

Probabilistic Specification Learning for Planning with Safety Constraints

Kandai Watanabe¹, Nicholas Renninger², Sriram Sankaranarayanan¹, and Morteza Lahijanian¹.

Abstract—This paper proposes a framework for learning task specifications from demonstrations, while ensuring that the learned specifications do not violate safety constraints. Furthermore, we show how these specifications can be used in a planning problem to control the robot under environments that can be different from those encountered during the learning phase. We formulate the specification learning problem as a grammatical inference problem, using probabilistic automata to represent specifications. The edge probabilities of the resulting automata represent the demonstrator’s preferences. The main novelty in our approach is to incorporate the safety property during the learning process. We prove that the resulting automaton always respects a pre-specified safety property, and furthermore, the proposed method can easily be included in any Evidence-Driven State Merging (EDSM)-based automaton learning scheme. Finally, we introduce a planning algorithm that produces the most desirable plan by maximizing the probability of an accepting trace of the automaton. Case studies show that our algorithm learns the true probability distribution most accurately while maintaining safety. Since, specification is detached from the robot’s environment model, a satisfying plan can be synthesized for a variety of different robots and environments including both mobile robots and manipulators.

I. INTRODUCTION

Autonomous robots make decisions against varying environments while maintaining key safety properties in domains such as deep-sea and space exploration; assistive and service domains (e.g., human-shared assembly lines); and surgical robotics. As robots are getting more autonomous, they are also expected to perform more complex tasks. Often times, however, the *explicit and precise specifications* of such tasks are unavailable, or too cumbersome to be provided by non-experts. However, these tasks can be easily demonstrated through various means – demonstrations through human operation or data collected from past manual operations. Given such demonstrations, the robot needs to be able to infer the task specification, so that it may autonomously satisfy the specification even against varying environments. In this regard, we identify four key challenges: (a) identifying a rich specification formalism that permits efficient and precise learning algorithms from the given demonstrations; (b) ensuring that key safety properties are satisfied by the resulting




Fig. 1: Schematic of an autonomous deep-sea science mission.

specification; (c) capturing “operator preferences” or “hidden costs” from the demonstrations; and (d) using the specifications to guide autonomous behavior under environments that may be different from those under which the demonstrations were provided. In this paper, we propose solutions to all four problems using the framework of probabilistic automata learned using grammatical inference. We focus on *safety-property constrained learning*; wherein the final learned specification must satisfy the safety properties. Our proposed learning approach ensures that the safety properties constrain the intermediate steps of the learning algorithm as opposed to an “after-the-fact” approach wherein the safety properties are intersected with the final specifications to rule out any unsafe behavior. We show that our safety-enabled learning approach clearly outperforms the after-the-fact approach.

For example, consider an underwater robot in a deep-sea science mission as schematically illustrated in Figure 1. The scientists want the robot to explore a shipwreck on the sea-floor, study the behavior of a school of fish, while keeping away from the dangerous coral reefs. The goal of this work is to enable autonomous execution of this task by just exposing the robot to demonstrations from tele-operations or from related past missions. From these demonstrations, the robot should be able to infer the task and generate the necessary plan to accomplish it. We call this problem *Specification Learning from Demonstrations*, which can be regarded as a new form of *Learning from Demonstration* (LfD) [1].

Many LfD studies focus on either learning a policy or a reward structure [1], [2]. For policy learning, techniques such as *reinforcement learning* (RL) [3] and Dynamic Movement Primitives [4], [5] are typically used to learn a function that maps agent states to actions. In reward learning, a scalar reward function that maps agent states to rewards is learned via, e.g., *inverse reinforcement learning* (IRL) [6]–[9], to later train a policy on an agent. While very powerful, these methods learn a function that is specific to the environment (and robot) model used during training and hence are fragile to the changes to those models. More importantly, those methods are restricted to Markovian tasks. For example, the

This work was supported in part by the University of Colorado Boulder Autonomous Systems Interdisciplinary Research Theme; NSF award numbers 1932189 and 1836900; JASSO Overseas Graduate Fellowship and Keio University Global Fellowship.

¹Authors are with the Departments of Computer Science and Aerospace Engineering Sciences at University of Colorado Boulder, Boulder, Colorado. {first.lastname}@colorado.edu

²Author is with MITRE Corporation. This work was done while he was a student in the Dept. of Aerospace Engineering Sciences at University of Colorado Boulder. nicholas.renninger@colorado.edu

robot task in Figure 1, which requires a visit to both the shipwreck and school of fish in any order, is non-Markovian. To achieve it, the robot has to maintain a memory of previous locations to decide the next location to visit. Hence, it is important to enable LfD for non-Markovian tasks, but such an extension is nontrivial for the existing LfD methods.

An alternative approach to expressing tasks is to use formal languages such as *linear temporal logic* (LTL) [10], which is widely used in formal verification and increasingly employed in robotics in recent years, e.g., [11]. Such languages enable formal expression of rich missions, including non-Markovian tasks [12] as well as *liveness* (“something good eventually happens”) and *safety* (“something bad never happens”) requirements. Other important benefits of formal languages is in their ease of interpretability and flexibility to compose multiple specifications. Such benefits have even led to their use in RL, e.g., [13]–[15]. Nevertheless, writing correct formal specifications requires domain knowledge.

In recent years, a new line of research has emerged with a focus on learning formal specifications from data [12], [16]–[19]. Most work has been concerned with learning temporal logic formulas with the purpose of classification and prediction from user data (in the supervised learning sense) [17], [18] or interpretation and planning for tasks [19]. Those studies restrict the exploration problem to a set of formula templates provided *a priori*. Recent work [12] overcomes this restriction by iterating over all combinations of formulas. The method is based on maximum a posterior learning and can account for noisy samples. It however is slow due to the large space of exploration for formulas. Another important issue with formula learning methods for the purpose of planning is that they typically need to be translated to an automaton, which could lead to the *state-explosion* problem [10], [11]. Work [20] overcomes this issue by directly learning a *Deterministic Finite Automaton* (DFA). They however assume the structure of the DFA is known and only learn the transitions between the DFA states while an oracle labels each sample with DFA states.

In this work, we propose a new approach to task specification learning where we infer formal task specifications as probabilistic automata. Our approach used ideas from the field of *grammatical inference* (GI) [21], by modifying existing “evidence-driven state merging” algorithms in GI to learn the specification as a *Probabilistic Deterministic Finite Automaton* (PDFA). We further extend this method to incorporate safety properties during the learning process, so that all runs of the final PDFA satisfy the safety properties. Furthermore, we propose a planning algorithm with the inferred PDFA that generates the most preferred plan to achieve the task for any robot that can be abstracted to a deterministic transition system.

To the best of our knowledge, this is the first work that employs PDFA as a learning model for task specifications in robotics. The key technical contributions include: (a) the derivation of the safety guaranteed PDFA learning algorithm that is compatible with any EDSM techniques; (b) a planning algorithm with a learned PDFA that can handle varying

environments; and (c) a set of experiments that show the efficacy of the proposed algorithms in both mobile and manipulator robots. We also provide a comparison case study against [12] to highlight the similarities and differences between the two approaches.

II. PROBLEM FORMULATION

In this section, we describe and formalize our problem of task specification learning with preferences and safety constraints. Our aim is to use an interpretable learning model that we can use for planning for a robot to perform the task according to the user’s preferences.

A. Task Specifications and demonstrations

We assume the demonstrator has a task in mind that needs to be achieved in finite time. Our goal is to learn this task in the form of a deterministic finite state automaton (DFA) that reflects the overall goals of the demonstrator. Such automata accommodate a large class of tasks, including the ones that can be specified using co-safe LTL [22] and *LTL over finite traces* [23] formulas. Nevertheless, to avoid a trivial solution to the problem (e.g., a trivial specification that accepts all possible behaviors), we consider two important caveats.

- We specify, on the side, a safety property specification that the robot must not violate. The safety property characterizes a potentially infinite set of negative examples for our learner consisting of those specifications that violate the safety property.
- We want to encode the demonstrator’s *preferences* between various ways of satisfying an intended specification. For instance, a demonstrator may prefer to avoid a collision with an obstacle by *steering left* preferentially since it may position the vehicle to avoid a future collision with less effort.

Below, we define the mathematical machinery needed to formalize this problem.

For a given task φ_{task} , let $\Pi = \{p_1, \dots, p_k\}$ be a set of atomic propositions (boolean predicates) that represent important facts about the world that is relevant to the task. The set of atomic propositions that is true at a state of the world is called a *symbol*. The set of all symbols is denoted by Σ . Note that $\Sigma = 2^\Pi$. A demonstration ω (also known as a *trace* or *word*) is a finite sequence of symbols, i.e., $\omega = \omega_1\omega_2 \dots \omega_n$, where $\omega_i \in \Sigma$ for all $1 \leq i \leq n$.

From demonstrations, we want to learn task φ_{task} . However, instead of directly learning it as a temporal logic formula, our goal is to learn it in the form of a DFA.

Definition 1 (DFA). A *deterministic finite automaton (DFA)* is a tuple $\mathcal{A} = (Q, \Sigma, q_0, \delta, F)$, where

- Q is a finite set of states,
- Σ is a finite set of input symbols,
- $q_0 \in Q$ is the initial state,
- $\delta : Q \times \Sigma \rightarrow Q$ is the transition function, and
- $F \subseteq Q$ is the set of final or accepting states.

The transition function δ can be also viewed as a relation $\delta \subseteq Q \times \Sigma \times Q$, where every transition is a tuple $(q, \sigma, q') \in \delta$

iff $q' = \delta(q, \sigma)$, where $\sigma \in \Sigma$. A *run* of a DFA \mathcal{A} on trace $\omega = \omega_1\omega_2\dots\omega_n$ is a sequence of states $z = z_0z_1\dots z_n$, where $z_0 = q_0$ and $z_i = \delta(z_{i-1}, \omega_i)$ for $i = 1, \dots, n$. A run z is called *accepting* if $z_n \in F$. A trace ω is accepted by \mathcal{A} if it induces an accepting run. The set of all traces that are accepted by DFA \mathcal{A} is called the language of \mathcal{A} and is denoted by $\mathcal{L}(\mathcal{A})$.

We call a demonstration ω valid if it is accepted by the DFA that represents φ_{task} , i.e., achieves task φ_{task} .

Definition 2 (Valid Demonstration). *A valid demonstration for task φ_{task} is a finite trace $\omega \in \Sigma^*$ such that $\omega \in \mathcal{L}(\mathcal{A}_{\varphi_{\text{task}}})$, where $\mathcal{A}_{\varphi_{\text{task}}}$ is the DFA that represents φ_{task} .*

B. Safety Specification

To express the safety constraints, we use safe LTL [22].

Definition 3 (Safe Syntax). *A syntactically safe LTL formula over Π is recursively defined as*

$$\varphi := p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathcal{X}\varphi \mid \mathcal{G}\varphi$$

where $p \in \Pi$, \neg (negation), \vee (disjunction), and \wedge (conjunction) are boolean operators, and \mathcal{X} (“next”) and \mathcal{G} (“globally”) are temporal operators.

Safe LTL formulas reason over infinite traces, but finite traces are sufficient to violate them [22]. We denote the set of finite traces that violate safety formula φ_{safe} by $\mathcal{L}(\neg\varphi_{\text{safe}})$. Hence, given valid demonstrations and safety formula φ_{safe} , we want to learn a DFA $\tilde{\mathcal{A}}_{\varphi_{\text{task}}}$ such that it does not accept any trace that violates safety, i.e., $\mathcal{L}(\tilde{\mathcal{A}}_{\varphi_{\text{task}}}) \cap \mathcal{L}(\neg\varphi_{\text{safe}}) = \emptyset$.

C. Preferences

On the learned DFA, we also want to encode the preferences of the demonstrator. We assume that the preferences are correlated with the number of demonstrations of the same trace. Hence, preferences can be quantified as weights over traces and normalized over the entire language, which can be viewed as a probability distribution over $\mathcal{L}(\mathcal{A}_{\varphi_{\text{task}}})$. The higher the probability of an accepting trace is, the more preferred it is to the demonstrator. Hence, our aim is to learn a probabilistic DFA (PDFA), which captures both the accepting traces and their corresponding probabilities.

Definition 4 (PDFA). *A probabilistic DFA (PDFA) is a tuple $\mathcal{A}^{\mathbb{P}} = (\mathcal{A}, \delta_{\mathbb{P}}, F_{\mathbb{P}})$, where \mathcal{A} is a DFA, and $\delta_{\mathbb{P}} : \delta \rightarrow [0, 1]$ assigns a probability to every transition in δ such that $\sum_{\sigma \in \Sigma} \delta_{\mathbb{P}}(q, \sigma, \delta(q, \sigma)) = 1$ for every $q \in Q$, and $F_{\mathbb{P}} : Q \rightarrow [0, 1]$ assigns a probability of terminating at each state, where $F_{\mathbb{P}}(q) = 0$ if $q \notin F$.*

Consider trace $\omega = \omega_1\omega_2\dots\omega_n$ and its induced run $z = z_0z_1\dots z_n$ on PDFA $\mathcal{A}^{\mathbb{P}}$. The probability of ω is given by

$$P(\omega) = \prod_{i=1}^n \delta_{\mathbb{P}}(z_{i-1}, \omega_i, z_i) \cdot F_{\mathbb{P}}(z_n).$$

We say $\mathcal{A}^{\mathbb{P}}$ accepts ω iff $P(\omega) > 0$, and the demonstrator prefers ω over $\omega' \in (2^{\Pi})^*$ iff $P(\omega) > P(\omega')$. The language

of $\mathcal{A}^{\mathbb{P}}$ is the set of traces with non-zero probabilities, i.e.,

$$\mathcal{L}(\mathcal{A}^{\mathbb{P}}) = \{\omega \in (2^{\Pi})^* \mid P(\omega) > 0\}.$$

Therefore, our goal becomes to learn a $\tilde{\mathcal{A}}^{\mathbb{P}}_{\varphi_{\text{task}}}$ that captures the task φ_{task} , the demonstrator’s preferences (probability distribution over traces), and never violates φ_{safe} .

D. Robot Model and Plans

Once a specification is learned, we want to synthesize a plan for a robot that can realize the specification. To do so, we assume that we are given an abstraction model of the robot as a deterministic transition system. Such an abstraction is commonly used and constructed in formal approaches to both mobile robotics [11], [24], [25] and robotic manipulators [26], [27].

Definition 5 (DTS). *A deterministic transition system (DTS) is a tuple $\mathcal{T} = (X, A, x_0, \delta_{\mathcal{T}}, \Pi, L)$, where*

- X is a finite set of states,
- A is a finite set of controls or actions,
- $x_0 \in X$ is the initial state,
- $\delta_{\mathcal{T}} : X \times A \rightarrow X$ is the (partial) transition function,
- Π is a finite set of atomic propositions (predicates), and
- $L : X \rightarrow 2^{\Pi}$ is a labeling function that maps each state to the set of predicates that are true at that state.

A *plan* $\gamma = \gamma_0\gamma_1\dots\gamma_{n-1}$ is a sequence of actions, where $\gamma_i \in A$ for all $0 \leq i \leq n-1$. By executing γ , the robot generates a *trajectory* $s = s_0s_1\dots s_n$, where $s_0 = x_0$ and $s_{i+1} = \delta_{\mathcal{T}}(s_i, \gamma_i)$. A *valid plan* is plan γ that respects the transition function $\delta_{\mathcal{T}}$, i.e., $\delta_{\mathcal{T}}(s_i, \gamma_i)$ exists for all $0 \leq i \leq n-1$. We denote the set of all valid plans by Γ .

The observation trace of the above trajectory is the sequence of observed labels, i.e., $\rho_{\gamma} = L(s_0)L(s_1)\dots L(s_n)$. We refer to the set of all observation traces that a robot can generate as the language of \mathcal{T} , i.e., $\mathcal{L}(\mathcal{T}) = \{\rho_{\gamma} \mid \gamma \in \Gamma\}$.

In the planning problem, we are interested in a plan γ that generates an observation trace ρ_{γ} that achieves task φ_{task} and is the most preferred behavior, i.e.,

$$\rho_{\gamma} = \arg \max_{\omega \in \mathcal{L}(\varphi_{\text{task}}) \cap \mathcal{L}(\mathcal{T})} P(\omega).$$

E. Problem

We are now able to formally define our problem.

Problem. *Given demonstrations $\Omega = [\omega_i]_{i=1}^{n_{\Omega}}$ that are sampled from a hidden task specification DFA $\mathcal{A}_{\varphi_{\text{task}}}$ according to some preferences (probability distribution) of the demonstrator and a safety specification φ_{safe} :*

- 1) *learn φ_{task} and the demonstrator’s preferences as a PDFA $\mathcal{A}^{\mathbb{P}}_{\varphi_{\text{task}}}$ such that safety is never violated, i.e., $\mathcal{L}(\tilde{\mathcal{A}}^{\mathbb{P}}_{\varphi_{\text{task}}}) \cap \mathcal{L}(\neg\varphi_{\text{safe}}) = \emptyset$;*
- 2) *furthermore, given a DTS robot model \mathcal{T} , compute a plan for the robot \mathcal{T} that generates the most preferred behavior that satisfies φ_{task} , i.e., generates the trace with the highest probability in $\mathcal{L}(\tilde{\mathcal{A}}^{\mathbb{P}}_{\varphi_{\text{task}}}) \cap \mathcal{L}(\mathcal{T})$.*

For Problem 1, we use grammatical inference [21] while incorporating the safety property during the learning process,

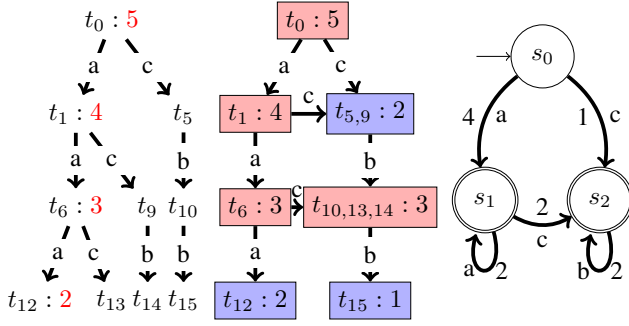


Fig. 2: Schematic illustration of evidence-driven state merging (EDSM) algorithm. **(Left)** A frequency prefix tree acceptor (FPTA) is constructed from the given demonstrations with frequencies greater than 1 shown in red; **(Middle)** intermediate automaton as states are merged according to criteria that differ across various algorithms with frequencies shown at each node; and **(Right)** the final frequency DFA (FDFA) that is learned is shown in red.

as described in Section III. We use this PDFA to solve Problem 2 as detailed in Section IV.

III. SAFETY GUARANTEED PDFA LEARNING

In this section, we explain how a PDFA can be learned from demonstrations and present our method that can embed safety specification to guarantee safety on the outcome. We first show a general PDFA learning algorithm, and then we describe our algorithms to incorporate safety.

A. Grammatical Inference: PDFA Learning

PDFA learning has been extensively studied as part of *grammatical inference* (GI) with existing algorithms such as ALERGIA, DSAI, and MDI, that can learn PDFAs from unlabeled demonstrations [21]. These algorithms are all based on a principle called *evidence-driven state-merging* (EDSM). At a high level, EDSM approaches find an appropriate structure for an automaton $\mathcal{A}^{\mathbb{P}}$ and simultaneously estimate the probability distribution parameters $F_{\mathbb{P}}$ and $\delta_{\mathbb{P}}$ given a set of sample traces. This is achieved by first constructing a large (prefix) tree from the samples, and repeatedly merging the states of the tree in order to form an automaton that is as simple as possible while continuing to accept the sample traces from the demonstration. The various algorithms (e.g., ALERGIA and MDI) differ on what states are merged.

Figure 2 shows a general scheme for an EDSM-based algorithm for learning a PDFA. The initial step is to construct a *frequency prefix tree acceptor* (FPTA) from the traces in Ω (Fig. 2-Left) and then incrementally merge states of the FPTA, two at a time, based on a *compatibility* criterion that varies depending on the actual algorithm. As two states are merged, so are their subtrees in the FPTA (Fig. 2-Middle). The nodes of the intermediate automata are variously colored red/blue using a coloring scheme to influence how states are selected for merging. Furthermore, algorithms also maintain frequencies alongside the nodes based on the number of demonstration traces that reach a particular node. These frequencies are also combined during the state merging process. The final result is a *frequency DFA* (FDFA) wherein frequencies along edges indicate how often they are taken

by a demonstration (Fig. 2-Right). The frequencies of all outgoing edges are normalized to yield a distribution.

The various PDFA learning algorithms such as ALERGIA or MDI differ on how they implement the *compatibility* check for whether two given nodes can be merged. For instance, the ALERGIA algorithm implements a statistical test based on frequencies to compare if two states are compatible, whereas the MDI approach first temporarily merges two states and their subtrees, while accepting the merge if a metric computed on automaton after the merge is smaller than that before the merge. We assume that the basic PDFA learning algorithm as a given, and our goal is to learn while respecting a safety property.

B. Learning with Safety Specification

We now consider two different approaches for learning with a safety specification. The first method is a *post-processing* technique that simply runs the PDFA learning algorithm on the given demonstration traces and then subsequently intersects the resulting PDFA with the automaton for the safety property. The second method incorporates the safety specification during the learning process by modifying the EDSM algorithm. In particular, the merges are defined so that the result continues to satisfy the safety specifications.

1) *Post-process Algorithm*: From φ_{safe} , we first construct a DFA $\mathcal{A}_{\neg\varphi_{\text{safe}}}$ that accepts precisely all those traces that violate the safety property [22]. Then, by complementing $\mathcal{A}_{\neg\varphi_{\text{safe}}}$, we obtain $\mathcal{A}_{\text{safe}} = (Q^s, \Sigma, q_0^s, \delta^s, F^s)$ that accepts all the traces that do not violate φ_{safe} . Let $\mathcal{A}^{\mathbb{P}} = (Q, \Sigma, q_0, \delta, F, \delta_{\mathbb{P}}, F_{\mathbb{P}})$ be the PDFA learned from the given demonstration traces without considering the safety property. We intersect the languages of $\mathcal{A}^{\mathbb{P}}$ and $\mathcal{A}_{\text{safe}}$ by constructing a product automaton $\mathcal{A}_{\text{safe}}^{\mathbb{P}} = \mathcal{A}^{\mathbb{P}} \otimes \mathcal{A}_{\text{safe}} = (Q_{\text{safe}}, \Sigma, q_{0,\text{safe}}, \delta_{\text{safe}}, F_{\text{safe}}, \delta_{\mathbb{P},\text{safe}}, F_{\mathbb{P},\text{safe}})$, where $Q_{\text{safe}} = Q \times Q^s$, $q_{0,\text{safe}} = (q_0, q_0^s)$, $F_{\text{safe}} = F \times F^s$,

$$\begin{aligned} \delta_{\text{safe}}((q, q^s), \sigma) &= (q', q^{s'}) \text{ if } q' = \delta(q, \sigma) \wedge q^{s'} = \delta^s(q^s, \sigma), \\ F_{\mathbb{P},\text{safe}}((q, q^s)) &= F_{\mathbb{P}}(q), \text{ and} \\ \delta_{\mathbb{P},\text{safe}}((q, q^s), \sigma, (q', q^{s'})) &= \begin{cases} \frac{\delta_{\mathbb{P}}(q, \sigma, q')}{N(q, q^s)} & \text{if } (q', q^{s'}) = \\ \delta_{\text{safe}}((q, q^s), \sigma) & \text{otherwise} \end{cases} \quad (1) \end{aligned}$$

where $N(q, q^s)$ is the normalizing function such that $\sum_{(\sigma, q'_{\text{safe}}) \in (\Sigma \times Q_{\text{safe}})} \delta_{\mathbb{P},\text{safe}}((q, q^s), \sigma, q'_{\text{safe}}) = 1$. The resulting PDFA is guaranteed to be safe due to the intersection of languages. However, this method of pruning (imposing safety) as a post-process step alters the probability distributions over the next-state transitions, since we remove the transitions that violate safety and renormalize the probability distribution at each state, as shown in (1). This overrides the probability distributions constructed by the original PDFA learning algorithm in an unpredictable manner. Therefore, while this method of imposing safety always succeeds, its probability distributions may not reflect the preferences embedded in the demonstrations accurately.

2) *Safety-Incorporated Learning Algorithm using “Pre-Processing”*: Whereas the *post-processing* approach enforces safety after the PDFA is learned, the *pre-processing* approach guarantees that the intermediate results also preserve safety, hence preventing alterations to the probability distributions due to unsafety. The main idea is to build the PDFA that generalizes the demonstrated traces but carries along with it information about how the generalization satisfies the safety property φ_{safe} at the same time in the form of a simulation relation with $\mathcal{A}_{\text{safe}}$.

Definition 6 (Simulation Relation). *A simulation relation R between two automata \mathcal{A} and \mathcal{B} is a relation between their states, $R \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$*

- (a) *Initial states of \mathcal{A} relate to the initial states of \mathcal{B} ;*
- (b) *If pair $(s, t) \in R$, where $s \in Q_{\mathcal{A}}$ and $t \in Q_{\mathcal{B}}$, and automaton \mathcal{A} can transition from s to $s' \in Q_{\mathcal{A}}$ on symbol σ , then there must exist a state $t' \in Q_{\mathcal{B}}$ such that automaton \mathcal{B} transitions from t to t' on the same symbol σ and $(s', t') \in R$;*
- (c) *For each $(s, t) \in R$, if s is final in \mathcal{A} then t must be final in \mathcal{B} .*

Theorem 1. *Let R be a simulation relation between automata \mathcal{A} and \mathcal{B} . It follows that $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$.*

The proof simply shows by induction that for any accepting run corresponding to an input trace ω in automaton \mathcal{A} from the initial state to a final state, there exists an accepting run in \mathcal{B} for the same trace ω from its initial state to the final state. The relation R allows us to construct such a run.

The key idea behind the *pre-process* approach is to maintain a simulation relation between the FDFFA and safety automaton $\mathcal{A}_{\text{safe}}$ at all intermediate states. The key is to restrict the merging of states so that we can guarantee that a simulation relation between the original automaton and $\mathcal{A}_{\text{safe}}$ before merging can be modified to yield a simulation relation between the merged automaton and $\mathcal{A}_{\text{safe}}$ afterwards.

Formally, we build the safety FPTA by augmenting the initial FPTA so that each state is now a tuple of the form (t_j, s_k) wherein t_j is a node in the original FPTA and s_k is the state in $\mathcal{A}_{\text{safe}}$ reached when the prefix that leads upto the state t_j is run through $\mathcal{A}_{\text{safe}}$. Thus, we ensure that every branch not only corresponds to a demonstration but also to a valid trace in $\mathcal{A}_{\text{safe}}$.

Let R be a relation between states of the FPTA and $\mathcal{A}_{\text{safe}}$ that contains all nodes (t_j, s_k) in the safety FPTA.

Lemma 1. *Assuming no demonstration trace violates the safety property φ_{safe} , then R is a simulation relation between the initial FPTA and the automaton $\mathcal{A}_{\text{safe}}$.*

We can represent any intermediate FDFFA state in the form (T, s) wherein T is a set of states from the initial FPTA, and s is state in $\mathcal{A}_{\text{safe}}$. Next, we modify the EDSM approach to allow a merge between two states (T_i, s_k) and (T_j, s_l) only if $s_k = s_l$. The result of the merge creates a state $(T_i \cup T_j, s_k)$.

Lemma 2. *Let \mathcal{A}_1 and \mathcal{A}_2 be the automata before and after an EDSM merge that is compatible with respect to the $\mathcal{A}_{\text{safe}}$*

states. Let R_1 be the relation between the states of \mathcal{A}_1 and those of $\mathcal{A}_{\text{safe}}$ that is a simulation relation. We can construct a simulation relation R_2 between the states of \mathcal{A}_2 and $\mathcal{A}_{\text{safe}}$.

Combining Lemmas 1 and 2, we conclude by induction on the number of merging steps that the final resulting PDFA must have a simulation relation to the safety automaton $\mathcal{A}_{\text{safe}}$. Since we have a simulation relation, we conclude that the language of the final resulting FDFFA and PDFA are contained in the that of $\mathcal{A}_{\text{safe}}$, i.e., the resulting PDFA does not accept a trace that violates φ_{safe} .

IV. PLANNING WITH PDFA

Once a task is learned as a PDFA, we are interested in synthesizing a plan that not only achieves the task but also respects the demonstrator’s preferences (Problem 2). Here, we introduce a method to do this for any robot with a DTS (abstract) model. We first construct a (product) graph that captures all the ways the robot can achieve the task by composing the DTS with the learned PDFA. Then, we reduce the optimal planning problem to a search on this graph.

The plans are computed over the product graph $\mathcal{P} = \mathcal{A}^{\mathbb{P}} \times \bar{\mathcal{T}}$, where $\bar{\mathcal{T}}$ is a DTS obtained by augmenting \mathcal{T} with a new initial state \bar{x}_0 with a transition to x_0 . Formally, $\bar{\mathcal{T}} = (\bar{X}, A, \bar{x}_0, \delta_{\bar{\mathcal{T}}}, \Pi, L)$, where $\bar{X} = X \cup \{\bar{x}_0\}$, and $\delta_{\bar{\mathcal{T}}}(x, a) = x_0$ if $x = \bar{x}_0$, otherwise $\delta_{\bar{\mathcal{T}}}(x, a) = \delta_{\mathcal{T}}(x, a) \forall a \in A$. This augmentation allows to correctly observe the label of x_0 and assign the edge weights in \mathcal{P} according to the probabilities in $\mathcal{A}^{\mathbb{P}}$ through the product rule below. Given learned $\mathcal{A}^{\mathbb{P}} = (Q, \Sigma, q_0, \delta, F, \delta_{\mathbb{P}}, F_{\mathbb{P}})$ and $\bar{\mathcal{T}}$, we construct *weighted product graph* $\mathcal{P} = (Q^{\mathcal{P}}, q_0^{\mathcal{P}}, E, W)$, where

- $Q^{\mathcal{P}} = (Q \times \bar{X}) \cup \{q_t^{\mathcal{P}}\}$ is a set of states, where $q_t^{\mathcal{P}}$ is the terminal state,
- $q_0^{\mathcal{P}} = (q_0, \bar{x}_0)$ is the initial state,
- $E \subseteq Q^{\mathcal{P}} \times Q^{\mathcal{P}}$ is a set of edges, and
- $W : E \rightarrow \mathbb{R}^{<0}$ is a weight function that assigns to each edge $e \in E$ a weight according to its probability in $\mathcal{A}^{\mathbb{P}}$.

The constructions of E and W are as follows.

- Edge $e = ((q, x), (q', x')) \in E$ if $q' = \delta(q, L(x'))$ and $x' = \delta_{\bar{\mathcal{T}}}(x, a)$ for some $a \in A$. Then, $W((q, x), (q', x')) = \log(\delta_{\mathbb{P}}(q, L(x'), q'))$.
- Edge $e = ((q, x), q_t^{\mathcal{P}}) \in E$ if $F_{\mathbb{P}}(q) > 0$. Then, its weight $W((q, x), q_t^{\mathcal{P}}) = \log(F_{\mathbb{P}}(q))$.

Product graph \mathcal{P} captures the constraints of both the robot and task along with the demonstrator’s preferences. Let $\lambda = (q_0, \bar{x}_0)(q_1, x_0) \dots (q_n, x_{n-1})q_t^{\mathcal{P}}$ be a path over \mathcal{P} . The projection of this path (with the deletion of $q_t^{\mathcal{P}}$) onto $\mathcal{A}^{\mathbb{P}}$ is an accepting run with the trace $\omega = \omega_1 \dots \omega_n$. The projection of λ on \mathcal{T} is the robot trajectory that generates the accepting trace ω . The probability of this trace is in fact the inverse logarithm of the total weight of λ , i.e.,

$$\begin{aligned} \sum_{i=0}^n W(\lambda_i, \lambda_{i+1}) &= \sum_{i=0}^{n-1} \log(\delta_{\mathbb{P}}(q_i, q_{i+1})) + \log(F_{\mathbb{P}}(q_n)) \\ &= \log \left(\prod_{i=0}^{n-1} \delta_{\mathbb{P}}(q_i, q_{i+1}) \cdot F_{\mathbb{P}}(q_n) \right) = \log(P(\omega)). \end{aligned}$$

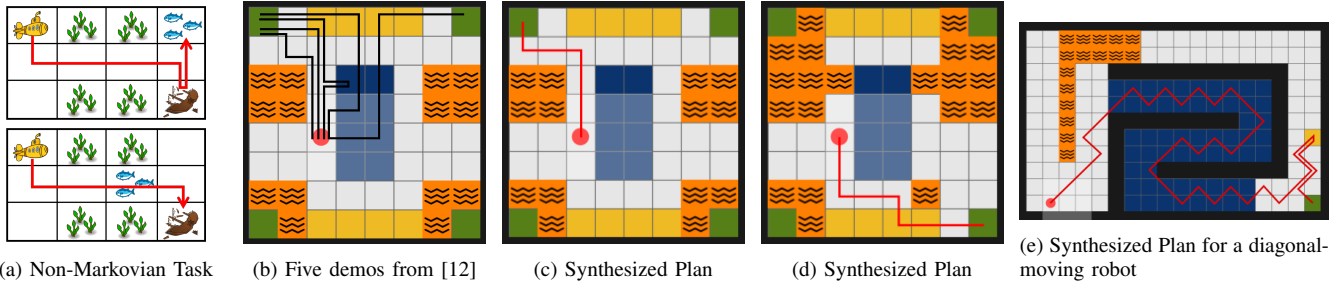


Fig. 3: Various environments and robots considered for the case studies. (a) Learning and planning for the non-Markovian task. (b) Environment and demonstrations from [12]. (c)-(e) Synthesized plans (shown in red) based on the learned task from (b).

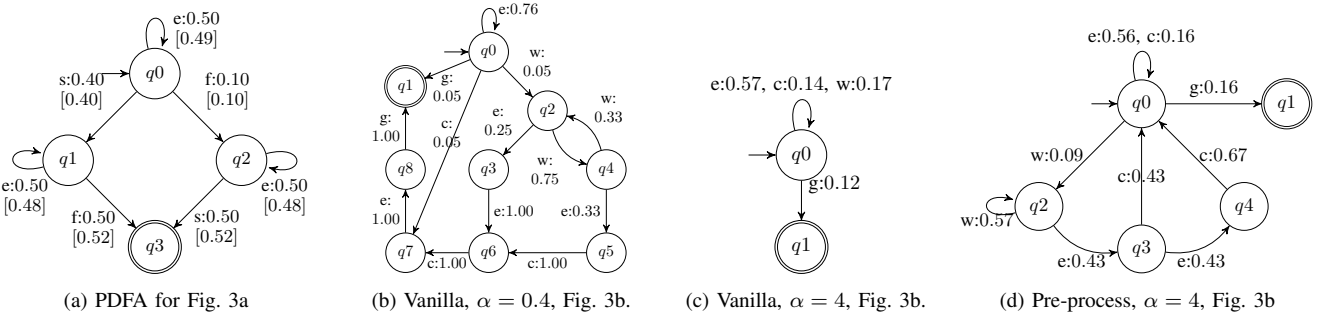


Fig. 4: The task specification and the learned PDFAs for the scenarios in Fig. 3a and 3b. Each letter represents a symbol with a single atomic proposition $s=\{ship\}$, $f=\{fish\}$, $b=\{blue\}$, $c=\{carpet\}$, $g=\{green\}$, $p=\{purple\}$, and $e=\emptyset$. The termination probability F_T of double-edged states is 1 and 0 at all other states.

Therefore, to compute a robot plan that produces an accepting trace with the highest probability in $\mathcal{L}(\mathcal{A}^P) \cap \mathcal{L}(\mathcal{T})$, it is enough to find a path on \mathcal{P} that reaches the terminal state q_t^P with the maximum total weight, i.e.,

$$\arg \max_{\omega \in \mathcal{L}(\mathcal{A}^P) \cap \mathcal{L}(\mathcal{T})} P(\omega) = \text{PROJ} \left(\arg \max_{\lambda \in \Lambda} \sum_{i=0}^{|\lambda|-1} W(\lambda_i, \lambda_{i+1}) \right),$$

where Λ is the set of paths of \mathcal{P} that terminate in q_t^P , and PROJ is the projection operator that maps λ to its corresponding trace ω . Note that the all the edge weights of \mathcal{P} are negative, hence, it is a simple graph search problem that can be performed using algorithms such as Dijkstra's.

V. CASE STUDIES AND EVALUATIONS

We illustrate the performance of the proposed algorithms in five case studies. Our implementation of the EDSM algorithm is based on the MDI method that is used in the *flexfringe* library [28]. We call the basic algorithm the *Vanilla* algorithm. All the case studies were run on a MacBook Pro with 2.3 GHz Dual-Core Intel Core i5 and 16 GB RAM. Videos of all case studies are available to view¹.

A. Learning and Planning for Non-Markovian Tasks

In this case study, we consider the robotic scenario in Figure 1. The task is to visit both the school of fish and the shipwreck in any order and always avoid coral reefs. The preference is to visit the shipwreck first. A PDFa representation of this specification is shown in Figure 4a.

To learn this task, we sampled 1000 traces from this PDFa on the gridworld environment in Figure 1. From these

demonstrations, the Vanilla algorithm learned a PDFa with the same exact structure as the true PDFa and probabilities within 0.02 of the true values (in square brackets in Fig. 4a).

As the PDFa shows, our method correctly learned the non-Markovian task of visiting both the shipwreck and the school of fish in both orders and favors going to the shipwreck first. Using this PDFa, our planner generated the robot trajectory shown in Figure 3a (top), which correctly visits the shipwreck first and then the school of fish. Next, we changed the environment by moving the location of the fish to be on the robot's way to the shipwreck as shown in Figure 3a (bottom). This figure also shows the synthesized plan in this environment using the same learned PDFa. Notice that the robot is not visiting the shipwreck first due to environmental constraint. Instead, it visits the fish and then the shipwreck, which is also a correct behavior. This generality is the strength of learning the specification rather than learning a policy that is strongly dependent on the environment.

B. Learning from Small Number of Samples with Safety

In this case study, we consider the environment and five demonstrations depicted in Figure 3b taken from [12] to learn the specification in a form of a PDFa as a comparison to the approach in [12], which is based on learning specification formulas. In this gridworld, each color represents an object, where orange is *lava*, blue is *water*, yellow is a *drying carpet*, white is an *empty* space, and green is a *charging station*. The task is to reach a charging station. However, the robot should not charge while it is wet. That is, once it gets wet (goes to water), the robot has to dry at the *drying carpet*.

1) *Small number of samples*: We first used the Vanilla algorithm with the five demonstrations, which learned the

¹<https://youtu.be/TU8MhPBDBBs>

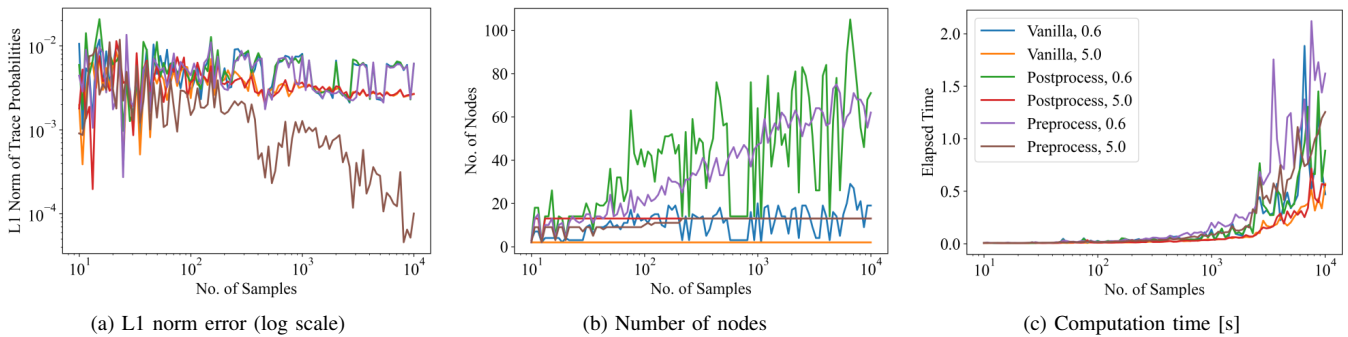


Fig. 5: Performance analysis for the proposed algorithm. Plots in (a)-(b) use the same legend as (c).

PDFA in Figure 4b. Note, in the learned PDFA, region green (*charging station*) must always be observed to reach the final state. This shows that the task of reaching the *charging station* is learned correctly. Next, on the right most branch of the PDFA, *carpet* is always observed when the robot gets wet. Again, the algorithm succeeded in learning the task of visiting *carpet* once the robots gets wet before reaching the *charging station*. One interesting observation is that the PDFA also learned that the robot has to go to the *charging station* in one step after leaving the *carpet*. This is in fact a bias in the samples since every shown demonstration that includes *carpet* has this property. If that is the intention of the demonstrator, then it is a correct behavior. If it is not, then it can be resolved by providing more samples.

Such one-step bias is not apparent in [12] because the “next” operator is not allowed in the syntax of the *language* they consider. In contrary, our method infers over regular language, which includes the next operator. Furthermore, in [12], it took 95 seconds to learn the specification from just 5 demonstrations whereas ours took less than 0.01 seconds.

2) *Hyperparameter choice and safety*: The PDFA in Figure 4b is the result of the Vanilla algorithm when the hyperparameter of α is set to 0.4. It is a knob of how aggressive we allow the merges. Higher the value of α is, the smaller the PDFA becomes. If we can tune the hyperparameter correctly, we can get a desirable result as described above. But, if we increase α too much, some merges could induce unsafe behavior. Unwanted merges occur because the algorithm is simply trying to *minimize* the size of the structure. In fact, the question of how to choose a correct value for α is an open problem. For $\alpha = 4$, the learned PDFA from the same demonstrations is shown in Figure 4c. This PDFA has no regards for safety and only requires to reach the *charging station*. We can mitigate this problem by embedding safety specification. We define the following safety formula:

$$\varphi_{\text{safe}} = \mathcal{G}\neg\text{lava} \wedge \mathcal{G}(\text{water} \rightarrow \mathcal{X}(\varphi(\neg\text{charge}, \text{carpet}, k)))$$

where φ is a formula recursively defined as: $\varphi(a, b, k) = a \wedge (b \vee \mathcal{X}(\varphi(a, b, k-1)))$ and $\varphi(a, b, 0) = a$, and is read, “visit a for k steps unless b is visited”. This formulas requires never going to *lava* and, if the robot enters *water*, it cannot *charge* unless it visits *carpet* or stays in *empty* for k consecutive steps to get dry. We set $k = 10$ in all experiments.

From the same five demonstrations, we now learn PDFAs using the Post-process and Pre-process algorithms with $\alpha =$

4 subject to φ_{safe} . The Post-process algorithm generates a large PDFA with 13 nodes and 36 edges since the safety DFA itself is large (12 nodes and 34 edges). Despite the size, it always guarantees no violation to φ_{safe} . The PDFA generated by the Pre-process algorithm is shown in Figure 4d. It is small and correctly embeds both safety and liveness. Further, all the demonstrations are accepted by both learned PDFA. As for probabilities, the average L1 norm error was 1.65×10^{-3} for the Post-process PDFA and 7.42×10^{-5} for the Pre-process PDFA, indicating better performance by the Pre-process algorithm. The larger error in the probabilities of the Post-process PDFA is due to the composition with the safety DFA, which prunes away the unsafe traces in the learned PDFA, corrupting the learned probability distributions.

Next, we perform a thorough comparison of the learning methods by increasing the number of samples.

C. Post-process versus Pre-process Algorithm

Here, the task is similar to the one above, but the goal is to quantitatively analyze and compare the performances of the proposed algorithms as the number of samples increases. We sampled demonstrations randomly from the true PDFA and used Post-process and Pre-process algorithms to learn PDFAs with hyperparameter values of $\alpha = 0.6$ (less aggressive merge) and $\alpha = 5.0$ (aggressive merge) to show the extreme results. We evaluated the resulting PDFAs with respect to the following metrics: L1 norm of the trace probability errors, number of states, and computation times. The results are shown in Figure 5 (all the plots share the same legend).

With respect to the error metric, the Pre-process algorithm (purple line for $\alpha = 0.6$ and brown line for $\alpha = 5$) consistently performs the best, followed by the Post-process and then the Vanilla algorithms. It indicates that the Pre-process algorithm is learning the correct distribution over the language of the target PDFA. As for the number of states and computation time metrics, the results indicate that Pre-process algorithm again performs better and faster than the others. From these results, we can say that the Pre-process algorithm is the best performing algorithm. Moreover, its output PDFA does not violate the safety across all the trials (checked but not shown in the figures).

D. Planning for Various Robots in different Environments

From the learned PDFAs above, we picked one with a small L1 error norm. Then, using this PDFA, we planned

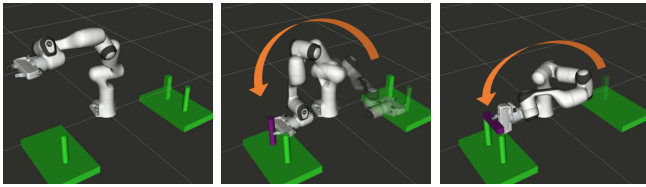


Fig. 6: Manipulator completing the learned task of building an arch.

for various robots and environments that are different from the one the demonstrations were shown in (see Figure 3b). In all the cases, the computed plans correctly meet the requirements and preferences. In the environment in Figure 3d, the *lava* forces the robot to go to the bottom-right *charging station*. Note that the robot avoids *water* by going through *carpet*, which is the preferred behavior. In Figure 3e, we modified the robot’s dynamics to only allow diagonal moves. The algorithm is again successful in generating a satisfying plan without violating the specification. Because the specification is independent from any robotic systems and any environments, our framework is robust against the changes in the environment and robot dynamics.

E. Learning and Planning for Manipulation tasks

To show that our method is not limited to mobile robots, we considered a manipulation example in Figure 6 (left). The robot is the Franka Emika Panda manipulator with 7 DoF, and the latent task is to build an arch with two cylindrical objects as columns and a rectangular box on top. The abstraction of the robot to a DTS was done according to [26], [27], which ended up with around 20,000 states. The robot was given nine demonstrations: five most preferred (fastest), three mid-length (1.4 times as many actions), and one very bad demonstration (3 times as many actions). A PDFA was learned with $\alpha = 1.8$. The learned PDFA has four states, and planning took 0.036 seconds. The execution of the plan by the robot is shown in Figure 6 (middle and right), which shows that robot successfully learned and executed the most preferred method of completing the task.

VI. CONCLUSION

In this paper, we presented a new approach to learning specifications from demonstrations in a form of a PDFA. Unlike existing work, this method does not require prior knowledge, is fast, and captures preferences. We showed how safety constraints can be incorporated in the learning algorithm, which significantly improves performance. We also introduced a planning algorithm for specifications learned in this form, which generates the most preferred (most probable) plan. Extensive evaluations illustrate the framework’s flexible and capability of robust knowledge transfer to various environments and robots.

REFERENCES

- [1] 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, pp. 297–330, 2020.
- [2] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, “Imitation learning: A survey of learning methods,” *ACM Computing Surveys (CSUR)*, vol. 50, no. 2, pp. 1–35, 2017.
- [3] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [4] S. Schaal, “Dynamic movement primitives—a framework for motor control in humans and humanoid robotics,” in *Adaptive motion of animals and machines*. Springer, 2006, pp. 261–280.
- [5] A. Paraschos, C. Daniel, J. Peters, G. Neumann, et al., “Probabilistic movement primitives,” *Neurips*, 2013.
- [6] A. Y. Ng, S. J. Russell, et al., “Algorithms for inverse reinforcement learning,” in *ICML*, vol. 1, 2000, p. 2.
- [7] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, “Maximum entropy inverse reinforcement learning,” in *AAAI*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.
- [8] M. Wulfmeier, P. Ondruska, and I. Posner, “Maximum entropy deep inverse reinforcement learning,” *arXiv:1507.04888*, 2015.
- [9] D. Ramachandran and E. Amir, “Bayesian inverse reinforcement learning,” in *IJCAI*, vol. 7, 2007, pp. 2586–2591.
- [10] C. Baier and J.-P. Katoen, *Principles of Model Checking*. Cambridge, MA: The MIT Press, 2008.
- [11] H. Kress-Gazit, M. Lahijanian, and V. Raman, “Synthesis for robots: Guarantees and feedback for robot behavior,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 211–236, May 2018.
- [12] M. Vazquez-Chanlatte, S. Jha, A. Tiwari, M. K. Ho, and S. Seshia, “Learning task specifications from demonstrations,” in *NeurIPS*, vol. 31, 2018.
- [13] A. Camacho, R. Toro Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith, “LTL and beyond: Formal languages for reward function specification in reinforcement learning,” in *Int’l Joint Conference on Artificial Intelligence*, 7 2019, pp. 6065–6073.
- [14] X. Li, C.-I. Vasile, and C. Belta, “Reinforcement learning with temporal logic rewards,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2017, pp. 3834–3839.
- [15] X. Li, Z. Serlin, G. Yang, and C. Belta, “A formal methods approach to interpretable reinforcement learning for robotic planning,” *Science Robotics*, vol. 4, no. 37, 2019.
- [16] M. Vazquez-Chanlatte, J. V. Deshmukh, X. Jin, and S. A. Seshia, “Logical clustering and learning for time-series data,” in *Computer Aided Verification*. Springer, 2017, pp. 305–325.
- [17] Z. Xu, S. Saha, B. Hu, S. Mishra, and A. A. Julius, “Advisory temporal logic inference and controller design for semiautonomous robots,” *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 1, pp. 459–477, 2018.
- [18] S. Jha, A. Tiwari, S. A. Seshia, T. Sahai, and N. Shankar, “Telex: Passive stl learning using only positive examples,” in *International Conference on Runtime Verification*. Springer, 2017, pp. 208–224.
- [19] A. J. Shah, P. Kamath, S. Li, and J. A. Shah, “Bayesian inference of temporal task specifications from demonstrations,” 2018.
- [20] B. Araki, K. Vodrahalli, T. Leech, C.-I. Vasile, M. D. Donahue, and D. L. Rus, “Learning to plan with logical automata,” 2019.
- [21] C. De la Higuera, *Grammatical inference: learning automata and grammars*. Cambridge University Press, 2010.
- [22] O. Kupferman and M. Y. Vardi, “Model checking of safety properties,” *Formal Methods in System Design*, vol. 19, pp. 291–314, 2001.
- [23] G. De Giacomo and M. Y. Vardi, “Linear temporal logic and linear dynamic logic on finite traces,” in *Int. Joint Conf. on Artificial Intelligence*, vol. 13, 2013, pp. 854–860.
- [24] H. Kress-Gazit, G. Fainekos, and G. J. Pappas, “Where’s Waldo? sensor-based temporal logic motion planning,” in *Int. Conf. on Robotics and Automation*. Rome, Italy: IEEE, 2007, pp. 3116–3121.
- [25] M. Lahijanian, M. Kloetzer, S. Itani, C. Belta, and S. Andersson, “Automatic deployment of autonomous cars in a robotic urban-like environment (RULE),” in *Int. Conf. on Robotics and Automation*. Kobe, Japan: IEEE, 2009, pp. 2055–2060.
- [26] K. He, M. Lahijanian, L. E. Kavradi, and M. Y. Vardi, “Towards manipulation planning with temporal logic specifications,” in *Int. Conf. Robotics and Automation*. IEEE, May 2015, pp. 346–352.
- [27] K. He, M. Lahijanian, E. Kavradi, Lydia, and Y. Vardi, Moshe, “Automated abstraction of manipulation domains for cost-based reactive synthesis,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 285–292, Apr. 2019.
- [28] S. Verwer and C. A. Hammerschmidt, “Flexfringe: a passive automaton learning package,” in *Intl. Conf. Software Maintenance and Evolution (ICSME)*. IEEE, 2017, pp. 638–642.