PSCC2024

Implicit layers: A toolkit for AI in power systems

Priya L. Donti
Assistant Professor
Dept. Electrical Engineering & Computer Science
Laboratory for Information & Decision Systems
Massachusetts Institute of Technology (MIT)

Power systems problems involve physics, hard constraints, and decision-making

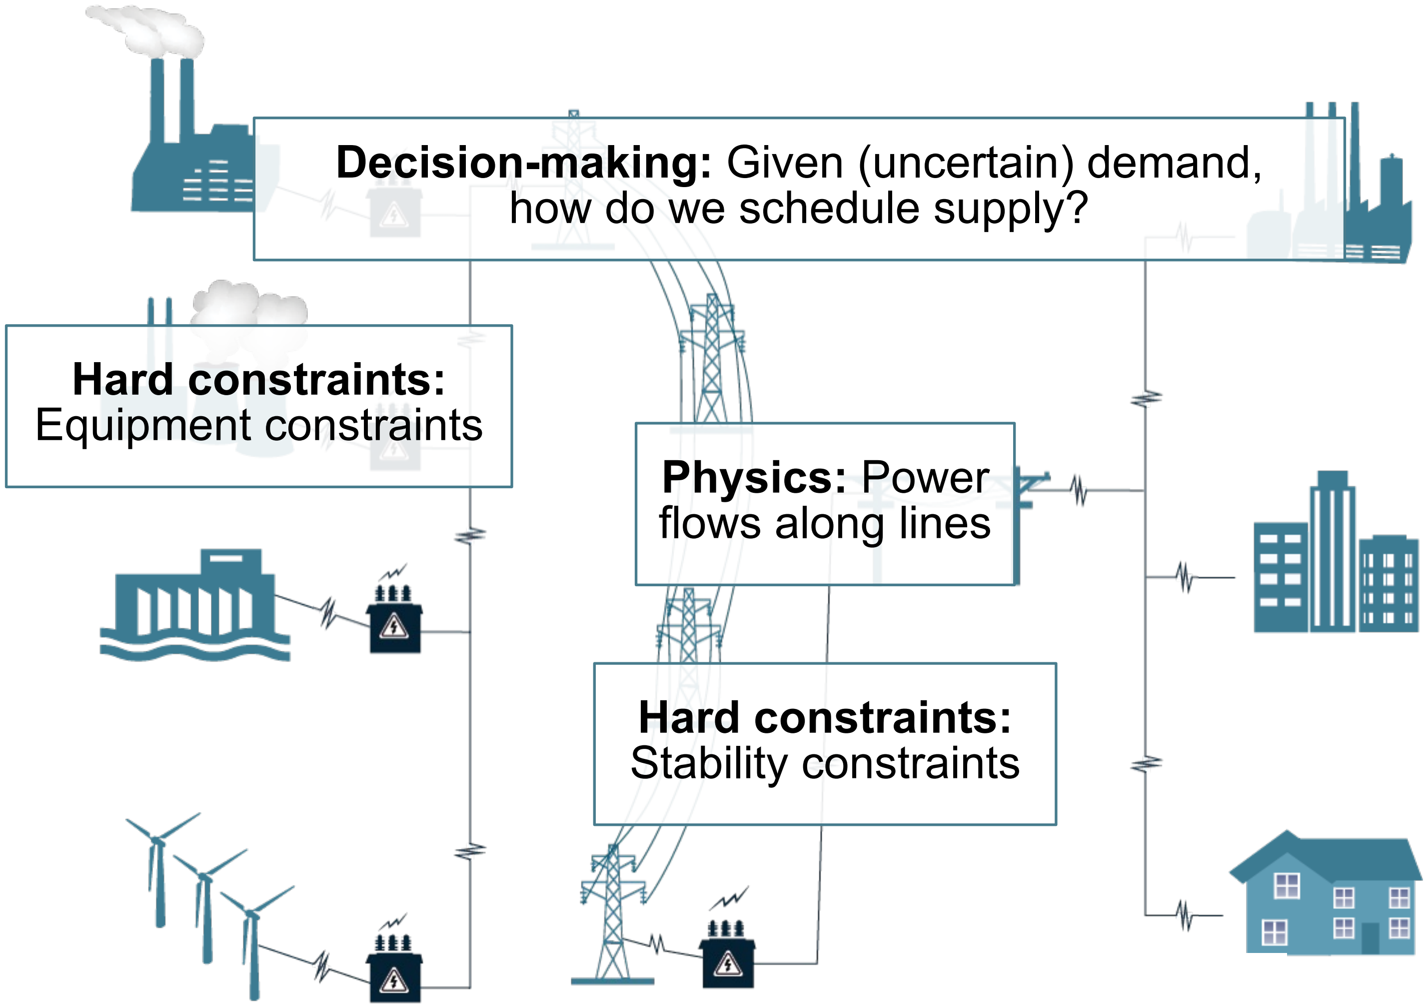

Trad. optimization & control
- Satisfies (many) constraints
- Struggles with speed / scale

Machine learning (ML)
- Fast and scalable
- Struggles with constraints

Figure adapted from: US Congressional Budget Office

Implicit layers in deep learning

Toolkit for developing deep learning methods that incorporate knowledge of physics, hard constraints, or downstream decision-making procedures

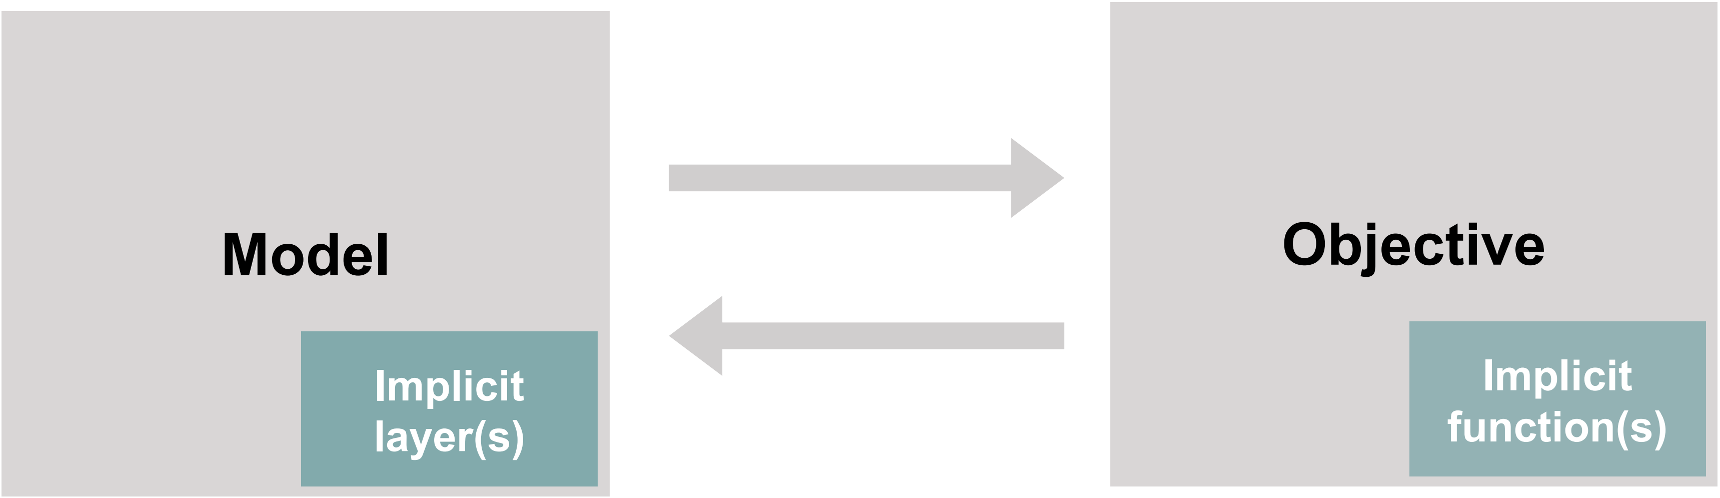

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Deep learning is differentiable function composition

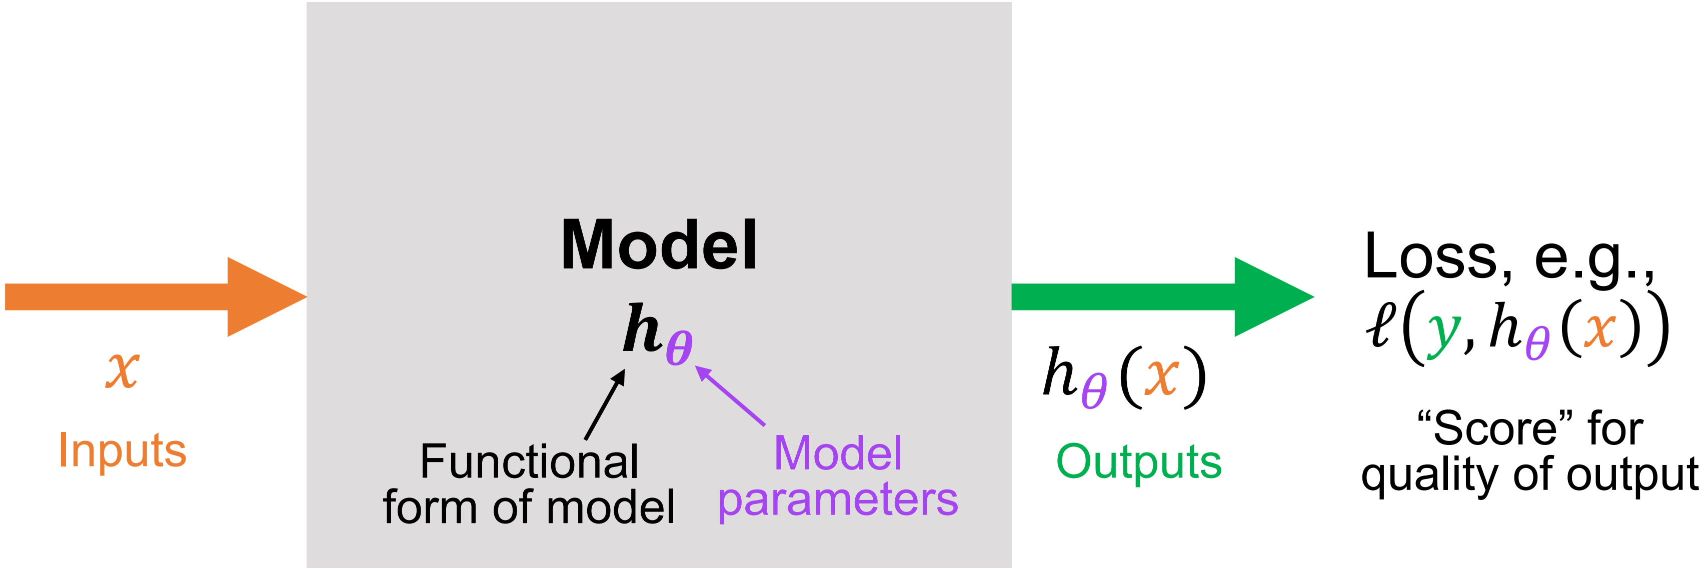

Deep learning is differentiable function composition

- 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**

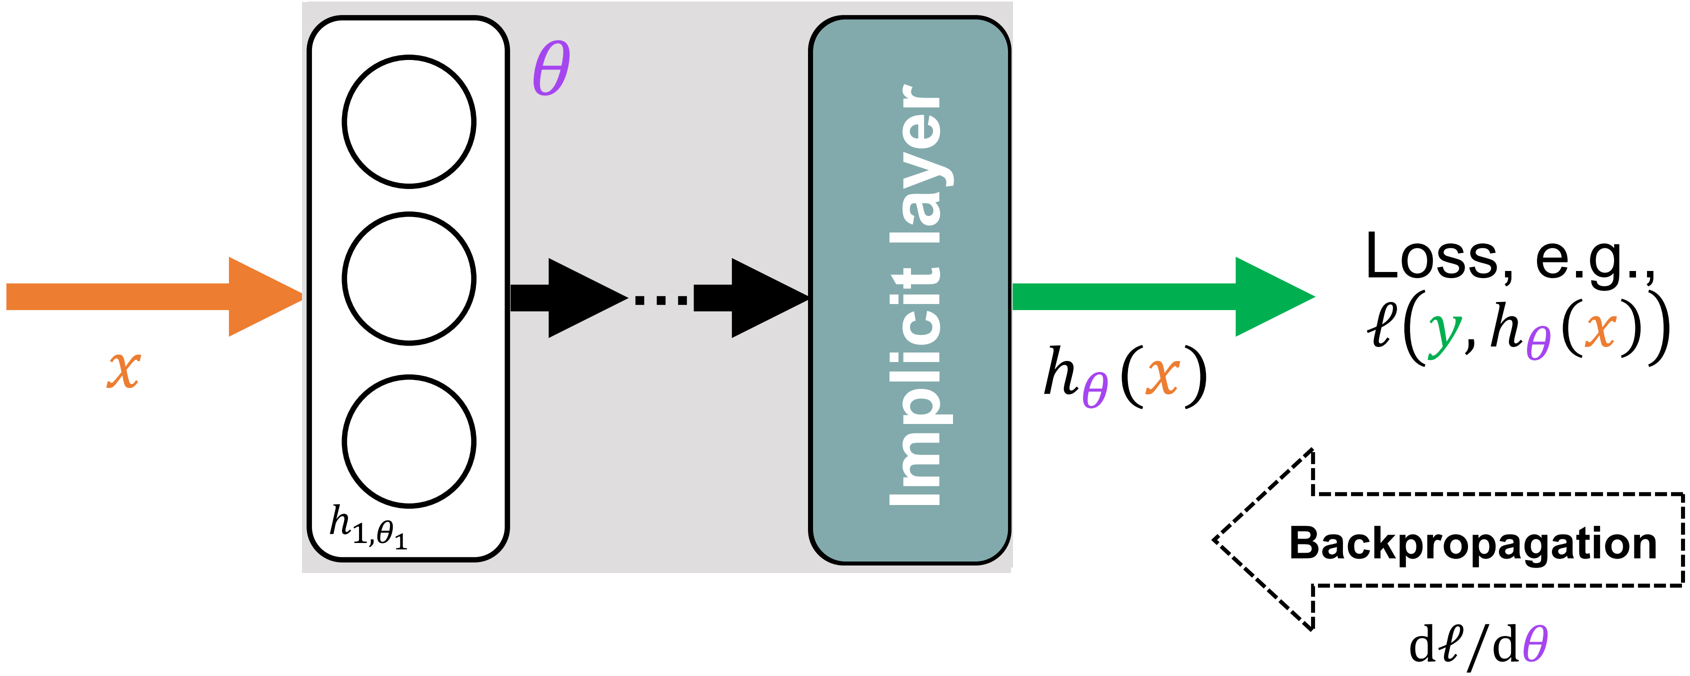

Backpropagation and gradient descent (one layer)

Forward pass:

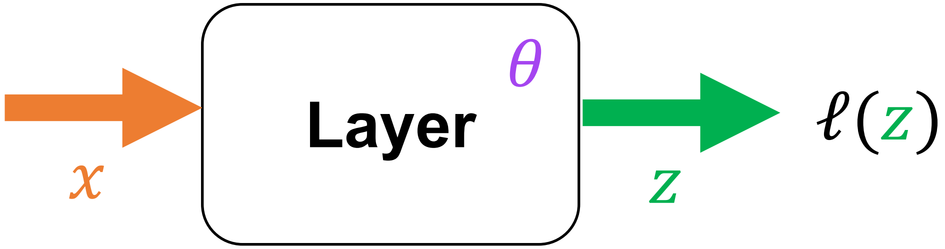

Backward pass:

$$\frac{d\ell}{dz^\star}, \frac{d\ell}{d\theta}, \frac{d\ell}{dx}$$

via chain rule

$$\frac{d\ell}{d\theta} = \frac{d\ell}{dz^\star}\frac{dz^\star}{d\theta}$$

$$\frac{d\ell}{dx} = \frac{d\ell}{dz^\star}\frac{dz^\star}{dx}$$

Update (gradient descent):

$$\theta \leftarrow \theta - \alpha \frac{d\ell}{d\theta}$$

Explicit vs. implicit layer

	Explicit layer	**Implicit layer**
Forward pass: 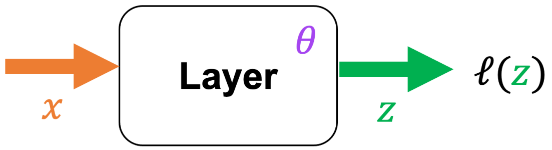	$z = f(x, \theta)$ [e.g., $z = \sigma(\theta^T x + \theta_0)$]	Find z such that $g(z, x, \theta) = 0$ [e.g., power flow]
Backward pass: $\frac{\mathrm{d}\ell}{\mathrm{d}\theta} = \frac{\mathrm{d}\ell}{\mathrm{d}z^\star}\frac{\mathrm{d}z^\star}{\mathrm{d}\theta}$ $\frac{\mathrm{d}\ell}{\mathrm{d}x} = \frac{\mathrm{d}\ell}{\mathrm{d}z^\star}\frac{\mathrm{d}z^\star}{\mathrm{d}x}$	$\frac{\mathrm{d}z^\star}{\mathrm{d}x} = \frac{\mathrm{d}f(x, \theta)}{\mathrm{d}x}$	Find $\mathrm{d}z^\star/\mathrm{d}x$ such that $\frac{\mathrm{d}g(z^\star, x, \theta)}{\mathrm{d}x} = 0$ by using implicit function theorem at a solution point

See also: Zico Kolter, David Duvenaud, and Matt Johnson. “Deep Implicit Layers - Neural ODEs, Deep Equilibirum Models, and Beyond.” Tutorial at *NeurIPS 2020*. https://implicit-layers-tutorial.org/

Example: Differentiable quadratic programming layer

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

QP layer (output z, all else are inputs/params)

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

Selected KKT optimality conditions

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

Step 1: Apply implicit function theorem to the KKT conditions

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

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.*

Many types of implicit layers

Insight: Apply implicit function theorem to equilibrium or optimality conditions (and use computational tricks to efficiently compute $\mathrm{d}\ell/\mathrm{d}\theta$ directly)

OptNet: Differentiable Optimization as a Layer in Neural Networks

Task-based End-to-end Model Learning in Stochastic Optimization

SATNet: Bridging deep learning and logical reasoning using a differentiable

...laximization

Differentiable Convex Optimization Layers

End-to-End Differentiable Physics for Learning and Control

Filipe de A. Belbute-Peres
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
filiped@cs.cmu.edu

Kevin A. Smith
Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139
k2smith@mit.edu

Kelsey R. Allen

Joshua B. Tenenbaum

Neural Ordinary Differential Equations

Ricky T. Q. Chen*, Yulia Rubanova*, Jesse Bettencourt*, David Duvenaud
University of Toronto, Vector Institute
{rtqichen, rubanova, jessebett, duvenaud}@cs.toronto.edu

Powerful toolkit for incorporating important structure into deep learning methods

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Overview: Provably robust control via RL

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

Approach: Use implicit layers in deep reinforcement learning (RL) to guarantee enforcement of hard constraints

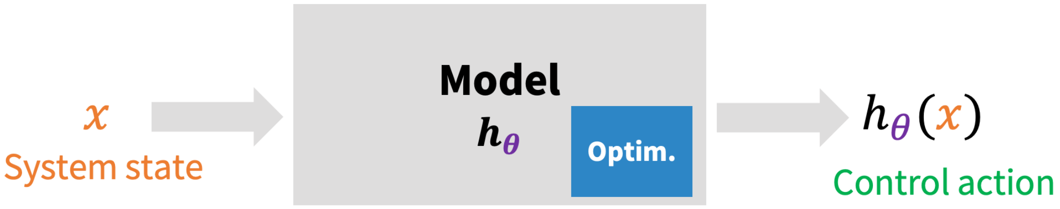

Settings:

- Asymptotic stability in power grids
- Realistic-scale building control

Deep reinforcement learning vs. robust control

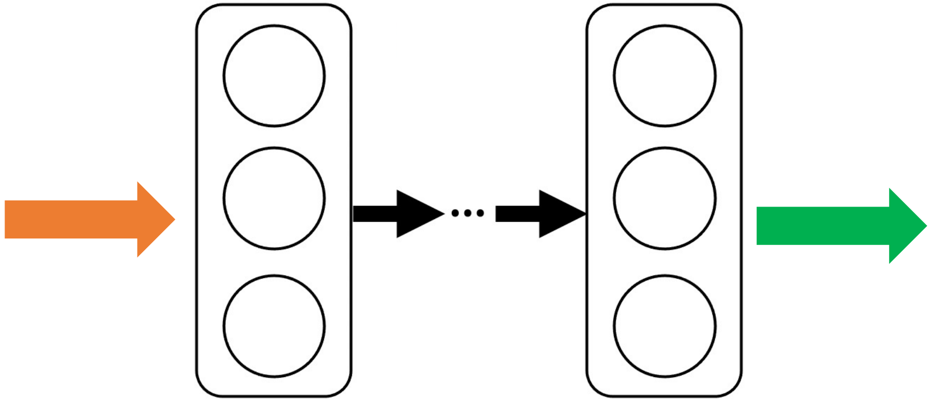

Deep RL

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

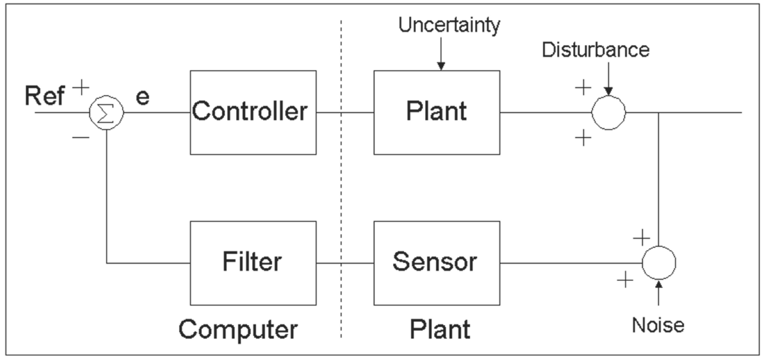

Robust control

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** (even for a randomly initialized neural network), trainable using reinforcement learning

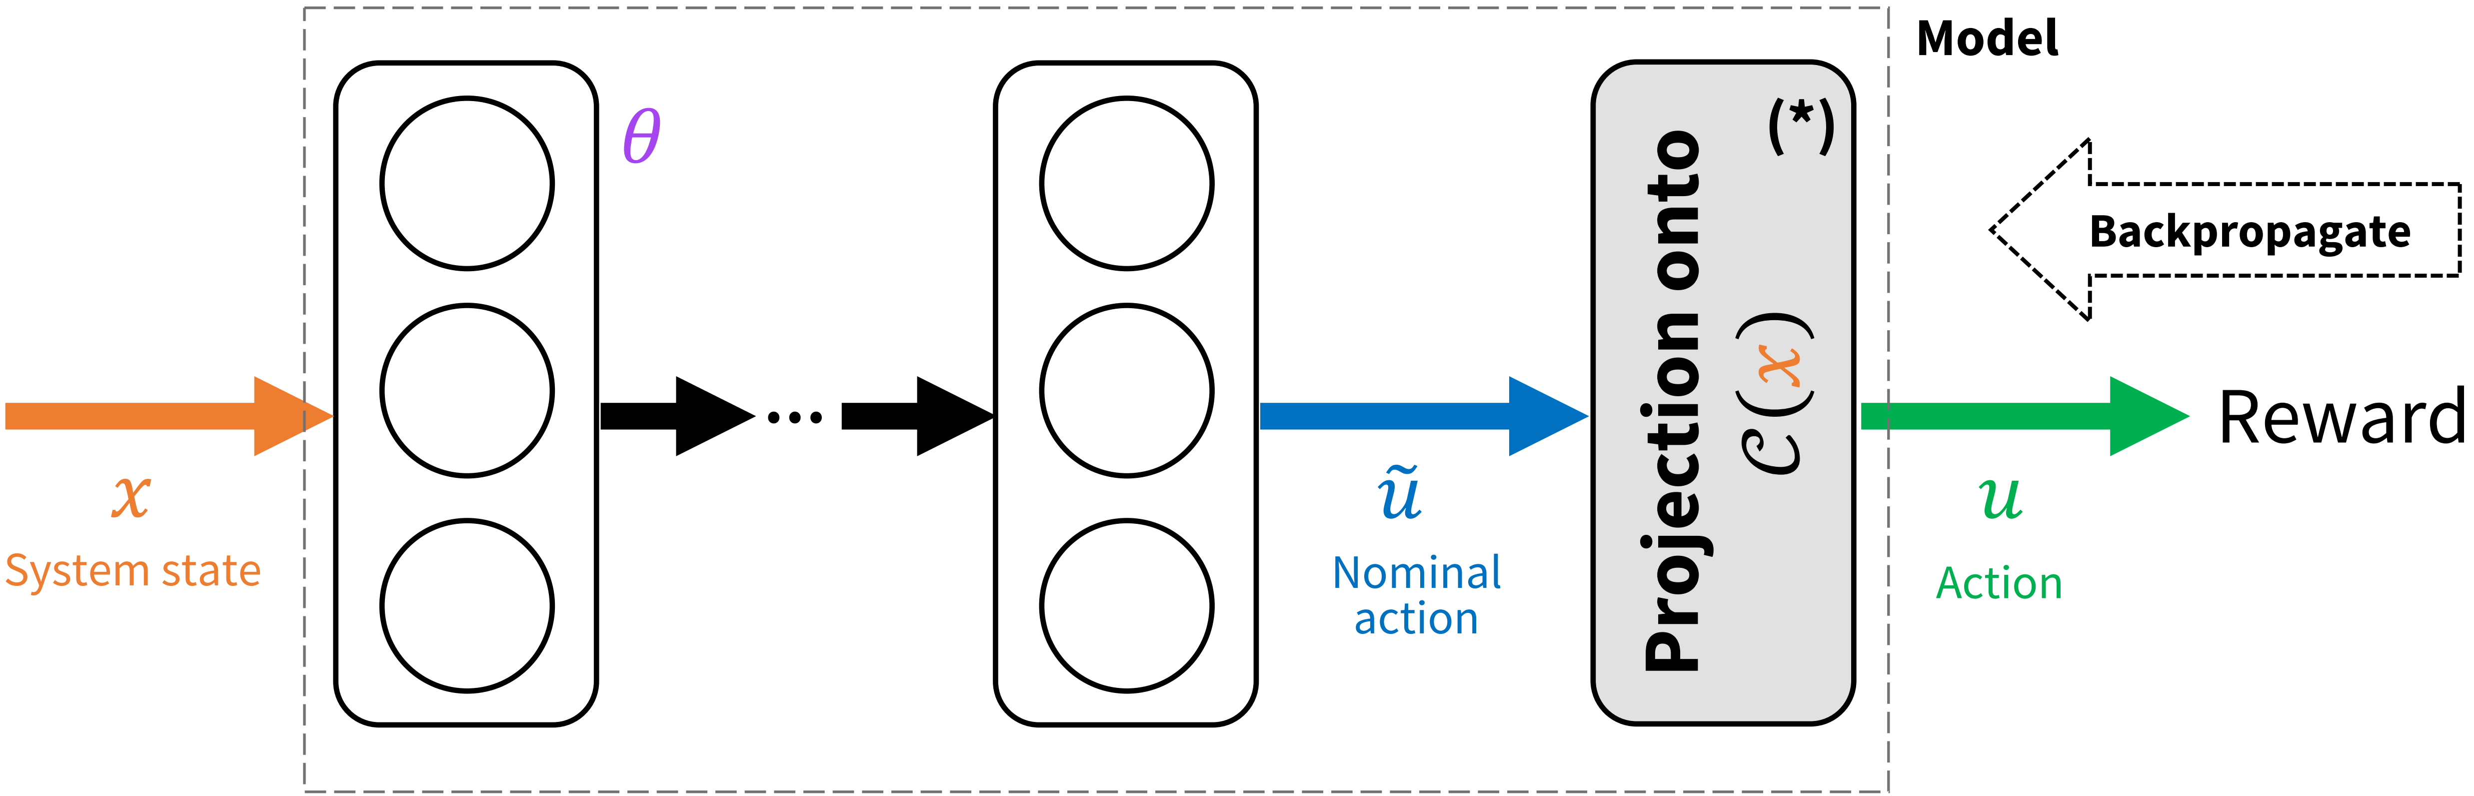

Finding a set of stabilizing actions (example)

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. $\|w(t)\|_2 \leq \|Cx(t) + Du(t)\|_2$
- Lyapunov function V obtained via robust control synthesis
- Exponential stability criterion: $\dot{V}(x(t)) \leq -\alpha V(x(t)), \forall x \neq 0$

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

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

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

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

Note: t-dependence has been dropped for brevity

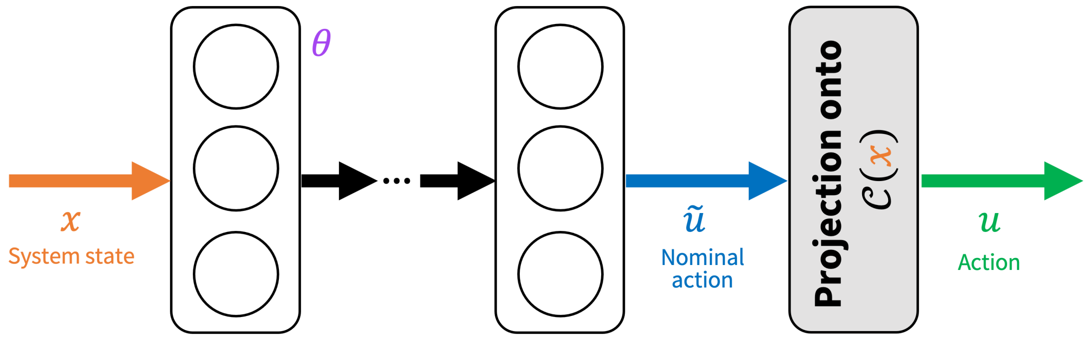

Illustrative results: Synthetic NLDI system

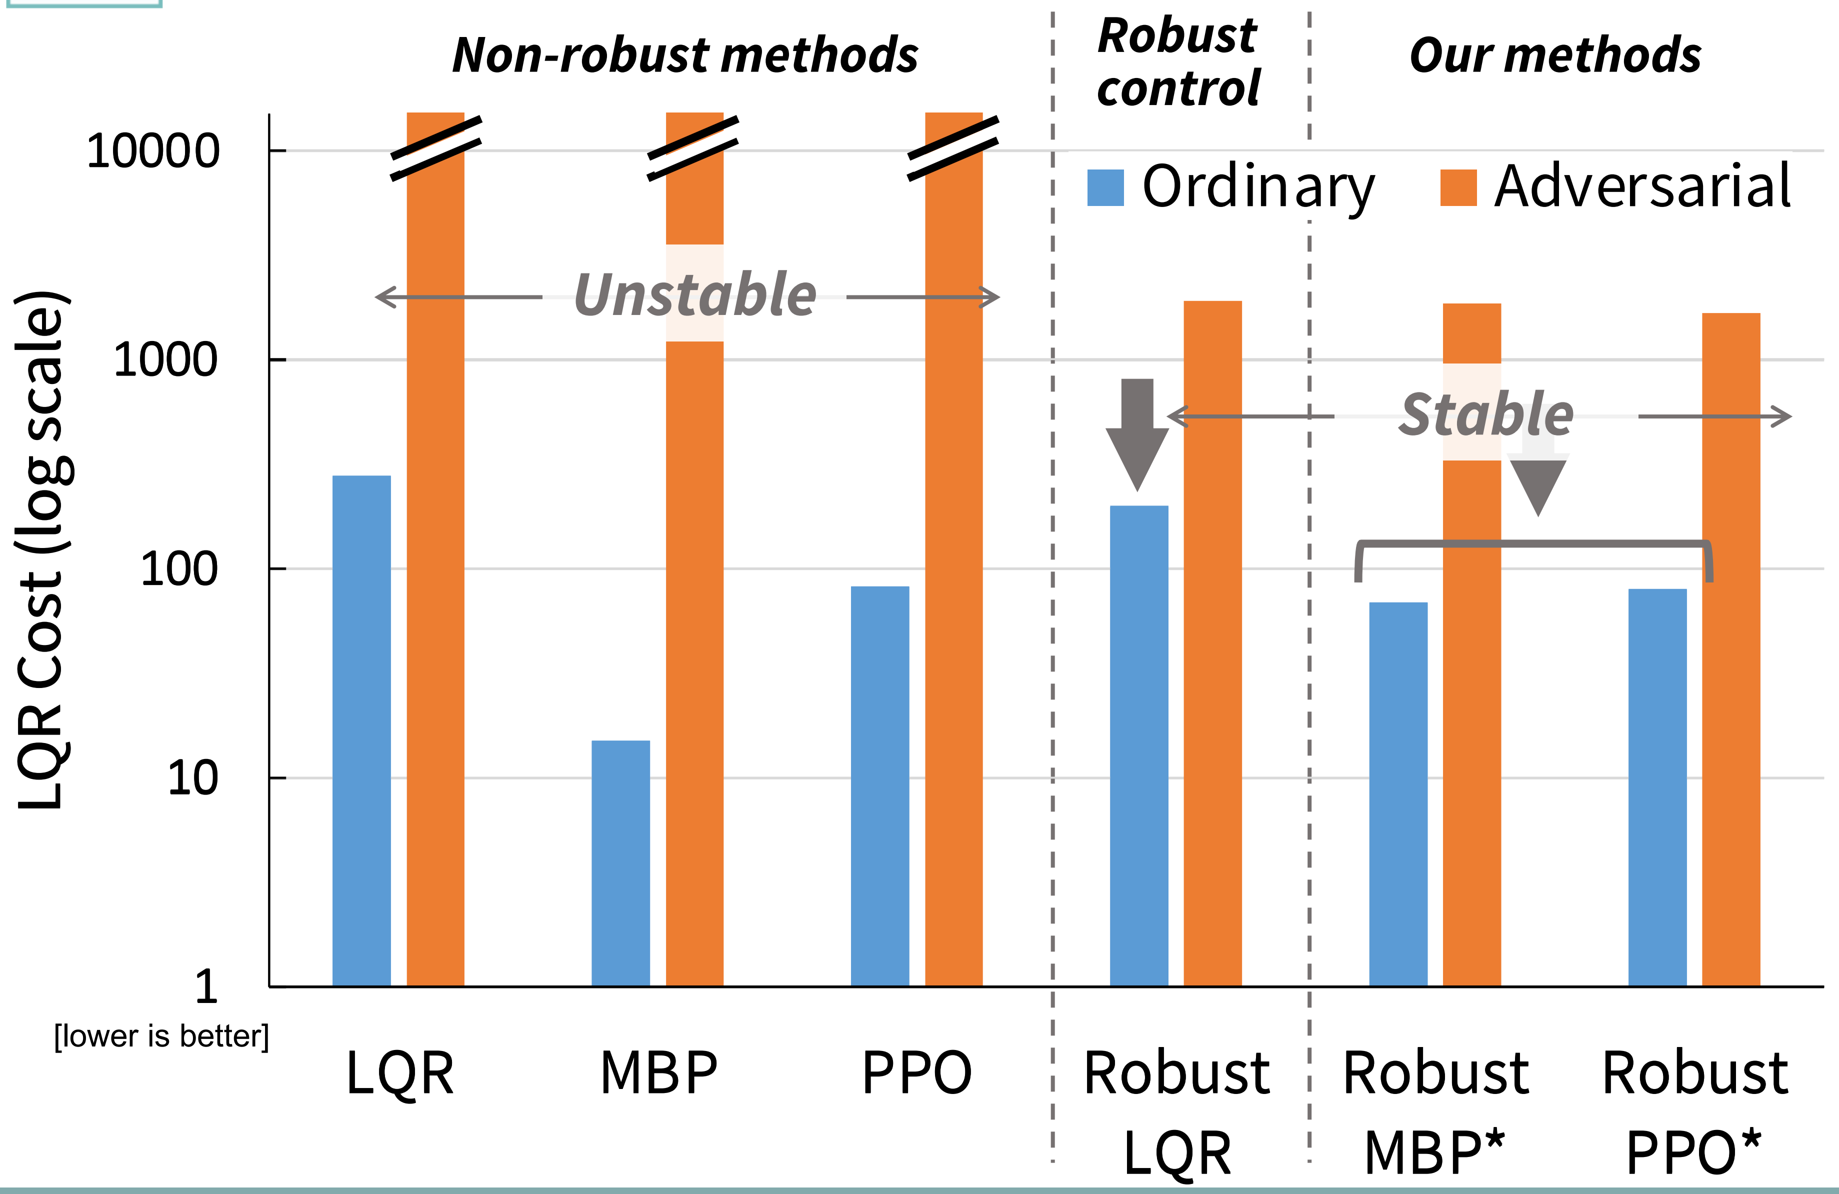

Improved "average-case" performance over robust baselines

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

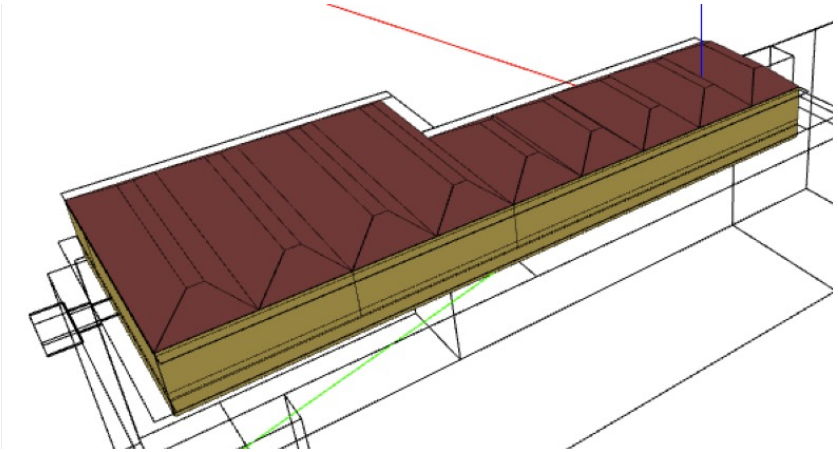

Intelligent Workplace
Margaret Morrison Hall, 4th Floor
(✤ Zhang & Lam, 2018)

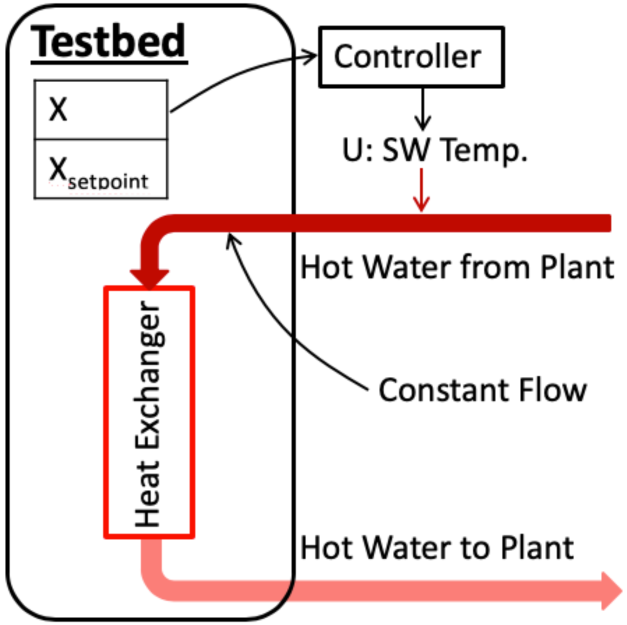

HVAC Schematic

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

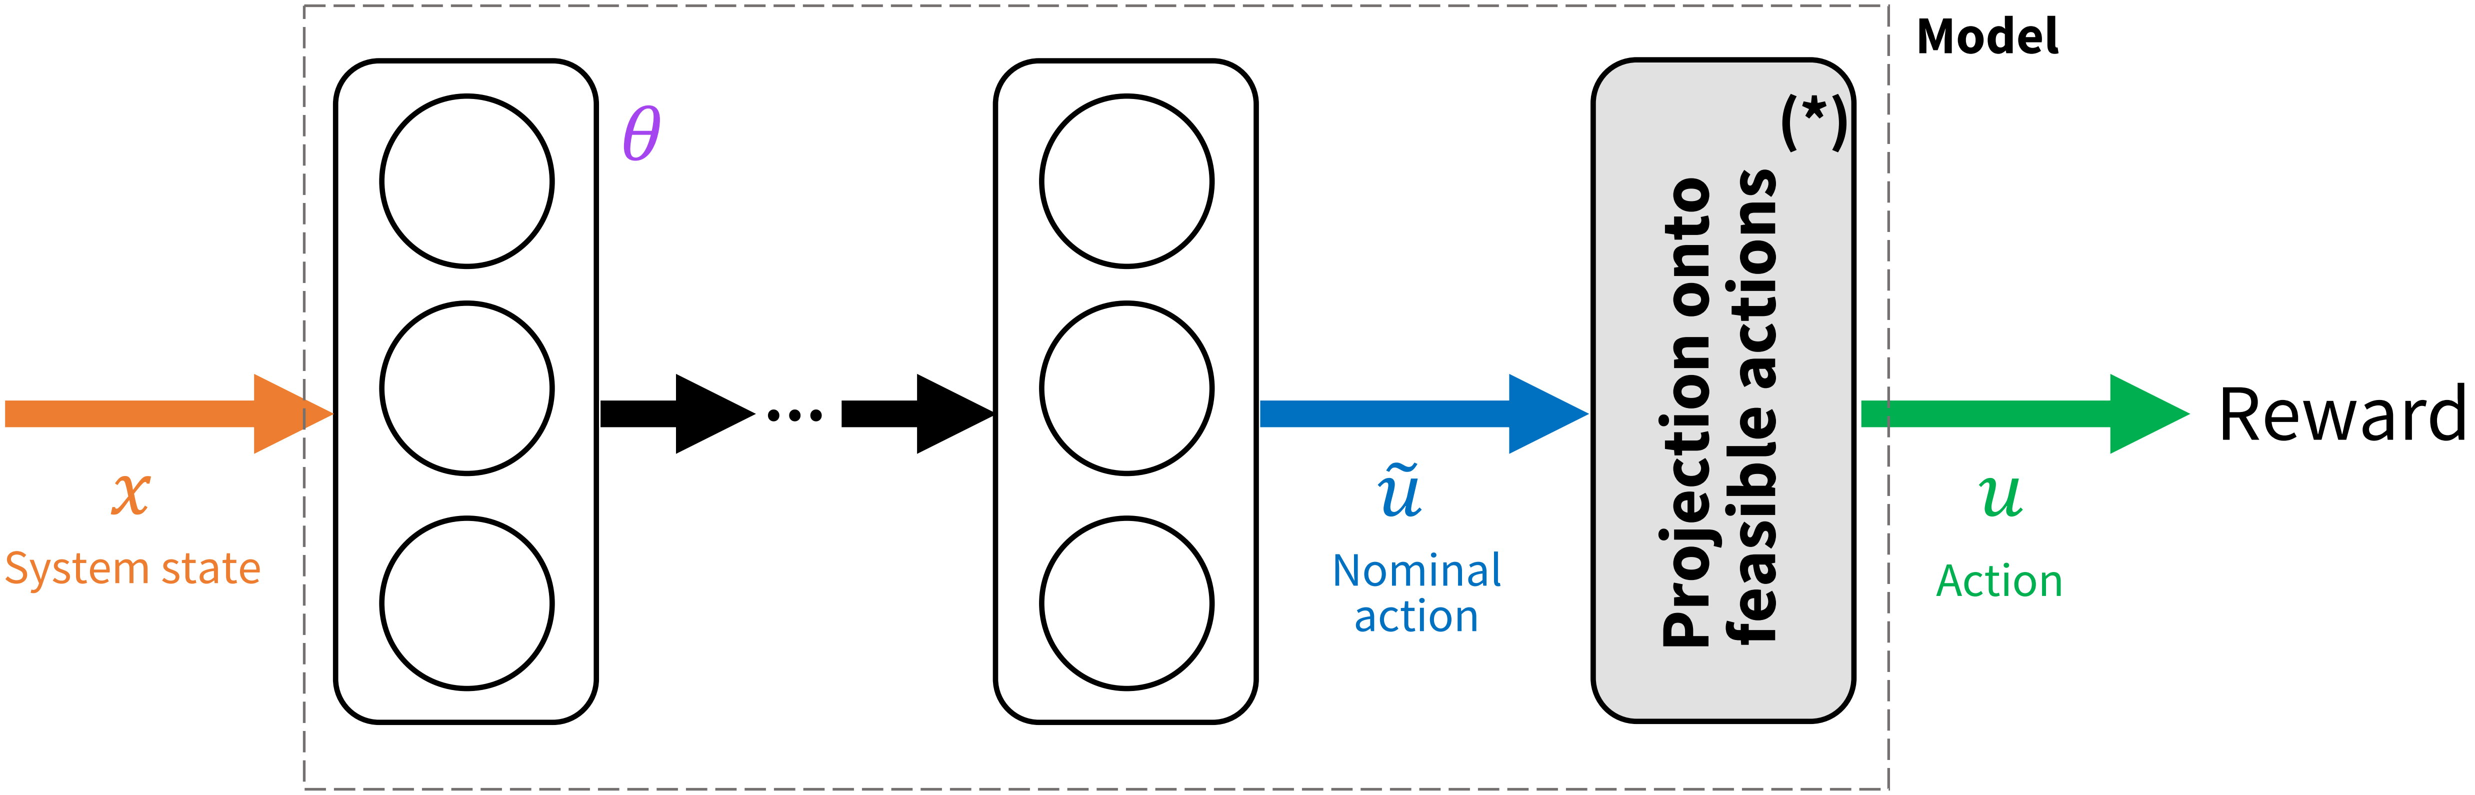

Summary: Provably robust control via RL

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

Settings:

- Asymptotic stability in power grids
- Realistic-scale building control

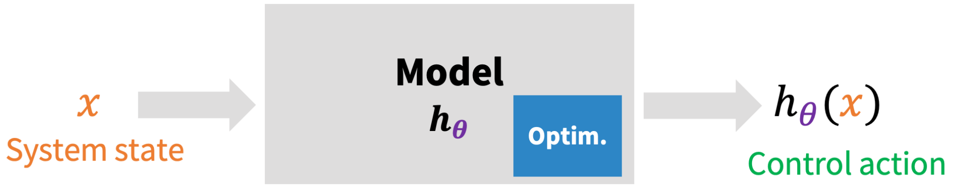

Insight: Project outputs of neural network onto a set of "safe" actions

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

Future directions:

- Leveraging more modern control theoretic formulations
- Improving computational costs / scaling to larger systems

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Overview: Decision-cognizant prediction

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

Approach: Construct decision-cognizant model using implicit function(s) in objective

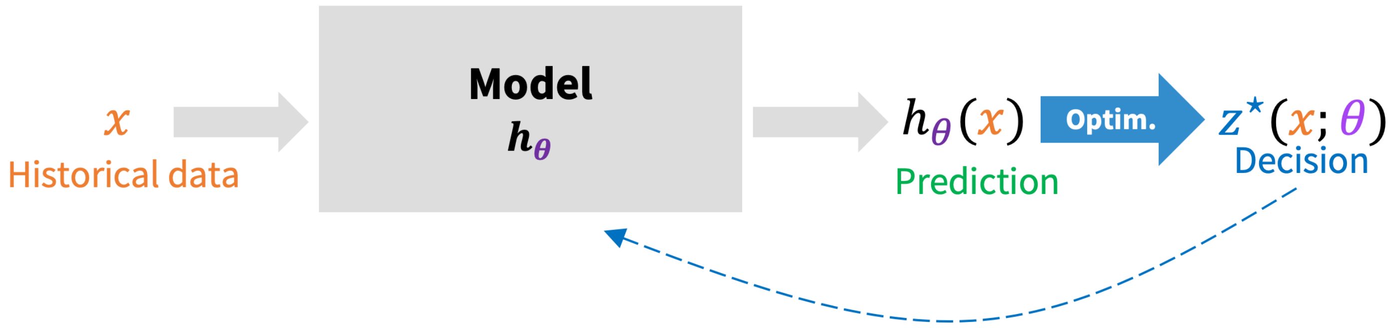

Setting: Decision-cognizant electricity demand forecasting

Decision-cognizant demand forecasting

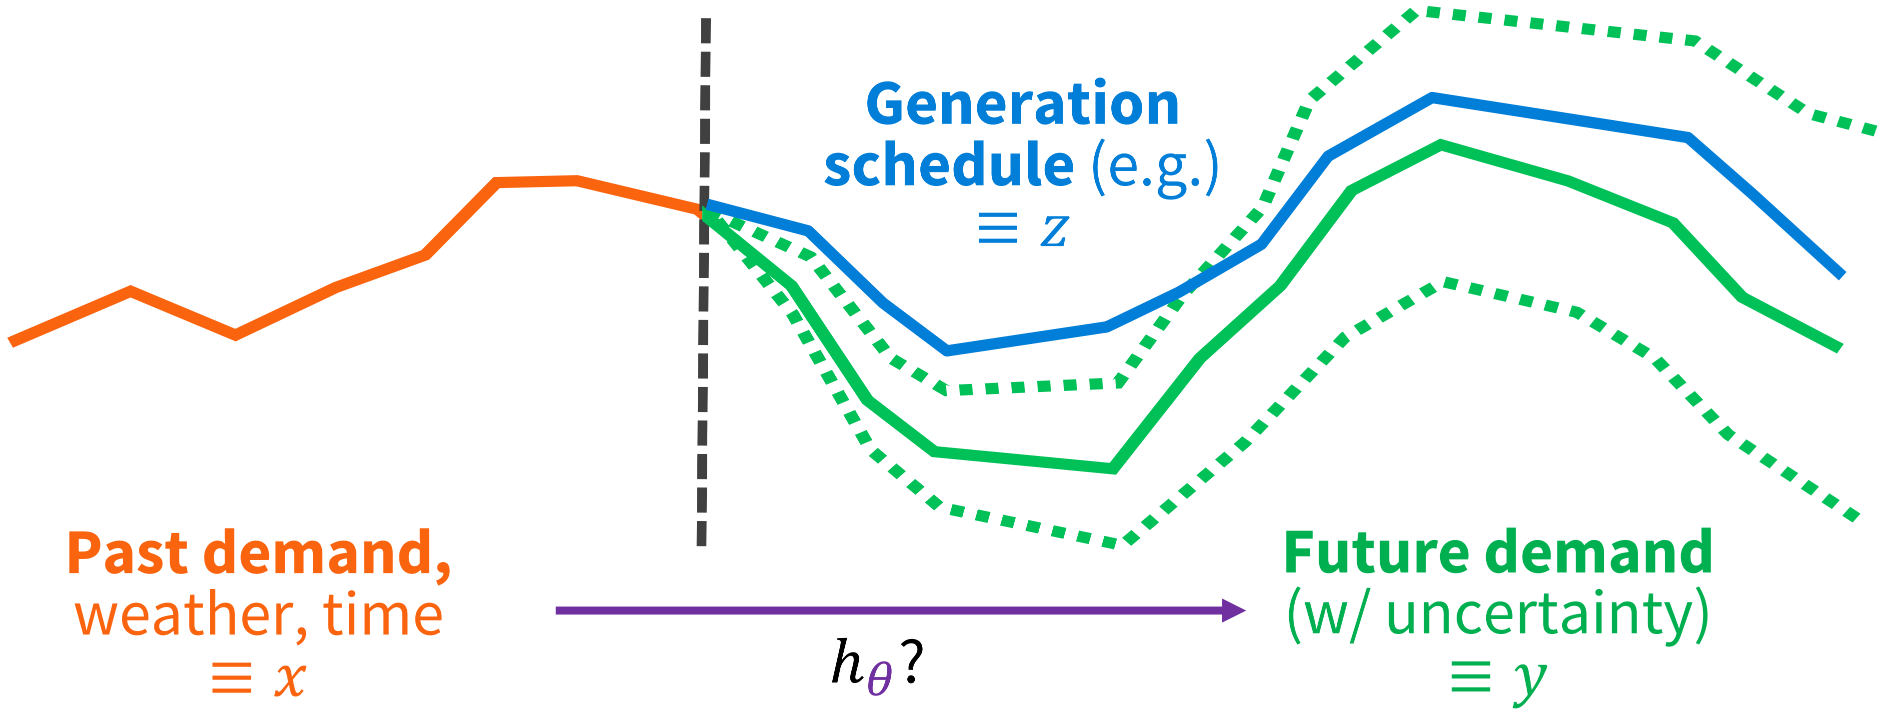

Goal: Optimize for quality of generation schedule when we observe actual demands

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

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

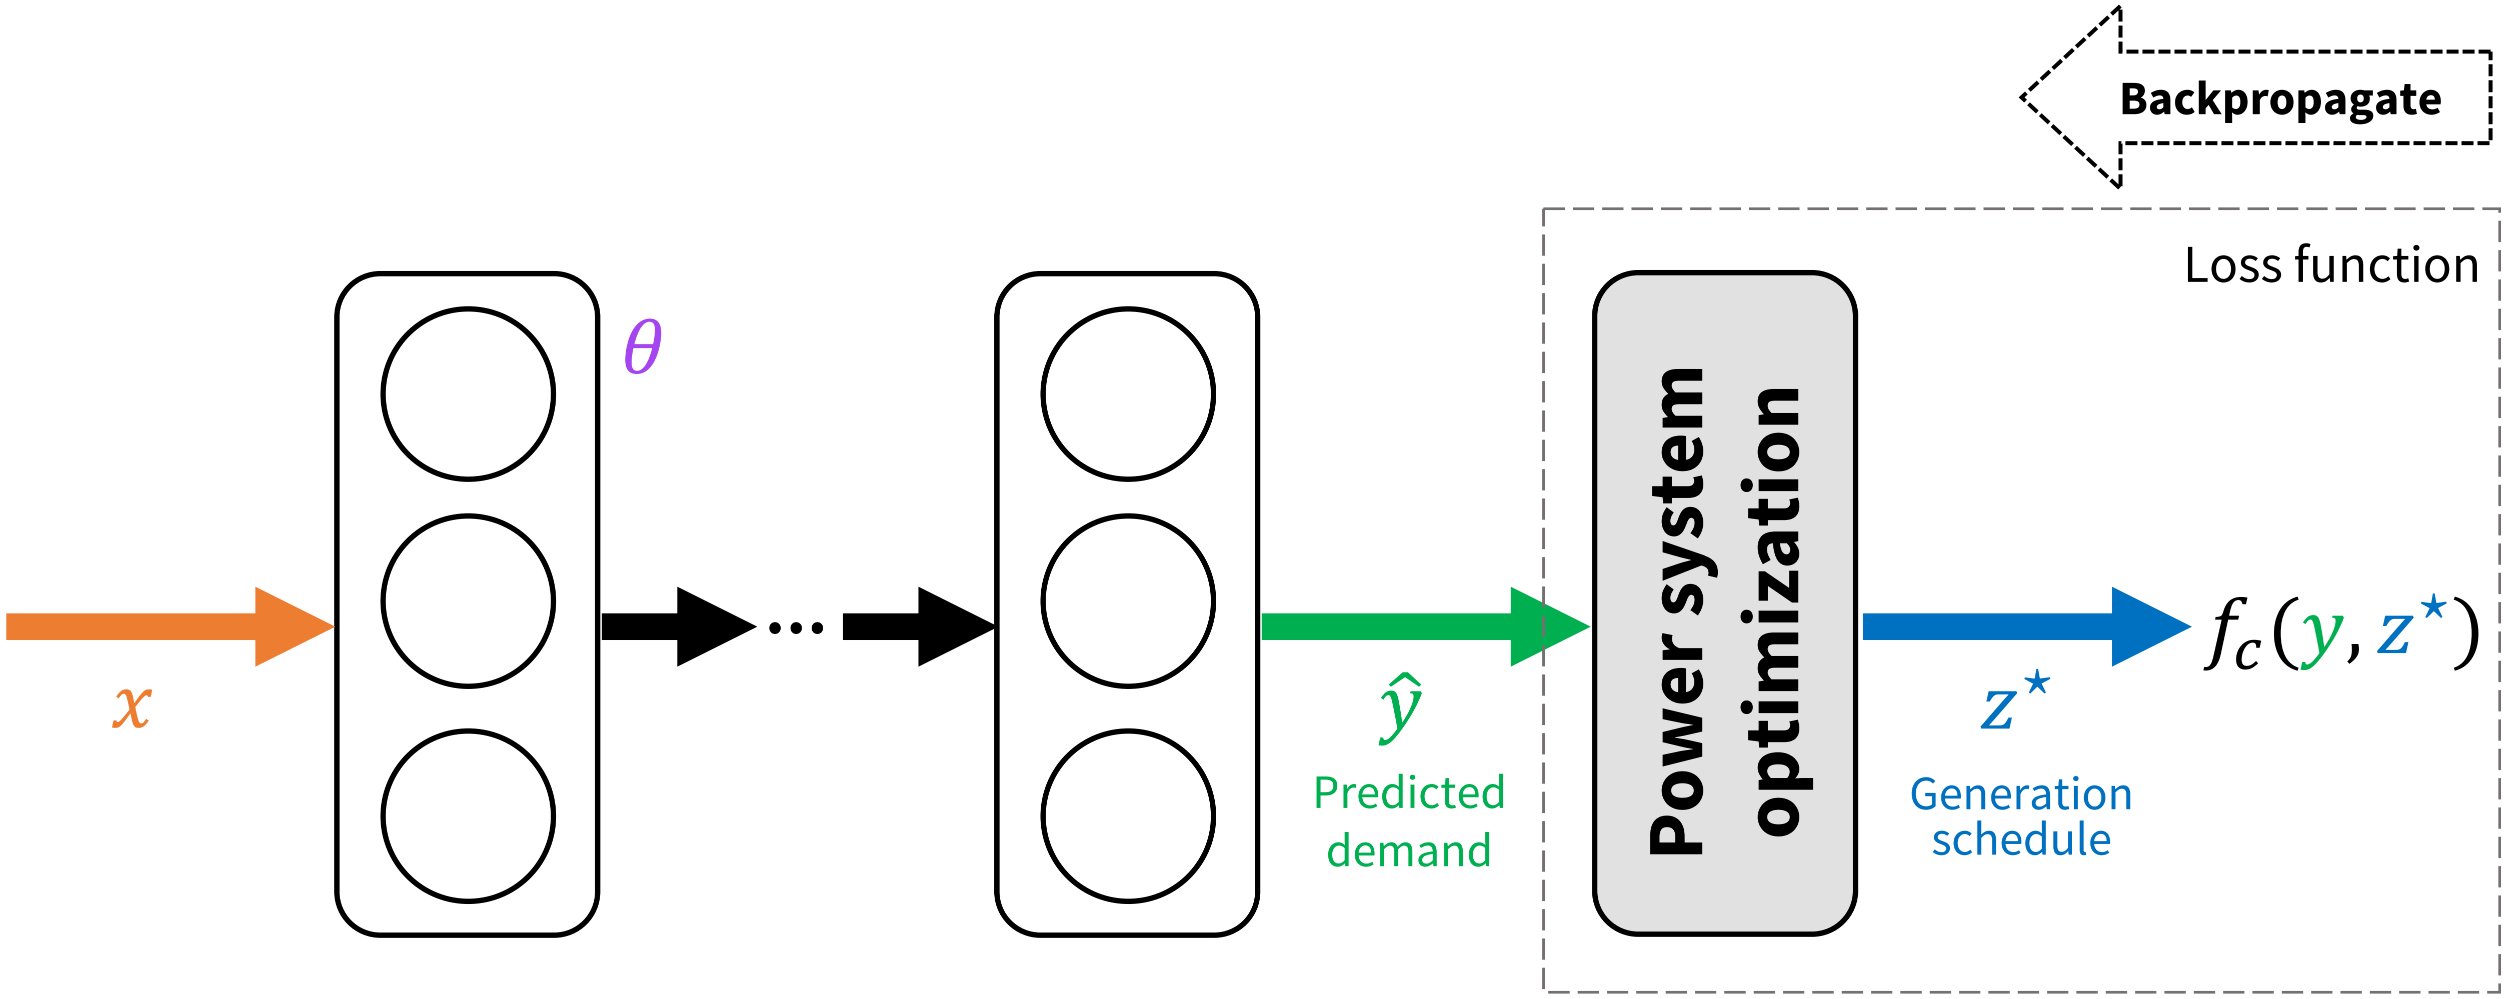

Decision-cognizant approach can dramatically improve generation scheduling outcomes

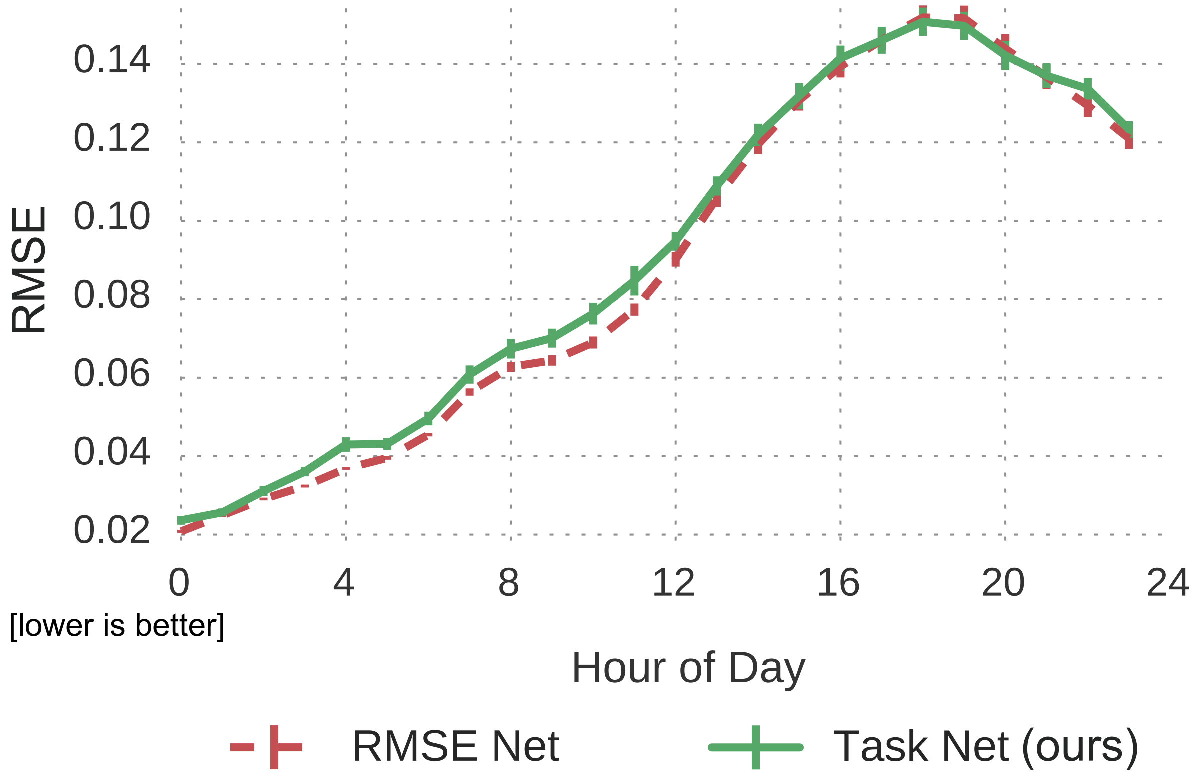

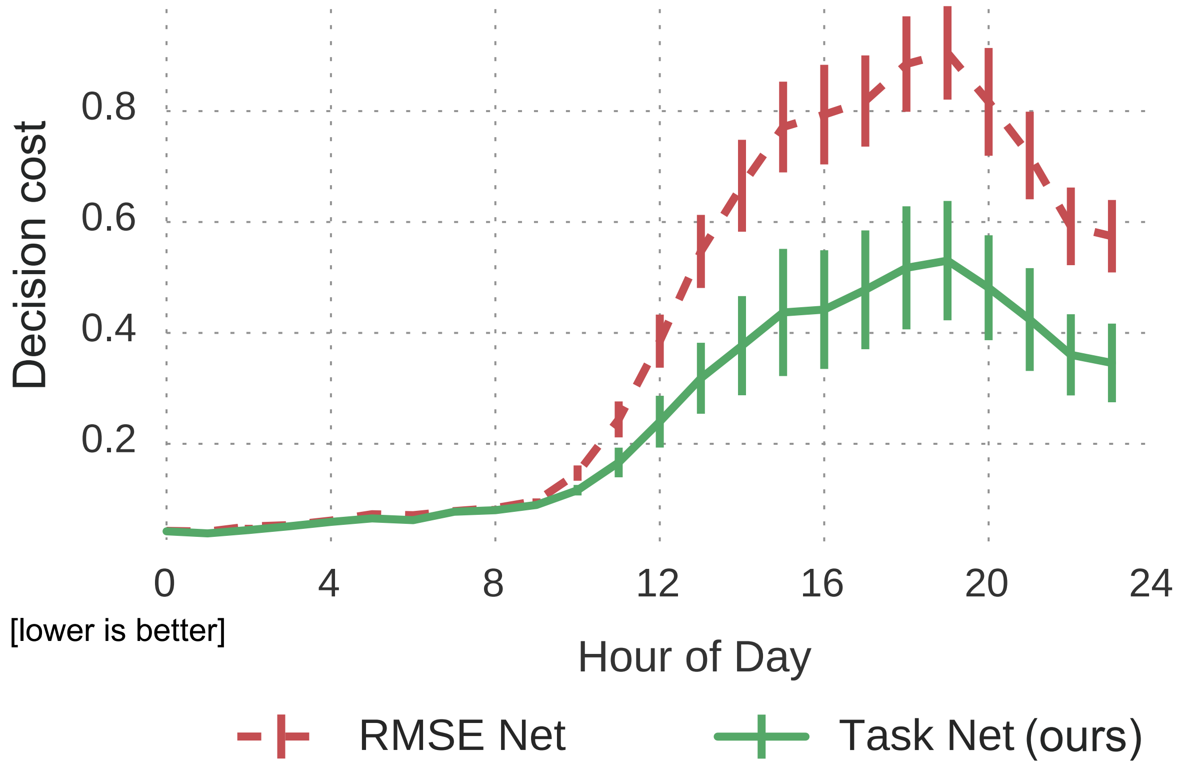

Decision-cognizant approach gives ~**39%** improvement in decision cost.

Summary: Decision-cognizant prediction

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

Setting: Electricity demand forecasting

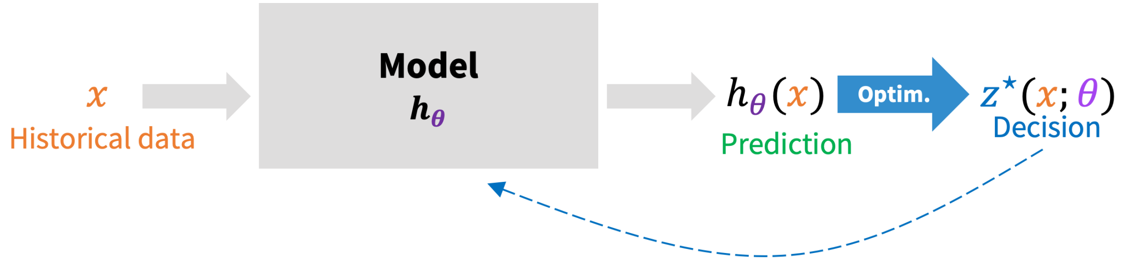

Insight: Incorporate knowledge of downstream decision-making process into the loss function, using implicit layers (differentiable optimization).

Future directions:

- Incorporating larger / more realistic decision-making procedures
- Extension to additional settings (e.g., end-to-end modeling + control)
- Understanding tradeoffs between decision-cognizant vs. decision-agnostic models

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Feasible optimization proxies

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

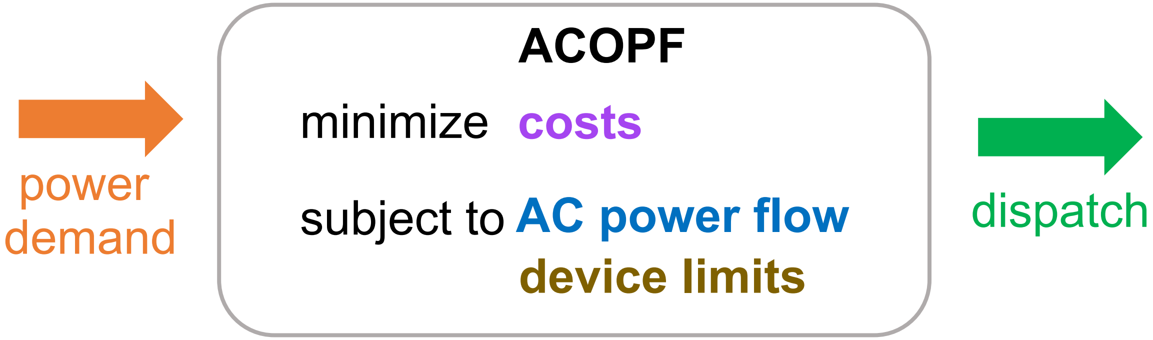

Approach:

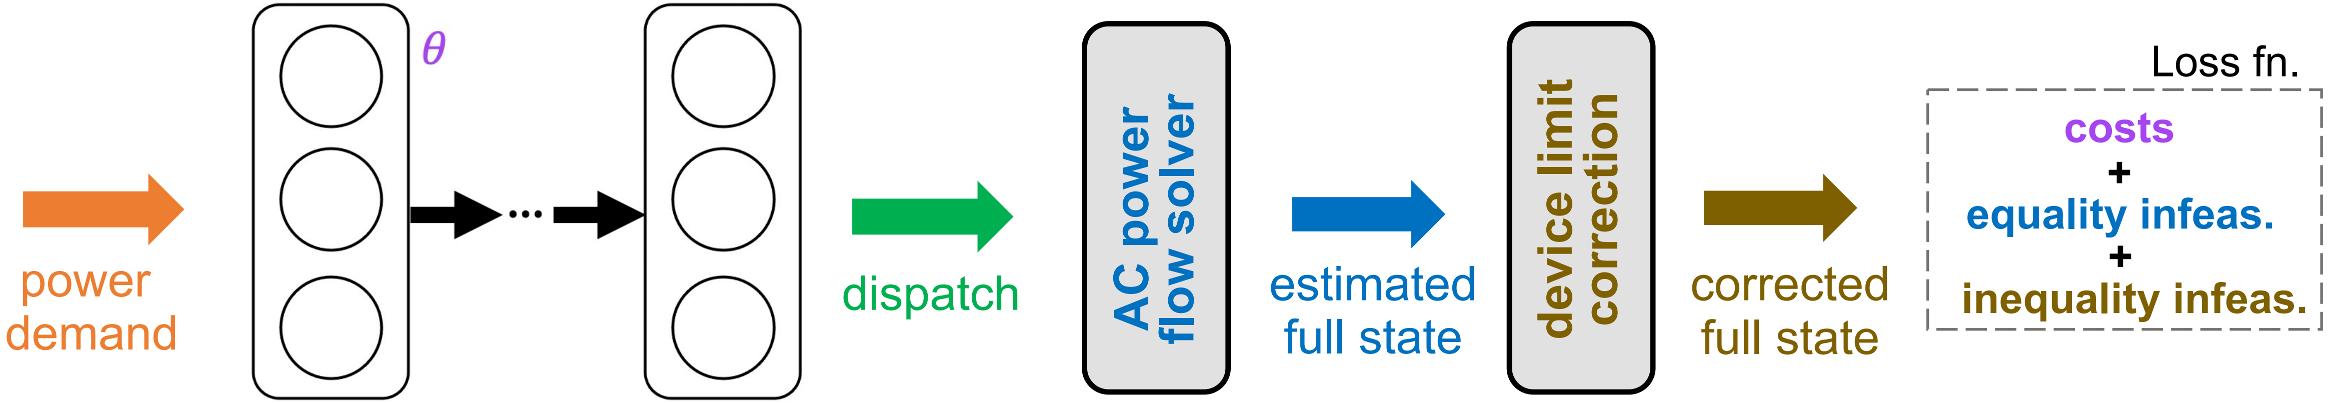

Note: Learns directly from problem specification (no training labels)

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.*

Approximating ACOPF: 57-bus test case

	Comparable objective value	Satisfies all constraints (unlike baselines)		10x faster than IPOPT
	Objective value	**Max equality violation**	**Mean equality violation**	**Time (s)**
IPOPT	3.81 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.949 ± 0.002
Baseline NN	—	0.19 + 0.01	0.03 ± 0.00	—
Our approach	3.82 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.089 ± 0.000

Future directions:

- Larger scale trials (fitting on a GPU)
- Mixed-integer problems (e.g., unit commitment)
- Generalization over topologies (e.g., via GNNs)
- Combinations with frameworks like PDL [PH2023]

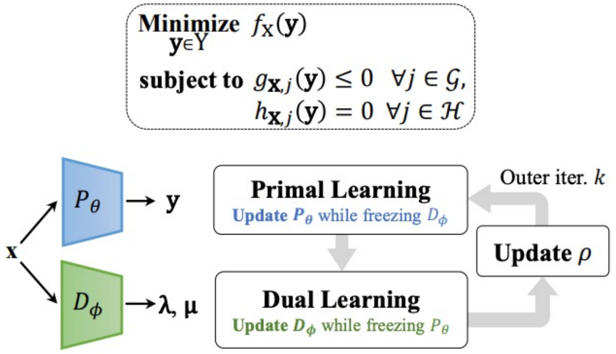

Image source: [PH2023]

[PH2023] Seonho Park, Pascal Van Hentenryck. "Self-Supervised Primal-Dual Learning for Constrained Optimization." *AAAI* (2023).

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Adversarially robust deep learning

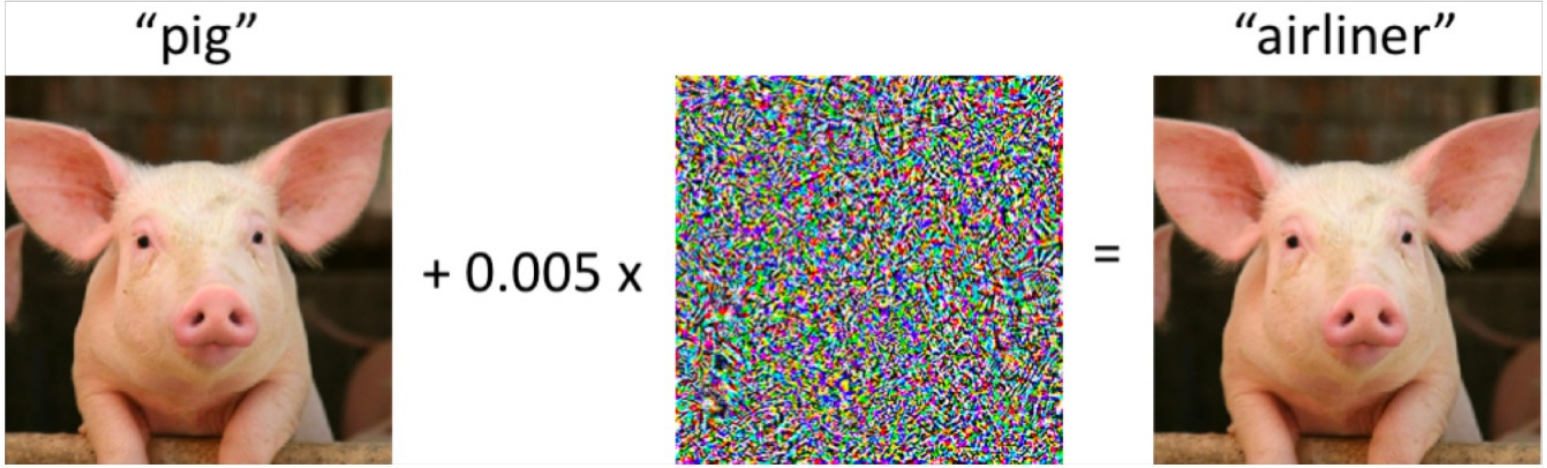

Part II: Training a robust classifier

$$\min_{\theta} \sum_{x,y \in S} \max_{\delta \in \Delta} \text{Loss}\,(x + \delta, y; \theta)$$

Part I: Creating an adversarial example (or ensuring one does not exist)

Slide adapted from: Zico Kolter and Aleksander Madry. "Adversarial Robustness - Theory and Practice." Tutorial at *NeurIPS 2018*. adversarial-ml-tutorial.org.

Drawing inspiration from adversarially robust DL

Adversarially robust deep learning: Pick neural network parameters to bound the cost of any worst-case perturbation

- Required scalable gradient-based optimization methods

Security-constrained OPF: Pick dispatch to bound the cost of worst-case contingencies

- Leverage similar scalable gradient-based optimization methods?

Adversarial robustness for N-k SCOPF

N-k security-constrained optimal power flow (SCOPF): Schedule power to be robust to potentially *k* simultaneous generator or line failures (contingencies)

DEFENSE (dispatch) **ATTACK** (contingencies)

$\text{minimize}_{x \in \mathcal{X}} \; \max_{y \in \mathcal{Y}}$ **Base case cost** + **Contingency case cost**

subject to $x \in$ **Base case grid and device constraints** (incl. nonlinear equality constraint)

$y \in$ **Second-stage adjustments to power and voltage** (optimization problem)

Step 0: Formulate as bilevel (attacker-defender) opt. over dispatch and continuous outer relaxation of contingencies

Step 1 ("attack stage"): Find worst-case contingency via a few steps of projected gradient ascent (with implicit diff.)

Step 2 ("defense stage"): Update dispatch to improve robustness against worst case contingency (e.g., via projected gradient descent or efficient Gauss-Seidel approach)

Priya L. Donti*, Aayushya Agarwal*, Neeraj Vijay Bedmutha, Larry Pileggi, and J. Zico Kolter. "Adversarially Robust Learning for Security-Constrained Optimal Power Flow." *Conference on Neural Information Processing Systems (NeurIPS)* 2021.

Illustrative results (4622-bus system)

3-4x improvement over OPF for N-2/N-3 SCOPF, **in only 21 minutes on a laptop**

Contingency type	N-1	N-2	N-3
Scenarios tested	6,133	359,712	428,730
OPF violations	59	10,572	4,086
CAN∂Y violations*	36	3,580	1,122

* our approach

Note: Comparable N-1 SCOPF performance, and superior N-2 and N-3 performance, to ARPA-E GO Competition baselines

See also: Results on stochastic OPF for 11,615-bus system (PSCC 2022)

Talk outline

Implicit layers in deep learning

Provably robust control via deep RL

Decision-cognizant prediction

Feasible optimization proxies

Bonus: Adversarial robustness for N-k SCOPF

Enablers for next-gen. optimization & control

More openness in data, beyond only bilateral agreements and limited access

- Can include sharing of synthetic data

Simulators and test beds, with realistic/diverse scenarios and easy-to-use interfaces

- Includes digital twins, but also simpler frameworks (e.g., Grid2Op)
- Need for *progression pathways* from basic to advanced simulators/test beds

Evaluation metrics / benchmarks: What does it mean for a method to succeed (or fail)?

Open-source software, enabling integration and evaluation of new methods

Internal research capacity with external exchange: Enables translation of ideas without sharing difficult-to-share information across organizational boundaries

Note: None of these enablers are machine learning-specific!

Closing thoughts

Implicit layers in deep learning: powerful paradigm for bridging ML with power systems specifications

Going from theory to practice requires deep interdisciplinary **collaboration** and research-to-deployment **infrastructure**

Reach out if you'd like to chat, and check out the Climate Change AI network (www.climatechange.ai)

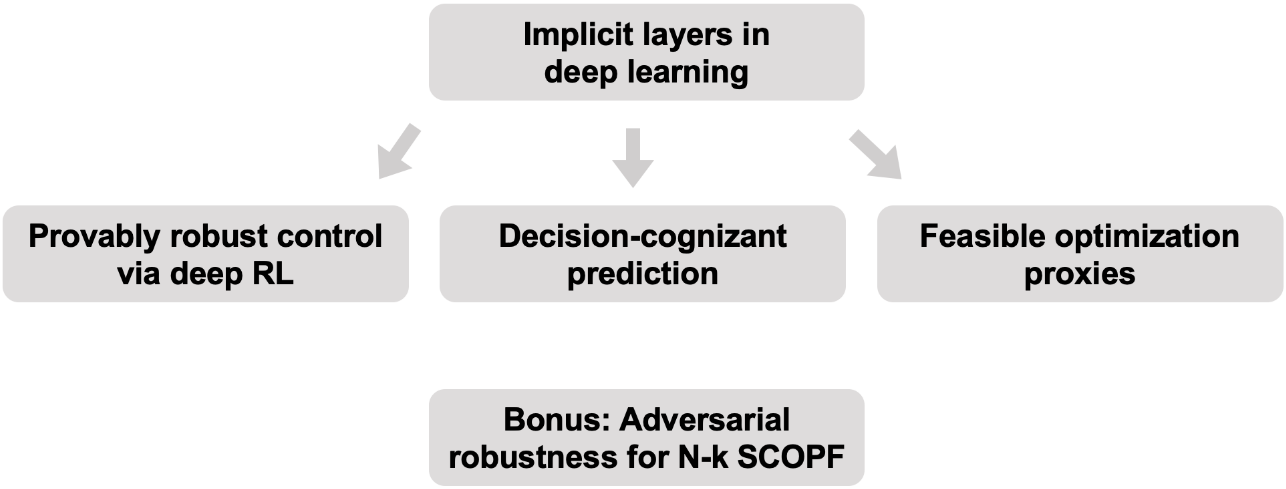

Priya L. Donti: donti@mit.edu

Backup slides

ML for power systems: Recurring themes

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

ML for power systems: Recurring themes

Distilling raw data into insights (GHG emissions, solar panels, vegetation)

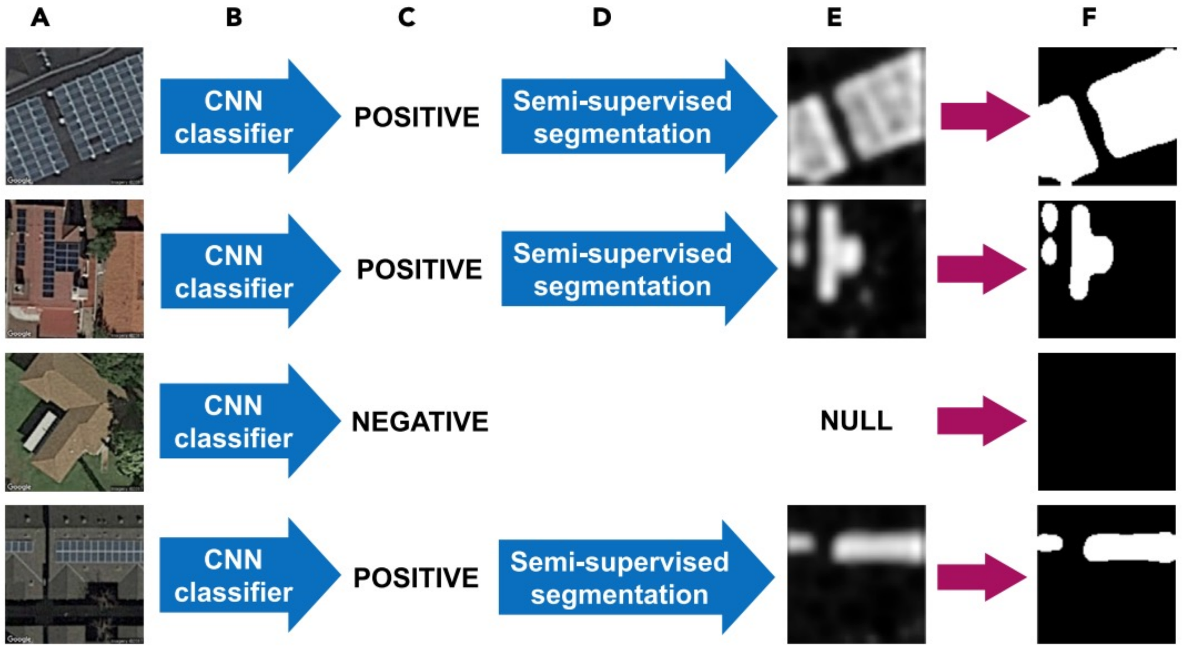

Image source: Yu, Wang, Majumdar, Rajagopal (2018)

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

ML for power systems: Recurring themes

Distilling raw data into insights (GHG emissions, solar panels, vegetation)

Forecasting (renewable energy, marginal/average emissions, prices)

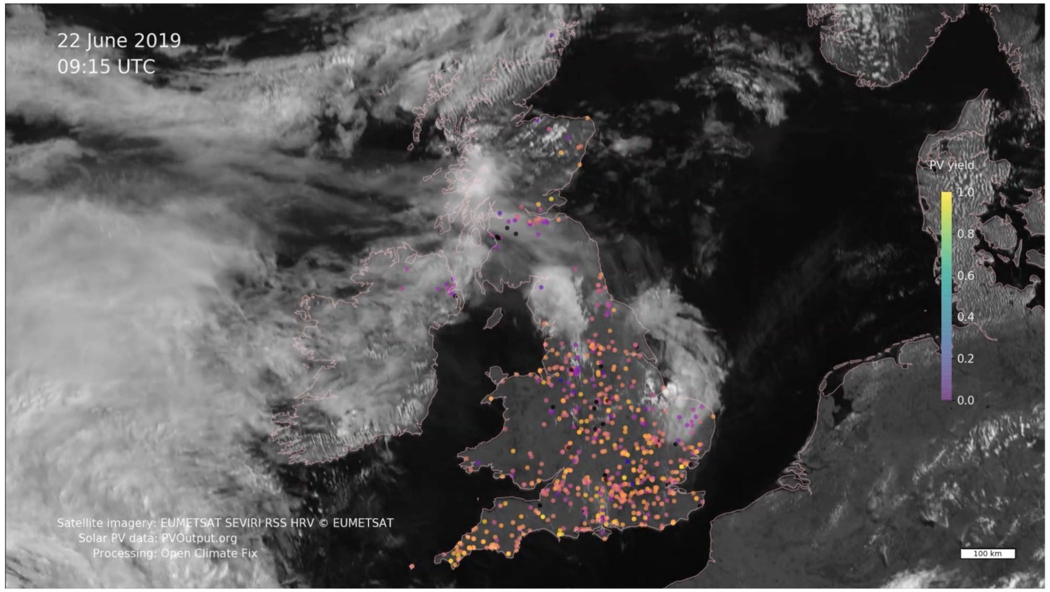

Image source: Open Climate Fix

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

ML for power systems: Recurring themes

Distilling raw data into insights (GHG emissions, solar panels, vegetation)

Forecasting (renewable energy, marginal/average emissions, prices)

Fast and dynamic optimization (power scheduling, MPPT)

Image source: L2RPN Challenge

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

ML for power systems: Recurring themes

Distilling raw data into insights (GHG emissions, solar panels, vegetation)

Forecasting (renewable energy, marginal/average emissions, prices)

Fast and dynamic optimization (power scheduling, MPPT)

Predictive maintenance
(resilient infrastructure, methane leaks)

Image source: [EPRI Journal (2019)]

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

ML for power systems: Recurring themes

Distilling raw data into insights (GHG emissions, solar panels, vegetation)

Forecasting (renewable energy, marginal/average emissions, prices)

Fast and dynamic optimization (power scheduling, MPPT)

Predictive maintenance
(resilient infrastructure, methane leaks)

Accelerated science
(batteries, solar, electrofuels, fusion)

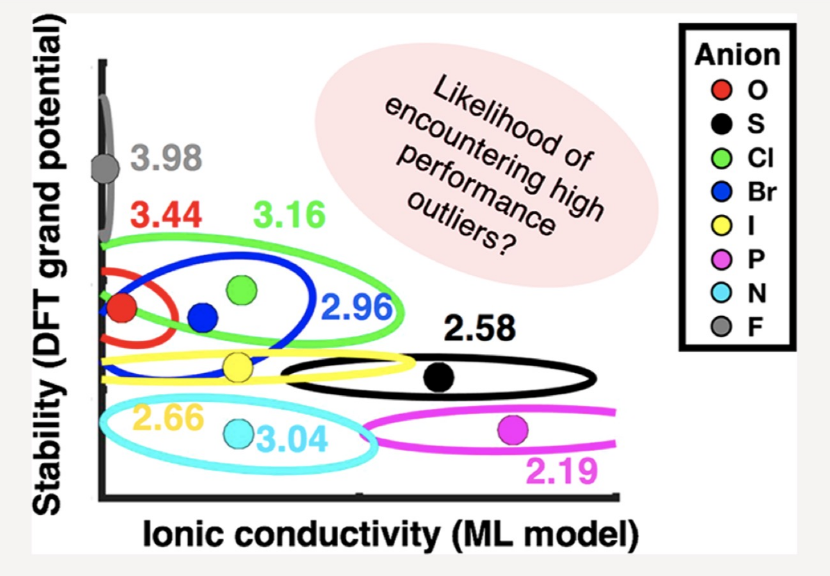

Image source: Sendek et al. (2020)

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

ML for power systems: Recurring themes

Distilling raw data into insights (GHG emissions, solar panels, vegetation)

Forecasting (renewable energy, marginal/average emissions, prices)

Fast and dynamic optimization (power scheduling, MPPT)

Predictive maintenance
(resilient infrastructure, methane leaks)

Accelerated science
(batteries, solar, electrofuels, fusion)

Data management
(data matching/fusion, data generation)

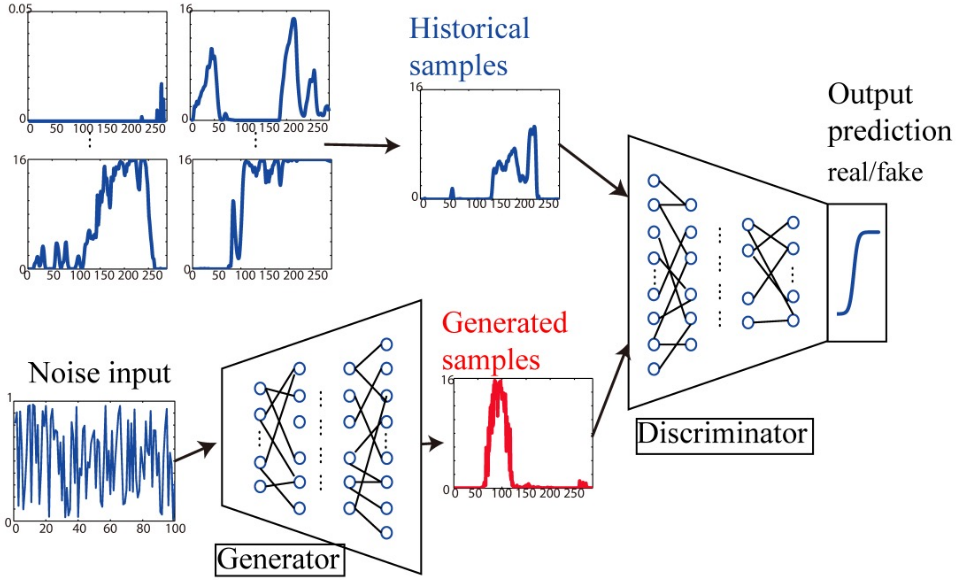

Image source: Chen, Wang, Kirschen, Zhang (2018)

See also: Donti, P. L. & Kolter, J. Z. (2021). Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, *46*, 719-747.

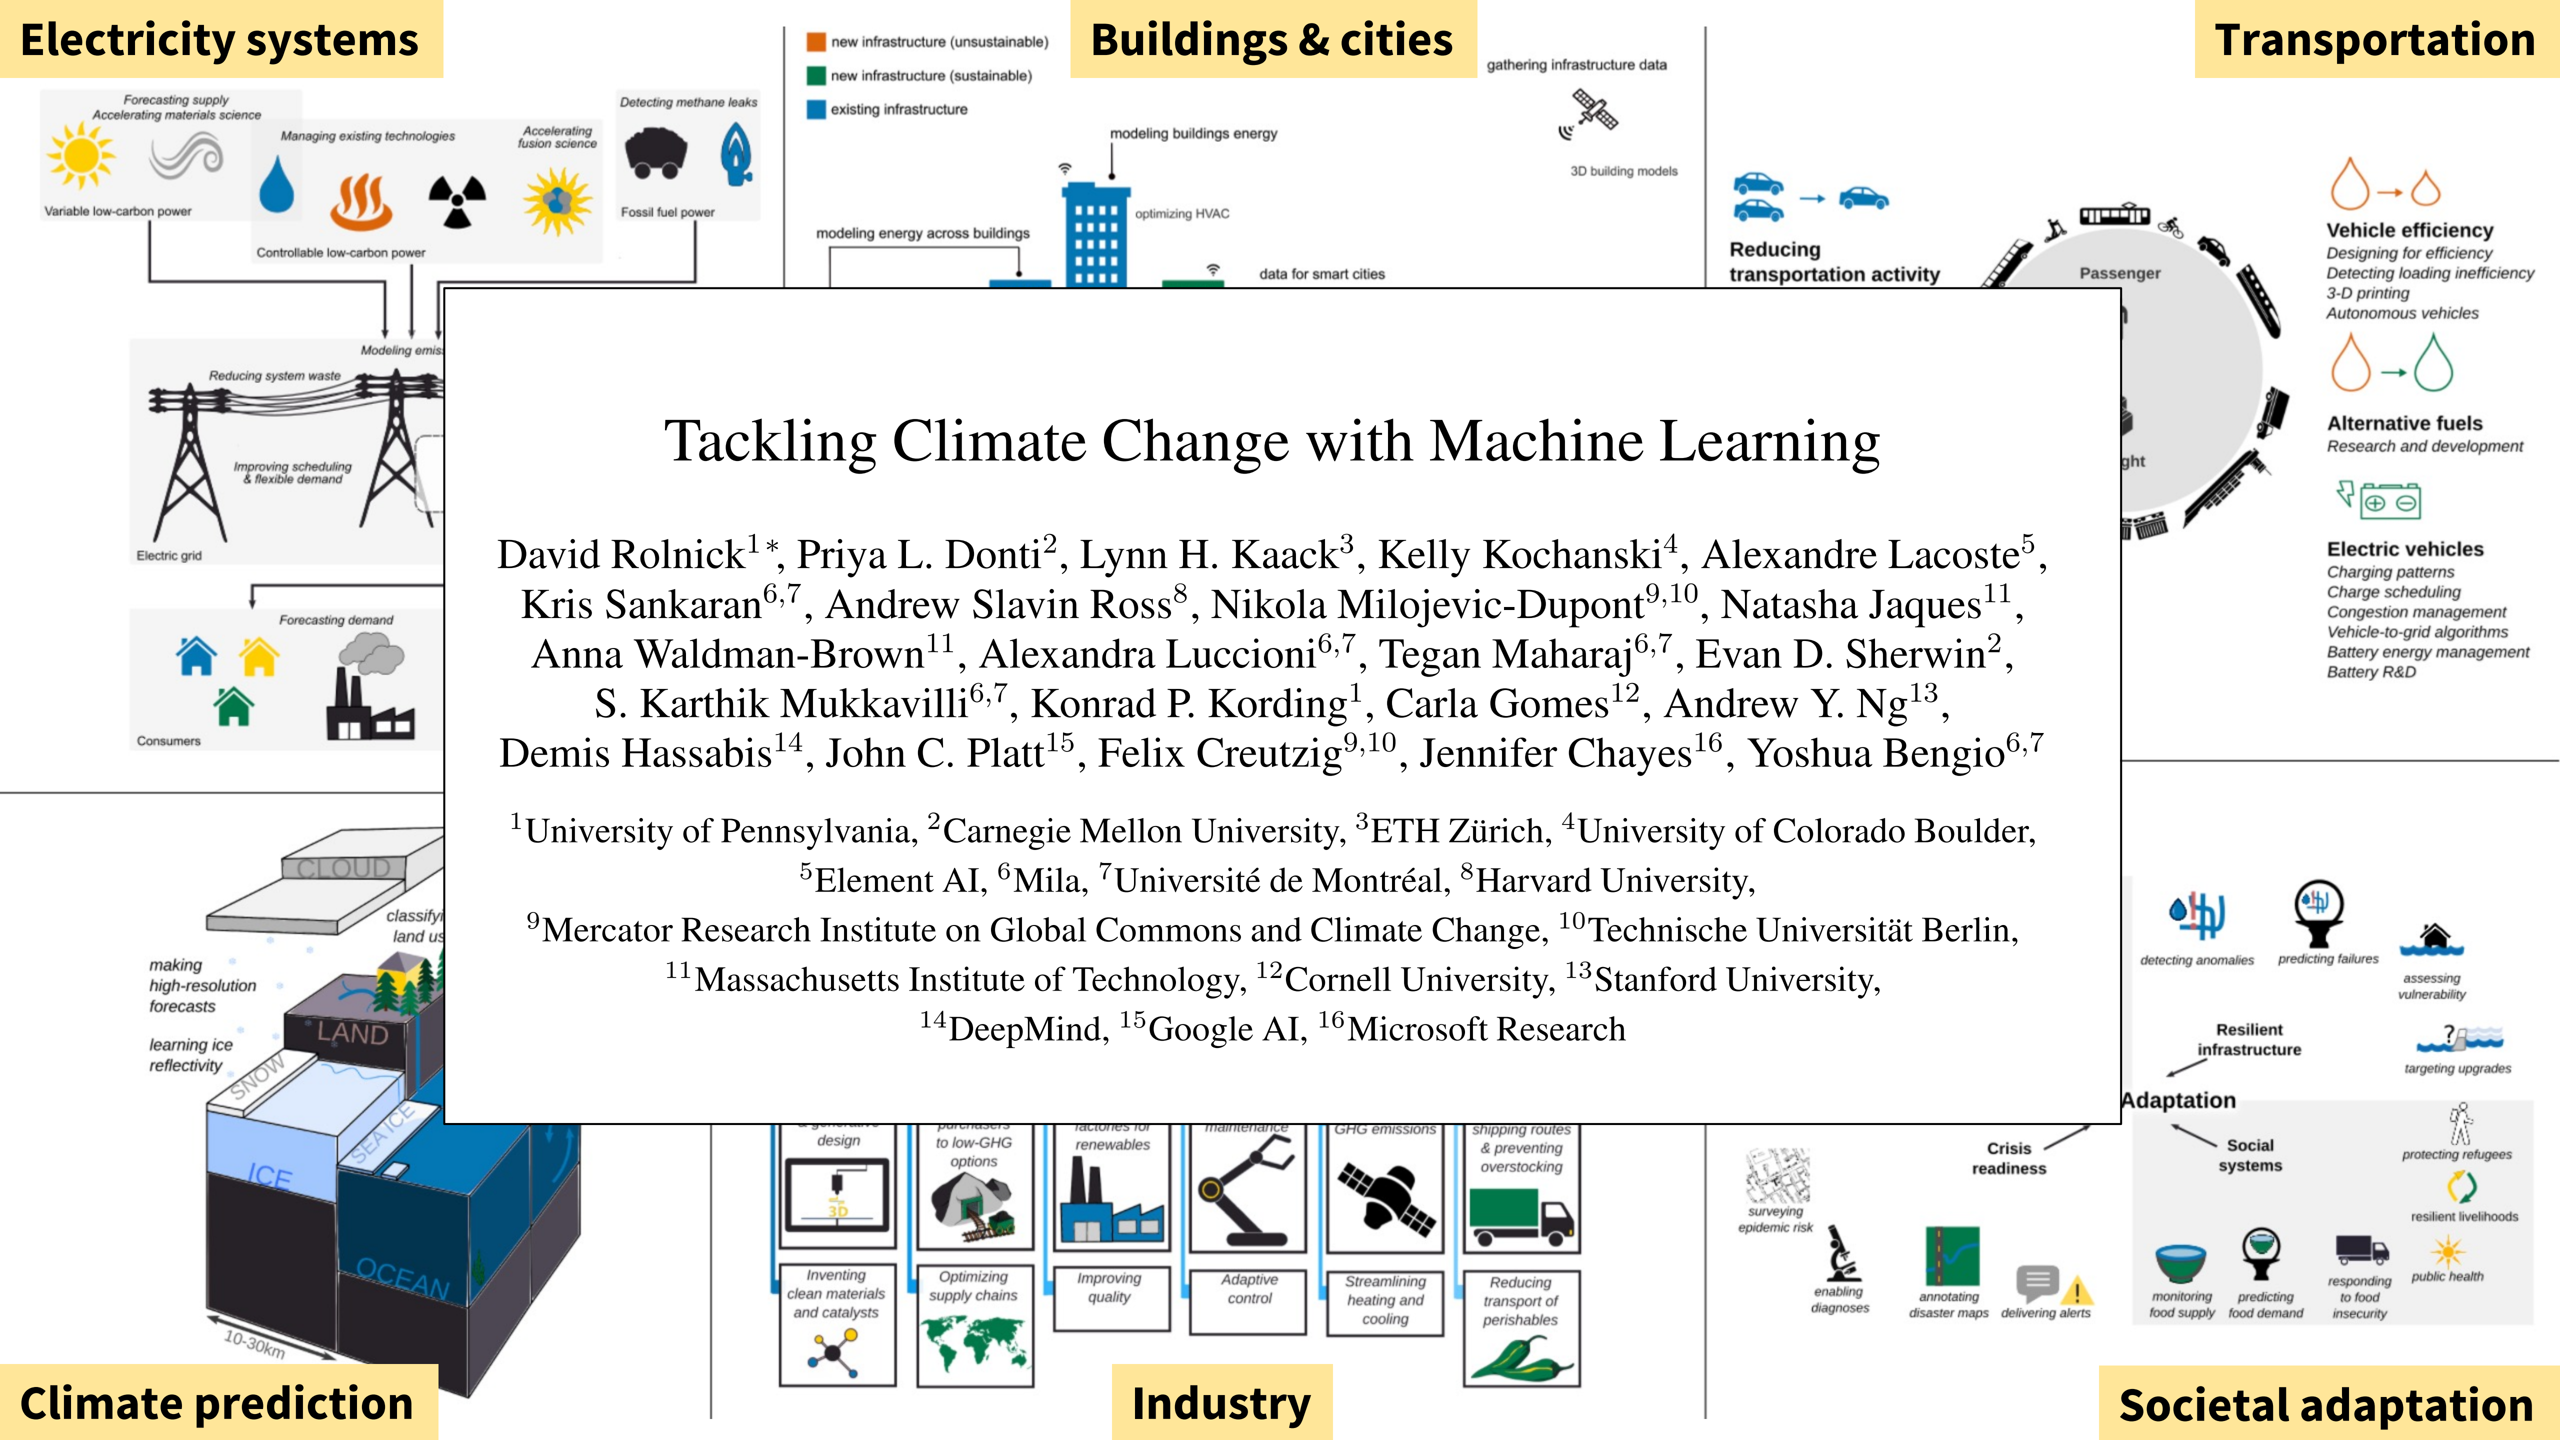

Tackling Climate Change with Machine Learning

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

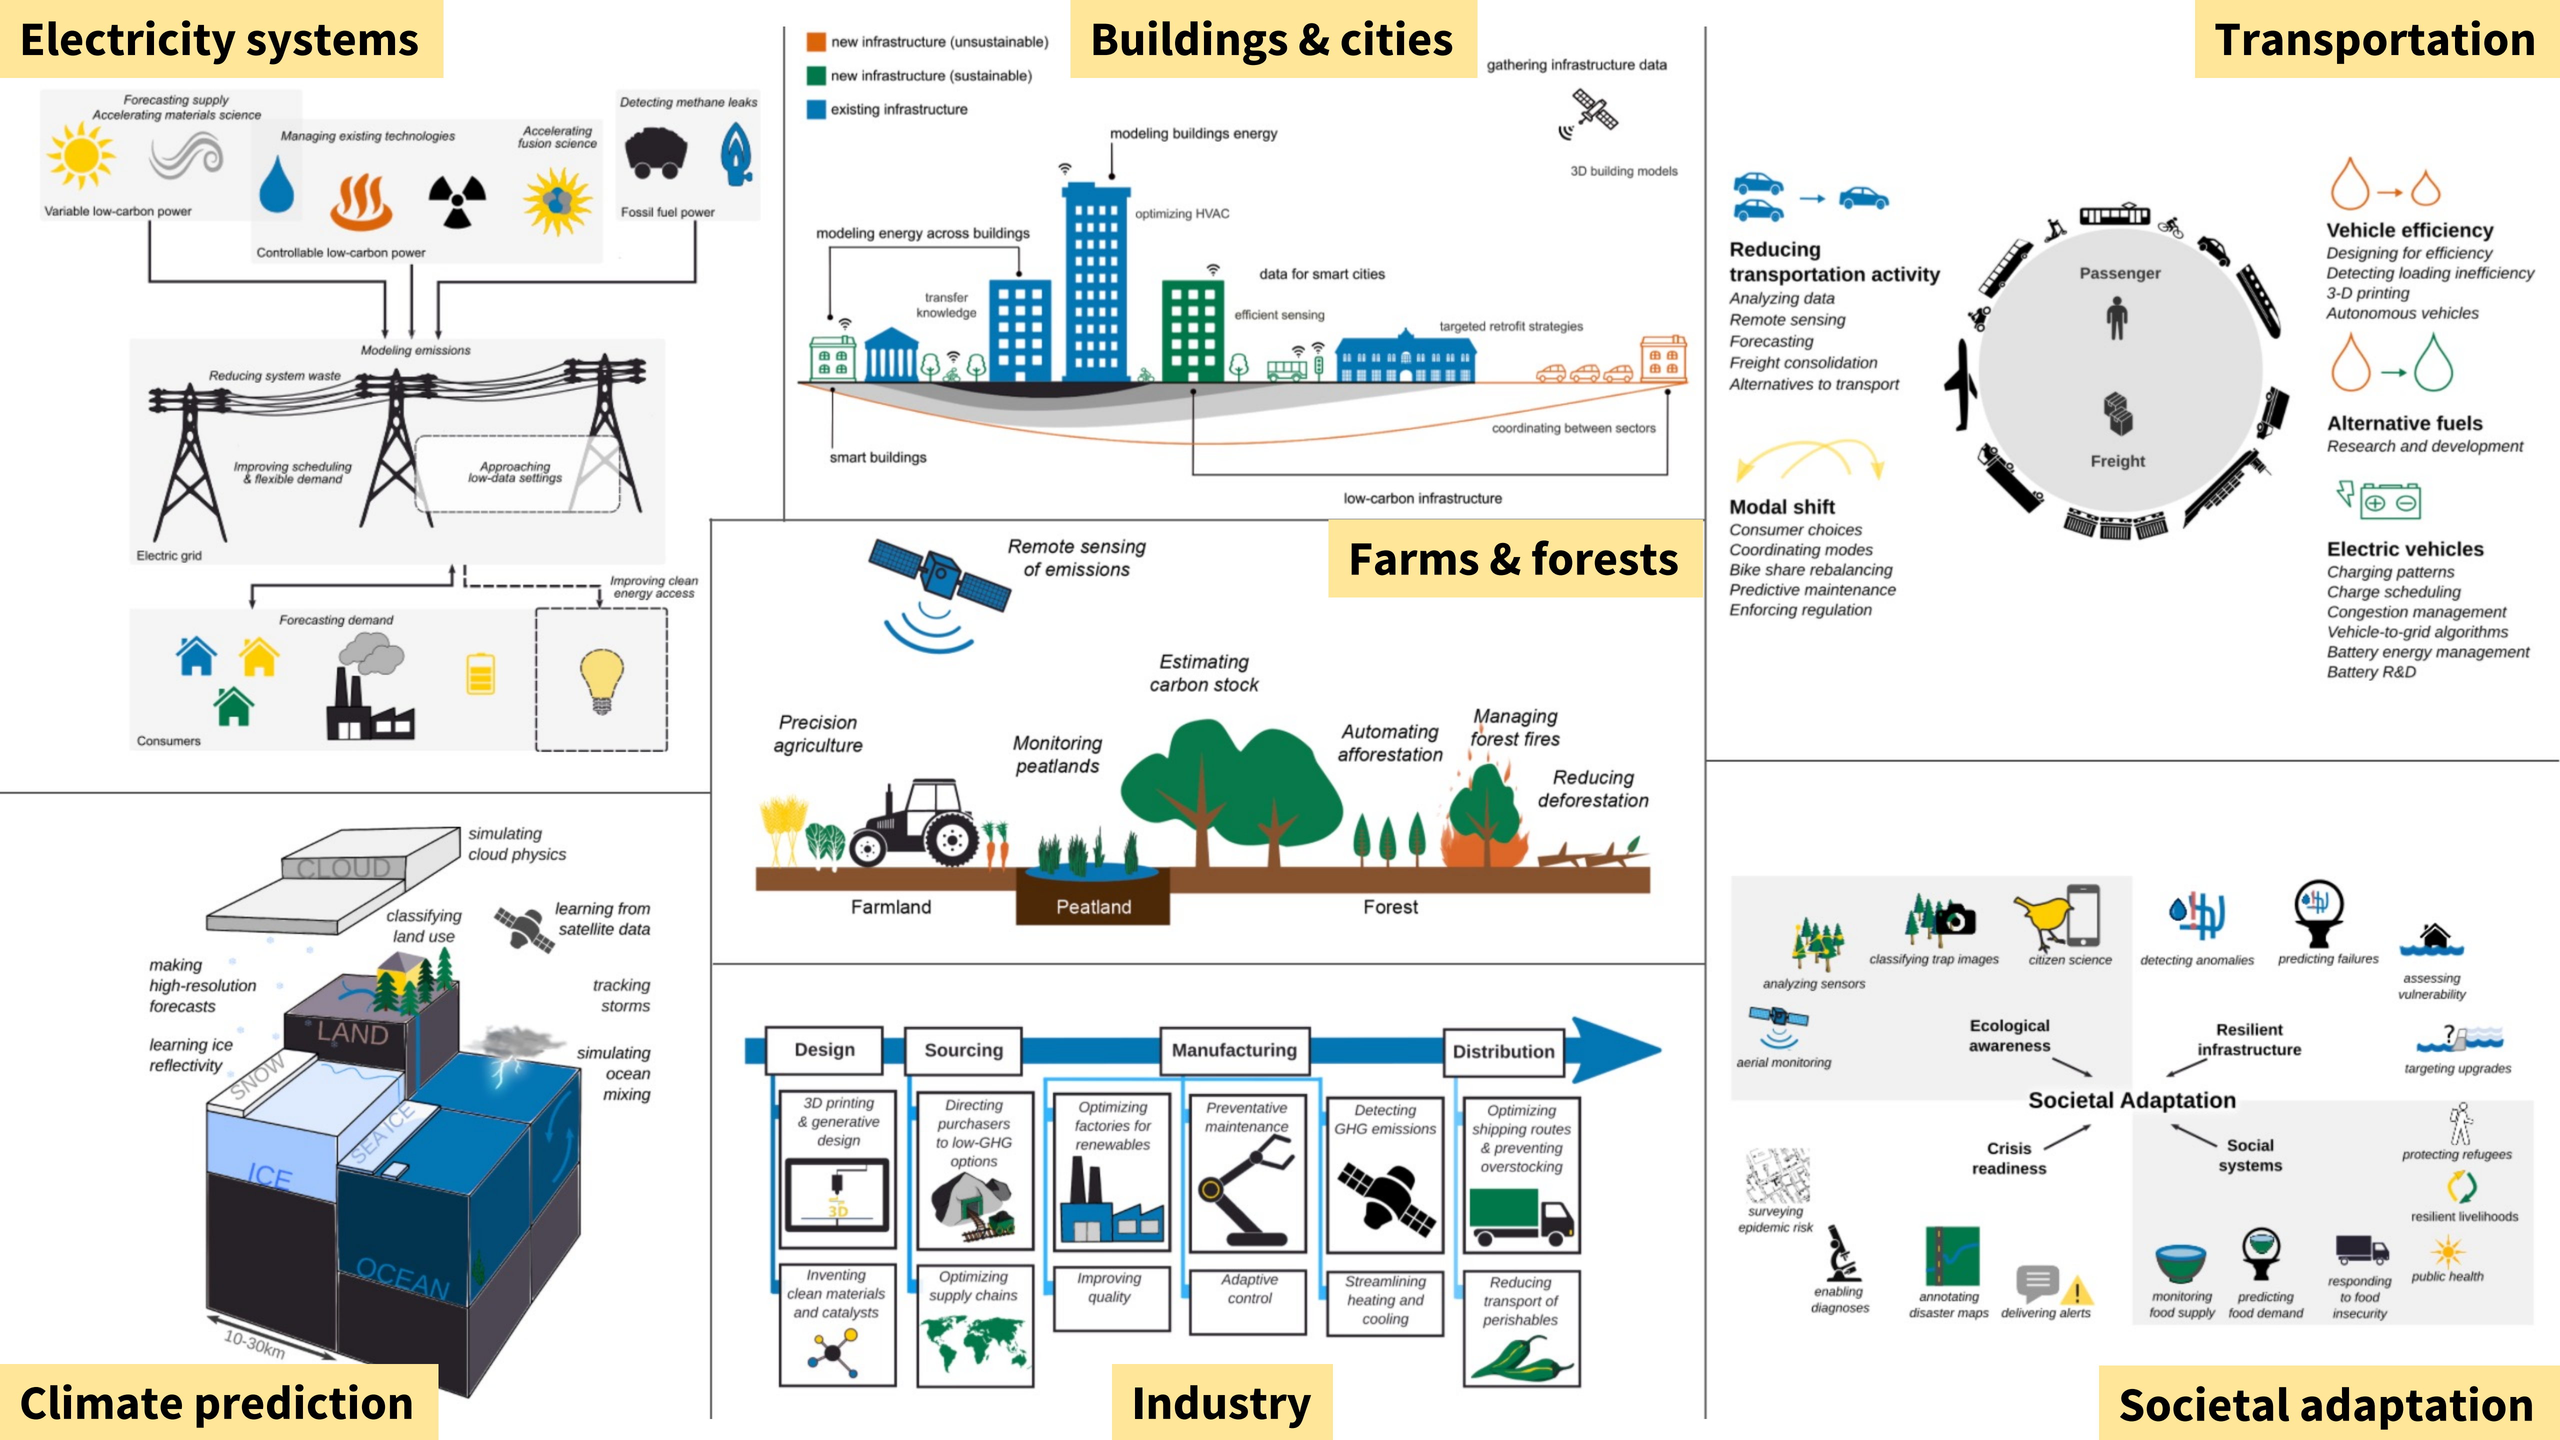

Electricity systems
Forecasting supply
Accelerating materials science
Managing existing technologies
Accelerating fusion science
Detecting methane leaks
Variable low-carbon power
Controllable low-carbon power
Fossil fuel power
Modeling emissions
Reducing system waste
Improving scheduling & flexible demand
Approaching low-data settings
Electric grid
Improving clean energy access
Forecasting demand
Consumers
Buildings & cities
new infrastructure (unsustainable)
new infrastructure (sustainable)
existing infrastructure
gathering infrastructure data
3D building models
modeling buildings energy
optimizing HVAC
modeling energy across buildings
transfer knowledge
data for smart cities
efficient sensing
targeted retrofit strategies
coordinating between sectors
smart buildings
low-carbon infrastructure
Transportation
Reducing transportation activity
Analyzing data
Remote sensing
Forecasting
Freight consolidation
Alternatives to transport
Modal shift
Consumer choices
Coordinating modes
Bike share rebalancing
Predictive maintenance
Enforcing regulation
Passenger
Freight
Vehicle efficiency
Designing for efficiency
Detecting loading inefficiency
3-D printing
Autonomous vehicles
Alternative fuels
Research and development
Electric vehicles
Charging patterns
Charge scheduling
Congestion management
Vehicle-to-grid algorithms
Battery energy management
Battery R&D
Farms & forests
Remote sensing of emissions
Estimating carbon stock
Precision agriculture
Monitoring peatlands
Automating afforestation
Managing forest fires
Reducing deforestation
Farmland
Peatland
Forest
Climate prediction
simulating cloud physics
CLOUD
classifying land use
learning from satellite data
making high-resolution forecasts
tracking storms
LAND
learning ice reflectivity
simulating ocean mixing
SNOW
SEA ICE
ICE
OCEAN
10-30km
Industry
Design
Sourcing
Manufacturing
Distribution
3D printing & generative design
Directing purchasers to low-GHG options
Optimizing factories for renewables
Preventative maintenance
Detecting GHG emissions
Optimizing shipping routes & preventing overstocking
Inventing clean materials and catalysts
Optimizing supply chains
Improving quality
Adaptive control
Streamlining heating and cooling
Reducing transport of perishables
Societal adaptation
analyzing sensors
classifying trap images
citizen science
detecting anomalies
predicting failures
assessing vulnerability
aerial monitoring
Ecological awareness
Resilient infrastructure
targeting upgrades
Societal Adaptation
Crisis readiness
Social systems
surveying epidemic risk
protecting refugees
resilient livelihoods
enabling diagnoses
annotating disaster maps
delivering alerts
monitoring food supply
predicting food demand
responding to food insecurity
public health