Probabilistic safety constraints for learned high relative degree system dynamics

Mohammad Javad Khojasteh*[1] Vikas Dhiman*[2] Massimo Franceschetti[2] Nikolay Atanasov[2]

[2]University of California San Diego, [1]California Institute of Technology

Objective

We study the problem of enforcing probabilistic safety when system dynamics are unknown and being learned from the samples,

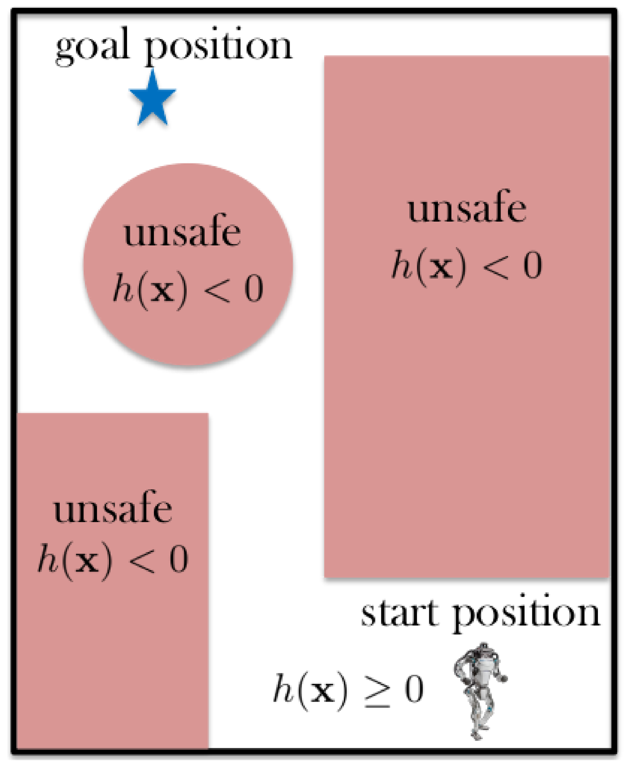

$$\min_{\mathbf{u}_k} \text{Task cost } (\mathbf{x}_k, \mathbf{u}_k)$$
$$\text{s.t. } \mathbb{P}(\text{ Safety } | \mathbf{x}_k, \mathbf{u}_k) \geq 1 - \text{ risk tolerance} \quad (1)$$

Contributions

Matrix Variate Gaussian Process We derive inference equations for Matrix Variate Gaussian Processes that preserve structure.

Safe-controller for higher relative degree systems We derive Cantelli-inequality based safety bound for higher relative degree systems and use it create QCQP based safe controller.

Inter-triggering time safety analysis We derive conditions to ensure safety between control computation times.

Notation

Symbol	Meaning
$\mathbf{x}_k \in \mathbb{R}^n$	System state at discrete time k
$\mathbf{x}(t)$	System state at cont. time t
$\mathbf{u} \in \mathbb{R}^m$	Control signal
$\underline{\mathbf{u}}$	$\triangleq (1; \mathbf{u})$
$f(\mathbf{x})$	drift term of system dynamics
$g(\mathbf{x})$	input gain term of system dynamics
$F(\mathbf{x})$	$\triangleq [f(\mathbf{x}), g(\mathbf{x})]$
vec(M)	Column-major vectorization of a matrix M
$h(\mathbf{x})$	Control barrier function defining the safety region as $h(\mathbf{x}) \geq 0$
$\pi_\epsilon(\mathbf{x})$	Reference controller whose trajectory we want to follow as close as possible

Problem formulation

For a control-affine system dynamics,

$$\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x})\mathbf{u} = [f(\mathbf{x}) + g(\mathbf{x})] \begin{bmatrix} 1 \\ \mathbf{u} \end{bmatrix} = F(\mathbf{x})\underline{\mathbf{u}},$$

assume the system dynamics $F(\mathbf{x})$ to be a Gaussian Process

$$\text{vec}(F(\mathbf{x})) \sim \mathcal{GP}(\text{vec}(\mathbf{M}_0(\mathbf{x})), \mathbf{K}_0(\mathbf{x}, \mathbf{x}')), \quad (2)$$

where vec$(F(\mathbf{x}))$ is column-major vectorization of $F(\mathbf{x})$. Design a safe controller with safe probability p_k,

$$\min_{\mathbf{u}_k \in \mathcal{U}} \quad \|\mathbf{u}_k - \pi_\epsilon(\mathbf{x}_k)\|$$
$$\text{s.t. } \mathbb{P}(\text{safety at all times}) \geq p_k \quad (3)$$

where the safety condition can be simply $\dot{h}(\mathbf{x}_k) > 0$ or a Control Barrier Condition [2], $\text{CBC}(\mathbf{x}, \mathbf{u}) \triangleq \dot{h}(\mathbf{x}) + \alpha h(\mathbf{x}) \geq 0$ with $\alpha > 0$.

Matrix Variate Gaussian Process

We define Matrix Variate Gaussian Process $\mathcal{MVGP}(\mathbf{M}(\mathbf{x}), \mathbf{A}, \mathbf{B}(\mathbf{x}, \mathbf{x}'))$ [4, 3].

$$\text{vec}(F(\mathbf{x})) \sim \mathcal{GP}(\text{vec}(\mathbf{M}_0(\mathbf{x})), \mathbf{B}_0(\mathbf{x}, \mathbf{x}') \otimes \mathbf{A})$$
$$\Leftrightarrow F(\mathbf{x}) \sim \mathcal{MVGP}(\mathbf{M}_0(\mathbf{x}), \mathbf{A}, \mathbf{B}_0(\mathbf{x}, \mathbf{x}')) \quad (4)$$

Advantages to alternative approaches:

Fewer parameters Only $(1+m)^2 + n^2$ parameters when learning $\mathbf{B}_0$ and $\mathbf{A}$ as compared to $(1+m)^2n^2$ parameters for $\mathbf{K}_0$. (m is control dimensions, n is state dimensions)

Captures correlation across output dimensions As compared to learning a GP per dimension, we capture correlation across output dimensions without excessive computational cost: $O((1+m)^3k^2) + O(k^3)$ vs $O((1+m)k^2) + O(k^3)$ where k is number of samples.

Preserves structure across inference Inference with k data samples of $\{\dot{\mathbf{x}}_i, \mathbf{x}_i, \mathbf{u}_i\}_{i=1}^k$, leads to another MVGP,

$$F_k(\mathbf{x}^*) \sim \mathcal{MVGP}(\mathbf{M}_k(\mathbf{x}^*), \mathbf{A}, \mathbf{B}_k(\mathbf{x}^*, \mathbf{x}^*)) \quad (5)$$

where $\mathbf{M}_k$ and $\mathbf{B}_k$ can be computed from the data samples.

Stochastic Control Barrier Condition

We consider the safety condition for system of relative degree r (defined as $\mathcal{L}_g\mathcal{L}_f^{r-1}h(\mathbf{x}) \neq 0$, but $\mathcal{L}_g\mathcal{L}_f^{j}h(\mathbf{x}) = 0$ for all $j = \{0, \dots, r-2\}$) as the exponential control barrier condition $\text{CBC}^{(r)}$ defined as [1],

$$\text{CBC}^{(r)}(\mathbf{x}, \mathbf{u}) := \mathcal{L}_f^{(r)}h(\mathbf{x}) + \mathcal{L}_g\mathcal{L}_f^{(r-1)}h(\mathbf{x})\mathbf{u} + \mathbf{k}_\alpha^\top \eta(\mathbf{x}),$$
$$\text{where } \eta(\mathbf{x}) \triangleq (h(\mathbf{x}); \mathcal{L}_f h(\mathbf{x}); \dots; \mathcal{L}_f^{(r-1)}h(\mathbf{x})) \quad (6)$$

We show that the mean and variance of $\text{CBC}^{(r)}(\mathbf{x}, \mathbf{u})$ are affine and quadratic in $\mathbf{u}$.

$$\mathbb{E}[\text{CBC}^{(r)}] = \left(\mathbb{E}[F(\mathbf{x})^\top \nabla \mathcal{L}_f^{(r-1)}h(\mathbf{x})] + \mathbb{E}[[\mathbf{k}_\alpha^\top \eta(\mathbf{x}), \mathbf{0}^\top]^\top]\right)^\top \underline{\mathbf{u}} \quad (7)$$

$$Var[\text{CBC}^{(r)}] = \underline{\mathbf{u}}^\top Var\left[\nabla \mathcal{L}_f^{(r-1)}h(\mathbf{x})^\top F(\mathbf{x}) + [\mathbf{k}_\alpha^\top \eta(\mathbf{x}), \mathbf{0}^\top]\right] \underline{\mathbf{u}} \quad (8)$$

Hence the safety condition $\mathbb{P}(\text{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k) \geq \zeta > 0) \geq \tilde{p}_k$ can be written as quadratic constraints using Cantelli's inequality and the controller thus becomes,

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi_\epsilon(\mathbf{x}_k)\|$$
$$\text{s.t.} \quad (\mathbb{E}[\text{CBC}_k^{(r)}] - \zeta)^2 \geq \frac{\tilde{p}_k}{1 - \tilde{p}_k} Var[\text{CBC}_k^{(r)}]$$
$$\mathbb{E}[\text{CBC}_k^{(r)}] - \zeta \geq 0 \quad (9)$$

While we derive closed form expression for $\mathbb{E}[\text{CBC}^{(r)}]$ and $Var[\text{CBC}^{(r)}]$ for $r = 1$ and $r = 2$, in general for $r \geq 3$ the mean and variance can be estimated using Monte Carlo estimators.

Inter-triggering time safety analysis

We assume sample trajectories from Gaussian Process dynamics are locally L_k-Lipchitz with large probability q_k, then we establish that

$$\mathbb{P}(\text{CBC}(\mathbf{x}_k, \mathbf{u}_k) \geq \zeta > 0 | \mathbf{x}_k, \mathbf{u}_k) \geq \tilde{p}_k$$
$$\text{and } \tau_k \leq \frac{1}{L_k} \ln(1 + \frac{L_k \zeta}{(\chi_k L_k + L_{\alpha \cdot h})\|\dot{\mathbf{x}}_k\|})$$
$$\implies \mathbb{P}(\text{CBC}(\mathbf{x}(t), \mathbf{u}_k) \geq 0) \geq p_k = \tilde{p}_k q_k \qquad \forall t \in [t_k, t_k + \tau_k) \quad (10)$$

Results

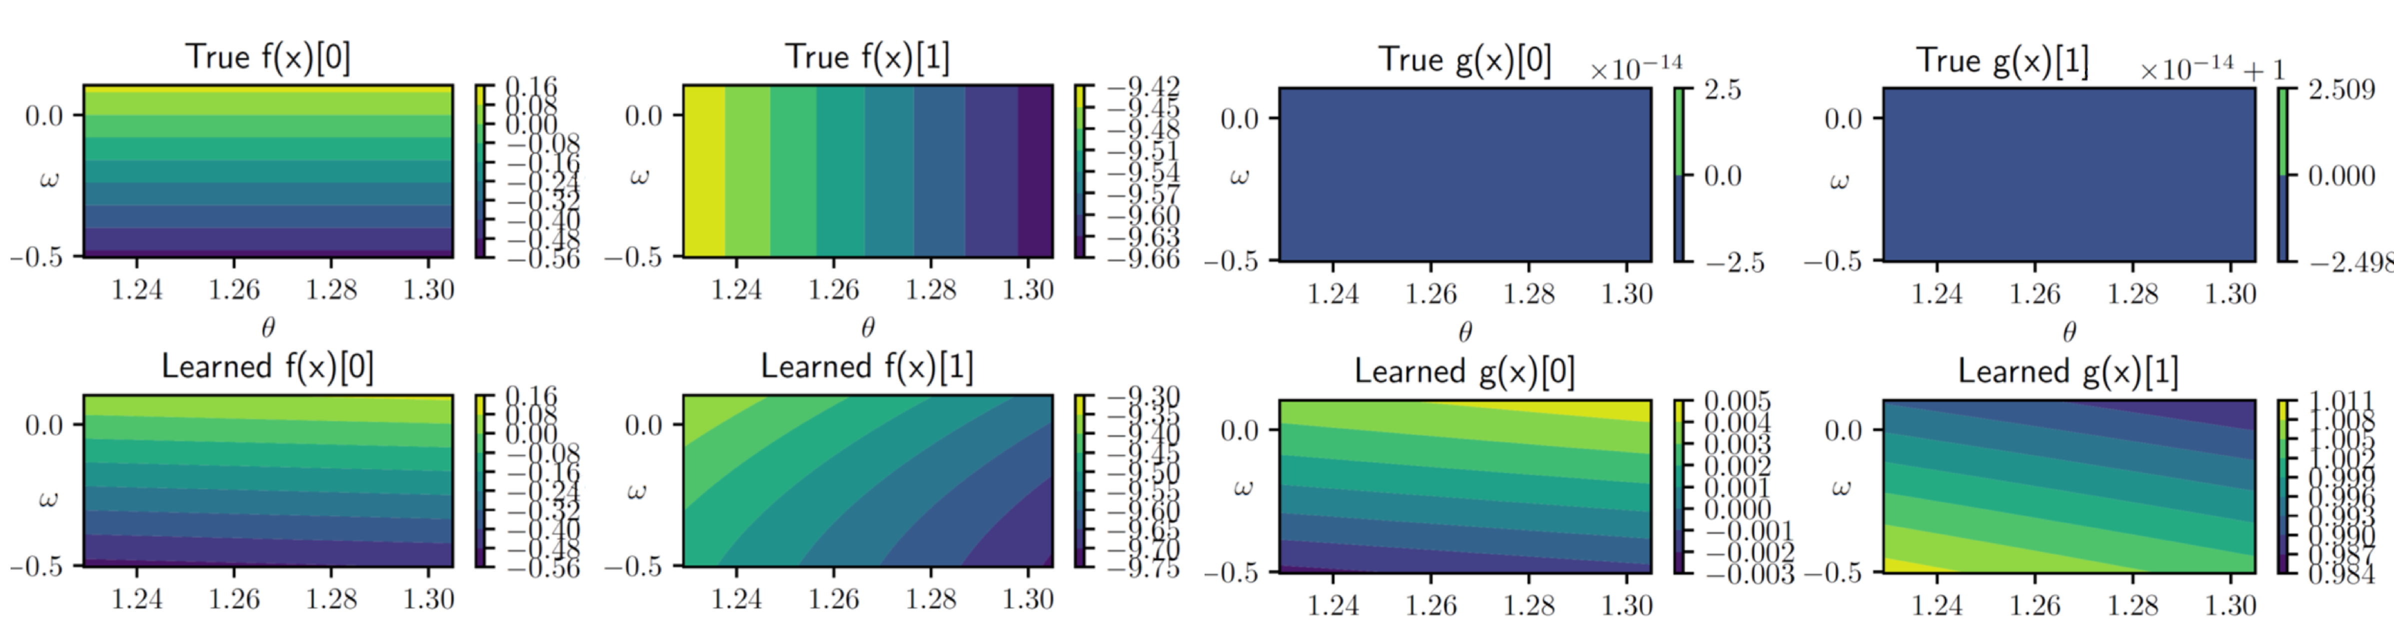

Figure 1:**Bottom row**: Learned vs true pendulum dynamics using matrix variate Gaussian Process regression

Experiment and Results

We evaluate the proposed approach on a pendulum with mass m and length l with state $\mathbf{x} = [\theta, \omega]$ and control-affine dynamics $f(\mathbf{x}) = [\omega, -\frac{g}{l}\sin(\theta)]$ and $g(\mathbf{x}) = [0, \frac{1}{ml}]$ as depicted in Fig 2. The control barrier function is chosen as $h(\mathbf{x}) = \cos(\Delta_{col}) - \cos(\theta - \theta_c)$.

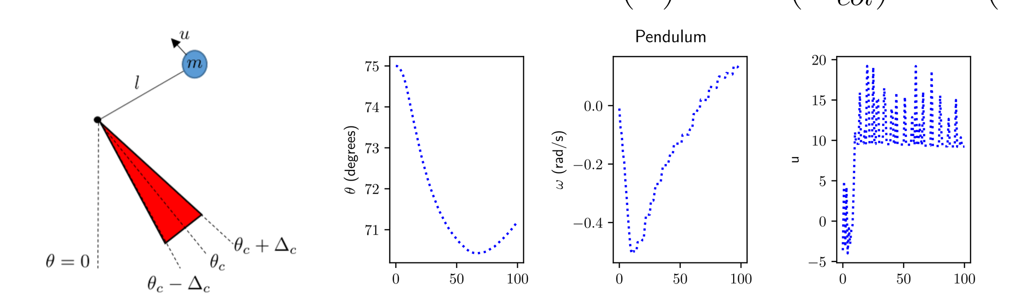

Figure 2:**Top left**: Pendulum simulation (left) with an unsafe (red) region. **Top right**: The pendulum trajectory (middle) resulting from the application of safe control inputs (right) is shown.

Conclusion and Ongoing work

- More experiments (closer to the Motivation).
- What if QCQP is non-convex?
- Entropy objective to pick optimal actions for reducing uncertainity.

References

[1] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada. Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431, June 2019.

[2] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE TAC*, 62(8):3861–3876, 2016.

[3] Shengyang Sun, Changyou Chen, and Lawrence Carin. Learning Structured Weight Uncertainty in Bayesian Neural Networks. In *AISTATS*, pages 1283–1292, 2017.

[4] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

Acknowledgements

We gratefully acknowledge support from NSF awards CNS-1446891 and ECCS-1917177, and support from ARL DCIST CRA W911NF-17-2-0181.

Contact Information

Paper URL: `https://arxiv.org/abs/1912.10116`

Mohammad Javad Khojasteh `mjkhojas@caltech.edu`

Vikas Dhiman `vdhiman@ucsd.edu`