

Safe and efficient control using neural networks: An interior point approach

Baosen Zhang

ECE, University of Washington

NREL Autonomous Energy Systems Workshop

July 15th , 2022

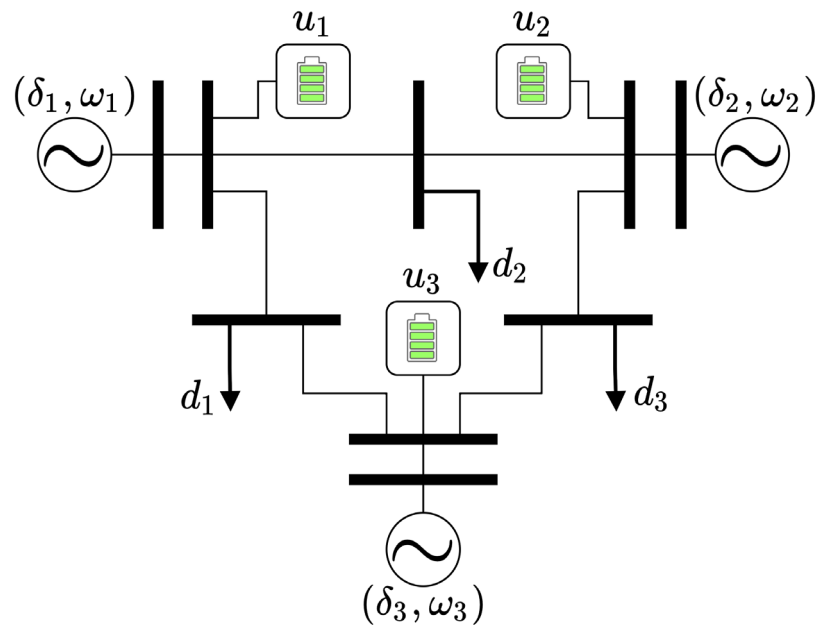
Joint work with
Daniel Tabas



- Power systems have several regulated quantities that must stay within prescribed ranges:
 - Bus voltage
 - Generator frequencies
 - Line flows
- Currently we have wide operating margins, but this is becoming increasingly difficult
- Adaptive control and decisions are useful, and we focus on the **computational aspects**

- Make fast, efficient, and safe decisions

E.g., Frequency regulation with storage



- Frequencies and angles should stay within some bound
- Storage have actuation constraints

- **Safety**: all constraints should be satisfied for a set of disturbances
- **Efficiency**: minimize cost, low computation time

$$\min_x c(x) + \boxed{o(x)} \quad \text{—————} \quad \min_y g(x, y)$$

subject to: subject to:

Energy adequacy constraints
Technology constraints
Transmission constraints

Power balance constraints
Generation constraints

- The problem size can get very large, especially if storage and multiple scenarios are considered

Computation Bottlenecks




- Common problem structure:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{d}_t)$$

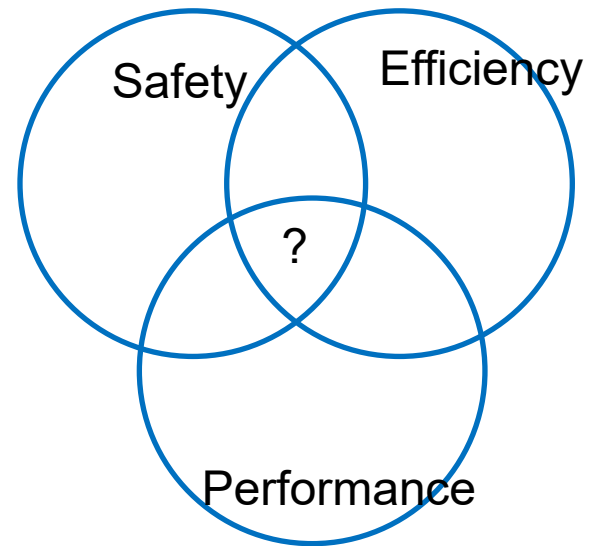

state control disturbance

Find a control sequence that minimize cost and satisfy all state and control constraints

- Sometimes explicitly solving the problem is too computationally expensive
- Use a neural network as a proxy



Key Issues and Design Goals




Safe: x and u stay in their constraints

Computationally efficient

Performance: minimize some cost

For a class of systems (e.g., linear)

- We provide a way to design neural networks that are always safe
- Easy to train, simple algebraic operations
- Based on geometry of convex sets and interior point algorithms



$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + \mathbf{E}d_t$$


$$\left. \begin{array}{l} \text{State: } \mathbf{x} \in \mathcal{X} \\ \text{Control: } \mathbf{u} \in \mathcal{U} \\ \text{Disturbance: } d \in \mathcal{D} \end{array} \right\} \begin{array}{l} \text{convex} \\ \text{polytopes} \end{array} \quad \mathcal{X} = \{\mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$$

Static Feedback Policy: $\mathbf{u}_t = \pi_{\theta}(\mathbf{x}_t)$

Safety: For all t , $\mathbf{u}_t \in \mathcal{U}$ and $\mathbf{x}_{t+1} \in \mathcal{X}$

Performance: $\min_{\theta} \text{cost}(\mathbf{x}, \pi_{\theta}(\mathbf{u}))$

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + \mathbf{E}\mathbf{d}_t$$



Operational Constraints

Invariant set

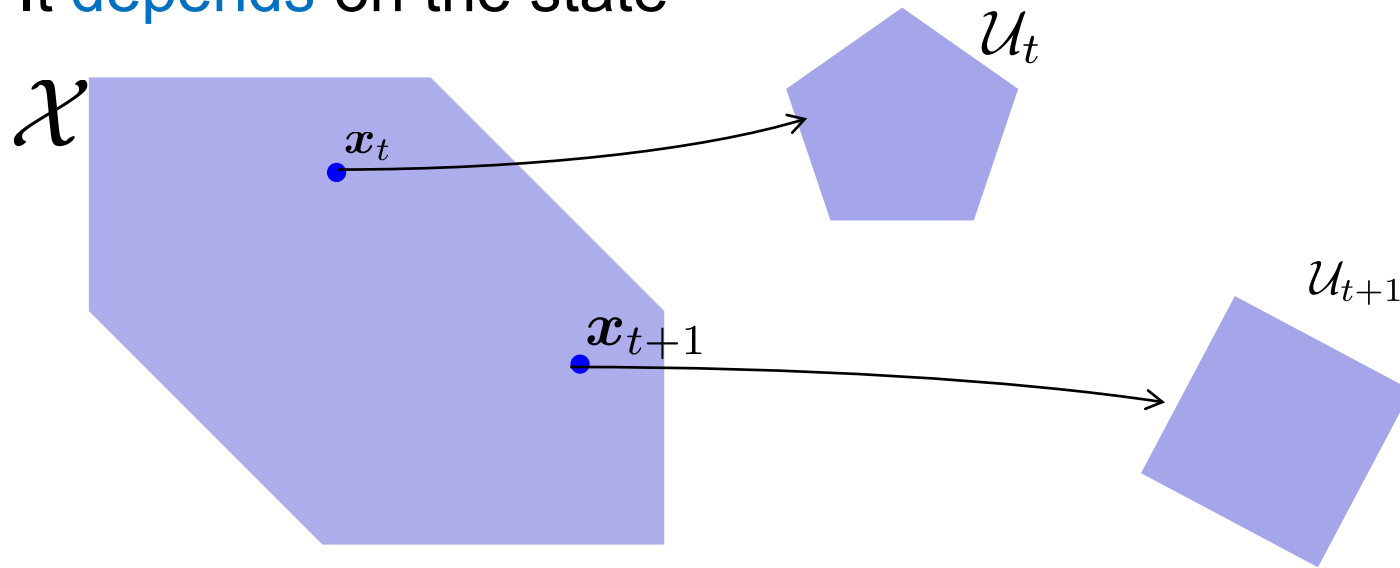
$$\mathcal{X} = \{ \mathbf{x} : \exists \mathbf{u}, \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{E}\mathbf{d} \in \mathcal{X} \forall \mathbf{d}, \mathbf{x} \in \bar{\mathcal{X}} \}$$

- Invariant set can be computed offline

Safe Actions



- Safe action set: $\mathcal{U}_t = \{u \in \mathcal{U} : x_{t+1} \in \mathcal{X}\}$
- It **depends** on the state



- Main challenge: how to get $\pi(x_t)$ to be in these shifting polytopes?

- It's not hard to get the output of a neural network to be in a **fixed** polytope
- The safe action set **changes** at every timestep


Current strategy: penalty or projection

Penalty-Based Methods



- High training cost for trajectories that violates the constraints
- Hard to balance safety and exploration
- At best probabilistic guarantees

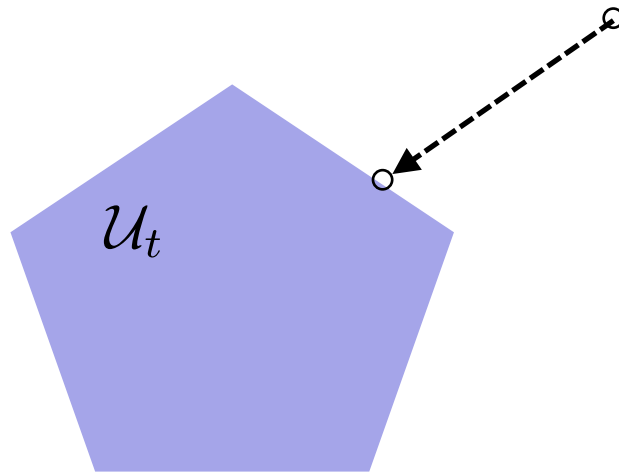
State Constraints



Projection-Based Methods



- Train a controller, if the control action falls outside, project it back

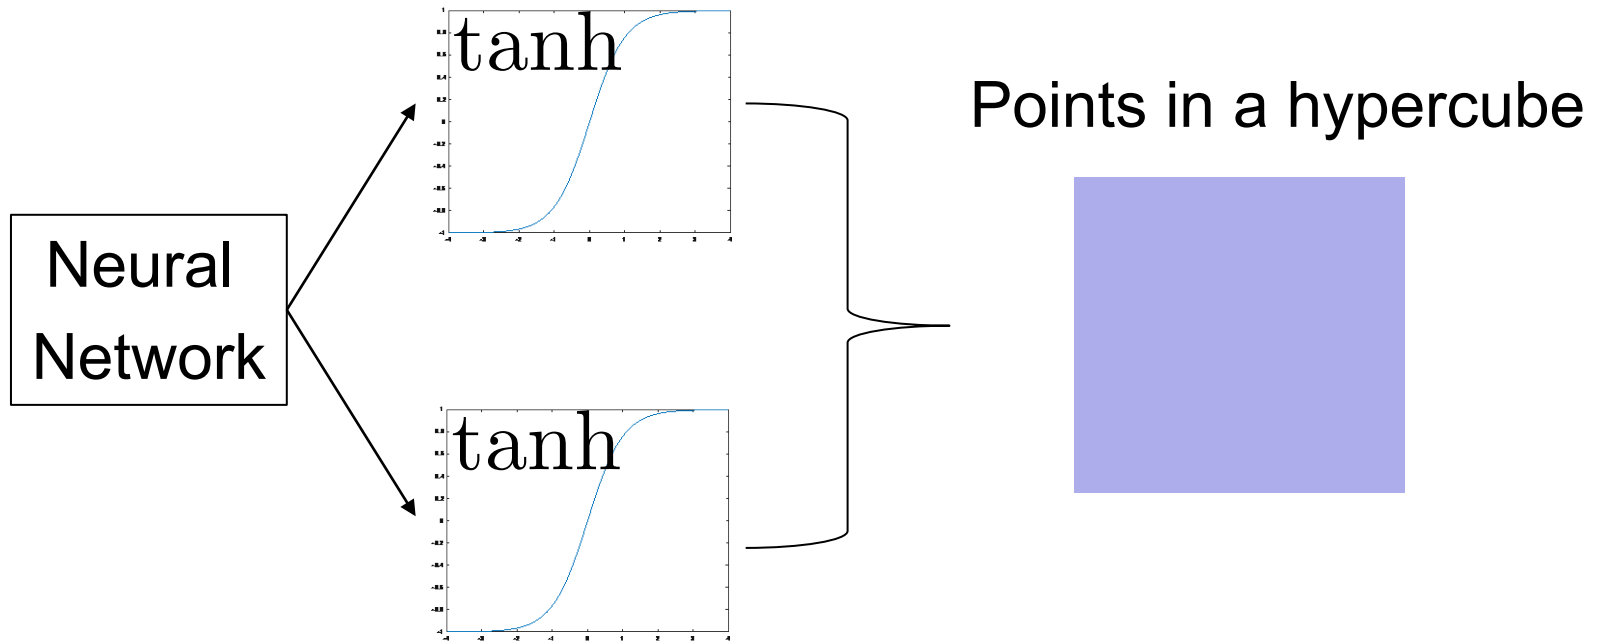


- Needs to solve an optimization problem
- Can over-explore the boundary


Restricting the Output



- Not all sets are difficult: Axis-aligned rectangles are easy



- Mapping a box to polytopes



- We want a mapping between polytopes that is
 - Bijective
 - Easy to compute
 - (sub) Differentiable
- We use the [gauge map](#)

Mapping Between Polytopes

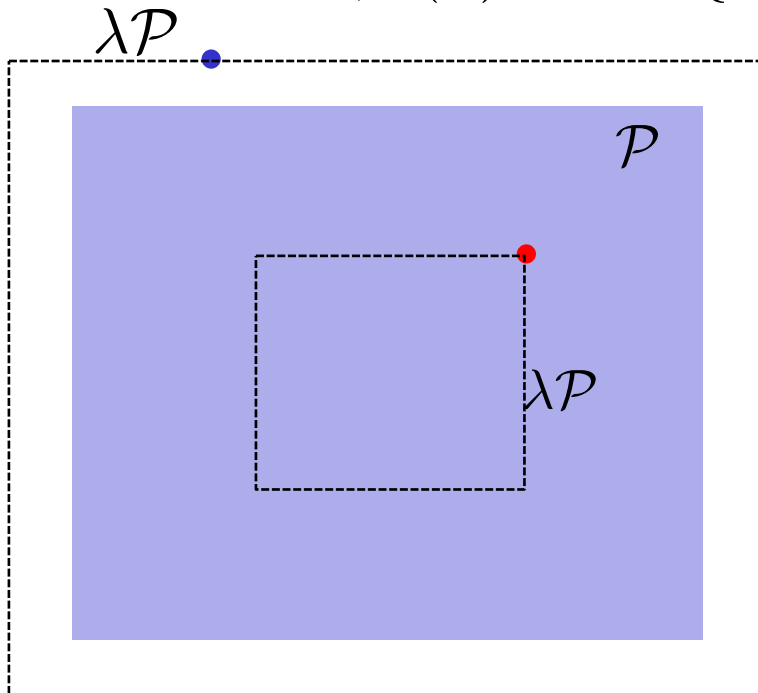


- Given a polytope that contains the origin,

$$\mathcal{P} = \{z \in \mathbb{R}^n : \mathbf{F}z \leq \mathbf{g}\}$$



- The gauge function

$$\gamma_{\mathcal{P}}(z) = \min\{\lambda \geq 0 : z \in \lambda\mathcal{P}\}$$



$$\begin{aligned}\gamma_{\mathcal{P}}(z) &= \min\{\lambda \geq 0 : \mathbf{F}z \leq \lambda\mathbf{g}\} \\ &= \min\{\lambda \geq 0 : \lambda \geq \frac{\mathbf{F}_i^T z}{g_i}, \forall i\} \\ &= \max_i \frac{\mathbf{F}_i^T z}{g_i}\end{aligned}$$

Gauge Map





- Bijection between the sets
- Closed-form formula
- Sub-differentiable in the parameters

Convex Sets




- A ray from the origin can only cross the boundary of a convex set once
- This provides a unique way of identifying the the points with respect to this set

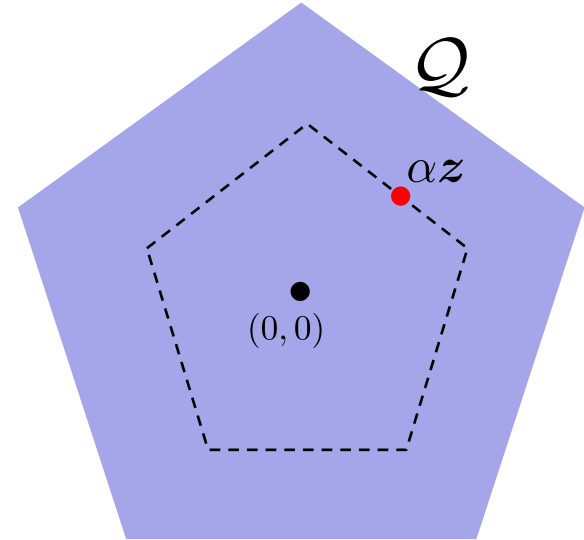


- For convex polytopes, the gauge function is easy to compute

Gauge Map



$$\mathcal{P} = \{z : \mathbf{F}z \leq \mathbf{g}\}$$




$$\mathcal{Q} = \{z : \mathbf{H}z \leq \mathbf{k}\}$$

Find $\alpha \geq 0$ such that $\gamma_{\mathcal{Q}}(\alpha z) = \gamma_{\mathcal{P}}(z)$

$$\alpha = \frac{\gamma_{\mathcal{P}}(z)}{\gamma_{\mathcal{Q}}(z)} = \frac{\max_i \mathbf{F}_i^T z / g_i}{\max_i \mathbf{H}_i^T z / k_i}$$

Only requires the half-space representation of polytopes

Gauge Policy Networks



- \mathcal{P} : hypercube
- \mathcal{U}_t : half-space representation $\mathcal{U}_t = \{u : \mathbf{H}u \leq \mathbf{h}(x_t)\}$
- Control policy:

$$\pi_{\theta}(\mathbf{x}_t) = G_t \circ \tanh \circ \phi_{\theta}(\mathbf{x}_t)$$

gauge map

clamping function


feedforward neural network

- Sub-differentiable, easy to train
- But \mathcal{U}_t may not contain the origin

Interior Points



- If we know an interior point, we can always shift a set to contain the origin



- \mathbf{u}^{int} some interior point

$$\pi_{\theta}(\mathbf{x}_t) = G_t \circ \tanh \circ \phi_{\theta}(\mathbf{x}_t) + \mathbf{u}^{int}$$

↑
gauge map to $\mathcal{U}_t - \mathbf{u}^{int}$

Finding Interior Points

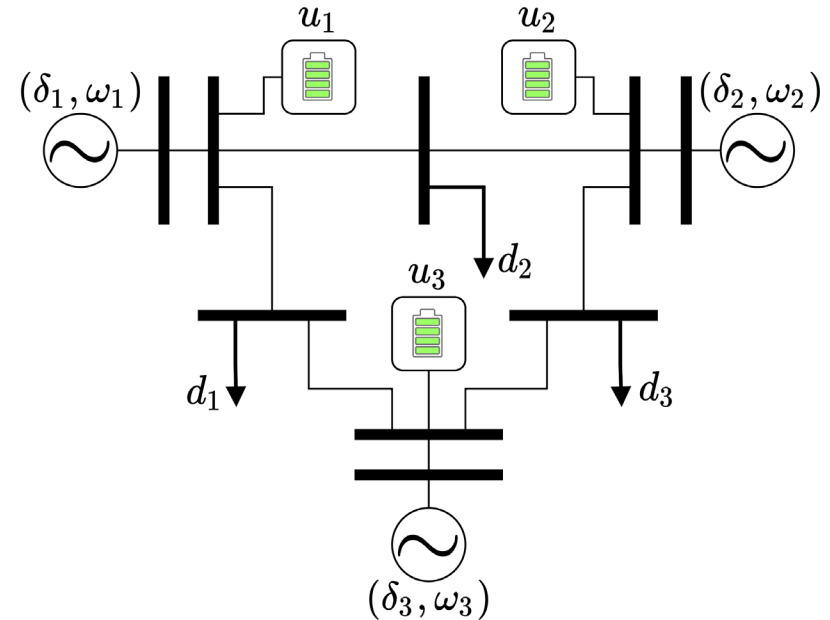


- In planning, there are often trivial, high cost, but feasible solutions
- In real-time control, there are often easy ways to find feasible control actions
 - There is a baseline feasible controller
 - A few iterations suffices

Frequency Control



- States: angle/frequency
- Control: battery output
- Cost: freq dev+bat deg



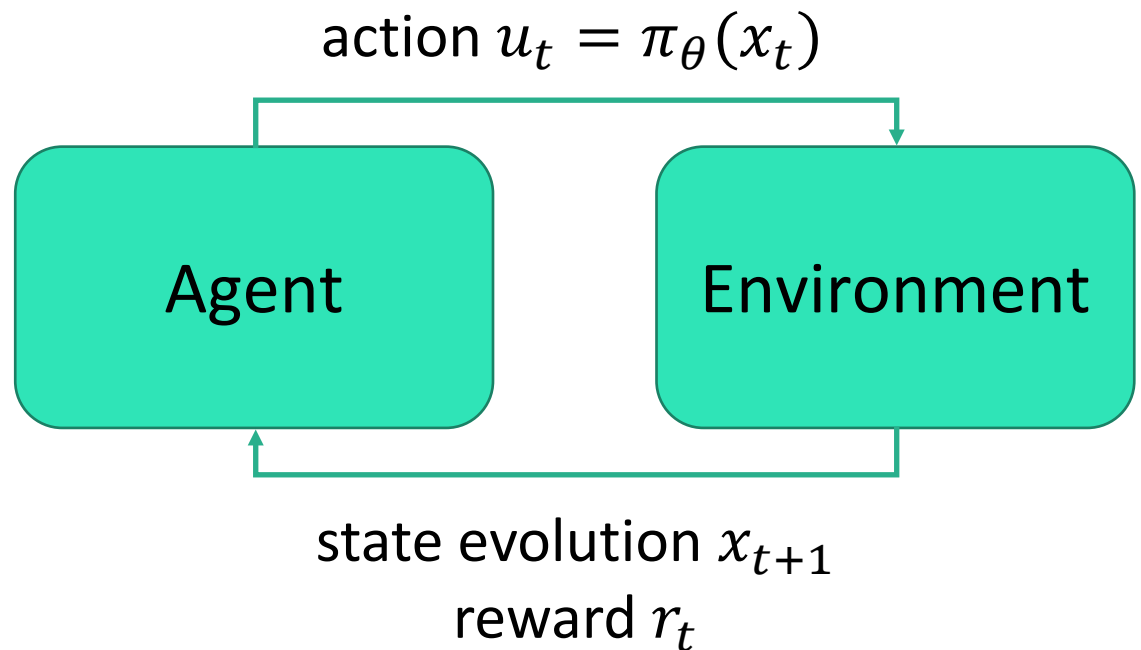
- There exist a linear safe controller:

$$\mathbf{u}_t^{int} = \mathbf{K} \mathbf{x}_t$$

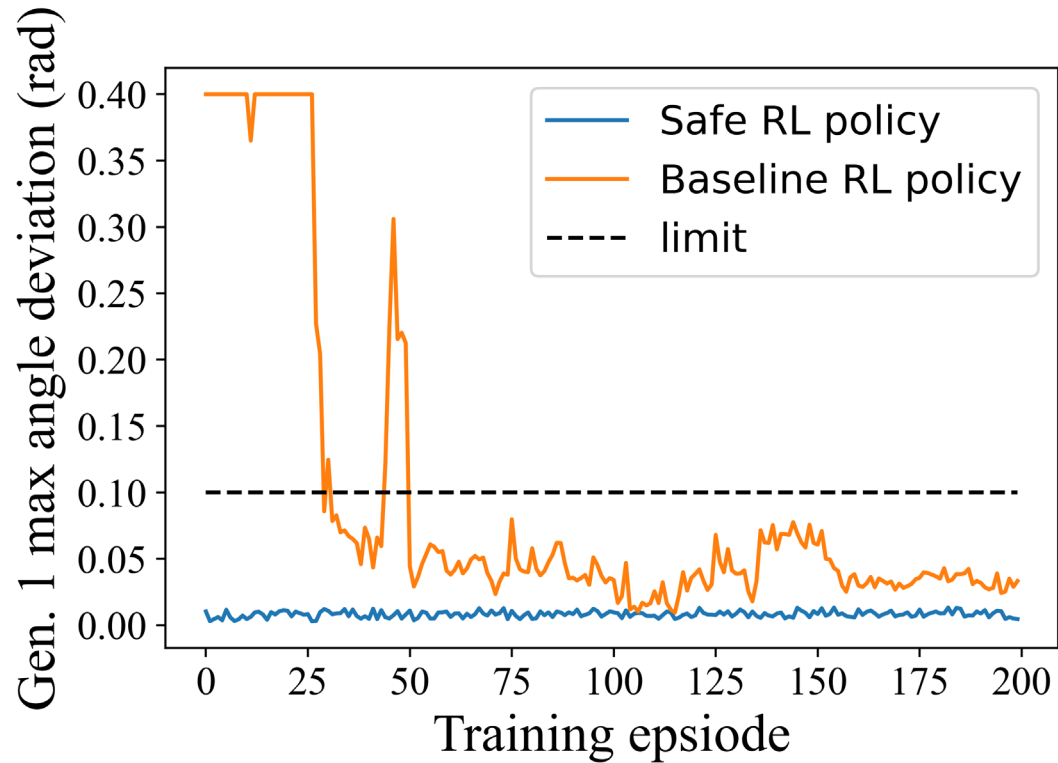
Policy Network



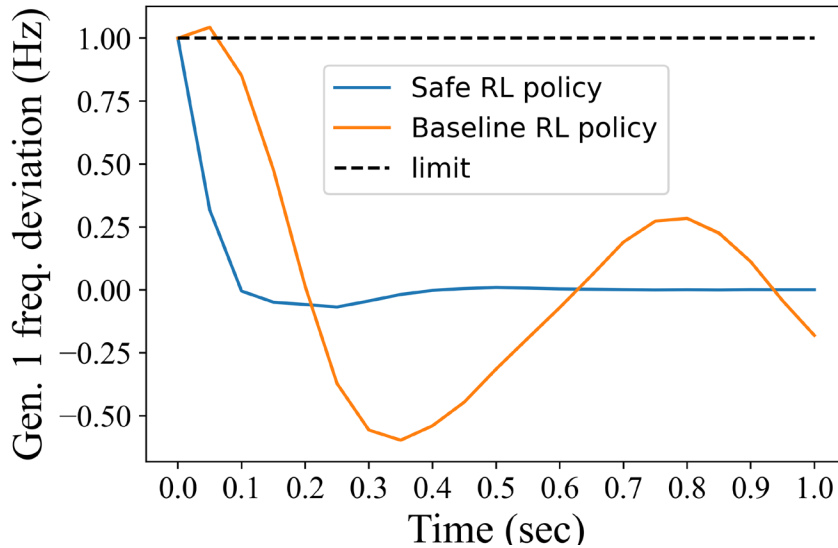
- Training: safely interact with environment to choose θ (policy gradient algorithms)
- Testing: θ fixed
- Benchmark: penalty-based approach



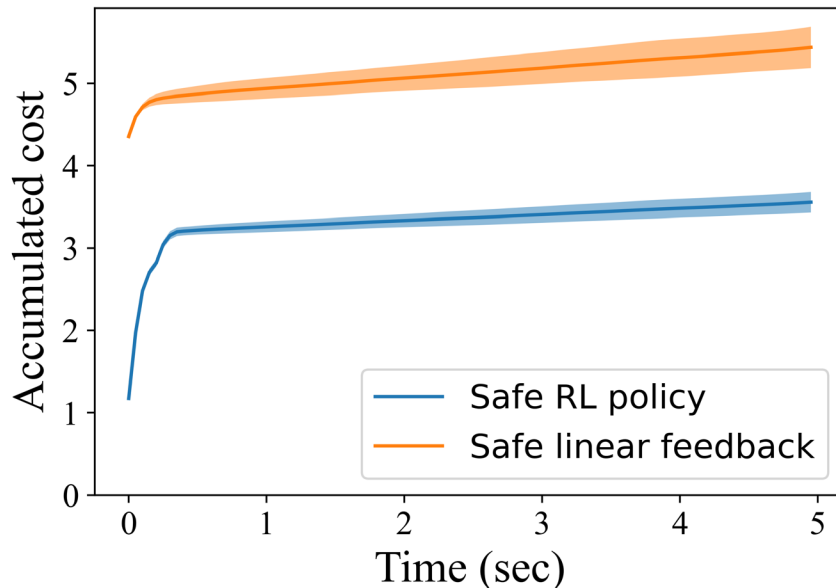
Safety During Training



Testing and Cost



Safety during testing



Lower operational costs than (safe) linear feedback

Explicit Model Predictive Control




- MPC is a very popular control strategy
- Exploit known dynamics and cost
- Constraint satisfaction and treat disturbances less conservatively
- Solving online MPC can be computationally intensive

$$\begin{aligned} \min_{\mathbf{u}_1, \dots, \mathbf{u}_T} \quad & \sum_{t=1}^T c(\mathbf{x}_t, \mathbf{u}_t) \\ \text{s.t.} \quad & \mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t \\ & \mathbf{x}_t \in \mathcal{X}, \mathbf{u}_t \in \mathcal{U} \quad \forall t \end{aligned}$$

- If the system is linear and the cost quadratic, then the optimal action are piecewise affine functions

$$\mathbf{u}_1^*, \dots, \mathbf{u}_T^* = f(\mathbf{x}_0)$$



- The number of pieces scales exponentially
- Lots of work on using neural network to approximate
- Projection is used to provide hard guarantees

Solution Concept



- Use gauge neural network to choose a sequence u_0, u_1, \dots, u_{T-1} inside the feasible set $\mathcal{F}(x_0)$
- **Guarantees safety** without oracles or projections

How do we get an interior point in $\mathcal{F}(x_0)$?

Same approach as interior point algorithms used in optimization.

Interior Point Algorithms



- Suppose we want to solve

$$\begin{aligned} \min_{\mathbf{x}} f(\mathbf{x}) \\ \text{s.t. } g_i(\mathbf{x}) \leq 0, \quad i = 1, \dots, m \end{aligned}$$

- Path-following interior point algorithm solves a sequence of problems

$$\mathbf{x}_\mu = \arg \min f(\mathbf{x}) + \mu \sum_i^m \log(-g_i(\mathbf{x}))$$

← Barrier function

$$\mathbf{x}_\mu \rightarrow \mathbf{x}^* \text{ as } \mu \rightarrow 0$$

- But this path requires a strictly feasible point to get started

Phase 1 of Interior Point Algorithms

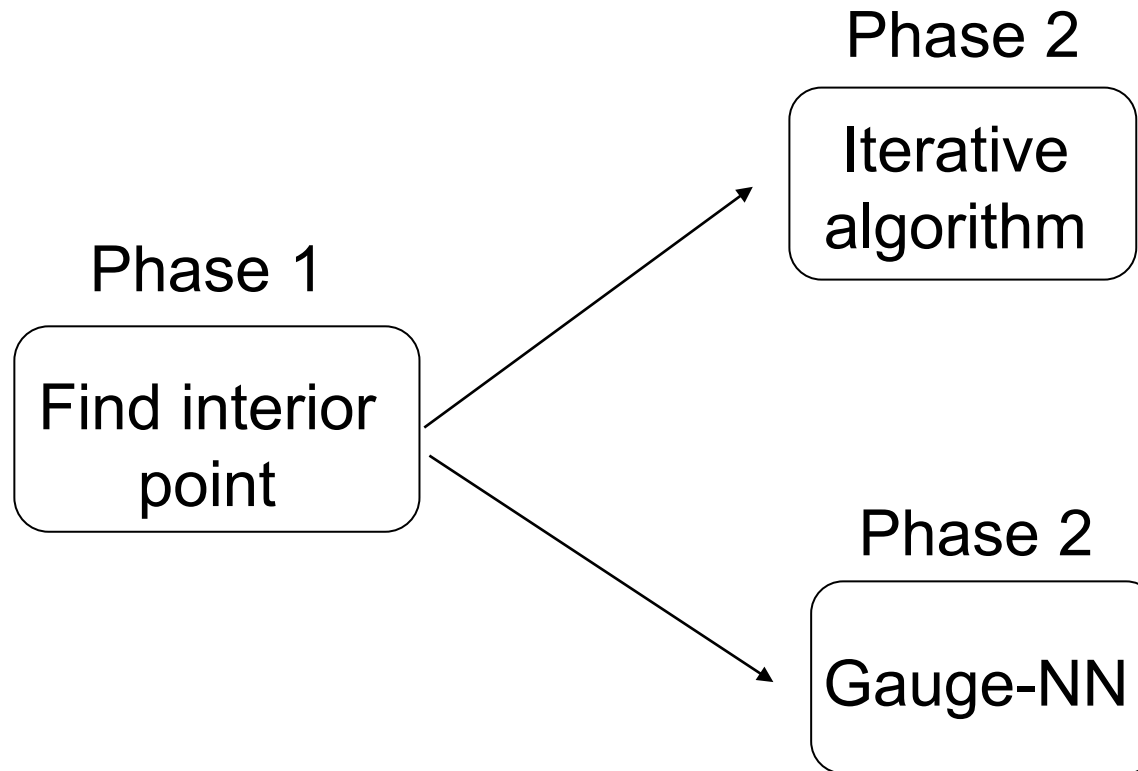


$$\begin{aligned} \min_{\mathbf{x}} f(\mathbf{x}) \\ \text{s.t. } g_i(\mathbf{x}) \leq 0, \quad i = 1, \dots, m \end{aligned}$$

- Sometimes a feasible point is easy to find
- If not, a phase 1 problem is solved

$$\begin{aligned} \min_{\mathbf{x}, s} s \\ \text{s.t. } g_i(\mathbf{x}) \leq s, \quad i = 1, \dots, m \end{aligned} \left. \begin{array}{l} \mathbf{x} = \mathbf{0} \\ s = \max_i g_i(\mathbf{0}) \end{array} \right\} \text{is feasible}$$

- This problem can be terminated once $s \leq 0$
- Finding a feasible point usually takes two or two iterations




- Gauge-NN can lead to significant speedup if phase 2 is more expensive than phase 1

- MPC with linear constraints and quadratic cost function
- Training:
 - Gauge neural network $u_\theta: R^n \rightarrow R^{mT}$
 - Data: $\{x_0^k\}_{k=1}^N$
 - Predict trajectories generated by u_θ
 - Training loss: average MPC cost over sample trajectories
- Testing:
 - Close the loop: policy $\pi_\theta(x_t)$ given by first action from $u_\theta(x_t)$

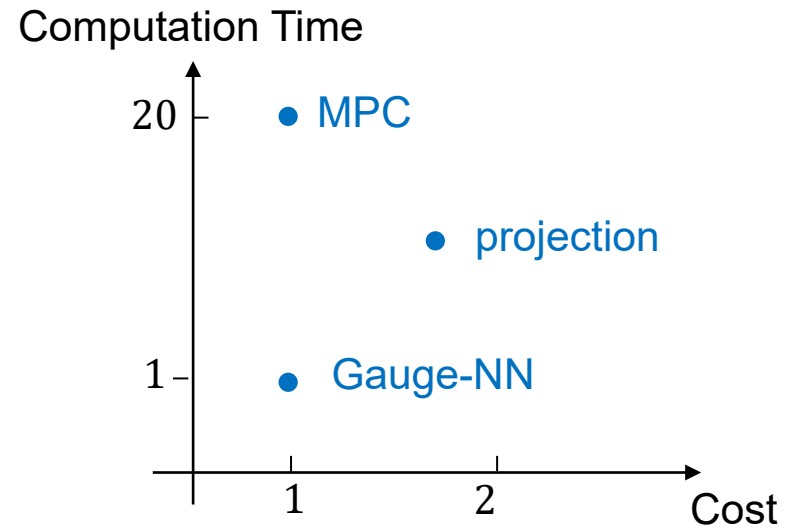
Simulation Results



Deterministic MPC



Scenario-based MPC



- A way to find safe policies through mapping between convex sets
- Can learn nonlinear policies
- Performance and computation speedup

Future work

- Convex but non-polytopic sets
- Non-convex costs (e.g., robust optimization)
- Multi-agent systems