Data-Driven Safety Verification using Reachability Analysis

Amr Alanwar

Technical University of Munich

Motivation

- Many accidents in using Tesla and Uber self-driving cars
- Dealing with industrial robots is a threat to human life

How to guarantee safety when dealing with CPS?

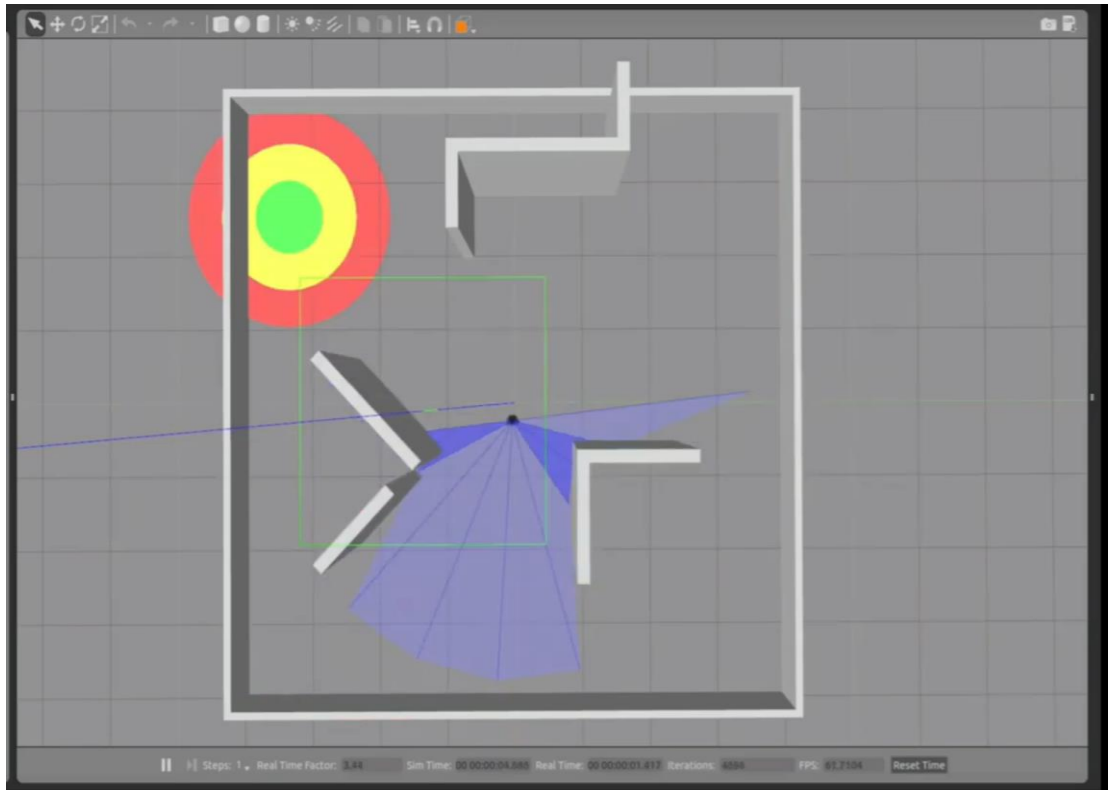

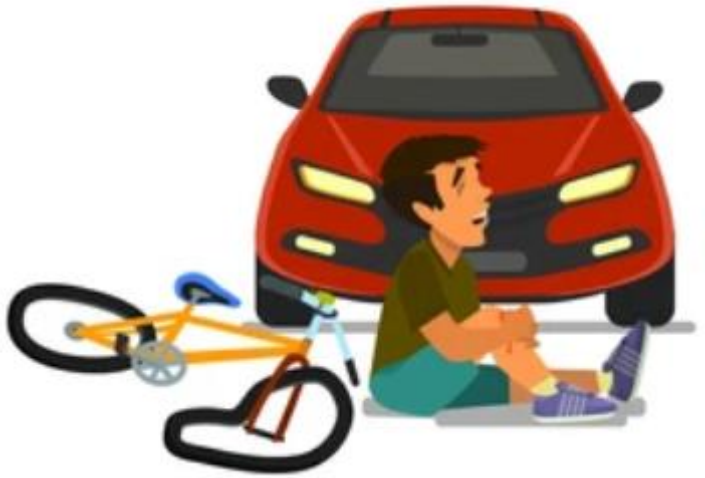

Can Testing Guarantee Safety?

- Consider a pedestrian crossing the road
- We might do n test scenarios
- Failure may happen at scenario n+1
- It is time for formal safety guarantees during learning and control

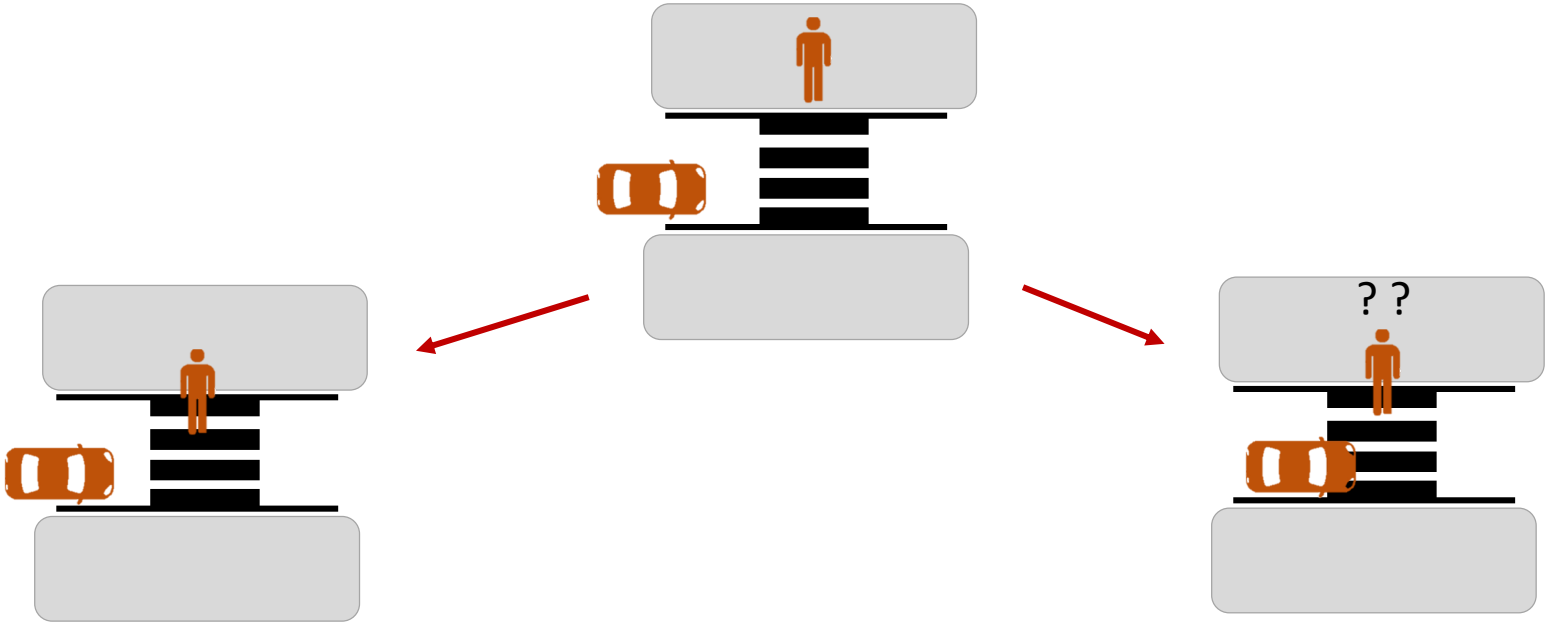

Safety Guarantees through Reachability Analysis

- State $x(k)$ can be position, speed, acceleration … etc.
- Reachability analysis computes the set of reachable states of a dynamical system with uncertain initial states, inputs, and parameters
- Reachability analysis traditionally requires a model

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

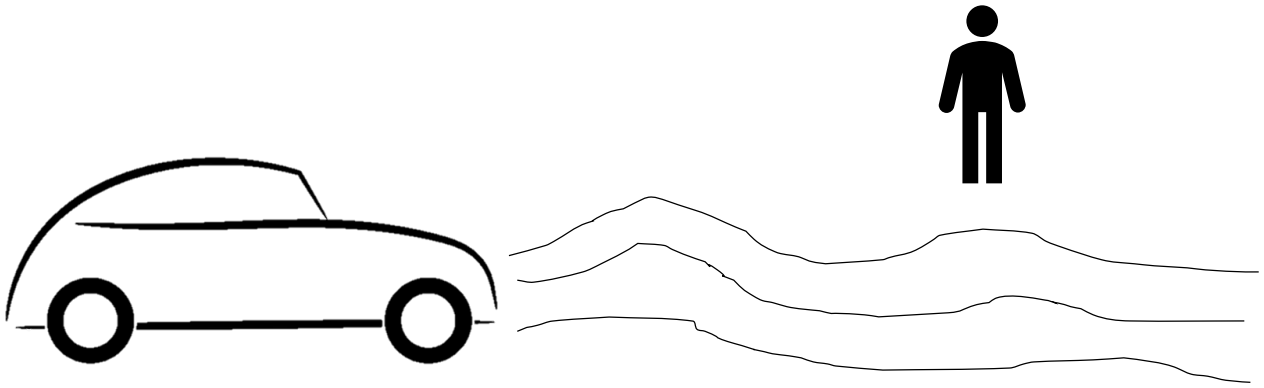

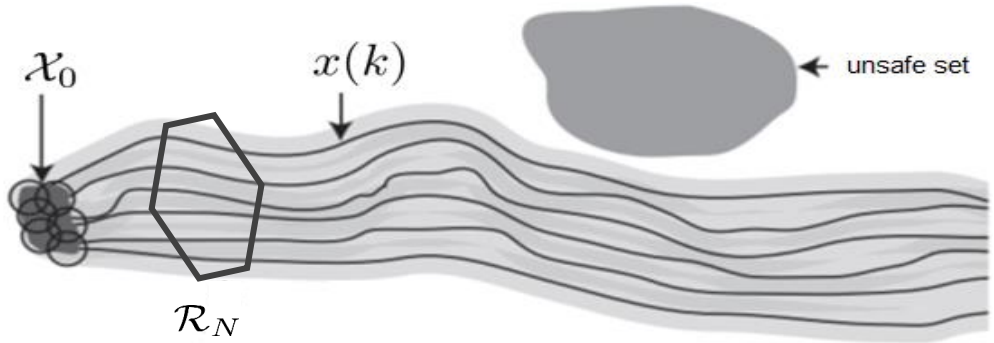

Effect of Model Choice

- CPS are becoming too complex to model
- One often has an abundance of data but no model to guarantee safety

Can we depend on data without trusting a single model?

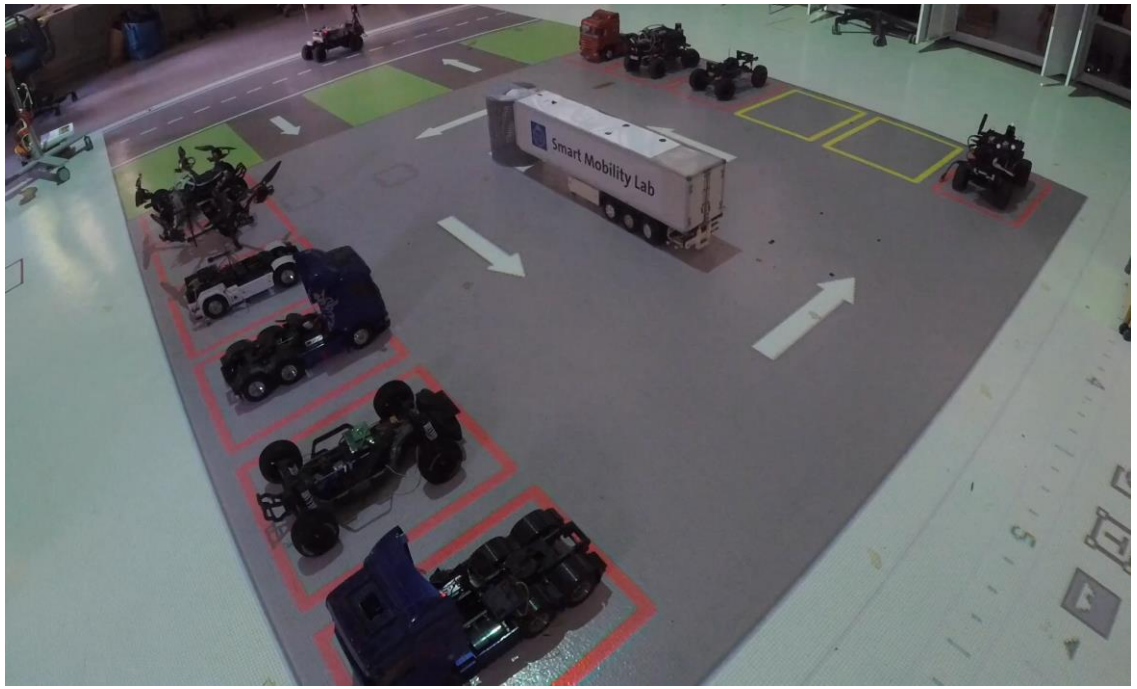

Untuned model

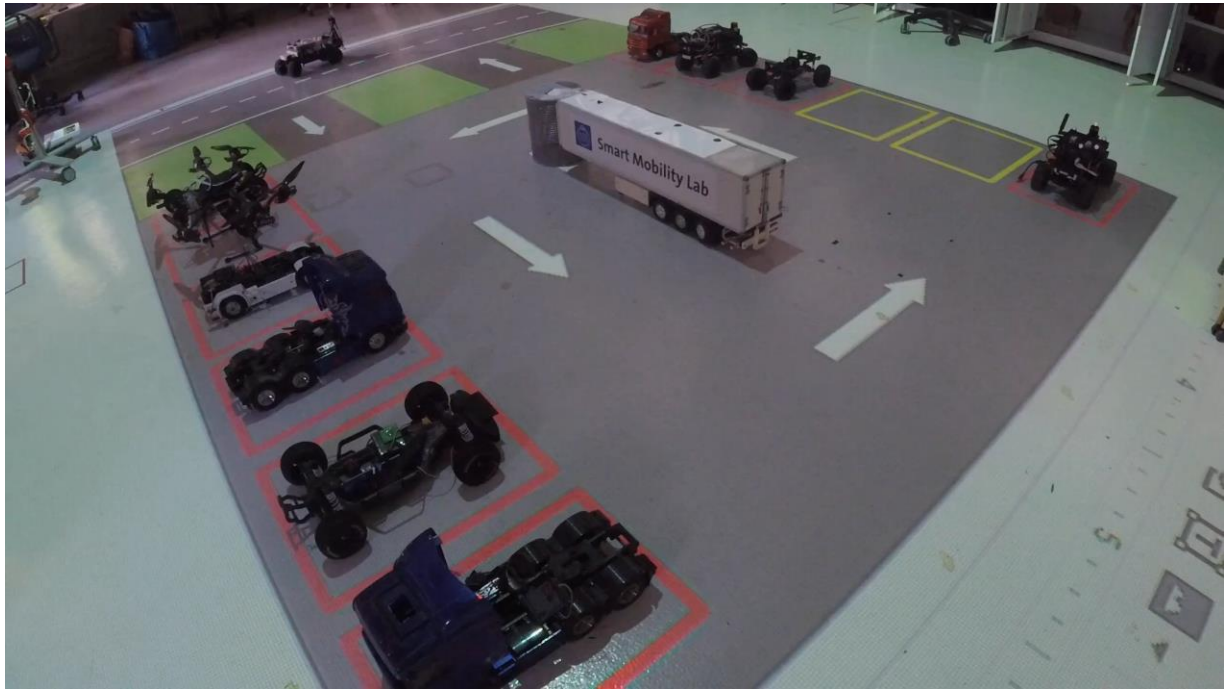

Tuned model

Problem Formulation

How to compute (over-approximate) reachable sets from data?

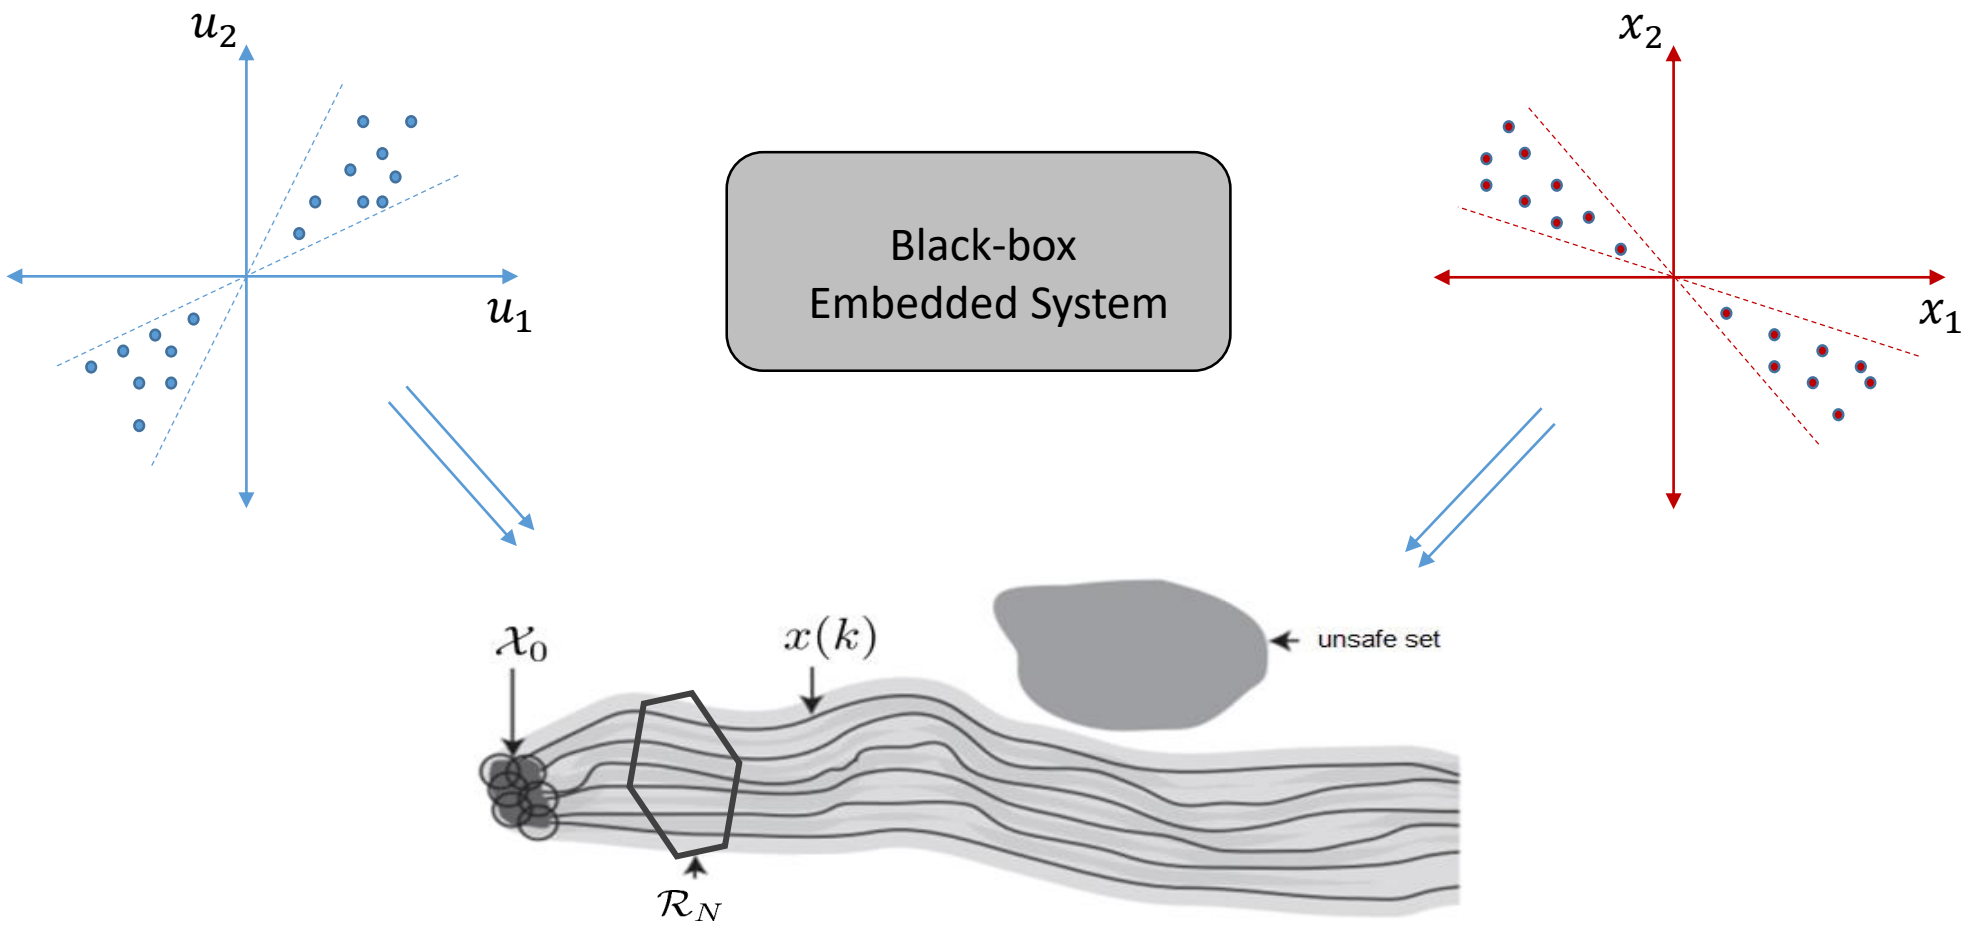

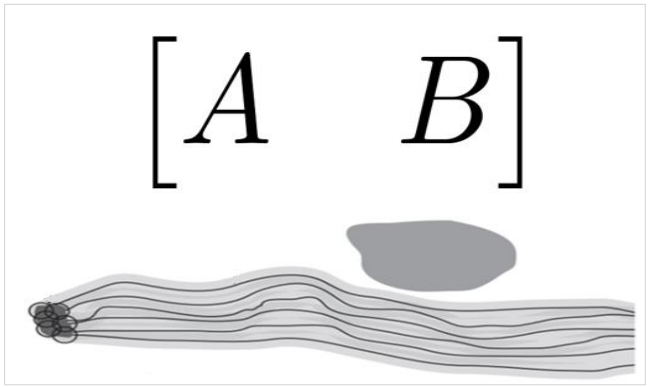

1- Basic Reachability

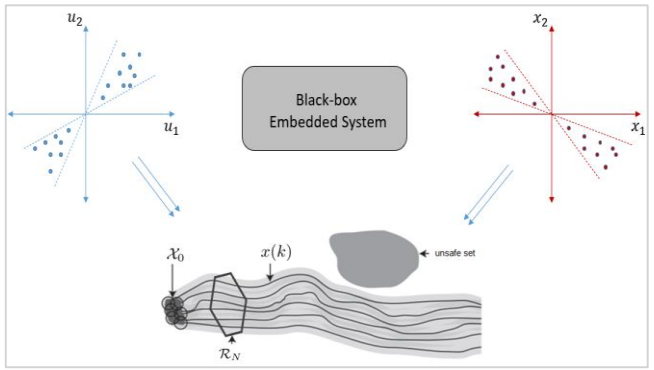

2- Data-driven Reachability

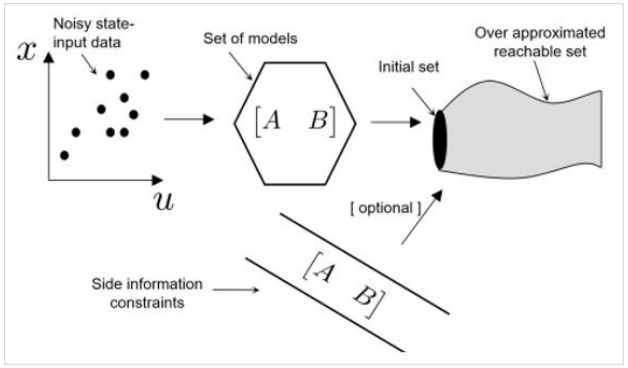

3- Side Information

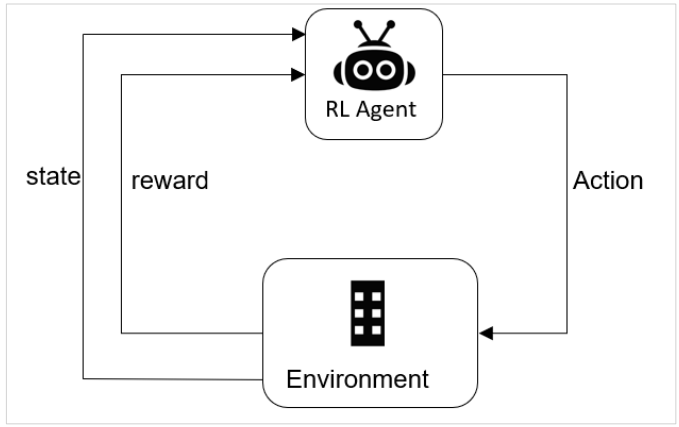

5- Safe Reinforcement Learning

Zonotope

A **zonotope** $\mathcal{Z} = \langle c, G \rangle$ is a set

$$\mathcal{Z} = \left\{ x \in \mathbb{R}^n \;\middle|\; x = c + \sum_{i=1}^{\gamma} \beta_i \, g_i \, , -1 \leq \beta_i \leq 1 \right\}$$

where $c \in \mathbb{R}^n$ is the center and $G = [g_1, ..., g_\gamma] \in \mathbb{R}^{n \times \gamma}$ the generator vectors

A matrix zonotope $\mathcal{M}$ consists of center matrix and list of generator matrices

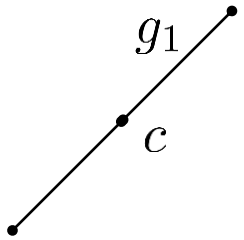

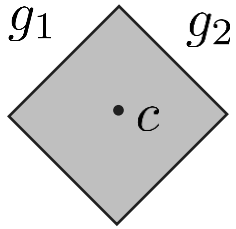

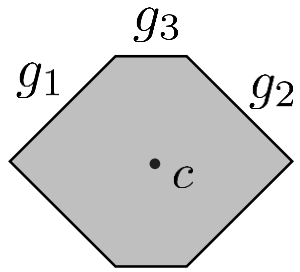

$\mathcal{Z} = \langle c, g_1 \rangle$ $\mathcal{Z} = \langle c, \begin{bmatrix} g_1 & g_2 \end{bmatrix} \rangle$ $\mathcal{Z} = \langle c, \begin{bmatrix} g_1 & g_2 & g_3 \end{bmatrix} \rangle$

Fedorov, Elemente der Gestaltenlehre. Zeitschrift für Krystallographie und Mineralogie, 1893

Zonotope Properties

- Zonotopes $\mathcal{Z} = \langle c, G \rangle$ are closed under linear maps L:

$$L\mathcal{Z} = \{Lz | z \in \mathcal{Z}\} = \langle Lc, LG \rangle$$

- Closed under Minkowski sum:

$$\mathcal{Z}_1 \oplus \mathcal{Z}_2 = \{z_1 + z_2 | z_1 \in \mathcal{Z}_1, z_2 \in \mathcal{Z}_2\} = \left\langle c_1 + c_2, [G_1, G_2] \right\rangle$$

- Cartesian product:

$$\mathcal{Z}_1 \times \mathcal{Z}_2 = \left\{ \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \middle| z_1 \in \mathcal{Z}_1, z_2 \in \mathcal{Z}_2 \right\} = \left\langle \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}, \begin{bmatrix} G_1 & 0 \\ 0 & G_2 \end{bmatrix} \right\rangle$$

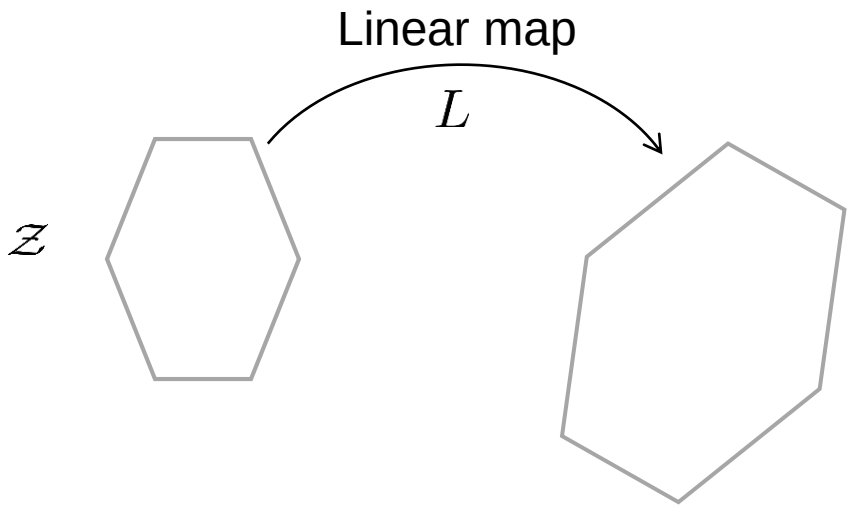

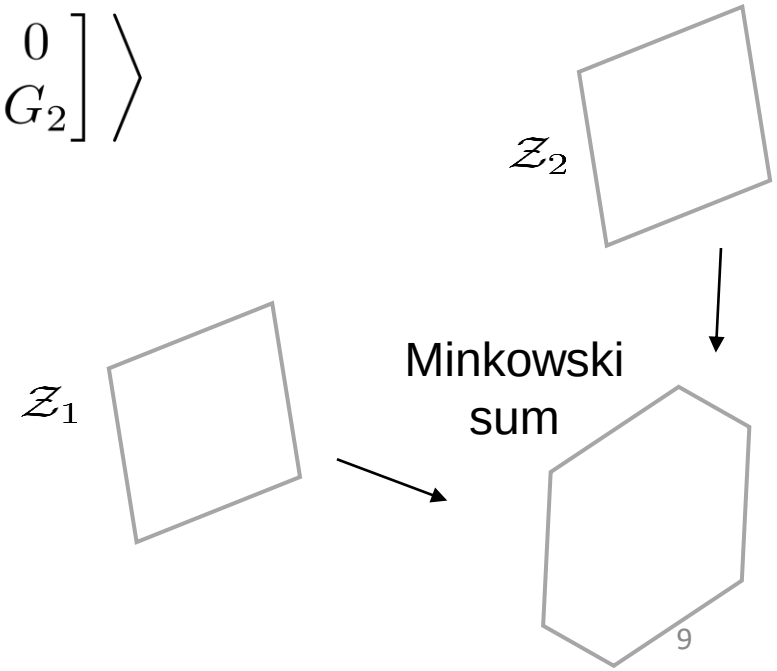

[*Kühn, 1998*]

Model-based Reachability using Zonotopes

Consider a system

$$x(k+1) = \bar{A}x(k) + \bar{B}u(k) + w(k)$$

with initial state set, control set, and noise set all are zonotopes.

Then, the reachable set can be computed as

$$\mathcal{R}_{k+1} = \bar{A}\mathcal{R}_k \oplus \bar{B}\mathcal{U}_k \oplus \mathcal{Z}_w = \begin{bmatrix} \bar{A} & \bar{B} \end{bmatrix} (\mathcal{R}_k \times \mathcal{U}_k) \oplus \mathcal{Z}_w$$

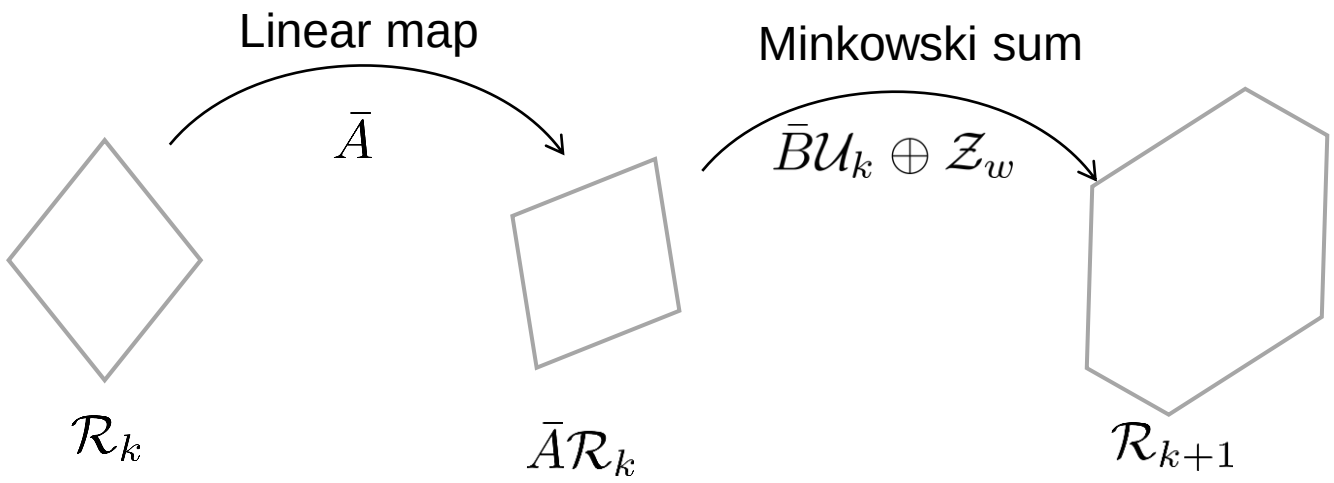

[*Girard, HSCC, 2005*]

Related Work

Reachability computation is a classical topic in systems and control theory:

- E.g., [*Pecsvaradi & Narendra, 1971*]

Many recent techniques for formal verification:

- Ellipsoidal techniques [*Kurzhanski & Varaiya, 2000*]
- Hamilton-Jacobi approaches [*Mitchell et al., 2001*]
- Barrier certificates [*Prajna & Jadbabaie, 2004*]
- Simulation-based approaches [*Girard & Pappas, 2006*]
- Zonotopes [*Girard, 2005*] *[Althoff, 2010]*
- Monte Carlo method [*Devonport*,*2020*]
- Intervals [*Djeumou*,*2021*]

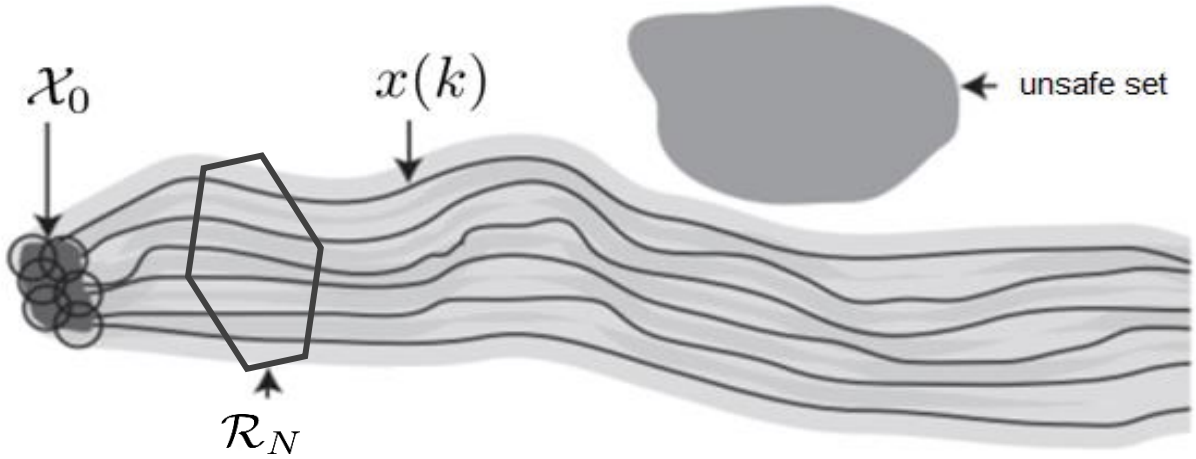

How to compute data-driven reachable sets?

Proposed Approach

- Given the noisy-data, there is no single model that can be trusted and fits the data
- We compute a set of models that is consistent with the data instead of depending on a single model
- We guarantee that the true model is inside the set of models

Linear System: $x(k+1) = \bar{A}x(k) + \bar{B}u(k) + w(k)$

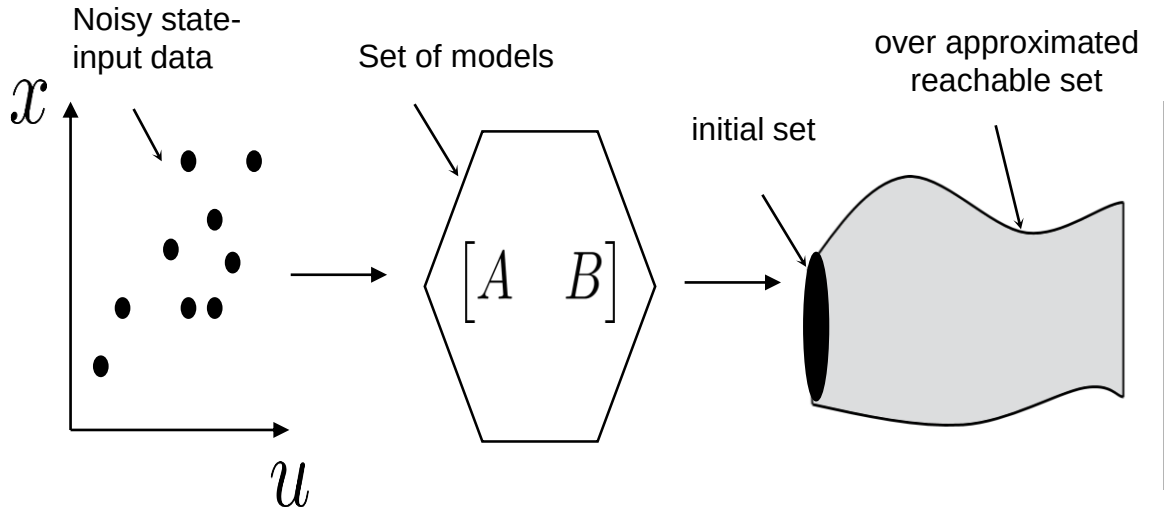

Set of Models Consistent with Data

$$x(k+1) = \bar{A}x(k) + \bar{B}u(k) + w(k)$$

The set of all models consistent with the input-state data is denoted

$$\mathcal{N}_\Sigma = \{ \begin{bmatrix} A & B \end{bmatrix} \mid X_+ = AX + BU + W, W \in \mathcal{M}_w \}$$

$$U = \begin{bmatrix} u(0) & u(1) & \dots & u(T-1) \end{bmatrix}$$
$$X = \begin{bmatrix} x(0) & x(1) & \dots & x(T-1) \end{bmatrix}$$
$$X_+ = \begin{bmatrix} x(1) & x(2) & \dots & x(T) \end{bmatrix}$$
$$W = \begin{bmatrix} w(0) & w(1) & \dots & w(T-1) \end{bmatrix}$$

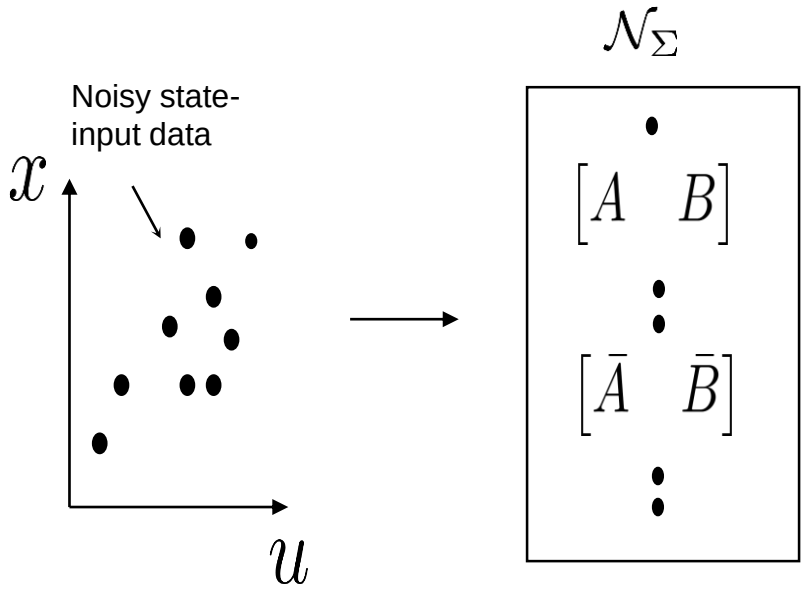

True model inside $\mathcal{N}_\Sigma$

Computation of Over-approximate Set of Models

We assume bounded noise. From system model, it follows that

$$X_+ = \begin{bmatrix} A & B \end{bmatrix} \begin{bmatrix} X \\ U \end{bmatrix} + W,$$

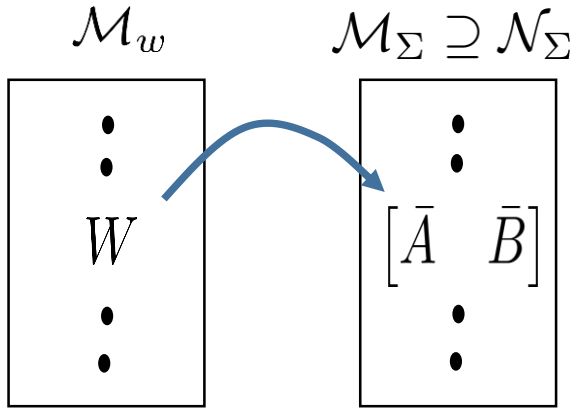

which can be rearranged as

$$\begin{bmatrix} A & B \end{bmatrix} = \left(X_+ - C_w - \sum_{i=1}^{\gamma_w} \hat{\beta}_{i,w} G_{i,w} \right) \begin{bmatrix} X \\ U \end{bmatrix}^\dagger$$

for some choice of $\hat{\beta}_{i,w}$. Hence, by considering all possible values of $\hat{\beta}_{i,w}$ where $-1 \leq \hat{\beta}_{i,w} \leq 1$, $i = 1, \ldots, \gamma_w$, the matrix zonotope

$$\mathcal{M}_\Sigma = (X_+ \ominus \mathcal{M}_w) \begin{bmatrix} X \\ U \end{bmatrix}^\dagger$$

contains all models in $\mathcal{N}_\Sigma$

$\mathcal{M}_\Sigma$ is an over-approximation of the set of models $\mathcal{M}_\Sigma \supseteq \mathcal{N}_\Sigma$

Data-Driven Reachable set

- The true system model is within the set of models $\begin{bmatrix}\bar{A} & \bar{B}\end{bmatrix} \in \mathcal{N}_\Sigma \subseteq \mathcal{M}_\Sigma$
- Compare with model-based

$$\mathcal{R}_{k+1} = \begin{bmatrix}\bar{A} & \bar{B}\end{bmatrix} (\mathcal{R}_k \times \mathcal{U}_k) \oplus \mathcal{Z}_w$$

Algorithm 1 LTI-Reachability

Input: input-state trajectories $D = (U_-, X)$, initial set $\mathcal{X}_0$, process noise zonotope $\mathcal{Z}_w$ and matrix zonotope $\mathcal{M}_w$, input zonotope $\mathcal{U}_k$, $\forall k = 0, \ldots, N-1$

Output: reachable sets $\hat{\mathcal{R}}_k, \forall k = 1, \ldots, N$

1: $\hat{\mathcal{R}}_0 = \mathcal{X}_0$

2: $\mathcal{M}_\Sigma = (X_+ \ominus \mathcal{M}_w) \begin{bmatrix}X \\ U\end{bmatrix}^\dagger$

3: **for** $k = 0 : N-1$ **do**

4: $\quad \hat{\mathcal{R}}_{k+1} = \mathcal{M}_\Sigma(\hat{\mathcal{R}}_k \times \mathcal{U}_k) \oplus \mathcal{Z}_w$

5: **end for**

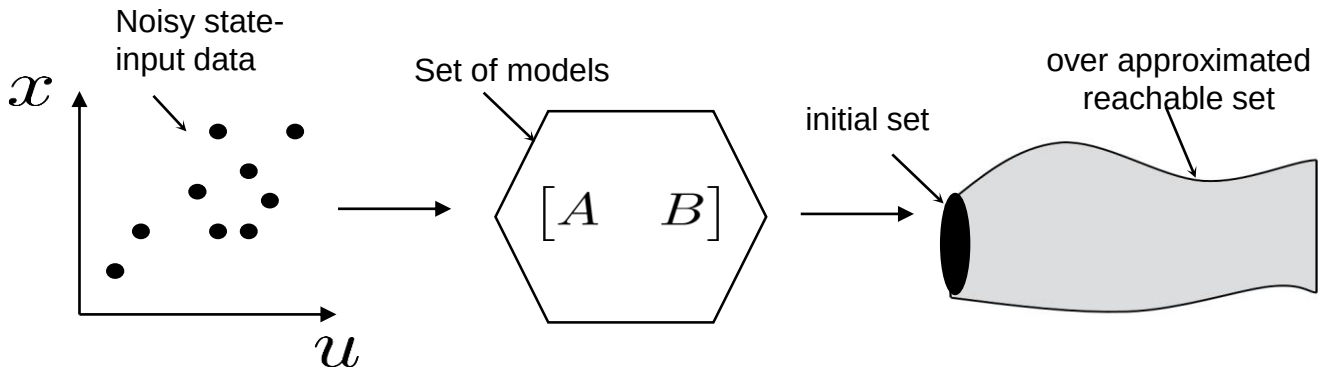

Example

- Linear system

$$x(k+1) = \underbrace{\begin{bmatrix} 0.9323 & -0.1890 & 0 & 0 & 0 \\ 0.1890 & 0.9323 & 0 & 0 & 0 \\ 0 & 0 & 0.8596 & 0.0430 & 0 \\ 0 & 0 & -0.0430 & 0.8596 & 0 \\ 0 & 0 & 0 & 0 & 0.9048 \end{bmatrix}}_{\bar{A}} x(k) + \underbrace{\begin{bmatrix} 0.0436 \\ 0.0533 \\ 0.0475 \\ 0.0453 \\ 0.0476 \end{bmatrix}}_{\bar{B}} u(k) + w(k)$$

- The data-driven reachable set over approximates the model based reachable set $\hat{\mathcal{R}}_k \supseteq \mathcal{R}_k$

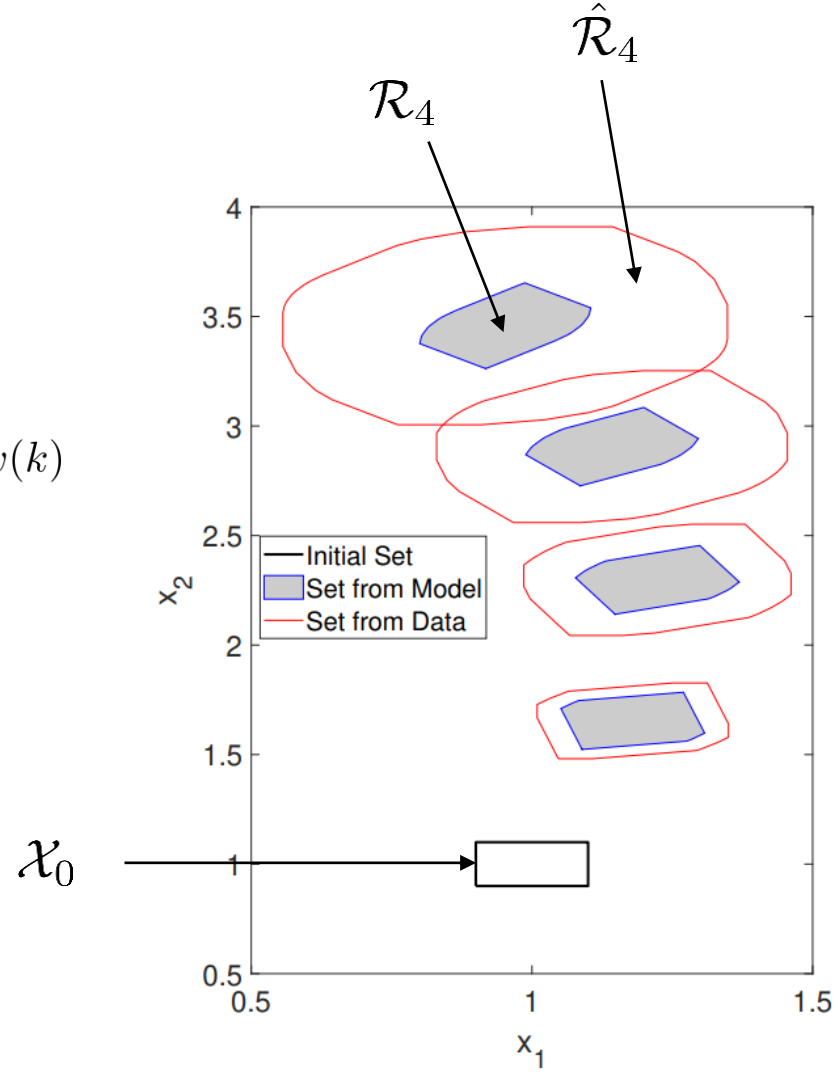

Can we have a tighter over approximate set?

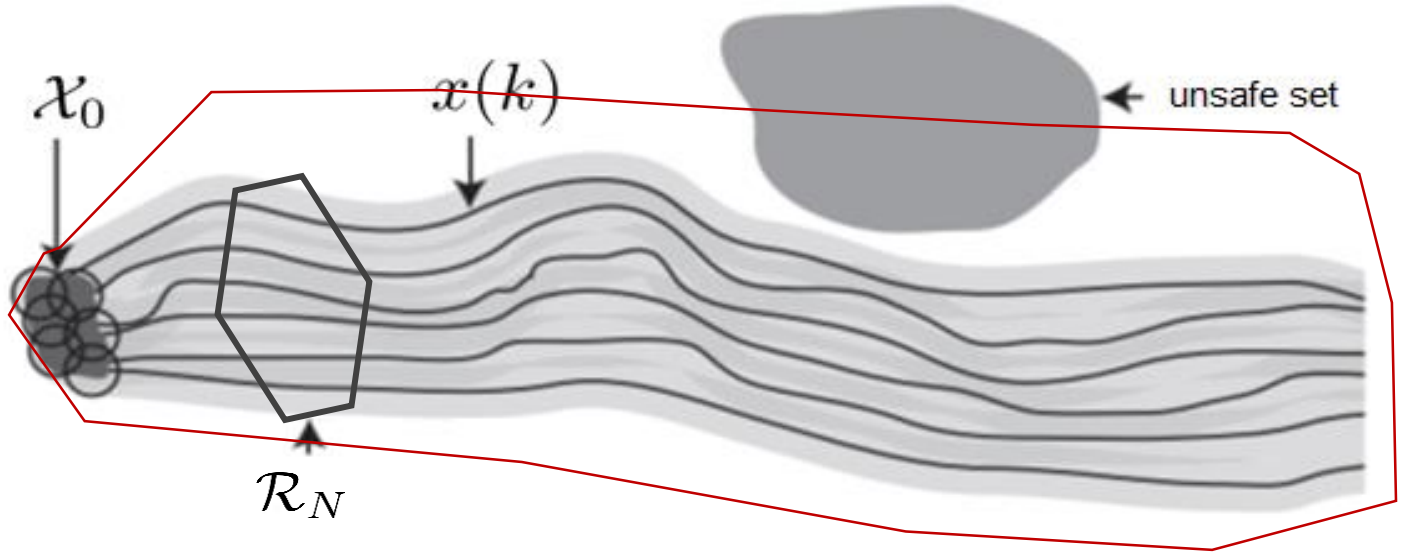

Exact Noise Description Consistent with Model

- We aim to find a tighter over approximate reachable set
- Our previous solution $\mathcal{M}_\Sigma$ is based on finding $[A \quad B]$ for each $W \in \mathcal{M}_w$
- There might not exist a solution $[A \quad B]$ for all $W \in \mathcal{M}_w$

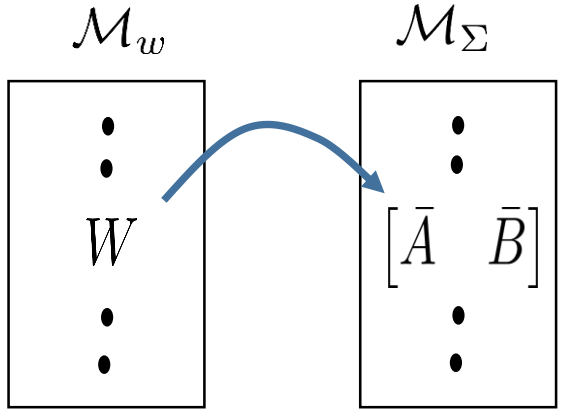

$$[A \quad B] \begin{bmatrix} X \\ U \end{bmatrix} = X_+ - W$$

- An exact description for all systems consistent with the data and the noise bound would therefore be the set [A. Koch, 2020]

$$\mathcal{N}_\Sigma = (X_+ \ominus \mathcal{N}_w) \begin{bmatrix} X \\ U \end{bmatrix}^\dagger \quad \text{with} \quad \mathcal{N}_w = \{W \in \mathcal{M}_w \mid (X_+ - W) \begin{bmatrix} X \\ U \end{bmatrix}^\perp = 0\}$$

Constrained Matrix Zonotope

- We propose constrained matrix zonotope as a new set representation
- A constrained matrix zonotope $\mathcal{C}$ is a set

$$\mathcal{C} = \left\{ X \in \mathbb{R}^{n \times p} \;\middle|\; X = C + \sum_{i=1}^{\gamma} \beta_i G_i, \sum_{i=1}^{\gamma} \beta_i A_i = B \;, -1 \leq \beta_i \leq 1 \right\}.$$

where, C is the center matrix

$\tilde{G} = [G_1, ..., G_\gamma]$ are the generator matrices, and

$\tilde{A} = [A_1, ..., A_\gamma]$ and B are constraining the factors $[\beta_1, ..., \beta_\gamma]$

How to incorporate system model side information?

Side Information on System Model

- Consider prior side information about the unknown models:
  - prior bounds on entries in the system matrices
  - decoupling in the dynamics
  - partial dynamics knowledge
- We propose a general framework $\left|\bar{Q}\begin{bmatrix}A & B\end{bmatrix} - \bar{Y}\right| \leq \bar{R}$

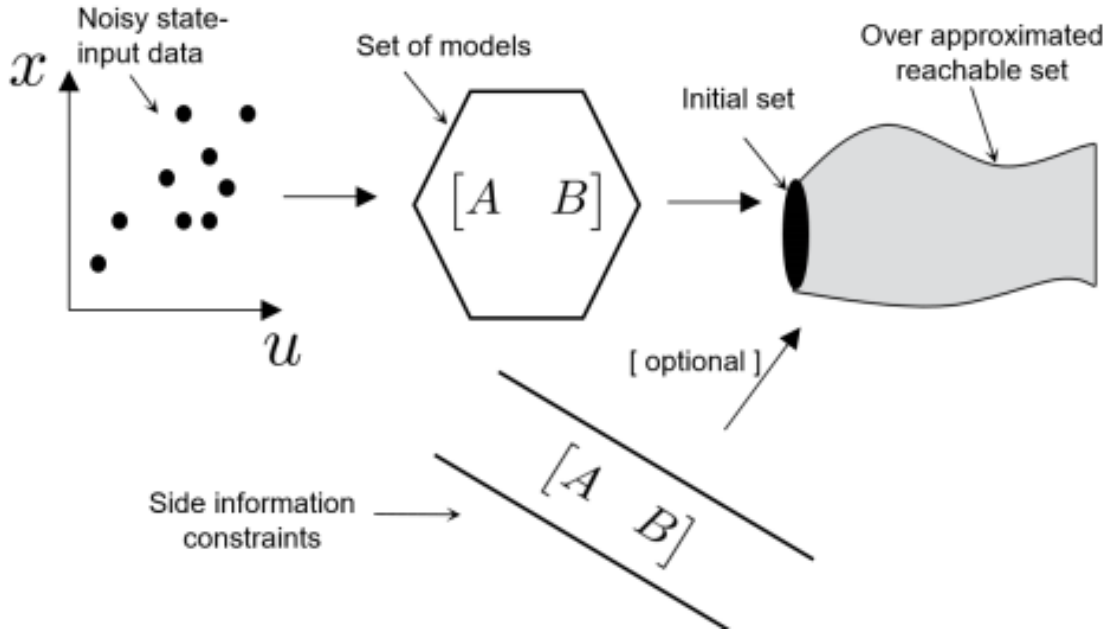

Example

$$x(k+1) = \underbrace{\begin{bmatrix} 0.9323 & -0.1890 & 0 & 0 & 0 \\ 0.1890 & 0.9323 & 0 & 0 & 0 \\ 0 & 0 & 0.8596 & 0.0430 & 0 \\ 0 & 0 & -0.0430 & 0.8596 & 0 \\ 0 & 0 & 0 & 0 & 0.9048 \end{bmatrix}}_{\bar{A}} x(k) + \underbrace{\begin{bmatrix} 0.0436 \\ 0.0533 \\ 0.0475 \\ 0.0453 \\ 0.0476 \end{bmatrix}}_{\bar{B}} u(k) + w(k)$$

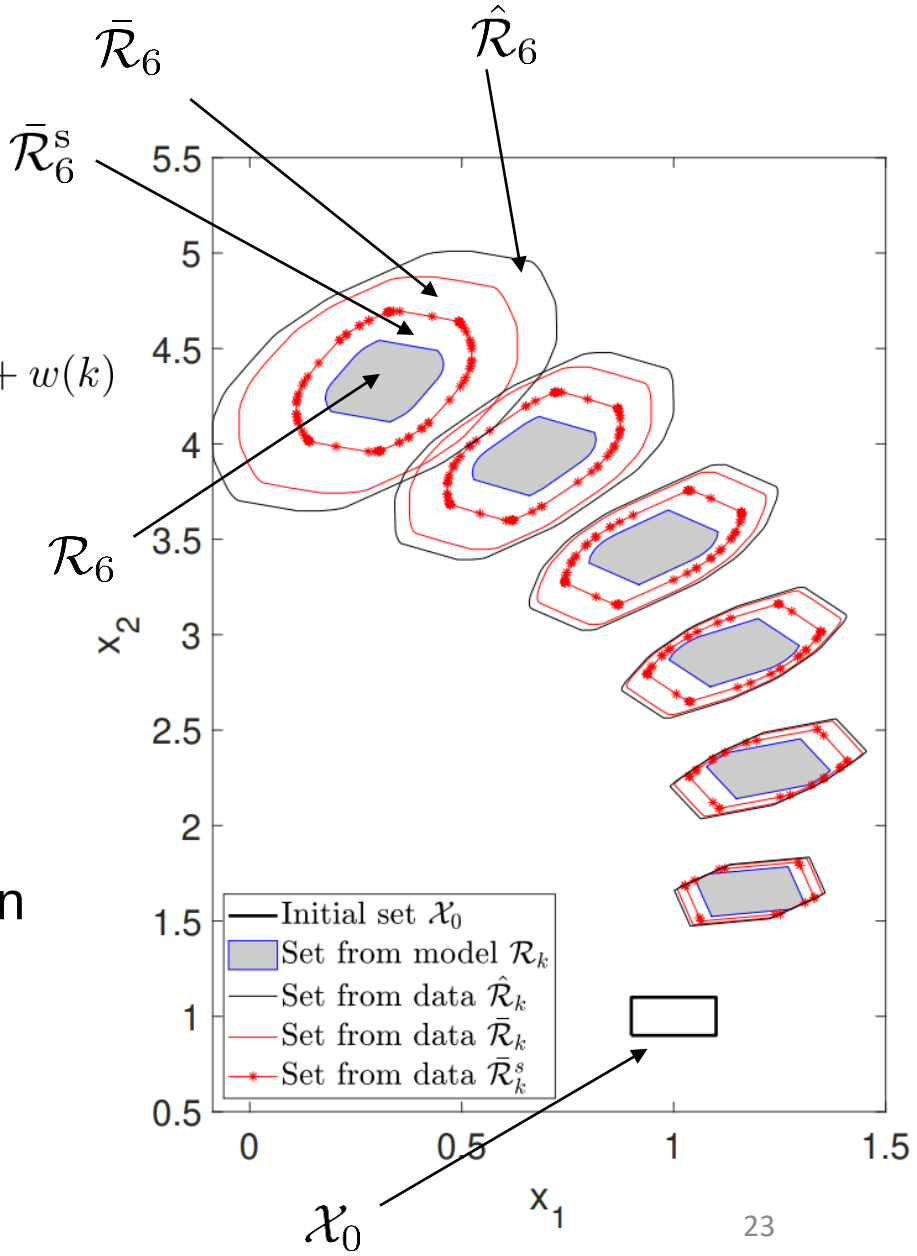

- $\hat{\mathcal{R}}_k$: Reachable sets computed based on matrix zonotopes
- $\bar{\mathcal{R}}_k$: Reachable sets computed using exact noise description
- $\bar{\mathcal{R}}_k^{\mathrm{s}}$: Reachable sets using state decoupling as side information
- We note that $\hat{\mathcal{R}}_k \supseteq \bar{\mathcal{R}}_k \supseteq \bar{\mathcal{R}}_k^{\mathrm{s}} \supseteq \mathcal{R}$

SVEA Reachable Sets

- The inputs to the vehicle are the steering angle and the velocity and the output is the position of the vehicle. The model of the vehicle is nonlinear in general

Small-Vehicles-for-Autonomy (SVEA) equipped with NVIDIA Jetson TX2 embedded computer and Qualisys motion capture system

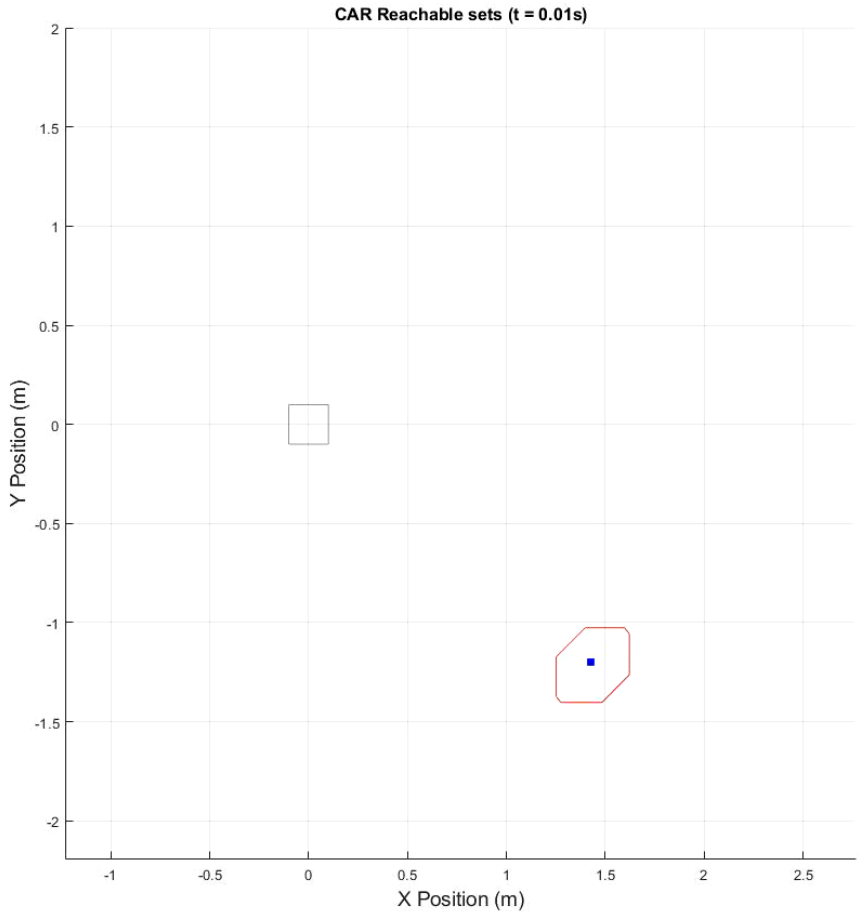

How to incorporate signal temporal logic side information?

STL Side Information

- In certain scenarios, the reachable sets can be conservative due to the amount of noise in the data
- We often have side information about the system or environment which can be generally described using **signal temporal logic (STL)**
- STL is a formal language for describing a broad range of real-valued, temporal properties in cyber-physical systems

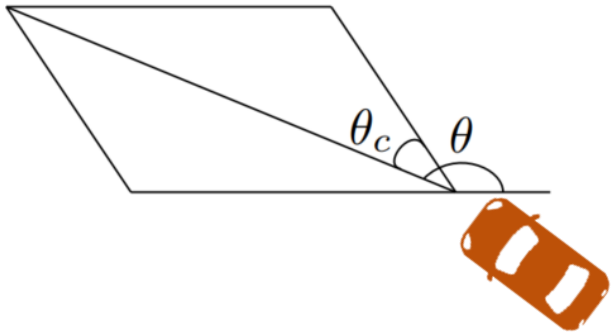

Incorporating STL Side Information

- We have data-driven zonotope $\hat{\mathcal{R}}_k$ and STL side information
- Construct a predicate function $\mathfrak{h}_k(x)$ from ϕ_k such that if $\mathfrak{h}_k(x) \geq 0$ then $x \models \phi_{i,k}$

[*Fainekos, 2009*]

- Represent $\mathfrak{h}_k(x) \geq 0$ as strip or nonlinear strip

STL

$$\phi_k = G_{[0,25]}(\mathcal{P}) \ \wedge \ G_{[25,40]}(\mathcal{S}) \xrightarrow{\text{Conversion}} \mathfrak{h}_k(x) \geq 0 \Longrightarrow$$

Linear and nonlinear Strips

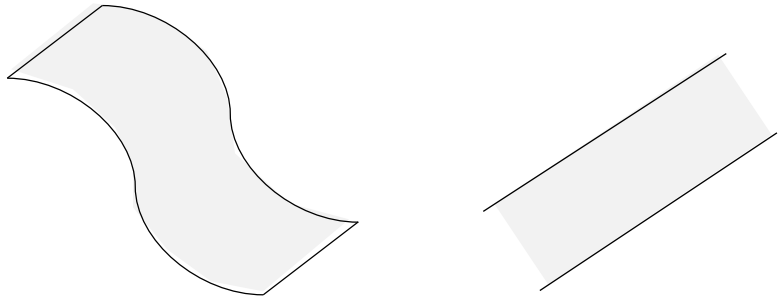

$\mathcal{S}_k^{\mathrm{n}}$ $\mathcal{S}_k$

Intersection with Strips

- We provide intersection with nonlinear strips by linearization and considering a zonotope that over approximates the linearization error

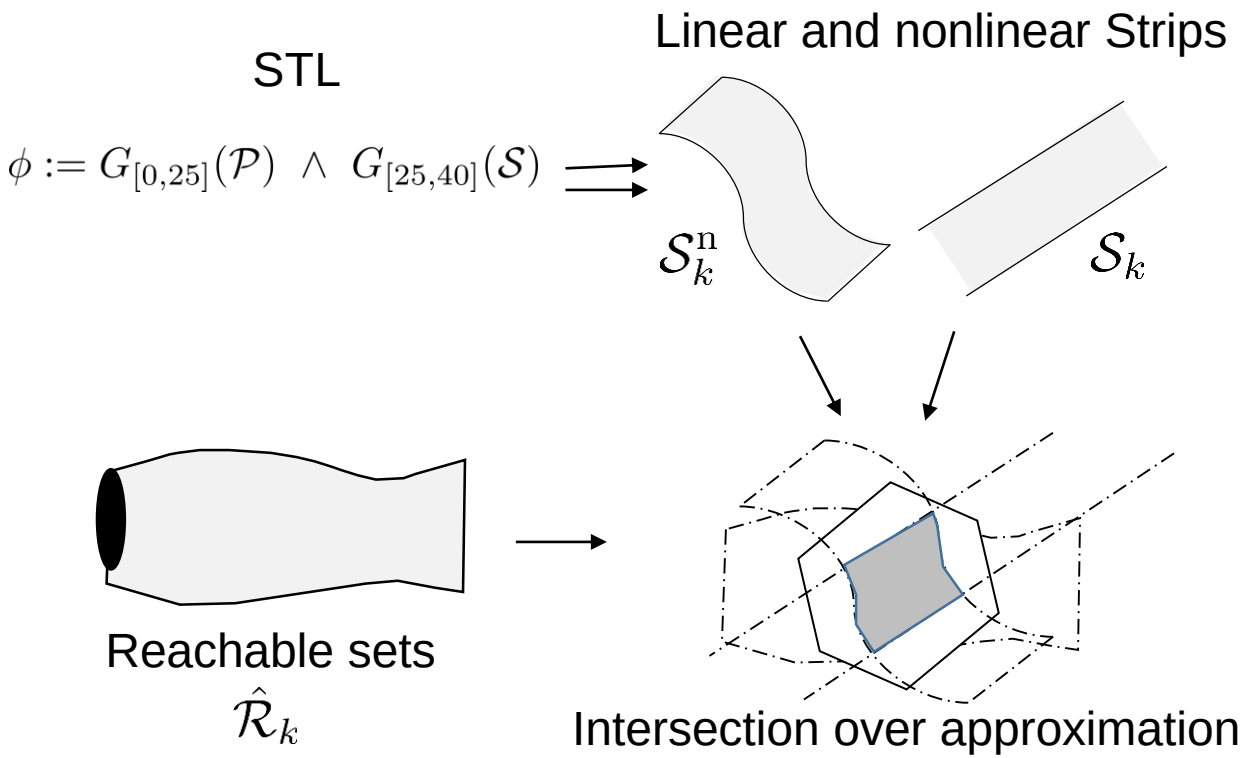

Algorithm 1 Reachability analysis under STL side information using zonotopes

Input: data-driven zonotope $\hat{\mathcal{Z}}_k = \langle \hat{c}_k, \hat{G}_k \rangle$, STL side information $\phi_{i,k}$, $\forall i = 1, \ldots, n_{\phi,k}$
Output: STL zonotope $\bar{\mathcal{Z}}_k = \langle \bar{c}_k, \bar{G}_k \rangle$

$\bar{\mathcal{Z}}_k = \hat{\mathcal{Z}}_k$
for $i = 1, \ldots, n_{\phi,k}$ **do**
  Construct $\mathfrak{h}_{i,k}(x)$ from $\phi_{i,k}$
  if $\mathfrak{h}_{i,k}(x)$ is linear **then**
    // $\mathfrak{h}_{i,k}(x) = r_{i,k} - |H_{i,k}x - y_{i,k}|$
    $\bar{c}_k = \bar{c}_k + \lambda_{i,k}(y_{i,k} - H_{i,k}\bar{c}_k)$
    $\bar{G}_k = \begin{bmatrix}(I - \lambda_{i,k}H_{i,k})\bar{G}_k & \lambda_{i,k}r_{i,k}\end{bmatrix}$
  else if $\mathfrak{h}_{i,k}(x)$ is nonlinear **then**
    // $\mathfrak{h}_{i,k}(x) = r_{i,k} - |h_{i,k}(x)|$
    $\bar{c}_k = \bar{c}_k - \lambda_{i,k}\left(h_{i,k}(x^*_{i,k}) + \frac{\partial h_{i,k}}{\partial x}|_{x^*_{i,k}}(\bar{c}_k - x^*_{i,k}) + c_{L,i,k}\right)$
    $\bar{G}_k = \begin{bmatrix}(I - \lambda_{i,k}\frac{\partial h_{i,k}}{\partial x}|_{x^*_{i,k}})\bar{G}_k & \lambda_{i,k}r_{i,k} & -\lambda_{i,k}G_{L,i,k}\end{bmatrix}$
  end if
end for

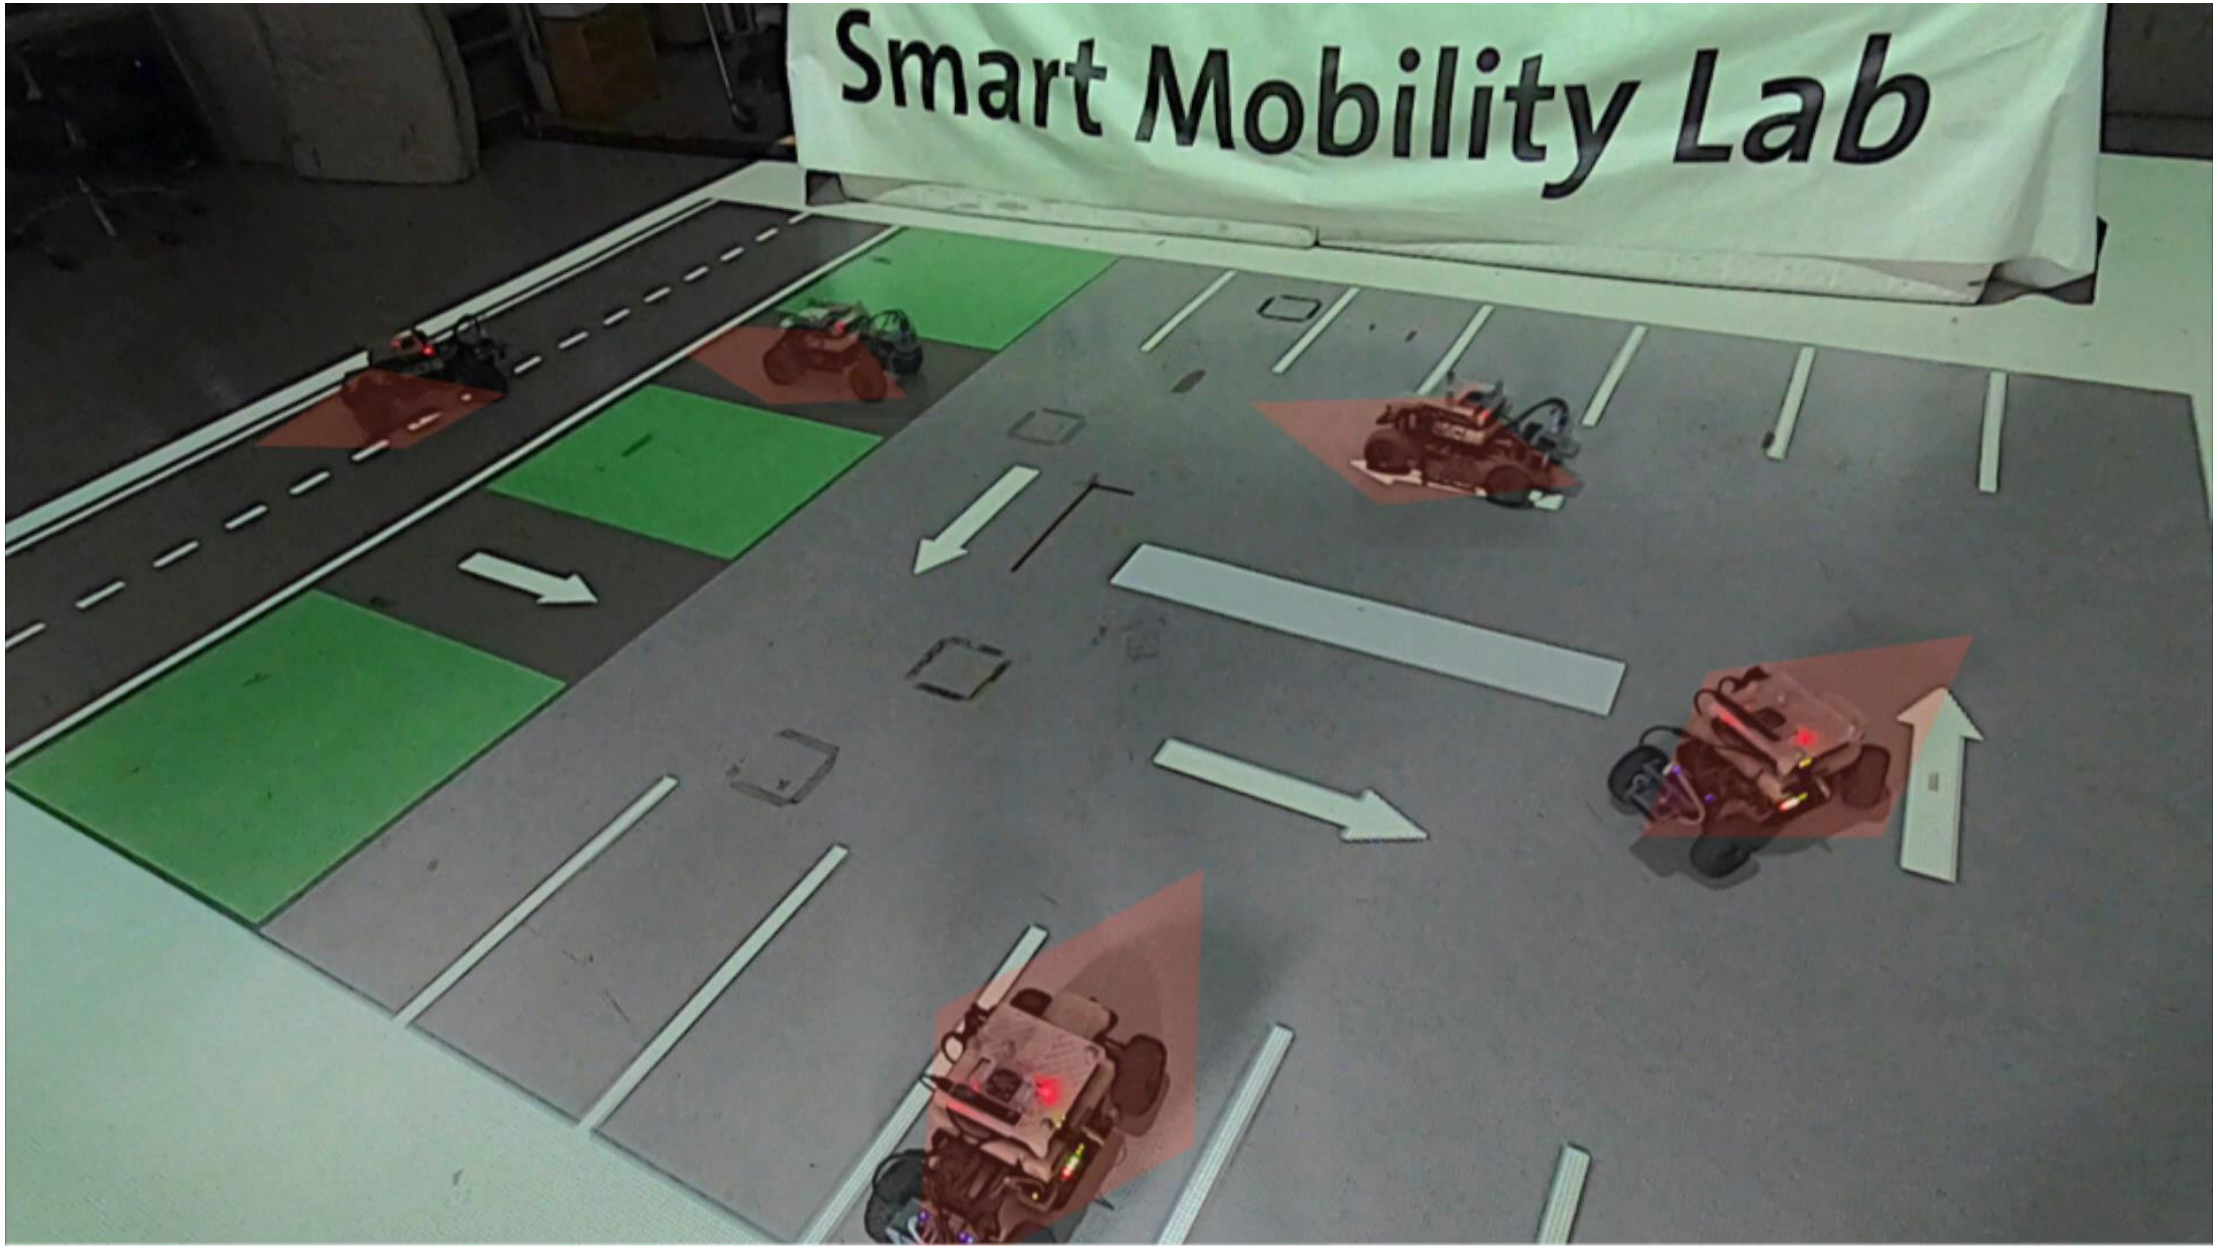
Smart Mobility Lab

How to obtain safety guaranteed reinforcement learning?

Reinforcement Learning (RL)

- Powerful in learning optimal policies
- Requires exploration and exploitation

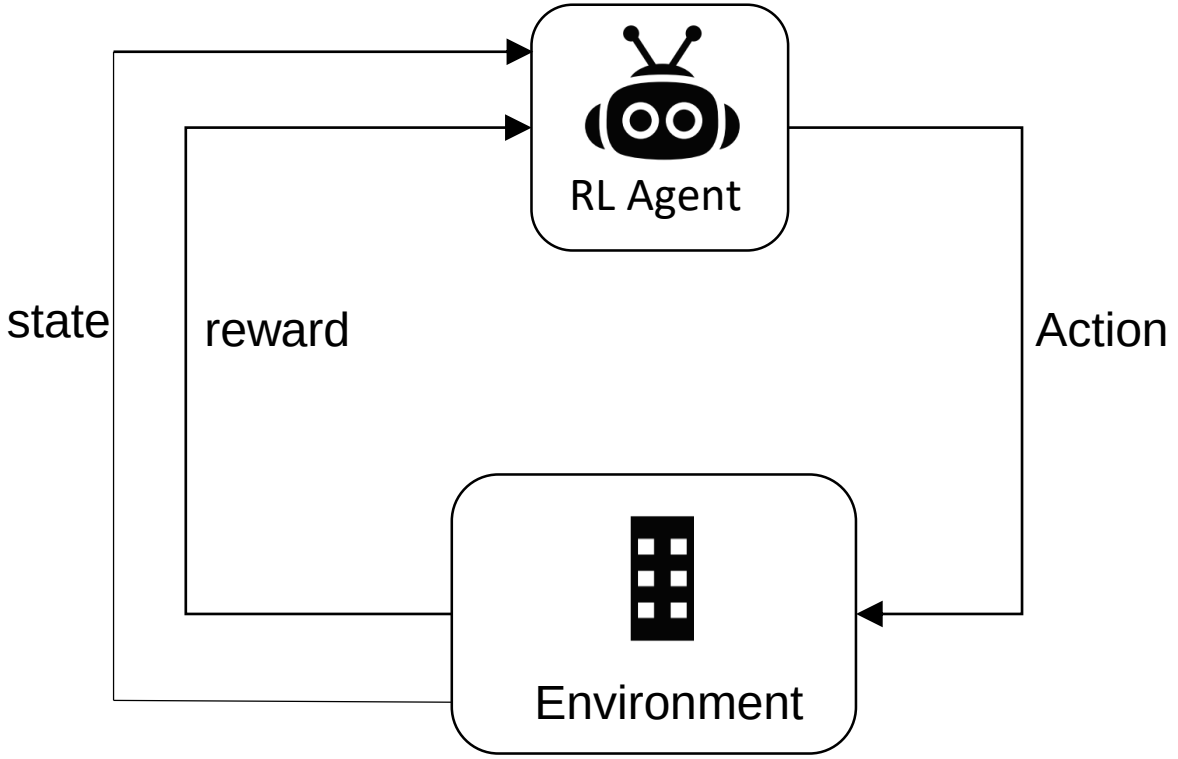

RL and Safety Critical Systems

- Most RL algorithms explore all possible actions, which is not safe for real-world
- RL is rarely applied to safety critical systems especially when models of the robot and environment are unknown

Black-box Reachability-based Safety Layer (BRSL)

- BRSL consists of three main components
  - Data-driven reachability analysis for a black-box robot and environment
  - Rollout trajectory planner
  - Differentiable collision check
- Rollout trajectory planner: the agent dreams a sequence of actions by using an ensemble of neural networks to plan a head

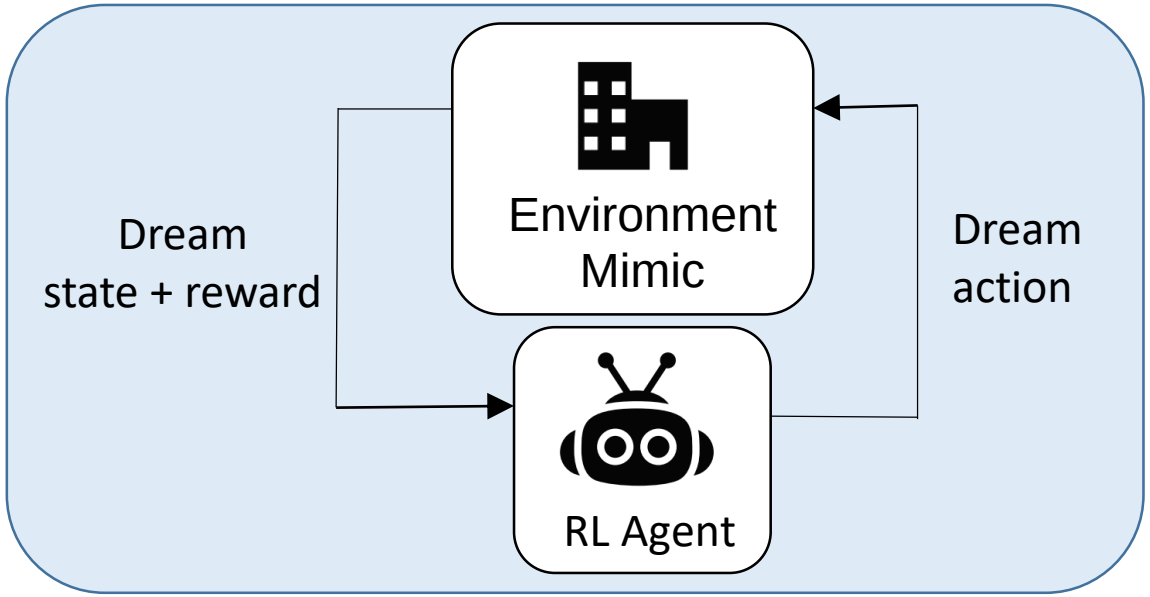

Is it safe action?

- After having a sequence of actions from the rollout trajectory planner, we verify their safety
- Check the intersection between the reachable sets and unsafe sets
- What if we have an unsafe action?

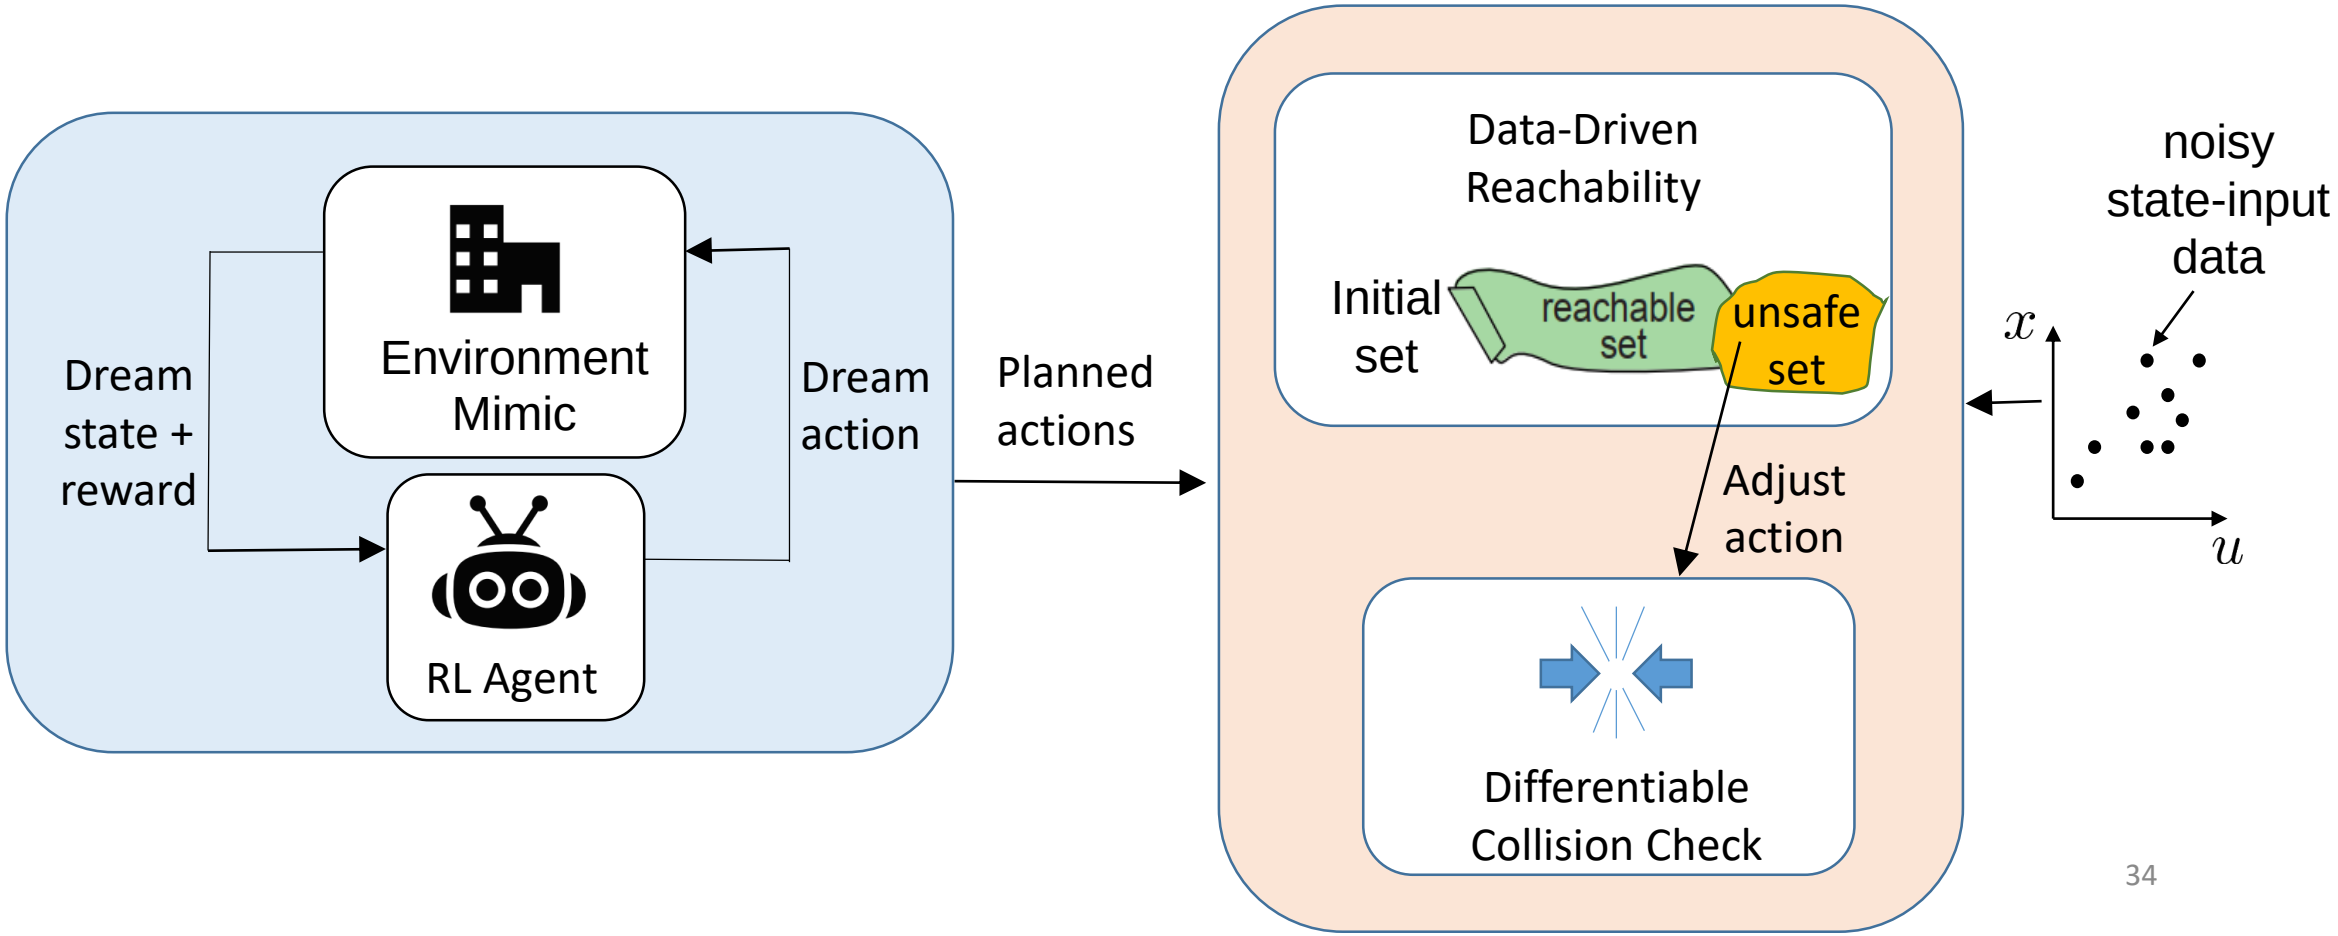

Differentiable Collision Check

- Adjust unsafe action by solving an optimization problem to push zonotopes out of intersection

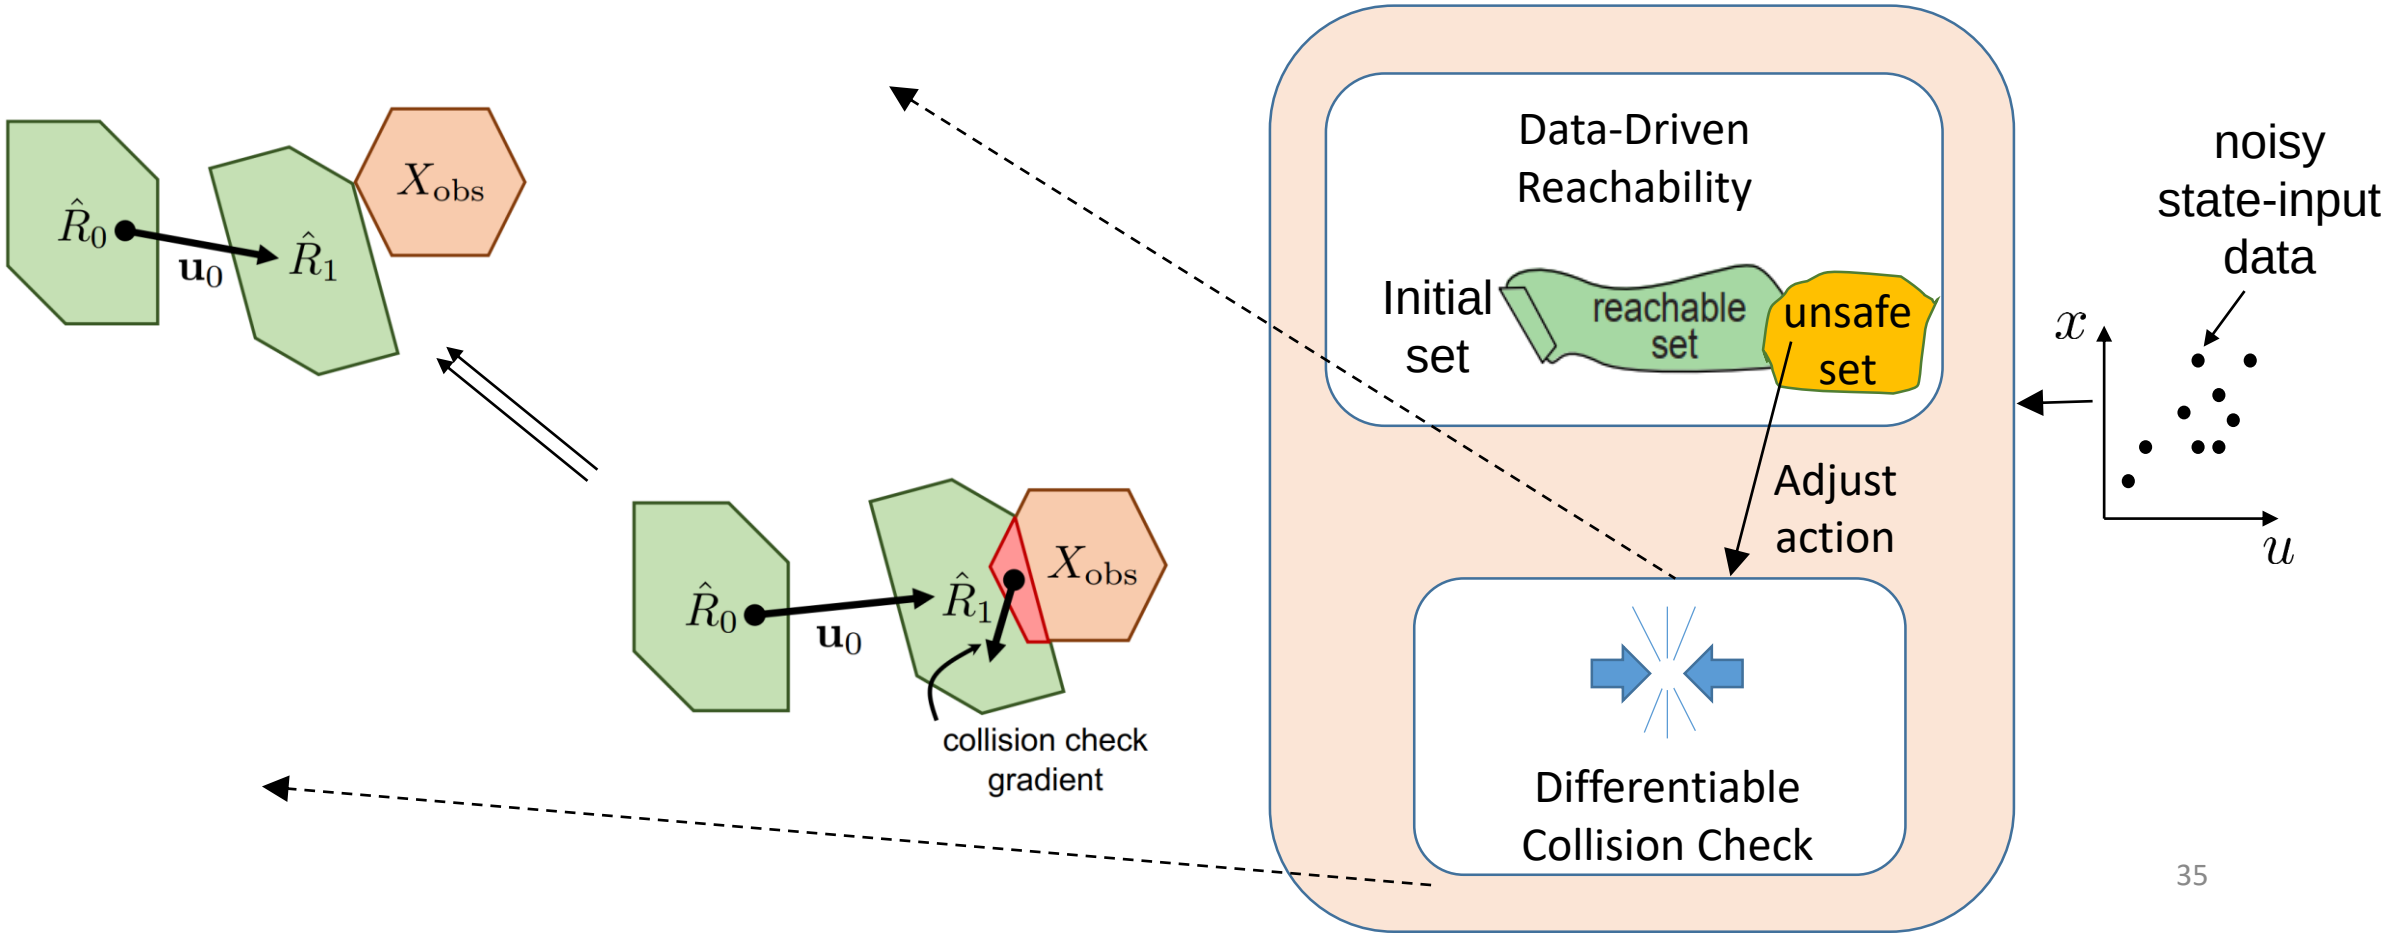

Apply Safe Action

- Apply safe/adjusted first action to the real environment
- Give penalty when adjusting an action

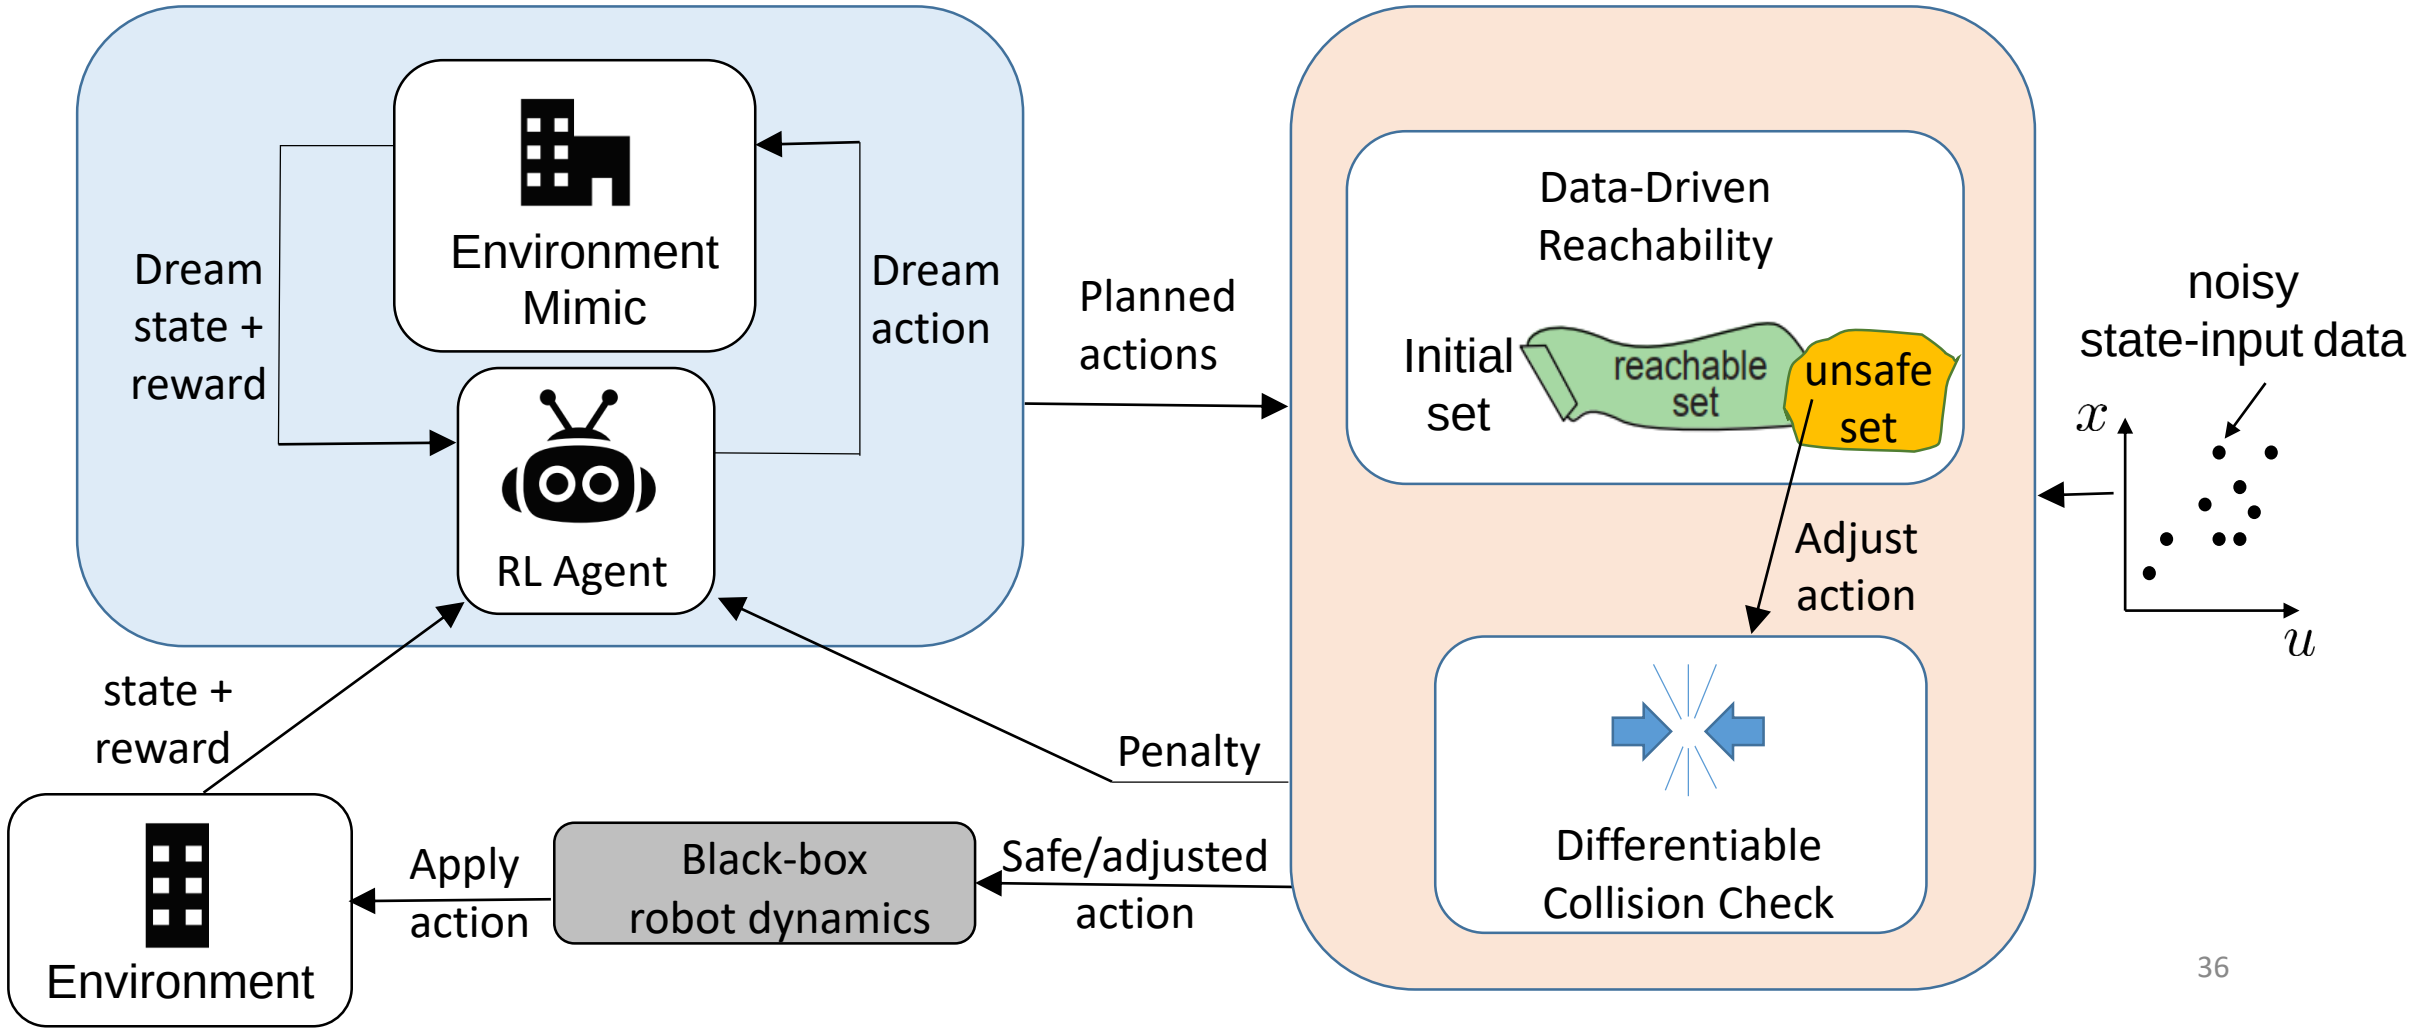

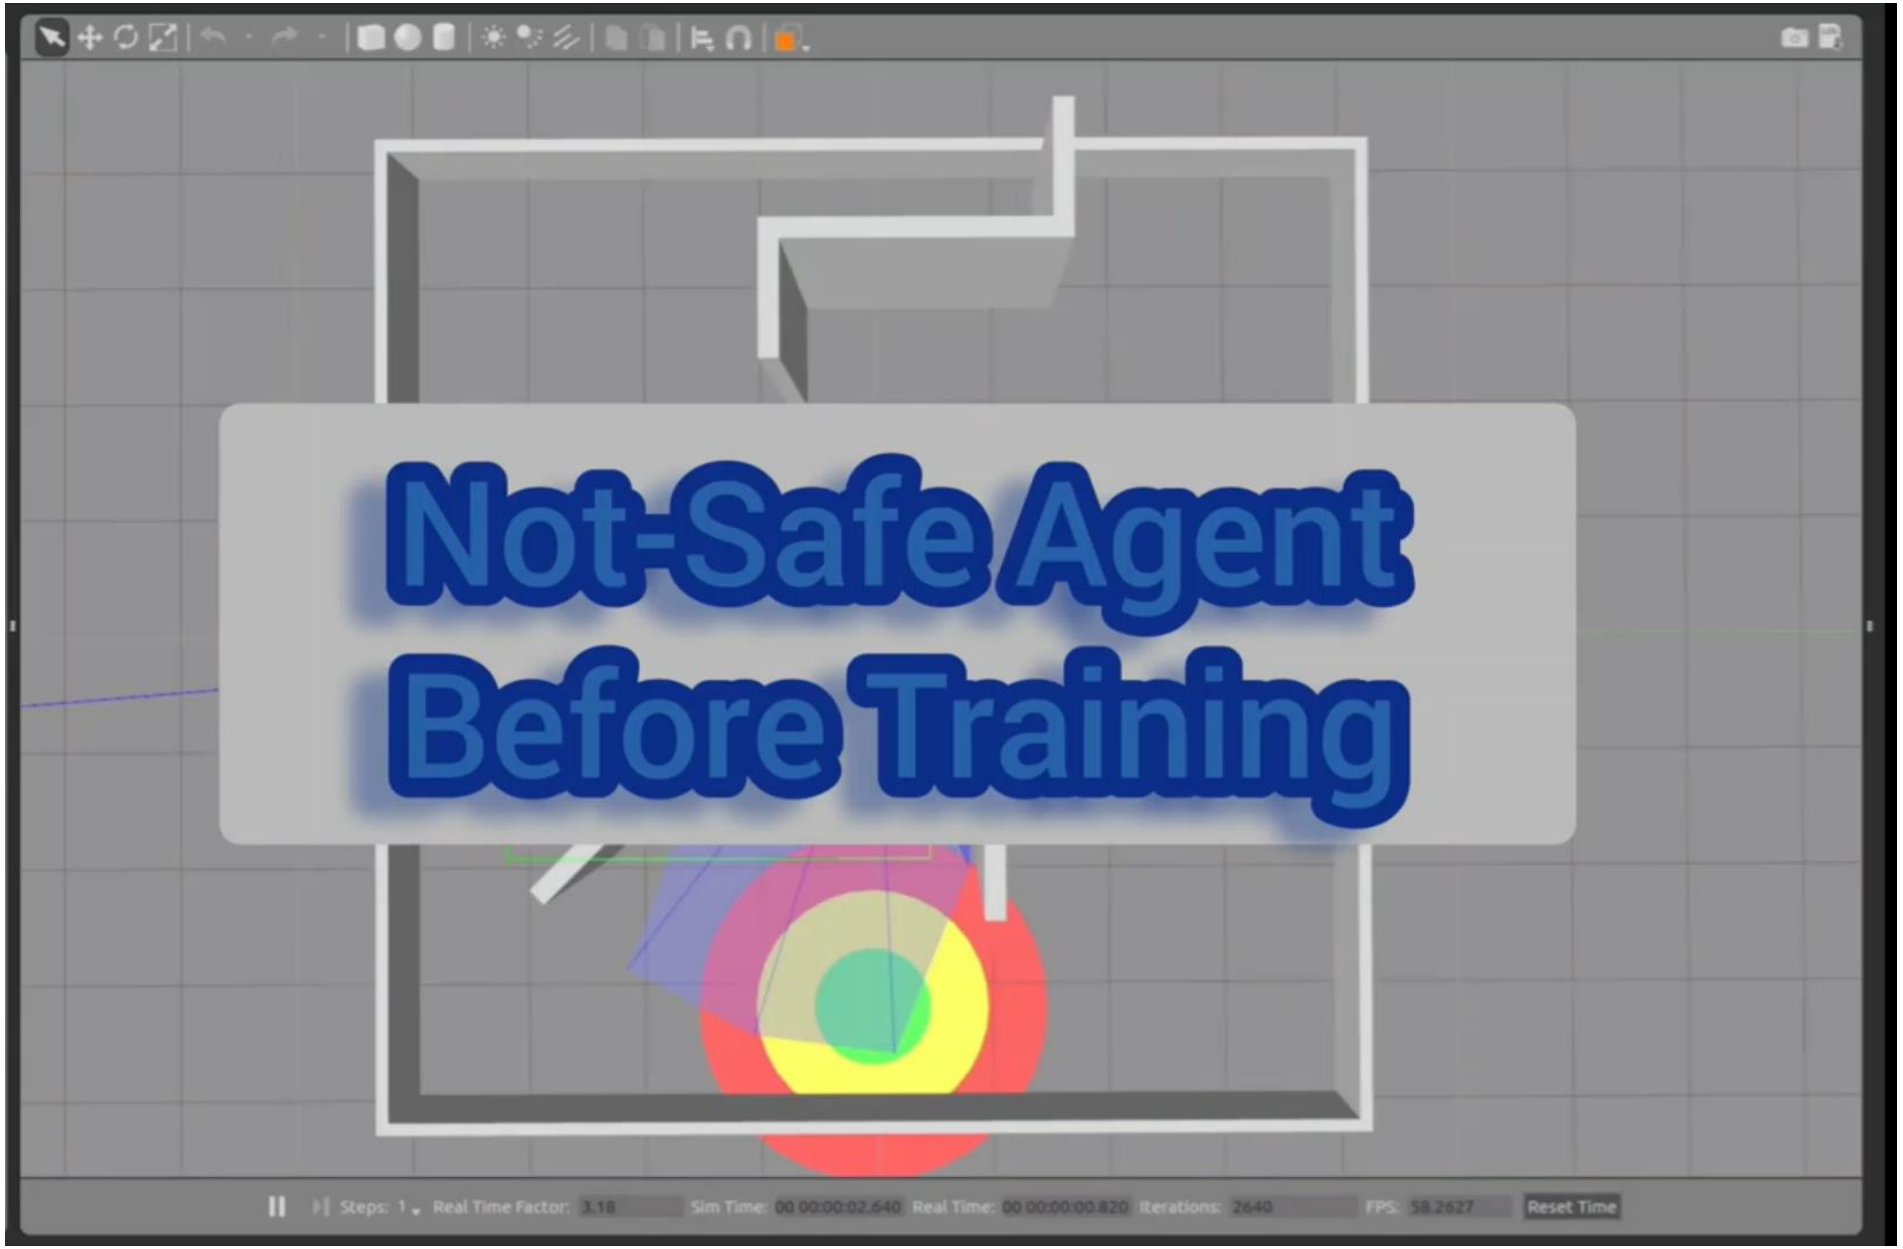
Not-Safe Agent
Before Training
Steps: 1
Real Time Factor: 3.18
Sim Time: 00 00:00:02.640
Real Time: 00 00:00:00.820
Iterations: 2640
FPS: 58.2627
Reset Time

Can We Enhance the learning?

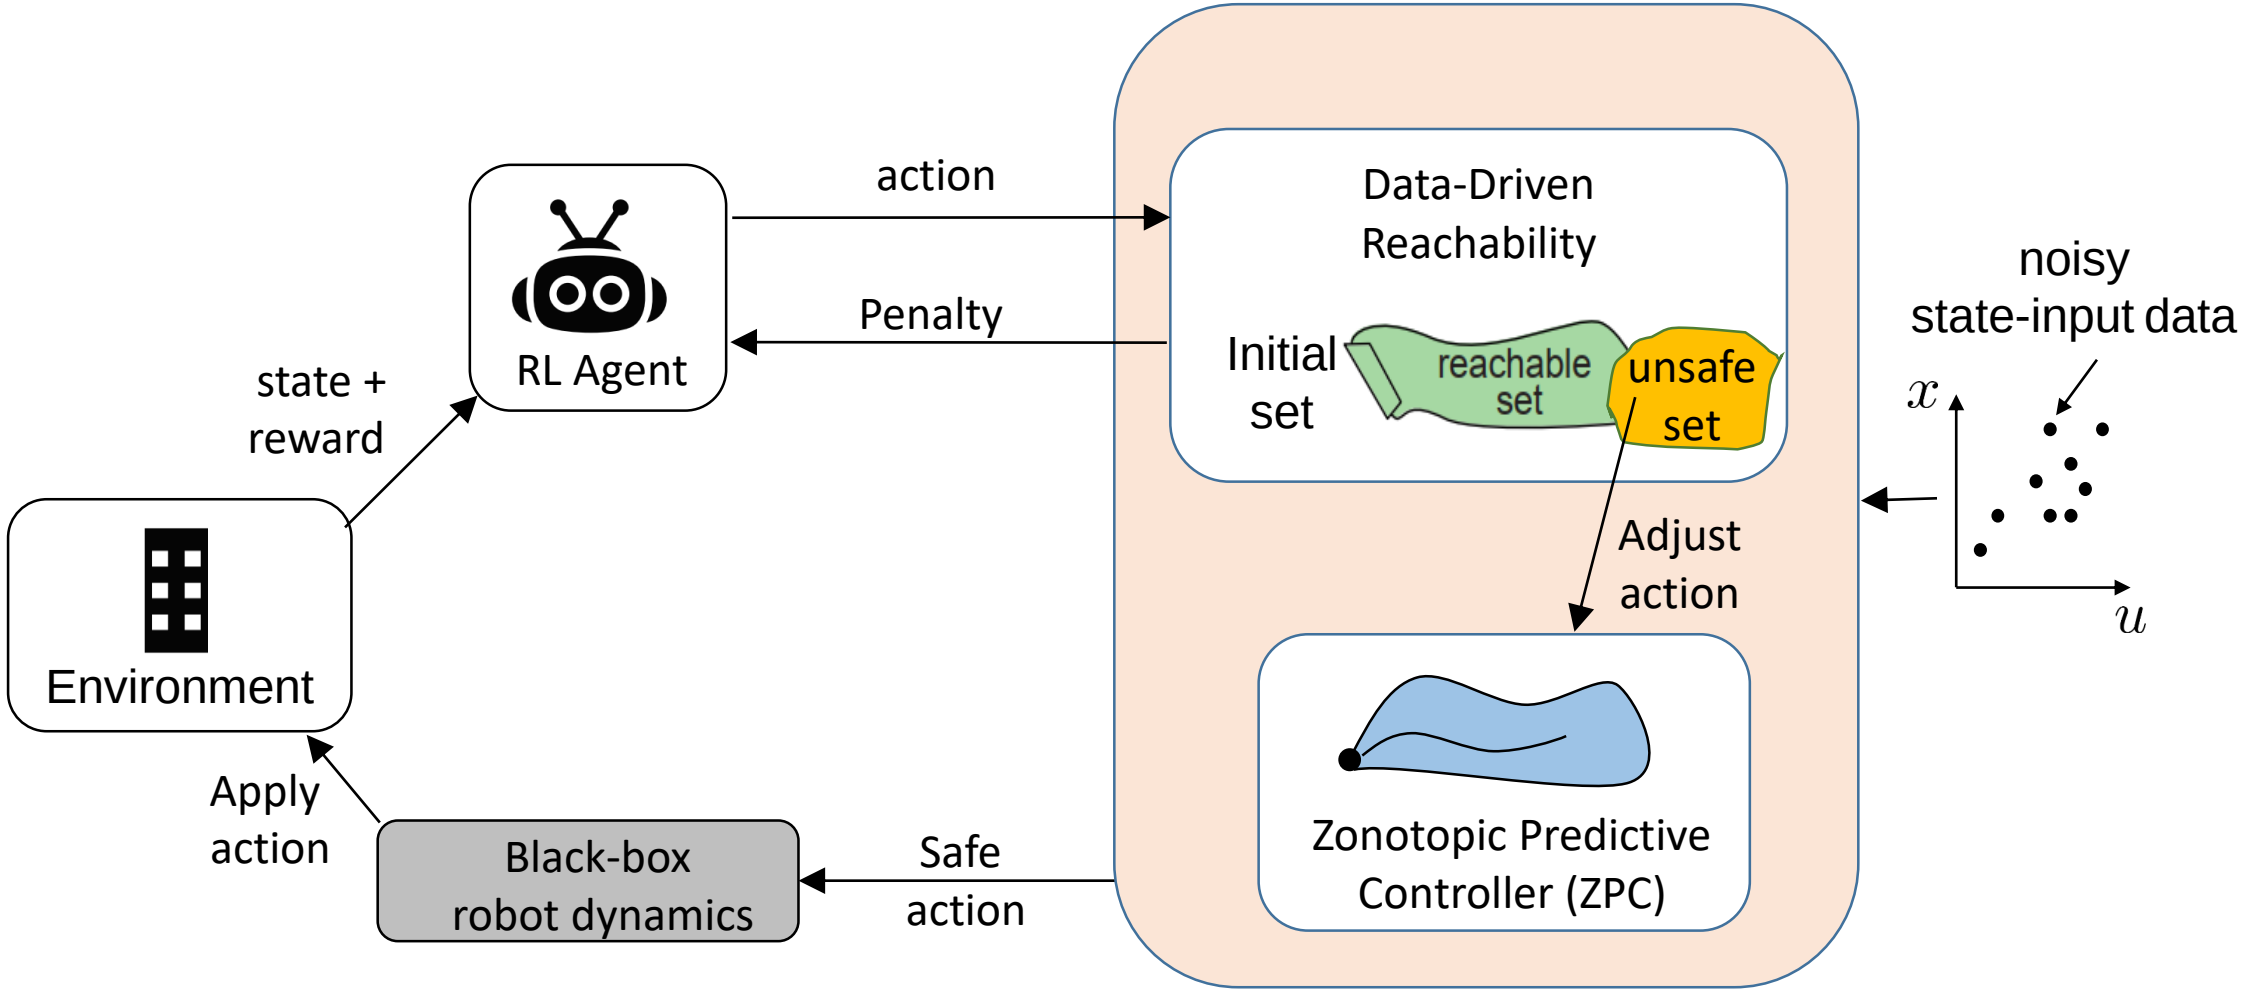

Zonotopic Predictive Controller (ZPC)

- We do not have the robot model
- The true model is within the set of system models $\mathcal{M}_\Sigma$
- Predict ahead using the set of system models $\mathcal{M}_\Sigma$
- Find a controller such that the output stays within the reachable region

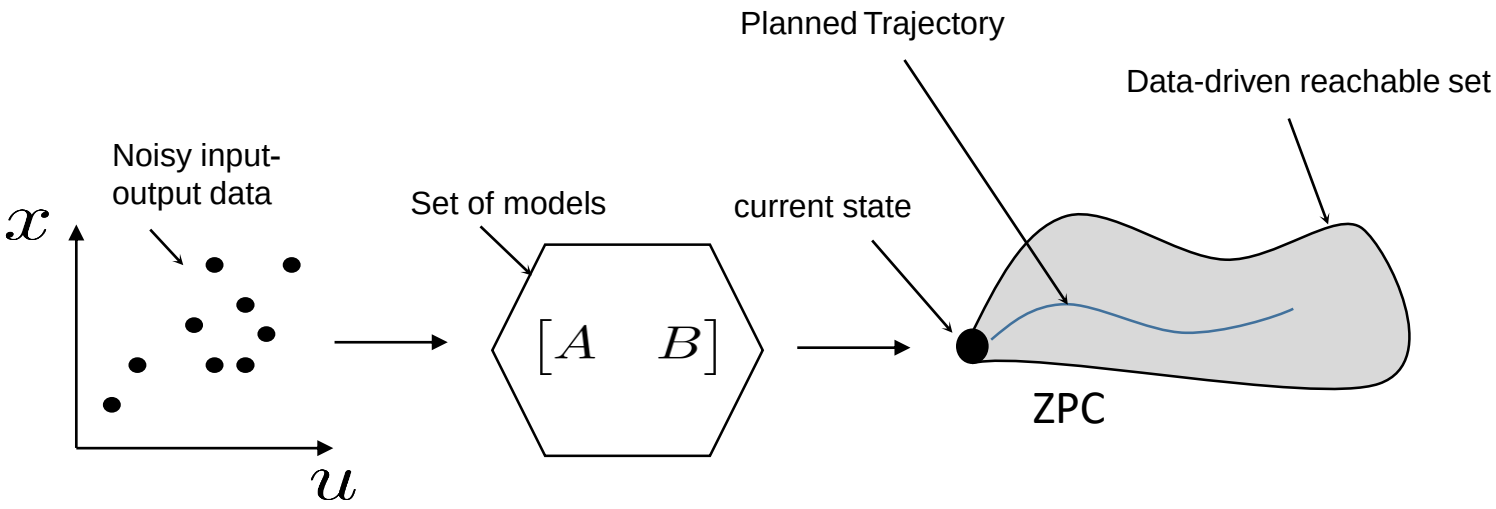

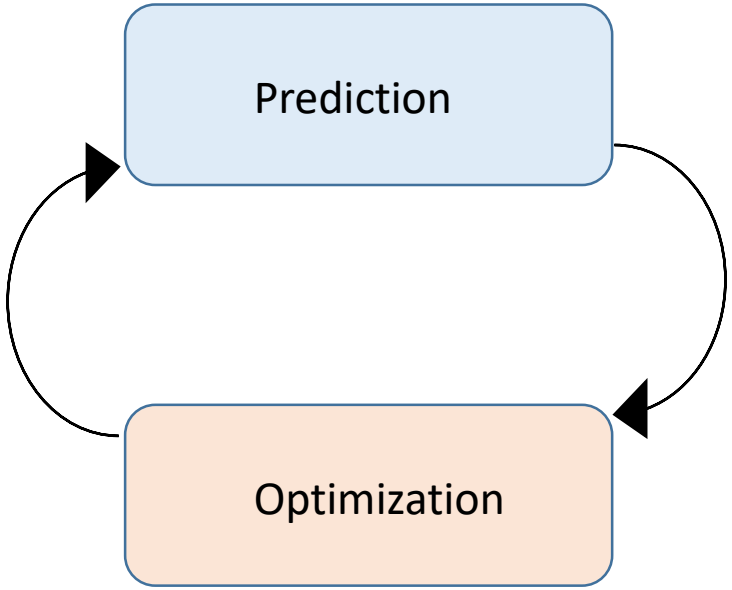

Adjust the Unsafe Action

- Adjust unsafe action by solving a new optimization problem

minimize	Difference between old action and new safe action
subject to	output stays within the reachable set
	input constraints
	output constraints
	bounded process noise
	bounded measurement noise

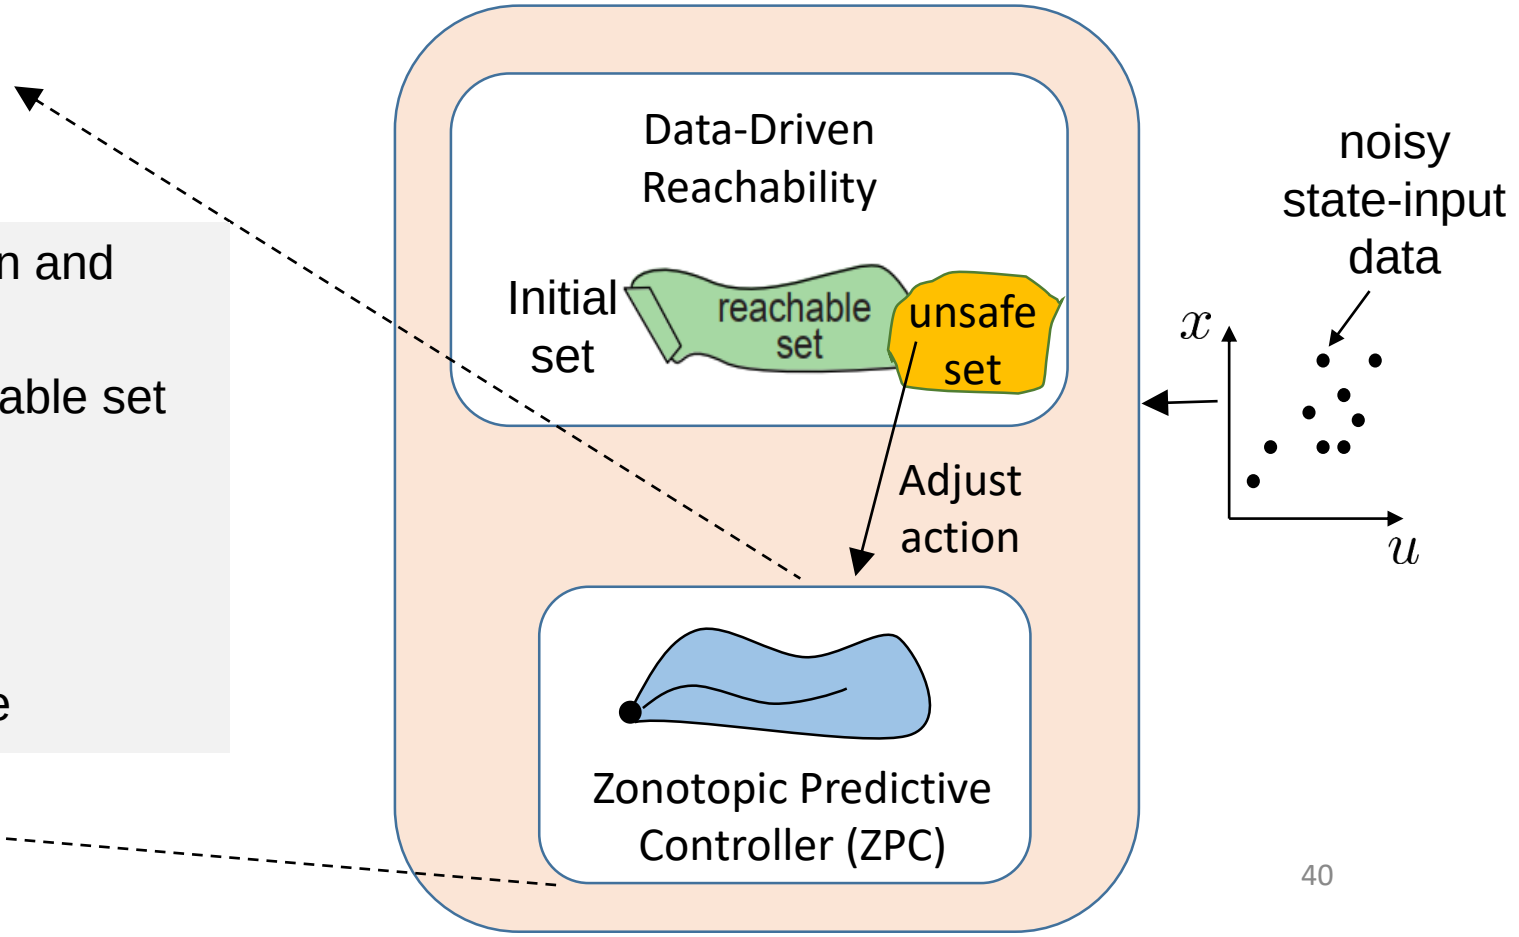

Adjust the Unsafe Action

- Adjust unsafe action by solving a new optimization problem

$$
\begin{aligned}
\min_{u,y,s_u,s_l} \quad & \|u_{\mathrm{RL}} - u_{t|t}\|_R^2 \\
\text{s.t.} \quad & \hat{\mathcal{R}}_{t+k+1|t} = \mathcal{M}_\Sigma(\hat{\mathcal{R}}_{t+k|t} \times \mathcal{Z}_{u,t+k}) + \mathcal{Z}_w + \mathcal{Z}_v - \mathcal{Z}_{Av}, \\
& u_{t+k|t} \in \mathcal{U}_{t+k}, \\
& y_{t+k+1|t} + s_{u,t+k+1|t} = \hat{\mathcal{R}}_{u,t+k+1}, \\
& y_{t+k+1|t} - s_{l,t+k+1|t} = \hat{\mathcal{R}}_{l,t+k+1}, \\
& y_{t+k+1|t} + s_{u,t+k+1|t} \leq \mathcal{Y}_{u,t+k+1}, \\
& y_{t+k+1|t} - s_{l,t+k+1|t} \geq \mathcal{Y}_{l,t+k+1}, \\
& s_{u,t+k+1|t} \geq 0, \\
& s_{l,t+k+1|t} \geq 0, \\
& y_{t|t} = y(t)
\end{aligned}
$$

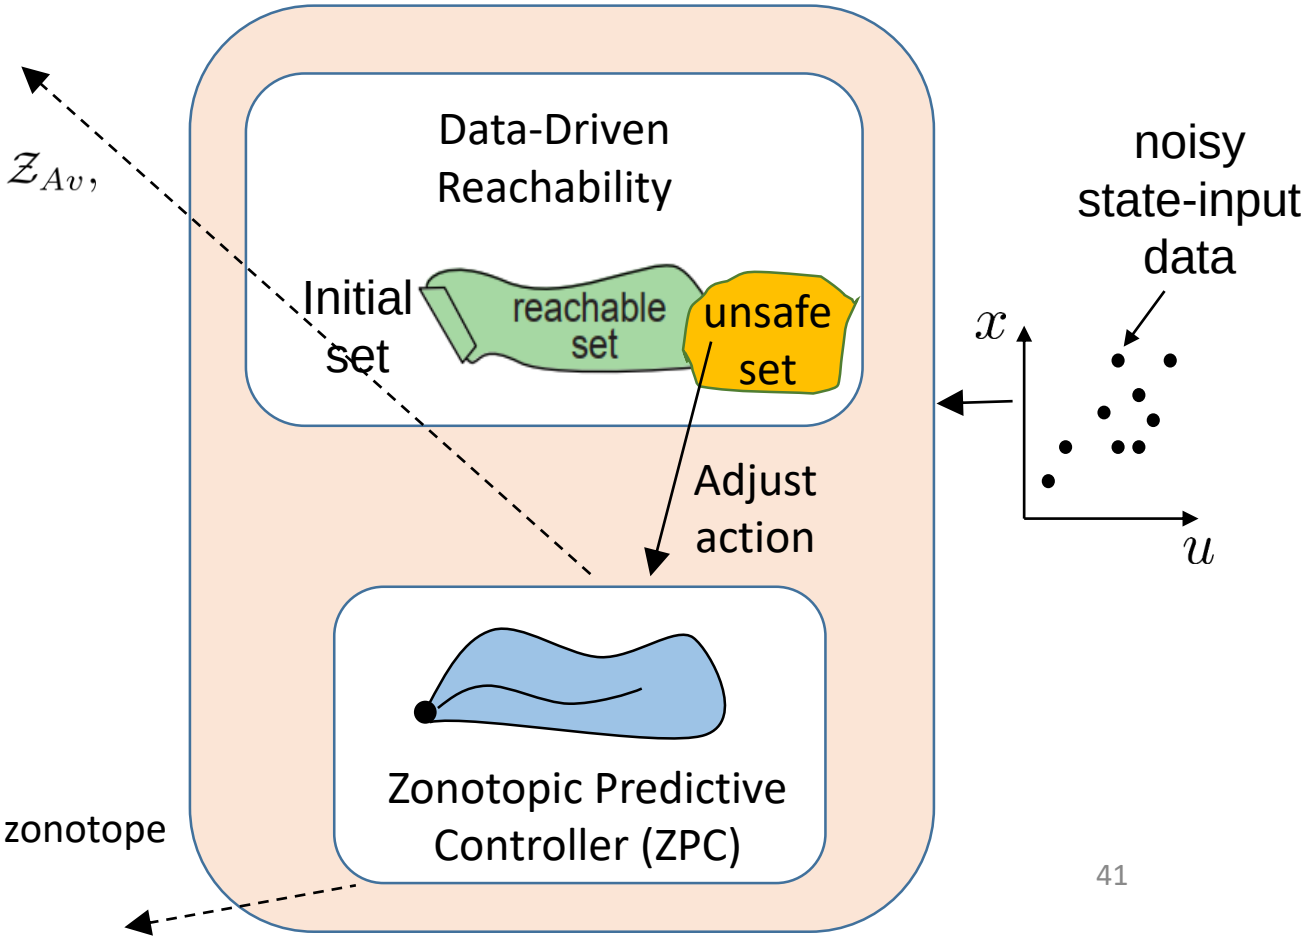

$\mathcal{Y}_{l,t+k+1}$, $\mathcal{Y}_{u,t+k+1}$: lower and upper bounds of the output constraint zonotope

$\hat{\mathcal{R}}_{l,t+k+1}$, $\hat{\mathcal{R}}_{u,t+k+1}$: lower and upper bounds of the reachable set

Example

$$x(k+1) = \begin{bmatrix} 0.9323 & -0.1890 & 0 & 0 & 0 \\ 0.1890 & 0.9323 & 0 & 0 & 0 \\ 0 & 0 & 0.8596 & 0.0430 & 0 \\ 0 & 0 & -0.0430 & 0.8596 & 0 \\ 0 & 0 & 0 & 0 & 0.9048 \end{bmatrix} x(k) + \begin{bmatrix} 0.0436 \\ 0.0533 \\ 0.0475 \\ 0.0453 \\ 0.0476 \end{bmatrix} u(k) + w(k).$$

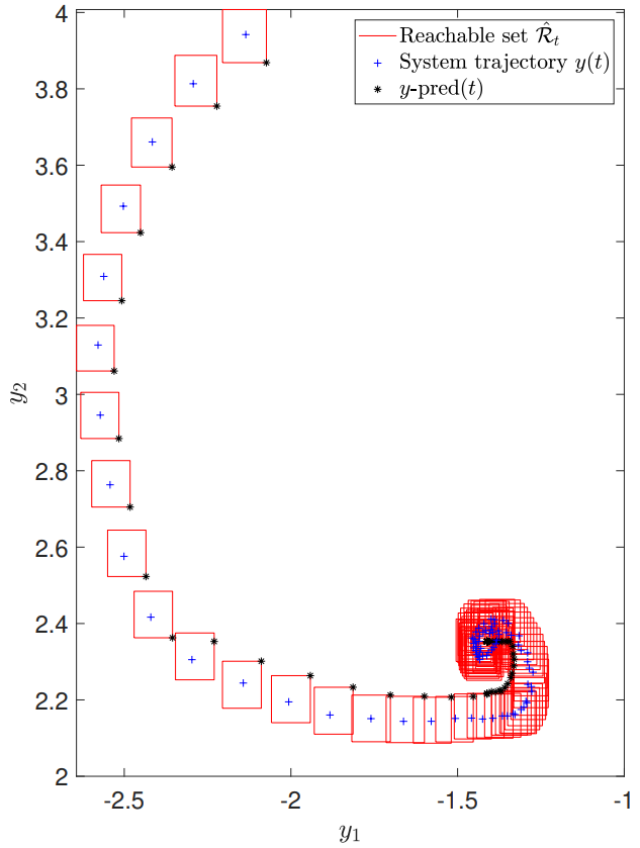

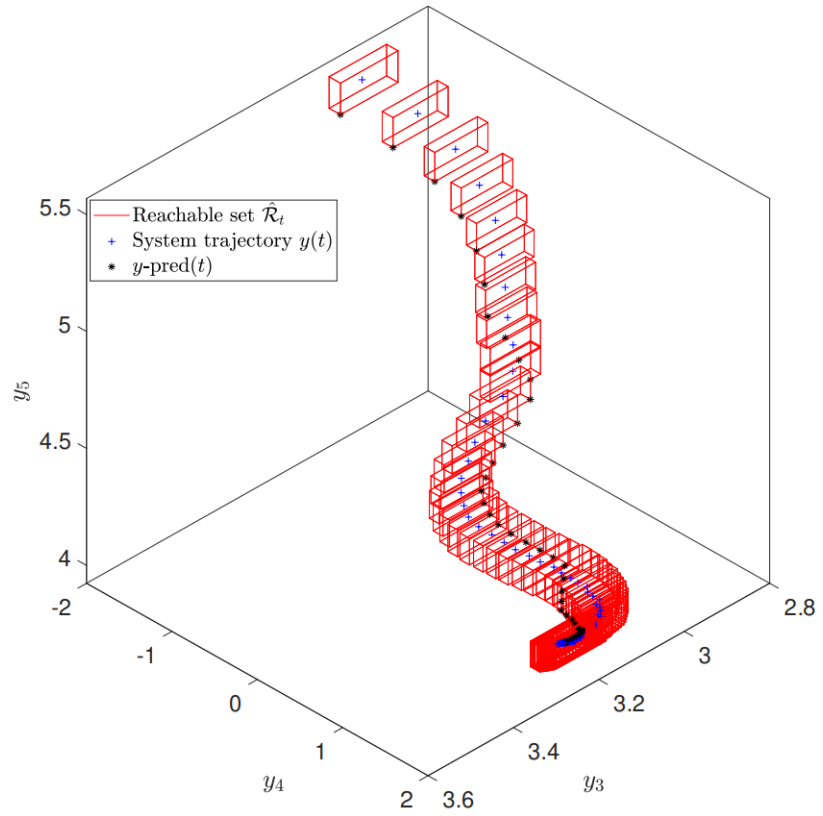

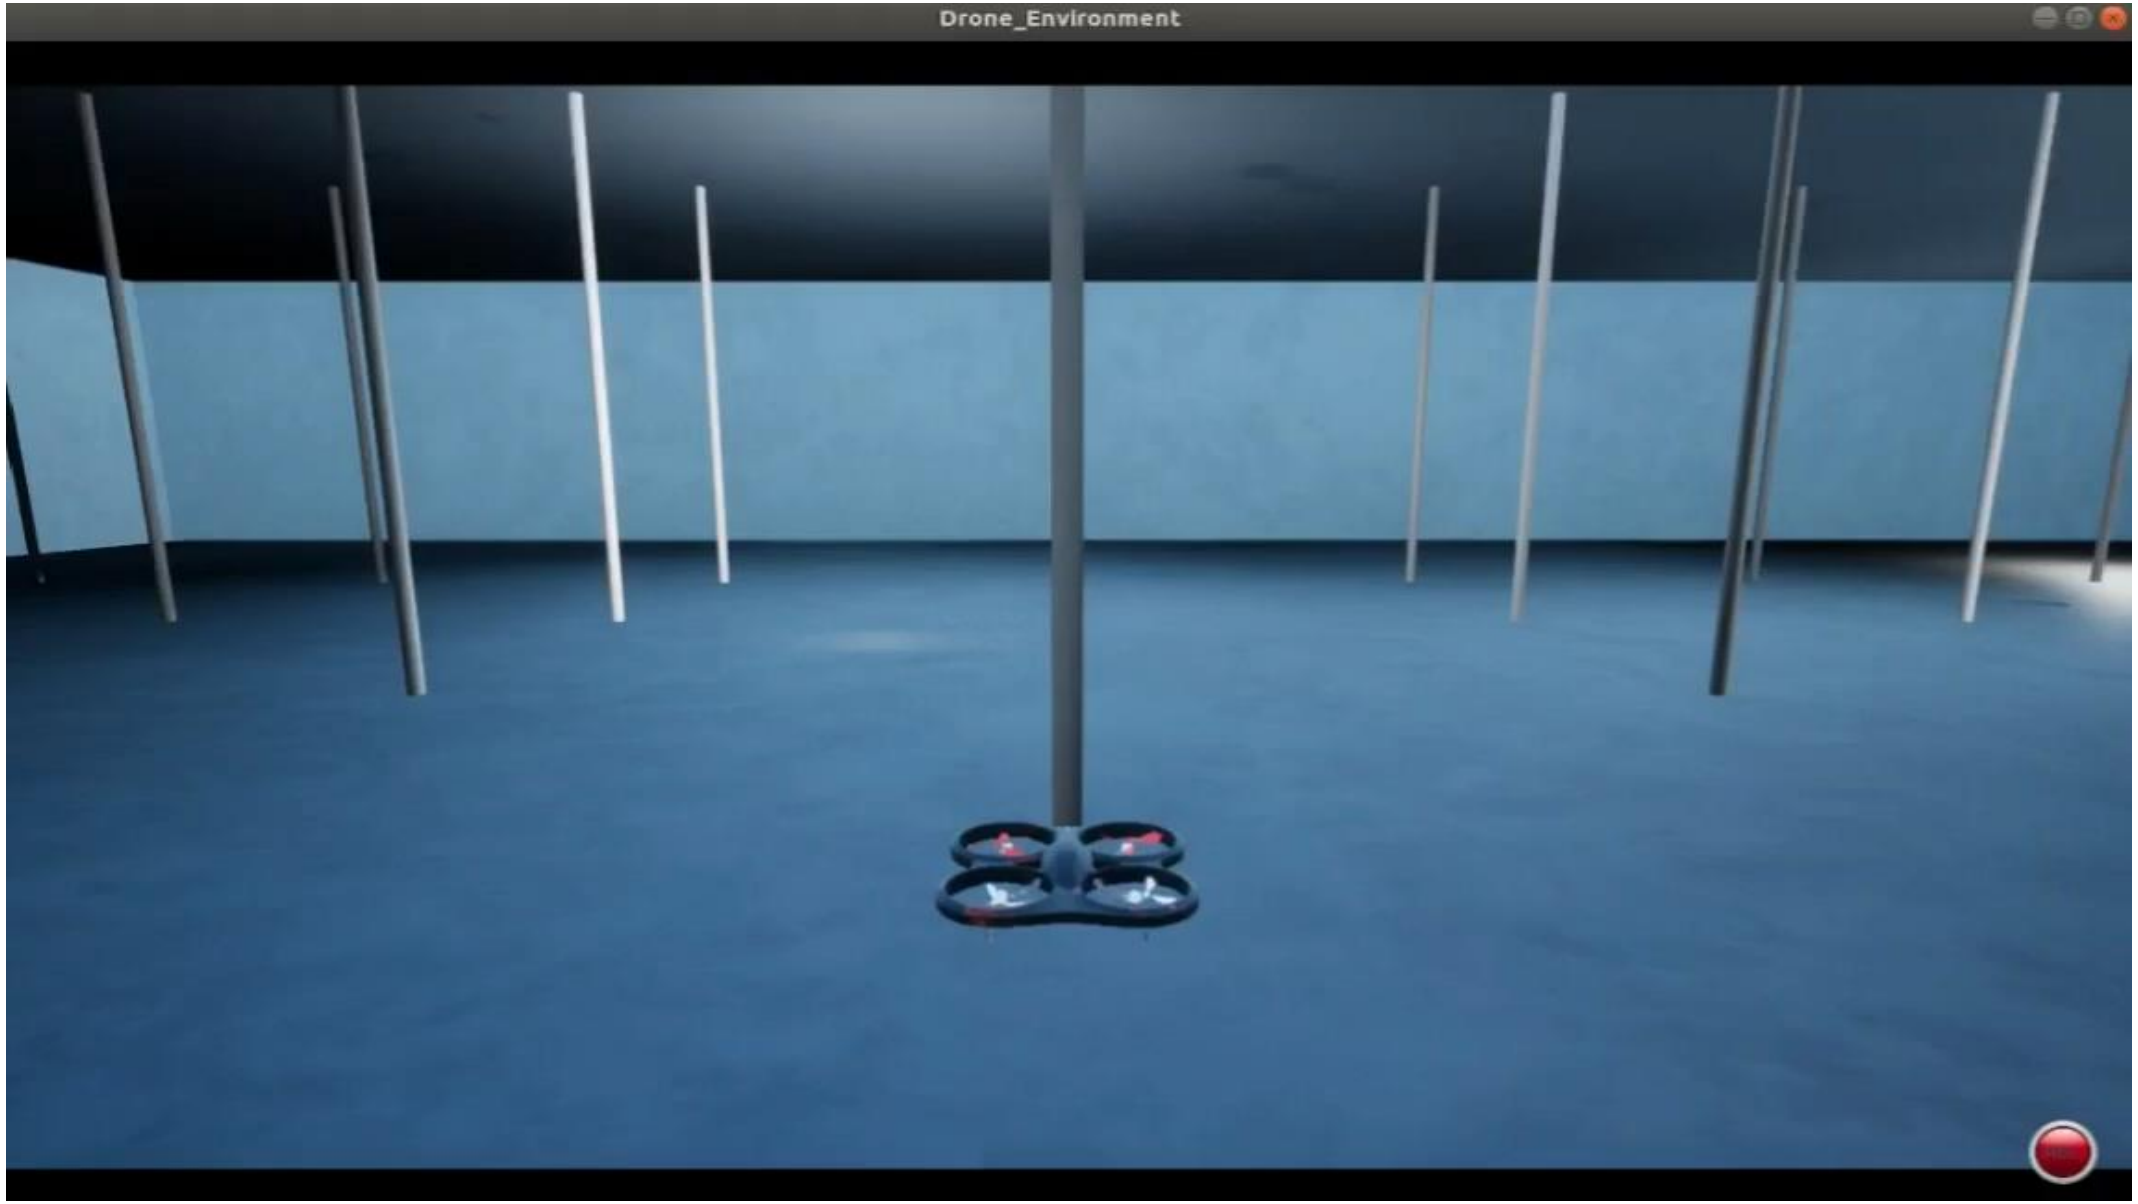

Safe Reinforcement Learning using Data-Driven Predictive Control M Selim, A Alanwar, MW El-Kharashi, HM Abbas, KH Johansson

Acknowledgments


Matthias Althoff


Karl H. Johansson


Henrik Sandberg


Anne Koch


Yvonne Stürz


Frank Allgöwer

- A. Alanwar, A. Koch, F. Allgöwer, and K. H. Johansson, “Data-Driven Reachability Analysis Using Matrix Zonotopes,” Learning for Dynamics and Control, L4DC 2021.
- A. Alanwar, A. Koch, F. Allgöwer, and K. H. Johansson, “Data-Driven Reachability Analysis from Noisy Data,” IEEE Transactions on Automatic Control, TAC, 2023.
- A. Alanwar, Y. Stürz, and K. H. Johansson, “Robust data-driven predictive control using reachability analysis,” European Control Conference, ECC 2022.
- A Berndt, A. Alanwar, K. H. Johansson and H. Sandberg “Data-driven set-based estimation using matrix zonotopes with set containment guarantees,” European Control Conference, ECC 2022.
- A. Alanwar, F. Jiang, M. Sharifi, D. V. Dimarogonas, K. H. Johansson, “Enhancing Data-Driven Reachability Analysis using Temporal Logic Side Information,” International Conference on Robotics and Automation, ICRA 2022.

Knut och Alice Wallenbergs Stiftelse

Vetenskapsrådet

VINNOVA

Conclusion

- Data-driven safety verification of an embedded system using reachability with formal guarantees
- Compute a set of models that is consistent with the data
- How to incorporate a model-based side information to get a tighter set

Ongoing work

- STL-based side information
- Guaranteed safe reinforcement learning

https://sites.google.com/view/amr-alanwar/

https://github.com/aalanwar/


THANKS