Optimization-in-the-loop ML for energy and climate

Priya L. Donti
Executive Director, Climate Change AI
Incoming Assistant Professor, MIT

Climate change warrants rapid action

Filippo Monteforte | AFP | Getty Images

David Mcnew | Getty Images

NASA

Piyaset | Shutterstock.com

Impacts felt globally

Disproportionate impacts on most disadvantaged populations

Need net-zero greenhouse gas emissions by 2050 (IPCC 2018)

- Across energy, transport, buildings, industry, agriculture, forestry, etc.

Can machine learning play a role?

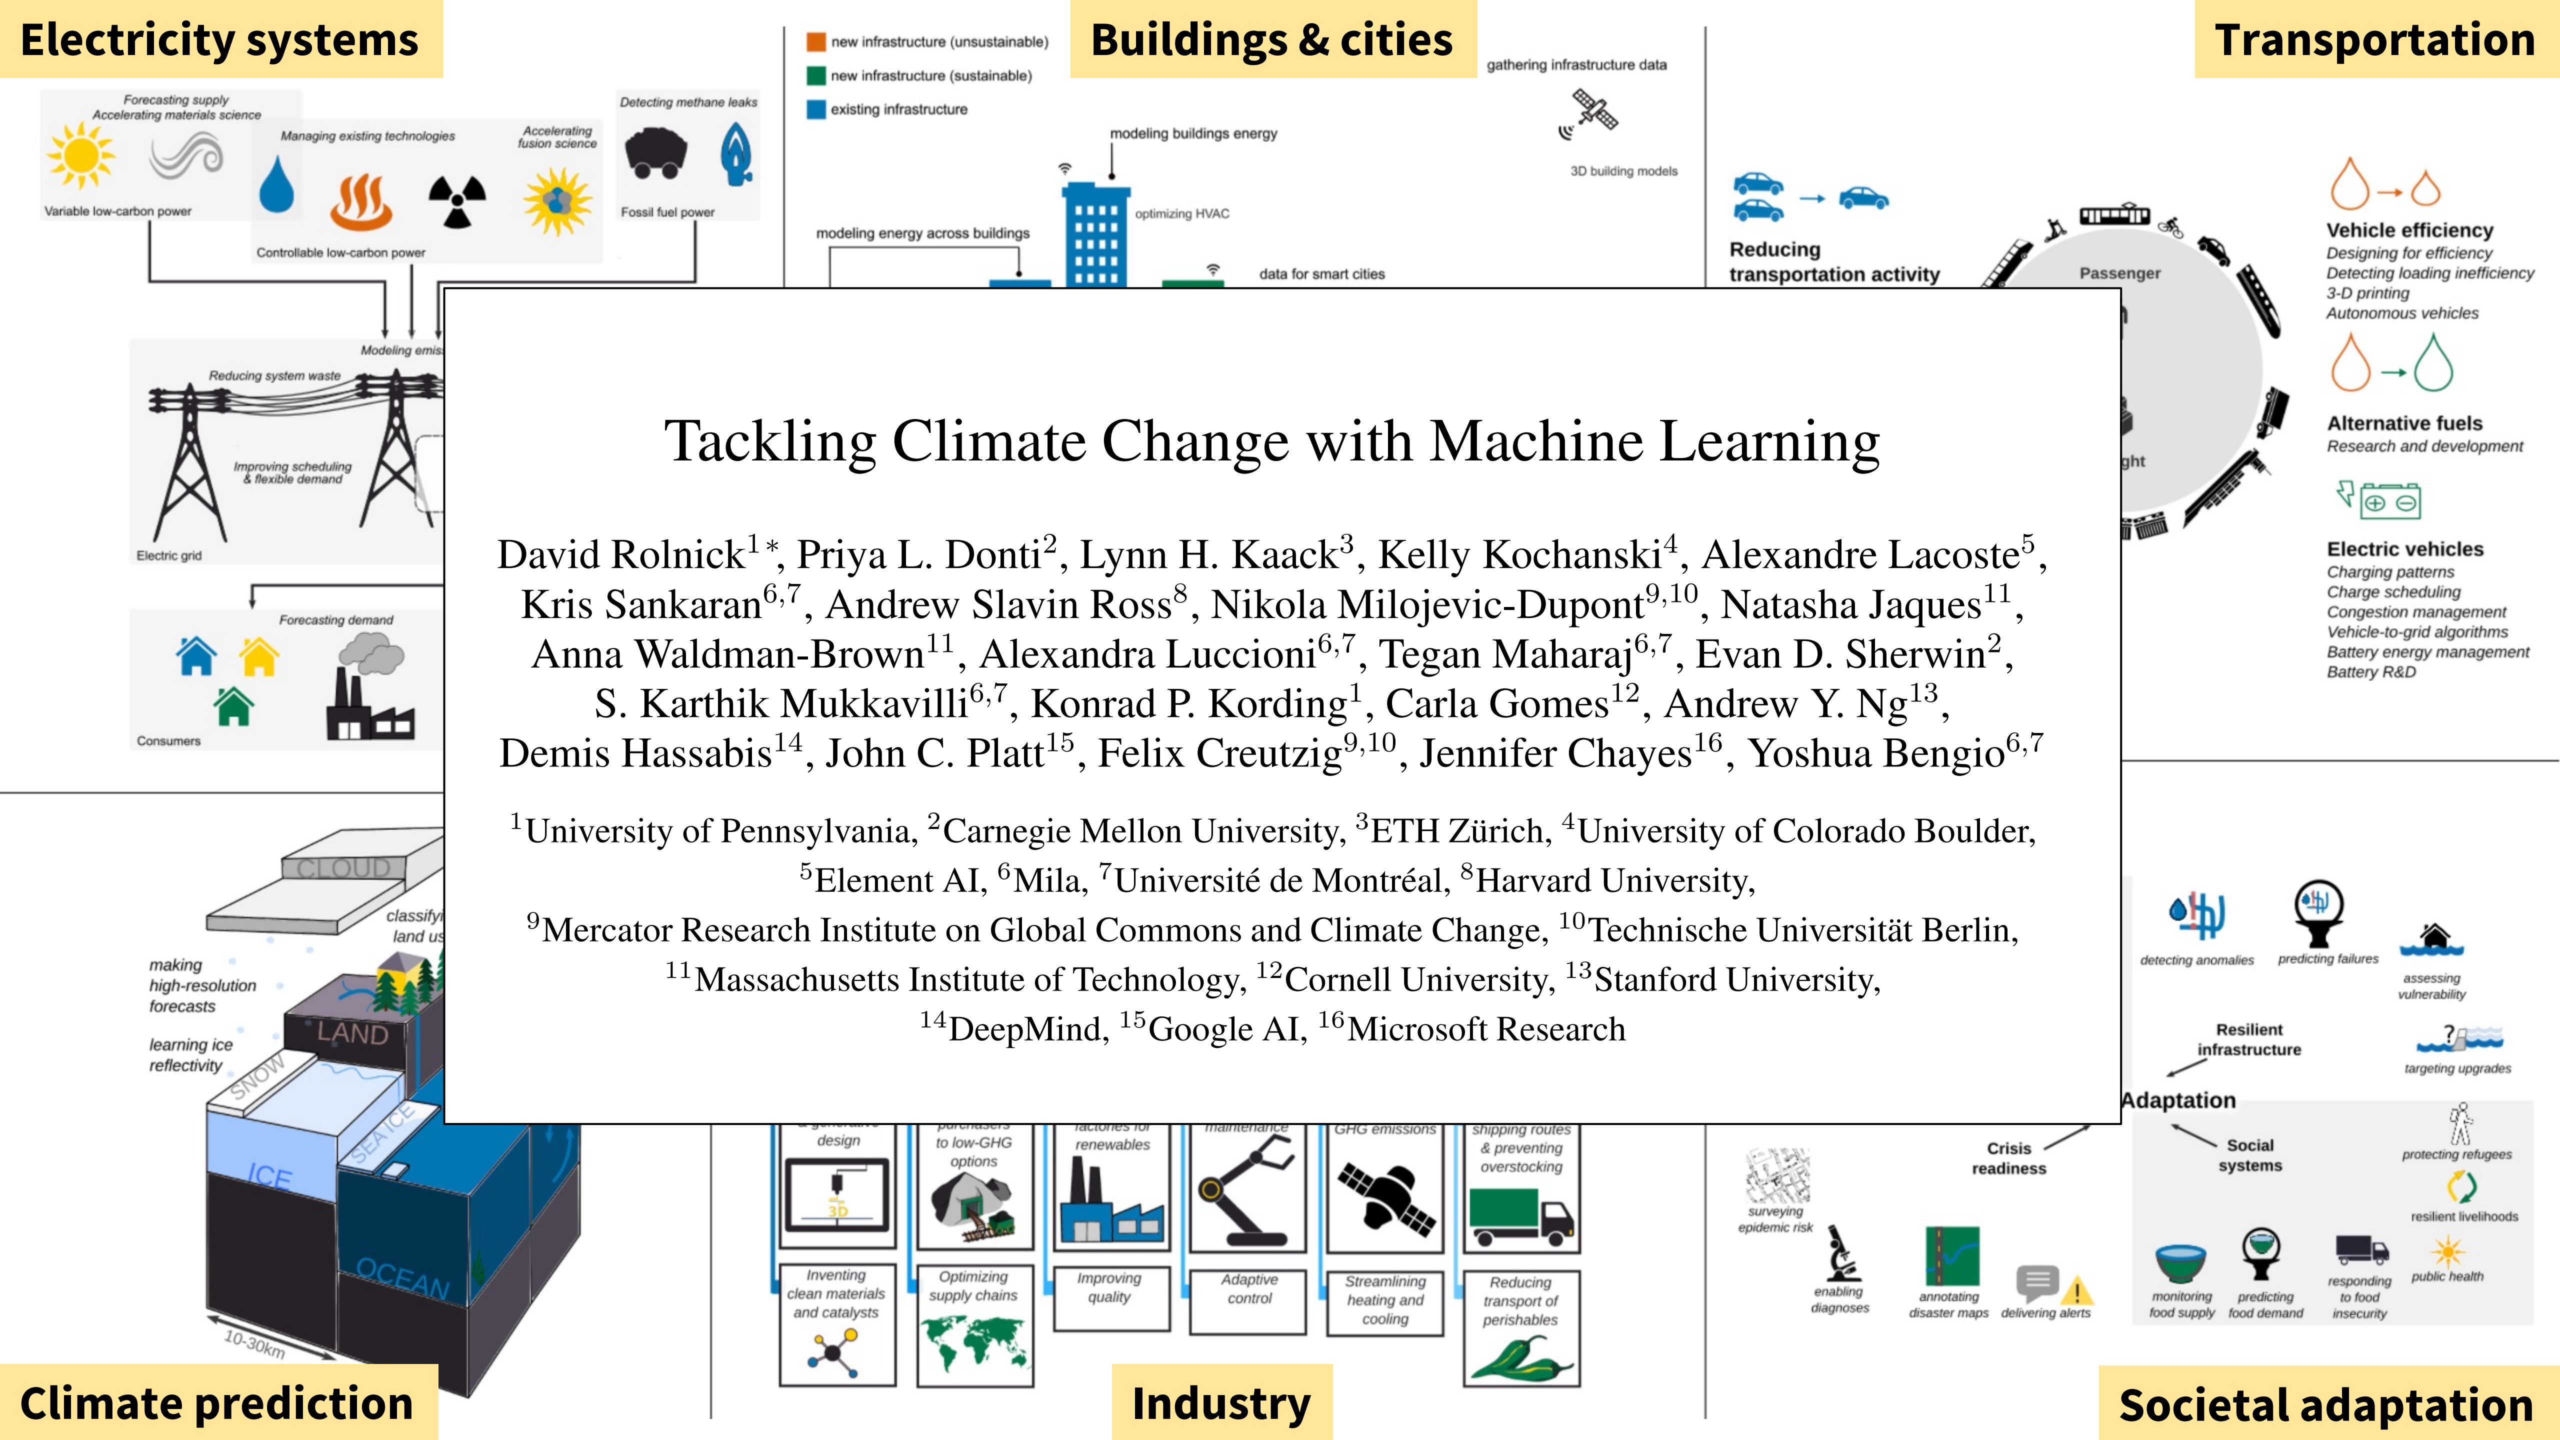

Tackling Climate Change with Machine Learning

David Rolnick[1]*, Priya L. Donti[2], Lynn H. Kaack[3], Kelly Kochanski[4], Alexandre Lacoste[5], Kris Sankaran[6,7], Andrew Slavin Ross[8], Nikola Milojevic-Dupont[9,10], Natasha Jaques[11], Anna Waldman-Brown[11], Alexandra Luccioni[6,7], Tegan Maharaj[6,7], Evan D. Sherwin[2], S. Karthik Mukkavilli[6,7], Konrad P. Kording[1], Carla Gomes[12], Andrew Y. Ng[13], Demis Hassabis[14], John C. Platt[15], Felix Creutzig[9,10], Jennifer Chayes[16], Yoshua Bengio[6,7]

[1]University of Pennsylvania, [2]Carnegie Mellon University, [3]ETH Zürich, [4]University of Colorado Boulder, [5]Element AI, [6]Mila, [7]Université de Montréal, [8]Harvard University, [9]Mercator Research Institute on Global Commons and Climate Change, [10]Technische Universität Berlin, [11]Massachusetts Institute of Technology, [12]Cornell University, [13]Stanford University, [14]DeepMind, [15]Google AI, [16]Microsoft Research

Power & energy problems involve physics, hard constraints, and decision-making

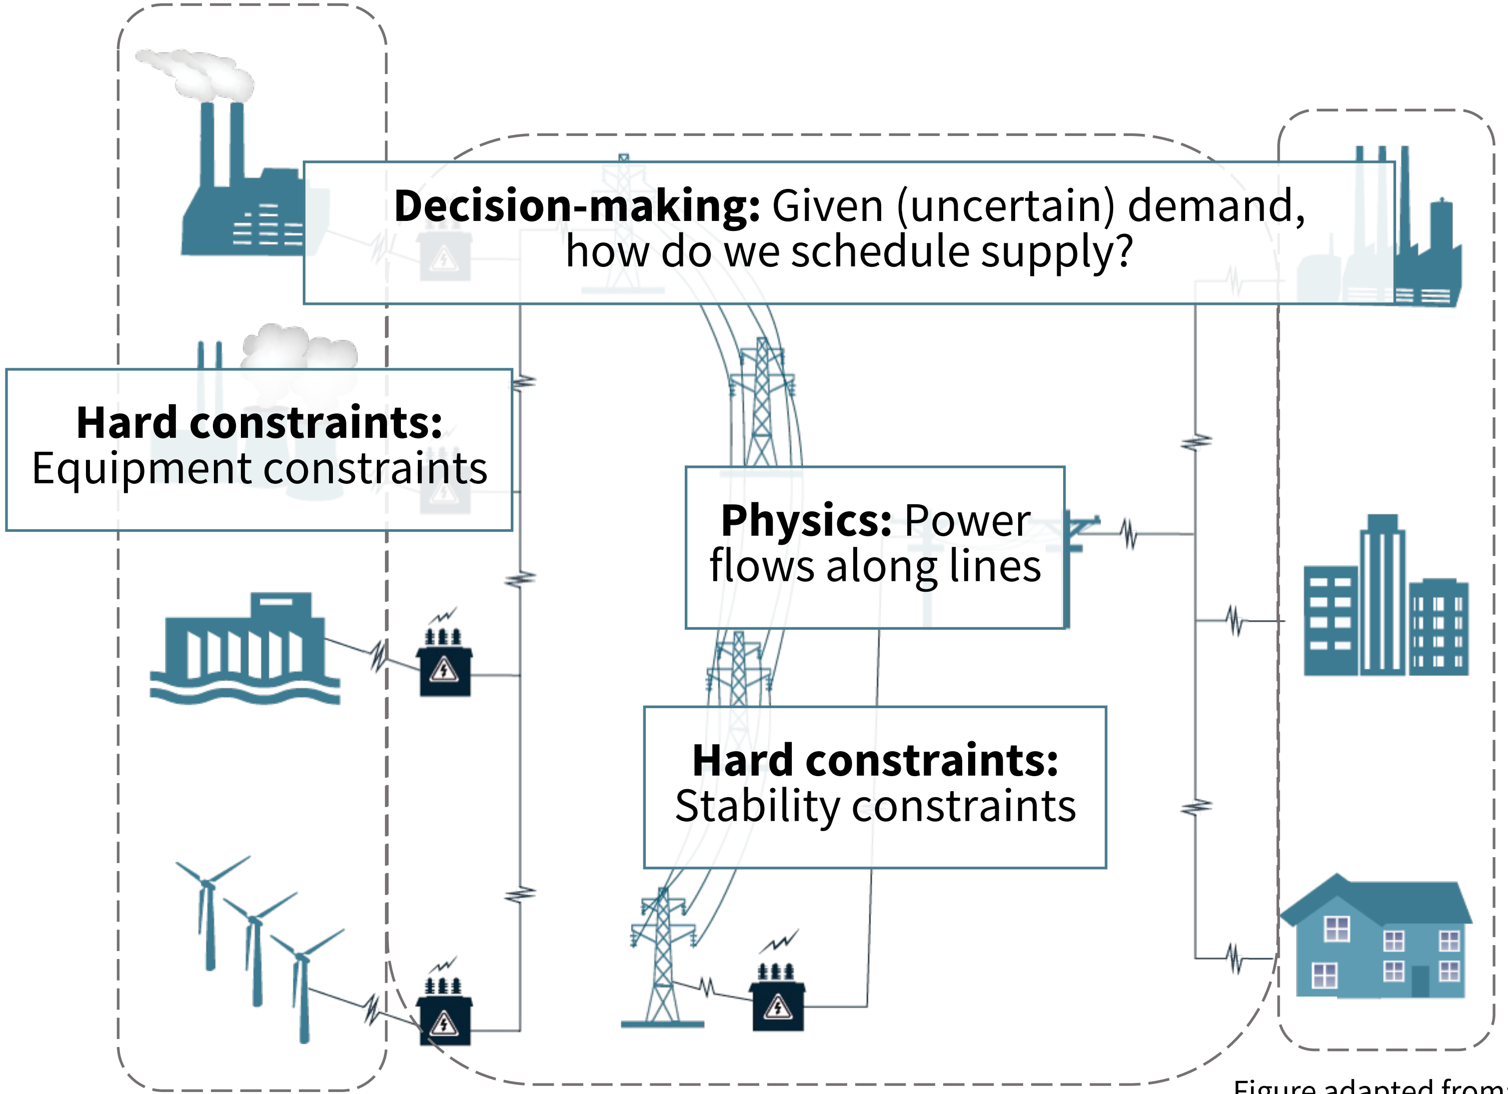

Figure adapted from: US Congressional Budget Office

Machine learning methods struggle with physics, hard constraints, and decision-making

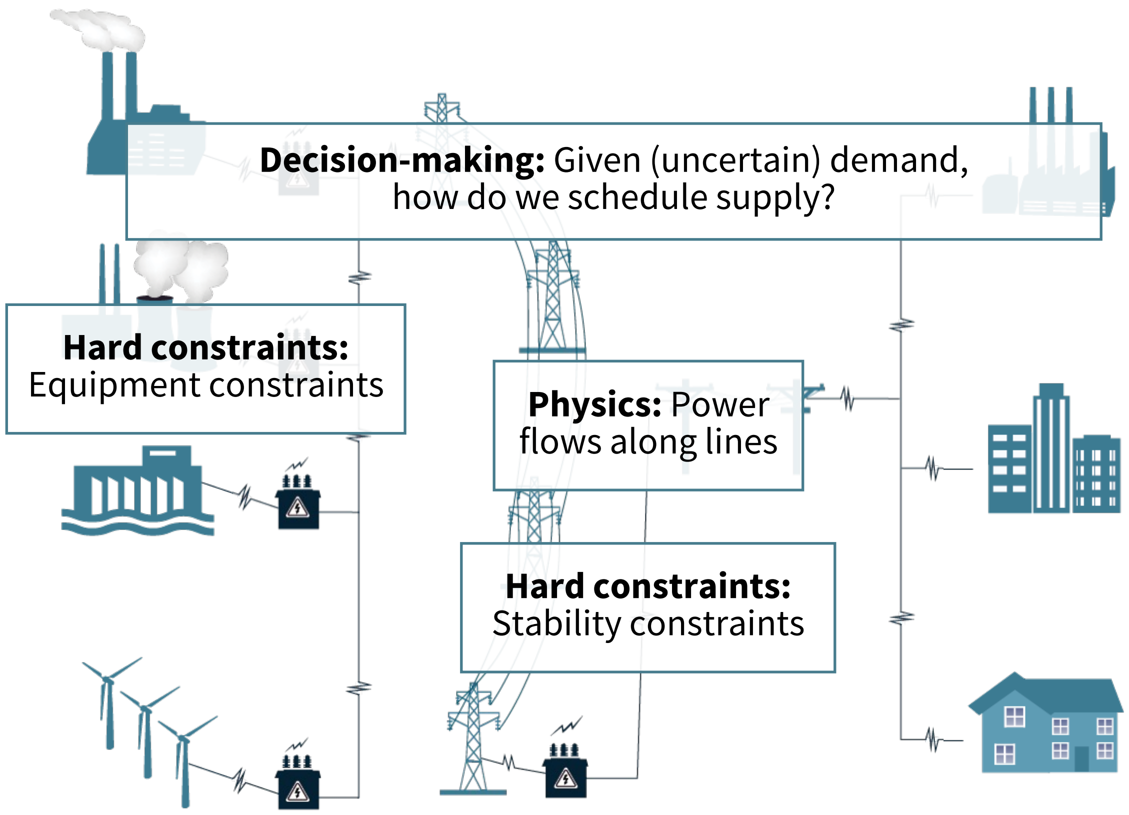

Need: Adaptive control of power generators, inverters, and batteries

- **ML:** Dynamic, data-driven control
- **Limitation:** Difficulty enforcing constraints (physics, equipment, stability)

Need: Electricity demand prediction

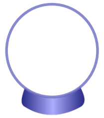

- **ML:** Time series forecasting
- **Limitation:** Difficulty making decision-cognizant error tradeoffs

How do we reap the benefits of ML methods while mitigating limitations?

Figure adapted from: US Congressional Budget Office

Optimization-in-the-loop ML

Framework for developing ML methods that incorporate knowledge of physics, hard constraints, or downstream decision-making procedures, via optimization problems

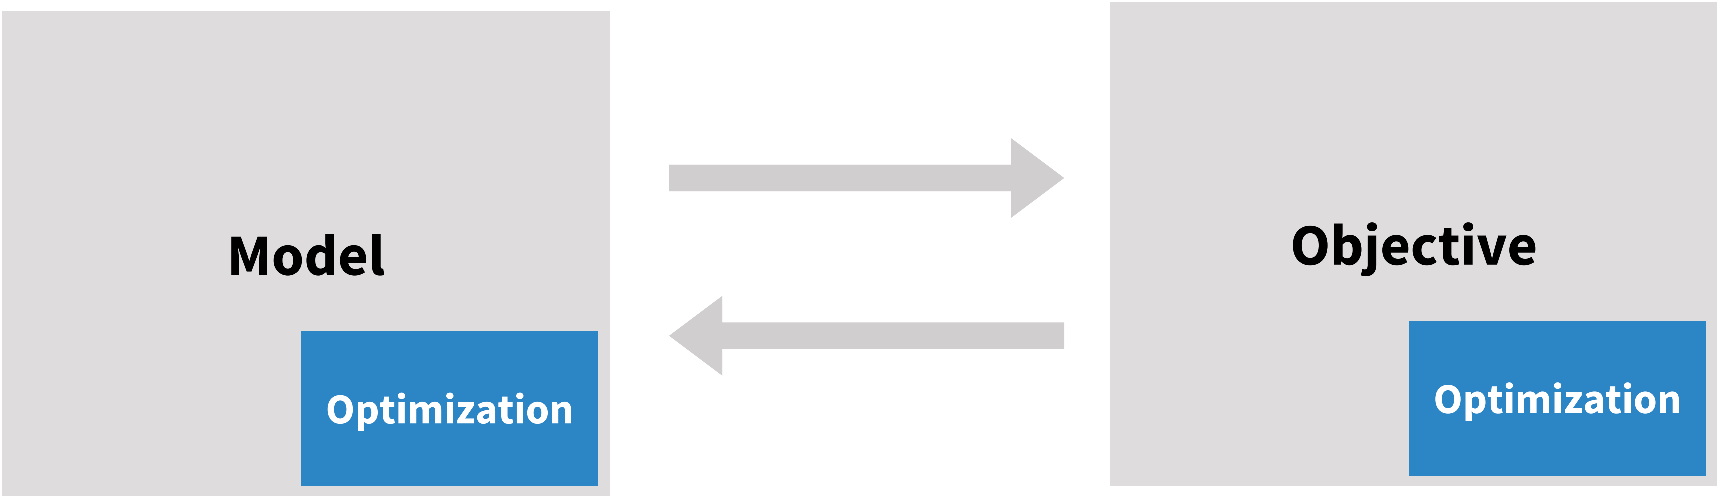

Optimization-in-the-loop ML

Framework for developing ML methods that incorporate knowledge of physics, hard constraints, or downstream decision-making procedures, via optimization problems

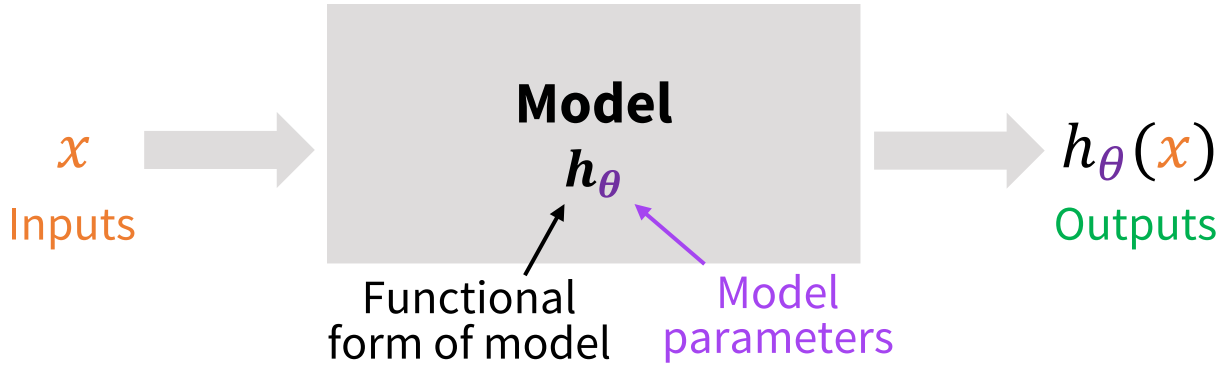

Optimization-in-the-loop ML

Framework for developing ML methods that incorporate knowledge of physics, hard constraints, or downstream decision-making procedures, via optimization problems

Example: Robust control (e.g., power system control with stability constraints)

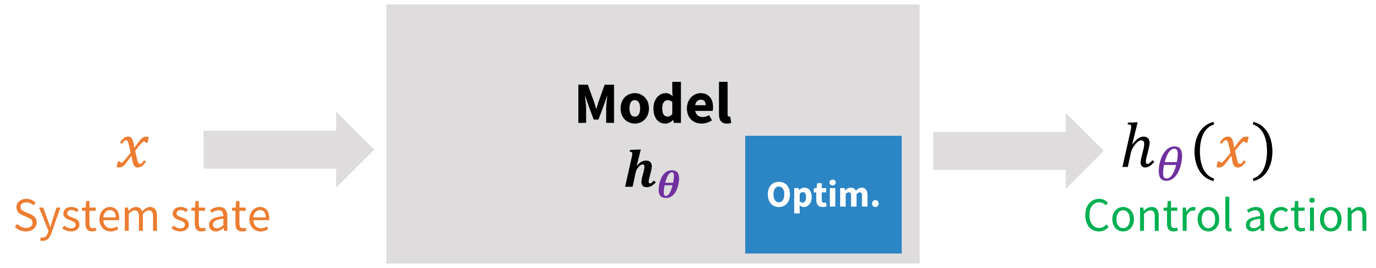

Example: Prediction (e.g., decision-cognizant demand forecasting)

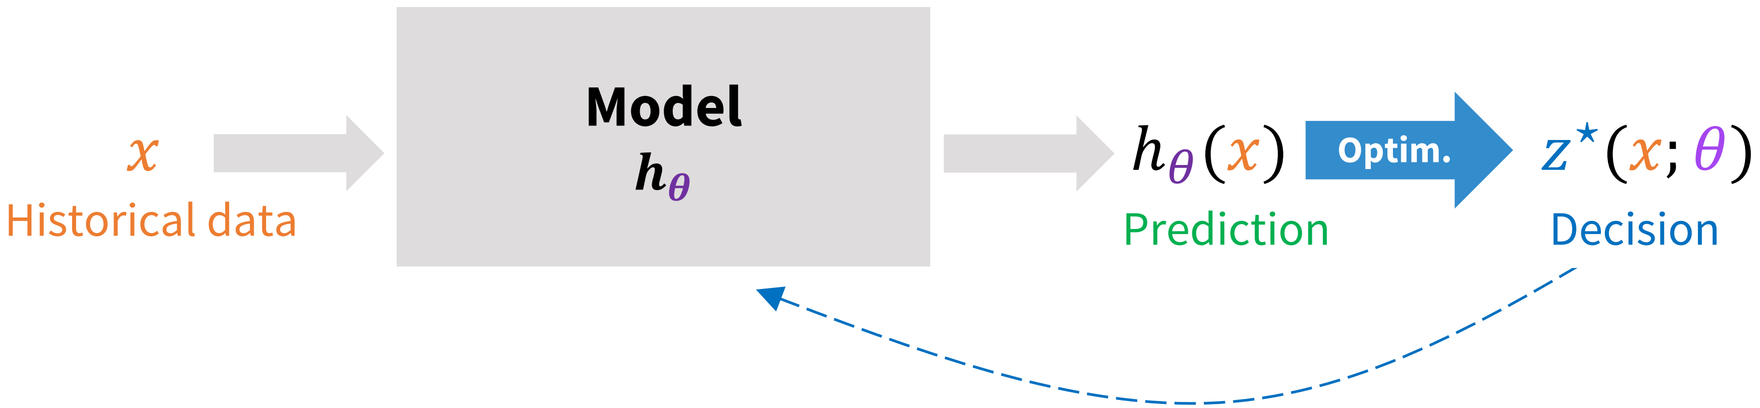

Talk outline

Optimization-in-the-loop ML

Toolkit: Differentiable optimization

Setting: Hard control constraints

Setting: Downstream decision-making

Talk outline

Optimization-in-the-loop ML

Toolkit: Differentiable optimization

Setting: Hard control constraints

Setting: Downstream decision-making

Overview: Differentiable optimization

Motivation: Need for tools to implement optimization-in-the-loop methods

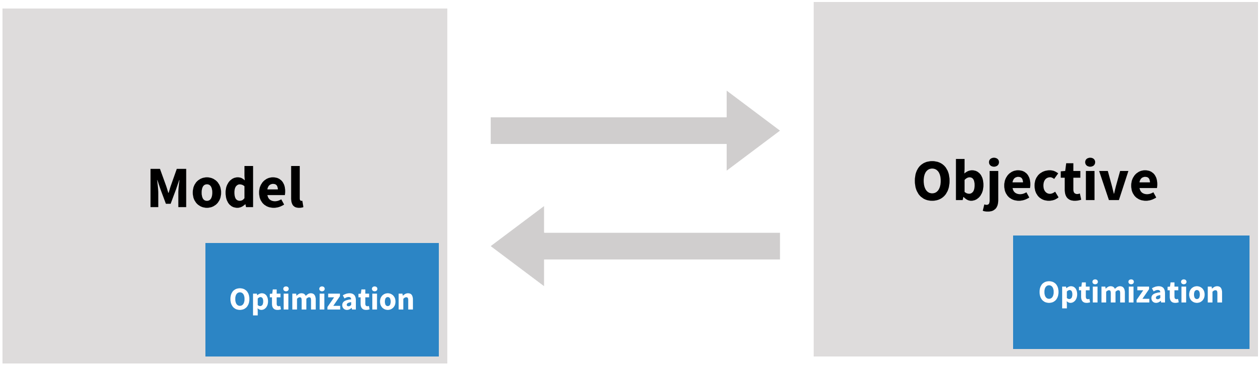

Approach: Differentiable optimization in deep learning

- General framework [GFCASG2016, AK2017, **D**AK2017]
- Additional tools

Background: Deep learning

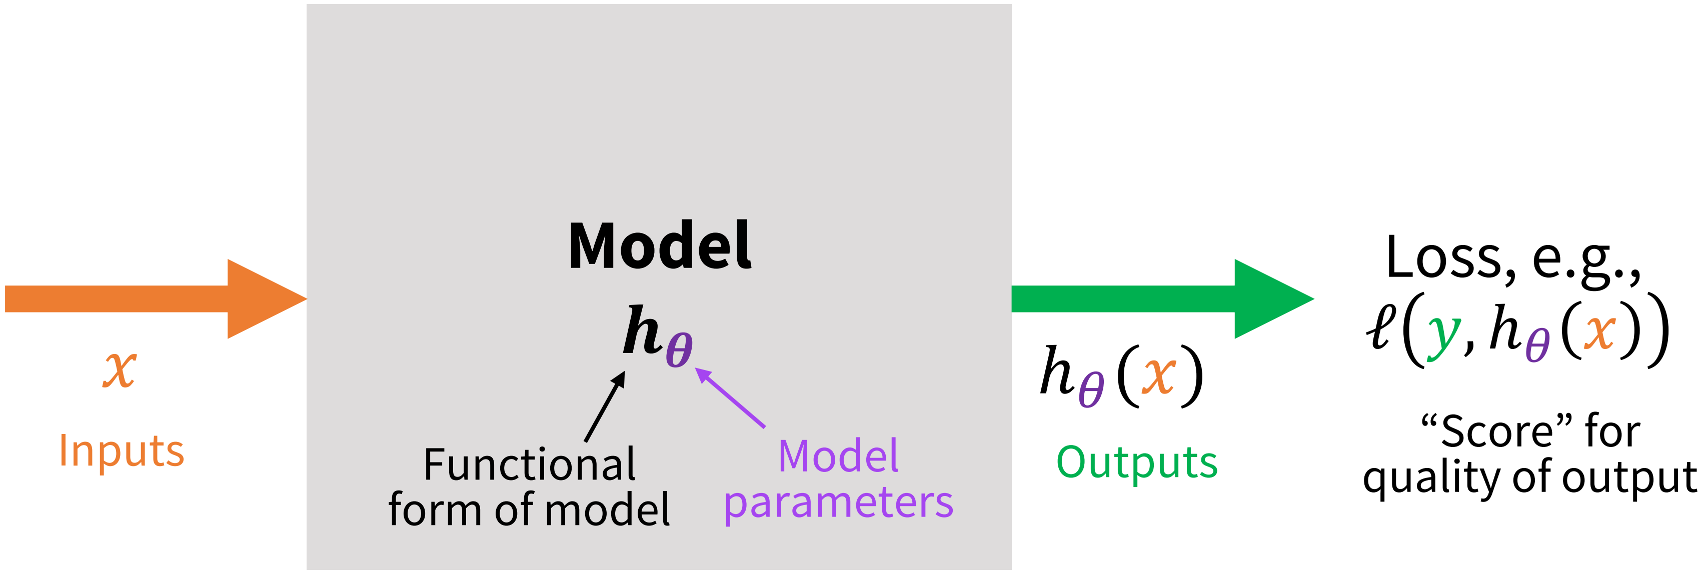

Background: Deep learning

- Neural network h_θ = composition of nonlinear, parameterized functions (*layers*)
- Update parameters θ to minimize loss ℓ using gradients from *backpropagation*
- All components (layers and loss) **must be differentiable**

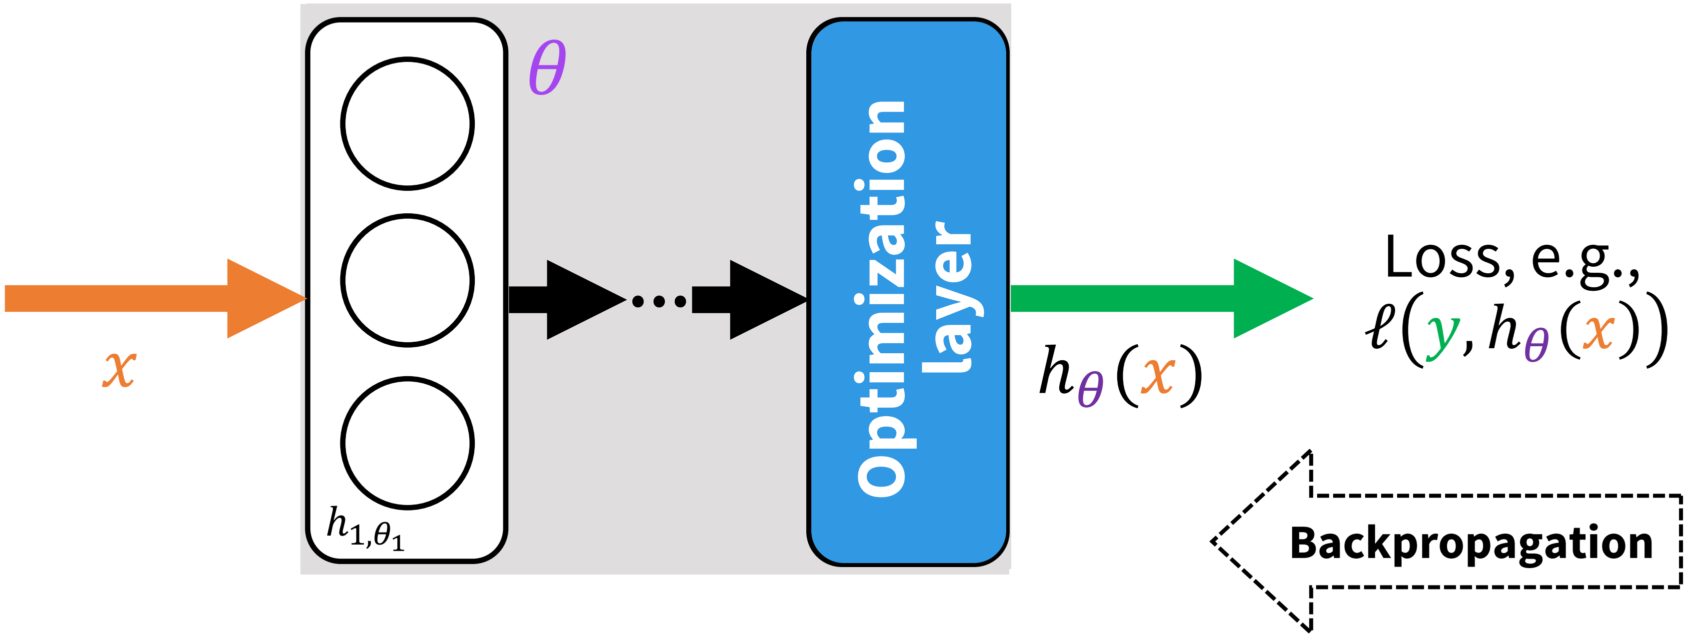

Differentiating through optimization problems

Insight: Apply the implicit function theorem to the KKT optimality conditions

Example optimization problem

$$\begin{array}{ll} \underset{z}{\text{minimize}} & \tfrac{1}{2} z^T Q z + q^T z \\ \text{subject to} & Az = b \\ & Gz \leq h \end{array}$$

Selected KKT optimality conditions

$$\begin{gathered} Qz^\star + q + A^T \nu^\star + G^T \lambda^\star = 0 \\ Az^\star - b = 0 \\ \operatorname{diag}(\lambda^\star)(Gz^\star - h) = 0 \end{gathered}$$

Step 1: Apply implicit function theorem to the KKT conditions

$$\underbrace{\begin{bmatrix} Q & G^T & A^T \\ \operatorname{diag}(\lambda^\star) G & \operatorname{diag}(Gz^\star - h) & 0 \\ A & 0 & 0 \end{bmatrix}}_{\text{Generalized Jacobian of KKT conditions}} \underbrace{\begin{bmatrix} \mathrm{d}z \\ \mathrm{d}\lambda \\ \mathrm{d}\nu \end{bmatrix}}_{\text{Desired gradients}} = - \underbrace{\begin{bmatrix} \mathrm{d}Qz^\star + \mathrm{d}q + \mathrm{d}G^T \lambda^\star + \mathrm{d}A^T \nu^\star \\ \operatorname{diag}(\lambda^\star)\mathrm{d}Gz^\star - \operatorname{diag}(\lambda^\star)\mathrm{d}h \\ \mathrm{d}Az^\star - \mathrm{d}b \end{bmatrix}}_{\text{Gradients of problem parameters}}$$

Step 2: Use "Jacobian-vector trick" for efficient backpropagation

Brandon Amos and J. Zico Kolter. "OptNet: Differentiable optimization as a layer in neural networks." *ICML 2017.*
Priya L. Donti, Brandon Amos, and J. Zico Kolter. "Task-based end-to-end model learning in stochastic optimization." *NeurIPS 2017.*

Follow-on work in differentiable optimization

[**D**AK2017, AK2017]: KKT differentiation techniques for convex optimization problems

Many additional tools since then:

- Combinatorial optimization [DK2017, TSK2018, WDT2018]
- AC optimal power flow [**D**AK2018]
- Disciplined convex programs [AABBDK2019]
- Maximum satisfiability problems [W**D**WK2019]
- Additional optimization problems [GHC2019]

Powerful toolkit for optimization-in-the-loop ML in the context of deep learning

Talk outline

Optimization-in-the-loop ML

Toolkit: Differentiable optimization

Setting: Hard control constraints

Setting: Downstream decision-making

Overview: Enforcing hard control constraints

Motivation: Need for well-performing control methods that also guarantee enforcement of hard constraints

Approach: Optimization-in-the-loop reinforcement learning (RL) techniques with guaranteed enforcement of hard constraints

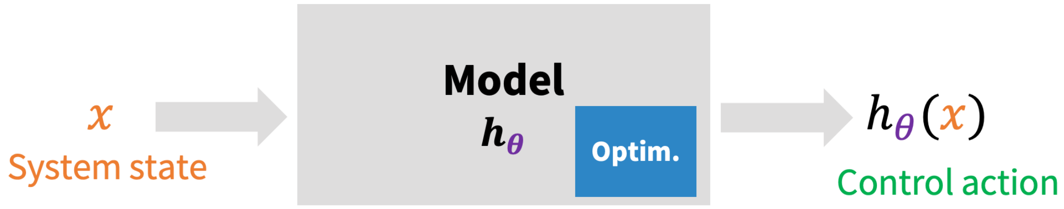

Settings:

- Asymptotic stability in power grids [**D**RFK2021]

- Realistic-scale building control [C**D**BKB2021]

Deep reinforcement learning vs. robust control

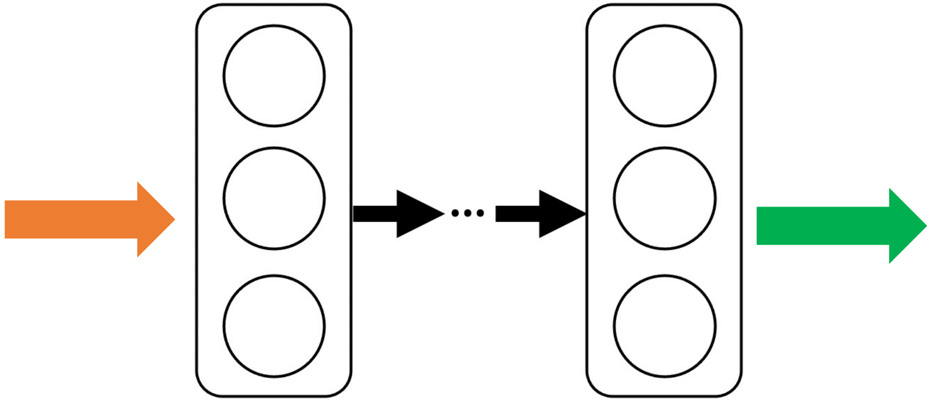

Deep RL

Pro: Expressive, well-performing policies
Con: Potential (catastrophic) failures

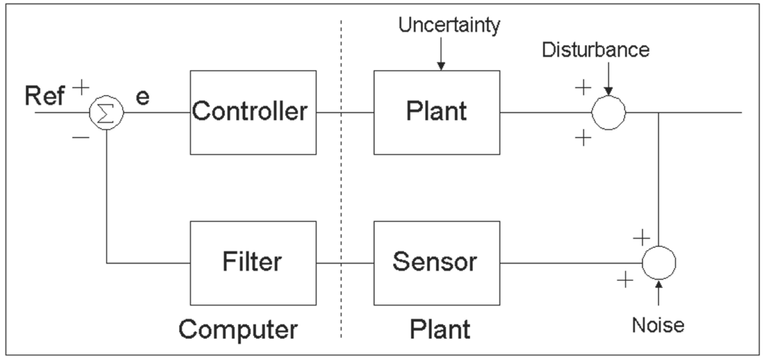

Robust control

Pro: Provable stability guarantees
Con: Simple policies (e.g., linear)

Can we improve performance while still guaranteeing stability?

Priya L. Donti, Melrose Roderick, Mahyar Fazlyab, and J. Zico Kolter. "Enforcing robust control guarantees within neural network policies." *International Conference on Learning Representations (ICLR) 2021.*

Differentiable projection onto stabilizing actions

Deep learning-based policy with **provable robustness guarantees**, trainable using standard reinforcement learning approaches

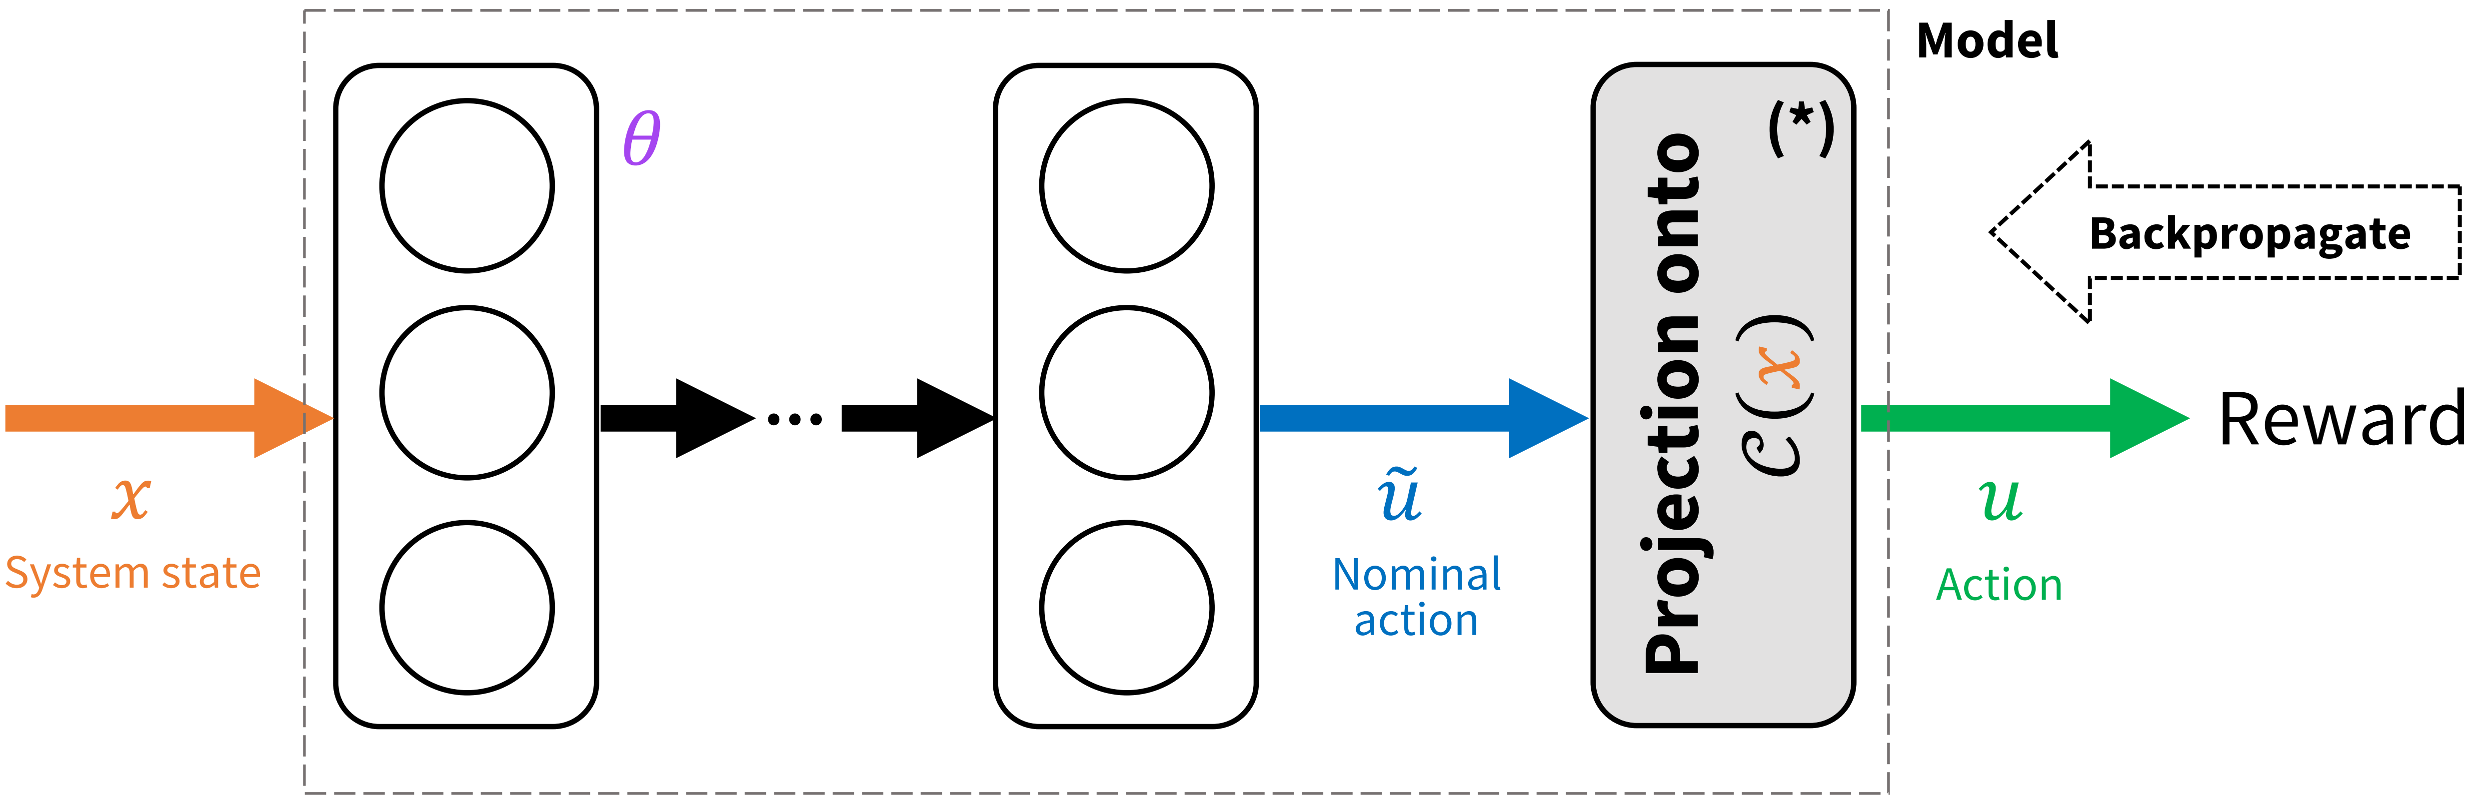

Details: Finding a set of stabilizing actions

Insight: Find a set of actions that are guaranteed to satisfy relevant Lyapunov stability criteria at a given state, even under worst-case conditions

Given the following (from robust control)**:**

- Uncertainty model: e.g., $\dot{x}(t) \in Ax(t) + Bu(t) + Gw(t)$ s.t. $\|w(t)\|_2 \le \|Cx(t) + Du(t)\|_2$
- Lyapunov function V obtained via robust control synthesis
- Exponential stability criterion: $\dot{V}\big(x(t)\big) \le -\alpha V\big(x(t)\big), \forall x \neq 0$

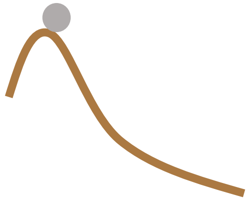

Find: For given x, set of actions satisfying exponential stability criterion even in worst case

$$\mathcal{C}(x) \equiv \left\{ u : \left(\sup_{w \,:\, \|w\|_2 \le \|Cx + Du\|_2} \dot{V}(x) \right) \le -\alpha V(x) \right\}$$

$$\Rightarrow \{ u : \|k_1(x) + Du\|_2 \le k_2(x) + k_3(x)^T u \}$$

Convex (non-empty) set in $u(t)$

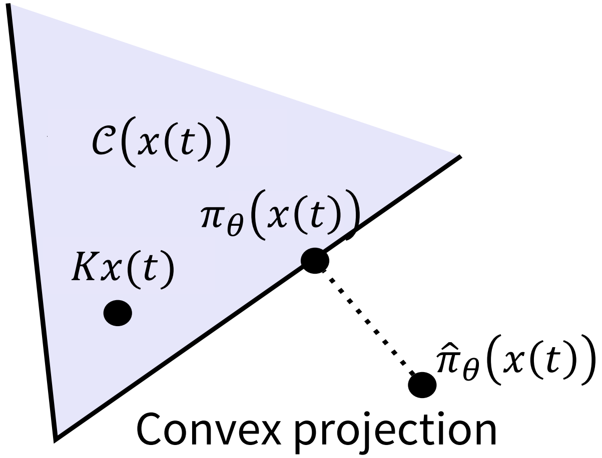

Note: t-dependence has been dropped for brevity

Illustrative results: Synthetic NLDI system

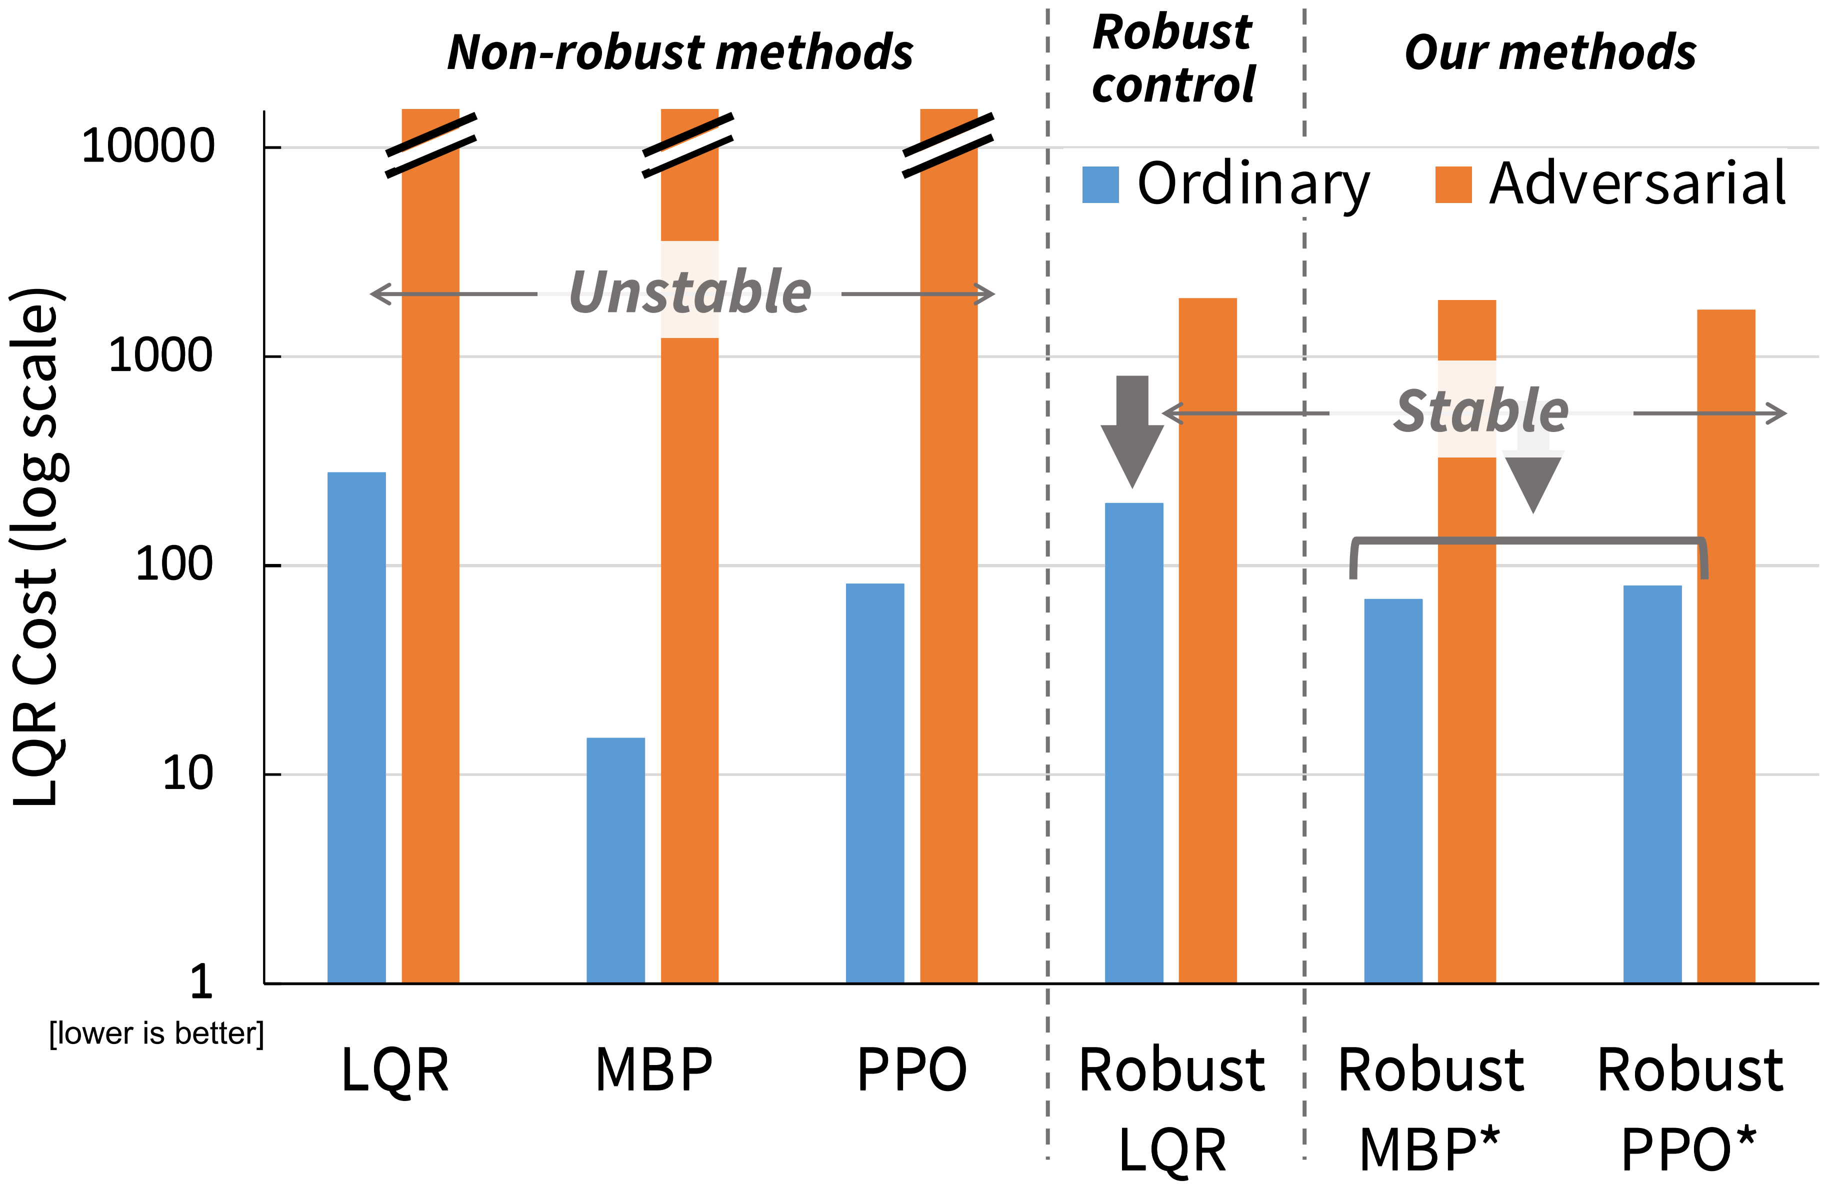

Improved "average-case" performance over robust baselines

Provably stable under "worst-case" dynamics (unlike non-robust baselines)

Downside: Speed / computational cost

Energy-efficient heating and cooling

Goal: Control the HVAC supply water temperature to minimize energy use, while respecting equipment constraints and maintaining thermal comfort

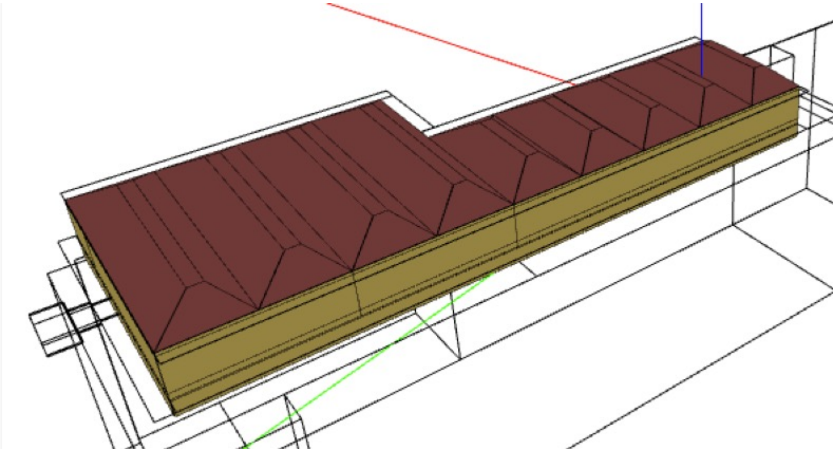

Intelligent Workplace
Margaret Morrison Hall, 4th Floor
(✣ Zhang & Lam, 2018)

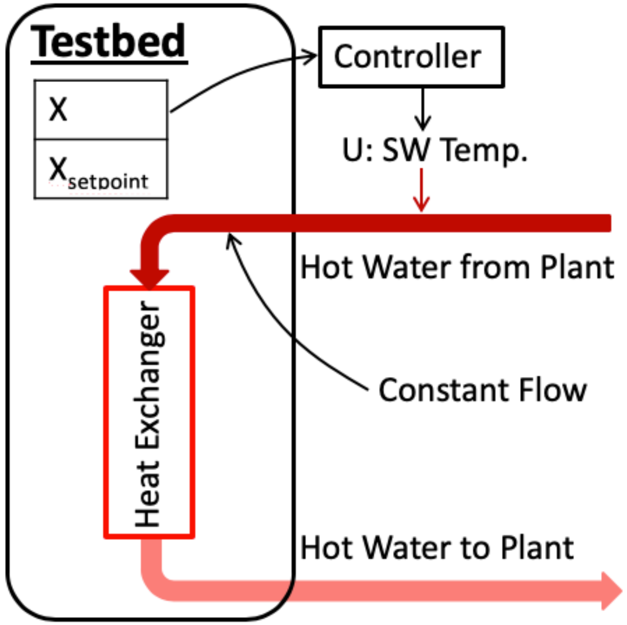

HVAC Schematic

Bingqing Chen*, **Priya L. Donti***, Kyri Baker, J. Zico Kolter, and Mario Berges. "Enforcing Policy Feasibility Constraints through Differentiable Projection for Energy Optimization." *ACM International Conference on Future Energy Systems (ACM e-Energy) 2021*.

Differentiable projection onto feasible actions

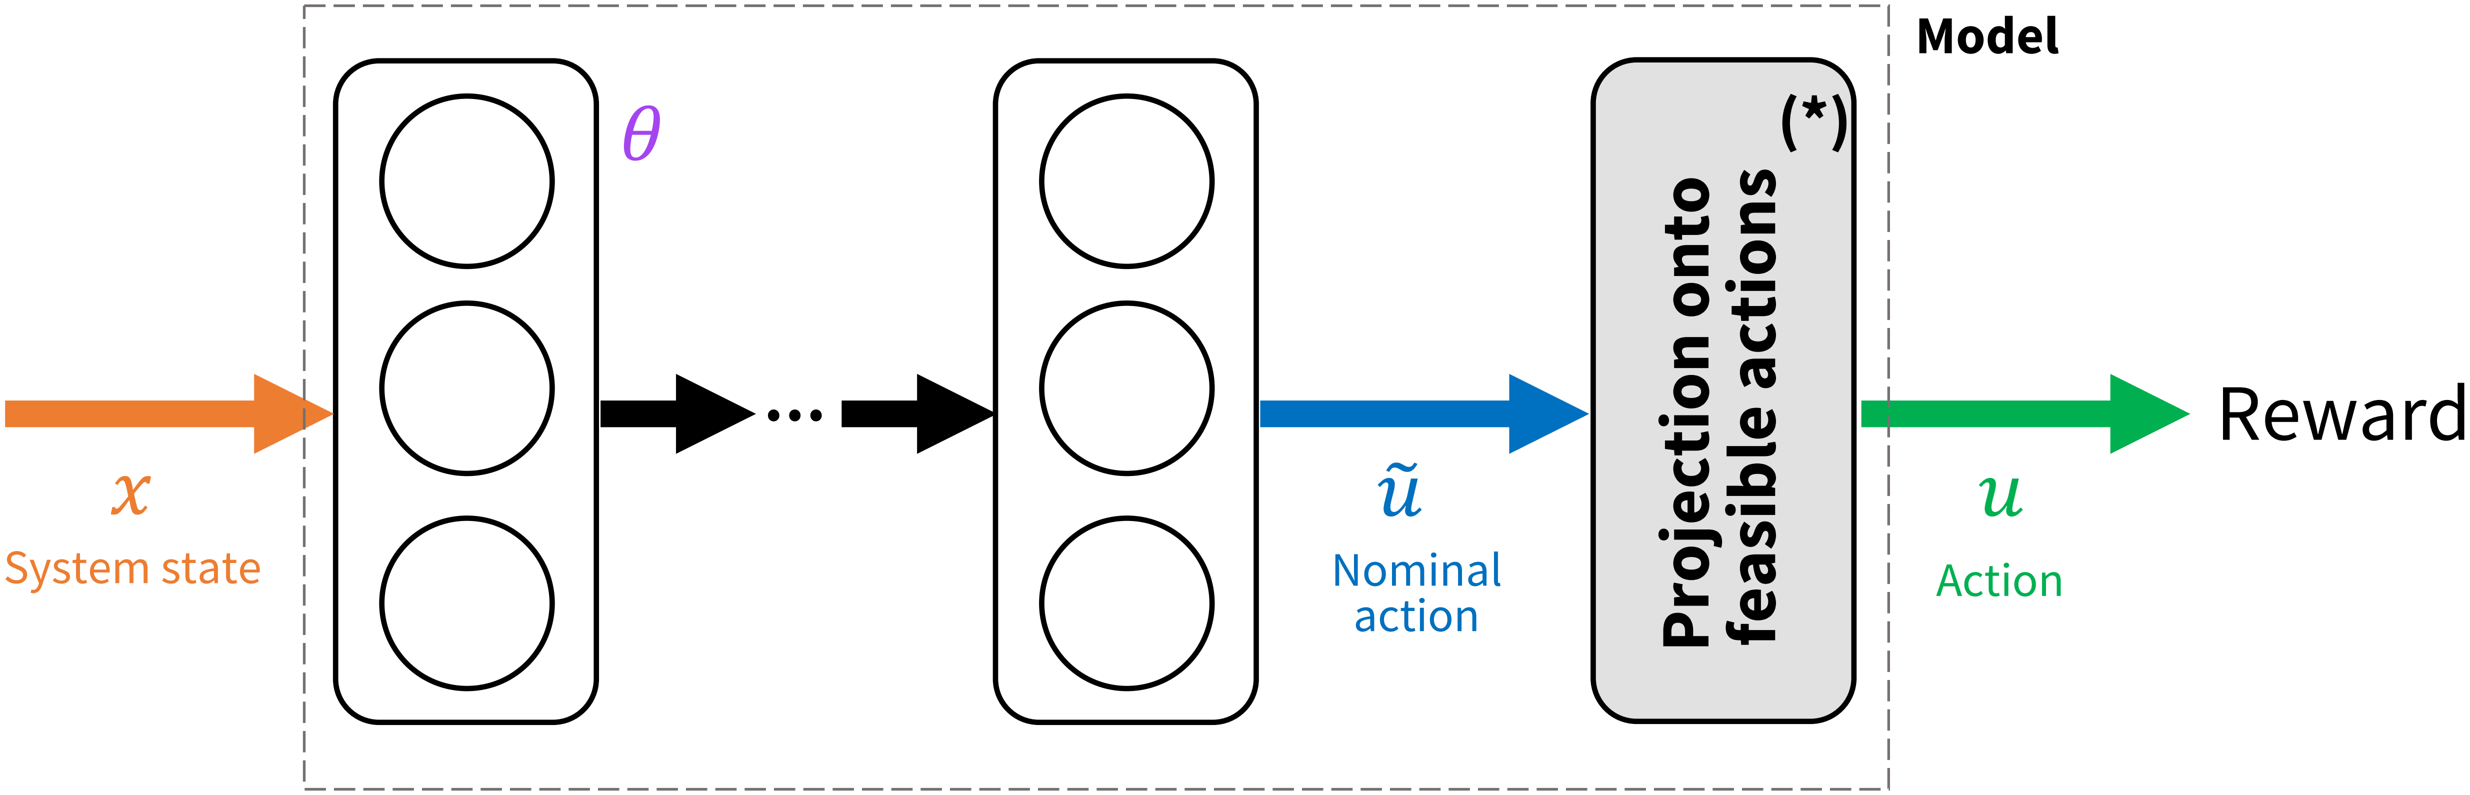

Summary: Enforcing hard control constraints

Motivation: Need for well-performing control methods that also guarantee enforcement of hard constraints

Settings:

- Asymptotic stability in power grids [**D**RFK2021]
- Realistic-scale building control [C**D**BKB2021]

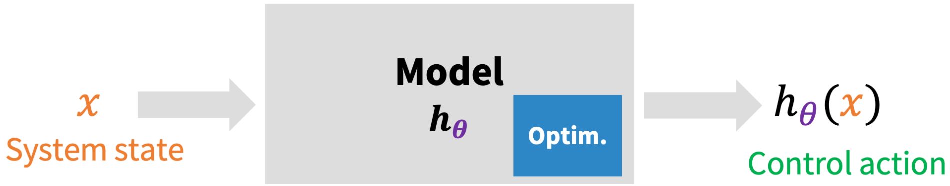

Insight: Project outputs of neural network onto a set of "safe" actions

- Obtain safe actions using domain knowledge
- Differentiable projection (optimization layer) = end-to-end training

Future directions:

- Additional paradigms for bridging RL and robust control
- Improving computational costs

Talk outline

Optimization-in-the-loop ML

Toolkit: Differentiable optimization

Setting: Hard control constraints

Setting: Downstream decision-making

Overview: Incorporating downstream decision-making

Motivation: Predictive methods operate within some larger decision-making process but do not often take this into account, potentially leading to critical mistakes.

Approach: Construction of decision-cognizant ("task-based") models via optimization-in-the-loop learning

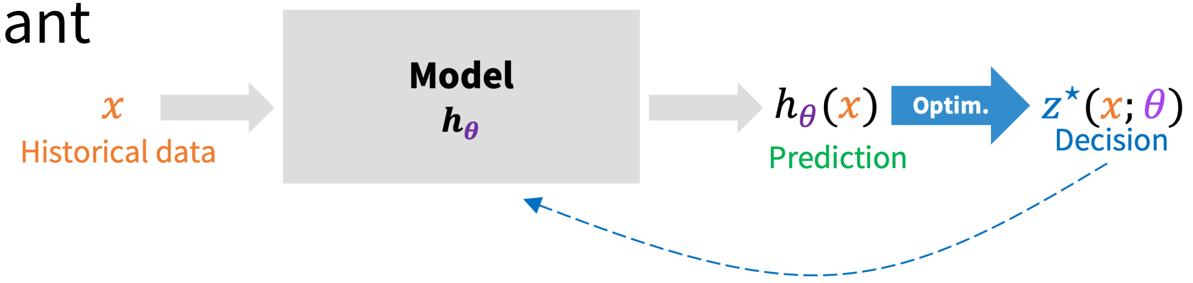

Settings:

- Decision-cognizant electricity demand forecasting [**D**AK2017]
- Approximating AC optimal power flow [**D**RK2021]

Decision-cognizant demand forecasting

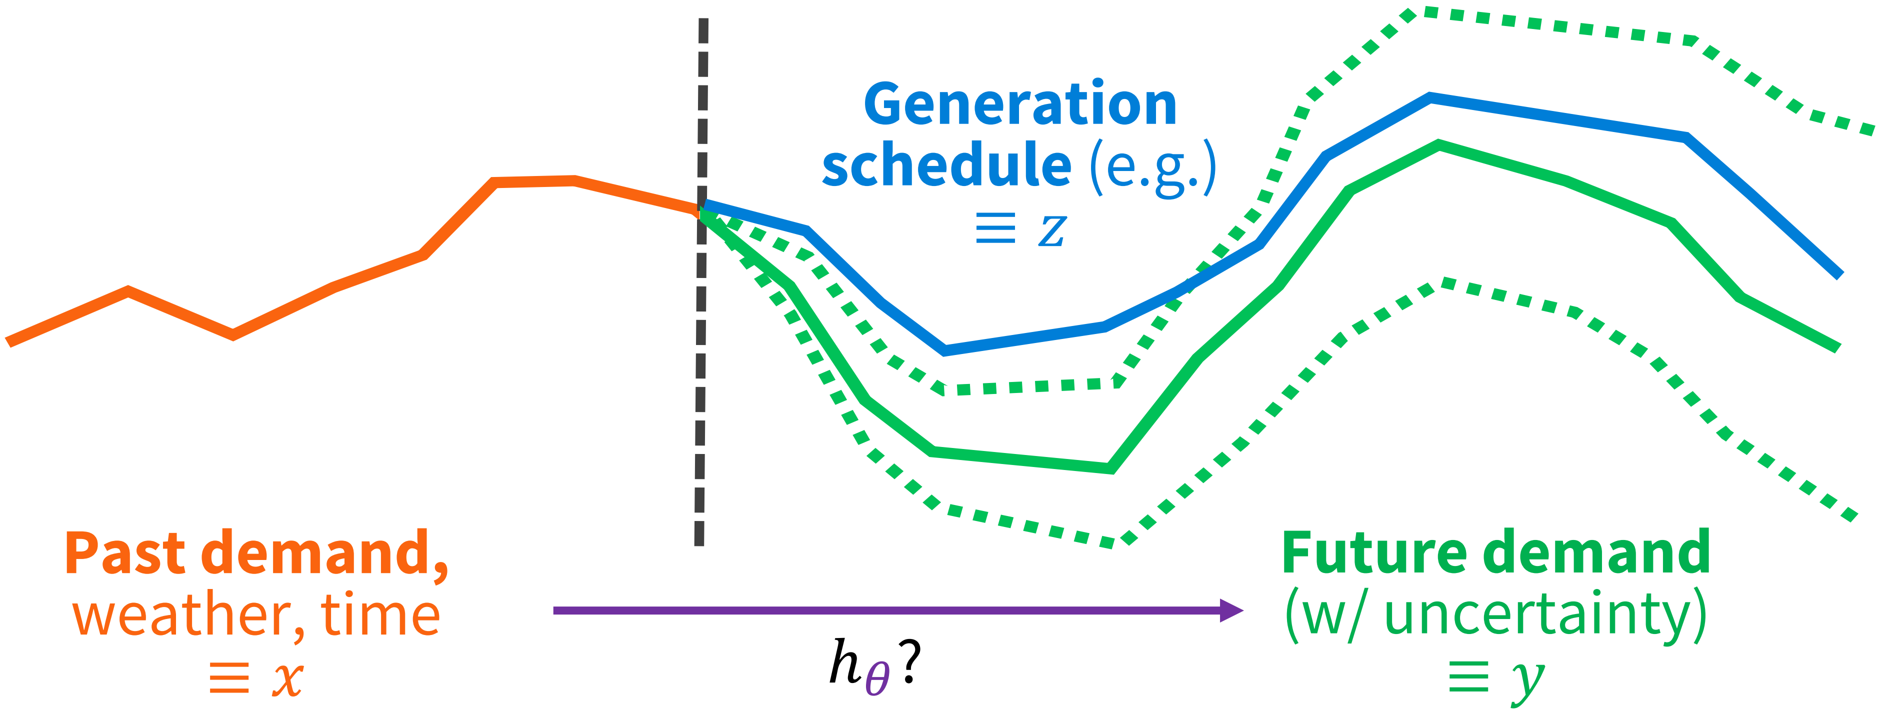

Goal: Optimize for quality of generation schedule when we observe actual demands

$$\underset{\theta}{\text{minimize}}\ f_c(y, z^\star(x; \theta))$$

Priya L. Donti, Brandon Amos, and J. Zico Kolter. "Task-based end-to-end model learning in stochastic optimization." *Conference on Neural Information Processing Systems (NeurIPS)* 2017.

Decision-cognizant model

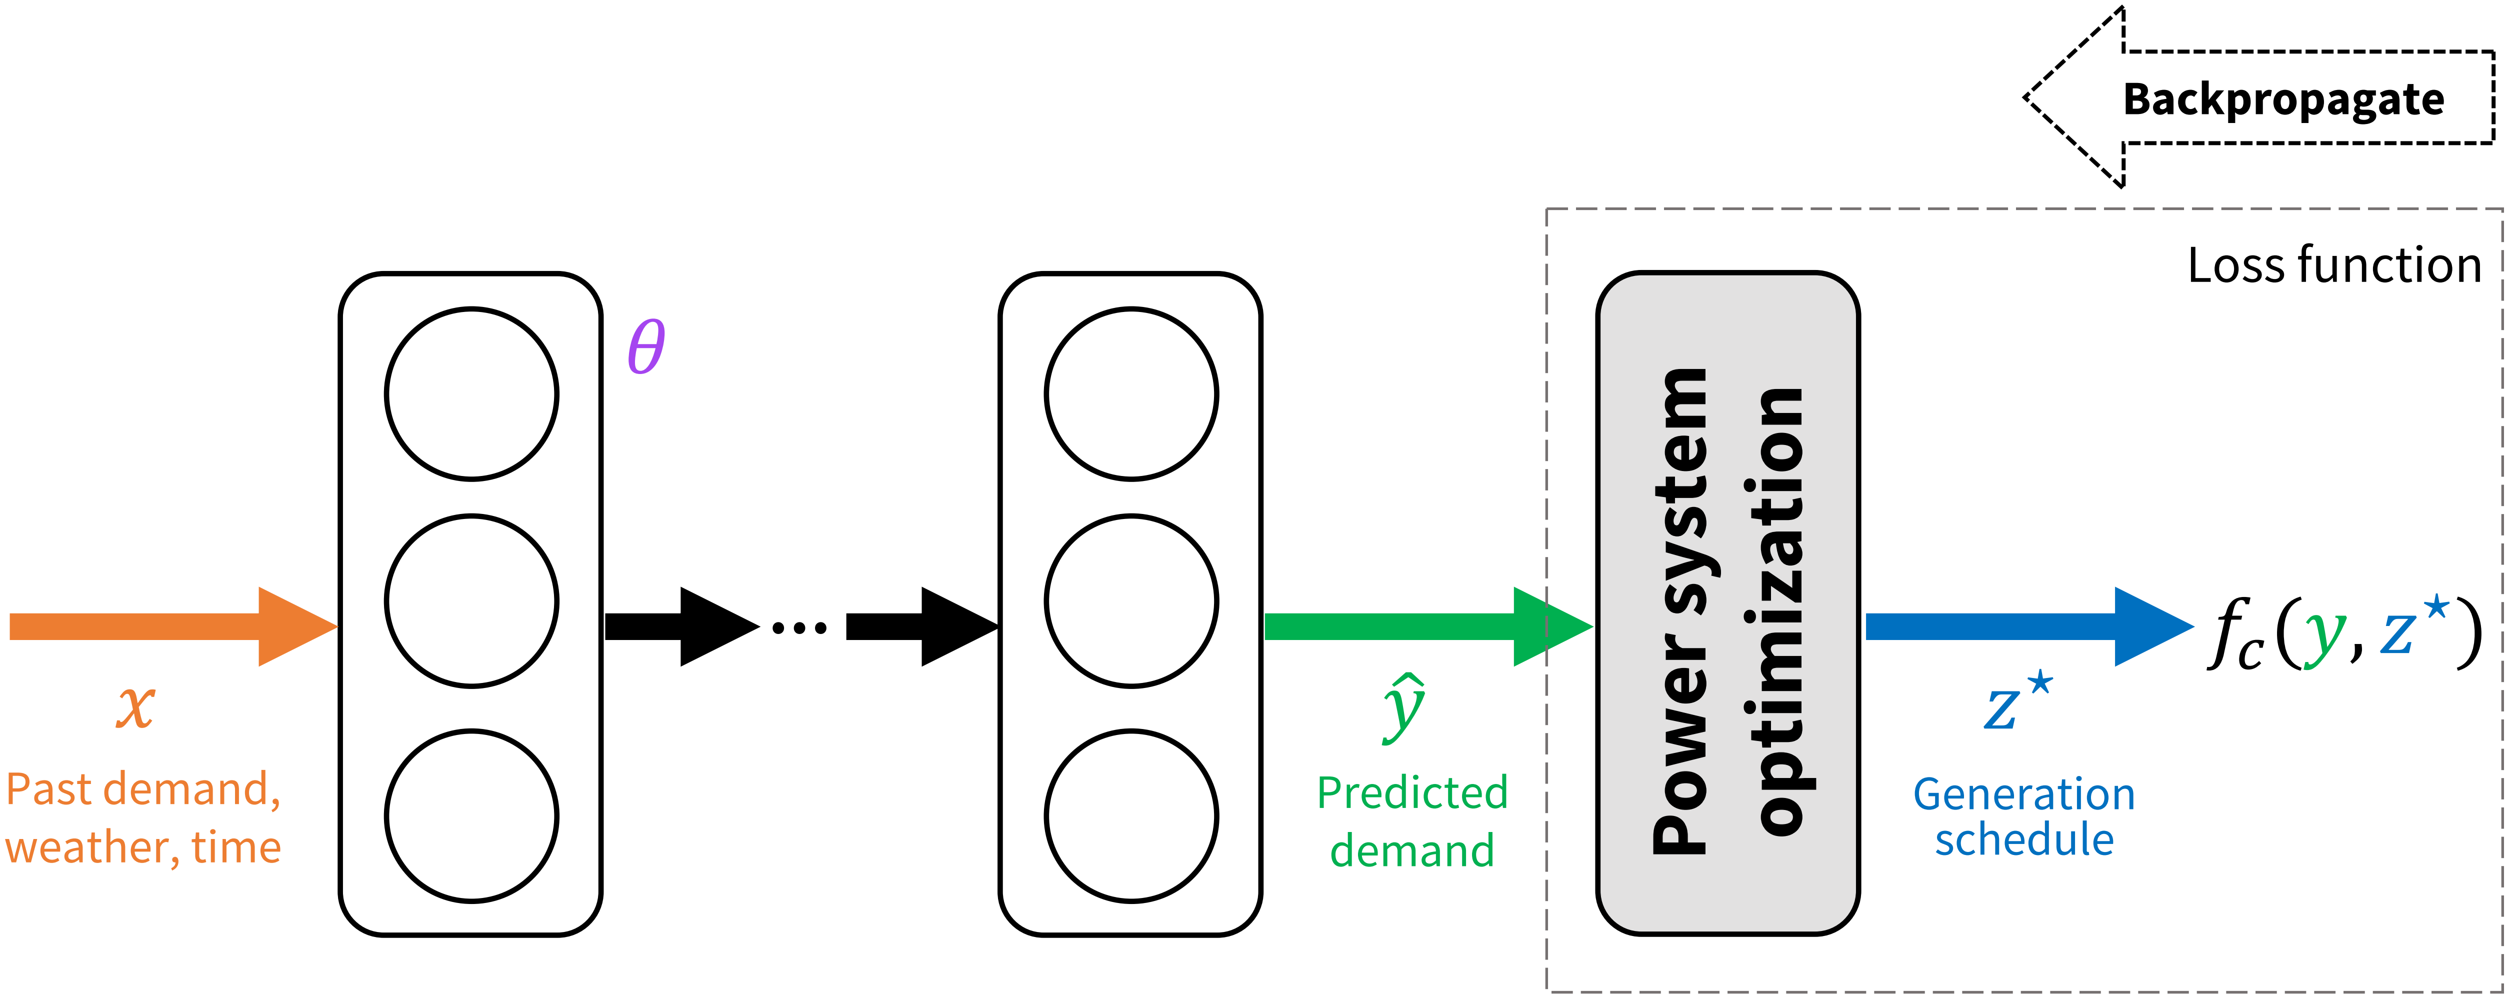

Decision-cognizant approach can dramatically improve generation scheduling outcomes

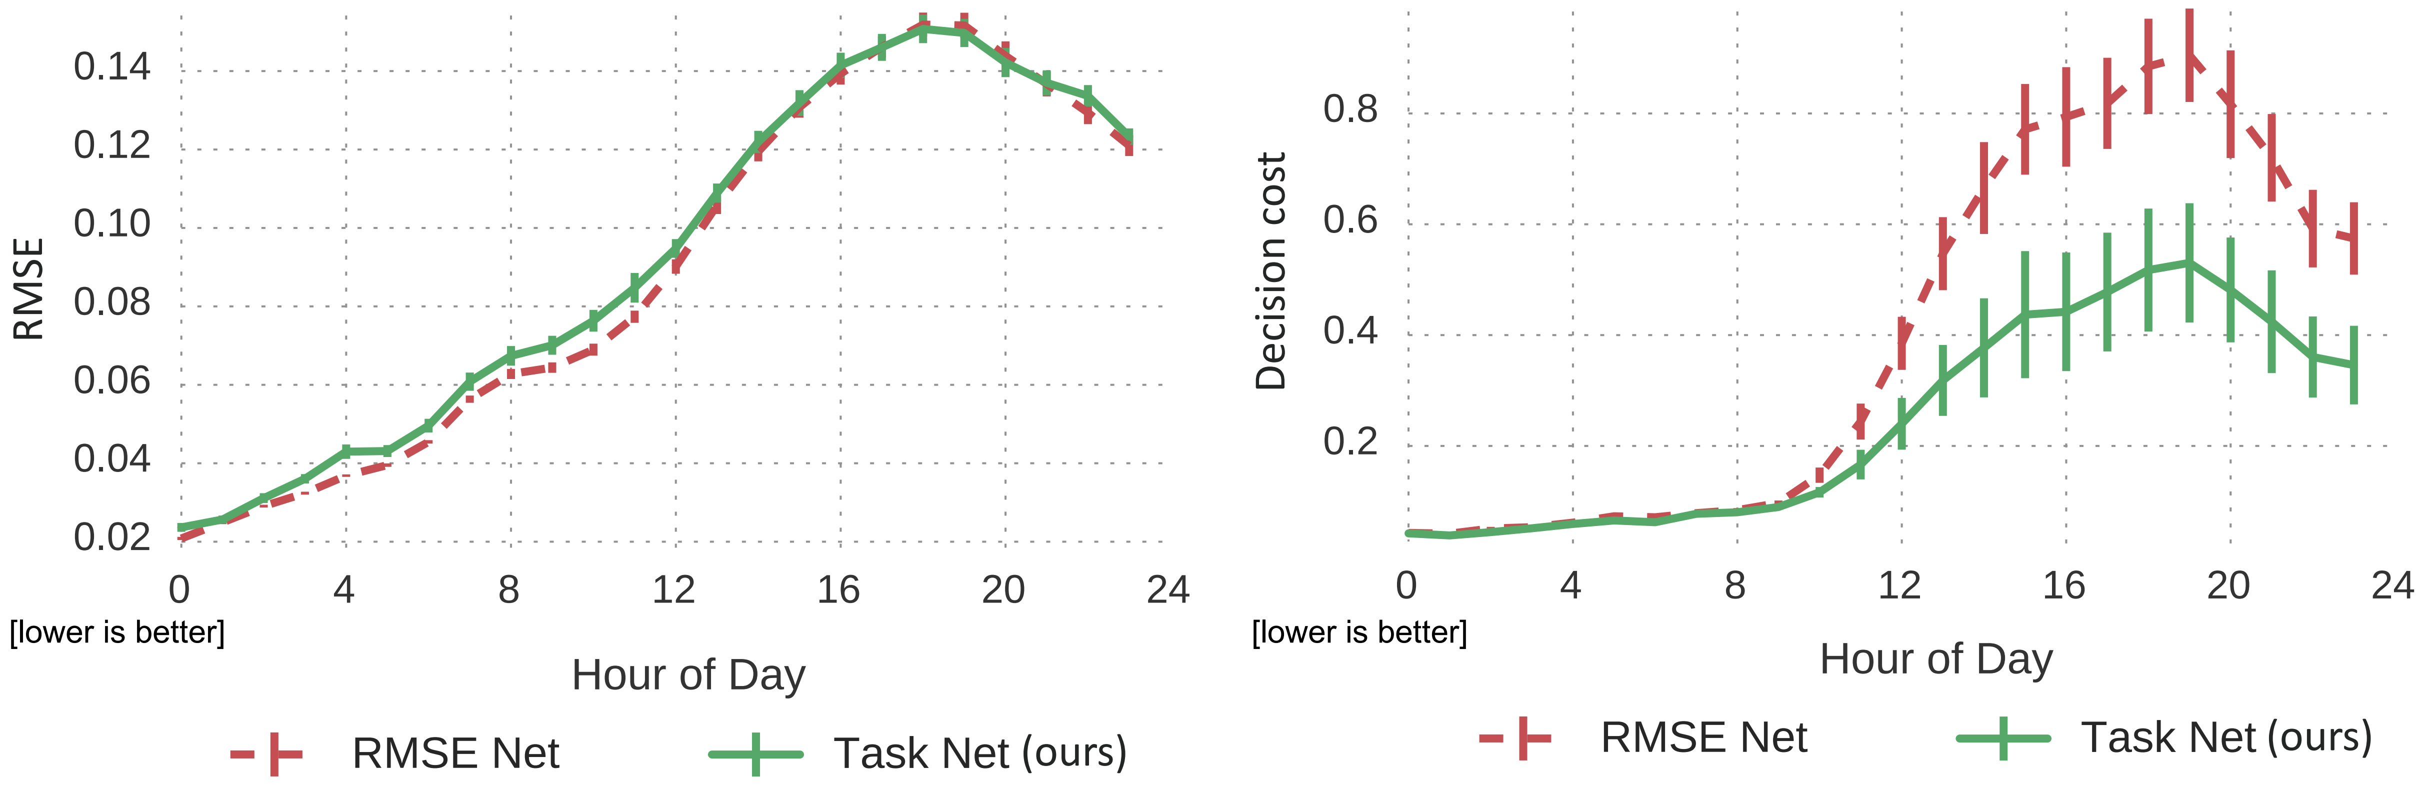

Decision-cognizant approach gives ~**39%** improvement in decision cost.

Approximating AC optimal power flow

Goal: Provide fast, feasible approximations to AC optimal power flow (ACOPF)

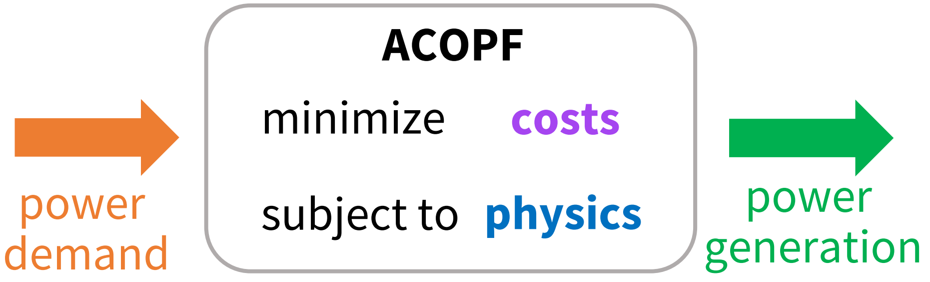

Approach:

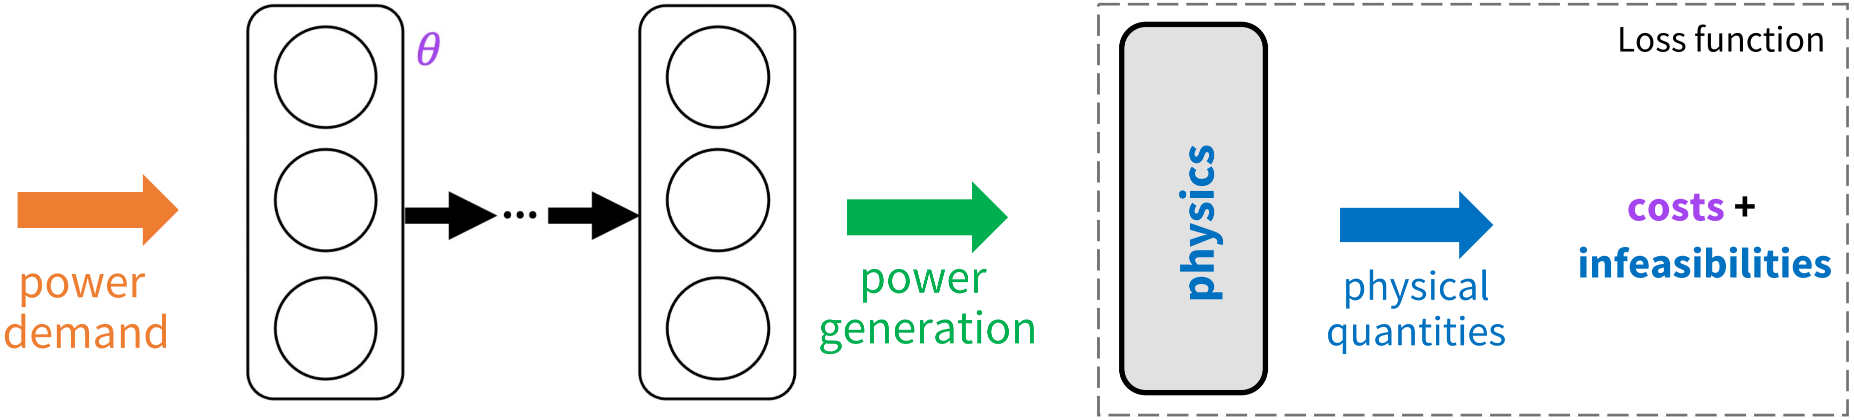

Results (57-bus test case): High-quality solutions 10x faster than baseline optimizer

Priya L. Donti*, David Rolnick*, and J. Zico Kolter. "DC3: A learning method for optimization with hard constraints.” *International Conference on Learning Representations (ICLR) 2021.*

Summary: Incorporating downstream decision-making

Motivation: Predictive methods operate within some larger decision-making process but do not often take this into account, potentially leading to critical mistakes.

Settings:

- Electricity demand forecasting [**D**AK2017]
- Approximating ACOPF [**D**RK2021]

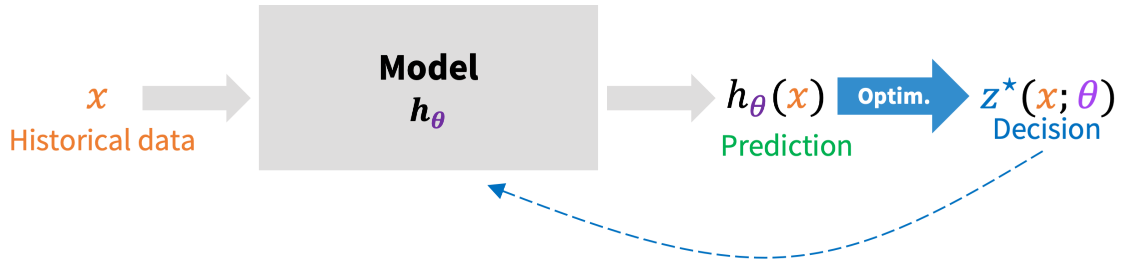

Insight: Incorporate knowledge of downstream decision-making (or physics) into the loss function, via differentiable optimization.

Future directions:

- Incorporating a wider range of decision-making paradigms
- Understanding tradeoffs between task-agnostic vs. task-based models

Summary

- **Optimization-in-the-loop ML (framework)**, via differentiable optimization in deep learning

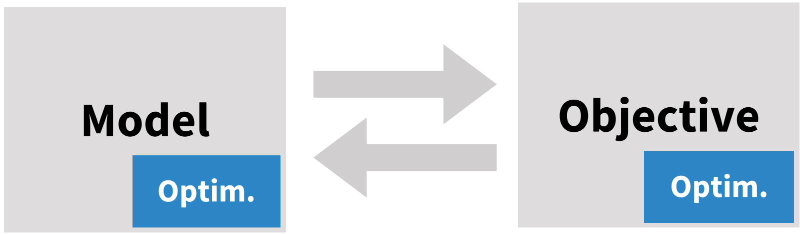

- **Enforcing hard control constraints:** RL with provable robustness / constraint enforcement
    - Asymptotic stability (power grids)
    - Operational constraints (HVAC in buildings)

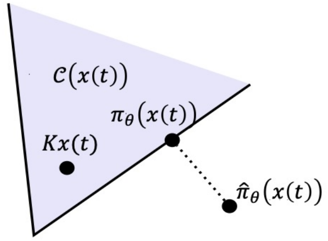

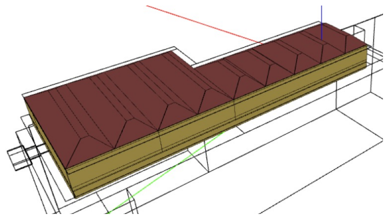

- **Incorporating downstream decision-making:** Decision-cognizant predictive models
    - Electricity demand forecasting
    - Approximate power system optimization

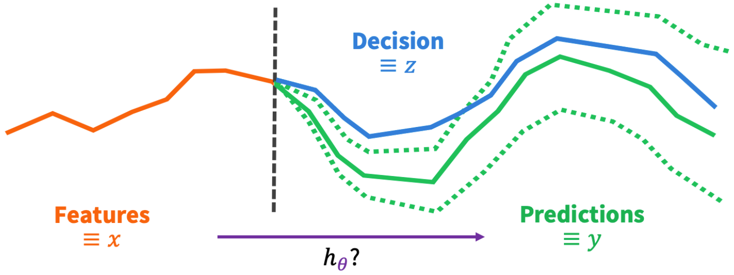

Priya L. Donti: donti@mit.edu