Probabilistic safety constraints for learned high relative degree system dynamics

Mohammad Javad Khojasteh*, Vikas Dhiman*,
Massimo Franceschetti, and Nikolay Atanasov

Conference on Learning for Dynamics and Control (L4DC), 2020

* These authors contributed equally

Taking robots into the real world

Brittle hand-designed dynamics models work for lab operation but fail to account for the complexity and uncertainty of real-world operation

Learning for dynamics and control

Cyber

learning online relying on streaming data

Physical

control objectives and guaranteeing safe operation

Problem formulation

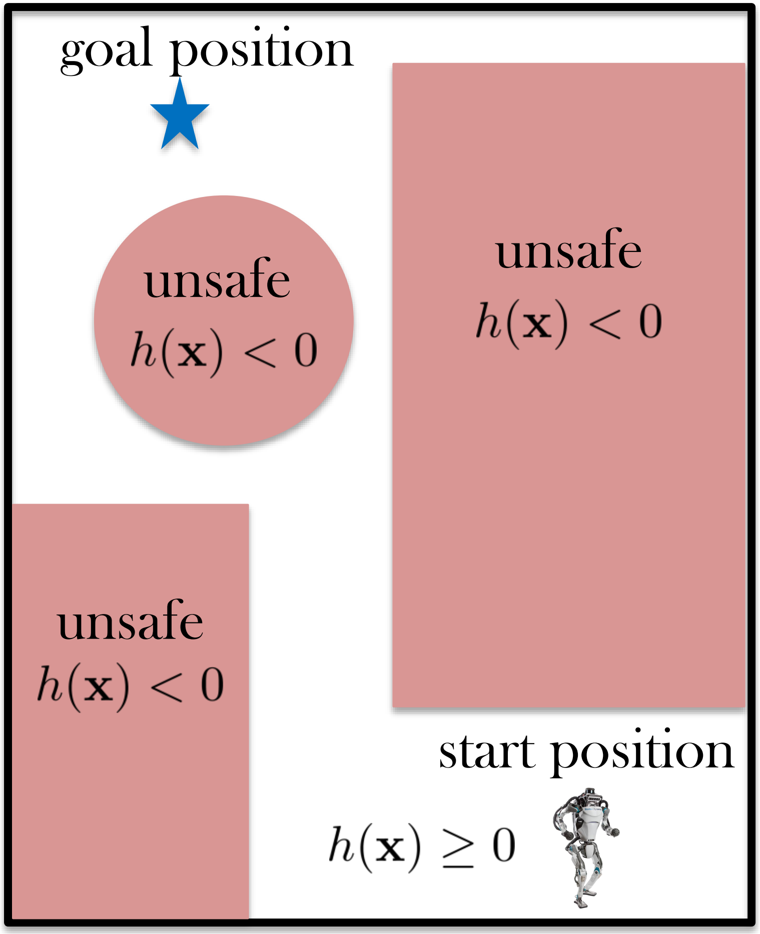

$$\begin{aligned}\dot{\mathbf{x}} &= f(\mathbf{x}) + g(\mathbf{x})\mathbf{u} \\ &= \begin{bmatrix} f(\mathbf{x}) & g(\mathbf{x}) \end{bmatrix} \begin{bmatrix} 1 \\ \mathbf{u} \end{bmatrix} \\ &= F(\mathbf{x})\underline{\mathbf{u}}\end{aligned}$$

drift term $f : \mathbb{R}^n \to \mathbb{R}^n$

input gain $g : \mathbb{R}^n \to \mathbb{R}^{n \times m}$

We study the problem of enforcing probabilistic safety when f and g are unknown

Problem formulation

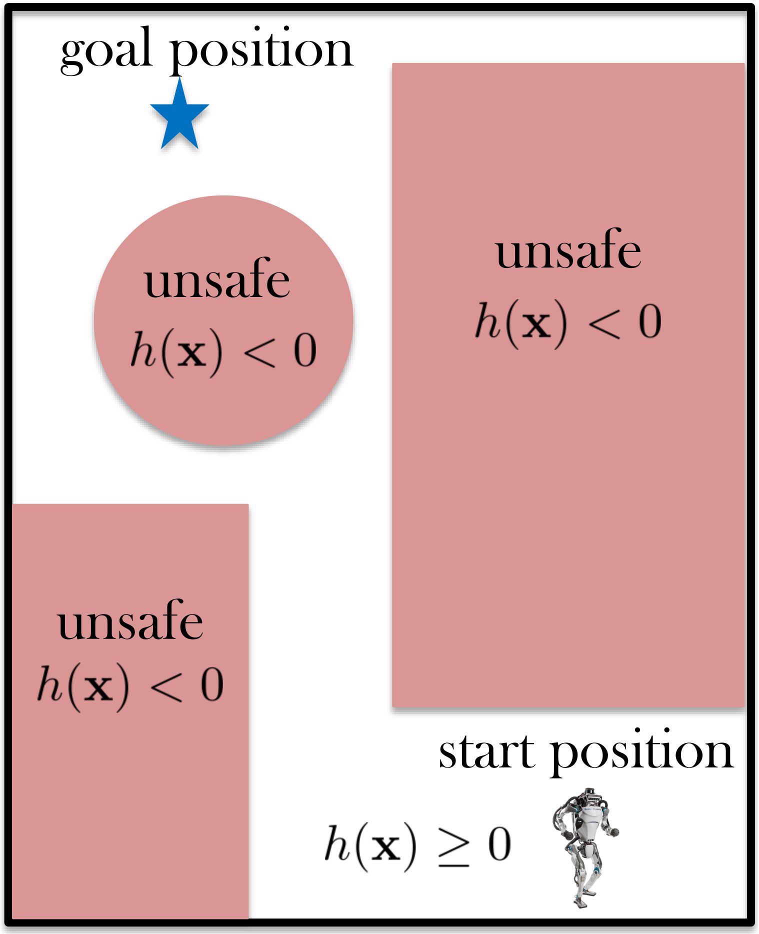

$$\dot{\mathbf{x}} = F(\mathbf{x})\underline{\mathbf{u}}$$

$$vec(F(\mathbf{x})) \sim \mathcal{GP}\left(vec(\mathbf{M}_0(\mathbf{x})), \mathbf{K}_0(\mathbf{x}, \mathbf{x}')\right)$$

baseline control policy

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi(\mathbf{x}_k)\|$$

$$\text{s.t.} \quad \mathbb{P}(\text{safety}) \geq p_k$$

user-specified risk tolerance

Approach

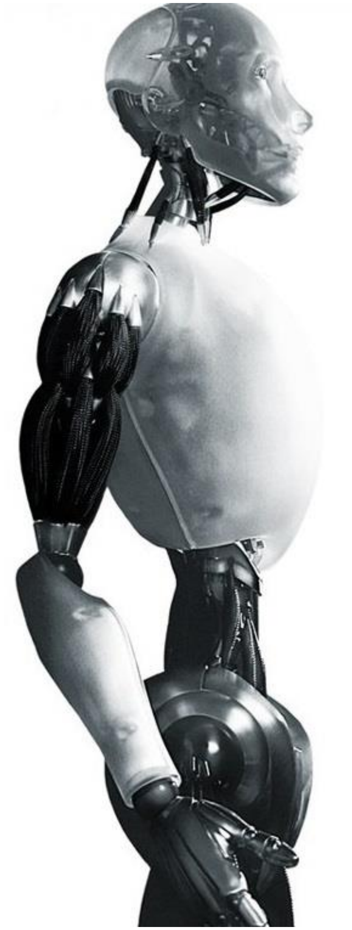

1. Bayesian learning
2. Propagate uncertainty to the safety condition
3. Self-triggered control: extension to continous time
4. Extension to higher relative degree systems

Gaussian processes for machine learning

$$\dot{\mathbf{x}} = F(\mathbf{x})\underline{\mathbf{u}}$$

$$vec(F(\mathbf{x})) \sim \mathcal{GP}\left(vec(\mathbf{M}_0(\mathbf{x})), \mathbf{K}_0(\mathbf{x}, \mathbf{x}')\right)$$

The controller observes
$$\mathbf{X}_{1:k} := [\mathbf{x}(t_1), \ldots, \mathbf{x}(t_k)]$$
$$\mathbf{U}_{1:k} := [\mathbf{u}(t_1), \ldots, \mathbf{u}(t_k)]$$
without noise,

but the measurements $\dot{\mathbf{X}}_{1:k} = [\dot{\mathbf{x}}(t_1), \ldots, \dot{\mathbf{x}}(t_k)]$ might be noisy.

In general, there may be a correlation among different components of f and g.

Thus, we need to develop an efficient factorization of $\mathbf{K}_0(\mathbf{x}, \mathbf{x}')$.

Matrix variate Gaussian processes (MVGP)

$$vec(F(\mathbf{x})) \sim \mathcal{GP}\left(vec(\mathbf{M}_0(\mathbf{x})), \underline{\mathbf{K}_0(\mathbf{x}, \mathbf{x}')}\right)$$

$$\mathbf{B}_0(\mathbf{x}, \mathbf{x}') \otimes \mathbf{A}$$ → Louizos and Welling (ICML 2016)
Sun et al. (AISTATS 2017)

The above parameterization is efficient because we need to learn smaller matrices $\mathbf{B}_0(\mathbf{x}, \mathbf{x}') \in \mathbb{R}^{(m+1)\times(m+1)}$ and $\mathbf{A} \in \mathbb{R}^{n\times n}$. Also, this parameterization preserves its structure during inference.

Inference

$$vec(F(\mathbf{x}_*)) \sim \mathcal{GP}(vec(\mathbf{M}_k(\mathbf{x}_*)), \mathbf{B}_k(\mathbf{x}_*, \mathbf{x}'_*) \otimes \mathbf{A})$$

$$F(\mathbf{x}_*)\underline{\mathbf{u}}_* = f(\mathbf{x}_*) + g(\mathbf{x}_*)\mathbf{u}_* \sim \mathcal{GP}(\mathbf{M}_k(\mathbf{x}_*)\underline{\mathbf{u}}_*, \underline{\mathbf{u}}_*^\top \mathbf{B}_k(\mathbf{x}_*, \mathbf{x}'_*)\underline{\mathbf{u}}_* \otimes \mathbf{A})$$

$\mathbf{M}_k(\mathbf{x}_*)$ and $\mathbf{B}_k(\mathbf{x}_*, \mathbf{x}'_*)$ are calculated in our paper

Two alternative approaches

1. Develop a decoupled GP regression per system dimension:

   Does not model the dependencies among different components of f and g

   Inference computational complexity:

   decoupled GP $O((1+m)k^2)+O(k^3)$ MVGP $O((1+m)^3k^2)+O(k^3)$

2. Coregionalization models [Alvarez et al. (FTML 2012)]:

$$\mathbf{K}_0(\mathbf{x},\mathbf{x}') = \boldsymbol{\Sigma}\underbrace{\kappa_0(\mathbf{x},\mathbf{x}')}_{\text{scalar state-dependent kernel}}$$

The nice matrix-times-scalar-kernel structure is not preserved in the posterior

Approach

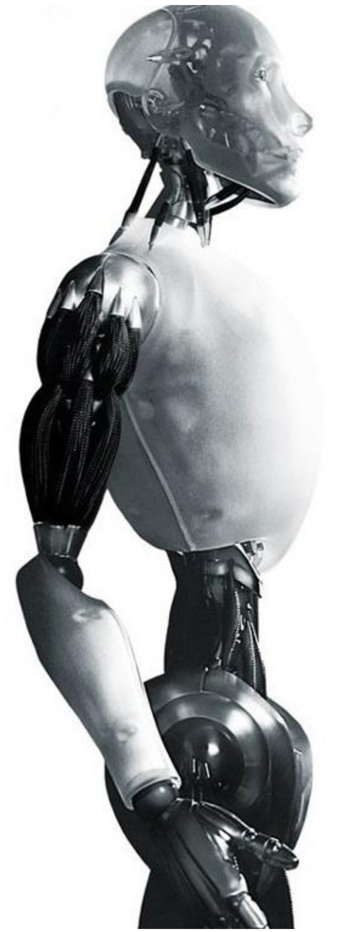

1. Bayesian learning
2. Propagate uncertainty to the safety condition
3. Self-triggered control: extension to continous time
4. Extension to higher relative degree systems

Control Barrier Functions (CBF)

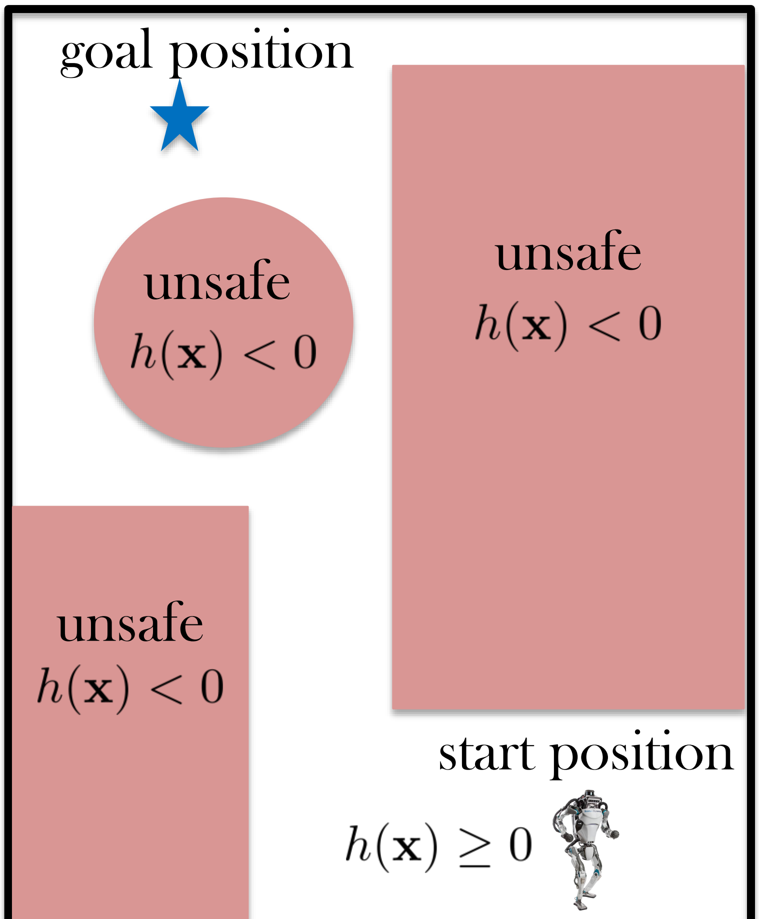

Previously, CBF are used to dynamically enforce the safety for known dynamics

Ames et al. (ECC 2019)

Control Barrier Condition (CBC)

$$\mathrm{CBC}(\mathbf{x}, \mathbf{u}) := \mathcal{L}_f h(\mathbf{x}) + \mathcal{L}_g h(\mathbf{x})\mathbf{u} + \alpha h(\mathbf{x}) \geq 0$$

$\nabla_{\mathbf{x}} h(\mathbf{x}) F(\mathbf{x}) \underline{\mathbf{u}}$ $\qquad$ $\alpha > 0$

A lower bound on the derivative

Uncertainity propagation to CBC

$$\mathrm{CBC}(\mathbf{x}, \mathbf{u}) = \underline{\mathcal{L}_f h(\mathbf{x}) + \mathcal{L}_g h(\mathbf{x})\mathbf{u}} + \underline{\alpha} h(\mathbf{x})$$

$$\nabla_{\mathbf{x}} h(\mathbf{x}) F(\mathbf{x}) \underline{\mathbf{u}} \qquad \alpha > 0$$

$$vec(F(\mathbf{x}_*)) \sim \mathcal{GP}(vec(\mathbf{M}_k(\mathbf{x}_*)), \mathbf{B}_k(\mathbf{x}_*, \mathbf{x}_*') \otimes \mathbf{A})$$

We have shown given $\mathbf{x}_k$ and $\mathbf{u}_k$, $\mathrm{CBC}(\mathbf{x}_k, \mathbf{u}_k)$ is a Gaussian random variable with the following parameters

$$\mathbb{E}[\mathrm{CBC}_k] = \nabla_{\mathbf{x}} h(\mathbf{x}_k)^\top \mathbf{M}_k(\mathbf{x}_k) \underline{\mathbf{u}}_k + \alpha h(\mathbf{x}_k)$$

$$\mathrm{Var}[\mathrm{CBC}_k] = \underline{\mathbf{u}}_k^\top \mathbf{B}_k(\mathbf{x}_k, \mathbf{x}_k) \underline{\mathbf{u}}_k \nabla_{\mathbf{x}} h(\mathbf{x}_k)^\top \mathbf{A} \nabla_{\mathbf{x}} h(\mathbf{x}_k)$$

Note: mean and variance are Affine and Quadratic in $\mathbf{u}$ respectively.

Deterministic condition for controller

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi(\mathbf{x}_k)\|$$

$$\text{s.t.} \quad \mathbb{P}(\text{CBC}(\mathbf{x}_k, \mathbf{u}_k) \geq \zeta > 0 | \mathbf{x}_k, \mathbf{u}_k) \geq \tilde{p}_k$$

$$\mathbb{E}[\text{CBC}(\mathbf{x}_k, \mathbf{u}_k)] - \zeta)^2 \geq 2\text{Var}[\text{CBC}(\mathbf{x}_k, \mathbf{u}_k)] \; (\text{erf}^{-1}(1 - 2\tilde{p}_k))^2$$

$$\mathbb{E}[\text{CBC}(\mathbf{x}_k, \mathbf{u}_k)] - \zeta \geq 0$$

A safe optimization-based controller which is a Quadratically Constrained Quadratic Program (QCQP)

Approach

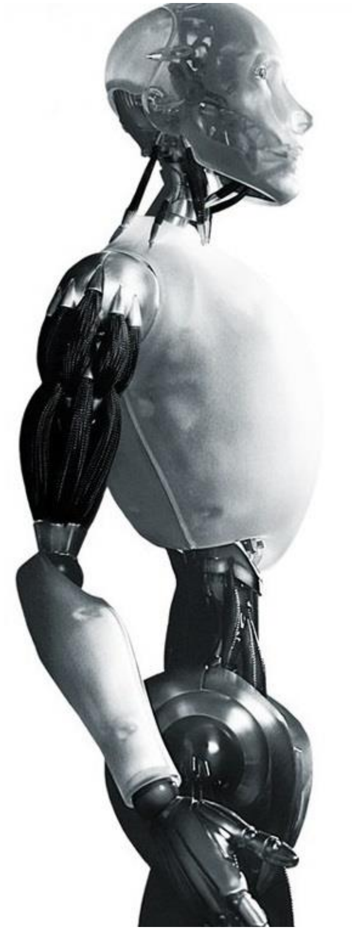

1. Bayesian learning
2. Propagate uncertainty to the safety condition
3. Self-triggered control: extension to continous time
4. Extension to higher relative degree systems

Safety beyond triggering times

Safety at triggering times

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi(\mathbf{x}_k)\|$$

$$\text{s.t.} \quad \mathbb{P}(\mathrm{CBC}(\mathbf{x}_k, \mathbf{u}_k) \geq \zeta > 0 | \mathbf{x}_k, \mathbf{u}_k) \geq \tilde{p}_k$$

Safety during the inter-triggering times

$\mathbf{u}(t) \equiv \mathbf{u}_k$ zero-order hold (ZOH) control mechanism $\forall t \in [t_k, t_k + \tau_k)$

$\tau_k = ?$ $\quad \mathbb{P}(\mathrm{CBC}(\mathbf{x}(t), \mathbf{u}_k) \geq 0) \geq p_k \quad \forall t \in [t_k, t_k + \tau_k)$

Self-triggered Control with Probabilistic Safety Constraints

We assume the sample paths of the GP used to model the dynamics are locally Lipschitz with sufficiently large probability q_k

This assumption is valid for a large class of GPs, e.g., squared exponential and some Matérn kernels → Srinivas et al. (TIT 2012) Shekhar and Javidi (EJS 2018)

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi(\mathbf{x}_k)\|$$

$$\text{s.t.} \quad \mathbb{P}(\mathrm{CBC}(\mathbf{x}_k, \mathbf{u}_k) \geq \zeta > 0 | \mathbf{x}_k, \mathbf{u}_k) \geq \tilde{p}_k$$

+

$$\tau_k \leq \frac{1}{L_k} \ln\left(1 + \frac{L_k \zeta}{(\chi_k L_k + L_{\alpha \circ h})\|\dot{\mathbf{x}}_k\|}\right)$$

→

$$\mathbb{P}(\mathrm{CBC}(\mathbf{x}(t), \mathbf{u}_k) \geq 0) \geq p_k = \tilde{p}_k q_k$$

$$\forall t \in [t_k, t_k + \tau_k)$$

The parameters are calculated in our paper

Approach

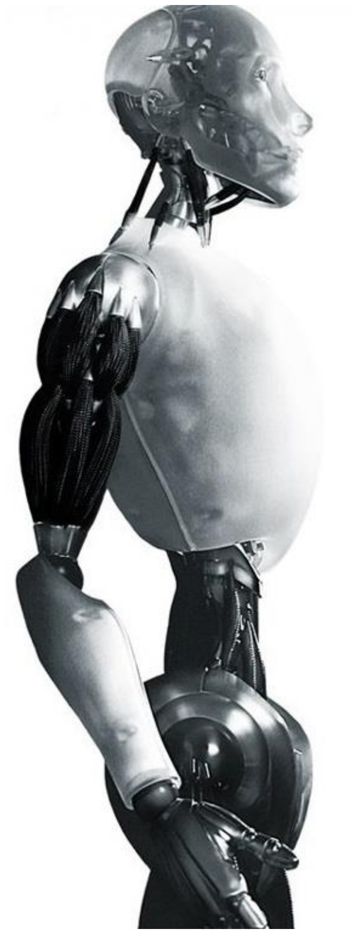

1. Bayesian learning
2. Propagate uncertainty to the safety condition
3. Self-triggered control: extension to continous time
4. Extension to higher relative degree systems

Higher relative degree CBFs

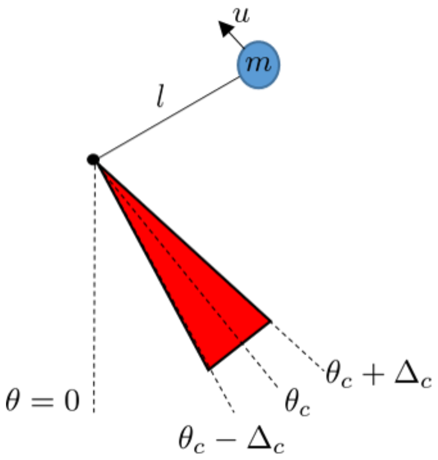

$$\mathbf{x} = [\theta, \omega]^\top$$

$$\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}$$

$$f(\mathbf{x}) = \left[\omega, -\tfrac{g}{l}\sin(\theta)\right]^\top \quad g(\mathbf{x}) = \left[0, \tfrac{1}{ml}\right]^\top$$

We want to avoid a radial region $[\theta_c - \Delta_c, \theta_c + \Delta_c]$

CBF: $h(\mathbf{x}) = \cos(\Delta_c) - \cos(\theta - \theta_c)$

Notice $\mathcal{L}_g h(\mathbf{x}) = \nabla h(\mathbf{x}) g(\mathbf{x}) = 0$

$\mathrm{CBC}(\mathbf{x}, \mathbf{u}) = \mathcal{L}_f h(\mathbf{x}) + \mathcal{L}_g h(\mathbf{x})\mathbf{u} + \alpha h(\mathbf{x})$ is independent of $\mathbf{u}$

Exponential Control Barrier Functions (ECBF)

Let $r \geq 1$ be the relative degree of $h(\mathbf{x})$, that is, $\mathcal{L}_g \mathcal{L}_f^{(r-1)} h(\mathbf{x}) \neq 0$ and $\mathcal{L}_g \mathcal{L}_f^{(k-1)} h(\mathbf{x}) = 0,\ \forall k \in \{1, \ldots, r-2\}$.

ECBC:

$$\mathrm{CBC}^{(r)}(\mathbf{x}, \mathbf{u}) := \mathcal{L}_f^{(r)} h(\mathbf{x}) + \mathcal{L}_g \mathcal{L}_f^{(r-1)} h(\mathbf{x}) \mathbf{u} + K_\alpha \begin{bmatrix} h(\mathbf{x}) \\ \mathcal{L}_f h(\mathbf{x}) \\ \vdots \\ \mathcal{L}_f^{(r-1)} h(\mathbf{x}) \end{bmatrix}$$

If K_α is chosen appropriately, $\mathrm{CBC}^{(r)} \geq 0$ enforce the safety for known dynamics. → Ames et al. (ECC 2019)
Nguyen and Sreenath (ACC 2016)

Chance constraint over ECBC

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi(\mathbf{x}_k)\|$$

$$\text{s.t.} \quad \mathbb{P}(\text{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k) \geq \zeta > 0 | \mathbf{x}_k, \mathbf{u}_k) \geq \tilde{p}_k$$

Cantelli's inequality

$$(\mathbb{E}[\text{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)] - \zeta)^2 \geq \frac{\tilde{p}_k}{1-\tilde{p}_k} \text{Var}[\text{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)]$$

$$\mathbb{E}[\text{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)] - \zeta \geq 0$$

A safe optimization-based controller which is a Quadratically Constrained Quadratic Program (QCQP)

Safe controller using ECBF

$$\min_{\mathbf{u}_k \in \mathcal{U}} \|\mathbf{u}_k - \pi(\mathbf{x}_k)\|$$

$$\text{s.t.} \quad (\mathbb{E}[\mathrm{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)] - \zeta)^2 \geq \frac{\tilde{p}_k}{1-\tilde{p}_k} \mathrm{Var}[\mathrm{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)]$$

$$\mathbb{E}[\mathrm{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)] - \zeta \geq 0$$

Solving this program requires the knowledge of the mean and variance of

$$\mathrm{CBC}^{(r)}(\mathbf{x}_k, \mathbf{u}_k)$$

In general, Monte Carlo sampling could be used to estimate these quantities.

We also explicitly quantified them in our paper for relative-degree-two systems. Bipedal and car-like robots are examples of these systems.

Toy example

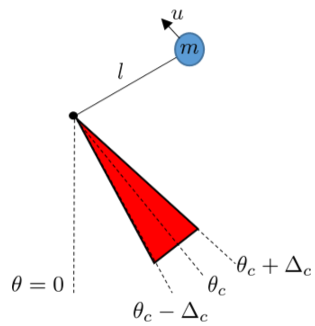

$$\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x})\mathbf{u}$$

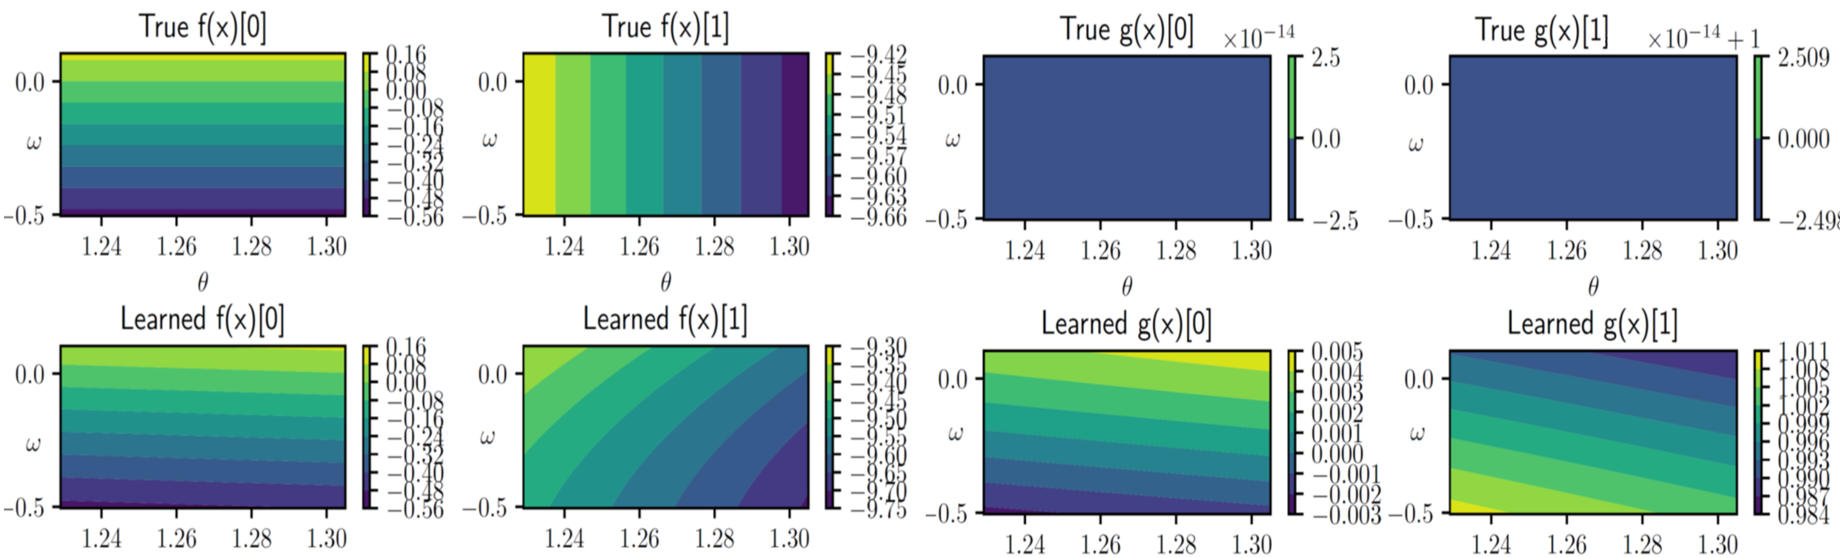

Thank You. Questions?

Paper URL: arxiv.org/abs/1912.10116

Mohammad Javad Khojasteh *
Caltech
mjkhojas@caltech.edu

Vikas Dhiman *
UC San Diego
vdhiman@ucsd.edu

Massimo Franceschetti
UC San Diego
massimo@ece.ucsd.edu

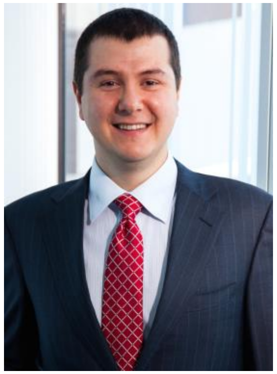

Nikolay Atanasov
UC San Diego
natanasov@eng.ucsd.edu

* These authors contributed equally