New approaches to synthesizing Stochastic Barrier Functions

Frederik Baymler Mathiesen

Delft Center for Systems and Control, TU Delft

July 9th, 2024

Agenda

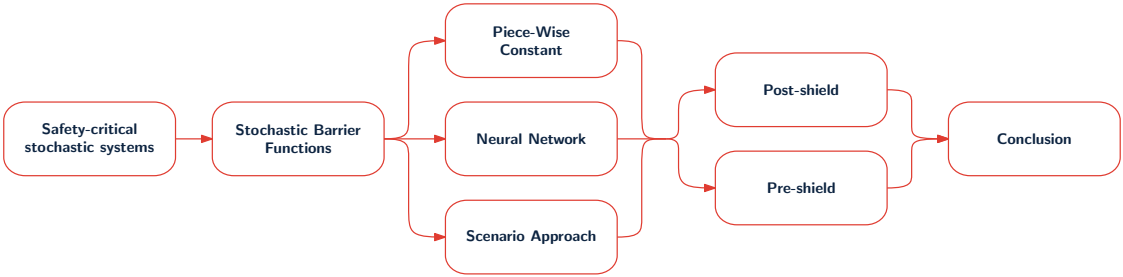
Safety-critical stochastic systems
Stochastic Barrier Functions
Piece-Wise Constant
Neural Network
Scenario Approach
Post-shield
Pre-shield
Conclusion

What are safety-critical systems?

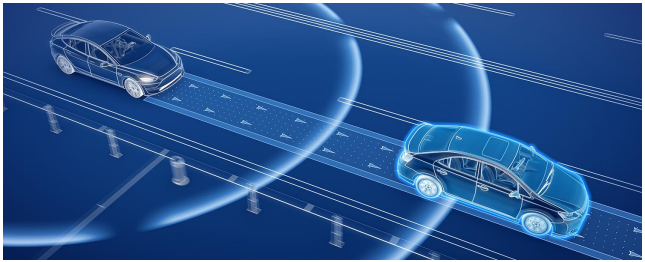

System model

$$x(k+1) = f(x(k), u(k), v(k))$$

Variations:

- No control $x(k+1) = f(x(k), v(k))$
- Additive control $x(k+1) = f(x(k), v(k)) + u(k)$
- Additive noise $x(k+1) = f(x(k)) + v(k)$
- Unknown noise distribution

Safety - formally

Probabilistic Safety

Let the following be given

- the dynamics f,
- a controller $\pi : X \to U$, and
- probability distribution of $v(k)$.

Then the safety probability is

$$P_{safe}(X_0, X_s, H) = \inf_{x_0 \in X_0} \mathbb{P}\left[x(k) \in X_s,\ \forall 0 \leq k \leq H \mid x(0) = x_0\right].$$

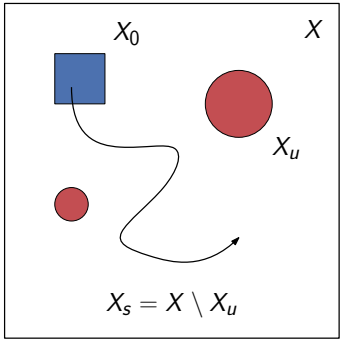

Problem statement

Given f, π, distribution of $v(k)$, X_0, X_s, and H, compute $P_{safe}(X_0, X_s, H)$.

Stochastic Barrier Functions

$$
\begin{aligned}
B(x) &\geq 0 \quad \forall x \in X \\
B(x) &\geq 1 \quad \forall x \in X_u \\
B(x) &\leq \gamma \quad \forall x \in X_0 \\
\mathbb{E}[B(f(x,v))] &\leq B(x) + \beta \quad \forall x \in X_s
\end{aligned}
$$

Safety certification

$$P_{safe}(X_0, X_s, H) \geq 1 - (\gamma + \beta H)$$

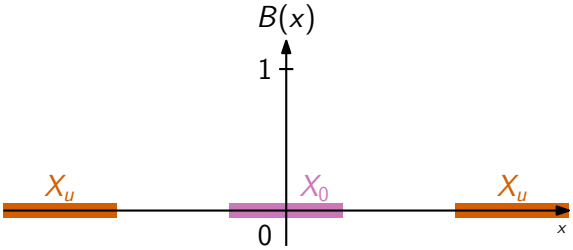

Stochastic Barrier Functions

$$
\begin{aligned}
B(x) &\geq 0 \quad \forall x \in X \\
B(x) &\geq 1 \quad \forall x \in X_u \\
B(x) &\leq \gamma \quad \forall x \in X_0 \\
\mathbb{E}[B(f(x,v))] &\leq B(x) + \beta \quad \forall x \in X_s
\end{aligned}
$$

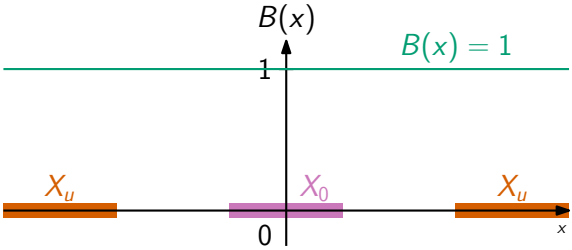

Safety certification

$$P_{safe}(X_0, X_s, H) \geq 1 - (\gamma + \beta H)$$

Stochastic Barrier Functions

$$
\begin{aligned}
B(x) &\geq 0 \quad \forall x \in X \\
B(x) &\geq 1 \quad \forall x \in X_u \\
B(x) &\leq \gamma \quad \forall x \in X_0 \\
\mathbb{E}[B(f(x,v))] &\leq B(x) + \beta \quad \forall x \in X_s
\end{aligned}
$$

Safety certification

$$P_{safe}(X_0, X_s, H) \geq 1 - (\gamma + \beta H)$$

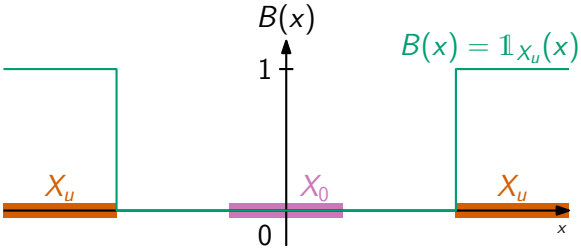

Stochastic Barrier Functions

$$
\begin{aligned}
B(x) &\geq 0 \quad \forall x \in X \\
B(x) &\geq 1 \quad \forall x \in X_u \\
B(x) &\leq \gamma \quad \forall x \in X_0 \\
\mathbb{E}[B(f(x,v))] &\leq B(x) + \beta \quad \forall x \in X_s
\end{aligned}
$$

Safety certification

$$P_{safe}(X_0, X_s, H) \geq 1 - (\gamma + \beta H)$$

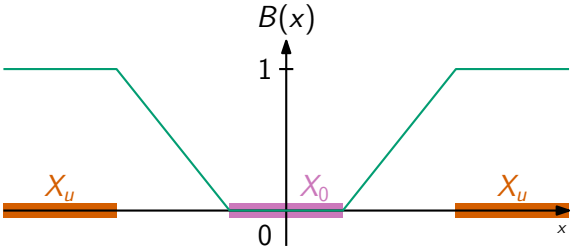

How can we guarantee safety using SBFs?

Bellman equation

If V_k is satisfying

$$V_H(x) = \mathbb{1}_{X_u}(x)$$
$$V_k(x) = \mathbb{1}_{X_u}(x) + \mathbb{1}_{X_s}(x)\mathbb{E}[V_{k+1}(f(x,v))]$$

then

$$P_{safe}(X_0, X_s, H) = 1 - \sup_{x \in X_0} V_0(x)$$

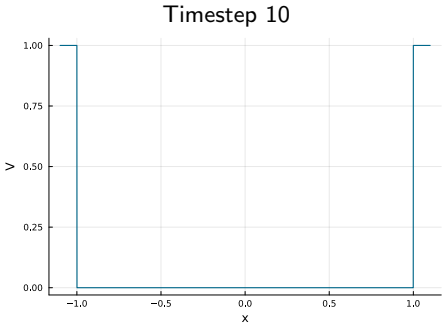

$$x(k+1) = 0.99x(k) + v(k),$$
$$v(k) \sim \mathcal{N}(0, 0.05)$$

How can we guarantee safety using SBFs?

Bellman equation

If V_k is satisfying

$$V_H(x) = \mathbb{1}_{X_u}(x)$$
$$V_k(x) = \mathbb{1}_{X_u}(x) + \mathbb{1}_{X_s}(x)\mathbb{E}[V_{k+1}(f(x,v))]$$

then

$$P_{safe}(X_0, X_s, H) = 1 - \sup_{x \in X_0} V_0(x)$$

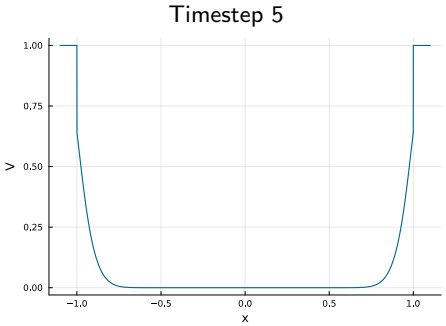

$$x(k+1) = 0.99x(k) + v(k),$$
$$v(k) \sim \mathcal{N}(0, 0.05)$$

How can we guarantee safety using SBFs?

Bellman equation

If V_k is satisfying

$$V_H(x) = \mathbb{1}_{X_u}(x)$$
$$V_k(x) = \mathbb{1}_{X_u}(x) + \mathbb{1}_{X_s}(x)\mathbb{E}[V_{k+1}(f(x,v))]$$

then

$$P_{safe}(X_0, X_s, H) = 1 - \sup_{x \in X_0} V_0(x)$$

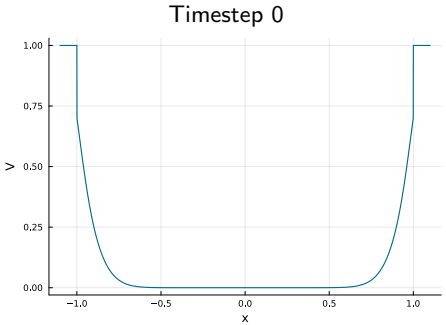

$$x(k+1) = 0.99x(k) + v(k),$$
$$v(k) \sim \mathcal{N}(0, 0.05)$$

How can we guarantee safety using SBFs?

Bellman equation

If $\mathcal{B}_k$ is satisfying

$$\mathcal{B}_H(x) = B(x)$$
$$\mathcal{B}_k(x) = \mathbb{1}_{X_u}(x) + \mathbb{1}_{X_s}(x)\mathbb{E}[\mathcal{B}_{k+1}(f(x,v))]$$

and $\beta \geq \sup_{x \in X_s} \mathbb{E}[B(f(x,v))] - B(x)$, then

$$V_k(x) \leq \mathcal{B}_k(x) \leq B(x) + \beta(H-k)$$
$$P_{safe}(X_0, X_s, H) \geq 1 - \sup_{x \in X_0}(B(x) + \beta H)$$

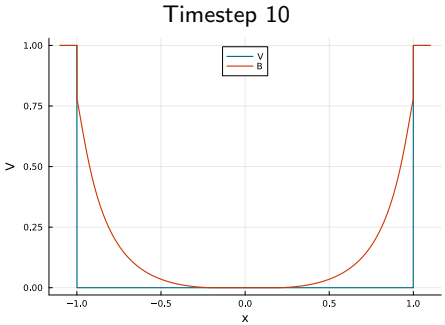

$$x(k+1) = 0.99x(k) + v(k),$$
$$v(k) \sim \mathcal{N}(0, 0.05)$$

How can we guarantee safety using SBFs?

Bellman equation

If $\mathcal{B}_k$ is satisfying

$$\mathcal{B}_H(x) = B(x)$$
$$\mathcal{B}_k(x) = \mathbb{1}_{X_u}(x) + \mathbb{1}_{X_s}(x)\mathbb{E}[\mathcal{B}_{k+1}(f(x,v))]$$

and $\beta \geq \sup_{x \in X_s} \mathbb{E}[B(f(x,v))] - B(x)$, then

$$V_k(x) \leq \mathcal{B}_k(x) \leq B(x) + \beta(H-k)$$
$$P_{safe}(X_0, X_s, H) \geq 1 - \sup_{x \in X_0}(B(x) + \beta H)$$

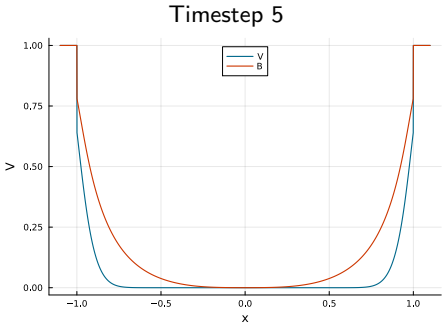

$$x(k+1) = 0.99x(k) + v(k),$$
$$v(k) \sim \mathcal{N}(0, 0.05)$$

How can we guarantee safety using SBFs?

Bellman equation

If $\mathcal{B}_k$ is satisfying

$$\mathcal{B}_H(x) = B(x)$$
$$\mathcal{B}_k(x) = \mathbb{1}_{X_u}(x) + \mathbb{1}_{X_s}(x)\mathbb{E}[\mathcal{B}_{k+1}(f(x,v))]$$

and $\beta \geq \sup_{x \in X_s} \mathbb{E}[B(f(x,v))] - B(x)$, then

$$V_k(x) \leq \mathcal{B}_k(x) \leq B(x) + \beta(H-k)$$
$$P_{safe}(X_0, X_s, H) \geq 1 - \sup_{x \in X_0}(B(x) + \beta H)$$

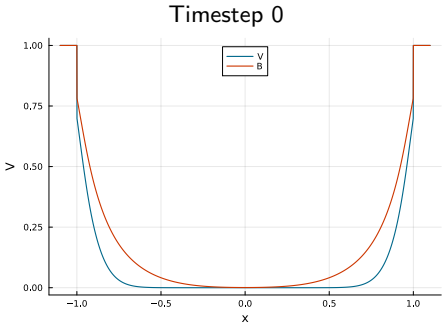

$$x(k+1) = 0.99x(k) + v(k),$$
$$v(k) \sim \mathcal{N}(0, 0.05)$$

"Traditional" Stochastic Barrier Function Synthesis

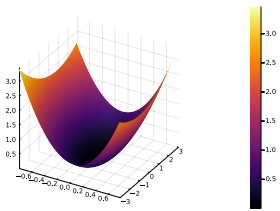

$$B(x) = x^T P x$$

- Easy to synthesize
- Conservative
- Symmetry

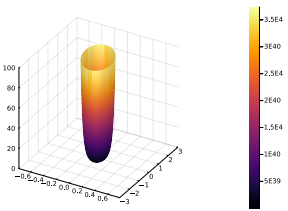

$$B(x) = e^{x^T P x} - 1$$

- Difficult to synthesize
- Conservative
- Symmetry

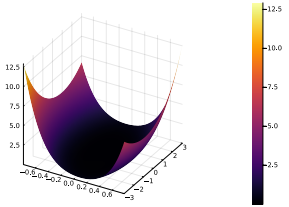

$$B(x) = m(x)^T P m(x)$$

- Does not scale
- Less conservative
- Partial symmetry

Piece-wise Constant Stochastic Barrier Functions[1]

- Optimal
- Closely related to Interval Value Iteration
- Synthesis methods:
  - Linear programming
  - Projected (sub)-gradient descent

$$
\begin{aligned}
B_j &\geq 0 \quad \forall j \\
B_j &\geq 1 \quad \forall j : X_j \cap X_u \neq \varnothing \\
B_j &\leq \gamma \quad \forall j : X_j \cap X_0 \neq \varnothing \\
\sum_{i=1}^{n} B_i p_{ij} &\leq B_j + \beta \quad \forall j : X_j \cap X_s \neq \varnothing
\end{aligned}
$$

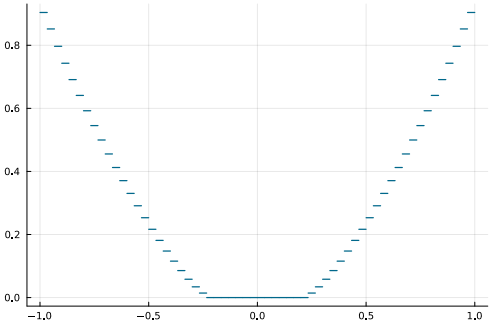

where $p_{ij} \in \left[\min_{x \in X_j} \mathbb{P}[x(k) \in X_i \mid x(k-1) = x],\ \max_{x \in X_j} \mathbb{P}[x(k) \in X_i \mid x(k-1) = x] \right]$.

[1] Mathiesen, F. B., Mazouz, R., Calvert, S. C., Lahijanian, M., & Laurenti, L., Piecewise Barrier Functions for Stochastic Systems, Under preparation.

Piece-wise Constant Stochastic Barrier Functions

Linear Programming

The following LP synthesizes a PWC SBF guaranteeing $P_{safe}(X_0, X_s, H) \geq 1 - (\gamma + \beta H)$.

$$
\begin{array}{lll}
\min & \gamma + \beta H & \\
\text{s.t.} & 0 \leq B_j \leq 1 & \forall j \\
& B_j = 1 & \forall j : X_j \cap X_u \neq \varnothing \\
& B_j \leq \gamma & \forall j : X_j \cap X_0 \neq \varnothing \\
& h_j^\top \lambda_j \leq B_j + \beta & \forall j : X_j \cap X_s \neq \varnothing \\
& H_j^\top \lambda_j = \begin{bmatrix} B_1 & \cdots & B_N \end{bmatrix}^\top & \forall j : X_j \cap X_s \neq \varnothing
\end{array}
$$

$$
H_j p_j = \begin{bmatrix} -1 & & \\ 1 & & \\ & \ddots & \\ & & -1 \\ & & 1 \\ -1 & \cdots & -1 \\ 1 & \cdots & 1 \end{bmatrix} p_j \leq \begin{bmatrix} -\underline{p}_{1j} \\ \bar{p}_{1j} \\ \vdots \\ -\underline{p}_{Nj} \\ \bar{p}_{Nj} \\ -1 \\ 1 \end{bmatrix} = h_j
$$

Piece-wise Constant Stochastic Barrier Functions

Projected Subgradient Descent

- Loss function

$$\mathcal{L}(B) = \gamma + \beta H$$

$$\text{where} \quad \gamma = \max_{\forall j : X_j \cap X_0 \neq \varnothing} B_j$$

$$\beta = \max_{\forall j : X_j \cap X_s \neq \varnothing} \sup_{p_j \in P_j} \max(B^\top p_j - B_j, 0)$$

- Project onto $[0,1]^N$ where all regions $j : X_j \cap X_u \neq \varnothing$ is forced 1.

Piece-wise Constant Stochastic Barrier Functions - Results

Table: n is the dimensionality of the system, $|K|$ the number of regions, and $\mathcal{I}$ the number of iterations. T_p and T_s is the time to compute transition probability and to synthesize a barrier, respectively. TO means computation time exceeded 1 hour.

				Linear Program		Gradient Descent			Sum-of-Squares			
Model	n	$\|K\|$	T_p (s)	P_{safe}	T_s (s)	P_{safe}	$\mathcal{I}$	T_s (s)	$\|K\|$	Deg	P_{safe}	T_s (s)
Linear	2	64	0.018	0.992	0.52	0.985	2	0.37	-	4	0.582	0.014
Convex		225	0.313	0.998	164.60	0.973	8	0.10	-	8	0.582	0.265
		900	8.849	0.999	$1.1 \cdot 10^3$	0.990	6	0.12	-	30	0.978	151.16
		2500	41.44	0.999	$2.9 \cdot 10^3$	0.991	10	0.49	-	35	0.988	458.21
Linear	2	900	5.04	0.494	$1.2 \cdot 10^3$	0.494	30	0.52	-	12	0	0.020
Non-Convex		1225	8.20	0.800	$1.4 \cdot 10^3$	0.800	20	0.64	-	20	0.08	141.20
		1444	9.18	0.920	$1.5 \cdot 10^3$	0.920	20	0.93	-	30	-	TO
Unicycle	4	1250	$1.1 \cdot 10^3$	0.750	$1.0 \cdot 10^3$	0.75	$3 \cdot 10^3$	5.68	1800	2	0	$3.1 \cdot 10^3$
		1800	$1.8 \cdot 10^3$	0.975	$1.7 \cdot 10^3$	0.975	$4 \cdot 10^4$	25.78	1800	4	0	$5.4 \cdot 10^3$
		2400	$2.0 \cdot 10^3$	0.998	$2.5 \cdot 10^3$	0.998	$2 \cdot 10^4$	55.59	1800	6	-	TO

Piece-wise Constant vs Sum-of-Squares

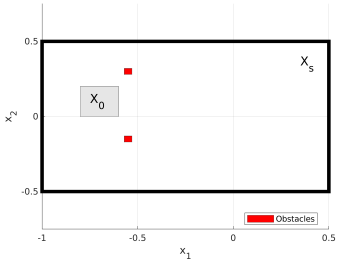

Figure: Initial and safe sets

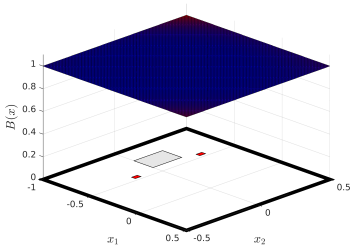

Figure: Degree-30 SOS SBF

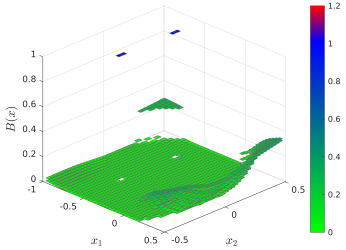

Figure: Piece-wise Constant SBF

Neural Stochastic Barrier Functions[2]

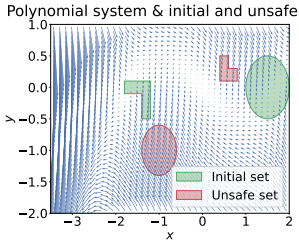

(a) Nominal dynamics

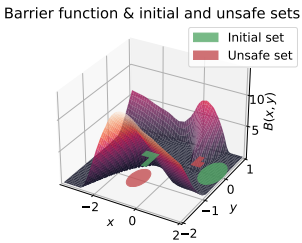

(b) Barrier

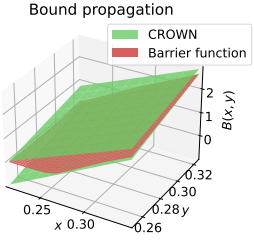

(c) CROWN

[2] Mathiesen, F. B., Calvert, S. C., & Laurenti, L. (2022). Safety certification for stochastic systems via neural barrier functions. IEEE Control Systems Letters.

How to train a Neural Stochastic Barrier Function

Training iteration

1. Sample n points from X, X_u, X_0, and X_s
2. Compute constraint violation in hyperrectangle around each point
3. Loss = sum of violation
4. Backprop + gradient step

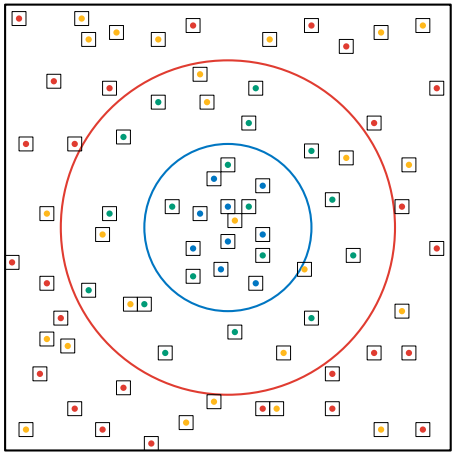

How to certify a Neural Stochastic Barrier Function

CROWN

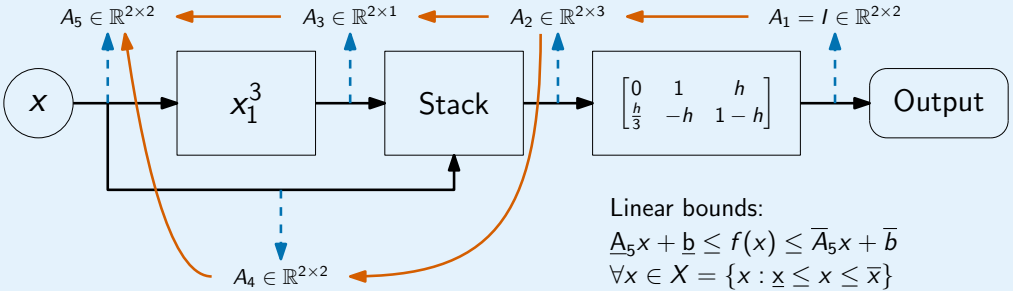

Linear bounds:

$$\underline{A}_5 x + \underline{b} \leq f(x) \leq \overline{A}_5 x + \overline{b}$$

$$\forall x \in X = \{x : \underline{x} \leq x \leq \overline{x}\}$$

How to certify a Neural Stochastic Barrier Function

Branch-and-bound

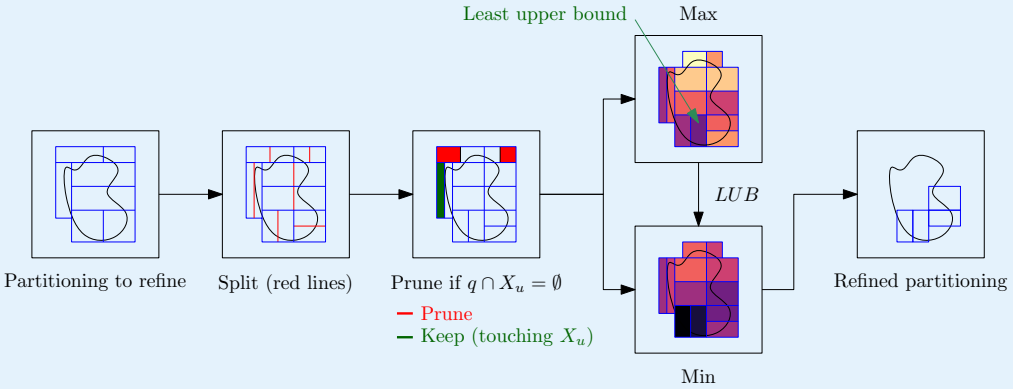

Neural Stochastic Barrier Functions - Results

Table: Certified $P_{safe}(X_0, X_s, H)$. Cells with "-" denotes that SoS failed to compute a barrier.

Method	Linear	2-D polynomial	Dubin's car
SoS (4)	0.690906	0.000000	-
SoS (8)	0.975079	0.232710	-
SoS (13)	0.998405	0.681383	-
SoS (15)	0.999761	-	-
Lipschitz	0.824654	0.792392	0
NBF (grid)	0	0	0
NBF (BaB)	**0.999969**	**0.991664**	**0.870272**

Scenario Stochastic Barrier Functions[3]

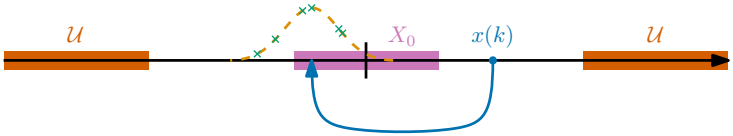

$$\mathbb{E}[B(f(x,v))] \leq B(x) + \beta$$

Expectation to sample-average constraint

Let δ be an error term, then

$$\frac{1}{|\mathcal{D}|} \sum_{v_i \in \mathcal{D}} B(f(x,v_i)) + \delta \leq B(x) + \beta, \quad \forall x \in X_s$$

[3] Mathiesen, F. B., Romao, L., Calvert, S. C., Abate, A., & Laurenti, L. (2023). Inner Approximations of Stochastic Programs for Data-driven Stochastic Barrier Function Design. Conference on Decision and Control.

Scenario Stochastic Barrier Functions[3]

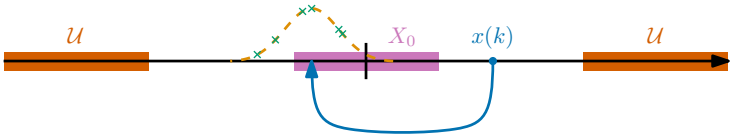

$$\mathbb{E}[B(f(x,v))] \leq B(x) + \beta$$

Expectation to chance constraint

Let $\nu \geq \frac{\epsilon M}{1-\epsilon}$ where $M = \sup_{x \in X_s} B(x)$, then

$$\mathbb{P}[v : B(f(x,v)) + \nu \leq B(x) + \beta,\ \forall x \in X_s] \geq 1 - \epsilon$$

[3] Mathiesen, F. B., Romao, L., Calvert, S. C., Abate, A., & Laurenti, L. (2023). Inner Approximations of Stochastic Programs for Data-driven Stochastic Barrier Function Design. Conference on Decision and Control.

Scenario Stochastic Barrier Functions[3]

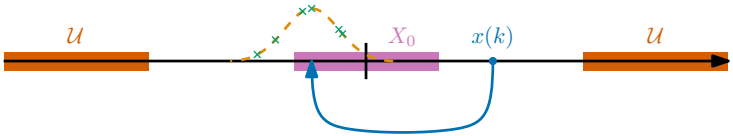

$$\mathbb{E}[B(f(x,v))] \leq B(x) + \beta$$

Expectation to scenario constraint

Let $\nu \geq \frac{\epsilon M}{1-\epsilon}$ where $M = \sup_{x \in X_s} B(x)$, then

$$B(f(x,v_i)) + \nu \leq B(x) + \beta, \quad \forall x \in X_s,\ v_i \in \mathcal{D}$$

[3] Mathiesen, F. B., Romao, L., Calvert, S. C., Abate, A., & Laurenti, L. (2023). Inner Approximations of Stochastic Programs for Data-driven Stochastic Barrier Function Design. Conference on Decision and Control.

Sample complexity

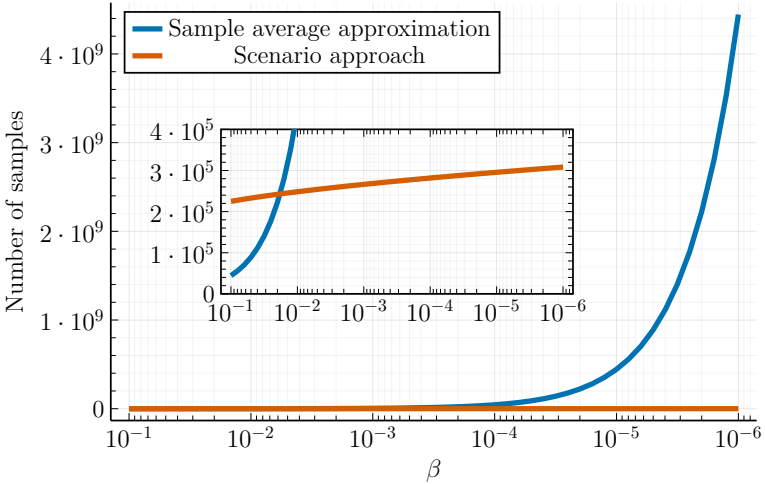

Scenario Stochastic Barrier Functions - Results

Table: n is the dimensionality of the system, $\bar{\ell}$ is the number of pieces of the PWA SBF B, and the confidence in the certificate is $1-\eta$.

System	n	Method	$\bar{\ell}$	η	$P_{safe}(X_0, X_s, H)$	Comp. time (s)
Martingale	1	Our	7	$\mathbf{10^{-9}}$	0.769	**0.096**
		SAA	-	10^{-6}	**0.910**	0.249
Drone	2	Our	33	$\mathbf{10^{-9}}$	**0.995**	4.84
		SAA	-	10^{-6}	0.950	**1.18**
Vehicle	2	Our	18	$\mathbf{10^{-9}}$	0.618	**1.44**
		Our	42	$\mathbf{10^{-9}}$	0.712	2.45
		Our	46	$\mathbf{10^{-9}}$	0.842	3.72
		Our	126	$\mathbf{10^{-9}}$	**0.994**	9.06
		SAA	-	10^{-6}	0.000	2.14

Comparison

Piece-wise Constant SBFs

- Works for discontinuous systems and controllers
- Assumes known transition probabilities
- Scales $O(n^2)$ in the number of regions

Neural SBFs

- Useful for complex continuous systems
- Requires verification of neural networks
- Scales $O(n^2)$ in the number of layers

Scenario SBFs

- Can be used for unknown distributions
- Works for complex and discontinuous systems
- Scales $O(n^2)$ in the number of regions

Proposal: A joint method for scalable SBF

Data

Neural Barrier Function

Scenario Approch

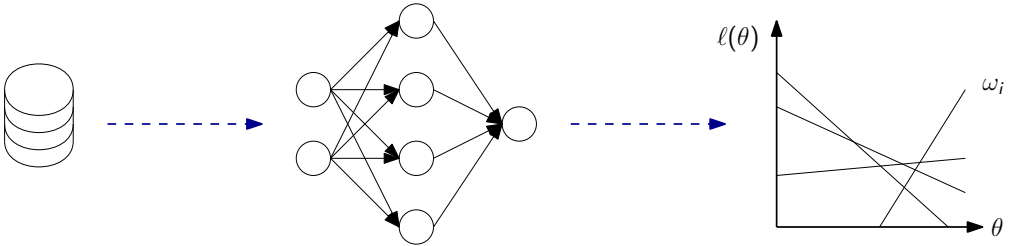

Safe control

Pre-shield

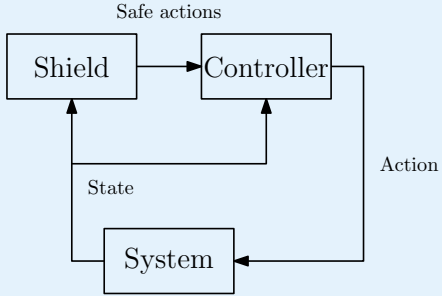

Post-shield

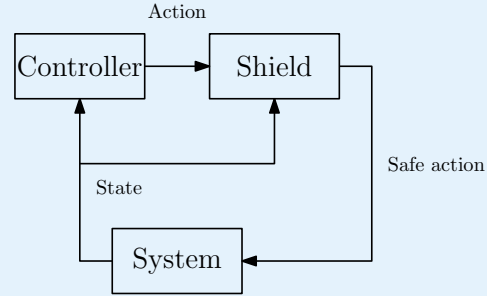

Minimally-invasive Controller Synthesis[4]

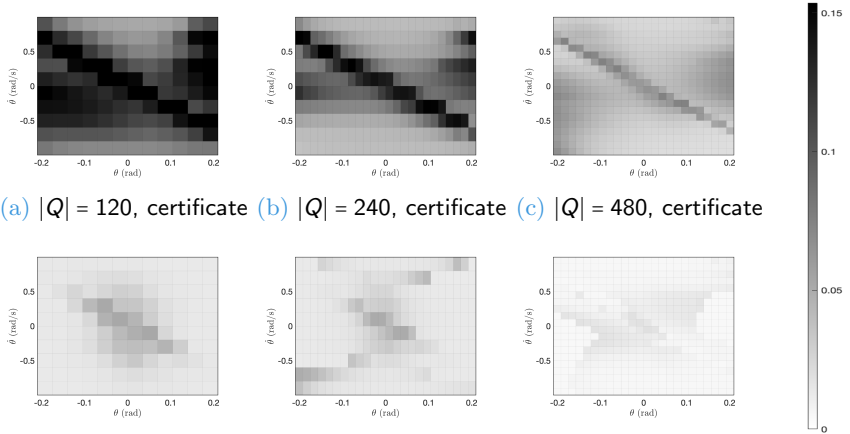

(a) $|Q| = 120$, certificate (b) $|Q| = 240$, certificate (c) $|Q| = 480$, certificate

(d) $|Q| = 120$, control (e) $|Q| = 240$, control (f) $|Q| = 480$, control

[4] Mazouz, R., Muvvala, K., Ratheesh Babu, A., Laurenti, L., & Lahijanian, M. (2022). Safety Guarantees for Neural Network Nynamic Systems via Stochastic Barrier Functions. Advances in Neural Information Processing Systems.

Control Invariant Sets[5]

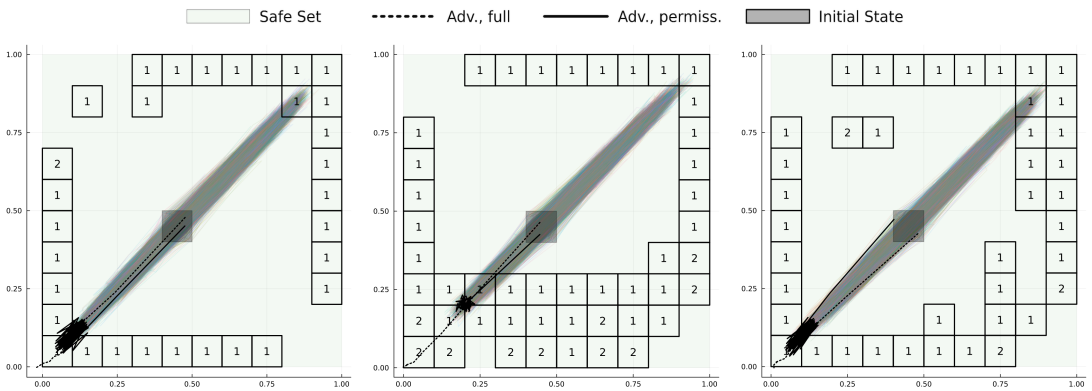

(a) Known
$P_{safe} \geq 1 - 5.83 \times 10^{-5}$

(b) GP 500
$P_{safe} \geq 1 - 2.39 \times 10^{-4}$

(c) GP 2000
$P_{safe} \geq 1 - 1.93 \times 10^{-5}$

[5] Mazouz, Rayan, et al. "Data-Driven Permissible Safe Control with Barrier Certificates." arXiv preprint arXiv:2405.00136 (2024).

Conclusion

Summary

- Verify safety for stochastic systems
- Rely on "worst-case" propagation
- Enables safe control synthesis
- Efficient and scalable synthesis remains an open problem

Future work

- Fully data-driven stochastic barrier functions
- Improving scalability
- Applications to autonomous driving

Collaborators

Licio Romao

Alessandro Abate

Simeon C. Calvert

Rayan Mazouz

Morteza Lahijanian

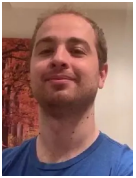

Luca Laurenti

Stochastic Barrier Functions for Neural Network Dynamics Systems[6]

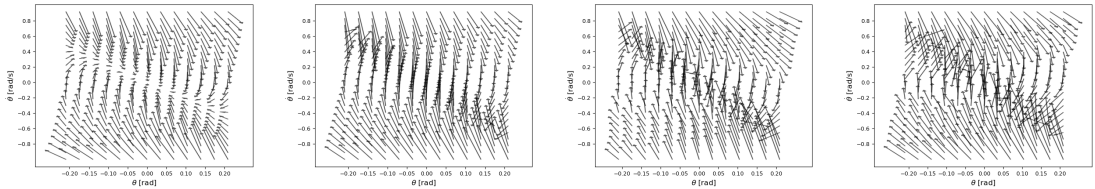

(a) 1 layer (b) 2 layers (c) 3 layers (d) 5 layers

Figure: Vector fields of the NNDMs for representing an inverted pendulum.

[6] Mazouz, R., Muvvala, K., Ratheesh Babu, A., Laurenti, L., & Lahijanian, M. (2022). Safety Guarantees for Neural Network Nynamic Systems via Stochastic Barrier Functions. Advances in Neural Information Processing Systems.

LogSumExp: real softmax

$$LSE(x_1, \ldots, x_n) = \log\left(\sum_{i=1}^{n} e^x\right)$$

Standard guarantees

$$\max(x_1, \ldots, x_n) \leq LSE(x_1, \ldots, x_n) \leq \max(x_1, \ldots, x_n) + \log(n)$$

Tempered

$$\max(x_1, \ldots, x_n) \leq \frac{1}{t} LSE(tx_1, \ldots, tx_n) \leq \max(x_1, \ldots, x_n) + \frac{\log(n)}{t}$$

Gaussian smoothing adversarial training[7]

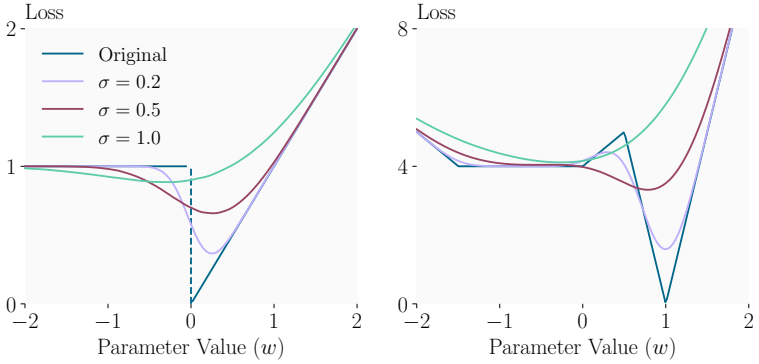

(a) Discontinuity (b) Sensitivity

[7] Balauca, Stefan, et al. "Overcoming the Paradox of Certified Training with Gaussian Smoothing." arXiv preprint arXiv:2403.07095 (2024).

IMDP vs SBF

If V_k is satisfying

$$V_H(s) = \mathbb{1}_{S_u}(s)$$
$$V_k(s) = \mathbb{1}_{S_u}(s) + \mathbb{1}_{S \setminus S_u}(s)\mathbb{E}[V_{k+1}(s')]$$

then

$$P_{safe}(X_0, X_s, H) \geq 1 - \sup_{s \in S_0} V_0(s)$$

Approach	Soundness	Optimality	Accuracy	Computational Effort	Nonlinear Dynamics
SBF	+	−	−	+	+
IMDP	+	+	+	−	+

Scenario approach

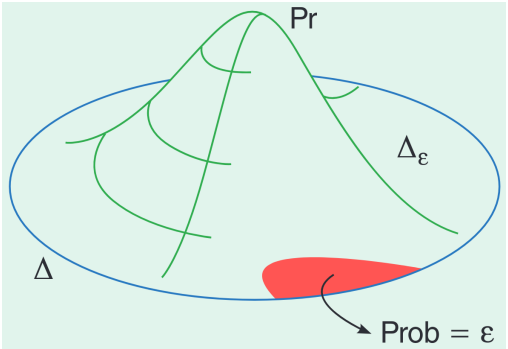

Chance-constrained problem

$$\begin{aligned} \min \quad & f(x) \\ s.t. \quad & \mathbb{P}[g(x,\omega) \leq 0] \geq 1-\epsilon \end{aligned}$$

Scenario problem

$$\begin{aligned} \min \quad & f(x) \\ s.t. \quad & g(x,\omega_i) \leq 0, \quad \forall \omega_i \in D \end{aligned}$$

Then $V(x^\star) \leq \epsilon$ with high probability.

Computational tricks for scenario barrier function

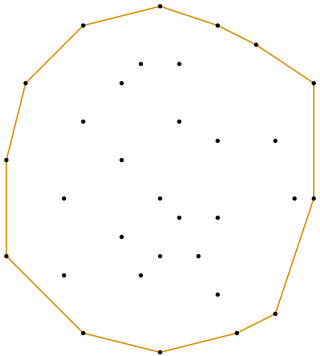

Figure: Convex hull over noise

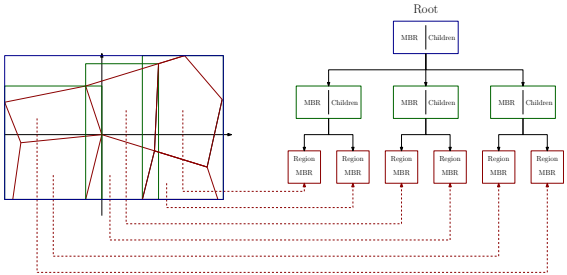

Figure: R-tree spatial indexing