Proceedings of Machine Learning Research vol 211:1–11, 2023 5th Annual Conference on Learning for Dynamics and Control

Probabilistic Invariance for Gaussian Process State Space Models

Paul Griffioen GRIFFIOEN@BERKELEY.EDU
Alex Devonport ALEX_DEVONPORT@BERKELEY.EDU
Murat Arcak ARCAK@BERKELEY.EDU
Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

Editors: N. Matni, M. Morari, G. J. Pappas

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 for Gaussian process state space models in the form of probabilistic invariant sets, where the state trajectory is guaranteed to lie within an invariant set for all time with a particular probability. We provide a sufficient condition in the form of a linear matrix inequality to evaluate the probabilistic invariance of the system, and we demonstrate our contributions with an illustrative example.

Keywords: Gaussian processes, probabilistic invariance, safety

1. Introduction

Gaussian process state space models (GPSSMs) are increasingly used to account for the inherent nonlinearities and unknown dynamics of physical systems (Frigola et al., 2014, 2013; Turner et al., 2010; Eleftheriadis et al., 2017; Svensson et al., 2016; Umlauft and Hirche, 2020). 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 (Schneider, 1996; Deisenroth et al., 2013).

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. Characterizing the effect of uncertainties on the safety of the system through invariant sets is an important problem in system analysis (Blanchini, 1999). Robust positively invariant sets have been used to describe a region of the state space in which the state is guaranteed to lie under an unknown but bounded disturbance (Brockman and Corless, 1995, 1998; Alessandri et al., 2004). Similarly, probabilistic invariant sets have been introduced in Kofman et al. (2012, 2016) for linear systems to describe a set in which the state trajectory lies at all times with a certain probability. The relationship between these two types of sets for linear systems has been examined in Hewing et al. (2018).

While Kofman et al. (2012, 2016); Hewing et al. (2018) examine probabilistic invariance for linear systems, the extension to nonlinear systems has not yet been addressed. Furthermore, while Brockman and Corless (1995, 1998); Alessandri et al. (2004) examine robust invariance for both linear and nonlinear systems, robust invariance does not apply to systems where probabilistic uncertainty exists. This is true even for works that consider robust invariance properties for systems learned using Gaussian process (GP) regression models, such as Wang et al. (2018); Taylor et al. (2020). These works provide robust invariance guarantees (via barrier functions), but do not treat the uncertainty as probabilistic; instead, uncertainty is treated as a fixed function for which the GP regression model provides a prediction with pointwise error bounds. In contrast, we consider GPSSMs which enable us to examine probabilistic invariance for nonlinear systems, where the probabilistic uncertainty is due to uncertainty in the functional model of the system as well as stochastic noise. In particular, we ensure that for GPSSMs the state trajectory lies within a particular invariant set with a certain probability for all time.

The remainder of this paper is organized as follows. Section 2 introduces the system model and a few properties of GPSSMs. In Section 3, we derive probability level sets for GPs. Section 4 introduces probabilistic invariance and sets forth a sufficient condition for evaluating the probabilistic invariance of the overall system. Section 5 illustrates the usefulness of this analysis for a smart water distribution system, and Section 6 concludes the paper.

2. 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 GPs, given by

$$x_{k+1} = Ax_k + g(x_k) + Bu_k + w_k, \tag{1}$$

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, $g(x_k) \triangleq \begin{bmatrix} g_1(x_k) & \cdots & g_n(x_k) \end{bmatrix}^T$, and $g_i(x_k) \sim \mathcal{GP}(0, k_i(x_k, x_k'))$ is a zero mean GP specified by its covariance function $k_i(x_k, x_k')$: $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$. A GP is a distribution over functions, assigning a joint Gaussian distribution to any finite subset of the state space (Rasmussen and Williams, 2006). 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 a state feedback control policy $u_k = Lx_k$ is implemented so that the system in (1) can be written as

$$x_{k+1} = A_{bl}x_k + g(x_k) + w_k, \quad A_{bl} \triangleq A + BL. \tag{2}$$

We also 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 $\bar{D} \triangleq \{\{\bar{x}_j, \bar{u}_j\}, \bar{x}_j^+\}_{j=1}^N$, where

$$\bar{x}_j^+ = A\bar{x}_j + g(\bar{x}_j) + B\bar{u}_j + w, \quad w \sim \mathcal{N}(0, Q). \tag{3}$$

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

$$g(x_k)|\{x_k, \bar{D}\} \sim \mathcal{N}(\mu(x_k), \Sigma(x_k)), \quad \mu(x_k) \triangleq \begin{bmatrix} \bar{k}_1(x_k)^T(K_1 + \sigma_1^2 I_N)^{-1}(y_1 - \bar{y}_1) \\ \vdots \\ \bar{k}_n(x_k)^T(K_n + \sigma_n^2 I_N)^{-1}(y_n - \bar{y}_n) \end{bmatrix}, \tag{4}$$

$$\Sigma(x_k) \triangleq \begin{bmatrix} \xi_1(x_k) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \xi_n(x_k) \end{bmatrix}, \quad \xi_i(x_k) \triangleq k_i(x_k, x_k) - \bar{k}_i(x_k)^T (K_i + \sigma_i^2 I_N)^{-1} \bar{k}_i(x_k),$$

$$\bar{k}_i(x_k) \triangleq \begin{bmatrix} k_i(\bar{x}_1, x_k) \\ \vdots \\ k_i(\bar{x}_N, x_k) \end{bmatrix}, K_i \triangleq \begin{bmatrix} k_i(\bar{x}_1, \bar{x}_1) & \cdots & k_i(\bar{x}_1, \bar{x}_N) \\ \vdots & \ddots & \vdots \\ k_i(\bar{x}_N, \bar{x}_1) & \cdots & k_i(\bar{x}_N, \bar{x}_N) \end{bmatrix}, y_i \triangleq \begin{bmatrix} \bar{x}_1^+(i) \\ \vdots \\ \bar{x}_N^+(i) \end{bmatrix}, \bar{y}_i \triangleq \begin{bmatrix} A_i \bar{x}_1 + B_i \bar{u}_1 \\ \vdots \\ A_i \bar{x}_N + B_i \bar{u}_N \end{bmatrix},$$

where $\bar{x}_j^+(i)$ denotes the i^{th} dimension of $\bar{x}_j^+$ while A_i and B_i denote the i^{th} rows of A and B, respectively. This implies that the system in (2) can be written equivalently as

$$x_{k+1} = A_{bl} x_k + \mathcal{I} \bar{w}_k(x_k), \tag{5}$$

where $\mathcal{I} \triangleq \begin{bmatrix} I_n & I_n & I_n \end{bmatrix}$, $\bar{w}_k(x_k) \triangleq \begin{bmatrix} \mu(x_k)^T & \bar{g}(x_k)^T & w_k^T \end{bmatrix}^T$, and

$$\bar{g}(x_k) | \{x_k, \bar{D}\} \sim \mathcal{N}(0, \Sigma(x_k)). \tag{6}$$

Remark 1 *Note that for stationary covariance functions, $k_i(x_k, x_k)$ is a constant because it is not a function of x_k. For example, the value of $k_i(x_k, x_k)$ for each of the following stationary covariance functions is respectively given by $\bar{\sigma}_i^2$, $\bar{\sigma}_i^2$, 1, $\bar{\sigma}_i^2$, $\bar{\sigma}_i^2$, and 1.*

$$k_i(x_k, x_k') = \begin{cases} \bar{\sigma}_i^2 & \textit{constant} \\ \bar{\sigma}_i^2 e^{-\frac{1}{2}(x_k - x_k')^T \bar{L}_i^{-2}(x_k - x_k')} & \textit{squared exponential} \\ \frac{(2\nu_i (x_k - x_k')^T \bar{L}_i^{-2}(x_k - x_k'))^{\nu_i/2} K_{\nu_i}((2\nu_i (x_k - x_k')^T \bar{L}_i^{-2}(x_k - x_k'))^{1/2})}{2^{\nu_i - 1}\Gamma(\nu_i)} & \textit{Mátern} \\ \bar{\sigma}_i^2 e^{-((x_k - x_k')^T \bar{L}_i^{-2}(x_k - x_k'))^{1/2}} & \textit{exponential} \\ \bar{\sigma}_i^2 e^{-((x_k - x_k')^T \bar{L}_i^{-2}(x_k - x_k'))^{\gamma_i/2}} & \gamma\textit{-exponential} \\ (1 + \frac{1}{2c_i}(x_k - x_k')^T \bar{L}_i^{-2}(x_k - x_k'))^{-c_i} & \textit{rational quadratic} \end{cases}$$

The parameters in each of these stationary covariance functions are given by $\bar{\sigma}_i^2 \in \mathbb{R}_{\geq 0}$, $\nu_i, c_i \in \mathbb{R}_{>0}$, $\gamma_i \in (0, 2]$, $\bar{L}_i = \mathit{Diag}(\ell_{i1}, \cdots, \ell_{in})$ with $\ell_{ij} \in \mathbb{R}_{>0}$, K_{ν_i} is a modified Bessel function of the second kind, and Γ is the gamma function.

Before setting forth a probabilistic invariant set for the system in (5), we note that the quantity $\bar{g}(x_k)^T \Sigma(x_k)^{-1} \bar{g}(x_k) | x_k$ follows a χ^2 distribution, as shown in Lemma 1.

Lemma 1

$$\bar{g}(x_k)^T \Sigma(x_k)^{-1} \bar{g}(x_k) | x_k \sim \chi^2(n), \tag{7}$$

where $\chi^2(n)$ represents the χ^2 distribution with n degrees of freedom.

Proof Follows directly from (6) and the properties of the χ^2 distribution. ■

3. Gaussian Process Probability Level Sets

To set forth a probabilistic invariant set for the system in (5), we first characterize sets in which $\mu(x_k)$ and $\bar{g}(x_k)$ each lie with at least probability p.

3.1. Probability Level Set for $\mu(x_k)$

Lemma 2 provides a set in which $\mu(x_k)$ lies with 100% probability.

Lemma 2 *For the system in* (5) *with stationary covariance functions and* $\forall p \in [0, 1]$,

$$Pr(\mu(x_k) \in X_1) \geq p, \quad X_1 \triangleq \left\{ s \middle| s^T s \leq \sum_{i=1}^{n} (\hat{k}_i^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2 \right\}, \tag{8}$$

where $\hat{k}_i \triangleq \begin{bmatrix} k_i(x_k, x_k) & \cdots & k_i(x_k, x_k) \end{bmatrix}^T \in \mathbb{R}^N$ *which is a constant for stationary covariance functions.*

Proof

$$\mu(x_k)^T \mu(x_k) = \sum_{i=1}^{n} (\bar{k}_i(x_k)^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2 \leq \sum_{i=1}^{n} (\hat{k}_i^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2,$$

where the inequality follows from the fact that $k_i(x_k, x_k) = \max_{x_k \in \mathbb{R}^n} k_i(\bar{x}_j, x_k)\ \forall j$ for stationary covariance functions. ■

3.2. Probability Level Sets for $\bar{g}(x_k)$

Using the results in Lemma 1, Lemma 3 provides a probability level set for $\bar{g}(x_k)$, and Lemma 4 provides a similar result that more explicitly depicts the relationship between the size of the confidence region, the probability level, and the degrees of freedom. However, the probability level set presented in Lemma 4 is larger than that presented in Lemma 3 since Lemma 4 uses the multidimensional Chebyshev inequality instead of directly computing the CDF of the χ^2 distribution.

Lemma 3 *For the system in* (5) *with stationary covariance functions and* $p \in [0, 1]$,

$$Pr(\bar{g}(x_k) \in X_2(p)) \geq p, \quad X_2(p) \triangleq \left\{ s \middle| s^T \bar{\Sigma}^{-1} s \leq F_{\chi^2}^{-1}(p, n) \right\}, \tag{9}$$

where $\bar{\Sigma} \triangleq Diag(k_1(x_k, x_k), \cdots, k_n(x_k, x_k))$*, which is a constant for stationary covariance functions, and* $F_{\chi^2}^{-1}(p, n)$ *is the value of the inverse CDF for the* χ^2 *distribution evaluated at probability* p *with* n *degrees of freedom.*

Proof According to Lemma 1,

$$\Pr(\bar{g}(x_k) \in \bar{X}_2(p) | x_k) = p, \quad \bar{X}_2(p) \triangleq \left\{ s \middle| s^T \Sigma(x_k)^{-1} s \leq F_{\chi^2}^{-1}(p, n) \right\}. \tag{10}$$

Note that $\bar{\Sigma} \succeq \Sigma(x_k)\ \forall x_k \in \mathbb{R}^n$, implying that $X_2(p) \supseteq \bar{X}_2(p)$, in turn implying that

$$\Pr(\bar{g}(x_k) \in X_2(p) | x_k) \geq \Pr(\bar{g}(x_k) \in \bar{X}_2(p) | x_k). \tag{11}$$

Combining (10) and (11) yields

$$\Pr(\bar{g}(x_k) \in X_2(p) | x_k) \geq p. \tag{12}$$

Applying the law of total probability with (12) yields the desired result:

$$\Pr(\bar{g}(x_k) \in X_2(p)) = \int_{x_k \in \mathbb{R}^n} \Pr(\bar{g}(x_k) \in X_2(p)|x_k) f(x_k) dx_k \geq p \int_{x_k \in \mathbb{R}^n} f(x_k) dx_k = p,$$

where $f(x_k)$ represents the probability density function (PDF) of x_k. ■

Lemma 4 *For the system in* (5) *with stationary covariance functions and* $p \in [0, 1]$,

$$Pr(\bar{g}(x_k) \in \mathcal{X}_2(p)) \geq p, \quad \mathcal{X}_2(p) \triangleq \left\{ s \middle| s^T \bar{\Sigma}^{-1} s \leq \frac{n}{1-p} \right\}. \tag{13}$$

Proof According to (6) and the multidimensional Chebyshev inequality,

$$\Pr(\bar{g}(x_k) \in \bar{\mathcal{X}}_2(p)|x_k) \geq p, \quad \bar{\mathcal{X}}_2(p) \triangleq \left\{ s \middle| s^T \Sigma(x_k)^{-1} s \leq \frac{n}{1-p} \right\}. \tag{14}$$

Since $\bar{\Sigma} \succeq \Sigma(x_k) \; \forall x_k \in \mathbb{R}^n$, this implies that $\mathcal{X}_2(p) \supseteq \bar{\mathcal{X}}_2(p)$, in turn implying that

$$\Pr(\bar{g}(x_k) \in \mathcal{X}_2(p)|x_k) \geq \Pr(\bar{g}(x_k) \in \bar{\mathcal{X}}_2(p)|x_k). \tag{15}$$

Combining (14) and (15) yields

$$\Pr(\bar{g}(x_k) \in \mathcal{X}_2(p)|x_k) \geq p. \tag{16}$$

Applying the law of total probability with (16) yields the desired result:

$$\Pr(\bar{g}(x_k) \in \mathcal{X}_2(p)) = \int_{x_k \in \mathbb{R}^n} \Pr(\bar{g}(x_k) \in \mathcal{X}_2(p)|x_k) f(x_k) dx_k \geq p \int_{x_k \in \mathbb{R}^n} f(x_k) dx_k = p.$$

■

4. Probabilistic Invariance Guarantees

Having quantified probability level sets for $\mu(x_k)$ and $\bar{g}(x_k)$, we now set forth the ellipsoid in which x_k lies with at least probability p for the system in (5). To do so, we leverage the notions of quadratic boundedness, robust positive invariance, and probabilistic positive invariance which are described in Definitions 1, 2, and 3, respectively. These definitions, along with Lemmas 5 and 6, have been modified and adapted from Alessandri et al. (2004); Hewing et al. (2018) in order to arrive at the results presented in Theorems 1 and 2.

Definition 1 (Alessandri et al. (2004)) *Let* $z_k \in \mathbb{R}^{n_z}$ *and* $d_k \in \mathbb{R}^{n_d}$ *represent state and disturbance vectors, respectively, and let* D *be a compact set. A system of the form*

$$z_{k+1} = \mathcal{A} z_k + \mathcal{B} d_k \tag{17}$$

is quadratically bounded with symmetric positive definite Lyapunov matrix $\mathcal{P}$ *if and only if*

$$z_k^T \mathcal{P} z_k \geq 1 \implies z_{k+1}^T \mathcal{P} z_{k+1} \leq z_k^T \mathcal{P} z_k \; \forall d_k \in D. \tag{18}$$

Definition 2 (Alessandri et al. (2004)) *The set $\mathcal{Z}$ is a robustly positively invariant set for (17) if and only if $z_k \in \mathcal{Z}$ implies that $z_{k+1} \in \mathcal{Z}\ \forall d_k \in D$.*

Definition 3 (Hewing et al. (2018)) *Let $d_k \sim \mathcal{Q}_k^d$ indicate that d_k is a random variable of distribution $\mathcal{Q}_k^d$. The set $\mathcal{Z}$ is a probabilistic positively invariant set of probability level p for (17) with $d_k \sim \mathcal{Q}_k^d$ if and only if $z_0 \in \mathcal{Z}$ implies that $Pr(z_k \in \mathcal{Z}) \geq p\ \forall k \geq 0$.*

Given these definitions, Lemma 5 provides a condition that is equivalent to robust positive invariance, Lemma 6 shows the correspondence between robust positive invariance and probabilistic positive invariance, and Lemma 7 combines these results to provide a sufficient condition for probabilistic positive invariance.

Lemma 5 (Alessandri et al. (2004)) *Let $d_k \in D \triangleq \left\{s \middle| s^T \mathcal{D} s \leq 1,\ \mathcal{D} \succ 0\right\}$. $\mathcal{Z} \triangleq \left\{s \middle| s^T \mathcal{P} s \leq 1\right\}$ is a robustly positively invariant set for (17) and the system in (17) is quadratically bounded with symmetric positive definite Lyapunov matrix $\mathcal{P}$ if and only if $\exists \alpha \geq 0$ such that*

$$\begin{bmatrix} (\alpha-1)\mathcal{P} + \mathcal{A}^T \mathcal{P} \mathcal{A} & \mathcal{A}^T \mathcal{P} \mathcal{B} \\ \mathcal{B}^T \mathcal{P} \mathcal{A} & \mathcal{B}^T \mathcal{P} \mathcal{B} - \alpha \mathcal{D} \end{bmatrix} \preceq 0. \tag{19}$$

Lemma 6 (Hewing et al. (2018)) *If $\mathcal{Z}$ is a robustly positively invariant set for (17) with $d_k \in D \triangleq \left\{s \middle| s^T \mathcal{D} s \leq 1,\ \mathcal{D} \succ 0\right\}$, then $\mathcal{Z}$ is also a probabilistic positively invariant set of probability level p for (17) with $d_k \sim \mathcal{Q}_k^d$, $\mathbb{E}[d_k] = 0$, $Cov[d_k] = \mathcal{D}^{-1}$, and $Pr(d_k \in D) \geq p$.*

Lemma 7 *Let $d_k \sim \mathcal{Q}_k^d$, $\mathbb{E}[d_k] = 0$, and $Cov[d_k] = \mathcal{D}^{-1}$. If $Pr(d_k \in D) \geq p$ and $\exists \alpha \geq 0$, $\mathcal{P} \succ 0$ such that (19) is satisfied, then $\mathcal{Z} \triangleq \left\{s \middle| s^T \mathcal{P} s \leq 1\right\}$ is a probabilistic positively invariant set of probability level p for (17).*

Proof Follows directly from Lemmas 5 and 6. ■

Given this sufficient condition for evaluating probabilistic invariance, Theorems 1 and 2 set forth an ellipsoid in which x_k lies with at least probability p for the system in (5), and Corollary 1 describes the volume of this ellipsoid with respect to p. Note that because $X_2(p) \subseteq \mathcal{X}_2(p)$ (as discussed immediately preceding Lemma 3), the ellipsoid defined in Theorem 1 is smaller than that presented in Theorem 2. However, the linear matrix inequality (LMI) presented in Theorem 2 is not a function of the probability p and therefore only needs to be evaluated once, unlike the LMI presented in Theorem 1 which is a function of p and must be evaluated each time for different values of p.

Theorem 1 *If $\exists \alpha \geq 0$, $P \succ 0$ such that*

$$\begin{bmatrix} (\alpha-1)P + A_{bl}^T P A_{bl} & A_{bl}^T P \mathcal{I} \\ \mathcal{I}^T P A_{bl} & \mathcal{I}^T P \mathcal{I} - \alpha \bar{R}(p) \end{bmatrix} \preceq 0, \tag{20}$$

where $\bar{R}(p) \triangleq \frac{1}{3} BlkDiag\left(\frac{1}{\sum_{i=1}^n (\hat{k}_i^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2} I_n, \frac{1}{F_{\chi^2}^{-1}(p,n)} \bar{\Sigma}^{-1}, \frac{1-p}{n} Q^{-1}\right)$, then

$$\bar{\mathcal{E}} \triangleq \left\{s \middle| s^T P s \leq 1\right\} \tag{21}$$

is a probabilistic positively invariant set of probability level p for (5) when stationary covariance functions are used for the GPs.

Proof According to the multidimensional Chebyshev inequality,

$$\Pr(w_k \in W(p)) \geq p, \quad W(p) \triangleq \left\{ s \middle| s^T Q^{-1} s \leq \frac{n}{1-p} \right\}. \tag{22}$$

According to Lemmas 2 and 3, when stationary covariance functions are used for the GPs, $\Pr(\mu(x_k) \in X_1) \geq p$ and $\Pr(\bar{g}(x_k) \in X_2(p)) \geq p$, which together with (22) yield

$$\Pr(\bar{w}_k(x_k) \in \bar{W}(p)) \geq p, \quad \bar{W}(p) \triangleq \left\{ s \middle| s^T \bar{R}(p) s \leq 1 \right\}.$$

It can be shown that applying Lemma 7 to (5) yields the desired result when noting that X_1 is centered at 0 with shape matrix I_n (Hewing et al., 2018). ■

Theorem 2 *If* $\exists \alpha \geq 0$, $P \succ 0$ *such that*

$$\begin{bmatrix} (\alpha - 1)P + A_{bl}^T P A_{bl} & A_{bl}^T P \mathcal{I} \\ \mathcal{I}^T P A_{bl} & \mathcal{I}^T P \mathcal{I} - \alpha R \end{bmatrix} \preceq 0, \tag{23}$$

where $R \triangleq \frac{1}{3}\mathit{BlkDiag}\left(\frac{1}{\sum_{i=1}^n (\hat{k}_i^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2} I_n, \frac{1}{n}\bar{\Sigma}^{-1}, \frac{1}{n} Q^{-1} \right)$, *then*

$$\mathcal{E}(p) \triangleq \left\{ s \middle| s^T P s \leq \frac{1}{1-p} \right\} \tag{24}$$

is a probabilistic positively invariant set of probability level p *for* (5) *when stationary covariance functions are used for the GPs.*

Proof According to Lemmas 2 and 4, when stationary covariance functions are used for the GPs, $\Pr(\mu(x_k) \in X_1) \geq p$ and $\Pr(\bar{g}(x_k) \in \mathcal{X}_2(p)) \geq p$, which together with (22) yield

$$\Pr(\bar{w}_k(x_k) \in \hat{W}(p)) \geq p, \quad \hat{W}(p) \triangleq \left\{ s \middle| s^T \hat{R}(p) s \leq 1 \right\}, \tag{25}$$

where $\hat{R}(p) \triangleq \frac{1}{3}\mathrm{BlkDiag}\left(\frac{1}{\sum_{i=1}^n (\hat{k}_i^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2} I_n, \frac{1-p}{n}\bar{\Sigma}^{-1}, \frac{1-p}{n} Q^{-1} \right)$. Note that

$$\Pr(\bar{w}_k(x_k) \in \hat{\hat{W}}(p)) \geq \Pr(\bar{w}_k(x_k) \in \hat{W}(p)), \quad \hat{\hat{W}}(p) \triangleq \left\{ s \middle| s^T \hat{\hat{R}}(p) s \leq 1 \right\}, \tag{26}$$

where $\hat{\hat{R}}(p) \triangleq \frac{1}{3}\mathrm{BlkDiag}\left(\frac{1-p}{\sum_{i=1}^n (\hat{k}_i^T (K_i + \sigma_i^2 I_N)^{-1} (y_i - \bar{y}_i))^2} I_n, \frac{1-p}{n}\bar{\Sigma}^{-1}, \frac{1-p}{n} Q^{-1} \right)$. This is due to the fact that $\hat{\hat{W}}(p) \supseteq \hat{W}(p)$ since $\hat{\hat{R}}(p) \preceq \hat{R}(p)$. Combining (25) and (26) yields

$$\Pr(\bar{w}_k(x_k) \in \hat{\hat{W}}(p)) \geq p.$$

It can be shown that applying Lemma 7 to (5) (when noting that X_1 is centered at 0 with shape matrix I_n (Hewing et al., 2018)) yields that if $\exists \alpha \geq 0$, $\bar{P} \succ 0$ such that

$$\begin{bmatrix} (\alpha - 1)\bar{P} + A_{bl}^T \bar{P} A_{bl} & A_{bl}^T \bar{P} \mathcal{I} \\ \mathcal{I}^T \bar{P} A_{bl} & \mathcal{I}^T \bar{P} \mathcal{I} - \alpha \hat{\hat{R}}(p) \end{bmatrix} \preceq 0,$$

then $\mathcal{E}(p) = \left\{ s \middle| s^T \bar{P} s \leq 1 \right\}$ is a probabilistic positively invariant set of probability level p for (5). Letting $\bar{P} \triangleq (1-p)P$ yields the desired result since $\hat{\hat{R}}(p) = (1-p)R$. ■

Corollary 1 *The volume of $\mathcal{E}(p)$ is proportional to $-\log\det((1-p)P) = -\log((1-p)^n\det(P))$ $= -n\log(1-p) - \log\det(P)$, which grows unbounded as p approaches* 1.

To minimize the volume of the ellipsoid $\bar{\mathcal{E}}$ or $\mathcal{E}(p)$, the parameter P can be designed according to the following optimization problem:

$$\arg\max_{\alpha,P} \log\det P \text{ s.t. } \alpha \geq 0,\ P \succ 0,\ (20) \text{ or } (23) \text{ is satisfied.} \tag{27}$$

Since minimizing $\bar{\mathcal{E}}$ or $\mathcal{E}(p)$ jointly over α and P in (27) is not convex, a suboptimal solution may be obtained by restricting the possible values of α to a finite set in $[0, 1]$ and maximizing $\max_{P\succ 0} \log\det P$ subject to (20) or (23) over that finite set.

5. Example

We consider a smart water distribution system used at a four-hectare wine estate in the south of England (Fu and McCann, 2020). The goal of the water distribution system is to stabilize the water levels of three district meter area tanks at predesigned constant reference levels. The system state is given by the difference between the reference levels and the current water levels of the three tanks, the control inputs are the open levels of the valves, and the sensors measure the current water levels of the tanks. The system model is linearized at a reference level of 3 m as presented in Fu and McCann (2020), and the nonlinearities and uncertainties of the system are accounted for with a GPSSM as modeled in (1). We discretize the system with a sampling rate of 1 s and use a state feedback controller L with three eigenvalues placed at 0.001. A squared exponential covariance function is used for each GP, and we train the hyperparameters of the covariance functions as well as Q by optimizing the marginal likelihood with training data recorded from a trajectory of the water distribution system. We solve for P by finding a suboptimal solution to the optimization problem in (27) according to the procedure described immediately following (27).

Figure 1 shows the state trajectory for tanks 2 and 3 for one GP sample over 1000 s from the initial state $x_0 = \begin{bmatrix}-2 & -2 & -2\end{bmatrix}^T$. As can be seen in the figure, the state quickly converges towards the origin before remaining within the 70% probability level set for the rest of the time, demonstrating that the empirical probability of remaining within each probability level set is greater than the lower bound set forth in Theorem 1. This was confirmed by running 1000 GP samples for 1000 s each from the initial state $x_0 = \begin{bmatrix}-2 & -2 & -2\end{bmatrix}^T$, where we found that across all GP samples the probability of the state remaining within the 50% probability level set was 99.9827%.

Figures 2 and 3 depict the sizes of the 50% probability level sets as a function of the variance parameter $\bar{\sigma}_i^2$ and the length parameter $\bar{L}_i$ in the squared exponential covariance function. Here we see that the size of the probability level set increases with increasing $\bar{\sigma}_i^2$ and decreasing $\bar{L}_i$, demonstrating that there is less confidence in where the state lies as more uncertainty is introduced into each GP.

Figure 4 depicts the volume of $\mathcal{E}(p)$, the ellipsoid in which x_k lies with at least probability $p\ \forall k$. As stated in Corollary 1, Figure 4 shows that the volume of $\mathcal{E}(p)$ grows unbounded as p approaches 1, indicating that larger ellipsoids are associated with greater confidence regions.

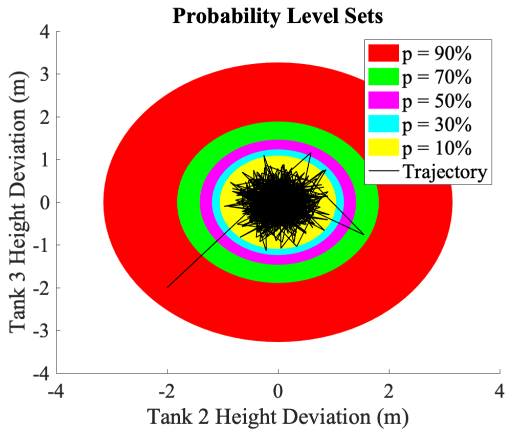

Figure 1: Example trajectory for tanks 2 and 3.

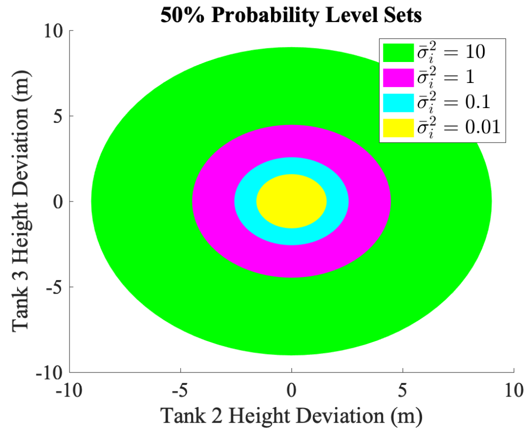

Figure 2: Probability level sets for $\bar{\sigma}_i^2$.

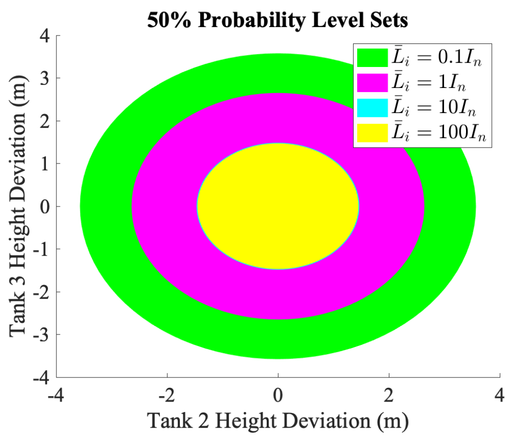

Figure 3: Probability level sets for $\bar{L}_i$.

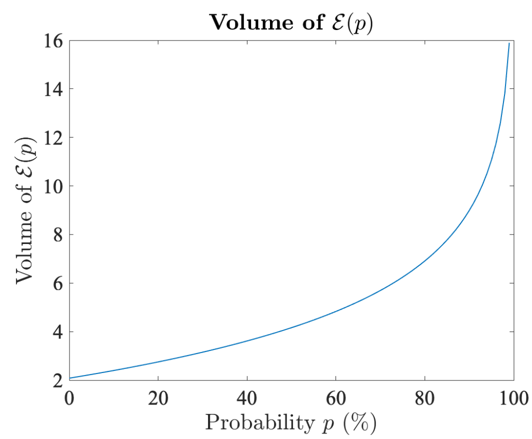

Figure 4: Volume of $\mathcal{E}(p)$ as a function of p.

6. Conclusion

This paper has presented safety guarantees in the form of probabilistic invariance for GPSSMs. A sufficient condition in the form of an LMI is provided to evaluate the probability of the state trajectory remaining within an invariant set for all time. Our results are illustrated with the example of a smart water distribution system. Future work includes investigating probabilistic invariance for GPSSMs with nonlinear means and/or nonstationary covariance functions. Another problem of interest includes examining probabilistic invariant sets that do not take ellipsoidal shapes and therefore may be less conservative. A further area of interest includes investigating how much the probabilistic invariant set contracts as more training data becomes available online.

Acknowledgments

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

References

Angelo Alessandri, Marco Baglietto, and Giorgio Battistelli. On Estimation Error Bounds for Receding-Horizon Filters Using Quadratic Boundedness. *IEEE Transactions on Automatic Control*, 49(8):1350–1355, 2004.

Franco Blanchini. Set Invariance in Control. *Automatica*, 35(11):1747–1767, 1999.

Mark L Brockman and Martin Corless. Quadratic Boundedness of Nonlinear Dynamical Systems. In *Proceedings of 1995 34th IEEE Conference on Decision and Control*, volume 1, pages 504–509. IEEE, 1995.

Mark L Brockman and Martin Corless. Quadratic Boundedness of Nominally Linear Systems. *International Journal of Control*, 71(6):1105–1117, 1998.

Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian Processes for Data-Efficient Learning in Robotics and Control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423, 2013.

Stefanos Eleftheriadis, Tom Nicholson, Marc Deisenroth, and James Hensman. Identification of Gaussian Process State Space Models. *Advances in Neural Information Processing Systems*, 30, 2017.

Roger Frigola, Fredrik Lindsten, Thomas B Schön, and Carl Edward Rasmussen. Bayesian Inference and Learning in Gaussian Process State-Space Models with Particle MCMC. *Advances in Neural Information Processing Systems*, 26, 2013.

Roger Frigola, Yutian Chen, and Carl Edward Rasmussen. Variational Gaussian Process State-Space Models. *Advances in Neural Information Processing Systems*, 27, 2014.

Anqi Fu and Julie A McCann. Dynamic Decentralized Periodic Event-Triggered Control for Wireless Cyber–Physical Systems. *IEEE Transactions on Control Systems Technology*, 29(4):1783–1790, 2020.

Lukas Hewing, Andrea Carron, Kim P Wabersich, and Melanie N Zeilinger. On a Correspondence Between Probabilistic and Robust Invariant Sets for Linear Systems. In *2018 European Control Conference (ECC)*, pages 1642–1647. IEEE, 2018.

Ernesto Kofman, José A De Doná, and Maria M Seron. Probabilistic Set Invariance and Ultimate Boundedness. *Automatica*, 48(10):2670–2676, 2012.

Ernesto Kofman, José A De Doná, Maria M Seron, and Noelia Pizzi. Continuous-Time Probabilistic Ultimate Bounds and Invariant Sets: Computation and Assignment. *Automatica*, 71:98–105, 2016.

Carl Edward Rasmussen and Christopher KI Williams. *Gaussian Processes for Machine Learning*, volume 2. MIT Press Cambridge, MA, 2006.

Jeff Schneider. Exploiting Model Uncertainty Estimates for Safe Dynamic Control Learning. *Advances in Neural Information Processing Systems*, 9, 1996.

Andreas Svensson, Arno Solin, Simo Särkkä, and Thomas Schön. Computationally Efficient Bayesian Learning of Gaussian Process State Space Models. In *Artificial Intelligence and Statistics*, pages 213–221. PMLR, 2016.

Andrew Taylor, Andrew Singletary, Yisong Yue, and Aaron Ames. Learning for Safety-Critical Control with Control Barrier Functions. In *Learning for Dynamics and Control*, pages 708–717. PMLR, 2020.

Ryan Turner, Marc Deisenroth, and Carl Rasmussen. State-Space Inference and Learning with Gaussian Processes. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 868–875. JMLR Workshop and Conference Proceedings, 2010.

Jonas Umlauft and Sandra Hirche. Learning Stochastically Stable Gaussian Process State–Space Models. *IFAC Journal of Systems and Control*, 12:100079, 2020.

Li Wang, Evangelos A Theodorou, and Magnus Egerstedt. Safe Learning of Quadrotor Dynamics Using Barrier Certificates. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2460–2465. IEEE, 2018.