

Jointly Learning Cost and Constraints from Demonstrations for Safe Trajectory Generation

Shivam Chaubey, Francesco Verdoja, Ville Kyrki

Abstract— Learning from Demonstration (LfD) allows robots to mimic human actions. However, these methods do not model constraints crucial to ensure safety of the learned skill. Moreover, even when explicitly modelling constraints, they rely on the assumption of a known cost function, which limits their practical usability for task with unknown cost. In this work we propose a two-step optimization process that allow to estimate cost and constraints by decoupling the learning of cost functions from the identification of unknown constraints within the demonstrated trajectories. Initially, we identify the cost function by isolating the effect of constraints on parts of the demonstrations. Subsequently, a constraint leaning method is used to identify the unknown constraints. Our approach is validated both on simulated trajectories and a real robotic manipulation task. Our experiments show the impact that incorrect cost estimation has on the learned constraints and illustrate how the proposed method is able to infer unknown constraints, such as obstacles, from demonstrated trajectories without any initial knowledge of the cost.

I. INTRODUCTION

Current robots are limited in their ability to perform complex tasks due to the difficulty of programmatically describing the desired behaviors. The main approach to address this challenge is Learning from Demonstration (LfD), a method in which robots learn to perform tasks effectively by imitating expert demonstrations, bypassing the need for explicit programming of complex behaviors [1].

However, learning generalizable policies with LfD methods can be challenging as these methods often do not explicitly guarantee constraint satisfaction. This is a significant issue, since constraints such as workspace obstacles or velocity limits in part of the workspace are used to represent essential safety requirements and performance specifications, and without guarantees of constraint satisfaction, the safety and efficacy of robotic systems cannot be ensured. Even outside LfD, the importance of constraints is visible in recent developments in Inverse Reinforcement Learning (IRL) where work has been devoted to methods to prioritize learning constraints over the traditional emphasis on reward functions.

One main limitation of existing methods for constraint learning is the assumption that the reward function (or cost) is known, as observed in [2]–[4]. However, incorrect user-defined cost can cause constraints to be erroneously identified. Therefore, identifying both the cost function and




Fig. 1: Starting from a human demonstration, we propose a novel method to jointly learn both cost and constraints to allow the robot to replicate the task. Here, a user demonstrates dropping a ball from a cup to a target; the retrieved constraint limits the cup tilt (θ) when the cup position (y) is not over target, as shown in red in the graph.

constraints simultaneously is necessary to learn skills while ensuring safety of the system.

In this paper, we introduce a method to jointly learn both cost function and constraints from human demonstrations. We achieve this with a novel method to estimate cost under the assumption that constraints are affecting parts of the demonstration and subsequently utilizing the retrieved cost to identify the constraints. The learned constraints can be applied to generate new trajectories. These trajectories not only align with the observed human demonstrations but also prioritize safety, ensuring that the robot’s actions are both effective and secure. The main contributions of this work are:

- A method jointly learning optimal cost and constraints for skills generated through human demonstrations in the presence of unknown constraints and cost.
- An experimental validation of our approach both in simulation and with a real-world system.
- Code for our method and experiments to facilitate reproducibility and further research in this area¹.

Our paper is structured as follows: In Sec. II, we discuss recent developments in the field of task learning and safety assurance, examining related work, and providing context

¹Source code for method and experimental setting of this paper is found at version.aalto.fi/gitlab/chaubey1/jointly-learning-cost-and-constraints.

This work was supported by Business Finland, decision 9249/31/2021. S. Chaubey, F. Verdoja, and V. Kyrki are with the School of Electrical Engineering, Aalto University, Espoo, Finland. {firstname.lastname}@aalto.fi

for our approach. Sec. III presents the problem formulation addressed by our method. Sec. IV, we introduce our proposed method for extracting cost functions and constraints, describing the techniques and algorithms used. Finally, Sec. V presents an evaluation of our approach, including its application in real-world scenarios, and conclusions are drawn in Sec. VI.

II. RELATED WORK

Recent studies [5]–[7], have identified three main categories within LfD: policy, plan, and cost or reward learning, based on the learning outcomes. Although these approaches offer powerful ways to train agents, they often struggle to ensure that important constraints are met [8].

Most of the methods proposed to infer constraints require prior knowledge of the cost function [2]–[4], [9], [10]. The method proposed in [3], [4] represents an initial step in learning arbitrary constraints. However, it assumes the reward function, which is required to calculate the policy, is known or calculated using an IRL method. In another approach [9], constraints are inferred by comparing demonstrated trajectories with those of higher performance from a nominal model. These constraints are found within the feature space of lower-cost trajectories not followed by the demonstrator, implying the presence of previously unidentified constraints. However, this approach relies on an already known cost function and lacks a mechanism for determining an optimal cost function.

The method proposed in [2] has limitations, including the extraction of only a single constraint and exclusive focus on soft constraints, which may not be suitable for safety-critical applications. The work by [10] recovers constraints provided reward function by introducing Bayesian Network Inverse Reinforcement Learning (BN-IRL) and its extension, Constraint-based BN-IRL (CBN-IRL). While BN-IRL organizes decision-making without strict constraints, CBN-IRL adds hard constraints but struggles with accurately identifying constrained tasks due to its flexible assignment variable. In the approach presented by [11], constraints are primarily extracted by applying Karush–Kuhn–Tucker (KKT) conditions, considering parametric uncertainty in cost up to a certain degree to handle locally-optimal demonstrations.

Our method contributes to the field of cost and constraint learning, addressing the limitations of previous approaches. Unlike prior methods that concentrate solely on soft constraints or require prior knowledge of the cost function, our approach segments demonstrations to isolate the impact of constraints on policy learning. This segmentation facilitates a more nuanced understanding of how policy influences decision-making without constraints, thereby allowing for a more accurate and robust inference of constraints directly from the learned policy and the demonstrations themselves.

III. PROBLEM FORMULATION

A. Optimal trajectory problem

Let us consider a set of optimal human demonstrations $\mathcal{Z} = (\zeta_1, \dots, \zeta_L)$, where each trajectory $\zeta_l =$

(x^l, u^l) consists of a sequence of states x^l and controls u^l . Within each trajectory, individual points in time are denoted $(x^l(k), u^l(k))$. For simplicity, we momentarily omit the superscript l .

Each demonstration is assumed to be the solution to constrained quadratic optimization problem

$$\min_{x(k), u(k)} J \quad (1a)$$

such that (s.t.)

$$\underbrace{\begin{aligned} x(k+1) - f(x(k), u(k)) &= 0 & \forall k \in [1, N-1] \\ x(k) - x_s &= 0 & k = 1 \\ x(k) - x_g &= 0 & k = N \end{aligned}}_{\text{Equality constraints}} \quad (1b)$$

$$\underbrace{\begin{aligned} g(x(k)) &\leq 0 & \forall k \in [1, N] \\ h(u(k)) &\leq 0 & \forall k \in [1, N-1] \end{aligned}}_{\text{Inequality constraints}} \quad (1c)$$

where J is the standard quadratic cost with parameters Q and R ,

$$J = \frac{1}{2} \left(\sum_{k=1}^N x(k)^\top Q x(k) + \sum_{k=1}^{N-1} u(k)^\top R u(k) \right), \quad (2)$$

and $x(k) \in \mathbb{R}^{n_x}$ represents the state vector, while $u(k) \in \mathbb{R}^{n_u}$ denotes the control vector. The parameter Q serves as the weight for penalizing states, contributing to the generation of the optimal trajectory. On the other hand, R is the weight applied to minimize control efforts. The inequality constraint, represented by $g(x(k)) : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_{c_x}}$ for the states and $h(u(k)) : \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_{c_u}}$ for the control, encloses all the known/unknown inequality constraints, respectively. Here, n_{c_x} and n_{c_u} denote the number of inequality constraints for states and control, respectively. The constraint $x(k+1) - f(x(k), u(k)) = 0$ represents known system dynamics. Lastly, conditions $x(1) - x_s = 0$ and $x(N) - x_g = 0$ correspond to the initial and final states, respectively.

B. Constraints formulation

The constraint set defined in (1b) and (1c) includes both equality and inequality constraints. The equality constraints are assumed to be known and account for the system dynamics governed by $f(x(k), u(k))$, and initial and final states, *i.e.*, x_s and x_g . These constraints ensure that the system starts and ends in specified states and adheres to the predefined dynamics throughout its operation. The inequality constraints for state $g(x(k))$ and control $h(u(k))$ can be known or unknown. These constraints limit the states and controls to prevent them from exceeding certain bounds, thereby ensuring safety and feasibility within the system's operational environment. The constraints are considered to be vector valued, and thus, an arbitrary number of constraints can be used.

Moreover, while not explicitly specified in (1c), inequality constraints can be divided into inclusive (■) and exclusive (□). Inclusive constraints ensure that the trajectory remains within

specified bounds, promoting safety and task compliance, while exclusive constraints are designed to keep the trajectory outside of certain bounds, thus avoiding unsafe regions. We also assume inequality constraints to be describable as axis-aligned bounds for inclusive constraints and axis-aligned shapes for exclusive constraints, noting that axis-aligned constraints can represent irregular shapes when utilized in large numbers. In the following, the vector $w(k) \in \mathbb{R}^{n_w}$ may represent either the state vector $x(k) \in \mathbb{R}^{n_x}$ or the action vector $u(k) \in \mathbb{R}^{n_u}$, depending on the problem context.

Given lower-bounds $w_{\square}^{lb}, w_{\square}^{ub} \in \mathbb{R}^{n_w}$ and upper-bounds $w_{\square}^{ub}, w_{\square}^{lb} \in \mathbb{R}^{n_w}$, inclusive constraints can be expressed as

$$w_{\square}^{lb} \leq w_{\square}(k) \leq w_{\square}^{ub} \quad \forall k \in [1, N], \quad (3)$$

while exclusive constraints can be expressed as

$$\bigvee_{i=1}^{n_w} w_{\square}(k, i) - w_{\square}^{lb}(i) \leq 0 \vee w_{\square}(k, i) + w_{\square}^{ub}(i) \leq 0 \quad \forall k \in [1, N], \quad (4)$$

where $i \in [1, n_w]$ refers to the i -th dimension out of the n_w state or action dimensions. The interpretation of (4) is that, in order for the demonstration to not violate an exclusive constraint, each time-step k needs to be outside the constraint's bounds on at least one dimension.

Please note that, while inclusive constraints are convex, exclusive constraints are not, and therefore they need to be relaxed when included in convex optimization. For this reason, in this work, while we present a method able to retrieve unknown constraints of both types, in the experiments we focus on exclusive constraints only.

C. Cost and constraints estimation problem

In the setting of this paper, some of the inequality constraints as well as the cost parameters are unknown. A significant challenge arises from the simultaneous determination of the cost function and multiple constraints, leading to the issue of unidentifiability: numerous configurations of cost and constraint functions may produce identical trajectories [12].

Therefore, in this work we propose a method to jointly identify both cost parameters (Q and R) as well as an arbitrary number of unknown inequality constraints affecting the demonstrations. Our method also addresses the relaxation of exclusive constraints.

IV. METHOD

To jointly estimate both cost parameters and constraints, our approach is twofold: first, we determine the cost associated with the demonstration under the assumption that the unknown constraints are active only at certain time steps. This step isolates the effect of the constraints on the cost function and captures the underlying cost structure that drives the system's behavior in the absence of external unknown constraints. Second, we infer the unknown constraints based on the previously estimated cost. This step relies on the understanding that deviations of the trajectory from the

unconstrained behavior can be attributed to the unknown constraints.

We will describe our approach for estimating the cost parameters in Sec. IV-B, followed by the estimation of the constraints in Sec. IV-C. But before, in Sec. IV-A we present a reformulation of the trajectory optimization problem presented in Sec. III-A according to the KKT conditions which is used in multiple parts of our method.

A. KKT formulation

The optimization problem presented in Sec. 2 is convex, differentiable, and satisfies constraint qualification criteria [13]. This enables us to utilize the first order KKT conditions [14], [15] to formulate optimization problems within a unified framework, capable of handling both equality and inequality constraints alongside the objective function. Given our particular interest in the inverse problem of 2, where solutions (trajectory ζ_t) are known and unknown costs and constraints need to be learned, we will adapt the KKT conditions, referred to as the Inverse KKT (IKKT) conditions, to infer these unknowns.

In the IKKT formulation, our goal is to derive unknown costs and constraints from given demonstration while considering some known constraints. We introduce specific vectors within our IKKT framework that represent various settings for extracting costs and constraints. Specifically, we have binary switching vectors for state ($\mathbf{v}_x^p, \mathbf{v}_x^d, \mathbf{v}_x^s$) and control ($\mathbf{v}_u^p, \mathbf{v}_u^d, \mathbf{v}_u^s$) for different conditions. In these vectors, an element of '1' signifies that the dimension is included in the optimization formulation, while a '0' means it is excluded. We define $\mathbf{v}_x^d = [\mathbf{v}_x^s, \mathbf{v}_x^p]^\top$ and $\mathbf{v}_u^d = [\mathbf{v}_u^s, \mathbf{v}_u^p]^\top$.

This approach in the IKKT formulation allows us to learn about costs and constraints without needing to alter the formulation. We reinterpret the KKT stationary condition; rather than setting it to zero, we represent this condition with a Lagrange derivative expression that needs to be minimized. This modification allows us to address cases where achieving the exact condition might not be possible or optimal, focusing on minimizing the deviation from the ideal stationary state. The IKKT formulation of our problem is then expressed as follows:

Primal Feasibility (p.f.):

$$\mathbf{v}_x^p \odot g(x(k)) \leq 0, \quad \mathbf{v}_u^p \odot h(u(k)) \leq 0 \quad (5)$$

Dual Feasibility (d.f.):

$$\mathbf{v}_x^d \odot \lambda_x(k) \geq 0, \quad \mathbf{v}_u^d \odot \lambda_u(k) \geq 0 \quad (6)$$

Complementary Slackness (c.s.):

$$\begin{aligned} \mathbf{v}_x^d \odot \lambda_x(k) \odot g(x(k)) &= 0, \\ \mathbf{v}_u^d \odot \lambda_u(k) \odot h(u(k)) &= 0 \end{aligned} \quad (7)$$

Lagrange Derivative (l.d.):

$$\begin{aligned} \left. \frac{\partial \mathcal{L}}{\partial x(k)} \right|_{k=1} &= Qx(k) + \mathbf{v}_x^s \odot \left[\lambda_x(k)^\top \frac{\partial g(x(k))}{\partial x(k)} \right]^\top \\ &\quad - \left[\vartheta_d(k)^\top \frac{\partial f(x(k), u(k))}{\partial x(k)} \right]^\top + \vartheta_s \end{aligned}$$

$$\begin{aligned}
\left. \frac{\partial \mathcal{L}}{\partial u(k)} \right|_{k=1} &= Ru(k) + \mathbf{v}_u^s \odot \left[\lambda_u(k)^\top \frac{\partial h(u(k))}{\partial u(k)} \right]^\top \\
&\quad - \left[\vartheta_d(k)^\top \frac{\partial f(x(k), u(k))}{\partial u(k)} \right]^\top \\
\left. \frac{\partial \mathcal{L}}{\partial x(k)} \right|_{k \in (1, N)} &= Qx(k) + \mathbf{v}_x^s \odot \left[\lambda_x(k)^\top \frac{\partial g(x(k))}{\partial x(k)} \right]^\top \\
&\quad + \vartheta_d(k-1) - \left[\vartheta_d(k)^\top \frac{\partial f(x(k), u(k))}{\partial x(k)} \right]^\top \quad (8) \\
\left. \frac{\partial \mathcal{L}}{\partial u(k)} \right|_{k \in (1, N)} &= Ru(k) + \mathbf{v}_u^s \odot \left[\lambda_u(k)^\top \frac{\partial h(u(k))}{\partial u(k)} \right]^\top \\
&\quad - \left[\vartheta_d(k)^\top \frac{\partial f(x(k), u(k))}{\partial u(k)} \right]^\top \\
\left. \frac{\partial \mathcal{L}}{\partial x(k)} \right|_{k=N} &= Qx(k) + \mathbf{v}_x^s \odot \left[\lambda_x(k)^\top \frac{\partial g(x(k))}{\partial x(k)} \right]^\top \\
&\quad + \vartheta_d(k-1) + \vartheta_g
\end{aligned}$$

Here, $\lambda_x(k) \in \mathbb{R}_{\geq 0}^{n_{c_x}}$ and $\lambda_u(k) \in \mathbb{R}_{\geq 0}^{n_{c_u}}$ represents the vector of Lagrange multipliers associated with the inequality constraints in state and control respectively. $\vartheta_d(k) \in \mathbb{R}^{n_x}$, $k \in [1, N]$ denotes the vector of Lagrange multipliers linked to system dynamics. $\vartheta_s \in \mathbb{R}^{n_x}$ and $\vartheta_g \in \mathbb{R}^{n_x}$ denotes the vector of Lagrange multipliers linked to initial point $k=1$ and final point $k=N$, respectively.

B. Cost extraction

We propose a two-step method to determine the cost parameters: initially, demonstrations are locally partitioned into segments based on whether the unknown constraints are active; subsequently, the inactive segments are utilized to learn the cost.

In cost learning setup, we set the vectors \mathbf{v}_x^p and \mathbf{v}_u^p to zeros. This step is taken to overlook the primal feasibility condition, as the demonstrations already meet this criterion for both known and unknown constraints. To understand the policy in the absence of unknown constraints, we introduce the vectors \mathbf{v}_x^s and \mathbf{v}_u^s . The elements within the vectors \mathbf{v}_x^s and \mathbf{v}_u^s are set to ‘0’ for unknown constraints and to ‘1’ for known constraints.

1) *Demonstration segmentation*: We want to partition a trajectory ζ into active segments $\zeta^\ominus \subset \zeta$ —i.e., where the unknown constraints are active—and inactive segments $\zeta^\oplus \subset \zeta$ —i.e., where the unknown constraints are inactive. Let us define the three-point segment of trajectory. Three points were chosen as the minimum required for solving the discrete KKT formulation. ζ starting at time-step j as $\zeta^j = \{\zeta(j), \zeta(j+1), \zeta(j+2)\}$. We propose a method to assign each segment ζ^j to either ζ^\ominus or ζ^\oplus leveraging the intuition that inactive segments will behave according to the optimal cost parameters and known constraints, while active segment—due to influence of unknown constraints—will showcase a noticeable variation in the trace of the normalized cost Ω^j , obtained by normalizing Q^j with regards to R^j . Therefore, we propose to estimate cost parameter for

each segment ζ^j with $j \in [1, N-2]$ and then perform an outlier detection step to identify the active indices.

Therefore, given a trajectory ζ and an index j , we solve

$$\min_{\substack{Q^j, R^j, \tilde{\lambda}_x^j, \tilde{\lambda}_u^j \\ \tilde{\vartheta}_d^j, \tilde{\vartheta}_s^j, \tilde{\vartheta}_g^j}} \|\mathcal{J}^j\|_2 \quad (9)$$

s.t. (6) and (7) satisfy, $Q^j \geq 0$, $\text{Tr}(Q^j) = 1$ and $R^j > 0$, and where

$$\begin{aligned}
\mathcal{J}^j &= \left[\frac{\partial \mathcal{L}}{\partial x(j)}, \frac{\partial \mathcal{L}}{\partial u(j)}, \frac{\partial \mathcal{L}}{\partial x(j+1)}, \frac{\partial \mathcal{L}}{\partial u(j+1)}, \frac{\partial \mathcal{L}}{\partial x(j+2)} \right]^\top \\
\tilde{\lambda}_x^j &= [\lambda_x(j), \lambda_x(j+1), \lambda_x(j+2)] \\
\tilde{\lambda}_u^j &= [\lambda_u(j), \lambda_u(j+1)] \\
\tilde{\vartheta}_d^j &= [\vartheta_d(j), \vartheta_d(j+1)] \\
\tilde{\vartheta}_s^j &= 0 \Leftarrow j \neq 1, \quad \tilde{\vartheta}_g^j = 0 \Leftarrow j \neq N-2.
\end{aligned}$$

Equation (9) gives Q^j , R^j , i.e., an estimate of the cost parameters for that segment. To detect outliers, the normalized cost $\text{Tr}(\Omega^j)$ is calculated for each segment at time-step j and the Generalized Extreme Studentized Deviate (GESD) [16] outlier procedure is performed on list of $\text{Tr}(\Omega^j) \forall j \in [1, N]$. We then include all ζ^j identified as outlier into ζ^\ominus and set $\zeta^\oplus = \zeta \setminus \zeta^\ominus$.

2) *Cost estimation from inactive segments*: To determine the unknown cost weights Q and R , we consider all indices j_l^\oplus of inactive segments ζ_l^\oplus within each trajectory ζ_l . This approach addresses the issue of unidentifiability by focusing exclusively on time steps that fully inform the demonstrated policy. It is important to note that unknown constraints only prompt behavioral changes when active. To learn a time-invariant cost that aligns with the entire set of demonstrations \mathcal{Z} , we solve

$$\min_{\substack{Q, R, \tilde{\lambda}_x, \tilde{\lambda}_u \\ \tilde{\vartheta}_d, \tilde{\vartheta}_s, \tilde{\vartheta}_g}} \|\mathcal{J}_1^\oplus, \mathcal{J}_2^\oplus, \dots, \mathcal{J}_L^\oplus\|_2 \quad (10)$$

s.t. (6) and (7) satisfy, $Q^j \geq 0$, $\text{Tr}(Q) = 1$ and $R^j > 0$, and where

$$\begin{aligned}
\mathcal{J}_l^\oplus &= \left[\frac{\partial \mathcal{L}_l}{\partial x(q_l)}, \frac{\partial \mathcal{L}_l}{\partial u(q_l)}, \dots \right] \\
\tilde{\lambda}_x &= [\lambda_x^l(q_l), \dots], \quad \tilde{\lambda}_u = [\lambda_u^l(q_l), \dots], \quad \tilde{\vartheta}_d = [\vartheta_d^l(q_l), \dots] \\
\tilde{\vartheta}_s &= [\vartheta_s^l(q_l), \dots], \quad \tilde{\vartheta}_s^l(q_l) = 0 \Leftarrow q_l \neq 1 \\
\tilde{\vartheta}_g &= [\vartheta_g^l(q_l), \dots], \quad \tilde{\vartheta}_g^l(q_l) = 0 \Leftarrow q_l \neq N_l \\
\forall q_l &\in j_l^\oplus
\end{aligned}$$

This optimization problem is solved once, incorporating all inactive segments from all demonstrated trajectories to estimate the cost weights Q and R . The computational cost of the problem is similar to that of a least squares problem. In the next section, we use these cost weights to identify the unknown constraints.

C. Constraint extraction

Our proposed method for extracting constraints from observed demonstrations comprises three essential steps: Firstly, we start with the formulation of inequality constraints. Secondly, we introduce additional complementary conditions to effectively manage both inclusive constraints and unobserved bounds. Finally, we present a constraint learning algorithm that utilizes the initial two steps.

1) *Inequality constraints representation:* As detailed in Sec. III-B, inequality constraints can be either inclusive or exclusive. We represent inclusive constraints as

$$D_{\blacksquare} w_{\blacksquare}(k) - F_{\blacksquare} \odot \bar{w}_{\blacksquare} \leq 0 \quad (11)$$

where

$$D_{\blacksquare} = \begin{pmatrix} -\mathbf{I}_{n_w \times n_w} \\ \mathbf{I}_{n_w \times n_w} \end{pmatrix}, F_{\blacksquare} = \begin{pmatrix} -\mathbf{1}_{n_w \times 1} \\ \mathbf{1}_{n_w \times 1} \end{pmatrix}, \bar{w}_{\blacksquare} = \begin{pmatrix} w_{\blacksquare}^{lb} \\ w_{\blacksquare}^{ub} \end{pmatrix}.$$

Exclusive constraints are non-convex and need to be relaxed when included in convex optimization. To relax the formulation presented in (4) into a convex set, we follow the method proposed by Schouwenaars *et al.* [17], please refer to that work for a complete explanation. Let us introduce a binary variable vector $p(k) \in \{0, 1\}^{2n_w}$. Each exclusive constraint can be represented as:

$$D_{\square} w_{\square}(k) - F_{\square} \odot \bar{w}_{\square} - M(1 - p(k)) \leq 0 \quad (12)$$

and

$$\|p(k)\|_1 \geq 1 \quad (13)$$

where M denotes a large positive number and

$$D_{\square} = \begin{pmatrix} \mathbf{I}_{n_w \times n_w} \\ -\mathbf{I}_{n_w \times n_w} \end{pmatrix}, F_{\square} = \begin{pmatrix} \mathbf{1}_{n_w \times 1} \\ -\mathbf{1}_{n_w \times 1} \end{pmatrix}, \bar{w}_{\square} = \begin{pmatrix} w_{\square}^{lb} \\ w_{\square}^{ub} \end{pmatrix}.$$

Equation (12) ensures compliance with the original constraints, distinguishing between active and inactive constraints, while (13) ensures that at least one constraint is always active. This relaxation allow exclusive constraints to be used in an Mixed-integer linear programming (MILP) optimization problem whose sub-problems are convex.

2) *Complementary exclusive constraints:* We introduce two additional complementary exclusive conditions. The first condition helps to identify bounds that remain inactive throughout the demonstration due to the trajectory discretization. This is a known problem with discrete state-spaces not solved in previous works [11]. This complementary constraint guarantees that any unknown inactive bound, orthogonal to an active constraint, is positioned at or beyond the active point orthogonally. This effectively defines the limits of movement in directions not explicitly restricted by the observed demonstration. For each exclusive constraint, let us define

$$\begin{pmatrix} e_{\square}^{lb} \\ e_{\square}^{ub} \end{pmatrix} = D_{\square} w_{\square}(k) - F_{\square} \odot \bar{w}_{\square}, \quad (14)$$

then, we construct the following complementary constraint:

$$\begin{aligned} \lambda_{\square}^{lb}(r) e_{\square}^{lb}(t) &\geq \epsilon, \lambda_{\square}^{lb}(r) e_{\square}^{ub}(t) \geq \epsilon, \\ \lambda_{\square}^{ub}(r) e_{\square}^{lb}(t) &\geq \epsilon, \lambda_{\square}^{ub}(r) e_{\square}^{ub}(t) \geq \epsilon \\ \forall r, t \in [1, n_w], r &\neq t \end{aligned} \quad (15)$$

In these equations, $\lambda_{\square}^{lb}(r), \lambda_{\square}^{ub}(r) \in \lambda_w$ denote Lagrange multipliers for the lower and upper bounds of the point w_{\square} in dimension r , respectively, and λ_w is either λ_x or λ_u depending on the context. Similarly, $e_{\square}^{lb}(t)$ and $e_{\square}^{ub}(t)$ represent the lower and upper bound constraints for dimension t as defined in (14). The equations assert that for all distinct pairs (r, t) , the product of each Lagrange multiplier with its orthogonal constraint must be greater than or equal to a small negative value ϵ . This formulation ensures that if any bound in a dimension is active for a point $w_{\square}(k)$ within the trajectory, non-active bounds in orthogonal dimensions remain within a reasonable margin from that point.

A second additional constraint ensures that Lagrange multipliers are nullified for constraints that are inactive, effectively isolating non-active constraints from the l.d. expression, d.f., and c.s. conditions. Formally, this idea is captured as:

$$(1 - p_r(k)) \odot \lambda_r(k) = 0. \quad (16)$$

The above formulation mandates that for a given dimension r and at time-step k , the value of the Lagrange multiplier $\lambda_r(k) \in \lambda_w$ must be zero if the associated constraints are inactive or non-binding—*i.e.*, $p_r(k) = 0$.

3) *Constraint learning formulation:* Given the estimated cost weights Q and R according to Sec. IV-B, we identify the set of n_b^2 unknown constraints acting on the demonstrations. Unlike to the cost learning approach outlined in Equation (10), which omits unknown constraints, our focus here shifts towards learning those constraints. Therefore, to ensure p.f., we assign a ‘1’ to the elements of vectors \mathbf{v}_x^p and \mathbf{v}_u^p that correspond to unknown constraints, and a ‘0’ for those representing known constraints. We set every element within the vectors $\mathbf{v}_x^s, \mathbf{v}_u^s, \mathbf{v}_x^d,$ and \mathbf{v}_u^d to one. The revised formulation is as follows:

$$\min_{\substack{\tilde{\lambda}_{x,k}, \tilde{\lambda}_{u,k} \\ \tilde{\vartheta}_{d,k}, \tilde{\vartheta}_s, \tilde{\vartheta}_g}} \left\| [\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_L]^{\top} \right\|_2 \quad (17)$$

s.t. (5), (6), (7), (13), (15), and (16) satisfy, and where

$$\begin{aligned} \mathcal{J}_l &= \left[\frac{\partial \mathcal{L}_l}{\partial x^l(k)}, \frac{\partial \mathcal{L}_l}{\partial u^l(k)}, \dots \right], \\ \tilde{\lambda}_x &= [\lambda_x^l(k), \dots], \tilde{\lambda}_u = [\lambda_u^l(k), \dots], \tilde{\vartheta}_d = [\vartheta_d^l(k), \dots], \\ \tilde{\vartheta}_s &= [\vartheta_s^l(k), \dots], \tilde{\vartheta}_s^l = 0 \Leftarrow k \neq 0 \\ \tilde{\vartheta}_g &= [\vartheta_g^l(k), \dots], \tilde{\vartheta}_g^l = 0 \Leftarrow k \neq N_l \\ \forall k \in [1, N_l], \forall l \in [1, L] \end{aligned}$$

The above formulation includes inequality constraints for both states, denoted by $g(x(k))$, and controls, represented as $h(u(k))$. These constraints are detailed according to the definitions provided in equations (11) and (12).

The bilinear variable terms in the model are simplified using the Matlab tools Yalmip [18]. Following this, problems (9) and (10) are solved using semidefinite solvers, while the problem (17) is addressed using MILP optimization solvers.

²The value of n_b can be set higher if the number of unknown constraints is not known, resulting in overlapping constraints or constraints lying within others. The union of these constraints will provide the final constraints.

V. EXPERIMENTS

The purpose of these experiments is to:

- study the importance of learning the cost weight in the context of constraint learning with an unknown policy;
- evaluate the ability of the proposed method to jointly learn both unknown cost and constraints with varying number of demonstrations and unknown constraints;
- demonstrate how the proposed method can be used on real trajectories in a robotic manipulation setting.

First, we will present a series of quantitative experiments from a set of eight simulated environments with varying number of demonstrations and unknown constraints. In simulation, we will present our full method together with an evaluation of the proposed method to extract constraints when faced with user-defined cost with different degrees of error in the cost parametrization. This aims to illustrate the implications of user-selected cost settings on constraint learning processes and the impact of jointly learning cost and constraints. Finally, a real manipulation task with a robotic arm is presented.

A. Simulation experiment setting

For our simulation setting, we consider an obstacle avoidance problem for a four-dimensional free-floating system, assuming unknown cost and obstacles but with known system dynamics, as well as velocity and acceleration constraints.

The initial and final position constraint for each demonstration ζ_l :

$$x^l(1) - x_s^l = 0, \quad x^l(N_l) - x_g^l = 0. \quad (18)$$

System dynamics constraint can be represented as:

$$x^l(k+1) = Ax^l(k) + Bu^l(k)$$

$$A = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0.5 dt^2 & 0 \\ 0 & 0.5 dt^2 \\ dt & 0 \\ 0 & dt \end{bmatrix}, \quad (19)$$

where, $x^l = [x_1^l, x_2^l, x_3^l, x_4^l]^\top$ represents the state vector, while $u^l = [u_1^l, u_2^l]^\top$ denotes the control vector. Specifically, x_1^l and x_2^l correspond to positions along the x -axis and y -axis, respectively, whereas x_3^l and x_4^l represent velocities in the x -axis and y -axis. Control inputs u_1^l and u_2^l are accelerations in the x -axis and y -axis directions, respectively.

In each of the eight simulation scenarios, shown in Fig. 2, we introduce a number of inequality constraints representing obstacles. Specifically, in each scenario we include a different number n_b of unknown axis-aligned exclusive constraints for the positional states x_1^l and x_2^l , aimed at avoiding obstacles, as well as known inclusive constraints for the velocity states and acceleration input x_3^l and x_4^l , which ensure the velocities and accelerations remain within predefined bounds. In each scenario we generate L optimal demonstrations according to a single Q and R per scenario and the system dynamics presented above; each demonstration has a different starting and goal point.

TABLE I: Average results over eight simulated scenarios

σ	Conv. %	RMSE	Prec.	Recall	F1	IoU
-0.010	37.5	0.024	0.667	0.494	0.994	0.494
-0.005	37.5	0.012	0.993	0.750	0.996	0.744
-0.001	50.0	0.002	0.750	0.744	0.996	0.744
0.000	100.0	0.000	0.994	1.000	0.997	0.994
0.001	75.0	0.002	0.869	0.875	0.997	0.869
0.005	37.5	0.012	0.565	0.500	0.994	0.494
0.010	37.5	0.024	0.391	0.250	0.987	0.244
0.020	12.5	0.047	0.391	0.247	0.983	0.242
Learned	100.0	7.13e-13	0.994	1.000	0.997	0.994

For our approach, in the context of cost learning, we define the vectors for states as $\mathbf{v}_x^p = [0, 0, 0, 0]^\top$ and $\mathbf{v}_x^s = [0, 0, 1, 1]^\top$, and for controls as $\mathbf{v}_u^p = [0, 0]^\top$ and $\mathbf{v}_u^s = [1, 1]^\top$. In the case of learning constraints, we set for states $\mathbf{v}_x^p = [1, 1, 0, 0]^\top$ and $\mathbf{v}_x^s = [1, 1, 1, 1]^\top$, and for controls $\mathbf{v}_u^p = [0, 0]^\top$ and $\mathbf{v}_u^s = [1, 1]^\top$.

B. Influence of cost weights on the retrieved constraints

To investigate how cost impacts the constraints extractions, we pair the method used for extracting constraints with eight different user-defined Q , where in each alter two components of the weight matrix Q . This modification is intended to mimic the real-world scenario of manually adjusting the weights, illustrating the practical challenges encountered in the demonstrations.

Initially, for each environment tested, we create distinct sets of cost Q and R , ensuring that $Q \geq 0$, $\text{Tr}(Q) = 1$, and R is strictly positive definite. Subsequently, L demonstrations are conducted using these predetermined Q and R values. To mimic the real-world practice of guessing weight for the demonstrations, we introduce perturbed versions of Q , denoted as \tilde{Q}_i , generated according to:

$$\tilde{Q}_i = \text{diag}([Q(1,1) + \sigma_i, Q(2,2) - \sigma_i, Q(3,3), Q(4,4)]),$$

$$\sigma_i \in \{-0.01, -0.005, -0.001, 0.0, 0.001, 0.005, 0.01, 0.02\}.$$

We then perform constraint extraction according to Sec. IV-C. Moreover, to evaluate our proposed method, we run the full proposed method, performing also cost extraction according to Sec. IV-B together with constraints extraction.

Tab. I presents the average results of these tests over the eight simulation scenario. We categorize the solution to the problem (17) as converged (Conv.) if it identifies a feasible solution within 10 minutes. The Root Mean Square Error (RMSE(y, y^*)) is used to quantify the deviation between known and perturbed costs. The vector $y(Q, R)$ is defined as $\left[\frac{Q(1,1)}{R(1,1)}, \frac{Q(2,2)}{R(2,2)}, \frac{Q(3,3)}{R(1,1)}, \frac{Q(4,4)}{R(2,2)} \right]^\top$, and similarly $y^*(Q, R)$ is formulated for the perturbed weights. This specific configuration is selected because the control input along the x -axis influences the position and velocity in the x -axis, and likewise, the control input along the y -axis affects the position and velocity in the y -axis. All other metrics refer to the area of the retrieved constraints when compared with the ground-truth constraints shown in Fig. 2. The performance of our constraint learning method on ground-truth cost is

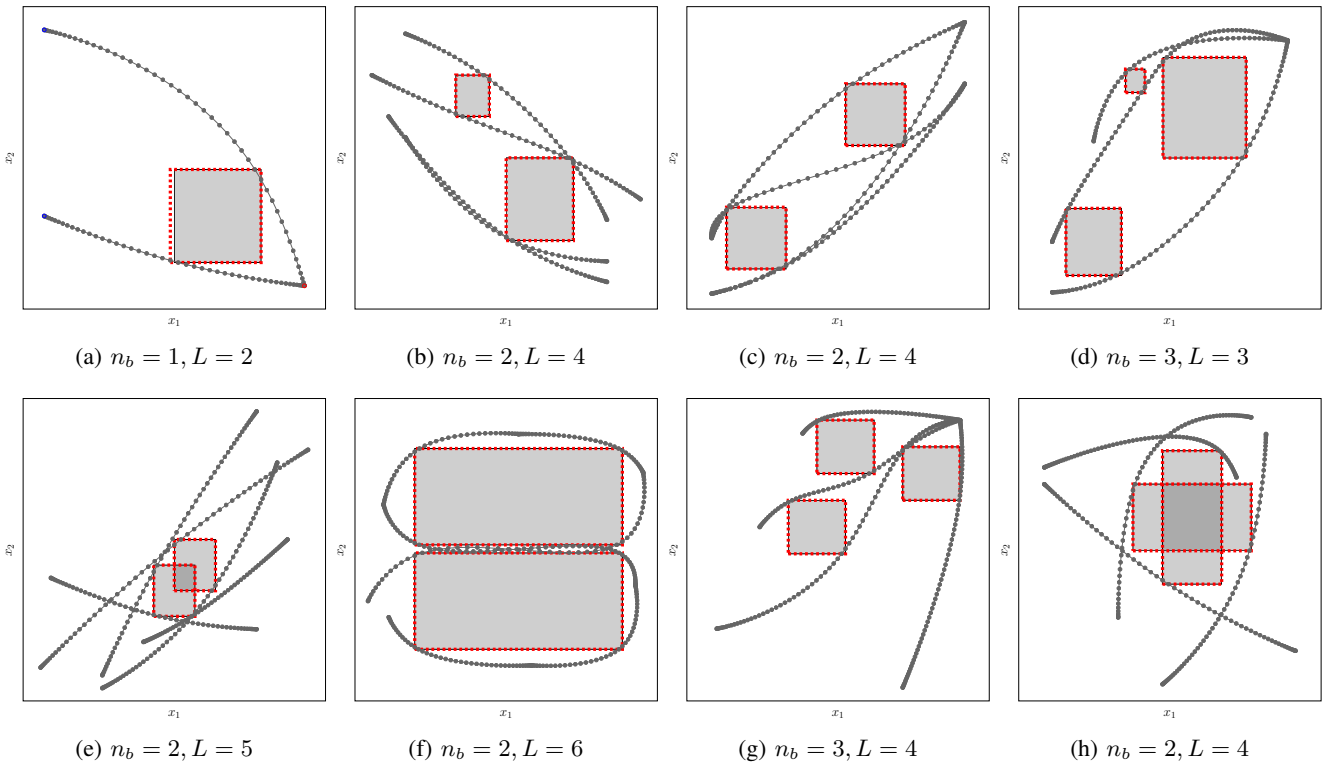


Fig. 2: The environments used in the simulation experiments. Each subcaption indicates the number of unknown constraints (n_b) and demonstrations (L) for that scenario. Legend: \square obstacles, $---$ optimal trajectory, and \square learned constraints.

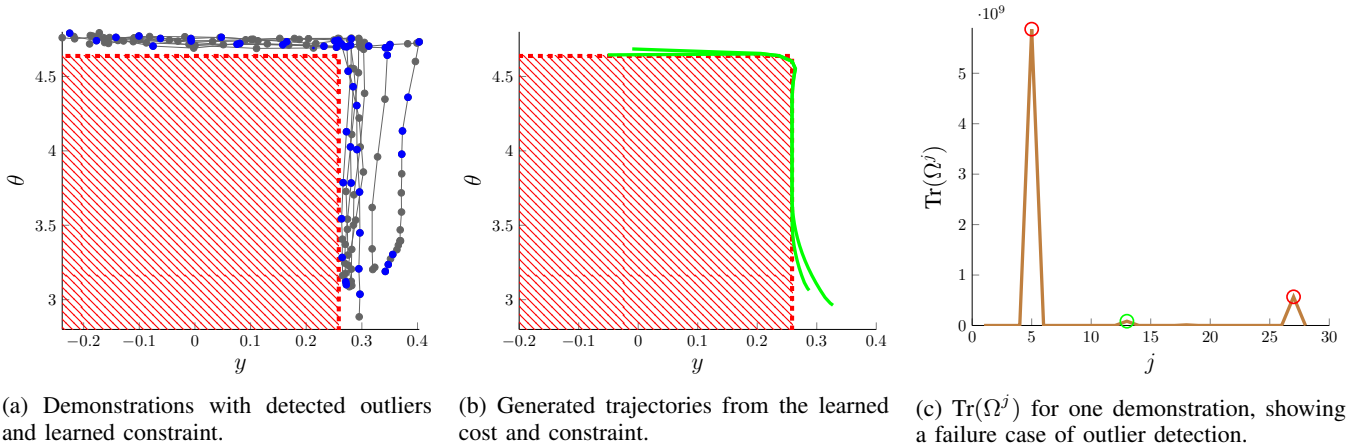


Fig. 3: Real manipulator experiment. Legend: $---$ collected demonstration, \bullet outlier, hatched learned constraint, $---$ generated trajectory, $---$ trace of normalized Q^j , \circ false positive, \circ true positive.

indicated in light gray, while row indicated as “Learned” denotes the results when both our constraints learning and our cost learning methods are used³.

These results reveal the significant influence of perturbed cost parameter Q on both convergence rates and overall solution quality within a constraint learning framework. Notably, when deviations in Q are confined to ± 0.005 ,

³It is worth mentioning that some trajectories intersect the borders. This occurs due to the sampling time used; smaller sampling times generate trajectories closer to the boundary. In practice, constraints can be expanded to enhance safety.

a mere 37.5% of the scenarios reach convergence. Additionally, among these scenarios, a marked reduction in recall to 0.75 is observed. Despite this, consistently high F1 scores in a majority of cases indicate the efficacy of the methodology in maintaining a balance between precision and recall. However, precision, recall, and Intersection over Union (IoU) metrics show a decline as the deviations in Q widen. This pattern highlights the sensitivity of the constraint learning process to slight alterations in specific Q parameters. Therefore, accurately learning the unknown cost is a critical step in our proposed method. It not only facilitates the ac-

quisition of skills from demonstrations but also significantly influences the constraints affecting those skills, highlighting that accurate estimation of cost parameters is pivotal for learning skills under constraints. The last row in Tab. I shows that jointly learning cost and constraints closely mirrors the known cost and constraints performance.

C. Robotic manipulation experiment

In this experiment, a user performs a set of demonstrations hand-guiding a Franka Emika Panda 7-DoF robotic manipulator holding a cup in the end effector. The task is to drop a ball contained in the cup into a bucket. This setup is depicted in Fig. 1. The trajectory followed during these demonstrations is illustrated in Fig. 3a.

In these demonstrations, x_1 corresponds to the displacement along the y-axis, and x_2 represents the angle of rotation, within the global frame. Additionally, x_3 and x_4 are associated with positional and angular velocity, respectively. All other settings follow the free-floating system presented in Sec. V-A.

In Fig. 3a, a high number of detected outlier can be noted. Many of the outliers detected are false positive due to the real demonstrations being locally suboptimal in those spots. An example demonstration where this phenomenon is visible is shown in Fig. 3c, where we can note how the trace of the normalized cost is not discriminative enough to differentiate the true outlier and the suboptimal spots (false positives). However, even under these conditions most of the true outliers are detected by the proposed algorithm and ultimately the unknown constraint is detected correctly. This is partly due to the availability of several demonstrations: the selected inactive points are sufficient to appropriately learn the cost and the retrieved constraint, which is shown with a red bounding box in the Fig. 3a. Finally, the learned cost and constraints are used to replicate the skill as shown by the trajectories in Fig. 3b.

VI. CONCLUSION

In this work, we presented a method to jointly estimate cost and constraints from demonstrations, lifting the reliance of previous methods on known cost function or their focus on soft constraints alone. Moreover, our experiments demonstrate that constraint estimation is sensitive to errors in cost parameters, and a real-world demonstration involving a Franka Panda manipulator performing a pouring task validates that the method can be applied in the presence of noisy real measurements, hinting at the practical usability of the presented method.

However, one of the main limitations of our method was also shown in the real experiment, where the outlier detection step used in the cost extraction was shown to be sensitive to regions of suboptimality in the demonstration, implying that further work should be devoted to improve the robustness of the detection of inactive segments. Moreover, the use of more generic cost functions and constraints could better capture complex tasks.

In conclusion, this work underscores the risk of user-defined cost and the importance of accurately determining the cost function when learning constraints. Our method is a first step in the direction of ensuring the constraints we learn better generalize from the given demonstrations. Moving forward, ensuring reliable learned constraints will be crucial for deploying robots that can safely and efficiently execute learned skills.

REFERENCES

- [1] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Survey: Robot programming by demonstration," Springer, Tech. Rep., 2008.
- [2] A. Gaurav, K. Rezaee, G. Liu, and P. Poupart, "Learning soft constraints from constrained expert demonstrations," *arXiv preprint arXiv:2206.01311*, 2022.
- [3] U. Anwar, S. Malik, A. Aghasi, and A. Ahmed, "Inverse constrained reinforcement learning," *CoRR*, vol. abs/2011.09999, 2020. [Online]. Available: <https://arxiv.org/abs/2011.09999>
- [4] D. R. R. Scobee and S. S. Sastry, "Maximum likelihood constraint inference for inverse reinforcement learning," *CoRR*, vol. abs/1909.05477, 2019. [Online]. Available: <http://arxiv.org/abs/1909.05477>
- [5] F. Ding and Y. Xue, "X-men: guaranteed xor-maximum entropy constrained inverse reinforcement learning," in *Uncertainty in Artificial Intelligence*. PMLR, 2022, pp. 589–598.
- [6] J. Fischer, C. Eyberg, M. Werling, and M. Lauer, "Sampling-based inverse reinforcement learning algorithms with safety constraints," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 791–798.
- [7] G. Kalweit, M. Huegle, M. Werling, and J. Boedecker, "Deep inverse q-learning with constraints," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 14 291–14 302.
- [8] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent Advances in Robot Learning from Demonstration," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, pp. 297–330, May 2020.
- [9] G. Chou, D. Berenson, and N. Ozay, "Learning constraints from demonstrations," in *Algorithmic Foundations of Robotics XIII: Proceedings of the 13th Workshop on the Algorithmic Foundations of Robotics 13*. Springer, 2020, pp. 228–245.
- [10] D. Park, M. Noseworthy, R. Paul, S. Roy, and N. Roy, "Inferring task goals and constraints using bayesian nonparametric inverse reinforcement learning," in *Conference on robot learning*. PMLR, 2020, pp. 1005–1014.
- [11] G. Chou, N. Ozay, and D. Berenson, "Learning constraints from locally-optimal demonstrations under cost function uncertainty," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3682–3690, 2020.
- [12] A. Y. Ng, S. Russell *et al.*, "Algorithms for inverse reinforcement learning," in *Icml*, vol. 1, 2000, p. 2.
- [13] B. Schölkopf and A. J. Smola, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- [14] W. Karush, "Minima of functions of several variables with inequalities as side constraints," *M. Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago*, 1939.
- [15] H. W. Kuhn and A. W. Tucker, "Nonlinear programming," in *Traces and emergence of nonlinear programming*. Springer, 2013, pp. 247–258.
- [16] B. Rosner, "Percentage points for a generalized esd many-outlier procedure," *Technometrics*, vol. 25, no. 2, pp. 165–172, 1983.
- [17] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *2001 European control conference (ECC)*. IEEE, 2001, pp. 2603–2608.
- [18] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in matlab," in *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.