Steer2Reach

Forward Trajectory Steering For HJ Reachability Analysis

Sungje Park, Stephen Tu

Ming Hsieh Department of Electrical and Computer Engineering, USC

USC Viterbi School of Engineering
Ming Hsieh Department of Electrical and Computer Engineering

Paper
arXiv:2608.11480

Hamilton-Jacobi Reachability

Which states are unrecoverable?

Hamilton–Jacobi reachability finds states where even the best controller can be driven to failure. Outside this set, a safety policy can avoid failure under worst-case disturbances over the planning horizon.

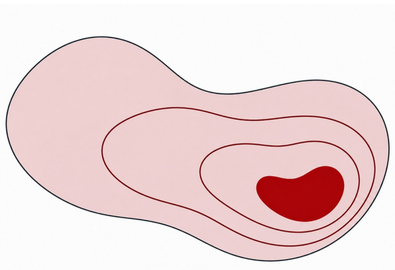

Failure can be forced
Pale region: V(x,0) ≤ 0

Immediate failure
Dark region: ℓ ≤ 0

Hamilton-Jacobi-Isaacs Variational Inequality

$$R_V(x,t) := \min\{\partial_t V + H,\ \ell - V\} = 0,$$

$$H = \max_{u\in\mathcal{U}} \min_{d\in\mathcal{D}} \langle \nabla_x V, f(x,u,d)\rangle,$$

$$V(x,T) = \ell(x).$$

Steer2Reach Method

Training algorithm

1 **Extract learned control and disturbance from the value network.**

$$u_\theta(x,t) = \arg\max_{u\in\mathcal{U}} \min_{d\in\mathcal{D}} \langle \nabla_x V_\theta, f(x,u,d)\rangle,$$

$$d_\theta(x,t) = \arg\min_{d\in\mathcal{D}} \langle \nabla_x V_\theta, f(x,u_\theta,d)\rangle.$$

2 **Roll forward SDE trajectories from sampled initial states.**

$$dX_t = f(X_t, u_\theta, d_\theta)\,dt + \sigma\,dB_t.$$

$$\mu_\theta = \mathrm{Law}\left((X_t^{x;u_\theta,d_\theta}, t)\right), \quad t\sim\mathcal{U}([0,T]),$$

$$\mu'_\theta = \mathrm{Law}\left(X_T^{x;u_\theta,d_\theta}\right), \quad x\sim\mathcal{U}(\Omega), \quad x\perp t.$$

3 **Update θ with the PINNs loss.**

$$\mathcal{L}_{\text{PINNs}}(\theta;\lambda) = \mathbb{E}_{(x,t)\sim\mu_\theta}\left[R_{V_\theta}(x,t)^2\right] + \lambda\,\mathbb{E}_{x\sim\mu'_\theta}\left[(V_\theta(x,T)-\ell(x))^2\right].$$

4 **Repeat steps 1–3.**

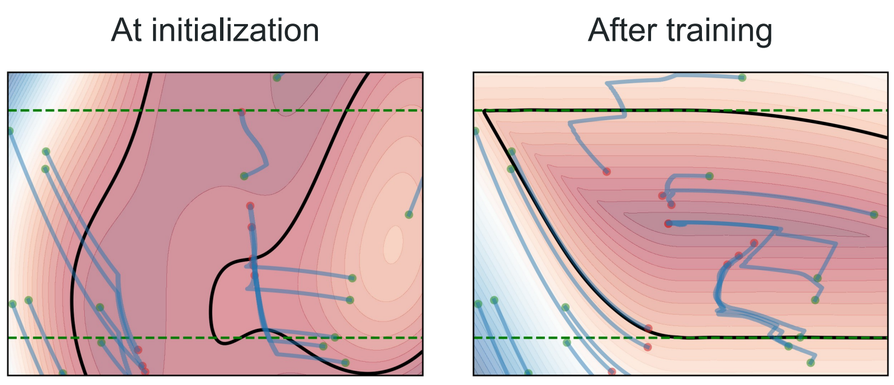

Velocity horizontally, height vertically. Blue paths. Black zero contour. Dashed floor/ceiling. Green/red: start/end.

Results

40D Publisher–Subscriber

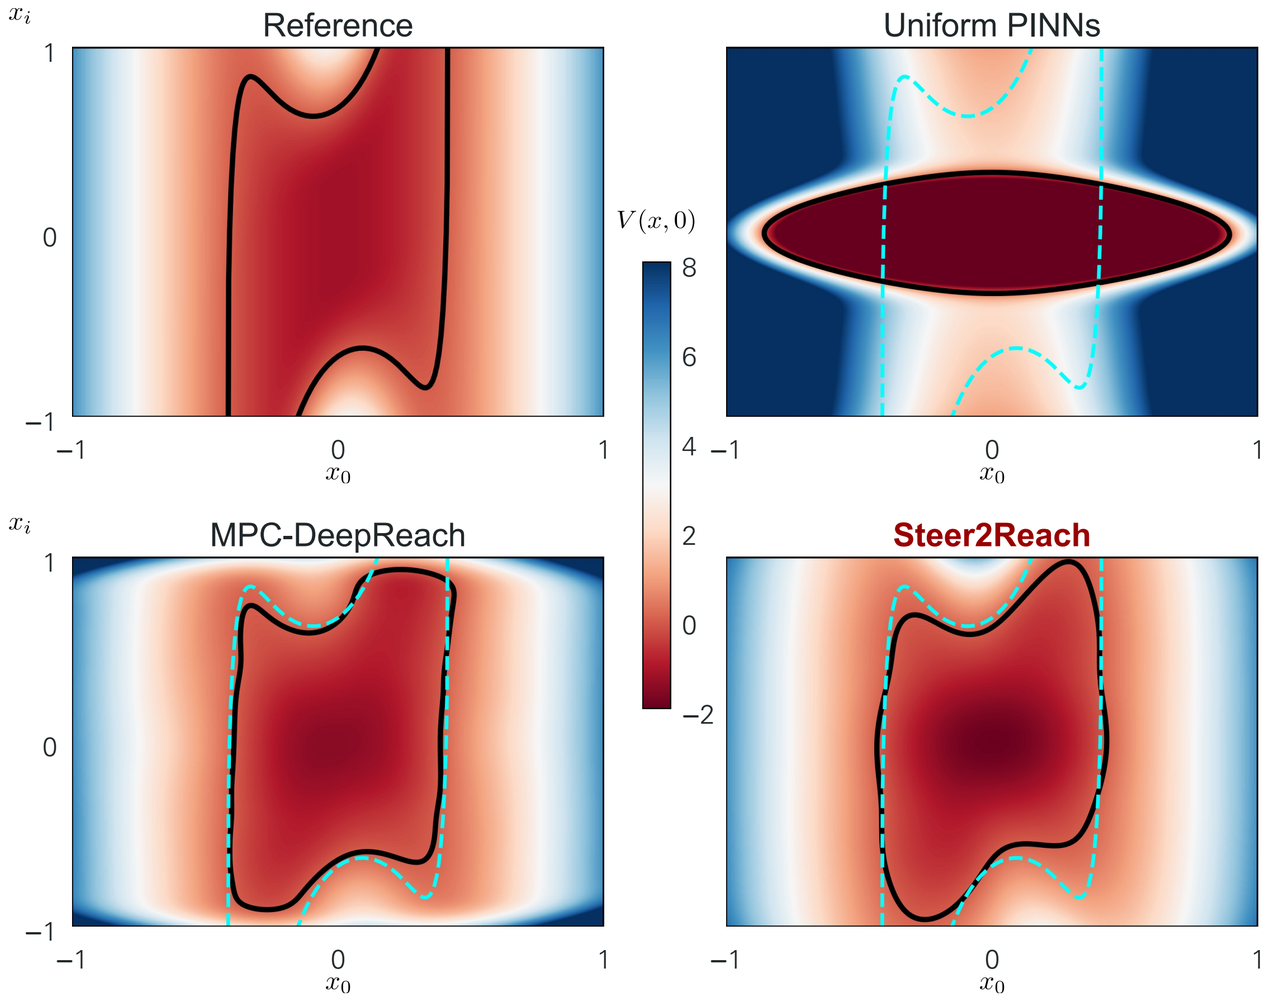

40D slice: $x = (x_0, x_i, \ldots, x_i)$ with all 39 subscriber coordinates equal.

MPC-DeepReach adds time curriculum, sampling-based MPC, joint PDE residual and MPC regression loss, and fine-tuning to uniform PINNs.

40D quantitative comparison

Method	Relative L² ↓	Safety IoU ↑	Time (min) ↓
PINNs	1.7727 ± 0.0756	0.0000 ± 0.0000	18.05
MPC-DeepReach	0.5676 ± 0.1539	0.9831 ± 0.0013	82.85
Steer2Reach	0.0983 ± 0.0266	0.9878 ± 0.0004	50.63

Mean ± SD, three initializations. L²: whole-domain value error. IoU: predicted safety vs. simulated outcomes. Training time: one H200 GPU.

2D Vertical Drone

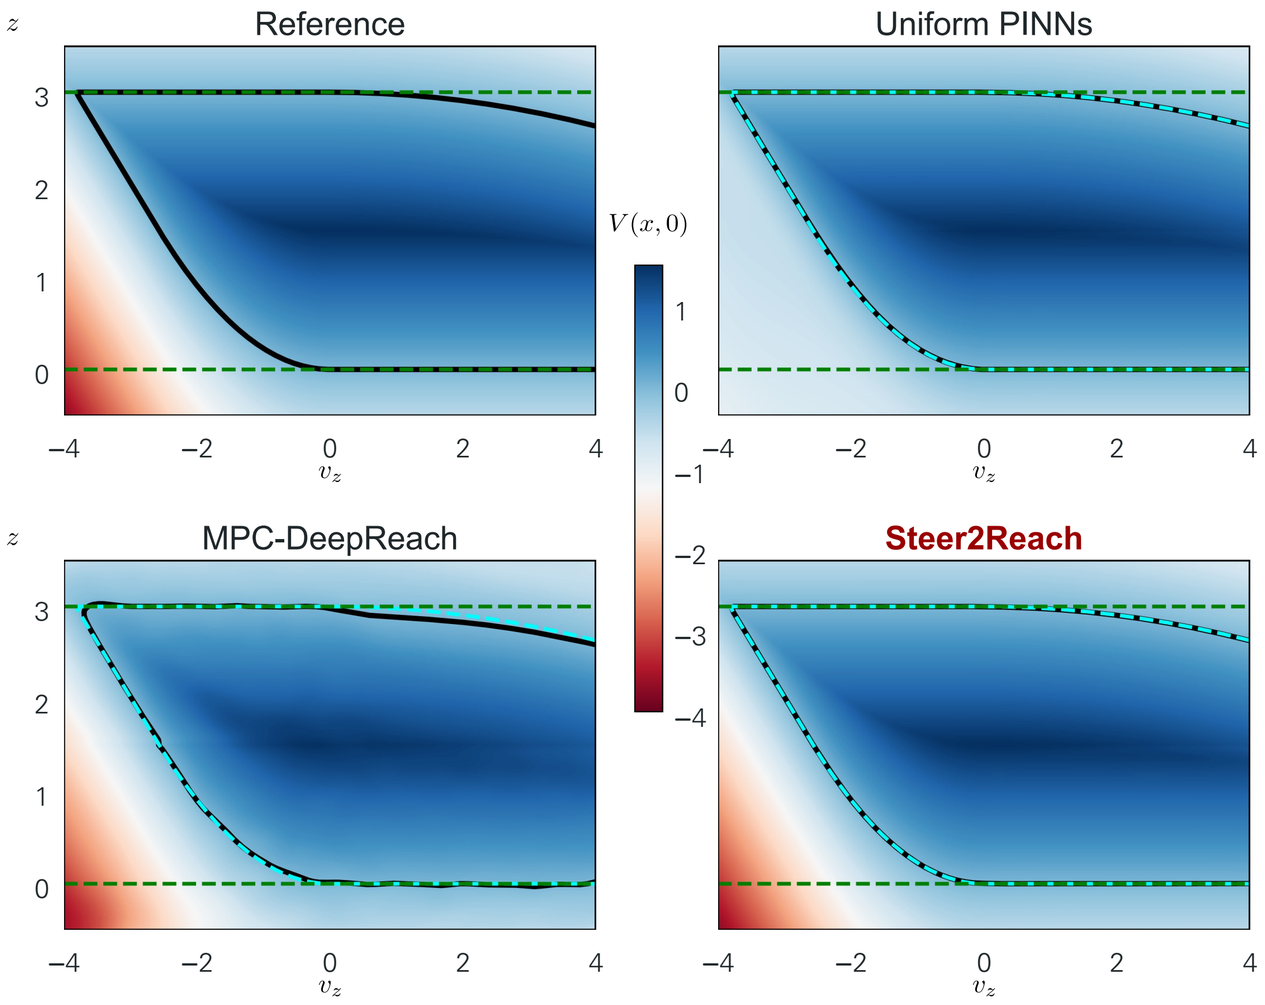

Black: model zero contour. Cyan dashes: reference. Green dashes: drone floor and ceiling.

Park & Tu, arXiv:2608.11480 (2026). Figures 1–2 and Table I. NSF CPS #2434460; NSF GRFP DGE-1842487.
sungjepa@usc.edu github.com/sungje-park/steer2reach