

Learning Minimally-Violating Continuous Control for Infeasible Linear Temporal Logic Specifications

Mingyu Cai¹, Makai Mann², Zachary Serlin², Kevin Leahy², Cristian-Ioan Vasile¹

Abstract—This paper explores continuous-time control synthesis for target-driven navigation to satisfy complex high-level tasks expressed in linear temporal logic (LTL). We propose a model-free framework using deep reinforcement learning (DRL) where the underlying dynamical system is unknown (an opaque box). Unlike prior work, we consider scenarios where the given LTL specification might be infeasible and therefore cannot be accomplished globally. Instead of modifying the given LTL formula, we provide a general DRL-based approach to satisfy it with minimal violation. To do this, we transform a previously multi-objective DRL problem, which requires simultaneous automata satisfaction and minimum violation cost, into a single objective. By guiding the DRL agent with a sampling-based path planning algorithm for the potentially infeasible LTL task, the proposed approach mitigates the myopic tendencies of DRL, which are often an issue when learning general LTL tasks that can have long or infinite horizons. This is achieved by decomposing an infeasible LTL formula into several reach-avoid sub-tasks with shorter horizons, which can be learned in a modular DRL architecture. Furthermore, we overcome the challenge of the exploration process for DRL in cluttered environments by using path planners to design rewards that are dense in the configuration space. The benefits of the presented approach are demonstrated through testing on various complex nonlinear systems and compared with state-of-the-art baselines. The video demonstration can be found here: <https://youtu.be/DqesqBsja9k>.

I. INTRODUCTION

Autonomous agents operating in complex environments must often accomplish high-level tasks while subject to various uncertainties. Possible uncertainties include dense obstacles or targets, and impassable or unexpected terrain. Additionally, a principled analytical model of the robot is typically (partially) unknown or hard to obtain. Motivated by these challenges, this work investigates model-free control of an autonomous system for satisfying a potentially infeasible high-level task with minimal violation while operating in a complex cluttered environment.

In the field of model-free navigation control, Reinforcement Learning (RL) is a popular unsupervised technique that optimizes long-term expected rewards to learn desired behaviors [1]. Recently, Deep Reinforcement Learning (DRL) techniques, such as the actor-critic method [2], have been shown to be capable of learning continuous controllers for

unknown (opaque-box) dynamics. To investigate complex navigation tasks, Linear Temporal Logic (LTL) is a formal language that has gained attention in recent years for expressing high-level, interpretable robotics tasks. Specifications have been shown to be effective at directing RL agents to learn desired policies. Typically, a discrete robot system is abstracted as a discrete Markov Decision Process (MDP) model and composed with an automaton representing the desired LTL formula to create a product automaton for learning or planning [3]–[5]. For general continuous control subject to LTL satisfaction over infinite horizons, prior works [6], [7] proposed a modular architecture by decomposing the task into sub-tasks using automata states. Crucially, to our knowledge, no previous work has investigated infeasible LTL tasks for continuous control. In this work, we focus on the problem jointly including two aspects: infeasible LTL tasks and cluttered environments using DRL for model-free continuous control.

Related works: Minimum violation of LTL specifications over finite horizons has been considered in [8]–[13] using sampling-based methods, graph-based optimization algorithms, and dynamic programming to solve the corresponding motion planning problem. For LTL specifications over infinite horizons, general violation measurements are proposed in [14] by quantitatively revising LTL automata. The approaches are applied in [15], [16], where they build product structures and encode potentially conflicting parts of a specification as a soft constraint. However, the prior literature generally assumes known dynamical systems with abstracted navigation controllers and only focuses on motion planning problems [8]–[11], [13], [15]. Our work improves on this by directly synthesizing a low-level controller for unknown dynamics in a model-free manner that minimally violate LTL satisfaction over both finite and infinite horizons.

Previous work [17], [18] has tackled RL while considering infeasible tasks in discrete state-action spaces. In particular, authors in [17] account for the infeasible solution case by visiting the automaton accepting sets as much as possible (if at all). But many LTL formulas only have one accepting state and this strategy often fails to generalize to those cases. The learning objectives in [18] ensure satisfaction of the automata acceptance condition and, thus, LTL formulae, and reduce the violation cost, resulting in a multi-objective RL (MORL) problem. However, providing minimum-violation performance guarantees for MORL continues to be an open problem [19].

Commonly in the literature for navigation in cluttered environments, motion planning algorithms are first used to find possible paths, and then path-tracking controllers are employed to follow that path [20]–[22]. This approach decomposes into a goal-reaching control problem and can

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited. This material is based upon work supported by the Under Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Under Secretary of Defense for Research and Engineering.

¹Mingyu Cai and Cristian-Ioan Vasile are with Mechanical Engineering, Lehigh University, Bethlehem, PA, 18015 USA. mic221@lehigh.edu, crv519@lehigh.edu

²Makai Mann, Zachary Serlin and Kevin Leahy are with MIT Lincoln Laboratory, Lexington, MA, 02421, USA. makai.mann@ll.mit.edu, Zachary.Serlin@ll.mit.edu, kevin.leahy@ll.mit.edu

be easily adapted for LTL satisfaction [23], [24]. Classic approaches however require known dynamic models and may be computationally expensive for nonlinear and high-dimensional systems. This fundamentally motivates our focus on employing model-free control techniques.

Contributions: In this paper, we translate the MORL problem first analyzed in [18] into a standard DRL problem and apply model-free geometric motion planning to guide the learning process. The generated trajectories are shown to minimally violate a given infeasible LTL specification. By learning a controller to reach waypoints in a geometric path, our approach automatically obtains an optimal control policy for minimally-violating the infeasible LTL task specification. This paper's main contributions are: (i) To the best of our knowledge, this is the first model-free learning-based continuous control for potentially infeasible LTL satisfaction with minimum-violation objectives. Such objectives facilitate meaningful action for infeasible LTL tasks, and the framework also generalizes discrete systems; (ii) Our framework can learn continuous control in cluttered environments and mitigates the myopic tendencies of DRL by decomposing the global task into sub-tasks, which can be trained compositionally; (iii) We demonstrate the benefits by comparing with several baselines in two nonlinear systems: a Dubins car and quadrotor.

II. PRELIMINARIES

Agent: We consider the general evolution of a continuous-time dynamical system \mathfrak{S} that is assumed to be unknown, starting from an initial state $s_0 \in S_0$ is given by

$$\dot{s} = f(s, a), \quad (1)$$

where $s \in S \subseteq \mathbb{R}^n$ is the state vector in the compact set S , $a \in A \subseteq \mathbb{R}^m$ is the control input. We assume the flow field $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is uniformly continuous in time and Lipschitz continuous in s for fixed a . Under these assumptions there exists a unique solution of \mathfrak{S} for a given a , providing trajectories of the system [25].

Definition 1. For a robot operating in an environment Env , the geometric workspace can be represented by a compact subset $X \subset \mathbb{R}^d$, $d \in \{2, 3\}$. The relation between dynamics \mathfrak{S} and workspace X is defined by the projection $Proj : S \rightarrow X$.

The space X contains regions of interest that are labeled by a set of atomic propositions AP , with the labeling function $L_X : X \rightarrow 2^{AP}$. Let $L : S \rightarrow 2^{AP}$ be a labeling function over \mathfrak{S} i.e., $L(s) = L_X(Proj(s))$.

Reinforcement Learning: The interactions between a robot with dynamics \mathfrak{S} and an environment Env can be captured as a continuous-labeled Markov Decision Process (cl-MDP). A cl-MDP is a tuple $\mathcal{M} = (S, S_0, A, p_S, R, \gamma, L)$, where $S \subseteq \mathbb{R}^n$ is a continuous state space, S_0 is a set of initial states, $A \subseteq \mathbb{R}^m$ is a continuous action space, and p_S captures the unknown system dynamics as a distribution. The distribution $p_S : \mathfrak{B}(\mathbb{R}^n) \times A \times S \rightarrow [0, 1]$ is a Borel-measurable conditional transition kernel, s.t. $p_S(\cdot | s, a)$ is a probability measure of the next state given current $s \in S$ and $a \in A$ over the Borel space $(\mathbb{R}^n, \mathfrak{B}(\mathbb{R}^n))$, where $\mathfrak{B}(\mathbb{R}^n)$ is the set of all Borel sets on \mathbb{R}^n . $R : S \times A \times S \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in (0, 1)$ is the discount factor. L is the labeling function in Def. 1.

Since the dynamics \mathfrak{S} are an opaque-box, the transition relation $p_S(\cdot | s, a)$ for any state-action pair is unknown.

Actor-critic RL algorithms [2] have been demonstrated as promising tools to solve continuous-control problems for the cl-MDP model, where each valid transition of cl-MDP follows f defined in (1) that is zero-order hold model for a continuous-time action.

Remark 1. Note it's intractable to explicitly construct the cl-MDP model with continuous state and action space to apply standard model-checking algorithm [26]. Inspired by prior work [6], we track it on-the-fly via a deep neural network.

Let Π denote the set of all policies for a cl-MDP, and $\pi \in \Pi$ denote a stochastic policy $\pi : S \times A \rightarrow [0, 1]$, that maps states to distributions over actions. The objective is to learn an optimal policy $\pi^*(a|s)$ that maximizes the expected discounted return $J(\pi) = \mathbb{E}^\pi \left[\sum_{k=0}^{\infty} \gamma^k \cdot R(s_k, a_k, s_{k+1}) \right]$.

Linear Temporal Logic (LTL): LTL is a formal language to describe complex properties and high-level specifications of a system. LTL formulae are built inductively from atomic propositions by applying Boolean and temporal operators to subformulae. The syntax in Backus-Naur form is given by:

$$\phi ::= \top \mid \mu \mid \phi_1 \wedge \phi_2 \mid \neg \phi_1 \mid \circ \phi \mid \phi_1 \mathcal{U} \phi_2 \mid \diamond \phi \mid \square \phi,$$

where $\mu \in AP$ is an atomic proposition, *true* \top , *negation* \neg , and *conjunction* \wedge are propositional logic operators, and *next* \circ and *until* \mathcal{U} are temporal operators [26]. Alongside the standard operators introduced above, other propositional logic operators, such as *false*, *disjunction* \vee , and *implication* \rightarrow , and temporal operators, such as *always* \square and *eventually* \diamond , are derived from the standard operators.

The semantics of an LTL formula are interpreted over words, where a word is an infinite (continuous) sequence $o = o_0 o_1 \dots$, with $o_i \in 2^{AP}$ for all $i \geq 0$, where 2^{AP} represents the power set of AP . The satisfaction of an LTL formula ϕ by the word o is denoted $o \models \phi$. The detailed semantics of LTL can be found in [26]. In this paper, we are interested in continuous-time control policies. The *next* operator is not always meaningful since it may require an immediate execution switch in the synthesized plans space. We can either exclude the *next* operator as is common in related work [23] or properly design practical LTL tasks [6] in continuous scenarios.

III. PROBLEM FORMULATION

Consider a cl-MDP $\mathcal{M} = (S, S_0, A, p_S, R, \gamma, L)$. The induced path under a policy π over \mathcal{M} is $s_\infty^\pi = s_0 \dots s_i s_{i+1} \dots$, where $p_S(s_{i+1} | s_i, a_i) > 0$ if $\pi(a_i | s_i) > 0$.

We extend the labeling function to traces such that $L(s_\infty^\pi) = o_0 o_1 \dots$ is the sequence of labels associated with s_∞^π . We denote the satisfaction relation of the induced trace with respect to ϕ by $L(s_\infty^\pi) \models \phi$. The probability of satisfying ϕ under the policy π , starting from an initial state $s_0 \in S_0$, is defined as

$$\Pr_M^\pi(\phi) = \Pr_M^\pi(L(s_\infty^\pi) \models \phi \mid s_\infty^\pi \in \mathcal{S}_M^\pi),$$

where \mathcal{S}_M^π is the set of admissible paths from the initial state s_0 , under the policy π [26].

Definition 2. Given a cl-MDP \mathcal{M} , an LTL task ϕ is fully feasible if and only if there exists a policy π s.t. $\Pr_M^\pi(\phi) > 0$.

Note that according to Def. 2, an infeasible case means no policy exists that satisfies the task, which can be interpreted as $\Pr_M^\pi(\phi) = 0, \forall \pi \in \Pi$.

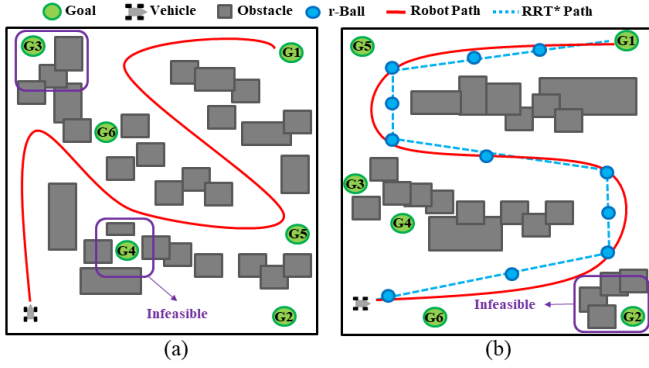


Fig. 1. The figure shows the example described in Example 1, where the given LTL task is infeasible since some regions of interest, e.g., $\mathcal{G}_3, \mathcal{G}_4$ of (a). (a) The illustration that learning navigation using DRL is challenging in complex cluttered environments. (b) Demonstration of the effective design reward in Section IV-C.

Definition 3. Given a cl-MDP \mathcal{M} , the expected discounted violation cost with respect to a given LTL task ϕ under the policy π is defined as

$$J_V(\mathcal{M}^\pi, \phi) = \mathbb{E}^\pi \left[\sum_{k=0}^{\infty} \gamma^k c_V(s_k, a_k, s_{k+1}, \phi) \right], \quad (2)$$

where $c_V(s, a, s', \phi)$ is formally defined in Def. 8 as the violation cost of a transition (s, a, s') with respect to ϕ , and a_i is the action generated based on the policy $\pi(s_i)$.

Problem 1. Given a robot with unknown dynamics \mathfrak{S} in an Env with workspace X containing regions of interest, a navigation task in the form of an LTL task ϕ over AP , their interactions can be captured as a cl-MDP. The DRL objective is to find an optimal policy π^* with the following capabilities: (i) if ϕ is feasible, $\Pr_{\mathcal{M}}^{\pi^*}(\phi) > 0$; (ii) if ϕ is infeasible s.t. $\Pr_{\mathcal{M}}^{\pi}(\phi) = 0, \forall \pi \in \Pi$, satisfy ϕ with minimum violation cost via minimizing $J_V(\mathcal{M}^\pi, \phi)$.

Example 1. In Fig. 1, consider an autonomous vehicle with unknown dynamics deployed in a complex, cluttered environment containing a set of labeled goal regions, $AP_G = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_5\}$, and labeled obstacles \mathcal{O} . The LTL specification is given by $\phi = \square \neg \mathcal{O} \wedge \square((\diamond \mathcal{G}_1 \wedge \diamond(\mathcal{G}_2 \wedge \diamond \dots \wedge \diamond \mathcal{G}_5))$. We observe that some goal regions are surrounded by obstacles such that ϕ is infeasible.

For LTL specifications, discrete goal-reaching may enable the automaton transitions, which is one requirement to receive sparse automata-based rewards [6], [27]. However, in cluttered environments such as those depicted in Fig. 1, it can be difficult for noisy policies to explore the sparse rewards and reach regions of interest at all during learning.

IV. SOLUTION

In section IV-A, we develop an automaton model with relaxed constraints to address infeasible tasks. In section IV-B, we show how to synthesize a model-free plan with minimum violation cost in continuous space. Finally, in section IV-C, we propose a novel DRL design to learn optimal policies that solve Problem 1.

A. Relaxed Product Automaton for Infeasible LTL

Let $dist : X \times X \rightarrow [0, \infty)$ define a metric function that computes the geometric Euclidean distance between two states. Prior work [7] introduced a transition system for an unknown dynamical system to capture the interactions between geometric space X and Env .

Definition 4. A generalized weighted transition system (G-WTS) of Env is a tuple $\mathcal{T} = (X, x_0, \rightarrow_{\mathcal{T}}, AP, L_X, C_{\mathcal{T}})$, where X is the configuration space of Env , x_0 is the initial state of the robot; $\rightarrow_{\mathcal{T}} \subseteq X \times X$ is the geometric transition relation s.t. $x \rightarrow_{\mathcal{T}} x'$ if $dist(x, x') \leq \eta$ and the straight line σ connecting x to x_{new} is collision-free; AP is the set of atomic propositions that label regions; $L_X : X \rightarrow 2^{AP}$ is the labeling function that returns an atomic proposition satisfied at location x ; and $C_{\mathcal{T}} : (\rightarrow_{\mathcal{T}}) \rightarrow \mathbb{R}^+$ is the geometric Euclidean distance, i.e.; $C_{\mathcal{T}}(x, x') = dist(x, x'), \forall (x, x') \in \rightarrow_{\mathcal{T}}$.

Let $\tau_{\mathcal{T}} = x_0 x_1 x_2 \dots$ denote a valid run of \mathcal{T} . As opposed to the standard WTS [23] which includes dynamic state space S and explicit dynamics $f(s, a)$ in transitions, only the geometric relations are available in our framework.

An LTL formula ϕ can be converted to a Non-deterministic Büchi Automaton (NBA) to evaluate satisfaction.

Definition 5. [28] An NBA over 2^{AP} is a tuple $\mathcal{B} = (Q, Q_0, \Sigma, \rightarrow_{\mathcal{B}}, Q_F)$, where Q is the set of states, $Q_0 \subseteq Q$ is the set of initial states, $\Sigma = 2^{AP}$ is the finite alphabet, $\rightarrow_{\mathcal{B}} \subseteq Q \times \Sigma \times Q$ is the transition relation, and $Q_F \subseteq Q$ is the set of accepting states.

A valid infinite run $\tau_{\mathcal{B}} = q_0 q_1 q_2 \dots$ of \mathcal{B} is called accepting, if it intersects with Q_F infinitely often. Infinite words $\tau_o = o_0 o_1 o_2 \dots, \forall o \in 2^{AP}$ generated from an accepting run satisfy the corresponding LTL formula ϕ . We use \mathcal{B}_{ϕ} to denote the NBA of LTL formula ϕ .

A common approach for synthesizing a geometric plan satisfying LTL formula ϕ is to construct the product Büchi automaton (PBA) between the G-WTS \mathcal{T} and the NBA \mathcal{B} [7], [23]. This approach assumes the given ϕ is feasible. Inspired by [14], [15] to handle infeasible cases, we introduce a relaxed PBA and define corresponding violation costs.

Definition 6. Given the G-WTS \mathcal{T} and the NBA \mathcal{B}_{ϕ} , the relaxed PBA is a tuple $P = \mathcal{T} \times \mathcal{B} = (Q_P, Q_P^0, \rightarrow_P, Q_P^F, c_P, L_P, \Sigma)$, where

- $Q_P = X \times Q$ is the set of infinite product states, $Q_P^0 = x_0 \times Q$ is the set of initial states, $Q_P^F = X \times Q_F$ is the set of accepting states, and $\Sigma = 2^{AP}$.
- $\rightarrow_P \subseteq Q_P \times Q_P$ is the transition function such that $q_P = (x, q) \rightarrow_P q'_P = (x', q')$ is a valid transition if and only if the following two conditions hold: $x \rightarrow_{\mathcal{T}} x'$ and $\exists \sigma_{AP} \in \Sigma$ s.t. $q \xrightarrow{\sigma_{AP}}_{\mathcal{B}_{\phi}} q'$,
- $c_P : (\rightarrow_P) \rightarrow \mathbb{R}^+$ is the geometric cost function defined as the cost in the configuration space, e.g., $c_P(q_P = (x, q), q'_P = (x', q')) = C_{\mathcal{T}}(x, x'), \forall (q_P, q'_P) \in \rightarrow_P$,
- $L_P : Q_P \rightarrow \Sigma$ is the labelling function s.t. $L_P(q_P) = L_X(x), \forall q_P = (x, q)$,
- $c_V : (\rightarrow_P) \rightarrow \mathbb{R}^+$ is the violation cost function for each transition with respect to the associated LTL formula ϕ .

The major novelty of the relaxed PBA is replacing the transition constraint $q \xrightarrow{L_X(x)}_{\mathcal{B}_{\phi}} q'$ of a standard PBA [26] with $\exists \sigma_{AP} \in \Sigma$ s.t. $q \xrightarrow{\sigma_{AP}}_{\mathcal{B}_{\phi}} q'$, to construct a more connected product graph. Next, we define the computation of the violation function c_V .

Suppose $AP = \{\mu_1, \mu_2, \dots, \mu_M\}$ and consider an evaluation function $Eval : \Sigma \rightarrow \{0, 1\}^M$, where $\Sigma = 2^{AP}$ and $M = |AP|$. $Eval(\sigma) = [v_i]_{i=1}^M$ and v_i is a binary variable s.t. $v_i = 1$ if $\mu_i \in \sigma$ and $v_i = 0$ if $\mu_i \notin \sigma$, where $i = 1, 2, \dots, M$ and $\sigma \in \Sigma$. To quantify the difference between two elements in Σ , consider $\rho(\sigma, \sigma') = \|v - v'\|_1 = \sum_{i=1}^M |v_i - v'_i|$,

where $v = \text{Eval}(\sigma)$, $v' = \text{Eval}(\sigma')$, $\sigma, \sigma' \in \Sigma$, and $\|\cdot\|_1$ is the L_1 norm. The distance from $\sigma \in \Sigma$ to a set $\mathcal{X} \subseteq 2^{\Sigma}$ is then defined as $D_V(\sigma, \mathcal{X}) = \min_{\sigma' \in \mathcal{X}} \rho(\sigma, \sigma')$ if $\sigma \notin \mathcal{X}$, and $D_V(\sigma, \mathcal{X}) = 0$ otherwise.

Definition 7. In a relaxed PBA, the violation cost of the transition $q_P = (x, q) \rightarrow_P q'_P = (x', q')$ imposed by ϕ can be computed with $c_V(q_P, q'_P) = D_V(L_P(q_P), \mathcal{X}(q, q'))$, where $\mathcal{X}(q, q') = \{\sigma_{AP} \in \Sigma \mid q \xrightarrow{\sigma_{AP}}_{B_\phi} q'\}$ is the set of input symbols enabling the transition from q to q' in B_ϕ .

Based on that, we can also measure the violation cost imposed on the dynamic system \mathfrak{S} as follows.

Definition 8. Given a transition (s, a, s') of dynamics and current automata state q , the next automata state q' can be generated on-the-fly as $q \xrightarrow{L(s)}_{B_\phi} q'$. The violation cost of the transition with respect to ϕ can be obtained as $c_V(s, a, s', \phi) = D_V(L(s), \mathcal{X}(q, q'))$.

Remark 2. Note that directly applying $c_V(s, a, s', \phi)$ into reward design for LTL satisfaction formulates a multi-objective DRL problem. The optimality performance of multi-objective DRL as a nonlinear regression process is hard to control. Instead, we use path planning over workspace X for DRL guidance.

Given a valid trace $\tau_P = q_P^0 q_P^1 q_P^2 \dots$ of a PBA, the total violation cost can be computed with $W_V(\tau_P) = \sum_{i=0}^{\infty} c_V(q_P^i, q_P^{i+1})$. The trace $\tau_P = q_P^0 q_P^1 q_P^2 \dots$ satisfies the acceptance conditions if it intersects with Q_P^E infinitely often. Its corresponding words are obtained as $\tau_o = o_0 o_1 o_2 \dots$, $\forall o_i = L_P(q_P^i)$. Now, the LTL satisfaction according to the accumulated violation cost can be measured as follows:

Lemma 1. [14] *The words τ_o^* of an optimal trace τ_P^* satisfy the corresponding potentially infeasible LTL specification ϕ with minimum-violation guarantees if and only the following two conditions hold: (i) τ_P^* satisfies the acceptance condition and (ii) $W_V(\tau_P^*)$ has the minimum total violation cost, i.e., $\tau_P^* = \arg \min_{\tau_P \in \tau_\infty} W_V(\tau_P)$, where τ_∞ denotes the set of all valid traces in the relaxed PBA P .*

The word τ_o^* of τ_P^* satisfies ϕ exactly if $W_V(\tau_P^*) = 0$. Let $\text{proj}|_X : Q_P \rightarrow X$ denote a projection s.t. $\text{proj}|_X(q_P) = x, \forall q_P = (x, q)$. We use this operator to extract the geometric trajectory $\tau_{\mathcal{T}}^* = \text{proj}|_X(\tau_P^*)$ for minimally-violating LTL satisfaction from the optimal trace τ_P^* .

Remark 3. This work considers unknown dynamic systems so that the generated path $\tau_{\mathcal{T}}^*$ is only optimal in the sense of geometric relations e.g., shortest euclidean distance.

B. Minimum-violation Synthesis and Decomposition

Minimum-violation Synthesis: Since this work considers continuous control and the geometric space X of \mathcal{T} is also continuous, it's impossible to explicitly construct the relaxed PBA, P , and find the optimal trace τ_P^* . Instead of discretizing space, we apply Temporal Logic Rapidly-Exploring Random Trees (TL-RRT*) [23], which is abstraction-free and builds a tree in the configuration space incrementally. This technique has the same properties as RRT* [24], [29]: it can generate probabilistically complete and asymptotically optimal solutions.

Formally, let τ_F denote any accepting run of PBA P . TL-RRT* leverages the fact that an accepting run τ_F is a lasso-type sequence with a prefix-suffix structure, i.e., $\tau_F =$

$\tau_P^{\text{pre}}[\tau_P^{\text{suf}}]^\omega$, where the prefix part $\tau_P^{\text{pre}} = q_P^0 q_P^1 \dots q_P^K$ is only executed once, and the suffix part $\tau_P^{\text{suf}} = q_P^{K+1} \dots q_P^{K+l}$ with $q_P^{K+1} = q_P^{K+l}$ loops infinitely. Following this idea, we build the prefix and suffix optimal trees, respectively. To ensure the acceptance condition, the set of goal states of the prefix tree $G_P^{\text{pre}} = (V_P^{\text{pre}}, E_P^{\text{pre}})$ is defined as $Q_{\text{goal}}^{\text{pre}} = \{q_P = (x, q) \in X_{\text{free}} \times Q \subseteq Q_P \mid q \in Q_F\}$, where X_{free} is the collision-free configuration space. The optimal goal states of the prefix tree are $Q_{\text{goal}}^* = V_P^{\text{pre}} \cap Q_{\text{goal}}^{\text{pre}}$ that is regarded as the roots of suffix tree $G_P^{\text{suf}} = (V_P^{\text{suf}}, E_P^{\text{suf}})$. The destination states of the suffix tree are

$$Q_{\text{goal}}^{\text{suf}}(q_P^*) = \{q_P = (x, q) \in X_{\text{free}} \times Q \subseteq Q_P \mid q_P \rightarrow_P q_P^* = (x^*, q^*), \forall q_P^* \in Q_{\text{goal}}^*\}.$$

We refer readers to [23] for more details.

Unlike [23], we consider infeasible cases where rapidly-exploring random trees should be built incrementally based on the relaxed PBA and our optimization objective includes both violation and geometric cost. To guarantee total minimum-violation, let the weight between any two product states in the tree be $w_P(q_P, q'_P) = c_P(q_P, q'_P) + \beta \cdot c_V(q_P, q'_P)$, where we select $\beta \gg 1$ to ensure the violation cost c_V has higher priority than geometric distance c_P . Finally, we can find the minimum-violation path in the form of $\tau_P^* = \tau_{\text{pre}}^*[\tau_{\text{suf}}^*]^\omega$, which can be projected into the configuration space $\tau_{\mathcal{T}}^* = \text{proj}|_X(\tau_P^*) = \tau_{\mathcal{T}, \text{pre}}^*[\tau_{\mathcal{T}, \text{suf}}^*]^\omega$.

Optimal decomposition: Since directly applying the whole geometric path $\tau_{\mathcal{T}}^* = \text{proj}|_X(\tau_P^*)$ with long or infinite horizons as guidance for DRL suffers from myopic tendencies (gathering intermediate rewards rather than reaching the desired region), we can decompose $\tau_{\mathcal{T}}^*$ into several sub-paths. To do so, we rewrite $\tau_P^* = \tau_0^* \tau_1^* \dots \tau_K^*[\tau_{K+1}^* \dots \tau_{K+l}^*]^\omega$ such that $\forall i \in \{0, 1, \dots, K+l\}$, τ_i^* is a sub-trajectory and the product states of each τ_i^* have the same automaton components i.e., $q = q_i, \forall q_P = (x, q) \in \tau_i^*$. Each segment τ_i^* can be projected into workspace as $\tau_{\mathcal{T}, i}^*$, and we have

$$\tau_{\mathcal{T}}^* = \text{proj}|_X(\tau_P^*) = \tau_{\mathcal{T}, 0}^* \tau_{\mathcal{T}, 1}^* \dots \tau_{\mathcal{T}, K}^*[\tau_{\mathcal{T}, K+1}^* \dots \tau_{\mathcal{T}, K+l}^*]^\omega.$$

Let $\text{proj}|_Q : Q_P \rightarrow Q$ denote a projection s.t. $\text{proj}|_Q(q_P) = q, \forall q_P = (x, q)$ to extract the automaton components, and we have

$$\tau_Q^* = \text{proj}|_Q(\tau_P^*) = q_0^* q_1^* \dots q_K^*[q_{K+1}^* \dots q_{K+l}^*]^\omega. \quad (3)$$

As a result, we can define each segment $\tau_{\mathcal{T}, i}^*$ as a reach-avoid path $\mathcal{R}_i(\mathcal{G}_i, \mathcal{O})$, where \mathcal{G}_i is its goal region and \mathcal{O} is a set of all obstacles. The lasso-type geometric reach-avoid path can be written

$$\mathcal{R}_{\mathcal{T}} = (\mathcal{R}_0 \mathcal{R}_1 \dots \mathcal{R}_K)(\mathcal{R}_{K+1} \dots \mathcal{R}_{K+l})^\omega. \quad (4)$$

For each reach-avoid task \mathcal{R}_i , let $x_{\mathcal{R}, i}$ denote the destination state, which is also the last state of the segment $\tau_{\mathcal{T}, i}^*$. From definition 7, we can obtain:

Lemma 2. *For a relaxed PBA, for a continuous trace $\tau_P = q_P^0 q_P^1 q_P^2 \dots$ satisfying the acceptance condition, we can extract an automaton trace $\tau_Q = \text{proj}|_Q(\tau_P) = q_0 q_1 \dots q_K[q_{K+1} \dots q_{K+l}^*]^\omega$ that satisfies the acceptance condition of the NBA. The total violation cost $W_V(\tau_P) = \sum_{k=0}^{\infty} c_V(q_P^k, q_P^{k+1})$ can be transformed into the finite discrete form using its*




Fig. 2. (Left) NBA \mathcal{B} of the LTL formula $\phi_{g_1} = \square \diamond G_1 \wedge \square \diamond G_2 \wedge \square \diamond G_3$ for $\phi = \square \neg \mathcal{O} \wedge \phi_{g_1}$; (Right) Decomposed reach-avoid paths (sub-tasks) that minimally violates ϕ , where \mathcal{G}_3 is not reachable.

automaton trace as $\sum_{k=0}^{K-1} D_V(L_X(x_{\mathcal{R},k+1}), \mathcal{X}(q_k, q_{k+1}))$
 $+\beta \cdot \sum_{k=K}^{K+l-1} D_V(L_X(x_{\mathcal{R},k+1}), \mathcal{X}(q_k, q_{k+1})).$

Lemma 2 states that we can measure the total violation cost of an infinite continuous path using a finite discrete form. The structure of $\mathcal{R}_{\mathcal{T}}$ can enable a modular structure of distributedly solving each \mathcal{R}_i with shorter horizons. Note that, the above process generates the asymptotically minimally-violating plan that is probabilistically complete.

Example 2. Fig. 2 shows an example of the optimal decomposition for the infeasible case. The LTL task is $\phi = \square \neg \mathcal{O} \wedge \phi_{g_1} = \square \diamond G_1 \wedge \square \diamond G_2 \wedge \square \diamond G_3$ over infinite horizons, whereas the goal regions \mathcal{G}_3 is not accessible. The resulting NBA and decomposed trajectories of TL-RRT* minimally-violating ϕ are shown in Fig. 2 (left) and (right), respectively, where decomposed reach-avoid paths (sub-tasks) are expressed as $\mathcal{R}_{\mathcal{T}} = \mathcal{R}_{red}(\mathcal{R}_{blue}\mathcal{R}_{brown})^\omega$.

C. Reward Design and Distributed DRL

Reward design: According to the optimal compositional plan $\mathcal{R}_{\mathcal{T}}$, this section designs a reward function for each \mathcal{R}_i that is dense in the configuration space to overcome the challenge of the complex environment. We can train every reach-avoid \mathcal{R}_i navigation control in the same way distributively. Even if we can straightforwardly use the geometric sub-path for reward design, there exist three issues: (a) without considering actual dynamical systems, controllers cannot strictly track the geometric path; (b) the robot may linger around previous waypoints resulting in sub-optimal solutions; (c) the waypoints might overlap with obstacles. We address them in the following descriptions, and the intuition of the reward design is shown in Fig. 1 (b), where the robot just needs to move towards the goal and is not required to strictly follow the geometric path.

(a): Each geometric path of \mathcal{R}_i is a sequence of waypoints $\tau_{\mathcal{T},i}^* = x_{i,0}x_{i,1}\dots x_{i,N}$. We construct an r-norm ball for each state $x_{i,j} \in \tau_{\mathcal{T},i}^*, \forall j = 0, 1, \dots, N$, as $Ball_r(x_{i,j}) = \{x \in X \mid dist(x, x_{i,j}) \leq r\}$. These balls allow the robot to pass them and move towards the goal as guidance, which we use to design a reward function. The consecutive balls can be regarded as a tube which is popular in the robust control community [30] to handle path tracking infeasibility.

(b): The intuitive approach is to track the minimum distance of the visited balls to the destination so far and utilize it as a constraint in the reward design. Since TL-RRT* is an extension of RRT* [29], it also provides the optimal distance function in the tree from each state to the global final destination denoted as $Cost(x)$. For a sub-path $\tau_{\mathcal{T},i}^* = x_{i,0}x_{i,1}\dots x_{i,N}$, we compute the distance from each

state $x_{i,j} \in \tau_{\mathcal{T},i}^*$ to the destination $x_{i,N}$ as $Dist(x_{i,j}) = Cost(x_{i,j}) - Cost(x_{i,N})$. We return the necessary distance information as:

$$D(x) = \begin{cases} Dist(x_i|x^*) & \text{if } x \in Ball_r(x_i|x^*) \\ \infty & \text{otherwise} \end{cases} \quad (5)$$

During each episode of training, a state-action sequence $s_0a_0s_1a_1\dots s_t$ up to current time t is split into the state and action sequences $\mathbf{s}_t = s_0s_1\dots s_t$ and $\mathbf{a}_t = a_0a_1\dots a_{t-1}$, respectively. We develop a progression function D_{min} to identify whether the next state is getting closer to the goal region as $D_{min}(s_t) = \min_{s \in \mathcal{S}_t} \{D(Proj(s))\}$.

The learning objective is to push D_{min} to decrease at each time-step. We design a constrained reward for each time t of an episode as $R(s_t^x) = r_+ > 0$ if $D(Proj(s_t)) < D_{min}(s_{t-1})$. Such a design also alleviates the first issue (a), since the RL agent is not required to visit every waypoint.

(c): We design the reward in different priorities such that if the RL-agent collides with obstacles, it will always return the negative reward and ignore other conditions. In summary, our reward design is the following with decreasing priorities:

$$\begin{aligned} R(s_t) &= r_- < 0, & \text{if } Proj(s_t) \in X_{\mathcal{O}}, \\ R(s_t) &= r_{++} > 0, & \text{if } D(Proj(s_t)) = 0, \\ R(s_t) &= r_+ > 0, & \text{if } D(Proj(s_t)) < D_{min}(s_{t-1}), \\ R(s_t) &= 0, & \text{otherwise,} \end{aligned} \quad (6)$$

Theorem 1. *If there exists at least one policy $\pi_{\mathcal{R},i}^*$ satisfying the reach-avoid task \mathcal{R}_i , by selecting r_{++} to be sufficiently larger than r_+ , i.e., $r_{++} \gg r_+$, any algorithm that optimizes the expected return $J(\pi)$ is guaranteed to find such an optimal policy $\pi_{\mathcal{R},i}^*$.*

Proof. The detailed proof can be found in [31]. \square

Note that even though training deep neural networks to optimize policies is a nonlinear regression process, we can use neural network verification techniques to certify and improve the learned policy $\pi_{\mathcal{R},i}^*$ for a reach-avoid task.

Remark 4. Employing $D_{min}(s_t)$ to design a reward function results in a non-Markovian property, since it relies on the past history \mathbf{s}_t , whereas the reward function (6) only depends on current state. To address this, we can leverage the structure of the product MDP [7], [26], and augment the current state s_t with the index of the closest visited waypoint, to construct a product state to recover the Markovian property.

Distributed DRL: For each \mathcal{R}_i , we can apply the above approach to an existing off-the-shelf DRL algorithms, e.g., SAC [2], to train the optimal policy $\pi_{\mathcal{R},i}^*$. The process can be repeated for all $\mathcal{R}_i, \forall i \in \{0, 1, \dots, K+l\}$ of $\mathcal{R}_{\mathcal{T}} = (\mathcal{R}_0\mathcal{R}_1\dots\mathcal{R}_K)(\mathcal{R}_{K+1}\dots\mathcal{R}_{K+l})^\omega$ in a distributed manner. We then concatenate them as the globally optimal policy

$$\pi_\phi^* = (\pi_{\mathcal{R},0}^*\pi_{\mathcal{R},1}^*\dots\pi_{\mathcal{R},K}^*)(\pi_{\mathcal{R},K+1}^*\dots\pi_{\mathcal{R},K+l}^*)^\omega, \quad (7)$$

where π_ϕ^* is the global policy satisfying LTL task ϕ and containing a set of neural network parameters. Note that if we train each reach-avoid task \mathcal{R}_i individually, we need to randomize the initial velocity and acceleration conditions of state S , since the ending condition of completing a policy $\pi_{\mathcal{R},i}^*$ results in different starting states for the next policy $\pi_{\mathcal{R},i+1}^*$. This enables a smooth concatenation instead of re-training the whole global policy π_ϕ^* .

Theorem 2. *If for all $i \in \{0, 1, \dots, K+l\}$, $\pi_{\mathcal{R},i}^*$ satisfies its corresponding reach-avoid task \mathcal{R}_i of $\mathcal{R}_{\mathcal{T}}$ in (4) under*

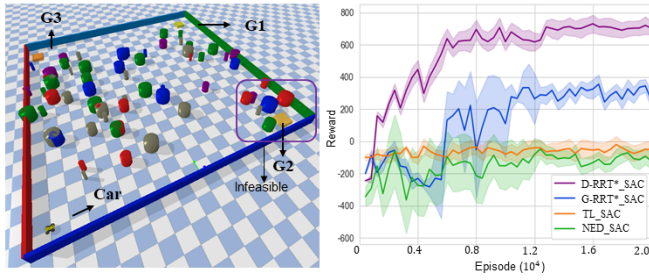


Fig. 3. Baselines comparison of the infeasible task $\phi_{1,inf}$ for a 2D Dubins car in the Pybullet environment with cluttered obstacles.

any initial states (velocities, accelerations), then their concatenation π_ϕ^* in (7) satisfies the LTL task ϕ with minimum-violation guarantees.

Proof. The intuition is to prove that every admissible path $s_\infty^{\pi_\phi^*} \in \mathcal{S}_\infty^{\pi_\phi^*}$ under the optimal policy π_ϕ^* satisfies the LTL specification ϕ with minimum total violation cost. The proof is based on the decomposition property in Section IV-B.

First, the geometric projection of every admissible path $s_\infty^{\pi_\phi^*} \in \mathcal{S}_\infty^{\pi_\phi^*}$ satisfies the lasso-type reach-avoid $\mathcal{R}_\mathcal{T}$ in (4), which is associated with a lasso-type automaton path τ_Q^* as (3). From Lemma 2 and the optimality of TL-RRT* [23], τ_Q^* satisfies the acceptance condition of the NBA of LTL formula ϕ and has the minimum total violation cost. \square

Remark 5. Due to the conclusion of minimally-violating the LTL task from Theorem 2, our framework generalizes to both feasible and infeasible cases for a given LTL task.

V. EXPERIMENTAL RESULTS

We implement the framework on two different nonlinear dynamical systems for various LTL tasks over both finite and infinite horizons. The algorithm test focuses on infeasible cases and cluttered environments, where dense cluttered obstacles are randomly sampled. We show that the framework can complete feasible LTL tasks exactly. We apply SAC [2] as a state-of-art DRL algorithm for all baselines.

Baseline Approaches: From the learning perspective, we refer to our distributed framework as "RRT*" or "D-RRT*", and compare it against three baselines: (i) The relaxed TL-based multi-objective rewards in [6], [18] referred to as "TL", for the single LTL task; (ii) For the goal-reaching task ϕ , the baseline referred to as "NED" designs the reward based on the negative Euclidean distance between the robot and destination; (iii) For a complex LTL task, instead of decomposition, this baseline directly applies the reward scheme (6) for the global trajectory $\tau_F^* = \tau_{pre}^* [\tau_{suf}^*]^\omega$ referred to as "G-RRT*".

From the perspective of infeasible LTL tasks, we compare with the work [17] of visiting as many automaton accepting sets as possible, and we empirically show the improvement over our prior work [7] that assumes the feasible cases.

Autonomous Vehicle We first implement the Dubins car model in the Pybullet physics engine shown in Fig. 3. We consider the surveillance LTL task $\phi_{1,inf} = \square \neg \mathcal{O} \wedge \diamond (\mathcal{G}_1 \wedge \diamond (\mathcal{G}_2 \wedge \diamond (\mathcal{G}_3 \wedge \diamond (\mathcal{G}_{init}))))$ that requires sequentially visiting regions labeled as $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$, and the robot's initial position infinitely often. Its finite-horizon version can be expressed as $\phi_{1,fin} = \square \neg \mathcal{O} \wedge \diamond (\mathcal{G}_1 \wedge \diamond (\mathcal{G}_2 \wedge \diamond (\mathcal{G}_3 \wedge \diamond (\mathcal{G}_{init}))))$ by removing the *always* operator \square . Both tasks are infeasible since the region \mathcal{G}_2 is surrounded by obstacles. Fig. 3 shows the learning curves of task $\phi_{1,inf}$ compared with different

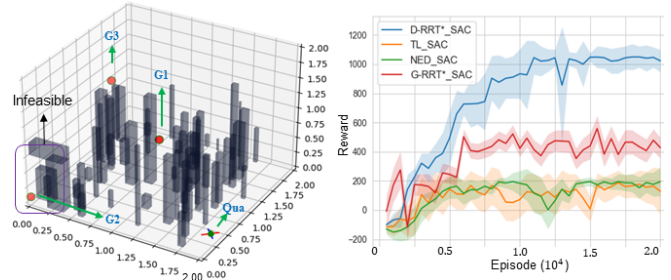


Fig. 4. Baselines comparison of the infeasible task $\phi_{2,inf}$ in the 3D complex cluttered environment.




Fig. 5. LDGBA of the LTL formula $\phi_{1,fin} = \square \neg \mathcal{O} \wedge \diamond (\mathcal{G}_1 \wedge \diamond (\mathcal{G}_2 \wedge \diamond (\mathcal{G}_3)))$ with only one accepting set $F = \{q_3\}$.

baselines, and $\phi_{1,fin}$ has the same comparison results. We can observe that our framework can provide better performance than other baselines under the challenge of complex cluttered environments.

Quadrotor Model We test our algorithms in a 3D environment with Quadrotor dynamics shown in Fig. 4. It demonstrates that our model-free framework is capable of handling complex cluttered environments with high dimensions. The LTL specification is given as $\phi_{2,inf} = \square \neg \mathcal{O} \wedge \diamond \mathcal{G}_1 \wedge \square \diamond \mathcal{G}_2 \wedge \square \diamond \mathcal{G}_3$, which requires the quadrotor to navigate regions of $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ infinitely often without specific orders. Its finite version can also be defined as $\phi_{2,fin} = \square \neg \mathcal{O} \wedge \diamond \mathcal{G}_1 \wedge \diamond \mathcal{G}_2 \wedge \diamond \mathcal{G}_3$. These two task are also infeasible and, their comparisons of learning performances are the same. Due to page limitation, we only show the learning results for the task $\phi_{2,inf}$ in Fig 4.

Feasibility Generalization According to theorem 2, the advantage of minimally violating a given task allows generalization of feasible cases. To show the metric, we remove the highlighted obstacles for the region \mathcal{G}_2 in both environments, and implement our algorithms. The feasible experimental results shown in the video demo demonstrate that our framework can satisfy feasible tasks exactly.

5.5 Infeasibility Analysis For each environment and dynamical system, we increase the task complexity by randomly sampling 12 obstacle-free goal regions in both environments and set the specifications as $\phi_{3,fin} = \square \neg \mathcal{O} \wedge ((\diamond \mathcal{G}_1 \wedge \diamond (\mathcal{G}_2 \wedge \diamond \dots \wedge \diamond \mathcal{G}_{12})))$, and $\phi_{3,inf} = \square \neg \mathcal{O} \wedge \square \diamond \mathcal{G}_1 \wedge \square \diamond \mathcal{G}_2 \dots \wedge \square \diamond \mathcal{G}_{12}$. We repeat the random sampling for 50 trials and record the success rates of generating valid plans for two environments. The results for these task $\phi_{3,fin}, \phi_{3,inf}$ compared with our prior work [7] referred to as "feasible" are shown in Fig. 6 (a). Due to the random process, the target regions may be surrounded by obstacles, and the baseline returns no solutions, whereas our framework always finds minimally-violating trajectories.

From the aspect of automaton-based rewards, the work [18]




Fig. 6. (a) Success rates of random goal regions compared with [7]. (b) Training time for each reach-avoid task compared with distributed baseline "NED" referred as "D-NED".

directly applies a linear combination of violation cost and automaton acceptance reward to formulate a MORL process, with no guarantees on satisfaction and violation. Our proposed framework leveraging planning methods as guidance improves the performance. This seen in the learning comparisons in Fig. 3 and Fig. 4. Moreover, we also compare the work [17] that addresses infeasible cases for discrete MDP model. It finds the policy that satisfies the given LTL tasks as much as possible by intersecting with the maximum number of automaton accepting sets. However, the LDGBAs of some LTL formulas only have one accepting set. As an example, Fig. 5 shows the LDGBA of $\phi_{1,fin} = \square \neg \mathcal{O} \wedge \diamond (\mathcal{G}_1 \wedge \diamond (\mathcal{G}_2 \wedge \diamond (\mathcal{G}_3)))$ for an autonomous vehicle that has only one accepting set $\{q_3\}$. The same case also applies to the LTL formula $\phi_{1,inf}, \phi_{3,fin}$. In such cases, the work [17] returns no solution since the only accepting set is not accessible.

5.6 Training efficiency We record the training time for each reach-avoid sub-task for the finite-horizon LTL tasks i.e., $\phi_{1,fin}, \phi_{2,fin}, \phi_{3,fin}$ in two environments. We compared with the baseline "NED" that distributedly solves each reach-avoid sub-task using distance based rewards. The results in Fig. 6 (b) shows that our framework is more efficient.

VI. CONCLUSION

This article proposes a model-free framework using DRL to learn continuous controllers for completing complex LTL navigation tasks. We overcome the challenge of infeasible LTL specifications in complex cluttered environments. To minimally violate the LTL task, we apply geometric path planning methods as guidance of DRL that can guarantee minimal total violation cost, which can decompose the task into sub-reach-avoid missions. Based on the advantage of minimal violation, our work generalizes both feasible and infeasible cases. Future work will explore safety-critical exploration and bridge the gap of sim-to-real to deploy the framework in real-world robots.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Intl Conference on Machine Learning*, 2018, pp. 1861–1870.
- [3] D. Sadigh, E. S. Kim, S. Coogan, S. S. Sastry, and S. A. Seshia, "A learning based approach to control synthesis of markov decision processes for linear temporal logic specifications," in *IEEE Conference on Decision and Control*, 2014, pp. 1091–1096.
- [4] Y. Kantaros, "Accelerated reinforcement learning for temporal logic control objectives," *Intl Conf on Intelligent Robots and Systems*, 2022.
- [5] M. Cai, S. Xiao, B. Li, Z. Li, and Z. Kan, "Reinforcement learning based temporal logic control with maximum probabilistic satisfaction," in *IEEE Intl Conf on Robotics and Automation*, 2021, pp. 806–812.

- [6] M. Cai, M. Hasanbeig, S. Xiao, A. Abate, and Z. Kan, "Modular deep reinforcement learning for continuous motion planning with temporal logic," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7973–7980, 2021.
- [7] M. Cai, E. Aasi, C. Belta, and C.-I. Vasile, "Overcoming exploration: Deep reinforcement learning for continuous control in cluttered environments from temporal logic specifications," *IEEE Robotics and Automation Letters*, 2023.
- [8] J. Tumova, G. C. Hall, S. Karaman, E. Frazzoli, and D. Rus, "Least-violating control strategy synthesis with safety rules," in *Proc. Int. Conf. Hybrid syst., Comput. Control*, 2013, pp. 1–10.
- [9] M. Lahijanian, M. R. Maly, D. Fried, L. E. Kavraki, H. Kress-Gazit, and M. Y. Vardi, "Iterative temporal planning in uncertain environments with partial satisfaction guarantees," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 583–599, 2016.
- [10] C.-I. Vasile, J. Tumova, S. Karaman, C. Belta, and D. Rus, "Minimum-violation sLTL motion planning for mobility-on-demand," in *IEEE Intl Conference on Robotics and Automation (ICRA)*, 2017, pp. 1481–1488.
- [11] B. Lacerda, F. Faruq, D. Parker, and N. Hawes, "Probabilistic planning with formal performance guarantees for mobile service robots," *Intl Journal of Robotics Research*, vol. 38, no. 9, pp. 1098–1123, 2019.
- [12] H. Rahmani and J. M. O’Kane, "What to do when you can’t do it all: Temporal logic planning with soft temporal logic constraints," in *Intl. Conference on Intelligent Robots and Systems*, 2020, pp. 6619–6626.
- [13] T. Wongpiromsarn, K. Slutsky, E. Frazzoli, and U. Topcu, "Minimum-violation planning for autonomous systems: Theoretical and practical considerations," in *American Control Conference*, 2021, pp. 4866–4872.
- [14] K. Kim, G. Fainekos, and S. Sankaranarayanan, "On the minimal revision problem of specification automata," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1515–1535, 2015.
- [15] M. Guo and D. V. Dimarogonas, "Multi-agent plan reconfiguration under local LTL specifications," *International Journal of Robotics Research*, vol. 34, no. 2, pp. 218–235, 2015.
- [16] M. Cai, S. Xiao, Z. Li, and Z. Kan, "Optimal probabilistic motion planning with potential infeasible ltl constraints," *IEEE Transactions on Automatic Control*, 2021.
- [17] M. Hasanbeig, Y. Kantaros, A. Abate, D. Kroening, G. J. Pappas, and I. Lee, "Reinforcement learning for temporal logic control synthesis with probabilistic satisfaction guarantees," in *IEEE Conference on Decision and Control (CDC)*, 2019, pp. 5338–5343.
- [18] M. Cai, H. Peng, Z. Li, and Z. Kan, "Learning-based probabilistic ltl motion planning with environment and motion uncertainties," *IEEE Tran on Automatic Control*, vol. 66, no. 5, pp. 2386–2392, 2021.
- [19] K. Van Moffaert and A. Nowé, "Multi-objective reinforcement learning using sets of pareto dominating policies," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 3483–3512, 2014.
- [20] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control engineering practice*, vol. 11, no. 7, pp. 733–764, 2003.
- [21] I. R. Manchester and J.-J. E. Slotine, "Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design," *IEEE Tran on Automatic Control*, vol. 62, no. 6, pp. 3046–3053, 2017.
- [22] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin, "Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking," *IEEE Transactions on Automatic Control*, vol. 66, no. 12, pp. 5861–5876, 2021.
- [23] X. Luo, Y. Kantaros, and M. M. Zavlanos, "An abstraction-free method for multirobot temporal logic optimal control synthesis," *IEEE Transactions on Robotics*, 2021.
- [24] C. I. Vasile, X. Li, and C. Belta, "Reactive Sampling-Based Path Planning with Temporal Logic Specifications," *International Journal of Robotics Research*, vol. 39, no. 8, pp. 1002–1028, June 2020.
- [25] E. A. Coddington and N. Levinson, "Theory of ordinary differential equations." *McGraw-Hill*, pp. 1–42, 1955.
- [26] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press, 2008.
- [27] 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, p. eaay6276, 2019.
- [28] M. Y. Vardi and P. Wolper, "An automata-theoretic approach to automatic program verification," in *1st Symposium on Logic in Computer Science*. IEEE Computer Society, 1986, pp. 322–331.
- [29] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [30] D. Q. Mayne, E. C. Kerrigan, E. Van Wyk, and P. Falugi, "Tube-based robust nonlinear model predictive control," *International journal of robust and nonlinear control*, vol. 21, no. 11, pp. 1341–1353, 2011.
- [31] M. Cai, M. Mann, Z. Serlin, K. Leahy, and C.-I. Vasile, "Learning minimally-violating continuous control for infeasible linear temporal logic specifications," *arXiv preprint arXiv:2210.01162*, 2022.