjecture has been evidenced in the second row and third row in Fig. 10, where the left and middle panels show that the magnitude-correction methods, although having learned a different weight vector, have both the cost regret and trajectory error converging to zero.

Also for the magnitude-correction methods [6], [7] (in blue lines), when $\|\mathbf{a}_{t_k}\| = 0.00125$ (first column), we observe that there is a sharp increase of the cost regret and trajectory error

near convergence after 30 iterations. This could be because the magnitude-correction methods [6], [7] are not robust against the large magnitude of corrections near the desired trajectory. Specifically, after 30 iterations near convergence, a relatively larger magnitude, say $\|\mathbf{a}_{t_k}\| = 0.00125$, can be an overcorrection and thus violating $J(\bar{\mathbf{u}}_{0:T}^{\theta_k}, \theta^*) < J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$. This explanation has been supported by the second column of Fig. 10, where with a smaller magnitude $\|\mathbf{a}_{t_k}\| = 0.001$, there is no such overcorrection phenomenon.

Combining all three columns in Fig. 10, one can conclude that for the magnitude-correction methods [6], [7], large correction magnitude, such as $\|\mathbf{a}_{t_k}\| = 0.00125$, will have faster convergence, but it can lead to overcorrections, as shown in Fig. 10(a), making the learning process unstable. Smaller correction magnitude, such as $\|\mathbf{a}_{t_k}\| = 0.0008$, can avoid the overcorrection issue, but it leads to slow convergence, as shown in Fig. 10(c). Note that the magnitude change from $\|\mathbf{a}_{t_k}\| = 0.00125$ (unstable) in Fig. 10(a) to $\|\mathbf{a}_{t_k}\| = 0.0008$ (stable) in Fig. 10(c) is very small, which could suggest that the magnitude-correction methods are sensitive to the selection of correction magnitudes. Thus, it could be difficult in practice to provide a valid correction magnitude. We will further show and analyze this in our following user study in Section VI.

In contrast, Fig. 10 has shown that the proposed method consistently learns the true θ^* and achieves a zero cost regret and zero trajectory error regardless of the correction magnitude levels. Also, compared to the magnitude-correction methods [6], [7], the proposed method requires less learning iterations (corrections) for convergence. Since the proposed method only leverages the direction of \mathbf{a}_{t_k} regardless of $\|\mathbf{a}_{t_k}\|$, there is no overcorrection issues near convergence. In practice, the choice of directional corrections is more flexible than choosing magnitude corrections, as analyzed in Fig. 3 in Section II. We will continue to show this advantage in the following user study in Section VI.

VI. USER STUDY

To show the effectiveness of the proposed method for learning from real human directional corrections, we have developed two human-robot simulation games based on which a user study is conducted. One game is robot arm reaching (Fig. 11) and the other is 6-DoF quadrotor maneuvering (Fig. 12). In each game, a human participant observes the visualization of robot motion, while applying directional corrections through a keyboard. The goal of each game is to teach a robot to learn an objective function, such that it can successfully navigate through an environment without the knowledge about the environment obstacles. We have released the accompanying codes of those two games for the readers to try themselves: <https://github.com/wanxinjin/Learning-from-Directional-Corrections>.

In what follows, Sections VI-A and VI-B describe the designs of the two games. Section VI-C provides the details of human participants and the procedure of the user study. The outcomes of the user study are presented in Section VI-D, and detailed analysis is given in Section VI-E.

TABLE I
KEYBOARD INTERFACE FOR THE ROBOT ARM GAME

Keys	Directional correction	Interpretation of correction
<i>up</i>	$\mathbf{a} = [1, 0]$	Counter clock wise torque to Joint 1
<i>down</i>	$\mathbf{a} = [-1, 0]$	Clock wise torque to Joint 1
<i>left</i>	$\mathbf{a} = [0, 1]$	Counter clock wise torque to Joint 2
<i>right</i>	$\mathbf{a} = [0, -1]$	Clock wise torque to Joint 2

*When implementing [6] and [7], to allow magnitude corrections via the same interface, we additionally detect the key pressing duration δt and interpret the correction magnitude u to be proportional to δt with saturation u_{\max} ; i.e., $u = \min\{\beta\delta t, u_{\max}\}$ with $\beta = 2$ and $u_{\max} = 1$. The magnitude correction using the above interface is $\mathbf{u} = u\mathbf{a}$. The correction time t_k is the beginning time of the key press.

A. Robot Arm Reaching Game

1) *Robot Setup*: The dynamics of a robot arm and its physical parameters follow Section V-B. The robot initial state is $\mathbf{x}_0 = [-\frac{\pi}{2}, 0, 0, 0]^T$, as in Fig. 11. The parameterized cost function $J(\boldsymbol{\theta})$ in (2) is

$$\boldsymbol{\phi} = [q_1^2, q_1, q_2^2, q_2, \|\mathbf{u}\|^2]^T \in \mathbb{R}^5, \quad (47a)$$

$$\boldsymbol{\theta} = [\theta_1, \theta_2, \theta_3, \theta_4, \theta_5]^T \in \mathbb{R}^5. \quad (47b)$$

Here, $\boldsymbol{\theta}^T \boldsymbol{\phi}$ is a general second-order polynomial. It is worth noting that, in practice, if prior knowledge about good features is not available, general polynomial features are a good option. The final cost $h(\mathbf{x}_{T+1})$ in (2) is

$$h(\mathbf{x}_{T+1}) = 100 \left(\left(q_1 - \frac{\pi}{2} \right)^2 + q_2^2 + \dot{q}_1^2 + \dot{q}_2^2 \right) \quad (48)$$

as the robot arm aims to finally reach and stop at the target pose: $q_1^{\text{target}} = \frac{\pi}{2}$ and $q_2^{\text{target}} = 0$, marked in Fig. 11. The time horizon of the game is set as $T = 50$ (that is, $50\Delta = 10$ s). Since we choose the polynomial features (47a), $\boldsymbol{\theta}$ will dictate how the robot reaches the target (i.e., the specific trajectory to the goal). By default, the initial search space Ω_0 is set

$$\Omega_0 = \{\boldsymbol{\theta} \mid \theta_1, \theta_3 \in [0, 1], \theta_2, \theta_4 \in [-3, 3], \theta_5 \in [0, 0.5]\}. \quad (49)$$

Without human corrections, the robot with a random $\boldsymbol{\theta}_0$ will move and crash into the obstacle in Fig. 11. The goal of the game is to let a human participant teach the robot arm to learn a cost function, such that it can move from the initial condition and reach the target pose while avoiding the obstacle. Note that we have no knowledge about the obstacle.

2) *Interface*: In the robot arm game, we use a keyboard as the interface for a human participant to apply directional corrections. We customize (*up*, *down*, *left*, and *right*) keys and associate them with corresponding directional corrections as listed in Table I. During the game, at each learning iteration, a human participant is allowed to press one or multiple keys from (*up*, *down*, *left*, and *right*), and the interface is listening to which key(s) the human player hits and recording the timing of the keystroke(s). The recorded information is translated into the directional correction \mathbf{a}_{t_k} as per Table I. For example, at iteration k , while the robot is executing the motion trajectory $\boldsymbol{\xi}_{\boldsymbol{\theta}_k}$, a human player hits the *up* and *left* keys simultaneously at the time step 10; then the corresponding correction information is translated into $\mathbf{a}_{t_k} =$

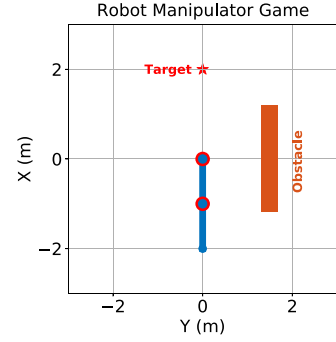


Fig. 11. Robot arm reaching game. A human participant teaches the robot arm to learn a cost function by applying directional corrections via a keyboard (see the instructions in Table I). The goal is to let the robot arm successfully move from the initial pose (the pose in the figure) to reach the target (the star) while avoiding the obstacle. Note that we have no knowledge about the location and size of the obstacle.

$[1, 1]^T$ with $t_k = 10$ according to Table I. We obtain averaged $\bar{\mathbf{a}}_k$ from \mathbf{a}_{t_k} via (5) and (6b).

B. Six-DoF Quadrotor Maneuvering Game

1) *Robot Setup*: The dynamics of a quadrotor is

$$\dot{\mathbf{r}}_I = \dot{\mathbf{v}}_I, \quad (50a)$$

$$m\dot{\mathbf{v}}_I = m\mathbf{g}_I + \mathbf{f}_I, \quad (50b)$$

$$\dot{\mathbf{q}}_{B/I} = \frac{1}{2}\Omega(\boldsymbol{\omega}_B)\mathbf{q}_{B/I}, \quad (50c)$$

$$J_B\dot{\boldsymbol{\omega}}_B = \boldsymbol{\tau}_B - \boldsymbol{\omega}_B \times J_B\boldsymbol{\omega}_B. \quad (50d)$$

Here, the subscripts B and I denote the quantities expressed in the quadrotor's body frame and world frame, respectively; m is the mass of the quadrotor; $\mathbf{r}_I \in \mathbb{R}^3$ and $\mathbf{v}_I \in \mathbb{R}^3$ are the position and velocity, respectively; $J_B \in \mathbb{R}^{3 \times 3}$ is the moment of inertia; $\boldsymbol{\omega}_B \in \mathbb{R}^3$ is the angular velocity; $\mathbf{q}_{B/I} \in \mathbb{R}^4$ is the unit quaternion [31] describing the attitude of the quadrotor's body frame with respect to the world frame; (50c) is the time derivative of the quaternion with $\Omega(\boldsymbol{\omega}_B)$ being the matrix form of $\boldsymbol{\omega}_B$ for quaternion multiplication [31]; $\boldsymbol{\tau}_B \in \mathbb{R}^3$ is the torque vector applied to the quadrotor; and $\mathbf{f}_I \in \mathbb{R}^3$ is the total force vector applied to the center of mass (COM). The total force magnitude $\|\mathbf{f}_I\| = f$ (along the z -axis of the quadrotor's body frame) and the torque $\boldsymbol{\tau}_B = [\tau_x, \tau_y, \tau_z]^T$ are generated by four rotor thrusts $[T_1, T_2, T_3, T_4]^T$ as follows:

$$\begin{bmatrix} f \\ \tau_x \\ \tau_y \\ \tau_z \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & -l_w/2 & 0 & l_w/2 \\ -l_w/2 & 0 & l_w/2 & 0 \\ \kappa & -\kappa & \kappa & -\kappa \end{bmatrix} \begin{bmatrix} T_1 \\ T_2 \\ T_3 \\ T_4 \end{bmatrix}, \quad (51)$$

with l_w denoting the wing length of the quadrotor and κ a fixed constant; here, $\kappa = 0.1$. In dynamics (50), the gravity constant $\|\mathbf{g}_I\|$ is set as 10 m/s^2 and the other physical parameters are units. We define the state vector of the quadrotor as

$$\mathbf{x} = [\mathbf{r}_I \quad \mathbf{v}_I \quad \mathbf{q}_{B/I} \quad \boldsymbol{\omega}_B] \in \mathbb{R}^{13} \quad (52)$$

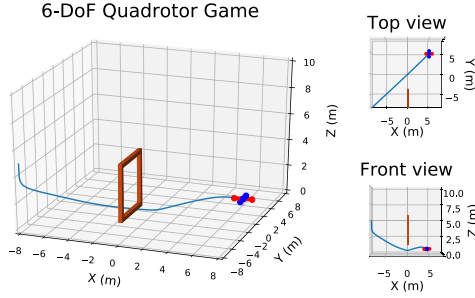


Fig. 12. Six-DoF quadrotor game. A human participant teaches the quadrotor to learn a cost function by applying directional corrections through a keyboard (see the instruction in Table II). The goal is that the quadrotor can fly from the initial position (in the bottom left), go through a gate (colored in brown), and land on the target (in the upper right). Note that we have no knowledge about the gate.

and the control input vector as

$$\mathbf{u} = [T_1 \ T_2 \ T_3 \ T_4]^T \in \mathbb{R}^4. \quad (53)$$

We discretize (50) by the Euler method with interval $\Delta = 0.1$ s. To achieve $SE(3)$ maneuvering, we define the attitude error between the quadrotor attitude \mathbf{q} and target $\mathbf{q}^{\text{target}}$ as [32]

$$e(\mathbf{q}, \mathbf{q}^{\text{target}}) = \frac{1}{2} \text{trace}(I - R^T(\mathbf{q}^{\text{target}})R(\mathbf{q})) \quad (54)$$

where $R(\mathbf{q}) \in \mathbb{R}^{3 \times 3}$ is the rotation matrix corresponding to \mathbf{q} .

In the cost function (2), we set the final cost term as

$$h(\mathbf{x}_{T+1}) = \|\mathbf{r}_I - \mathbf{r}_I^{\text{target}}\|^2 + 10\|\mathbf{v}_I\|^2 + 100e(\mathbf{q}_{B/I}, \mathbf{q}_{B/I}^{\text{target}}) + 10\|\mathbf{w}_B\|^2 \quad (55)$$

because the quadrotor aims to finally land on a target position $\mathbf{r}_I^{\text{target}}$ with a target attitude $\mathbf{q}_{B/I}^{\text{target}}$. Here, $\mathbf{r}_I = [r_x, r_y, r_z]^T$ is the position of the quadrotor expressed in the world frame. We set the weight-feature cost term as

$$\phi = [r_x^2 \ r_x \ r_y^2 \ r_y \ r_z^2 \ r_z \ \|\mathbf{u}\|^2]^T \in \mathbb{R}^7, \quad (56a)$$

$$\theta = [\theta_1 \ \theta_2 \ \theta_3 \ \theta_4 \ \theta_5 \ \theta_6 \ \theta_7]^T \in \mathbb{R}^7. \quad (56b)$$

Here, feature vector ϕ consists of general polynomial features and θ will determine the specific trajectory of the quadrotor to the target. By default, the initial weight search space is

$$\Omega_0 = \{\theta \mid \theta_1, \theta_3, \theta_5 \in [0, 1], \theta_2, \theta_4, \theta_6 \in [-8, 8], \theta_7 \in [0, 0.5]\}. \quad (57)$$

As in Fig. 12, the goal of this quadrotor game is to let a human participant teach the quadrotor to learn a cost function, such that the quadrotor can successfully fly from the initial position $\mathbf{r}_I(0) = [-8, -8, 5]^T$ (in the bottom left), go through a gate (in brown color), and finally land on a target position $\mathbf{r}_I^{\text{target}} = [8, 8, 0]^T$ (in the upper right) with the target attitude $\mathbf{q}_{B/I}^{\text{target}} = [1, 0, 0, 0]^T$. The initial attitude of the quadrotor is $\mathbf{q}_{B/I}(0) = [1, 0, 0, 0]^T$ and initial velocities zeros. The time horizon for this game is $T = 50$, that is, $T\Delta = 5$ s. Note that we have no knowledge about the gate.

TABLE II
KEYBOARD INTERFACE FOR 6-DOF QUADROTOR GAME

Keys	Directional correction	Interpretation of correction
“up”	$T_1=1, T_2=1,$ $T_3=1, T_4=1$	Upward force applied at COM
“down”	$T_1=-1, T_2=-1,$ $T_3=-1, T_4=-1$	Downward force applied at COM
“w”	$T_1=0, T_2=1$ $T_3=0, T_4=-1$	Negative torque along body-axis x
“s”	$T_1=0, T_2=-1$ $T_3=0, T_4=1$	Positive torque along body-axis x
“a”	$T_1=1, T_2=0$ $T_3=-1, T_4=0$	Negative torque along body-axis y
“d”	$T_1=-1, T_2=0$ $T_3=1, T_4=0$	Positive torque along body-axis y

*When implementing [6] and [7], to allow magnitude corrections via the same interface, we additionally detect key pressing duration δt and interpret the correction magnitude u to be $u = \min\{\beta\delta t, u_{\max}\}$ with $\beta = 2$ and $u_{\max} = 1$ in our implementation. Thus, the magnitude correction using the above interface is $\mathbf{u} = u[T_1, T_2, T_3, T_4]^T$. The correction time t_k is the beginning time of the key press.

2) *Interface*: The keyboard interface for applying directional corrections is in Table II. Specifically, we customize the (“up,” “down,” “w,” “s,” “a,” and “d”) keys and map them to specific directional correction signals. During a learning iteration, a human participant is allowed to press one or multiple keys in Table II. The interface listens to the keystrokes and translates the keystrokes into the directional corrections based on Table II. The time step at which a key is hit is the correction time t_k . For example, if a human participant presses “s” key at time step 5, then the directional correction will be $\mathbf{a}_{t_k} = [0, -1, 0, 1]^T$ with $t_k = 5$. Based on (5) and (6b), we obtain $\bar{\mathbf{a}}_k$ from \mathbf{a}_{t_k} .

C. Participants and Procedure

A total of 17 volunteers from Purdue College of Engineering have participated in our user study. Among these 17 participants, 1 was female and 16 were male, with ages from 20 to 37 years (25.812 ± 3.936). Five of those 17 participants had no robotics or related background, and all participants were novices to the two games. This user study had been reviewed and approved by the Institutional Review Board of Purdue University, and all participants had signed the consent forms.

Each participant was instructed to play the above two games. In each game, a participant played five successive rounds with the proposed method and five successive rounds with the magnitude-correction method [6], [7]. A *game round* is defined as a complete run of a learning method until the success or failure of the round. A failure of a game round is identified if the robot fails to accomplish the task after the total number of human corrections has exceeded the maximum correction count. In our user study, the maximum correction count is 20 for the robot arm game and 15 for the quadrotor game. We set the maximum correction count because the participants are expected to try their best to teach the robot with as few corrections as possible.

We used a within-subjects design: The order of the two games (i.e., which game goes first) and learning methods used (i.e., which method is used first) were counterbalanced across all

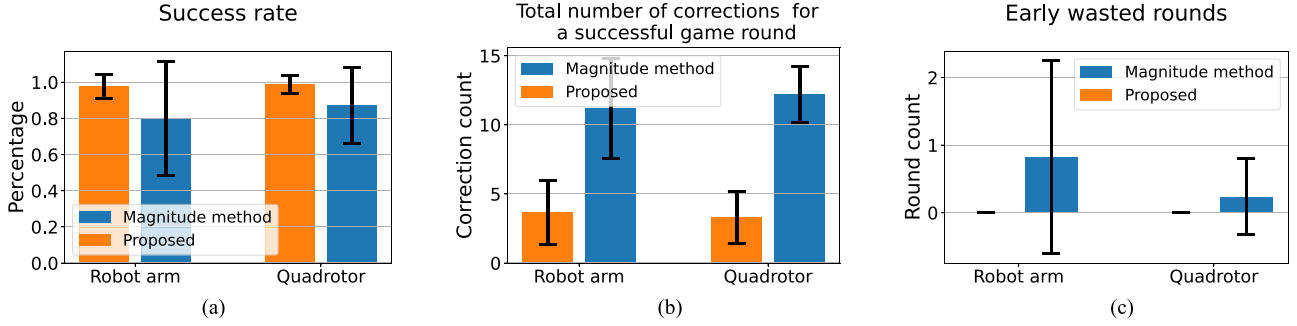


Fig. 13. Outcomes of all 17 participants for playing different robot games (robot-arm or quadrotor) using different learning methods (the proposed method and magnitude-correction method [6], [7]). (a) Success rate, i.e., the ratio of successful game rounds over total five rounds attempted. (b) Total number of human corrections needed for one successful game round. (c) Number of the early wasted (failed) game rounds before a participant constantly play successfully. The bars denote the mean values over all participants, and the top line segments are the standard deviations. Analysis based on statistical tests (with significance level $\alpha = 0.05$) is in Section VI-E. One illustrative round for the robot arm game from one participant is in Fig. 14 and Table III, and one illustrative quadrotor game round from a participant is in Fig. 15 and Table IV.

participants in our user study. Also, the participants were not informed which methods were being used for the current game. Before starting a game, each participant was instructed how to play the game and was given one hands-on game round to get familiar with the interface before the experiment recording starts. After the games, each participant was asked to finish a postgame survey about her/his opinions of the game experience (which will be reported in Section in VI-E4).

D. Outcome Measurements

For each game (robot arm or quadrotor) with each learning method (the proposed method or magnitude-correction method [6], [7]), we measure the following outcomes for one participant in his/her five successive game rounds.

- 1) *Success rate.* This is the ratio of successful game rounds over all rounds attempted (which is five here). This outcome indicates the **efficacy** of a learning method.
- 2) *Total number of corrections for a successful game round.* This measures how many human corrections are needed for a successful game round. This measure can show the **efficiency/effortlessness** of a learning method—fewer corrections mean less human effort in teaching a robot.
- 3) *Early wasted rounds.* This is to measure how many early rounds of a game are wasted (failed) before a participant begins to constantly play successful game rounds. This can show the **accessibility** of a learning method—fewer early wasted rounds indicate a more accessible experience for a participant to successfully teach a robot.

The outcomes of all participants are organized according to the method type (the proposed method and magnitude-correction method [6], [7]) and the game type (robot-arm and quadrotor). We report all outcomes in Fig. 13 and provide a detailed analysis of the results based on statistical tests.

E. Results and Analysis

The statistics of three outcomes over all 17 participants are shown in Fig. 13. Each outcome is presented with respect to each learning method (the proposed method or magnitude-correction

method [6], [7]) on each game (robot arm or quadrotor). All three outcomes here are averaged over all participants with the error bar denoting the standard deviation over all participants. We have the following analysis of the results based on statistical tests with significance level $\alpha = 0.05$.

1) *Success Rate:* In Fig. 13(a), for the robot arm game, the proposed method obtains a success rate of $97.65\% \pm 6.64\%$ compared to the magnitude-correction method's [6], [7] $80.00\% \pm 31.62\%$. The proposed method yields a significantly higher success rate than the magnitude-correction method with t -score $t = 2.25$ and p -value $p = 0.031 < \alpha$. For the quadrotor game in Fig. 13(a), the proposed method attains $94\% \pm 12.8\%$ versus the magnitude-correction method's $82\% \pm 20.9\%$. It also shows a significantly higher success rate than the magnitude-correction method [6], [7], with t -score $t = 2.24$ and p -value $p = 0.032 \leq \alpha$.

We provide the following comments for the above results. Recall that the magnitude-correction method [6], [7] assumes $J(\mathbf{u}_{0:T}^{\theta_k} + \bar{\mathbf{a}}_k, \theta^*) < J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$, while our method assumes $\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \bar{\mathbf{a}}_k \rangle > 0$, regardless of magnitude $\|\bar{\mathbf{a}}_k\|$. As shown in Fig. 3, *the allowable region of directional corrections is much larger than that of magnitude corrections*. Since all participants are novices to both games, giving valid magnitude correction could be difficult for them. By contrast, novice participants are more likely to give valid directional corrections. This leads to higher success rate and lower variance for the proposed method than the magnitude-correction method.

The difficulty of giving valid magnitude corrections can also be seen from the number of early wasted/failed rounds in Fig. 13(c). For the magnitude method, before a participant masters a game, 0.82 ± 1.42 early game rounds are wasted/failed for the robot arm game, and 0.24 ± 0.56 early rounds are failed for the quadrotor game. In contrast, the proposed method has 0 ± 0 early wasted rounds in both games. This means that a novice participant needs more practice to master the skill of giving valid magnitude corrections.

2) *Total Number of Corrections for Successful Game:* In Fig. 13(b), for a successful robot arm game, the proposed method needs a total of 3.66 ± 2.29 corrections, while the

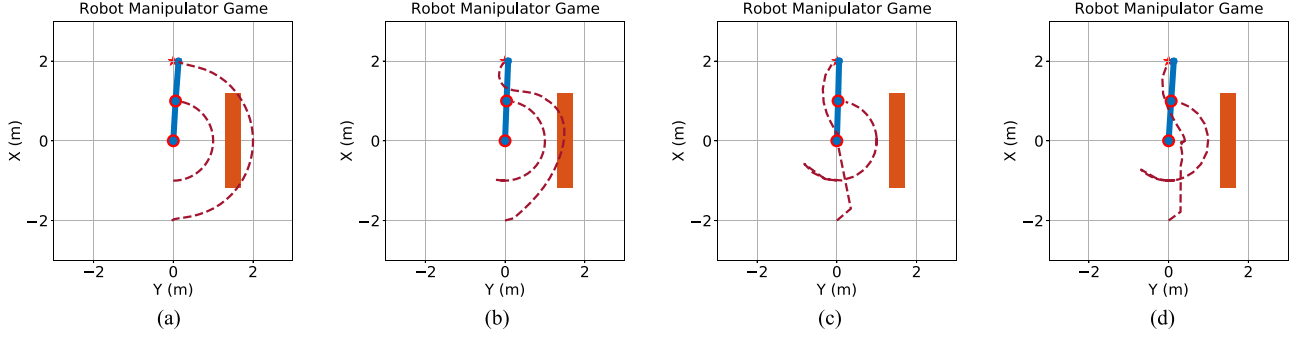


Fig. 14. One illustrative round for the robot arm game from one participant. (a) Iteration $k = 1$. (b) Iteration $k = 2$. (c) Iteration $k = 3$. (d) Iteration $k = 4$. At each iteration k , the robot weight vector guess θ_k and the participant directional correction \mathbf{a}_{t_k} are reported in Table III.

TABLE III
ILLUSTRATIVE ROUND FOR THE ROBOT ARM GAME FROM ONE PARTICIPANT

Iteration k	Current weight vector guess θ_k	Participant's directional correction \mathbf{a}_{t_k} and correction time t_k
$k = 1$	$\theta_k = [0.50, 0.00, 0.50, -0.00, 0.25]^T$	$\mathbf{a}_{t_k} = [0, 1]$ (i.e., <i>left</i> key pressed) and $t_k = 11$
$k = 2$	$\theta_k = [0.50, 0.00, 0.50, -1.50, 0.25]^T$	$\mathbf{a}_{t_k} = [0, 1]$ (i.e., <i>left</i> key pressed) and $t_k = 16$
$k = 3$	$\theta_k = [0.50, 0.00, 0.34, -2.03, 0.25]^T$	$\mathbf{a}_{t_k} = [-1, 0]$ (i.e., <i>down</i> key pressed) and $t_k = 34$
$k = 4$	$\theta_k = [0.50, 1.48, 0.36, -2.00, 0.25]^T$	<i>Game success!</i>

magnitude-correction method [6], [7] requires 11.21 ± 3.62 corrections. Thus, the proposed method requires significantly fewer corrections than the magnitude-correction method, with t -score $t = 7.27$ and p -value $p = 2.94 \times 10^{-8} \leq \alpha$. In the quadrotor game, the proposed method takes a total of 3.27 ± 1.88 corrections for a successful round, while the magnitude-correction method needs 12.20 ± 2.02 corrections. This again shows that the proposed method is significantly more efficient (fewer corrections) than the magnitude-correction method, with t -score $t = 13.32$ and p -value $p = 1.35 \times 10^{-14} < \alpha$.

We have the following interpretations for the above results. Recall that the assumption for a valid magnitude correction $\bar{\mathbf{a}}_k$ is $J(\mathbf{u}_{0:T}^{\theta_k} + \bar{\mathbf{a}}_k, \theta^*) < J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*)$. Since all participants are novice to the games, the magnitude corrections made by a participant were not always valid in a game round; in other words, some magnitude corrections can be overcorrection. Because of those invalid corrections, the magnitude-correction method takes longer (more corrections) to succeed. In contrast, the proposed method only leverages directional corrections, which are more likely to satisfy $\langle -\nabla J(\mathbf{u}_{0:T}^{\theta_k}, \theta^*), \bar{\mathbf{a}}_k \rangle > 0$; it requires fewer human corrections (effort).

The total number of corrections for a successful game indicates how much effort a participant needs to successfully teach a robot. Thus, the above results in Fig. 13(b) show that the proposed method is significantly more effortless than the magnitude-correction method [6], [7].

3) *Early Wasted Rounds*: In Fig. 13(c), in the robot arm game (five rounds in total), the early wasted (failed) rounds is 0 ± 0 for the proposed method and 0.82 ± 1.42 for the magnitude-correction method [6], [7]; thus, the proposed method requires significantly fewer rounds of practice (thus is more accessible) than the magnitude-correction method, with t -score $t =$

2.38 and p -value $p = 0.02 < \alpha$. For the quadrotor game (five rounds in total), a participant wasted 0 ± 0 early rounds using the proposed method, while 0.24 ± 0.56 early rounds using the magnitude-correction method. The early wasted rounds of the proposed method are not significantly lower than that of the magnitude-correction method, with t -score $t = 1.73$ and p -value $p = 0.09 > \alpha$.

We have the following explanations for the above results. For the magnitude-correction method, since not all magnitude corrections are valid, a participant needs more early trials to realize the effects of correction magnitude on the robot motion and how to apply valid magnitude. Instead, for the proposed method, a participant only gives directional commands, which can be more intuitive, as shown in Figs. 14 and 15. Fig. 13(c) shows that the proposed method is potentially more accessible than the magnitude-correction method [6], [7]

4) *Other Observations and Subjective Survey*: We also observe that the number of the directional corrections in the user study is generally smaller than that in the simulation cases in Section V. This is because the convergence of the robot motion trajectory (to the desired motion in the human's perspective) is empirically faster than the convergence of the cost function weight vector (to the true θ^*). This has been shown in Fig. 10, where the convergence of $\|\theta_k - \theta^*\|^2$ requires around 20 iterations, while the convergence of $\|\xi_{\theta_k} - \xi_{\theta^*}\|^2$ requires only 10 iterations. Since a game is deemed successful as long as the robot motion can successfully navigate through the environment, a participant usually takes fewer corrections. We also find that there is no unique solution for the choice of directional corrections and that different participants applied different corrections to manage the games.




Fig. 15. One illustrative round for the quadrotor game from one participant. At each iteration k , the quadrotor weight vector guess θ_k and the participant directional correction \mathbf{a}_{t_k} are in Table IV. Iteration $k = 1$ is shown in Fig. 12. (a) Iteration $k = 2$. (b) Iteration $k = 3$ (c) Iteration $k = 4$. (d) Iteration $k = 5$.

TABLE IV
ILLUSTRATIVE ROUND FOR THE QUADROTOR GAME FROM ONE PARTICIPANT (ITERATION $k = 1$ IS SHOWN IN FIG. 12)

Iteration k	Current weight vector guess θ_k	A participant's directional corrections \mathbf{a}_{t_k} and correction time t_k
$k = 1$	$\theta_k = [0.50, 0.00, 0.50, -0.00, 0.50, -0.00, 0.25]^T$	$\mathbf{a}_{t_k} = [1, 1, 1, 1]$ (i.e., <i>up</i> key pressed) and $t_k = 8$ $\mathbf{a}_{t_k} = [1, 1, 1, 1]$ (i.e., <i>up</i> key pressed) and $t_k = 20$
$k = 2$	$\theta_k = [0.50, -0.00, 0.50, -0.00, 0.50, -3.99, 0.25]^T$	$\mathbf{a}_{t_k} = [-1, -1, -1, -1]$ (i.e., <i>down</i> key pressed) and $t_k = 14$
$k = 3$	$\theta_k = [0.50, -1.70, 0.50, -1.70, 0.52, -1.89, 0.25]^T$	$\mathbf{a}_{t_k} = [0, -1, 0, 1]$ (i.e., <i>s</i> key pressed) and $t_k = 13$
$k = 4$	$\theta_k = [0.50, -2.76, 0.50, -2.45, 0.60, -2.22, 0.25]^T$	$\mathbf{a}_{t_k} = [0, -1, 0, 1]$ (i.e., <i>s</i> key pressed) and $t_k = 19$
$k = 5$	$\theta_k = [0.50, -3.11, 0.50, 4.89, 0.65, -2.67, 0.25]^T$	<i>Game success!</i>

We have also performed a postgame survey about how participants perceived the two learning methods. Ten out of 17 participants said that they felt *the robot with the proposed method is much "smarter" than the robot with the magnitude-correction method* [6], [7]. Those opinions are also consistent with the results in Fig. 13, as we have analyzed above.

5) *Summary*: In summary of all statistical analyses above, we conclude that compared to the state-of-the-art magnitude-correction method [6], [7], the proposed method is:

- 1) *significantly more effective* ($p < 0.033$)—it attains higher success rate for teaching robots;
- 2) *significantly more efficient* ($p < 10^{-7}$)—fewer human corrections are needed for successful robot learning;
- 3) *potentially more accessible* ($p \leq 0.09$) for human users—a novice human user takes fewer trials before constantly providing successful corrections to a robot.

The postgame survey reports that a majority of the user-study participants indicated that a robot with the proposed method is much "smarter" than with the magnitude-correction method.

VII. REAL-WORLD EXPERIMENT

In this section, we test the proposed method in a real-world experiment using a Parrot Mambo quadrotor.

A. Experiment Setup and Procedure

Our real-world experiment is a Parrot Mambo quadrotor task shown in Fig. 16. The goal of the experiment is to let a human user teach the quadrotor to accomplish the following task: starting from the initial position (red circle), flying through the gate (labeled in yellow), and landing on the target position (green circle). Without human correction, the quadrotor with a random initial cost function fails to accomplish the task. The procedure of learning from human directional corrections is as follows.




Fig. 16. Setup of the real-world experiment. A human user teaches the quadrotor to fly from an initial position (red circle), go through the gate (with yellow label), and finally, land on a target position (green circle). Without human corrections, the quadrotor is initialized with a random cost function and fails to accomplish the task. Note that we have no knowledge about the gate. (a) Main camera view. (b) Back camera view.

- 1) First, perform the quadrotor motion planning in a remote computer by solving a trajectory optimization with the current cost function.
- 2) Second, the quadrotor executes the plan (tracking control), and at the same time, a human user watches the quadrotor's motion and applies directional corrections via the keyboard of the remote computer.
- 3) Third, update the quadrotor's cost function in the remote computer according to Algorithm 1; then repeat the above steps until the success of the task.

In the above procedure, the interface for a human user to apply directional corrections is still a keyboard, following Table II. The communication frequency between the quadrotor and the remote computer is around 10 Hz. Note that because the cost function in some learning iterations may lead the quadrotor to collide with the gate during execution, we have included a collision stop mechanism: if the quadrotor detects that it is too close to the gate, an emergence stop is triggered. In the iteration where an emergence stop is triggered, human corrections, if applied before




Fig. 17. Results of a real-world quadrotor learning from human directional corrections. Here, the quadrotor trajectory is marked in red lines. The initial position (red circle) and target position (green circle) are also marked. (a) Iteration $k = 1$. (b) Iteration $k = 2$. (c) Iteration $k = 3$. (d) Iteration $k = 4$. In (a), during the quadrotor flying, the human user applied a torque correction in the opposite of the x -axis, shown by the green arrow. In (b), during the quadrotor flying, the human user applied an upward thrust correction (in green arrow) in z -axis. In (c), the quadrotor detected a potential collision with the left pillar of the gate and thus triggered an emergence stop. But before the emergence stop, the human user had already applied a positive x -axis (in green arrow) torque to correct the quadrotor. In (d), the quadrotor successfully learned a cost function to fly from the initial position, go through the gate, and finally reach the target position.

an emergence stop, can still count for updating the cost function for the next iteration. We will discuss the damage protection of robot learning in detail in Appendix B-C.

We have invited a human user (not the coauthors themselves), who is a novice to our work, to perform the above experiment. The human user first had a successful warm-up training in our previous quadrotor game in order to get familiar with the interface (Table II) and then successfully conducted this real-world quadrotor experiment in just one run.

B. Results and Analysis

The results in all learning iterations are in Fig. 17. At each iteration, the quadrotor's trajectory is highlighted in the red lines, and the human directional corrections are labeled in the green arrows. The initial and target positions are also marked. We provide the following analysis for the results in Fig. 17.

In the first iteration in Fig. 17(a), since the quadrotor was initialized with a random cost function, the planned and executed trajectory (labeled in the red line) are off the gate—the quadrotor missed the gate and flew from its right side. Thus, when observing the quadrotor's motion, the human user applied a directional torque to correct the quadrotor, which is in the direction of the negative x -axis, as shown by the green arrow in Fig. 17(a). After receiving the correction, the quadrotor updated its cost function and planned a new trajectory for the next iteration in Fig. 17(b). In the second iteration in Fig. 17(b), while the quadrotor was executing the new motion, the human user observed that its motion was too below the gate and thus applied a directional force in the positive z -axis. Using this directional correction, the quadrotor updated its cost function for the third iteration in Fig. 17(c).

In the third iteration in Fig. 17(c), while executing the planned motion, the quadrotor detected a potential collision with the left pillar of the gate and thus immediately triggered an emergence stop. However, because humans are able to predict the potential collision, the human user had already applied a directional torque in the positive x -axis (the green arrow) *before the emergence stop*. Thus, although the quadrotor had stopped in emergence, the directional correction still counted in the update of the cost function. In the fourth iteration in Fig. 17(d), the quadrotor successfully flew through the gate and reached the target position.

Thus, it took only three human directional corrections for the quadrotor to learn a cost function to accomplish the task in an unknown environment. Based on the above results and analysis, we can conclude that the proposed method is effective and efficient for learning an objective function for a desired robot motion from human directional corrections.

VIII. CONCLUSION

This article has proposed a new method to enable a robot to learn an objective function from human directional corrections. A human directional correction can be any input change to the robot as long as it is in a direction that improves the robot's current motion relative to an implicit objective function. The proposed learning method only uses the direction of a correction to update the estimate of the objective function. The learning process is based on the cutting plane method and has straightforward geometric interpretations. We have established the theoretical results to show the convergence of the estimate of the objective function toward the true one induced by all human corrections. The proposed approach has been validated by numerical examples, a user study on two human–robot games, and a real-world quadrotor experiment. The results confirm the convergence and efficacy of the proposed method and show that the proposed method is significantly more effective (higher success rate), efficient/effortless (less correction needed), and potentially more accessible (fewer early wasted trials) than the state-of-the-art human–robot learning methods.

APPENDIX A PROOF OF LEMMA 2

First, we prove (25). From Step 2 in the main algorithm, we know that the robot's current trajectory $\xi_{\theta_k} = \{x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k}\}$ is a result of minimizing the cost function $J(\theta_k)$ subject to dynamics constraint in (1). This means that ξ_{θ_k} must satisfy the optimality condition (i.e., first-order condition). Following a similar derivation from (12) to (19) in the proof of Lemma 1, we can obtain the optimality condition

$$\begin{aligned} 0 = \nabla J(u_{0:T}^{\theta_k}, \theta_k) &= H_1(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k})\theta_k \\ &+ H_2(x_{0:T+1}^{\theta_k}, u_{0:T}^{\theta_k})\nabla h(x_{T+1}^{\theta_k}). \end{aligned} \quad (58)$$

It is worth mentioning that the above optimality condition (58) is also derived in [5]. As a result

$$0 = \left\langle \nabla J(\mathbf{u}_{0:T}^{\theta_k}, \boldsymbol{\theta}_k), \bar{\mathbf{a}}_k \right\rangle \quad (59)$$

$$= \left\langle \mathbf{H}_1(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k}) \boldsymbol{\theta}_k, \bar{\mathbf{a}}_k \right\rangle + \left\langle \mathbf{H}_2(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k}) \nabla h(\mathbf{x}_{T+1}^{\theta_k}), \bar{\mathbf{a}}_k \right\rangle \quad (60)$$

$$= \langle \mathbf{h}_k, \boldsymbol{\theta}_k \rangle + b_k \quad (61)$$

where the third equality is due to the definition of hyperplane in (8). This completes the proof of (25).

Next, we prove (26) by induction. In the main algorithm, we know $\boldsymbol{\theta}^* \in \Omega_0$ for $k = 0$. Assume that $\boldsymbol{\theta}^* \in \Omega_{k-1}$ holds at the $(k-1)$ th iteration. By Step 3 in the main algorithm, we have

$$\Omega_k = \Omega_{k-1} \cap \{ \boldsymbol{\theta} \in \mathbb{R}^r \mid \langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k < 0 \}. \quad (62)$$

In order to prove $\boldsymbol{\theta}^* \in \Omega_k$, we only need to show that

$$\langle \mathbf{h}_k, \boldsymbol{\theta}^* \rangle + b_k < 0 \quad (63)$$

which is true according to (7) in Lemma 1. Thus, $\boldsymbol{\theta}^* \in \Omega_k$ also holds at the k th iteration. Thus, we conclude that (26) holds. This completes the proof of Lemma 2. \square

APPENDIX B DISCUSSION

A. Learning Nonconvex Cost Functions

In our problem formulation and method development, we have not imposed any restriction on the convexity of the cost function (2). However, the proposed method can sufficiently handle both convex and nonconvex cost functions. But for nonconvex cost functions, we have the following comments.

For a general nonconvex cost function (2), the robot desired motion can be a *local* minimum of this nonconvex cost function. As long as all human corrections aim to drive the robot toward the same local minimum (i.e., the same desired motion), all theories and methods developed in the article still apply. In the otherwise case, human corrections may not have a consistent goal. For example, in one iteration, the human intends to drive the robot toward one local minimum, while in the other iteration, the human aims to drive the robot to a *different* local minimum—a different desired motion. In those nonconsistent cases, assumption (4) would be violated, thus potentially leading to the failure of the proposed method.

In sum, if the cost function (2) is nonconvex, the desired robot motion $\boldsymbol{\xi}_{\theta^*}$ can be a local minimum of the nonconvex cost function $J(\boldsymbol{\theta}^*)$. The proposed method in this article still applies if all human corrections aim to drive the robot toward the same local minimum $\boldsymbol{\xi}_{\theta^*}$.

B. On-the-Fly Implementation of the Proposed Method

In the previous experiments, each learning iteration requires a robot to start over the task. However, this is not necessary, and one can readily implement the proposed method in an on-the-fly manner. An on-the-fly implementation is given in Algorithm 2.

Algorithm 2: Learning from directional corrections (an on-the-fly implementation version)

Input: Specify a termination threshold ϵ and use it to compute the maximum iteration K by (35).

Initialization: Initial weight search space Ω_0 in (21), and initial robot state \mathbf{x}_{t_0} with $t_0 = 0$

for $k = 1, 2, \dots, K$ **do**

 Choose a weight vector guess $\boldsymbol{\theta}_k \in \Omega_{k-1}$ by Lemma 3;

 Starting from the current robot state $\mathbf{x}_{t_{k-1}}$, plan a new robot motion $\boldsymbol{\xi}_{\theta_k}$ by solving a trajectory optimization with the cost function $J(\boldsymbol{\theta}_k)$ and dynamics (1);

 Starting from $\mathbf{x}_{t_{k-1}}$, the robot executes the planned motion trajectory $\boldsymbol{\xi}_{\theta_k}$ while receiving human directional corrections $\bar{\mathbf{a}}_{t_k}$;

 Compute the matrices $\mathbf{H}_1(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$ and $\mathbf{H}_2(\mathbf{x}_{0:T+1}^{\theta_k}, \mathbf{u}_{0:T}^{\theta_k})$, and then compute the hyperplane and half space $\langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k < 0$ by (7) and (8);

 Update the weight search space by $\Omega_k = \Omega_{k-1} \cap \{ \boldsymbol{\theta} \in \mathbb{R}^r \mid \langle \mathbf{h}_k, \boldsymbol{\theta} \rangle + b_k < 0 \}$ by (22);

end

Output: $\boldsymbol{\theta}_K$.

Specifically, in Algorithm 2, after the robot receives a human directional correction \mathbf{a}_{t_k} at time step t_k , instead of starting over the task (i.e., returning to the very early initial condition \mathbf{x}_0), we can update the robot cost function immediately and plan the motion by starting from the robot *current* state \mathbf{x}_{t_k} (on which the correction \mathbf{a}_{t_k} is made). Thus, the robot will continue to execute the task from the latest state without starting over. For this on-the-fly implementation, all theories and other properties of our method remain unchanged.

C. Damage Prevention During Learning

Since an intermediate cost function during learning iterations can lead a robot to collision or damage, we briefly discuss how to avoid such scenarios.

One effective way to prevent damage and collision during robot learning is to add emergence stop mechanisms in the robot's lower level control, as we have adopted for our real-world quadrotor experiment in Section VII. Fortunately, as long as a human user applies directional corrections before the emergence stop, those corrections still count for the update of the cost function. Based on the previous user study and real-world experiment, we observe that humans usually have a good ability to predict potential collisions and are able to make preemptive corrections *before the robot triggers emergence stop*. Thus, it is a usual case that a robot has already received human directional corrections before triggering an emergence stop, and, hence, the update of the cost function continues. This has been shown in Fig. 17(c) in our real-world experiment in Section VII.

D. Choice of Ω_0 and ϵ in Algorithm 1

The initial weight search space Ω_0 in (21) should include any possible true $\boldsymbol{\theta}^*$, i.e., $\boldsymbol{\theta}^* \in \Omega_0$. Although this is hard to verify as $\boldsymbol{\theta}^*$ is usually unknown in practice, a good practice is to choose as large Ω_0 as possible (note that Ω_0 also needs to

ensure the existence of solution ξ_{θ} ; this is the reason why in our experiments, the range of the weight for second-order terms in a polynomial is always positive). Alternatively, one can also use a trial-and-error procedure to set the initial Ω_0 : first, try a small Ω_0 ; under this small Ω_0 , if the final (converged) robot motion is not desired, increase the size of Ω_0 and repeat this process until satisfied. Our previous experiment experience has shown that choosing a good Ω_0 is not difficult. This is because even a small size of Ω_0 has enough expressiveness power to represent a good variety of robot motions. This empiricism is consistent with the recent results of learning implicit models [33]–[35] in the machine learning community, where it shows that simple objective or energy functions can have enough representation power.

In Algorithm 1, ϵ determines the accuracy of weight vector convergence, as stated in Theorem 1. The choice of ϵ depends on specific accuracy requirements, and a large ϵ will terminate the algorithm early. In fact, it is always easy to set ϵ using a reasonably small value because our previous experiment, such as in Fig. 10, has shown that the convergence of the robot motion trajectory is much faster than the convergence of the objective function itself. This means that it is a usual case that one observes a good convergence of robot motion trajectory before the convergence of cost function reaches termination condition. This empiricism has also been reported in the literature on inverse optimal control such as [4].

E. Free From Noisy Robot Execution

Finally, we would point out that, in implementation, the noise in robot states and controls cannot directly enter into the proposed algorithm. Specifically, as stated in Algorithm 1 and our real-world experiment, the proposed algorithm decouples motion planning (i.e., plan ξ_{θ_k} by solving a trajectory optimization) and the robot execution of ξ_{θ_k} . The noise of the robot states and inputs can only enter into the robot execution stage, not the motion planning stage. While a human observes the robot's noisy execution of ξ_{θ_k} and gives directional corrections \bar{a}_k , the proposed algorithm updates the cost function using the *noise-free* ξ_{θ_k} taken from the motion planner instead of from the robot's execution. Thus, during the entire learning process, the noise in the robot's actual execution will never enter into the algorithm and, thus, will not influence the learning results.

APPENDIX C OTHER CHOICES OF θ_k

For the weight search space $\Omega_{k-1} \subset \mathbb{R}^r$, we choose θ_k as the center of MVE inscribes Ω_{k-1} . Other choices for θ_k could be the center of gravity [36], the Chebyshev center [37], the analytic center [38], etc.

1) *Center of Gravity*: The center of gravity for a polytope Ω is defined as

$$\theta_{cg} = \frac{\int_{\Omega} \theta d\theta}{\int_{\Omega} d\theta}. \quad (64)$$

Following [39], the volume reduction rate using the center of gravity is

$$\frac{\text{Vol}(\Omega_{k+1})}{\text{Vol}(\Omega_k)} \leq 1 - \frac{1}{e} \approx 0.63 \quad (65)$$

which may lead to faster convergence than the rate $(1 - 1/r)$ using the center of MVE. However, for a polytope described by a set of linear inequalities, it is more expensive to compute the center of gravity in (65) than to compute the center of MVE [27].

2) *Chebyshev Center*: The Chebyshev center is defined as the center of the largest Euclidean ball that lies inside the polytope Ω . Chebyshev center for a polytope can be efficiently computed by solving a linear program [28]. But Chebyshev center is not affinely invariant to the transformations of coordinates [27]. Thus, a linear mapping of features may lead to an inconsistent weight vector estimation.

3) *Analytic Center*: Given a polytope $\Omega = \{\theta \mid \langle h_i, \theta \rangle + b_i < 0, i = 1, \dots, m\}$, the analytic center is defined as

$$\theta_{ac} = \min_{\theta} - \sum_{i=1}^m \log(b_i - h_i^T \theta). \quad (66)$$

As shown by [40] and [41], the analytic center achieves a good tradeoff in terms of simplicity and practical performance. However, it might not be friendly for analyzing the volume reduction compared with using the center of MVE.

ACKNOWLEDGMENT

The authors would like to thank Tianyu Zhou for his effort in helping organize and collect the data in the user study.

REFERENCES

- [1] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 3, pp. 297–330, 2020.
- [2] M. Kuderer, S. Gulati, and W. Burgard, "Learning driving styles for autonomous vehicles from demonstration," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 2641–2646.
- [3] P. Englert, N. A. Vien, and M. Toussaint, "Inverse KKT: Learning cost functions of manipulation tasks from demonstrations," *Int. J. Robot. Res.*, vol. 36, no. 13/14, pp. 1474–1488, 2017.
- [4] W. Jin, T. D. Murphey, D. Kulić, N. Ezer, and S. Mou, "Learning from sparse demonstrations," *IEEE Trans. Robot.*, to be published, doi: [10.48550/arXiv.2008.02159](https://doi.org/10.48550/arXiv.2008.02159).
- [5] W. Jin, D. Kulić, S. Mou, and S. Hirche, "Inverse optimal control from incomplete trajectory observations," *Int. J. Robot. Res.*, vol. 40, no. 6/7, pp. 848–865, 2021.
- [6] A. Jain, S. Sharma, T. Joachims, and A. Saxena, "Learning preferences for manipulation tasks from online coactive feedback," *Int. J. Robot. Res.*, vol. 34, no. 10, pp. 1296–1313, 2015.
- [7] A. Bajcsy, D. P. Losey, M. K. O'Malley, and A. D. Dragan, "Learning robot objectives from physical human interaction," in *Proc. Conf. Robot Learn.*, vol. 78, 2017, pp. 217–226.
- [8] J. Y. Zhang and A. D. Dragan, "Learning from extrapolated corrections," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 7034–7040.
- [9] D. P. Losey and M. K. O'Malley, "Including uncertainty when learning from human corrections," in *Proc. Conf. Robot Learn.*, 2018, pp. 123–132.
- [10] A. Jain, B. Wojcik, T. Joachims, and A. Saxena, "Learning trajectory preferences for manipulators via iterative improvement," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 575–583.
- [11] P. Shivaswamy and T. Joachims, "Online structured prediction via coactive learning," in *Proc. IEEE Int. Conf. Mach. Learn.*, 2012, pp. 59–66.
- [12] P. Moylan and B. Anderson, "Nonlinear regulator theory and an inverse optimal control problem," *IEEE Trans. Autom. Control*, vol. 18, no. 5, pp. 460–465, Oct. 1973.

- [13] A.-S. Puydupin-Jamin, M. Johnson, and T. Bretl, "A convex approach to inverse optimal control and its application to modeling human locomotion," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2012, pp. 531–536.
- [14] W. Jin, D. Kulić, J. F.-S. Lin, S. Mou, and S. Hirche, "Inverse optimal control for multiphase cost functions," *IEEE Trans. Robot.*, vol. 35, no. 6, pp. 1387–1398, Dec. 2019.
- [15] A. Y. Ng et al., "Algorithms for inverse reinforcement learning," in *Proc. IEEE Int. Conf. Mach. Learn.*, vol. 1, 2000, pp. 663–670.
- [16] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. Conf. Assoc. Advance. Artif. Intell.*, 2008, pp. 1433–1438.
- [17] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, "Maximum margin planning," in *Proc. IEEE Int. Conf. Mach. Learn.*, 2006, pp. 729–736.
- [18] W. Jin, Z. Wang, Z. Yang, and S. Mou, "Pontryagin differentiable programming: An end-to-end learning and control framework," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 7979–7992.
- [19] Z. Liang, W. Jin, and S. Mou, "An Iterative Method for Inverse Optimal Control," in *Proc. Asian Control Conf.*, 2022, pp. 959–964.
- [20] C. Moro, G. Nejat, and A. Mihailidis, "Learning and personalizing socially assistive robot behaviors to aid with activities of daily living," in *Proc. Int. Conf. Mach. Learn.*, vol. 7, no. 2, 2018, pp. 1–25.
- [21] A. Bobu, A. Bajcsy, J. F. Fisac, and A. D. Dragan, "Learning under misspecified objective spaces," in *Proc. Conf. Robot Learn.*, 2018, pp. 796–805.
- [22] A. D. Dragan, K. Muelling, J. A. Bagnell, and S. S. Srinivasa, "Movement primitives via optimization," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 2339–2346.
- [23] D. P. Losey and M. K. O'Malley, "Learning the correct robot trajectory in real-time from physical human interactions," *ACM Trans. Hum.-Robot Interact.*, vol. 9, no. 1, pp. 1–19, 2019.
- [24] R. Tedrake, "Underactuated robotics: Algorithms for walking, running, swimming, flying, and manipulation," *From Course Notes for MIT 6.832*, 2022.
- [25] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *Proc. Int. Conf. Informat. Control, Automat. Robot.*, 2004, pp. 222–229.
- [26] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Math. Program. Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [27] S. Boyd and L. Vandenberghe, "Localization and cutting-plane methods," From Stanford EE 364b Lecture Notes, 2007.
- [28] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [29] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *J. Mach. Learn. Res.*, vol. 17, no. 83, pp. 1–5, 2016.
- [30] S. P. Tarasov, "The method of inscribed ellipsoids," *Sov. Math.-Doklady*, vol. 37, no. 1, 1988, pp. 226–230.
- [31] J. B. Kuipers *Quaternions and Rotation Sequences*, vol. 66. Princeton, NJ, USA: Princeton Univ. Press, 1999.
- [32] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor UAV on SE(3)," in *Proc. IEEE Conf. Decis. Control*, 2010, pp. 5420–5425.
- [33] B. Amos and J. Z. Kolter, "OptNet: Differentiable optimization as a layer in neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 136–145.
- [34] W. Jin, S. Mou, and G. J. Pappas, "Safe Pontryagin differentiable programming," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, pp. 16034–16050.
- [35] W. Jin, A. Aydinoglu, M. Halm, and M. Posa, "Learning linear complementarity systems," in *Proc. Learn. Dyn. Control Conf.*, 2022, pp. 1137–1149.
- [36] D. J. Newman, "Location of the maximum on unimodal surfaces," *J. ACM*, vol. 12, no. 3, pp. 395–398, 1965.
- [37] J. Elzinga and T. G. Moore, "A central cutting plane algorithm for the convex programming problem," *Math. Program.*, vol. 8, no. 1, pp. 134–145, 1975.
- [38] J.-L. Goffin and J.-P. Vial, "On the computation of weighted analytic centers and dual ellipsoids with the projective algorithm," *Math. Program.*, vol. 60, no. 1–3, pp. 81–92, 1993.
- [39] B. Grünbaum et al., "Partitions of mass-distributions and of convex bodies by hyperplanes," *Pacific J. Math.*, vol. 10, no. 4, pp. 1257–1261, 1960.
- [40] D. S. Atkinson and P. M. Vaidya, "A cutting plane algorithm for convex programming that uses analytic centers," *Math. Program.*, vol. 69, no. 1–3, pp. 1–43, 1995.
- [41] Y. Nesterov, "Cutting plane algorithms from analytic centers: Efficiency estimates," *Math. Program.*, vol. 69, no. 1, pp. 149–176, 1995.



Wanxin Jin received the Ph.D. degree in autonomy and control from Purdue University, West Lafayette, IN, USA, in 2021.

He is currently a Postdoctoral Researcher with the GRASP Laboratory, University of Pennsylvania, Philadelphia, PA, USA. From 2016 to 2017, he was a Research Assistant with Technical University Munich, Munich, Germany. His research interests include robotics, control, machine learning, and optimization, with emphasis on learning, planning, and control of robots as they interact with the world and humans.



Todd D. Murphey received the B.S. degree in mathematics from the University of Arizona, Tucson, AZ, USA, in 1997, and the Ph.D. degree in control and dynamical systems from the California Institute of Technology, Pasadena, CA, USA, in 2002.


He is currently a Professor of mechanical engineering with Northwestern University, Evanston, IL, USA. His laboratory is part of the Center for Robotics and Biosystems. His research interests include robotics, control, machine learning in physical systems, and computational neuroscience.



Zehui Lu received the M.S. degree in mechanical engineering from the University of Michigan, Ann Arbor, MI, USA, in 2019. He is currently working toward the Ph.D. degree in aeronautics and astronautics engineering with Purdue University, West Lafayette, IN, USA.

Between January 2020 and July 2020, he was a Research Engineer with UM Ford Center for Autonomous Vehicles (FCAV) and ROAHM LAB, all with the University of Michigan. His current research interests include autonomy, control, and optimization,

and their applications in robotics and multiagent systems.



Shaoshuai Mou (Member, IEEE) received the Ph.D. degree in electrical engineering from Yale University, Yale University New Haven, CT, USA, in 2014.

He is currently an Associate Professor with the School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN, USA. Before joining Purdue, he worked as a Postdoc Researcher with MIT for a year after that. His research interests include multi-agent system, control and learning, robotics control, human–robot teaming, resilient autonomy, and also experimental research involving autonomous air and ground vehicles. He codirects Purdue University's Center for Innovation in Control, Optimization and Networks (ICON), which aims to integrate classical theories in control/optimization/networks with recent advances in machine learning/AI/data science to address fundamental challenges in autonomous and connected systems.