

Co-learning Planning and Control Policies Constrained by Differentiable Logic Specifications

Zikang Xiong, Daniel Lawson, Joe Eappen, Ahmed H. Qureshi, and Suresh Jagannathan

Abstract—Synthesizing planning and control policies in robotics is a fundamental task, further complicated by factors such as complex logic specifications and high-dimensional robot dynamics. This paper presents a novel reinforcement learning approach to solving high-dimensional robot navigation tasks with complex logic specifications by co-learning planning and control policies. Notably, this approach significantly reduces the sample complexity in training, allowing us to train high-quality policies with much fewer samples compared to existing reinforcement learning algorithms. In addition, our methodology streamlines complex specification extraction from map images and enables the efficient generation of long-horizon robot motion paths across different map layouts. Moreover, our approach also demonstrates capabilities for high-dimensional control and avoiding suboptimal policies via policy alignment. The efficacy of our approach is demonstrated through experiments involving simulated high-dimensional quadruped robot dynamics and a real-world differential drive robot (TurtleBot3) under different types of task specifications.

I. INTRODUCTION

Synthesizing planning and control policies is among the core tasks of robotics. The planning policy determines a path comprising a sequence of robot configurations between the given start and goal. In contrast, the control policy helps robots interact with the physical environment allowing them to follow the given plan to reach a goal. However, synthesizing planning and control policies is challenging since the high-dimensional robot dynamics introduce a vast space of possible movements and reactions, making it difficult to predict and control every potential outcome. In addition, the planning policy may be subject to complex logic specifications, which dictate how a robot should behave under certain conditions or constraints. These specifications can range from simple commands like “reaching A”, and “avoiding B” to intricate sets of rules like “surveilling crucial spots and stopping after a certain condition is triggered”. Merging planning and control while adhering to these logic rules and managing the inherent complexities of robot dynamics presents a significant challenge, even for state-of-the-art approaches.

We leverage Reinforcement Learning (RL) for high-dimensional robot planning and control under complex dynamics and sensor observation. To ensure that robots adhere to behavioral rules, temporal logic can be incorporated into RL [1]–[5]. This offers a structured way to define and check a robot’s behavior over time, ensuring tasks are completed in order. However, incorporating temporal logic

into RL presents its own challenges. For example, previous approaches [1]–[10] that have used temporal logic for reward shaping, involves the construction of complicated reward functions that require large amounts of samples [11]. Furthermore, handcrafting these temporal logic specifications can be challenging. For instance, situations where map layouts are represented as images or have irregular obstacle shapes present numerous challenges in defining precise, actionable logical rules. Although existing solvers [12] can find robot paths under a given logical specification, they lack scalability and can only handle linear and, to some extent, quadratic constraints, which is not ideal for practical scenarios, as highlighted in our experiments. Furthermore, these methods only account for path planning without considering errors from the underlying controller or vice versa, which often leads to failure in executions.

Therefore, we propose a novel approach, called Differentiable Specifications Constrained Reinforcement Learning (DSCRL), to address the above challenges using the following key features:

Lower Sample Complexity: We introduce a new methodology that integrates differentiable specifications into constrained RL, which significantly lowers the sample complexity of training.

Control and Planning Alignment: Learning both planning and control policies separately faces the challenges of suboptimal policies due to lack of alignment. We analyze this issue and demonstrate that our approach allows for better alignment between planning and control policies.

Efficient Planning: Our approach generates long horizon plans in cluttered environments, which is typically hard for other RL algorithms as well as existing Signal Temporal Logic (STL) solvers to handle.

Specifications from Image: Our approach can extract irregular obstacle specifications directly from different map images, and solve them efficiently with a novel neural planning policy built upon pre-trained backbones.

High-dimensional Policy: Our approach has been effectively tested on simulated high-dimensional quadruped robot dynamics and a real-world differential drive robot (TurtleBot3) with LiDAR observation.

II. RELATED WORK

Temporal Logic and RL: This work considers symbolic specifications written in temporal logic [13], [14], which are widely used in specifying various planning and control tasks [15]–[18]. Along with the advances of RL in stochastic, model-free settings, a line of work [1]–[10] have considered

Authors affiliate to Computer Science Department, Purdue University, IN 47906, USA. (E-mail: xiong84@purdue.edu; lawson95@purdue.edu; jeappen@purdue.edu; ahqureshi@purdue.edu; suresh@cs.purdue.edu)

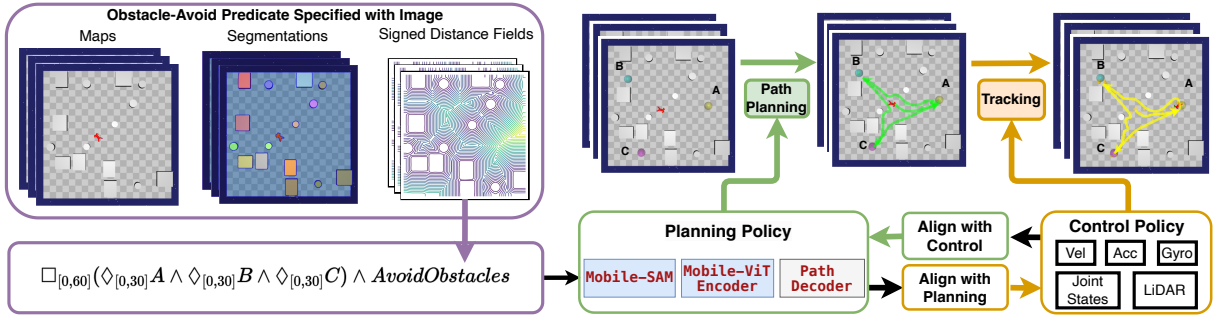


Fig. 1: The differentiable specification encodes a task visiting location A, B, C repeatedly, and the predicate AvoidObstacles is a Signed Distance Field (SDF) directly specified from a map image. Differentiable symbolic specifications constrain policy updates and notably reduce the *sample complexity* of training policies. Based on a different initial position and map layouts, the planning policy crafts paths in compliance with the given specification. The control policy is a goal-conditioned policy, following the paths laid out by the planner. It learns to meet goals generated by the planning policy. The planning policy and control policy are aligned with each other during training.

combining LTL-specified tasks with reinforcement learning. However, these approaches formulate such integration using temporal logic as rewards, leading to inefficient or even intractable algorithms [19].

Constrained RL: Constrained reinforcement learning [20]–[22] provides an efficient approach to enforce constraints on policies. Given a differentiable specification, these methods can enforce constraints only on the policy without shaping rewards. However, prior work did not consider how to synergistically take advantage of policies with differentiable formal constraints [23].

Gradient and Neural Path Planning: Our work also relates to gradient-based motion planning [23]–[27]. While these techniques are typically only applicable when the dynamics are known and deterministic, our focus is on planning in the face of unknown and stochastic dynamics. Our high-level planning policy is related to planning networks [28]. However, learning a planning neural network in an RL loop with task specifications, which aligns both the planning policy and control policy, has not been explored previously.

III. BACKGROUND

A. Policy Gradient

Consider a trajectory τ sampled from policy π_θ . The expected return $R(\tau)$ for a policy parameterized by θ is given by $J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau)]$. The policy gradient theorem [29] states that the gradient of the expected return is given by $\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [\nabla_\theta \mathcal{P}(\pi_\theta, \tau) R(\tau)]$, where $\mathcal{P}(\pi_\theta, \tau)$ is the log probability of trajectory τ under policy π_θ . This implies that the probability of generating a trajectory τ with a high return is maximized compared to those with a lower return. In practice, a baseline $b(\tau)$ is also subtracted from the reward to reduce the variance of the gradient estimator [30].

B. Signal Temporal Logic

The syntax of STL contains both first-order logic operators \wedge (and), \neg (not), \vee (or), \Rightarrow (implies), and temporal operators \bigcirc (next), $\Diamond_{[a,b]}$ (eventually), $\Box_{[a,b]}$ (globally), $\mathcal{U}_{[a,b]}$ (until). The initial time a and end time b “truncate” a path. For example, $\Box_{[a,b]}$ qualifies property that globally holds during

time a and b . The syntax of STL is recursively defined via the following grammar:

$$\begin{aligned} \phi := & \top \mid \perp \mid \mathcal{F} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \Rightarrow \psi \\ & \mid \bigcirc\phi \mid \Diamond_{[a,b]}\phi \mid \Box_{[a,b]}\phi \mid \phi \mathcal{U}_{[a,b]}\psi, \end{aligned} \quad (1)$$

where $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{R}$ is a predicate function mapping a state to a real value (e.g., a function computes the distance to an obstacle surface). The STL semantics $\rho(\tau, \phi) : \mathbb{R}^{T \times n} \rightarrow \mathbb{R}$ is defined over a formula ϕ and a trajectory τ with length T , and maps them to a real value [23]. For convenience, we denote $\phi(\tau) := \rho(\tau, \phi)$. As an example, the semantics of $\Box_{[a,b]}\phi$ is:

$$\rho(\tau, \Box_{[a,b]}\phi) = \min_{t \in [a,b]} \rho(\tau_{[t:b]}, \phi), \quad (2)$$

meaning that the formula $\Box_{[a,b]}\phi$ holds if and only if ϕ holds at every time step ($\forall t \in [a,b], \phi(\tau_{[t:b]}) > 0$). $\tau_{[t:b]}$ is a path slice from t to b . The full semantics can be found in previous works such as [23], [31]. In practice, \min, \max in STL semantics can cause gradient vanishing, which can be alleviated by using soften techniques [31].

IV. APPROACH

This section presents the key components of our method for solving complex robot navigation tasks. In Sec. IV-A, we will introduce a novel planning policy π_ϕ^h which generates paths based on task specification ϕ and map images, while also aligns with the control policy π^l . Then, we introduce the control policy π^l which follows paths and aligns itself with the planning policy π_ϕ^h in Sec. IV-B.

A. Planning Policy

The planning policy π_ϕ^h is trained to meet a task specification ϕ . Given an initial position x_0 and an arbitrary map image I , the planned path τ is sampled from the distribution predicted by the planning policy, where $\tau = (g_0, g_1, \dots, g_T)$, $g_0 = x_0$, and T is the max time horizon, i.e., $\tau \sim \pi_\phi^h(x_0, I)$. STL has the following definition:

$$\phi(\tau) = \begin{cases} > 0, & \text{if } \tau \text{ satisfies } \phi, \\ < 0, & \text{if } \tau \text{ does not satisfy } \phi. \end{cases} \quad (3)$$

1) *Differentiable STL and Control Feedback*: The quantitative semantics of ϕ is provided by STL [14]. Because ϕ is differentiable [23], we can directly use ϕ as a part of the training objective, and turn the policy search into a gradient optimization problem. Meanwhile, to align the planning policy π_ϕ^h with the feedback from the control policy π^l , while also satisfying ϕ , we formulate the problem by minimizing the following loss function:

$$\mathcal{L}_{\pi_\phi^h} = -\mathbb{E}_{\substack{I \sim \mathcal{I}, x_0 \sim \mathcal{X}_0, \\ \tau \sim \pi_\phi^h(x_0, I)}} \left(\underbrace{\mathcal{P}(\tau, \pi_\phi^h) \cdot r^h(\tau, \pi^l)}_{\text{PG}} + \lambda \phi(\tau) \right). \quad (4)$$

The reward $r^h(\tau, \pi^l)$ here is the negative of the steps to reach the goal g_T . When following the path τ with the control policy π^l , fewer steps mean a higher reward. The $\mathcal{P}(\tau, \pi_\phi^h)$ is the log probability of the path τ under policy π_ϕ^h . The task specification ϕ is formulated as a Lagrangian constraint with a multiplier λ . Optimizing (4) will maximize the expectation of reward, which is positively correlated to the expectation of $\mathcal{P}(\tau, \pi_\phi^h) \cdot r^h(\tau, \pi^l)$, and generate paths satisfied ϕ (i.e., $\phi(\tau) > 0$). In practice, $r^h(\tau, \pi^l)$ will subtract a baseline to reduce the policy gradient variance [30]. The Lagrangian multiplier λ can be dynamically updated for better performance [32].

One may only train the planning policy π_ϕ^h with the task specification ϕ without control feedback (i.e., ignoring $\mathcal{P}(\tau, \pi_\phi^h) \cdot r^h(\tau, \pi^l)$). However, due to the complexity of the high-dimensional dynamics we are handling, it is impractical to encode dynamics constraints in STL. Thus, despite the fact that the planner trained with $\phi(\tau)$ can generate paths satisfying ϕ , these paths are not necessary to be easy to follow by the control policy. For example, it may generate a path with lots of sharp turns. More importantly, choosing the goal for a high-dimensional, non-linear goal-conditioned control policy π^l is non-trivial. Naively choosing a path based on Euclidean distance does not necessarily mean the minimum in the steps. Thus, optimizing the π_ϕ^h with control feedback is necessary to avoid suboptimal policies.

Previous RL approaches [1]–[10] have investigated training policies with STL. However, these approaches are based on reward shaping, which is different from our constraint-based formalization. The reward-shaping formulation minimizes the following loss function:

$$\mathcal{L}_{RS} = -\mathbb{E}_{\substack{I \sim \mathcal{I}, x_0 \sim \mathcal{X}_0, \\ \tau \sim \pi_\phi^h(x_0, I)}} \left(\underbrace{\mathcal{P}(\tau, \pi_\phi^h) \cdot \left(r^h(\tau, \pi^l) + \lambda \phi(\tau) \right)}_{\text{PG}} \right). \quad (5)$$

The policy gradient part of (4) and (5) is marked as PG, which backpropagates through $\mathcal{P}(\tau, \pi_\phi^h)$. Getting an accurate policy gradient typically requires many samples, and a more complex reward function typically requires more samples to estimate the policy gradient. Thus, as illustrated in (4), decoupling the complex STL specification ϕ from the policy gradient part and directly backpropagating through ϕ (as well

as $\mathcal{P}(\tau, \pi_\phi^h)$) is expected to reduce the sample complexity in training (i.e., converge faster). In addition, the estimated policy gradients are only meaningful around existing samples [20], [22]. However, during iterative updates, the policy can quickly shift to a region with few samples and make the policy gradient updates random. In contrast, when backpropagating through ϕ , as the gradient is not estimated from “local” samples, but computed from ϕ , our objective (4) also demonstrates more stable performance during training.




Fig. 2: The planning policy model has three parts. The Mobile-SAM [33], whose input is an arbitrary map image, outputs a mask of obstacles. The mask image is then fed into a Mobile-ViT encoder [34] that outputs the map embedding. The transformer path decoder takes the map embedding and the initial position as input and outputs the deviation between subgoals. Mobile-SAM and Mobile-ViT are pre-trained models with frozen weights (denoted by ❄️), and the decoder is trained with the loss function in (4). Both the Mobile-SAM and Mobile-ViT are optimized for the edge devices, the path can be generated in 0.2 seconds on a consumer-grade GPU.

2) *Model Design of Planning Policy*: Given an arbitrary map layout image I , the planning policy model processes it in stages to make path planning efficient. Initially, to reduce the noise inherent in real images, the MobileSAM [33] extracts an obstacle mask M from the map image I . Subsequently, the MobileViT encoder [34] maps this mask image M into a map embedding E . Lastly, $Decoder_\theta$, parameterized by θ (the component we exclusively trained) computes subgoal deviations, denoted as $\frac{d\tau}{dt}$, using the map embedding E and the robot’s initial position x_0 :

$$\frac{d\tau}{dt} = \tanh(Decoder_\theta(E, x_0)) \times \Delta g. \quad (6)$$

Here, Δg is the maximum deviation between subgoals. The deviation is scaled by a \tanh function to ensure that the subgoals are within a reasonable range. The subgoal in a planned path $\tau = [g_0, g_1, \dots, g_T]$ is generated by integrating the deviation with the initial position x_0 : $g_t = x_0 + \int_0^t \frac{d\tau}{dt} dt$. The decoder is a transformer decoder [35] which will be trained with the loss function in (4).

3) *Extract Obstacle Specification from Map Image*: One limitation of STL is that all specifications have to be encoded by users, including all obstacles on the map. This can be challenging as the map layout becomes complicated and implicit (e.g., the map is given as an image). In addition, state-of-the-art STL solvers, such as [12], can only encode

obstacles with linear or limited quadratic constraints, making it challenging to handle real-world obstacles with irregular shapes.

Our approach extends STL with specifications directly generated from images, by converting them into a Signed Distance Field (SDF) [36]. The SDF is a function that returns the distance to the nearest obstacle surface for any given point in an obstacle mask M :

$$SDF(g, M) = \text{sign}(g) \cdot \min_{p \in M} \|g - p\|_2, \quad (7)$$

Here, $\text{sign}(g)$ is negative if g is inside an obstacle, and positive if g is outside an obstacle. We directly use SDF as an obstacle avoidance predicate in STL, which is consistent with the semantics of STL (3). As an example, the predicate $AvoidObstacles$ in Fig. 1 is defined as:

$$AvoidObstacles = \square_{[0, T]} SDF(\cdot, M). \quad (8)$$

The $\square_{[0, T]}$ is a min operator defined in the path $\tau = [g_0, g_1, \dots, g_T]$, which returns the minimum value of $SDF(g_t, M)$ for all g_t . In other words, $AvoidObstacles(\tau)$ is positive if and only if $\forall g_t \in \tau, SDF(g_t, M) > 0$, which means that all g_t are outside the obstacles.

The mask M is stored as a binary image, where the obstacle pixels are 1 and the free space pixels are 0. We can extract an SDF from the mask image M with a KD-Tree [37]. Given a mask image with $W \times H$ pixels, we build a KD-Tree with the outline points of the obstacles, where each point is a 2D index of a pixel in the mask image. A $KDTree_M$ constructed specifically for a mask M , when input a pixel index, returns the nearest pixel that is part of an obstacle outline within the mask M . If the field size is $m \times n$ meters in the real world, we can compute a transformation \mathcal{T} from the pixel index to the real-world coordinate, and its inversion \mathcal{T}^{-1} , with W, H, m, n . An SDF is implemented as:

$$SDF(g, M) = \text{sign}(g) \|\mathcal{T} \cdot KDTree_M(\mathcal{T}^{-1} \cdot g) - g\|_2. \quad (9)$$

The gradient can be backpropagated through (9) to update the planning policy.

B. Control Policy

The low-level control policy $\pi^l(o_t | g)$ is a goal-conditioned policy [38] trained with PPO [22], and o_t is the robot observation at time t . The control reward r_t^l at time t is defined as

$$r^l(o_t, o_{t+1}, g) = \|g - pos(o_t)\|_2 - \|g - pos(o_{t+1})\|_2, \quad (10)$$

where $pos(o_t)$ and $pos(o_{t+1})$ are the positions of the robot at time t and $t + 1$, respectively. The reward is positive if and only if the robot gets closer to the planned goal g .

The planning policy π_ϕ^h generates the goal distribution \mathcal{G} for the control policy. As the planning policy is updated, \mathcal{G} will change. Thus, the control policy π^l needs to be updated to follow the new goal distribution. Formally, the general

goal for the goal-conditioned RL is to find an optimal policy π^* with

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}_{g \sim \mathcal{G}} \mathbb{E}_{\substack{o_0 \sim \mathcal{O}_0, a_t \sim \pi(o_t | g), \\ o_{t+1} \sim P(o_{t+1} | o_t, a_t)}} \left(\sum_{t=0}^T \gamma^t r^l(o_t, o_{t+1}, g) \right), \quad (11)$$

where \mathcal{O}_0 is the initial observation distribution, and $P(o_{t+1} | o_t, a_t)$ is the transition function. The γ is the discount factor.

Let $V^\pi(g) = \mathbb{E}_{\substack{o_0 \sim \mathcal{O}_0, a_t \sim \pi(o_t | g), \\ o_{t+1} \sim P(o_{t+1} | o_t, a_t)}} \left(\sum_{t=0}^T \gamma^t r^l(o_t, o_{t+1}, g) \right)$ be the value function of the control policy π with goal g , the objective function in (11) can be rewritten as:

$$\pi^* = \operatorname{argmax}_{\pi} \sum_g P_G(g) V^\pi(g). \quad (12)$$

The $P_G(g)$ is the probability of the goal g under the goal distribution \mathcal{G} . The change of the goal distribution \mathcal{G} (as planning policy updates) will change P_G , and thus the optimal policy π^* shifts. Thus, the control policy π^l needs to be updated for the new goal distribution. This can be achieved by training the control policy with any RL algorithm by sampling the goal g with the new goal distribution [38].

C. Overall Pipeline

We summarize the training and deployment pipeline of our approach in this section. First, users express their goals without delving into intricate maps, as exemplified by $\square_{[0, 60]}(\diamond_{[0, 30]}A \wedge \diamond_{[0, 30]}B \wedge \diamond_{[0, 30]}C)$ in Fig. 1. Next, the $AvoidObstacles$ specification is derived from map images with the method discussed in Sec. IV-A.3. For policy training, we initially train π^l using goals sampled from π_ϕ^h while keeping π_ϕ^h fixed. Then, we train the planning policy π_ϕ^h using (4) with the control policy π^l fixed. This back-and-forth is important as simultaneous updates can introduce instability due to the evolving objectives in (4) and (12). We iterate through this process until both policies converge. During deployment, the planning policy π_ϕ^h generates a path τ from the initial position x_0 and an arbitrary map image I . The control policy π^l follows τ to achieve the task ϕ .

V. EXPERIMENTS

A. Setup

1) *Tasks*: The tasks evaluated are depicted in Fig. 3. The STL specifications of the *Sequence* task is $\bigwedge_{i=1}^N \diamond_{[t_i, t_{i+1}]} \phi_i$, where ϕ_i is a predicate that corresponds to reaching each blue region. The robot must visit the region specified by $\phi_1, \phi_2, \dots, \phi_N$ in sequence. The *Cover* task is specified by $\bigwedge_{i=1}^N \diamond_{[0, T]} \phi_i$. Here the robot is asked to cover all the specified regions in any given order. The *Branch* task is specified by $\bigvee_{i=1}^{\frac{N}{2}} (\diamond_{[0, T]} \phi_i \wedge \diamond_{[0, T]} \phi_{\frac{N}{2}+i})$, the robot must either visit the branch of blue regions (ϕ_1, ϕ_3) or the branch of green regions (ϕ_2, ϕ_4). The *Loop* task $\square_{[0, T - \lfloor \frac{T}{M} \rfloor]} (\bigwedge_{i=1}^N \diamond_{[0, \lfloor \frac{T}{M} \rfloor]} \phi_i)$ asks the robot to cyclically traverse the blue regions specified by ϕ_i for M times. Lastly, the *Signal* task is specified by $\square_{[0, T - \lfloor \frac{T}{M} \rfloor]} (\bigwedge_{i=1}^N \diamond_{[0, \lfloor \frac{T}{M} \rfloor]} \phi_i) \mathcal{U}_{[0, 1]} \psi_1 \wedge \diamond_{[0, T]} \psi_2$. The robot should loop among the three blue regions (ϕ_i) until




Fig. 3: In the *Sequence* task, the *Turtlebot3* must visit the three blue regions sequentially. The *Cover* task requires the robot to traverse all three blue regions in any sequence. For the *Branch* task, the robot chooses between visiting either the two blue regions or the two green regions. The *Loop* task asks the robot to cyclically traverse the three blue regions. Lastly, the *Signal* task requires the robot to loop among three blue regions and exit the loop after visiting the lower blue region twice, and finally visiting the yellow region. The yellow tape in the center is used for calibrating the coordinates.

it visits the lower blue region twice (ψ_1) and finally visits the yellow region (ψ_2). All these specifications are appended with $\wedge \square_{[0,T]} SDF(\cdot, M)$ extracted from map M .

2) *Observation, Action, and Dynamics*: The *Doggo* (shown in Fig. 1) features a 74-dimensional observation space: it has joint states (44D), IMU data (12D), position (2D), and a 16-beam LiDAR for obstacle detection. Its action space is 12-dimensional, comprising 8 hip and 4 ankle control commands. On the other hand, *TurtleBot3* (shown in Fig. 3) utilizes 4-dimensional IMU data (including x-axis position, y-axis position, and the sin and cos of rotation) and a 36-dimensional down-sampled LiDAR for measuring obstacle distances. Its action commands are dedicated to controlling the robot’s linear and angular velocity.

B. Policy Performance and Alignment

1) *Metrics*: We evaluate the planning and control policies using Success Rate (SR) and Time to Reach (TtR) metrics. SR quantifies the proportion of successful plans. A plan is considered successful if the path τ satisfies the condition ($\phi(\tau) > 0$) without any collisions when followed by the control policy. TtR calculates the duration (in seconds) needed to achieve the final destination. Each learned policy is evaluated on 1000 episodes with different map layouts. There are 10 obstacles for *TurtleBot3* with 2.42×2.42 meters maps and 15 obstacles for *Doggo* with 3×3 meters maps.

TABLE I: Alignment Experiments. The Success Rate (SR) and Time to Reach (TtR) of the control policy are evaluated on 1000 episodes with randomly sampled map layouts. The performance of aligned policies is reported in the column “w.”, while the unaligned policies is reported in the column “w.o.”. The aligned policies perform significantly better than the unaligned policies. All the aligned policies have SR above 96%.

Robot	<i>Doggo</i>				<i>Turtlebot3</i>			
Metrics	SR \uparrow		TtR \downarrow		SR \uparrow		TtR \downarrow	
Align	w.o.	w.	w.o.	w.	w.o.	w.	w.o.	w.
<i>Seq</i>	87.1%	97.7%	64.6	46.3	79.1%	98.2%	105.6	92.7
<i>Cover</i>	83.9%	96.4%	73.7	57.0	73.1%	96.1%	92.5	81.9
<i>Branch</i>	87.2%	98.2%	56.4	44.6	91.1%	97.2%	81.4	73.8
<i>Loop</i>	83.1%	98.5%	70.6	51.3	85.1%	98.2%	114.3	94.5
<i>Signal</i>	79.9%	97.1%	72.1	57.3	79.6%	97.8%	118.1	106.5

2) *Performance and Alignment*: Experiments in Table I show the performance of our final trained policies and the effect of policy alignment. The unaligned planning policy π_ϕ^h is trained with STL objective ϕ without control feedback, and the unaligned control policy is trained on random goals. The aligned policies are trained with our approach. All the aligned policies have SR above 96%, which is significantly higher than the unaligned policies. The aligned policies also have lower TtR than the unaligned policies. The results demonstrate that our approach can work effectively on high-dimensional control tasks with various specifications and maps with high SR and low TtR. Moreover, the results also show that the alignment between planning and control policies is essential for good performance, justifying our algorithm design.

TABLE II: Deploy trained policy to real TurtleBot3. The first row shows the specs. The second and third rows are identical to TABLE I. The results are computed over 6 sampled episodes each.

Spec	<i>Seq</i>				<i>Signal</i>			
Metrics	SR \uparrow		TtR \downarrow		SR \uparrow		TtR \downarrow	
Align	w.o.	w.	w.o.	w.	w.o.	w.	w.o.	w.
Stat.	3/6	5/6	142.7	108.3	1/6	6/6	144.7	109.1

3) *Policy Performance In A Real TurtleBot3*: We evaluated our policies with and without alignment on a real Turtlebot3 robot for the *Sequence* and *Signal* tasks. The results in Table II show that our policies can work effectively in the real world with high SR. Compared with the unaligned policies, the aligned policies have higher SR and lower TtR.

C. Sample Efficiency and Comparison with Reward Shaping

We focused on the number of samples required for our policies to reach convergence. Each sample represents a transition, moving from one observation o_t to another o_{t+1} upon taking an action a_t and acquiring the corresponding reward r_t^l . The high-level reward r^h is determined after we collect a full trajectory τ .

To evaluate the efficacy of our methodology, we compare it with the Reward Shaping (RS) strategy [2] illustrated in (5). Additionally, we benchmarked against the Reward Machine (RM) [5] strategy that shapes rewards with the progression of a pre-defined planned path. For example, given a planned

path A, B, C , if the planning policy did not reach A , the reward is shaped to encourage the planning policy to reach A ; if the planning policy reached A but did not reach B , the reward is shaped to encourage the planning policy to reach B , and so on. Notice that the RM strategy requires an additional solver to generate the planned path, which is a stronger requirement than ours. Our experimentation spanned two distinct settings:

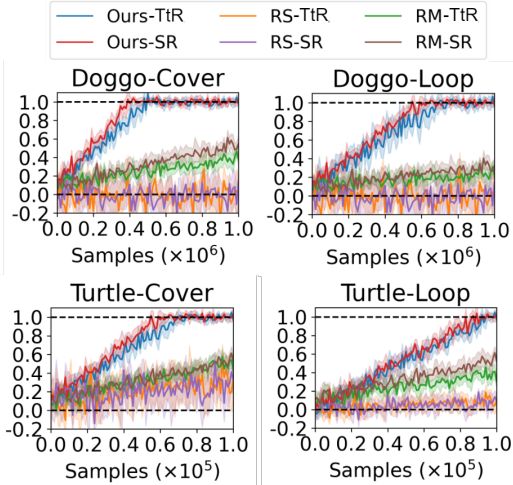


Fig. 4: Align pretrained policies on *Cover* and *Loop* tasks. The x-axis is the number of samples, and the y-axis is the normalized score. The inverse of the TtR is plotted where a value of 1.0 represents the smallest TtR achieved. The upper black dashed line is the score of the aligned policies (“w.” columns in Table I), and the lower black dashed line is the score of the unaligned policies (“w.o.” columns in Table I). The values are computed over 5 different random seeds. The results show that our approach can converge with fewer samples than RS and RM.

1) *Two-Stage Training*: In this setup, we first train the planning policy π_ϕ^h with STL objective ϕ without control feedback, while the control policy is trained on random goals. Later, in our approach, π_ϕ^h was finetuned using (4). In contrast, RS trains the π_ϕ^h with (5), while RM trains π_ϕ^h with rewards crafted from a path satisfying ϕ . Following this, for all methods, π^l is trained on goals sampled from π_ϕ^h . Unlike RS and RM, our method decouples the STL objective from the reward and converges faster (Fig. 4). The trade-off between sample efficiency and reward function complexity is also evident in RS and RM, with RM outperforming RS with simplified reward functions.

TABLE III: Converging samples when training from scratch. The data shows the number of samples required for each algorithm to converge to the scores near Table I ($\pm 5\%$). The data is scaled by the coefficient c next to the robot name. A dash “-” means that it does not converge in $100 \times c$ samples. The results are computed over 5 random seeds.

Robot	<i>Doggo</i> ($\times 10^7$)			<i>Turtlebot3</i> ($\times 10^6$)		
Algo	Ours	RS	RM	Ours	RS	RM
<i>Cover</i>	5.1 ± 1.2	-	27.4 ± 7.1	7.9 ± 1.9	69.2 ± 17.1	30.9 ± 5.4
<i>Loop</i>	5.7 ± 0.9	-	-	9.1 ± 2.8	-	48.7 ± 4.9

2) *End-to-End Training*: Alternatively, we conduct evaluations by training from scratch where both π_ϕ^h and π^l were initialized with random weights. In Table III, we observe that our approach consistently requires fewer samples to converge when compared to the other two approaches. The results also show RM performs better than RS, which is consistent with the results in Fig. 4.

D. More Scalable STL Planning

One alternative approach for our problem is solving the STL specifications with an STL solver (such as STLPY [12]) and tracking the solved path with a goal-conditioned control policy. However, this approach has its limitations. First, manual obstacle specification is limited to linear or quadratic constraints. Second, paths produced by STLPY, while sound and complete w.r.t specifications, can still lead to collisions or stalls when tracking with an unaligned control policy. In the *Loop* task with 100 timesteps, STLPY paths succeed 14 out of 20 times for *Doggo* and 13 out of 20 times for *TurtleBot3* with the optimal robustness parameter [12]. The planning policy’s SR is 98.5% for *Doggo* and 98.2% for *TurtleBot3* on *Loop* task. Third, as noted by [12], the STL




Fig. 5: Planning time against time horizon and obstacle count. The spec for the left is $\square_{[0, T - \lfloor \frac{T}{3} \rfloor]} (\bigwedge_{i=1}^3 \diamond_{[0, \lfloor \frac{T}{3} \rfloor]} \phi_i)$, where T is the time horizon. The spec for the right is $\square_{[0, 27]} (\bigwedge_{i=1}^3 \diamond_{[0, 13]} \phi_i) \wedge \bigwedge_{i=1}^n Obs_i$, where $\bigwedge_i Obs_i$ represents the rectangle obstacle predicates with linear constraints, and n is the number of obstacles. Results from 20 runs per data point show STLPY’s solving time grows exponentially with horizon and obstacles, but the neural planning policy remains consistently fast (< 0.5 secs).

solver has scalability issues shown in Fig. 5. The planning time for the STL solver increases exponentially with the time horizon and the number of obstacles.

VI. CONCLUSIONS

In this paper, we present a new method combining differentiable specifications with constrained RL, which reduces training sample complexity and aligns planning and control policies more efficiently than current reward-shaping techniques. Our method performs well in long-horizon planning in cluttered environments, extracts complex obstacle details from images, and uses them efficiently with our neural planning policy. It also demonstrates advantages in high-dimensional dynamics. Although our framework can adapt to different maps, one limitation is that every planning policy is trained only for one specification. Encoding the specification semantics, with a large language model (LLM) backbone, into the planning policy and training a single policy for multiple specifications will be a path to explore further.

REFERENCES

- [1] J. Fu and U. Topcu, "Probably approximately correct mdp learning and control with temporal logic constraints," *arXiv preprint arXiv:1404.7073*, 2014.
- [2] 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 (IROS)*. IEEE, 2017, pp. 3834–3839.
- [3] 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 *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 5338–5343.
- [4] K. Jothimurugan, R. Alur, and O. Bastani, "A composable specification language for reinforcement learning tasks," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [5] R. T. Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith, "Reward machines: Exploiting reward function structure in reinforcement learning," *Journal of Artificial Intelligence Research*, vol. 73, pp. 173–208, 2022.
- [6] M. Hasanbeig, D. Kroening, and A. Abate, "LCRL: Certified policy synthesis via logically-constrained reinforcement learning," in *International Conference on Quantitative Evaluation of Systems*. Springer, 2022.
- [7] Y. Jiang, S. Bharadwaj, B. Wu, R. Shah, U. Topcu, and P. Stone, "Temporal-logic-based reward shaping for continuing learning tasks," *arXiv preprint arXiv:2007.01498*, 2020.
- [8] A. K. Bozkurt, Y. Wang, M. M. Zavlanos, and M. Pajic, "Control synthesis from linear temporal logic specifications using model-free reinforcement learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 349–10 355.
- [9] Z. Xu, I. Gavran, Y. Ahmad, R. Majumdar, D. Neider, U. Topcu, and B. Wu, "Joint inference of reward machines and policies for reinforcement learning," *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 30, no. 1, pp. 590–598, Jun. 2020.
- [10] H. Zhang and Z. Kan, "Temporal logic guided meta q-learning of multiple tasks," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8194–8201, 2022.
- [11] G. Dulac-Arnold, D. Mankowitz, and T. Hester, "Challenges of real-world reinforcement learning," *arXiv preprint arXiv:1904.12901*, 2019.
- [12] V. Kurtz and H. Lin, "Mixed-integer programming for signal temporal logic with fewer binary variables," *IEEE Control Systems Letters*, 2022.
- [13] A. Pnueli, "The temporal logic of programs," in *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, 1977, pp. 46–57.
- [14] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*. Springer, 2004, pp. 152–166.
- [15] G. E. Fainekos, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for mobile robots," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE, 2005, pp. 2020–2025.
- [16] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 53, no. 1, pp. 287–297, 2008.
- [17] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [18] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning," *IEEE Transactions on Automatic Control*, vol. 57, no. 11, pp. 2817–2830, 2012.
- [19] C. Yang, M. L. Littman, and M. Carbin, "On the (in)tractability of reinforcement learning for LTL objectives," in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, 2022, pp. 3650–3658.
- [20] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.
- [21] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *International conference on machine learning*. PMLR, 2017, pp. 22–31.
- [22] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [23] K. Leung, N. Aréchiga, and M. Pavone, "Backpropagation through signal temporal logic specifications: Infusing logical structure into gradient-based methods," *Int. Journal of Robotics Research*, 2022.
- [24] K. Leung, N. Aréchiga, and M. Pavone, "Back-propagation through signal temporal logic specifications: Infusing logical structure into gradient-based methods," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2020, pp. 432–449.
- [25] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 489–494.
- [26] M. Campana, F. Lamiroux, and J.-P. Laumond, "A gradient-based path optimization method for motion planning," *Advanced Robotics*, vol. 30, no. 17-18, pp. 1126–1144, 2016.
- [27] C. Dawson and C. Fan, "Robust counterexample-guided optimization for planning from differentiable temporal logic," *arXiv preprint arXiv:2203.02038*, 2022.
- [28] A. H. Qureshi, A. Simeonov, M. J. Bency, and M. C. Yip, "Motion planning networks," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2118–2124.
- [29] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, pp. 229–256, 1992.
- [30] C. Wu, A. Rajeswaran, Y. Duan, V. Kumar, A. M. Bayen, S. Kakade, I. Mordatch, and P. Abbeel, "Variance reduction for policy gradient with action-dependent factorized baselines," *arXiv preprint arXiv:1803.07246*, 2018.
- [31] Y. Gilpin, V. Kurtz, and H. Lin, "A smooth robustness measure of signal temporal logic for symbolic control," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 241–246, 2021.
- [32] A. Stooke, J. Achiam, and P. Abbeel, "Responsive safety in reinforcement learning by pid lagrangian methods," in *International Conference on Machine Learning*. PMLR, 2020, pp. 9133–9143.
- [33] C. Zhang, D. Han, Y. Qiao, J. U. Kim, S.-H. Bae, S. Lee, and C. S. Hong, "Faster segment anything: Towards lightweight sam for mobile applications," *arXiv preprint arXiv:2306.14289*, 2023.
- [34] S. Mehta and M. Rastegari, "Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer," *arXiv preprint arXiv:2110.02178*, 2021.
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [36] R. Malladi, J. A. Sethian, and B. C. Vemuri, "Shape modeling with front propagation: A level set approach," *IEEE transactions on pattern analysis and machine intelligence*, vol. 17, no. 2, pp. 158–175, 1995.
- [37] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Commun. ACM*, vol. 18, no. 9, p. 509–517, sep 1975. [Online]. Available: <https://doi.org/10.1145/361002.361007>
- [38] T. Schaul, D. Horgan, K. Gregor, and D. Silver, "Universal value function approximators," in *International conference on machine learning*. PMLR, 2015, pp. 1312–1320.