2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Data-Driven Reachability Analysis for Gaussian Process State Space Models

Paul Griffioen and Murat Arcak

Abstract— **Gaussian process state space models are becoming common tools for the analysis and design of nonlinear systems with uncertain dynamics. When designing control policies for these systems, safety is an important property to consider. In this paper, we provide safety guarantees by computing finite-horizon forward reachable sets for Gaussian process state space models. We use data-driven reachability analysis to provide exact probability measures for state trajectories of arbitrary length, even when no data samples are available. We investigate two numerical examples to demonstrate the power of this approach, such as providing highly non-convex reachable sets and detecting holes in the reachable set.**

I. Introduction

Gaussian process state space models (GPSSMs) are increasingly used to account for the inherent nonlinearities and unknown dynamics of physical systems [1]–[6]. In contrast to models like recurrent neural networks, GPSSMs are inherently regularized by a prior model, mitigating the tendency to overfit, and are therefore more effective in situations where data is not abundant. GPSSMs also possess useful probabilistic properties in quantifying uncertainty and modeling errors as a distribution over functions, ensuring that the model is not overconfident in regions of the state space where data is scarce [7], [8].

When using GPSSMs for the design and control of dynamical systems, an important property to consider is safety, ensuring that unsafe regions of the state space will be avoided under a particular control policy. An effective way to analyze the safety of a dynamical system in the face of uncertainty is reachability analysis, a set-based method that characterizes all possible evolutions of the state trajectory over a finite time horizon [9]. Many algorithms in reachability analysis use detailed system information to compute an overapproximation or an underapproximation of the reachable set. However, this detailed system information is not available in many important applications such as complex cyber-physical systems where only partial knowledge of the system dynamics is available through simulations or experiments.

Consequently, these applications motivate data-driven reachability analysis, methods that use data obtained from simulations and experiments to estimate reachable sets. In contrast to traditional reachability methods, data-driven reachability analysis typically cannot provide as tight of an approximation to the reachable set due to incomplete or nonexistent prior information about the system dynamics. In contrast, data-driven reachability analysis with GPSSMs allows probabilistic reachable sets to be computed with an exact measure due to the prior information the GPSSM contains about the system dynamics.

While a number of works have investigated data-driven reachability analysis, these works do not use a GPSSM for modeling general nonlinear system dynamics, instead focusing on a more restrictive class of system dynamics such as linear systems [10]–[12], polynomial systems [10], Lipschitz nonlinear systems [10], [11], [13], control-affine systems [14], [15], and mixed monotone systems [16]. Furthermore, these works do not provide probabilistic reachable sets with exact measures, instead providing overapproximations or underapproximations to the reachable sets. A number of works have provided probabilistic reachability guarantees based only on data with no assumptions about the form of the system dynamics, but these guarantees require a minimum number of samples to provide the probabilistic reachable set [17]–[22]. Furthermore, there is a confidence level in the accuracy of the probability measure for the probabilistic reachable set that is less than 100% due to there only being a finite number of samples. These works also only provide probabilistic reachable sets for an arbitrary instant in time, not an entire trajectory.

In contrast to these works, this paper provides probabilistic reachable sets having exact measures, with no overapproximations or underapproximations, for general nonlinear systems that can be modeled by GPSSMs. These probabilistic reachable sets apply to entire trajectories instead of particular instances in time. Since GPSSMs include a prior distribution over the possible state transition functions, this work does not require any samples to be taken in order to obtain the probabilistic guarantees. Furthermore, this work always has perfect confidence in the accuracy of the probabilistic reachable sets being generated.

The remainder of this paper is organized as follows. Section II introduces probabilistic reachability analysis and sets forth the problem statement for the paper. In Section III, we introduce GPSSMs and some of their properties. Section IV sets forth the main result of the paper by introducing probabilistic reachable sets with exact measures for GPSSMs that can be used in safety analysis. Section V demonstrates the power of this reachability analysis for a few highly nonlinear examples, and Section VI concludes the paper.

P. Griffioen and M. Arcak are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA, USA 94720. Email: {griffioen|arcak}@berkeley.edu

This work was supported in part by the Air Force Office of Scientific Research under Grant FA9550-21-1-0288 and the National Science Foundation under Grant CNS-2111688.

II. Preliminaries

A. Probabilistic Reachability Analysis

Before setting forth the problem statement, we first introduce the general area of probabilistic reachability. Consider a discrete-time dynamical system with a state transition function $\Phi(k_f, k_0, x_0, d)$ that maps an initial state $x_{k_0} = x_0 \in \mathbb{R}^n$ at time step k_0 to a unique final state at time step k_f under a disturbance $d : [k_0, k_f] \to \mathbb{R}^n$. For example, when the system's state dynamics $x_{k+1} = g(k, x_k, d_k)$ are known and have unique solutions on the interval $[k_0, k_f]$, then $\Phi(k_f, k_0, x_0, d)$ is just x_{k_f}, where x is the solution of the state dynamics with initial condition $x_{k_0} = x_0$. In addition to representing exogenous disturbances, the disturbance signal d may account for functional uncertainty in the system's state dynamics and deviations from a nominal control law.

In forward reachability analysis, we are interested in computing the set where all possible evolutions of the state trajectory lie during the time period k_0 to k_f given a set $X_0 \subseteq \mathbb{R}^n$ in which the initial state lies, a set D in which the disturbance lies, and a time range $[k_0, k_f]$. The *forward reachable set* is defined as the set of all states to which the system can transition in the time range $[k_0, k_f]$ with initial states in X_0 and disturbances in D, given by

$$R_{k_0:k_f} = \bigcup_{k=k_0}^{k_f} \{\Phi(k, k_0, x_0, d) | x_0 \in X_0,\ d \in D\}. \quad (1)$$

In many cases, the disturbance d cannot be limited to a finite set D and is better described as a random variable following a probability distribution $d \sim \mathcal{Q}^d$. In addition, some information about the system's state dynamics $x_{k+1} = g(k, x_k, d_k)$ may be known, but the functional uncertainty in these dynamics may be modeled with a probability distribution $g \sim \mathcal{Q}^g$, as in the case with GPSSMs. Modeling the disturbance and the unknown portion of the state dynamics as random variables that follow probability distributions $\mathcal{Q}^d$ and $\mathcal{Q}^g$, respectively, results in $\Phi(k_f, k_0, x_0, d)$ being a random variable over the forward reachable set, whose probability measure we denote as φ.

If we take samples of the random variables $d \sim \mathcal{Q}^d$ and $g \sim \mathcal{Q}^g$, then the vector $\Phi(k_f, k_0, x_0, d)$ lies in $X \subseteq \mathbb{R}^n$ with probability $\varphi(X)$, and the support of $\Phi(k_f, k_0, x_0, d)$ is the reachable set. We can then view $\varphi(X)$ as a measure of probabilistic accuracy. If a set $X \subseteq \mathbb{R}^n$ has a greater measure $\varphi(X)$ than a set $Y \subseteq \mathbb{R}^n$, then X is a more accurate approximation of the reachable set than Y since it "misses" less of the probability mass than Y does. In probabilistic forward reachability, the goal is to find a tight reachable set such that $\varphi(R_{k_0:k_f})$ is close to 1. More formally, this can be described as follows.

B. Problem Statement

Problem 1: Given the state transition function $\Phi(k_f, k_0, x_0, d)$, time range $[k_0, k_f]$, initial set X_0, disturbance distribution $\mathcal{Q}^d$, state dynamics function distribution $\mathcal{Q}^g$, and a probability level p, compute a set $R_{k_0:k_f}$ such that $\varphi(R_{k_0:k_f}) = p$.

Note that we would like to compute a set $R_{k_0:k_f}$ in which the state trajectory $x_{k_0:k_f}$ lies with an exact probability measure of p, not an over-approximation or an under-approximation. In solving this problem, therefore, there will be no drawbacks in terms of approximation or conservativeness.

III. Gaussian Process State Space Models

We consider a discrete time system model with a continuous-valued state, where uncertainty in the model is captured by n independent Gaussian processes (GPs), given by

$$x_{k+1} = g(x_k, u_k) + w_k, \quad (2)$$

where $x_k \in \mathbb{R}^n$ represents the system state at time step k, $u_k \in \mathbb{R}^m$ is the control input vector, $w_k \sim \mathcal{N}(0, Q)$ with $Q \triangleq \text{Diag}(\sigma_1^2, \cdots, \sigma_n^2)$ is independent and identically distributed (i.i.d.) GP noise,

$$g(x_k, u_k) = \begin{bmatrix} g_1(x_k, u_k) \\ \vdots \\ g_n(x_k, u_k) \end{bmatrix}, \quad (3)$$

and

$$g_i(x_k, u_k) \sim \mathcal{GP}(m_i(\hat{x}_k), k_i(\hat{x}_k, \hat{x}_k')),\ \hat{x}_k \triangleq \begin{bmatrix} x_k \\ u_k \end{bmatrix}, \quad (4)$$

is a GP specified by its mean function $m_i(\hat{x}_k)$: $\mathbb{R}^{n+m} \to \mathbb{R}$ and covariance function $k_i(\hat{x}_k, \hat{x}_k')$: $\mathbb{R}^{n+m} \times \mathbb{R}^{n+m} \to \mathbb{R}$, given by

$$m_i(\hat{x}_k) = \mathbb{E}[g_i(\hat{x}_k)], \quad (5)$$

$$k_i(\hat{x}_k, \hat{x}_k') = \mathbb{E}[(g_i(\hat{x}_k) - m_i(\hat{x}_k))(g_i(\hat{x}_k') - m_i(\hat{x}_k'))]. \quad (6)$$

A GP is a distribution over functions, assigning a joint Gaussian distribution to any finite subset of the state and control input space [23]. The covariance function of a GP is also called the kernel function of the process, which determines the class of functions over which the distribution is defined.

We assume that N measurements of the state are taken, either through recorded trajectory data or simply by sampling the state transition function at various points in the state and control input space. This training data set, composed of N data pairs, is given by $\mathcal{D} \triangleq \{\{\bar{x}_j, \bar{u}_j\}, \bar{x}_j^+\}_{j=1}^N$, where

$$\bar{x}_j^+ = g(\bar{x}_j, \bar{u}_j) + w_j, \quad w_j \sim \mathcal{N}(0, Q). \quad (7)$$

The training data can be used to determine the values of the hyperparameters for the mean function and the covariance function by optimizing the marginal likelihood. Given input training data $\{\bar{x}_j, \bar{u}_j\}_{j=1}^N$ and output training data $\{\bar{x}_j^+\}_{j=1}^N$, $g(x_k, u_k)$ conditioned on x_k, u_k, and $\mathcal{D}$ follows a Gaussian distribution, given by

$$g(x_k, u_k) | \{x_k, u_k, \mathcal{D}\} \sim \mathcal{N}(\mu(\hat{x}_k), \Sigma(\hat{x}_k)), \quad (8)$$

$$\mu(\hat{x}_k) \triangleq \begin{bmatrix} m_1(\hat{x}_k) + \bar{k}_1(\hat{x}_k)^T (K_1 + \sigma_1^2 I_N)^{-1} (y_1 - \bar{y}_1) \\ \vdots \\ m_n(\hat{x}_k) + \bar{k}_n(\hat{x}_k)^T (K_n + \sigma_n^2 I_N)^{-1} (y_n - \bar{y}_n) \end{bmatrix},$$

$$\Sigma(\hat{x}_k) \triangleq \begin{bmatrix} \xi_1(\hat{x}_k) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \xi_n(\hat{x}_k) \end{bmatrix},$$

$$\xi_i(\hat{x}_k) \triangleq k_i(\hat{x}_k, \hat{x}_k) - \bar{k}_i(\hat{x}_k)^T (K_i + \sigma_i^2 I_N)^{-1} \bar{k}_i(\hat{x}_k),$$

$$K_i \triangleq \begin{bmatrix} k_i(\hat{\bar{x}}_1, \hat{\bar{x}}_1) & \cdots & k_i(\hat{\bar{x}}_1, \hat{\bar{x}}_N) \\ \vdots & \ddots & \vdots \\ k_i(\hat{\bar{x}}_N, \hat{\bar{x}}_1) & \cdots & k_i(\hat{\bar{x}}_N, \hat{\bar{x}}_N) \end{bmatrix}, \ \hat{\bar{x}}_j \triangleq \begin{bmatrix} \bar{x}_j \\ \bar{u}_j \end{bmatrix},$$

$$\bar{k}_i(\hat{x}_k) \triangleq \begin{bmatrix} k_i(\hat{\bar{x}}_1, \hat{x}_k) \\ \vdots \\ k_i(\hat{\bar{x}}_N, \hat{x}_k) \end{bmatrix}, \ y_i \triangleq \begin{bmatrix} \bar{x}_1^+(i) \\ \vdots \\ \bar{x}_N^+(i) \end{bmatrix}, \ \bar{y}_i \triangleq \begin{bmatrix} m_i(\hat{\bar{x}}_1) \\ \vdots \\ m_i(\hat{\bar{x}}_N) \end{bmatrix},$$

where $\bar{x}_j^+(i)$ denotes the i^{th} dimension of $\bar{x}_j^+$.

The conditional Gaussian distribution that $g(x_k, u_k)$ follows will be used in the next section to develop exact probability measures of the forward reachable set for the system in (2).

IV. Data-Driven Reachability Analysis

A. χ^2 Distribution

Before introducing probabilistic reachable sets for the system in (2), we first derive a preliminary result, showing in Lemma 1 how a quadratic function of $x_{k+1} - \mu(\hat{x}_k)$ follows a χ^2 distribution when conditioned on $\hat{x}_k$ and $\mathcal{D}$.

Lemma 1: Let $\chi^2(n)$ represent the χ^2 distribution with n degrees of freedom, and let

$$h(\hat{x}_k, x_{k+1}) = (x_{k+1} - \mu(\hat{x}_k))^T (\Sigma(\hat{x}_k) + Q)^{-1} (x_{k+1} - \mu(\hat{x}_k)). \tag{9}$$

Then

$$h(\hat{x}_k, x_{k+1}) | \{\hat{x}_k, \mathcal{D}\} \sim \chi^2(n). \tag{10}$$

Proof: According to (2), (8), and the fact that $w_k \sim \mathcal{N}(0, Q)$,

$$x_{k+1} | \{x_k, u_k, \mathcal{D}\} \sim \mathcal{N}(\mu(\hat{x}_k), \Sigma(\hat{x}_k) + Q).$$

The result then follows directly from the properties of the χ^2 distribution. ∎

B. Probabilistic Forward Reachable Set

We now formally introduce the notion of a probabilistic forward reachable set for trajectories of arbitrary length, and Theorem 1 derives this set for GPSSMs, ensuring that state trajectories lie within this set with probability p.

Definition 1: A *(p, T)-reachable set* of probability level p and trajectory length T is a set $\mathcal{X}_T(p) \subseteq \mathbb{R}^{nT}$ such that

$$\Pr\left(x_{1:T} \in \mathcal{X}_T(p) | x_0 \in X_0\right) = p. \tag{11}$$

In other words, a (p, T)-reachable set is one in which the trajectory $x_{1:T}$ lies with probability p, given that the initial state lies in $X_0 \subseteq \mathbb{R}^n$.

Note that in Definition 1, $\mathcal{X}_T(p)$ differs from a reachable tube since it is not necessarily decomposable into a Cartesian product of T subsets of $\mathbb{R}^n$. This is illustrated in Figure 1 for the system in (2). Also note that the initial set X_0 is not restricted in any way and can be a highly nonconvex disjoint set, enabling the result in Theorem 1 to be applied to a wide variety of systems.

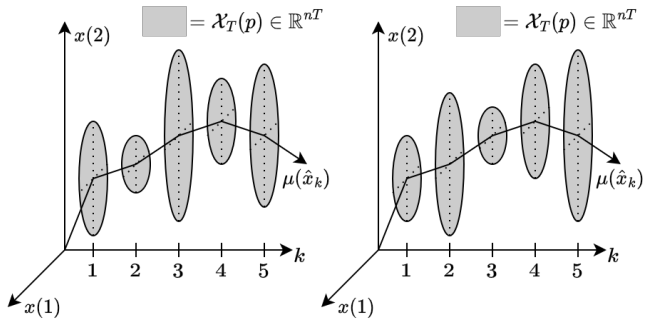

Fig. 1. The probabilistic reachable set $\mathcal{X}_T(p) \in \mathbb{R}^{nT}$ for the system in (2) is defined by a bound on the cumulative deviation from the mean trajectory $\mu(\hat{x}_k)$. In other words, $\mathcal{X}_T(p)$ is defined as a set of ellipsoids centered on the mean trajectory whose sum of areas is bounded by the inverse CDF of the χ^2 distribution. Two possibilities, each of which have the same sum of areas, are shown here for an $n = 2$ dimensional system.

Theorem 1: The (p, T)-reachable set for the system in (2) is given by

$$\mathcal{X}_T(p) = \left\{ x_{1:T} \middle| \sum_{k=0}^{T-1} h(\hat{x}_k, x_{k+1}) \leq F_{\chi^2}^{-1}(p, nT), \ x_0 \in X_0 \right\}, \tag{12}$$

where $h(\hat{x}_k, x_{k+1})$ is given in (9) and $F_{\chi^2}^{-1}(p, nT)$ is the value of the inverse cumulative distribution function (CDF) for the χ^2 distribution evaluated at probability p with nT degrees of freedom.

Proof: According to Lemma 1,

$$\Pr\left(h(\hat{x}_k, x_{k+1}) \leq \eta | \hat{x}_k, \mathcal{D}\right) = F_{\chi^2}(\eta, n), \tag{13}$$

where $F_{\chi^2}(\eta, n)$ is the value of the CDF for the χ^2 distribution evaluated at η with n degrees of freedom. Applying the law of total probability with (13) yields

$$\begin{aligned} &\Pr(h(\hat{x}_k, x_{k+1}) \leq \eta | \mathcal{D}) \\ &= \int_{\hat{x}_k \in \mathbb{R}^{n+m}} \Pr\left(h(\hat{x}_k, x_{k+1}) \leq \eta | \hat{x}_k, \mathcal{D}\right) f(\hat{x}_k | \mathcal{D}) d\hat{x}_k \\ &= F_{\chi^2}(\eta, n) \int_{\hat{x}_k \in \mathbb{R}^{n+m}} f(\hat{x}_k | \mathcal{D}) d\hat{x}_k = F_{\chi^2}(\eta, n), \end{aligned} \tag{14}$$

where $f(\hat{x}_k | \mathcal{D})$ represents the probability density function (PDF) of $\hat{x}_k$ given $\mathcal{D}$. Note that the CDF in (13) conditioned on $\hat{x}_k$ and the CDF in (14) not conditioned on $\hat{x}_k$ are equal $\forall \eta$, implying that $h(\hat{x}_k, x_{k+1}) | \mathcal{D}$ follows the same distribution as $h(\hat{x}_k, x_{k+1}) | \hat{x}_k, \mathcal{D}$. Together with Lemma 1, this implies that

$$h(\hat{x}_k, x_{k+1}) | \mathcal{D} \sim \chi^2(n).$$

According to the properties of the χ^2 distribution, this in turn implies that

$$\sum_{k=0}^{T-1} h(\hat{x}_k, x_{k+1}) | \mathcal{D} \sim \chi^2(nT).$$

The result then follows directly from the properties of the χ^2 distribution and the fact that $x_0 \in X_0$:

$$\Pr\left(\sum_{k=0}^{T-1} h(\hat{x}_k, x_{k+1}) \leq F^{-1}_{\chi^2}(p, nT) \middle| \mathcal{D}\right) = p.$$

■

Note that in [17]–[22], a minimum number of samples are needed in order to obtain probabilistic reachable sets like the one given in Theorem 1. However, Theorem 1 does not require any samples for its results to still hold because GPSSMs provide some prior knowledge about the system dynamics, namely a prior distribution over the possible state transition functions. This also allows us to have 100% confidence in the probabilistic statement made about the reachable set, whereas an infinite number of samples would be required to have perfect confidence in cases where no prior knowledge about the system dynamics is available.

As seen from the quadratic form of $h(\hat{x}_k, x_{k+1})$, Theorem 1 states that the state trajectory deviates from the mean trajectory $\mu(\hat{x}_k)$ with a certain probability for a particular time horizon. More specifically, Theorem 1 states that the cumulative deviation from the mean trajectory over a particular time horizon is bounded by the inverse CDF of the χ^2 distribution, as illustrated in Figure 1. This inverse CDF is itself a function of the time horizon T and is strictly increasing in T [24]. Consequently, for a given probability level, the cumulative deviation from the mean trajectory increases as the time horizon T increases. This can also be intuited since the state trajectory will have some deviation from the mean trajectory at each time step.

Corollary 1 states an implication of this fact, showing how the size of the probabilistic reachable set for any subset of the trajectory increases in T, since a larger time horizon means larger cumulative deviation from the mean trajectory.

Corollary 1: Let $\mathcal{X}^k_T(p) \subseteq \mathbb{R}^{nk}$ represent the projection of $\mathcal{X}_T(p)$ onto $x_{1:k}$, where $k \leq T$. Then

$$\mathcal{X}^k_T(p) \subseteq \mathcal{X}^k_{T'}(p) \ \forall T \leq T'. \tag{15}$$

Proof: The projection $\mathcal{X}^k_T(p)$ is given by

$$\mathcal{X}^k_T(p) \triangleq \left\{ x_{1:k} \middle| \sum_{j=0}^{k-1} h(\hat{x}_j, x_{j+1}) \leq F^{-1}_{\chi^2}(p, nT), \ x_0 \in X_0 \right\}.$$

$F^{-1}_{\chi^2}(p, nT)$ is strictly increasing in T [24], thus (15) follows. ■

Note that $\mathcal{X}^k_T(p)$ is a projection onto a subset of the trajectory, not a projection onto dimensions of the state space since the subscript in $x_{1:k}$ is a time index, not a dimensional index.

Corollary 2 also states a property that arises from the fact that the cumulative deviation from the mean trajectory increases with increases in the time horizon. Corollary 2 states that the probability of any subset of the trajectory lying within the probabilistic reachable set will be greater than or equal to the probability of the full trajectory lying within the probabilistic reachable set.

Corollary 2: The projection $\mathcal{X}^k_T(p)$ has the following property:

$$\Pr\left(x_{1:k} \in \mathcal{X}^k_T(p) \middle| x_0 \in X_0\right) \geq p. \tag{16}$$

Proof: According to Corollary 1,

$$\mathcal{X}^k_T(p) \supseteq \mathcal{X}^k_k(p) = \mathcal{X}_k(p) \ \forall k \leq T,$$

implying that

$$\Pr\left(x_{1:k} \in \mathcal{X}^k_T(p) \middle| x_0 \in X_0\right) \geq \Pr(x_{1:k} \in \mathcal{X}_k(p) | x_0 \in X_0). \tag{17}$$

From Definition 1 and Theorem 1, we know that

$$\Pr\left(x_{1:k} \in \mathcal{X}_k(p) | x_0 \in X_0\right) = p. \tag{18}$$

Combining (17) with (18) yields the desired result. ■

The results provided in Corollaries 1 and 2 are therefore useful in analyzing subsets of the overall trajectory.

C. Representing Probabilistic Forward Reachable Sets

Algorithm 1 shows how to use the result obtained in Theorem 1 to graphically depict probabilistic forward reachable sets for GPSSMs. These sets are graphically depicted by plotting M sample trajectories that satisfy the inequality in (12). For each sample trajectory, Algorithm 1 first draws a sample of the state at random from the initial set (line 4). Then the algorithm iterates through each time step in the trajectory, using the state feedback control policy to compute the control input (line 6). After that, a sample of the state at the next time step is drawn at random from $\mathcal{E}^k_T(p)$, an ellipsoid defined by (12) to ensure that the sample does not lie outside the boundary of $\mathcal{X}_T(p)$. After M sample trajectories have been generated, they can be plotted to graphically depict an approximation of $\mathcal{X}_T(p)$.

Algorithm 1 Data-Driven (p, T)-Reachable Set Representation

1: Initialize M, the number of sample trajectories
2: Initialize X_0, the initial set
3: **for** $i = 1 : M$
4: Draw a sample x^i_0 from the initial set X_0
5: **for** $k = 1 : T$
6: Compute u^i_{k-1} from the control policy
7: $\hat{x}^i_{k-1} = \begin{bmatrix} x^i_{k-1} \\ u^i_{k-1} \end{bmatrix}$
8: Draw a sample x^i_k from the ellipsoid $\mathcal{E}^k_T(p) \triangleq$

$$\left\{ x^i_k \middle| h(\hat{x}^i_{k-1}, x^i_k) \leq F^{-1}_{\chi^2}(p, nT) - \sum_{j=0}^{k-2} h(\hat{x}^i_j, x^i_{j+1}) \right\}$$

9: **end for**
10: **end for**
11: Approximate the (p, T)-reachable set with the sample trajectories $\{x^1_{1:k}, \cdots, x^M_{1:k}\}$

V. Examples

In the following two examples, we demonstrate how Theorem 1 accurately quantifies (p, T)-reachable sets, even with systems that are highly nonlinear and result in quite nonconvex forward reachable sets. We also validate the results from Corollaries 1 and 2, showing how the cumulative deviation from the mean trajectory increases as the time horizon increases.

A. Chaotic Nonlinear Map

We first examine the Tinkerbell map, a nonlinear map that exhibits chaotic behavior [25]. The Tinkerbell map is used as the mean function in the GPSSM, given by

$$m_1(\hat{x}_k) = x_k(1)^2 - x_k(2)^2 + ax_k(1) + bx_k(2), \quad (19)$$
$$m_2(\hat{x}_k) = 2x_k(1)x_k(2) + cx_k(1) + rx_k(2), \quad (20)$$

where $x_k(i)$ denotes the i^{th} dimension of x_k and the system parameters are chosen to be $a = 0.9$, $b = -0.6013$, $c = 2$, and $r = 0.5$. We choose a squared exponential kernel function and train the hyperparameters of the kernel function by optimizing the marginal likelihood with training data recorded from a trajectory of the Tinkerbell map. After training, the observation noise covariance is given by $Q = \text{Diag}(1.2211, 2.6842) \times 10^{-5}$. The initial set X_0 is defined such that $x_0(1) \in [-0.85, -0.6]$ and $x_0(2) \in [-0.75, -0.5]$.

Figure 2a depicts the (p, T)-reachable set as defined in (12) for the Tinkerbell map with probability $p = 99\%$ and time horizon $T = 100$ time steps. This approximate representation of the (p, T)-reachable set was generated according to Algorithm 1 with $M = 1000$ sample trajectories. As can be seen from the figure, the probabilistic reachable set is highly nonconvex, is not simply connected, and contains numerous holes. This demonstrates the power of Theorem 1 in characterizing probabilistic reachable sets for general nonlinear systems.

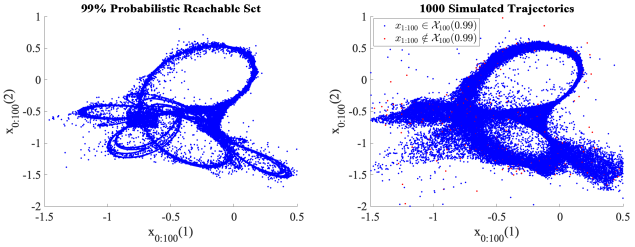

(a) (b)

Fig. 2. For the Tinkerbell map with a time horizon of $T = 100$ time steps, (a) 99% probabilistic reachable set and (b) 1000 simulated trajectories, where blue trajectories lie inside and red trajectories lie outside the 99% probabilistic reachable set.

To verify the accuracy of our approach, Figure 2b plots 1000 trajectories for the Tinkerbell map over a time horizon of $T = 100$ time steps simulated from the GPSSM dynamics of the system in (2). Trajectories that lie within the (p, T)-reachable set of probability level $p = 99\%$ are plotted in blue, and those that do not lie within this set are plotted in red. As can be seen in the figure, 98.2% of the simulated trajectories lie within the 99% probabilistic reachable set, demonstrating the accuracy of the probabilistic reachable set derived in Theorem 1. Note that many red trajectories contain points that lie near the mean trajectory in the middle of the blue regions, but they do not lie within the 99% probabilistic reachable set due to another point in the trajectory deviating far away from the mean trajectory.

Figure 3 validates the results from Corollaries 1 and 2, showing that for a 99% probabilistic reachable set, the cumulative deviation from the mean trajectory increases as the time horizon increases.

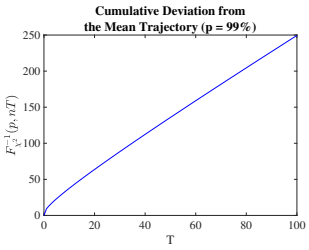

Fig. 3. For the 99% probabilistic reachable set of the Tinkerbell map, cumulative deviation from the mean trajectory as a function of the time horizon T.

B. Traffic Model

We next examine a traffic model used as a reachability benchmark in [26]. This traffic model investigates the density of traffic on a highway using a discretization of the cell transmission model that divides the highway into n equal segments. The spatially discretized model has n states, where $x_k(i)$ represents the number of vehicles on segment i at time step k. Traffic enters through segment 1 and flows through each successive segment before leaving through segment n, where some traffic leaves the highway through exit ramps located at the end of each segment. The traffic model dynamics are used as the mean function in the GPSSM, given by

$$m_1(\hat{x}_k) = x_k(1) - \min\left(c, vx_k(1), \frac{q}{\beta}(z - x_k(2))\right) + u_k, \quad (21)$$

$$m_i(\hat{x}_k) = x_k(i) + \min(\beta c, \beta v x_k(i-1), q(z - x_k(i))) - \min\left(c, vx_k(i), \frac{q}{\beta}(z - x_k(i+1))\right), \quad i = 2, \cdots, n-1, \quad (22)$$

$$m_n(\hat{x}_k) = x_k(n) + \min(\beta c, \beta v x_k(n-1), q(z - x_k(n))) - \min(c, vx_k(n)), \quad (23)$$

where c is the capacity, v is the free-flow speed, q is the congestion-wave speed, and z is the jam occupancy of the segment. The parameter β is defined such that a fraction $1 - \beta$ of vehicles on a segment exit the highway at the end of that segment. The input u_k represents the number of vehicles arriving on segment 1 from upstream. In this example, we examine a 10-mile portion of the highway discretized into $n = 10$ 1-mile segments with $c = 40$ vehicles/time step, $v = 0.5$ segments/time step, $q = 1/6$ segments/time step, $z = 320$ vehicles, and $\beta = 0.75$. We model u_k as random input sampled uniformly between 40 and 60 vehicles at each time step. We choose a squared exponential kernel function and train the hyperparameters of the kernel function by optimizing the marginal likelihood with training data recorded from a trajectory of the traffic model. After training, the observation noise covariance is given by $Q = \text{Diag}(0.6907, 1.0889, 0.0854, 0.4202, 0.5491, 1.7253, 0.1427, 0.2729, 0.9272, 0.7745)$. The initial set X_0 is defined such that $x_0(i) \in [100, 150]$ $\forall i \in \{1, 3, 5, 7, 9\}$ and $x_0(i) \in [200, 250]$ $\forall i \in \{2, 4, 6, 8, 10\}$.

Figure 4a depicts the (p, T)-reachable set as defined in (12) for segments 9 and 10 of the highway with probability $p = 99\%$ and a time horizon of 10 minutes ($T = 20$ time steps, with each time step being 0.5 minutes). This approximate representation of the (p, T)-reachable set was generated according to Algorithm 1 with $M = 1000$ sample trajectories. As can be seen from the figure, the probabilistic reachable set is not convex and is useful in characterizing probabilistic reachable sets for large-scale nonlinear systems in the real world.

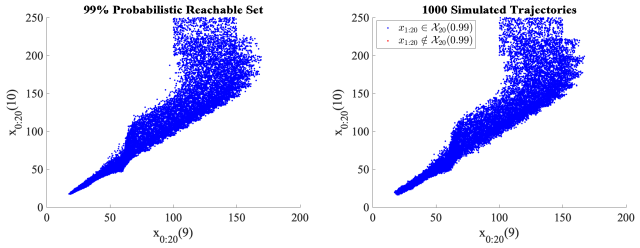

(a) (b)

Fig. 4. For the traffic on segments 9 and 10 of the highway with a time horizon of 10 minutes ($T = 20$ time steps), (a) 99% probabilistic reachable set and (b) 1000 simulated trajectories, where blue trajectories lie inside and red trajectories lie outside the 99% probabilistic reachable set.

To verify the accuracy of our approach, Figure 4b plots 1000 trajectories for segments 9 and 10 of the highway over a time horizon of 10 minutes simulated from the GPSSM dynamics of the system in (2). Trajectories that lie within the (p, T)-reachable set of probability level $p = 99\%$ are plotted in blue, and those that do not lie within this set are plotted in red. As can be seen in the figure, 100% of the simulated trajectories lie within the 99% probabilistic reachable set, demonstrating the accuracy of the probabilistic reachable set derived in Theorem 1.

VI. Conclusion

This paper investigates data-driven probabilistic reachable sets for GPSSMs. We provide exact probability measures for state trajectories of arbitrary length, even when no data samples are available, and we also investigate a few properties of subsets of the overall trajectory. We demonstrate the accuracy of these probability measures and their power in characterizing highly nonlinear reachable sets with the examples of a chaotic nonlinear map and a real-world traffic model. Future work includes providing exact probability measures for state trajectories when the state is not directly observable as well as when both the GPs and the process noise are not independent. Future work also includes investigating the relationship between the amount of training data and the size of the probabilistic reachable set.

References

[1] R. Frigola, Y. Chen, and C. E. Rasmussen, "Variational Gaussian Process State-Space Models," *Advances in Neural Information Processing Systems*, vol. 27, 2014.

[2] R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen, "Bayesian Inference and Learning in Gaussian Process State-Space Models with Particle MCMC," *Advances in Neural Information Processing Systems*, vol. 26, 2013.

[3] R. Turner, M. Deisenroth, and C. Rasmussen, "State-Space Inference and Learning with Gaussian Processes," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 868–875.

[4] S. Eleftheriadis, T. Nicholson, M. Deisenroth, and J. Hensman, "Identification of Gaussian Process State Space Models," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[5] A. Svensson, A. Solin, S. Särkkä, and T. Schön, "Computationally Efficient Bayesian Learning of Gaussian Process State Space Models," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 213–221.

[6] J. Umlauft, A. Lederer, and S. Hirche, "Learning Stable Gaussian Process State Space Models," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 1499–1504.

[7] J. Schneider, "Exploiting Model Uncertainty Estimates for Safe Dynamic Control Learning," *Advances in Neural Information Processing Systems*, vol. 9, 1996.

[8] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian Processes for Data-Efficient Learning in Robotics and Control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 408–423, 2013.

[9] M. Althoff, "Reachability Analysis and its Application to the Safety Assessment of Autonomous Cars," Ph.D. dissertation, Technische Universität München, 2010.

[10] A. Alanwar, A. Koch, F. Allgöwer, and K. H. Johansson, "Data-Driven Reachability Analysis from Noisy Data," *IEEE Transactions on Automatic Control*, 2023.

[11] ——, "Data-Driven Reachability Analysis Using Matrix Zonotopes," in *Learning for Dynamics and Control*. PMLR, 2021, pp. 163–175.

[12] A. Alanwar, Y. Stürz, and K. H. Johansson, "Robust Data-Driven Predictive Control Using Reachability Analysis," *European Journal of Control*, vol. 68, p. 100666, 2022.

[13] A. Alanwar, F. J. Jiang, M. Sharifi, D. V. Dimarogonas, and K. H. Johansson, "Enhancing Data-Driven Reachability Analysis using Temporal Logic Side Information," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6793–6799.

[14] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, and C. J. Tomlin, "Reachability-Based Safe Learning with Gaussian Processes," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 1424–1431.

[15] S. Gorantla, J. Chatrola, J. Bhagiya, A. Saoud, and P. Jagtap, "Funnel-Based Reachability Control of Unknown Nonlinear Systems Using Gaussian Processes," *arXiv preprint arXiv:2209.14015*, 2022.

[16] M. E. Cao, M. Bloch, and S. Coogan, "Estimating High Probability Reachable Sets Using Gaussian Processes," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 3881–3886.

[17] A. Devonport and M. Arcak, "Estimating Reachable Sets with Scenario Optimization," in *Learning for dynamics and control*. PMLR, 2020, pp. 75–84.

[18] ——, "Data-Driven Reachable Set Computation Using Adaptive Gaussian Process Classification and Monte Carlo Methods," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 2629–2634.

[19] A. Devonport, F. Yang, L. El Ghaoui, and M. Arcak, "Data-Driven Reachability Analysis with Christoffel Functions," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 5067–5072.

[20] A. Devonport, F. Yang, L. E. Ghaoui, and M. Arcak, "Data-Driven Reachability Analysis and Support Set Estimation with Christoffel Functions," *arXiv preprint arXiv:2112.09995*, 2021.

[21] A. J. Thorpe, K. R. Ortiz, and M. M. Oishi, "State-Based Confidence Bounds for Data-Driven Stochastic Reachability Using Hilbert Space Embeddings," *Automatica*, vol. 138, p. 110146, 2022.

[22] ——, "SReachTools Kernel Module: Data-Driven Stochastic Reachability Using Hilbert Space Embeddings of Distributions," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 5073–5079.

[23] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. MIT Press Cambridge, MA, 2006, vol. 2, no. 3.

[24] Y. Sun, Á. Baricz, and S. Zhou, "On the Monotonicity, Log-Concavity, and Tight Bounds of the Generalized Marcum and Nuttall Q-Functions," *IEEE Transactions on Information Theory*, vol. 56, no. 3, pp. 1166–1186, 2010.

[25] S. Yuan, T. Jiang, and Z. Jing, "Bifurcation and Chaos in the Tinkerbell Map," *International Journal of Bifurcation and Chaos*, vol. 21, no. 11, pp. 3137–3156, 2011.

[26] S. Coogan and M. Arcak, "A Benchmark Problem in Transportation Networks," *arXiv preprint arXiv:1803.00367*, 2018.