

Follow The Rules: Online Signal Temporal Logic Tree Search for Guided Imitation Learning in Stochastic Domains

Jasmine Jerry Aloor^{*1}, Jay Patrikar^{*2}, Parv Kapoor², Jean Oh² and Sebastian Scherer²

Abstract—Seamlessly integrating rules in Learning-from-Demonstrations (LfD) policies is a critical requirement to enable the real-world deployment of AI agents. Recently, Signal Temporal Logic (STL) has been shown to be an effective language for encoding rules as spatio-temporal constraints. This work uses Monte Carlo Tree Search (MCTS) as a means of integrating STL specification into a vanilla LfD policy to improve constraint satisfaction. We propose augmenting the MCTS heuristic with STL robustness values to bias the tree search towards branches with higher constraint satisfaction. While the domain-independent method can be applied to integrate STL rules online into any pre-trained LfD algorithm, we choose goal-conditioned Generative Adversarial Imitation Learning as the offline LfD policy. We apply the proposed method to the domain of planning trajectories for General Aviation aircraft around a non-towered airfield. Results using the simulator trained on real-world data showcase 60% improved performance over baseline LfD methods that do not use STL heuristics. [Code]¹ [Video]²

I. INTRODUCTION

In many real-world domains, the ability to learn from multiple interactions with the environment is either prohibitively expensive or comes with safety concerns. Providing embodied AI agents with the ability to learn effectively from offline datasets of human demonstrations is thus critical in our pursuit to deploy them reliably in real-world domains. Prior works have demonstrated limited success when Learning from Demonstrations (LfD) is treated as a purely supervised learning task [1]. Expert demonstrations are often noisy, incomplete, or both. Thus, pure LfD policies, like Behavior Cloning (BC), perform suboptimally when deployed in the real world. They also often fail to *distill* the underlying rules and constraints that guide the behaviors of the expert agent. These issues are especially more pronounced when using LfD to train policies for safety critical systems when a goal is to be achieved while adhering to rules. We thus ask the question:

Can we encode rules as high-level task specifications to improve the online performance of the LfD policies?

For continuous dynamical systems, there often exist spatio-temporal constraints that define the rules of operation of the

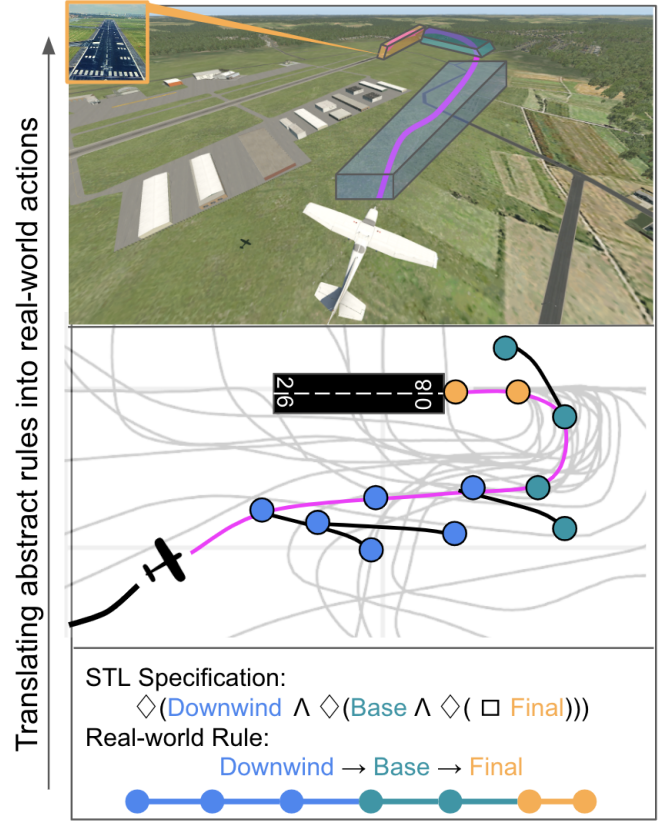


Fig. 1: Figure shows the proposed approach in a prototypical scenario to plan for an aircraft landing in a non-towered airfield. The expert/pilot demonstrations (grey) are used offline to train an LfD policy; online, we use the Signal Temporal Logic specifications in the MCTS expansion to ensure rule compliance.

system. However, traditional LfD methods do not provide a formal way to bias the system behavior for the satisfaction of rules. Moreover, the rules are often expressed in ambiguous natural language with partial specifications. Temporal logics such as Signal Temporal Logic (STL) provide a mathematically robust representation to encode such spatio-temporal constraints. They can be used to logically specify desired behavior translated from requirements expressed in natural language. STL is used to specify properties over real-valued dense time signals often generated by continuous dynamical systems [2]. Quantitative semantics associated with STL provides a real value called robustness which quantifies the degree of satisfaction or violation. This property enables the designer to encode domain-specific constraints and quanti-

^{*}equal contribution

This work was supported by Army Futures Command.

¹Jasmine J. Aloor is with the Department of Aerospace Engineering, Indian Institute of Technology, Kharagpur, WB, India (work done while at Carnegie Mellon University) jasminejerry@iitkgp.ac.in

²Authors are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. {jpatrika, parvk, hyaejino, basti}@andrew.cmu.edu

¹Codebase: <https://github.com/castacks/mcts-stl-planning>

²Video: <https://youtu.be/tiFCwc57MQs>

tatively measure their satisfaction with motion planning and control.

In order to infuse the STL specifications to improve the base LfD, we propose using a variant of Monte Carlo Tree Search (MCTS) that uses an offline pre-trained network to generate simulated roll-outs. MCTS is a powerful heuristic search algorithm often deployed for long-horizon decision-making tasks [3]. MCTS, when used with a UCT heuristic, [4] has properties like anytime convergence to the best action. This makes it well suited to be used as a long-horizon goal-directed planner within large state spaces with a time budget.

In this work, we present a novel decision-making method that uses MCTS to build a tree structure that guides the offline pre-trained LfD policies online with STL specifications. We achieve this by biasing the MCTS heuristic sampling towards branches with higher STL satisfactions. The primary insight is that while the LfD policy decides on the low-level executions in line with the expert demonstrations, STL encourages the satisfaction of high-level objectives. This hierarchical approach provides a method to encode rules while allowing the agent to choose how to satisfy the constraints in a learning-enabled framework. The STL specifications thus provide a guide rail against the low-bandwidth noise in the expert demonstrations.

In order to showcase the efficacy of the proposed algorithm, we use the prototypical case study of planning for a General Aviation (GA) aircraft operating at a non-towered airfield. GA aircraft operating under Visual Flight Rules (VFR) in non-towered airspaces are allowed to operate without a central authority while following some general rules established by the FAA. While not enforced, pilots are expected to adhere to these rules, but deviations arise due to a multitude of factors. Transponder data is available to observe and learn from this behavior, but it is often noisy as each pilot follows a variation of the rules. The proposed algorithm uses this data to train a behavior cloning policy and ensures rule compliance by augmenting the MCTS using STL specifications derived from the rules. Our contributions are as follows:

- 1) We propose a novel method to incorporate high-level rules expressed in STL specifications in online MCTS simulation that augments any base pre-trained LfD policy.
- 2) We showcase results on a challenging real-world problem that uses human demonstration data with experimental evaluations performed on a simulator and show improvements over the base policy.

The organization of the remainder of the paper is as follows: In Section II, we provide an overview of the previous approaches to solve sequential decision-making using imitation learning, MCTS, and STL. In Section III, we introduce the approach and our algorithm. In Section IV, we provide implementation details for the GA problem statement. Section V provides comparative results with established baselines. Section VI provides concluding remarks and outlines future work.

II. RELATED WORK

Guiding LfD robot behavior using specifications expressed in formal logic has been previously explored in literature. Previous methods have leveraged LTL [5], [6] or, more recently, STL [7]–[10] to encode the desired robot behavior to enable planning for autonomous systems. Most prior work focuses on offline backpropagating STL robustness along with imitation learning loss to improve the trained policy’s constraint satisfaction. These proposed offline methods that learn from either a margin based on the lower bound of STL satisfaction [11], reward functions [12], [13], vector representation [14], or risk metrics [15].

While offline learning has led to improved STL satisfaction, there are no guarantees that the resulting controller will produce satisfying trajectories [16] nor can it accommodate post hoc specification changes. To this end, methods that use constraints online in the form of Control Barrier Functions [17] and one-step state feedback [9] have been proposed, but neither uses expert demonstrations. [18] takes an alternate approach to improving the efficiency and applicability of LfD-generated policies using beam search with goal generation as a hierarchical approach. The closest to our work is the recently proposed method [16] that uses STL and expert demonstrations to synthesize a trajectory-feedback controller. The offline component uses an LSTM-based controller whose parameters are modified on-the-fly. Our proposed method has no such requirements for using a particular architecture, nor does it require updating the base policy’s parameters.

While a majority of the algorithms showcase results on either grid-based worlds [19], simple reach-avoid problems [17], or deterministic settings [8], our work showcases results using real-world, noisy expert demonstrations.

III. METHODOLOGY

This section details the problem statement and the proposed framework.

A. Preliminaries

Given a continuous-space dynamic system of the form $\dot{s} = f(s, a)$, we define a discrete-time Markov Decision Process without rewards, (MDP \setminus R). Let $\mathcal{M} = (S, A, T, \rho_0, G)$ where S is the set of states $s \in S$, $a \in A$ is the discrete set of actions or motion primitives that follow $f(\cdot)$, $T : S \times A \Rightarrow S$ is the transition function, $\rho_0 \in S$ is an initial state distribution, and G is the goal distribution.

The task is to produce a policy $\pi(\theta)$ from a start location $s_0 \in \rho_0$ to goal location $g \in G$ that leads to a trajectory $\tau = (s_0, a_0, s_1, a_1, \dots, g)$. We also assume access to expert trajectories $\mathcal{D} = \{(s_0^j, a_0^j, s_1^j, a_1^j, \dots, g^j)\}_{j=0}^D$ and high-level STL specification Φ that encodes any rules we expect the system to follow. An STL formula Φ can be built recursively from predicates using the following grammar

$$\Phi := \top \mid \mu_c \mid \neg\Phi \mid \Phi \wedge \Psi \mid \diamond_{[a,b]}\Phi \mid \square_{[a,b]}\Phi \mid \Phi_1 \mathcal{U}_{[a,b]} \Phi_2 \quad (1)$$

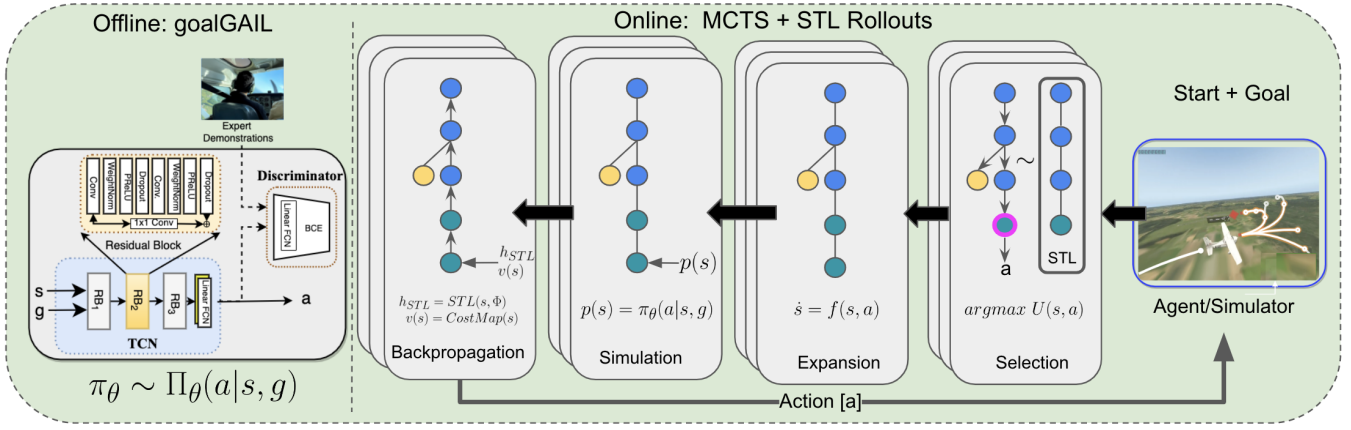


Fig. 2: Overview of the approach: Offline, we train an LfD policy using datasets, which are used Online in a Monte-Carlo Tree Search (MCTS). The online expansion uses a modified heuristic that uses robustness values from Signal Temporal Logic (STL) specification to guide the search toward higher rule conformance.

where Φ_1, Φ_2 are STL formulas, \top is the Boolean True, μ_c is a predicate of the form $\mu(s) > c$, \neg and \wedge the Boolean negation and AND operators, respectively, and $0 \leq a \leq b < \infty$ denote time intervals. The temporal operators \diamond , \square and \mathcal{U} are called “eventually”, “always”, and “until” respectively. The quantitative semantics of a formula with respect to a signal \vec{x}_t can be used to compute robustness values [20] for the specifications used in our application.

B. Framework

The overall framework is shown in Fig. 2 and Algorithm 1. We first train an LfD policy by formulating the problem as finding a distribution of future actions \hat{a}_t conditioned on the past trajectories $s_{t-t_{obs}:t}$ and the goal g where t_{obs} is the observation time horizon.

$$\hat{\pi}_\theta \sim \Pi_\theta(\hat{a}_t | s_{t-t_{obs}:t}, g) \quad (2)$$

The MCTS (see Algorithm 2) uses the policy $\hat{\pi}_\theta$ to generate simulations. Each simulation starts from the root state and iteratively selects moves that maximize the STL-modified UCT heuristic.

For each state transition, we maintain a directed edge in the tree with an action value $Q(s, a)$, prior probability $P(s, a)$, STL heuristic $H(s, a)$, and a visit count $N(s, a)$. The total heuristic value is calculated as

$$U(s, a) = Q(s, a') + \frac{c_1 P(s, a') \sqrt{N(s)}}{1 + N(s, a')} + c_2 H(s, a') \quad (3)$$

where c_1 and c_2 are the hyperparameters controlling the degree of exploration and STL heuristic’s weight. Starting with the initial state $s(0)$, at each time step, we calculate the action to take, which maximizes $U(s, a)$ (Line 12). If the next state already exists in the tree, we continue our simulation, else a new node is created in our tree, and we initialize its $P(s, \cdot) = \vec{p}_\theta(s)$ from our policy $\hat{\pi}_\theta$ (Line 8). The expected reward $v = v_\theta(s)$ can be provided by the user as a learned value function or as a cost-map (Line 9). The heuristic h_{STL} is calculated using the STL specification

(Line 14). We initialize $Q(s, a)$, $H(s, a)$ and $N(s, a)$ to 0 for all a . We then propagate the cost v and the STL heuristic h_{STL} back up the MCTS tree, updating all the $Q(s, a)$ and $H(s, a)$ values seen during the simulation, and start again from the root.

After running forward simulations of the MCTS, the $N(s, a)$ values provide a good approximation for the optimal stochastic process from each state. Hence, the action we take is randomly sampled from a distribution of actions with probability proportional to $N(s, a)$. We expand the search space until we exhaust the planning budget time $planHorizon$. After each action is selected, the MCTS tree is reinitialized from the actual trajectory followed by the agent. The planner is terminated when the goal is reached, or the maximum number of steps is reached, whichever comes first.

Algorithm 1: Plan (θ, Φ)

```

1  $s \leftarrow Sample(\rho_0)$ 
2  $g \leftarrow Sample(G)$ 
3 while  $s \neq g$  and not  $maxSteps$  do
4   while  $timeElapsed \leq planHorizon$  do
5      $N(\cdot) \leftarrow MCTS(s_{0:t}, g, \theta, \Phi, 0)$ 
6   end
7    $a \leftarrow choice_{a'}(N(s, a'))$ 
8    $s \leftarrow T(s, a)$ 
9 end

```

IV. EXPERIMENTS

In our experiments, we tackle the case of planning for a general aviation aircraft at a non-towered airport. This section provides the necessary background, problem definition, and implementation details.

A. Background

In an environment without a centralized controlling authority, such as near a non-towered airspace, the FAA establishes rectangular airport traffic patterns. These patterns facilitate the smooth flow of aircraft entering and leaving the airspace,

Algorithm 2: MCTS ($s_{0:t}, g, \theta, \Phi, h_{STL}$)

```

1 if  $s \in G$  then
2   return  $s == g, h_{STL}$ 
3 end
4 if  $s \notin Tree$  then
5    $Tree \leftarrow Tree \cup s$ 
6    $Q(s, a) \leftarrow 0$ 
7    $N(s, a) \leftarrow 0$ 
8    $P(s, a) \leftarrow \hat{\pi}_\theta(s)$ 
9    $v(s) \leftarrow CostMap(s)$ 
10  return  $v(s), h_{STL}$ 
11 else
12   $a \leftarrow \operatorname{argmax}_{a'} \left[ Q(s, a') + \frac{c_1 P(s, a') \sqrt{N(s)}}{1 + N(s, a')} + \right.$ 
13     $\left. c_2 H(s, a') \right]$ 
14   $s' \leftarrow T(s, a)$ 
15   $h_{STL} \leftarrow STL(s' + s_{0:t}, \Phi)$ 
16   $v, h_{STL} \leftarrow MCTS(s' + s_{0:t}, g, \theta, \Phi, h_{STL})$ 
17   $N(s, a) \leftarrow N(s, a) + 1$ 
18   $Q(s, a) \leftarrow \frac{N(s, a)Q(s, a) + v}{1 + N(s, a)}$ 
19   $H(s, a) \leftarrow \frac{N(s, a) * H(s, a) + h_{STL}}{1 + N(s, a)}$ 
20 return  $v, h_{STL}$ 

```

similar to roadways for automobiles. However, different pilots follow distinct paths [21], leading to variations in the expert demonstration trajectories. This variation arises due to the unique combinations of aircraft type, weather, pilots' experience, visual observations, and the non-strict regulations for following the pattern. A single fixed set of waypoints to follow would not suffice to cover all possibilities for the aircraft. With such a broad set of rules, a range of possible trajectories can be followed to land an aircraft safely.

To illustrate with an example, the high-level goal of landing on a runway through traffic patterns can usually be divided into three stages [22]. The three stages can be represented by three individual regions to be occupied by the aircraft in a particular order (See Fig. 3). The first stage is termed the downwind leg, which involves flying in the opposite direction to the intended landing runway at 1000 ft above ground level. Once the pilot is at a 45° angle from the approach end of the runway, the second stage begins, in which a perpendicular turn is made to the base leg until it reaches the runway approach's extended center line. Once the turn is complete, the third and final stage begins, which is a descending path to the point of touchdown. We provide a mechanism to encode such traffic patterns and utilize LfD to implement the sequence of actions that can land an aircraft from a given start to the runway, just like a human pilot.

B. Problem Definition

Consider a fixed-wing aircraft at time t , let $s_t = (x_t, y_t, z_t, \chi_t) \in \mathcal{R}^3 \times SO(2)$ denote the position and heading of the agent. We also define the system $\dot{s} = f(s, a)$




Fig. 3: Goal representation is in the form of a one-hot vector where each region is the respective goal element in the goal vector G . The maroon rectangle shows airport traffic patterns.

(4) as:

$$\dot{x} = v_{2D} \cos \chi \quad (4a)$$

$$\dot{y} = v_{2D} \sin \chi \quad (4b)$$

$$\dot{z} = v_h \quad (4c)$$

$$\dot{\chi} = \frac{g_e \tan \phi}{v_{2D}} \quad (4d)$$

$$v = \sqrt{v_{2D}^2 + v_h^2} \quad (4e)$$

where v is the aircraft's inertial speed, v_{2D} is the speed in the 2-D plane, ϕ is the roll-angle, and v_h is vertical speed. Finally, g_e is the acceleration due to gravity. We ignore the effects of wind.

The action space $A = \{(v^j, v_h^j, \phi^j)\}_{j=0}^A$ is a fixed library of 30 motion primitives that discretize each of the control inputs. We use the inertial velocities (v) 70 and 90 knots, the vertical velocities (v_h) +500 ft/min and -500 ft/min, and the bank angle (ϕ) discretization such that χ changes by 45° and 90° heading over the chosen time-horizon of 20 sec. The goal distribution G is a one-hot vector representation of the final goal of a particular agent as the eight cardinal directions along with two runway ends, as shown in Fig. 3. The set G is represented as $G = \{N, NE, E, SE, S, SW, W, R08, R26\}$ with each element representing the final region the airplane is desired to reach as shown in Fig. 3. For simplicity, we also set the start states to one of these regions, i.e. $\rho_0 = G$.

C. Implementation Details

The implementation details are split between online and offline components.

1) *Dataset*: We use the *TrajAir*³ dataset, which provides recorded trajectories of aircraft operating at the Pittsburgh-Butler Regional Airport (ICAO:KBTP) [21]. The dataset contains 111 days of transponder data, processed to obtain the local x, y , and z coordinates of aircraft at every second. This

³<http://theairlab.org/trajair/>

is used to extract expert demonstrations from pilots as they navigate the un-towered airspace. The dataset trajectories are smoothed using a B-spline (basis-spline) approximate representation of order 2. The trajectories are not biased toward satisfying the STL rules as there are variations in landing patterns exhibited by the dataset.

2) *Offline LfD Policy Details:* The LfD policy takes as input the past trajectories of the agent to predict its possible action distribution. While the method can use any LfD policy, we use a goal-conditioned generative adversarial imitation learning (GoalGAIL) method [23]. The GoalGAIL is modified to use Temporal Convolutional Layers (TCNs) to process the sequential trajectory data. TCN layers encode a trajectory’s spatio-temporal information into a latent vector without losing the underlying data’s temporal (causal) relations [24]. We use TCNs as an alternative to using LSTMs [25] for encoding the trajectories.

We break the trajectories in a scene into sequences of length $t_{obs} + t_{pred}$ where $t_{obs} = 11sec$ and $t_{pred} = 20sec$. In a given scene, the raw trajectory in absolute coordinates of the agent is encoded using the TCN layers as h_{obs} . The agent’s goal $g \sim G$ is encoded through an MLP layer, φ_1 , and is concatenated with the encoded trajectory vector. Equations 5 show the encoding sequence.

$$h_{obs} = TCN_{obs}(s_{1:t_{obs}}) \quad (5a)$$

$$h_g = \varphi_1(g) \quad (5b)$$

$$h_{enc} = h_{obs} \oplus h_g \quad (5c)$$

$$\hat{s}_{t_{obs}:t_{obs}+t_{pred}} = \varphi_2(h_{enc}) \quad (5d)$$

The \mathcal{L}_{act} measures how close the predicted trajectory is to the expert trajectory using a mean squared error (MSE) loss.

$$\mathcal{L}_{act} = MSE(s_{t_{obs}:t_{obs}+t_{pred}}, \hat{s}_{t_{obs}:t_{obs}+t_{pred}}) \quad (6)$$

A discriminator D_ψ is trained to distinguish expert transitions $(s, a, g) \sim \tau_{expert}, E(s_{1:t_{obs}})$, from policy transitions $(s, a, g) \sim \tau_{policy}, \hat{s}_{t_{obs}:t_{obs}+t_{pred}}$. The discriminator is trained to minimize,

$$\mathcal{L}_{goalGAIL}(D_\psi) = \mathbb{E}_{(s,a,g) \sim \text{policy}} [\log D_\psi(a, s, g)] + \mathbb{E}_{(s,a,g) \sim \text{expert}} [\log(1 - D_\psi(a, s, g))] \quad (7)$$

The combination of these two loss functions is used to train the model.

$$\mathcal{L}_{total} = \mathcal{L}_{act} + \mathcal{L}_{goalGAIL} \quad (8)$$

In order to convert \hat{s} to \hat{a} , we match the generated trajectories from the control inputs in the library A using a weighted L2 Euclidean error distance over (x, y, z) points on the trajectory. For training, we use the AdamW optimizer with a learning rate of $3e - 3$.

3) *Signal Temporal Logic Specifications:* We evaluate the performance of our agent based on reaching the goal while adhering to the airport traffic pattern. The goal objective, as well as traffic pattern compliance, are both encoded using STL specifications. We use the three stages for the landing pattern as defined in (IV-A). Φ_1 , Φ_2 , and Φ_3 represent the

STL formulas encoding occupancy of regions corresponding to the downwind, base, and final stages, respectively. The landing STL specification becomes:

$$\Phi_L = \diamond(\Phi_1 \wedge (\diamond(\Phi_2 \wedge (\diamond \square \Phi_3)))) \quad (9)$$

$\diamond(\Phi)$ can be interpreted as ‘‘Eventually’’ being in a region represented by Φ . The nested \diamond operators encode a sequential visit of regions represented by Φ_1 , Φ_2 , and Φ_3 . Similarly, the takeoff STL specification is defined based on the goal regions reached by the aircraft. By defining reaching a goal region $g \sim G$ by an STL formula Φ_4 , we get the takeoff specification:

$$\Phi_T = \diamond(\Phi_4) \quad (10)$$

The robustness values of the STL specifications are evaluated on the state trajectory traces generated by the search tree. The first element of the traces is the tree root node, and the last element of the trace is the node whose value is currently computed.

4) *Online Monte Carlo Tree Search:* The MCTS is implemented as a recursive function where each iteration ends with a new leaf that corresponds to an action in the trajectory library. The implementation uses a normalized costmap $v(s)$ that is built by counting the frequencies of the aircraft in the TrajAir dataset at particular states s .

V. EVALUATIONS

Evaluation of the proposed approach is done using a custom simulator that follows the dynamics defined in Eq. 4. The network implementations are done on PyTorch. For calculating STL robustness values, we formulate our specifications using the `rtamt` [26] package, an online monitoring library. To showcase real-time online evaluations, simulations are performed on an Intel NUC computer with Intel® Core™ i7-8559U CPU @ 2.70GHz \times 8. The complete implementation details and parameter details are in the open-sourced code-base⁴ and the associated Readme.

A. Qualitative results

Figure 4 shows an example scenario where the aircraft starts from the South-West and is tasked with landing at R26 while following the standard FAA traffic pattern. The rules of the traffic pattern are encoded as STL specifications. The white-marked line shows the aircraft’s position at each step. At every step, the MCTS replans, and the resulting tree is shown in magenta. The STL sub-specifications are shown as rectangles. As can be seen, the aircraft manages to reach the runway while satisfying the specifications. The size of the search tree is a function of available *planHorizon*.

B. Comparative results

In order to perform quantitative studies, we compare the performance of the proposed algorithm with vanilla LfD policies. We uniformly sample 100 start-goal pairs randomly from ρ_0, G . G is truncated to N, S, E, W, R where R represents both R08 and R26 to condense the results. We then

⁴Codebase: <https://github.com/castacks/mcts-stl-planning>

Algorithm	Takeoff Specification Φ_T				Landing Specification Φ_L				Total
	N	S	E	W	N	S	E	W	
BC + MCTS	0.2 / 0.2	0.0 / 0.1	0.3 / 0.1	0.3 / 0.1	0.0 / 0.0	0.1 / 0.4	0.1 / 0.4	0.3 / 0.4	0.1 / 0.2
GoalGAIL + MCTS	0.0 / 0.3	0.1 / 0.4	0.0 / 0.3	0.0 / 0.4	0.3 / 0.6	0.1 / 0.2	0.3 / 0.7	0.1 / 0.5	0.1 / 0.4
BC + MCTS + STL	1.0 / 0.9	1.0 / 0.9	0.6 / 0.3	0.7 / 0.4	0.2 / 0.5	0.2 / 0.5	0.2 / 0.5	0.2 / 0.5	0.5 / 0.6
GoalGAIL + MCTS+STL	1.0 / 0.9	1.0 / 0.9	0.8 / 0.7	1.0 / 0.9	0.7 / 0.9	0.8 / 0.9	0.3 / 0.8	0.5 / 0.9	0.7 / 0.8

TABLE I: Table shows the quantitative results with randomly sampled start and goal states for two LfD policies with ablation studies with the STL heuristic. Results show the Success Rate \uparrow / STL Score \uparrow for two vanilla LfD policies and their ablations with the STL heuristic. Results show both Landing $X \Rightarrow R$ and Takeoff $R \Rightarrow X$ scenarios for each of the cardinal directions X .




Fig. 4: Figure shows a qualitative example from one of the cases where the aircraft starts from the South-West and needs to land at one of the runways (R26). The specifications Φ_1 , Φ_2 and Φ_3 are shown as rectangles. White marked lines show the aircraft trajectory, and the magenta shows the MCTS tree. The runway threshold for R08 (+x-axis) is at the center.

provide these to Algorithm 1, which plans for the agent. Based on the selected start and goal pair, a relevant STL specification is chosen. Each episode ends when the agent reaches the goal or if the maximum steps are exceeded. In addition to the goalGAIL policy, we also train a pure Behavior Cloning (BC) policy. The BC policy uses a TCN to encode the history and outputs an action without a goal vector or a discriminator. Comparisons were carried out with both these policies integrated into the MCTS framework with ablation on the STL heuristic to show the impact of the STL specification.

We define our evaluation metrics as follows:

- Success Rate: Fraction of episodes that were successful in reaching their goal locations.
- STL Score: Average of the normalized fraction of the STL robustness value satisfied over all the episodes. A higher value indicates better satisfaction.

Table I shows the quantitative results. Our proposed algorithm performs significantly better than the baselines for all start-goal pairs. We get an almost perfect success rate in the aircraft takeoff scenarios. The aircraft landing cases are more challenging due to following the specific landing patterns when incoming from different sides of the runway, which is reflected in the success rates shown. Additionally, we observe the baseline BC performs similarly to GAIL on the success metric but not on the STL robustness. This indicates that while BC policies are comparable in reaching the goals, the

transient performance of GoalGAIL is better. Incorporating STL improves the robustness values for both LfD policies.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we present a novel method that improves the online robustness of offline pre-trained LfD policies using MCTS to fuse STL specifications. To the best of the authors' knowledge, this is the first method that combines high-level STL specifications with low-level LfD policies through MCTS. Our experimental evaluation targets the real-world problem of autonomous aircraft planning and exhibits the feasibility of our techniques for similar challenging decision-making problems. We show our method outperforms vanilla LfD methods on a number of successful missions for multiple complex objectives using real-world data.

While the results are promising, future work involves producing theoretical guarantees on the satisfaction of the STL specifications. The current work biases the search tree towards satisfying STL constraints but provides no guarantees. Another line of work is to validate the algorithm on a high-fidelity simulator and real-world flight tests.

VII. ACKNOWLEDGMENT

The authors would like to thank Joao Dantas for his help with the costmaps.

REFERENCES

- [1] J. A. Bagnell, "An invitation to imitation," Carnegie-Mellon Univ Pittsburgh Pa Robotics Inst, Tech. Rep., 2015.
- [2] 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.
- [3] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [4] L. Kocsis and C. Szepesvári, "Bandit Based Monte-Carlo Planning," in *Machine Learning: ECML 2006*, J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 282–293.
- [5] C. Innes and S. Ramamoorthy, "Elaborating on learned demonstrations with temporal logic specifications," *arXiv preprint arXiv:2002.00784*, 2020.
- [6] H. Wang, H. He, W. Shang, and Z. Kan, "Temporal logic guided motion primitives for complex manipulation tasks with user preferences," *arXiv preprint arXiv:2202.04375*, 2022.
- [7] 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.
- [8] P. Kapoor, A. Balakrishnan, and J. V. Deshmukh, "Model-based reinforcement learning from signal temporal logic specifications," *arXiv preprint arXiv:2011.04950*, 2020.
- [9] S. Yaghoubi and G. Fainekos, "Worst-case satisfaction of stl specifications using feedforward neural network controllers: a lagrange multipliers approach," in *2020 Information Theory and Applications Workshop (ITA)*. IEEE, 2020, pp. 1–20.
- [10] V. Raman, A. Donzé, M. Maasoumy, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, "Model predictive control with signal temporal logic specifications," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 81–87.
- [11] K. Cho and S. Oh, "Learning-based model predictive control under signal temporal logic specifications," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7322–7329.
- [12] A. G. Puranic, J. V. Deshmukh, and S. Nikolaidis, "Learning from demonstrations using signal temporal logic," *arXiv preprint arXiv:2102.07730*, 2021.
- [13] —, "Learning from demonstrations using signal temporal logic in stochastic and continuous domains," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6250–6257, 2021.
- [14] W. Hashimoto, K. Hashimoto, and S. Takai, "Stl2vec: Signal temporal logic embeddings for control synthesis with recurrent neural networks," *arXiv preprint arXiv:2109.04636*, 2021.
- [15] X. Li, J. DeCastro, C. I. Vasile, S. Karaman, and D. Rus, "Learning a risk-aware trajectory planner from demonstrations using logic monitor," in *Conference on Robot Learning*. PMLR, 2022, pp. 1326–1335.
- [16] K. Leung and M. Pavone, "Semi-supervised trajectory-feedback controller synthesis for signal temporal logic specifications," *arXiv preprint arXiv:2202.01997*, 2022.
- [17] W. Liu, N. Mehdipour, and C. Belta, "Recurrent neural network controllers for signal temporal logic specifications subject to safety constraints," *IEEE Control Systems Letters*, vol. 6, pp. 91–96, 2021.
- [18] M. Igl, D. Kim, A. Kuefler, P. Mougín, P. Shah, K. Shiarlis, D. Anguelov, M. Palatucci, B. White, and S. Whiteson, "Symphony: Learning realistic and diverse agents for autonomous driving simulation," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 2445–2451.
- [19] K. C. Kalagarla, R. Jain, and P. Nuzzo, "Model-free reinforcement learning for optimal control of markovdecision processes under signal temporal logic specifications," *arXiv preprint arXiv:2109.13377*, 2021.
- [20] A. Donzé, T. Ferrere, and O. Maler, "Efficient robust monitoring for stl," in *International Conference on Computer Aided Verification*. Springer, 2013, pp. 264–279.
- [21] J. Patrikar, B. Moon, S. Ghosh, J. Oh, and S. Scherer, "Trajair: A general aviation trajectory dataset," Jun 2021. [Online]. Available: <https://doi.org/10.1184/R1/14866251.v1>
- [22] F. A. A. (FAA), *Airplane Flying Handbook: FAA-H-8083-3C*. Aviation Supplies & Academics, Incorporated, 2021.
- [23] Y. Ding, C. Florensa, M. Phielipp, and P. Abbeel, "Goal-conditioned imitation learning," *arXiv preprint arXiv:1906.05838*, 2019.
- [24] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [25] Z. Zhao, W. Zeng, Z. Quan, M. Chen, and Z. Yang, "Aircraft trajectory prediction using deep long short-term memory networks," in *CICTP 2019*, 2019, pp. 124–135.
- [26] D. Ničković and T. Yamaguchi, "Rtam: Online robustness monitors from stl," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2020, pp. 564–571.