

Reactive Temporal Logic-based Planning and Control for Interactive Robotic Tasks

Farhad Nawaz^{1*}, Shaoting Peng¹, Lars Lindemann², Nadia Figueroa¹ and Nikolai Matni¹

Abstract—Robots interacting with humans must be safe, reactive and adapt online to unforeseen environmental and task changes. Achieving these requirements concurrently is a challenge as interactive planners lack formal safety guarantees, while safe motion planners lack flexibility to adapt. To tackle this, we propose a modular control architecture that generates both safe and reactive motion plans for human-robot interaction by integrating temporal logic-based discrete task level plans with continuous Dynamical System (DS)-based motion plans. We formulate a reactive temporal logic formula that enables users to define task specifications through structured language, and propose a planning algorithm at the task level that generates a sequence of desired robot behaviors while being adaptive to environmental changes. At the motion level, we incorporate control Lyapunov functions and control barrier functions to compute stable and safe continuous motion plans for two types of robot behaviors: (i) complex, possibly periodic motions given by autonomous DS and (ii) time-critical tasks specified by Signal Temporal Logic (STL). Our methodology is demonstrated on the Franka robot arm performing wiping tasks on a whiteboard and a mannequin that is compliant to human interactions and adaptive to environmental changes.

I. INTRODUCTION

There is a growing presence of robot assistants in households, workplaces, and industrial plants. In such applications, robots should adapt to environmental changes and safely interact with humans while performing a diverse range of tasks that conform to sequential and time-critical behaviors. Existing end-to-end frameworks [1]–[3] have demonstrated robots performing diverse tasks, while model-based methods [4], [5] focus on predefined motions that are stable and safe. However, they do not address the following challenges jointly: (i) formal guarantees on safety and satisfying task specifications, and (ii) reactivity to online environmental events and unforeseen disturbances – which are critical for deploying robots in human-centric environments.

Consider the scenario in Fig. 1, where the robot is wiping a whiteboard. We want the robot to switch between periodic wiping motions on the left and right side of the board as commanded by the user. If the camera detects a blue stain, then the robot should wipe that off quickly. If there is an external perturbation where the robot drops the eraser, then it should pick it up and continue the wiping motion. This involves switching between different motions (left and right wiping), reacting to environmental events (stain), time-critical tasks (wiping stain), and human interactions (eraser dropped).

¹Farhad Nawaz, Shaoting Peng, Nadia Figueroa and Nikolai Matni are with the GRASP Lab, University of Pennsylvania, PA 19104, USA. ²Lars Lindemann is with the University of Southern California, CA 90007, USA.

*Corresponding author: farhadn@seas.upenn.edu

The key challenge for robots in such interactive scenarios is maintaining safety and compliance while reacting to online environmental events and switching between tasks. To alleviate this, we develop a computationally efficient online planning strategy that generates safe and reactive motions for human-robot interaction, by integrating a 200 Hz discrete planner based on temporal logic and a 1 KHz continuous Dynamical System (DS)-based planner.

Task specifications: At the task level, one way to describe desired robot behaviors is via temporal logic specifications [6], previously used in a variety of tasks such as cooking [7], mobile manipulation [8], search and rescue missions [9]. Such specifications allow robots to understand and interpret commands involving temporal relationships using structured language. For example, a high-level specification for the scenario in Fig. 1 can be given as

$$\begin{aligned} &(\textit{left} \rightarrow \textit{wipe left}) \text{ and } (\textit{right} \rightarrow \textit{wipe right}) \text{ and} \\ &(\textit{blue stain} \rightarrow \textit{wipe stain in 5 seconds}) \text{ and} \\ &(\textit{no eraser} \rightarrow \textit{pick up eraser}). \end{aligned} \quad (1)$$

Specifications like (1) are typically satisfied by integrating high-level discrete actions and low-level continuous motions via formal methods for control [10]–[12]. Such methods compute a control signal for predefined scenarios but lack reactivity to online events and disturbances. Here, we focus on reactive planning, i.e., the robot should adapt its behaviors depending on the changing environmental observations online. Prior work using Signal Temporal Logic (STL) [13] relies on computationally intensive mixed-integer linear programs [14], [15] for online adaptation, while methods based on Linear Temporal Logic (LTL) [9], [16]–[18] require computationally taxing planners on discrete transition systems. To tackle such computational burdens and inflexibility, we instead translate STL specifications into constraints on the continuous motion level enforced via efficient Quadratic Programs (QPs) for real-time planning as in [19].

Discrete task plan: We adapt ideas from [20], and formally define *controllable* and *uncontrollable* propositions in Section III-A to integrate high level reactive temporal logic specifications with continuous low level motions. Intuitively, *controllable* propositions, highlighted in *blue* in (1), represent robot behaviors that can always be controlled by appropriate choice of motion planners. *Uncontrollable* propositions, highlighted in *red* in (1), are discrete environmental events that cannot always be controlled by the robot. We define a reactive temporal logic formula in Section III-B that builds upon *controllable* and *uncontrollable* propositions. We then construct a Büchi automaton [21] using existing tools [22]

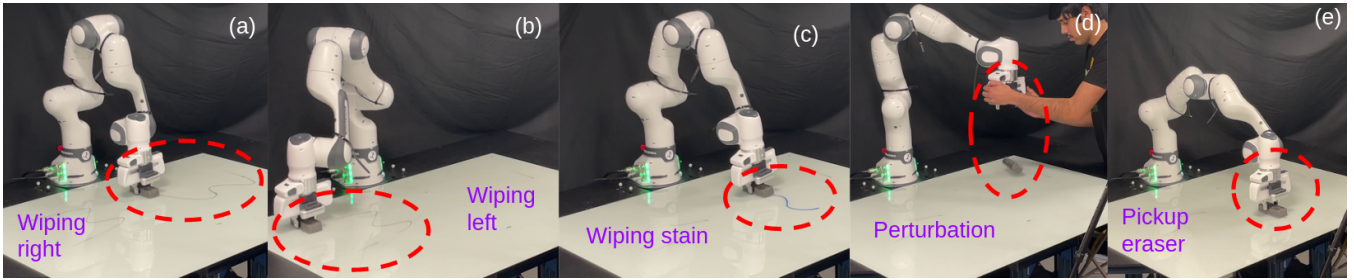


Fig. 1: Periodic wiping motion on the (a) right side and (b) left side of the board, (c) wiping the blue stain detected by a camera, (d) eraser dropped because of human perturbation and (e) robot picks up the dropped eraser.

and employ it within a discrete task planner that decides the sequence of desired behaviors the robot should follow to always satisfy the task specification, despite environmental changes such as the blue stain and perturbation in Fig 1.

Continuous motion plan: We integrate our discrete task planner with a continuous DS-based motion planner that executes each desired robot behavior. While our method is agnostic to the underlying DS motion policy, we focus on two types that encapsulate diverse tasks: (i) nominal motions given by an autonomous DS, and (ii) time-critical tasks specified by STL. Nominal motions could either be provided as a library of autonomous DS, or learned from demonstrations [7], [23]. We utilize our prior work [24] that uses Neural ODEs and tools from Lyapunov theory to learn nominal (possibly periodic) motions from only 3-4 demonstrations while also guaranteeing safety and stability. We model time-critical tasks as STL specifications and time-varying Control Barrier Functions (CBFs) [19]. Since our prior work [24] also uses CBFs for safety and Control Lyapunov Functions (CLFs) for stability, we endow robots the capability to follow both time-critical behaviors and complex nominal motions in a single framework by solving an efficient QP online. Thus enabling the robot to react to environmental events and follow the desired behavior, guaranteeing task satisfaction even in the face of disturbances.

Contributions: First, we adapt the notion of controllable propositions from [20] to model more general continuous time robot behaviors such as nominal DS-based motions and time-critical STL tasks. We use such propositions as the basic building blocks of a reactive temporal logic formula to represent task specifications in structured language. Second, we propose a discrete task planning algorithm running online at 200 Hz that decides the sequence of robot behaviors to execute to satisfy the task specification despite environmental changes online. Third, we integrate our task planner with a continuous DS-based motion planner that uses CLFs and CBFs to follow nominal motion plans and satisfy time-critical STL tasks in the presence of unforeseen disturbances, by solving a QP online at 1 KHz. Finally, we validate our approach on the Franka robot arm that performs reactive wiping tasks on a whiteboard and human mannequin using environmental observations from a RealSense camera, while also being compliant to human interactions.

II. PRELIMINARIES

We introduce the background to define our reactive temporal logic specification, and tools from control theory to

satisfy time-critical tasks and guarantee disturbance rejection. Let $x \in \mathcal{X} \subset \mathbb{R}^d$ and $u \in \mathcal{U} \subset \mathbb{R}^m$ be the state and control input for the nonlinear control affine dynamical system

$$\dot{x} = f(x) + g(x)u, \quad (2)$$

where, $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m}$.

A. Linear Temporal Logic and Büchi Automaton

We use LTL to specify in structured language the desired reaction of the robot to discrete environmental events. Consider a finite set of atomic propositions AP , where each $p \in AP$ can either be true or false. An LTL formula φ is constructed using the following syntax.

$$\varphi ::= \text{true} \mid p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 U \varphi_2 \quad (3)$$

The symbol $::=$ means that φ in (3) is assigned to be one of the expressions from the right hand side separated by vertical bars $|$. The different symbols mean the following: \neg is negation, \wedge is conjunction and U is the until operator that expresses that φ_1 should be true until φ_2 becomes true, where φ_1, φ_2 are well-defined LTL formulas.

The set of discrete environmental observations and robot behaviors that dictate the LTL formula φ is given by AP . The specification φ is a combination of Boolean and temporal operations between the elements of AP . For example, consider a task where the robot should add oil to a pan once the pan is hot. The set $AP = \{hot, oil\}$ which corresponds to whether the pan is hot ($hot = 1$) and oil is poured on the pan ($oil = 1$). The LTL formula is $\varphi := G(hot \Rightarrow Foil)$, where G denote *always* and F denote *eventually* [11].

An input signal is defined as $\sigma : \mathbb{Z}^+ \rightarrow \{0, 1\}^{|AP|}$, which is a discrete time signal that represents the boolean values of the propositions AP . We are interested in signals that satisfy a task specification φ . The symbol $(\sigma, k) \models \varphi$ indicates that the signal σ satisfies a formula φ at time step k . We refer the reader to [19] for details on LTL semantics.

Büchi automaton: A Büchi automaton [12] is a tool that gives a formal way of deciding whether signals satisfy LTL specifications. Given φ , we can construct a nondeterministic Büchi automaton $\mathcal{B}_\varphi = (\mathcal{S}, s_0, AP, \Delta, \mathcal{F})$, where \mathcal{S} is a finite set of states, $s_0 \in \mathcal{S}$ is the initial state, AP is the finite set of atomic propositions, $\Delta \subseteq \mathcal{S} \times \{0, 1\}^{|AP|} \rightarrow \mathcal{S}$ is a transition relation such that $(s', d, s'') \in \Delta$ indicates a transition from s' to s'' when the input label is $d \in \{0, 1\}^{|AP|}$, and $\mathcal{F} \subseteq \mathcal{S}$ is a set of accepting states. Given an input signal $\sigma : \mathbb{Z}^+ \rightarrow \{0, 1\}^{|AP|}$, a run of the Büchi automaton \mathcal{B}_φ is an infinite state sequence $s := \{s_0, s_1, s_2, \dots\}$ that satisfies the

transition relation $(s_i, \sigma(i), s_{i+1}) \in \Delta$ for each $i \in \mathbb{Z}^+$. An accepting run on \mathcal{B}_φ is a sequence s for which $\text{inf}(s) \cap \mathcal{F} \neq \emptyset$, where $\text{inf}(s)$ is the set of states that appear infinitely often in s . If the sequence s is an accepting run on \mathcal{B}_φ , then the input signal σ satisfies φ [12].

B. Signal Temporal Logic and Control Barrier Functions

We use STL to specify time-critical tasks, time-varying CBFs to satisfy STL tasks and CLFs to enforce stability.

Instead of atomic propositions in LTL, predicates μ describe truth values in STL using a predicate function $h(x) : \mu = \text{true} \Leftrightarrow h(x) \geq 0$ and $\mu = \text{false}$ otherwise, where $x \in \mathbb{R}^d$ is the robot state. The syntax and semantics of STL closely resemble those of LTL, but can encode explicit quantitative time constraints. We direct the reader to [19] for details on STL syntax and semantics.

Control Barrier Functions: We use time-varying CBFs [19] to satisfy STL specifications. Satisfying a time-critical task specification is defined as the forward invariance of time-varying safety set $\mathcal{C}(t) \subseteq \mathcal{X} \forall t \geq 0$ for the system (2). The set $\mathcal{C}(t)$ is defined as the super-level set of a continuously differentiable function $B : \mathbb{R}^d \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$: $\mathcal{C}(t) = \{x \in \mathcal{X} : B(x, t) \geq 0\}$. Our objective is to find a control input u such that the states x that evolve according to the dynamics (2) always stay inside the set $\mathcal{C}(t)$. Such an objective is formalized using *forward invariance* of the set $\mathcal{C}(t)$. The safe set $\mathcal{C}(t)$ is forward invariant for a given control law u if for every initial point $x(0) \in \mathcal{C}(0)$, the future states $x(t) \in \mathcal{C}(t)$ for all $t \geq 0$. A continuously differentiable function $B : \mathbb{R}^d \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a *time-varying CBF* for (2) if there exists an extended class \mathcal{K}_∞ function $\gamma(\cdot)$ such that for all $(x, t) \in \mathcal{X} \times \mathbb{R}_{\geq 0}$,

$$\sup_{u \in \mathcal{U}} \frac{\partial B(x, t)}{\partial x}^\top (f(x) + g(x)u) + \frac{\partial B(x, t)}{\partial t} \geq -\gamma(B(x, t)) \quad (4)$$

The set of all control inputs that satisfy the condition in (4) for each $(x, t) \in \mathcal{X} \times \mathbb{R}_{\geq 0}$ is $K(x, t) := \{u \in \mathcal{U} : \frac{\partial B(x, t)}{\partial x}^\top (f(x) + g(x)u) + \frac{\partial B(x, t)}{\partial t} \geq -\gamma(B(x, t))\}$.

The formal result on forward invariance using time-varying CBFs follows from [19].

Theorem 1: Let $u(x, t) \in K(x, t)$ be a feedback control law that is Lipschitz in x and piecewise continuous in t . Then, the set $\mathcal{C}(t)$ is forward invariant for the control law $u(x, t)$ if $B(x, t)$ is a time-varying control barrier function.

STL specifications can be enforced via time-varying CBFs following the construction from [19]. For example, consider the STL formula $\phi ::= F_{[0,5]} \|x - x^*\| \leq \epsilon$ which specifies that eventually (F), the robot x should be ϵ -close to a target point x^* within the time interval $[0, 5]$. Similar to [19], we can define a CBF $B(x, t) = \gamma(t) - \|x(t) - x^*\|$, where $\gamma(t)$ defines the rate at which the robot reaches x^* such that $\gamma(5) = \epsilon, \gamma(0) \geq \|x(0) - x^*\|$.

We use time-independent CBFs and CLFs to guarantee safety and stability of nominal motion plans given by autonomous DS models, which is detailed in prior work [24], [25]. The difference between a time-independent CBF defined in [25] and the time-varying CBF is the presence of

the term $\frac{\partial B(x, t)}{\partial t}$ in (4), which makes constraint satisfaction more difficult if $\frac{\partial B(x, t)}{\partial t} < 0$. A *Control Lyapunov Function (CLF)* $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is a special case of a time-independent CBF $B : \mathbb{R}^d \rightarrow \mathbb{R}$ with $B(\cdot) = -V(\cdot)$ and the time-independent safe set $\mathcal{C} = \{0\}$, where Theorem 1 guarantees asymptotic stability: see [25] for more details.

III. PROBLEM FORMULATION

We first define a reactive temporal logic formula to describe the task specification using structured language, and formally write the problem statement.

A. Controllable Propositions

We introduce a general notion of controllable propositions that is adapted from [20] to model discrete events in our task specification that depend only on the state signal $x(t)$. Prior work [20] model only STL tasks, but we also focus on autonomous DS-based motions and describe how observations from the continuous time state signal x translate to a discrete time signal σ of LTL in Section II-A.

We define a generalized version of the predicates μ described in II-B as controllable propositions.

Definition 1: The truth value of a controllable proposition $p_c(t)$ at time t is defined below using a tuple $(\mathcal{L}, \mathcal{M}, \mathcal{Q})$, where $\mathcal{L} : \mathcal{P} \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{Q}$ is an operator from the space of state signals $\mathcal{P} = \{x \mid x : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^d\}$ and time $\mathbb{R}_{\geq 0}$ to a space of signals $\mathcal{Q} = \{y \mid y : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^q\}$ in cartesian space \mathbb{R}^q , $\mathcal{M} : \mathcal{P} \times \mathbb{R}_{\geq 0} \rightarrow 2^{\mathcal{Q}}$, where $2^{\mathcal{Q}}$ is the power set of \mathcal{Q} .

$$p_c(t) = \begin{cases} \text{true} & \Leftrightarrow \mathcal{L}(x, t) \in \mathcal{M}(x, t) \\ \text{false} & \Leftrightarrow \mathcal{L}(x, t) \notin \mathcal{M}(x, t) \end{cases} \quad (5)$$

Intuitively, the operator \mathcal{L} describes the current behavior of the robot relevant to $p_c(t)$ at time t , \mathcal{M} describes the desired behavior relevant to $p_c(t)$, and \mathcal{Q} is the space of signals that is applicable for the controllable proposition p_c . For the example in Fig. 1, we represent all motions in 3D ($d = 3$) and wiping motions as autonomous DS $\dot{x}(t) = f(x(t))$ so that $q = d$, $\mathcal{L}(x, t) = \dot{x}$ for all $t \geq 0$ and $\mathcal{M}(x, t) = \{y \mid y : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^3 \text{ s.t. } y(t) = f(x(t))\}$.

We can model STL tasks using controllable propositions. First, for the predicates μ defined in Section II-B, $q = 1$, $\mathcal{L}(x, t) = h \circ x$ for all $t \geq 0$ and $\mathcal{M}(x, t) = \{y \mid y : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R} \text{ s.t. } y(t) \geq 0\}$ for all $x \in \mathcal{P}$. Then, an STL specification $F_{[0,5]} \|x - x^*\| \leq \epsilon$ can be modeled as $\mathcal{L}(x, t) = h \circ x$ for all $t \geq 0$, where $h(x(t)) = -\|x(t) - x^*\| + \epsilon$, $q = 1$ and $\mathcal{M}(x, t) = \{y \mid \exists t' \in [t, t + 5] \text{ s.t. } y(t') \geq 0\} \forall x \in \mathcal{P}$.

We denote the propositions p_c in Definition 1 as *controllable* because their truth value entirely depends on the state signal x which can be controlled by u in (2), but we do not state any formal equivalence to controllability [26]. While the notion of controllable propositions is in general abstract, we focus on nominal DS-based motions and STL tasks.

Continuous to discrete time: We consider a finite set of controllable propositions \mathcal{A}_c , where each $p_c(t) \in \mathcal{A}_c$ can either be `true` or `false` as defined by (5). An LTL formula can be defined with appropriate syntax and semantics as

given in Section II-A with $AP = \mathcal{A}_c$. However, the signal σ in Section II-A that dictate the boolean values of AP is a discrete time signal, but controllable propositions depend on the continuous time state signal x . In practice, we fully observe the state $x(t)$ at each continuous time $t \geq 0$, but update the boolean values of the controllable propositions only at discrete times t_k , where $k \in \mathbb{Z}^+$. We connect the different time scales by denoting $\{t_0, t_1, \dots\}$ to be a possibly irregular sampling of the continuous time interval $[0, \infty)$ with $t_0 = 0$, where each t_k corresponds to the discrete time step k . We define a *controllable signal* $\sigma_c : \mathbb{Z}^+ \rightarrow \{0, 1\}^{|\mathcal{A}_c|}$ such that for all $k \in \mathbb{Z}^+$, $\sigma_c(k)$ depends on the state signal x at time t_k . Now we can use a Büchi automaton as described in II-A for any signal $\sigma_c : \mathbb{Z}^+ \rightarrow \{0, 1\}^{|\mathcal{A}_c|}$ to verify satisfiability of LTL formulae over \mathcal{A}_c .

B. Reactive temporal logic

We define the syntax of a *Reactive Temporal Logic* (RTL) formula that the robot should satisfy, using LTL formulae over controllable and uncontrollable propositions. Such a formula describes how the robot should react to environmental events by following different desired behaviors. In addition to \mathcal{A}_c described in III-A, we define *uncontrollable propositions* \mathcal{A}_u that represent discrete environmental observations which cannot always be controlled by the robot, such as *left*, *right*, *blue stain*, *no eraser* described in Section I. The truth values of \mathcal{A}_u are determined by both the state signal x and an observed *uncontrollable signal* $\sigma_u : \mathbb{Z}^+ \rightarrow \{0, 1\}^{|\mathcal{A}_u|}$. The syntax of RTL formulas Ψ is

$$\Psi ::= \psi \mid G(\varphi \Rightarrow \Psi) \mid G(\psi \Rightarrow \Psi) \mid \Psi_1 \wedge \Psi_2 \quad (6)$$

where, ψ is an LTL formula defined in (3) with $AP = \mathcal{A}_c$, φ is an LTL formula with $AP = \mathcal{A}_u$, and G indicates the always operator [11]. Our work always focus on reacting to environmental observations represented by \mathcal{A}_u by following a desired behavior dictated by \mathcal{A}_c . The semantics of Ψ follow that of LTL in Section II-A with $AP = \mathcal{A}_c \cup \mathcal{A}_u$ and $\mathcal{A}_c \cap \mathcal{A}_u = \emptyset$. We can construct Büchi Automaton \mathcal{B}_Ψ to verify if an input signal $\sigma : \mathbb{Z}^+ \rightarrow \{0, 1\}^{|\mathcal{A}^P|}$ satisfies Ψ . Similar to Section II-A, the symbol $([\sigma_c, \sigma_u], k) \models \Psi$ indicates that the controllable signal σ_c and uncontrollable signal σ_u satisfies a formula Ψ at time step $k \in \mathbb{Z}^+$. For example, consider a simple reactive cooking task where the robot should keep stirring a pan if it is hot. Otherwise, the robot should increase the temperature of the stove by pressing a button. The RTL specification would be $\Psi := G(\text{hot} \Rightarrow \text{stir}) \wedge G(\text{-hot} \Rightarrow F(\text{press button}))$, where $\mathcal{A}_u = \{\text{hot}\}$, $\mathcal{A}_c = \{\text{stir}, \text{press button}\}$.

C. Task and Motion Plan

Given a set of nominal autonomous DS and an RTL specification Ψ , we aim to develop a reactive planning strategy to satisfy Ψ while being adaptive to environmental events and stable with respect to disturbances. We use a DS-based motion plan for a robotic manipulator defined as

$$\dot{x} = \hat{f}_{n(t)}(x) + u(x, t), \quad (7)$$

where, $\hat{f}_{n(t)}(x)$ encodes a nominal motion $n(t) \in \{1, 2, \dots, N\}$ at time t , while $u(x, t)$ is an

auxiliary control signal used to satisfy the task specification and enforce disturbance rejection properties. We assume that the N autonomous dynamical systems $\mathcal{D} = \{\hat{f}_i\}_{i=1}^N$ that govern the nominal motions are given to us; these may, for example, be learned from demonstrations [24]. We choose $u(x, t)$ and $n(t)$ for the robot to follow at each time step based on the discrete controllable signal σ_c that satisfy Ψ , which further depends on the uncontrollable signal σ_u . We model the controllable signal as following low level continuous motions and assume that the robot can follow at-most one motion at a given time, since we allow switching online between different motions.

We are given an RTL specification Ψ and a set of nominal autonomous dynamical systems \mathcal{D} specifying a DS-based family of motion plans (7). Our task is formalized in the following problem statement:

Develop a reactive control strategy such that at each discrete time t_k , and for any uncontrollable signal σ_u the solution $x(t)$ to (7) generates a controllable signal $\{\sigma_c(k), \sigma_c(k+1), \dots\}$ such that $([\sigma_c, \sigma_u], k) \models \Psi$ for all $t_k \geq 0$.

The above problem statement requires the robot to react online to uncontrollable environmental changes given by σ_u while also adapt to unforeseen external disturbances via the auxiliary input $u(x, t)$.

IV. PROPOSED METHODOLOGY

We present our modular control architecture that integrates a discrete task planner to satisfy the specification Ψ and a continuous DS-based motion planner to reject external disturbances while adapting to dynamic environmental changes. Unlike prior approaches [3]–[5] that do not jointly show reactivity to online environmental events and stable task execution, we propose a single framework that demonstrates reactive robot tasks with formal stability guarantees for physical robot interaction. We will use CLFs and CBFs described in Section II-B for our parameterization (7), which is a control-affine dynamical system (2) for the time period when $n(t)$ is constant with $f(x) := \hat{f}_{n(t)}(x)$ and $g(x) := I$.

A schematic of the proposed control flow is presented in Fig. 2. The blue blocks represent the offline components and the green blocks are the online planning modules. The autonomous DS models $\hat{f}_i(x)$ could either be given by a library of motion primitives [27], or can be learned from demonstrations [23], [24]. The user specifies a reactive temporal specification Ψ that the robot should satisfy and we construct the corresponding Büchi automaton \mathcal{B}_Ψ using automated tools [22]. At deployment, given an observation of the external environmental events, our task planner decides the type of robot behavior $p_m \in \mathcal{A}_c$ to follow at each time step using \mathcal{B}_Ψ so that Ψ is satisfied. The details of the task planner on how to decide p_m are given in Section IV-A. Then, the motion planner computes the reference velocity \dot{x}_{ref} to follow p_m based on (7) using the *virtual control input* $u(x, t)$, which enforces disturbance rejection properties and satisfies STL specifications included in Ψ . We denote the reference velocity for the low level controller as \dot{x}_{ref} ,

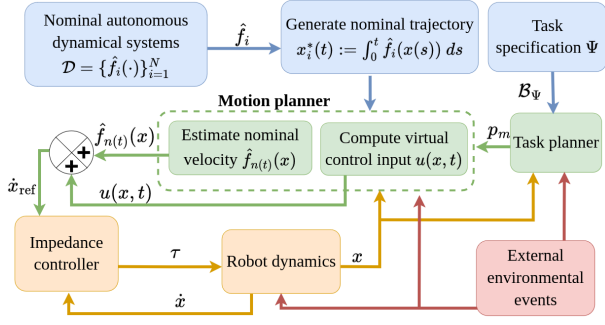


Fig. 2: The control flow of our proposed pipeline. The state of the robot is x , the low level control input are the joint torques τ , and the desired velocity in state space is \dot{x}_{ref} .

which in general may be different from the real velocity \dot{x} of the robot. The reference velocity \dot{x}_{ref} is given as input to an impedance controller [28] that computes the low level control input τ (joint torques) for the physical robotic system. We emphasize that the virtual control input $u(x, t)$ is different from the low-level control inputs τ given in Fig. 2, and $u(x, t)$ is a component of the motion planning DS (7).

A. Task Planner

We propose Algorithm 1 that outputs the behavior p_m the robot should follow at each time step to satisfy Ψ , even when the environmental events are dynamically changing.

Description of Algorithm 1: Given a task specification Ψ , we construct the Büchi automaton $\mathcal{B}_\Psi = (\mathcal{S}, s_0, AP, \Delta, \mathcal{F})$ using Spot 2.0 [22]. We denote s and s' to be the current and previous state of the automaton, respectively, and k is the current discrete time step. Lines 2-15 are implemented online at 200 Hz in the task planner block of Fig. 2. In line 6, we form a directed graph $\mathcal{G}(\mathcal{S}, E)$ from \mathcal{B}_Ψ , which is implemented only at $k = 0$ or if there is a change in the values of the uncontrollable propositions $\sigma_u(k)$. The nodes of \mathcal{G} is \mathcal{S} , while the edges E reflect valid transitions on the automaton that comply with the current value of uncontrollable propositions $\sigma_u(k)$, i.e., for all $s_1, s_2 \in \mathcal{S}$, $(s_1, s_2) \in E$ if and only if there exists a $p \in \{0, 1\}^{|\mathcal{A}_c|}$ such

Algorithm 1: Task Planner

```

Initialize:
 $\mathcal{B}_\Psi := (\mathcal{S}, s_0, \mathcal{A}_c \cup \mathcal{A}_u, \Delta, \mathcal{F}) \leftarrow \text{Automaton}(\Psi)$ ;
 $s' \leftarrow s_0, k \leftarrow 0$ ;
1 while task running do
2    $[\sigma_c(k), \sigma_u(k)] \leftarrow \text{UpdateSensor}()$ ;
3    $s \in \{s'' : (s', [\sigma_c(k), \sigma_u(k)], s'') \in \Delta\}$ ;
4   if  $k = 0$  or  $\sigma_u(k) \neq \sigma_u(k-1)$  then
5      $j \leftarrow 0, s_0 \leftarrow s$ ;
6      $\mathcal{G}(\mathcal{S}, E) \leftarrow \text{Graph}(\mathcal{B}_\Psi, \sigma_u(k))$ ;
7      $\{s_0, s_1, s_2, \dots\} \leftarrow \text{ShortestPath}(\mathcal{G}, s_0, \mathcal{F})$ ;
8   else
9     if  $s \neq s'$  then
10      |  $j \leftarrow j + 1$ ;
11    end
12  end
13   $p_m \leftarrow p_c(k)$ , where  $p_c(k) \in \mathcal{A}_c$  such that
     $(s_j, [p_c^j, \sigma_u(k)], s_{j+1}) \in \Delta$  and  $p_c^j[p_c(k)] \models \text{true}$ ;
14   $s' \leftarrow s, k \leftarrow k + 1$ ;
15  MotionPlan( $p_m$ );
16 end

```

that $(s_1, [p, \sigma_u(k)], s_2) \in \Delta$. In line 6, we plan a shortest path on $\mathcal{G}(\mathcal{S}, E)$ using breadth-first search [29] that reaches an accepting state in \mathcal{F} and stays there infinitely often. Such a path always satisfies Ψ as described in Section II-A. If there is no change in the values of uncontrollable propositions, we implement lines 9-11 so that the robot moves along the path, where index j represents the progress along the path to satisfy Ψ . Finally, in line 13, we decide the type of motion plan p_m to be implemented.

Decide motion plan p_m : Let the shortest path obtained in line 6 be $w = \{s_0, s_1, s_2, \dots\}$, where $s_0 = s$, $(s_j, [p_c^j, \sigma_u(k)], s_{j+1}) \in \Delta$ for all $j \in \mathbb{Z}^+$, and $p_c^j \in \{0, 1\}^{|\mathcal{A}_c|}$ is the value of controllable propositions when transitioning from s_j to s_{j+1} . Although there could be multiple such p_c^j , we observed from our experience that the transitions on the shortest path will have exactly one p_c^j for each unique $\sigma_u(k)$. In line 13, we denote $p_c(k) \in \mathcal{A}_c$ to be the controllable proposition that is true in p_c^j , where p_c^j is the label of the transition along the path that satisfies Ψ . We thus choose the type of motion plan p_m to be the current controllable proposition $p_c(k)$ that has to be true to move along the shortest path. For example, the automaton \mathcal{B}_Ψ is given in Fig. 3 for the specification $\Psi := G(h \Rightarrow s) \wedge G(\neg h \Rightarrow F(p))$ described in Section III-B, where $h := \text{hot}$, $s := \text{stir}$, $p := \text{press button}$. If the initialization is such that $s_0 = 0$; $h = 1, s = 1, p = 0$, then the shortest path is just the self-transition $0 \rightarrow 0$. Then, from the label of the transition $0 \rightarrow 0$ and $\sigma_u(k) = \{h, \neg p\}$, the motion plan to be implemented is $p_m = s$ so that the robot follows the behavior *stir*. If $h = 0$ while the robot is stirring ($s = 1$), then the system transitions to state 1. Now, the shortest path is computed again and the path is $1 \rightarrow 0 \rightarrow 0$. From the label of the transitions $1 \rightarrow 0 \rightarrow 0$ and $h = 0$, the motion plan to be implemented is $p_m = p$ so that the robot follows the behavior *press button*. In the next section, we describe how the desired behavior p_m is achieved by implementing a DS-based motion plan.

Remark: In this work, we assume that for each uncontrollable signal σ_u , there exists a possibly different controllable signal σ_c that satisfies the given RTL specification Ψ . Furthermore, we assume that at each discrete time step k , if $\bar{\sigma}_u(k) \neq \sigma_u(k)$ and $\bar{\sigma}_u(\tau) = \sigma_u(\tau)$ for all $\tau \neq k$, then there exists a controllable signal $\bar{\sigma}_c$ such that $\bar{\sigma}_c(i) = \sigma_c(i)$ for all $i \leq k$ and $([\bar{\sigma}_c, \bar{\sigma}_u], k) \models \Psi$. This assumption asserts the existence of a possibly different controllable signal that satisfies Ψ , even if the uncontrollable propositions change online. Without this assumption, there might not exist a controllable signal that satisfies Ψ if the environmental observations represented as uncontrollable propositions change at task execution. Such scenarios are known as deadlock modes [30], where a supervisory signal [31] is required to recover the system, which we reserve for future work.

B. Motion Planner

Given the output p_m of the discrete task planner at each time step k , we solve a QP online at 1 KHz that (i) guarantees satisfaction of STL specifications, and (ii) enforces safety




Fig. 3: Automaton \mathcal{B}_Ψ for $\Psi := G(h \Rightarrow s) \wedge G(\neg h \Rightarrow F(p))$, where $!p \equiv \neg p$.

and stability properties while following nominal motions given by autonomous DS. We also propose a simple yet efficient switching strategy so that robots switch between different behaviors when reacting to environmental observations, unlike previous approaches [14], [15] that solve expensive optimization problems to re-plan.

1) *Nominal motion plans*: We compute safe and stable motions that converge to a nominal target trajectory which is generated from an autonomous dynamical system. Although one could assume that the set of nominal dynamical systems \mathcal{D} are given to the user from an existing library of motion primitives [27], we use the CLF-CBF-NODE from our prior work [24] that learns from demonstrations DS models \hat{f}_i for each behavior $i \in \{1, 2, \dots, N\}$ that includes complex, possibly periodic motions. For example, the *stir* behavior mentioned in the previous section can be represented by a periodic trajectory $x_1^*(t)$ using a DS model $\dot{x} = \hat{f}_1(x)$ that is learned from demonstrations.

CLF-CBF-QP: We abuse notation and use $x^*(t)$ and $\hat{f}(t)$ to denote the target trajectory and the autonomous DS for a single nominal motion, but the method is the same for any other nominal motion. Let the error be $e(t) = x(t) - x^*(t)$, and for ease of notation, we drop the explicit dependence on time t , and write e , x , and x^* for the current error, state, and target point at time t , respectively. The control-affine error dynamics is $\dot{e} = \hat{f}(x) - \dot{x}^* + u(x)$ from (7). The explicit dependence of $u(x)$ and $\hat{f}(x)$ on time is removed since we focus only on time-independent CBFs to model safety sets, and CLFs for the time period when $n(t)$ is constant. In the next section, we describe how we switch between different behaviors. Then, by Theorem 1, if there exists a CLF V for the dynamics \dot{e} , then, any feedback control law $u(x) \in K(x)$ will drive the error asymptotically to zero. Similarly, if there exists a time-independent CBF B corresponding to a safety set $\mathcal{C} \subseteq \mathcal{X}$ for the dynamics (7), then, any feedback control law $u(x) \in K(x)$ will render the system safe. We combine both the properties of safety and stability and solve the below QP to compute $u(x)$, where we prioritize safety over stability.

Given the current state of the robot x and the target point x^* , the QP that guarantees a safe motion plan is

$$\begin{aligned} (u(x), -) = \arg \min_{\{v, \eta\}} & \|v\|_2^2 + \lambda \eta^2 \\ \text{s.t.} & \frac{\partial B(x)}{\partial x}^\top (\hat{f}(x) + v) \geq -\gamma(B(x)) \\ & \frac{\partial V(e)}{\partial e}^\top (\hat{f}(x) - \dot{x}^* + v) \leq -\alpha(V(e)) + \eta, \end{aligned} \quad (8)$$

where α is a class \mathcal{K} function, η is a relaxation variable to ensure feasibility of (8) and is penalized by $\lambda > 0$. We use Algorithm 1 from our prior work [24] to choose x^* and $\dot{x}^* =$

$\hat{f}(x^*)$ that ensures tracking of complex periodic trajectories in spite of external disturbances. If the safety specifications do not disturb the nominal motion, then $\eta = 0$ in (8), which guarantees stability.

2) *STL tasks*: We use time-varying control barrier functions to satisfy time-critical tasks specified by STL [19]. For example, consider the STL task $\phi ::= F_{[a,b]} \|x - x^*\| \leq \epsilon$ which encodes that the robot's end-effector position x should eventually be ϵ close to a goal point x^* within the time interval $[a, b]$. Such a task ϕ could represent the *press button* behavior mentioned in Section IV-A, where the robot should be in close proximity to the button within the time interval $[a, b]$ so that the button is pressed and the pan remains hot. A CBF for ϕ is $B(x, t) = \epsilon^2 - \|x - x^*\|^2 + \gamma(t)$, where $\gamma(\cdot)$ controls the rate at which the robot reaches x^* such that $\gamma(b) = 0, B(x, a) \geq 0$. For example, we can use $\gamma(t) = \max\{0, \frac{\|x(a) - x^*\|^2}{a-b}(t-b)\}$ so that $B(x, a) = \epsilon^2$. If there exists a control law $u(x, t)$ so that the solution of (7) satisfies $B(x(t), t) \geq 0$ for all $t \in [a, b]$, then ϕ is satisfied. Since $\gamma(t)$ is monotonically decreasing until $t = b$, the size of the set $\mathcal{C}(t)$ defined in Section II-B will decrease until $t = b$ and the control law $u(x, t) \in K(x, t)$ will drive the state x ϵ -close to x^* . We can formulate a QP similar to [19] and compute $u(x, t)$ that renders the set $\mathcal{C}(t)$ forward invariant for the dynamical system (7) with the CBF $B(x, t)$:

$$\begin{aligned} u(x, t) = \arg \min_v & \|v\|_2^2 \\ \text{s.t.} & \frac{\partial B(x, t)}{\partial x}^\top (\hat{f}_{n(t)}(x) + v) + \frac{\partial B(x, t)}{\partial t} \geq -\gamma(B(x, t)) \end{aligned} \quad (9)$$

The above problem is always feasible if $B(x, t)$ is a valid CBF and [19] presents a closed-form expression under the additional assumption that $B(x, t)$ is either non-decreasing or non-increasing in t . Both problems (8) and (9) are QPs that we solve efficiently in real-time using OSQP solver [32].

3) *Switching between behaviors*: The RTL specification requires the robot to switch between different behaviors depending on the changing environmental observations. Consider a scenario from our running example in Fig. 3, where the pan is hot ($h = 1$) at all time $t \leq t_1$, but is not hot enough ($h = 0$) at time $t_2 > t_1$. Then, the robot would be performing the *stirring* motion at time $t \leq t_1$ which involves commanding the virtual control $u(x)$ computed from (8) with $\hat{f} := \hat{f}_{n(t)}$ modeling the stirring motion. Now, at time t_2 , since $h = 0$, the robot should follow the *press button* behavior to satisfy Ψ , which could be satisfied by commanding $u(x, t)$ from (9). This scenario involves switching between two types of motion planning strategies that correspond to $p_m = s$ and $p_m = p$.

Mixing strategy: We propose a mixing strategy that connects two successive motion planners when the robot has to switch between behaviors so that there is a smooth transition. Let the robot follow a behavior $a \in \mathcal{A}_c$ during the time $t \leq t_1$ and that it has to switch to behavior $b \in \mathcal{A}_c$ when time $t > t_1$. Let the nominal vector field be $\hat{f}_a(x)$ for behavior a and $\hat{f}_b(x)$ for behavior b . Let the virtual control input be $u_a(x, t)$ for behavior a and $u_b(x, t)$ for behavior

b. Then, the commanded reference velocity at time t_1 is $\dot{x}_{\text{ref}}(t_1) = \hat{f}_a(x) + u_a(x, t_1)$. We define the mixing strategy when the robot switches from behavior $a \rightarrow b$ at time t_1 as

$$\begin{aligned} \dot{x}_{\text{ref}}(t) &= (1 - \beta(t))\dot{x}_{\text{ref}}(t_1) + \beta(t) \left(\hat{f}_b(x) + u_b(x, t) \right), \\ \beta(t) &= \max \left(0, \min \left(1, \frac{t - t_1}{\Delta t} \right) \right), \end{aligned}$$

where Δt is the during for which the mixing strategy is implemented. The mixing factor $\beta(t)$ linearly increases from 0 to 1 over the time interval $t \in [t_1, t_1 + \Delta t]$ so that the commanded reference velocity smoothly varies from $\dot{x}_{\text{ref}}(t_1)$ (behavior a) to $\dot{x}_{\text{ref}}(t_1 + \Delta t)$ (behavior b). If we increase Δt , the robot will take longer to switch between behaviors and vice-versa.

V. EXPERIMENTAL RESULTS

We demonstrate our method on a Franka robot arm performing two experiments: (i) wiping a whiteboard and (ii) wiping a mannequin, both while adapting to environmental changes and being compliant to human interactions.

A. Wiping a whiteboard

The task execution by the robot using our method is given in Fig. 4, for the same scenario that was introduced in Fig. 1. On a high level, the robot should switch between two different nominal wiping motions, react to the environment by wiping off any stain on the board, and adapt to external disturbances such as the eraser being dropped off from the end effector. The uncontrollable propositions are (i) $\text{eraser} = 1 \Leftrightarrow$ the end-effector is holding an eraser, (ii) $\text{stain} = 1 \Leftrightarrow$ there is a stain on the whiteboard, (iii) $\text{left} = 1 \Leftrightarrow$ nominal wiping motion on the left side of the board and (iv) $\text{right} = 1 \Leftrightarrow$ a different wiping motion on the right side of the board. The boolean values for left and right are controlled by the user using a GUI, the stain is detected by an Intel Realsense camera that returns the bounding box of the stain, and we assume the eraser is dropped if the gripper of the end-effector is open. The controllable propositions that model different desired

behaviors are given in Table I. The RTL specification that the robot should satisfy is

$$\begin{aligned} \Psi ::= & G((\text{eraser} \text{ and } \text{left} \text{ and } \neg \text{stain}) \Rightarrow w_{\text{left}}) \text{ and} \\ & G((\text{eraser} \text{ and } \text{right} \text{ and } \neg \text{stain}) \Rightarrow w_{\text{right}}) \text{ and} \\ & ((\text{eraser} \text{ and } \text{stain}) \Rightarrow s_p) \text{ and } (\neg \text{eraser} \Rightarrow e_p) \text{ and} \\ & G((\text{eraser} \text{ and } \text{stain}) \Rightarrow (s_p \text{ U } w_{\text{stain}})). \end{aligned}$$

B. Wiping a human mannequin

We present the experimental details of another wiping task, this time on a human mannequin, where the robot motions are given in Fig. 5. The robot should switch between two different nominal motions on the legs and the hands, wipe the back if the human gestures in that region, and wet the towel at regular intervals where a human squeezes the towel after wiping. The uncontrollable propositions are: (i) $\text{leg} = 1 \Leftrightarrow$ wipe the legs, (ii) $\text{hand} = 1 \Leftrightarrow$ wipe the hands, (iii) $\text{back} = 1 \Leftrightarrow$ human gesturing on the back, (iv) $\text{wet} = 1 \Leftrightarrow$ towel is wet, and (v) $\text{human} = 1 \Leftrightarrow$ human is inside the workspace to squeeze the towel. The boolean values for leg and hand are controlled by the user using a GUI, the human gesture on the back is detected by an Intel Realsense camera, and the human being inside the robot workspace is detected by Optitrack markers on the human hand. We assume that the towel becomes dry, i.e., $\text{wet} = 0$, every 30 secs. The controllable propositions are given in Table II, and RTL specification is

$$\begin{aligned} \Psi ::= & G((\text{wet} \text{ and } \text{leg} \text{ and } \neg \text{back}) \Rightarrow w_{\text{leg}}) \text{ and} \\ & G((\text{wet} \text{ and } \text{hand} \text{ and } \neg \text{back}) \Rightarrow w_{\text{hand}}) \text{ and} \\ & G((\text{wet} \text{ and } \text{back}) \Rightarrow w_{\text{back}}) \text{ and } G(\neg \text{wet} \Rightarrow d_p) \text{ and} \\ & G((\text{human} \text{ and } \text{wet}) \Rightarrow b_p) \text{ and } G(\neg \text{human} \text{ and } \text{wet}) \Rightarrow \text{true}). \end{aligned}$$

Implementation details: For both the experiments, we use Spot 2.0 [22] to construct the automaton \mathcal{B}_Ψ that has 43 states, 269 transitions for the whiteboard task; and 64 states, 510 transitions for the mannequin task. The nominal DS \hat{f}_a in Table I and II are learned from 3 demonstrations for each motion using Neural ODEs [24] that efficiently capture periodic motions. For wiping off the stain (w_{stain}), we generate a nominal trajectory online modeled as \hat{f}_{stain} that covers the bounding box of the detected stain, and the

| Sym. | Boolean condition | Description |
|--------------------|--|--|
| w_a | $w_a = 1 \Leftrightarrow \dot{x}(t) = \hat{f}_a(x(t))$ | Wiping motion by nominal autonomous DS \hat{f}_a , $a \in \{\text{left, right}\}$ indicates the side on the board. |
| w_{stain} | $w_{\text{stain}} = 1 \Leftrightarrow \dot{x}(t) = \hat{f}_{\text{stain}}(x(t))$ | Motion to wipe off the detected stain online given by a nominal autonomous DS \hat{f}_{stain} |
| s_p | $s_p = 1 \Leftrightarrow \exists F \in [0, t_s] \ x(t) - x_s\ \leq \epsilon$ | The end-effector should eventually (F) be ϵ close to the stain x_s within time interval $[0, t_s]$. |
| e_p | $e_p = 1 \Leftrightarrow \exists F \in [0, t_e] \ x(t) - x_e\ \leq \epsilon$ | The end-effector should eventually be ϵ close to the eraser e_p within $[0, t_e]$. |




TABLE I: Controllable propositions (Whiteboard)

Fig. 4: Whiteboard wiping: (a) switching between nominal motions (b) going to the stain position (c) wiping off stain (d) human perturbation and eraser pick up.

| Symbol | Boolean condition | Description |
|--------|---|---|
| w_a | $w_a = 1 \Leftrightarrow \dot{x}(t) = \hat{f}_a(x(t))$ | Wiping motion on $a \in \{\text{leg, hand, back}\}$ given by a nominal autonomous DS \hat{f}_a |
| d_p | $d_p = 1 \Leftrightarrow \exists F \in [0, t_d] \ x(t) - x_d\ \leq \epsilon$ | The end-effector should eventually (F) be ϵ -close to the dip position x_d within the time interval $[0, t_d]$. |
| b_p | $b_p = 1 \Leftrightarrow \exists F \in [0, t_b] \ x(t) - x_b\ \leq \epsilon$ | The end-effector should eventually be ϵ -close to the bowl position x_b within $[0, t_b]$. |




TABLE II: Controllable propositions (mannequin)

Fig. 5: Mannequin wiping: (a) switching between nominal motions (b) going to the back (c) wiping the back (d) dipping the towel in water (e) human squeezing the wet towel and (f) continuing wiping.

trajectory is tracked using CLFs. We compute the time within which the robot should satisfy STL tasks as $t^* = \frac{\|x(0) - x^*\|}{v_u}$, where $(t^*, x^*) \in \{(t_s, x_s), (t_e, x_e), (t_d, x_d), (t_b, x_b)\}$, v_u is analogous to a uniform velocity for the robot and $x(0)$ is the state when the robot starts to follow the relevant behavior. We satisfy the STL tasks using time-varying CBFs as described in Section IV-B.2 with $\gamma(t) = \max \left\{ 0, \frac{\|x(0) - x^*\|_2^2 (e^{-t} - e^{-t^*})}{1 - e^{-t^*}} \right\}$ and $\epsilon = 5$ cm. Such an exponential γ function generates a motion plan that is faster at the beginning with smoother convergence to the target point x^* , when compared to a linear γ function which is typically used in prior work [19], [20]. We use $\Delta t = 0.67$ secs for the mixing strategy that generates smooth switching behaviors while also reacting quickly to environmental changes. Additional experiment videos and code-base is available at <https://sites.google.com/view/rtl-plan-control-hri>.

VI. CONCLUSION & FUTURE WORK

We propose a modular reactive planner which integrates discrete temporal logic based plans and continuous DS-based motion plans that guarantees task satisfaction and stability while adapting to unforeseen disturbances and external observations. We demonstrate our method on the Franka robot arm for reactive wiping tasks involving safe human interactions.

One limitation of our work is not incorporating deadlock scenarios [30] when robots cannot make task progress. Our future work aims to address such issues by integrating Large Language Models [33] into LTL specifications that can act at a high level as a supervisory control [31]. Future work will also investigate planning strategies to satisfy task specifications on $\mathcal{SO}(3)$ and $\mathcal{SE}(3)$ manifolds.

REFERENCES

- [1] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” *arXiv preprint arXiv:2204.01691*, 2022.
- [2] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, “Code as policies: Language model programs for embodied control,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 9493–9500.
- [3] J. Mendez-Mendez, L. P. Kaelbling, and T. Lozano-Pérez, “Embodied lifelong learning for task and motion planning,” in *Conference on Robot Learning*. PMLR, 2023, pp. 2134–2150.
- [4] J. J. Kuffner, S. Kagami, K. Nishiwaki, M. Inaba, and H. Inoue, “Dynamically-stable motion planning for humanoid robots,” *Autonomous robots*, vol. 12, pp. 105–118, 2002.
- [5] R. Holladay, T. Lozano-Pérez, and A. Rodriguez, “Robust planning for multi-stage forceful manipulation,” *The International Journal of Robotics Research*, vol. 43, no. 3, pp. 330–353, 2024.
- [6] E. Plaku and S. Karaman, “Motion planning with temporal-logic specifications: Progress and challenges,” *AI communications*, vol. 29, no. 1, pp. 151–162, 2016.
- [7] Y. Wang, N. Figueroa, S. Li, A. Shah, and J. Shah, “Temporal logic imitation: Learning plan-satisficing motion policies from demonstrations,” *arXiv preprint arXiv:2206.04632*, 2022.
- [8] V. Vasilopoulos, Y. Kantaros, G. J. Pappas, and D. E. Koditschek, “Reactive planning for mobile manipulation tasks in unexplored semantic environments,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6385–6392.
- [9] 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.
- [10] 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.
- [11] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas, “Temporal logic motion planning for dynamic robots,” *Automatica*, vol. 45, no. 2, pp. 343–352, 2009.
- [12] C. Belta and S. Sadraddini, “Formal methods for control synthesis: An optimization perspective,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 115–140, 2019.
- [13] O. Maler and D. Nickovic, “Monitoring temporal properties of continuous signals,” in *International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*. Springer, 2004, pp. 152–166.
- [14] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia, “Reactive synthesis from signal temporal logic specifications,” in *Proceedings of the 18th international conference on hybrid systems: Computation and control*, 2015, pp. 239–248.
- [15] M. Guo and D. V. Dimarogonas, “Multi-agent plan reconfiguration under local ltl specifications,” *The International Journal of Robotics Research*, vol. 34, no. 2, pp. 218–235, 2015.
- [16] A. I. M. Ayala, S. B. Andersson, and C. Belta, “Temporal logic motion planning in unknown environments,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 5279–5284.
- [17] C. I. Vasile, X. Li, and C. Belta, “Reactive sampling-based path planning with temporal logic specifications,” *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 1002–1028, 2020.
- [18] F. Nawaz and M. Ornik, “Explorative probabilistic planning with unknown target locations,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 2732–2737.
- [19] L. Lindemann and D. V. Dimarogonas, “Control barrier functions for signal temporal logic tasks,” *IEEE control systems letters*, vol. 3, no. 1, pp. 96–101, 2018.
- [20] D. Gundana and H. Kress-Gazit, “Event-based signal temporal logic synthesis for single and multi-robot tasks,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3687–3694, 2021.
- [21] L. Lindemann and D. V. Dimarogonas, “Efficient automata-based planning and control under spatio-temporal logic specifications,” in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 4707–4714.
- [22] A. Duret-Lutz, A. Lewkowicz, A. Fauchille, T. Michaud, E. Renault, and L. Xu, “Spot 2.0—a framework for ltl and-automata manipulation,” in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2016, pp. 122–129.
- [23] A. Billard, S. Mirrazavi, and N. Figueroa, *Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach*. MIT Press, 2022.
- [24] F. Nawaz, T. Li, N. Matni, and N. Figueroa, “Learning safe and stable motion plans with neural ordinary differential equations,” *arXiv preprint arXiv:2308.00186*, 2023.
- [25] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control barrier functions: Theory and applications,” in *2019 18th European control conference (ECC)*. IEEE, 2019, pp. 3420–3431.
- [26] O. Katsuhiko, *Modern control engineering*. Editorial Félix Varela, 2009.
- [27] M. Saveriano, F. Abu-Dakka, A. Kramberger, and L. Peternel, “Dynamic movement primitives in robotics: A tutorial survey. arxiv 2021,” *arXiv preprint arXiv:2102.03861*.
- [28] K. Kronander and A. Billard, “Passive interaction control with dynamical systems,” *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 106–113, 2016.
- [29] A. Bundy and L. Wallen, “Breadth-first search,” *Catalogue of artificial intelligence tools*, pp. 13–13, 1984.
- [30] J. Alonso-Mora, J. A. DeCastro, V. Raman, D. Rus, and H. Kress-Gazit, “Reactive mission and motion planning with deadlock resolution avoiding dynamic obstacles,” *Autonomous Robots*, vol. 42, pp. 801–824, 2018.
- [31] C. G. Cassandras and S. Lafortune, Eds., *Supervisory Control*. Boston, MA: Springer US, 2008, pp. 133–221.
- [32] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “OSQP: an operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.
- [33] J. X. Liu, Z. Yang, I. Idrees, S. Liang, B. Schornstein, S. Tellex, and A. Shah, “Grounding complex natural language commands for temporal tasks in unseen environments,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1084–1110.