

Synergistic Offline-Online Control Synthesis via Local Gaussian Process Regression

John Jackson¹, Luca Laurenti², Eric Frew¹, and Morteza Lahijanian¹

Abstract—Autonomous systems often have complex and possibly unknown dynamics due to, e.g., black-box components. This leads to unpredictable behaviors and makes control design with performance guarantees a major challenge. This paper presents a data-driven control synthesis framework for such systems subject to linear temporal logic on finite traces (LTLf) specifications. The framework combines a baseline (offline) controller with a novel online controller and refinement procedure that improves the baseline guarantees as new data is collected. The baseline controller is computed offline on an uncertain abstraction constructed using Gaussian process (GP) regression on a given dataset. The offline controller provides a lower bound on the probability of satisfying the LTLf specification, which may be far from optimal due to both discretization and regression errors. The synergy arises from the online controller using the offline abstraction along with the current state and new data to choose the next best action. The online controller may improve the baseline guarantees since it avoids the discretization error and reduces regression error as new data is collected. The new data are also used to refine the abstraction and offline controller using local GP regression, which significantly reduces the computation overhead. Evaluations show the efficacy of the proposed offline-online framework, especially when compared against the offline controller.

I. INTRODUCTION

The accelerating development of autonomous systems technology has led to their proliferation in many facets of our society. Examples include manufacturing robots on factory floors, autonomous cars on the roads, and delivery drones in urban airspaces. Such systems often have complex and possibly unknown dynamics due to, e.g., black-box components, making accurate modeling unattainable. This introduces a major challenge for analysis and control design for such systems. This challenge is especially important to confront in *safety-critical* domains, where assurances are required. This work considers this challenge and aims to provide a control synthesis framework that uses data in lieu of a system model to provide performance guarantees.

A powerful approach to control design with assurances is formal synthesis, e.g., [1]–[5]. In this approach, the system specification is expressed in a formal language such as *linear temporal logic* (LTL) [6], and the system evolution is abstracted to a finite model—called *abstraction*—with a simulation relation. Then, using automated verification techniques, a controller is synthesized on the abstraction with

performance guarantees. The controller and its guarantees are then refined onto the underlying system. These frameworks provide strong assurances but rely on accurate dynamics model to construct the abstraction, which may be unavailable for complex autonomous systems.

In recent years, data-driven control approaches have been increasingly studied, e.g., [7]–[11]. These methods typically consist of a *machine learning* component that constructs a variety of regression models, e.g., polynomial functions and neural networks. Among them, *Gaussian process* (GP) regression [12] offers a unique advantage of providing quantified error bounds between the regressed model and the underlying system [13]–[16]. This has given rise to a wide use of GPs in safe learning frameworks [8]–[10]. In particular, safe reinforcement learning approaches are developed to take advantage of the GP modeling to learn a control policy with guarantees such as stability [8], [11] and respecting safety constraints [9], [10]. A shortcoming of classical GP regression, however, is that it leads to computational intractability as the data grows, making it difficult to use for online refinement of policies.

Sparse GP approximation has been explored extensively to handle large datasets [17]. These approximations range from choosing subset of data, to making conditional assumptions on the joint prior using inducing datapoints. While powerful, sparse GP approximations can lack analytical guarantees [18]. The use of localized GP regressions to model local dynamics has been studied for model learning and control [19], [20], and it is shown that they can outperform sparse GP approximations if intelligent data partitioning is used [21]. Although they have been used in various control settings, formal online synthesis using GPs remains an open challenge.

In previous work [22], we introduce a data-driven formal control synthesis framework for unknown stochastic systems. The framework takes a set of input-output data of the system and a desired property in *LTL on finite traces* (LTLf) [23] as input and synthesizes a *robust* control policy with guaranteed lower-bound on the probability of achieving the LTLf property. The method employs GP regression and a discretization to construct an abstraction in the form of an uncertain Markov process for the unknown system with a simulation relation. It then computes a robust policy on the abstraction with a quantified bound on the error (distance) to the optimal solution. This error is due to both the discretization and regression errors in generating the abstraction. Given the inherent computational complexity, the framework is completely offline and cannot take advantage of runtime data to reduce its error.

In this work, we build on [22], [24] and propose a

This work was supported in part by the NSF grant 2039062 and NSF Center for Unmanned Aircraft Systems under award IIP-1650468.

¹ Smead Aerospace Engineering Sciences, University of Colorado Boulder, CO, USA. {john.m.jackson, eric.frew, morteza.lahijanian}@colorado.edu

² Center of Systems and Control, TU Delft University, The Netherlands. l.laurenti@tudelft.nl

synergistic offline-online framework to approach an optimal solution. We combine the baseline offline controller in [22] with a novel online controller and refinement procedure that improves the baseline guarantees as new data is collected. The synergy arises from the online controller using the offline abstraction along with the current state and new data to choose the next best action. The online controller improves the baseline guarantees by avoiding the discretization error and reducing regression error as new data is collected. The new data are also used to refine the abstraction and offline controller via local GP regression, which significantly reduces the computation overhead, enabling online computations.

The contributions of this work are threefold. First, we introduce an online procedure for control synthesis that improves offline control policy's probability of satisfaction of the LTLf specification. Specifically, we show that the online method works in synergy with the offline framework to monotonically decrease the error to the optimal solution at runtime. This is more general than the safety problem considered in [9], [10]. Second, we reduce computational overhead to enable refinement of the abstraction online, and hence, improving the baseline controller via local GPs. We show that a global GP is too expensive to use online, and local GPs enable comparatively quick computations. Thirdly, we compare the performance of the offline method against the proposed synergistic framework in simulation, illustrating the efficacy of the proposed method.

II. PROBLEM FORMULATION

We consider a stochastic system given by

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k \quad (1)$$

where $k \in \mathbb{N}$, $\mathbf{x}_k \in \mathbb{R}^n$, and $\mathbf{u}_k \in U = \{u_1, \dots, u_{|U|}\}$ is a finite set of controls or actions. \mathbf{x}_k is a discrete-time controlled stochastic process governed by a fully- or partially-unknown, non-linear function $f: \mathbb{R}^n \times U \rightarrow \mathbb{R}^n$ driven by an additive noise term $\mathbf{w}_k \in \mathbb{R}^n$. We assume that \mathbf{w}_k follows a stationary and independent sub-Gaussian distribution $p_{\mathbf{w}}$. This includes Gaussian distributions as well as all distributions with bounded support [25].

To reason about Process (1) without having full knowledge of f , we assume to have a set of state-action-state measurements $D = \{(x_i, u_i, x_i^+)_{i=1}^m\}$ generated by Process (1), where $x_i^+ \in \mathbb{R}^n$ is a sample of one-step evolution of Process (1) initialized at $x_i \in \mathbb{R}^n$ with control $u_i \in U$. Our goal is to use D as well as the data collected at runtime to infer $f(\cdot, u)$ for each $u \in U$. The following assumption guarantees that $f(\cdot, u)$ can be learned arbitrarily well via GP regression.

Assumption 1. For a compact set $X \subset \mathbb{R}^n$, let $\kappa: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a given kernel and $\mathcal{H}_\kappa(X)$ the reproducing kernel Hilbert space (RKHS) of functions over X corresponding to κ with norm $\|\cdot\|_\kappa$ [13]. Then, for each $u \in U$ and $i \in \{1, \dots, n\}$, $f^{(i)}(\cdot, u) \in \mathcal{H}_\kappa(X)$ and for a constant $B_i > 0$, it holds that $\|f^{(i)}(\cdot, u)\|_\kappa \leq B_i$, where $f^{(i)}$ is the i -th component of f .

For instance, assuming that κ is the widely used squared exponential kernel (as in our experiments), then $\mathcal{H}_\kappa(X)$ is

a space of functions that is dense with respect to the set of continuous functions on X , i.e., members of $\mathcal{H}_\kappa(X)$ can approximate any continuous function on X arbitrarily well [26].

We denote by $\omega_{\mathbf{x}} = x_0 \xrightarrow{u_0} x_1 \xrightarrow{u_1} \dots$ a path or trajectory of \mathbf{x}_k and use $\omega_{\mathbf{x}}(k) = x_k$ to indicate the state of $\omega_{\mathbf{x}}$ at time k . Further, we denote by $\Omega_{\mathbf{x}}^{\text{fin}}$ the set of all sample paths with finite length, i.e., the set of trajectories $\omega_{\mathbf{x}}^k = x_0 \xrightarrow{u_0} x_1 \xrightarrow{u_1} \dots \xrightarrow{u_{k-1}} x_k$ for all $k \in \mathbb{N}$. With a slight abuse of notation, given a path $\omega_{\mathbf{x}}$, we denote by $\omega_{\mathbf{x}}^k$ the prefix of $\omega_{\mathbf{x}}$ up to step k . A control strategy $\pi_{\mathbf{x}}: \Omega_{\mathbf{x}}^{\text{fin}} \rightarrow U$ is a function that chooses the next control $u \in U$ given a finite path $\omega_{\mathbf{x}}^k \in \Omega_{\mathbf{x}}^{\text{fin}}$.

For $u \in U$, a Borel measurable set $X \subseteq \mathbb{R}^n$, and $x \in \mathbb{R}^n$, call

$$T(X | x, u) = \int \mathbf{1}_X(f(x, u) + \bar{v})p_{\mathbf{w}}(\bar{v})d\bar{v},$$

the stochastic transition function induced by Process (1), where

$$\mathbf{1}_X(x) = \begin{cases} 1 & \text{if } x \in X \\ 0 & \text{otherwise} \end{cases}$$

is the indicator function. Kernel $T(X | x, u)$ describes the probability of \mathbf{x} ending in set X in one-step evolution given the current state x and control u . Given a control strategy $\pi_{\mathbf{x}}$ and a time horizon $[0, N]$, it is possible [27] to define a probability measure P over the paths of \mathbf{x}_k uniquely generated by T and a (fixed) initial condition $x_0 \in \mathbb{R}^n$ such that

$$P[\omega_{\mathbf{x}}^N(0) \in X] = \mathbf{1}_X(x_0),$$

and for $k \in \{1, \dots, N\}$,

$$P[\omega_{\mathbf{x}}^N(k) \in X | \omega_{\mathbf{x}}^N(k-1) = x, \pi_{\mathbf{x}}] = T(X | x, \pi_{\mathbf{x}}(\omega_{\mathbf{x}}^{k-1})).$$

Furthermore, for $N = \infty$, P is still uniquely defined by T by the Ionescu-Tulcea extension theorem [28].

A. Specification Language

We are interested in the properties of Process (1) in a compact set $X \subset \mathbb{R}^n$ with respect to a finite set of closed regions of interest $R = \{r_1, \dots, r_{|R|}\}$, where $r_i \subseteq X$. To this end, we associate to each region r_i the atomic proposition \mathbf{p}_i such that $\mathbf{p}_i = \top$ (i.e., \mathbf{p}_i is true) if $x \in r_i$; otherwise $\mathbf{p}_i = \perp$ (i.e., \mathbf{p}_i is false). Let $AP = \{\mathbf{p}_1, \dots, \mathbf{p}_{|R|}\}$ denote the set of all atomic propositions and $L: X \rightarrow 2^{AP}$ be the labeling function that assigns to state x the set of atomic propositions that are true at x . Then, we define the trace of path $\omega_{\mathbf{x}}^k = x_0 \xrightarrow{u_0} x_1 \xrightarrow{u_1} \dots \xrightarrow{u_{k-1}} x_k$ to be

$$\rho = \rho_0 \rho_1 \dots \rho_k,$$

where $\rho_i = L(x_i) \in 2^{AP}$ for all $i \leq k$. With an abuse of notation, we use $L(\omega_{\mathbf{x}}^k)$ to denote the trace of $\omega_{\mathbf{x}}^k$.

We use a temporal logic to express desired properties of Process (1). Classical temporal logics such as LTL specifications [6], however, are interpreted over infinite behaviors (traces), and given the high levels of uncertainty of Process (1) (unknown dynamics as well as noise), its infinite behaviors have trivial probabilities. Therefore, we instead

employ recently developed *LTL over finite traces* (LTLf) [23], which has the same syntax as LTL but its semantics is defined over finite traces.

Definition 1 (LTLf Syntax). *An LTLf formula φ over a set of atomic propositions AP is recursively defined as*

$$\varphi := \top \mid \mathbf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathcal{X}\varphi \mid \varphi \mathcal{U}\varphi$$

where $\mathbf{p} \in AP$, \neg (negation) and \wedge (conjunction) are Boolean operators, and \mathcal{X} (next) and \mathcal{U} (until) are temporal operators.

Definition 2 (LTLf Semantics). *The semantics of an LTLf formula φ are defined over finite traces. Let trace $\rho \in (2^{AP})^*$, $|\rho|$ denote its length, and ρ_i be the i -th symbol of ρ . Then, the satisfaction of φ by the i -th step of ρ , denoted by $\rho, i \models \varphi$, is recursively defined as*

- $\rho, i \models \top$,
- $\rho, i \models \mathbf{p} \Leftrightarrow \mathbf{p} \in \rho_i$,
- $\rho, i \models \neg\varphi \Leftrightarrow \rho, i \not\models \varphi$,
- $\rho, i \models \varphi_1 \wedge \varphi_2 \Leftrightarrow \rho, i \models \varphi_1$ and $\rho, i \models \varphi_2$,
- $\rho, i \models \mathcal{X}\varphi \Leftrightarrow |\rho| > i + 1$ and $\rho, i + 1 \models \varphi$,
- $\rho, i \models \varphi_1 \mathcal{U}\varphi_2 \Leftrightarrow \exists j$ s.t. $i \leq j < |\rho|$ and $\rho, j \models \varphi_2$ and $\forall k, i \leq k < j, \rho, k \models \varphi_1$.

Finite trace ρ satisfies φ , denoted by $\rho \models \varphi$, if $\rho, 0 \models \varphi$.

Similar to LTL, the temporal operators \mathcal{F} (eventually) and \mathcal{G} (globally) are defined as: $\mathcal{F}\varphi = \top \mathcal{U}\varphi$ and $\mathcal{G}\varphi = \neg\mathcal{F}\neg\varphi$.

An LTLf formula φ defines a language $\mathcal{L}(\varphi) = \{\rho \in (2^{AP})^* \mid \rho \models \varphi\}$, which is in fact a regular language. Similar to [29], we say that a path ω_x of Process (1) satisfies LTLf formula φ if there exists a prefix of ω_x that lies entirely in X and its trace is in $\mathcal{L}(\varphi)$, i.e.,

$$\omega_x \models \varphi \Leftrightarrow \exists k \in \mathbb{N} \text{ s.t. } L(\omega_x^k) \in \mathcal{L}(\varphi) \text{ and } \omega_x^k(k') \in X \quad \forall k' \leq k, \quad (2)$$

B. Problem Statement

Given an LTLf specification φ and dataset D , our goal is to synthesize a control strategy π_x^* under which Process (1) attempts to satisfy φ with the maximum probability.

Problem 1. [Offline-Online Synthesis] *Given an initial dataset D and the capability of collecting data of Process (1) at runtime, a compact set X , and an LTLf property φ defined over the regions of interest in X , find a control strategy π_x^* that maximizes the probability of satisfying φ , i.e.,*

$$\pi_x^* = \arg \max_{\pi_x} P[\omega_x \models \varphi \mid \pi_x, \omega_x(0) = x_0].$$

The fact that Process (1) is unknown makes solving Problem 1 particularly challenging. In this work, we develop a data-driven method that achieves the optimality objective of the problem at the limit. In particular, we develop an offline-online synergistic framework, where the offline module uses GP regression on D to build an initial abstraction of Process (1) in terms of a finite uncertain Markov model and synthesizes a robust control strategy on the resulting abstraction. This abstraction and control strategy suffer from two types of errors: discretization and regression errors. Via

the online module, we iteratively reduce these errors as explained below.

The offline strategy serves as a baseline controller for the online module and provides an initial performance guarantee in terms of the lower bound on the probability of satisfying φ . Online, while executing the system, we use the observed data to refine the abstraction, potentially improving the control strategy and its probability of satisfaction of φ in two ways: 1) we reduce the space discretization error in the abstraction by reducing the uncertainty in the transition probabilities of the abstraction, which can be done by leveraging our continuous state knowledge at every time step, and 2) using the data collected at each time step, we improve our learned model of Process (1), hence, reducing the regression error. We show that updates to the abstraction and baseline controller at runtime is computationally feasible by learning a series of local GPs instead of a single global GP. Furthermore, we show that the framework formally accounts for the uncertainty in the learning process, and the online guarantees see monotonic improvement over those generated offline.

III. PRELIMINARIES

A. Gaussian Process Regression

Gaussian Process (GP) regression is a non-parametric Bayesian machine learning method [12] that aims to infer an unknown function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ from noisy data. A standard assumption of GP regression is that f is a sample from a GP with zero mean and covariance $\kappa: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$. Let $D = \{(x_i, y_i), i \in \{1, \dots, m\}\}$ be a dataset, where y_i is a sample of an observation of $f(x_i)$ with independent zero-mean noise v , which is assumed to be normally distributed with variance σ^2 and X and Y be ordered vectors with all points in D such that $X_i = x_i$ and $Y_i = y_i$. Further, call $K(X, X)$ the matrix with $K_{i,j}(X_i, X_j) = \kappa(x_i, x_j)$, $K(x, X)$ the vector such that $K_i(x, X) = \kappa(x, X_i)$, and $K(X, x)$ defined accordingly. Then, the predictive distribution of f at a test point x is given by the conditional distribution of f given D , which is Gaussian with mean μ_D and variance σ_D^2 given by

$$\mu_D(x) = K(x, X)(K(X, X) + \sigma^2 I_m)^{-1} Y \quad (3)$$

$$\sigma_D^2(x) = \kappa(x, x) - K(x, X)(K(X, X) + \sigma^2 I_m)^{-1} K(X, x), \quad (4)$$

where I_m is the identity matrix of size $m \times m$. For stationary kernels, such as the squared-exponential kernel, $\sigma_D(\cdot)$ sees monotonic decay as m increases. As $m \rightarrow \infty$, the posterior covariance $\sigma_D(\cdot) \rightarrow 0$ and $\mu_D \rightarrow g$ [12].

In this work, contrary to the standard GP regression setting described above, we do not assume that f is sampled from a given GP or that the observation noise is Gaussian. To quantify the estimation error in this more agnostic setting, under Assumption 1, we can rely on Theorem 2 from [15], which bounds the regression error with high-probability in the form:

$$P[\forall x \in X, |\mu_D(x) - f(x)| < \beta(\delta)\sigma_D(x)] \geq 1 - \delta, \quad (5)$$

where $\delta \in (0, 1)$ specifies the upper bound, and β is a compensator that depends on the choice of κ , the dataset D , and δ . We refer to [15] for further information on β and (5).

B. Interval Markov Decision Processes (IMDPs)

We use *interval Markov decision process* (IMDPs) as the abstraction model for Process (1). IMDPs generalize Markov decision processes (MDPs) by allowing an interval of values for transition probabilities [30].

Definition 3 (IMDP). *An interval Markov decision process (IMDP) is a tuple $\mathcal{I} = (Q, A, \check{P}, \hat{P}, AP, L)$, where*

- Q is a finite set of states,
- A is a finite set of actions, and $A(q)$ denotes the set of available actions at state $q \in Q$,
- $\check{P} : Q \times A \times Q \rightarrow [0, 1]$ is a function, where $\check{P}(q, a, q')$ defines the lower bound of the transition probability from state $q \in Q$ to state $q' \in Q$ under action $a \in A(q)$,
- $\hat{P} : Q \times A \times Q \rightarrow [0, 1]$ is a function, where $\hat{P}(q, a, q')$ defines the upper bound of the transition probability from state $q \in Q$ to state $q' \in Q$ under action $a \in A(q)$,
- AP is a finite set of atomic propositions,
- $L : Q \rightarrow 2^{AP}$ is a labeling function that assigns to each state $q \in Q$ a subset of AP .

For all $q, q' \in Q$ and $a \in A(q)$, it holds that $\check{P}(q, a, q') \leq \hat{P}(q, a, q')$ and $\sum_{q' \in Q} \check{P}(q, a, q') \leq 1 \leq \sum_{q' \in Q} \hat{P}(q, a, q')$.

A path of an IMDP is a sequence of states $\omega_{\mathcal{I}} = q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \xrightarrow{a_2} \dots$ such that $a_k \in A(q_k)$ and $\hat{P}(q_k, a_k, q_{k+1}) > 0$ for all $k \in \mathbb{N}$. We denote the last state of a finite path $\omega_{\mathcal{I}}^{\text{fin}}$ by $\text{last}(\omega_{\mathcal{I}}^{\text{fin}})$ and the set of all finite and infinite paths by $\text{Paths}^{\text{fin}}$ and Paths , respectively.

Definition 4 (Strategy). *A strategy π of an IMDP model \mathcal{I} is a function $\pi : \text{Paths}^{\text{fin}} \rightarrow A$ that maps a finite path $\omega_{\mathcal{I}}^{\text{fin}}$ of \mathcal{I} onto an action in $A(\text{last}(\omega_{\mathcal{I}}^{\text{fin}}))$. If a strategy depends only on $\text{last}(\omega_{\mathcal{I}}^{\text{fin}})$, it is called a *memoryless (stationary) strategy*. The set of all strategies is denoted by Π .*

Once an action is chosen by strategy π , we evolve from the current state to the next state according to a probability distribution that respects the transition probability bounds of the IMDP. There exist possibly infinitely many such distributions, and an *adversary* chooses this distribution.

Definition 5 (Adversary). *For an IMDP \mathcal{I} , an adversary is a function $\xi : \text{Paths}^{\text{fin}} \times A \rightarrow \mathcal{D}(Q)$ that, for each finite path $\omega_{\mathcal{I}}^{\text{fin}} \in \text{Paths}^{\text{fin}}$, state $q = \text{last}(\omega_{\mathcal{I}}^{\text{fin}})$, and action $a \in A(\text{last}(\omega_{\mathcal{I}}^{\text{fin}}))$, assigns a feasible distribution γ_q^a which satisfies $\check{P}(q, a, q') \leq \gamma_q^a(q') \leq \hat{P}(q, a, q')$. The set of all adversaries is denoted by Ξ .*

Given a strategy π and an adversary ξ , a probability measure over the paths of IMDP \mathcal{I} can be defined using the induced Markov chain [4].

C. Deterministic Finite Automaton (DFA)

Given LTLf formula φ , a DFA can be constructed that precisely accepts the language of φ [23].

Definition 6 (DFA). *A deterministic finite automaton (DFA) constructed from an LTLf formula φ defined over atomic propositions AP is a tuple $\mathcal{A}_{\varphi} = (Z, 2^{AP}, \delta, z_0, Z_F)$, where Z is a finite set of states, 2^{AP} is a finite set of input symbols, $\delta : Z \times 2^{AP} \rightarrow Z$ is the transition function, $z_0 \in Z$ is the initial state, and $Z_F \subseteq Z$ is the set of accepting (final) states.*

A finite run on \mathcal{A}_{φ} is a sequence of states $z_0 z_1 \dots z_{n+1}$ induced by trace $\rho = \rho_0 \rho_1 \dots \rho_n$, where $\rho_i \in 2^{AP}$ and $z_{i+1} = \delta(z_i, \rho_i)$. A finite run on ρ is accepting if $z_n \in Z_F$. If the run is accepting, then trace ρ is accepted by \mathcal{A}_{φ} . The set of all traces that are accepted by \mathcal{A}_{φ} is call the language of \mathcal{A}_{φ} , denoted by $\mathcal{L}(\mathcal{A}_{\varphi})$. This language is equal to the language of φ , i.e., $\mathcal{L}(\varphi) = \mathcal{L}(\mathcal{A}_{\varphi})$.

IV. OFFLINE-ONLINE SYNTHESIS

In this section, we describe our synergistic offline-online control synthesis framework, which consists of two modules. The offline module constructs an IMDP abstraction for Process (1) and generates the baseline control strategy with a lower-bound guarantee on the probably of satisfying φ . The online module uses the offline abstraction and lower-bound probabilities to refine the control strategy and abstraction, which results in monotonic improvements to the guarantees. We first briefly describe the offline module, which is based on previous work [22], and then detail the online module, which is the main contribution of this work.

A. Offline Module

The offline module first performs GP regression to estimate the unknown function $f(\cdot, u)$ for each $u \in U$ from the given initial dataset D . This is achieved by using (3) and (4) on each component of f . Note that the corresponding regression error $|\hat{f}^{(i)}(x, u) - f^{(i)}(x, u)|$, where $\hat{f}^{(i)}$ is the estimate (GP mean) of the i -th component of f obtained from (3), is given by (5).

Given \hat{f} , the method constructs the IMDP abstraction $\mathcal{I} = (Q, A, \check{P}, \hat{P}, AP, L)$ by partitioning of the compact set X . Each resulting discrete region is associated with an IMDP state $q \in Q$. One state of the IMDP is also used to represent the rest of the continuous space $\mathbb{R}^n \setminus X$. The labeling function L of the IMDP states is defined according to the labels of continuous states within each region. We remark that, to ensure correct labeling, the space discretization performed above must respect the regions of interest in R . This guarantees that the label of the discrete regions correctly hold for every continuous state in that region, i.e., for discrete region $q \subset \mathbb{R}^n$, $L(q) = L(x)$ for every $x \in q$.

Next, the set of IMDP actions is defined to be the set of controls of Process (1), i.e., $A = U$. The transition probability bounds are then given by the following proposition, which uses the discrete regions, the regressed GP, and its corresponding error.

Proposition 1 ([22], Theorem 1). *For a region $q \subset \mathbb{R}^n$ and control $u \in A$, let $\|h\|_{\infty}^{q,u} \equiv \sup_{x \in q} |h(x, u)|$. Given a control $u \in A$, regions $q, q' \subset \mathbb{R}^n$, dataset D , regression \hat{f} ,*

and positive real vectors $\epsilon \in \mathbb{R}^n$ and $\eta \in \mathbb{R}^n$, then

$$\begin{aligned} \hat{P}(q, u, q') &= \max_{x \in q} T(q' \mid x, u) \\ &\leq \mathbf{1}_{\bar{q}'(\epsilon+\eta)}^{Im(q,u)} \prod_{i=1}^n P[\|\hat{f}^{(i)} - f^{(i)}\|_{\infty}^{q,u} \leq \epsilon_i] \prod_{i=1}^n P[|w^{(i)}| \leq \eta_i] \\ &\quad + \prod_{i=1}^n (1 - P[\|\hat{f}^{(i)} - f^{(i)}\|_{\infty}^{q,u} \leq \epsilon_i]), \end{aligned} \quad (6)$$

$$\begin{aligned} \check{P}(q, u, q') &= \min_{x \in q} T(q' \mid x, u) \\ &\geq \mathbf{1}_{X \setminus \underline{q}'(\epsilon+\eta)}^{Im(q,u)} \prod_{i=1}^n P[\|\hat{f}^{(i)} - f^{(i)}\|_{\infty}^{q,u} \leq \epsilon_i] \\ &\quad \prod_{i=1}^n P[|w^{(i)}| \leq \eta_i]. \end{aligned} \quad (7)$$

where $\mathbf{1}_A^B$ returns one if $A \cap B \neq \emptyset$ otherwise zero, $\bar{q}'(c)$ and $\underline{q}'(c)$ are the regions obtained by expanding and shrinking (each dimension of) q' by the scalars in vector c , respectively, and $Im(q, u)$ is the image of q under \hat{f} defined as $Im(q, u) = \{\hat{f}(x, u) \mid x \in q\}$.

Further details of Proposition 1 can be found in [22]. Note that the discretization (volume of q and q') plays a significant role in the transition probability bounds as the image of q under the learned dynamics can be very large (especially when it is over-approximated). In addition, the regression errors dictate both the $P[\|\hat{f}^{(i)} - f^{(i)}\|_{\infty}^{q,u} \leq \epsilon_i]$ terms along with affecting the magnitude of the expanded and reduced sets. Hence, the larger these errors, the looser the bounds \check{P} and \hat{P} , leading to larger uncertainty in the abstraction.

Given abstraction \mathcal{I} and specification φ , a control strategy can be generated that is robust against these errors and maximizes the lower bound probability of satisfying φ . This is achieved via a product construction between abstraction \mathcal{I} and DFA \mathcal{A}_φ , which represents φ . The result is the product IMDP (PIMDP) which is defined below.

Definition 7 (Product IMDP). Given an IMDP $\mathcal{I} = (Q, A, \check{P}, \hat{P}, AP, L)$ and DFA $\mathcal{A}_\varphi = (Z, 2^{AP}, \delta, z_0, Z_F)$, the product IMDP (PIMDP) $\mathcal{P} = (S, A, \check{P}^{\mathcal{P}}, \hat{P}^{\mathcal{P}}, S_0, S_F)$, where $S = Q \times Z$, $S_F = Q \times Z_F$,

$$S_0 = \{(q, z_{init}) \mid q \in Q, z_{init} = \delta(z_0, L(q))\},$$

and

$$\begin{aligned} \check{P}^{\mathcal{P}}((q, z), u, (q', z')) &= \begin{cases} \check{P}(q, u, q') & \text{if } z' = \delta(z, L(q)) \\ 0 & \text{otherwise} \end{cases} \\ \hat{P}^{\mathcal{P}}((q, z), u, (q', z')) &= \begin{cases} \hat{P}(q, u, q') & \text{if } z' = \delta(z, L(q)) \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The PIMDP accepting states S_F encapsulate the satisfaction of the specification φ , i.e., a path that reaches S_F satisfies φ . Hence, we formulate the following optimization problem to compute a robust strategy $\pi_{\text{off}} : S \rightarrow A$ that maximizes the lower bound of the probability of reaching a state in S_F :

$$\pi_{\text{off}}(s) = \arg \max_{\pi \in \Pi} \min_{\xi \in \Xi} P[\omega_{\mathcal{P}} \models \varphi \mid \pi, \omega_{\mathcal{P}}(0) = s]. \quad (8)$$

We solve this optimization problem using an interval-value iteration as detailed in [4], which results in a memoryless strategy on \mathcal{P} . The interval-value iteration procedure also returns the lower- and upper-bounds of the probability of satisfaction of φ under π_{off} , i.e.,

$$\underline{p}_{\pi_{\text{off}}}(s) = \min_{\xi \in \Xi} P[\omega_{\mathcal{P}} \models \varphi \mid \pi_{\text{off}}, \omega_{\mathcal{P}}(0) = s], \quad (9)$$

$$\bar{p}_{\pi_{\text{off}}}(s) = \max_{\xi \in \Xi} P[\omega_{\mathcal{P}} \models \varphi \mid \pi_{\text{off}}, \omega_{\mathcal{P}}(0) = s], \quad (10)$$

respectively. These satisfaction intervals let us identify regions where success is more likely (a lower-bound near one), regions where violation is more likely (an upper-bound near zero), and regions with high uncertainty (large intervals). By deploying the system with π_{off} at q , the system is guaranteed to satisfy φ with a probability at least $\underline{p}_{\pi_{\text{off}}}(q, z_{init})$. The goal of the online module is to synthesize controllers at runtime to improve this guarantee towards the optimal probability, which is evidently most helpful for the regions with high uncertainty.

B. Online Module

In constructing the PIMDP \mathcal{P} and strategy π_{off} offline, we are limited by the uncertainty introduced by the finite amount of data in D and by the discretization of X . Our goal is to use the information collected at runtime to improve π_{off} by reducing both sources of uncertainty. In Alg. 1, we present a high-level description of our synergistic control framework that obtains both of these goals.

Online, at each time step, we observe the current state of the system x_k and build a set of (local) GPs to make predictions for the transitions of Process (1) to the states of the abstraction in one-time step starting from x_k (Line 5). Note that this set of GPs ($\mathcal{GP}_{x_k}^A$) are employed exclusively to make local predictions starting from x_k . Hence, to train them, we can simply use only data local to x_k , thus making the training process efficient enough to be employed at runtime. The resulting GPs are then employed to refine the PIMDP \mathcal{P} and get the optimal action u^* using the values computed offline as explained below. Until we reach an accepting state of \mathcal{P} , we continue the process and keep enlarging D with

Algorithm 1 Online control loop

```

1: procedure CONTROLLOOP( $\mathcal{P}, \mathcal{GP}^A, x_0, D$ )
2:    $k \leftarrow 0$ 
3:    $s_0 \leftarrow \text{INITIALPIMDPSTATE}(\mathcal{P}, x_0)$ 
4:   while  $s_k \notin S_F$  do
5:      $\mathcal{GP}_{x_k}^A \leftarrow \text{CREATELOC GP}(\mathcal{GP}^A, x_k, D)$ 
6:      $\mathcal{P} \leftarrow \text{UPDATEPIMDP}(\mathcal{P}, \mathcal{GP}_{x_k}^A)$ 
7:      $u^* \leftarrow \text{GETOPTIMALACTION}(\mathcal{P}, x_k)$ 
8:      $x_{k+1} \leftarrow f(x_k, u^*) + w_k$ 
9:      $s_{k+1} \leftarrow \text{UPDATEPIMPDPSTATE}(\mathcal{P}, x_{k+1})$ 
10:     $D \leftarrow D \cup \{(x_k, u^*, x_{k+1})\}$ 
11:     $k \leftarrow k + 1$ 
12:   end while
13: end procedure

```

the runtime observations. Below, we give further details of how we update \mathcal{P} and we select the optimal action as well as how we train local GPs $\mathcal{GP}_{x_k}^A$.

1) *Optimal Actions*: To choose the optimal action, we first augment the PIMDP with the current state and then choose the optimal next action (Lines 5 and 7 in Alg. 1). Given the current state x_k , we begin by defining a new \mathcal{P} state $s' = (x_k, z)$, where z is the current state of the DFA. By reasoning over a singleton (x_k instead of region q that contains x_k), we eliminate the discretization error over the next step. In particular, for $z' = \delta(z, L(x_k))$ we can compute $\check{P}^{\mathcal{P}}((x_k, z), u, (q', z')) = P[\mathbf{x}_{k+1} \in q' \mid \omega_{\mathbf{x}}(k-1) = x_k, u]$, (11)

and similarly for the upper bound. The bound in (11) can be computed directly using Proposition 1 and is guaranteed to improve compared to $\hat{P}^{\mathcal{P}}((q, z), u, (q', z'))$ as we can compute the image $Im(\{x_k\})$ under the learned dynamics exactly and mitigate this discretization error.

We now focus on how to choose the best action in order to reach a state in S_F . In particular, similar to the goal of the optimization problem in (8), we want to find the action u^* that maximizes (with an abuse of notation)

$$\underline{p}(s', u) = \min_{\xi \in \Xi} \sum_{s \in S^a(s')} \underline{p}_{\pi_{\text{off}}}(s) \xi(\omega_{\mathcal{P}}^{\text{fin}}, u)(s), \quad (12)$$

where $S^a(s') = \{s \in S \mid \hat{P}^{\mathcal{P}}(s', u, s) > 0\}$. Note that $\underline{p}(s', u)$ is the *worst-possible* expected lower-bound probability under action u . We pick the best control as

$$u^* = \arg \max_{u \in U} \underline{p}(s', u). \quad (13)$$

In practice, it is common to have multiple actions that maximize $\underline{p}(s', u)$ e.g. when it is 0 for all actions. We introduce secondary and tertiary metrics grounded encouraging progression, though maximizing $\underline{p}(s', u)$ is always preferred.

a) *Satisfaction Upper Bound*: As a secondary criteria, we consider the upper-bound of the probability of satisfaction from s' , which is defined as

$$\bar{p}(s', u) = \min_{\xi \in \Xi} \sum_{s \in S^a(s')} \bar{p}_{\pi_{\text{off}}}(s) \xi(\omega_{\mathcal{I}}^{\text{fin}}, u)(s). \quad (14)$$

In particular, we select the action $u^* = \arg \max_{u \in U^*} \bar{p}(s', u)$, where U^* is the set of actions returned by (13). Hence, u^* is the one with the greatest lower- and upper-bound probability of satisfying φ . Note that we are again selecting actions against the worst adversary, which is required in order to be robust against uncertainty.

Even with these primary metrics defined, in practice they are not always sufficient to realize good online behavior, and there may be multiple u^* optimizing both criteria. As we are concerned with choosing u^* that best improves the performance for the next step, we introduce two tertiary metrics that encourage one-step progression.

b) *Sink-State Metric*: This metric prefers actions that do not violate φ by using the inherent DFA sink state, from which reaching the accepting state is impossible. For each $s \in S^a(s')$, we check if the DFA transitions to the sink state and prefer actions with a low chance of doing so.

c) *Progression Metric*: We consider a progression metric that defines progress towards reaching S_F . Given s' , we can measure its distance from an accepting state in S_F by using the minimum distance to S_F from each $s \in S^a(s')$, which is easily computed offline. Once online, we choose the action that gives the shortest expected distance to S_F .

2) *Local GP and PIMDP Updates*: At runtime, we collect data that were not available offline. By using this new data, we can refine the regression error and in turn obtain a less uncertain IMDP abstraction. Rather than updating the full GP corresponding to action u , \mathcal{GP}^u , we construct a local GP $\mathcal{GP}_{x_k}^u$ using the ℓ -nearest datapoints in D to x_k . Training then proceeds as for standard GP regression using the same hyperparameters as \mathcal{GP}^u . The hyperparameter ℓ (number of data employed to train our local GP) induces a trade-off between regression accuracy and computation overhead. For more insight on picking ℓ we refer to [31], [32].

Our motivation for using local GP regression stems from computational tractability in the online setting. Performing regression with local data is $O(\ell^3)$ in the worst case. Thus, by choosing ℓ much smaller than the total number of available data, we can have a polynomial speed up that allows for runtime computations. Note that once we have $\mathcal{GP}_{x_k}^u$, we can update the abstraction and product PIMDP by recalculating $\hat{P}(q, u^*, q')$ and $\hat{P}(q, u^*, q')$ for all q in a neighbourhood of x_k . The transitions are then updated in \mathcal{I} and consequently in \mathcal{P} only if the resulting interval is tighter. Once enough transitions are updated in the abstraction, the offline strategy π_{off} can be updated using the refined PIMDP, denoted π_k . This is done again using interval-value iteration, which is polynomial time (cubic) in the size of \mathcal{P} .

C. Correctness

The following theorem guarantees that Alg. 1 monotonically reduces the uncertainty in the system predictions.

Theorem 1. *Let x_k be the state of Process (1) at the k -th step of Alg. 1. Similarly, let π_k be the refinement of π_{off} with k additional datapoints computed according to (8). Then, for any $s \in S$ it holds that*

$$[\underline{p}_{\pi_{k+1}}(s), \bar{p}_{\pi_{k+1}}(s)] \subseteq [\underline{p}_{\pi_k}(s), \bar{p}_{\pi_k}(s)]. \quad (15)$$

The proof of Theorem 1 relies on the fact that, at runtime, we modify only the transition probabilities of \mathcal{P} that are improved by the new data. Hence, the set of adversaries of the updated PIMDP are a subset of those of the PIMDP at the previous time step.

Finally, we remark that in the limit as D becomes dense in X , the regression error in (5) goes to zero with probability one. As a consequence, it is possible to show that as the size of the discretization goes to 0 and as the coverage of D approaches X , Alg. 1 converges to the optimal strategy for Process (1).

V. EVALUATIONS

We evaluated our framework on the nonlinear system

$$\mathbf{x}_{k+1} = \mathbf{x}_k + g(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k$$

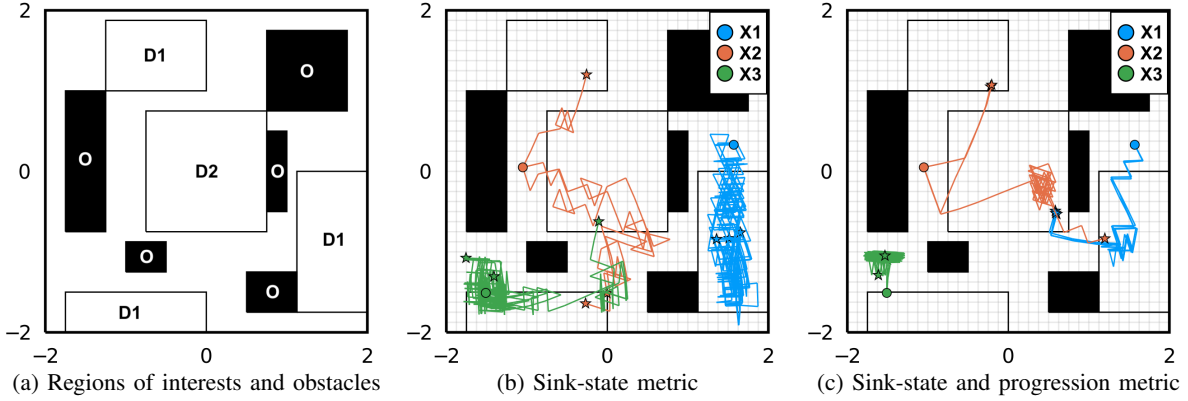


Fig. 1: The state space labels and simulations showing the resultant behavior of the two tertiary metrics used for benchmarking for optimal action selection.

TABLE I: Benchmark results from 500 simulations comparing offline method against the offline-online framework with three types of GP regression methods. The comparison metrics are: the probabilities of violating and satisfying φ , the average computation time of each control step, and the mean number and value of PIMDP transitions updates per run.

x_0	Tertiary	Global GP (static)			Local GP (updates with $\ell = 75$ neighbors)					Local GP (static)		
		P_{Violate}	P_{Satisfy}	Time	P_{Violate}	P_{Satisfy}	Time	# \mathcal{P} -updates	Val. \mathcal{P} -updates	P_{Violate}	P_{Satisfy}	Time
1	Offline	1.0	0.0	—	1.0	0.0	—	—	—	1.0	0.0	—
1	Sink+Prog	0.0	1.0	0.0002	0.0	1.0	0.0135	714.8	1.0	0.0	1.0	0.0003
1	Sink	0.778	0.216	0.0002	0.326	0.322	0.0323	6258	0.9996	0.77	0.22	0.0004
2	Offline	1.00	0.00	—	1.00	0.00	—	—	—	1.00	0.00	—
2	Sink+Prog	0.0	0.76	(0.0001)	0.0	0.996	0.0222	110.8	1.0	0.0	0.808	0.0001
2	Sink	0.0	1.0	0.0002	0.0	1.0	0.0115	521.6	0.9998	0.0	1.0	0.0002
3	Offline	0.348	0.652	—	0.348	0.652	—	—	—	0.348	0.652	—
3	Sink+Prog	0.098	0.0	0.0001	0.102	0.864	0.0130	1814	0.9957	0.074	0.0	0.0004
3	Sink	0.0	0.702	0.0001	0.088	0.6140	0.0120	5890	0.9970	0.00	0.704	0.0004

where $\mathbf{u}_k \in \{u_1, u_2, u_3, u_4\}$, and g is a priori unknown with $g(\mathbf{x}_k, \mathbf{u}_k) =$

$$\begin{cases} [0.25 + 0.05 \sin \mathbf{x}_k^{(2)}, 0.1 \cos \mathbf{x}_k^{(1)}]^T & \text{if } \mathbf{u}_k = u_1 \\ [-0.25 + 0.05 \sin \mathbf{x}_k^{(2)}, 0.1 \cos \mathbf{x}_k^{(1)}]^T & \text{if } \mathbf{u}_k = u_2 \\ [0.1 \cos \mathbf{x}_k^{(2)}, 0.25 + 0.05 \sin \mathbf{x}_k^{(1)}]^T & \text{if } \mathbf{u}_k = u_3 \\ [0.1 \cos \mathbf{x}_k^{(2)}, -0.25 + 0.05 \sin \mathbf{x}_k^{(1)}]^T & \text{if } \mathbf{u}_k = u_4, \end{cases}$$

and \mathbf{w}_k is Gaussian noise drawn from $\mathcal{N}(0, 0.01I)$.

The compact set $X = [-2, 2]^2$ and its regions of interest are shown in Fig. 1a. The atomic propositions are O (obstacle), $D1$ (destination 1) and $D2$ (destination 2). The specification is “to visit destinations 1 and 2 in any order and always avoid the obstacle.” It translates to the LTLf formula

$$\varphi = \mathcal{G}(-O) \wedge \mathcal{F}(D1) \wedge \mathcal{F}(D2).$$

The offline abstraction, PIMDP, and π_{off} were computed in an hour with $m = 200$ uniformly random datapoints for each action. We intentionally used a small D to demonstrate the efficacy of the proposed offline-online control framework.

In the online module, local GPs were created at every step using the $\ell = 75$ nearest datapoints and the same hyperparameters as the offline GPs (which were trained using a zero-mean prior and the squared-exponential kernel). We chose three initial states from the states with $p_{\pi_{\text{off}}} = 0$, and deployed the system with the proposed framework. Each simulation ran until the specification was satisfied, violated or the trajectory length reached a predefined bound of 500.

Note that φ is an unbounded-time property, and hence, the trajectories need to be extended to infinite length though this is impossible in practice. Figure 1 shows multiple simulated trajectories generated with the online framework using the sink-state (“Sink”) metric, and the combined sink-state and progression metric (“Sink-Prog”) (see Sec. IV-B).

Table I contains benchmark results for different GP cases and tertiary metrics. We compare the empirical probabilities of violating and satisfying φ using *static global GPs*, *static local GPs*, and *local GPs with online updates* over 500 simulations. The framework is evaluated using the tertiary metrics for each GP case and compared to the performance of π_{off} , which has $p_{\pi_{\text{off}}}(s) = 0$ for all possible initial states. In other words, π_{off} does not provide any *a priori* satisfaction guarantees.

First, we compare the efficacy of our online framework and the different tertiary metrics against the offline strategy. In all cases, the online framework decreases the number of φ -violating runs, and in many cases increases the number of φ -satisfying runs. Notably for the third initial state, π_{off} results in a 35-65 split in violating and satisfying runs while the online framework sees consistent drops in the percentage of violating runs. When the sum of the fractions of violating and satisfying runs is less than one, the remainder runs terminated in a Zeno-like behavior as observed in Figure 1b. This may be a beneficial behavior as it leads to collection of more data and reduction in the regression error.

The three major columns of Table I compare the outcomes using three types of GPs. In many cases, particularly for

the second initial state, *local GPs with updates* provides a benefit to the number of φ -satisfying runs. The proportion of φ -violating runs is worse for the third initial state, but the most successful runs are realized using the local GPs and the combined tertiary metrics. The same benchmark was run on the global GP with online updates (not included in the table), but all the trials timed out after 24 hours with no result clearly illustrating the advantage of local GPs for online computations.

Constructing the local GPs without updating them or the PIMDP is only marginally slower than using the global GP directly. The additional time to update the local GP and update the PIMDP transitions is significant compared to the static methods, but are still reasonable for online scenarios. Collecting data online and updating the PIMDP abstraction takes time, but it results in improved transition intervals between states. The additional data collected while using the sink-state metric shows many intervals are improved due to long trajectories that neither violate nor satisfy φ .

These evaluations exhibit the advantage of using the proposed offline-online control framework. The framework was able to increase the probability of satisfaction, in some cases from 0% to 100%, and the use of local GPs enables us to overcome computational limitations. Further investigation is needed to address the aforementioned hyperparameter choices, e.g., the value of ℓ , and number of \mathcal{P} states to update.

VI. CONCLUSION

In this work, we presented a synergistic offline-online control framework for stochastic systems with an unknown component via local GP regression. The offline module provides a baseline controller and guarantees for the online module. Online, the controller and its guarantees are iteratively refined as more data is collected. Evaluations illustrate the advantage of the framework over just the offline method. Future directions include investigation into hyperparameter choices for the online method and experiments (deployments) of actual physical platforms with this framework.

REFERENCES

- [1] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.
- [2] C. Belta, B. Yordanov, and E. A. Gol, *Formal methods for discrete-time dynamical systems*, vol. 89. Springer, 2017.
- [3] L. Doyen, G. Frehse, G. J. Pappas, and A. Platzer, "Verification of hybrid systems," in *Handbook of Model Checking*, pp. 1047–1110, Springer, 2018.
- [4] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Transactions on Automatic Control*, vol. 60, pp. 2031–2045, Aug. 2015.
- [5] L. Laurenti, M. Lahijanian, A. Abate, L. Cardelli, and M. Kwiatkowska, "Formal and efficient synthesis for continuous-time linear stochastic hybrid processes," *IEEE Transactions on Automatic Control*, 2020.
- [6] C. Baier, J.-P. Katoen, et al., *Principles of model checking*, vol. 26202649. MIT press Cambridge, 2008.
- [7] P. Jagtap, G. J. Pappas, and M. Zamani, "Control barrier functions for unknown nonlinear systems using gaussian processes," in *Conference on Decision and Control*, pp. 3699–3704, IEEE, 2020.
- [8] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Advances in neural information processing systems*, pp. 908–918, 2017.
- [9] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, "Safe Learning of Regions of Attraction for Uncertain, Nonlinear Systems with Gaussian Processes," *Conference on Decision and Control*, 2016.
- [10] A. Devonport, H. Yin, and M. Arcaç, "Bayesian safe learning and control with sum-of-squares analysis and polynomial kernels," in *Conference on Decision and Control*, pp. 3159–3165, IEEE, 2020.
- [11] M. Maiworm, D. Limon, and R. Findeisen, "Online learning-based model predictive control with gaussian process models and stability guarantees," *Int'l Journal of Robust and Nonlinear Control*, 2021.
- [12] C. E. Rasmussen and C. K. I. Williams, "Gaussian Processes for Machine Learning," 2005.
- [13] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.
- [14] P. Germain, F. Bach, A. Lacoste, and S. Lacoste-Julien, "Pac-bayesian theory meets bayesian inference," in *Advances in Neural Information Processing Systems*, pp. 1884–1892, 2016.
- [15] S. R. Chowdhury and A. Gopalan, "On kernelized multi-armed bandits," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 844–853, JMLR. org, 2017.
- [16] A. Lederer, J. Umlauft, and S. Hirche, "Uniform error bounds for gaussian process regression with application to safe control," in *Adv. in Neural Information Processing Systems*, pp. 657–667, 2019.
- [17] H. Liu, Y.-S. Ong, X. Shen, and J. Cai, "When gaussian process meets big data: A review of scalable gps," *IEEE transactions on neural networks and learning systems*, vol. 31, no. 11, pp. 4405–4423, 2020.
- [18] J. Quinero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate gaussian process regression," *The Journal of Machine Learning Research*, vol. 6, pp. 1939–1959, 2005.
- [19] D. Nguyen-Tuong, J. Peters, and M. Seeger, "Local gaussian process regression for real time online model learning and control," in *Proceedings of the 21st International Conference on Neural Information Processing Systems*, pp. 1193–1200, 2008.
- [20] A. Capone, G. Noske, J. Umlauft, T. Beckers, A. Lederer, and S. Hirche, "Localized active learning of gaussian process state space models," in *Learning for Dynamics and Control*, pp. 490–499, PMLR, 2020.
- [21] A. Lederer, A. J. O. Conejo, K. Maier, W. Xiao, and S. Hirche, "Real-Time Regression with Dividing Local Gaussian Processes," *arXiv*, 2020.
- [22] J. Jackson, L. Laurenti, E. Frew, and M. Lahijanian, "Strategy synthesis for partially-known switched stochastic systems," in *Conf. on Hybrid Systems: Computation and Control*, ACM, May 2021. (to appear).
- [23] G. De Giacomo and M. Y. Vardi, "Linear temporal logic and linear dynamic logic on finite traces," in *IJCAI'13 Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pp. 854–860, Association for Computing Machinery, 2013.
- [24] J. Jackson, L. Laurenti, E. Frew, and M. Lahijanian, "Safety verification of unknown dynamical systems via Gaussian process regression," in *Conference on Decision and Control*, IEEE, Dec. 2020.
- [25] P. Massart, *Concentration inequalities and model selection*, vol. 6. Springer, 2007.
- [26] I. Steinwart, "On the influence of the kernel on the consistency of support vector machines," *Journal of machine learning research*, vol. 2, no. Nov, pp. 67–93, 2001.
- [27] D. P. Bertsekas and S. Shreve, *Stochastic optimal control: the discrete-time case*. 2004.
- [28] A. Abate, F. Redig, and I. Tkachev, "On the effect of perturbation of conditional probabilities in total variation," *Statistics & Probability Letters*, vol. 88, pp. 1–8, 2014.
- [29] A. M. Wells, M. Lahijanian, L. E. Kavragi, and M. Y. Vardi, "LTLf synthesis on probabilistic systems," in *11th International Symposium on Games, Automata, Logics, and Formal Verification (J.-F. Raskin and D. Bresolin, eds.)*, vol. 326 of *Electronic Proceedings in Theoretical Computer Science*, pp. 166–181, Open Publishing Association, 2020.
- [30] R. Givan, S. Leach, and T. Dean, "Bounded-parameter Markov decision processes," *Artificial Intelligence*, vol. 122, no. 1-2, pp. 71–109, 2000.
- [31] D. Nguyen-Tuong, M. Seeger, and J. Peters, *Real-Time Local GP Model Learning*, pp. 193–207. Springer Berlin Heidelberg, 2010.
- [32] S. Das, S. Roy, and R. Sambasivan, "Fast Gaussian process regression for big data," *Big data research*, vol. 14, pp. 12–26, 2018.