Lifting as an Abstraction

Necmiye Ozay

Electrical Engineering and Computer Science & Robotics,
University of Michigan

CDC 2025 Pre-conference Workshop

Data-Driven Control of Autonomous Systems with Provable Guarantees

December 9, 2025

Control of safety-critical systems

- Many engineering systems are complex and safety critical

Power networks, drones, self-driving cars,…

- Violating constraints may lead to catastrophic events

Damages, accidents,…

- Need for systematic methods to design provably-safe policies

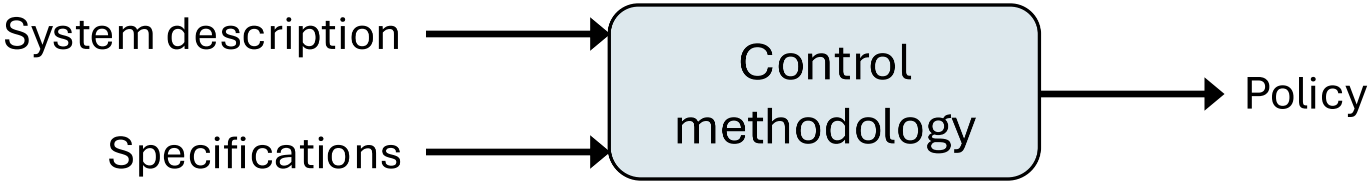

Problem statement

Given

- a nonlinear system

$$x_{t+1} = f(x_t, u_t)$$

over a domain $(x_t, u_t) \in \mathcal{X} \times \mathcal{U}$

- a specification (e.g., reach a region while avoiding others)

Find a policy

$$\pi: \mathcal{X} \to \mathcal{U}$$

such that the trajectories of the closed-loop system

$$x_{t+1} = f(x_t, \pi(x_t))$$

satisfy the specification.

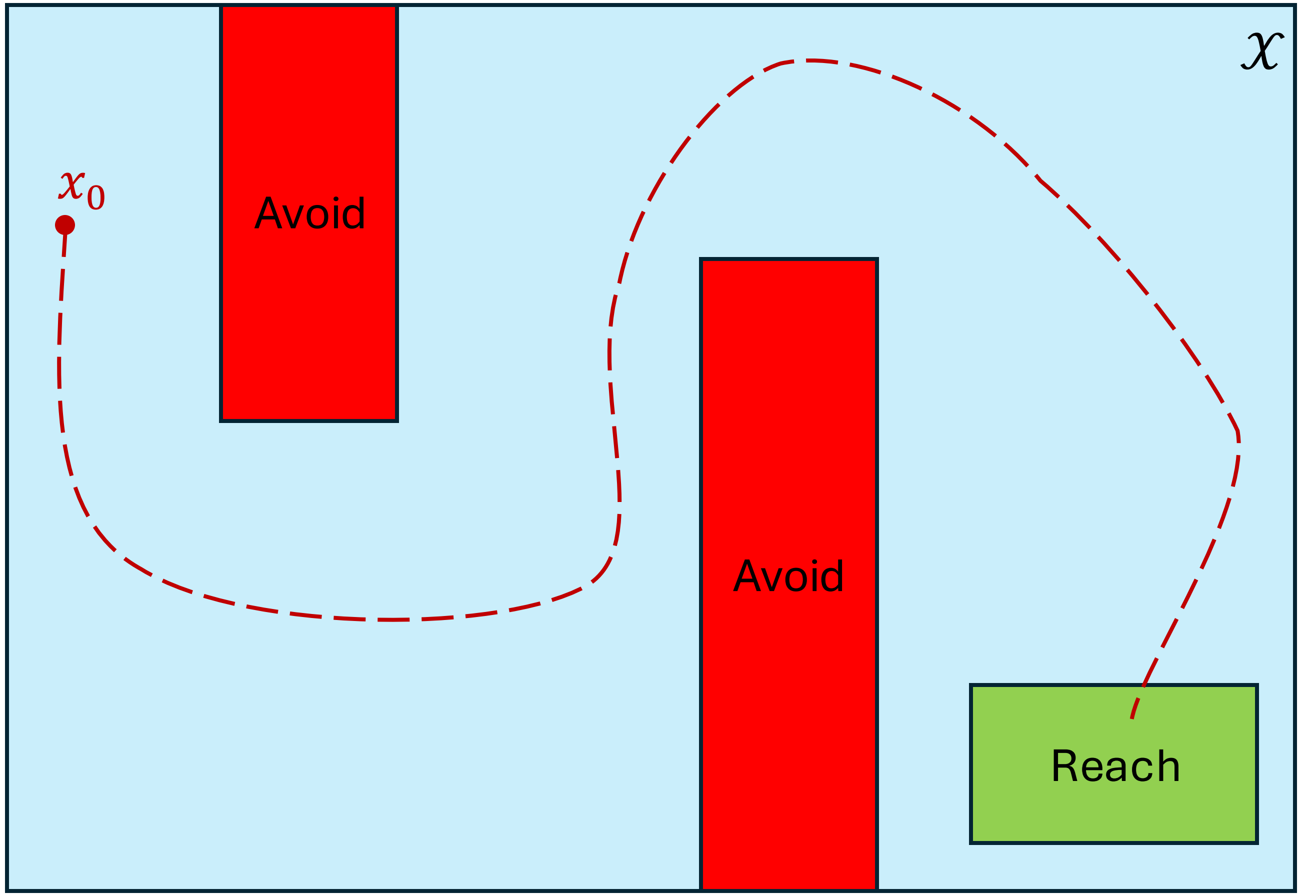

Outline

1. Introduction
2. **Over-approximations**
3. Lifted over-approximations
4. Conclusion

How do we over-approximate? Abstractions

Nonlinear dynamics: $x_{t+1} = f(x_t, u_t) \quad (x_t, u_t) \in \mathcal{X} \times \mathcal{U}$

Discrete-absractions:

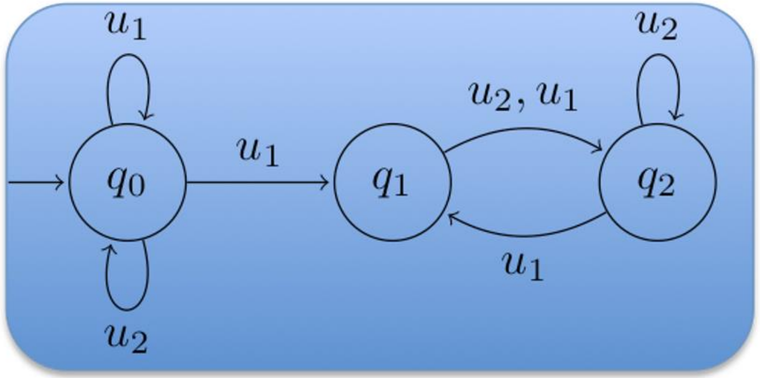

Theorem (informal)

If the abstraction function $A_x: \mathcal{X} \to Q, A_u: \mathcal{U} \to \mathcal{U}_a$ is such that for any x,
$f(x,u) \in A_x^{-1}(Post(A_x(x), A_u(u)))$
then any policy that enforces a reach-avoid specification can be concretized such that the trajectory of the original system also does so.

How do we over-approximate? Abstractions

Nonlinear dynamics: $x_{t+1} = f(x_t, u_t) \qquad (x_t, u_t) \in \mathcal{X} \times \mathcal{U}$

Discrete-absractions:

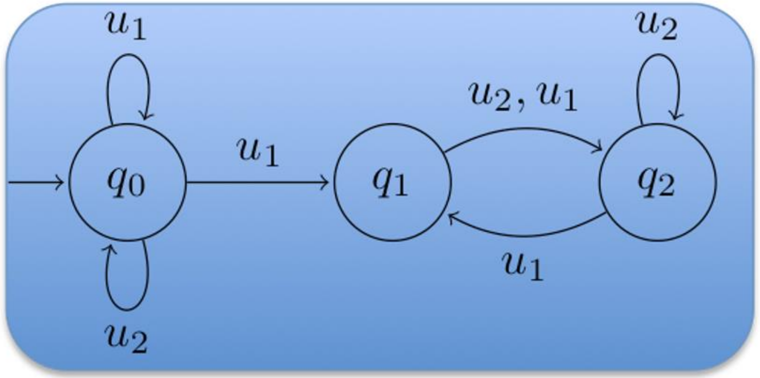

Abstraction can be learned: Abate, Soudjani, Mazo, Lavaei, etc…

Theorem (informal)

If the abstraction function $A_x: \mathcal{X} \to Q, A_u: \mathcal{U} \to \mathcal{U}_a$ is such that for any x,
$f(x,u) \in A_x^{-1}(Post(A_x(x), A_u(u)))$
then any policy that enforces a reach-avoid specification can be concretized such that the trajectory of the original system also does so.

System over-approximation

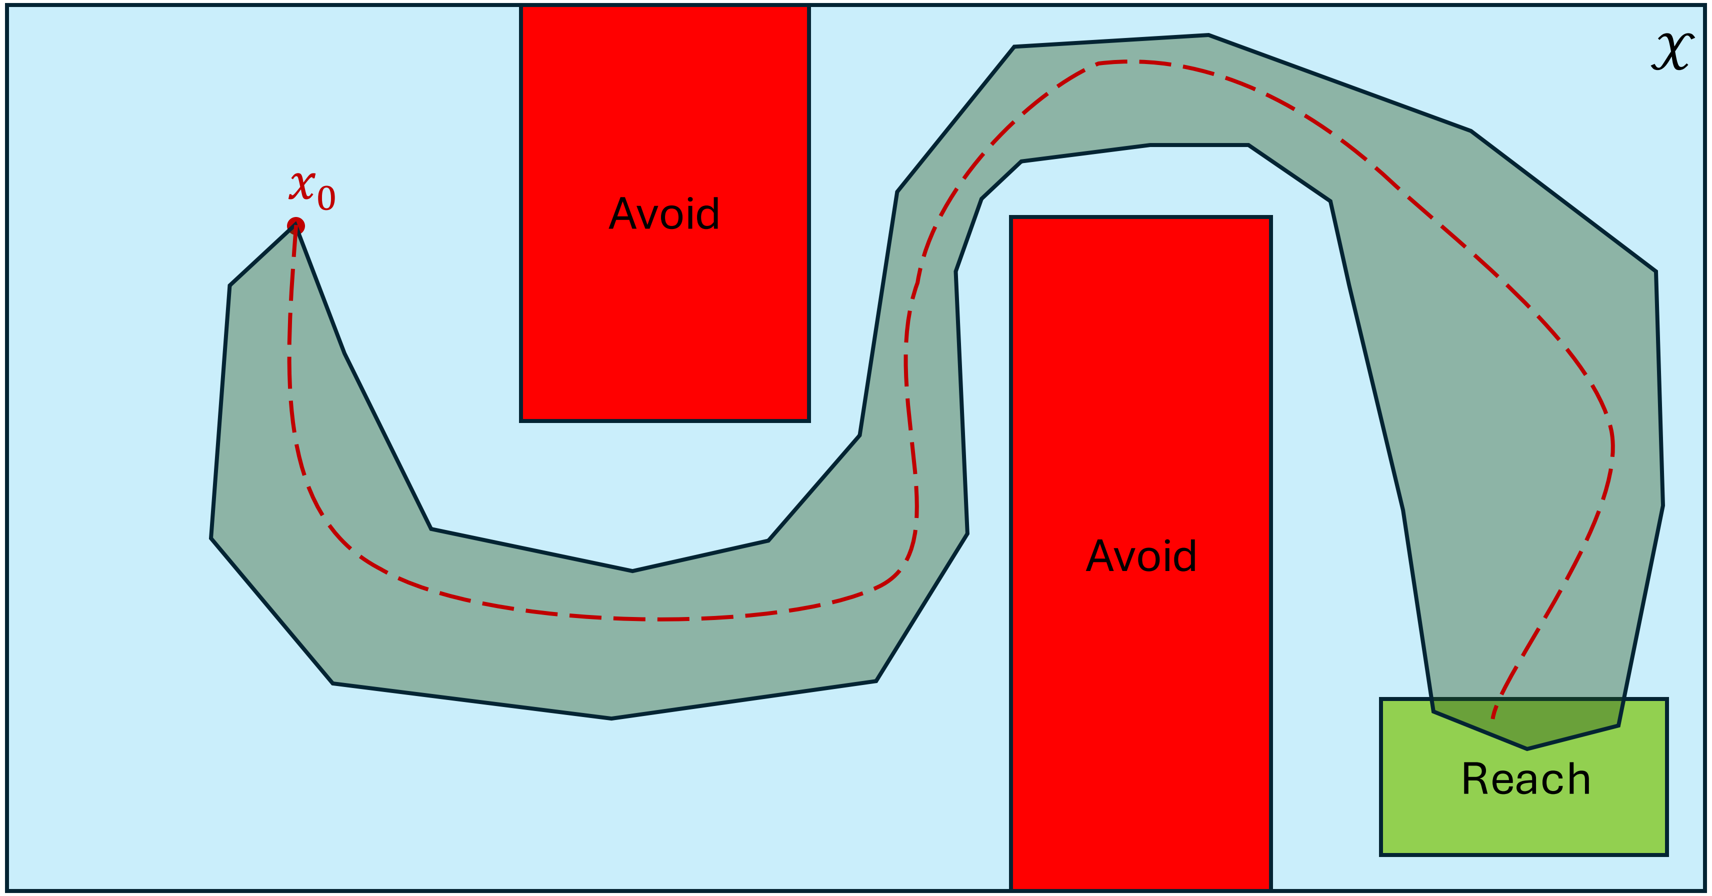

We trade nonlinearity for non-determinism
Much smaller state-space, dynamics on a graph!

System over-approximation

Nonlinear dynamics: $x_{t+1} = f(x_t, u_t) \qquad (x_t, u_t) \in \mathcal{X} \times \mathcal{U}$

Linear approximation: $f(x,u) \approx \hat{f}(x,u) \coloneqq Ax + Bu$

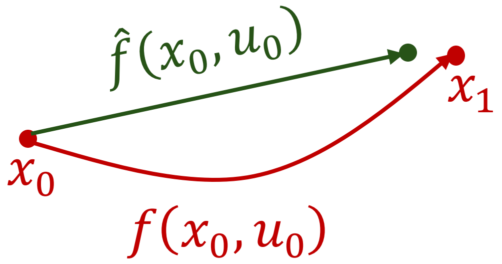

No formal guarantees…💡**Include the error in the dynamics**

$\forall (x,u) \in \mathcal{X} \times \mathcal{U}\colon\ f(x,u) - \hat{f}(x,u) \in \mathcal{E}$

$x_{t+1} \in \hat{f}(x_t, u_t) \oplus \mathcal{E} \stackrel{\text{def}}{=} \{\hat{f}(x_t, u_t) + e \mid e \in \mathcal{E}\}$

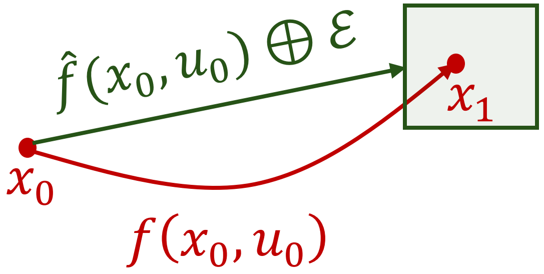

Theorem (informal)

For any x_0, if all trajectories of

$$x_{t+1} \in \hat{f}\big(x_t, \pi(x_t)\big) \oplus \mathcal{E}$$

satisfy a reach-avoid specification, then so does the trajectory of

$$x_{t+1} = f\big(x_t, \pi(x_t)\big).$$

System over-approximation

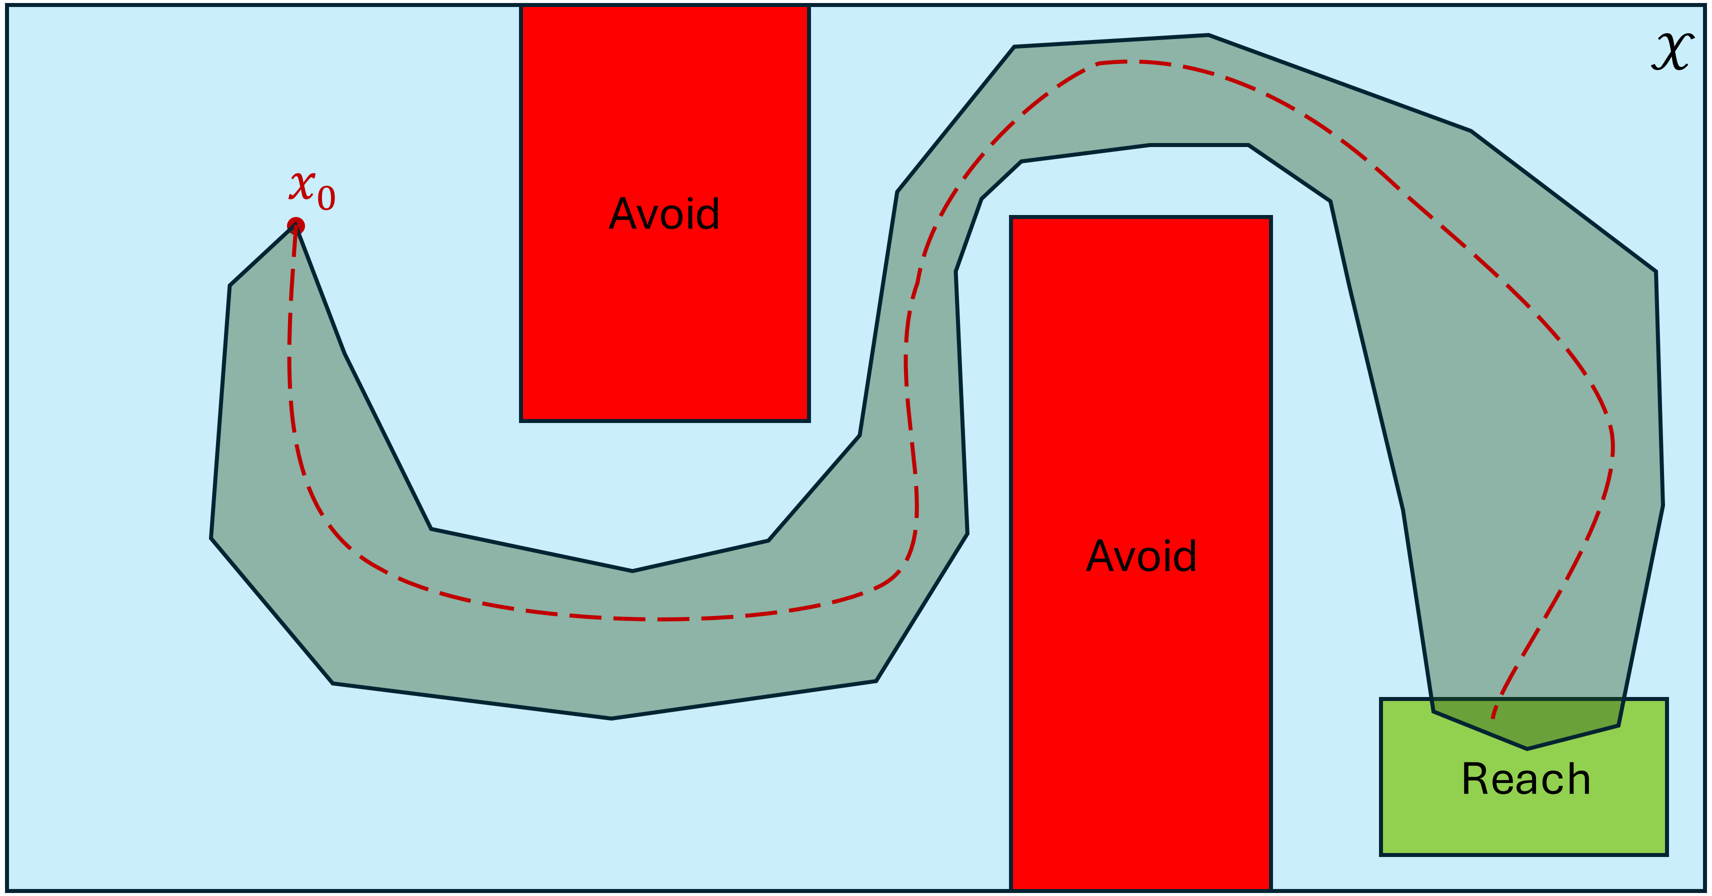

We trade nonlinearity for non-determinism

Hybridization [1]

Over a large domain $\mathcal{X}$, the error set $\mathcal{E}$ can be large → conservative results

Idea 💡: Piece-wise affine over-approximation

$\mathcal{X}$

$\mathcal{X}_1$ $x_{t+1} \in A_1 x + B_1 u \oplus \mathcal{E}_1$	$\mathcal{X}_2$ $x_{t+1} \in A_2 x + B_2 u \oplus \mathcal{E}_2$	…	…
…	…	…	$\mathcal{X}_K$ $x_{t+1} \in A_K x + B_K u \oplus \mathcal{E}_K$

Hybrid dynamics $\quad x_{t+1} \in \hat{f}_{\sigma(x)}(x, u) \oplus \mathcal{E}_{\sigma(x)} := A_{\sigma(x)} x + B_{\sigma(x)} u \oplus \mathcal{E}_{\sigma(x)}$

with partition function $\quad \sigma: \mathcal{X} \to \{1, \dots, K\}$

Universal approximator [1] 🙂… but complexity of synthesizing a policy increases with K 🙁

[1] E. Asarin, T. Dang & A. Girard, "Hybridization methods for the analysis of nonlinear systems," Acta Informatica, 2007

Outline

1. Introduction
2. Over-approximations
3. **Lifted over-approximations**
4. Conclusion

Lifted over-approximation

Idea — Immersions, Koopman,... [Mezic, Brunton, Korda,...]:

Lift the state to a higher dimensional space:

$$z_t = \psi(x_t) \in \mathbb{R}^{n_z}$$

Linear approximation: $\psi\big(f(x,u)\big) \approx \hat{f}(\psi(x),u) \coloneqq A\psi(x) + Bu$

No formal guarantees...💡**Include the error in the dynamics**

Linear over-approximation: $z_{t+1} \in \hat{f}(z_t,u_t) \oplus \mathcal{E}$ with

$\forall (x,u) \in \mathcal{X} \times \mathcal{U}$: $\psi\big(f(x,u)\big) - \hat{f}(\psi(x),u) \in \mathcal{E}$

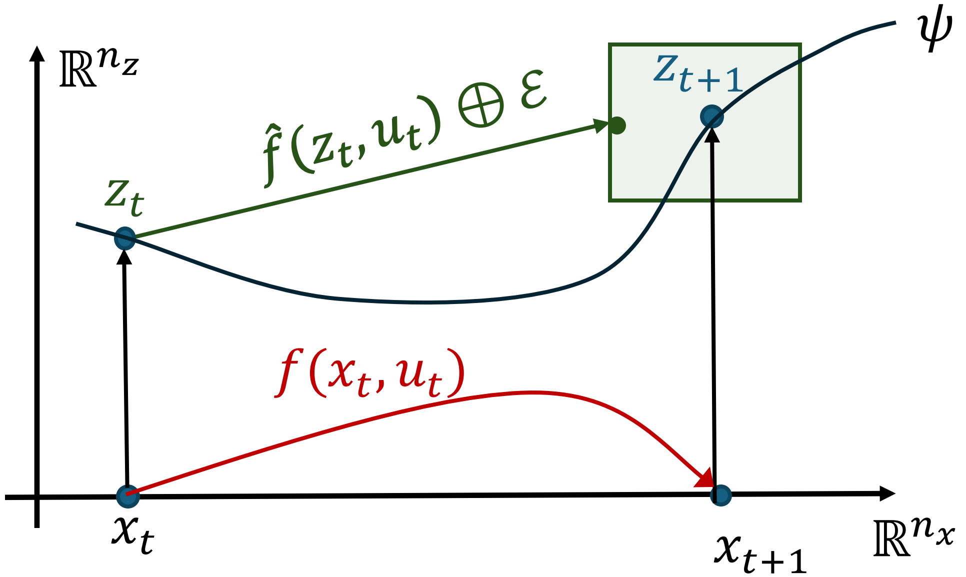

Include the original state in the lifting:

$$\psi(x) = \begin{bmatrix} x \\ \phi(x) \end{bmatrix}$$

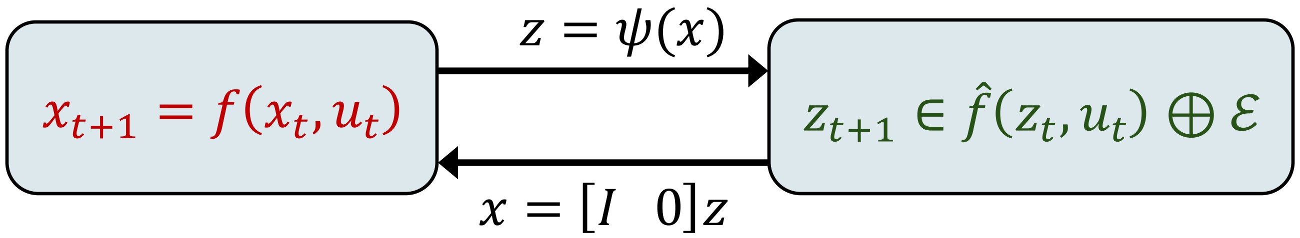

Lifting specifications

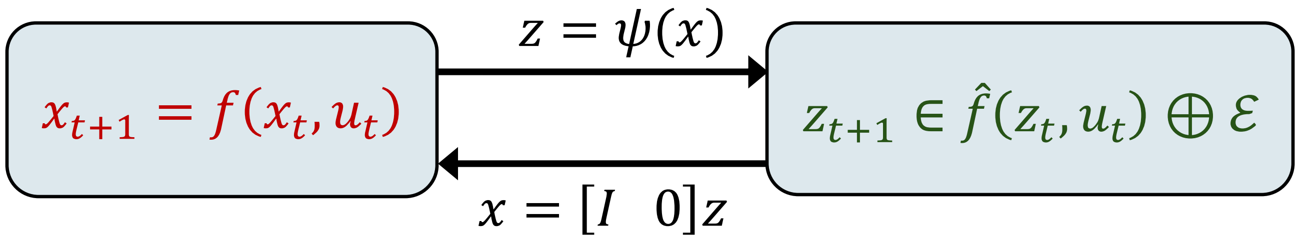

Reach-avoid specifications are in $\mathbb{R}^{n_x}$ → They need to be lifted to $\mathbb{R}^{n_z}$

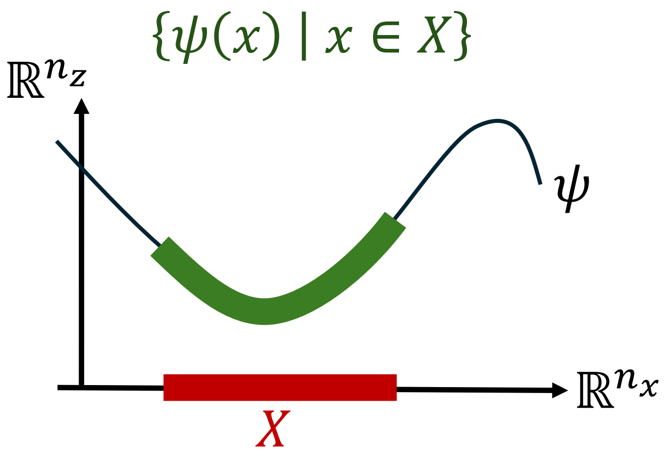

Not convex ☹
(even when X is)

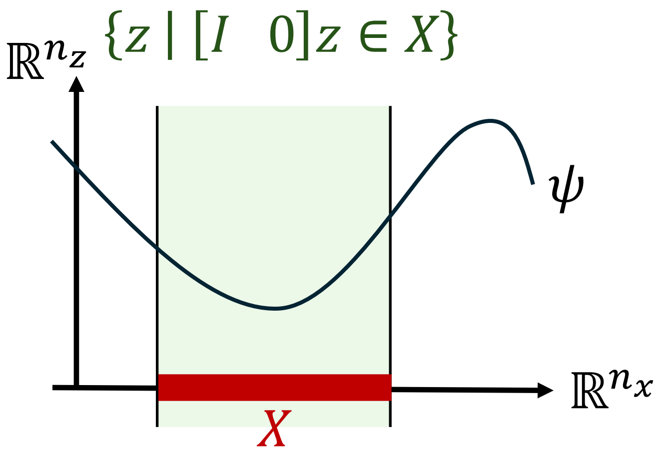

Preserves convexity 🙂
Preserves polyhedrality 🙂

Lifted over-approximation

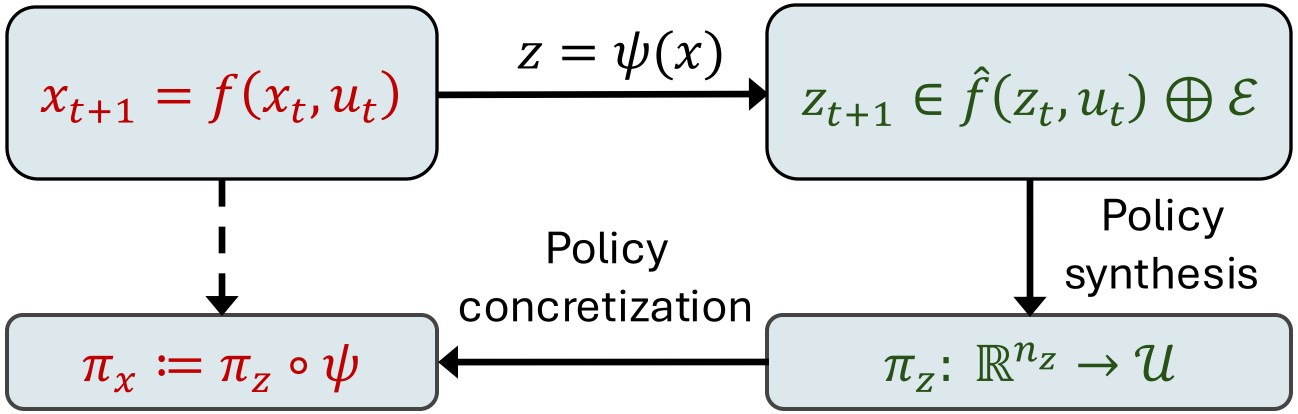

$$u_t = \pi_z(\psi(x_t))$$

Theorem (informal)

For any x_0, if all $(z_t)_{t=0}^{T}$ given by

$$z_{t+1} \in \hat{f}(z_t, \pi(z_t)) \oplus \mathcal{E}$$
$$z_0 = \psi(x_0)$$

satisfy a **lifted** reach-avoid specification, then the trajectory of

$$x_{t+1} = f\left(x_t, \pi\big(\psi(x_t)\big)\right)$$

satisfies the unlifted specification.

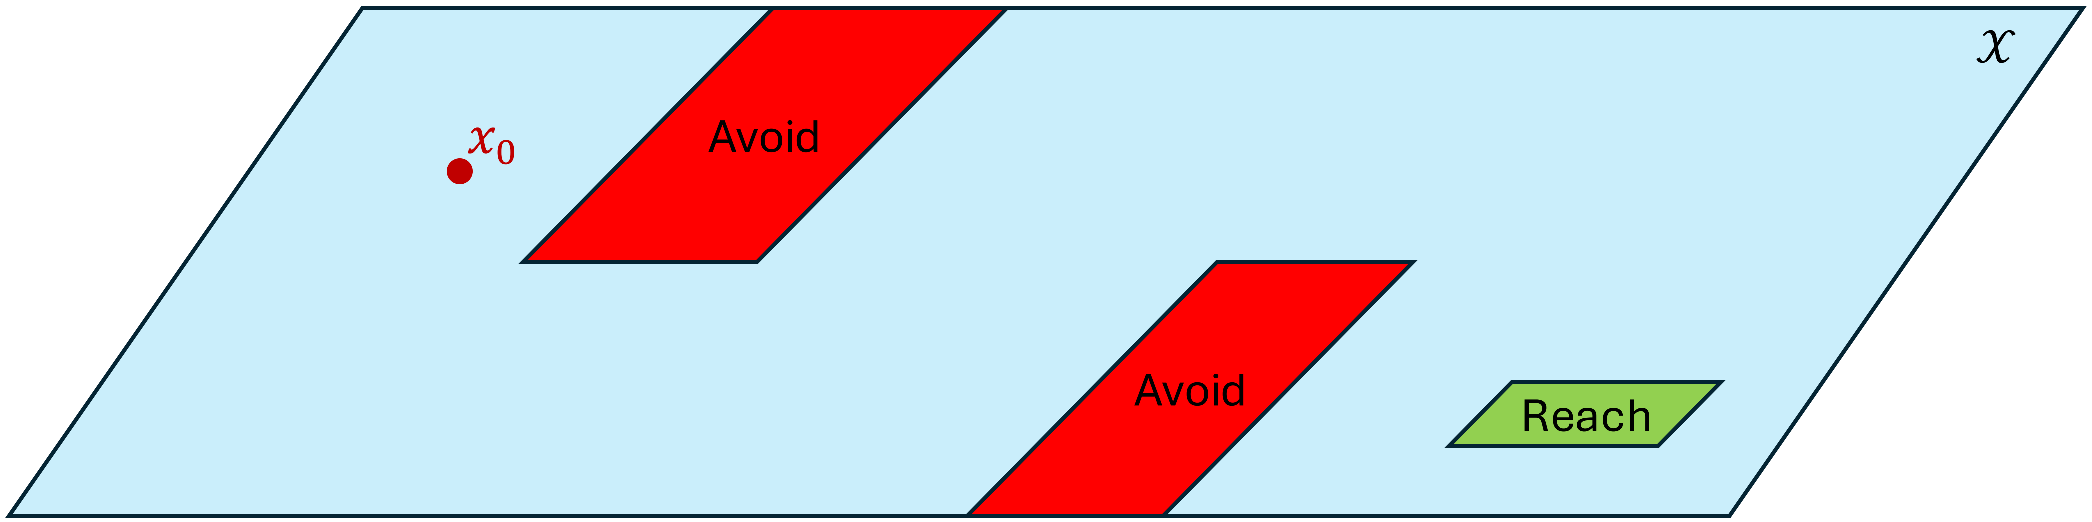
x_0
Avoid
Avoid
Reach
$\mathcal{X}$

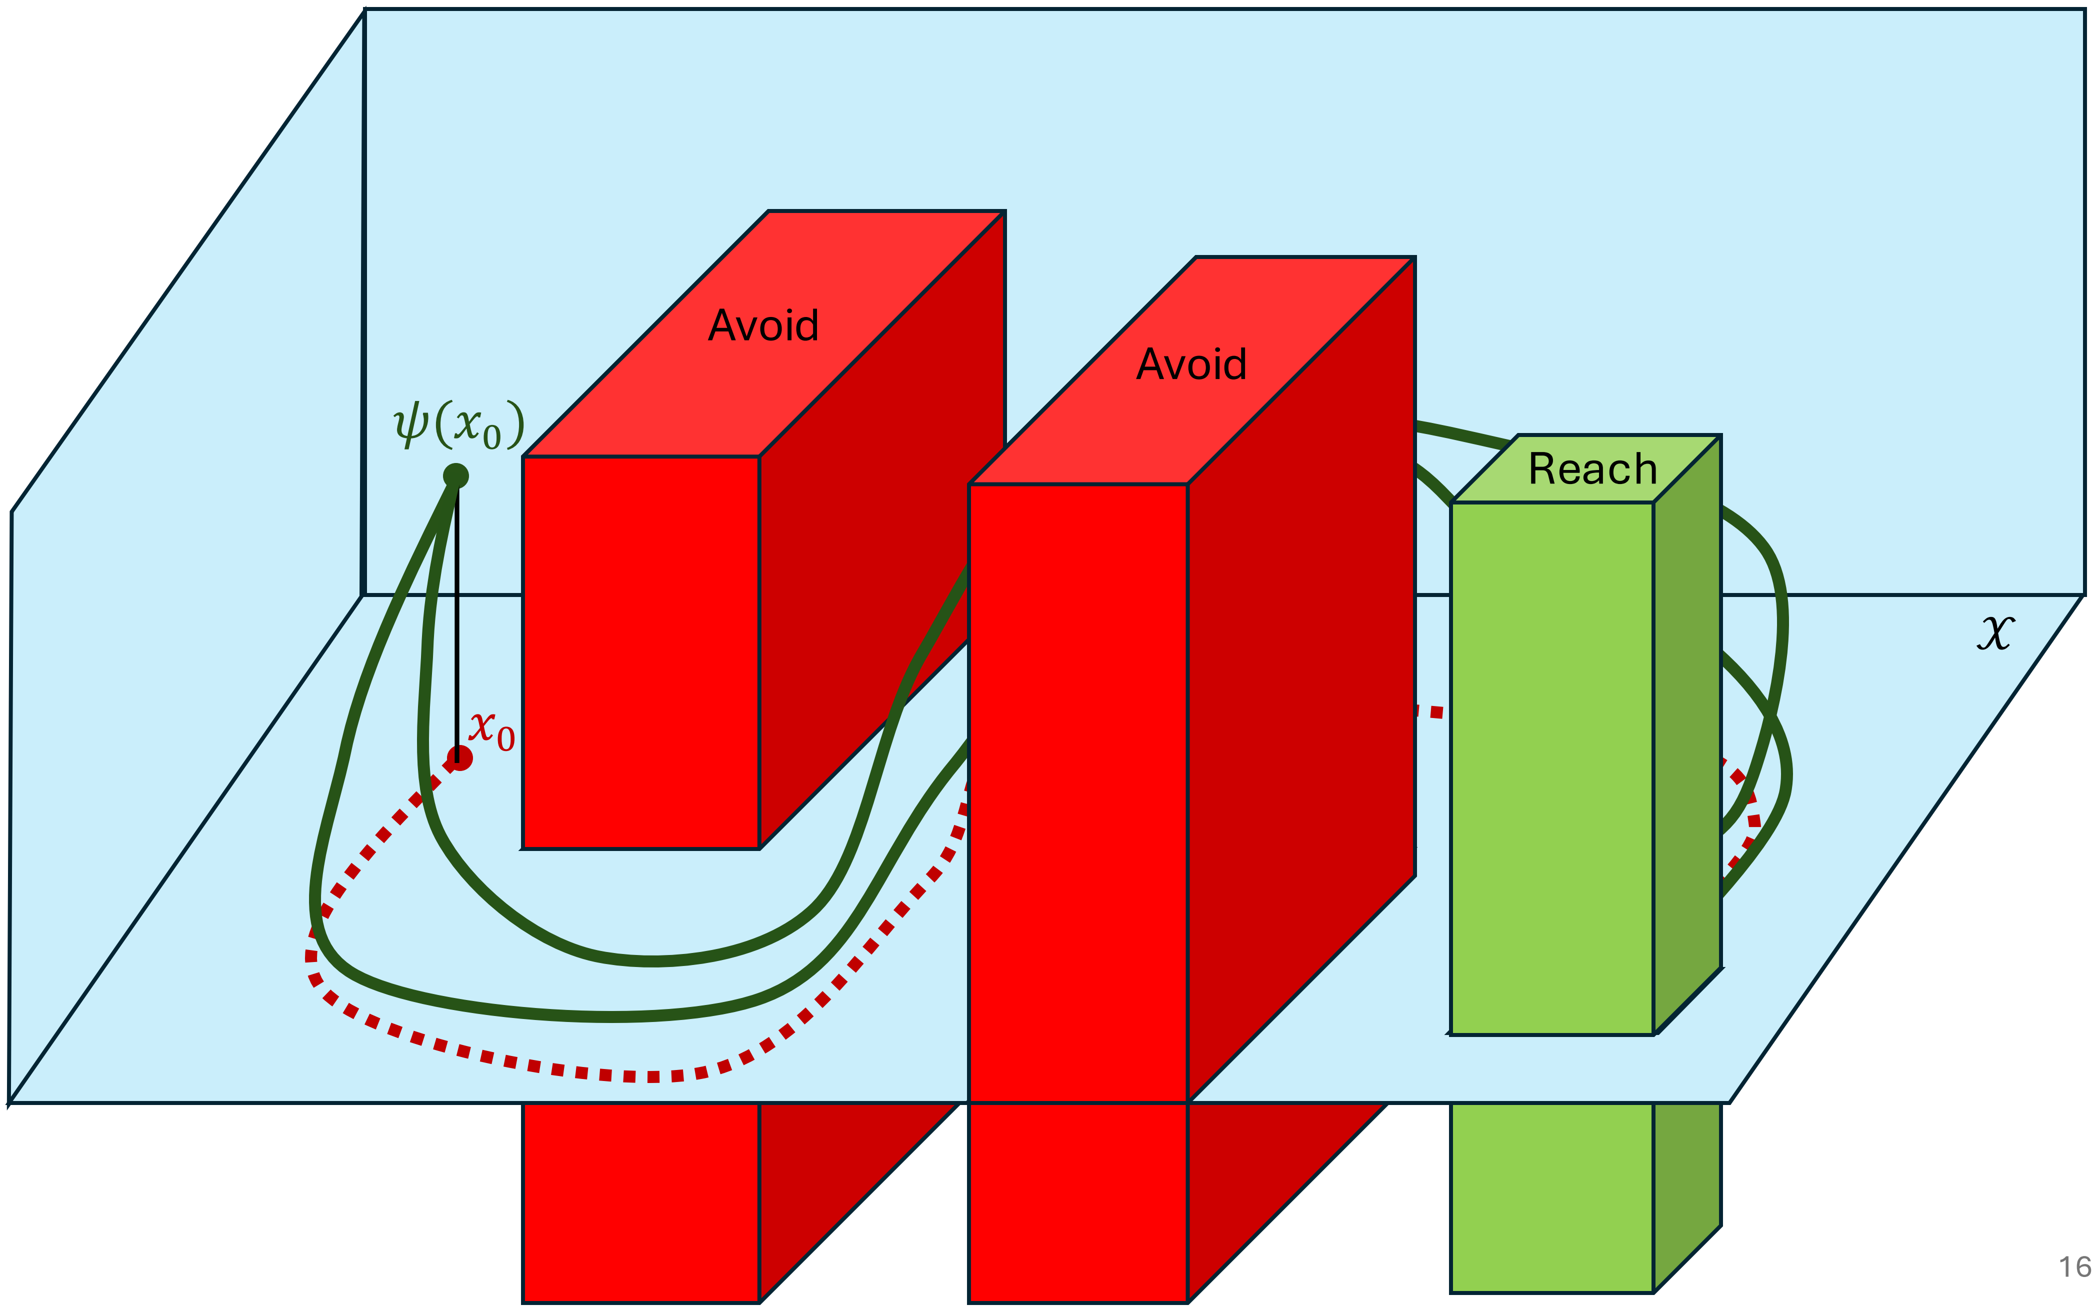
Avoid
Avoid
Reach
$\psi(x_0)$
x_0
$\mathcal{X}$

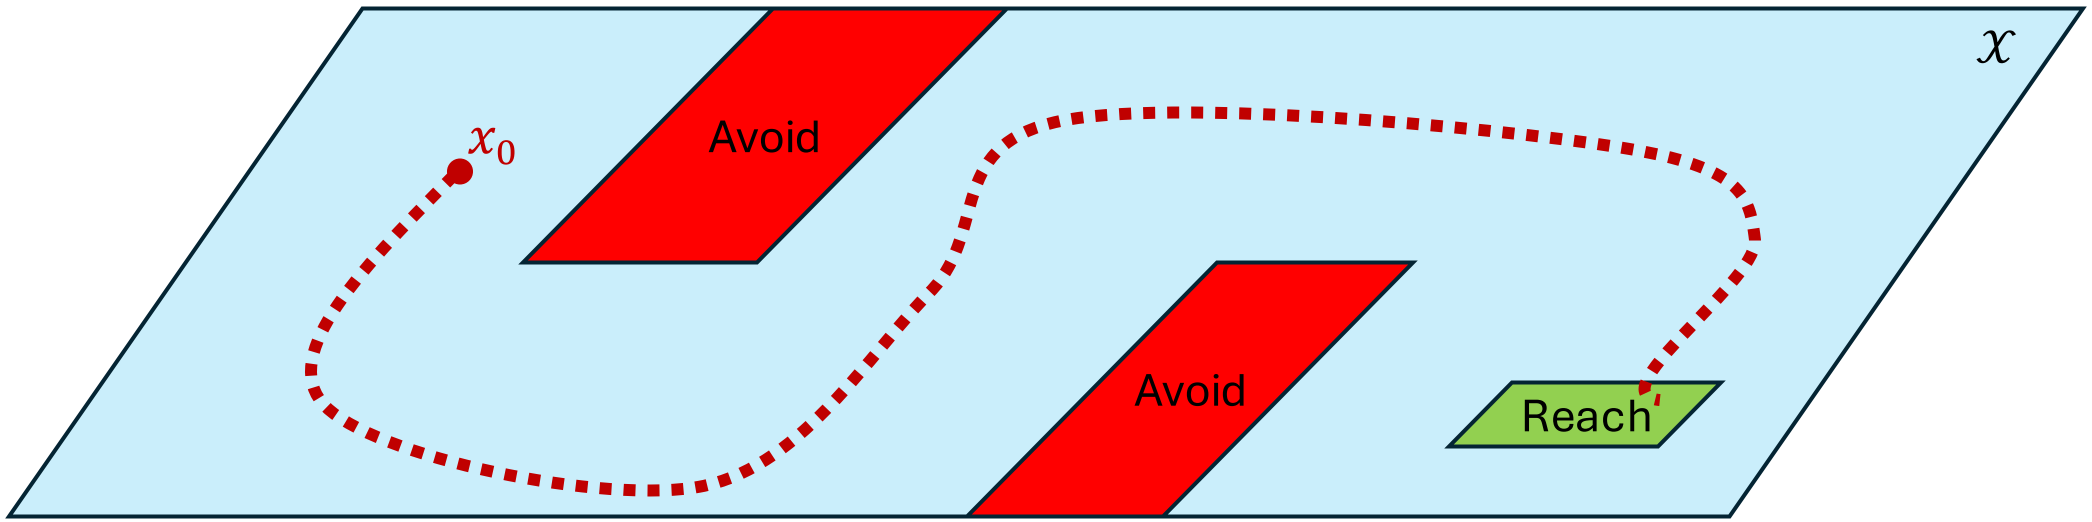
$\mathcal{X}$
x_0
Avoid
Avoid
Reach

Computing linear lifted over-approximations

Find A, B and $\mathcal{E}$ such that

$$\forall (x,u) \in \mathcal{X} \times \mathcal{U}: \quad \psi\big(f(x,u)\big) - A\psi(x) - Bu \in \mathcal{E}$$

with $\mathcal{E} = \times_{i=1}^{n_z} [m_i, M_i]$

$$\min_{A,B,m,M} \sum_{i=1}^{n_z} M_i - m_i \quad \text{s.t.} \quad \forall (x,u) \in \mathcal{X} \times \mathcal{U}: \quad m \leq \psi\big(f(x,u)\big) - A\psi(x) - Bu \leq M$$

Minimize nondeterminism

Infinite dimensional 🙁

Assumptions

1. f and ψ are polynomials
2. $\mathcal{X}$ and $\mathcal{U}$ are polytopes (or semi-algebraic sets)

→ Can be handled using Sum-of-Squares optimization 🙂

Learning linear lifted over-approximations

Find A, B and $\mathcal{E}$ such that

$$\forall (x,u) \in \mathcal{X} \times \mathcal{U}: \quad \psi\big(f(x,u)\big) - A\psi(x) - Bu \in \mathcal{E}$$

with $\mathcal{E} = \times_{i=1}^{n_z} [m_i, M_i]$

Minimize nondeterminism

$$\min_{A,B,m,M} \sum_{i=1}^{n_z} M_i - m_i \quad \text{s.t.} \quad \forall (x,u) \in \mathcal{X} \times \mathcal{U}: \quad m \leq \psi\big(f(x,u)\big) - A\psi(x) - Bu \leq M$$

Infinite dimensional 🙁

Assumptions

1. Lipschitz constants of f and ψ
2. Samples of $\mathcal{X} \times \mathcal{U}$

→ Can be handled by solving a Linear Program 🙂

Experiments : Backward reachable sets

Inverted pendulum
$\ddot{\theta} = 15 \sin\theta + 30u$

Discretized with Euler $dt = 0.1$

Input set $\mathcal{U} = [-0.35, 0.35]$

Lifting function $\psi\left(\begin{bmatrix}\theta \\ \dot{\theta}\end{bmatrix}\right) = \begin{bmatrix}\theta \\ \dot{\theta} \\ \sin\theta\end{bmatrix}$

Lifted hybridization with partition of $\mathcal{X}$ only

	# cells K	Comp. Time [s]
No lifting	640	952.6
With lifting	118	245.4

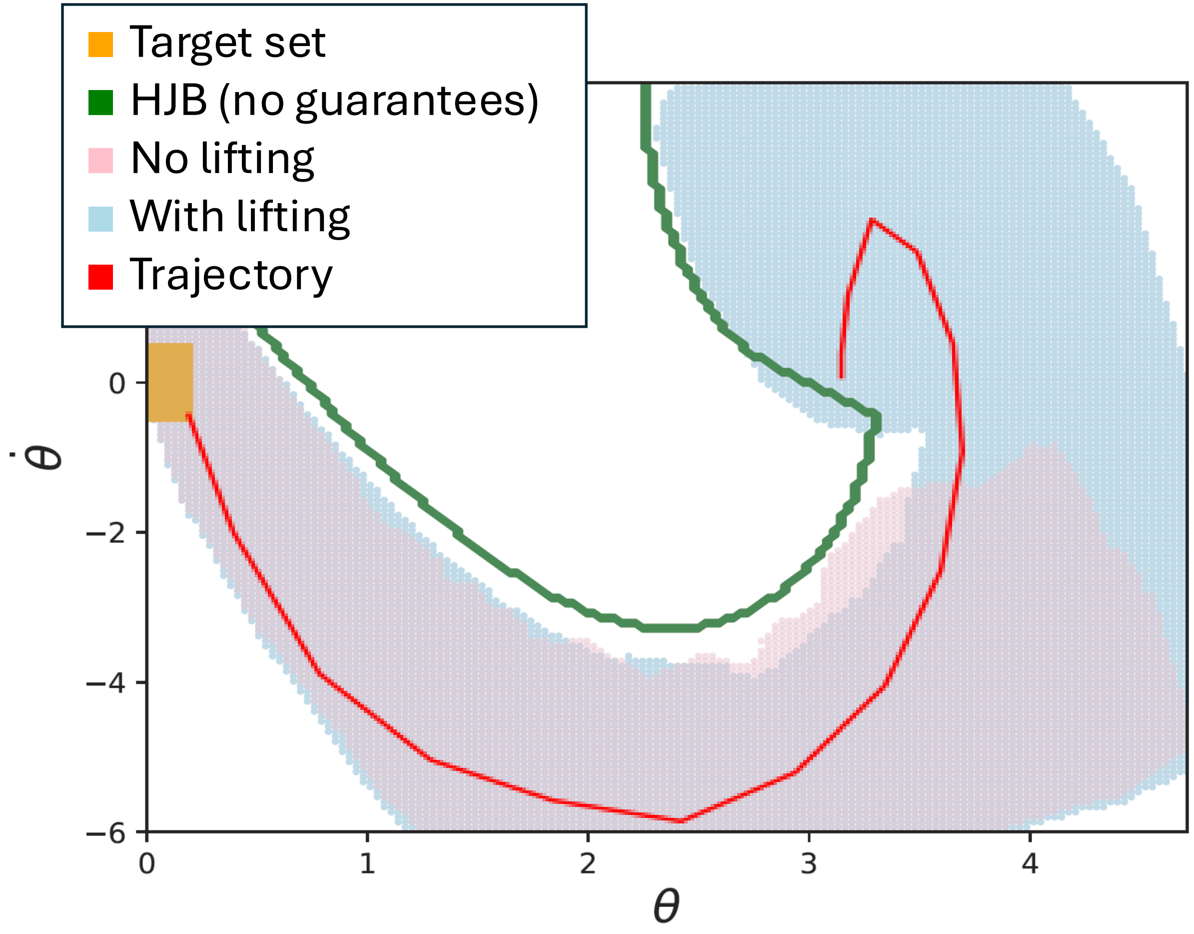

Lifting reduces both conservatism and computation time

t = 0

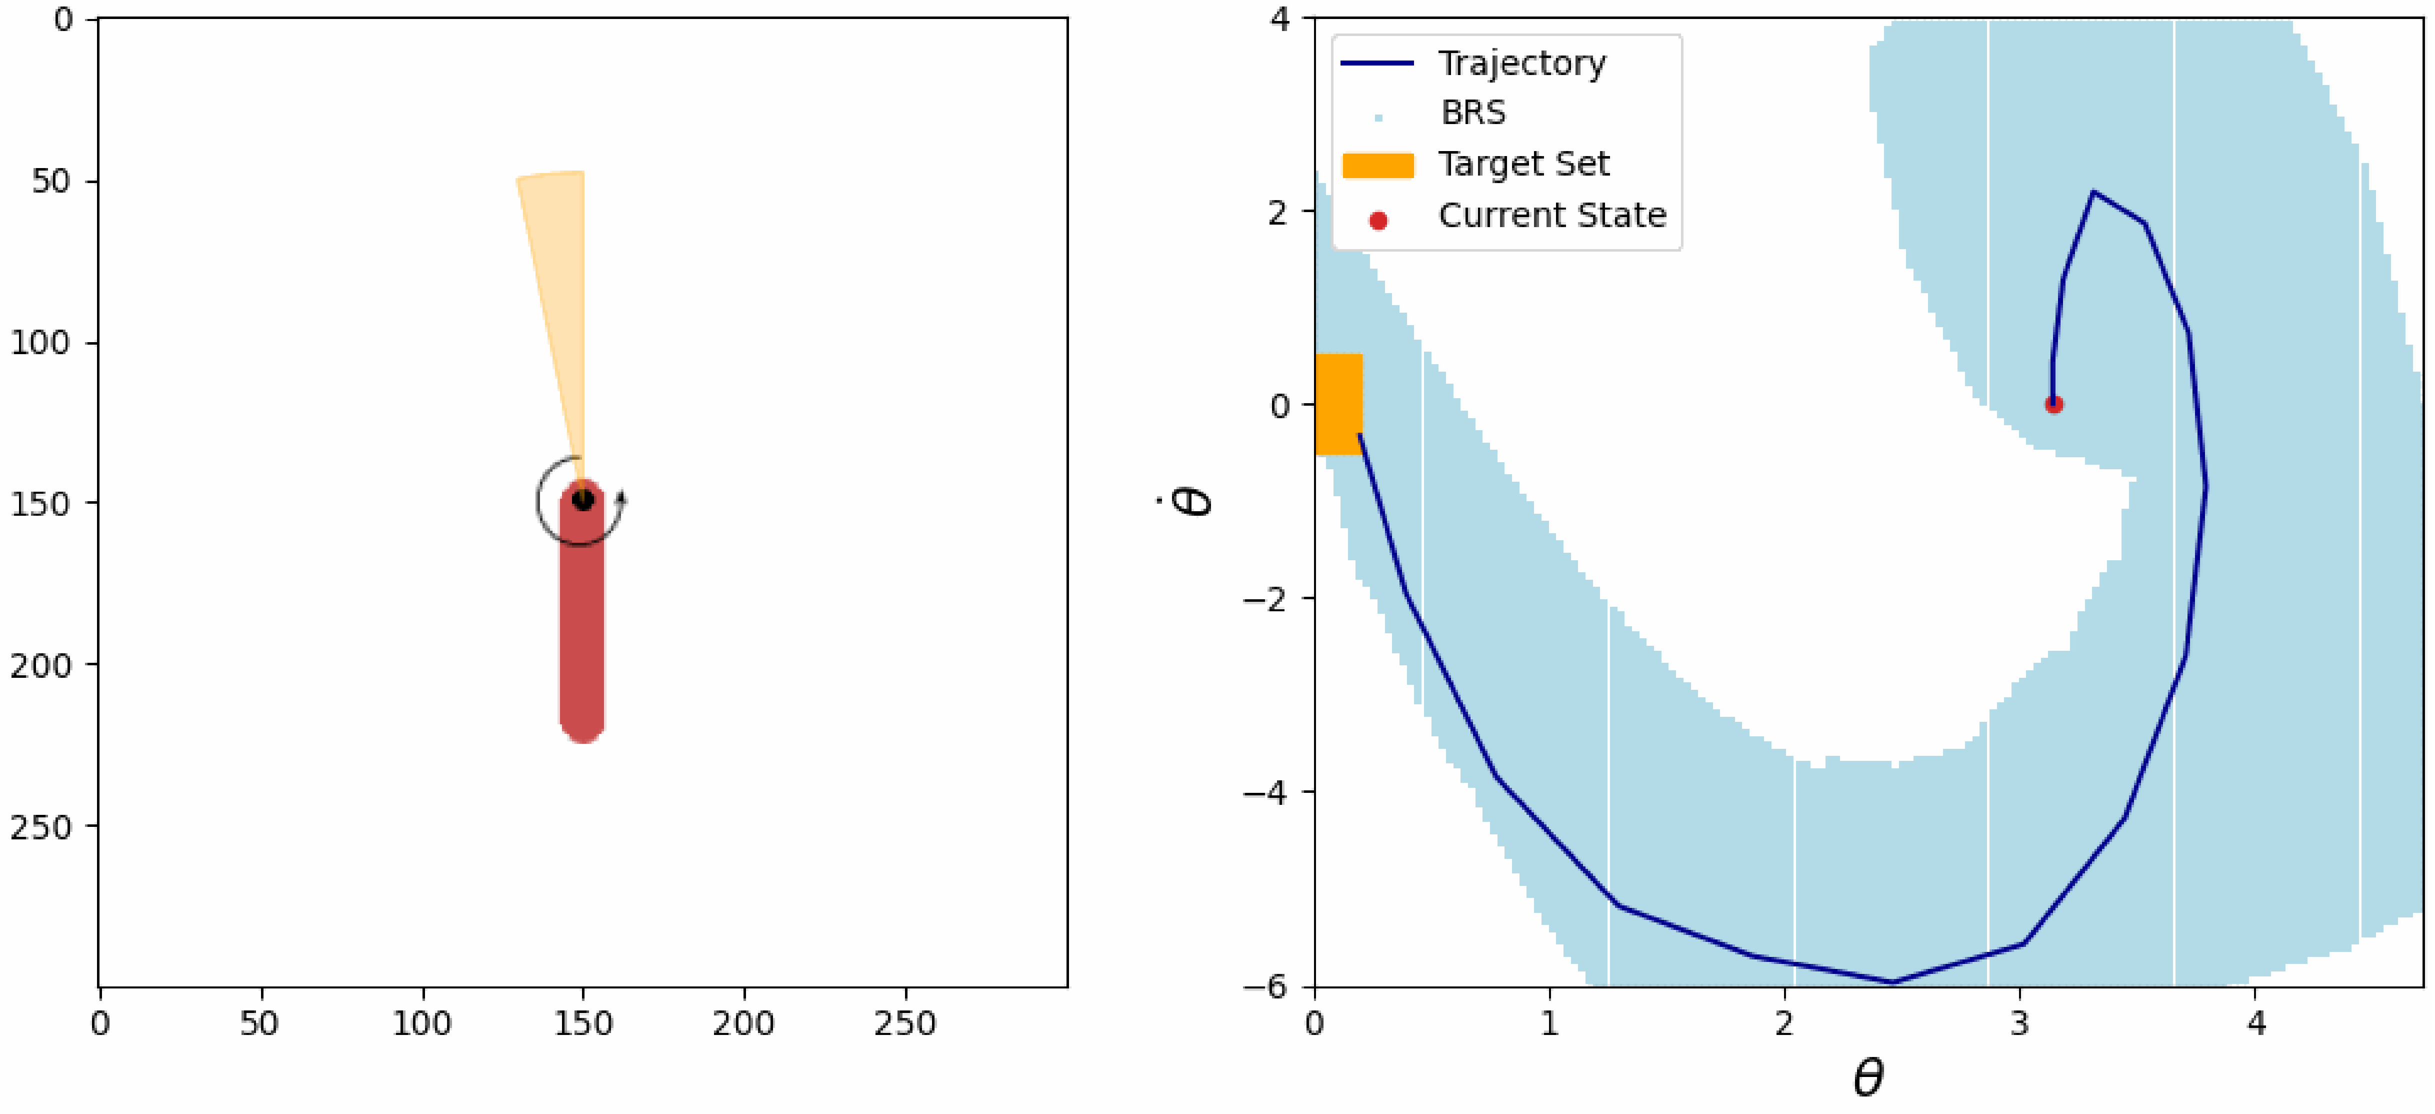

0
50
100
150
200
250
0
50
100
150
200
250
Trajectory
BRS
Target Set
Current State
4
2
0
−2
−4
−6
$\dot{\theta}$
0
1
2
3
4
θ

Outline

1. Introduction
2. Over-approximations
3. Lifted over-approximations
4. **Conclusion**

Lifting as abstractions

Nonlinear system
$x^+ = f(x,u)$

$z = \psi(x)$ →
← $x = Cz$

Koopman over-approximation
$z^+ \in A_i z + B_i u + W_i$ if $x \in D_i$

Hard ↓ Backward Reachable Set X_t

Easy ↓ Backward Reachable Set Z_t

Inner approximates: Backward Reachable Set X_t ← Backward Reachable Set Z_t

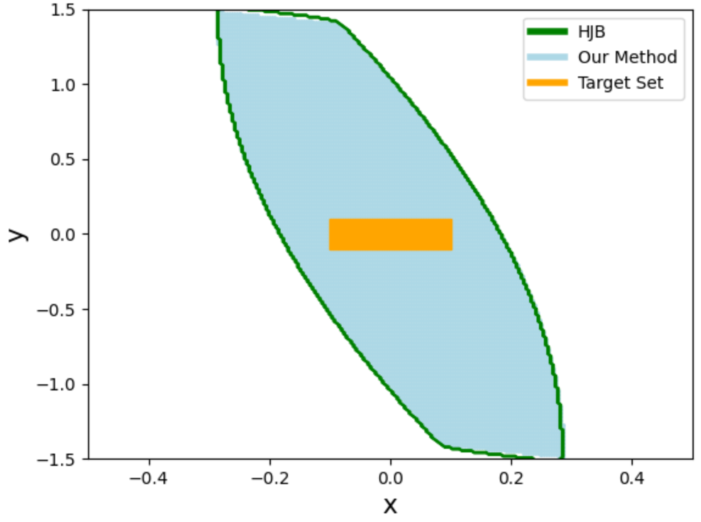

Inverted Pendulum

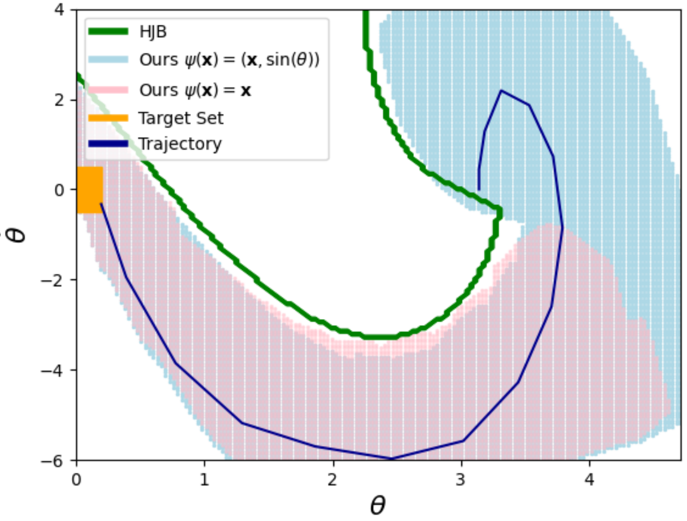

Balim, Aspeel, Liu, Ozay, L-CSS'23

Liftings give an alternative way of doing abstractions:

- Generalizes hybridization.
- Single linearization can be conservative or complex.
- Learning different over-approximations are learned over local subdomains (leading to a PWA system) for better accuracy:
  - Experiments show that to obtain BRSs with similar sizes, the Koopman over-approximation requires less pieces than direct linearization (hybridization).

Why do we need hybridization in the lifted space? See also work on non-existence of linear immersions for systems with multiple omega limit sets (Liu, Ozay, Sontag, Automatica'25)

References & collaborators

- Haldun Balim, Antoine Aspeel, Zexiang Liu, Necmiye Ozay (2023). Koopman-inspired implicit backward reachable sets for unknown nonlinear systems. *IEEE Control Systems Letters*.
- Antoine Aspeel, Necmiye Ozay (2024). A simulation preorder for koopman-like lifted control systems. *IFAC-PapersOnLine*.
- Zexiang Liu, Necmiye Ozay, Eduardo Sontag (2025). “Properties of Immersions for Systems with Multiple Limit Sets with Implications to Learning Koopman Embeddings”, *Automatica*.

Haldun Balim

Zexiang Liu

Antoine Aspeel

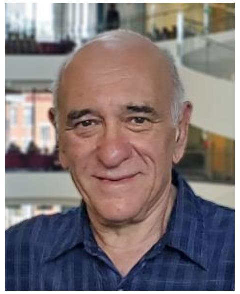

Eduardo Sontag

UNIVERSITY OF MICHIGAN