Hybrid Systems Laboratory

On-line out-of-distribution detection using simulation-informed deep Gaussian process state-space models

Alonso Marco and Claire J. Tomlin

Motivation: Quadrupedal navigation in uncertain terrains

▷ **Goal**

Detect out-of-distribution (OoD) environments on-the-fly [1]

▷ **Idea:** Compare predictions vs observations

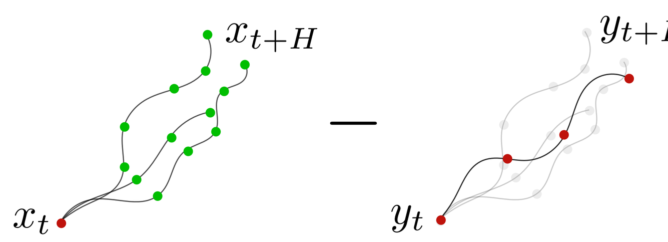

▷ **Challenges**

1. Well-calibrated uncertainties (no overconfidence): epistemic+aleatoric
2. On-line deployment requires fast predictions
3. Learning requires data-efficiency for re-training
4. OoD metric computationally efficient and probability-based

Approach

▷ 1. Represent real dynamics using a Gaussian process state-space model (GPSSM) [4]

$$x_{t+1} = f_{\text{real}}(x_t, u_t) + x_t$$
$$f_{\text{real}}(x_t, u_t) \sim \mathcal{GP}(0, k(\cdot,\cdot))$$

▷ 2. Use equivalent series expansion of a Gaussian process

Standard GP ✗	
Inference	$O(N^3)$
Predictions	$> O(HN^3)$

Karhunen–Loève expansion GP ✓	
Inference	$O(N)$
Predictions	$O(H)$

▷ 3. Increase data-efficiency by constructing a **simulation-informed prior**

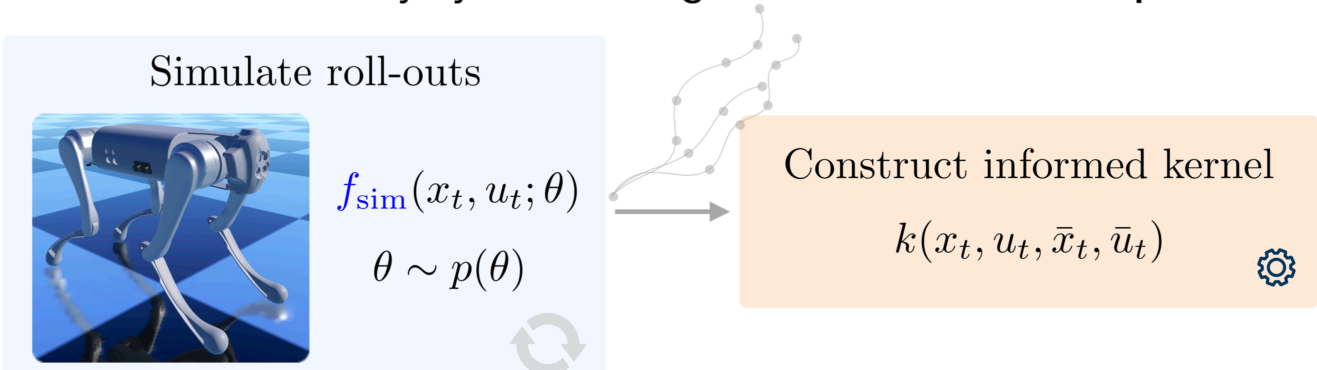

▷ 4.Compare the distribution over predicted states before and after observations

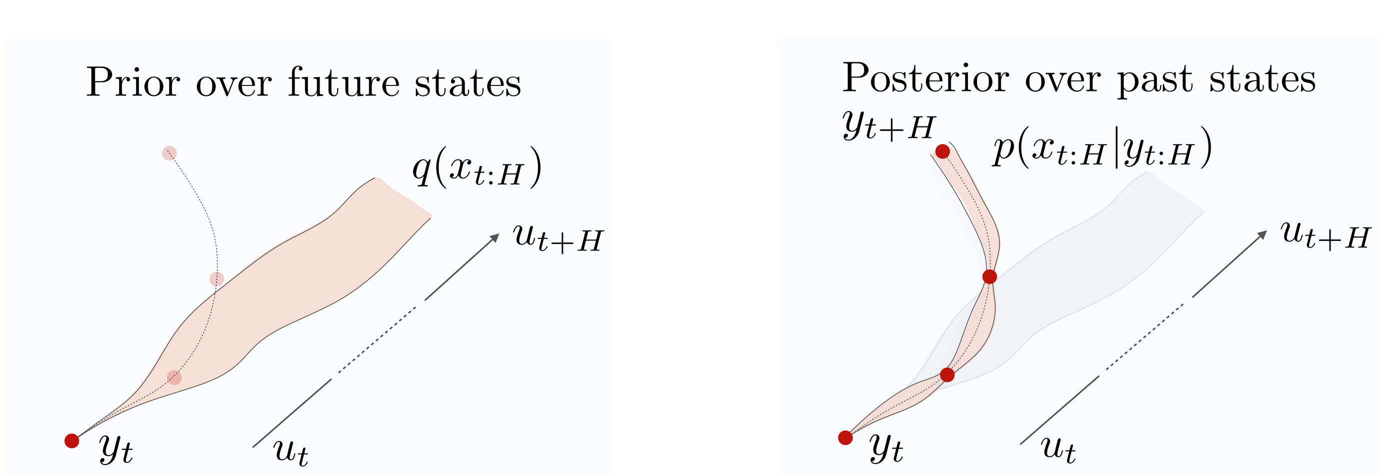

$$\mathcal{L}_{\text{OoD}}(y_{t:H}, u_{t:H}) = \text{KL}(q||p) \longrightarrow \text{ELBO}$$

Take Home

- ☑ Deployed on-line out-of-distribution detection on a real quadruped
- ☑ Improved long-term prediction capabilities by informing kernel with simulation data
- ☑ Proposed novel kernel design framework, simulator-agnostic

Karnhunen-Loève expansion of Gaussian process

▷ Bayesian linear model with deterministic features

$$f(x) = \sum_{j=1}^{M} \beta_j \phi_j(x) \qquad \beta_j \sim \mathcal{N}(m_j, \nu_j) \qquad \text{Cov}[\beta_i, \beta_j] = \begin{cases} \nu_j & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

▷ Samples of posterior dynamics are "callable"

$$\hat{f}(x) = \sum_{j=1}^{M} \hat{\beta}_j \phi_j(x)$$

▷ Posterior at cost $O(NM^3)$

$$\hat{\beta}|\mathcal{D} \sim \mathcal{N}(\mu, \Sigma)$$

Embedding prior information via Mercer kernel

▷ Kernel given as a finite sum of features [2]

$$k(x, \bar{x}) = \sum_{j=1}^{M} \nu_j \phi_j(x) \phi_j(\bar{x}) \ , \ \nu_j > 0 \quad \text{decreasing sequence}$$

▷ Architecture: Input embedding and Fourier features

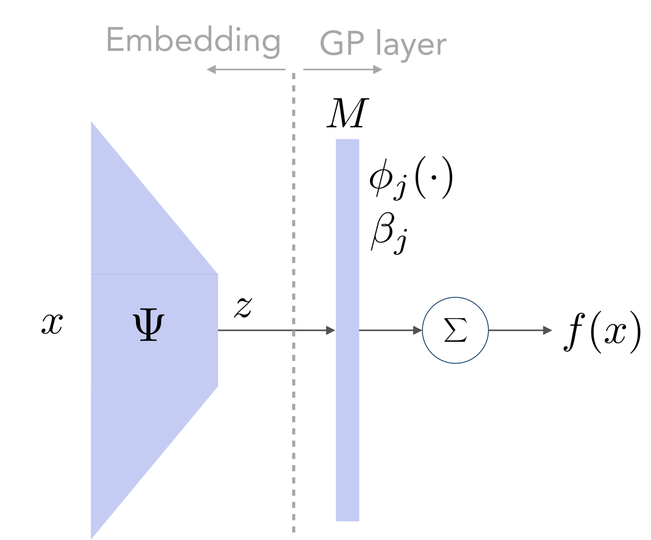

▷ Fourier features

$$\phi_j(x) = \cos(\omega_j^\top \Psi(x) + \varphi_j)$$

▷ Spectral density $\omega_j \sim S(\omega)$

▷ Prior variance proportional to wave amplitude [1] $\nu_j \propto S(\omega_j)$

▷ Phase $\varphi_j \sim \text{U}(-\pi, \pi)$

▷ Mean $m_j = 1/M$

▷ Train by moment-matching simulation data

$$\arg\min_{\Psi} \int_{x \in \mathcal{X}} ||\mathbb{E}[f_{\text{sim}}(x)] - \sum_j m_j \phi_j(x)|| + \lambda ||\mathbb{V}\text{ar}[f_{\text{sim}}(x)] - \sum_j \nu_j \phi_j^2(x)||$$

Simulation-informed priors are overconfident/stiff

▷ Fix by adding a standard kernel [3]

$$f_{\text{sim}}(x; \theta) = \theta x^2$$
$$\theta \sim \text{U}(-1, 1)$$

$$f_{\text{real}}(x) = \sum_j \beta_j \phi_j(x), \ \beta_j \sim \mathcal{N}(1/M, \nu/M) \ , \ \phi_j(x) = x^2$$

Informed kernel $k(x, \bar{x}) = \nu x^2 \bar{x}^2$

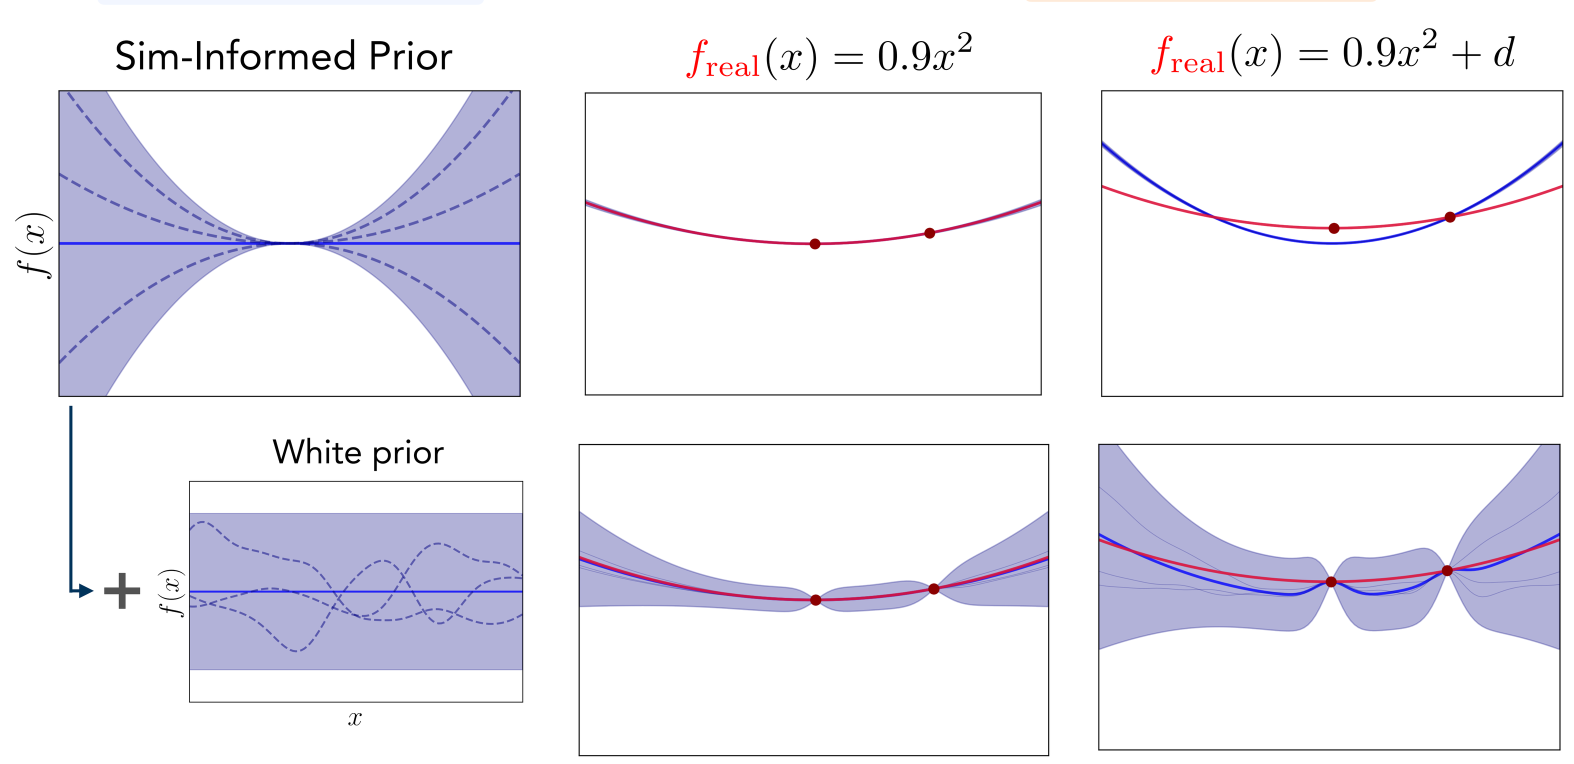

Results: Real quadruped detects OoD terrains

▷ **Simulation-informed Gaussian process state-space model**

$$f_d(x_t, u_t) = \beta_d^\top \Phi_d(x_t, u_t)$$
$$x_{t+1} \sim \mathcal{N}(f(x_t, u_t), Q)$$
$$y_t \sim \mathcal{N}(x_t, R)$$
$$f(\cdot) = [f_1(\cdot), \ldots, f_D(\cdot)]^\top$$

x_t: position, orientation

u_t: com velocity

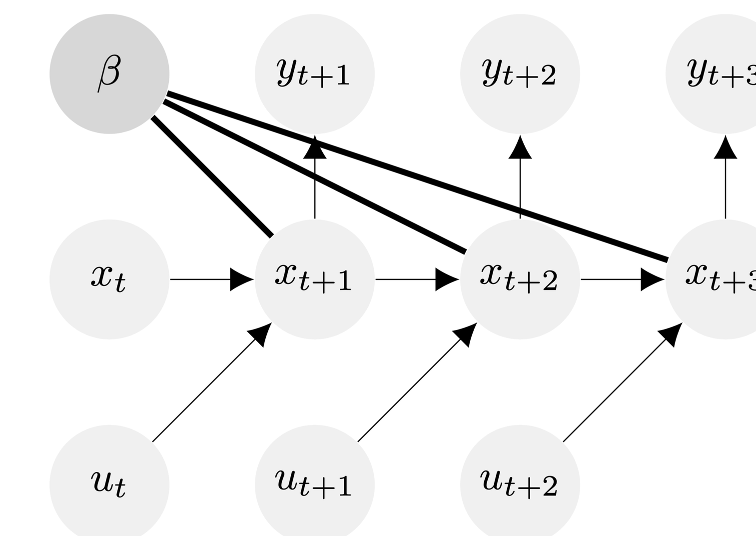

▷ **Training phase**

Kernel informed with walking circular trajectories

▷ **Test phase**

Robot deployed on a variety of terrains to test its OoD-detection capabilities

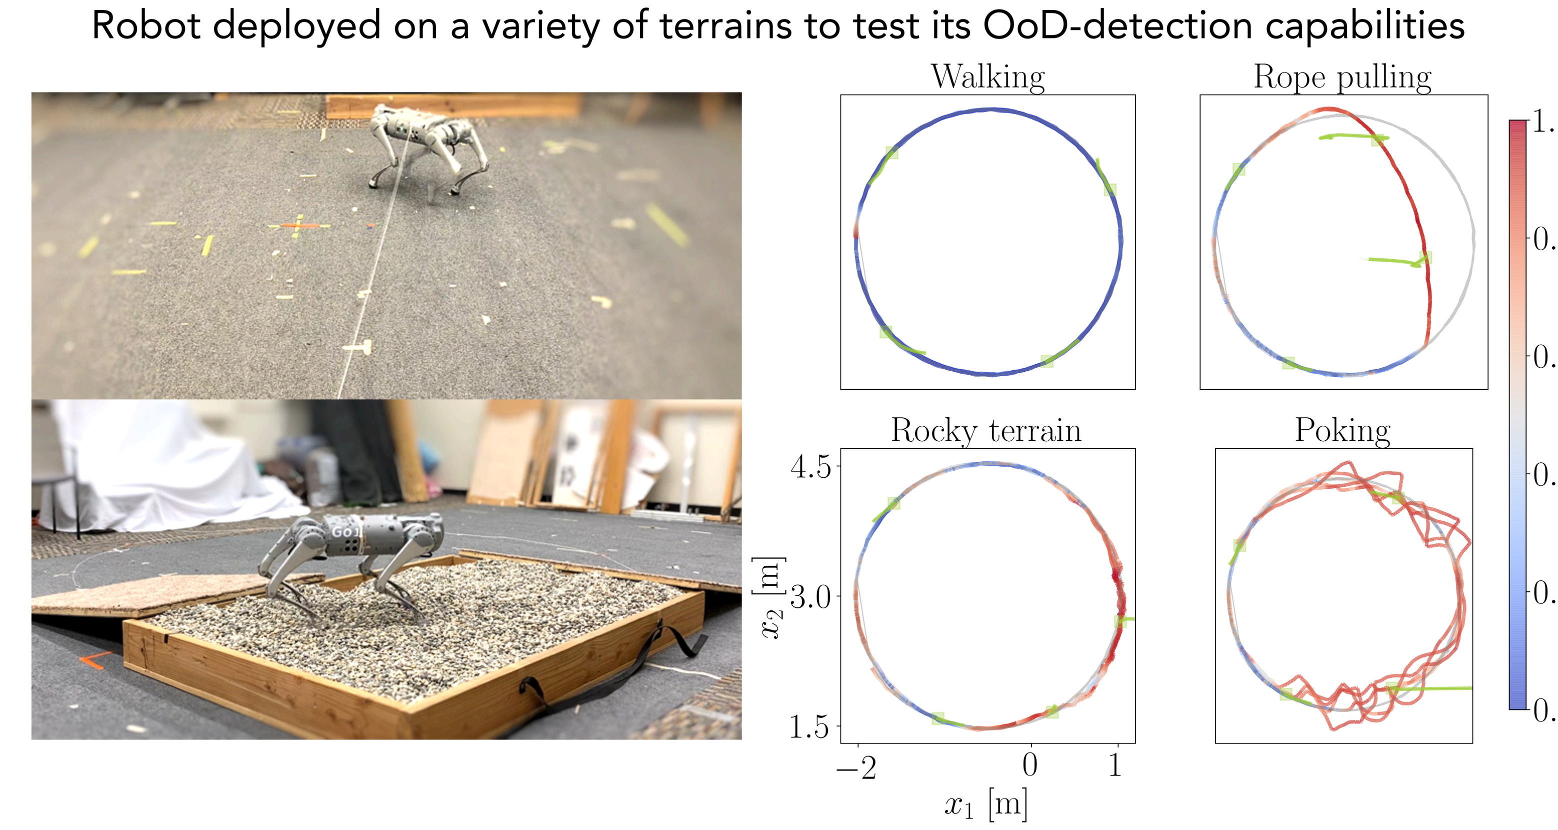

▷ **Empirical validation**

- ☑ **Well-calibrated uncertainties**: *If the observations don't match the predictions, it's because the environment is OoD, and not because the model is wrong*
- ☑ Our model consistently reflects OoD scenarios, outperforming GPSSM models with standard kernels

Frequency of out of distribution detection $\mathcal{L}_{\text{OoD}}(\cdot) > 0.5$

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

▷ **Future work**

- ☐ Integrate stochastic MPC and planning
- ☐ Use OoD detection to behave safely and trigger new model learning
- ☐ Learn a dictionary of models, one for each environment

References

[1] **Marco A.**, Morley E., Tomlin C. J. (2023). Out of Distribution Detection via Domain Informed Gaussian Process State Space Models. IEEE 62nd Conference on Decision and Control (CDC), (under review).

[2] Solin, A. and Särkkä, S., 2020. Hilbert space methods for reduced-rank Gaussian process regression. Statistics and Computing, 30(2), pp.419-446.

[3] **Marco, A.**, Hennig, P., Schaal, S. and Trimpe, S., 2017, December. On the design of LQR kernels for efficient controller learning. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC) (pp. 5193-5200). IEEE.

[4] Frigola, R., Chen, Y. and Rasmussen, C.E., 2014. Variational Gaussian process state-space models. Advances in neural information processing systems, 27.

Alonso Marco | Hybrid Systems Lab - University of California, Berkeley

The NASA University Leadership initiative (grant #80NSSC20M0163) provided funds to assist the authors with their research, but this article solely reflects the opinions and conclusions of its authors and not any NASA entity.