

Learning Constraints from Demonstrations

Glen Chou, Dmitry Berenson, Necmiye Ozay

Dept. of Electrical Engineering and Computer Science,
University of Michigan, Ann Arbor, MI, 48109, USA
{gchou,dmitryb,necmiye}@umich.edu

Abstract. We extend the learning from demonstration paradigm by providing a method for learning unknown constraints shared across tasks, using demonstrations of the tasks, their cost functions, and knowledge of the system dynamics and control constraints. Given safe demonstrations, our method uses hit-and-run sampling to obtain lower cost, and thus unsafe, trajectories. Both safe and unsafe trajectories are used to obtain a consistent representation of the unsafe set via solving an integer program. Our method generalizes across system dynamics and learns a guaranteed subset of the constraint. We also provide theoretical analysis on what subset of the constraint can be learnable from safe demonstrations. We demonstrate our method on linear and nonlinear system dynamics, show that it can be modified to work with suboptimal demonstrations, and that it can also be used to learn constraints in a feature space.

Keywords: learning from demonstration, machine learning, motion and path planning

1 Introduction

Inverse optimal control and inverse reinforcement learning (IOC/IRL) [2,6,22,26] have proven to be powerful tools in enabling robots to perform complex goal-directed tasks. These methods learn a cost function that replicates the behavior of an expert demonstrator when optimized. However, planning for many robotics and automation tasks also requires knowing constraints, which define what states or trajectories are safe. For example, the task of safely and efficiently navigating an autonomous vehicle can naturally be described by a cost function trading off user comfort and efficiency and by the constraints of collision avoidance and executing only legal driving behaviors. In some situations, constraints can provide a more interpretable representation of a behavior than cost functions. For example, in safety critical environments, recovering a hard constraint or an explicit representation of an unsafe set in the environment is more useful than learning a “softened” cost function representation of the constraint as a penalty term in the Lagrangian. Consider the autonomous vehicle, which absolutely must avoid collision, not simply give collisions a cost penalty. Furthermore, learning global constraints shared across many tasks can be useful for generalization. Again consider the autonomous vehicle, which must avoid the scene of a car accident: a shared constraint that holds regardless of the task it is trying to complete.

While constraints are important, it can be impractical for a user to exhaustively program into a robot all the possible constraints it should obey when performing its repertoire of tasks. To avoid this, we consider in this paper the problem of recovering the latent constraints within expert demonstrations that are shared across tasks in the environment. Our method is based on the key insight that each safe, optimal demonstration induces a set of lower-cost trajectories that must be unsafe due to violation of an unknown constraint. Our method samples these unsafe trajectories, ensuring they are also consistent with the known constraints (system dynamics, control constraints, and start/goal constraints), and uses these unsafe trajectories together with the safe demonstrations as constraints in an “inverse” integer program which recovers a consistent unsafe set. Our contributions are fourfold:

- We pose the novel problem of learning a shared constraint across tasks.
- We propose an algorithm that, given known constraints and boundedly sub-optimal demonstrations of state-control sequences, extracts unknown constraints defined in a wide range of constraint spaces (not limited to the trajectory or state spaces) shared across demonstrations of different tasks.
- We provide theoretical analysis on the limits of what subsets of a constraint can be learned, depending on the demonstrations, the system dynamics, and the trajectory discretization. We also show that our method can recover a guaranteed underapproximation of the constraint.
- We provide experiments that justify our theory and show that our algorithm can recover an unsafe set with few demonstrations, across different types of linear and nonlinear dynamics, and can be adapted to work with boundedly suboptimal demonstrations. We also demonstrate that our method can learn constraints in the state space and a feature space.

2 Related Work

Inverse optimal control [14, 16] (IOC) and **inverse reinforcement learning** (IRL) [22] aim to recover an objective function consistent with the received expert demonstrations, in the sense that the demonstrations (approximately) optimize the cost function. Our method is complementary to these approaches; if the demonstration is solving a constrained optimization problem, we are finding its constraints, given the objective function; IOC/IRL finds the objective function, given its constraints. For example, [12] attempts to learn the cost function of a constrained optimization problem from optimal demonstrations by minimizing the residuals of the KKT conditions, but the constraints themselves are assumed known. Another approach [5] can represent a state-space constraint shared across tasks as a penalty term in the reward function of an MDP. However, when viewing a constraint as a penalty, it becomes unclear if a demonstrated motion was performed to avoid a penalty or to improve the cost of the trajectory in terms of the true cost function (or both). Thus, learning a constraint which generalizes between tasks with different cost functions becomes difficult. To avoid this issue, we assume a known cost function to explicitly reason about the constraint.

One branch of **safe reinforcement learning** aims to perform exploration while minimizing visitation of unsafe states. Several methods for safe exploration

in the state space [3, 28, 29] use a Gaussian process (GP) to explore safe regions in the state space. These approaches differ from ours in that they use exploration instead of demonstrations. Some drawbacks to these methods include that unsafe states can still be visited, Lipschitz continuity of the safety function is assumed, or the dynamics are unknown but the safe set is known. Furthermore, states themselves are required to be explicitly labeled as safe or unsafe, while we only require the labeling of whole trajectories. Our method is capable of learning a binary constraint defined in other spaces, using only state-control trajectories.

There exists prior work in learning **geometric constraints** in the workspace. In [7], a method is proposed for learning Pfaffian constraints, recovering a linear constraint parametrization. In [25], a method is proposed to learn geometric constraints which can be described by the classes of considered constraint templates. Our method generalizes these methods by being able to learn a nonlinear constraint defined in any constraint space (not limited to the state space).

Learning **local trajectory-based constraints** has also been explored in the literature. The method in [19] samples feasible poses around waypoints of a single demonstration; areas where few feasible poses can be sampled are assumed to be constrained. Similarly, [20] performs online constraint inference in the feature space from a single trajectory, and then learns a mapping to the task space. The methods in [9, 10, 23, 30] also learn constraints in a single task. These methods are inherently local since only one trajectory or task is provided, unlike our method, which aims to learn a global constraint shared across tasks.

3 Preliminaries and Problem Statement

The goal of this work is to recover unknown constraints shared across a collection of optimization problems, given boundedly suboptimal solutions, the cost functions, and knowledge of the dynamics, control constraints, and start/goal constraints. We discuss the forward problem, which generates the demonstrations, and the inverse problem: the core of this work, which recovers the constraints.

3.1 Forward optimal control problem

Consider an agent described by a state in some state space $x \in \mathcal{X}$. It can take control actions $u \in \mathcal{U}$ to change its state. The agent performs tasks Π drawn from a set of tasks \mathcal{P} , where each task Π can be written as a constrained optimization problem over state trajectories in state trajectory space $\xi_x \in \mathcal{T}^x$ and control trajectories $\xi_u \in \mathcal{T}^u$ in control trajectory space:

Problem 1 (Forward problem / “task” Π).

$$\begin{aligned} & \underset{\xi_x, \xi_u}{\text{minimize}} && c_{\Pi}(\xi_x, \xi_u) \\ & \text{subject to} && \phi(\xi_x, \xi_u) \in \mathcal{S} \subseteq \mathcal{C} \\ & && \phi(\xi_x, \xi_u) \in \bar{\mathcal{S}} \subseteq \bar{\mathcal{C}} \\ & && \phi_{\Pi}(\xi_x, \xi_u) \in \mathcal{S}_{\Pi} \subseteq \mathcal{C}_{\Pi} \end{aligned} \tag{1}$$

where $c_{\Pi}(\cdot) : \mathcal{T}^x \times \mathcal{T}^u \rightarrow \mathbb{R}$ is a cost function for task Π , $\phi(\cdot, \cdot) : \mathcal{T}^x \times \mathcal{T}^u \rightarrow \mathcal{C}$ is a known feature function mapping state-control trajectories to some constraint

space \mathcal{C} . $\bar{\phi}(\cdot, \cdot) : \mathcal{T}^x \times \mathcal{T}^u \rightarrow \bar{\mathcal{C}}$ and $\phi_{\Pi}(\cdot, \cdot) : \mathcal{T}^x \times \mathcal{T}^u \rightarrow \mathcal{C}_{\Pi}$ are known and map to potentially different constraint spaces $\bar{\mathcal{C}}$ and \mathcal{C}_{Π} , containing a known shared safe set $\bar{\mathcal{S}}$ and a known task-dependent safe set \mathcal{S}_{Π} , respectively. \mathcal{S} is an unknown safe set, and the inverse problem aims to recover its complement, $\mathcal{A} \doteq \mathcal{S}^c$, the “unsafe” set. In this paper, we focus on constraints separable in time: $\phi(\xi_x, \xi_u) \in \mathcal{A} \Leftrightarrow \exists t \in \{1, \dots, T\} \phi(\xi_x(t), \xi_u(t)) \in \mathcal{A}$, where we overload ϕ so it applies to the instantaneous values of the state and the input. An analogous definition holds for the continuous time case. Our method easily learns non-separable trajectory constraints as well¹.

A demonstration, $\xi_{xu} \doteq (\xi_x, \xi_u)$, is a state-control trajectory which is a boundedly suboptimal solution to Problem 1, i.e. the demonstration satisfies all constraints and its cost is at most a factor of δ above the cost of the optimal solution ξ_{xu}^* , i.e. $c(\xi_x^*, \xi_u^*) \leq c(\xi_x, \xi_u) \leq (1 + \delta)c(\xi_x^*, \xi_u^*)$. Furthermore, let T be a finite time horizon which is allowed to vary. If ξ_{xu} is a discrete-time trajectory ($\xi_x = \{x_1, \dots, x_T\}$, $\xi_u = \{u_1, \dots, u_T\}$), Problem 1 is a finite-dimensional optimization problem, while Problem 1 becomes a functional optimization problem if ξ_{xu} is a continuous-time trajectory ($\xi_x : [0, T] \rightarrow \mathcal{X}$, $\xi_u : [0, T] \rightarrow \mathcal{U}$). We emphasize this setup does not restrict the unknown constraint to be defined on the trajectory space; it allows for constraints to be defined on any space described by the range of some known feature function ϕ .

We assume the trajectories are generated by a dynamical system $\dot{x} = f(x, u, t)$ or $x_{t+1} = f(x_t, u_t, t)$ with control constraints $u_t \in \mathcal{U}$, for all t , and that the dynamics, control constraints, and start/goal constraints are known. We further denote the set of state-control trajectories satisfying the unknown shared constraint, the known shared constraint, and the known task-dependent constraint as $\mathcal{T}_{\mathcal{S}}$, $\mathcal{T}_{\bar{\mathcal{S}}}$, and $\mathcal{T}_{\mathcal{S}_{\Pi}}$, respectively. Lastly, we also denote the set of trajectories satisfying all known constraints but violating the unknown constraint as $\mathcal{T}_{\mathcal{A}}$.

3.2 Inverse constraint learning problem

The goal of the inverse constraint learning problem is to recover an unsafe set, $\mathcal{A} \subseteq \mathcal{C}$, using N_s provided safe demonstrations $\xi_{s_j}^*$, $j = 1, \dots, N_s$, known constraints, and N_{-s} inferred unsafe trajectories, ξ_{-s_k} , $k = 1, \dots, N_{-s}$, generated by our method, which can come from multiple tasks. These trajectories can together be thought of as a set of constraints on the possible assignments of unsafe elements in \mathcal{C} . To recover a gridded approximation of the unsafe set \mathcal{A} that is consistent with these trajectories, we first discretize \mathcal{C} into a finite set of G discrete cells $\mathcal{Z} \doteq \{z_1, \dots, z_G\}$ and define an occupancy function, $\mathcal{O}(\cdot)$, which maps each cell to its safeness: $\mathcal{O}(\cdot) : \mathcal{Z} \rightarrow \{0, 1\}$, where

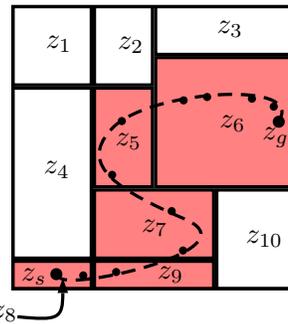


Fig. 1: Discretized constraint space with cells z_1, \dots, z_{10} . The trajectory’s constraint values are assigned to the red cells.

¹ Write Problem 2 constraints as sums over partially separable/inseparable feature components instead of completely separable components.

$\mathcal{O}(z_i) = 1$ if $z_i \in \mathcal{A}$, and 0 otherwise. Continuous space trajectories are gridded by concatenating the set of grid cells z_i that $\phi(x_1), \dots, \phi(x_T)$ lie in, which is graphically shown in Figure 1 with a non-uniform grid. Then, the problem can be written down as an integer feasibility problem

Problem 2 (Inverse feasibility problem).

$$\begin{aligned} & \text{find } \mathcal{O}(z_1), \dots, \mathcal{O}(z_G) \in \{0, 1\}^G \\ & \text{subject to } \sum_{\substack{z_i \in \{\phi(\xi_{s_j}^*(1)), \dots, \\ \phi(\xi_{s_j}^*(T_j))\}}} \mathcal{O}(z_i) = 0, \quad \forall j = 1, \dots, N_s \\ & \sum_{\substack{z_i \in \{\phi(\xi_{-s_k}(1)), \dots, \\ \phi(\xi_{-s_k}(T_k))\}}} \mathcal{O}(z_i) \geq 1, \quad \forall k = 1, \dots, N_{-s} \end{aligned} \quad (2)$$

Inferring unsafe trajectories, i.e. obtaining $\xi_{-s_k}, k = 1, \dots, N_{-s}$, is the most difficult part of this problem, since finding lower-cost trajectories consistent with known constraints that complete a task is essentially a planning problem. Much of the next section shows how to efficiently obtain ξ_{-s_k} . Further details on Problem 2, including conservativeness guarantees, incorporating a prior on the constraint, and a continuous relaxation can be found in Section 4.4.

4 Method

The key to our method lies in finding lower-cost trajectories that do not violate the known constraints, given a demonstration with boundedly-suboptimal cost satisfying all constraints. Such trajectories must then violate the unknown constraint. Our goal is to determine an unsafe set in the constraint space from these trajectories using Problem 2. In the following, Section 4.1 describes lower-cost trajectories consistent with the known constraints; Section 4.2 describes how to sample such trajectories; Section 4.3 describes how to get more information from unsafe trajectories; Section 4.4 describes details and extensions to Problem 2; Section 4.5 discusses how to extend our method to suboptimal demonstrations. The complete flow of our method is described in Algorithm 2.

4.1 Trajectories satisfying known constraints

Consider the forward problem (Problem 1). We define the set of unsafe state-control trajectories induced by an optimal, safe demonstration $\xi_{xu}^*, \mathcal{T}_A^{\xi_{xu}^*}$, as the set of state-control trajectories of lower cost that obey the known constraints:

$$\mathcal{T}_A^{\xi_{xu}^*} \doteq \{\xi_{xu} \mid c(\xi_x, \xi_u) < c(\xi_x^*, \xi_u^*), \xi_{xu} \in \mathcal{T}_{\mathcal{S}}, \xi_{xu} \in \mathcal{T}_{\mathcal{S}H}\}. \quad (3)$$

In this paper, we deal with the known constraints from the system dynamics, the control limits, and task-dependent start and goal state constraints. Hence, $\mathcal{T}_{\mathcal{S}} = \mathcal{D}^{\xi_{xu}} \cap \mathcal{U}^{\xi_{xu}}$, where $\mathcal{D}^{\xi_{xu}}$ denotes the set of dynamically feasible trajectories and $\mathcal{U}^{\xi_{xu}}$ denotes the set of trajectories using controls in \mathcal{U} at each time-step. $\mathcal{T}_{\mathcal{S}H}$

denotes trajectories satisfying start and goal constraints. We develop the method for discrete time trajectories, but analogous definitions hold in continuous time. For discrete time, length T trajectories, $\mathcal{U}^{\xi_{xu}}$, $\mathcal{D}^{\xi_{xu}}$, and \mathcal{T}_{S_H} are explicitly:

$$\begin{aligned}\mathcal{U}^{\xi_{xu}} &\doteq \{ \xi_{xu} \mid u_t \in \mathcal{U}, \forall t \in \{1, \dots, T-1\} \}, \\ \mathcal{D}^{\xi_{xu}} &\doteq \{ \xi_{xu} \mid x_{t+1} = f(x_t, u_t), \forall t \in \{1, \dots, T-1\} \}, \\ \mathcal{T}_{S_H} &\doteq \{ \xi_{xu} \mid x_1 = x_s, x_T = x_g \}.\end{aligned}\tag{4}$$

4.2 Sampling trajectories satisfying known constraints

Dynamics	Cost function	Control constraints	Sampling method
Linear	Quadratic	Convex	Ellipsoid hit-and-run (Section 4.2.1)
Linear	Convex	Convex	Convex hit-and-run (Section 4.2.2)
	Else		Non-convex hit-and-run (Section 4.2.3)

Table 1: Sampling methods for different dynamics/costs/feasible controls.

We sample from $\mathcal{T}_{\mathcal{A}}^{\xi_{xu}^*}$ to obtain lower-cost trajectories obeying the known constraints. For the most part, we use hit-and-run sampling [18] over the set $\mathcal{T}_{\mathcal{A}}^{\xi_{xu}^*}$, a method guaranteeing convergence to a uniform distribution of samples over $\mathcal{T}_{\mathcal{A}}^{\xi_{xu}^*}$ in the limit; the method is detailed in Algorithm 1 and an illustration is shown in Figure 2. Hit-and-run starts from an initial point within the set, chooses a direction uniformly at random, moves a random amount in that direction such that the new point remains within the set, and repeats.

Depending on the convexity of the cost function and the control constraints and the form of the dynamics, different sampling techniques can be used, organized in Table 1. The following sections describe each sampling method.

Algorithm 1: Hit-and-run

Output: $\text{out} \doteq \{\xi_1, \dots, \xi_{N-s}\}$
Input : $\mathcal{T}_{\mathcal{A}}^{\xi_{xu}^*}, \xi_{xu}^*, N-s$
1 $\xi_{xu} \leftarrow \xi_{xu}^*$; $\text{out} \leftarrow \{\}$;
2 **for** $i = 1:N-s$ **do**
3 $r \leftarrow \text{sampleRandDirection}$;
4 $\mathcal{L} \leftarrow \mathcal{T}_{\mathcal{A}}^{\xi_{xu}^*} \cap \{\xi'_{xu} \in \mathcal{T} \mid \xi'_{xu} = \xi_{xu} + \beta r\}$;
5 $L_-, L_+ \leftarrow \text{endpoints}(\mathcal{L})$;
6 $\xi_{xu} \sim \text{Uniform}(L_-, L_+)$;
7 $\text{out} \leftarrow \text{out} \cup \xi_{xu}$;
8 **end**

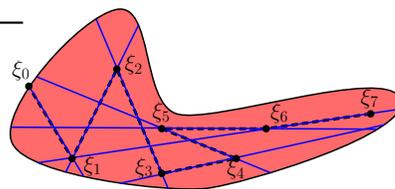


Fig. 2: Illustration of hit-and-run. Blue lines denote sampled random directions, black dots denote samples.

4.2.1 Ellipsoid hit-and-run When we have a linear system with quadratic cost and convex control constraints - a very common setup in the optimal control literature - $\mathcal{T}_{\mathcal{A}}^{\xi_{xu}^*} \doteq \{ \xi_{xu} \mid c(\xi_{xu}) < c(\xi_{xu}^*) \} \cap \mathcal{D}^{\xi_{xu}} \equiv \{ \xi_{xu} \mid \xi_{xu}^\top V \xi_{xu} < \xi_{xu}^{*\top} V \xi_{xu}^* \} \cap \mathcal{D}^{\xi_{xu}}$ is an ellipsoid in the trajectory space, which can be efficiently sampled via a specially-tailored hit-and-run method. Here, the quadratic cost is written

as $c(\xi_{xu}) \doteq \xi_{xu}^\top V \xi_{xu}$, where V is a matrix of cost parameters, and we omit the control and task constraints for now. Without dynamics, the endpoints of the line \mathcal{L} , L_- , L_+ , (c.f. Alg. 1), can be found by solving a quadratic equation $(\xi_{xu} + \beta r)^\top V (\xi_{xu} + \beta r) = \xi_{xu}^{*\top} V \xi_{xu}^*$. We show that this can still be done with linear dynamics by writing $\mathcal{T}_A^{\xi_{xu}}$ in a special way. $\mathcal{D}^{\xi_{xu}}$ can be written as an eigenspace of a singular “dynamics consistency” matrix, D_1 , which converts any arbitrary state-control trajectory to one that satisfies the dynamics, one time-step at a time. Precisely, if the dynamics can be written as $x_{t+1} = Ax_t + Bu_t$, we can write a matrix D_1 :

$$\underbrace{\begin{bmatrix} x_1 \\ u_1 \\ x_2 \\ u_2 \\ \tilde{x}_3 \\ \vdots \\ u_{T-1} \\ \tilde{x}_T \end{bmatrix}}_{\hat{\xi}_{xu}} = \underbrace{\begin{bmatrix} I & 0 & 0 & 0 & 0 & \cdots & \cdots & 0 \\ 0 & I & 0 & 0 & 0 & \cdots & \cdots & 0 \\ A & B & 0 & 0 & 0 & \cdots & \cdots & 0 \\ 0 & 0 & 0 & I & 0 & \cdots & \cdots & 0 \\ 0 & 0 & A & B & 0 & \cdots & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \cdots & 0 & I & 0 \\ 0 & 0 & 0 & \cdots & \cdots & A & B & 0 \end{bmatrix}}_{D_1} \underbrace{\begin{bmatrix} x_1 \\ u_1 \\ \tilde{x}_2 \\ u_2 \\ \tilde{x}_3 \\ \vdots \\ u_{T-1} \\ \tilde{x}_T \end{bmatrix}}_{\xi_{xu}} \quad (5)$$

that fixes the controls and the initial state and performs a one-step rollout, replacing the second state with the dynamically correct state. In Eq. 5, we denote by \tilde{x}_{t+1} a state that cannot be reached by applying control u_t to state x_t . Multiplying the one-step corrected trajectory $\hat{\xi}_{xu}$ by D_1 again changes \tilde{x}_3 to the dynamically reachable state x_3 . Applying D_1 to the original T -time-step infeasible trajectory $T - 1$ times results in a dynamically feasible trajectory, $\xi_{xu}^{\text{feas}} = D_1^{T-1} \hat{\xi}_{xu}$. Further, note that the set of dynamically feasible trajectories is $\mathcal{D}^{\xi_{xu}} \doteq \{\xi_{xu} \mid D_1 \xi_{xu} = \xi_{xu}\}$, which is the span of the eigenvectors of D_1 associated with eigenvalue 1. Thus, obtaining a feasible trajectory via repeated multiplication is akin to finding the eigenspace via power iteration [13]. One can also interpret this as propagating through the dynamics with a fixed control sequence. Now, we can write $\mathcal{T}_A^{\xi_{xu}}$ as another ellipsoid which can be efficiently sampled by finding L_- , L_+ by solving a quadratic equation:

$$\mathcal{T}_A^{\xi_{xu}} \doteq \{\xi_{xu} \mid \xi_{xu}^\top D_1^{T-1} V D_1^{T-1} \xi_{xu} \leq \xi_{xu}^{*\top} V \xi_{xu}^*\}. \quad (6)$$

We deal with control constraints separately, as the intersection of $\mathcal{U}^{\xi_{xu}}$ and Eq. 6 is in general not an ellipsoid. To ensure control constraint satisfaction, we reject samples with controls outside of $\mathcal{U}^{\xi_{xu}}$; this works if $\mathcal{U}^{\xi_{xu}}$ is not measure zero. For task constraints, we ensure all sampled rollouts obey the goal constraints by adding a large penalty term to the cost function: $\tilde{c}(\cdot) \doteq c(\cdot) + \alpha_c \|x_g - x_T\|_2^2$, where α_c is a large scalar, which can be incorporated into Eq. 6 by modifying V and including x_g in ξ_{xu} ; all trajectories sampled in this modified set satisfy the goal constraints to an arbitrarily small tolerance ε , depending on the value of α_c . The start constraint is satisfied trivially: all rollouts start at x_s . Note the demonstration cost remains the same, since the demonstration satisfies the start and goal constraints; this modification is made purely to ensure these constraints hold for sampled trajectories.

4.2.2 Convex hit-and-run For general convex cost functions, the same sampling method holds, but L_+ , L_- cannot be found by solving a quadratic function. Instead, we solve $c(\xi_{xu} + \beta r) = c(\xi_{xu}^*)$ via a root finding algorithm or line search.

4.2.3 Non-convex hit-and-run If $\mathcal{T}_A^{\xi_{xu}^*}$ is non-convex, \mathcal{L} can now in general be a union of disjoint line segments. In this scenario, we perform a “backtracking” line search by setting β to lie in some initial range: $\beta \in [\underline{\beta}, \bar{\beta}]$; sampling β_s within this range and then evaluating the cost function to see whether or not $\xi_{xu} + \beta_s r$ lies within the intersection. If it does, the sample is kept and hit-and-run proceeds normally; if not, then the range of possible β values is restricted to $[\beta_s, \bar{\beta}]$ if β_s is negative, and $[\underline{\beta}, \beta_s]$ otherwise. Then, new β_s are re-sampled until either the interval length shrinks below a threshold or a feasible sample is found. This altered hit-and-run technique still converges to a uniform distribution on the set in the limit, but has a slower mixing time than for the convex case, where mixing time describes the number of samples needed until the total variation distance to the steady state distribution is less than a small threshold [1]. Further, we accelerate sampling spread by relaxing the goal constraint to a larger tolerance $\hat{\varepsilon} > \varepsilon$ but keeping only the trajectories reaching within ε of the goal.

4.3 Improving learnability using cost function structure

Algorithm 2: Overall method	
Output: $\mathcal{O} \doteq \mathcal{O}(z_1), \dots, \mathcal{O}(z_G)$ Input : $\xi_s = \{\xi_1^*, \dots, \xi_{N_s}^*\}$, $c_{\Pi}(\cdot)$, known constraints}	
1 $\xi_u \leftarrow \{\}$; 2 for $i = 1:N_s$ do <i>/* Sample unsafe ξ */</i> 3 if <i>lin., quad., conv.</i> then 4 $\xi_u \leftarrow \xi_u \cap \text{ellipsoidHNR}(\xi_i^*)$; 5 else if <i>lin., conv., conv.</i> then 6 $\xi_u \leftarrow \xi_u \cap \text{convexHNR}(\xi_i^*)$; 7 else 8 $\xi_u \leftarrow \xi_u \cap \text{nonconvexHNR}(\xi_i^*)$; 9 end <i>/* Constraint recovery */</i> 10 if <i>prior, continuous</i> then 11 $\mathcal{O} \leftarrow \text{Problem 4}(\xi_s, \xi_u)$ 12 else if <i>prior, binary</i> then 13 $\mathcal{O} \leftarrow \text{Problem 3}(\xi_s, \xi_u)$ 14 else 15 $\mathcal{O} \leftarrow \text{Problem 2}(\xi_s, \xi_u)$ 16 end	Naïvely, the sampled unsafe trajectories may provide little information. Consider an unsafe, length- T discrete-time trajectory ξ , with start and end states in the safe set. This only says there exists at least one intermediate unsafe state in the trajectory, but says nothing directly about which state was unsafe. The weakness of this information can be made concrete using the notion of a version space. In machine learning, the version space is the set of consistent hypotheses given a set of examples [27]. In our setting, hypotheses are possible unsafe sets, and examples are the safe and unsafe trajectories. Knowing ξ is unsafe only disallows unsafe sets that mark every element of the constraint space that ξ traverses as safe: $(\mathcal{O}(z_2) = 0) \wedge \dots \wedge (\mathcal{O}(z_{T-1}) = 0)$. If \mathcal{C} is gridded into G cells, this information invalidates at most 2^{G-T+2} out of 2^G possible unsafe

sets. We could do exponentially better if we reduced the number of cells that ξ implies could be unsafe.

We can achieve this by sampling sub-segments (or sub-trajectories) of the larger demonstrations, holding other portions of the demonstration fixed. For ex-

ample, say we fix all but one of the points on ξ when sampling unsafe lower-cost trajectories. Since only one state can be different from the known safe demonstration, the unsafeness of the trajectory can be uniquely localized to whatever new point was sampled: then, this trajectory will reduce the version space by at most a factor of 2, invalidating at most $2^G - 2^{G-1} = 2^{G-1}$ unsafe sets. One can sample these sub-trajectories in the full-length trajectory space by fixing appropriate waypoints during sampling: this ensures the full trajectory has lower cost and only perturbs desired waypoints. However, to speed up sampling, sub-trajectories can be sampled directly in the lower dimensional sub-trajectory space if the cost function $c(\cdot)$ that is being optimized is strictly monotone [21]: for any costs $c_1, c_2 \in \mathbb{R}$, control $u \in \mathcal{U}$, and state $x \in \mathcal{X}$, $c_1 < c_2 \Rightarrow h(c_1, x, u) < h(c_2, x, u)$, for all x, u , where $h(c, x, u)$ represents the cost of starting with initial cost c at state x and taking control u . Strictly monotone cost functions include separable cost functions with additive or multiplicative stage costs, which are common in motion planning and optimal control. If the cost function is strictly monotone, we can sample lower-cost trajectories from sub-segments of the optimal path; otherwise it is possible that even if a new sub-segment with lower cost than the original sub-segment were sampled, the full trajectory containing the sub-segment could have a higher cost than the demonstration.

4.4 Integer program formulation

After sampling, we can solve Problem 2 to find an unsafe set consistent with the safe and unsafe trajectories. We now discuss the details of this process. **Conservative estimate:** One can obtain a conservative estimate of the unsafe set \mathcal{A} from Problem 2 by intersecting all possible solutions: if the unsafeness of a cell is shared across all feasible solutions, that cell must be occupied. In practice, it may be difficult to directly find all solutions to the feasibility problem, as in the worst case, finding the set of all feasible solutions is equivalent to exhaustive search in the full gridded space [24]. A more efficient method is to loop over all G grid cells and set each one to be safe, and see if the optimizer can still find a feasible solution. Cells where there exists no feasible solution are guaranteed unsafe. This amounts to solving G binary integer feasibility problems, which can be trivially parallelized. Furthermore, any cells that are known safe (from demonstrations) do not need to be checked. We use this method to compute the “learned guaranteed unsafe set”, $\mathcal{A}_l^{\text{rec}}$, in Section 6.

A prior on the constraint: As will be further discussed in Section 5.1, it may be fundamentally impossible to recover a unique unsafe set. If we have some prior on the nature of the unsafe set, such as it being simply connected, or that certain regions of the constraint space are unlikely to be unsafe, we can make the constraint learning problem more well-posed. Assume that this prior knowledge can be encoded in some “energy” function $E(\cdot, \dots, \cdot) : \{0, 1\}^G \rightarrow \mathbb{R}$ mapping the set of binary occupancies to a scalar value, which indicates the desirability of a particular unsafe set configuration. Using E as the objective function in Problem 2 results in a binary integer program, which finds an unsafe set consistent with the safe and unsafe trajectories, and minimizes the energy:

Problem 3 (Inverse binary minimization constraint recovery).

$$\begin{aligned}
& \underset{\mathcal{O}(z_1), \dots, \mathcal{O}(z_G) \in \{0,1\}^G}{\text{minimize}} && E(\mathcal{O}(z_1), \dots, \mathcal{O}(z_G)) \\
& \text{subject to} && \sum_{\substack{z_i \in \{\phi(\xi_{s_j}^*(1)), \dots, \\ \phi(\xi_{s_j}^*(T_j))\}}} \mathcal{O}(z_i) = 0, \quad \forall j = 1, \dots, N_s \\
& && \sum_{\substack{z_i \in \{\phi(\xi_{-s_k}(1)), \dots, \\ \phi(\xi_{-s_k}(T_k))\}}} \mathcal{O}(z_i) \geq 1, \quad \forall k = 1, \dots, N_{-s}
\end{aligned} \tag{7}$$

Probabilistic setting and continuous relaxation: A similar problem can be posed for a probabilistic setting, where grid cell occupancies represent beliefs over unsafeness: instead of the occupancy of a cell being an indicator variable, it is instead a random variable Z_i , where Z_i takes value 1 with probability $\tilde{\mathcal{O}}(Z_i)$ and value 0 with probability $1 - \tilde{\mathcal{O}}(Z_i)$. Here, the occupancy probability function maps cells to occupancy probabilities $\tilde{\mathcal{O}}(\cdot) : \mathcal{Z} \rightarrow [0, 1]$.

Trajectories can now be unsafe with some probability. We obtain analogous constraints from the integer program in Section 4.4 in the probabilistic setting. Known safe trajectories traverse cells that are unsafe with probability 0; we enforce this with the constraint $\sum_{Z_i \in \phi(\xi_{s_j}^*)} \tilde{\mathcal{O}}(Z_i) = 0$: if the unsafeness probabilities are all zero along a trajectory, then the trajectory must be safe. Trajectories that are unsafe with probability p_k satisfy $\sum_{Z_i \in \phi(\xi_{-s_k})} \tilde{\mathcal{O}}(Z_i) = \mathbb{E}[\sum_{Z_i \in \phi(\xi_{-s_k})} Z_i] = (1 - p_k) \cdot 0 + p_k \cdot S_k \geq p_k$ where we denote the number of unsafe grid cells $\phi(\xi_{-s_k})$ traverses when the trajectory is unsafe as S_k , where $S_k \geq 1$. The following problem directly optimizes over occupancy probabilities:

Problem 4 (Inverse continuous minimization constraint recovery).

$$\begin{aligned}
& \underset{\mathcal{O}(Z_1), \dots, \mathcal{O}(Z_G) \in [0,1]^G}{\text{minimize}} && E(\mathcal{O}(Z_1), \dots, \mathcal{O}(Z_G)) \\
& \text{subject to} && \sum_{\substack{Z_i \in \{\phi(\xi_{s_j}^*(1)), \dots, \\ \phi(\xi_{s_j}^*(T_j))\}}} \tilde{\mathcal{O}}(Z_i) = 0, \quad \forall j = 1, \dots, N_s \\
& && \sum_{\substack{Z_i \in \{\phi(\xi_{-s_k}(1)), \dots, \\ \phi(\xi_{-s_k}(T_k))\}}} \tilde{\mathcal{O}}(Z_i) \geq p_k, \quad \forall k = 1, \dots, N_{-s}
\end{aligned} \tag{8}$$

When $p_k = 1$, for all k (i.e. all unsafe trajectories are unsafe for sure), this probabilistic formulation coincides with the continuous relaxation of Problem 3. This justifies interpreting the solution of the continuous relaxation as occupancy probabilities for each cell. Note that Problem 3 and 4 have no conservativeness guarantees and use prior assumptions to make the problem more well-posed. However, we observe that they improve constraint recovery in our experiments.

4.5 Bounded suboptimality of demonstrations

If we are given a δ -suboptimal demonstration $\hat{\xi}$, where $c(\xi^*) \leq c(\hat{\xi}) \leq (1+\delta)c(\xi^*)$, where ξ^* is an optimal demonstration, we can still apply the sampling techniques discussed in earlier sections, but we must ensure that sampled unsafe trajectories

are truly unsafe: a sampled trajectory ξ' of cost $c(\xi') \geq c(\xi^*)$ can be potentially safe. Two options follow: one is to only keep trajectories with cost less than $\frac{c(\hat{\xi})}{1+\delta}$, but this can cause little to be learned if δ is large. Instead, if we assume a distribution on suboptimality, i.e. given a trajectory of cost $c(\hat{\xi})$, we know that a trajectory of cost $c(\xi') \in [\frac{c(\hat{\xi})}{1+\delta}, c(\hat{\xi})]$ is unsafe with probability p_k . We can then use these values of p_k to solve Problem 4. We implement this in the experiments.

5 Analysis

Due to space, the proofs/additional remarks can be found in the appendix.

5.1 Learnability

We provide analysis on the learnability of unsafe sets, given the known constraints and cost function. Most analysis assumes unsafe sets defined over the state space: $\mathcal{A} \subseteq \mathcal{X}$, but we extend it to the feature space in Corollary A.4. We provide some definitions and state a result bounding \mathcal{A}_l , the set of all states that can be learned guaranteed unsafe. We first define the signed distance:

Definition 1 (Signed distance). *Signed distance from point $p \in \mathbb{R}^m$ to set $\mathcal{S} \subseteq \mathbb{R}^m$, $sd(p, \mathcal{S}) = -\inf_{y \in \partial \mathcal{S}} \|p - y\|$ if $p \in \mathcal{S}$; $\inf_{y \in \partial \mathcal{S}} \|p - y\|$ if $p \in \mathcal{S}^c$.*

Theorem 1 (Learnability (discrete time)). *For trajectories generated by a discrete time dynamical system satisfying $\|x_{t+1} - x_t\| \leq \Delta x$ for all t , the set of learnable guaranteed unsafe states is a subset of the outermost Δx shell of the unsafe set: $\mathcal{A}_l \subseteq \{x \in \mathcal{A} \mid -\Delta x \leq sd(x, \mathcal{A}) \leq 0\}$ (see Section A.1 for illustration).*

Corollary 1 (Learnability (continuous time)). *For continuous trajectories $\xi(\cdot) : [0, T] \rightarrow \mathcal{X}$, the set of learnable guaranteed unsafe states shrinks to the boundary of the unsafe set: $\mathcal{A}_l \subseteq \{x \in \mathcal{A} \mid sd(x, \mathcal{A}) = 0\}$.*

Depending on the cost function, \mathcal{A}_l can become arbitrarily small: some cost functions are not very informative for recovering a constraint. For example, the path length cost function used in many of the experiments (which was chosen due to its common use in the motion planning community), prevents any lower-cost sub-trajectories from being sampled from straight sub-trajectories. The system's controllability also impacts learnability: the more controllable the system, the more of the Δx shell is reachable. We present a theorem quantifying when the dynamics allow unsafe trajectories to be sampled in Theorem A.2.

5.2 Conservativeness

We discuss conditions on \mathcal{A} and discretization which ensure our method provides a conservative estimate of \mathcal{A} . For analysis, we assume \mathcal{A} has a Lipschitz boundary [11]. We begin with notation (explanatory illustrations are in Section A.2):

Definition 2 (Set thickness). *Denote the outward-pointing normal vector at a point $p \in \partial \mathcal{A}$ as $\hat{n}(p)$. Furthermore, at non-differentiable points on $\partial \mathcal{A}$, $\hat{n}(p)$ is replaced by the set of normal vectors for the sub-gradient of the Lipschitz function describing $\partial \mathcal{A}$ at that point [4]. The set \mathcal{A} has a thickness larger than d_{thick} if $\forall x \in \partial \mathcal{A}, \forall d \in [0, d_{thick}], sd(x - d\hat{n}(x), \mathcal{A}) \leq 0$.*

Definition 3 (γ -offset padding). Define the γ -offset padding $\partial\mathcal{A}_\gamma$ as: $\partial\mathcal{A}_\gamma = \{x \in \mathcal{X} \mid \forall y \in \partial\mathcal{A} : x = y + d\hat{n}(y), d \in [0, \gamma]\}$.

Definition 4 (γ -padded set). We define the γ -padded set of the unsafe set \mathcal{A} , $\mathcal{A}(\gamma)$, as the union of the γ -offset padding and \mathcal{A} : $\mathcal{A}(\gamma) \doteq \partial\mathcal{A}_\gamma \cap \mathcal{A}$.

Corollary 2 (Conservative recovery of unsafe set). A sufficient condition ensuring that the set of learned guaranteed unsafe states $\mathcal{A}_l^{\text{rec}}$ is contained in \mathcal{A} is that \mathcal{A} has a set thickness greater than or equal to Δx (c.f. Definition 1).

If we use continuous trajectories directly, the guaranteed learnable set \mathcal{A}_l shrinks to a subset of the boundary of the unsafe set, $\partial\mathcal{A}$ (c.f. Corollary 1). However, if we discretize these trajectories, we can learn unsafe states lying in the interior, at the cost of conservativeness holding only for a padded unsafe set.

Theorem 2 (Continuous-to-discrete time conservativeness). Let ξ be a continuous trajectory: $\xi : [0, T] \rightarrow \mathcal{X}$. The system dynamics are described by $\dot{x} = f(x, u, t)$. The trajectory is discretized in time, potentially non-uniformly, resulting in a discretized trajectory $\{\xi(t_1), \dots, \xi(t_N)\}$, $t_i \in [0, T]$ for all i . Assume the maximum discretization time is Δt_{\max} . Denote

$$f_{\Delta x}^{\max} \doteq \sup_{x \in \mathcal{X}, u \in \mathcal{U}, t' \in [t, t + \Delta t_{\max})} \|f(x, u, t')\| \Delta t_{\max} \quad (9)$$

Then, our method recovers a subset of the $3f_{\Delta x}^{\max}$ -padded unsafe set, $\mathcal{A}(3f_{\Delta x}^{\max})$.

Corollary 3 (Continuous-to-discrete time and space conservativeness). Let the largest grid cell in the constraint space be contained by a ball of radius r_z . Then, if trajectories are discretized both in space and time, our method recovers a subset of the $2r_z + 3f_{\Delta x}^{\max}$ -padded unsafe set, $\mathcal{A}(2r_z + 3f_{\Delta x}^{\max})$.

Corollary 4 (Continuous-to-discrete feature space conservativeness). Let the feature mapping $\phi(x)$ from the state space to the constraint space be Lipschitz continuous with Lipschitz constant L . Then, our method recovers a subset of the $3Lf_{\Delta x}^{\max}$ -padded unsafe set in the feature space, $\mathcal{A}(3Lf_{\Delta x}^{\max})$.

6 Evaluations

We provide an example showing the importance of using unsafe trajectories, and experiments showing that our method generalizes across system dynamics, that it works with discretization and suboptimal demonstrations, and that it learns a constraint in a feature space from a single demonstration. See Appendix B for parameters, cost functions, the dynamics, control constraints, and timings.

Version space example: Consider a simple 5×5 8-connected grid world in which the tasks are to go from a start to a goal, minimizing Euclidean path length while staying out of the unsafe ‘‘U-shape’’, the outline of which is drawn in black (Fig. 3). Four demonstrations are provided, shown in Fig. 3 on the far left. Initially, the version space contains 2^{25} possible unsafe sets. Each safe trajectory

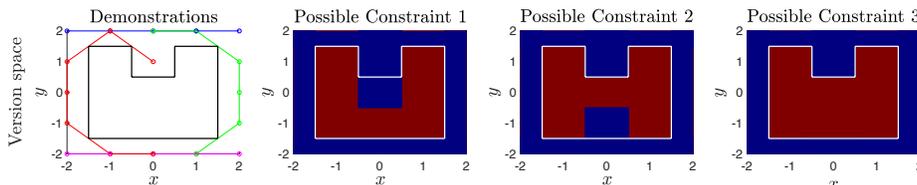


Fig. 3: *Leftmost:* Demonstrations and unsafe set. *Rest:* Set of possible constraints. Postulated unsafe cells are plotted in red, safe states in blue.

of length T reduces the version space at most by a factor of 2^T , invalidating at most $2^{25} - 2^{25-T}$ possible unsafe sets. Unsafe trajectories are computed by enumerating the set of trajectories going from the start to the goal at lower cost than the demonstration. The numbers of unsafe sets consistent with the safe and unsafe trajectories for varying numbers of safe trajectories are given in Table 2.

Ultimately, it is impossible to distinguish between the three unsafe sets on the right in Fig. 3. This is because there exists no task where a trajectory with cost lower than the demonstration can be sampled which only goes through one of the two uncertain states. Further, though the uncertain states are in the Δx shell of the constraint, due to the limitations of the cost function, we can only learn a subset of that shell (c.f. Theorem 1).

	1	2	3	4
Safe	262144	4096	1024	256
Safe & unsafe	11648	48	12	3

Table 2: Number of consistent unsafe sets, varying the no. of demonstrations, using/not using unsafe trajectories.

There are two main takeaways from this experiment. First, by generating unsafe trajectories, we can reduce the uncertainty arising from the ill-posedness of constraint learning: after 4 demonstrations, using unsafe demonstrations enables us to reduce the number of possible constraints by nearly a factor of 100, from 256 to 3. Second, due to limitations in the cost function, it may be impossible to recover a unique unsafe set, but the version space can be reduced substantially by sampling unsafe trajectories.

Dynamics and discretization: In the experiments in Fig. 4, we show that our method can be applied to several types of system dynamics, can learn non-convex/multiple unsafe sets, and can use continuous trajectories. The dynamics, control constraints, and cost functions for each experiment are given in Table 5 in Appendix B. All unsafe sets \mathcal{A} are open sets. We solve Problems 3 and 4, with an energy function promoting smoothness by penalizing squared deviations of the occupancy of a grid cell z_i from its 4-connected neighbors $N(z_i)$: $\sum_{i=1}^G \sum_{z_j \in N(z_i)} \|\mathcal{O}(z_i) - \mathcal{O}(z_j)\|_2^2$. In all experiments, the mean squared error (MSE) is computed as $\frac{1}{G} \sqrt{\sum_{i=1}^G \|\mathcal{O}(z_i)^* - \mathcal{O}(z_i)\|_2^2}$, where $\mathcal{O}(z_i)^*$ is the ground truth occupancy. The demonstrations are color-matched with their corresponding number on the x-axis of the MSE plots. For experiments with more demonstrations, only those causing a notable change in the MSE were color-coded. The learned guaranteed unsafe states $\mathcal{A}_l^{\text{rec}}$ are colored red on the left column.

We recover a non-convex “U-shaped” unsafe set in the state space using trivial 2D single-integrator dynamics (row 1 of Fig. 4). The solutions to both Problems 4

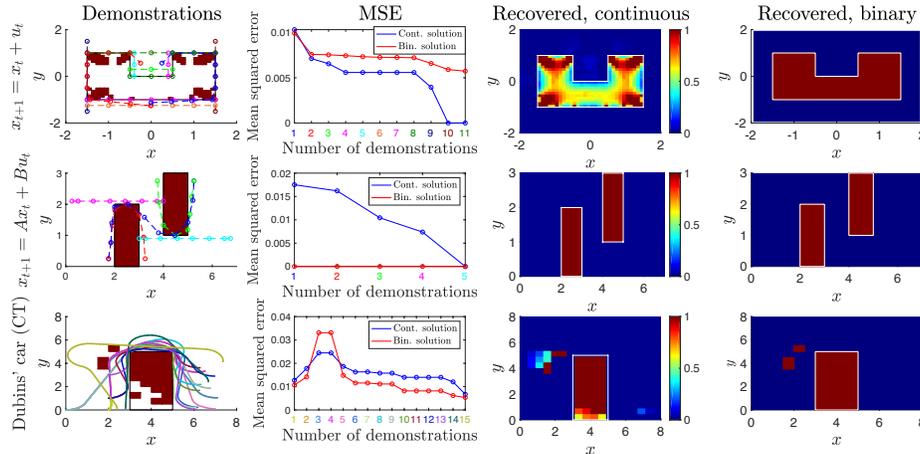


Fig. 4: Results across dynamics, discretization. **Rows (top-to-bottom):** Single integrator; double integrator; Dubins' car (CT). **Columns, left-to-right:** Demos., \mathcal{A} , $\mathcal{A}_l^{\text{rec}}$; MSE; Problem 4 solution, all demos.; Problem 3 solution, all demos.

and 3 return reasonable results, and the solution of Problem 3 achieves zero error. The second row shows learning two polyhedral unsafe sets in the state space with 4D double integrator linear dynamics, yielding similar results. We note the linear interpolation of some demonstrations in row 1 and 2 enter \mathcal{A} ; this is because both sets of dynamics are in discrete time and only the discrete waypoints must stay out of \mathcal{A} . The third row shows learning a polyhedral unsafe set in the state space, with time-discretized continuous, nonlinear Dubins' car dynamics, which has a 3D state $x \doteq [\chi \ y \ \theta]^\top$. These dynamics are more constrained than the previous cases, so sampling lower cost trajectories becomes more difficult, but despite this we can still achieve near zero error solving Problem 3. Some over-approximation results from some sampled unsafe trajectories entering regions not covered by the safe trajectories. For example, the cluster of red blocks to the top left of \mathcal{A} is generated by lower-cost trajectories that trade off the increased cost of entering the top left region by entering \mathcal{A} . This phenomenon is consistent with Theorem 3 of the appendix; we recover a set that is contained within the appropriate padded unsafe set (the max discretization time was 4.5 seconds). Learning curve spikes occur when overapproximation occurs. Overall, we note $\mathcal{A}_l^{\text{rec}}$ tends to be a significant underapproximation of \mathcal{A} due to the chosen cost function and limited demonstrations. For example, in row 1 of Fig. 4, $\mathcal{A}_l^{\text{rec}}$ cannot contain the portion of \mathcal{A} near long straight edges, since there exists no shorter path going from any start to any goal with only one state within that region. For row 3 of Fig. 4, we learn less of the bottom part of \mathcal{A} due to most demonstrations' start and goal locations making it harder to sample feasible control trajectories going through that region; with more demonstrations, this issue becomes less pronounced.

Suboptimal human demonstrations: We demonstrate our method on suboptimal demonstrations collected via a driving simulator, using a car model with CT Dubins' car dynamics. Human steering commands were recorded as demonstrations, where the task was to navigate around the orange box and drive be-

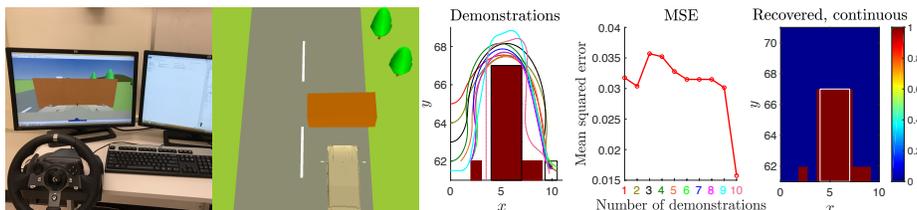


Fig. 5: *Suboptimal demonstrations: left: setup, center: demonstrations, \mathcal{A} , \mathcal{A}_l^{rec} , center-right: MSE, right: solution to Problem 4.*

tween the trees (Fig. 5). For a demonstration of cost c , trajectories with cost less than $0.9c$ were believed unsafe with probability 1. Trajectories with cost c' in the interval $[0.9c, c]$ were believed unsafe with probability $1 - ((c' - 0.9c)/0.1c)$. MSE for Problem 4 is shown in Fig. 5 (Problem 3 is not solved since the probabilistic interpretation is needed). The maximum discretization time is 6.9 seconds; hence, despite suboptimality, the learned guaranteed unsafe set is a subset of $\mathcal{A}(3f_{\Delta x}^{\max})$. While the MSE is highest here of all experiments, this is expected, as trajectories may be incorrectly labeled safe/unsafe with some probability.

Feature space constraint: We demonstrate that our framework is not limited to the state space by learning a constraint in a feature space. Consider the scenario of planning a safe path for a mobile robot with continuous Dubins' car dynamics through hilly terrain, where the magnitude of the terrain's slope is given as a feature map (i.e. $\phi(x) = \|\partial H(\hat{x})/\partial \hat{x}\|_2$, where $\hat{x} = [\chi \ y]^\top$ and $H(\hat{x})$ is the elevation map). The robot will slip if the magnitude of the terrain slope is too large, so we generate a demonstration which obeys the ground truth constraint $\phi(x) < 0.05$; hence, the ground truth unsafe set is $\mathcal{A} \doteq \{x \mid \phi(x) \geq 0.05\}$. From one safe trajectory (Fig. 6) generated by RRT* [15] and gridding the feature space as $\{0, 0.005, \dots, 0.145, 0.15\}$, we recover the constraint $\phi(x) < 0.05$ exactly.

7 Conclusion

In this paper we propose an algorithm that learns constraints from demonstrations, which acts as a complementary method to IOC/IRL algorithms. We analyze the properties of our algorithm as well as the theoretical limits of what subset of an unsafe set can be learned from safe demonstrations. The method works well on a variety of system dynamics and can be adapted to work with suboptimal demonstrations. We further show that our method can also learn constraints in a feature space. The largest shortcoming of our method is the constraint space gridding, which yields a complex constraint represen-

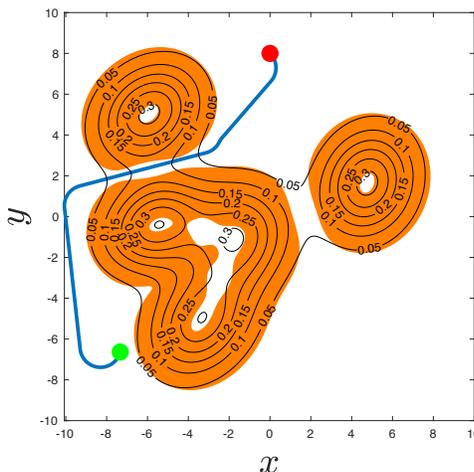


Fig. 6: *Demonstration (red: start, green: goal). Unsafe set \mathcal{A} is plotted in orange. Terrain isocontours $H(x) = \text{const}$ are overlaid.*

tation and causes the method to scale poorly to higher dimensional constraints. We aim to remedy this issue in future work by developing a grid-free counterpart of our method for convex unsafe sets, which can directly describe standard pose constraints like task space regions [8].

Acknowledgements

This work was supported in part by a Rackham first-year graduate fellowship, ONR grants N00014-18-1-2501 and N00014-17-1-2050, and NSF grants CNS-1446298, ECCS-1553873, and IIS-1750489.

References

1. Y. Abbasi-Yadkori, P. L. Bartlett, V. Gabillon, and A. Malek. Hit-and-run for sampling and planning in non-convex spaces. In *AISTATS 2017*, 2017.
2. P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004.
3. A. Akametalu, J. Fisac, J. Gillula, S. Kaynama, M. Zeilinger, and C. Tomlin. Reachability-based safe learning with gaussian processes. In *CDC*, Dec 2014.
4. G. Allaire, F. Jouve, and G. Michailidis. Thickness control in structural optimization via a level set method. *Struct. and Multidisciplinary Optimization*, 2016.
5. K. Amin, N. Jiang, and S. P. Singh. Repeated inverse reinforcement learning. In *NIPS*, pages 1813–1822, 2017.
6. B. Argall, S. Chernova, M. M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
7. L. Armesto, J. Bosga, V. Ivan, and S. Vijayakumar. Efficient learning of constraints and generic null space policies. In *ICRA*, pages 1520–1526. IEEE, 2017.
8. D. Berenson, S. S. Srinivasa, and J. J. Kuffner. Task space regions: A framework for pose-constrained manipulation planning. *IJRR*, 30(12):1435–1460, 2011.
9. S. Calinon and A. Billard. Incremental learning of gestures by imitation in a humanoid robot. In *HRI 2007*, pages 255–262, 2007.
10. S. Calinon and A. Billard. A probabilistic programming by demonstration framework handling constraints in joint space and task space. In *RSJ*, 2008.
11. B. Dacorogna. *Introduction to the calculus of variations*. Imp. College Press, 2015.
12. P. Englert, N. A. Vien, and M. Toussaint. Inverse kkt: Learning cost functions of manipulation tasks from demonstrations. *IJRR*, 36(13-14):1474–1488, 2017.
13. G. H. Golub and C. F. Van Loan. *Matrix Computations (3rd Ed.)*. 1996.
14. R. E. Kalman. When is a linear control system optimal? *Journal of Basic Engineering*, 86(1):51–60, Mar 1964.
15. S. Karaman and E. Frazzoli. Incremental sampling-based algorithms for optimal motion planning. In *RSS*, 2010.
16. A. Keshavarz, Y. Wang, and S. P. Boyd. Imputing a convex objective function. In *ISIC*, pages 613–619. IEEE, 2011.
17. H. K. Khalil. *Nonlinear systems*. Prentice-Hall, Upper Saddle River, NJ, 2002.
18. S. Kiatsupaibul, R. L. Smith, and Z. B. Zabinsky. An analysis of a variation of hit-and-run for uniform sampling from general regions. *TOMACS*, 2011.
19. C. Li and D. Berenson. Learning object orientation constraints and guiding constraints for narrow passages from one demonstration. In *ISER*. Springer, 2016.

20. N. Mehr, R. Horowitz, and A. D. Dragan. Inferring and assisting with constraints in shared autonomy. In *(CDC)*, pages 6689–6696, Dec 2016.
21. T. L. Morin. Monotonicity and the principle of optimality. *Journal of Mathematical Analysis and Applications*, 88(2):665 – 674, 1982.
22. A. Y. Ng and S. J. Russell. Algorithms for inverse reinforcement learning. In *ICML '00*, pages 663–670, San Francisco, CA, USA, 2000.
23. A. L. Pais, K. Umezawa, Y. Nakamura, and A. Billard. Learning robot skills through motion segmentation and constraints extraction. HRI, 2013.
24. C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice Hall, Englewood Cliffs, NJ, 1982.
25. C. Pérez-D’Arpino and J. A. Shah. C-LEARN: learning geometric constraints from demonstrations for multi-step manipulation in shared autonomy. In *ICRA*, 2017.
26. N. D. Ratliff, J. A. Bagnell, and M. Zinkevich. Maximum margin planning. In *ICML 2006*, pages 729–736, 2006.
27. S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. 2003.
28. J. Schreiter, D. Nguyen-Tuong, M. Eberts, B. Bischoff, H. Markert, and M. Toussaint. Safe exploration for active learning with gaussian processes. In *ECML*, 2015.
29. M. Turchetta, F. Berkenkamp, and A. Krause. Safe exploration in finite markov decision processes with gaussian processes. In *NIPS*, pages 4305–4313, 2016.
30. G. Ye and R. Alterovitz. Demonstration-guided motion planning. In *ISRR*, 2011.

A Analysis

A brief overview of the most important results in this section:

- Theorem A.1 shows that all states that can be guaranteed unsafe must lie within some distance to the boundary of the unsafe set. Corollary A.1 shows that the set of guaranteed unsafe states shrinks to a subset of the boundary of the unsafe set when using a continuous demonstration directly to learn the constraint.
- Corollary A.2 shows that for the discrete time case and the continuous, non-discretized case, our estimate of the unsafe set is a guaranteed under-approximation of the true unsafe set if the unsafe set is sufficiently “thick”.
- For continuous trajectories that are then discretized, Theorem A.3 shows us that the guaranteed unsafe set can be made to contain states on the interior of the unsafe set, but at the cost of potentially labeling states within some distance outside of the unsafe set as unsafe as well.

For convenience, we repeat the definitions, along with some illustrations for the sake of visualization.

A.1 Learnability

In this section, we will provide analysis on the learnability of unsafe sets, given the known constraints and cost function. Most of the analysis will be based off unsafe sets defined over the state space, i.e. $\mathcal{A} \subseteq \mathcal{X}$, but we will extend it to the feature space in Corollary A.4. If a state x can be learned to be guaranteed unsafe, then we denote that $x \in \mathcal{A}_l$, where \mathcal{A}_l is the set of all states that can be learned guaranteed unsafe.

We begin our analysis with some notation.

Definition A.1 (Signed distance). *Signed distance from point $p \in \mathbb{R}^m$ to set $\mathcal{S} \subseteq \mathbb{R}^m$, $sd(p, \mathcal{S}) = -\inf_{y \in \partial \mathcal{S}} \|p - y\|$ if $p \in \mathcal{S}$; $\inf_{y \in \partial \mathcal{S}} \|p - y\|$ if $p \in \mathcal{S}^c$.*

The following theorem describes the nature of \mathcal{A}_l :

Theorem A.1 (Learnability (discrete time)). *For trajectories generated by a discrete time dynamical system satisfying $\|x_{t+1} - x_t\| \leq \Delta x$ for all t , the set of learnable guaranteed unsafe states is a subset of the outermost Δx shell of the unsafe set: $\mathcal{A}_l \subseteq \{x \in \mathcal{A} \mid -\Delta x \leq sd(x, \mathcal{A}) \leq 0\}$.*

Proof. Consider the case of a length T unsafe trajectory $\xi = \{x_1, \dots, x_N\}$, $x_1 \in \mathcal{A} \vee \dots \vee x_T \in \mathcal{A}$. For a state to be learned guaranteed unsafe, $T - 1$ states in ξ must be learned safe. This implies that regardless of where that unsafe state is located in the trajectory, it must be reachable from some safe state within one time-step. This is because if multiple states in ξ differ from the original safe trajectory ξ^* , to learn that one state is unsafe with certainty means that the others should be learned safe from some other demonstration. Say that $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_T \in \mathcal{S}$, i.e. they are learned safe. Since $(\|x_{i+1} - x_i\| \leq$

$\Delta x) \wedge (\|x_i - x_{i-1}\| \leq \Delta x)$ and $x_{i-1}, x_{i+1} \in \mathcal{S}$, x_i must be within Δx of the boundary of the unsafe set: $-\min_{y \in \partial \mathcal{A}} \|x_i - y\| \geq \Delta x$, implying $-\Delta x \leq \text{sd}(x_i) \leq 0$.

□

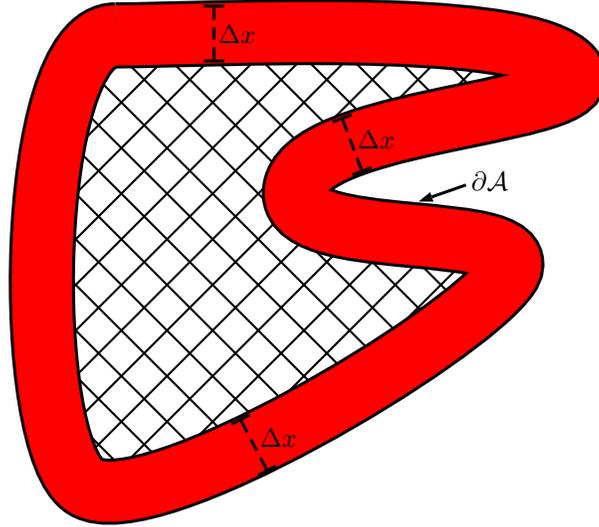


Fig. 7: Illustration of the outermost Δx shell (shown in red) of the unsafe set \mathcal{A} . The hatched area cannot be learned guaranteed safe.

Remark. For linear dynamics, Δx can be found exactly via

$$\underset{x \in \mathcal{X}, u \in \mathcal{U}}{\text{maximize}} \quad \|Ax + Bu - x\| \quad (10)$$

where convexity depends on the convexity of \mathcal{X} and \mathcal{U} .

In the case of general dynamics, an upper bound on Δx can be found via

$$\Delta x \leq \sup_{x \in \mathcal{X}, u \in \mathcal{U}, t \in [t_0, T]} \|f(x, u, t)\| \quad (11)$$

Corollary A.1 (Learnability (continuous time)). For continuous trajectories $\xi(\cdot) : [0, T] \rightarrow \mathcal{X}$, the set of learnable guaranteed unsafe states shrinks to the boundary of the unsafe set: $\mathcal{A}_l \subseteq \{x \in \mathcal{A} \mid \text{sd}(x, \mathcal{A}) = 0\}$.

Proof. The output trajectory of a continuous time system can be seen as the output of a discrete time system in the limit as the time-step is taken to 0. In this case, as long as the dynamics are locally Lipschitz continuous, $\Delta x \doteq \lim_{\Delta t \rightarrow 0} \|x(t + \Delta t) - x(t)\| \rightarrow 0$ [17], and via Theorem 1, the corollary is proved. □

It is worth noting that depending on the cost function chosen, \mathcal{A}_l can become arbitrarily small; in other words, some cost functions are more informative than

others in recovering a constraint. An interesting avenue of future work is to investigate the properties of cost functions that enable more to be learned about the constraints and how this knowledge can help inform reward (or cost) shaping.

A.1.1 Learnability (dynamics) Depending on the dynamics of the system, it may be impossible to obtain sub-trajectories with few perturbed waypoints from sampling, due to there only being one feasible control sequence that takes the system from a start to a goal state. We formalize this intuition in the following theorem:

Definition A.2 (Forward reachable set). *The forward reachable set $FRS(x_s, \mathcal{U}, T_1, T_2)$ is the set of all states that a dynamical system can reach at time $t = T_2$ starting from x_s at time $t = T_1$, using controls drawn from an admissible set of controls \mathcal{U} :*

$$FRS(x_s, \mathcal{U}, T_1, T_2) \doteq \{z \in \mathcal{X} \mid \exists u(t) : [T_1, T_2] \rightarrow \mathcal{U}, x_{T_1} = x_s, x_{T_2} = z\} \quad (12)$$

Theorem A.2 (Learnability (dynamics)). *Let x_1^*, \dots, x_M^* be consecutive waypoints on a safe trajectory ξ^* at times t_1, \dots, t_M , with time discretization Δt_i between states x_i^* and x_{i+1}^* , where all but x_1^*, x_M^* are free to move. Then, a necessary condition for being able to sample unsafe trajectories is that $\exists x_2 \in FRS(x_1^*, \mathcal{U}, t_1, t_1 + \Delta t_1), \dots, \exists x_{M-1} \in FRS(x_{M-2}, \mathcal{U}, t_{M-2}, t_{M-2} + \Delta t_{M-2}), x_M^* \in FRS(x_{M-1}, \mathcal{U}, t_{M-1}, t_{M-1} + \Delta t_{M-1})$ such that $\exists i \in \{2, \dots, M-1\} : x_i^* \neq x_i$: i.e. there exists at least one state that the dynamics allow to be moved from the demonstrated trajectory.*

Proof. Proof by contradiction. Assume that there does not exist an $i \in \{2, \dots, M-1\}$ such that $x_i \neq x_i^*$. Then, there exists no alternate sequence of controls taking the system from x_1^* to x_M^* ; hence no trajectories satisfying the start and goal constraints can be satisfied.

Additionally, the same analysis can be used for continuous trajectories in the limit as the time-step between consecutive waypoints, Δt , goes to 0. \square

Remark. This implies that when the dynamics are highly restrictive, less of the unsafe set can be learned to be guaranteed unsafe, and the learnable subset of the Δx -shell of the unsafe set (as described in Theorem 1) can become small.

A.2 Conservativeness

For the analysis in this section, we will assume that the unsafe set has a Lipschitz boundary; informally, this means that $\partial\mathcal{A}$ can be locally described by the graph of a Lipschitz continuous function. A formal definition can be found in [11]. We define some notation:

Definition A.3 (Set thickness). *Denote the outward-pointing normal vector at a point $p \in \partial\mathcal{A}$ as $\hat{n}(p)$. Furthermore, at non-differentiable points on $\partial\mathcal{A}$, $\hat{n}(p)$ is replaced by the set of normal vectors for the sub-gradient of the Lipschitz function describing $\partial\mathcal{A}$ at that point [4]. The set \mathcal{A} has a thickness larger than d_{thick} if $\forall x \in \partial\mathcal{A}, \forall d \in [0, d_{thick}], sd(x - d\hat{n}(x), \mathcal{A}) \leq 0$.*

Definition A.4 (γ -offset padding). Define the γ -offset padding $\partial\mathcal{A}_\gamma$ as: $\partial\mathcal{A}_\gamma = \{x \in \mathcal{X} \mid \forall y \in \partial\mathcal{A} : x = y + d\hat{n}(y), d \in [0, \gamma]\}$.

Definition A.5 (γ -padded set). We define the γ -padded set of the unsafe set \mathcal{A} , $\mathcal{A}(\gamma)$, as the union of the γ -offset padding and \mathcal{A} : $\mathcal{A}(\gamma) \doteq \partial\mathcal{A}_\gamma \cup \mathcal{A}$.

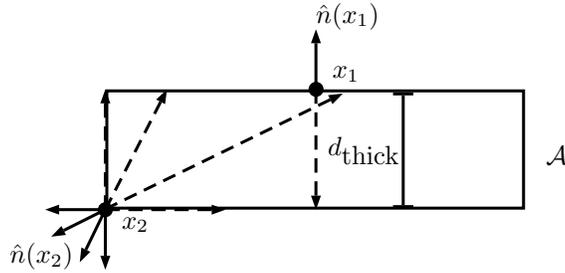


Fig. 8: Illustration of thickness, c.f. Definition 2.

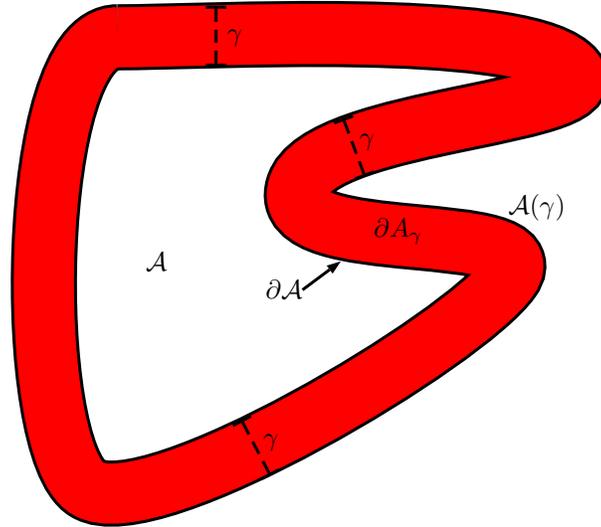


Fig. 9: Illustration of the γ -padded set $\mathcal{A}(\gamma)$, marked in red. The γ -offset padding is displayed in red. The original set \mathcal{A} is shown in white.

Corollary A.2 (Conservative recovery of unsafe set). A sufficient condition ensuring that the set of recovered guaranteed unsafe states $\mathcal{A}_1^{\text{rec}}$ is contained in \mathcal{A} is that \mathcal{A} has a set thickness greater than or equal to Δx (c.f. Definition A.1).

Proof. Via Theorem 1, our method will not determine that any state further inside than the outer Δx -shell is unsafe for sure. If \mathcal{A} has thickness at least Δx , then our method will only determine states that are within the unsafe set to be

guaranteed unsafe. This holds for discrete time dynamics and continuous time dynamics as well as $\Delta x \rightarrow 0$. \square

Note that if we deal with continuous trajectories directly, the guaranteed learnable set shrinks to a subset of the boundary of the unsafe set, $\partial\mathcal{A}$. However, if we discretize these trajectories, we can learn unsafe states lying in the interior, at the cost of conservativeness holding only for a padded unsafe set.

Theorem A.3 (Continuous-to-discrete time conservativeness). *Let ξ be a continuous trajectory: $\xi : [0, T] \rightarrow \mathcal{X}$. The system dynamics are described by $\dot{x} = f(x, u, t)$. The trajectory is discretized in time, potentially non-uniformly, resulting in a discretized trajectory $\{\xi(t_1), \dots, \xi(t_N)\}, t_i \in [0, T]$ for all i . Assume the maximum discretization time is Δt_{max} . Denote*

$$f_{\Delta x}^{max} \doteq \sup_{x \in \mathcal{X}, u \in \mathcal{U}, t' \in [t, t + \Delta t_{max})} \|f(x, u, t')\| \Delta t_{max} \quad (13)$$

Then, our method recovers a subset of the $3f_{\Delta x}^{max}$ -padded unsafe set, $\mathcal{A}(3f_{\Delta x}^{max})$.

Proof. Consider a sequence of discretized states on an unsafe, continuous trajectory $\xi : [0, T] \rightarrow \mathcal{X}$: $\xi(t_1), \dots, \xi(t_M)$. Any state $\xi(t_i), t_i \in [0, T]$ could have caused the trajectory to be unsafe.

Forgetting about the continuous nature of the trajectory, say we consider just the discretized states and learn that one discretized state $\xi(t_i)$ is guaranteed unsafe. Since the trajectory is continuous, any $\xi(t_j), t_j \in (\xi(t_{i-1}), \xi(t_{i+1}))$ could have been unsafe. From the dynamics, the distance d from that true unsafe state $\xi(t_u)$ to the unsafe state $\xi(t_i)$ is bounded by $\max(f_{\Delta x_{i-1}}, f_{\Delta x_i})$, where $f_{\Delta x_i} \doteq \sup_{x \in \mathcal{X}, u \in \mathcal{U}, t \in [t_i, t_{i+1})} \|f(x, u, t)\| (t_{i+1} - t_i)$. Furthermore, the set of unsafe states that may be obtained by moving state $\xi(t_i)$ is contained by the open ball $\mathcal{B}_{r=f_{\Delta x_{i-1}} + f_{\Delta x_i}}(\xi(t_i))$. Then, for this trajectory, our estimate is conservative if we pad the unsafe set by $f_{\Delta x_{i-1}} + f_{\Delta x_i} + \max(f_{\Delta x_{i-1}}, f_{\Delta x_i})$.

Lastly, for this to guarantee to hold for all possible trajectories with maximum discretization time Δt_{max} , the previous expression can be upper bounded by $3f_{\Delta x}^{max}$. \square

Corollary A.3 (Continuous-to-discrete time and space conservativeness). *Let the largest grid cell in the constraint space be contained by a ball of radius r_z . Then, if trajectories are discretized both in space and time, our method recovers a subset of the $2r_z + 3f_{\Delta x}^{max}$ -padded unsafe set, $\mathcal{A}(2r_z + 3f_{\Delta x}^{max})$.*

Proof. From Theorem A.3, the unsafe set estimate is a subset of the $3f_{\Delta x}^{max}$ -expanded estimate in the continuous space case. Since gridding enforces all continuous states lying within the cell have the same occupancy and the largest grid is contained within a ball of radius r_z , further padding the unsafe set by $2r_z$ ensures the unsafe set is conservative for states learned unsafe due to grid consistency. \square

Corollary A.4 (Continuous-to-discrete feature space conservativeness).

Let the feature mapping $\phi(x)$ from the state space to the constraint space be Lipschitz continuous with Lipschitz constant L . Then, our method recovers a subset of the $3Lf_{\Delta x}^{\max}$ -padded unsafe set in the feature space, $\mathcal{A}(3Lf_{\Delta x}^{\max})$.

Proof. From the definition of Lipschitz continuity, $\|\phi(x) - \phi(y)\| \leq L\|x - y\|$. From Theorem A.3, the unsafe set estimate is a subset of the $3f_{\Delta x}^{\max}$ -expanded estimate in the continuous space case. Using Lipschitz continuity, the value in the feature can at most change by $3Lf_{\Delta x}^{\max}$ from the boundary of the true constraint set to the boundary of the padded set; hence, the statement holds. \square

B Experimental details

Figure	Dynamics	Ctrl. constraints	Cost function
Fig. 4, Row 1	$x_{t+1} = x_t + u_t$	$\ u_t\ \leq 0.5$	$\sum_{t=1}^{T-1} \ u_t\ _2^2$
Fig. 4, Row 2	$x_{t+1} = Ax_t + Bu_t$, $A \doteq \exp\left(\text{diag}\left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\right)\right)$, $B \doteq \int_0^1 \exp(A\tau) d\tau \begin{bmatrix} 0 & 1 & 0 & 1 \end{bmatrix}^\top$	$ u_t \leq [20 \ 10]^\top$	$\sum_{i=1}^{T-1} \ x_i - x_{i+1}\ _2^2$
Fig. 4, Row 3	$\dot{x} = \begin{bmatrix} \cos(\theta) & \sin(\theta) & u \end{bmatrix}^\top$	$ u \leq 1$	$\sum_i \tau_{u_i}$
Fig. 5	$\dot{x} = \begin{bmatrix} \cos(\theta) & \sin(\theta) & u \end{bmatrix}^\top$	$ u \leq 1$	$\sum_i \tau_{u_i}$
Fig. 6	$\dot{x} = \begin{bmatrix} \cos(\theta) & \sin(\theta) & u \end{bmatrix}^\top$	$ u \leq 1$	$\sum_i \tau_{u_i}$

Table 3: Dynamics, control constraints, and cost functions used in experiments.

Figure	Timing (sampling trajectories)	Timing (constraint recovery)
Fig. 4, Row 1	11.5 min	3 min
Fig. 4, Row 2	4.5 min	4.5 min
Fig. 4, Row 3	2 hrs	4 min
Fig. 5	1 hr	2 min
Fig. 6	30 min	4 min

Table 4: Approximate runtime.

Here, $\sum_i \tau_{u_i}$ is the total time duration of applied control input (i.e. the time it took to go from start to goal). All experiments were conducted on a 4-core 2017 Macbook Pro with a 3.1 GHz Core i7. All code was implemented in MATLAB.

	Fig. 4, Row 1	Fig. 4, Row 2	Fig. 4, Row 3	Fig. 5	Fig. 6
Space discretization	0.1	0.25	0.5	1	1
Number of trajectories	300000	150000	10000	10000	10000
ε	n/a	n/a	10^{-3}	10^{-3}	10^{-3}
$\hat{\varepsilon}$	n/a	n/a	10^{-2}	10^{-2}	10^{-2}
α_c	10^{10}	10^4	1	1	1
Minimum \mathcal{L} length	n/a	n/a	10^{-10}	10^{-10}	10^{-10}
$f_{\Delta x}^{\max}$	n/a	n/a	4.53	6	6.96

Table 5: Parameters for each experiment.