

Out of Distribution Detection via Domain-Informed Gaussian Process State Space Models

Alonso Marco, Elias Morley and Claire J. Tomlin

Abstract—In order for robots to safely navigate in unseen scenarios using learning-based methods, it is important to accurately detect out-of-training-distribution (OoD) situations online. Recently, Gaussian process state-space models (GPSSMs) have proven useful to discriminate unexpected observations by comparing them against probabilistic predictions. However, the capability for the model to correctly distinguish between in- and out-of-training distribution observations hinges on the accuracy of these predictions, primarily affected by the class of functions the GPSSM kernel can represent. In this paper, we propose (i) a novel approach to embed existing domain knowledge in the kernel and (ii) an OoD online runtime monitor, based on receding-horizon predictions. Domain knowledge is provided in the form of a dataset, collected either in simulation or by using a nominal model. Numerical results show that the informed kernel yields better regression quality with smaller datasets, as compared to standard kernel choices. We demonstrate the effectiveness of the OoD monitor on a real quadruped navigating an indoor setting, which reliably classifies previously unseen terrains.

I. INTRODUCTION

As machine learning (ML) becomes increasingly integrated into autonomous robotic systems like service robotics [1], [2] and self-driving cars [3], the need for reliable and safe ML models has never been more critical. Particularly in human-centric and safety-critical settings, even minor representation errors in ML models can have severe consequences [4]. Hence, it is crucial to identify the shortcomings of these models. One key issue is their potential unreliability when faced with data that differs from what they were trained on, commonly known as out-of-training-distribution (OoD).

This challenge has motivated the emerging research area of OoD detection [5]. The field encompasses a variety of methodologies, ranging from invariant representation learning [6], [7] to causal learning for distribution shifts [8]. While these techniques have been extensively studied in image classification [9], [10], their application to all levels of the full autonomy stack is still unfolding [11].

In the field of robotics, real-time OoD detection is crucial for both real-time decision making and long-term reliability [12], especially in uncertain environments. Algorithms need to be uncertainty-aware and equipped with fallback strategies for OoD scenarios, either issuing warnings or adopting more

All the authors are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, 2594 Hearst Ave, Berkeley, CA 94720, USA. amarco@berkeley.edu, emorley@berkeley.edu, tomlin@eecs.berkeley.edu.

This work is supported by the DARPA Assured Autonomy program, the NASA ULI Project in Safe Aviation Autonomy, and the Rafael del Pino Foundation.




Fig. 1. A quadruped is exposed to different scenarios, unseen at training time, in order to assess the capability for out-of-training-distribution detection.

conservative behaviors. However, research avenues in real-time OoD detection are rather sparse. In [13], an OoD detector is trained on grasping datasets to have a provable low false negative rate. A similar idea is presented in [14], where a failure predictor with guaranteed bounds is deployed on a real drone. However, these are not real-time, but rather sequential OoD monitors, since a full episode needs to be observed before OoD can be assessed. The same episodic scheme is followed in [15], where OoD is measured using a task-specific performance cost. Our approach monitors OoD in real-time, which is crucial for decision-making in safety-critical situations, while remaining task-agnostic.

Real-time OoD runtime monitors offer significant advantages in the field of legged locomotion, where recognizing changes in environmental conditions is paramount to anticipate hazards. Yet, deploying these monitors in highly dynamical systems, such as quadrupeds, poses major computational and design challenges. In this work, we propose an OoD runtime monitor that tracks prediction errors in a receding horizon manner using Gaussian Process State-Space Models (GPSSMs) [16]–[21]. These models provide a probabilistic representation of the state-space dynamics and are well-suited for predicting future states and their uncertainties, making them ideal for safe navigation [22], [23].

Although GPSSMs have gained popularity for their ability to forecast state trajectories, they come with a few drawbacks. On one hand, the distribution of forecast states

is often intractable and requires numerical approximations. On the other hand, they require large datasets for training, which is often impractical in real-world robotics where data is scarce [21]. To address the first issue, variational inference approaches have been proposed to also reduce the computational complexity [19], [21], yet large datasets are still required. One effective strategy to alleviate the second issue is to incorporate prior knowledge into the GPSSM’s kernel function. Physics-aware models serve this purpose by using first-principles as regularizers [24]. Yet, these models often employ off-the-shelf kernels, overlooking the potential benefits of embedding existing knowledge in the kernel design. Recent work suggests that customizing the kernel for specific control problems can improve the model’s performance [22], [25]. In this work, we tackle both issues by formulating the GPSSM as a Karhunen–Loève (KL) series expansion [20], [26]–[28], which poses two key advantages: (i) it enables efficient sampling schemes for state distribution forecasting, and (ii) it provides a flexible framework for designing kernels that integrate prior knowledge.

Contributions: Our first contribution is a data-driven method for constructing GPSSM kernels that incorporates existing knowledge in the form of a queryable nominal dynamics model, such as a simulator or a first-principles representation. The key step is embedding the nominal model into the kernel via its Fourier representation. To this end, we minimize an autoencoder loss that reconstructs the nominal model. The resulting GPSSM requires less data to achieve the same prediction accuracy as with standard kernels.

Our second contribution is an OoD runtime monitor that automatically detects scenarios unseen at training time. Our experiments on a real quadruped show that by leveraging the proposed GPSSM, the OoD runtime monitor can more accurately detect new terrains as compared to other GPSSMs that use standard non-informed kernels (e.g., Gaussian, or Matérn [29]). To the best of our knowledge, our work is the first to explicitly design a deployable online runtime monitor for OoD detection on a real quadrupedal robot.

The paper is structured as follows. In Sec. II, we explain the proposed GPSSM as a KL series expansion. In Sec. III, we detail the proposed methodology to construct a kernel by embedding domain knowledge, given as a nominal dynamics model. In Sec. IV, we explain how state forecasting can be used to detect OoD situations. Finally, in Sec. V, we evaluate empirically the capabilities of the proposed model for OoD detection on a real quadruped, which navigates a room while being exposed to various changes in the environment, such as rocky terrain and external forces.

II. GAUSSIAN PROCESS STATE-SPACE MODEL USING MERCER KERNELS

Let $x_t \in \mathcal{X} \subseteq \mathbb{R}^D$ be the system state and $u_t \in \mathcal{U} \subseteq \mathbb{R}^{D_u}$ the control input at time t . The transition to the next state x_{t+1} follows the system dynamics $f_{\text{true}} : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$. Whereas the true dynamics are unknown, they can be represented by a dynamics model $f(\cdot) = [f_1(\cdot), \dots, f_D(\cdot)]^\top$, which is described as a Gaussian process state-space model

(GPSSM) [19]

$$\begin{aligned} f_d(\cdot) &\sim \mathcal{GP}(m_d(\cdot), k_d(\cdot, \cdot)), \quad d = 1, \dots, D \\ x_{t+1} &\sim \mathcal{N}(f(x_t, u_t), Q) \\ \hat{x}_{t+1} &\sim p(\hat{x}_{t+1} | x_{t+1}), \end{aligned} \quad (1)$$

where each component $f_d : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ follows an independent Gaussian process (GP) with kernel $k_d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and mean function $m_d : \mathcal{X} \rightarrow \mathbb{R}$, and $Q = \text{diag}([\sigma_1^2, \dots, \sigma_D^2])$. For simplicity, we assume a Gaussian observation model¹ $\hat{x}_{t+1} \sim \mathcal{N}(x_{t+1}, \sigma_n^2 I)$.

In the following, we focus on modeling a single component of the dynamics f_d without loss of generality. For ease of presentation, we denote $z_t = [x_t^\top, u_t^\top]^\top$ and $\mathcal{Z} = \mathcal{X} \times \mathcal{U}$. We drop the subindex d in both f_d and σ_d^2 and define the d component of a state observation as $X_t = [\hat{x}_t]_d$.

A. Karhunen–Loève expansion

The Gaussian process (GP) representing the dynamics f can be defined in weight space as [29, Sec. 2.1]

$$f(z_t) = \sum_{j=1}^M \beta_j \phi_j(z_t), \quad (2)$$

with $M > 0$ independently distributed random weights $\beta_j \sim \mathcal{N}(m_j, v_j)$, mean $m_j \in \mathbb{R}$ and variance $v_j > 0$, and deterministic features $\phi_j : \mathcal{Z} \rightarrow \mathbb{R}$. This formulation is known as the *Karhunen–Loève (KL) series expansion* of a GP [26]. The covariance function associated with (2) is given as [30]

$$k(z_t, z'_t) = \sum_{j=1}^M v_j \phi_j(z_t) \phi_j(z'_t), \quad (3)$$

which is a (positive definite) *Mercer kernel* [26]. Informally, the KL expansion (2) is equivalent to a standard GP formulation $\mathcal{GP}(m(z_t), k(z_t, z'_t))$ with a positive definite kernel k that admits a series representation (3).

B. Predictive posterior

We are interested in the above formulation to predict the dynamics at new locations $f(z_*)$. Let us assume access to an approximate description of the dynamics through a nominal model $f_{\text{nom}} : \mathcal{Z} \rightarrow \mathbb{R}$ that captures the general behavior of the system, e.g., a simulator. We use f_{nom} to acquire a dataset of T state-control-state tuples $\mathcal{D} = \{\hat{z}_t, X_{t+1}\}_{t=0}^{T-1} \sim \nu(z_t)$, where the measure $\nu(z_t)$ is given implicitly as a distribution over dynamically feasible trajectories in \mathcal{Z} , and the observations are $X_{t+1} = f_{\text{nom}}(\hat{z}_t)$. The predictive posterior is solely determined by the posterior over the weights $p(\beta | \mathcal{D}) = \mathcal{N}(\mu, \Sigma)$, which is given analytically [31, Sec. 9.3] by

$$\begin{aligned} \mu &= \Sigma(V^{-1}m + \sigma^{-2}\Phi_Z X) \\ \Sigma &= (V^{-1} + \sigma^{-2}\Phi_Z \Phi_Z^\top)^{-1}, \end{aligned} \quad (4)$$

with feature matrix $\Phi_Z = [\Phi(\hat{z}_0), \dots, \Phi(\hat{z}_{T-1})]$, feature vector $\Phi(\hat{z}_t) = [\phi_1(\hat{z}_t), \dots, \phi_M(\hat{z}_t)]^\top$, observations $X =$

¹While integrating non-Gaussian observation models is possible, this is known to be challenging [19], [21] and out of the scope of this paper.

$[X_1, \dots, X_T]^\top$, and prior parameters $m = [m_1, \dots, m_M]^\top$ and $V = \text{diag}([v_1, \dots, v_M])$. The predictive posterior over $f(z_*)$ at a new location z_* is

$$f(z_*)|\mathcal{D} \sim \mathcal{N}(\mu^\top \Phi_*, \Phi_*^\top \Sigma \Phi_*), \quad \Phi_* = \Phi(z_*). \quad (5)$$

C. Computational complexity and model expressiveness

In general, a larger number of features M in (3) produces a more expressive model [25]. Specifically, if $M \rightarrow \infty$, the series in (3) converges to a closed-form expression [25], [30] for particular feature choices (e.g., Gaussian-shaped features give rise to the squared exponential kernel [29, Sec. 4.2.1]). Since such closed-form kernels entangle infinitely many features, they yield more expressive models than their truncated counterpart. However, the GP posterior with such kernels suffers from $O(T^3)$ complexity. On the contrary, the truncated series formulation (2) poses two main benefits: (i) it has cost $O(M^3T)$, i.e., linear with the number of datapoints² and (ii) it yields flexibility in constructing Mercer kernels that integrate existing domain knowledge. We discuss this idea next.

III. CONSTRUCTING A DATA-DRIVEN DOMAIN INFORMED KERNEL

In this section we propose a methodology to integrate existing domain knowledge into a Mercer kernel (3).

A. Integrating domain knowledge via Fourier features

The nominal model f_{nom} introduced in Sec. II-B holds prior information about the dynamics. Our approach is integrating such model into the GPSSMs as the prior mean function [32] through the mean coefficients of the KL expansion (2), i.e.,

$$\mathbb{E} \left[\sum_{j=1}^M \beta_j \phi_j(z_t) \right] = \sum_{j=1}^M m_j \phi_j(z_t) = f_{\text{nom}}(z_t). \quad (6)$$

While there exist many viable function families to represent $\phi_j(z_t)$, we choose a Fourier series expansion [33] due to its well-studied connection with the frequency domain [29].

Let f_{nom} be square integrable with respect to ν . Then, its truncated multivariate Fourier series expansion can be written as a neural network with one hidden layer [34] and a cosine activation function

$$f_{\text{nom}}(z_t) = \sum_{j=1}^M S(\omega_j) \cos(\omega_j^\top z_t + \varphi(\omega_j)), \quad (7)$$

where the frequencies $\omega_j \in \Omega \subseteq \mathbb{R}^{D+D_u}$ are spaced in a regular grid $\omega_j = \tilde{\omega}[j_1, \dots, j_{D_z}]^\top$, $j_i \in \mathbb{Z}$, and $\tilde{\omega} > 0$ is a base frequency. The modulus $S(\cdot) = |\mathcal{F}[f_{\text{nom}}](\cdot)|$ is known as *spectral density* [35] and $\varphi(\cdot) = \angle \mathcal{F}[f_{\text{nom}}](\cdot)$ is the *system phase*. They both depend on the Fourier series coefficients $\mathcal{F}[f_{\text{nom}}](\omega_j)$, defined through the multivariate Fourier transform [35]

$$\mathcal{F}[f_{\text{nom}}](\omega_j) = \int_{\mathcal{Z}} f_{\text{nom}}(z_t) e^{-i\omega_j^\top z_t} d\nu(z_t). \quad (8)$$

²In practice, we keep M as large as computational resources allow.

This approach poses two main computational caveats. First, having the frequencies in (7) regularly distributed is problematic because (i) the grid size M increases exponentially with the dimensionality, and (ii), an uncaredful choice of grid size may result in suboptimal placement of frequencies in regions where the influence of $S(\omega_j)$ and $\varphi(\omega_j)$ is negligible³. Second, the integrand in (8) is only observable at sparse locations, sampled from $\nu(z_t)$ (see Sec. II-B). In the following, we propose a numerical approach to mitigate both issues. We denote $S_j = S(\omega_j)$ and $\varphi_j = \varphi(\omega_j)$.

To alleviate the first issue, we propose an irregular grid that increases resolution in areas where S_j is large. Specifically, we follow [34], where the optimal parameters $\kappa^* = \{S_j^*, \varphi_j^*, \omega_j^*\}_{j=1}^M$ are obtained by minimizing the loss

$$\min_{\kappa} \frac{1}{T} \sum_{t=1}^{T-1} \|X_{t+1} - \sum_{j=1}^M S_j \cos(\omega_j^\top \hat{z}_t + \varphi_j)\|_2^2 + \lambda_\omega \|W\|_2^2, \quad (9)$$

where $W = [\omega_1, \dots, \omega_M]$ regularizes the frequencies. With this approach, the decoder network (7) shall place the learned frequencies ω_j^* in a non-uniform grid with higher resolution in areas where S_j^* and φ_j^* are most influential. The reconstruction loss is evaluated on the dataset \mathcal{D} (see Sec. II-B).

We tackle the second issue by numerically approximating (8) as a quadrature $\sum_{t=1}^{T-1} X_{t+1} e^{-i\omega_j^\top \hat{z}_t} \eta_t$, which can be seen as a one-layer encoder network, where the integration steps $\eta_t > 0$ are jointly learned with κ .

The sought mean coefficients and features in (6) are

$$m_j = S_j^*, \quad \phi_j(z_t) = \cos((\omega_j^*)^\top z_t + \varphi_j^*). \quad (10)$$

The reconstruction quality is affected by the size M of the irregular grid. We illustrate in Fig. 2 the impact of different grid sizes using the following toy example:

Example 1. Let us consider the scalar ‘‘elbow’’ dynamical system $f(z_t) = 0.8 + (z_t + 0.2)(1 - 0.5(1 + \exp(-2z_t))^{-1})$ [37], which we wish to reconstruct using the aforementioned autoencoder. To this end, a dataset of size $T = 50$ is collected from this dynamical system within the interval $z_t \in [-10, 10]$. The first two rows show the reconstructed f_{nom} after optimizing the autoencoder. The reconstruction improves as we increase the number of features from $M = 5$ to $M = 20$ for a small number of optimization steps. We also depict the true spectral density $S(\omega)$ and system phase $\varphi(\omega)$ that uniquely characterize this dynamical system. After optimization, the optimal frequencies ω_j^* are found at locations where both S_j^* and φ_j^* contribute the most to the reconstruction.

The proposed formulation allows us to embed information about f_{nom} in the kernel (3) through the dataset \mathcal{D}_n , which is discussed next.

³The Riemann–Lebesgue lemma [36] establishes that $S(\omega)$ vanishes as $|\omega| \rightarrow \infty$.

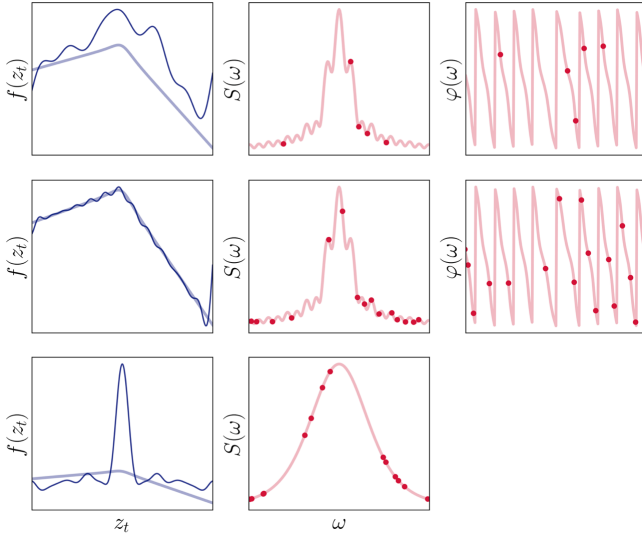


Fig. 2. Reconstruction of the “elbow” dynamical system. Top row, first column: True dynamical system (light blue line) and reconstructed dynamical system (dark blue line) with $M = 5$ features. Second column: Spectral density $S(\omega)$ of the true dynamical system. Third column: System phase $\varphi(\omega)$ of the true dynamical system. The red dots indicate the optimal values S_j^* , φ_j^* found after optimizing (9). Middle row: Analogous to top row with $M = 20$. Bottom row: Reconstruction using the spectral density of a Matérn 3/2 kernel.

B. Non-stationary domain-informed Mercer kernel

Given the informed features and mean coefficients (10), the proposed kernel follows from (3)

$$k(z_t, z'_t) = \sum_{j=1}^M v_j \cos((\omega_j^*)^\top z_t + \varphi_j^*) \cos((\omega_j^*)^\top z'_t + \varphi_j^*). \quad (11)$$

A key difference between this kernel and commonly used kernels (e.g., Gaussian or Matérn [29, Sec. 4.2]) or similar kernels with harmonic-based series representations [28], [38], [39], is that it is generally non-stationary, i.e., $k(z_t, z'_t) \neq k(z_t - z'_t)$. We find stationarity to be a restricting assumption, as it implies that the statistical properties of f_{nom} do not change over the input space, which is not necessarily true for arbitrary dynamical systems. The non-stationarity in (11) adds flexibility to the model without necessarily overfitting to a specific function class. We illustrate the shortcomings of stationarity in Fig. 2 (third row), where we attempt to reconstruct f_{nom} by using a Matérn kernel, which is an isotropic stationary kernel. To this end, we first transform (11) into a stationary kernel by fixing $v_j = \sqrt{2}$ and $\varphi_j \sim \mathcal{U}(0, 2\pi)$, as in [39]. Then, the optimal frequencies ω_j^* are obtained by solving (9), and setting $S_j^* = S_{\text{Matérn}}(\omega_j^*)$, where $S_{\text{Matérn}}(\cdot)$ is the spectral density of a Matérn kernel, given in [29, Sec. 4.2]. As shown, a non-informed stationary kernel alone is unable to capture the properties of the true dynamical system.

The proposed kernel enables embedding prior information into the proposed GPSSM (1), which shall influence the long-term state predictions. In the following, we discuss how to

use informed receding-horizon predictions to design an out-of-training-distribution runtime monitor.

IV. OOD RUNTIME MONITOR WITH GPSSM

We are interested in quantifying how likely it is that the robot navigates a scenario that is out-of-training-distribution (OoD). In the following, we propose a runtime monitor that compares the distribution of predicted states with the posterior after conditioning on collected observations.

A. Prediction-based OoD detection

Let $u_{t:H-1} = \{u_t, \dots, u_{t+H-1}\}$ be a control sequence that rolls out the dynamical model forward in time H steps, departing from the current state x_t . The resulting latent state sequence $x_{t+1:H}$, assumed to be Markovian, is distributed as

$$q(x_{t+1:H} | x_t, u_{t:H-1}) = \prod_{h=0}^{H-1} p(x_{t+h+1} | x_{t+h}, u_{t+h}). \quad (12)$$

Such distribution determines the future model predictions from the current state x_t . We let the system evolve for H time steps under $u_{t:H-1}$, and collect a set of observations $\hat{x}_{t+1:H}$, which induce a posterior distribution over the states $p(x_{t+1:H} | \hat{x}_{t+1:H}, x_t, u_{t:H-1})$. Our goal is to use the Kullback–Leibler divergence $D_{\text{KL}}(q||p) > 0$ as a proxy to detect OoD. The intuition behind this approach is as follows: In-distribution scenarios should be accurately predicted, resulting in a small divergence, while OoD scenarios should report a large divergence. Furthermore, it is well known that a low divergence $D_{\text{KL}}(q||p)$ yields a large evidence lower bound (ELBO) while the (unknown) log-evidence remains constant, i.e., $D_{\text{KL}}(q||p) + \mathbb{E}_q[\log p(\hat{x}_{t+1:H} | x_{t+1:H})] = \log p(\hat{x}_{t+1:H})$. Therefore, we define the OoD loss as the negative evidence lower bound

$$\mathcal{L}_{\text{OoD}} = -\mathbb{E}_q[\log p(\hat{x}_{t+1:H} | x_{t+1:H})]. \quad (13)$$

The distribution over the predicted state sequence (12) is non-Gaussian [21], which renders the expectation in (13) analytically intractable. In practice, we approximate it using a Monte Carlo sum [20], [27]. Since the likelihood (1) factorizes, the approximated runtime monitor is

$$\mathcal{L}_{\text{OoD}} \approx \frac{1}{R} \sum_{r=1}^R \sum_{h=1}^H \frac{1}{2\sigma_n^2} \| \hat{x}_{t+h} - \tilde{x}_{t+h}^{(r)} \|^2 + \text{const.}, \quad (14)$$

with rollouts $\tilde{x}_{t+1:H}^{(r)} \sim q(x_{t+1:H} | x_t, u_{t:H-1})$. We detail the sampling approach next.

B. Sampling scheme

Given a control sequence $u_{t:H-1}$, our goal is sampling a future state sequence from (12). While this is generally challenging, a key advantage of (1) over vanilla GPSSM is that β_d , $d = \{1, \dots, D\}$ constitutes sufficient statistics for predicting future states (see Fig. 3). Specifically, a posterior sample $\tilde{\beta}_d \sim \mathcal{N}(\mu_d, \Sigma_d)$ constitutes a deterministic instance of the unknown dynamics $\tilde{f}_d(x_t, u_t) = \tilde{\beta}_d^\top \Phi_d(x_t, u_t)$, $d = 1, \dots, D$. In practice, we use the reparametrization trick $\tilde{\beta}_d^{(r)} = \mu_d + L_d \tilde{\eta}^{(r)}$, $r = \{1, \dots, R\}$ to pre-sample and




Fig. 3. Graphical model for the proposed Gaussian process state-space model. Given a sequence of control inputs and the current state x_t , the distribution over future states is induced by β .

store R model instantiations at a fixed cost, where L_d is obtained through the Cholesky decomposition $\Sigma_d = L_d L_d^\top$ and $\tilde{\eta}^{(r)} \sim \mathcal{N}(0, 1)$. Under the Gaussian observation model (1), the next state is given deterministically as

$$\tilde{x}_{t+1}^{(r)}(x_t, u_t) = \begin{bmatrix} (\tilde{\beta}_1^{(r)})^\top \Phi_1(x_t, u_t) \\ \vdots \\ (\tilde{\beta}_D^{(r)})^\top \Phi_D(x_t, u_t) \end{bmatrix} + \tilde{\gamma}^{(r)}, \quad (15)$$

with $\tilde{\gamma}^{(r)} \sim \mathcal{N}(0, Q + \sigma_n^2 I)$. Hence, a sequence of H states is obtained by recursively calling (15). The overall runtime cost of sampling R trajectories is $O(HDR)$ per time step.

V. RESULTS

In this section, we present two sets of results. First, we assess the capability of the proposed method to detect out-of-training-distribution (OoD) situations with a real quadruped. We compare the performance of our GPSSM with another model that uses standard non-informed kernels. Second, we present a brief ablation study where we show that our domain-informed kernel is indeed more data efficient than existing GPSSMs that use standard kernels.

Parameter choices: The prior variance v_j over the weights (1) is a design choice that affects the quality of the regression. In practice, one can choose any $v_j > 0$ set of eigenvalues, e.g., constant [38]. Herein, we follow [28] and set $v_j = \xi m_j$, with $\xi > 0$, i.e., proportional to the spectral density value (see (10)). Intuitively, this choice assigns high confidence ($v_j \approx 0$) to weakly weighted features (i.e., $m_j \approx 0$), which causes them to have low impact in the model, while strongly weighted features remain uncertain, yet flexible and adaptable to the posterior. We utilize GPyTorch [40] to train the non-informed GPSSMs that employ standard kernels, conducting hyperparameter optimization for the same number of epochs as we do for the proposed model.

A. OoD runtime monitor in quadrupedal locomotion

In order to assess the efficacy of the proposed OoD runtime monitor, we conduct experiments using the quadrupedal robot Go1 [41] in an indoor environment. Our goal is to detect various changes in the environment as the robot traverses a circular path. To this end, we first gather motion

data as the robot walks in circles, without any environmental perturbations. Subsequently, we train the domain-informed GPSSM (Sec. II and III) using the collected data. Finally, we deploy the OoD runtime monitor described in Sec. IV, which leverages the trained model and detects environment changes. The goal is for the OoD runtime monitor to correctly identify the portions of the path affected by a variety of environment changes, which we detail next.

Rope pulling: A rope is attached to the robot's body while the other end of the rope is anchored, inducing a pulling force that deviates the robot from its intended circular path for a certain portion of the trajectory (see Fig. 1).

Rocky terrain: The robot is required to navigate through rocky terrain for a section of the trajectory.

Poking: The robot's body is gently prodded with a pole at various stages during its movement.

In the following, we describe our experimental setup, design choices, and discuss the results.

1) *Design choices:* We use the same state and control input representation as the ones for unicycle dynamics [42]. The state $x_t = [x_t, y_t, \alpha_t]^\top$ concatenates the X-Y robot position in the room and the heading angle $\alpha_t \in [-\pi, \pi)$, i.e., the robot's orientation in the X-Y plane. The control input $u_t = [v_t, \dot{\alpha}_t]^\top$ gathers the desired forward velocity v_t and the desired rotation speed $\dot{\alpha}_t$. Due to several real-world effects, the unicycle dynamics [42] are affected by several nonlinearities that are difficult to model, for example, communication delays and sensor noise. Also, inherent dynamics in the manufacturer's walking algorithms forbid from accurately tracking the desired speeds. The proposed GPSSM (1) absorbs such nonlinearities. We learn each component f_d using $M = 1500$ features and $\sigma_d = 0.01$, $d = \{1, 2, 3\}$.

Autoencoder training: As detailed in Sec. III-A, the autoencoder reconstructs a given nominal dynamics model f_{nom} . In this setting, the nominal model is implicitly represented through a dataset \mathcal{D} of state-control-state tuples, collected directly on the real system. Specifically, we collect 10 minutes of data at 10 Hz of the robot walking in circles on flat ground and use this dataset to train the autoencoder with $\lambda_\omega = 0.1$. We train for 100,000 epochs with a cosine decay learning rate. The frequencies ω_j are initialized by sampling from a uniform distribution within the domain $[-3, 3]^5$.

Sampling parameters: We follow the sampling scheme proposed in Sec. IV-B to compute (14). At each time step, we sample $R = 20$ system rollouts with a time horizon of $H = 30$ timesteps (i.e., 3 seconds lookahead).

2) *Experimental setup:* The position and orientation of the robot are estimated using a VICON motion tracker. The desired speed commands are computed online using a proportional controller that steers the robot toward a set of 20 waypoints distributed along a circle measuring 3 meters in diameter. The desired forward velocity is given as $v_t^{\text{des}} = 0.1 + K_p \| (x_{\text{next}}, y_{\text{next}}) - (x_t, y_t) \|_2$, where $(x_{\text{next}}, y_{\text{next}})$ is the position of the next waypoint. The desired angular velocity is computed as $\dot{\alpha}_t^{\text{des}} = 0.2(\theta_{\text{next}} - \theta_t)$. This controller is fixed across all environments. The OoD loss is normalized to be $\mathcal{L}_{\text{OoD}} \in [0, 1]$ across all environments. We send desired

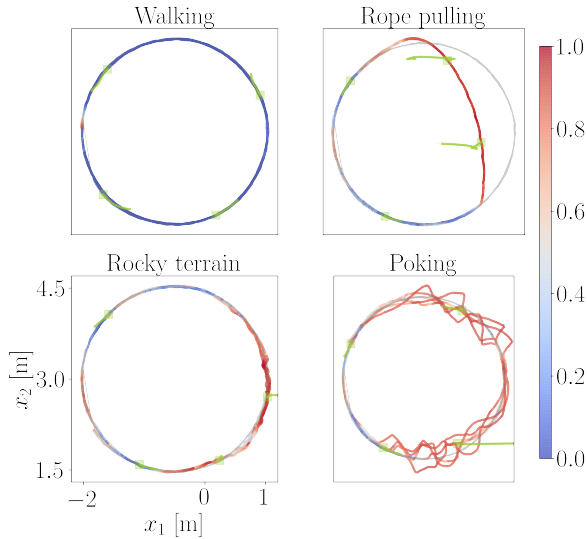


Fig. 4. A quadrupedal robot follows an anti-clockwise circular motion, while being exposed to three changes in the environment. Top-left: The robot walks in flat terrain without environment changes. Top-right: The robot is tied to a rope that creates a pulling force. Bottom-left: The robot walks on rocky terrain. Bottom-right: the robot is poked in random directions as it walks. The colormap transitions from blue (low OoD loss values) to red (high OoD loss values). Future state predictions are overlaid in light green at four different parts of the trajectory.

velocity commands to the robot’s on-board computer at 100 Hz via wireless communication. Our implementation uses ROS in Python and C++ on a standard commercial laptop.

3) *Results discussion:* In Fig. 4, we show that \mathcal{L}_{OoD} keeps low values when the robot passes through parts of the trajectory that were seen at training time. On the contrary, \mathcal{L}_{OoD} reaches high values when the robot is exposed to changes in the environment not captured at training time. In the rope pulling case, we see that the sampled state trajectories deviate significantly from the trajectory when the rope is tense, but remain in-distribution when it is loose. We see similar situations in the other two settings.

These results indicate that our model consistently remains in-distribution when the environment is significantly similar to the one it was exposed to at training time. In table I we show the percentage of time steps in which the robot was out of distribution. To determine this, we normalize the data offline and classify OoD situations with $\mathcal{L}_{\text{OoD}} > 0.5$. We compare these results with a Gaussian process state space model that uses a standard Matérn 3/2 kernel, trained with the same data from circular motions. As shown in table I, this model’s predictions remain in distribution less often. In

TABLE I
PERCENTAGE OF TIME STEPS IN WHICH $\mathcal{L}_{\text{OoD}} > 0.5$ ACROSS ALL TRAJECTORIES FOR EACH ENVIRONMENT.

	Walking	Rope	Rocky	Poking
Ours	1.8%	66.7%	64.4%	79.0%
GPSSM	87%	92.5%	98.7%	97.3%

TABLE II
DATA EFFICIENCY ANALYSIS.

	25%	50%	75%	100%
SE	-3.68 (0.09)	-3.72 (0.03)	-3.75 (0.01)	-3.47 (0.20)
Matern	-3.74 (0.03)	-3.71 (0.05)	-3.69 (0.08)	-3.58 (0.17)
Ours	-4.35 (0.04)	-4.42 (0.02)	-4.49 (0.04)	-4.54 (0.03)

addition, the model has difficulties predicting reliably the data it was trained on (i.e., first column). This indicates that for such a small dataset, informing the kernel is beneficial to yield more accurate predictions. We found the number of features $M = 1500$ to be a key factor in obtaining accurate state predictions. While a larger M did not yield better results, smaller M affected the quality of the long-term predictions, which made the OoD runtime monitor unreliable.

B. Data efficiency assessment

It is desirable to analyze the extent to which the domain-informed kernel increases data efficiency. To assess this, we collect 10 minutes of data in the same setting described above, but moving the robot through a set of randomly chosen waypoints across the room. We leave 10% of the dataset for testing and use the remaining portion as a training set. Then, we slice the training dataset in batches of increasing size (25%, 50%, 75%) and train the kernel of the informed GPSSM with each of them. As a performance metric we use the root mean squared error (RMSE) between the predictive mean of the GPSSM and the actual observations to assess the quality of the regression fit. We compare the performance of our model with a standard GPSSM that uses (i) a squared exponential kernel (SE) and (ii) a Matérn kernel. In table II, we see that our model systematically achieves a lower RMSE, which indicates that constructing the kernel by integrating existing data does increase data efficiency.

VI. CONCLUSION

In this paper, we have proposed (i) a general methodology to embed domain knowledge, given as a nominal model, into the kernel of a Gaussian process state-space model (GPSSM), and (ii) a runtime monitor for online out-of-training-distribution (OoD) detection using GPSSMs.

The resulting non-stationary kernel encapsulates domain knowledge using Fourier features, which are learned by minimizing an autoencoder loss using data from an implicit nominal model. Our results indicate that such a domain-informed kernel needs less training data to achieve the same regression quality as standard non-informed stationary kernels. In addition, we show that the proposed OoD runtime monitor, validated on a quadruped robot navigating in an indoor setting, can reliably detect previously unseen scenarios.

Whereas this work focuses on detecting OoD situations online, we plan to integrate the OoD runtime monitor with sampling-based model predictive control methods [43] for online decision-making, to enable safe navigation.

ACKNOWLEDGMENT

The authors of this paper express their gratitude to Jason J. Choi and Chams E. Mballo for insightful discussions, and to Prof. Koushil Sreenath for his support in providing computational resources.

REFERENCES

- [1] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” *arXiv preprint arXiv:2307.15818*, 2023.
- [2] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian, *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Conference on Robot Learning*, pp. 287–318, PMLR, 2023.
- [3] J. Nitsch, M. Itkina, R. Senanayake, J. Nieto, M. Schmidt, R. Siegwart, M. J. Kochenderfer, and C. Cadena, “Out-of-distribution detection for automotive perception,” in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 2938–2943, 2021.
- [4] M. Salem, G. Lakatos, F. Amirabdollahian, and K. Dautenhahn, “Towards safe and trustworthy social robots: Ethical challenges and practical issues,” in *Social Robotics: 7th International Conference*, pp. 584–593, Springer, 2015.
- [5] J. Yang, K. Zhou, Y. Li, and Z. Liu, “Generalized out-of-distribution detection: A survey,” *arXiv preprint arXiv:2110.11334*, 2021.
- [6] A. Zhang, R. T. McAllister, R. Calandra, Y. Gal, and S. Levine, “Learning invariant representations for reinforcement learning without reconstruction,” in *International Conference on Learning Representations*, 2020.
- [7] Y. Guo, K. Guo, X. Cao, T. Wu, and Y. Chang, “Out-of-distribution generalization of federated learning via implicit invariant relationships,” in *International Conference on Machine Learning*, PMLR, 2023.
- [8] S. S. G. Bagi, Z. Gharraee, O. Schulte, and M. Crowley, “Generative causal representation learning for out-of-distribution motion forecasting,” *arXiv preprint arXiv:2302.08635*, 2023.
- [9] M. Yang, L. Wang, C. Deng, and H. Zhang, “Bootstrap your own prior: Towards distribution-agnostic novel class discovery,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3459–3468, 2023.
- [10] A. De Silva, R. Ramesh, C. Priebe, P. Chaudhari, and J. T. Vogelstein, “The value of out-of-distribution data,” in *International Conference on Machine Learning*, pp. 7366–7389, PMLR, 2023.
- [11] A. Sharma, N. Azizan, and M. Pavone, “Sketching curvature for efficient out-of-distribution detection for deep neural networks,” in *Uncertainty in Artificial Intelligence*, pp. 1958–1967, PMLR, 2021.
- [12] R. Sinha, A. Sharma, S. Banerjee, T. Lew, R. Luo, S. M. Richards, Y. Sun, E. Schmerling, and M. Pavone, “A system-level view on out-of-distribution data in robotics,” *arXiv preprint arXiv:2212.14020*, 2022.
- [13] R. Luo, S. Zhao, J. Kuck, B. Ivanovic, S. Savarese, E. Schmerling, and M. Pavone, “Sample-efficient safety assurances using conformal prediction,” in *International Workshop on the Algorithmic Foundations of Robotics*, pp. 149–169, Springer, 2022.
- [14] A. Farid, D. Snyder, A. Z. Ren, and A. Majumdar, “Failure prediction with statistical guarantees for vision-based robot control,” *arXiv preprint arXiv:2202.05894*, 2022.
- [15] A. Farid, S. Veer, and A. Majumdar, “Task-driven out-of-distribution detection with statistical guarantees for robot learning,” in *Conference on Robot Learning*, pp. 970–980, PMLR, 2022.
- [16] A. Svensson, T. B. Schön, A. Solin, and S. Särkkä, “Nonlinear state space model identification using a regularized basis function expansion,” in *IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 481–484, 2015.
- [17] H. Salimbeni and M. Deisenroth, “Doubly stochastic variational inference for deep Gaussian processes,” *Advances in neural information processing systems*, vol. 30, 2017.
- [18] M. Buisson-Fenet, F. Solowjow, and S. Trimpe, “Actively learning Gaussian process dynamics,” in *Learning for dynamics and control*, pp. 5–15, PMLR, 2020.
- [19] R. Frigola, Y. Chen, and C. E. Rasmussen, “Variational Gaussian process state-space models,” *Advances in neural information processing systems*, vol. 27, 2014.
- [20] A. Svensson, A. Solin, S. Särkkä, and T. Schön, “Computationally efficient Bayesian learning of Gaussian process state space models,” in *Artificial Intelligence and Statistics*, pp. 213–221, PMLR, 2016.
- [21] 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.
- [22] D. Sun, M. J. Khojasteh, S. Shekhar, and C. Fan, “Uncertain-aware safe exploratory planning using Gaussian process and neural control contraction metric,” in *Learning for Dynamics and Control*, pp. 728–741, PMLR, 2021.
- [23] Z. Zhu, E. Biyik, and D. Sadigh, “Multi-agent safe planning with Gaussian processes,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6260–6267, 2020.
- [24] R. Wang and R. Yu, “Physics-guided deep learning for dynamical systems: A survey,” *arXiv preprint arXiv:2107.01272*, 2021.
- [25] A. Marco, P. Hennig, S. Schaal, and S. Trimpe, “On the design of LQR kernels for efficient controller learning,” in *IEEE 56th Annual Conference on Decision and Control*, pp. 5193–5200, IEEE, 2017.
- [26] M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur, “Gaussian processes and kernel methods: A review on connections and equivalences,” *arXiv preprint arXiv:1807.02582*, 2018.
- [27] A. Svensson and T. B. Schön, “A flexible state-space model for learning nonlinear dynamical systems,” *Automatica*, vol. 80, pp. 189–199, 2017.
- [28] A. Solin and S. Särkkä, “Hilbert space methods for reduced-rank Gaussian process regression,” *Statistics and Computing*, vol. 30, no. 2, pp. 419–446, 2020.
- [29] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [30] T. P. Minka, “Deriving quadrature rules from Gaussian processes,” tech. rep., Technical report, Statistics Department, Carnegie Mellon University, 2000.
- [31] M. P. Deisenroth, A. A. Faisal, and C. S. Ong, *Mathematics for machine learning*. Cambridge University Press, 2020.
- [32] 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.
- [33] A. Zygmund, *Trigonometric series*, vol. 1. Cambridge university press, 2002.
- [34] M. Ngom and O. Marin, “Fourier neural networks as function approximators and differential equation solvers,” *Statistical Analysis and Data Mining: The ASA Data Science Journal*, vol. 14, no. 6, pp. 647–661, 2021.
- [35] R. Bracewell, “The Fourier transform and its applications,” *American Journal of Physics*, vol. 34, no. 8, pp. 712–712, 1966.
- [36] Y. Katznelson, *An introduction to harmonic analysis*. Cambridge University Press, 2004.
- [37] A. D. Ialongo, M. Van Der Wilk, J. Hensman, and C. E. Rasmussen, “Overcoming mean-field approximations in recurrent Gaussian process models,” in *International Conference on Machine Learning*, pp. 2931–2940, PMLR, 2019.
- [38] J. Hensman, N. Durrande, A. Solin, *et al.*, “Variational Fourier features for Gaussian processes,” *Journal of Machine Learning Research*, vol. 18, no. 1, pp. 5537–5588, 2017.
- [39] A. Rahimi and B. Recht, “Random features for large-scale kernel machines,” *Advances in neural information processing systems*, vol. 20, 2007.
- [40] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, “Gpytorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration,” in *Advances in Neural Information Processing Systems*, 2018.
- [41] “Unitree robotics.” <https://www.unitree.com/>.
- [42] T. G. Molnar and A. D. Ames, “Safety-critical control with bounded inputs via reduced order models,” *arXiv preprint arXiv:2303.03247*, 2023.
- [43] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, “Aggressive driving with model predictive path integral control,” in *IEEE International Conference on Robotics and Automation*, pp. 1433–1440, 2016.