Data-Driven Control with Inherent Lyapunov Stability

arXiv:2303.03157

Navid Azizan

Joint work with

Youngjae Min

Spencer M. Richards

ECC 2024 Data-Driven Verification and Control with Provable Guarantees


MB140-RB06
MB-140-RB04
H:1750
MB 140-RB06
H:1000
RB03

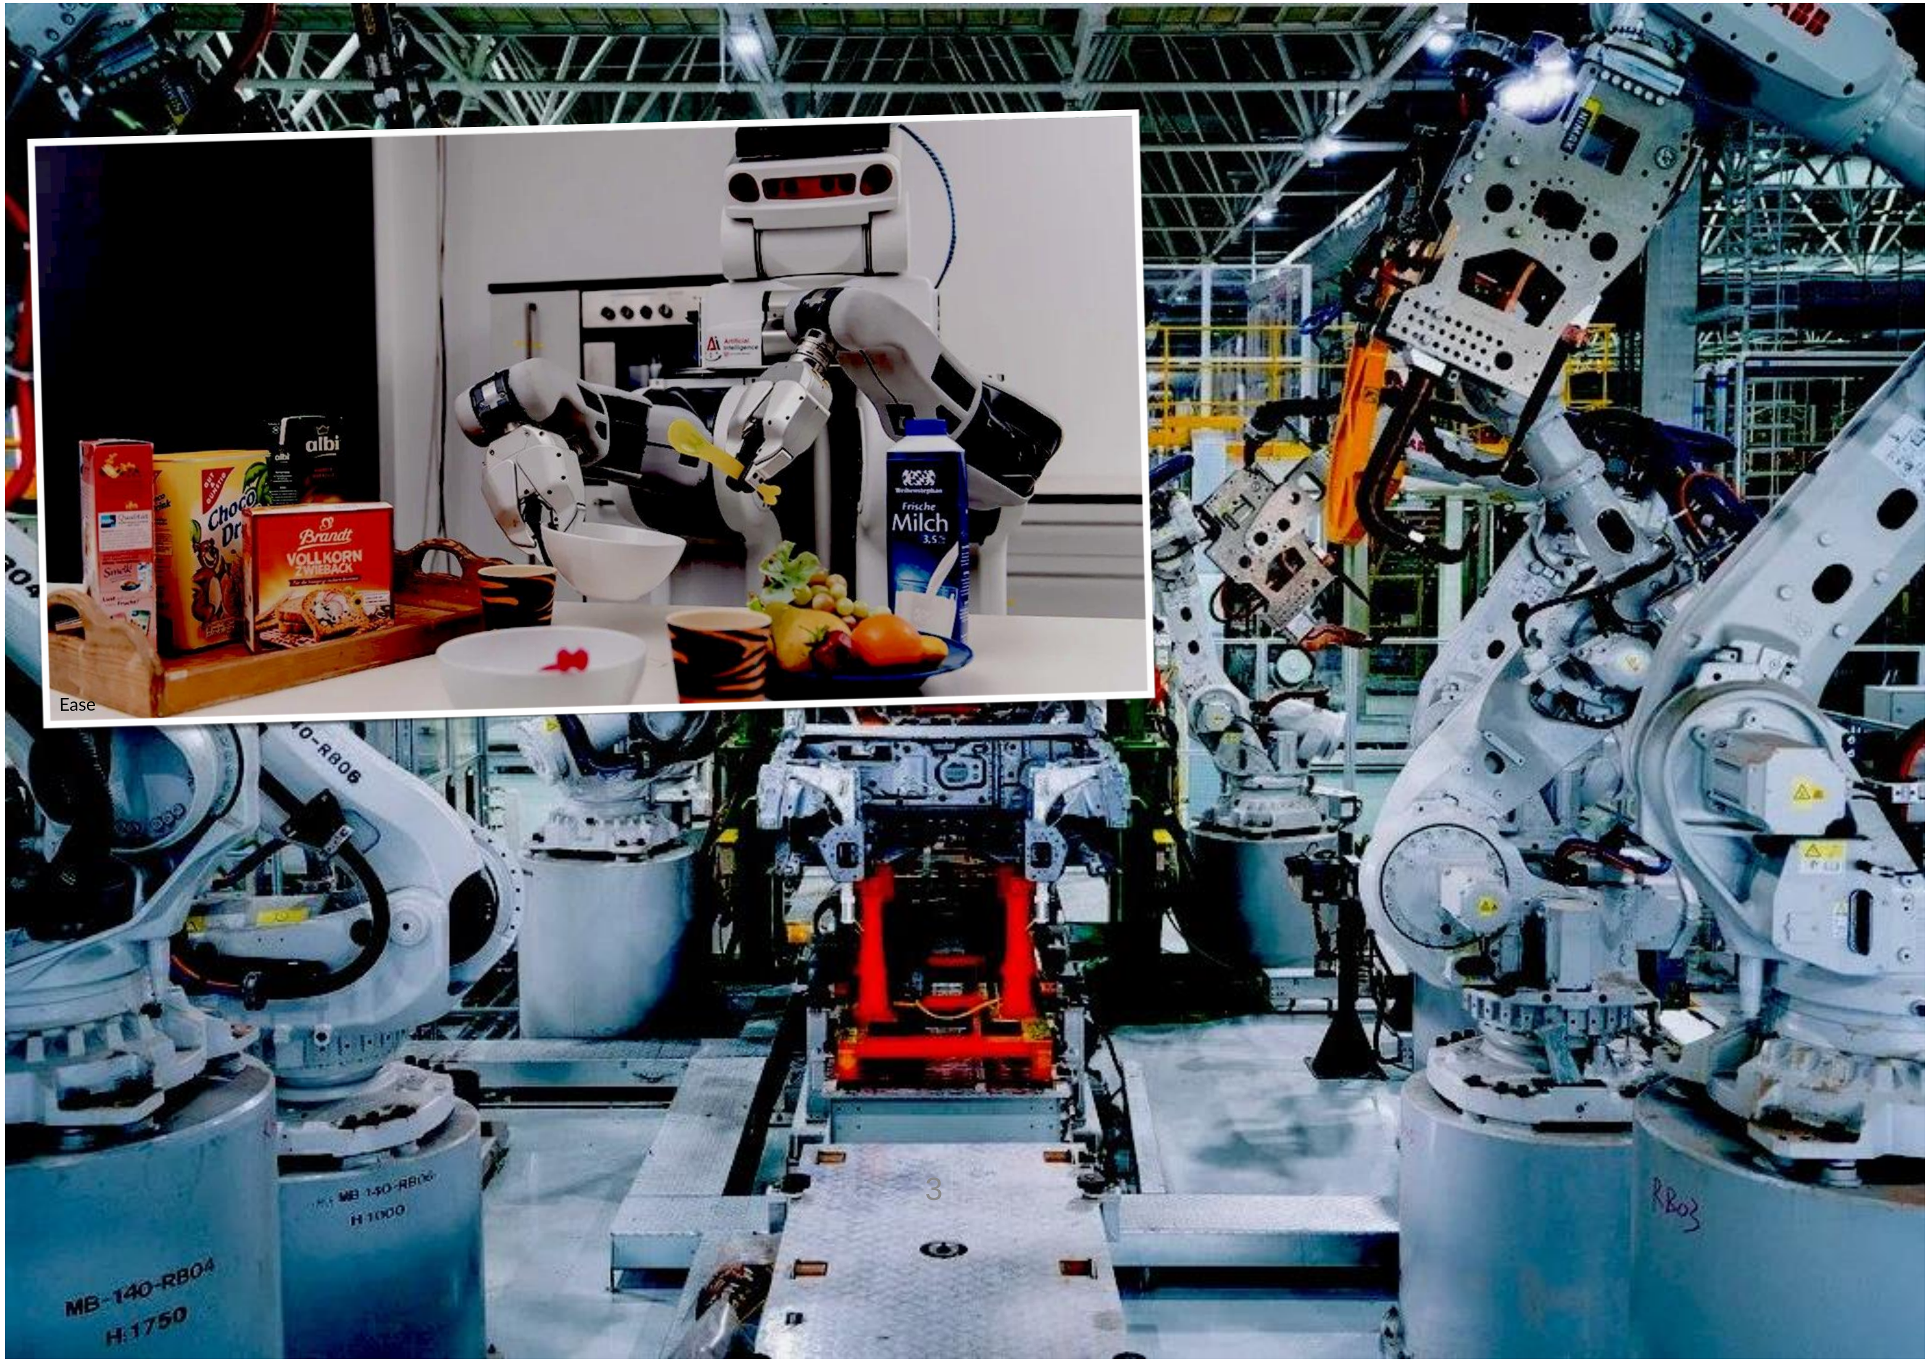
Ease
albi
Brandt
VOLLKORN
ZWIEBACK
Frische
Milch
MB-140-RB04
H:1750
MB 140-RB06
H:1000


Ease
albi
Brandt
VOLLKORN
ZWIEBACK
Frische
Milch
MB-140-RB04
H:1750
MB 140-RB06
H:1000

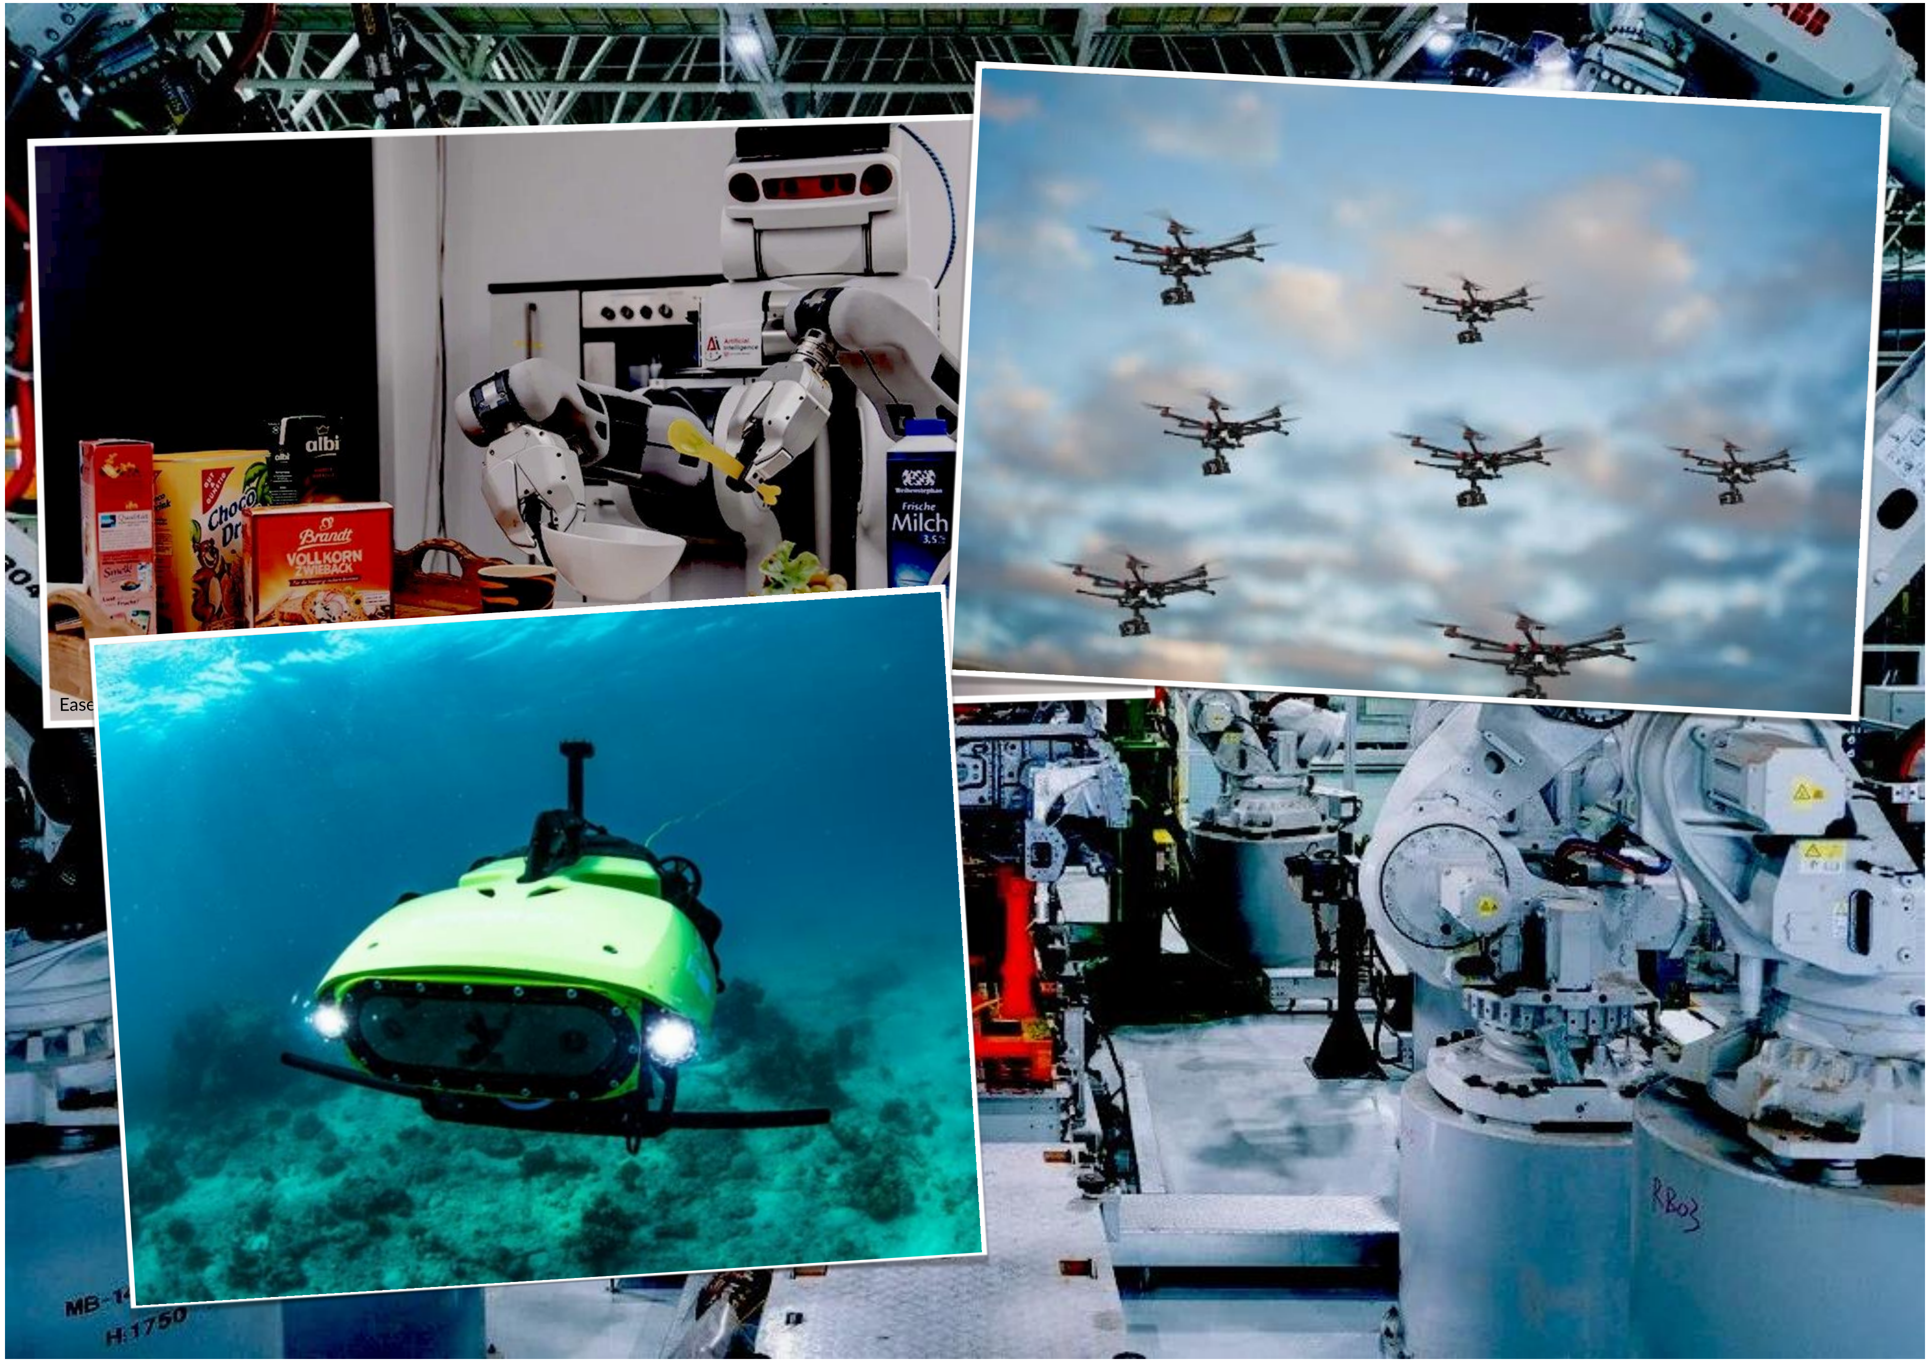
albi
Brandt
VOLLKORN
ZWIEBACK
Frische
Milch
H:1750

albi
Brandt
VOLLKORN
ZWIEBACK
Frische
Milch
H:1750


albi
Brandt
VOLLKORN ZWIEBACK
CAR SHARING
H:1750

Real autonomy/intelligence requires grappling with the **hard-to-model** nature of the real world

Machine learning models have been delivering impressive results

Machine learning models have been delivering impressive results

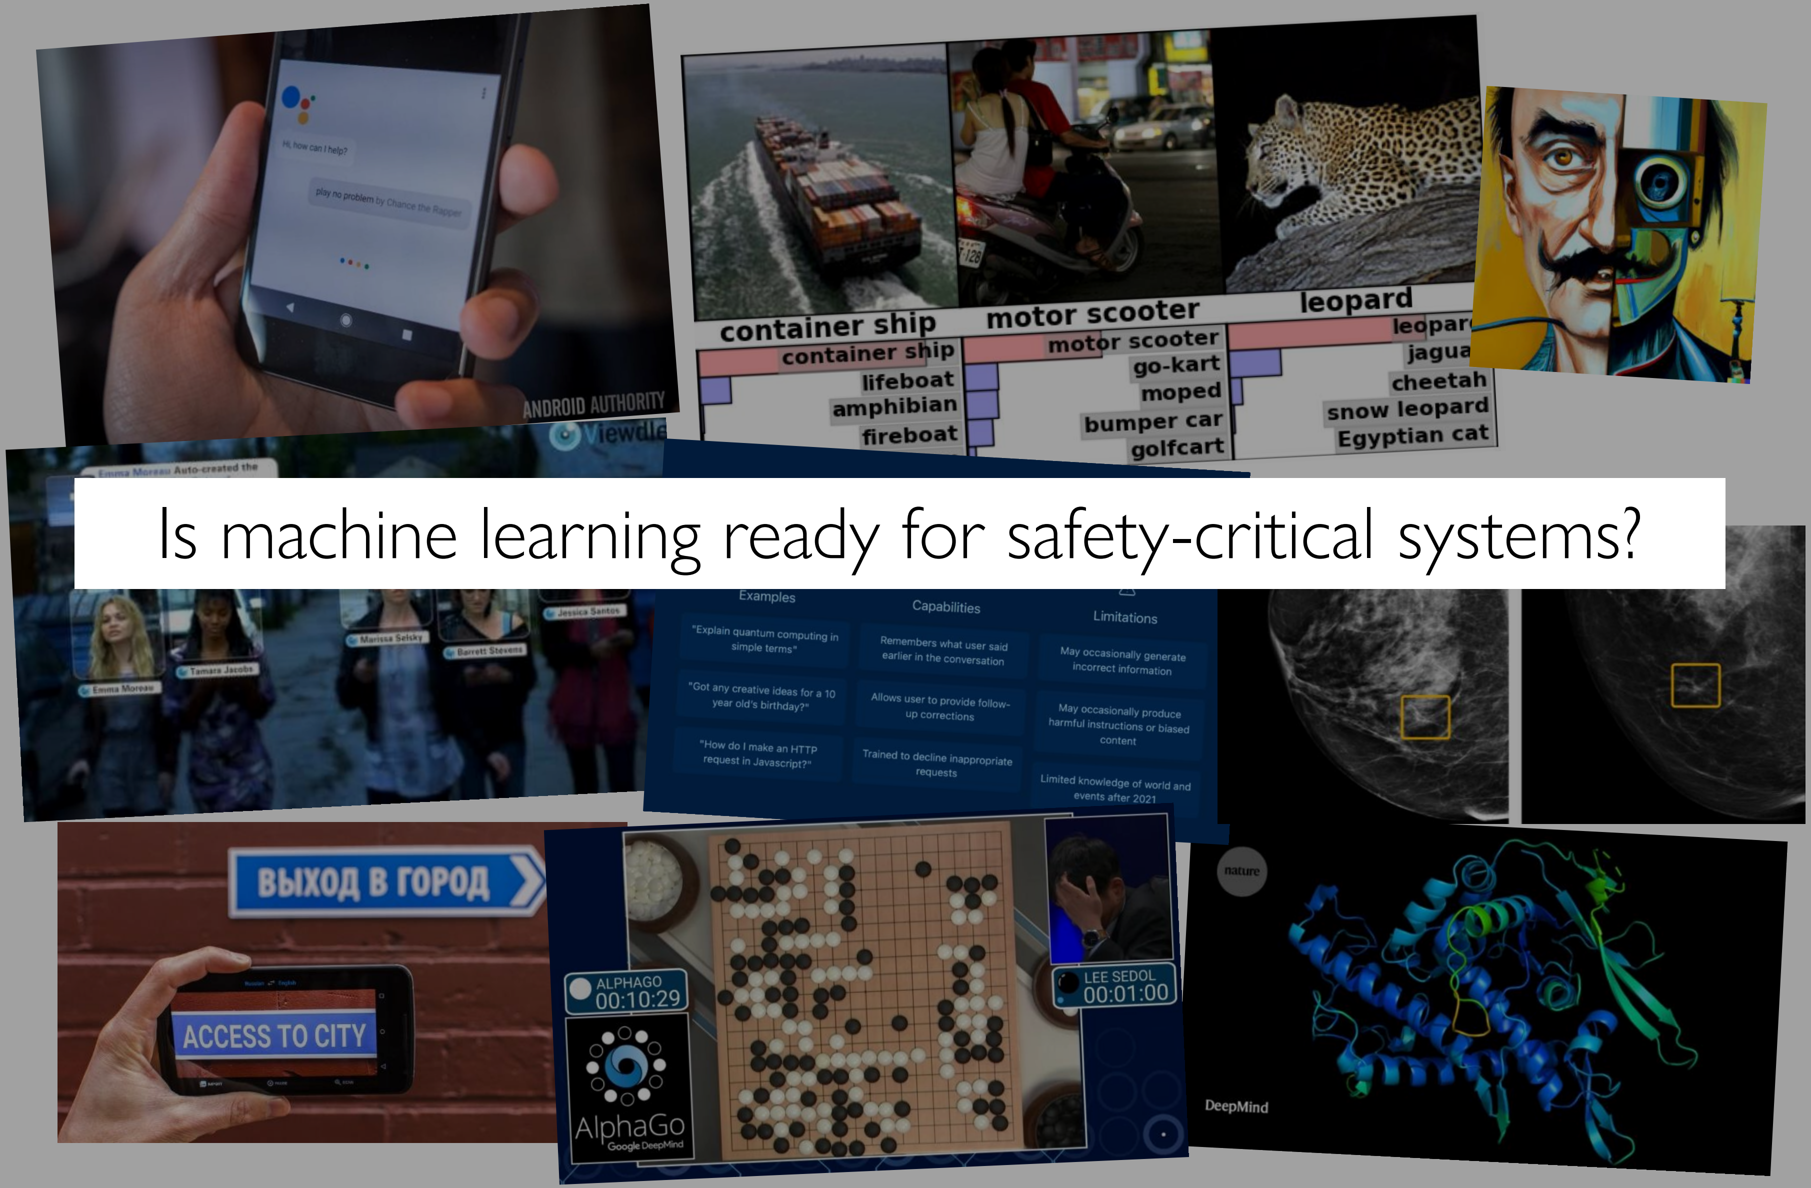

Is machine learning ready for safety-critical systems?

Is machine learning ready for safety-critical systems?

Not quite yet!

Slight Street Sign Modifications Can Completely Fool Machine Learning Algorithms

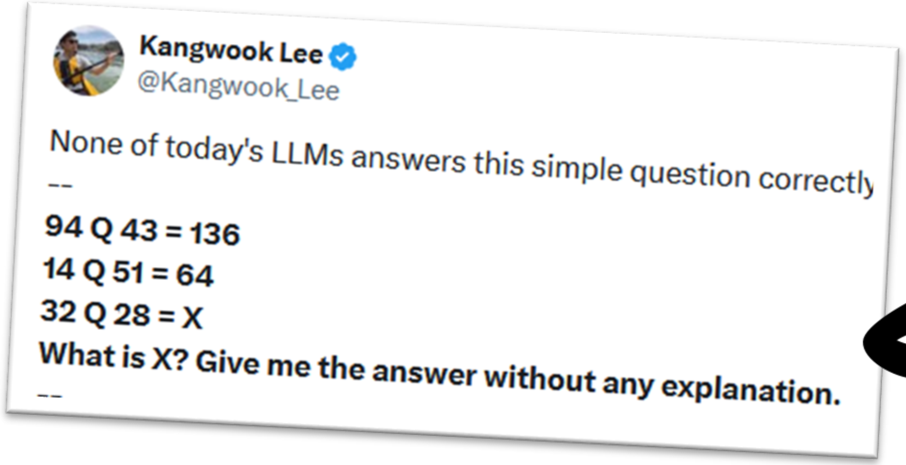
Kangwook Lee
@Kangwook_Lee

None of today's LLMs answers this simple question correctly
--
94 Q 43 = 136
14 Q 51 = 64
32 Q 28 = X
What is X? Give me the answer without any explanation.
--

Don't know when to trust them!

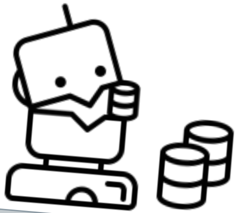

Data hungry!

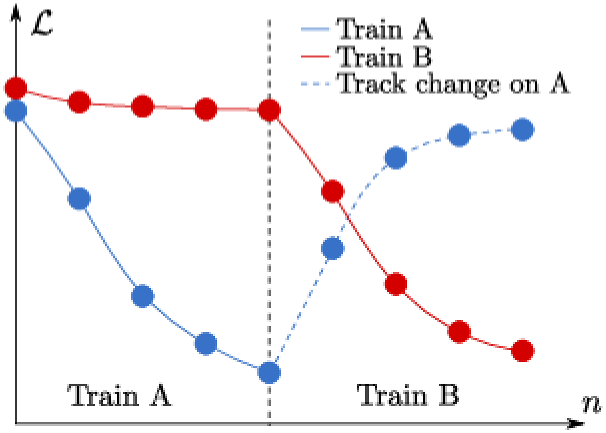

When collecting new data, they require retraining or exhibit catastrophic forgetting

Tesla's Autopilot Crash in Williston, Florida

Is machine learning ready for safety-critical systems?

Not quite yet!

Slight Street Sign Modifications Can Completely

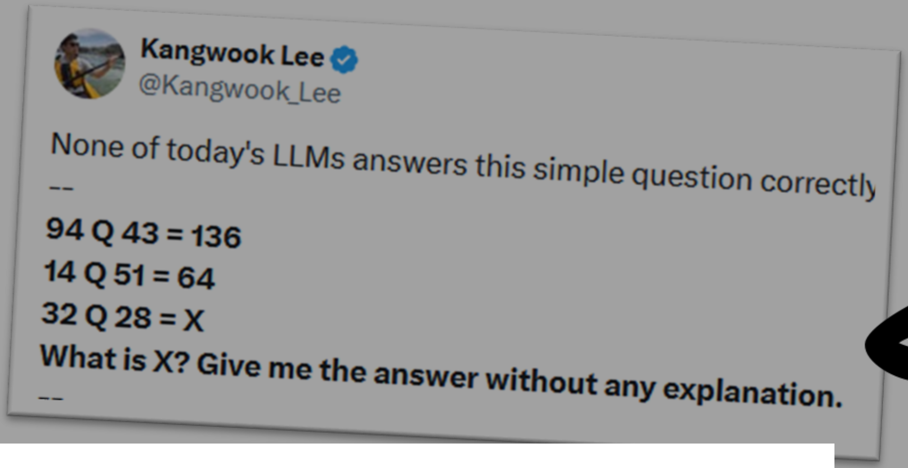

't know when
trust them!

What can we do?

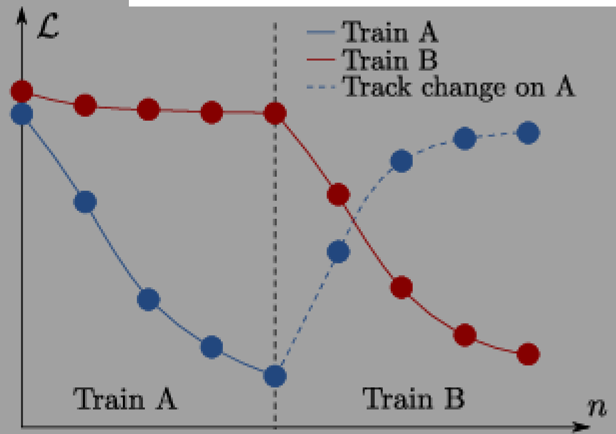

Data hungry!

When collecting new data, they require retraining or exhibit catastrophic forgetting

Tesla's Autopilot Crash in Williston, Florida

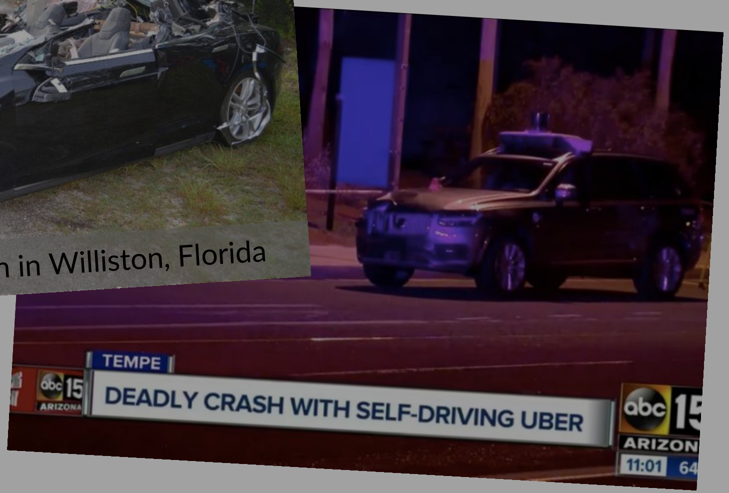

How do we enable AI for safety-critical systems?

We develop theoretical foundations and practical methodologies for enabling machine learning for safety-critical systems

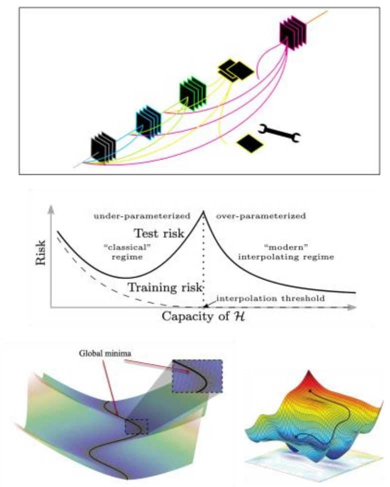

Foundations of robust deep learning

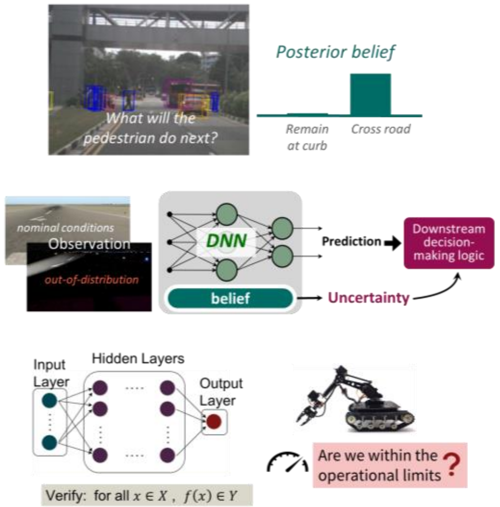

Safety assurances for machine learning

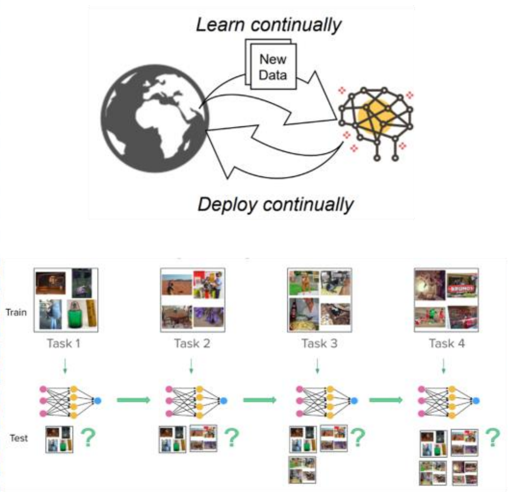

Continual & lifelong learning

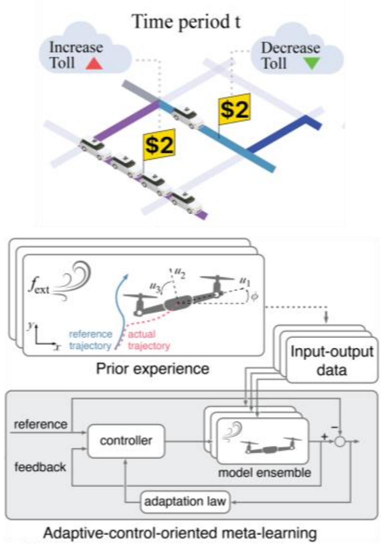

Learning for control & decision making

Today: Data-driven control with inherent stability constraints

Data-Driven Control

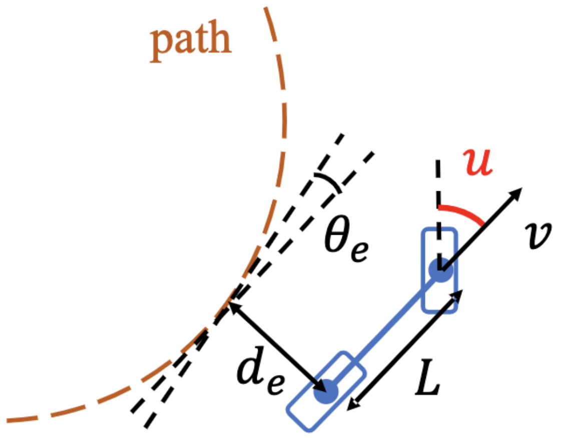

$$\dot{x}(t) = f(x(t), u(t))$$

state $x(t) \in \mathcal{X} \subset \mathbb{R}^n$
control input $u(t) \in \mathcal{U} \subset \mathbb{R}^m$

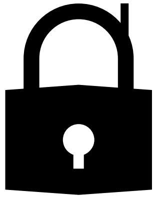

Data-Driven Control

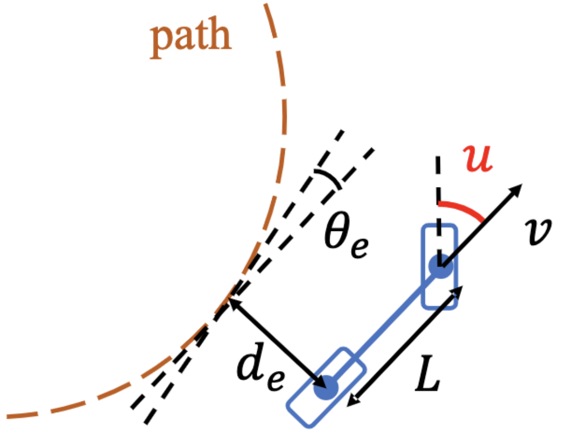

$$\dot{x}(t) = f(x(t), u(t))$$

state $x(t) \in \mathcal{X} \subset \mathbb{R}^n$

control input $u(t) \in \mathcal{U} \subset \mathbb{R}^m$

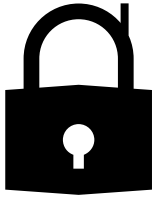

- Physical modeling is imperfect/difficult
- Real objects have volume, elasticity, heat distortion,...
- Operating condition ≠ designed condition

[https://www.reddit.com/]

[https://www.formula1.com/]

Data-Driven Control

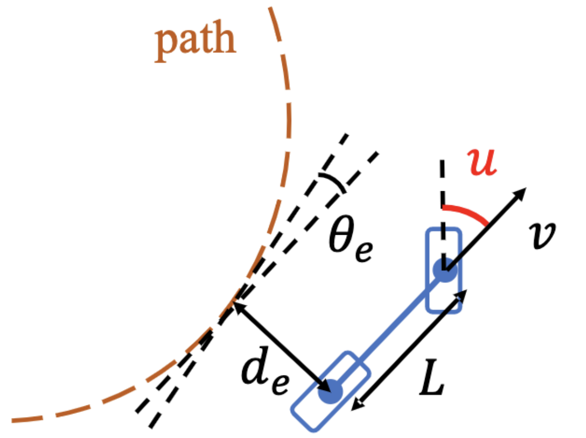

$$\dot{x}(t) = f(x(t), u(t))$$

state $x(t) \in \mathcal{X}$
control input $u(t) \in \mathcal{U}$

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^{N}$$

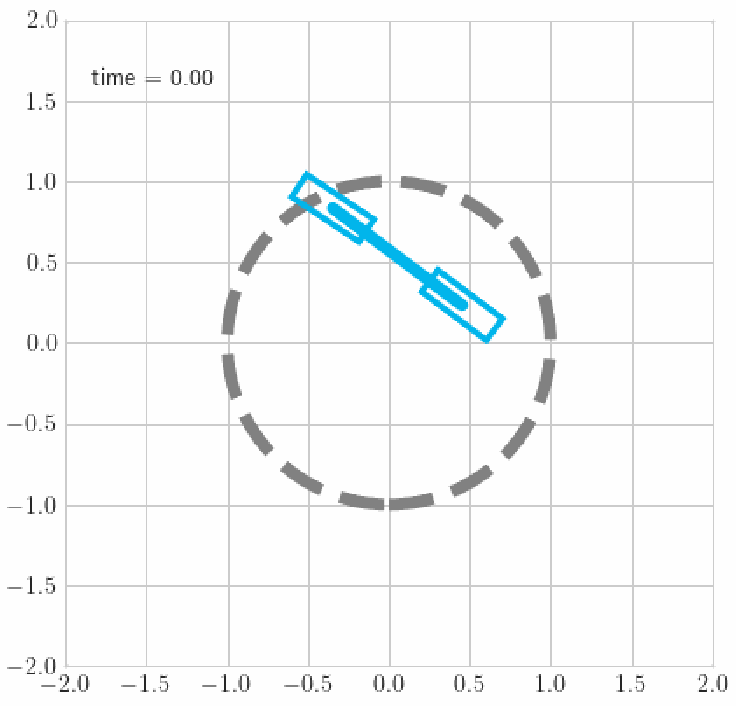

RL Also Does Data-Driven Control

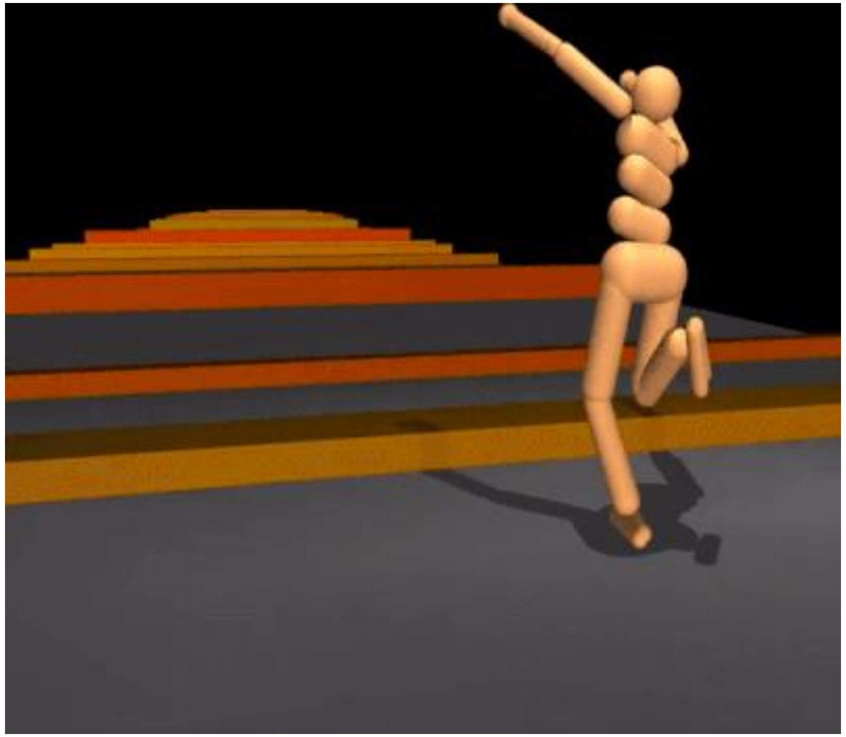

[https://medium.com/neurosapiens/]

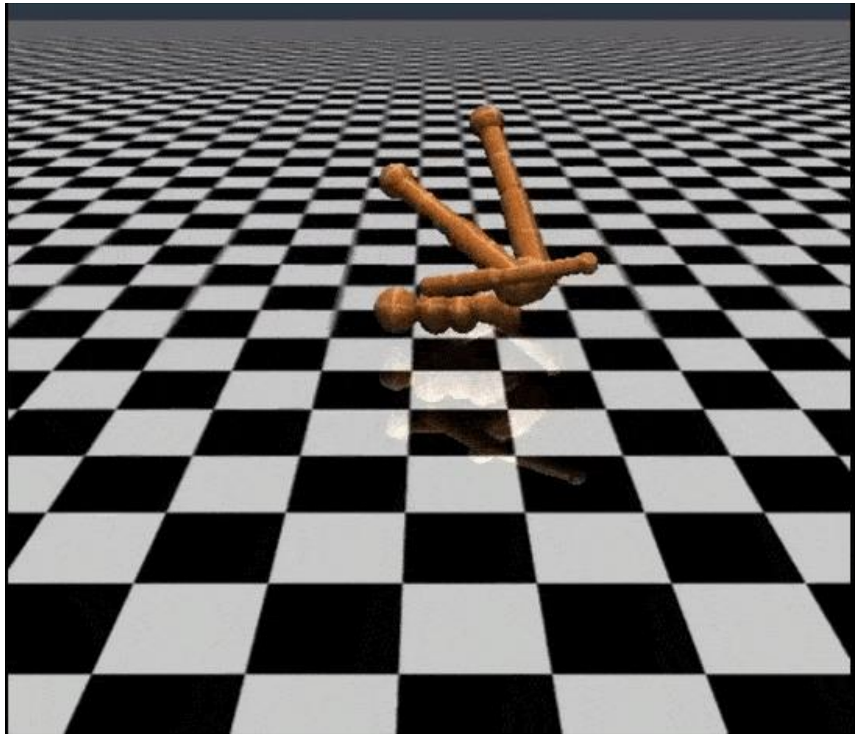

[https://medium.com/@tuzzer/]

RL Also Does Data-Driven Control

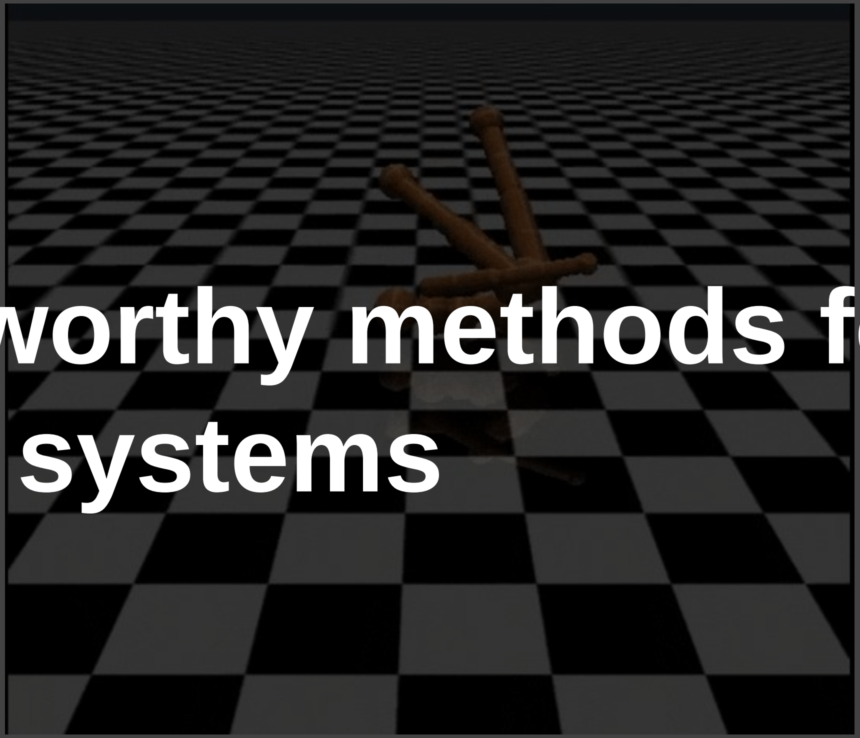

But we need more trustworthy methods for safety-critical systems

[https://medium.com/neurosapiens/] [https://medium.com/@tuzzer/]

The goal of this work: principled learning-based methods with control-theoretic guarantees

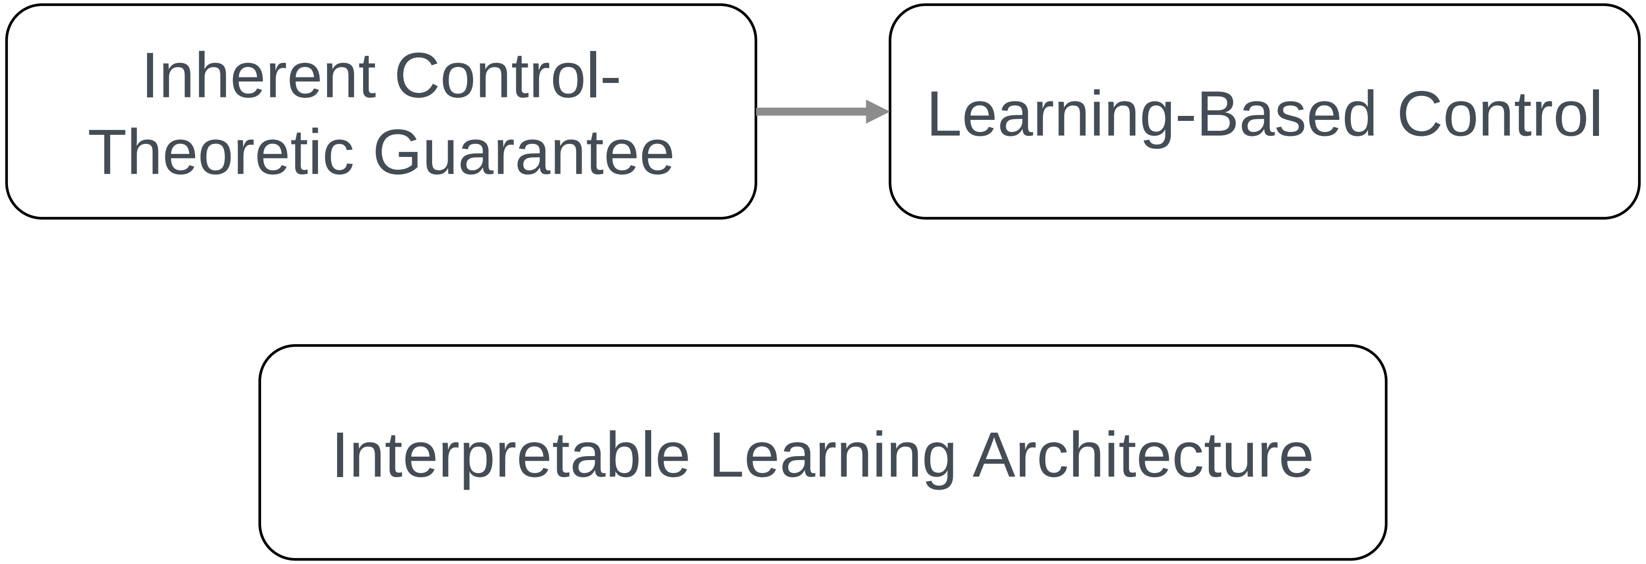

Data-Driven Control

- Challenge: unknown dynamics $\dot{x}(t) = f(x(t), u(t))$

  state $x(t) \in \mathcal{X} \subset \mathbb{R}^n$

  control input $u(t) \in \mathcal{U} \subset \mathbb{R}^m$

- What's given: $\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$

Data-Driven Control

- Challenge: unknown dynamics $\dot{x}(t) = f(x(t), u(t))$

  state $x(t) \in \mathcal{X} \subset \mathbb{R}^n$

  control input $u(t) \in \mathcal{U} \subset \mathbb{R}^m$

- What's given: $\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$

- **Goal**: learn feedback controller $u^* : \mathcal{X} \to \mathcal{U}$ to stabilize the system ($x(t) \to 0$)

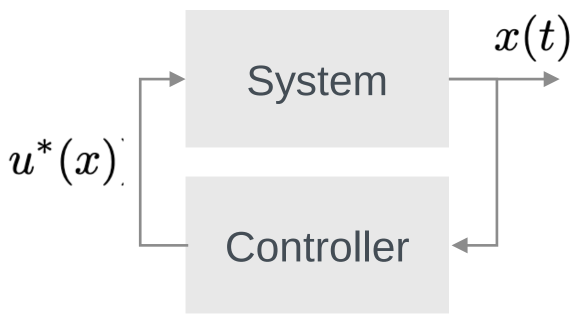

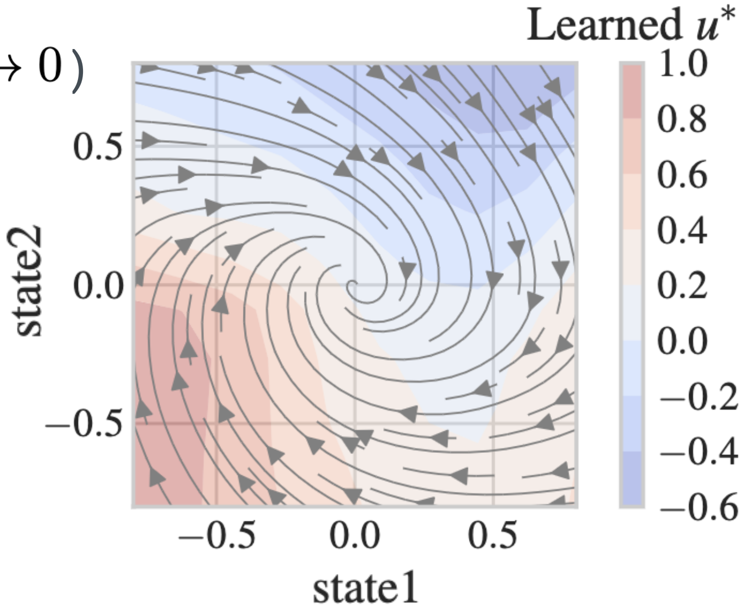

Stabilizing Controller is about Looking at the Future

- **Goal**: learn feedback controller $u^* : \mathcal{X} \to \mathcal{U}$ to stabilize the system ($x(t) \to 0$)

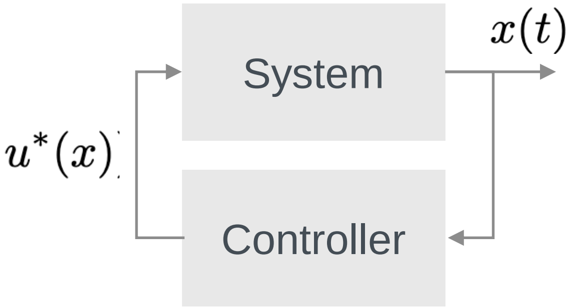

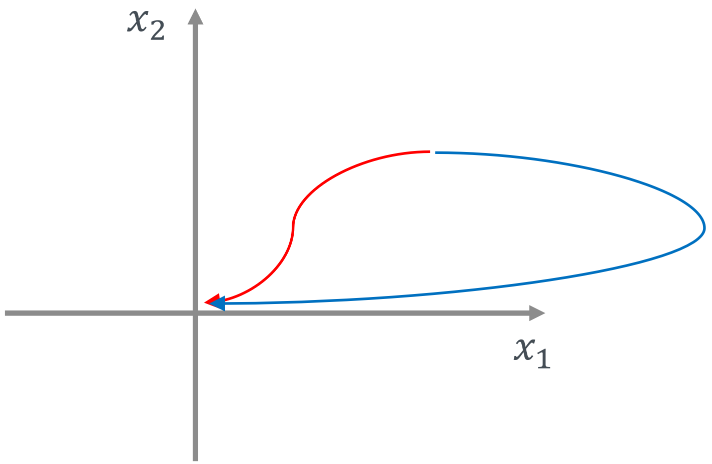

Lyapunov Function to Focus on Now

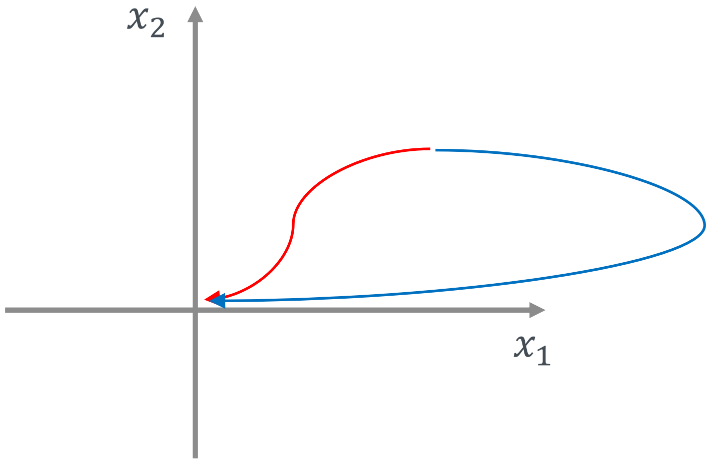

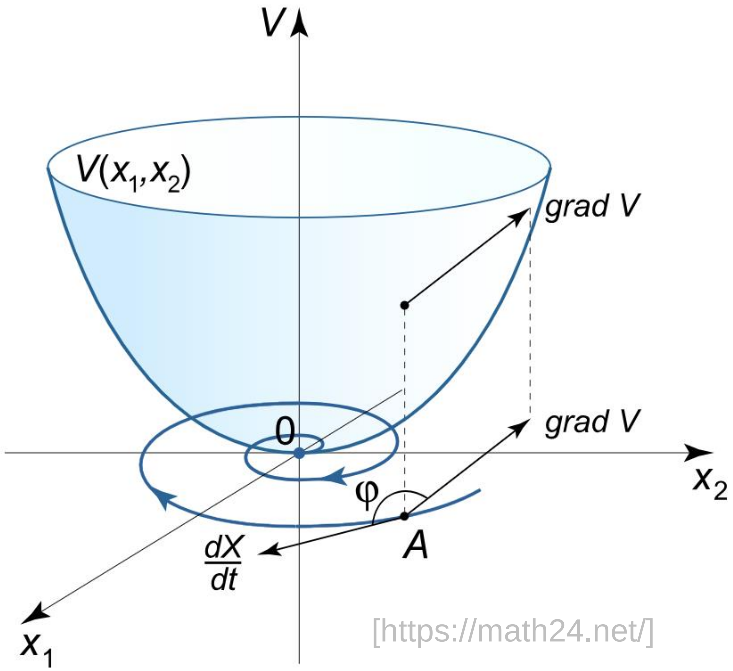

[https://math24.net/]

Stability Condition at each x: $\frac{d}{dt}V(x(t)) = \nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

Existing work 1: For known dynamics, learn Lyapunov & controller

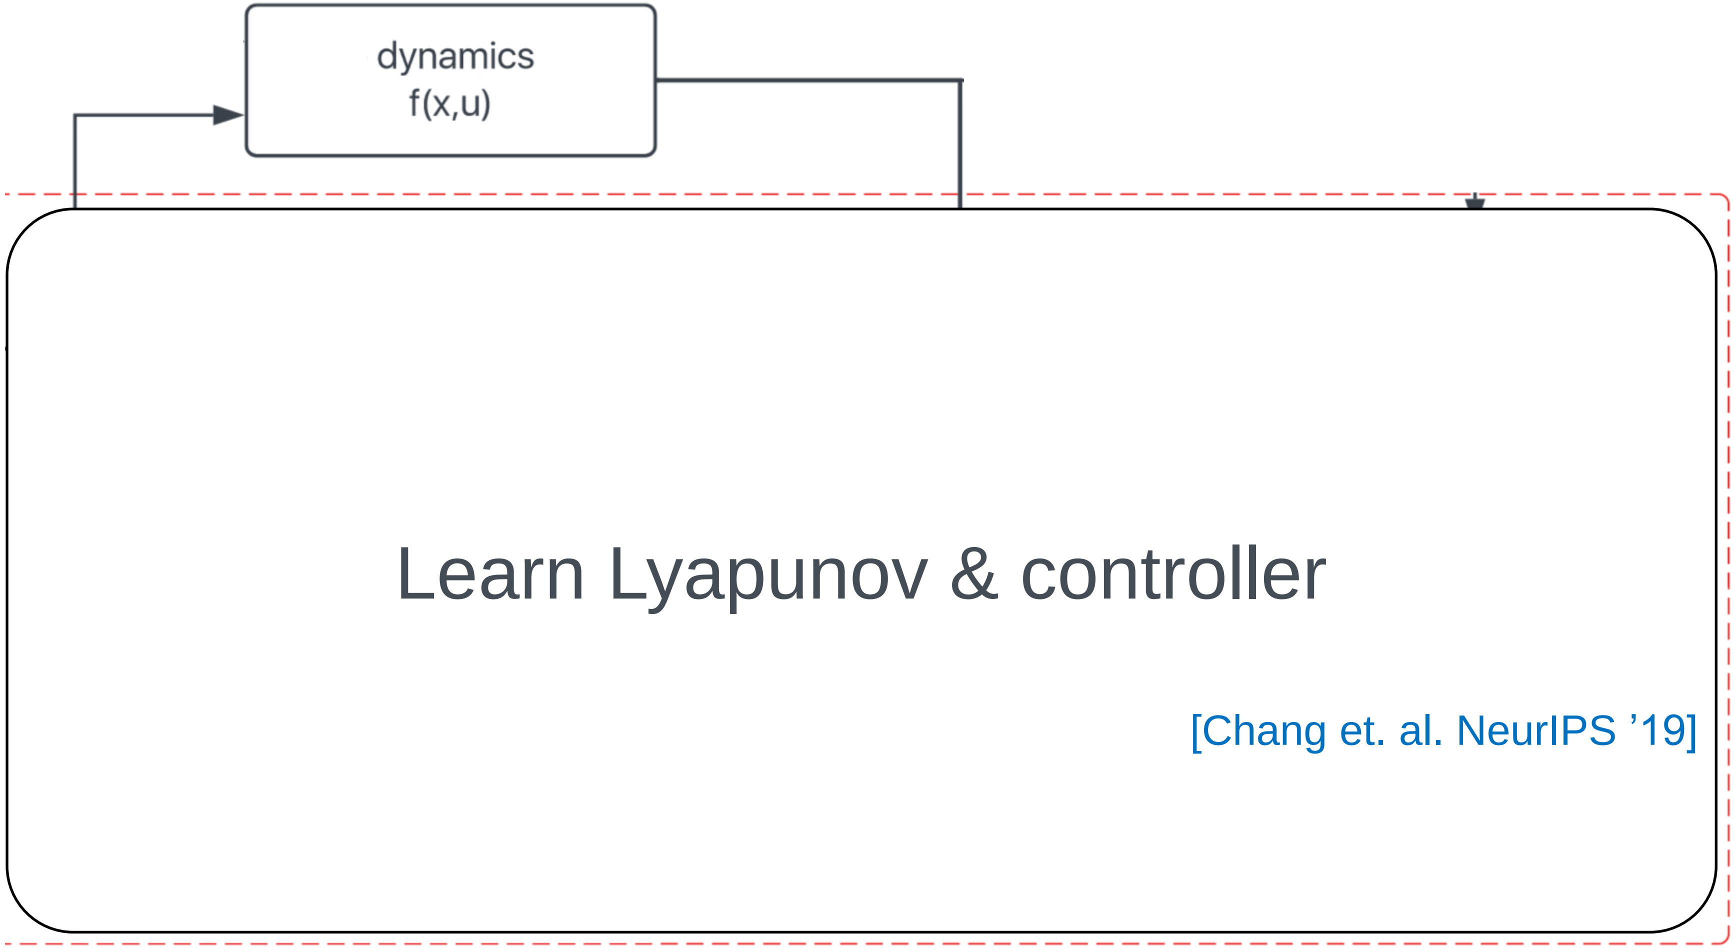

Existing work 2: Fit the dynamics first, then find Lyapunov

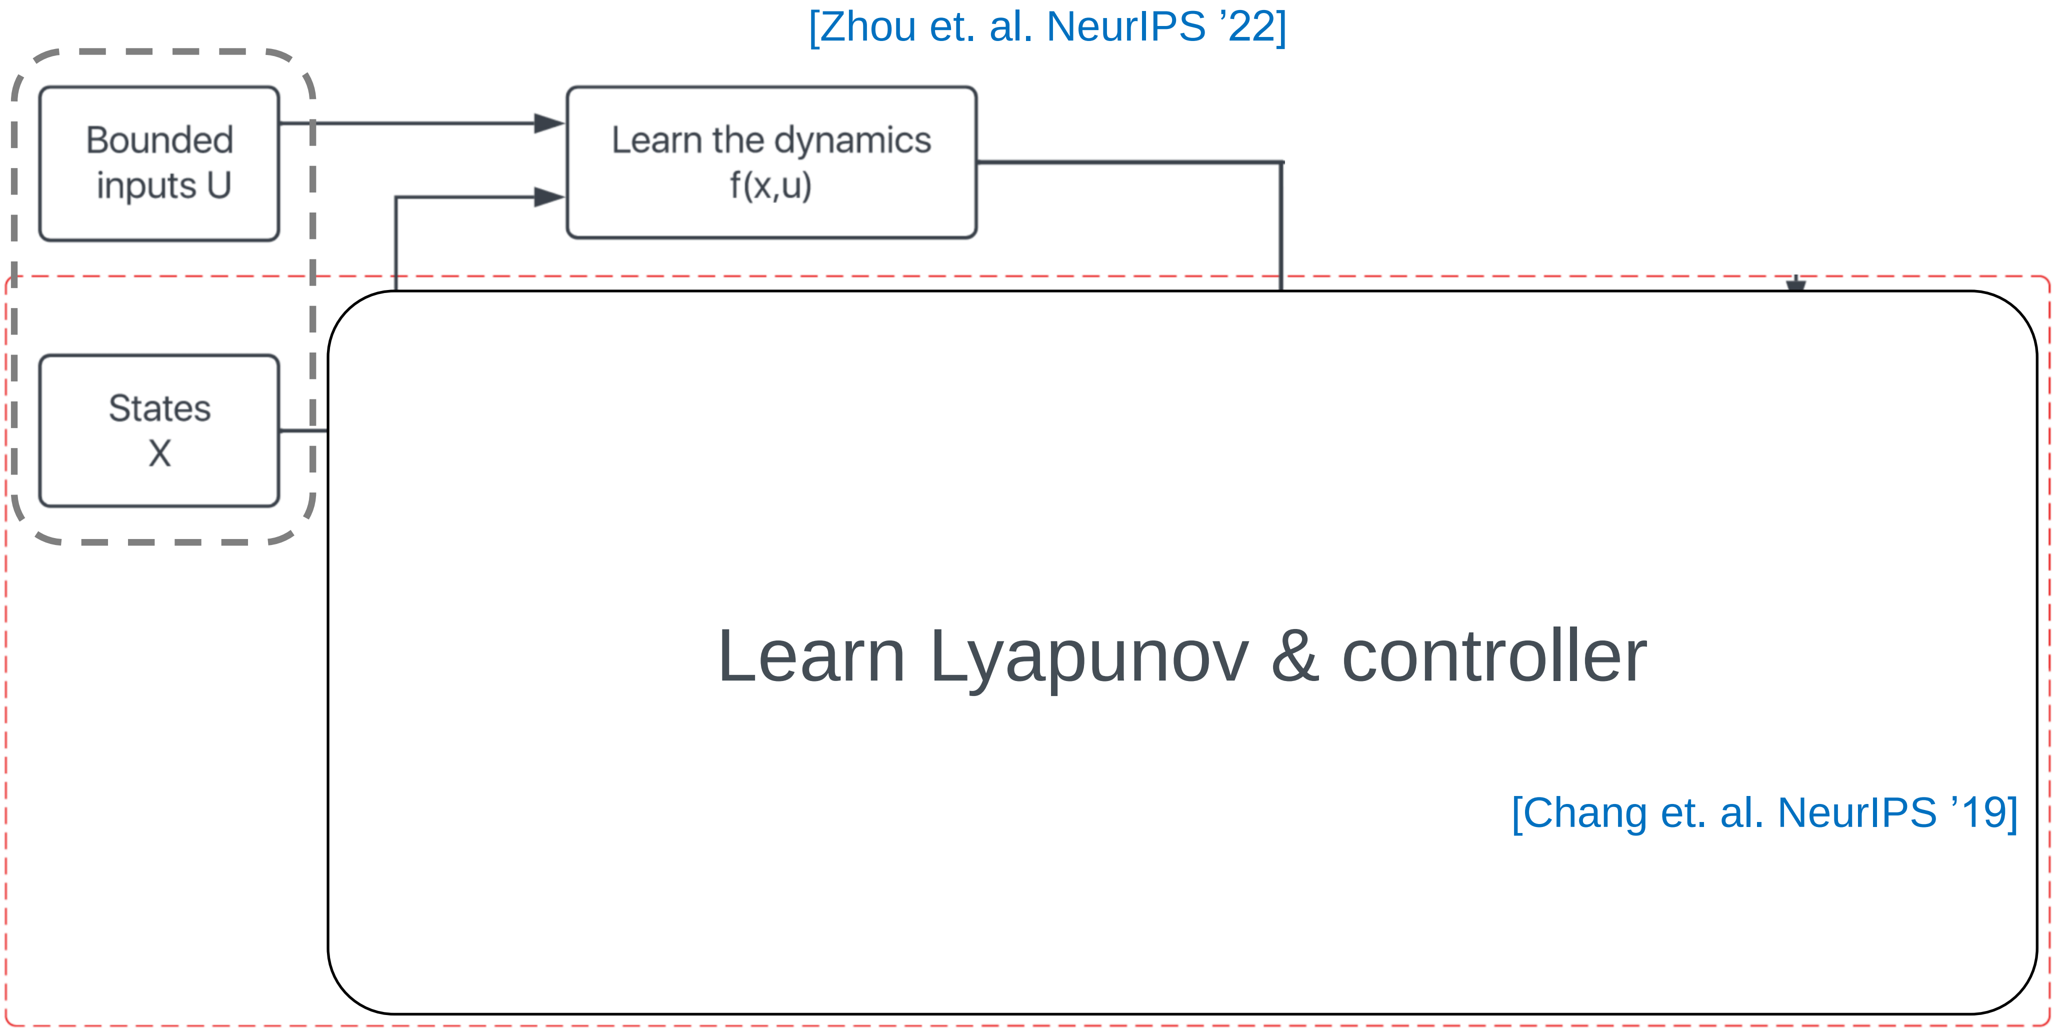

Existing work 2: Fit the dynamics first, then find Lyapunov

[Zhou et. al. NeurIPS ’22]

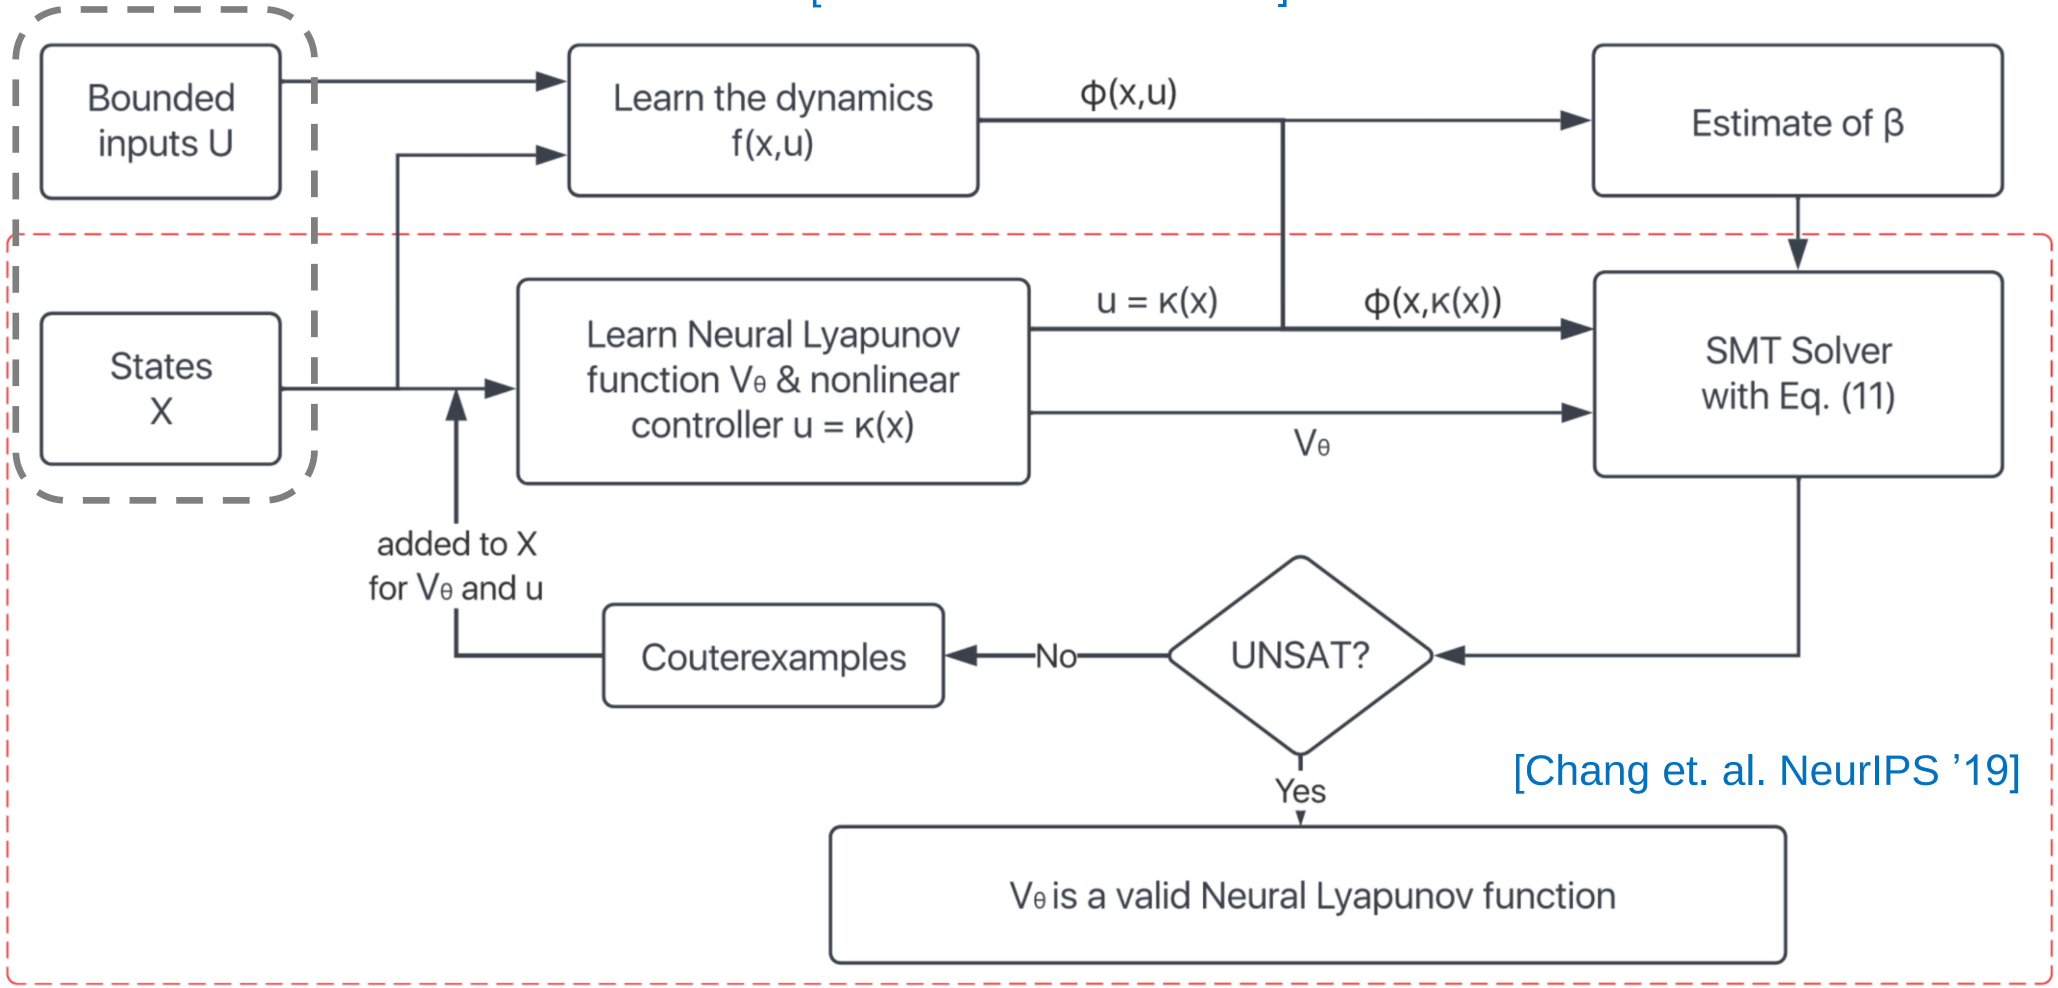

[Chang et. al. NeurIPS ’19]

Existing work 2: Fit the dynamics first, then find Lyapunov

[Zhou et. al. NeurIPS ’22]

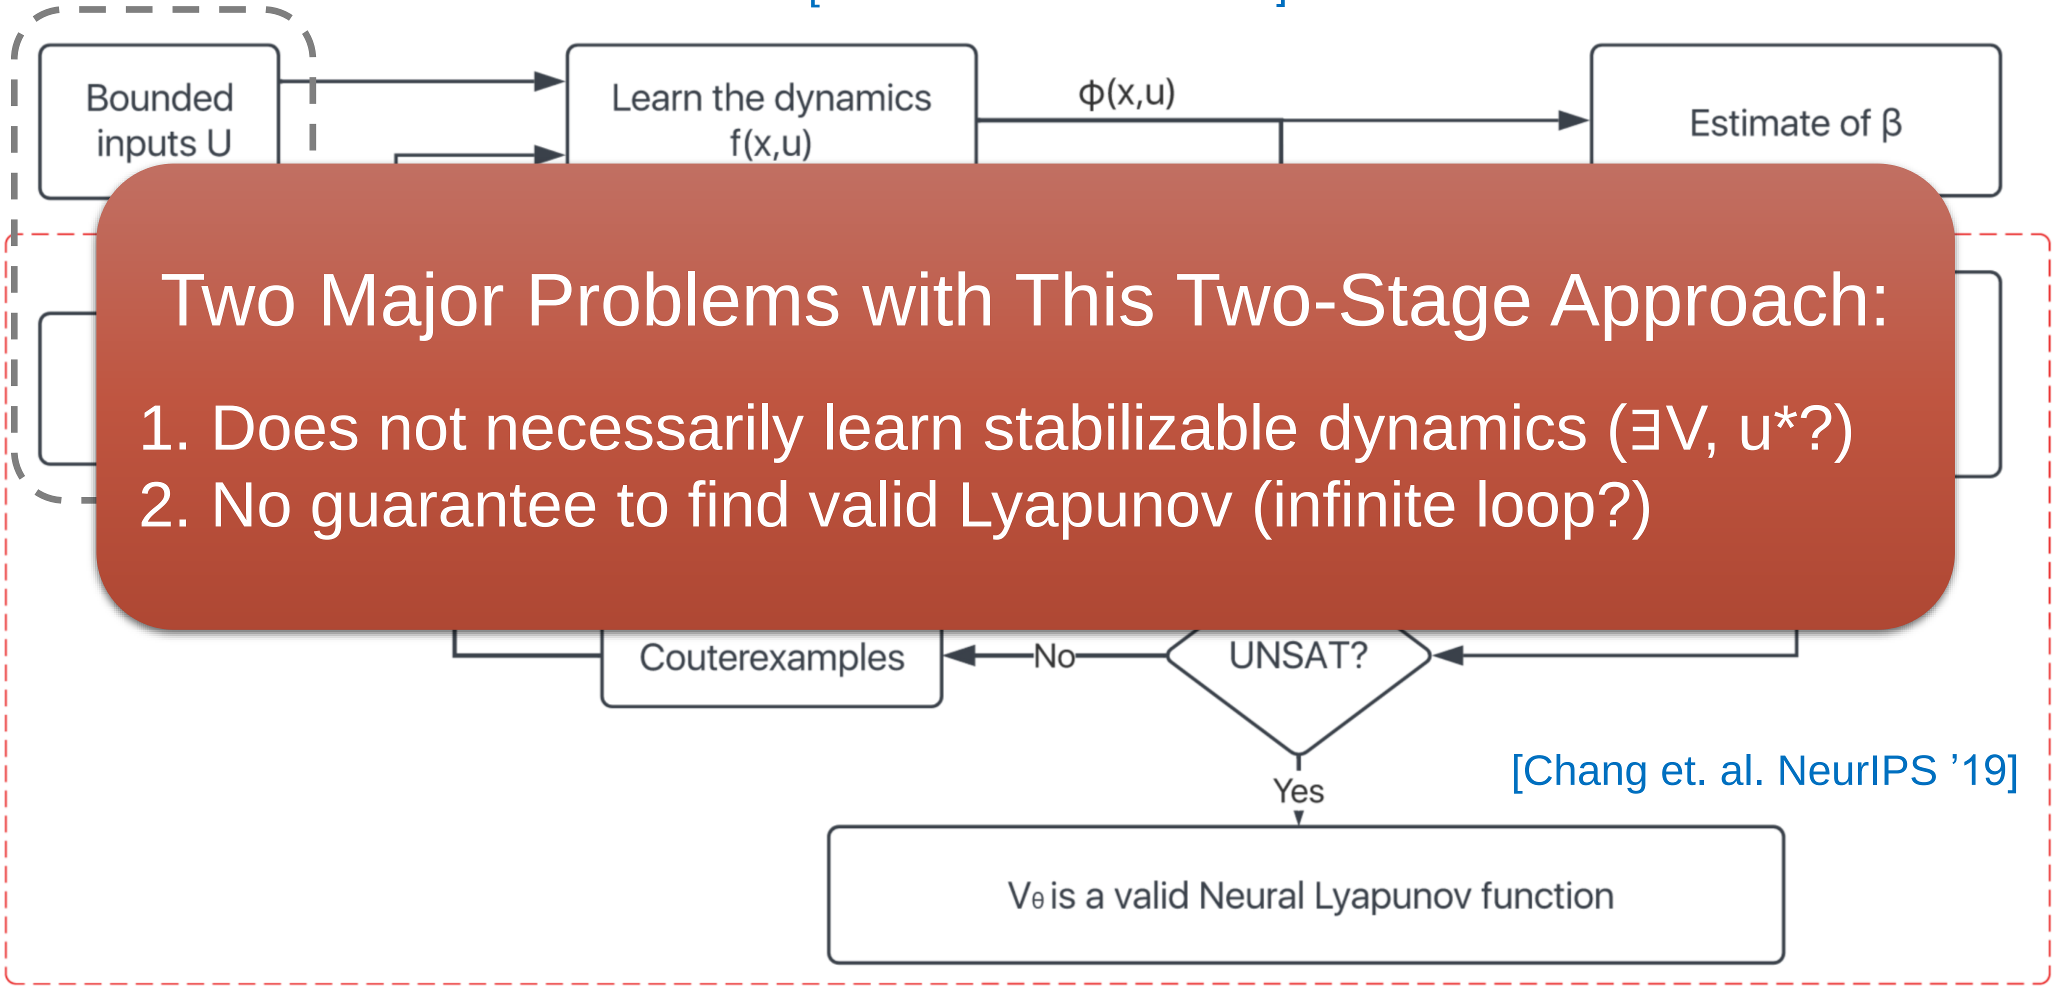

[Chang et. al. NeurIPS ’19]

Two Major Problems with This Two-Stage Approach:

1. Does not necessarily learn stabilizable dynamics (∃V, u*?)
2. No guarantee to find valid Lyapunov (infinite loop?)

This work: Data-Driven Control with Inherent Lyapunov Stability (CoILS)

Idea: Choose One among Certifiable Models

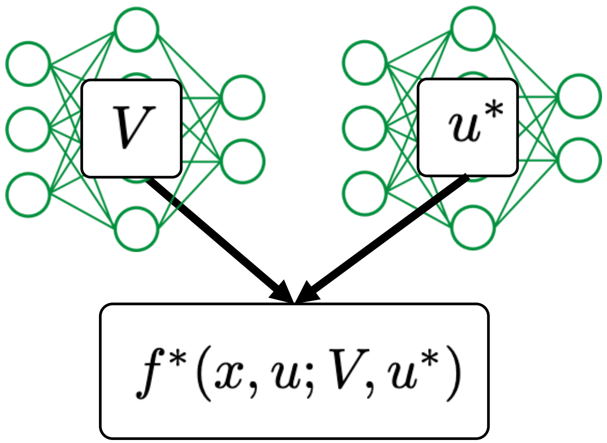

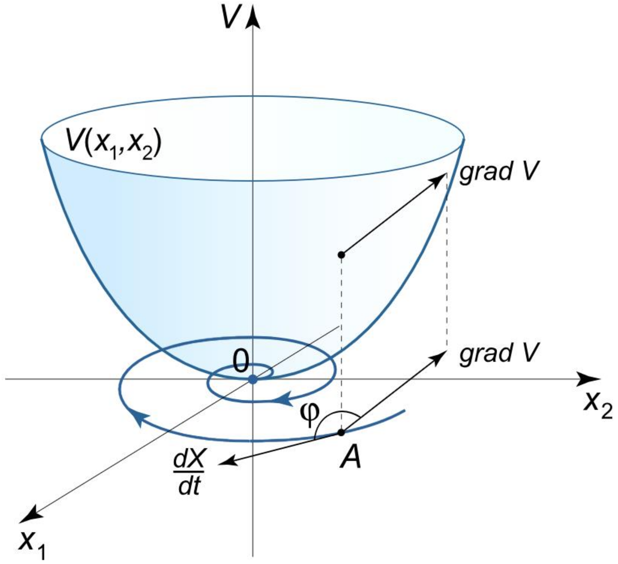

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

Idea: Choose One among Certifiable Models

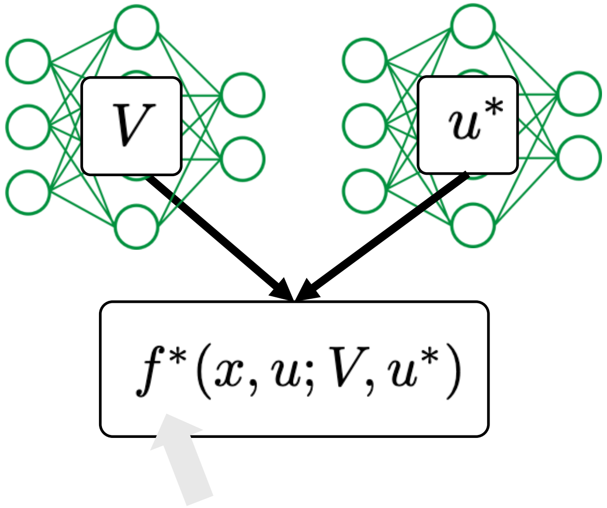

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$$

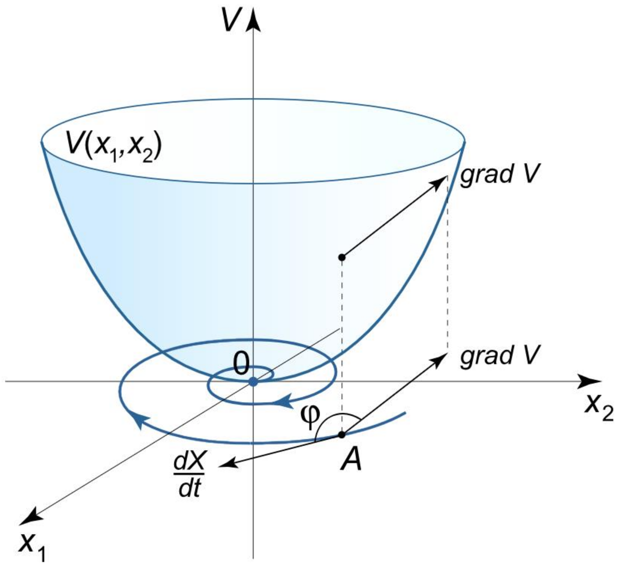

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

Idea: Choose One among Certifiable Models

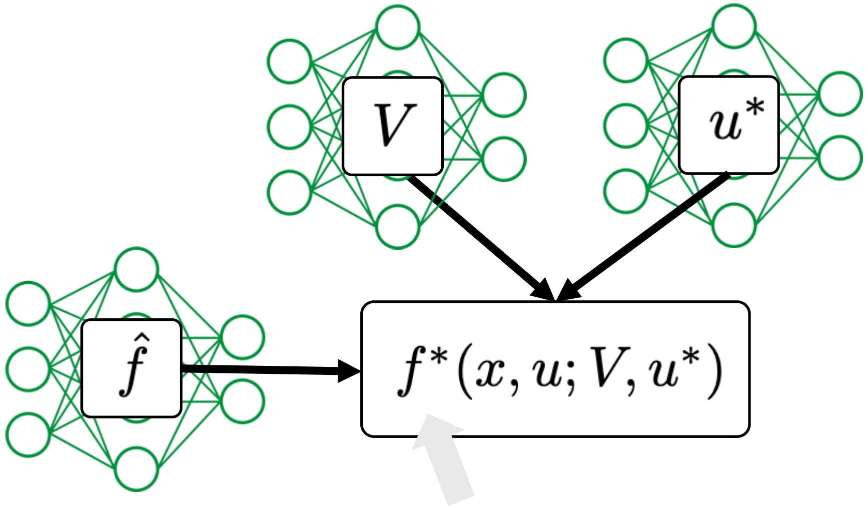

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$$

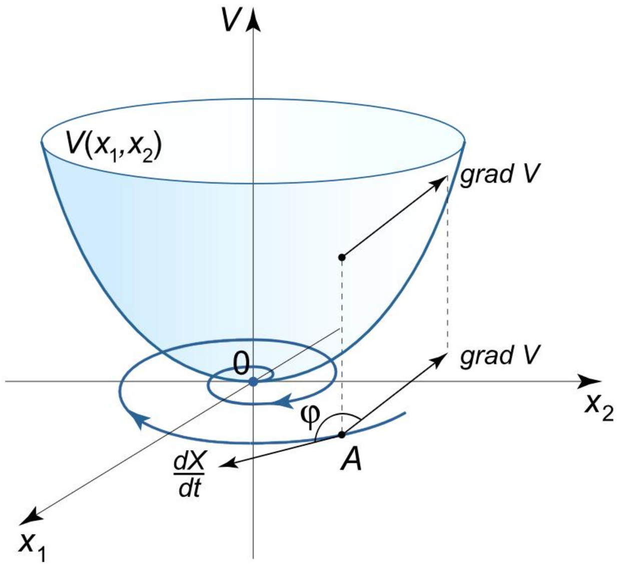

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

Idea: Choose One among Certifiable Models

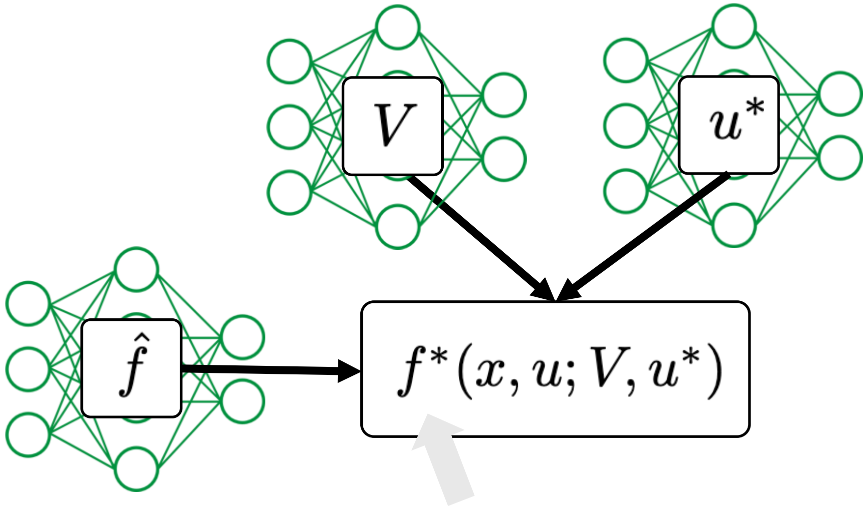

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$$

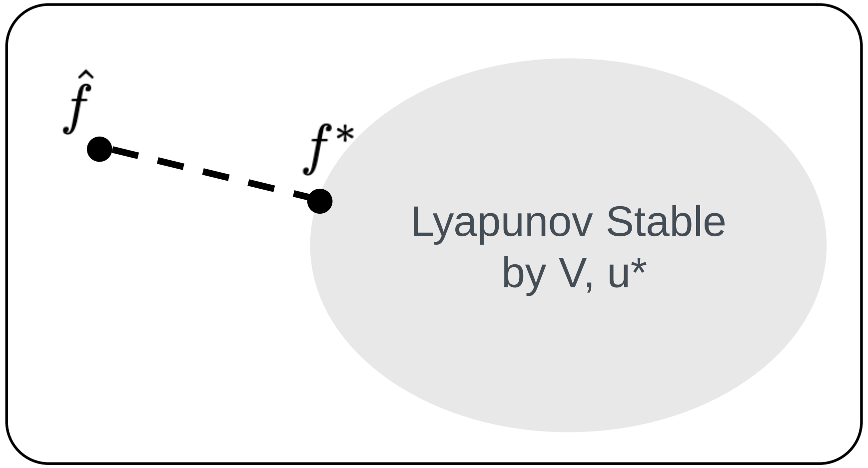

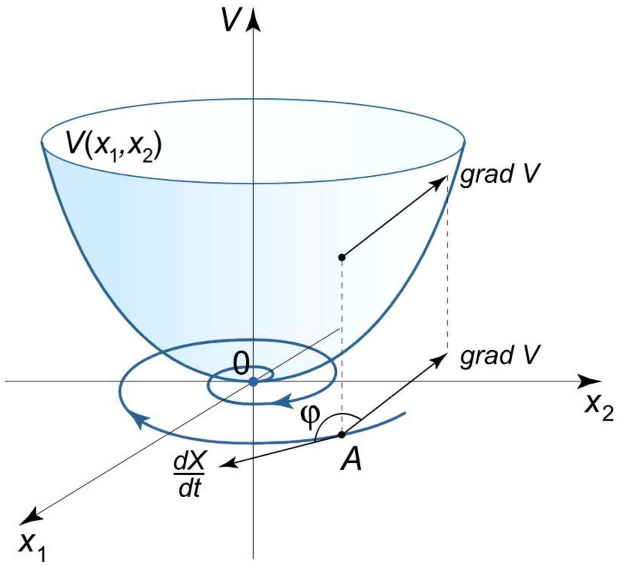

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

Idea: Choose One among Certifiable Models

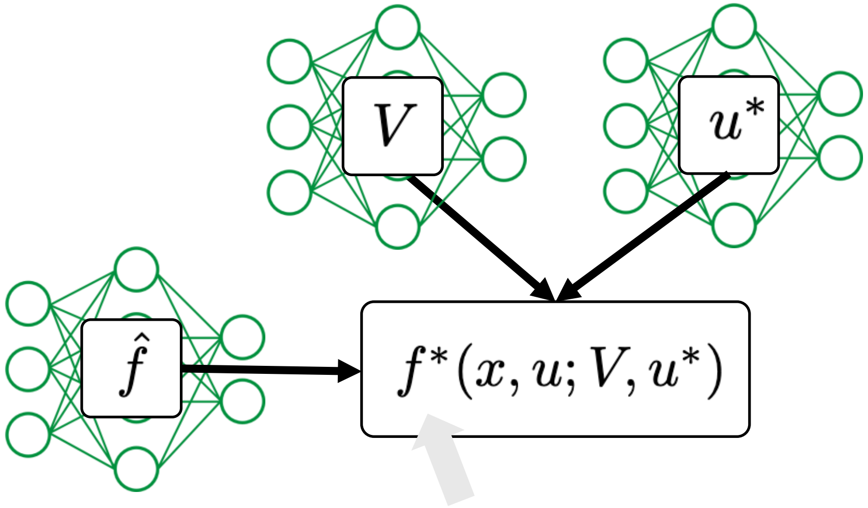

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$$

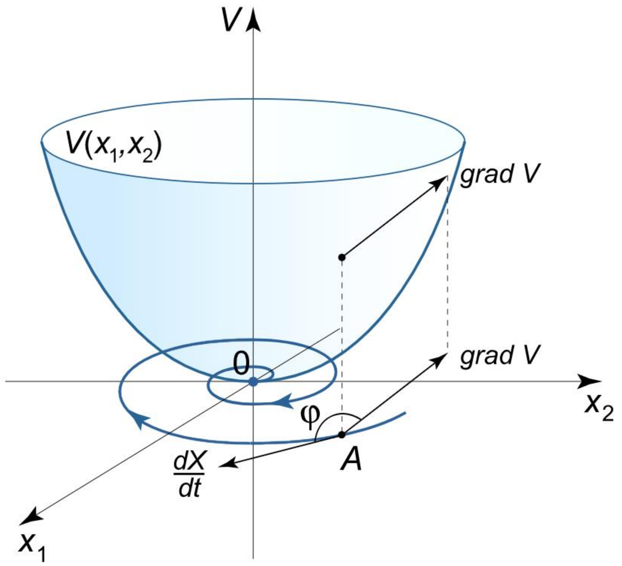

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

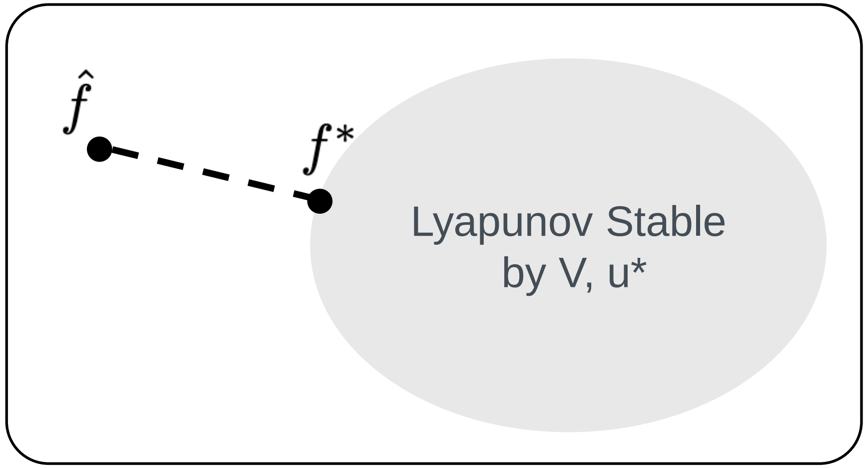

$$f^*(x, \cdot) := \hat{f}(x, \cdot) + \begin{array}{ll} \underset{\Delta f \in \mathbb{R}^n}{\arg\min} & \|\Delta f\|_2 \\ \text{s.t.} & \nabla V(x)^\top \left(\hat{f}(x, u^*(x)) + \Delta f\right) \leq -\alpha V(x) \end{array}$$

Idea: Choose One among Certifiable Models

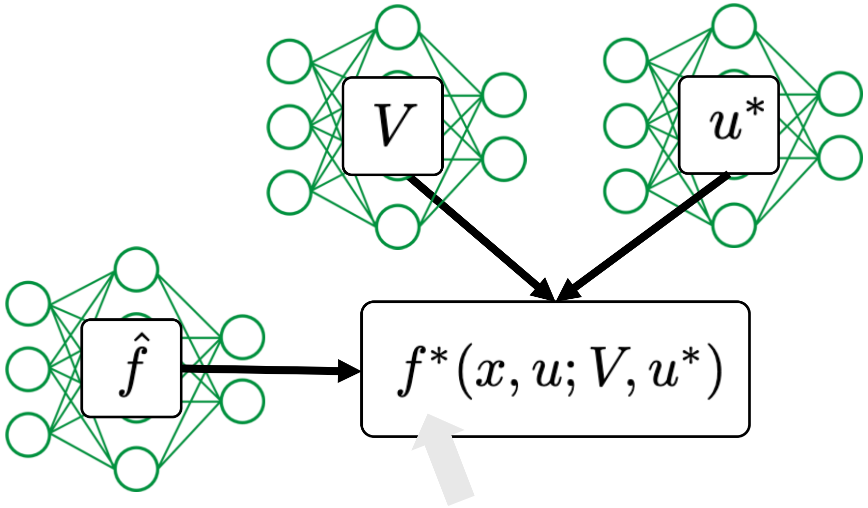

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$$

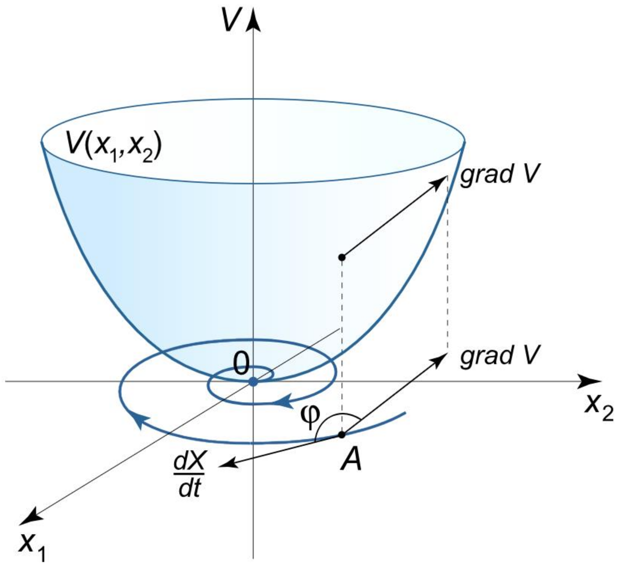

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

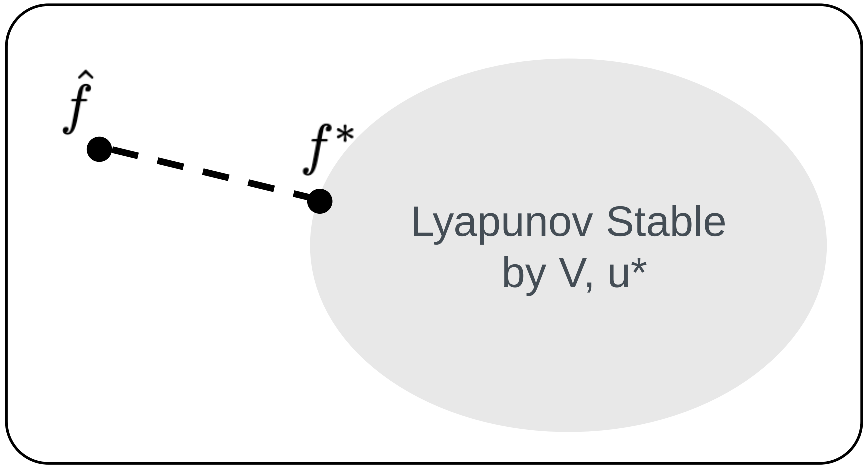

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\text{ReLU}(x) = \max(0, x)$$

Idea: Choose One among Certifiable Models

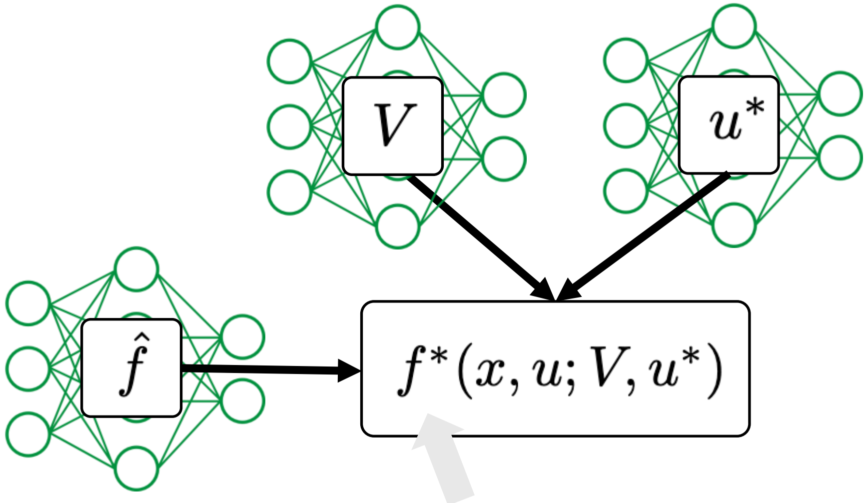

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^{N}$$

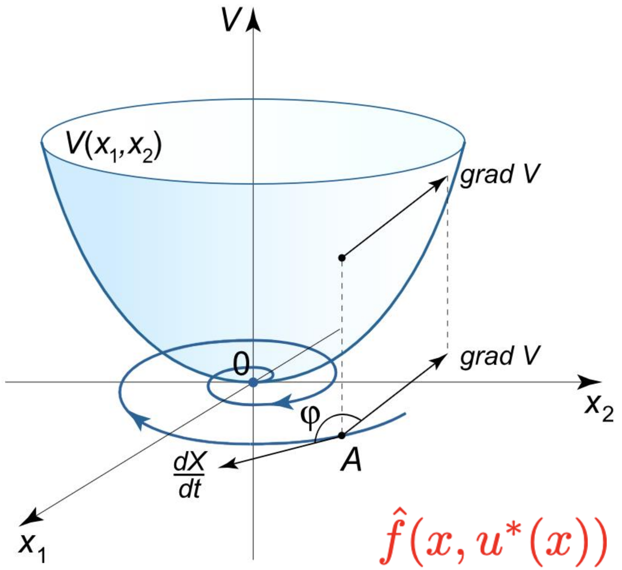

$\hat{f}(x, u^*(x))$

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

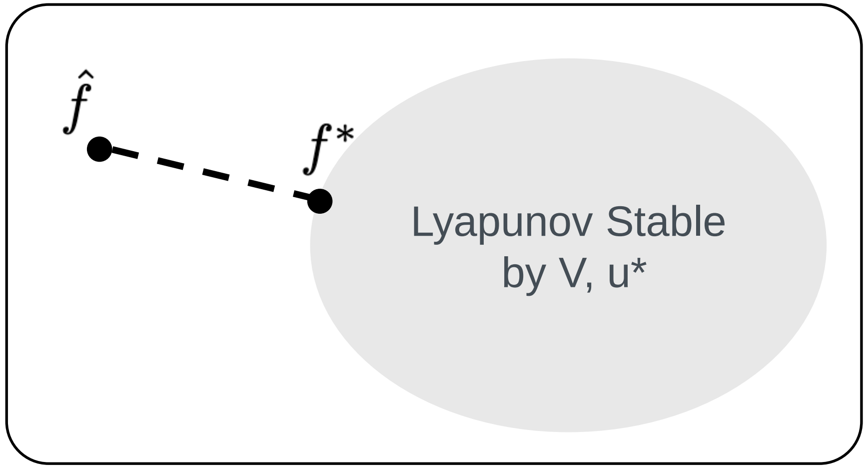

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\text{ReLU}(x) = \max(0, x)$$

Idea: Choose One among Certifiable Models

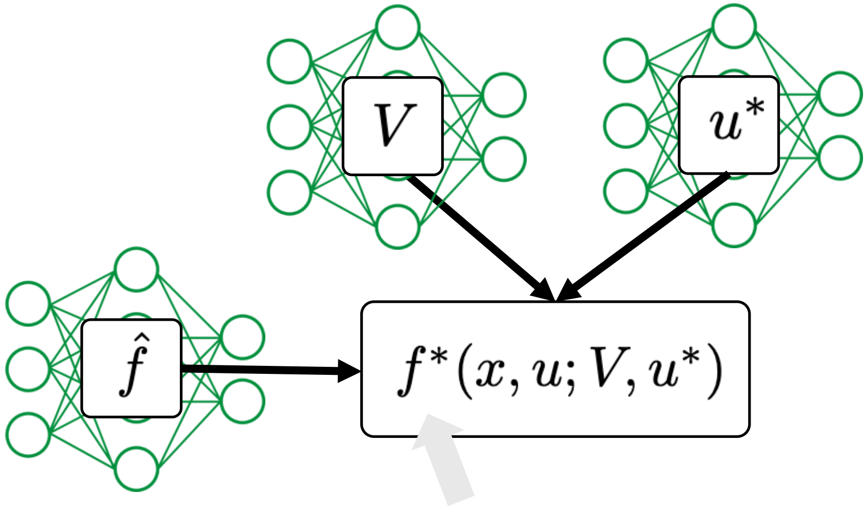

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$$

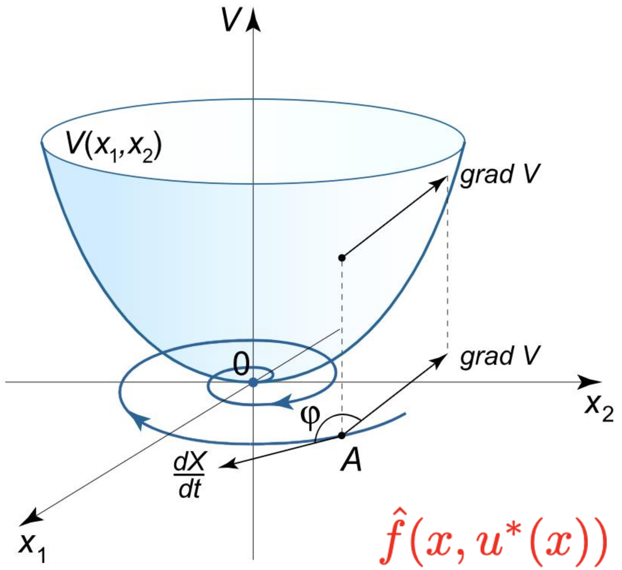

$\hat{f}(x, u^*(x))$

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

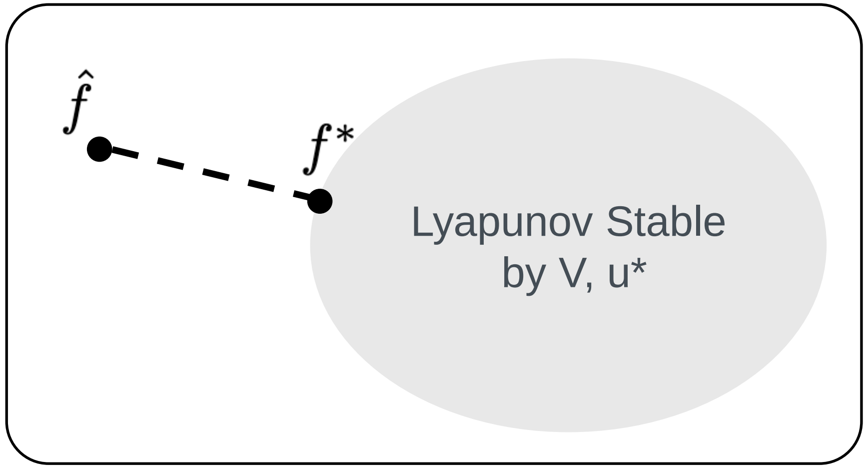

0

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\text{ReLU}(x) = \max(0, x)$$

Idea: Choose One among Certifiable Models

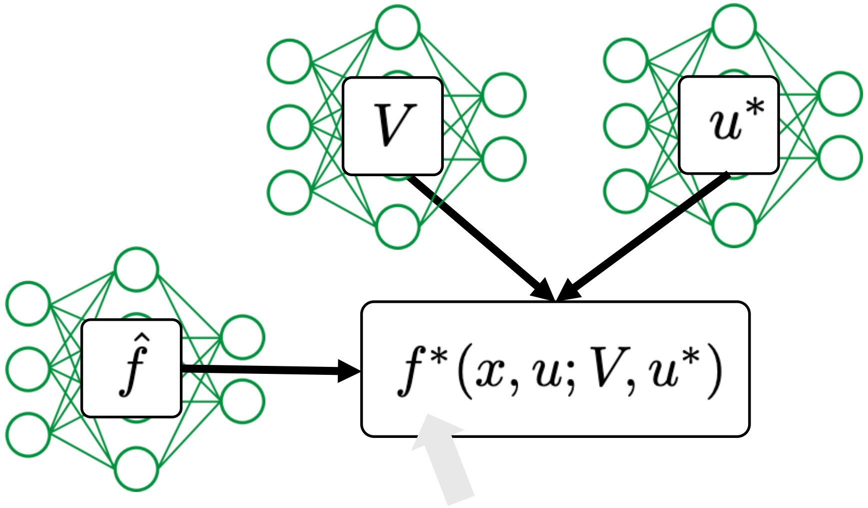

$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^N$

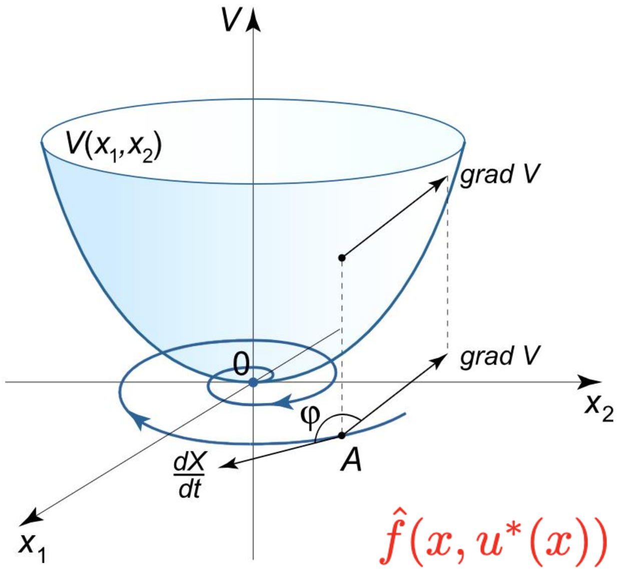

$\hat{f}(x, u^*(x))$

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

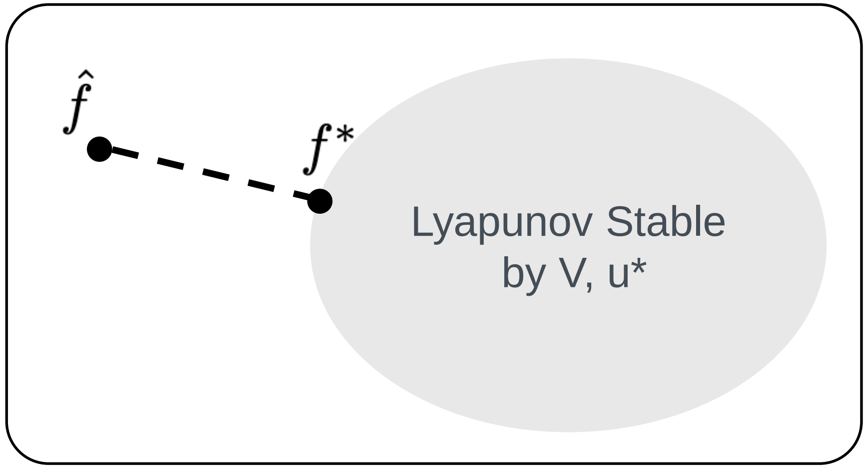

0

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$f^* = \hat{f}$

$\text{ReLU}(x) = \max(0, x)$

Idea: Choose One among Certifiable Models

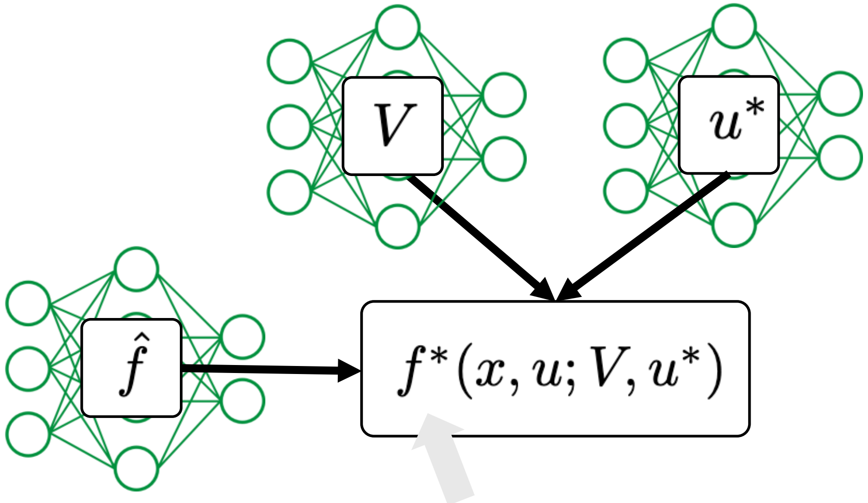

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^{N}$$

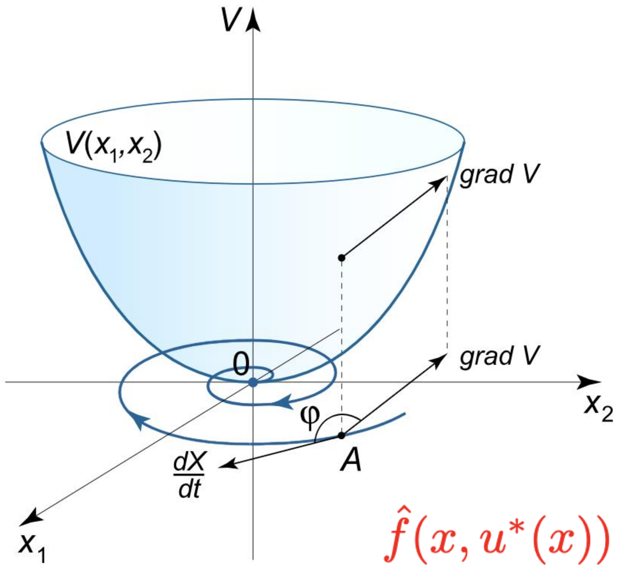

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

$\hat{f}(x, u^*(x)) >$

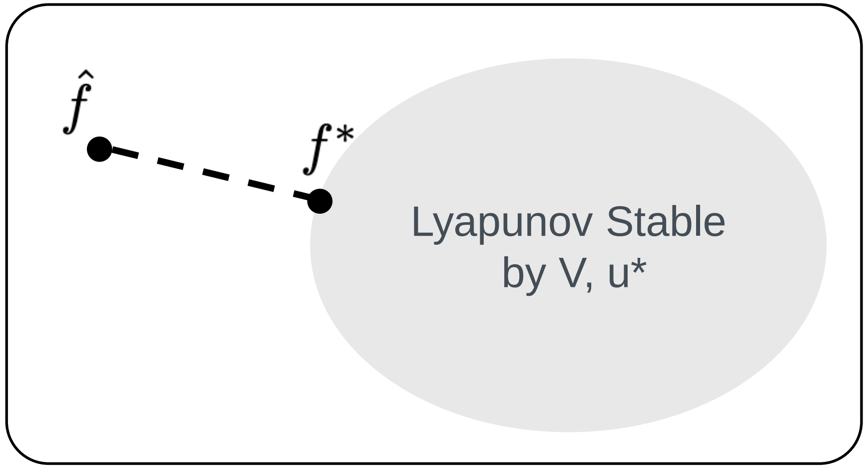

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\text{ReLU}(x) = \max(0, x)$$

Idea: Choose One among Certifiable Models

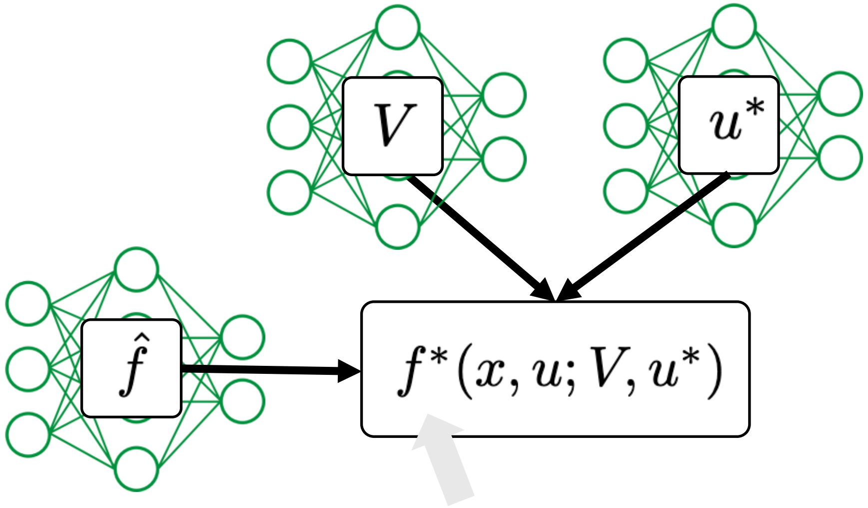

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^{N}$$

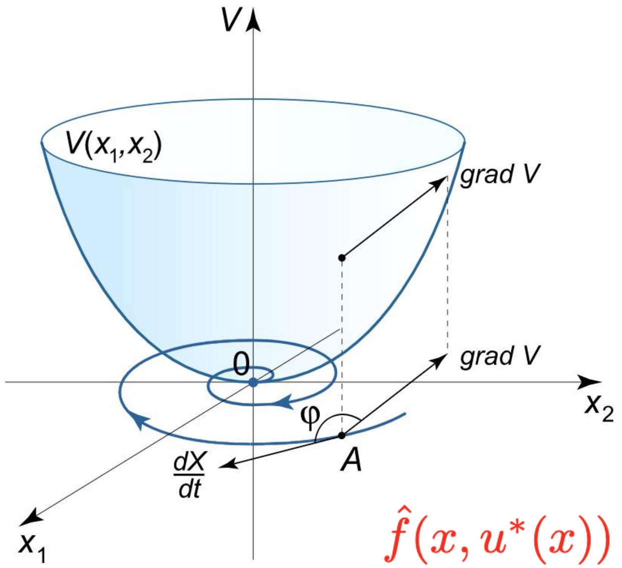

$\hat{f}(x, u^*(x)) >$

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

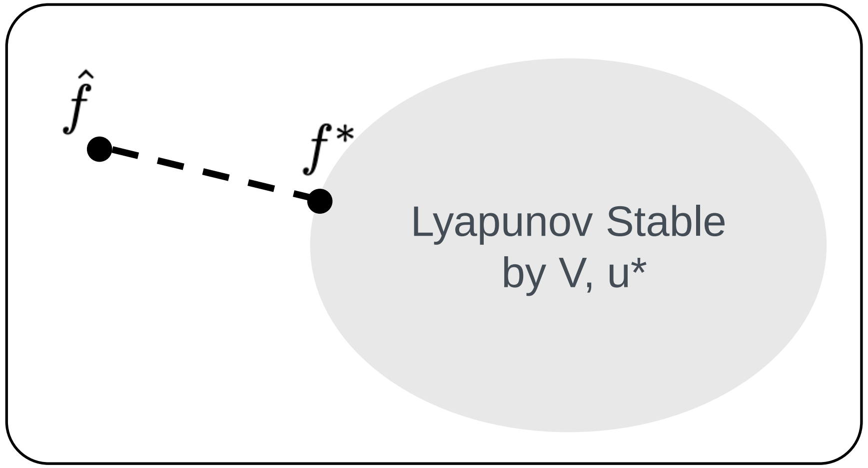

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\text{ReLU}(x) = \max(0, x)$$

Idea: Choose One among Certifiable Models

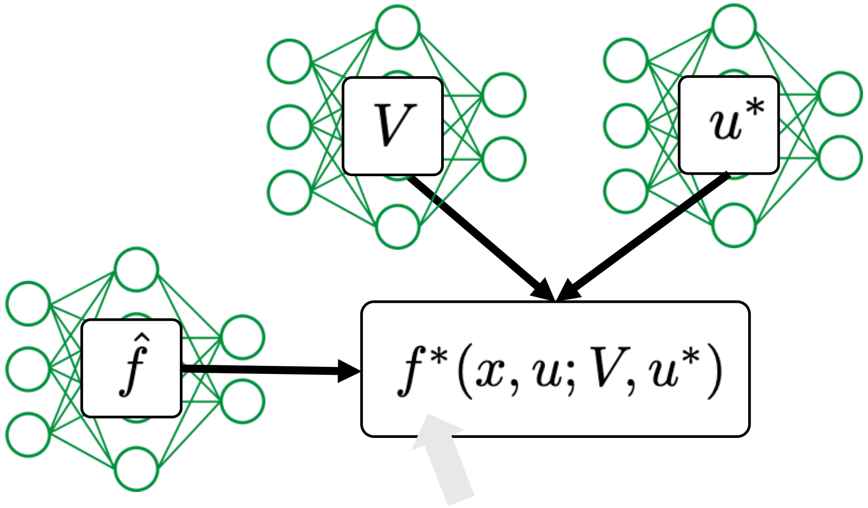

$$\mathcal{D} = \{(x_i, u_i, \dot{x}_i)\}_{i=1}^{N}$$

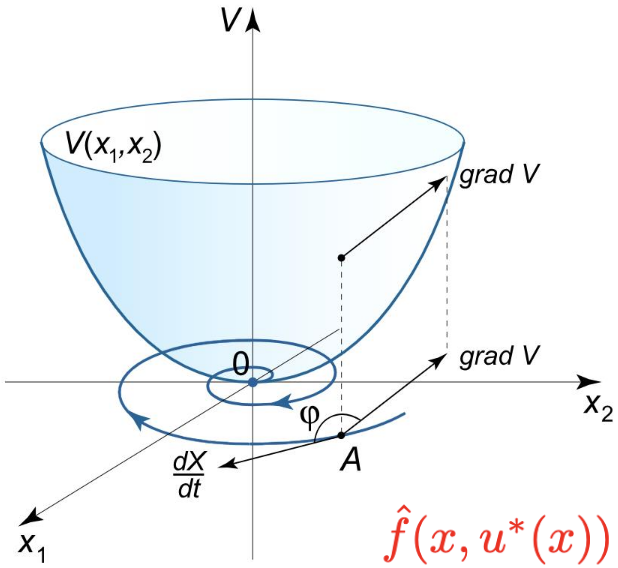

$\hat{f}(x, u^*(x)) >$

Stability Condition: $\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$

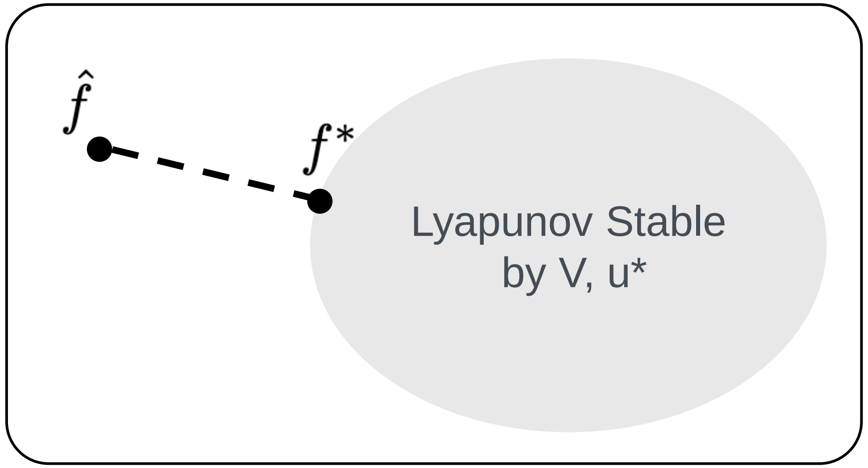

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x, u^*(x)) = -\alpha V(x)$$

$$\text{ReLU}(x) = \max(0, x)$$

Model Architecture

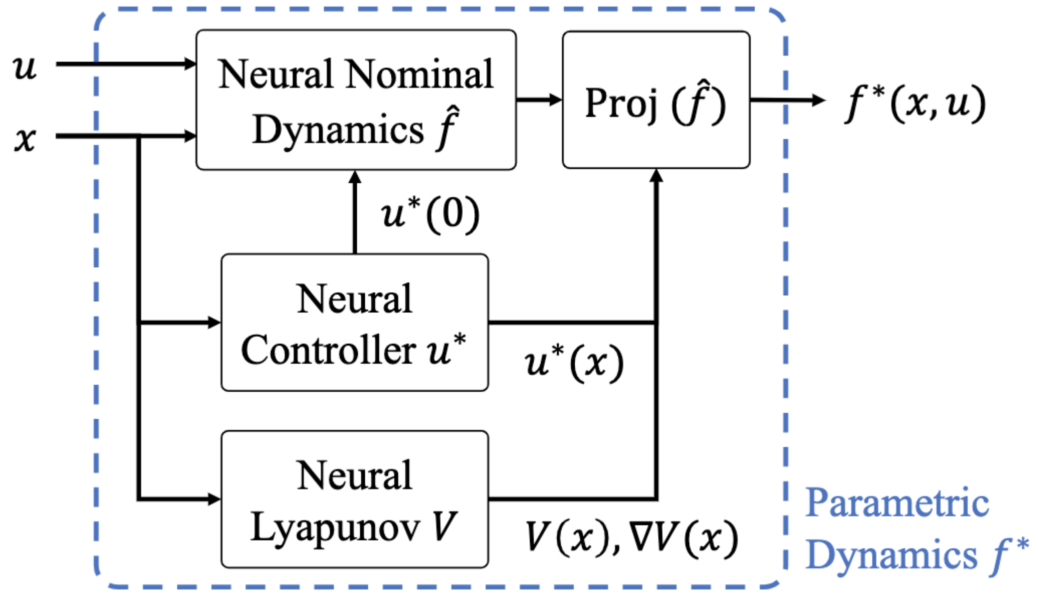

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$$

Model Architecture

- Nominal model: $\hat{f}(x, u) = g_f(x, u) - g_f(0, u^*(0))$
  - Origin as equilibrium point

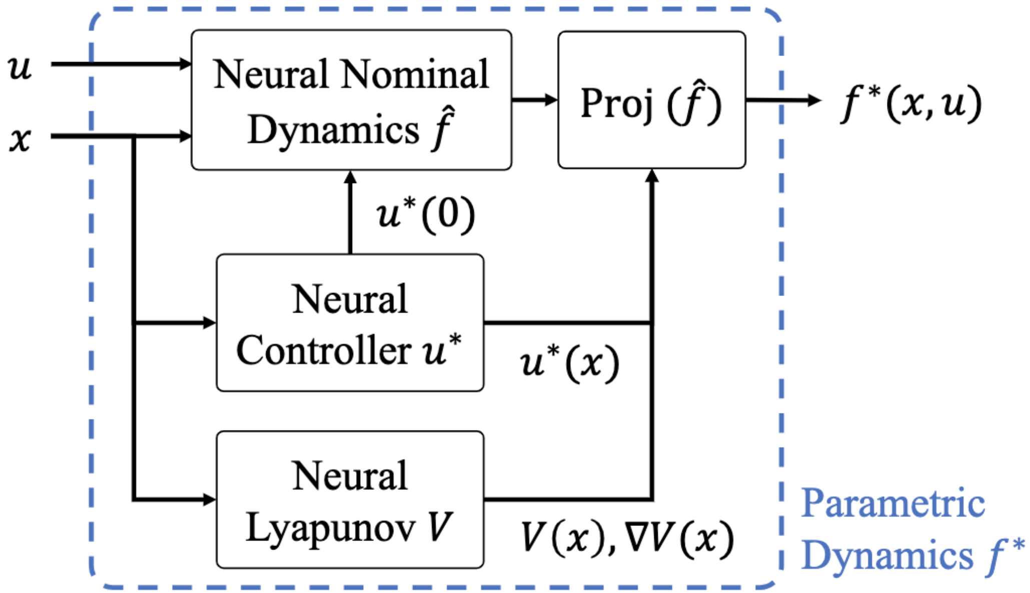

$$f^*(x, u) = \hat{f}(x, u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$$

Model Architecture

- Nominal model: $\hat{f}(x,u) = g_f(x,u) - g_f(0,u^*(0))$
  - Origin as equilibrium point
- Lyapunov function: $V(x) = \sigma(g_V(x) - g_V(0)) + \epsilon_{\text{pd}}\|x\|^2$
  - Positive definite
  - Continuously differentiable

σ : Smoothed ReLU

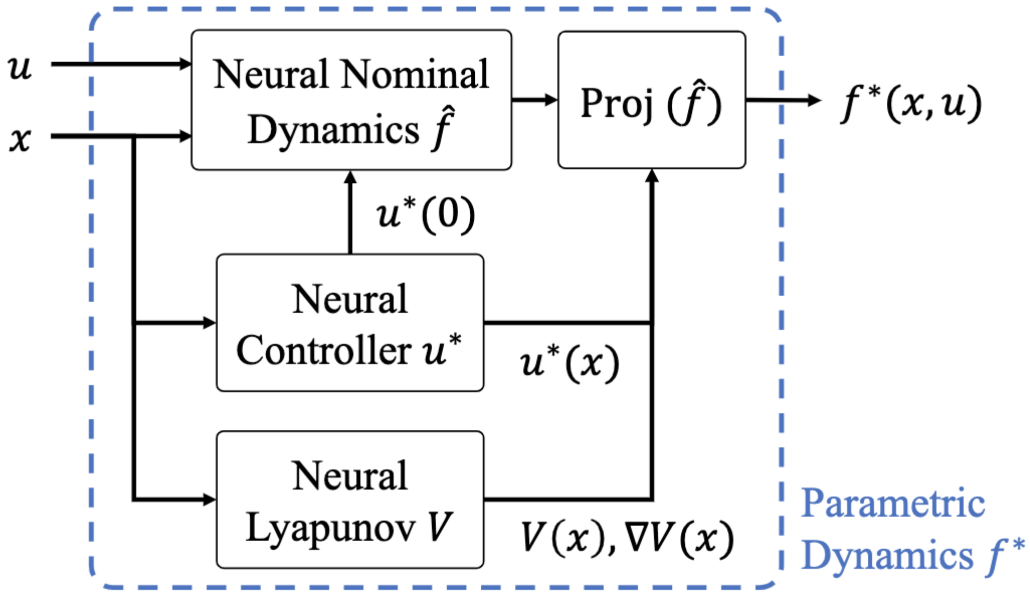

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x)\frac{\text{ReLU}(\nabla V(x)^\top \hat{f}(x,u^*(x)) + \alpha V(x))}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x,u^*(x)) \leq -\alpha V(x)$$

Model Architecture

- Nominal model: $\hat{f}(x,u) = g_f(x,u) - g_f(0, u^*(0))$
  - Origin as equilibrium point
- Lyapunov function: $V(x) = \sigma(g_V(x) - g_V(0)) + \epsilon_{\text{pd}}\|x\|^2$
  - Positive definite
  - Continuously differentiable

  σ : Smoothed ReLU
- Controller: $u^*(x) = \textbf{diag}(u_{\text{lim}})\tanh(g_u(x))$
  - Actuation limit

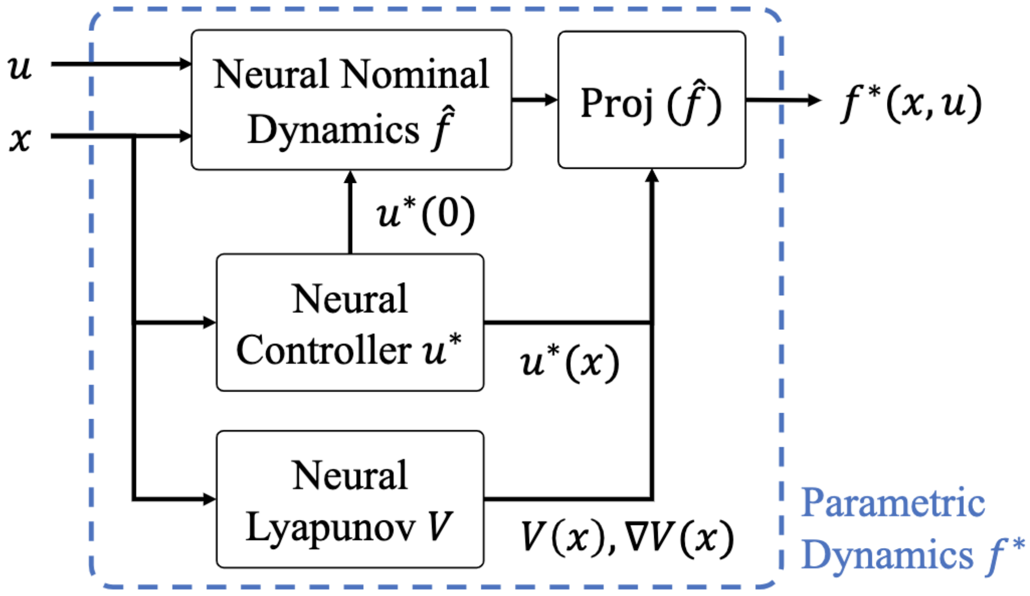

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x)\frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x, u^*(x)) \le -\alpha V(x)$$

Theoretical Guarantees on Stability

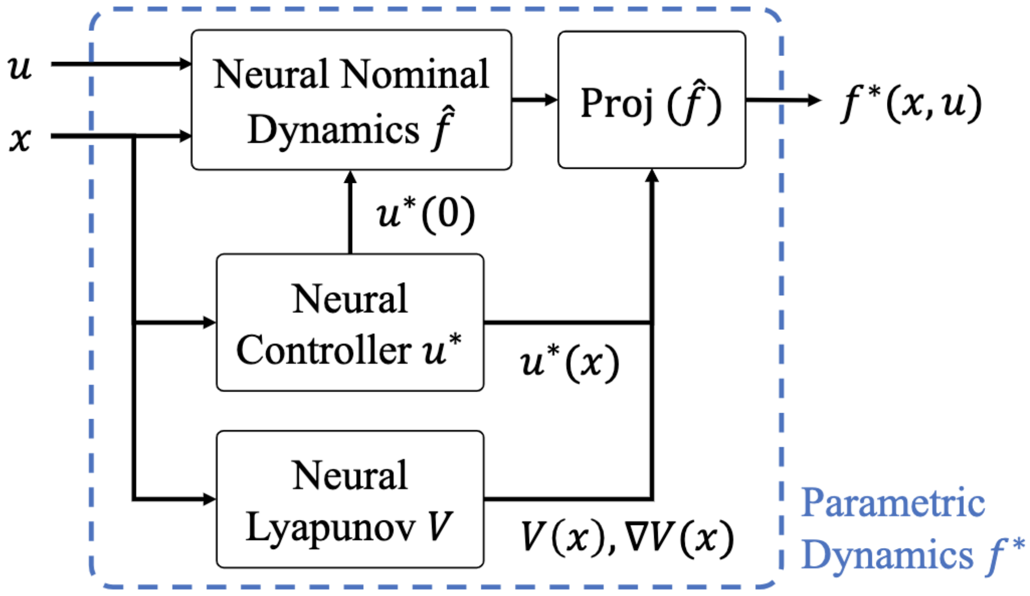

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$$

Theoretical Guarantees on Stability

Theorem 1 (informal): Learned controller u* exponentially stabilizes learned model f*.

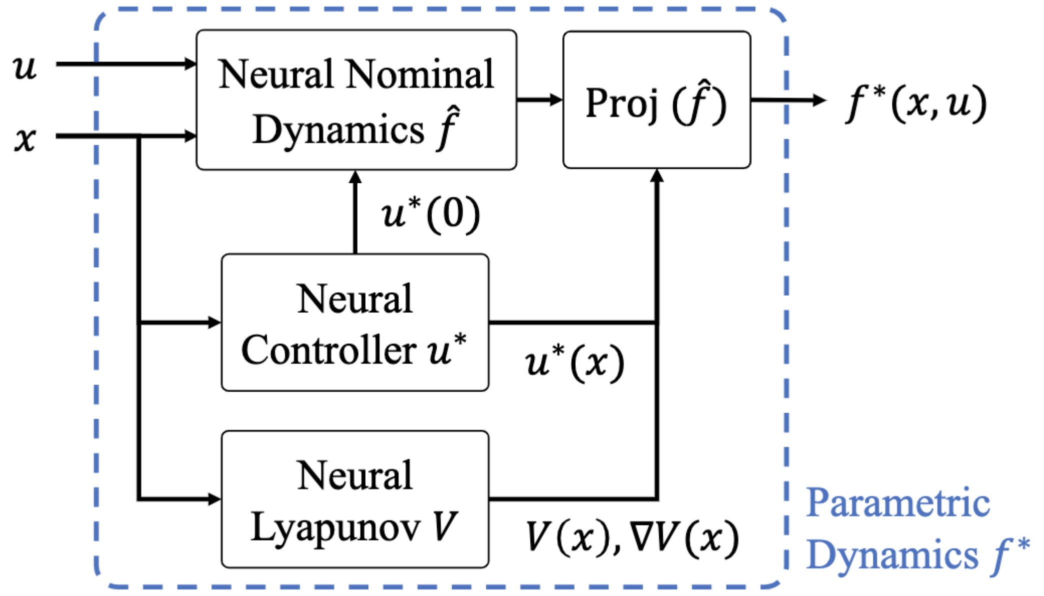

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\mathrm{ReLU}\left(\nabla V(x)^\top \hat{f}(x,u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x,u^*(x)) \leq -\alpha V(x)$$

Theoretical Guarantees on Stability

Theorem 1 (informal): Learned controller u* exponentially stabilizes learned model f*.

Theorem 2 (informal): For the true dynamics f, any trajectory generated by the learned controller u* reaches a small neighborhood of the origin if

1) Data is sampled with sufficient density and
2) The training error is small

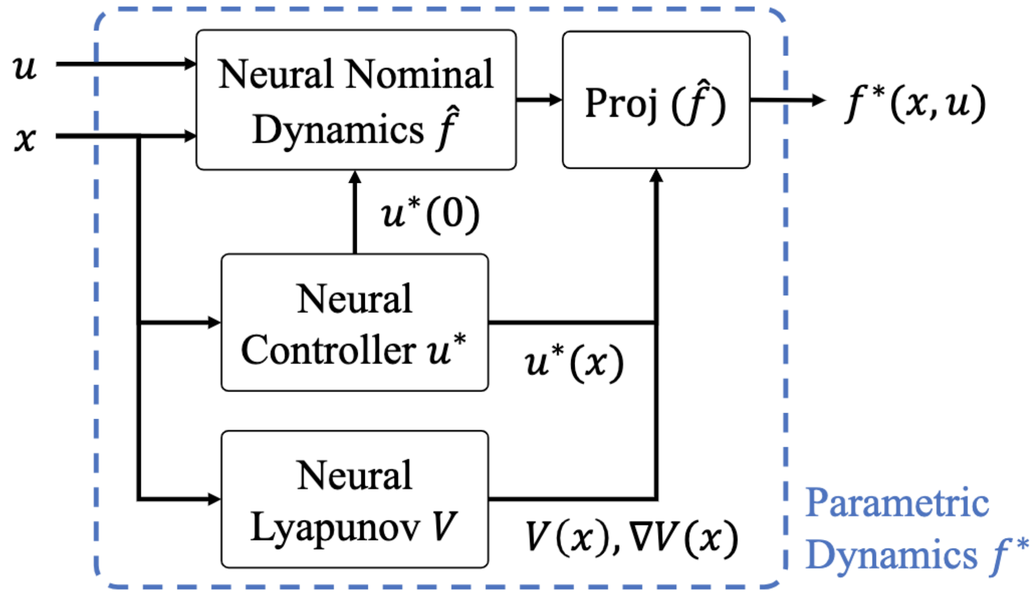

$$f^*(x,u) = \hat{f}(x,u) - \nabla V(x) \frac{\text{ReLU}\left(\nabla V(x)^\top \hat{f}(x, u^*(x)) + \alpha V(x)\right)}{\|\nabla V(x)\|_2^2}$$

$$\nabla V(x)^\top f^*(x, u^*(x)) \leq -\alpha V(x)$$

Experiment 1: Validate Construction with Random Networks

- Stability condition satisfied w/ random initialization

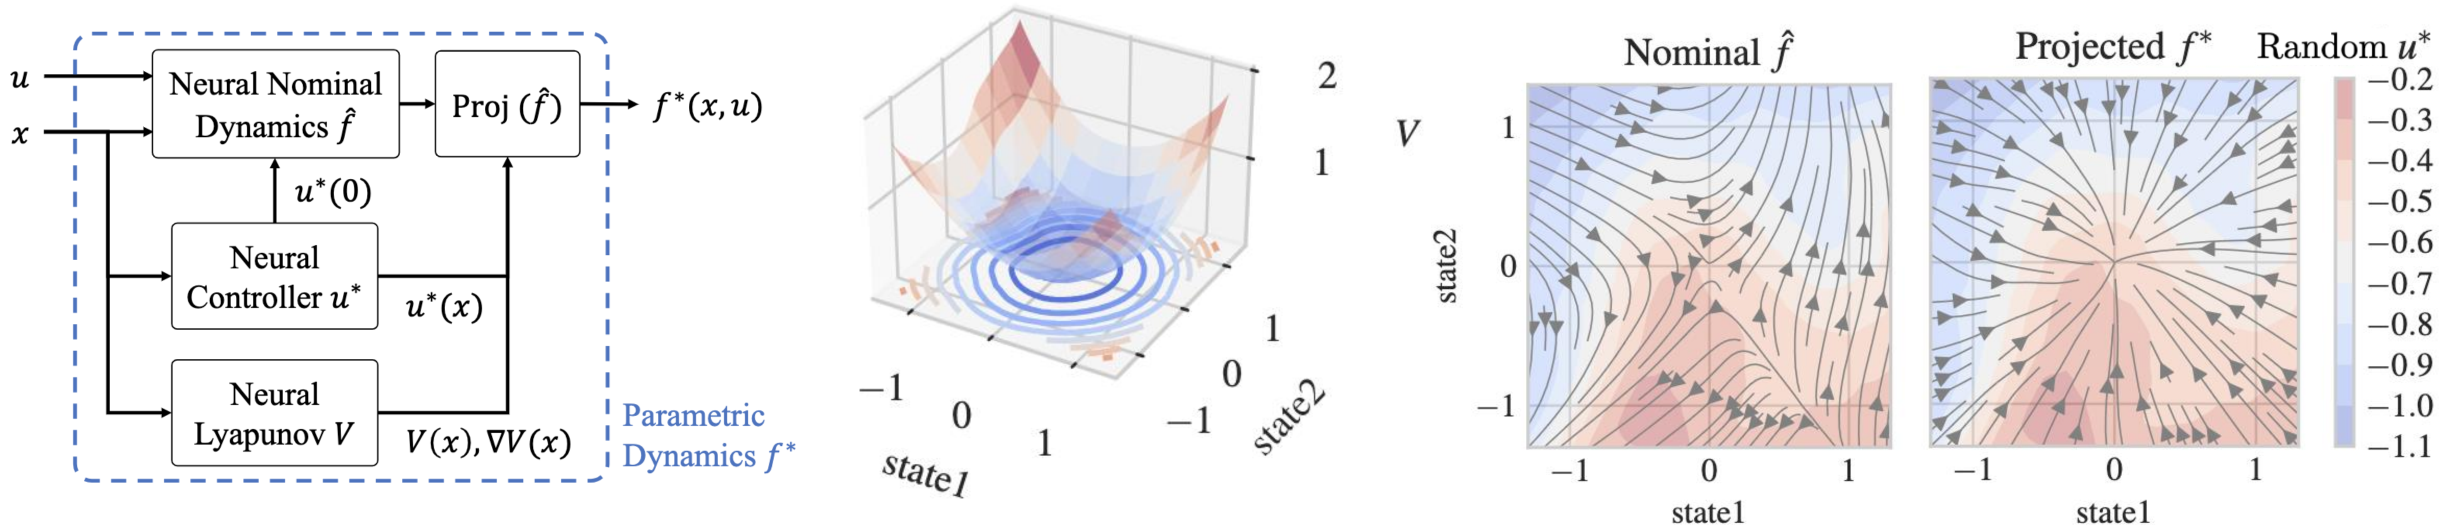

Experiment 2: Van der Pol Oscillator

- True model: $\ddot{z} = u - z + \mu(1 - z^2)\dot{z}$
- State: $x = [z, \dot{z}] \in \mathbb{R}^2$
- Control input: $u \in [-5, 5]$

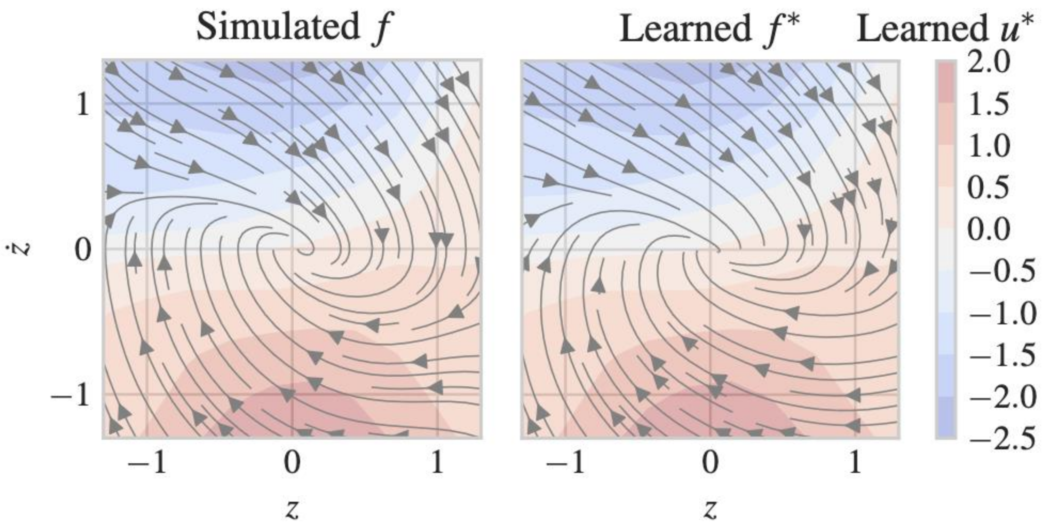

Experiment 2: Van der Pol Oscillator

- True model: $\ddot{z} = u - z + \mu(1 - z^2)\dot{z}$
- State: $x = [z, \dot{z}] \in \mathbb{R}^2$
- Control input: $u \in [-5, 5]$

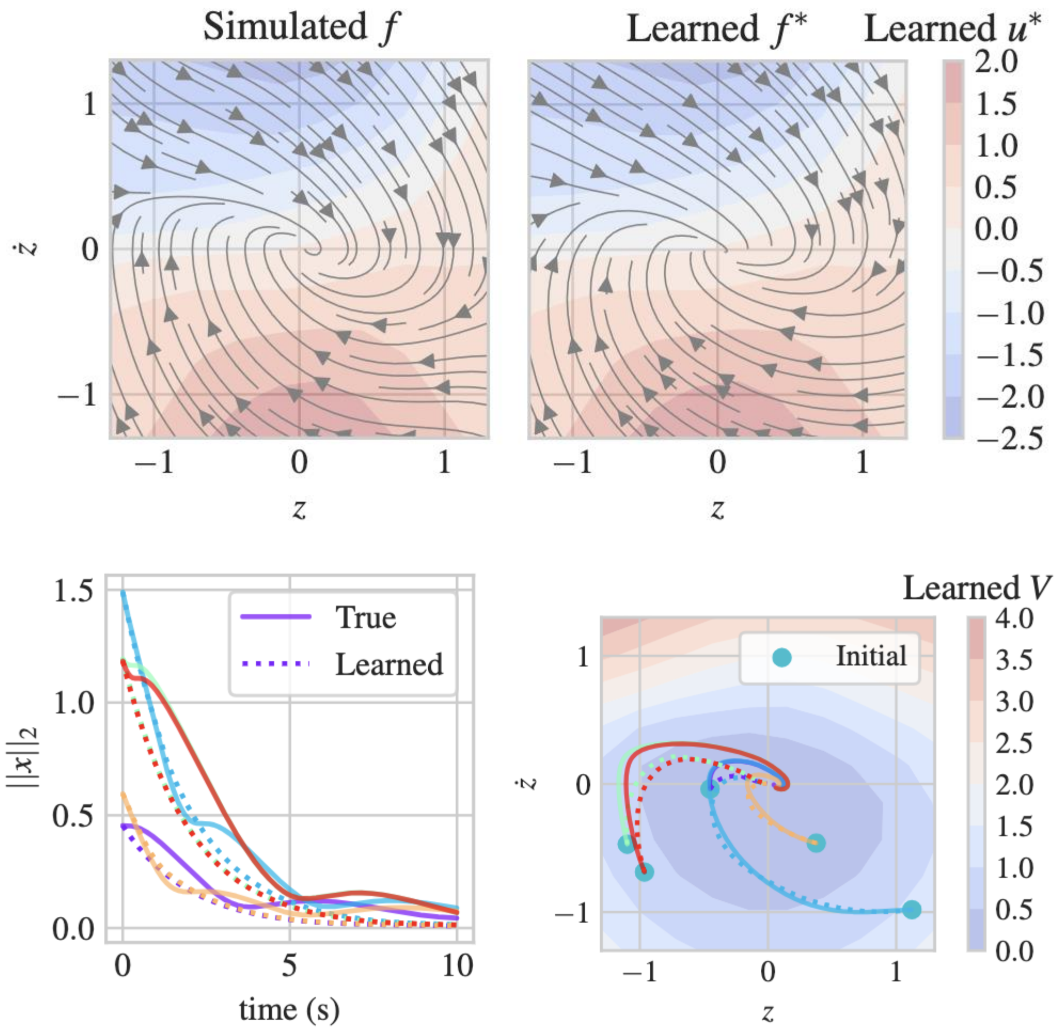

Experiment 3: Bicycle Path Following

- True model: $\dot{d}_e = v \sin\theta_e, \quad \dot{\theta}_e = \frac{v \tan u}{L} - \frac{v \cos\theta_e}{1 - d_e}$
- State: $x = [d_e, \theta_e] \in \mathbb{R}^2$
- Control input: $u \in [-0.4\pi, 0.4\pi]$

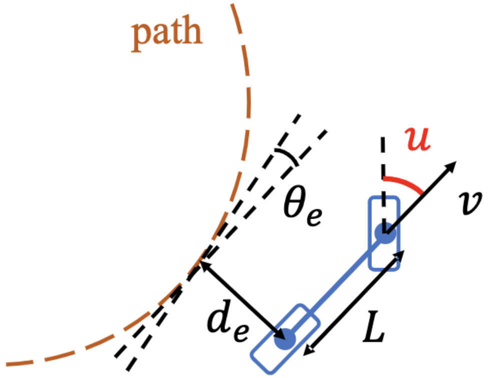

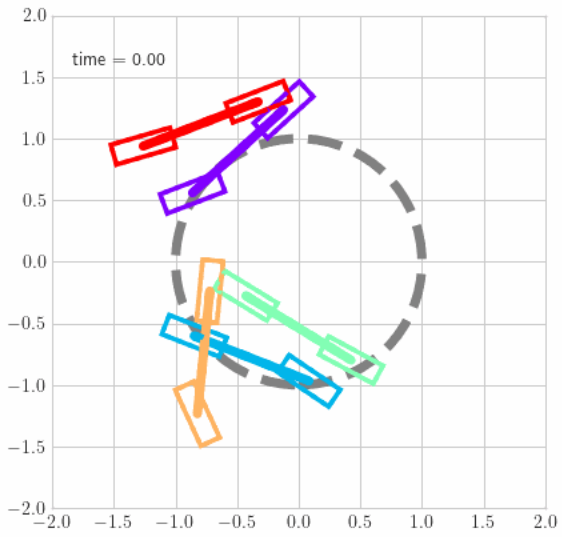

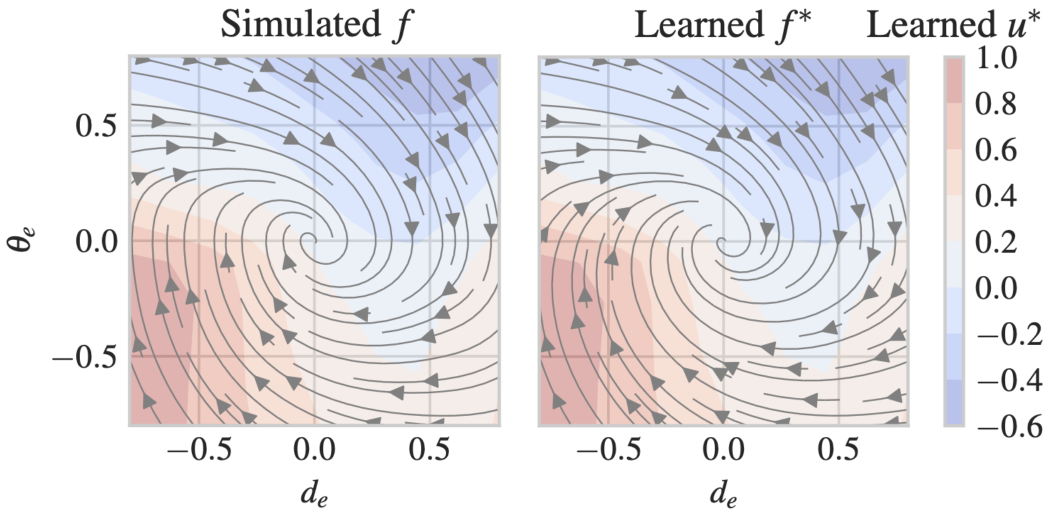

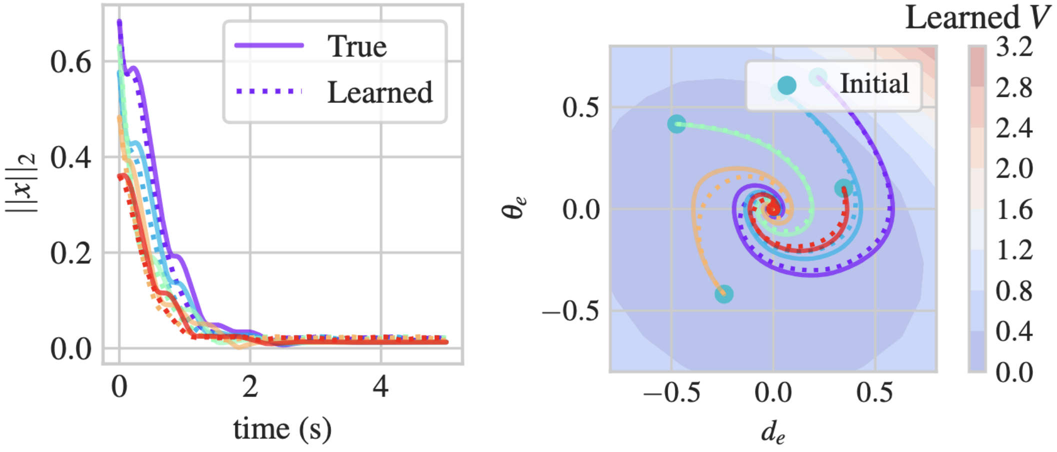

Interesting observation:
The constrained model still seems very expressive!

Interesting observation:
The constrained model still seems very expressive!

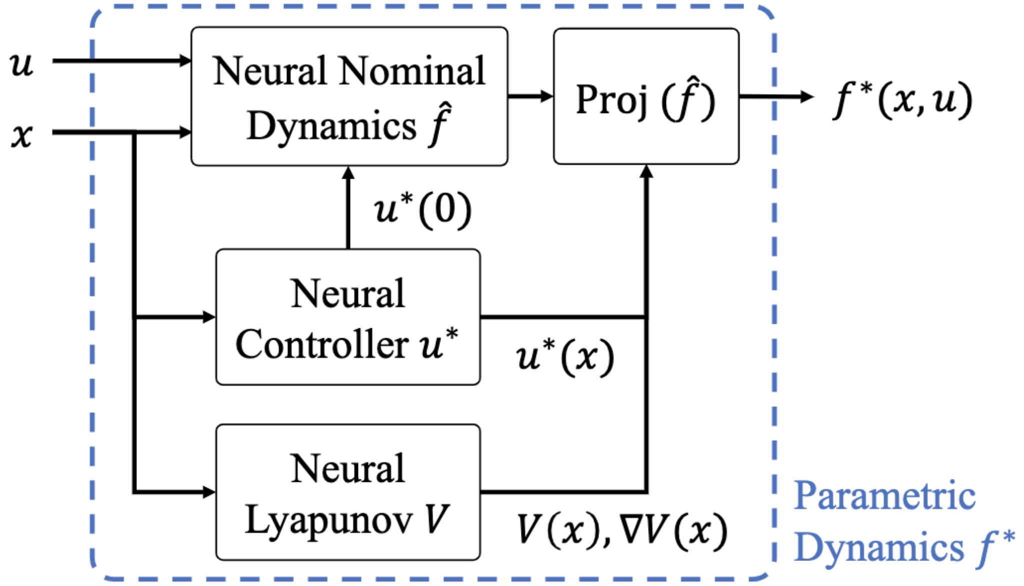

Interesting observation: The constrained model still seems very expressive!

Neural network is universal function approximator

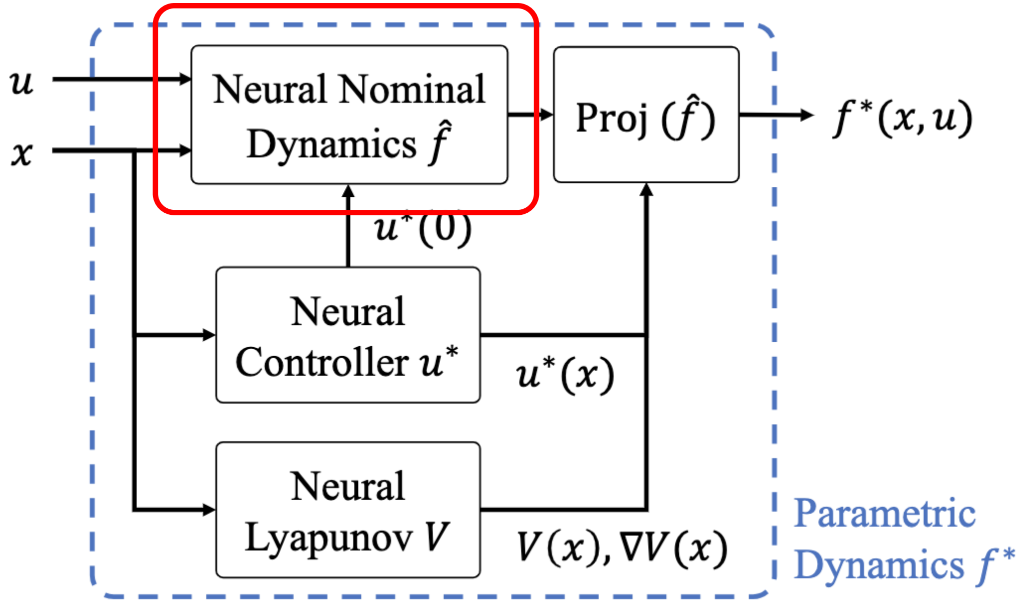

Interesting observation:
The constrained model still seems very expressive!

Neural network is universal function approximator

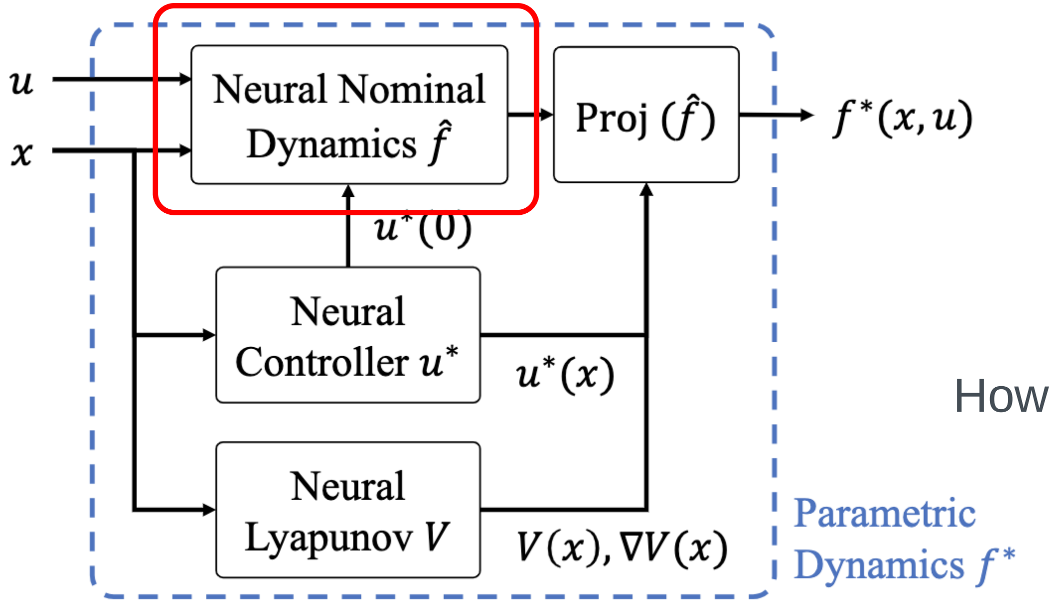

How about projected model?

New result: the constrained model is a universal approximator given the constraint

New result: the constrained model is a universal approximator given the constraint

- General input-dependent affine constraint: $a(x)^\top f(x) + b(x) \leq 0 \;\; \forall x \in \mathcal{X}$.

New result: the constrained model is a universal approximator given the constraint

- General input-dependent affine constraint: $a(x)^\top f(x) + b(x) \leq 0 \;\; \forall x \in \mathcal{X}$.

$$\mathcal{P}(f_\theta)(x) = \arg\min_{z \in \mathbb{R}^m} \|z - f_\theta(x)\|_2 \;\; \text{s.t. } a(x)^\top z + b(x) \leq 0 \quad \forall x \in \mathcal{X}$$

$$= f_\theta(x) - \frac{a(x)}{\|a(x)\|^2} \text{ReLU}\big(a(x)^\top f_\theta(x) + b(x)\big) \;\; \forall x \in \mathcal{X},$$

New result: the constrained model is a universal approximator given the constraint

- General input-dependent affine constraint: $a(x)^\top f(x) + b(x) \leq 0 \;\; \forall x \in \mathcal{X}$.

$$\mathcal{P}(f_\theta)(x) = \arg\min_{z \in \mathbb{R}^m} \|z - f_\theta(x)\|_2 \;\; \text{s.t. } a(x)^\top z + b(x) \leq 0 \quad \forall x \in \mathcal{X}$$

$$= f_\theta(x) - \frac{a(x)}{\|a(x)\|^2} \mathsf{ReLU}\big(a(x)^\top f_\theta(x) + b(x)\big) \;\; \forall x \in \mathcal{X},$$

Theorem 3 (informal): For any L^p function f with $p \in [1, \infty)$ that satisfies the constraint, there exists a deep neural network such that its projected model is arbitrarily close to f.

Conclusion

Conclusion

- Proposed a novel approach for control with inherent Lyapunov stability (CoILS)

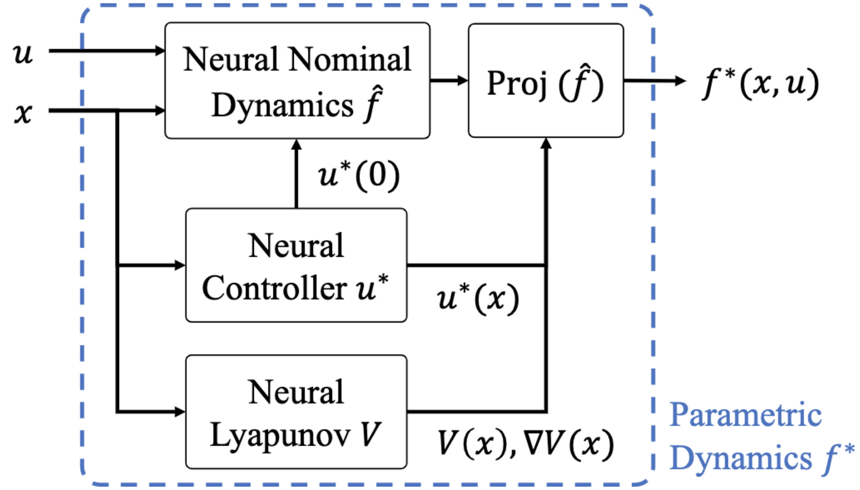

Conclusion

- Proposed a novel approach for control with inherent Lyapunov stability (CoILS)
- By fitting data, CoILS co-optimizes dynamics, controller, and Lyapunov function

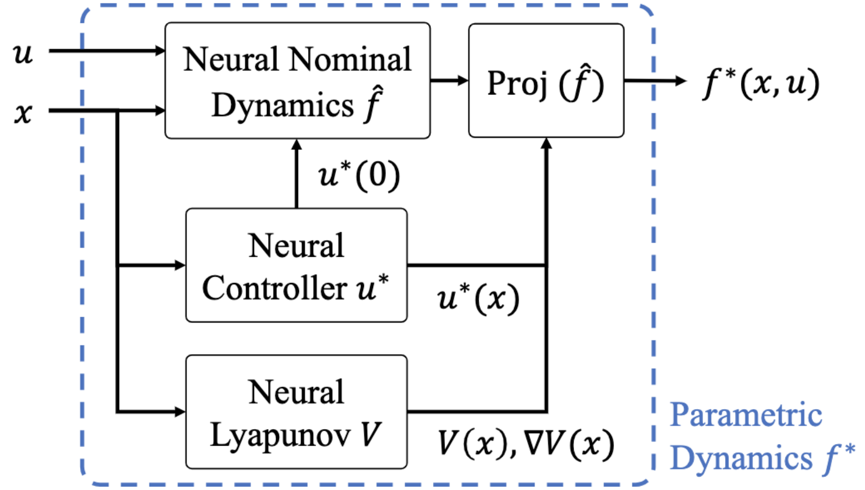

Conclusion

- Proposed a novel approach for control with inherent Lyapunov stability (CoILS)
- By fitting data, CoILS co-optimizes dynamics, controller, and Lyapunov function
- More generally: Alternative to neural net verification!
  Allows for incorporating constraints into the neural net

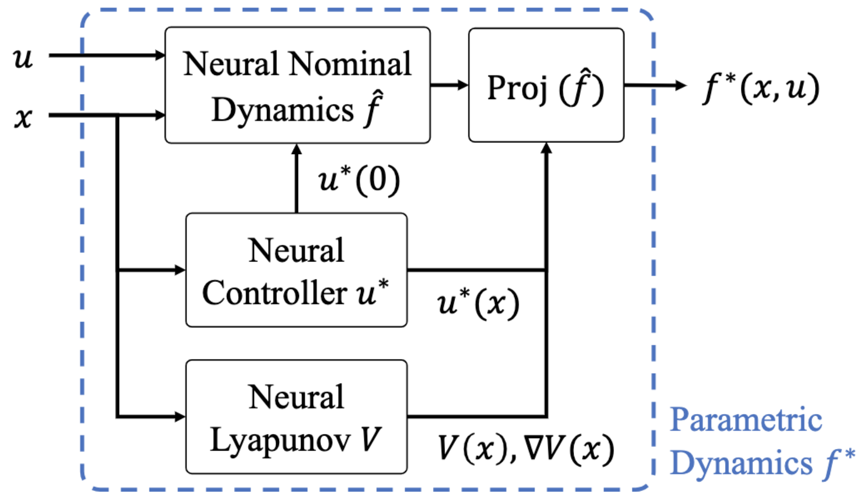

Conclusion

- Proposed a novel approach for control with inherent Lyapunov stability (CoILS)
- By fitting data, CoILS co-optimizes dynamics, controller, and Lyapunov function
- More generally: Alternative to neural net verification!
  Allows for incorporating constraints into the neural net

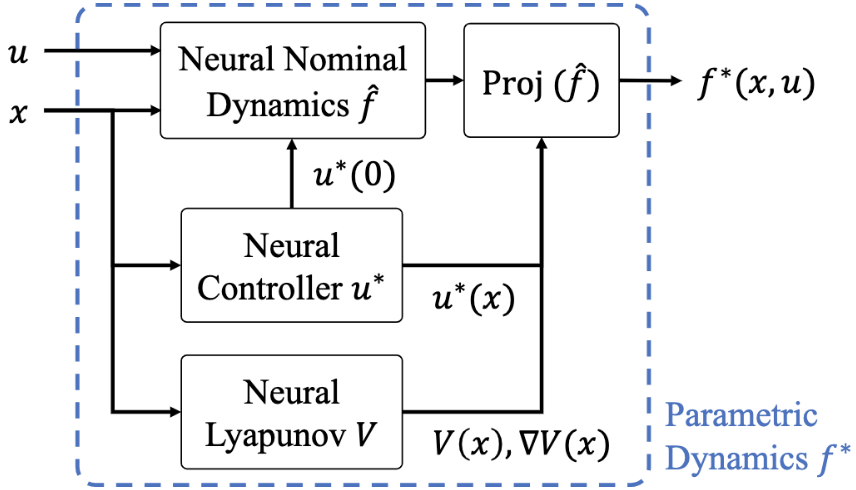

Future Work

- Include control-oriented metric in loss functions
- Extend to finding output-feedback controller for partially observable systems
- Scaling to high-dimensional systems
- Less conservative guarantees

Conclusion

- Proposed a novel approach for control with inherent Lyapunov stability (CoILS)
- By fitting data, CoILS co-optimizes dynamics, controller, and Lyapunov function
- More generally: Alternative to neural net verification!
  Allows for incorporating constraints into the neural net

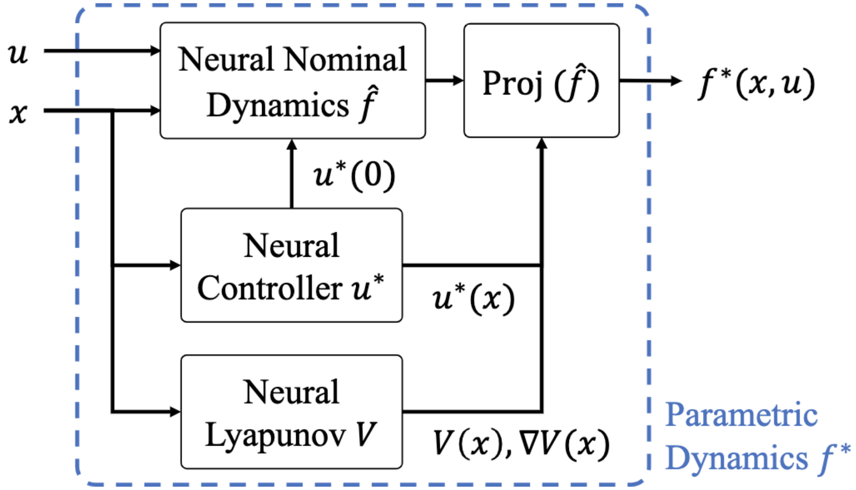

Future Work

- Include control-oriented metric in loss functions
- Extend to finding output-feedback controller for partially observable systems
- Scaling to high-dimensional systems
- Less conservative guarantees

arXiv:2303.03157