Learning HJ Solutions for Smooth and Tunable CBFs

Matthew Kim, William Sharpless, Hyun Joe Jeong, Sander Tonkens, Somil Bansal, Sylvia Herbert • (UC San Diego + Stanford)

Motivation

Control barrier functions (CBFs) are hard to construct.

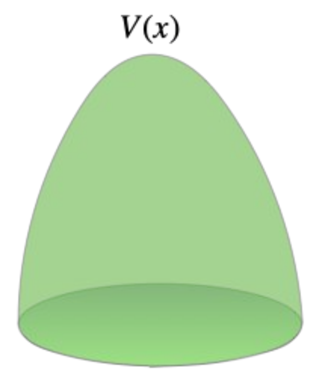

$$\max_{u \in U} \nabla B(x) \cdot f(x,u) \geq -\alpha(B(x)),$$

$$\forall x \in C,\ C = \{x \in \mathbb{R}^n : B(x) \geq 0\}$$

Hamilton-Jacobi reachability (HJR) is limited by the curse of dimensionality.

$$0 = \min\left\{l(x) - V(x,t),\ \frac{\partial V}{\partial t} + H(x, \nabla_x V, t)\right\},\quad V(x, t_f) = l(x)$$

HJR also yields jerky control behavior when used with the **CBF-QP**.

$$u^*(x) = \arg\min_{u \in \mathcal{U}} \|u - u_{\text{nom}}(x)\|_2^2$$

Control barrier value functions (CBVFs) [Choi 2021] can be made via HJR. But this is subjected to dimensional limits and non-differentiability.

$$V_\gamma(x,t) := \sup_{u \in \mathbb{U}} \min_{s \in [t, t_f]} e^{\gamma(s-t)} l(x(s))$$

$$0 = \min\left\{l(x) - V_\gamma(x,t),\ \frac{\partial V_\gamma}{\partial t} + H(x, \nabla_x V_\gamma, t) + \gamma V_\gamma(x,t)\right\},\quad V_\gamma(x, t_f) = l(x)$$

DeepReach [Bansal 2019] is a PINN used to learn the HJ-PDE viscosity solution.

$$\dot{V} + \min\{0, H(x, \nabla_x V, \tau)\} = 0$$
$$V(x, t_f) = J_T(x)$$

HJB-PDE

HJB-PINN
DeepReach

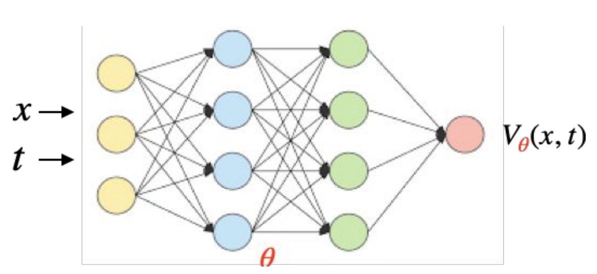

$$\mathcal{L}_{PDE}(\theta) = \mathbb{E}_{\mathbb{X},\mathbb{T}}\left[V_\theta(x,t) + H(x, \nabla_x V_\theta, t)\right]$$

Contribution: Learn CBFs via HJR

- **Structured construction** and **differentiable.**
- **Curse of complexity**, not dimensionality.
- Parameterized discount factor adds **tunability.**

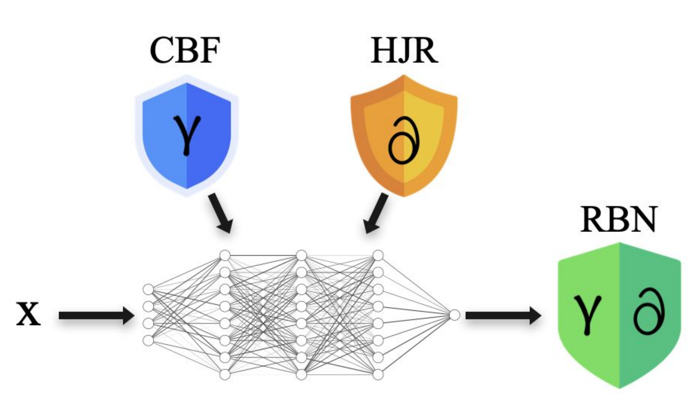

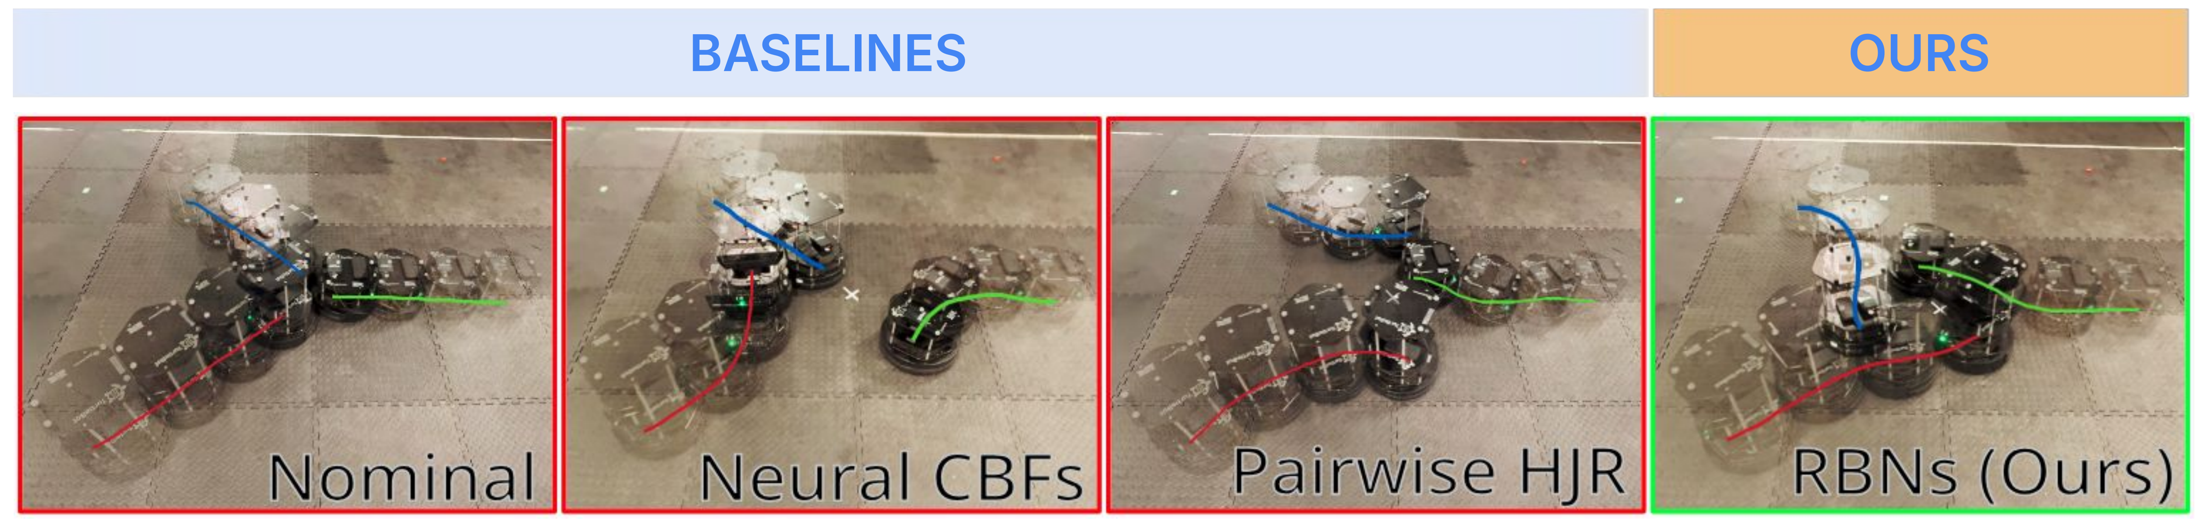

Reachability Barrier Networks (RBNs)

$$\mathcal{L}(\theta) = \mathbb{E}_{x,t,\gamma}\left[\left\| \min\left\{l(x) - V_\theta(x,t,\gamma),\ \frac{\partial V_\theta}{\partial t} + H(x, \nabla_x V_\theta, t) + \gamma V_\theta(x,t,\gamma)\right\}\right\|\right]$$

$$l(x) + (t_f - t) \cdot \mathrm{NN}_\theta(x,t,\gamma)$$

We replace the value function with a learned surrogate, enabling **smooth** CBF approximations and **tunable** γ-based control.

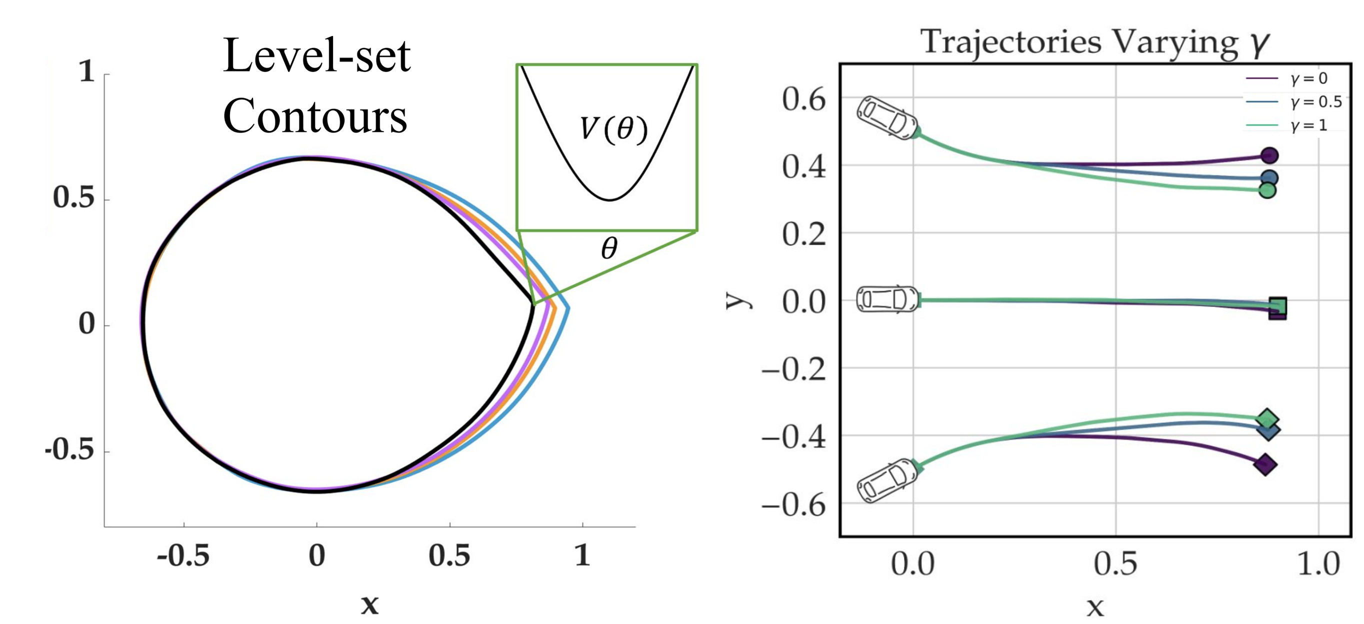

Low γ (more conservative) → High γ (less conservative / more aggressive safety).

Conformal Prediction for Probabilistic Safety

Estimate upper-bound:

$$\mathbb{P}\left(V(x,t,\gamma) \leq 0 \mid V_\theta(x,t,\gamma) > \delta\right) < \epsilon$$

Shift value function:

$$V_\theta^\delta(x,t,\gamma) = V_\theta(x,t,\gamma) - \delta$$

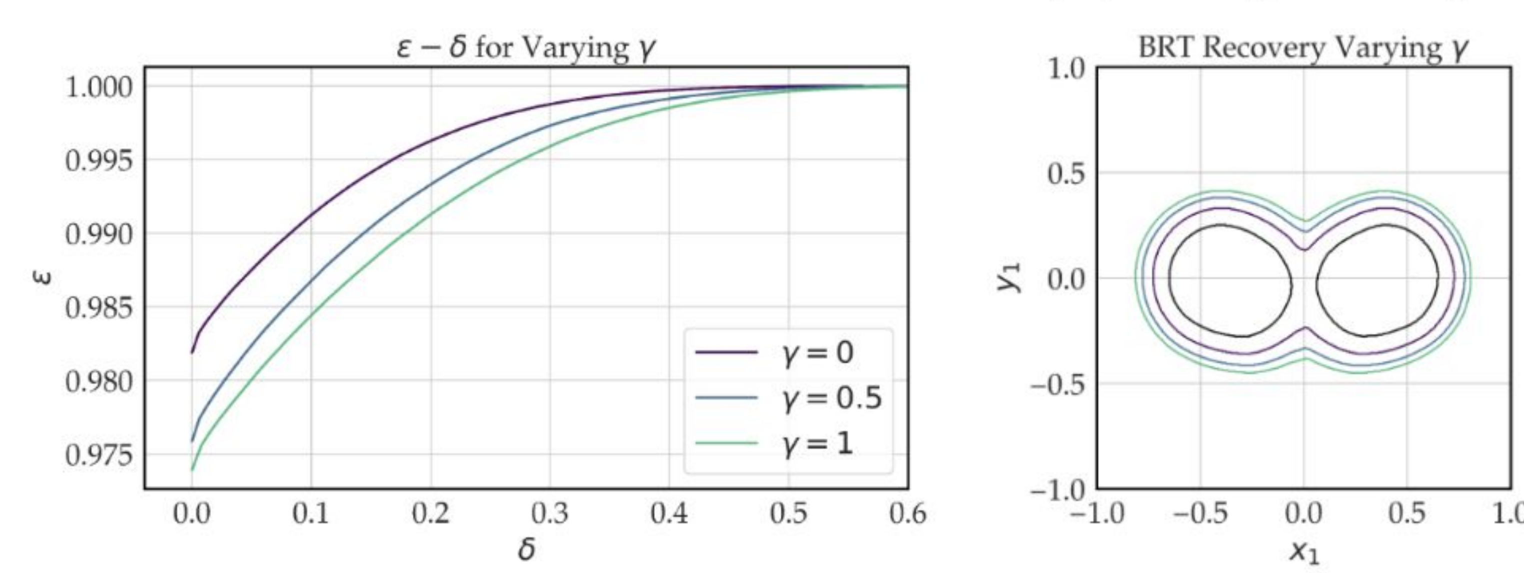

Low γ (small BRT volume recovery) → High γ (larger BRT volume recovery / lower learning accuracy).

Simulation

- Simulated rollouts comparing **Neural CBFs [Dawson 2021]** to RBNs as safety filters.

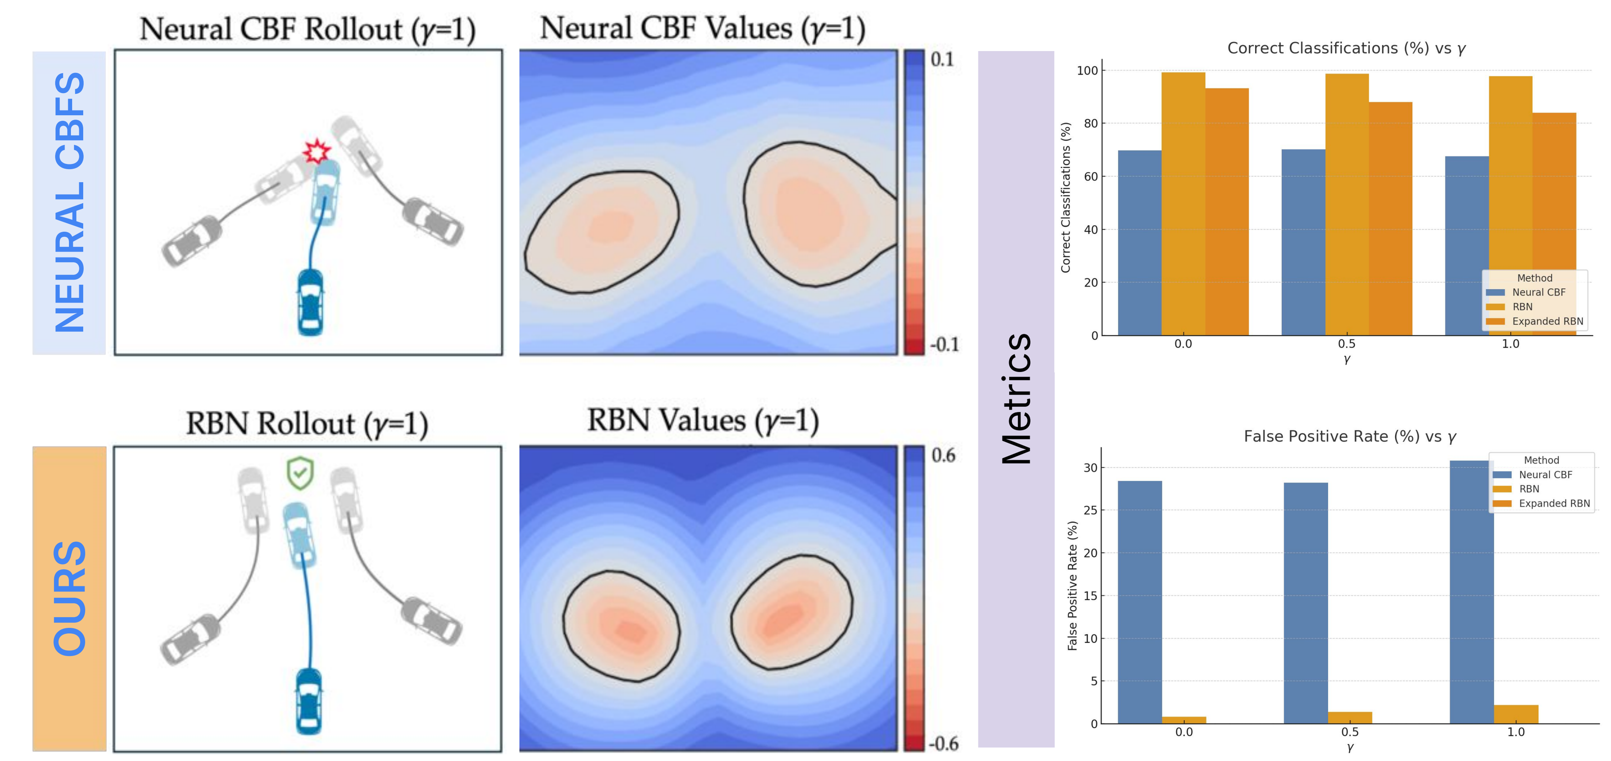

RBNs have 22% FEWER FP and 29% MORE CC than Neural CBFs.

Hardware

- **60 second** time-horizon experiments.
- Agents pursue **independent goals.**
- Agent gets a new goal upon reaching the current one or after 10 seconds.

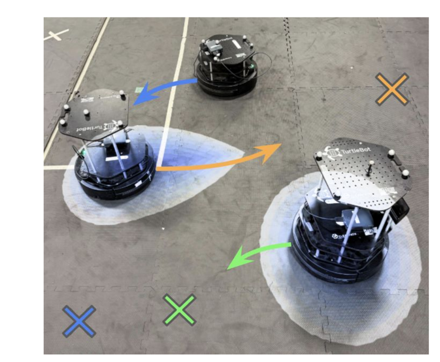

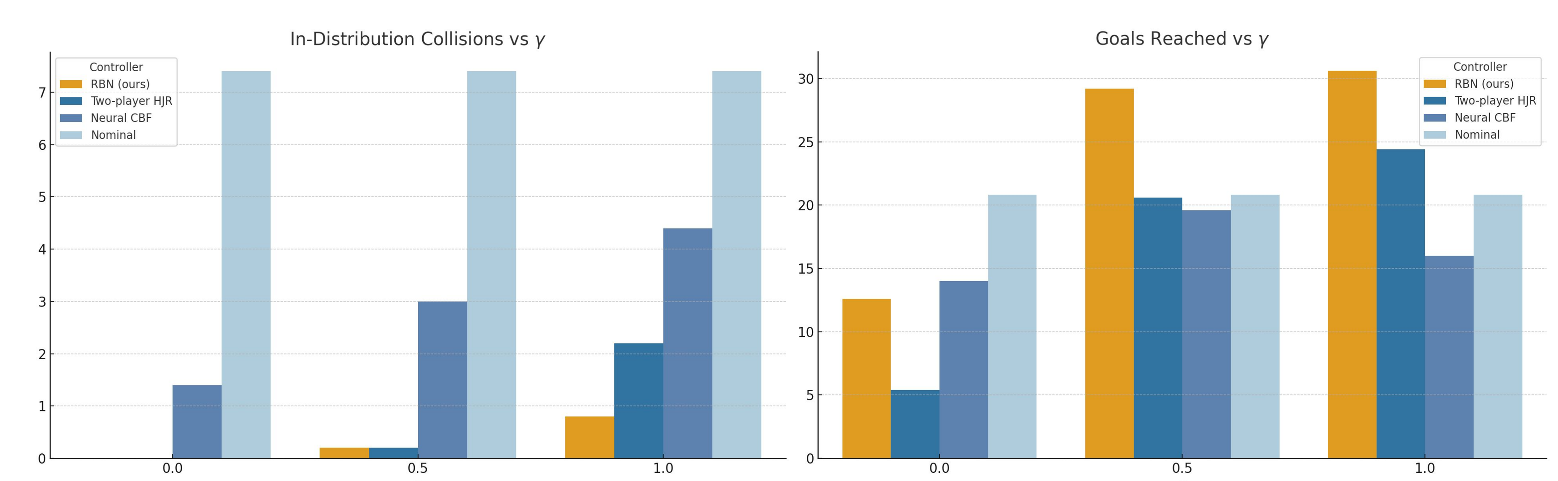

RBNs incur 5.5x FEWER COLLISIONS and reach 1.9x MORE GOALS than Neural CBFs.

Conclusion & Future Work

- **Scalable method for computing smooth and tunable CBFs.**
- **Insights of probabilistic guarantees for varying γ.**
- **Hardware experiments showcasing RBN efficacy.**

- Uncertainty quantification for OOD scenarios.
- Abstracting away modeling assumptions.
- Warmstarting to improve scalability.
- Multimodal data for better state estimation.