

Machine Learning for Environmental Grand Challenges

Shakir Mohamed

Research Scientist, DeepMind

Principles to Products

Applications

**Assistive
Technology**

**Advancing
Science**

**Climate and
Energy**

Healthcare

**Fairness and
Safety**

**Autonomous
systems**

Reasoning

Planning

Explanation

Rapid Learning

**World
Simulation**

**Objects and
Relations**

Information

Uncertainty

Information Gain

Causality

Prediction

Principles

**Probability
Theory**

**Bayesian
Analysis**

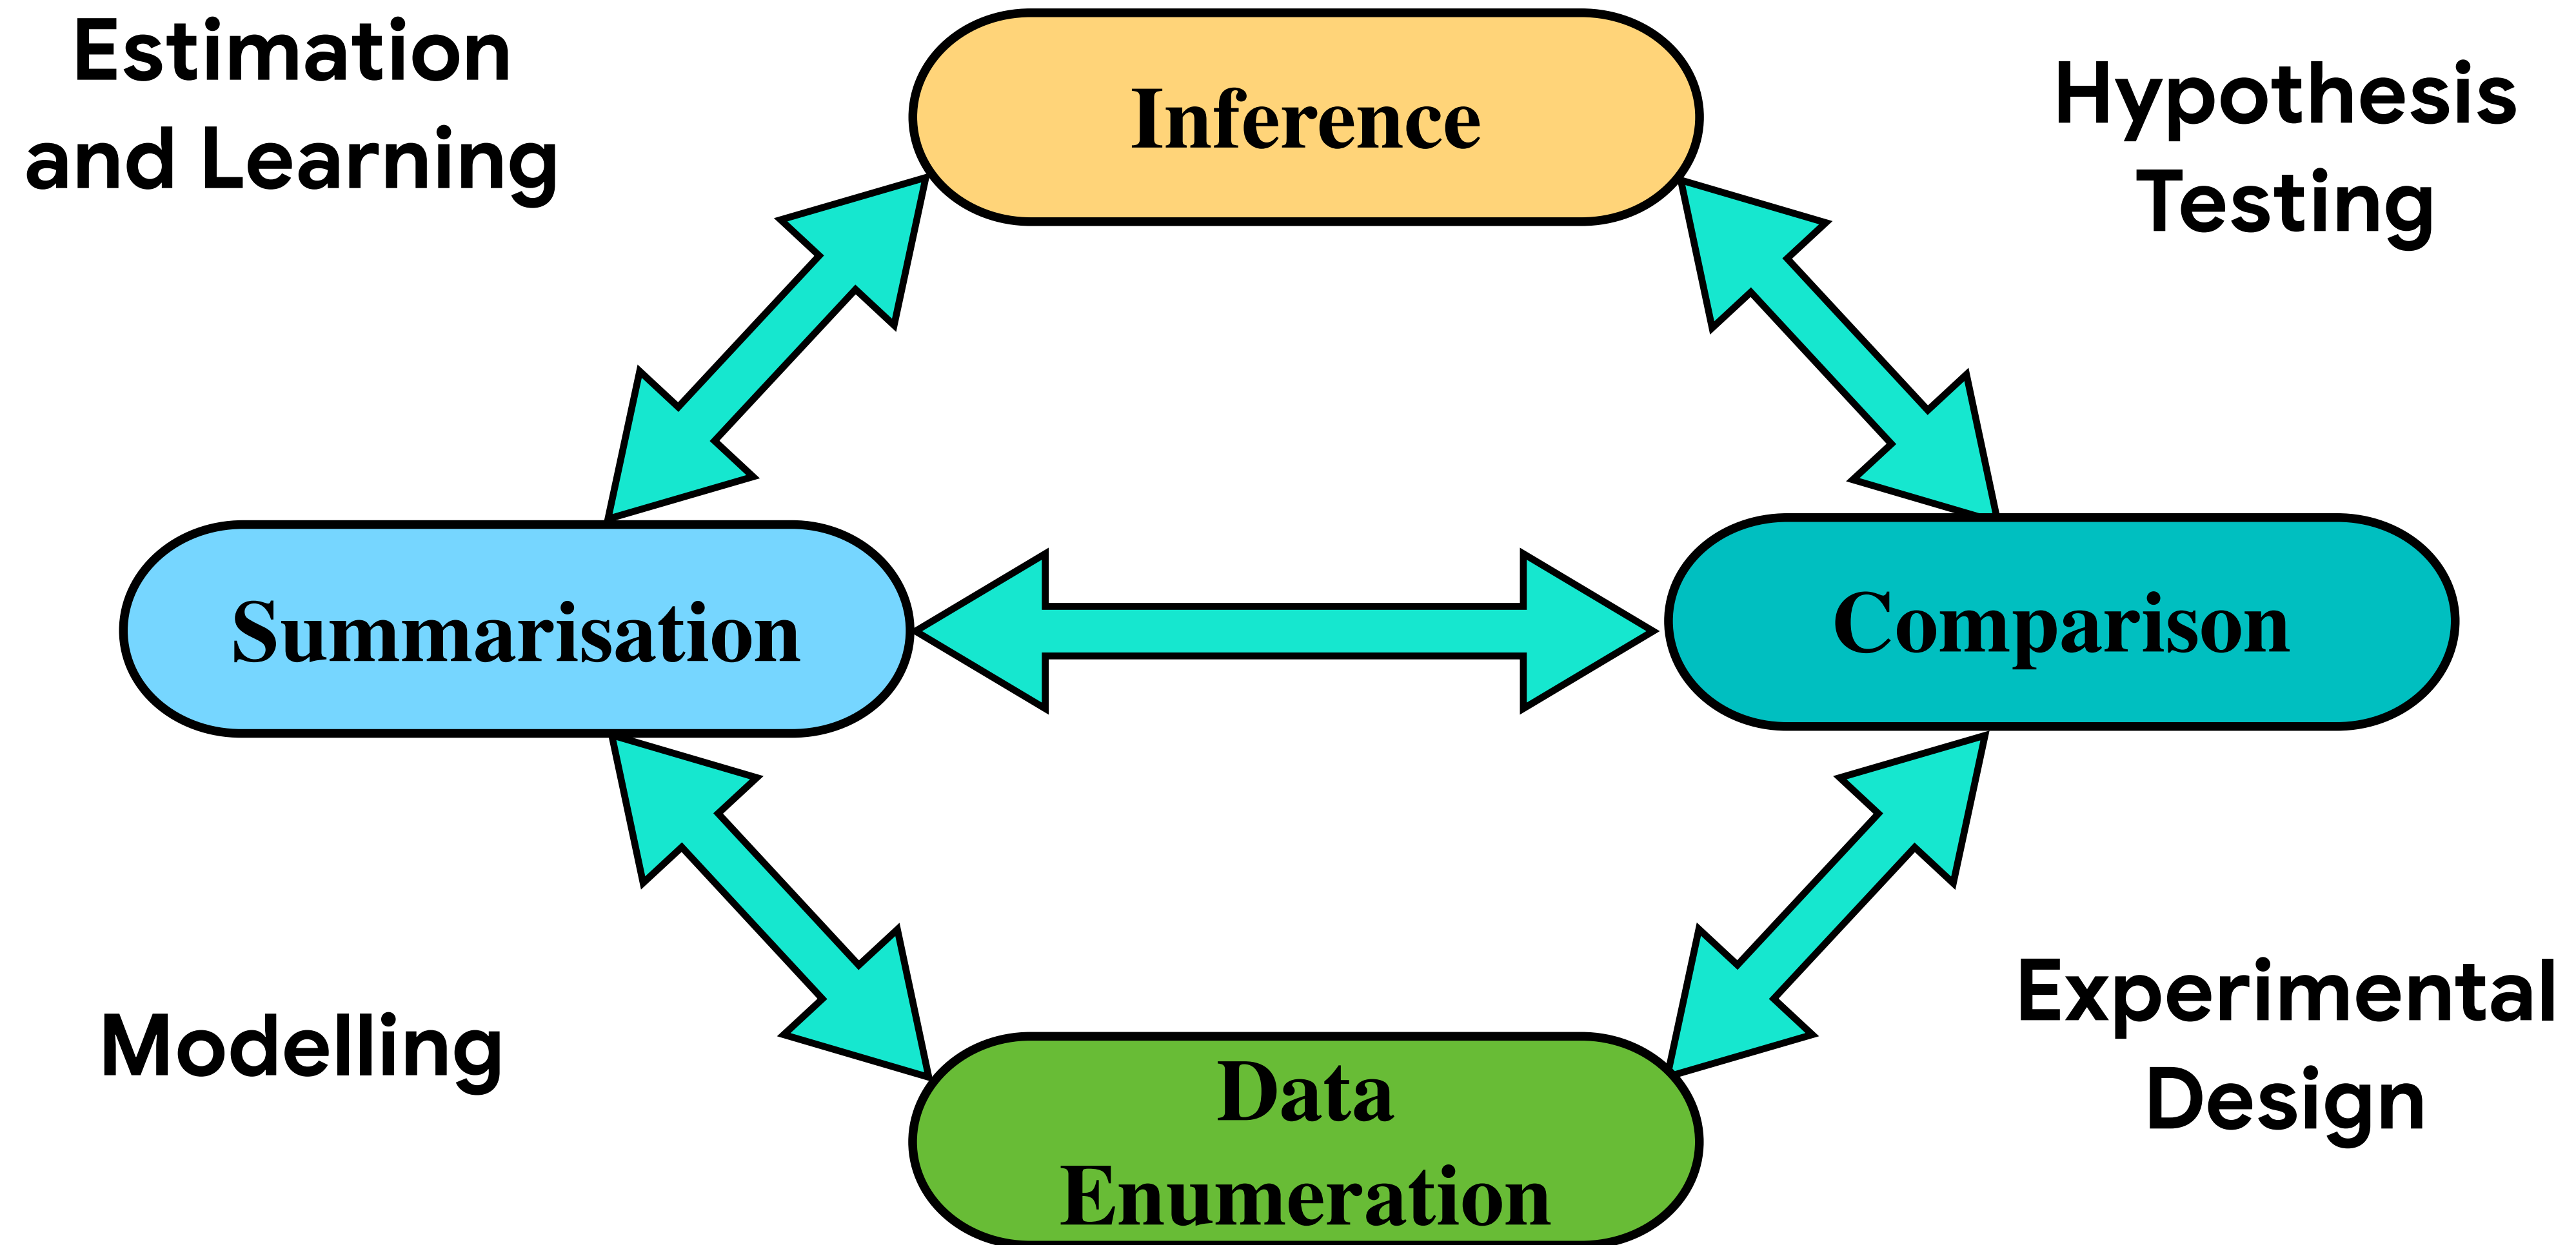
**Hypothesis
Testing**

**Estimation
Theory**

Asymptotics



Statistical Operations



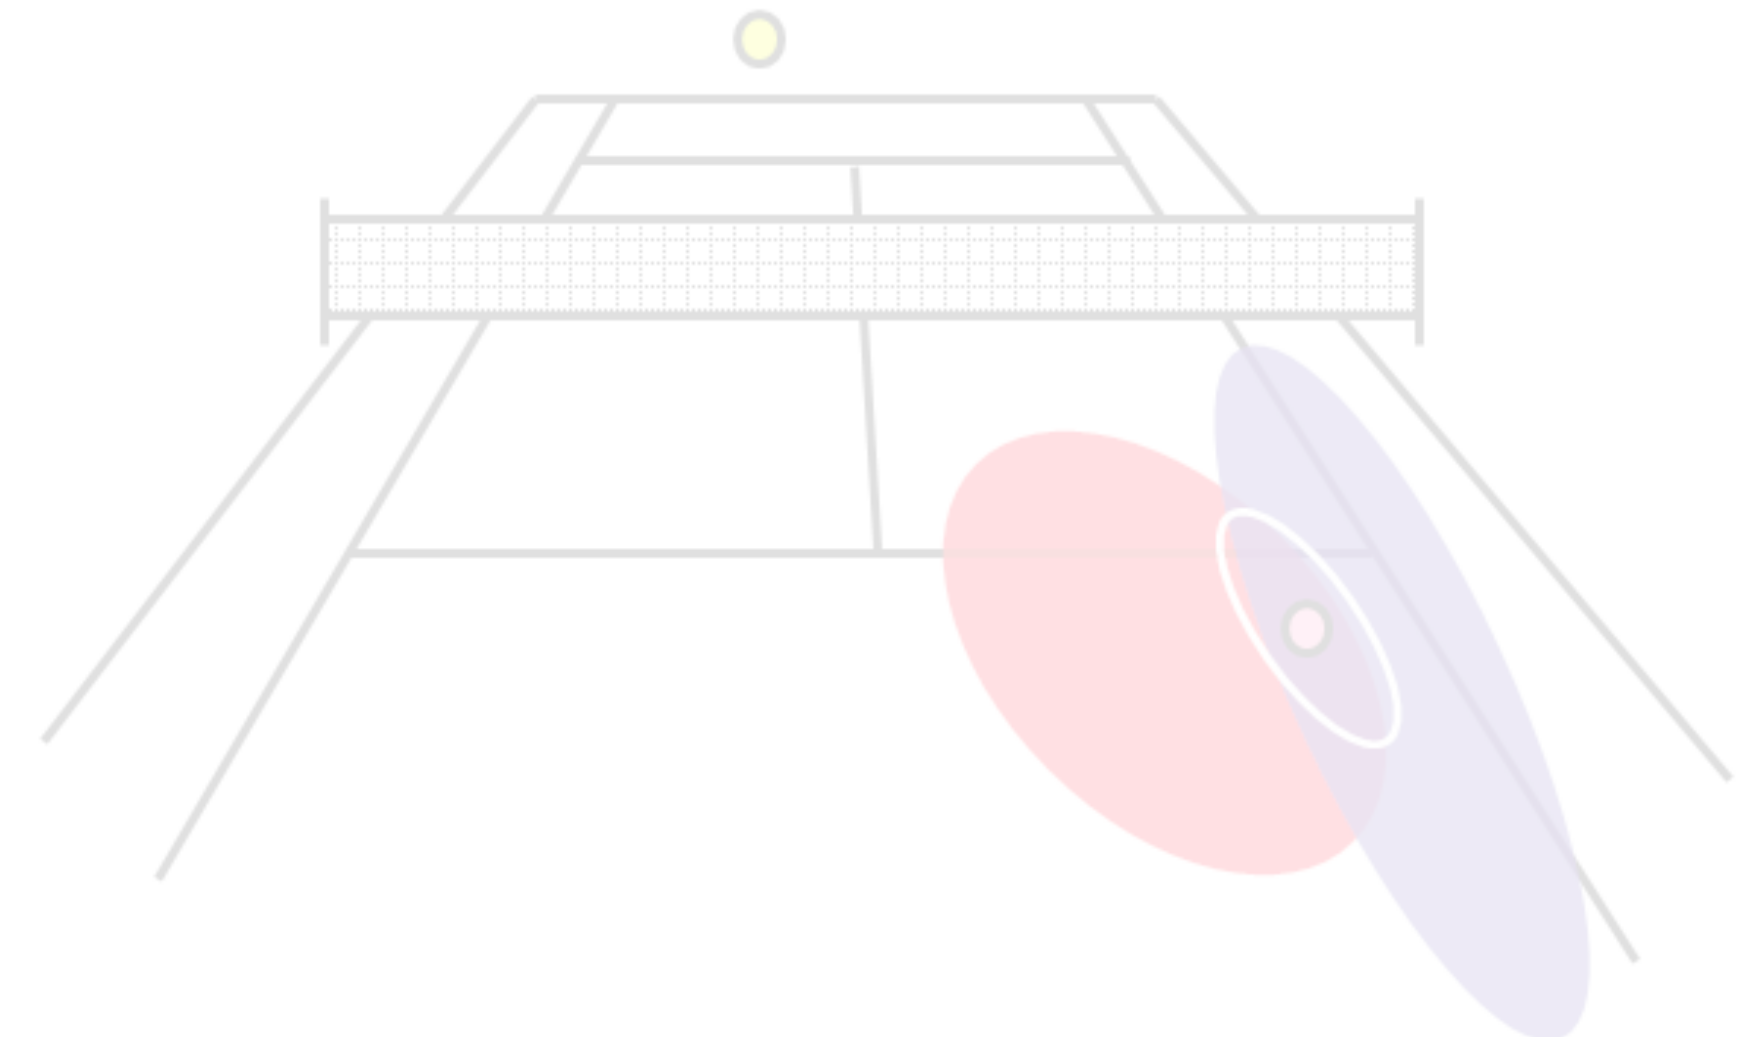
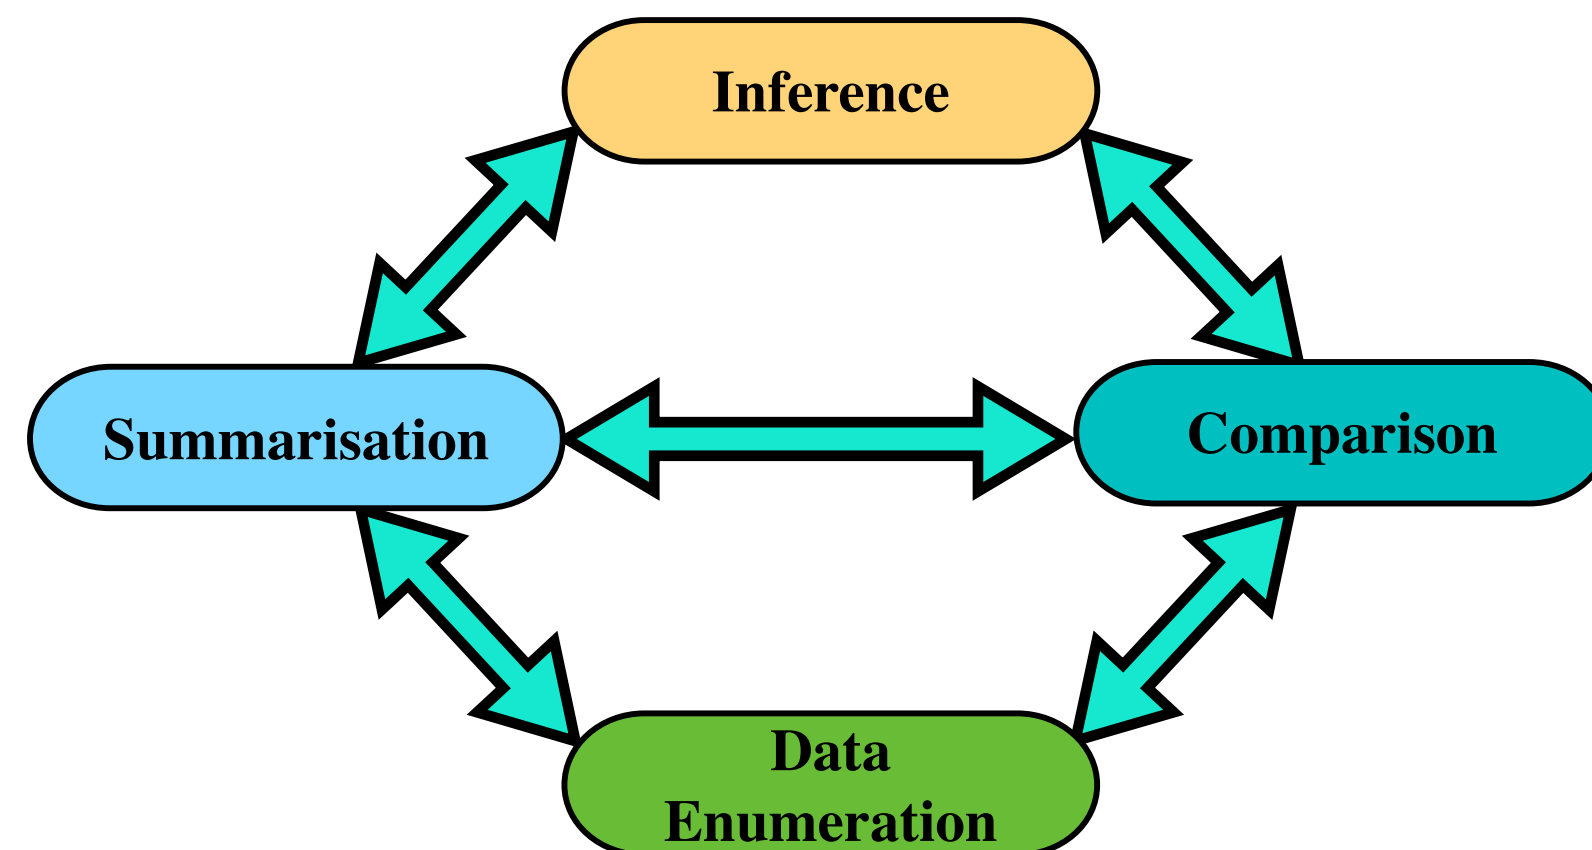
Statistical Operations

Inference

What we can
know about our data

Decision-making

What we can
do with our data.



Part I: Pathways in Machine Learning

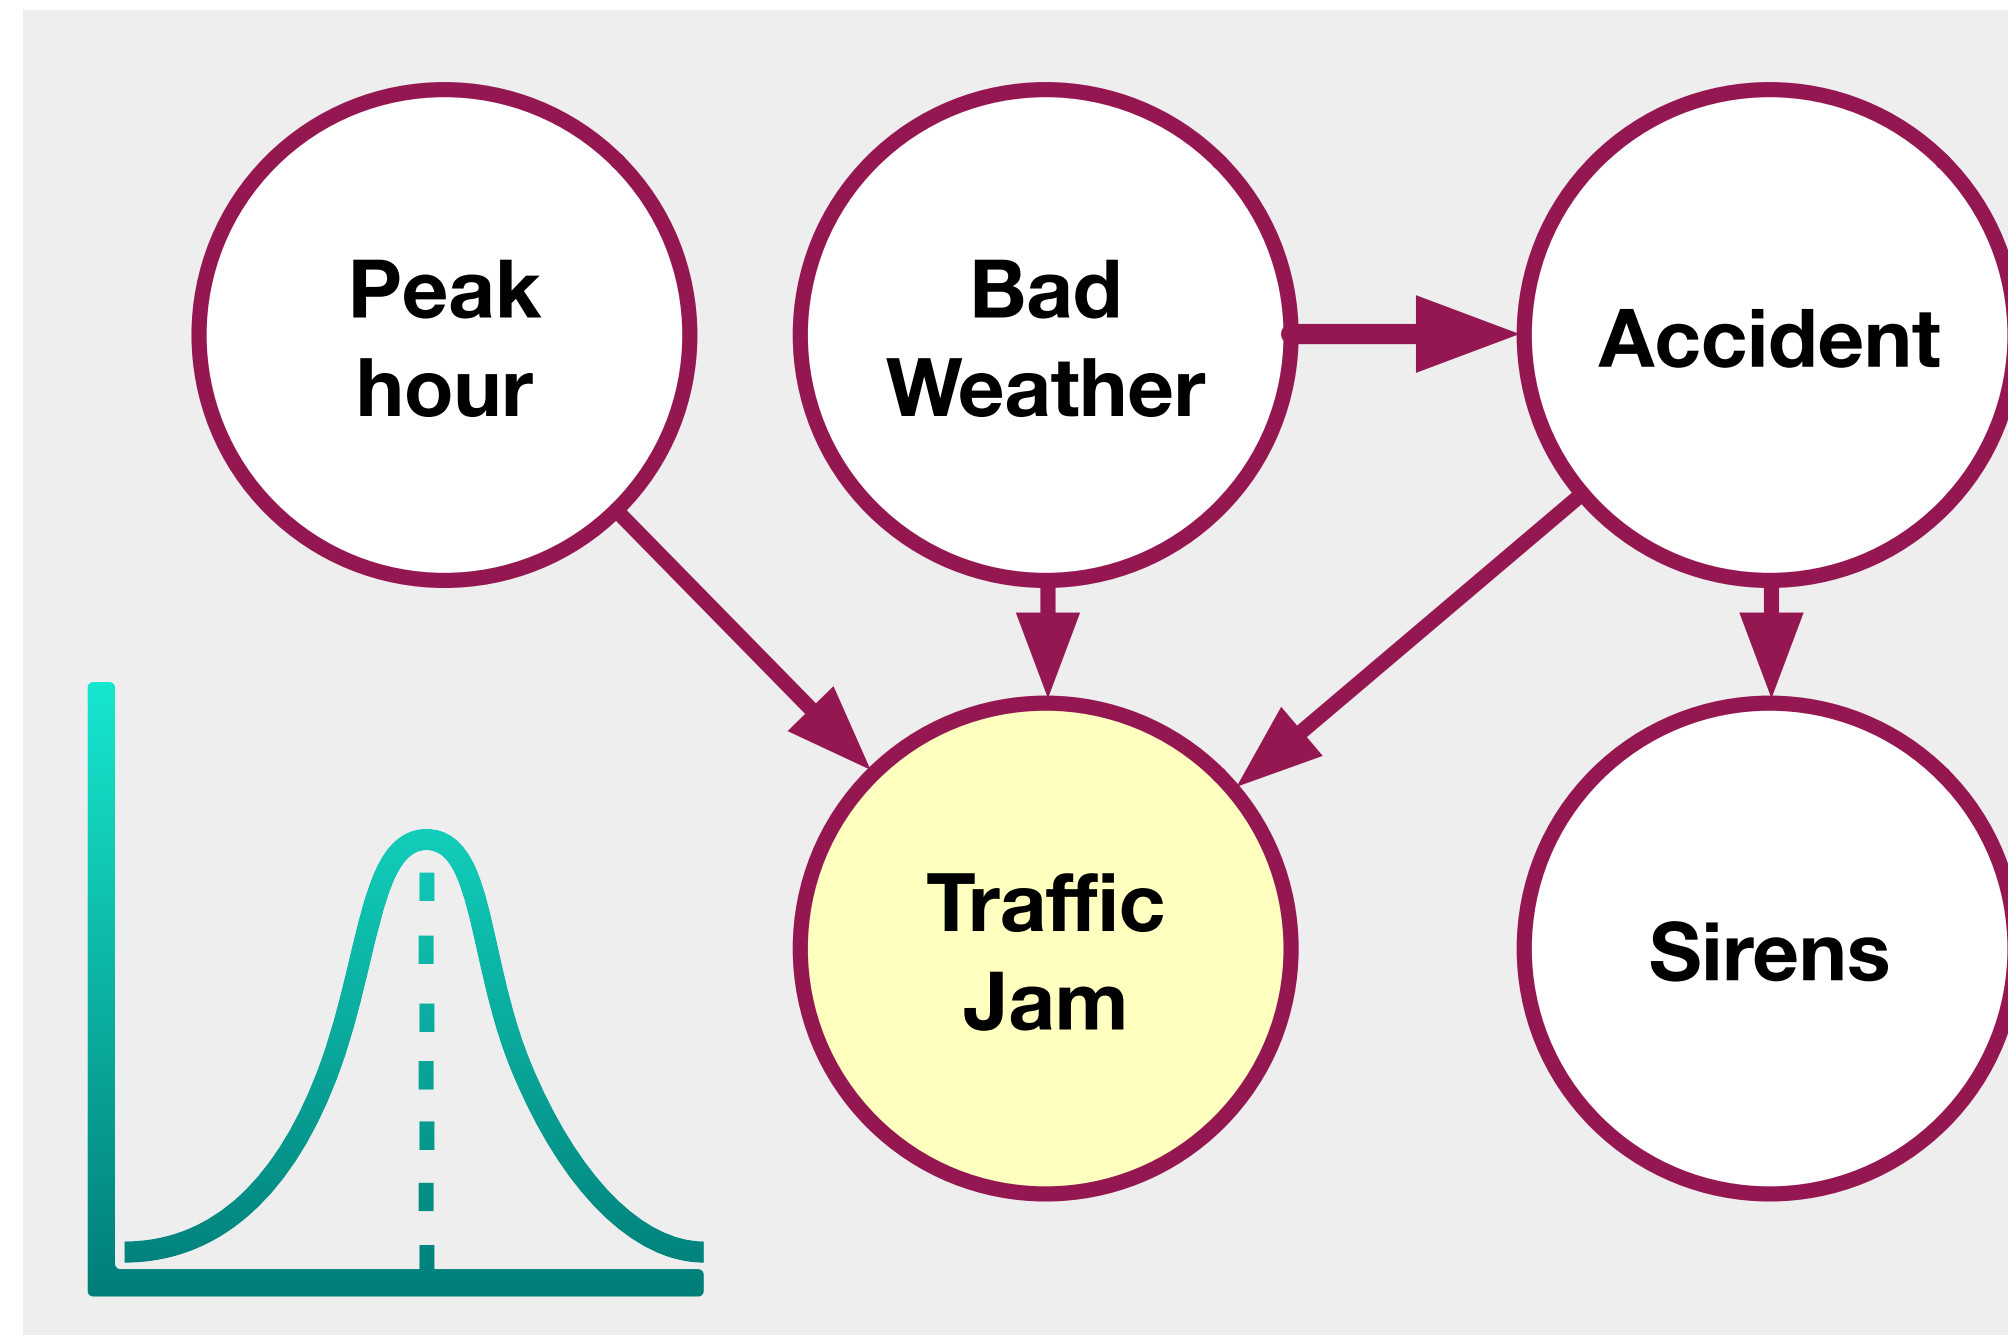
Shakir Mohamed

Research Scientist, DeepMind

On Models

Model: Description of the world, of data, of potential scenarios, of processes.

A probabilistic model writes out these models using the language of probability



Equations of motion (ECWMF model)

$$\frac{\partial U}{\partial t} + \frac{1}{a \cos^2 \theta} \left\{ U \frac{\partial U}{\partial \lambda} + v \cos \theta \frac{\partial U}{\partial \theta} \right\} + \eta \frac{\partial U}{\partial \eta} \quad \text{East-west wind}$$

$$(-fv) + \frac{1}{a} \left\{ \frac{\partial \Phi}{\partial \lambda} + R_{dy} T_v \frac{\partial}{\partial \lambda} (\ln p) \right\} = \underline{P_U + K_U}$$

$$\frac{\partial V}{\partial t} + \frac{1}{a \cos^2 \theta} \left\{ U \frac{\partial V}{\partial \lambda} + V \cos \theta \frac{\partial V}{\partial \theta} + \sin \theta (U^2 + V^2) \right\} + \eta \frac{\partial V}{\partial \eta} \quad \text{North-south wind}$$

$$+ fU + \frac{\cos \theta}{a} \left\{ \frac{\partial \Phi}{\partial \theta} + R_{dy} T_v \frac{\partial}{\partial \theta} (\ln p) \right\} = \underline{P_V + K_V}$$

$$\frac{\partial T}{\partial t} + \frac{1}{a \cos^2 \theta} \left\{ U \frac{\partial T}{\partial \lambda} + V \cos \theta \frac{\partial T}{\partial \theta} \right\} + \eta \frac{\partial T}{\partial \eta} - \frac{\kappa T_v \omega}{(1 + (\delta - 1)q)p} = \underline{P_T + K_T} \quad \text{Temperature}$$

$$\frac{\partial q}{\partial t} = \frac{1}{a \cos^2 \theta} \left\{ U \frac{\partial q}{\partial \lambda} + V \cos \theta \frac{\partial q}{\partial \theta} \right\} = \eta \frac{\partial q}{\partial \eta} = \underline{P_q + K_q} \quad \text{Humidity}$$

$$\frac{\partial}{\partial t} \left(\frac{\partial p}{\partial \eta} \right) + \nabla \cdot \left(\mathbf{v}_H \frac{\partial p}{\partial \eta} \right) + \frac{\partial}{\partial \eta} \left(\eta \frac{\partial p}{\partial \eta} \right) = 0 \quad \text{Continuity of mass}$$

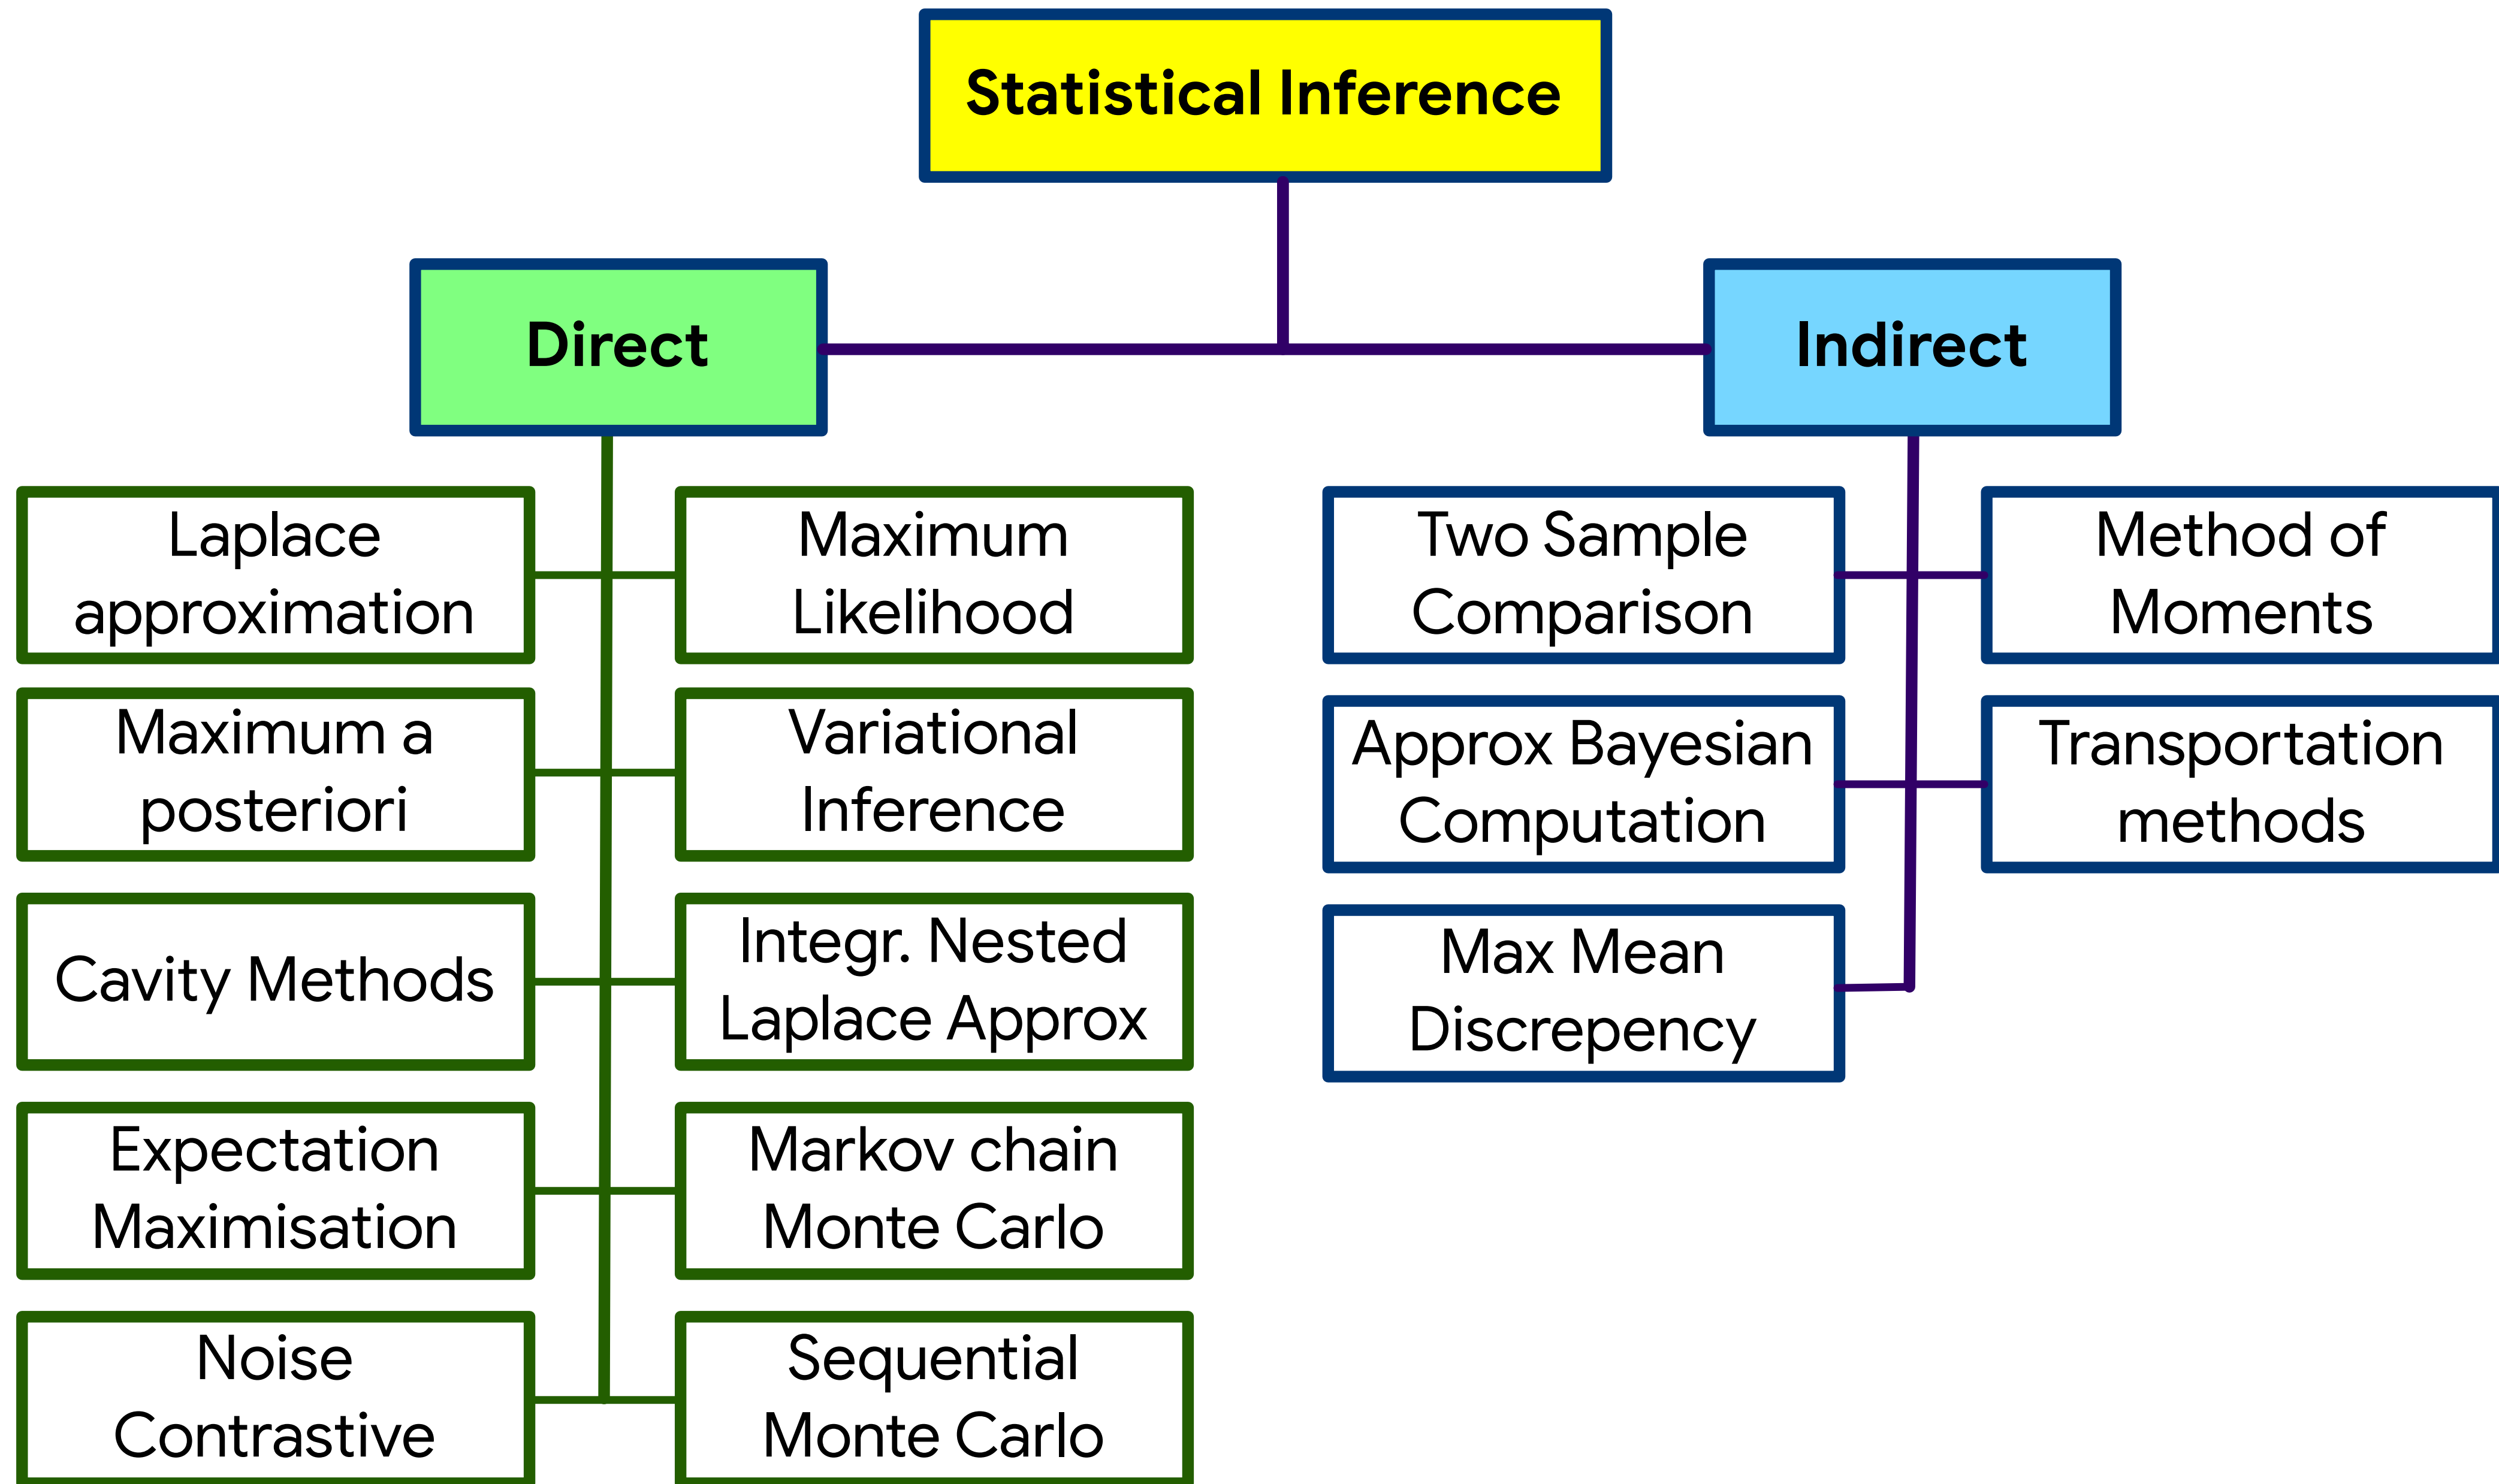
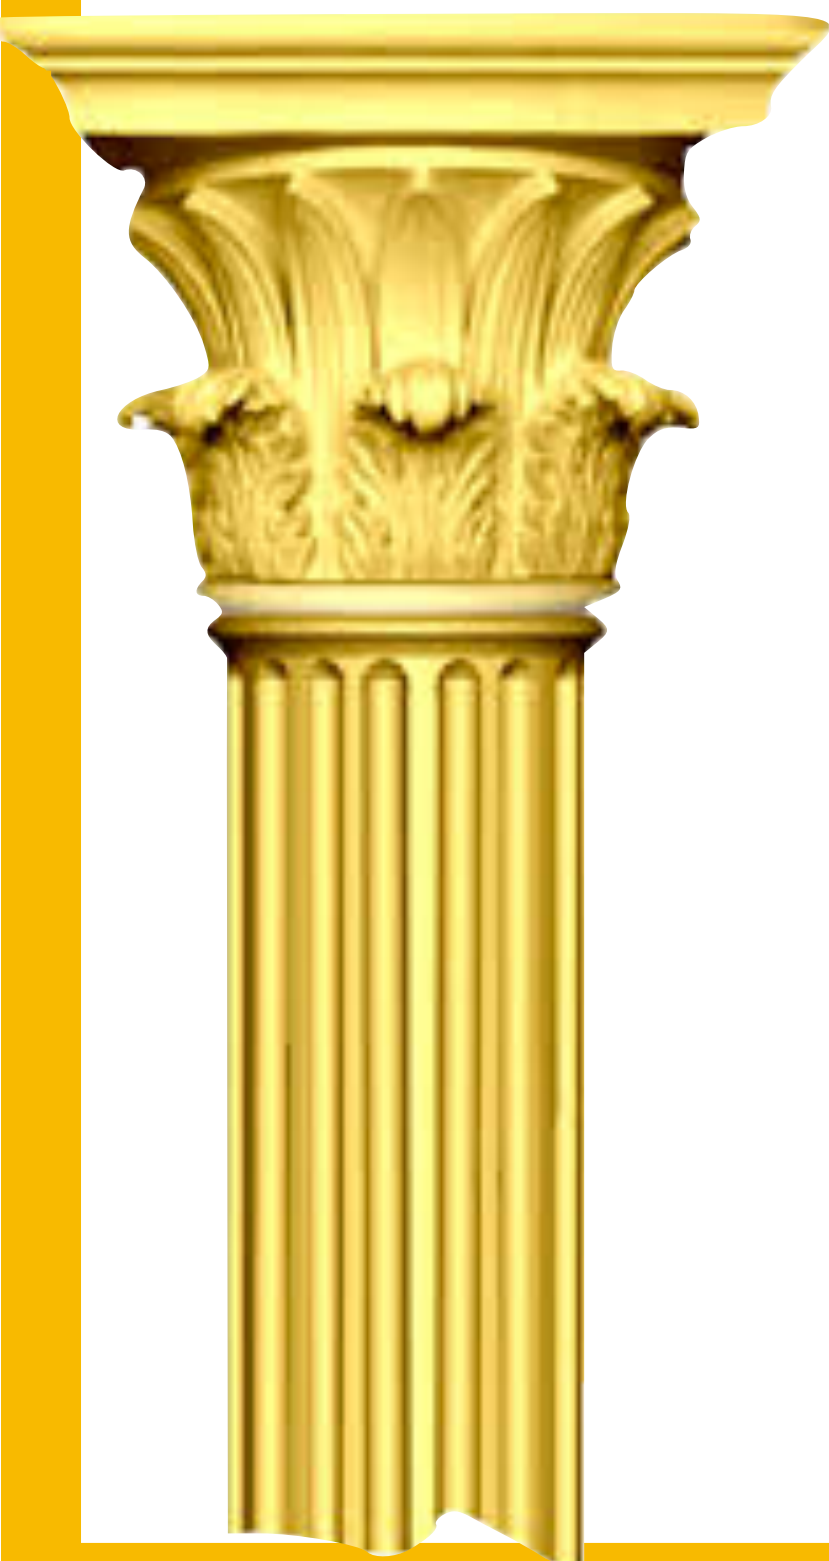
$$\frac{\partial p_{surf}}{\partial t} = - \int_0^1 \nabla \cdot \left(\mathbf{v}_H \frac{\partial p}{\partial \eta} \right) d\eta \quad \text{Surface pressure}$$

Probabilistic models let you learn probability distributions of data.

Most models in machine learning are probabilistic.



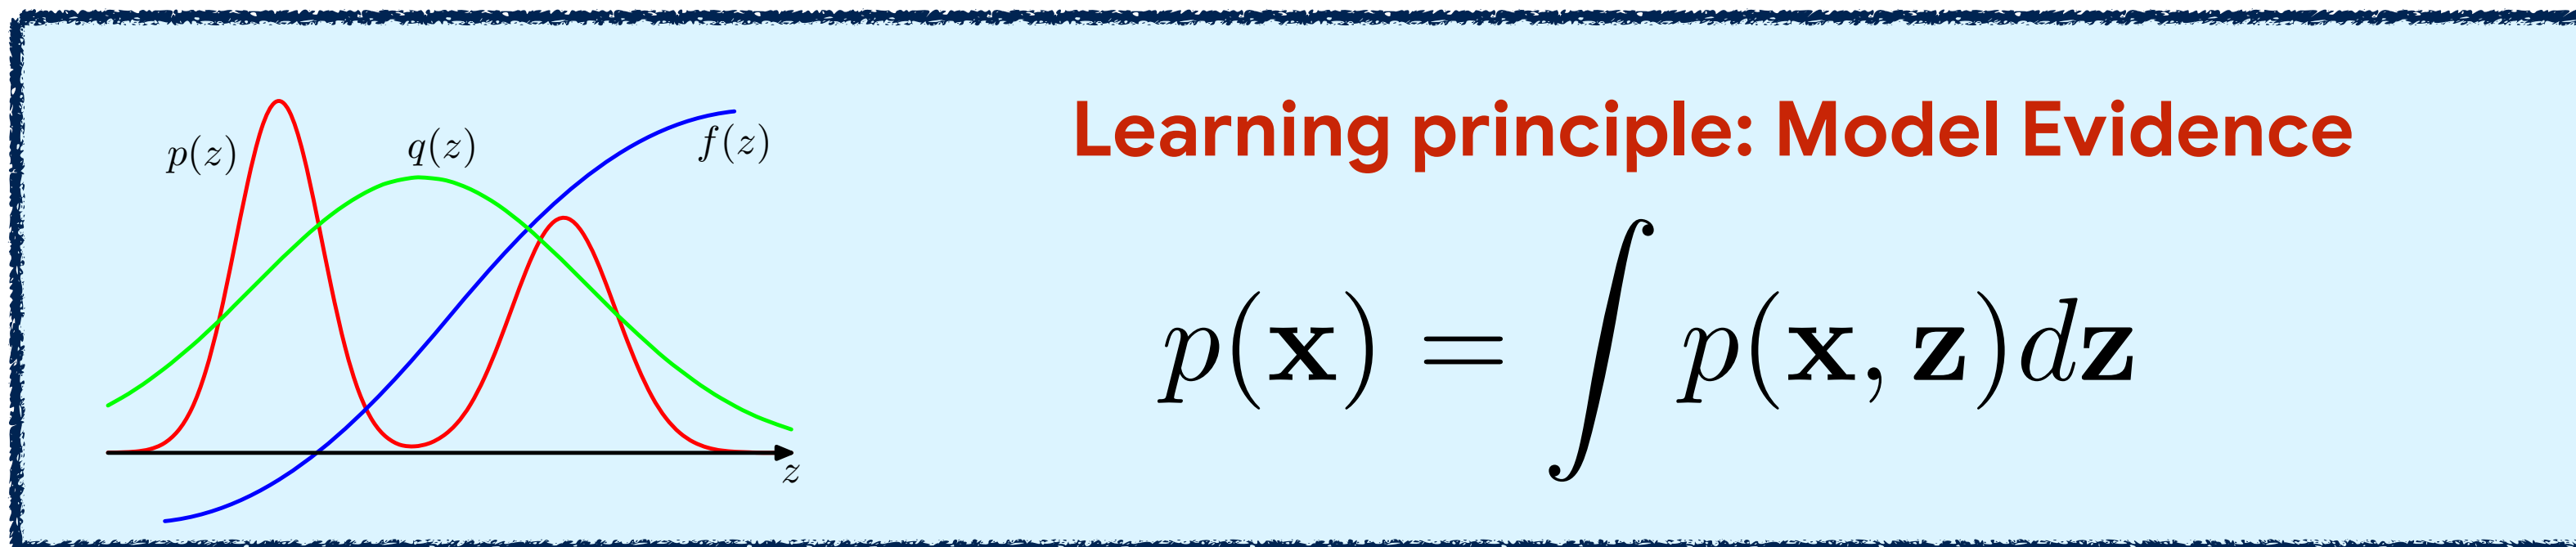
Learning Principles



Model Evidence

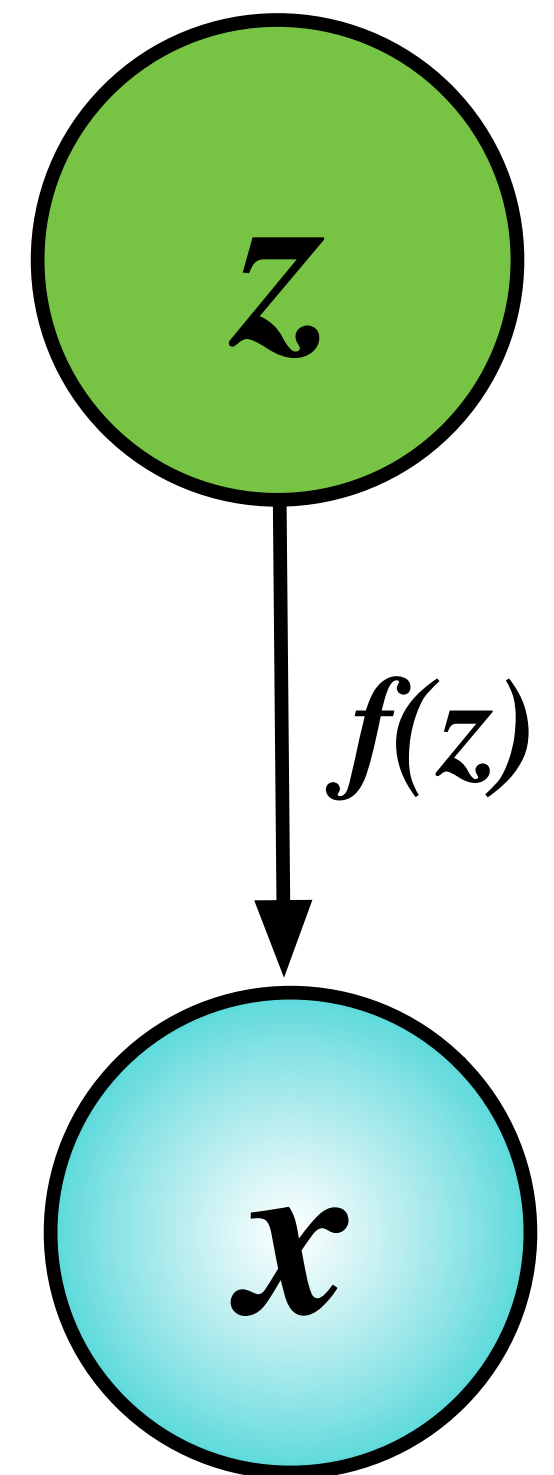
Model evidence (or marginal likelihood, partition function):

Integrating out any global and local variables enables model scoring, comparison, selection, moment estimation, normalisation, posterior computation and prediction.

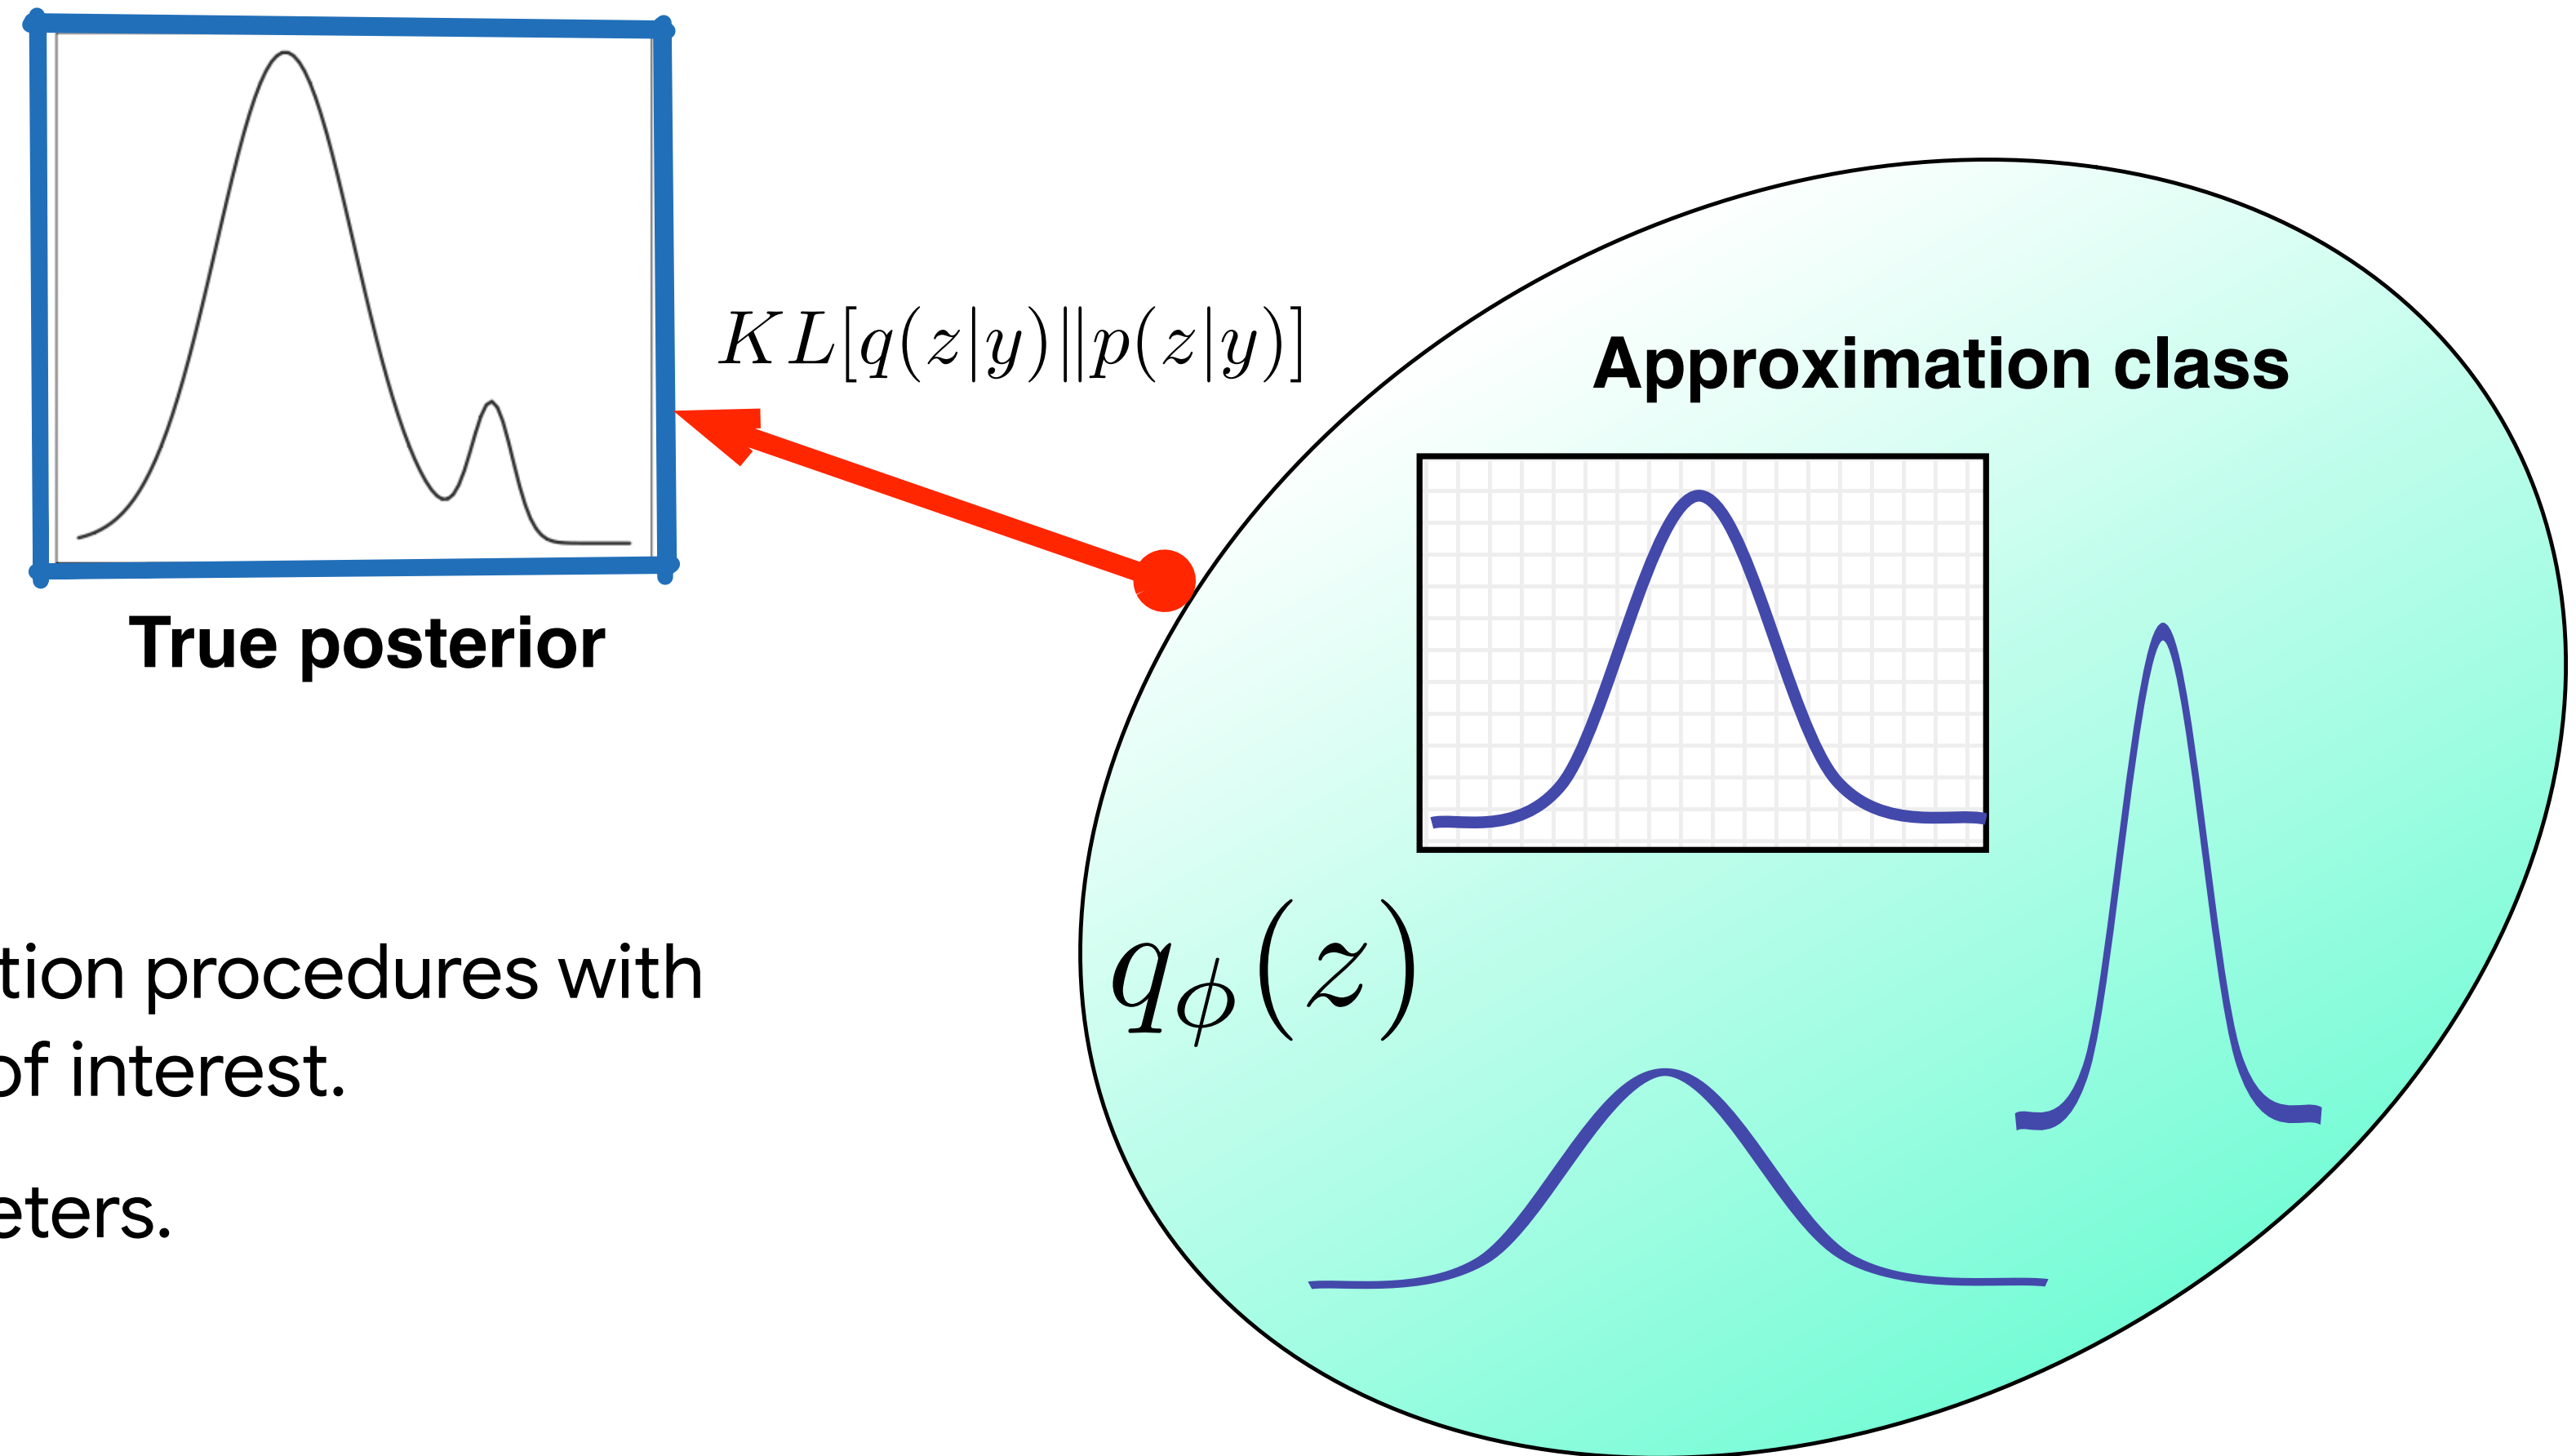


Integral is intractable in general and requires approximation.

Basic idea: Transform the integral into an expectation over a simple, known distribution.



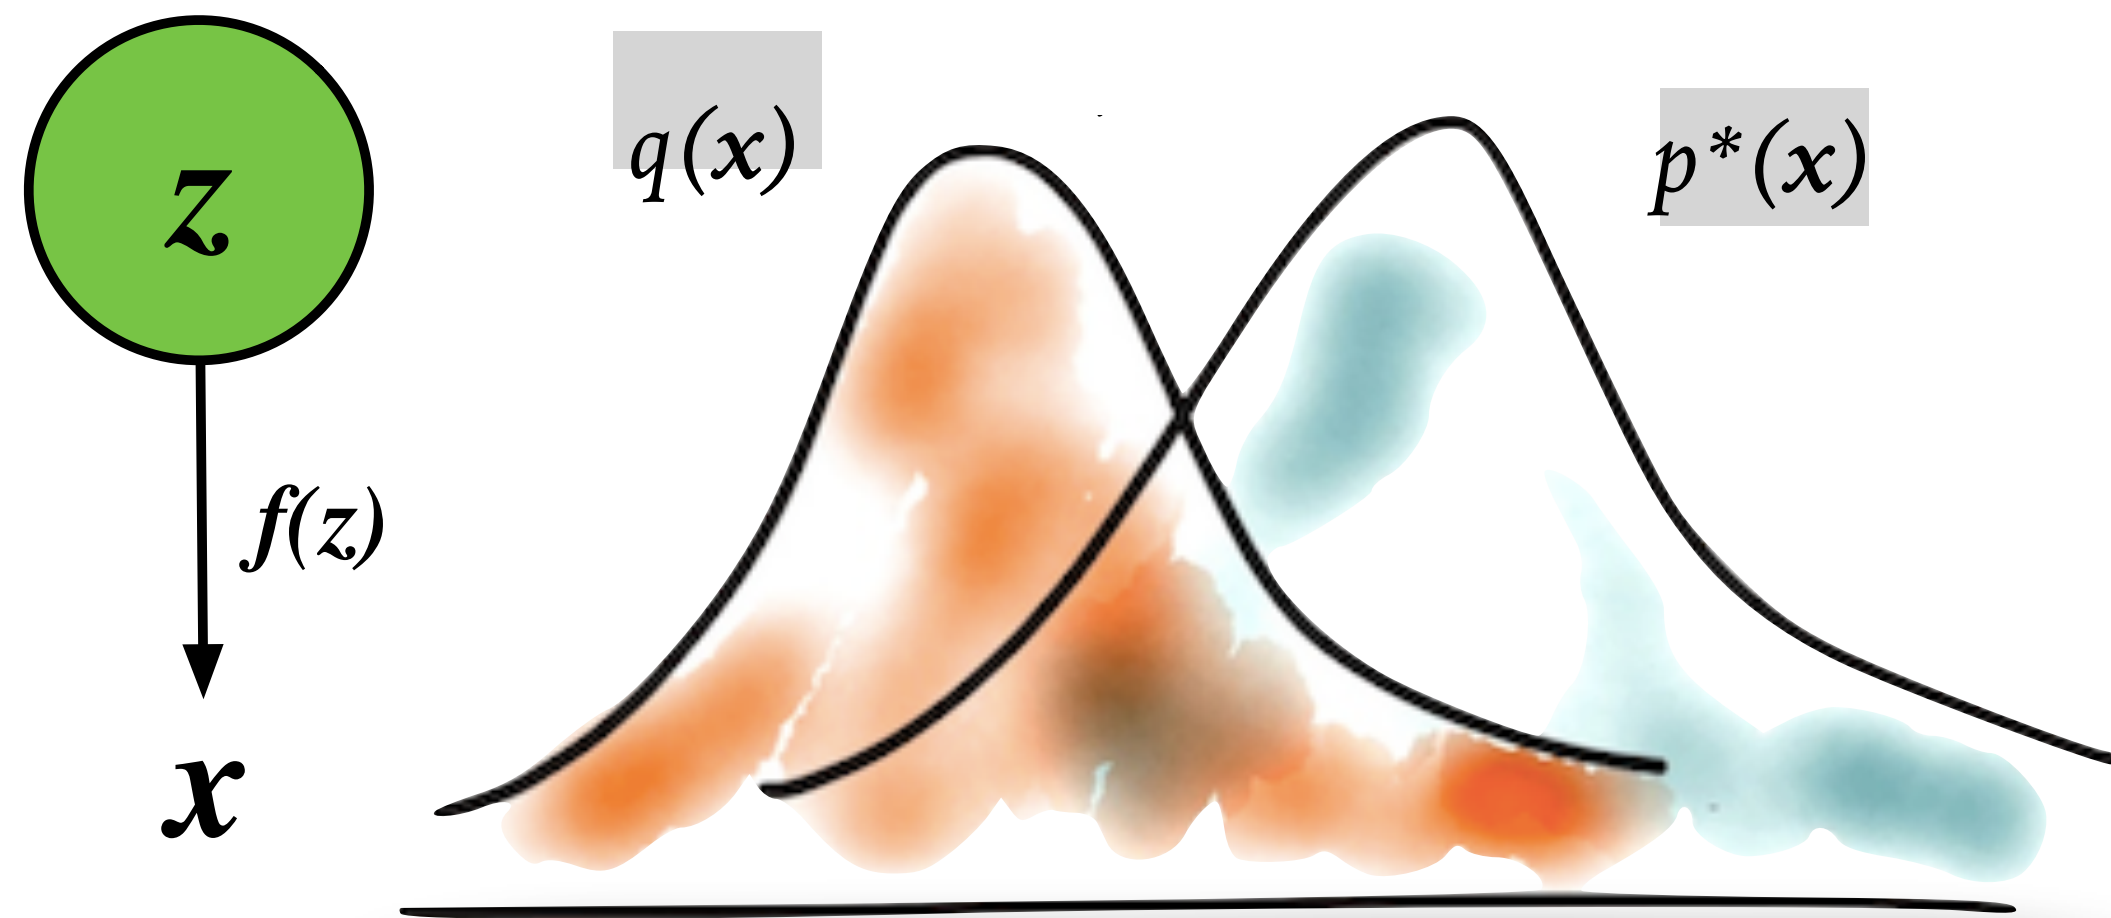
Variational Methods



Deterministic approximation procedures with bounds on probabilities of interest.

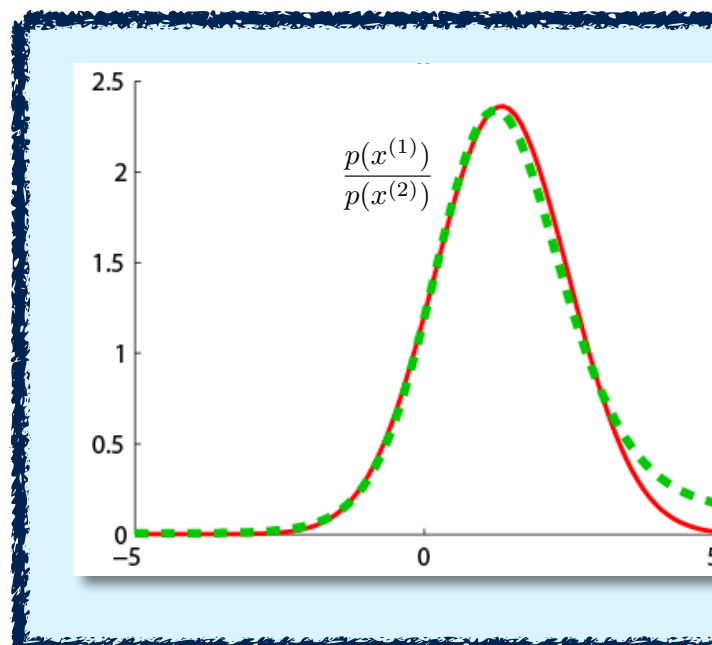
Fit the variational parameters.

Learning by Comparison



We compare the estimated distribution $q(x)$ to the true distribution $p^*(x)$ using samples.

Basic idea:
Transform into learning a model of the density ratio.



Learning principle: Two-sample tests

$$\frac{p^*(\mathbf{x})}{q(\mathbf{x})} = 1 \quad p^*(\mathbf{x}) = q(\mathbf{x})$$

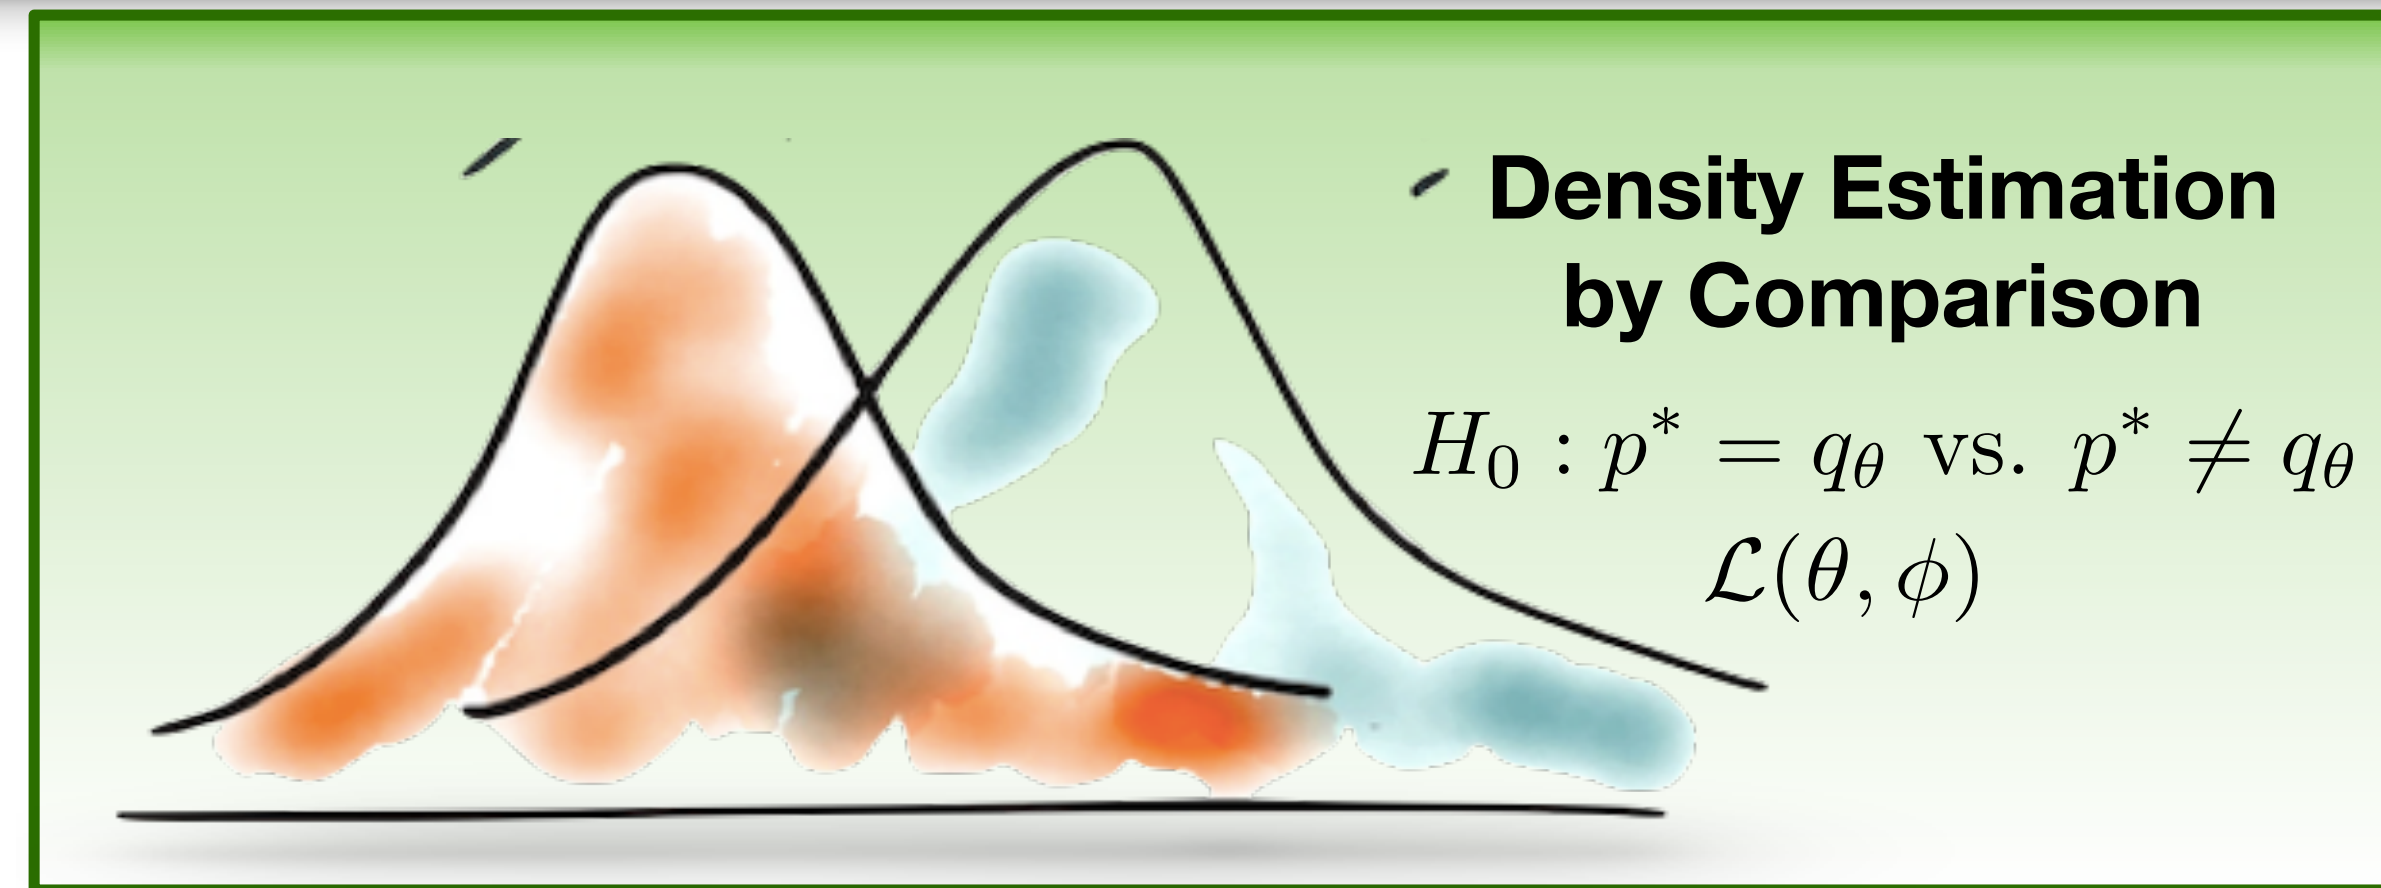
Interest is not in estimating the marginal probabilities, only in how they are related.

Estimation by Comparison

Two steps

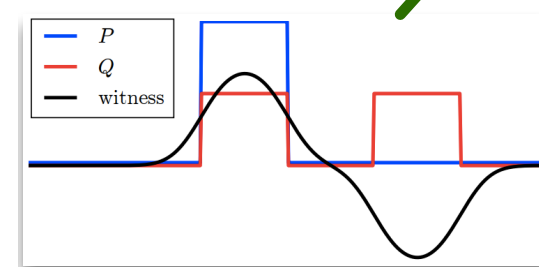
1. Use a hypothesis **test or comparison** to obtain some model to tells how data from our model differs from observed data.

2. **Adjust model** to better match the data distribution using the comparison model from step 1.

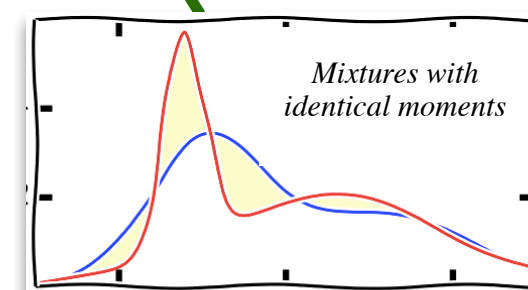


Density Difference

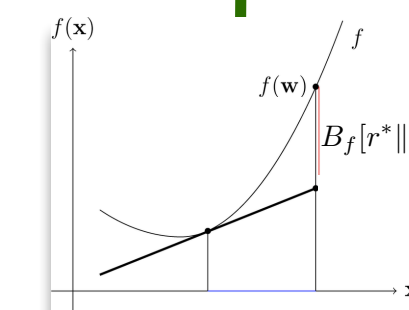
$$r_\phi = p^* - q_\theta$$



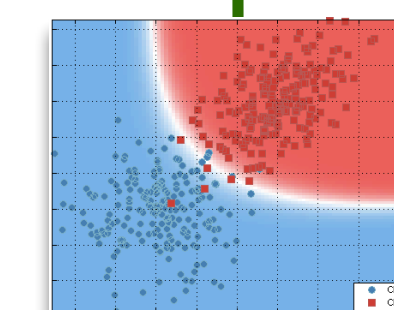
Max Mean Discrepancy



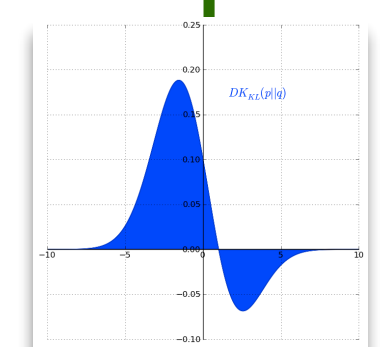
Moment Matching



Bregman Divergence



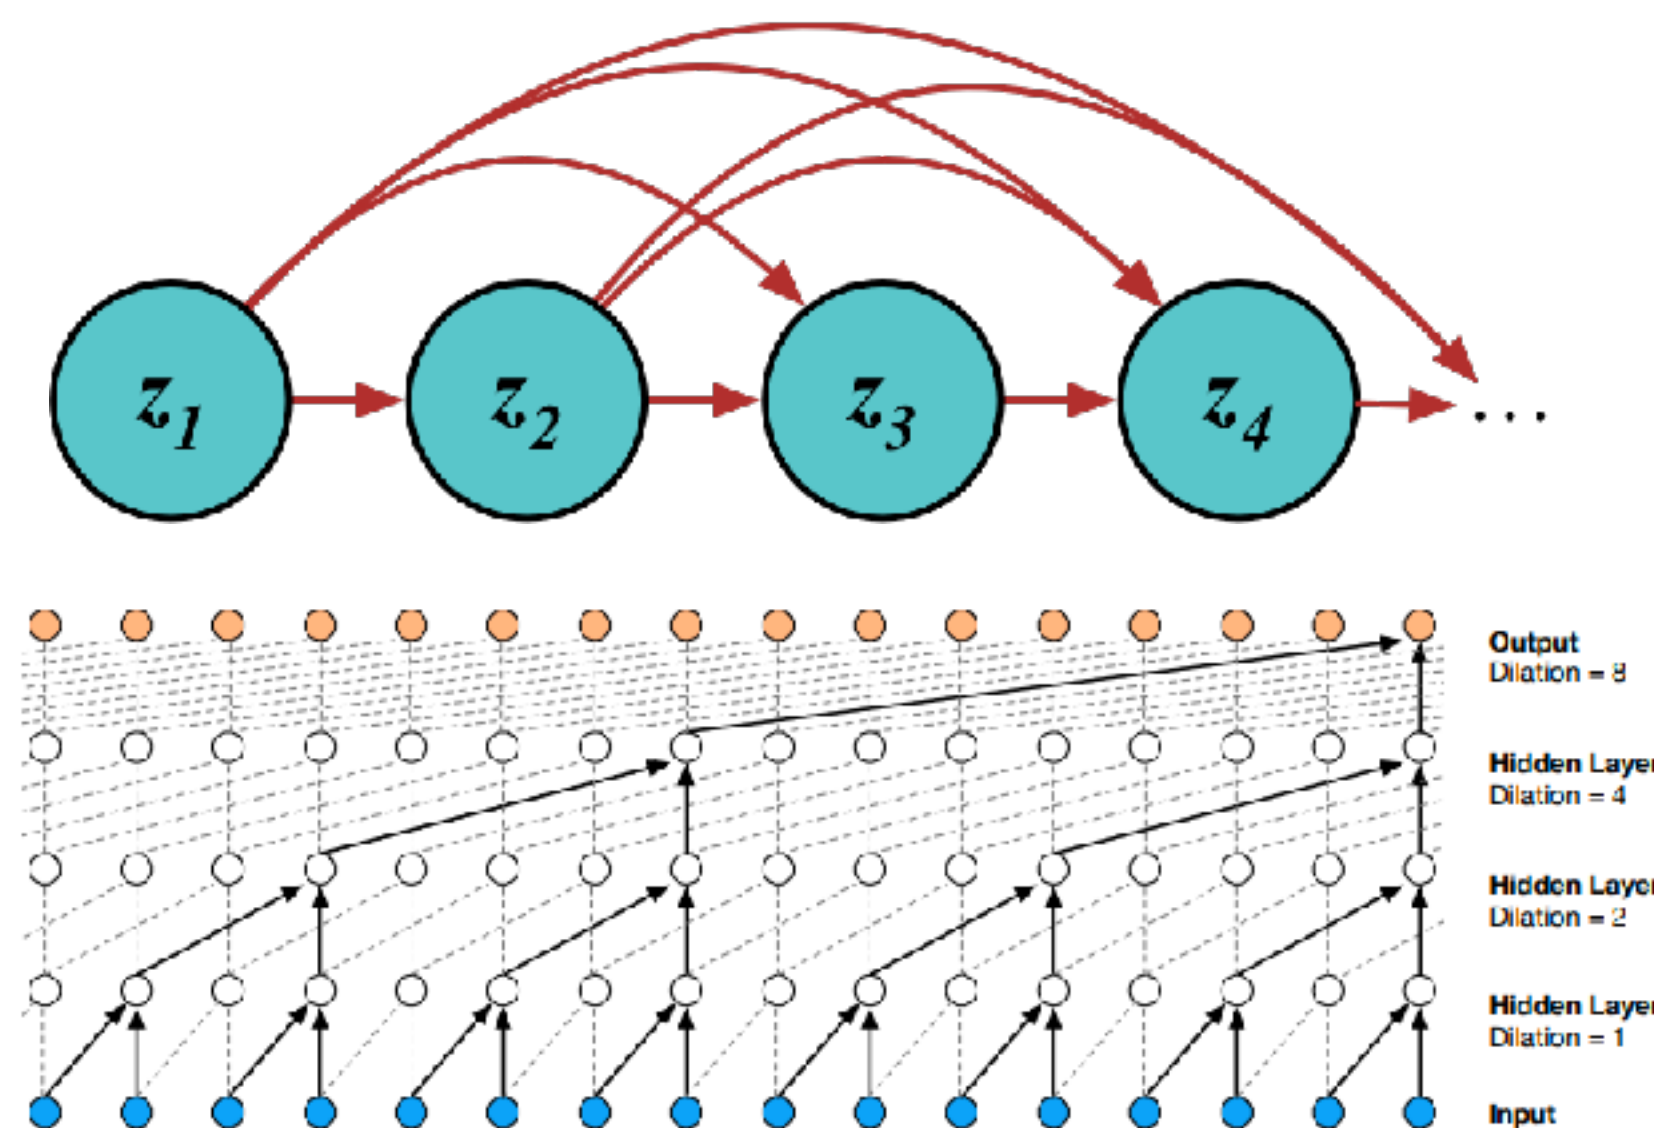
Class Probability Estimation



f -Divergence

$f(u) = u \log u - (u + 1) \log(u + 1)$

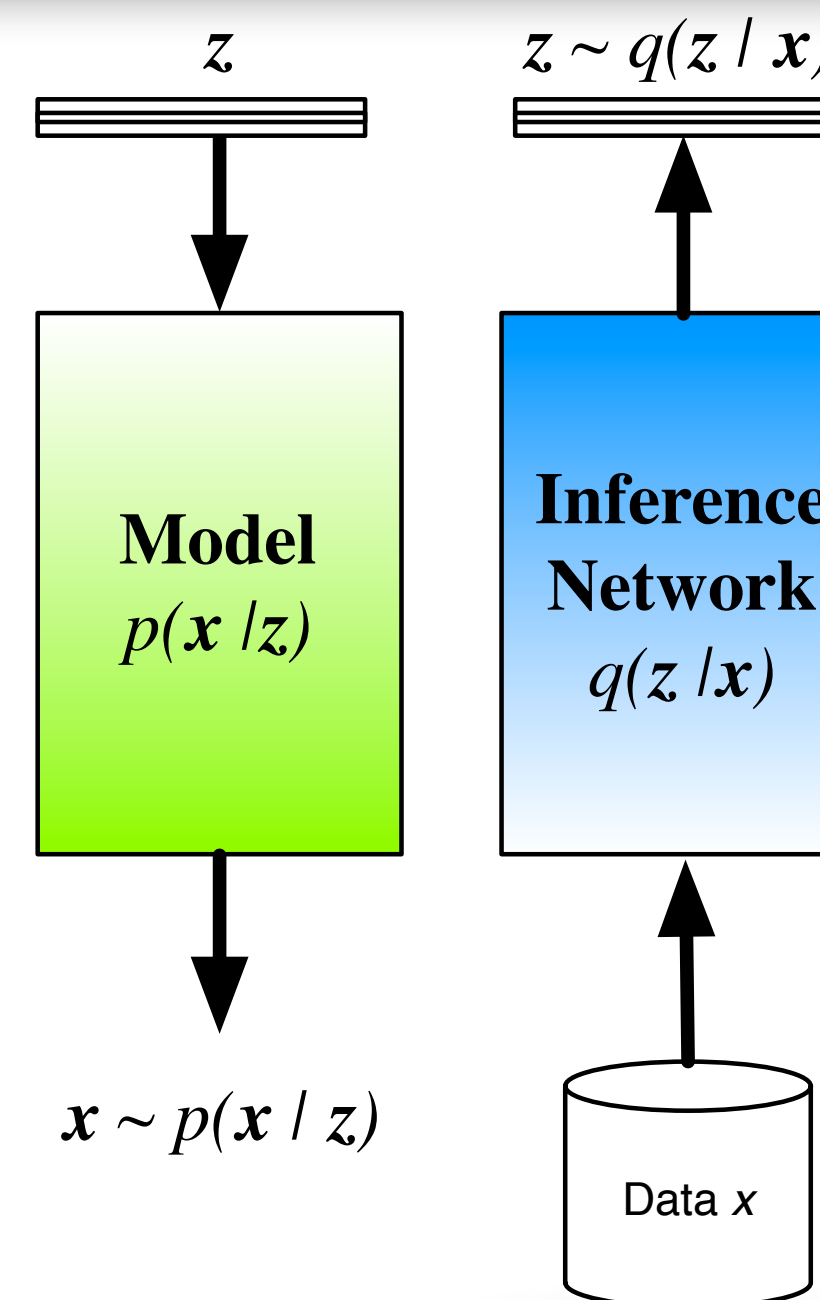
Algorithms for Generative Models



Fully-observed auto-regressive models

$$p(\mathbf{x})$$

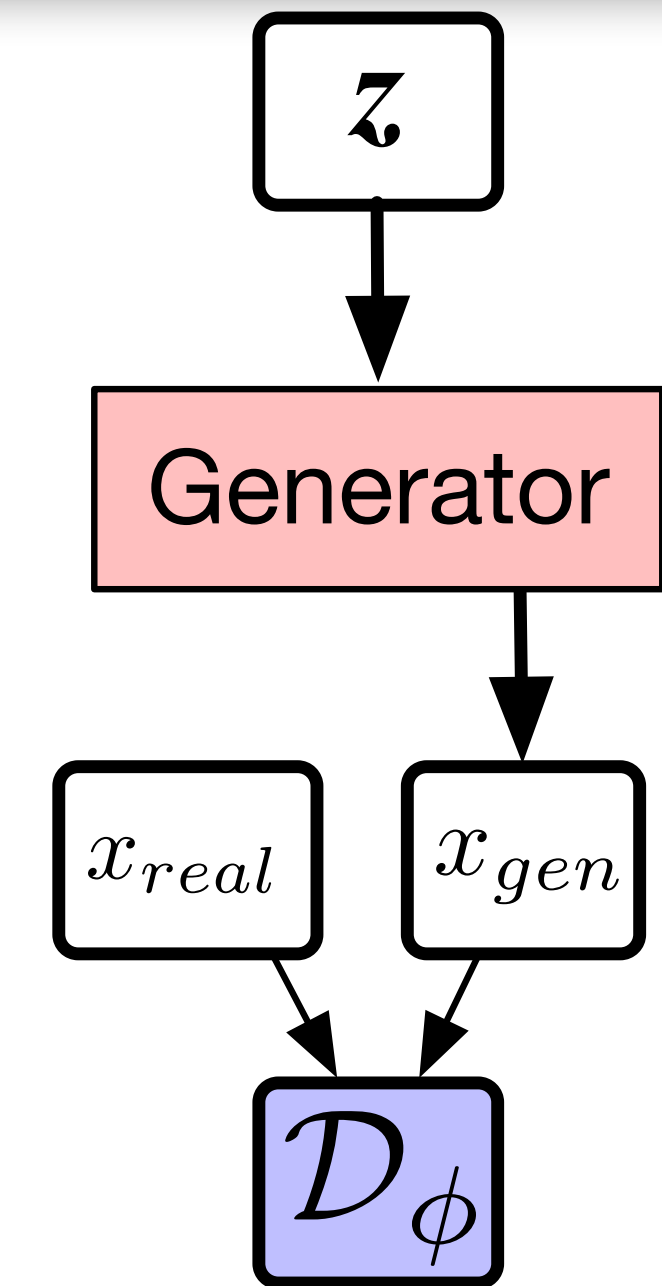
**PixelCNN and
Wavenet**



Prescribed latent variable models and variational inference

$$\tilde{p}(\mathbf{x}) \leq p(\mathbf{x})$$

**Variational
Autoencoders**



Implicit latent variable models and estimation-by-comparison

$$r(\mathbf{x}) = \frac{p^*(\mathbf{x})}{p(\mathbf{x})}$$

**Generative
Adversarial Networks**

Stochastic Optimisation

Common gradient problem

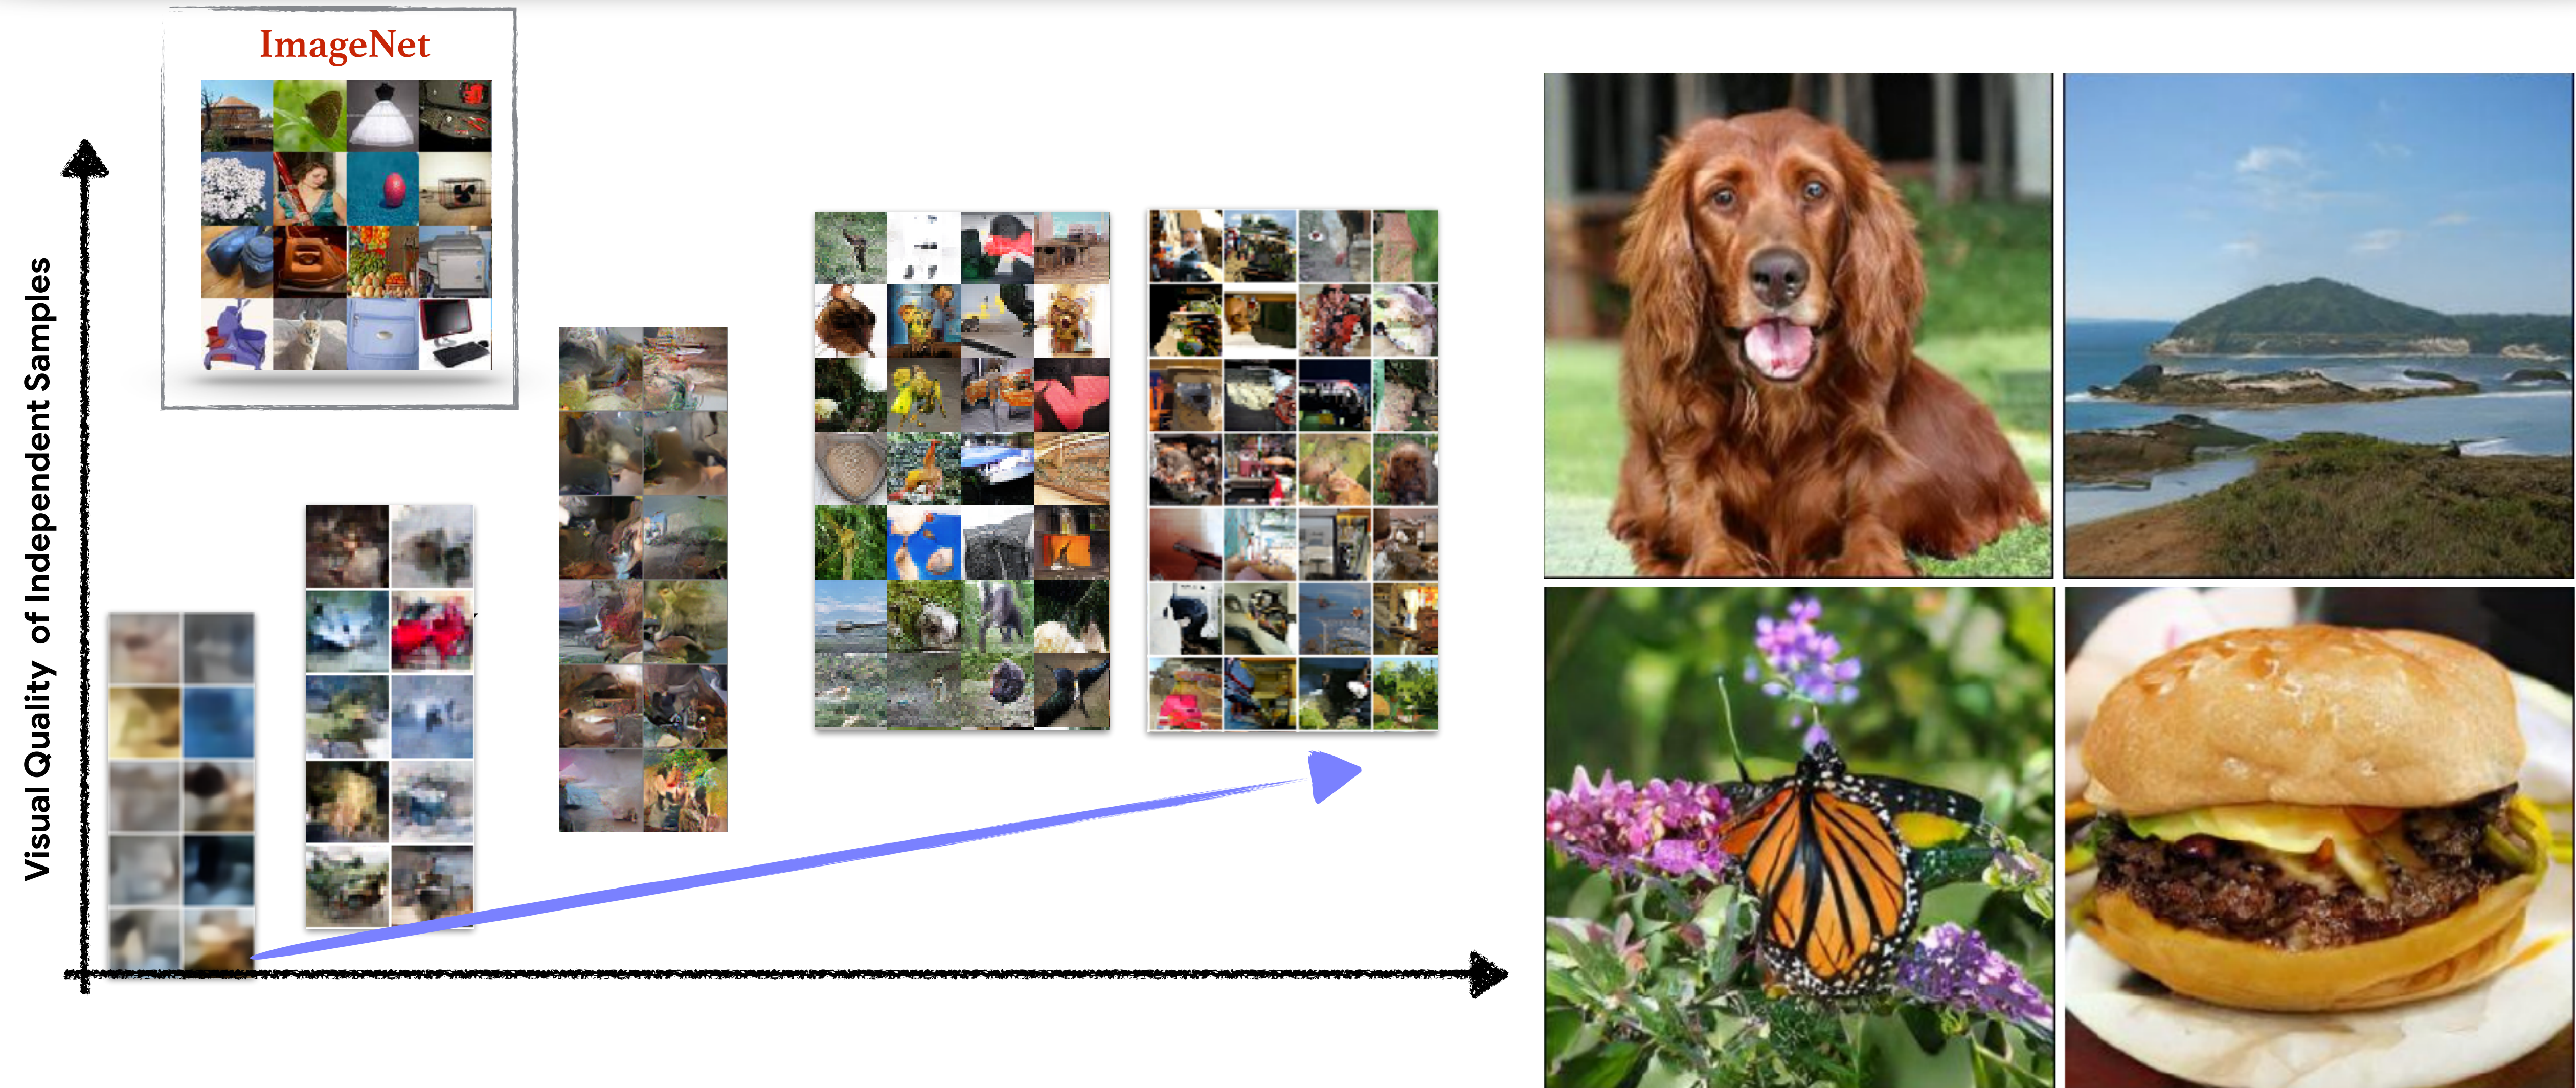
$$\nabla_{\phi} \mathbb{E}_{q_{\phi}(\mathbf{z})} [f_{\theta}(\mathbf{z})] = \nabla \int q_{\phi}(\mathbf{z}) f_{\theta}(\mathbf{z}) d\mathbf{z}$$

1. **Pathwise estimator**: Differentiate the function $\mathbf{f}(\mathbf{z})$
2. **Score-function estimator**: Differentiate the density $\mathbf{q}(\mathbf{z}|\mathbf{x})$

Typical problem areas

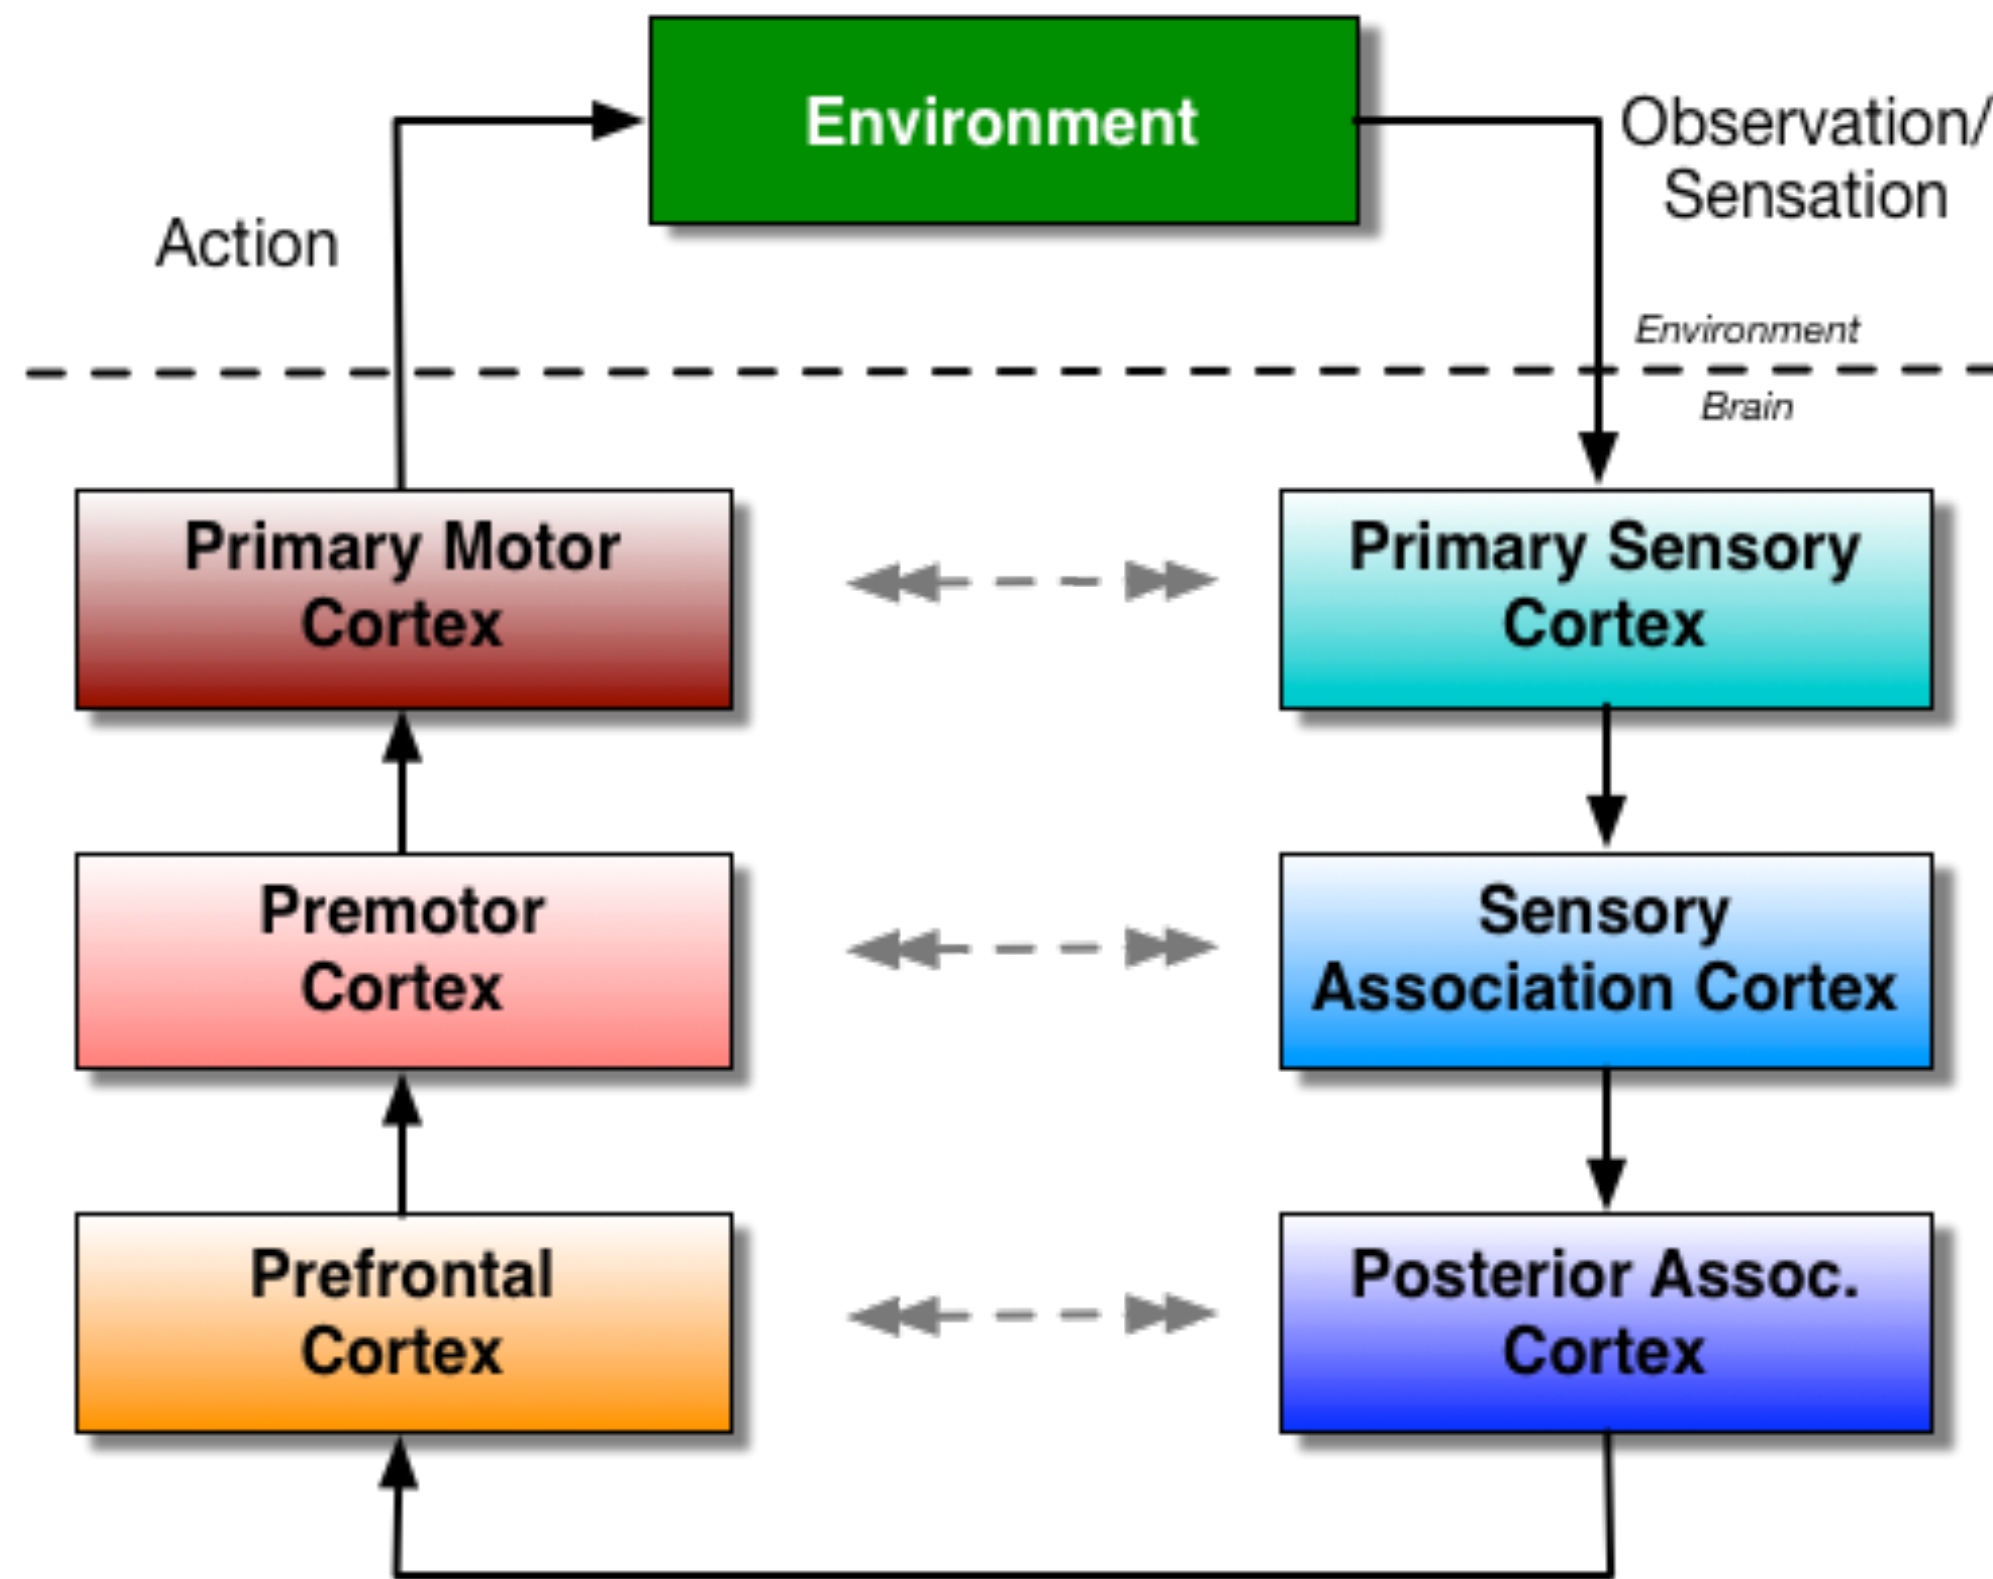
- Sensitivity analysis
- Generative models and inference
- Reinforcement learning and control
- Operations research and inventory control
- Monte Carlo simulation
- Finance and asset pricing

Progress in Generative Models

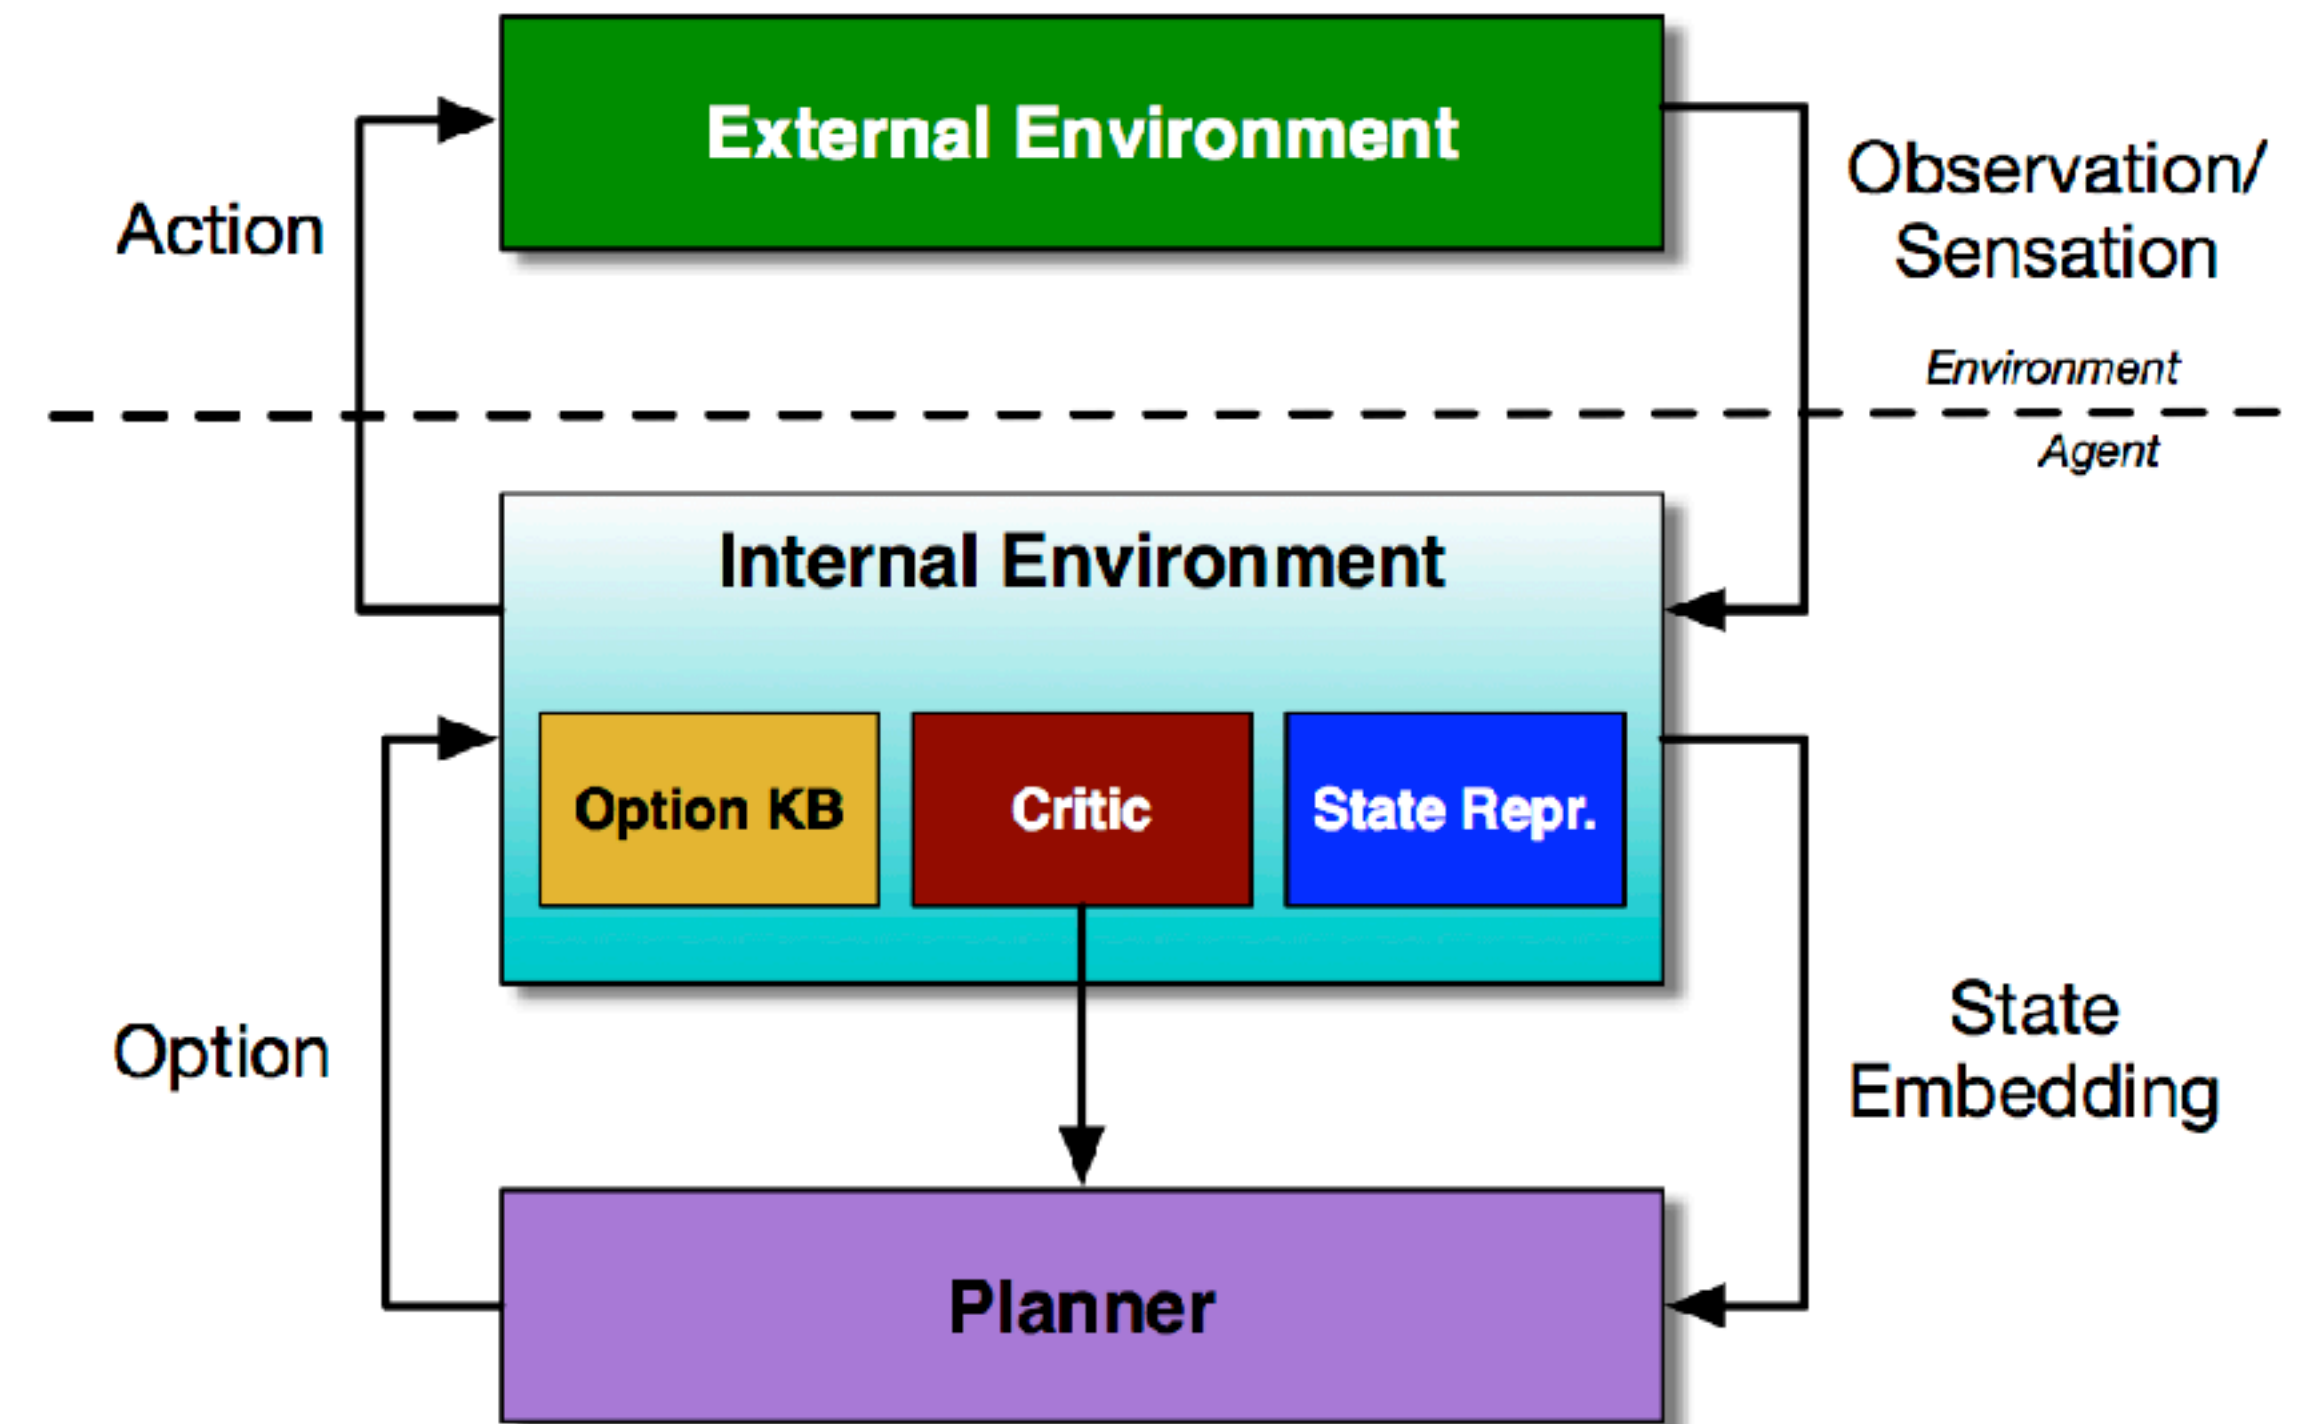


Perception-Action Loops

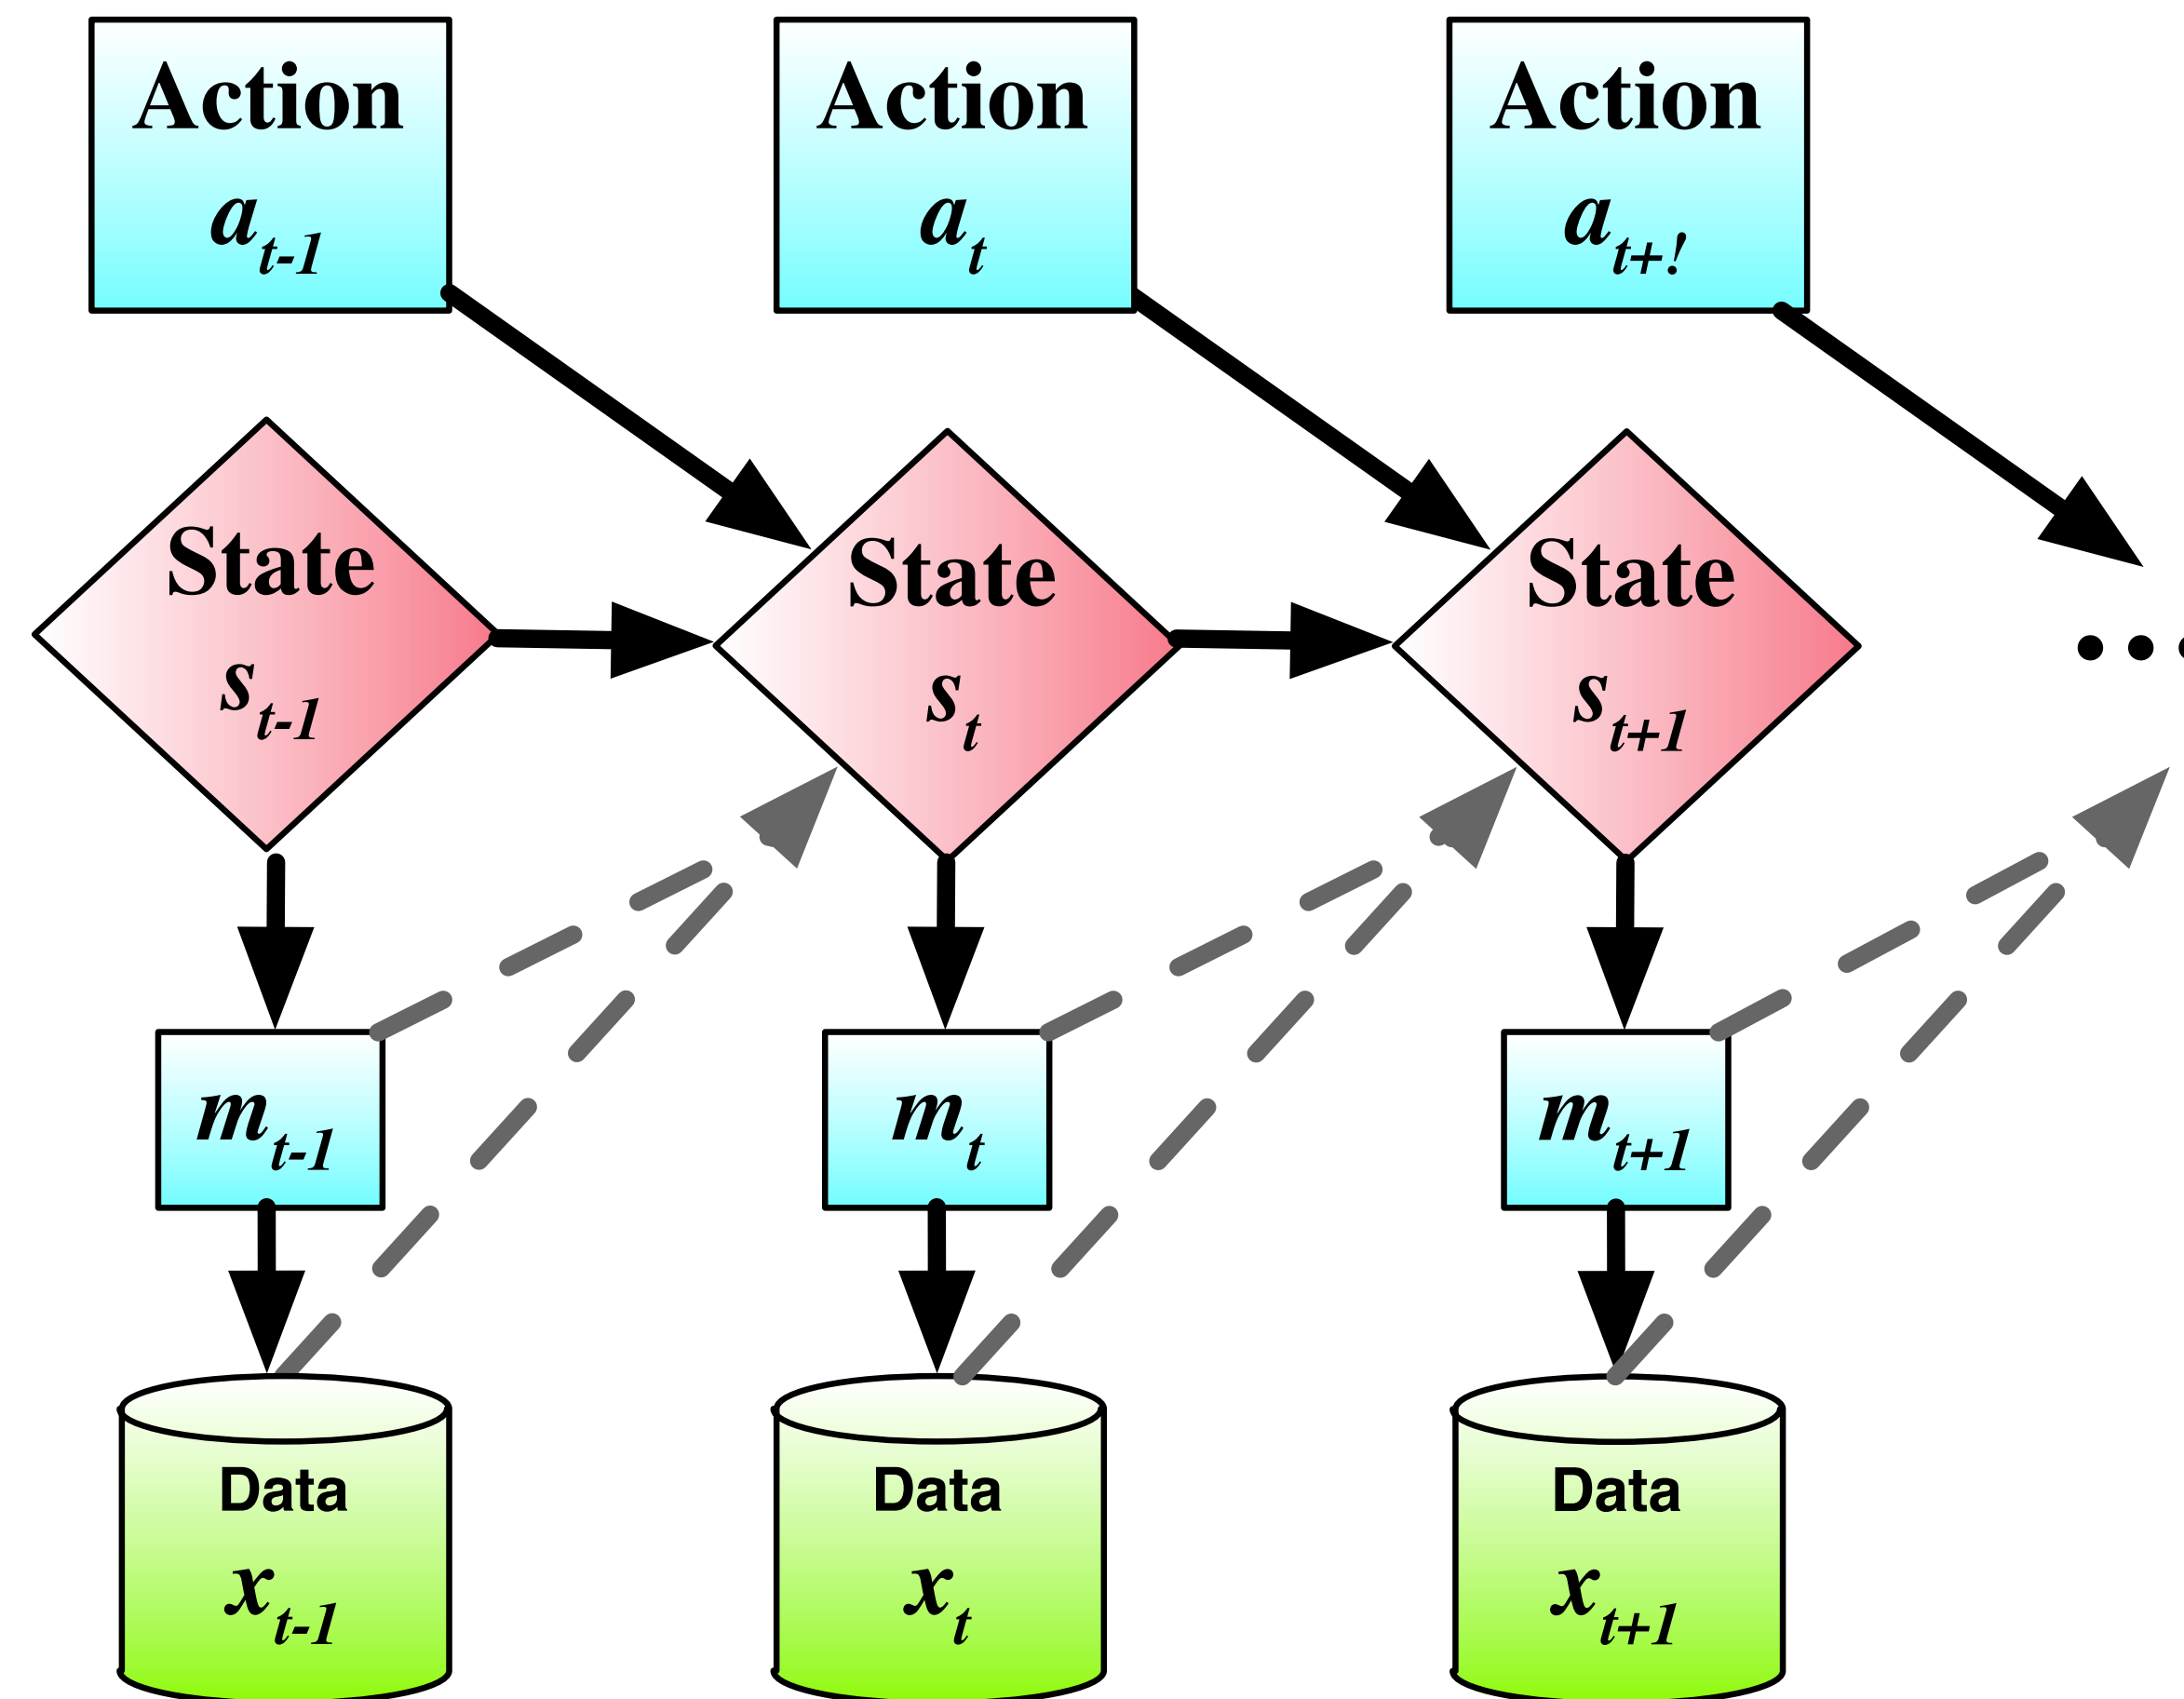
Biological
perception-action loop



Computational
perception-action loop



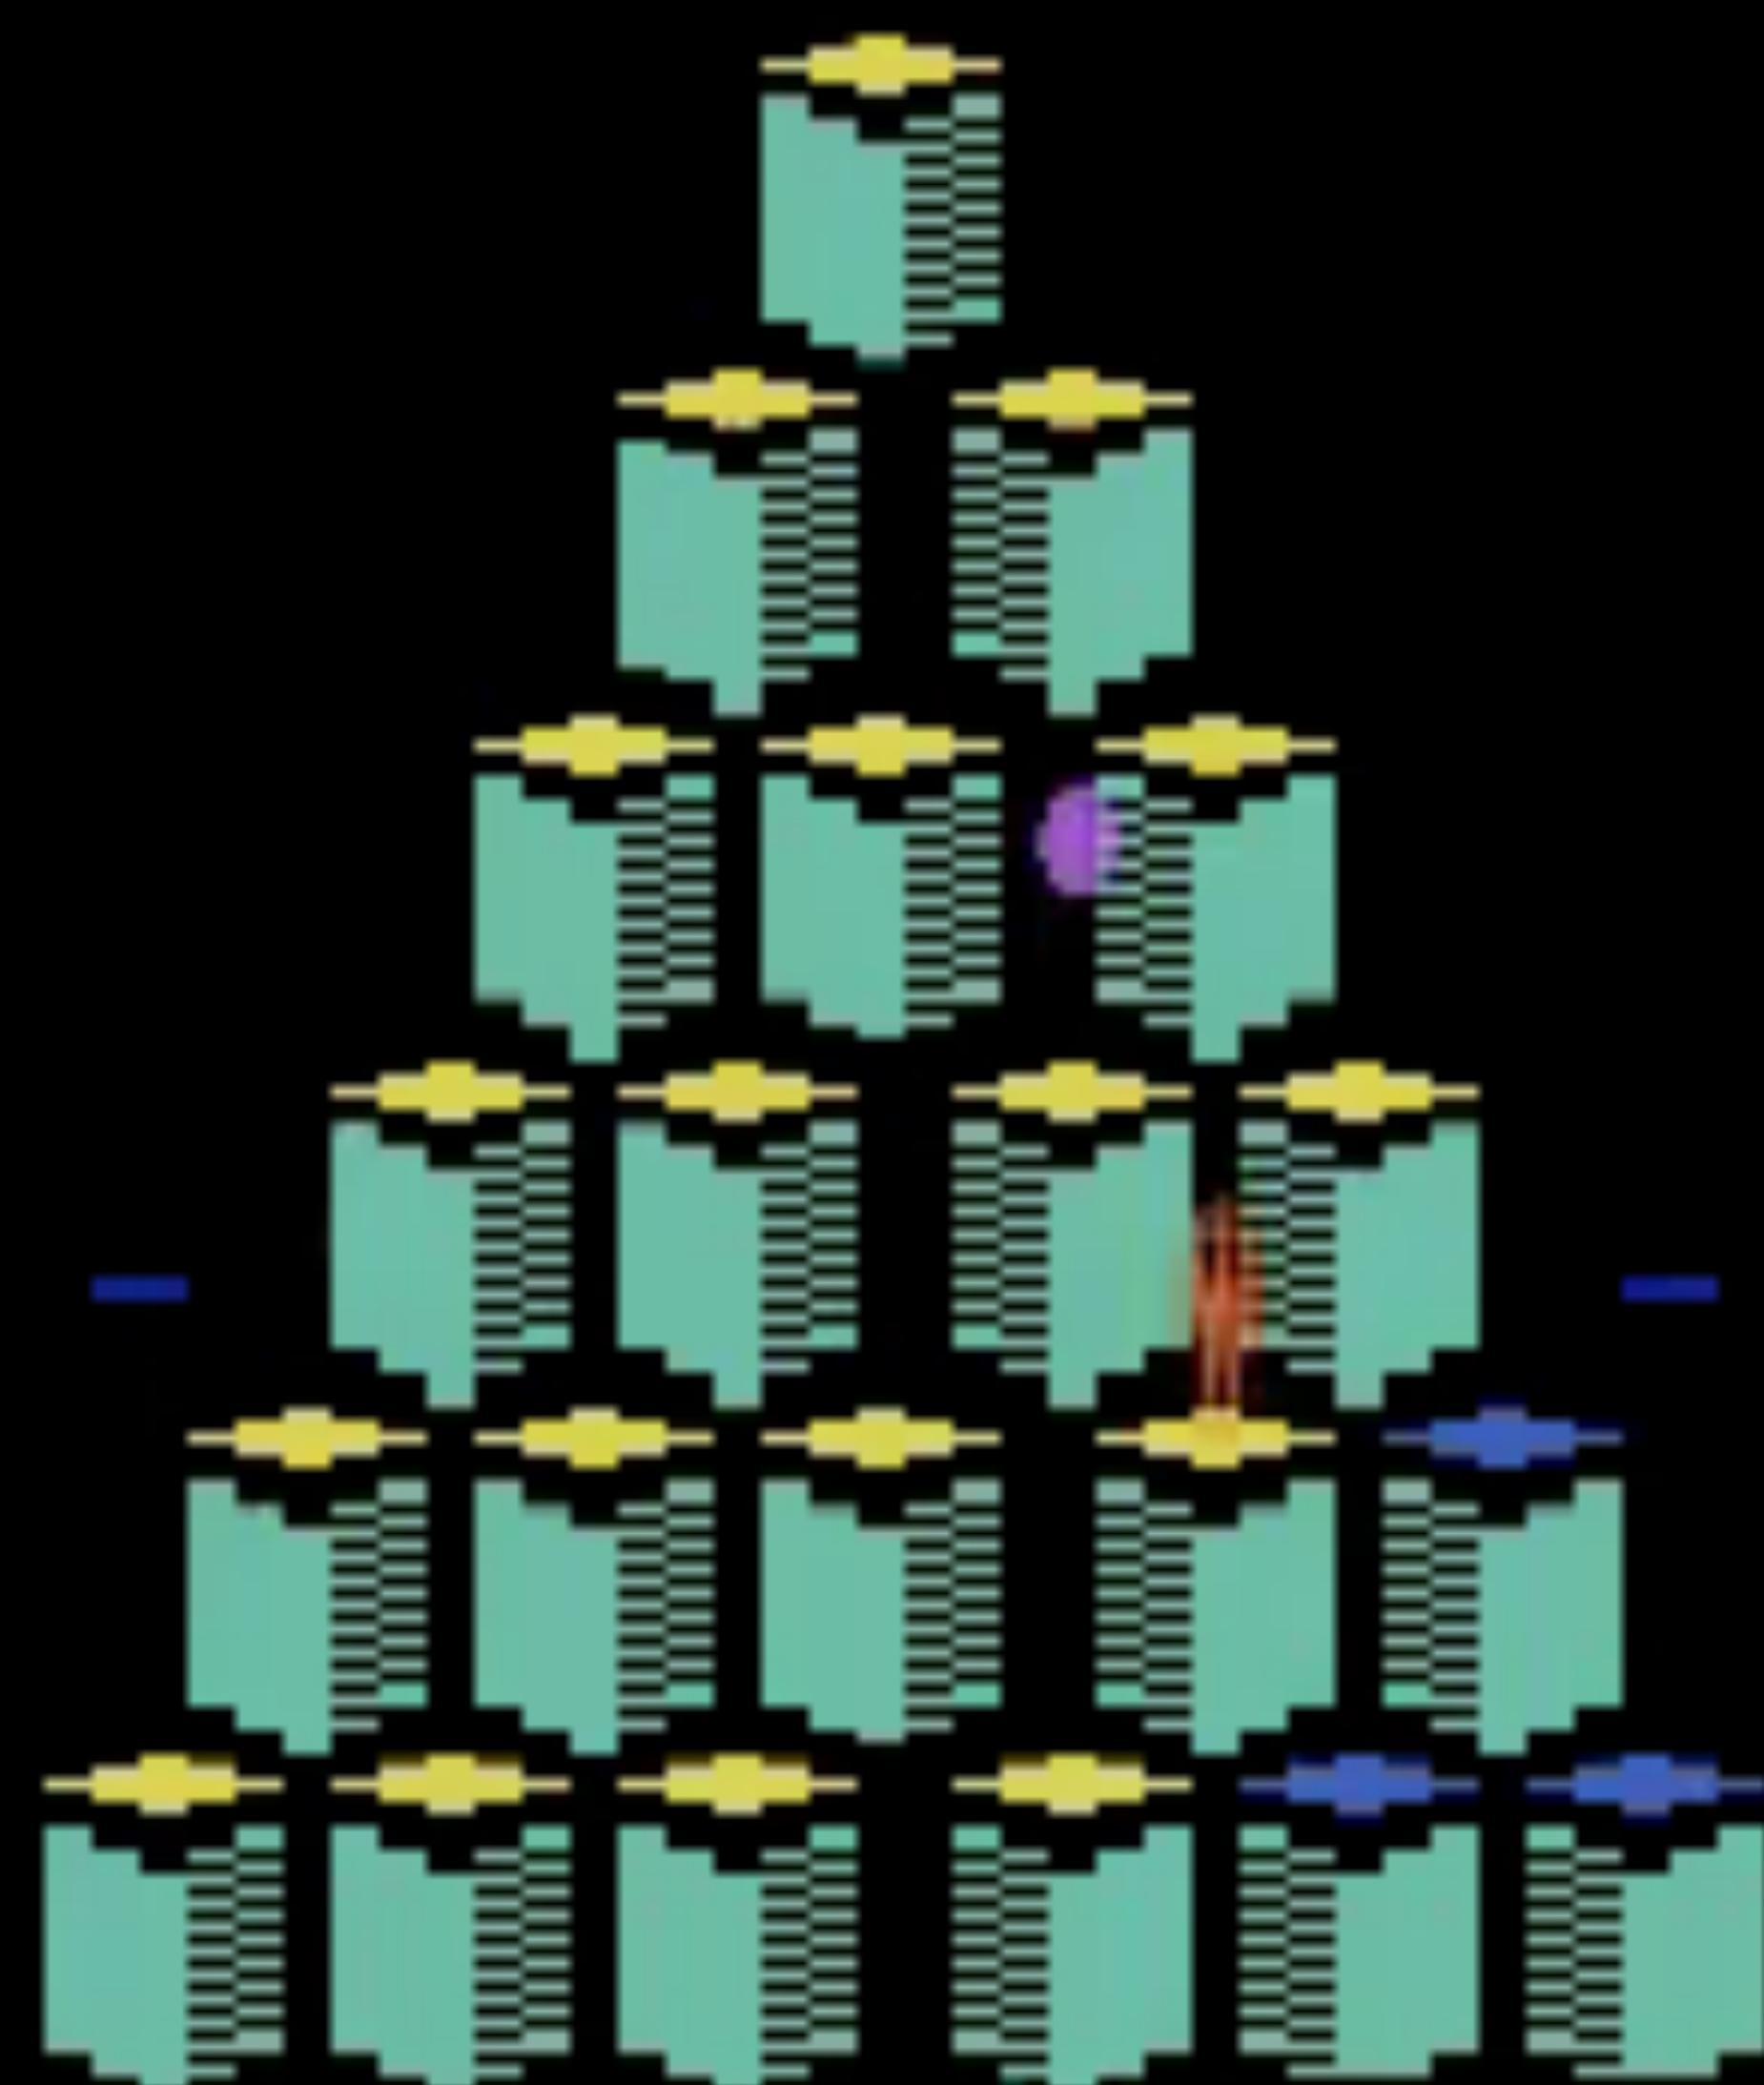
Environment Simulation



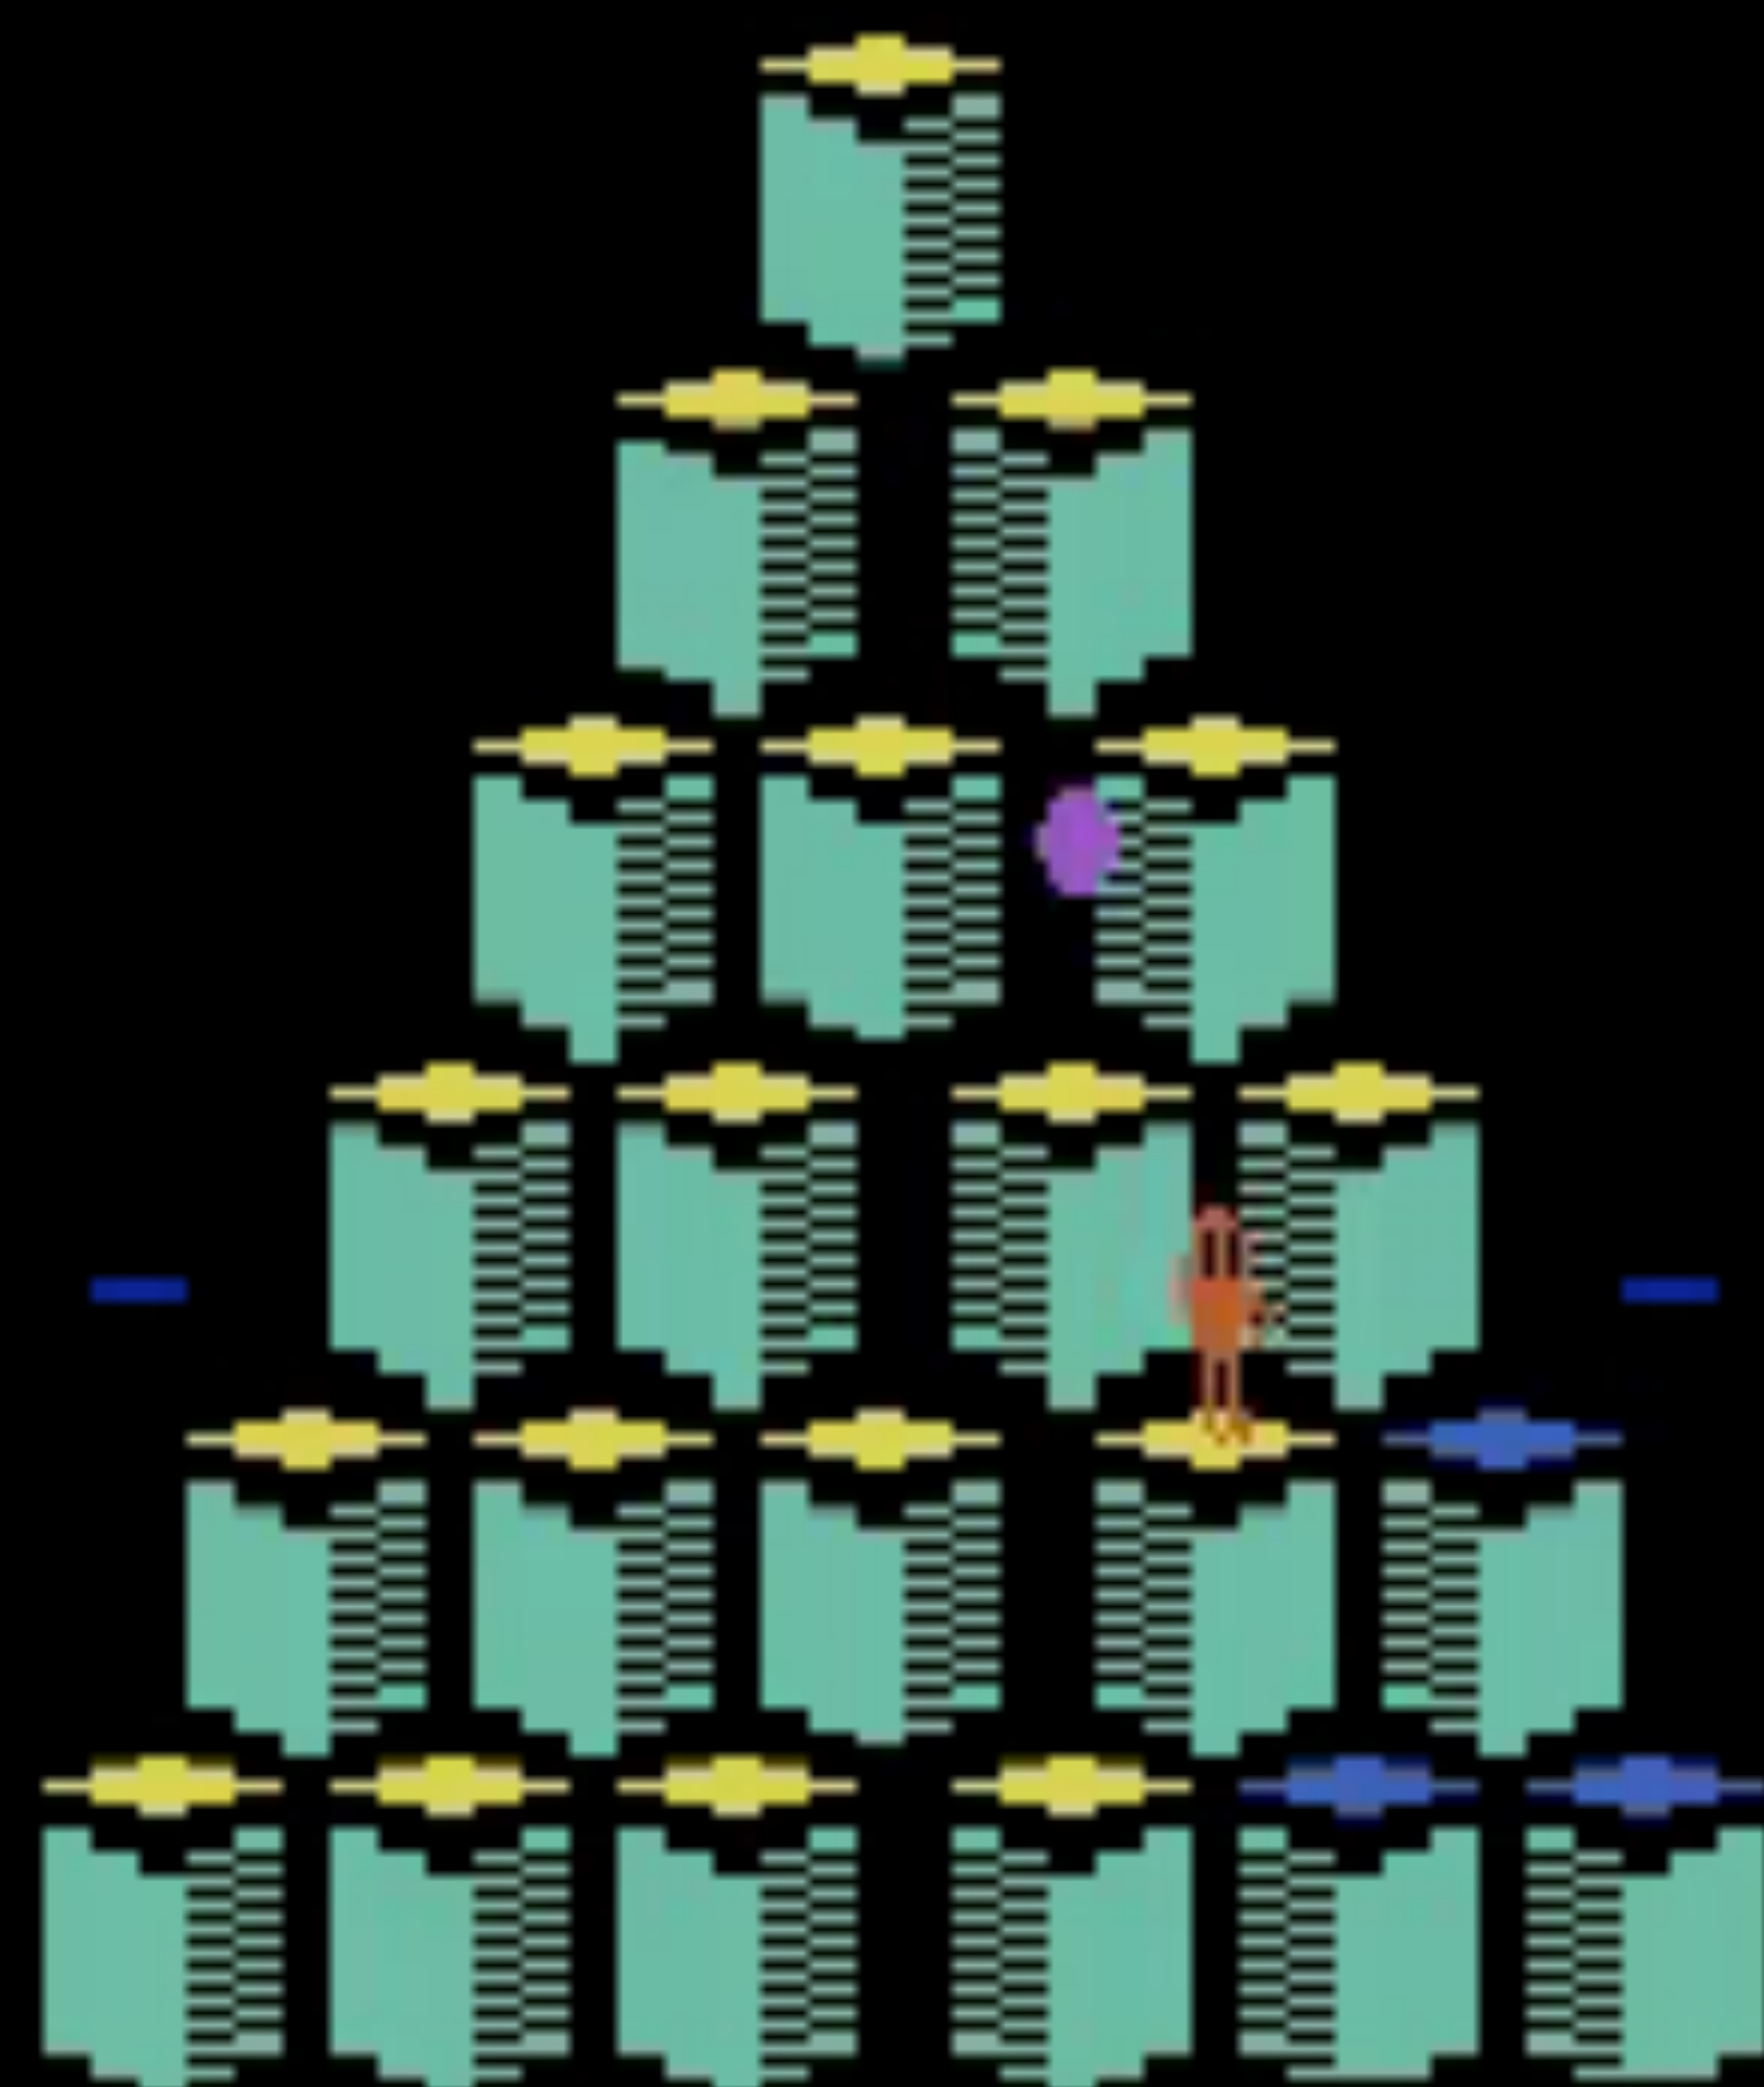
Action-conditional and latent-only transitions.

Grounded representations in actions and observations, using simulation to support grounding.

00475



00475

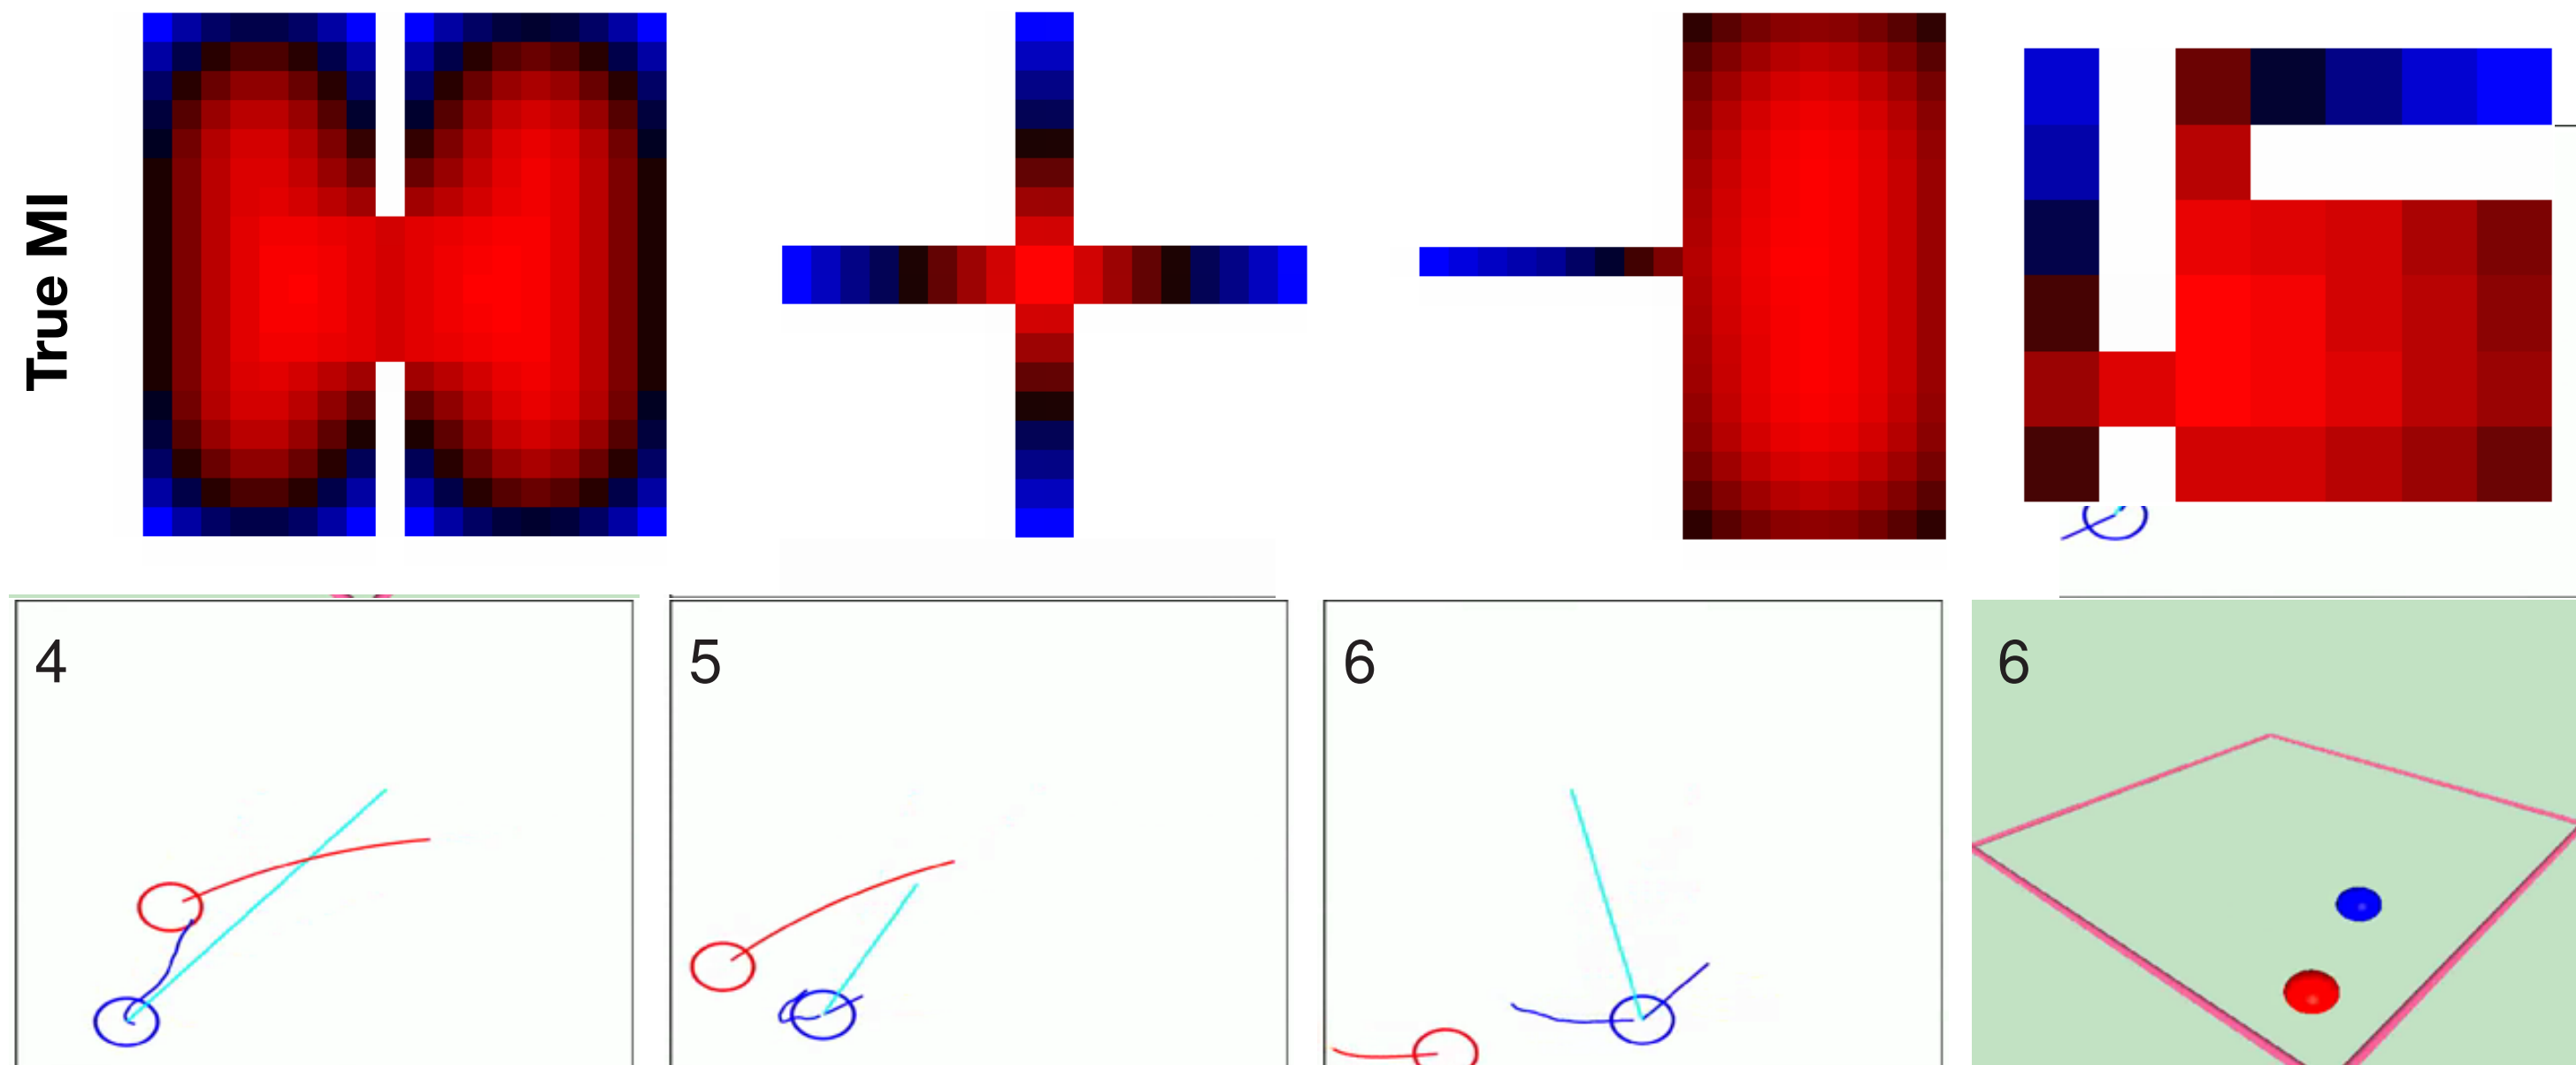


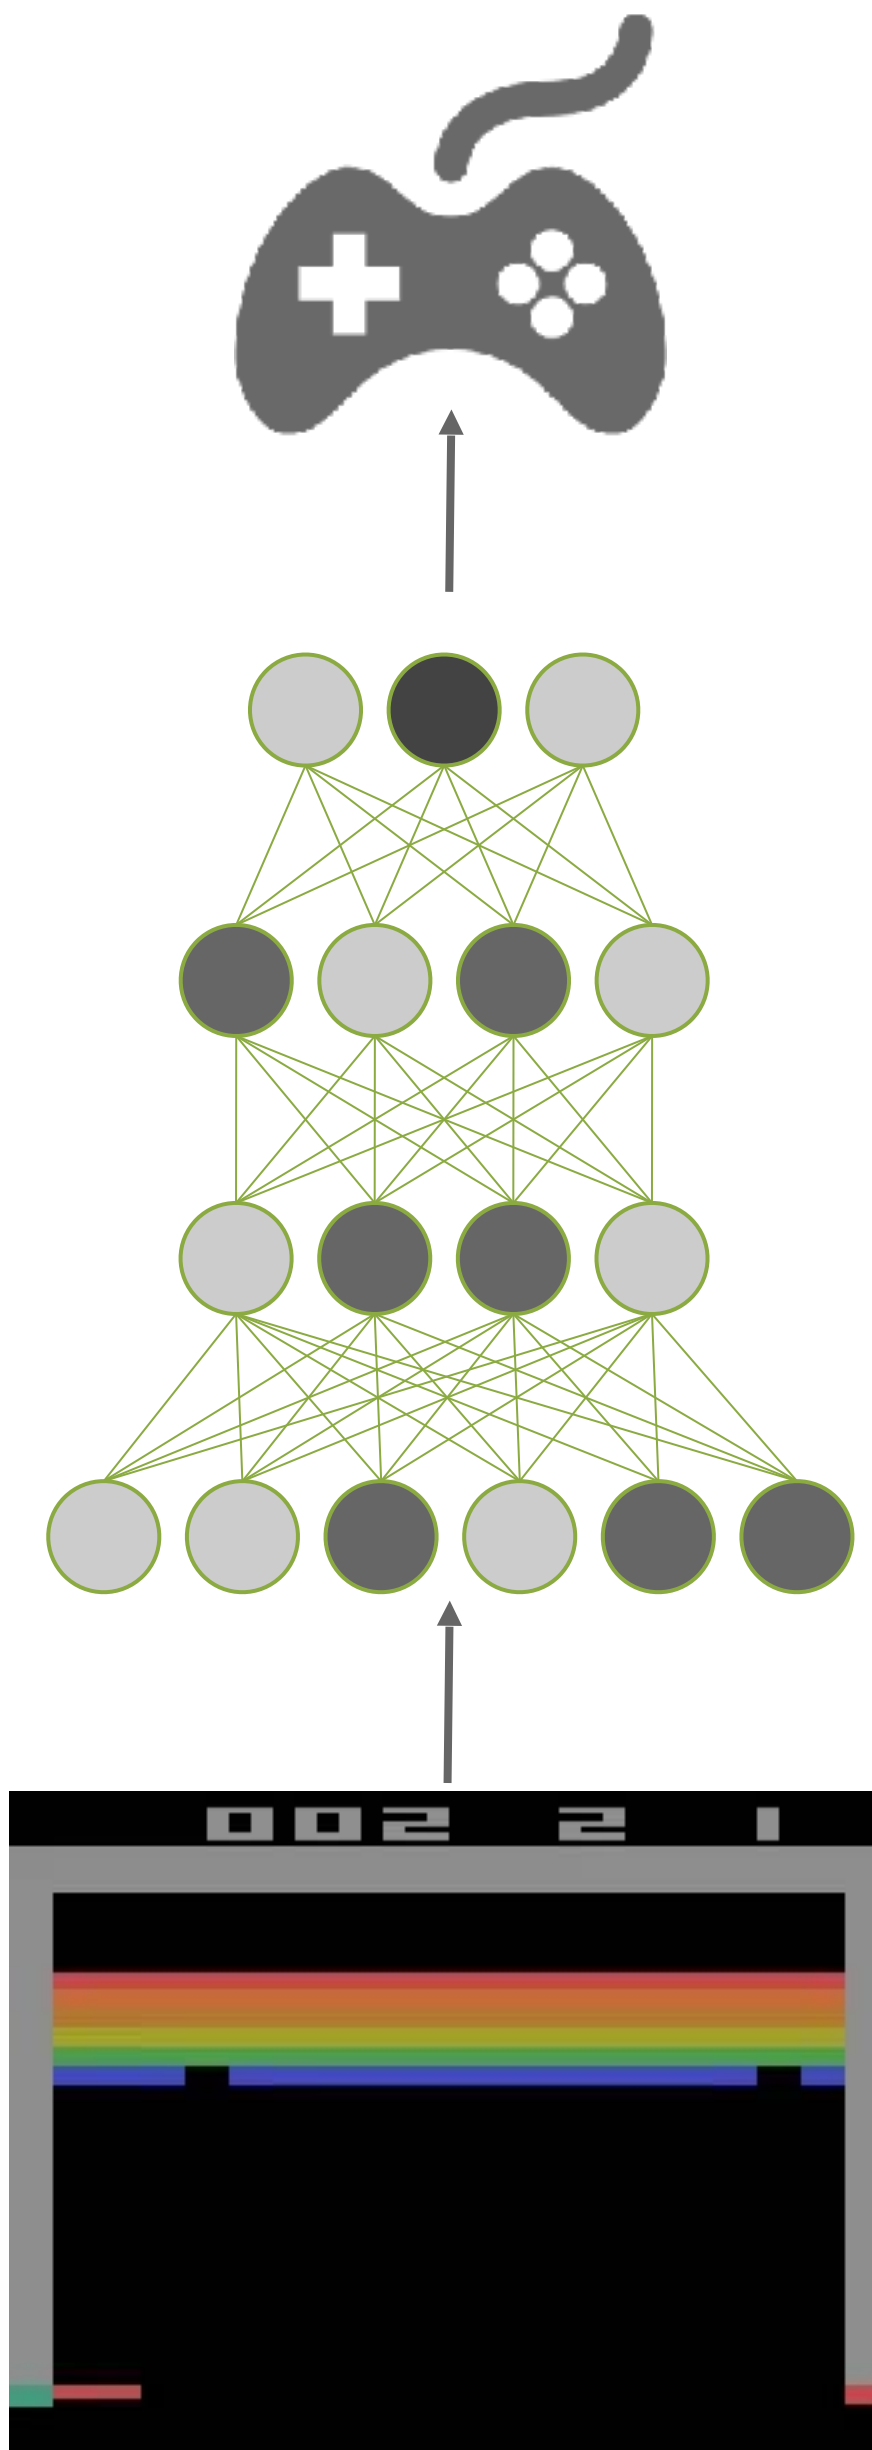
Intrinsic Motivation

Equip agents with mechanisms to produce and learn from internal rewards that can guide behaviour, when external rewards are absent.

$$\mathcal{E}(\mathbf{s}) = \max_{\omega} \mathcal{I}^{\omega}(\mathbf{a}, \mathbf{s}' | \mathbf{s}) \quad \max_{\omega} \mathbb{E}_{p(\mathbf{s}' | \mathbf{a}, \mathbf{s}) \omega(\mathbf{a} | \mathbf{s})} \left[\log \frac{p(\mathbf{s}', \mathbf{a} | \mathbf{s})}{p(\mathbf{s}' | \mathbf{s}) \omega(\mathbf{a} | \mathbf{s})} \right]$$

Escaping a Predator





AlphaZero

Generalising AlphaGo to any 2-player game

Fully general; No opening book; No endgame database; No heuristics; Starts from random

All learned **without any reference to past human games**

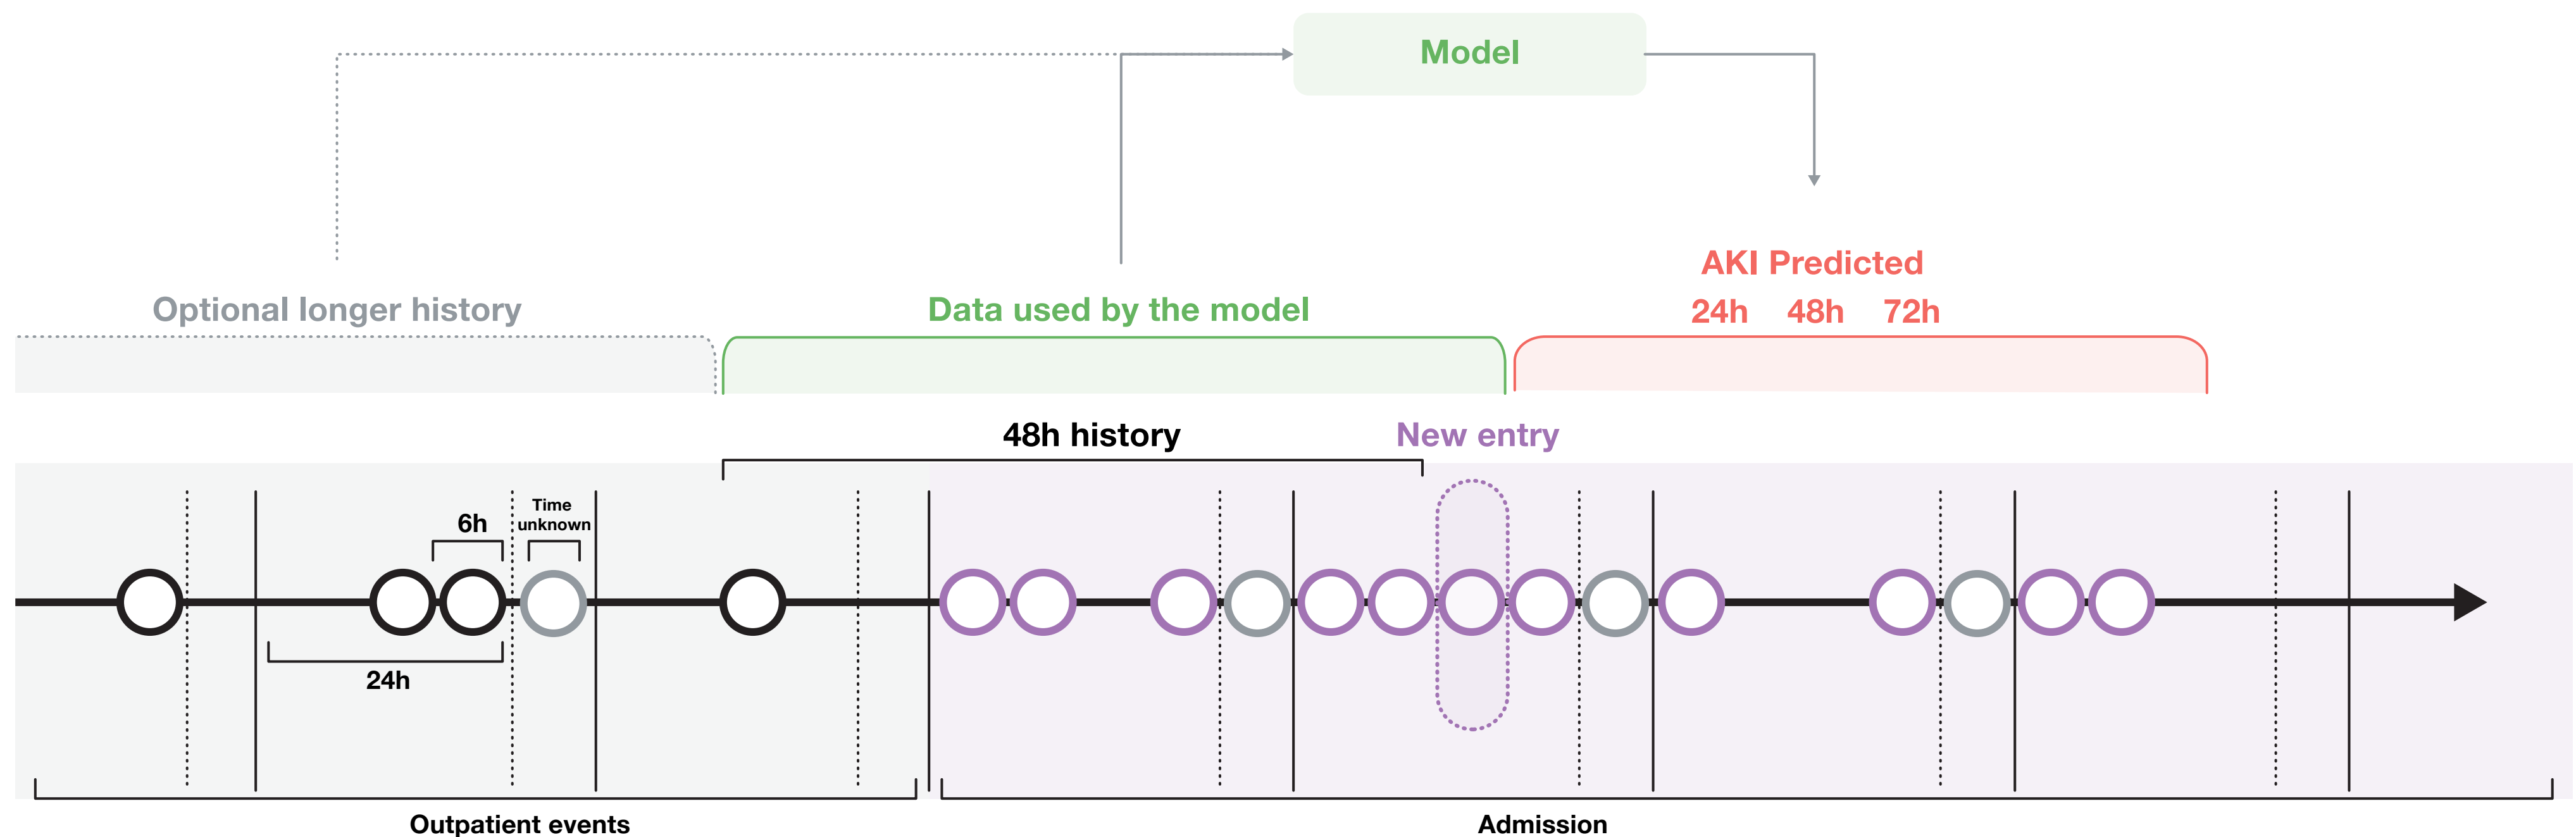
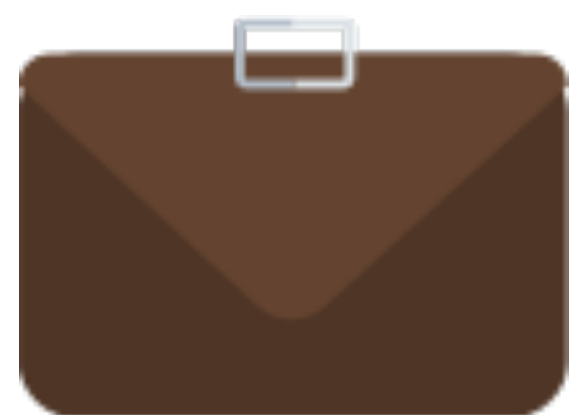


Applications in Healthcare

1 Better clinical outcomes

2 Enhance patient and clinician experience

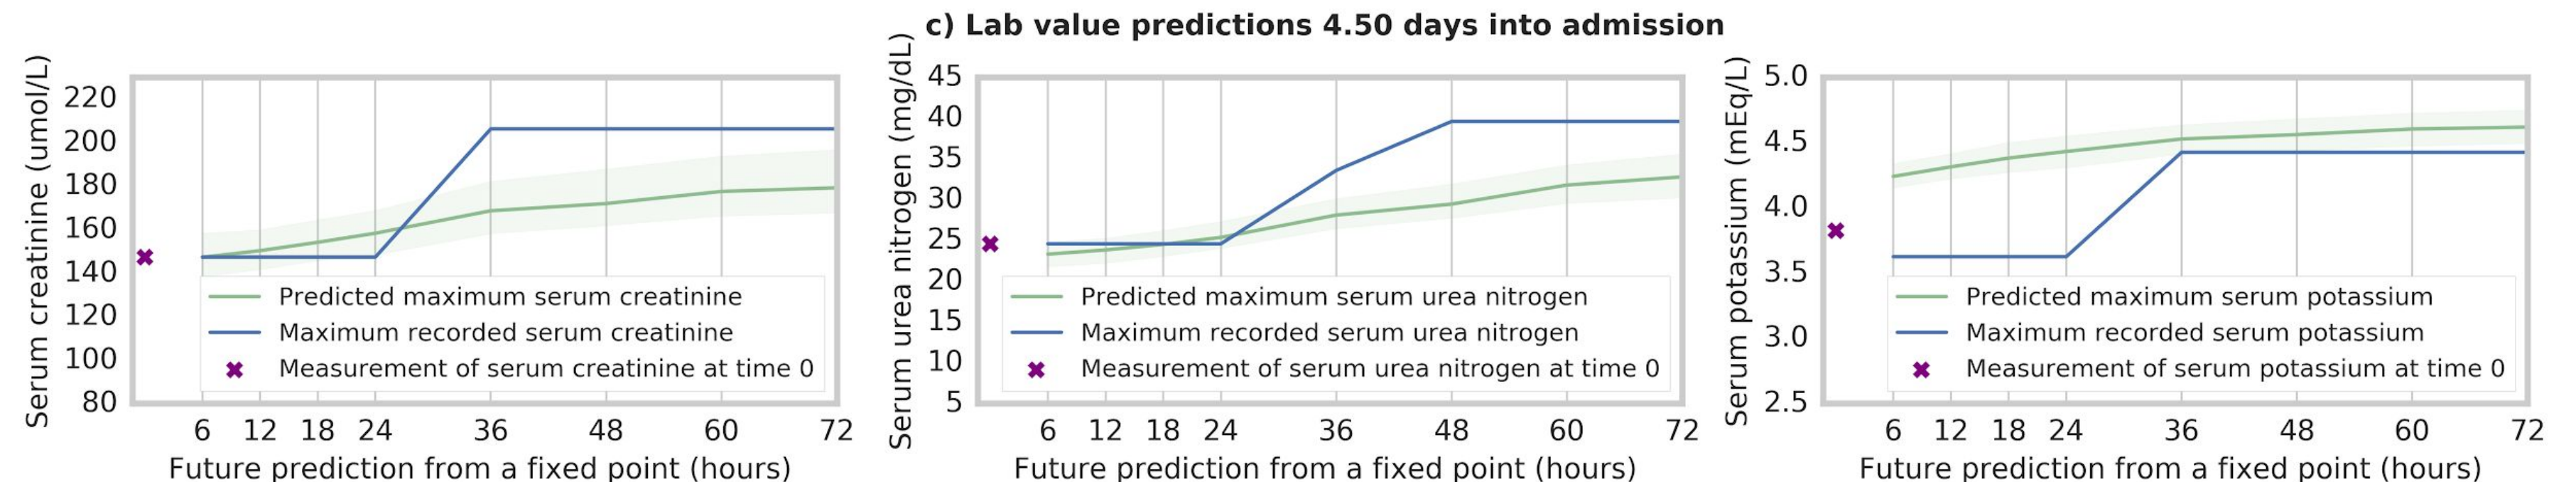
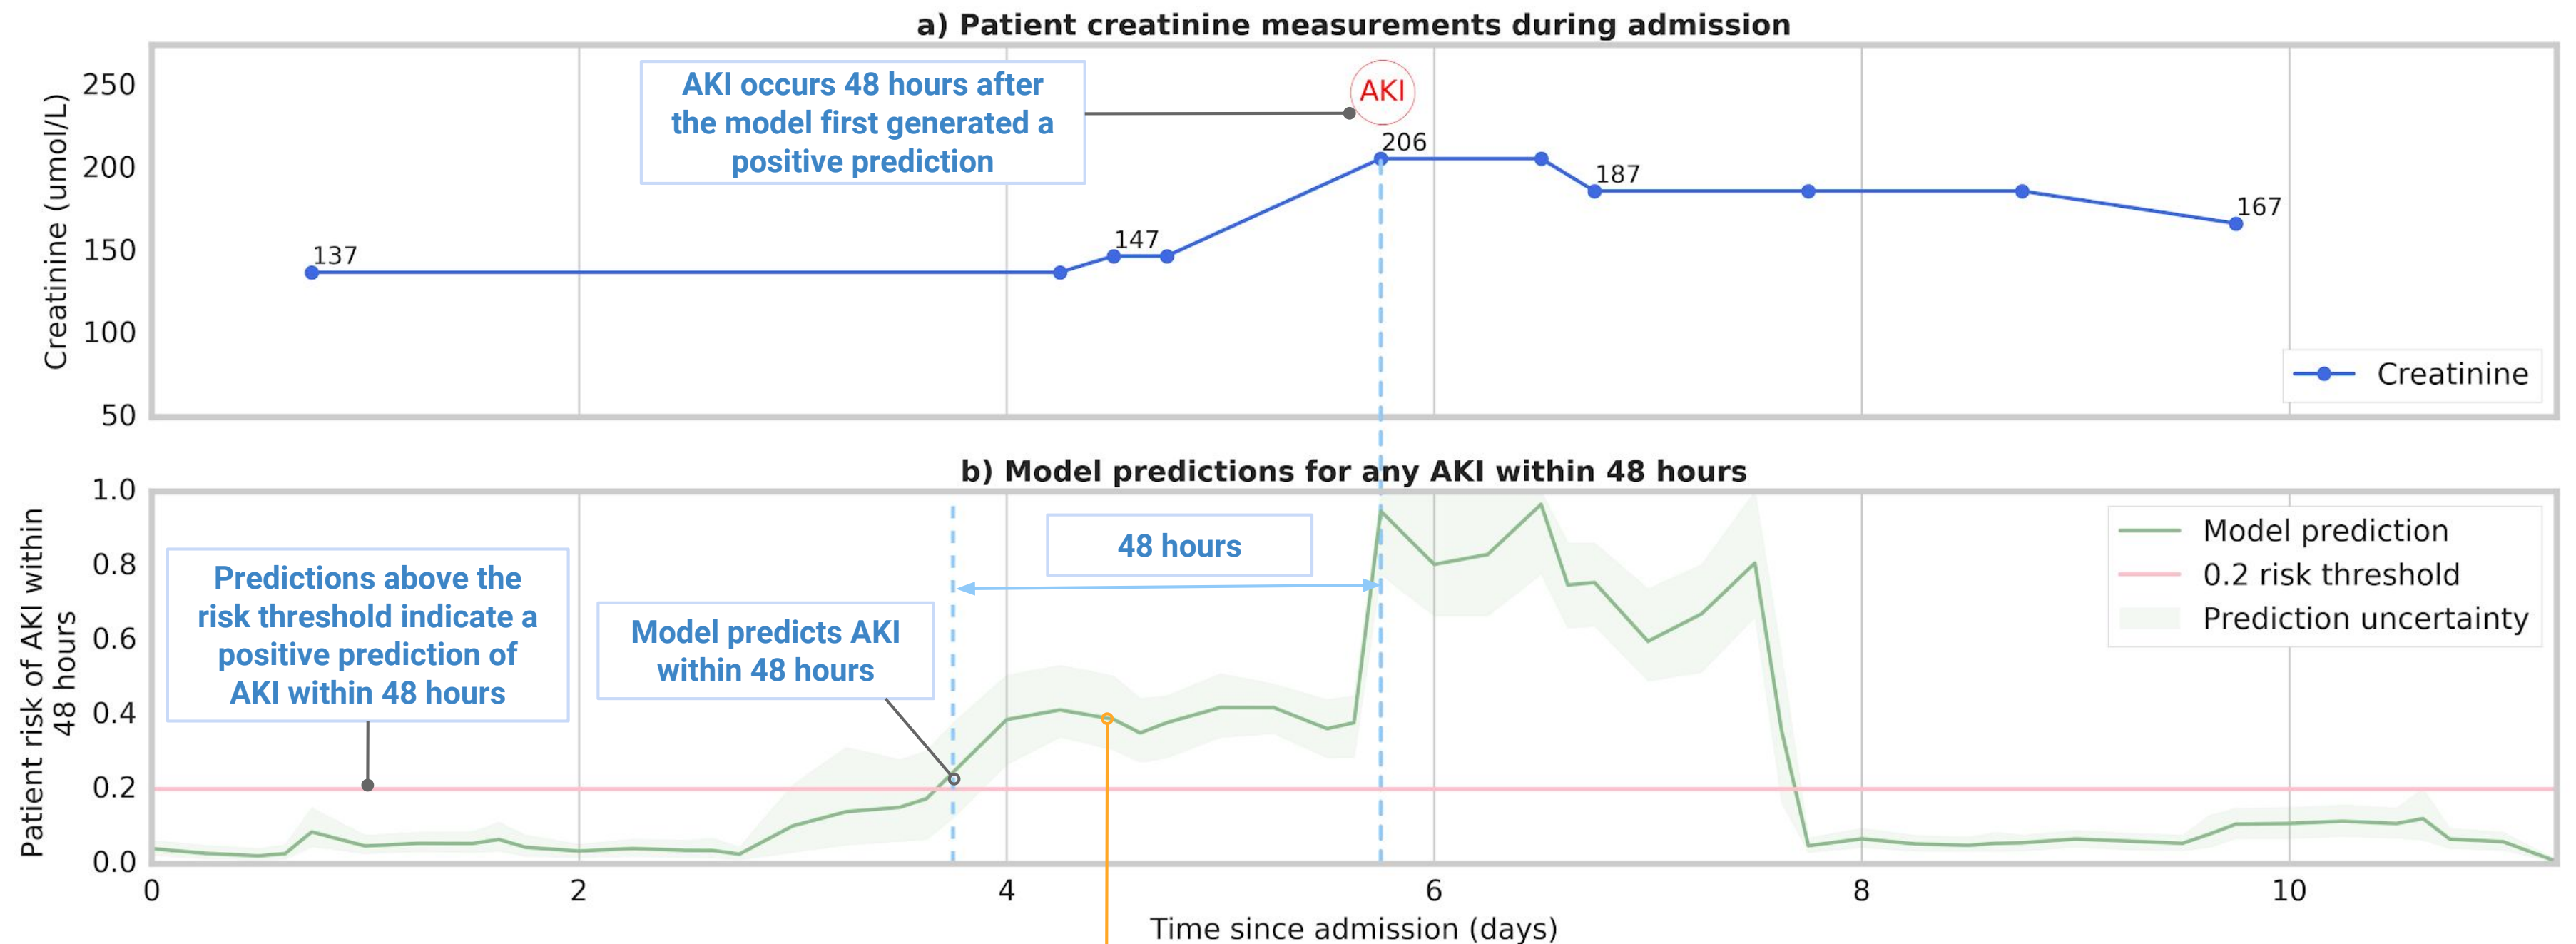
3 Reduce costs



Predicting Organ Failure

Make predictions of AKI up to 48hr ahead.
Provide a window for meaningful action.

For the most severe cases, can detect up to 90% of cases.



Critical Practice for ML



Consider the uses of our models.

What are the dual uses of generative models. How do we think critically about these uses, educate, regulate, co-design these tools.

[illegible]

Neutrality and Universality

Neutrality Traps

- **The Portability Trap:** Failure to understand how repurposing algorithmic solutions designed for one social context may be inaccurate / do harm when applied to a different context.
- **The Formalism Trap:** Failure to account for the full meaning of social concepts such as fairness, which be resolved through mathematical formalisms.
- **The Ripple Effect Trap:** Failure to understand how the insertion of technology into an existing social system changes the behaviours and embedded values of the pre-existing system .
- **The Solutionism Trap:** Failure to recognise the possibility that the best solution to a problem may not involve technology.

Universality

‘A mono-cultural view of ethics conceives itself as the only valid one. In order to avoid this kind of ethical chauvinism and colonialism it is necessary that transcultural ethics arise from an intercultural dialogue instead of thinking of itself as universal without noticing its own cultural bias.’ Capurro, 2004



Part II: AI for Environmental Risk

A faint, light-yellow globe is centered in the background, showing the continents of Africa and Europe.

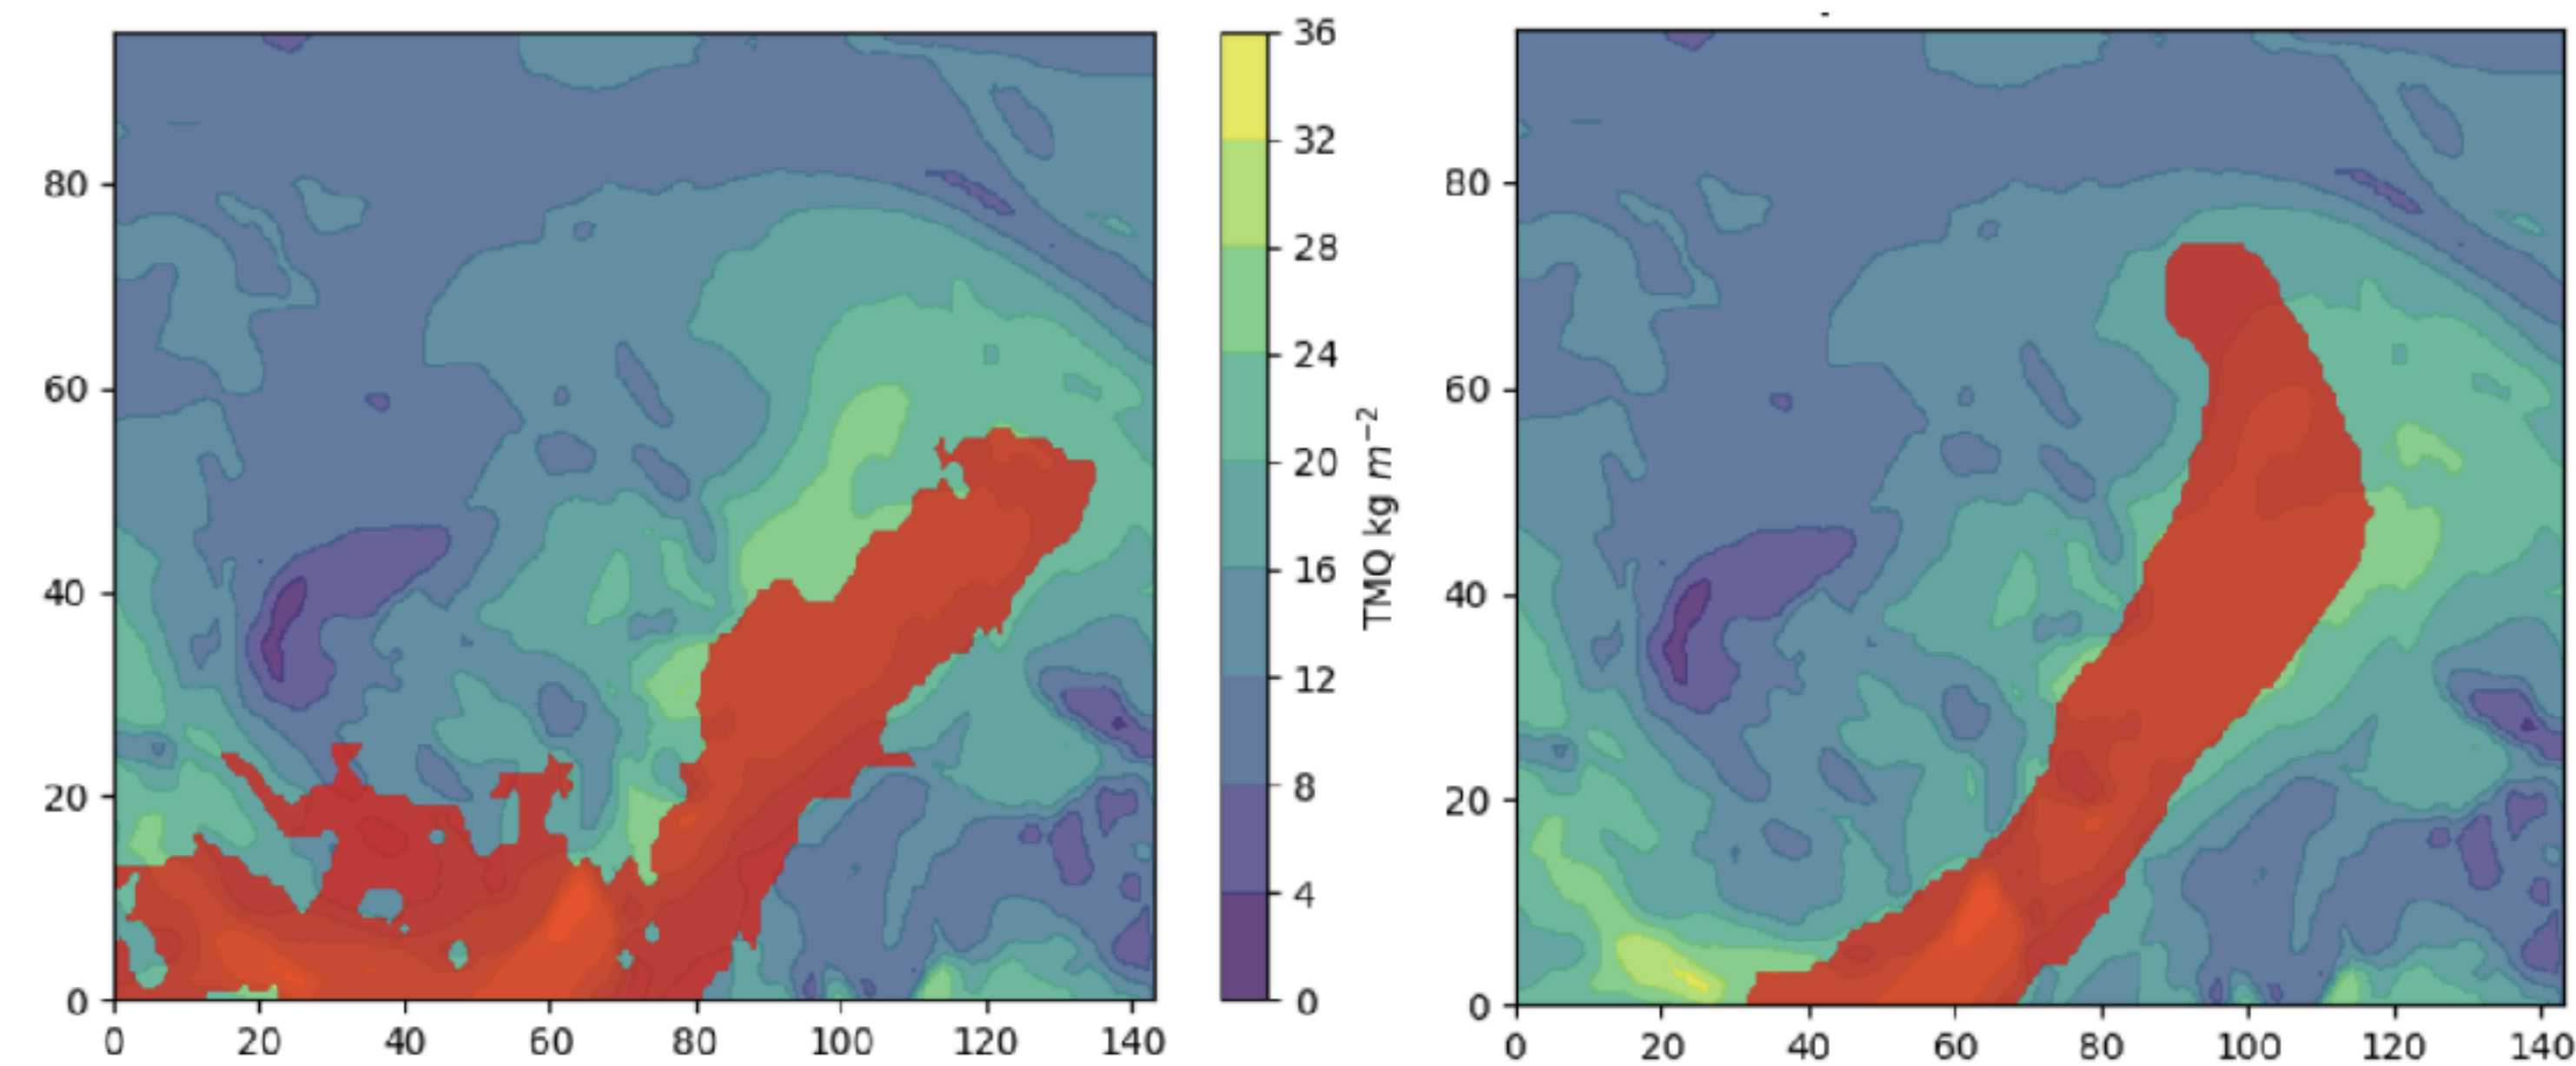
Shakir Mohamed

Research Scientist, DeepMind

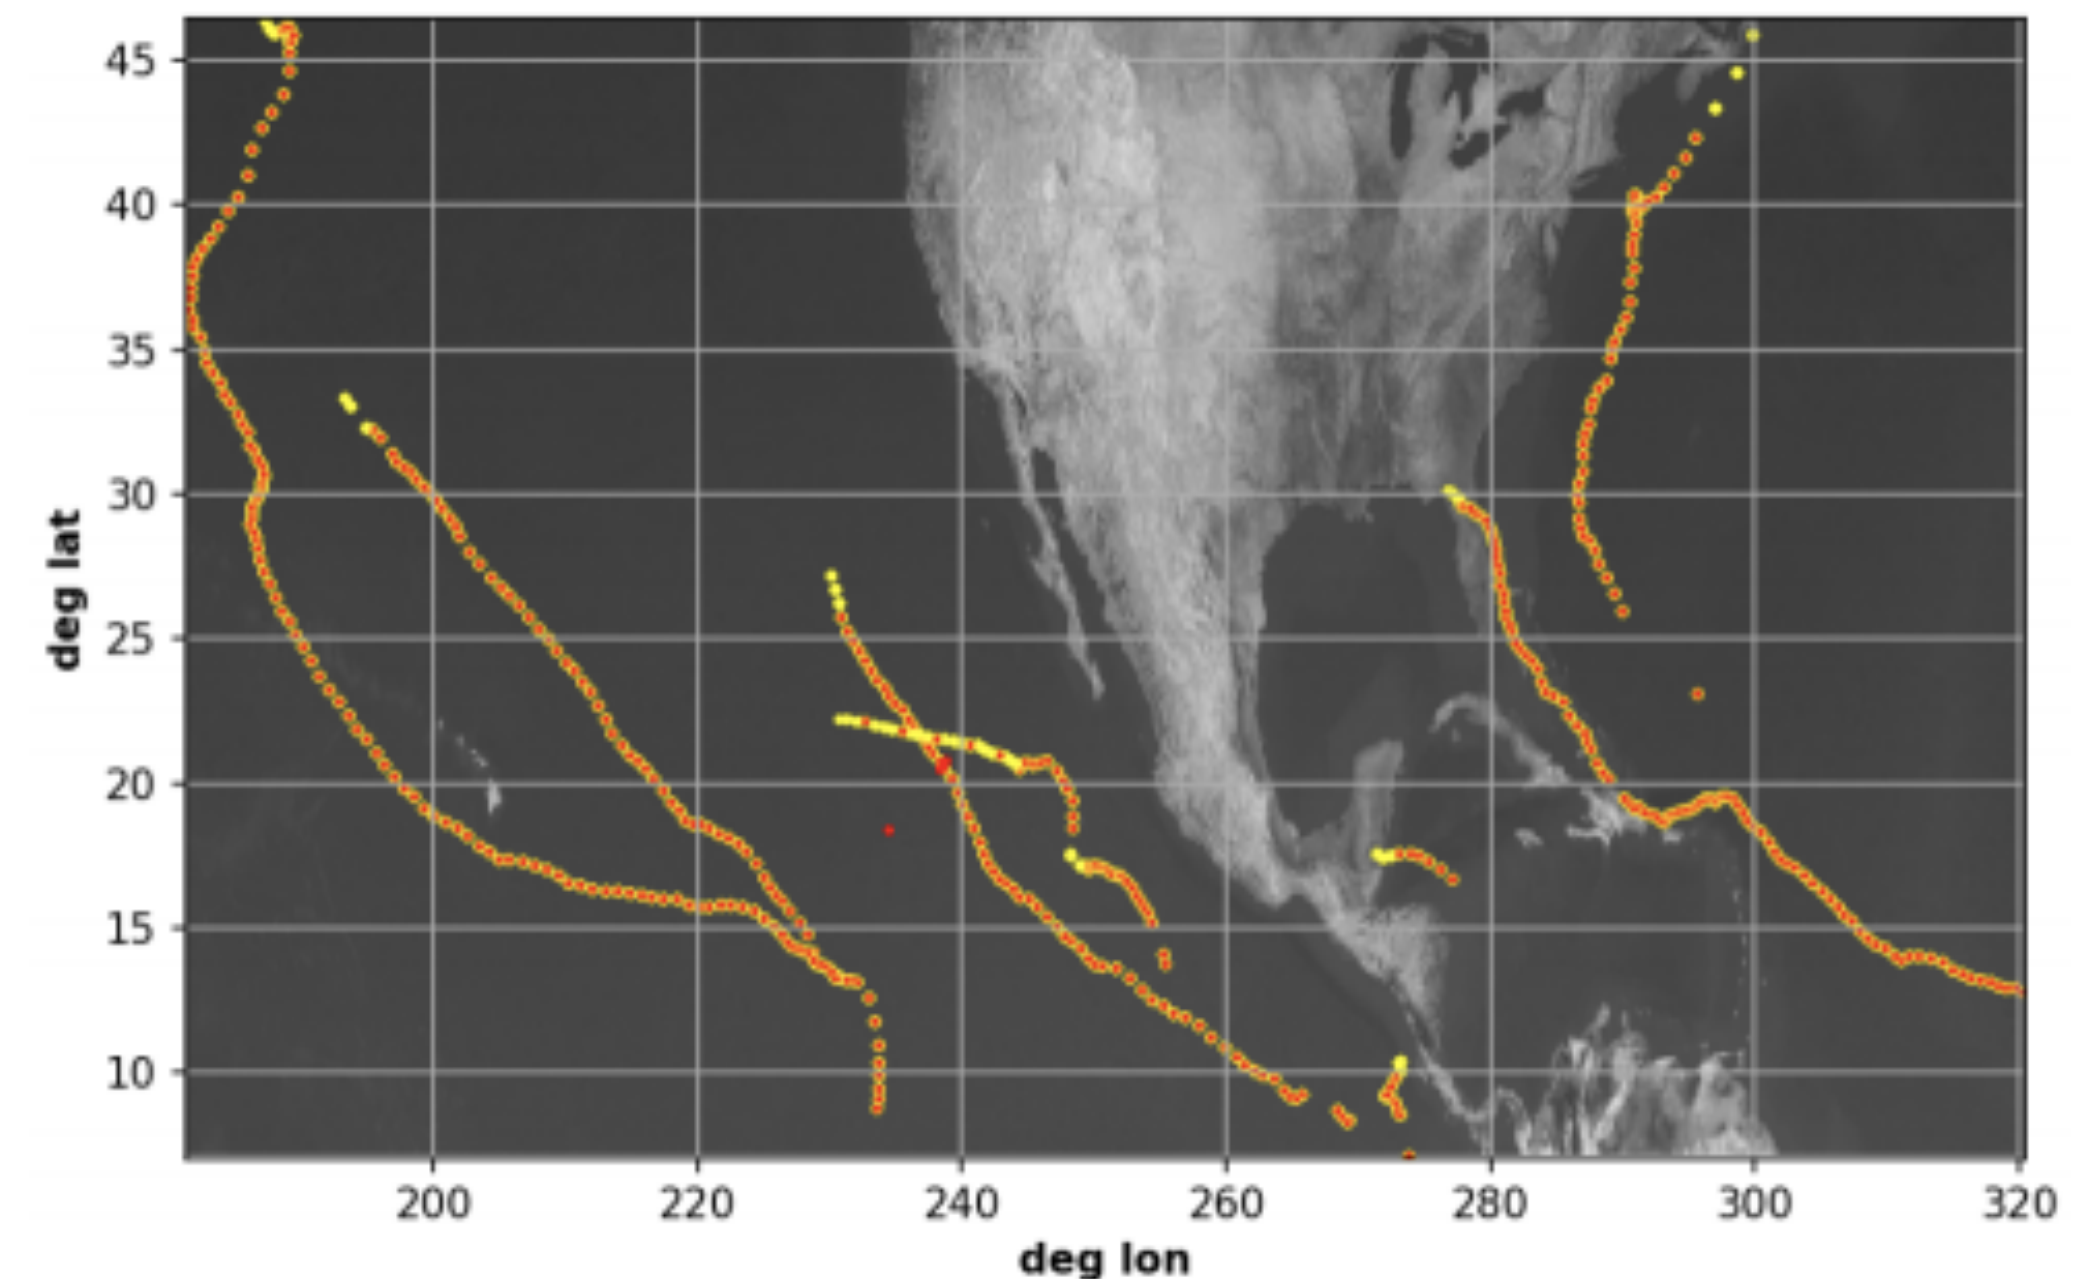


Extreme Weather Events

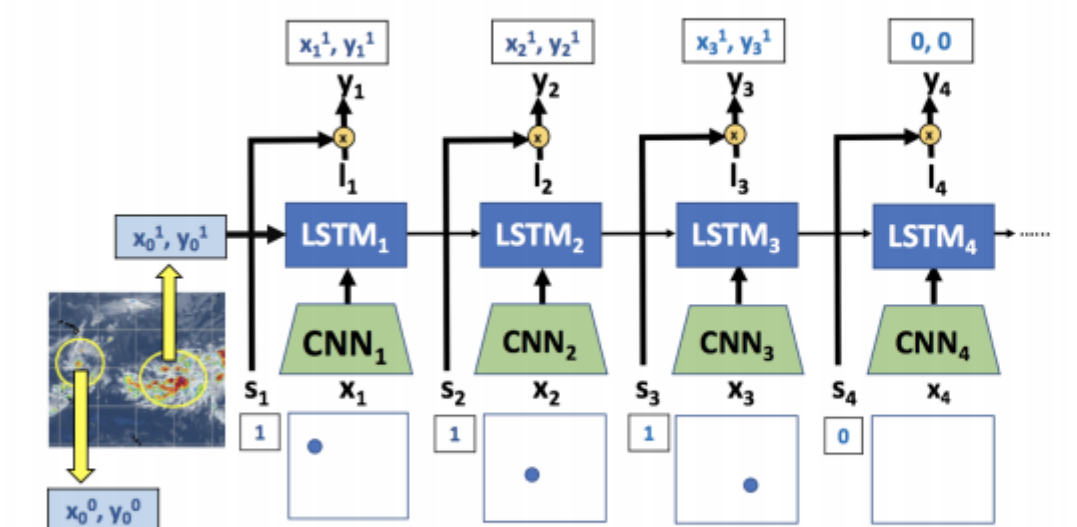
Segment Tropical Cyclones, Atmospheric Rivers from background



Given CAM5 outputs of a tropical cyclone and its initial position, track its trajectory.



Tools for data assimilation, analysis of NWP simulations, and new types of decision support.



Hybrid Physical Process Modelling

Predict future sea surface temperature (SST) from previous synthetic SST data from NEMO (Nucleus for European Modeling of the Ocean)

Physical Model: Advection-Diffusion Equation

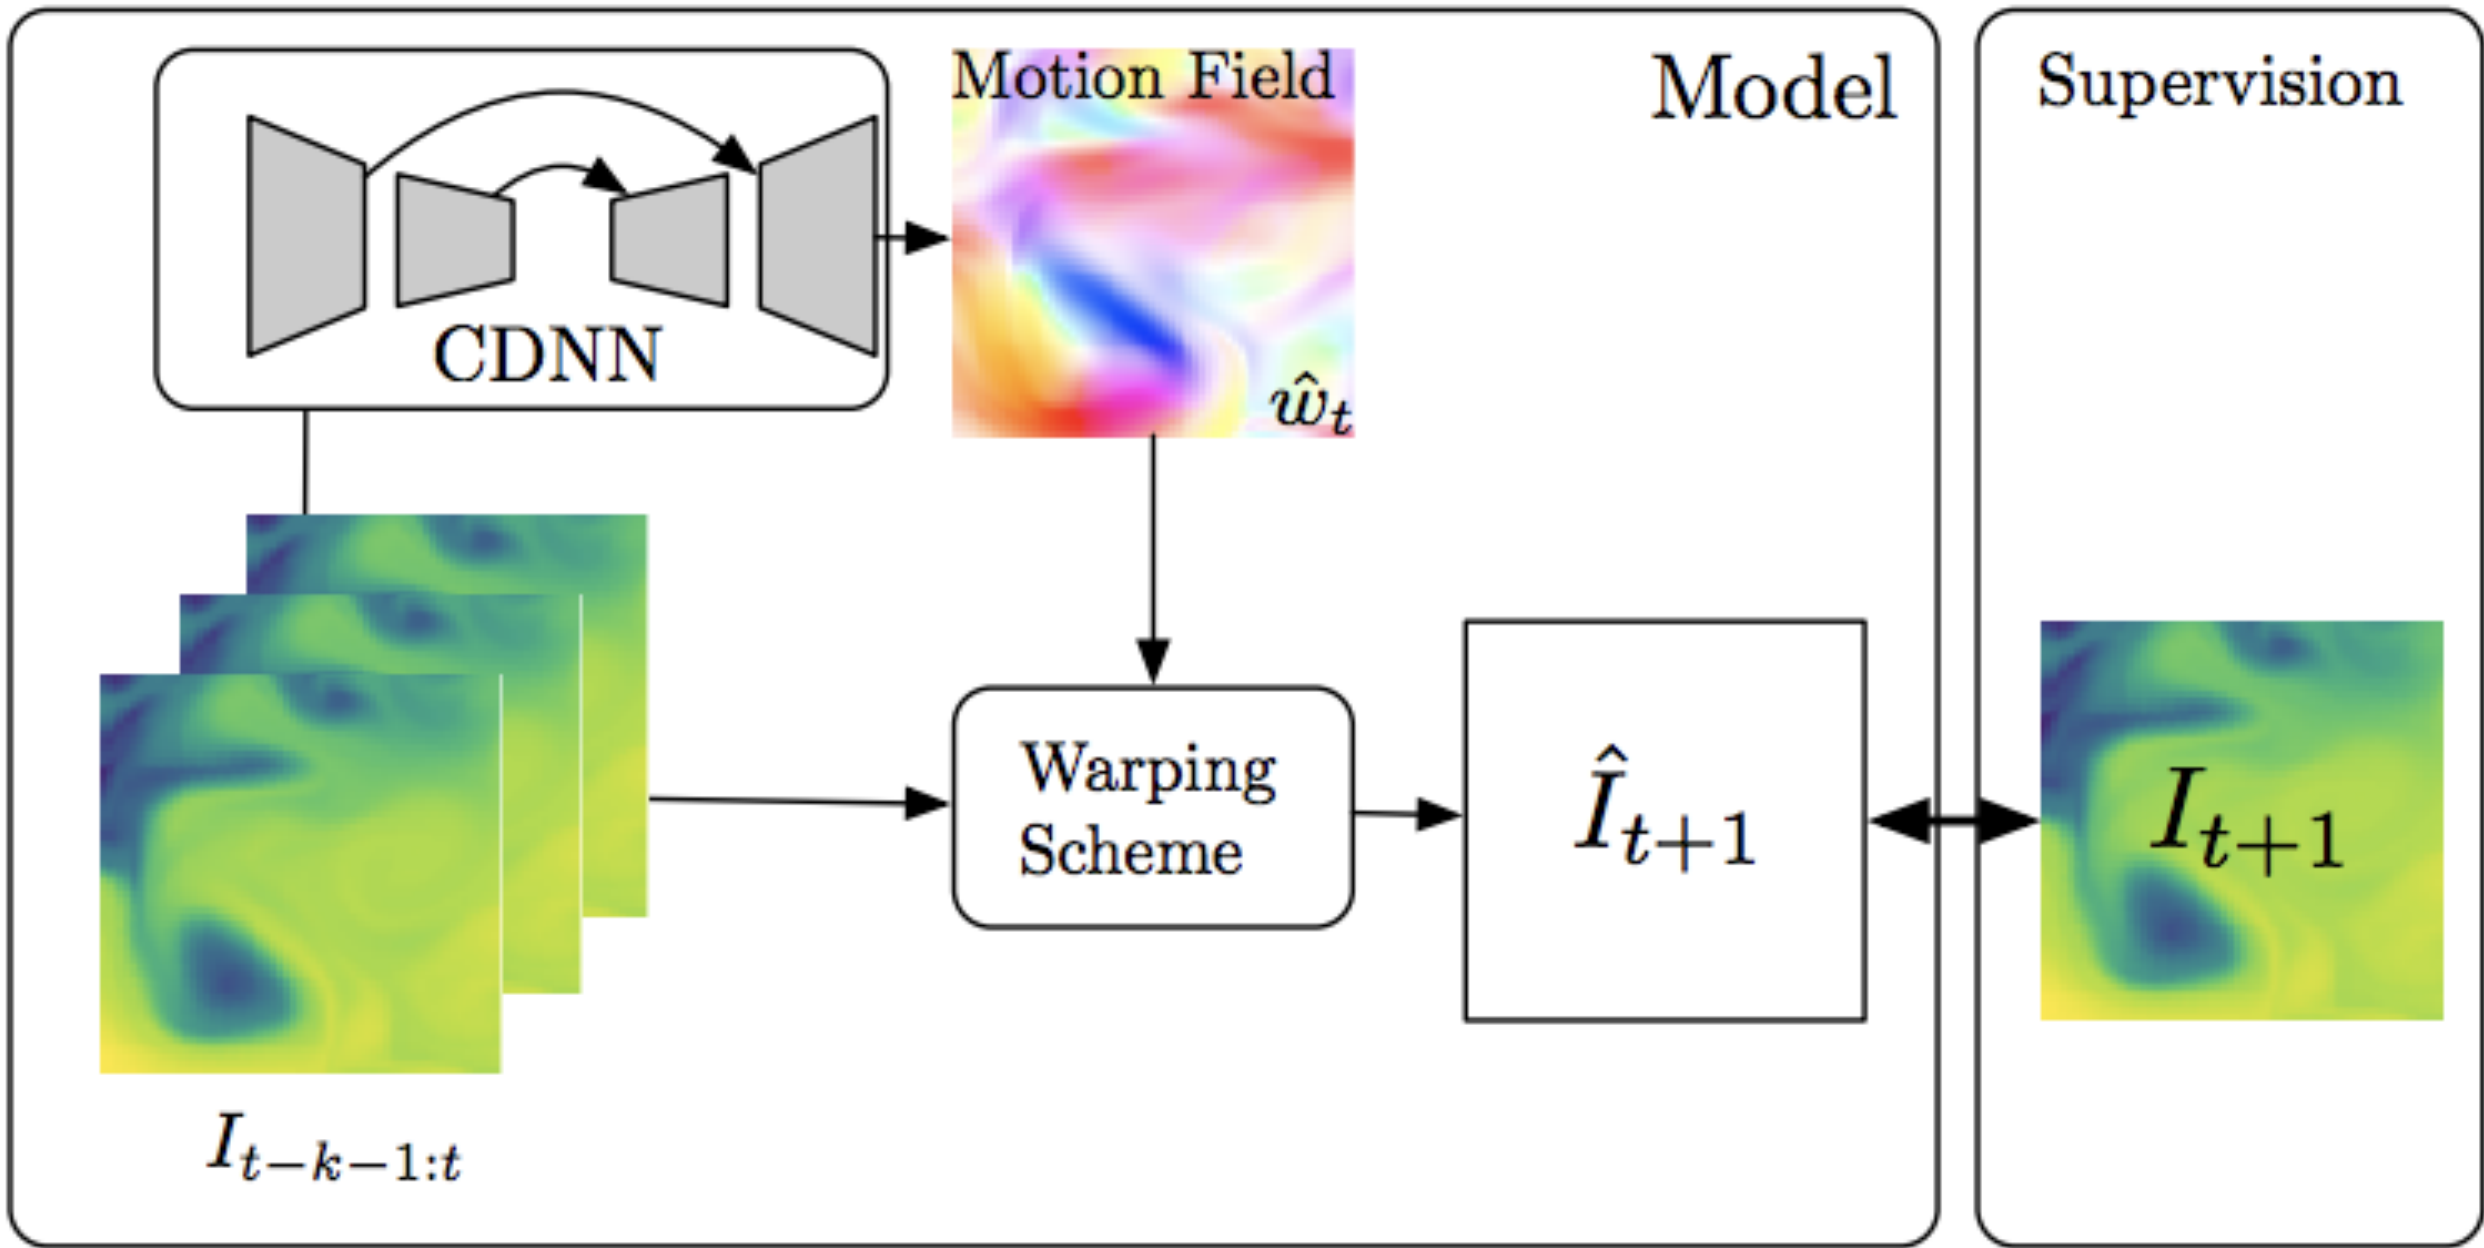
$$\frac{\partial I}{\partial t} + (w \cdot \nabla) I = D \nabla^2 I$$

Solution

$$I(x, t) = \int_{\mathbb{R}^2} k(x - w, y) I_0(y) dy$$

$$\hat{I}_{t+1}(x) = \sum_{y \in \Omega} k(x - \hat{w}(x), y) I_t(y)$$

Key Idea: Predict w



$$L_t = \sum_{x \in \Omega} \left\| \hat{I}_{t+1}(x) - I_{t+1}(x) \right\|^2 + \lambda_{\text{div}} (\nabla \cdot \hat{w}_t(x))^2 + \lambda_{\text{magn}} \left\| \hat{w}_t(x) \right\|^2 + \lambda_{\text{grad}} \left\| \nabla \hat{w}_t(x) \right\|^2$$

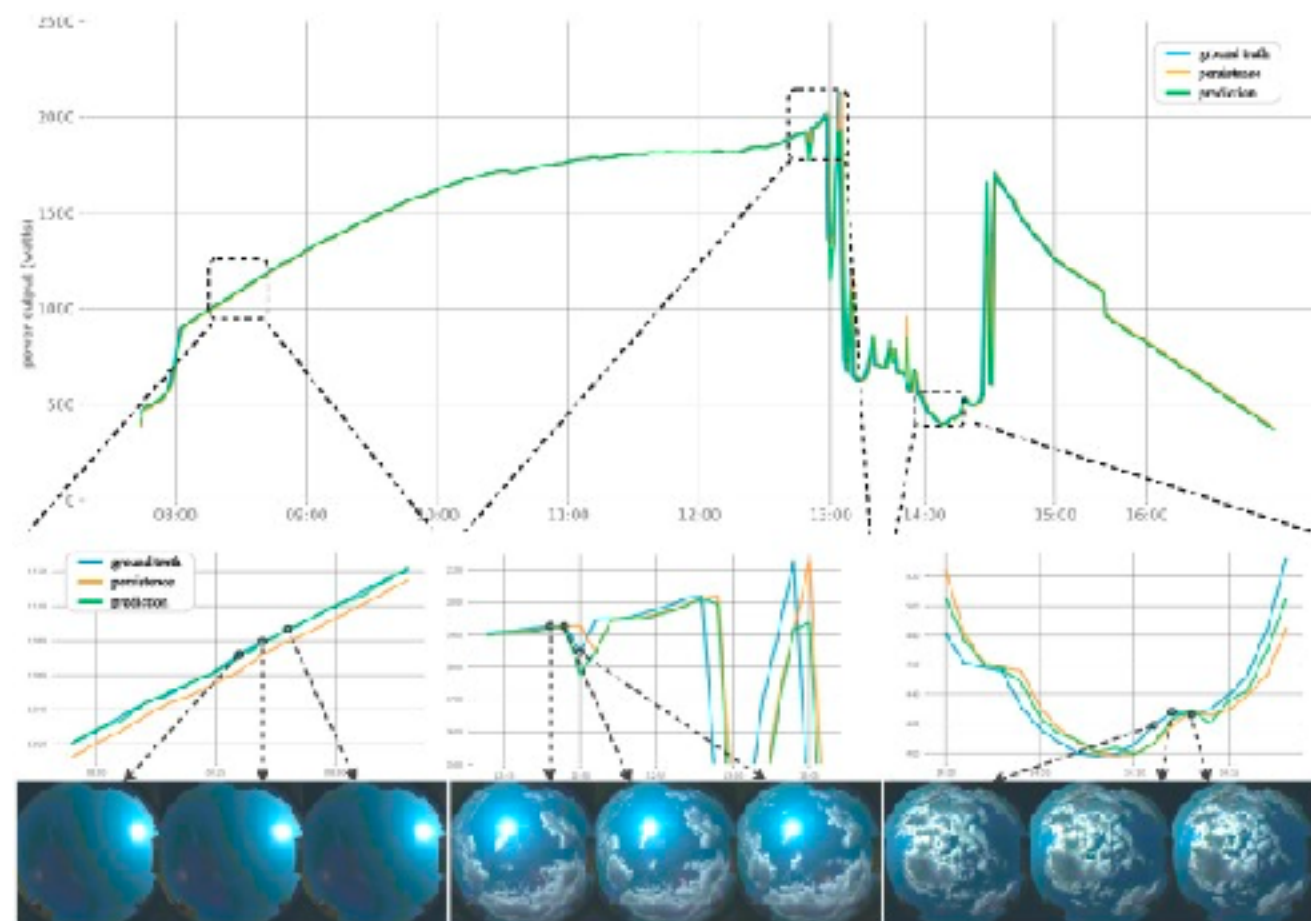
Model	Average Score (MSE)	Average Time
Numerical model [1]	1.99	4.8 s
ConvLSTM [9]	5.76	0.018 s
ACNN	15.84	0.54 s
GAN Video Generation ([7])	4.73	0.096 s
Proposed model with regularization	1.42	0.040 s
Proposed model without regularization	2.01	0.040 s



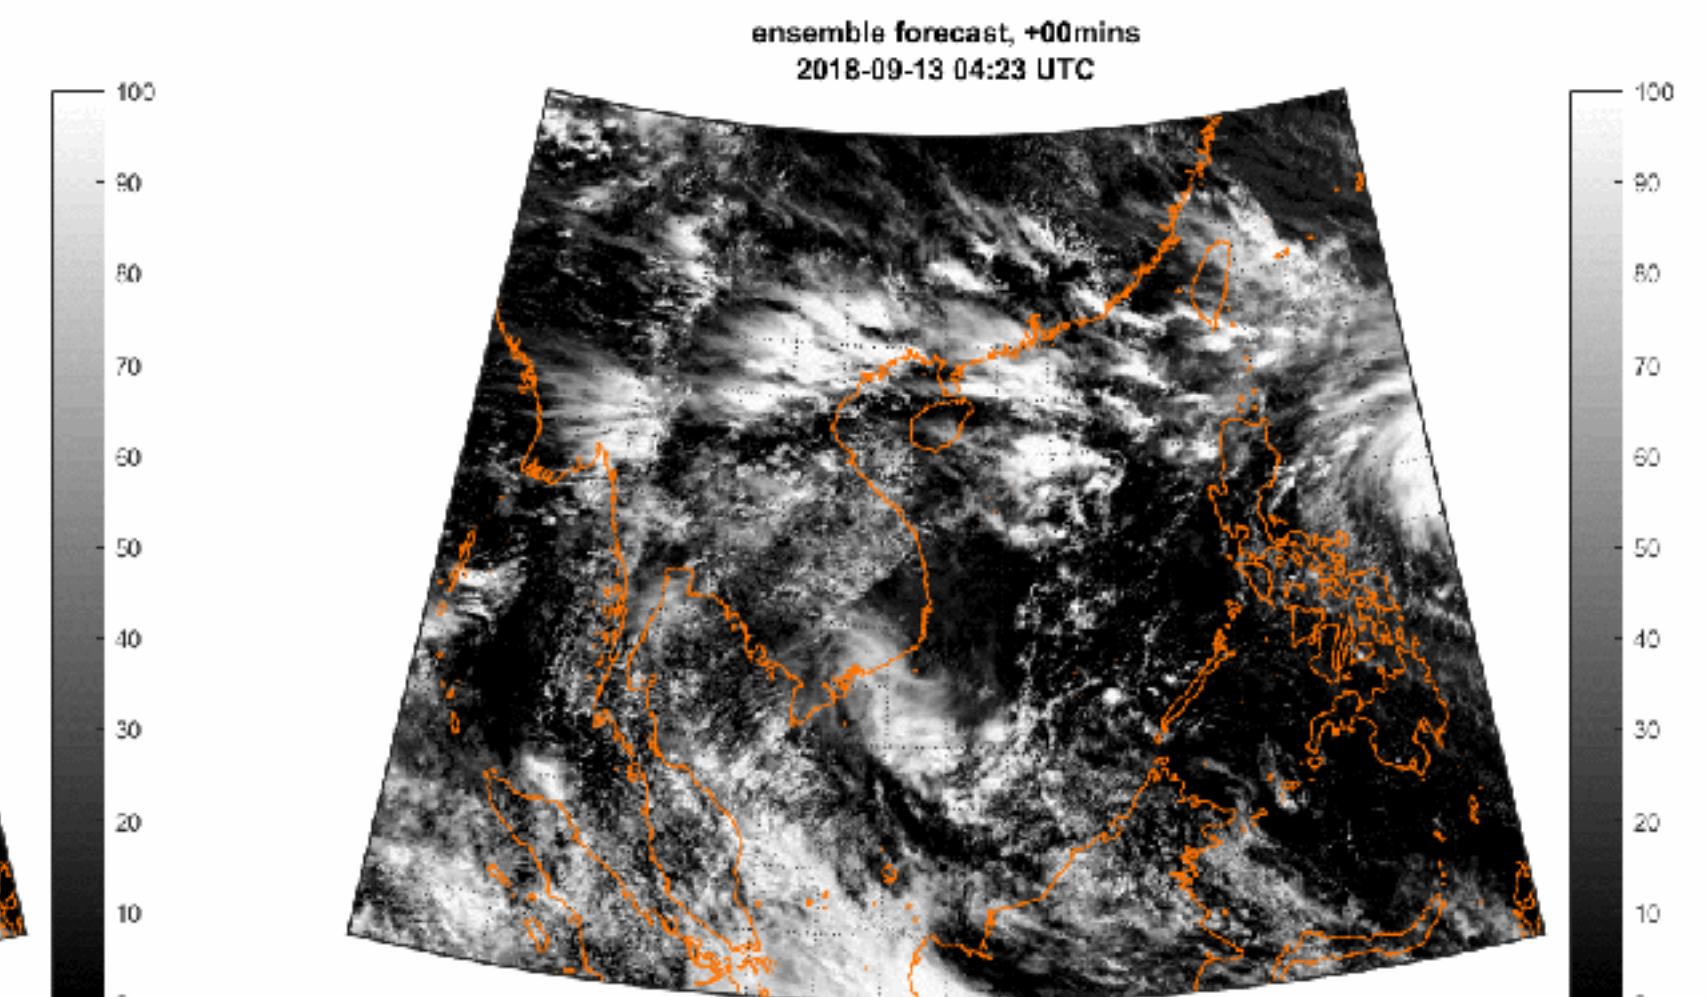
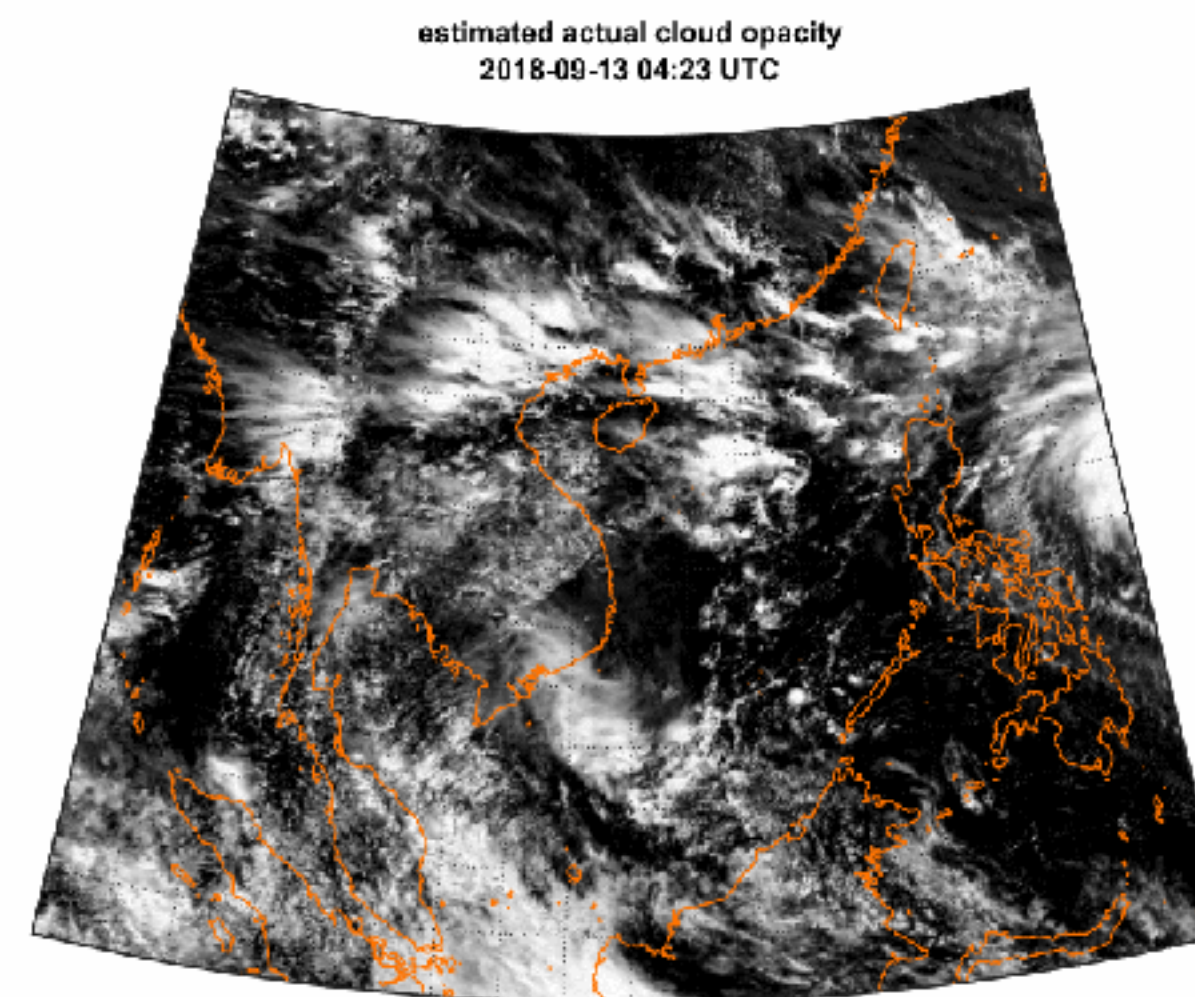
Solar Nowcasting

Predict solar irradiance, accounting for clouds.

- Numerical weather models become out of date with respect to the most recent observations.
- Solar irradiance is greatly affected by clouds; operational numerical weather models can't resolve clouds.
- Radiative transfer codes in numerical weather models are some of the most computationally expensive bits of numerical weather models.



SOLCAST



Energy Consumption

**Dramatically increase efficiency of
existing systems
Application to Google Data Centres**



Data Centre Energy Usage

Data centres across the world use around **3% of the world's electricity**

Cooling energy is the largest non-server load (up to 40% of total energy usage)

State

- ✓ Incoming IT load
- ✓ Power meters
- ✓ Pressure sensors
- ✓ Temperature sensors
- ✓ Water flow meters
- ✓ Pump and fan speeds
- ✓ Fault alarms
- ✓ Weather conditions

Actions

- Number of cooling towers ✓
- Number of chillers ✓
- Number of pumps ✓
- Temperature setpoints ✓
- Pressure setpoints ✓
- Flow setpoints ✓
- Valve positions ✓

Over 1,200 state variables and 20 actions



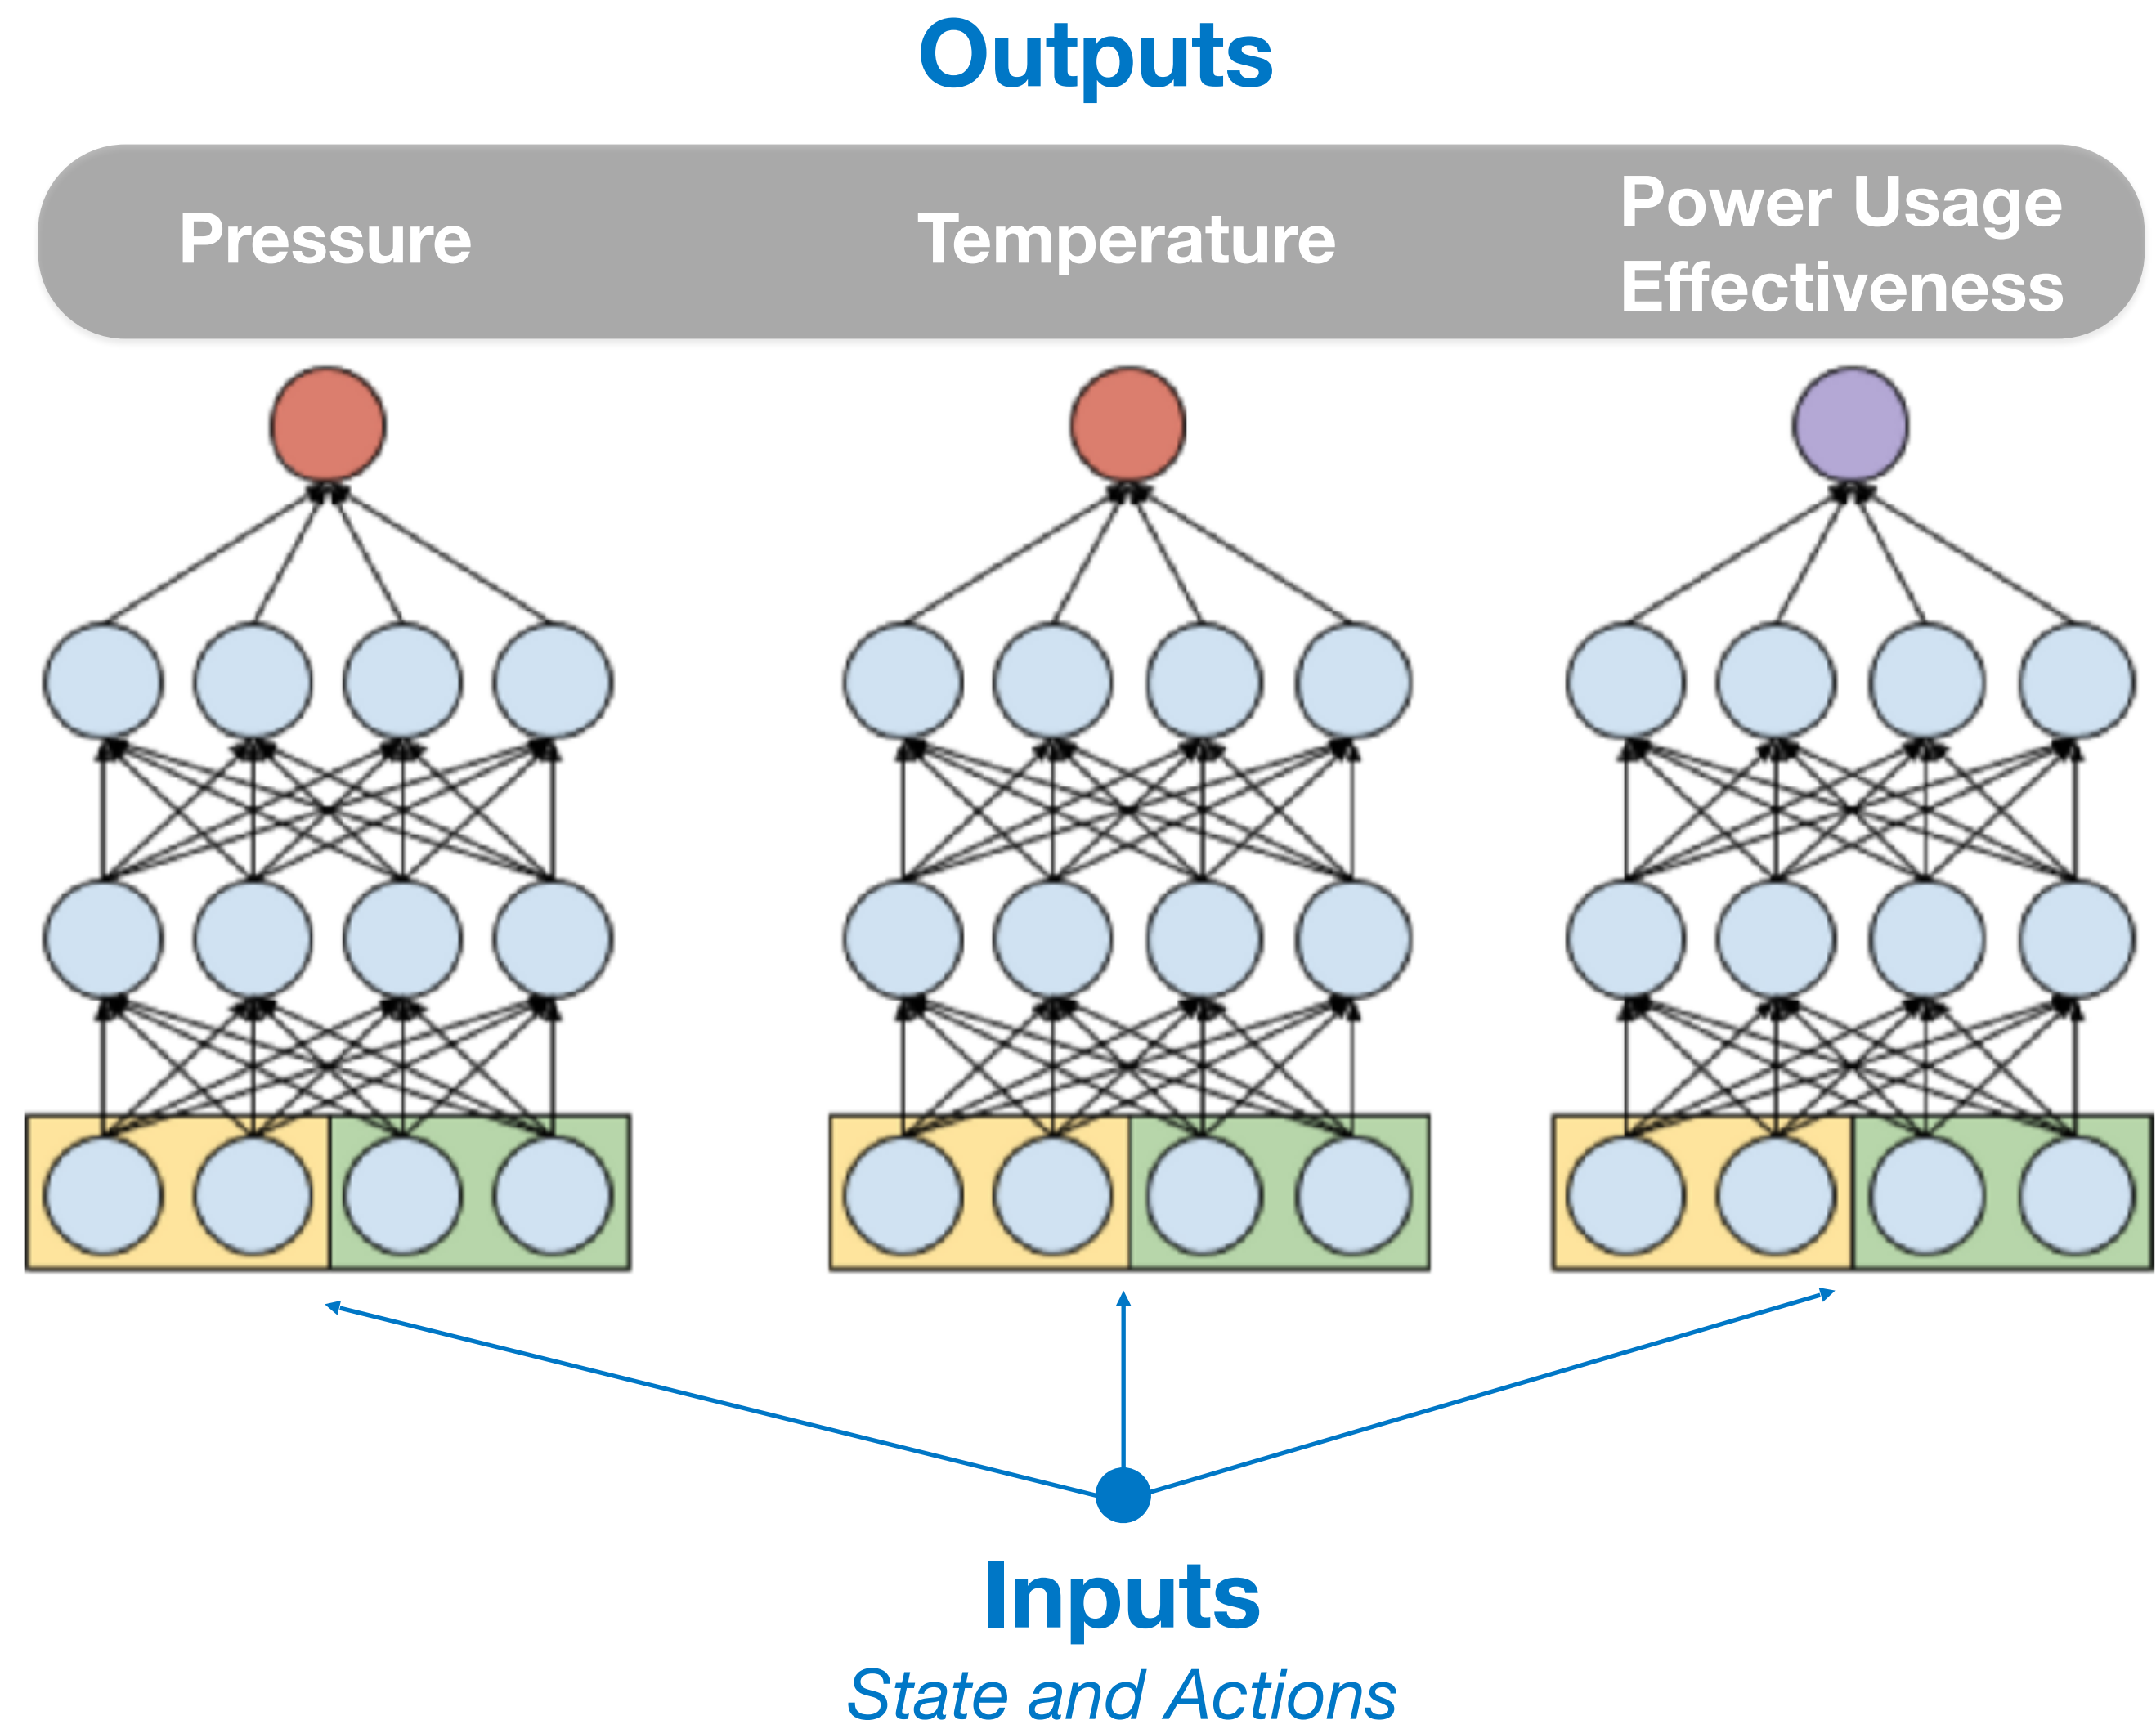
General Learning Framework for DC Operations

State inputs

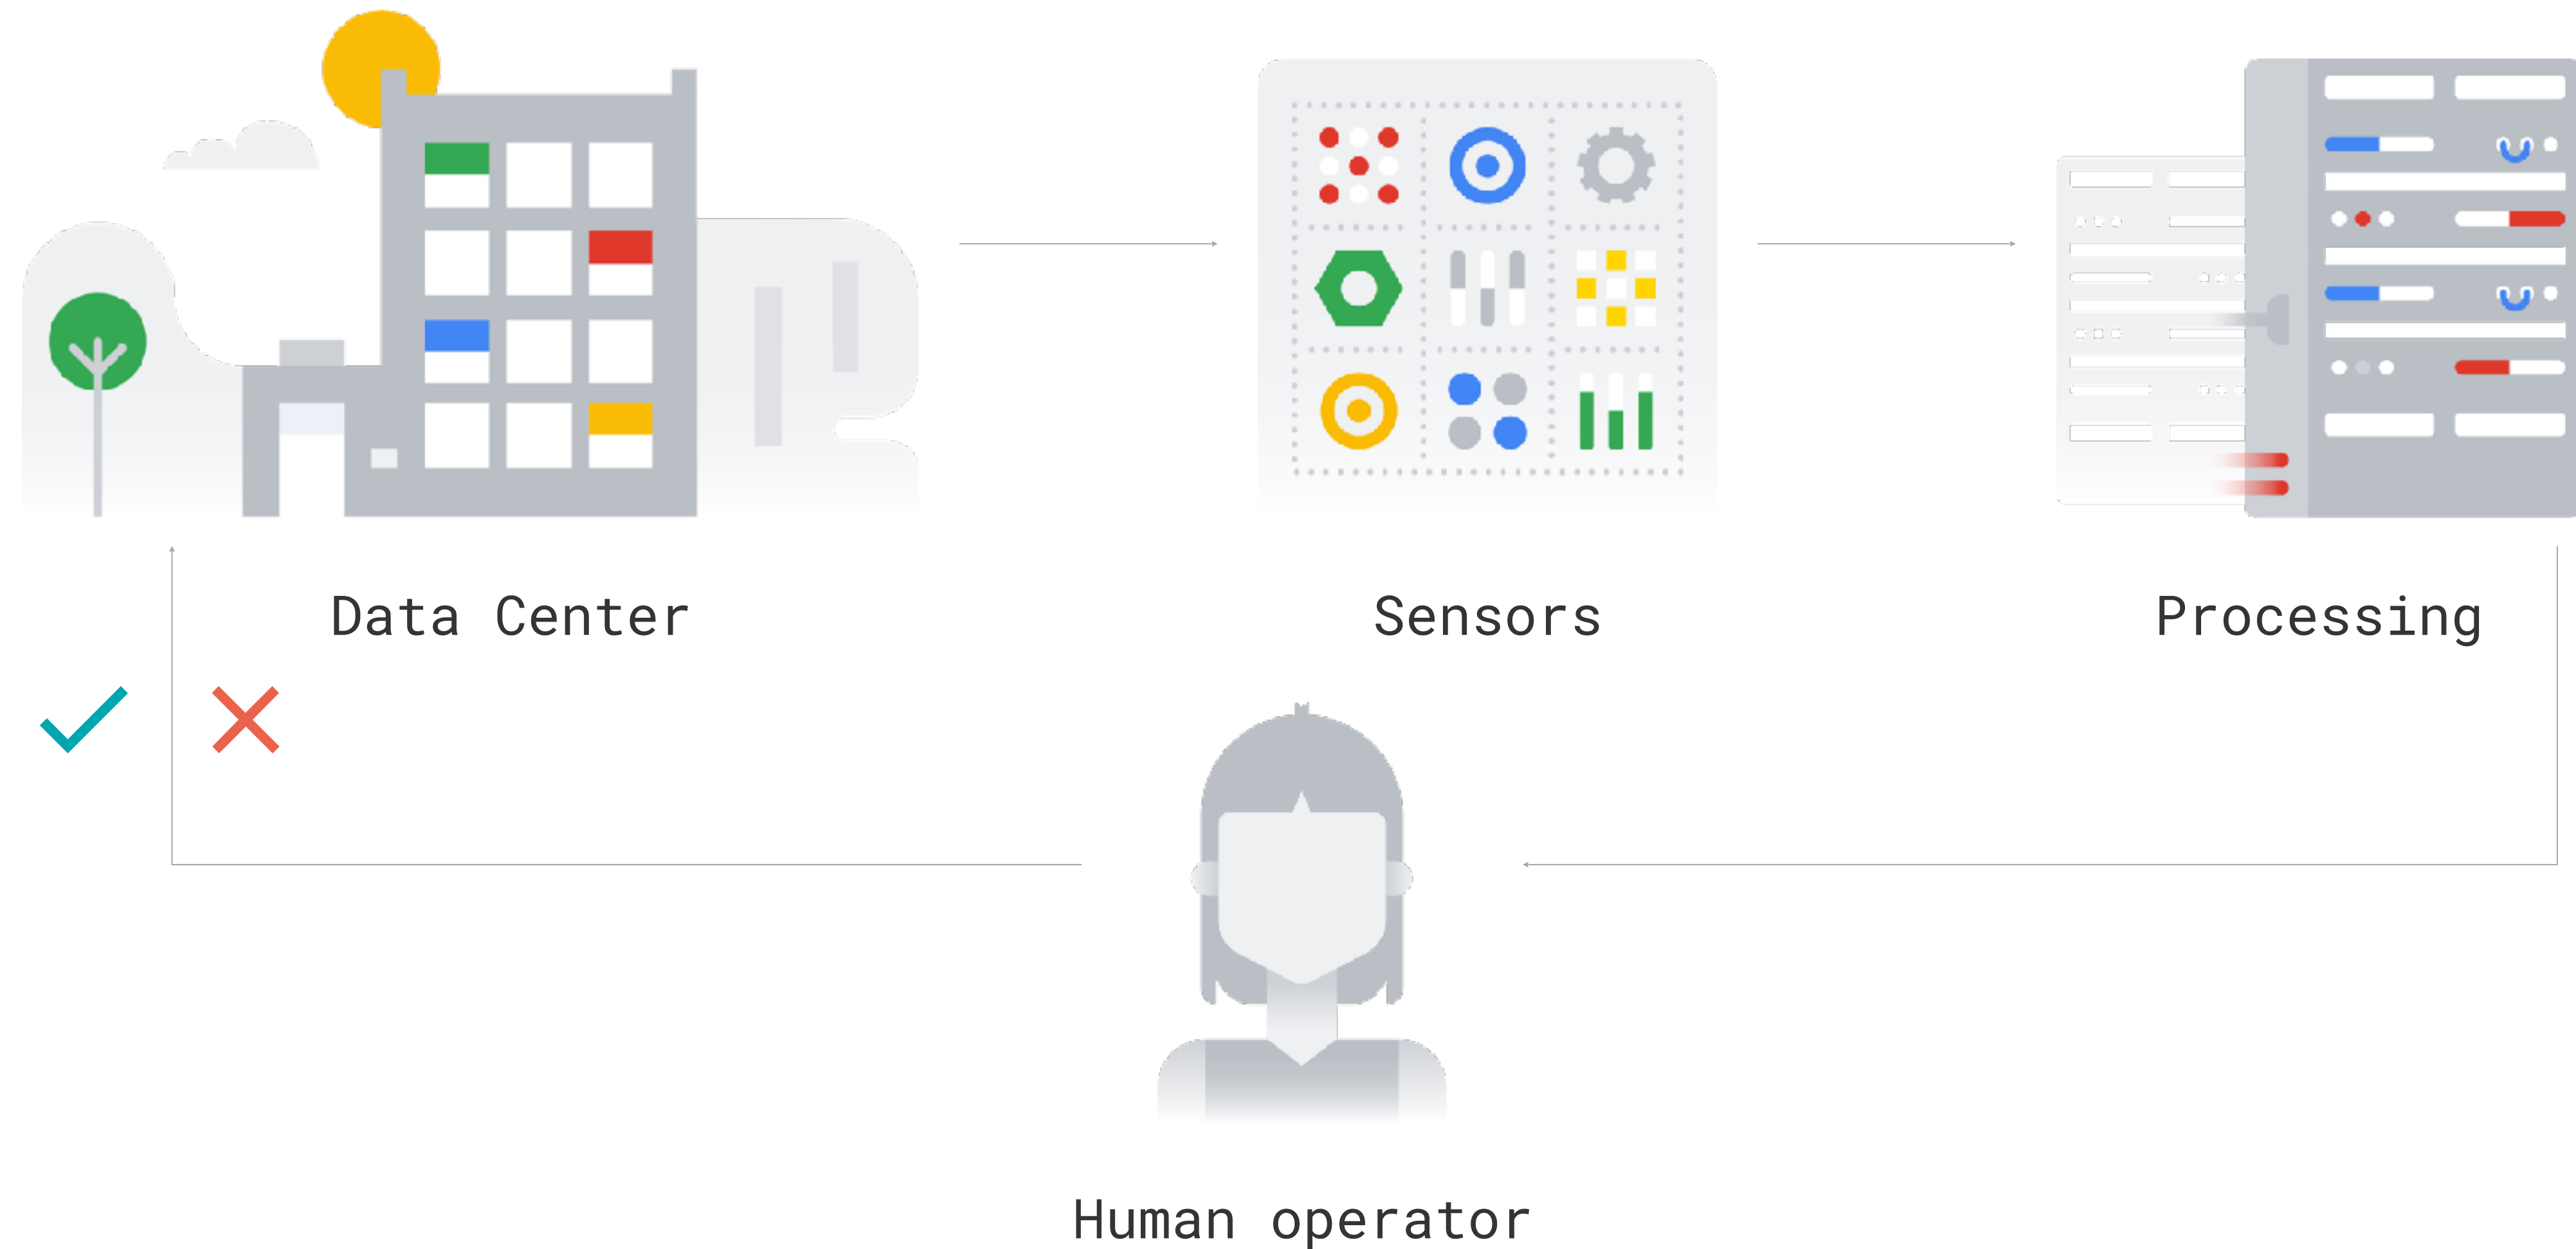
- Current IT load
- Power meters
- Pressure sensors
- Temp sensors
- Weather
- Fan speeds
- ...

Actions

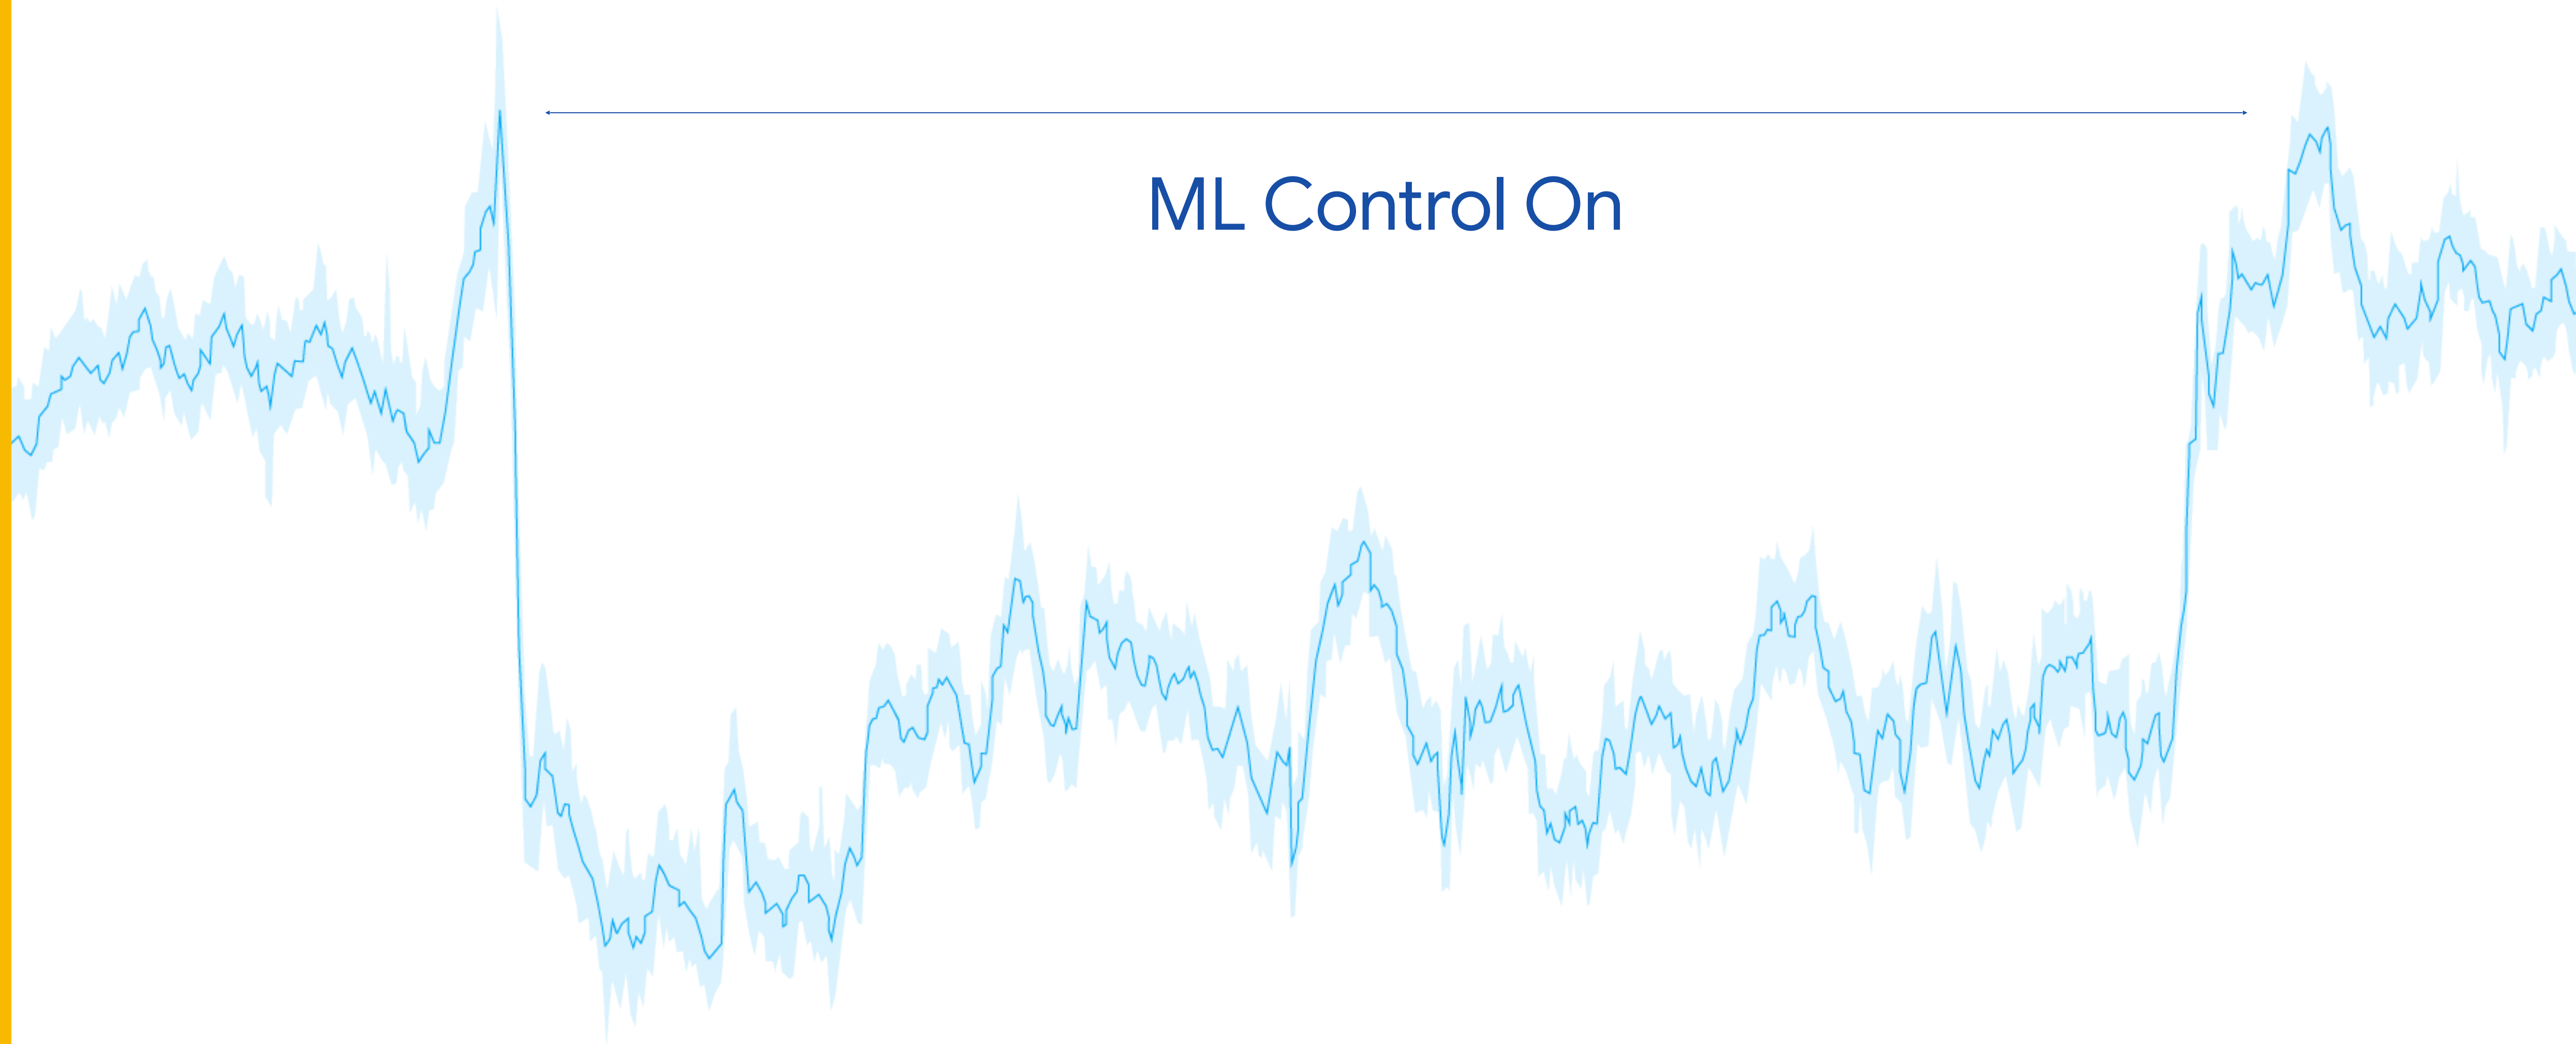
- # active coolers
- # chillers
- Pumps on/off
- Temp setpoints
- Valve setpoints
- Pressure setpoints
- ...



Every five minutes: generate recommendations, send to a human operator for implementation



40% reduction in data center cooling energy



System Insights

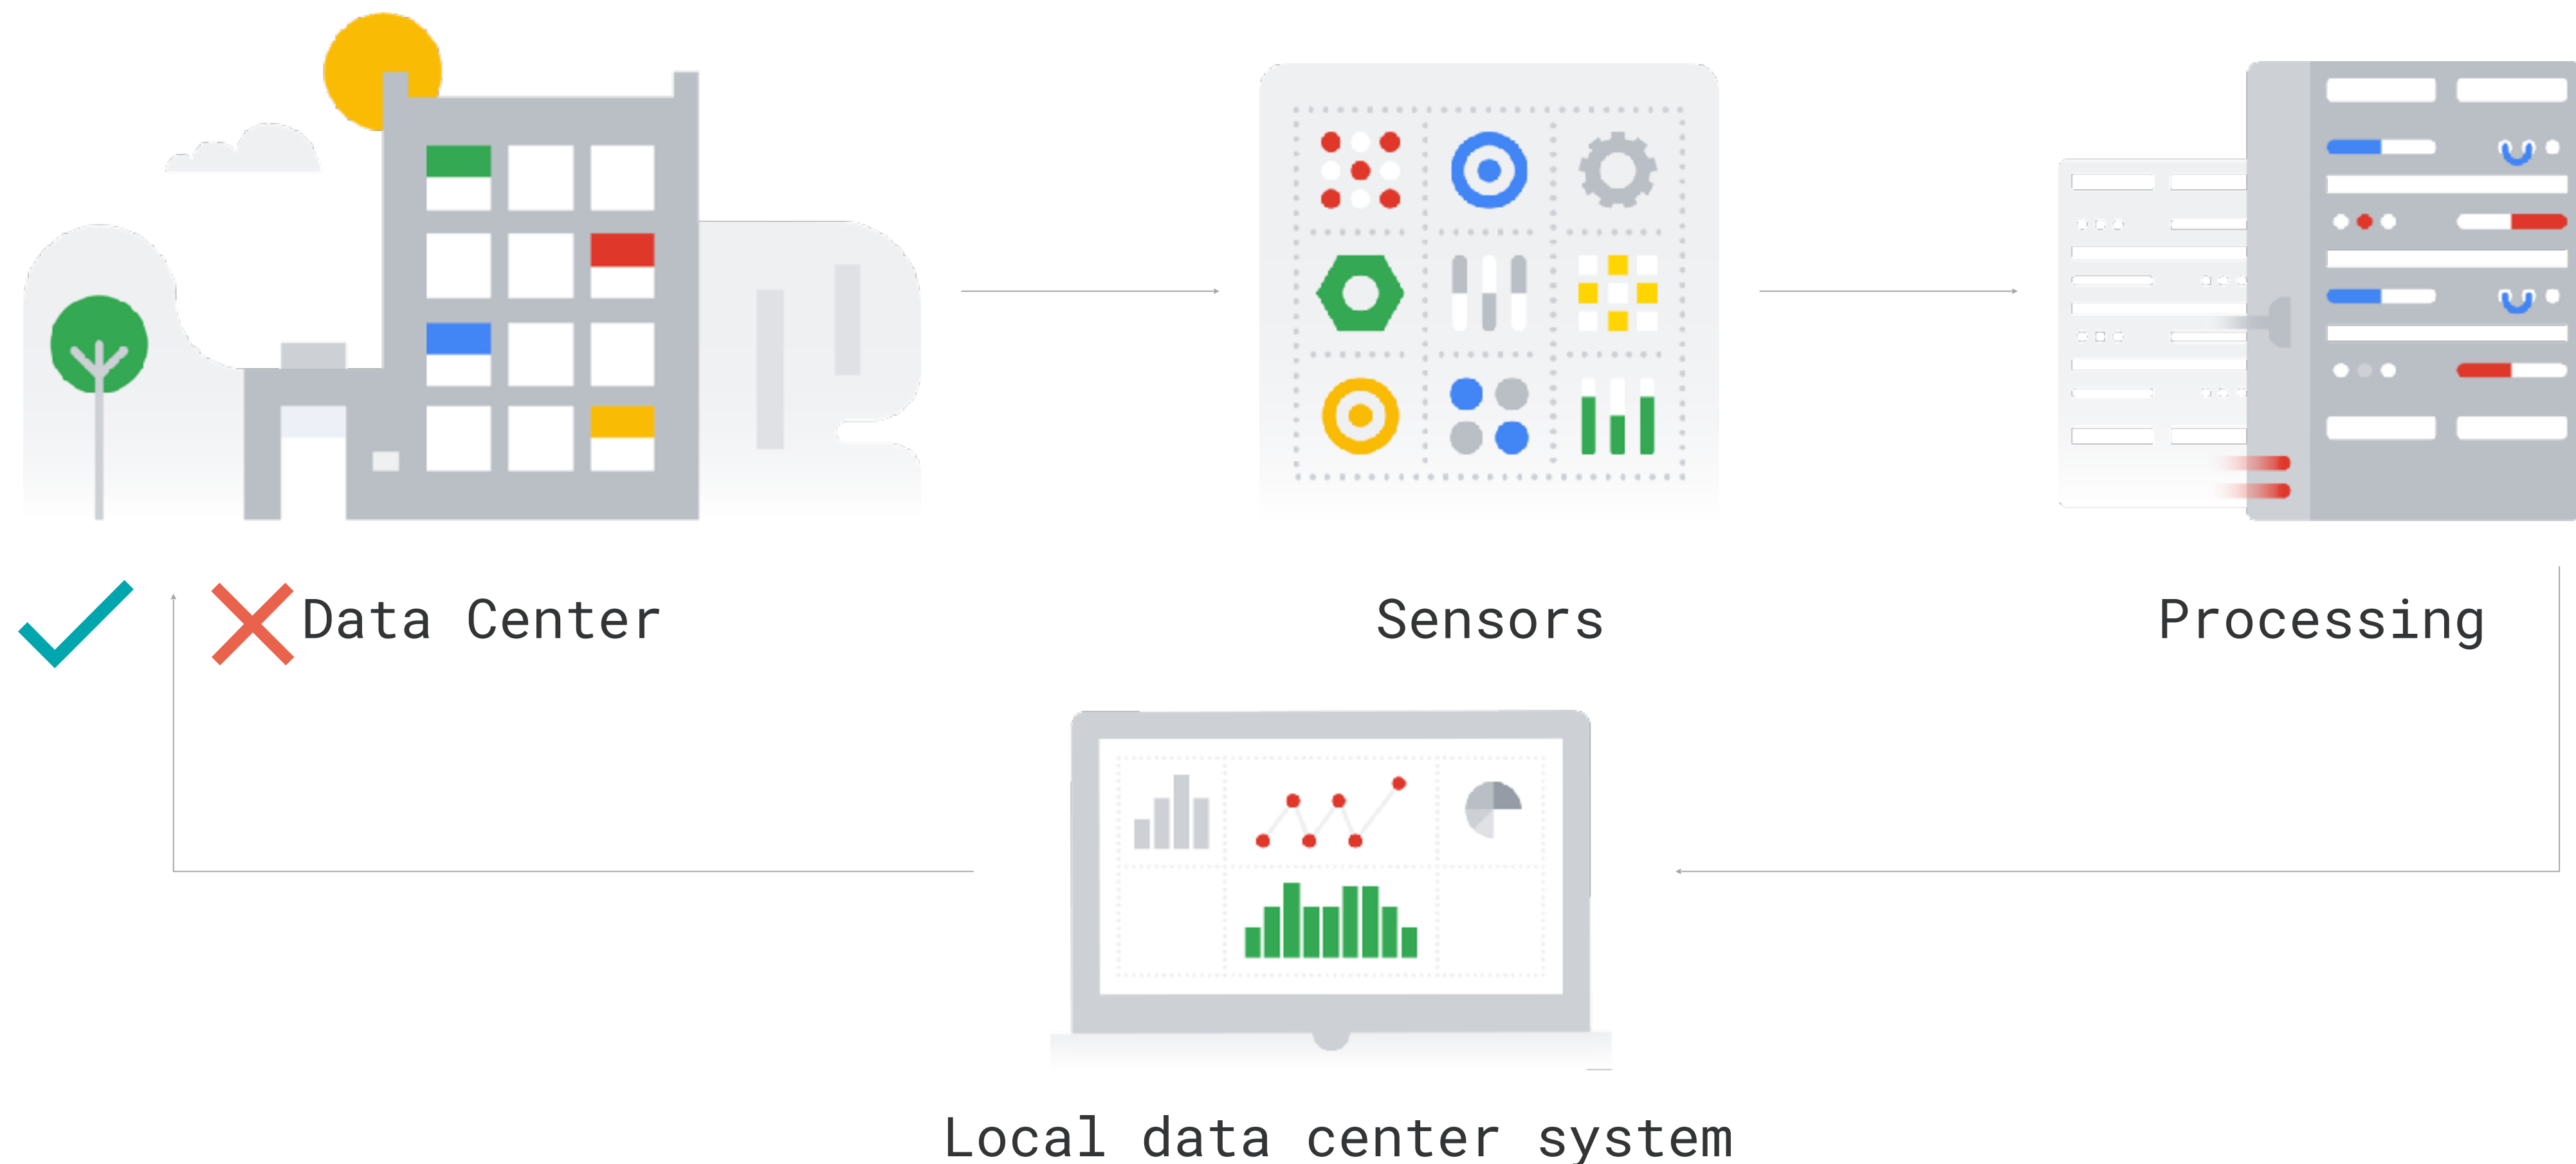
Spread the load across more equipment.
Local v. Global trade-offs.

Higher flow is not always better.
Reduced water flow to chillers in some weather conditions.

Shifting the loads.
Learned to shift cooling load to components that were more or less efficient at different times of year.

After three quarters of operation, scaling it up and getting it into production using a safety-first automation approach

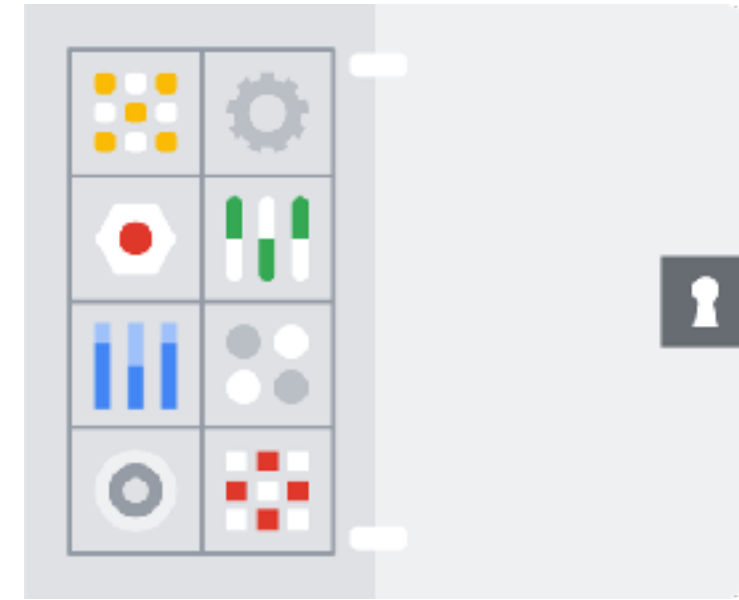
Recommendations are sent directly to the data centre, to be verified by the local controls system for safety before implementation.



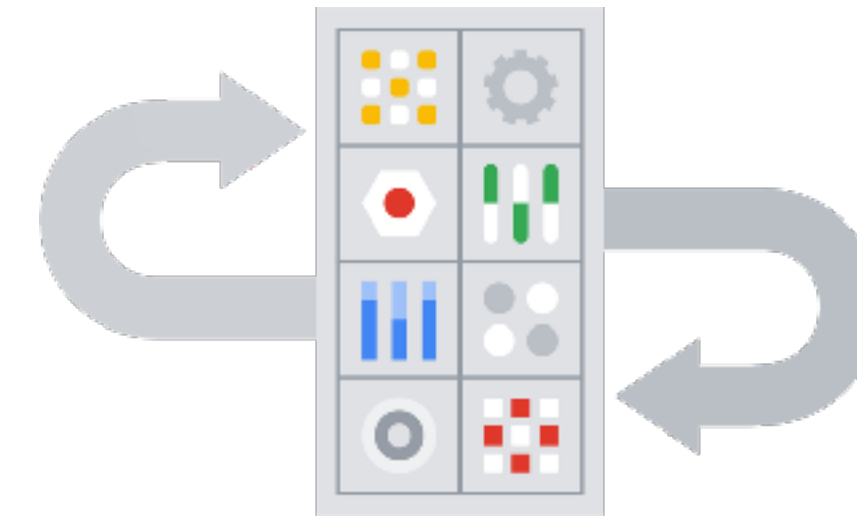
Safety-first for direct AI control



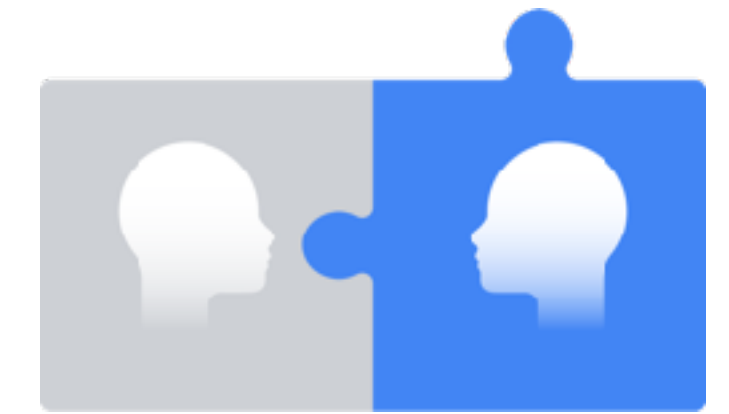
Continuous monitoring



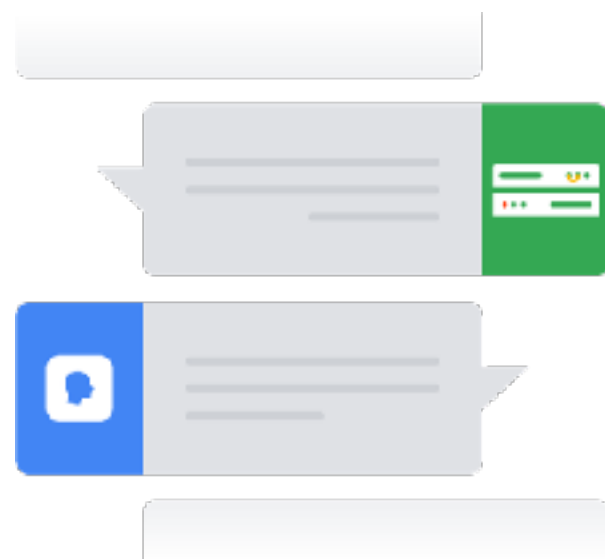
Automatic failover



Smooth transfer



Two-layer verification



Constant communication



Uncertainty estimation



Rules and heuristics as backup



Human in the loop

Managing Energy Generation

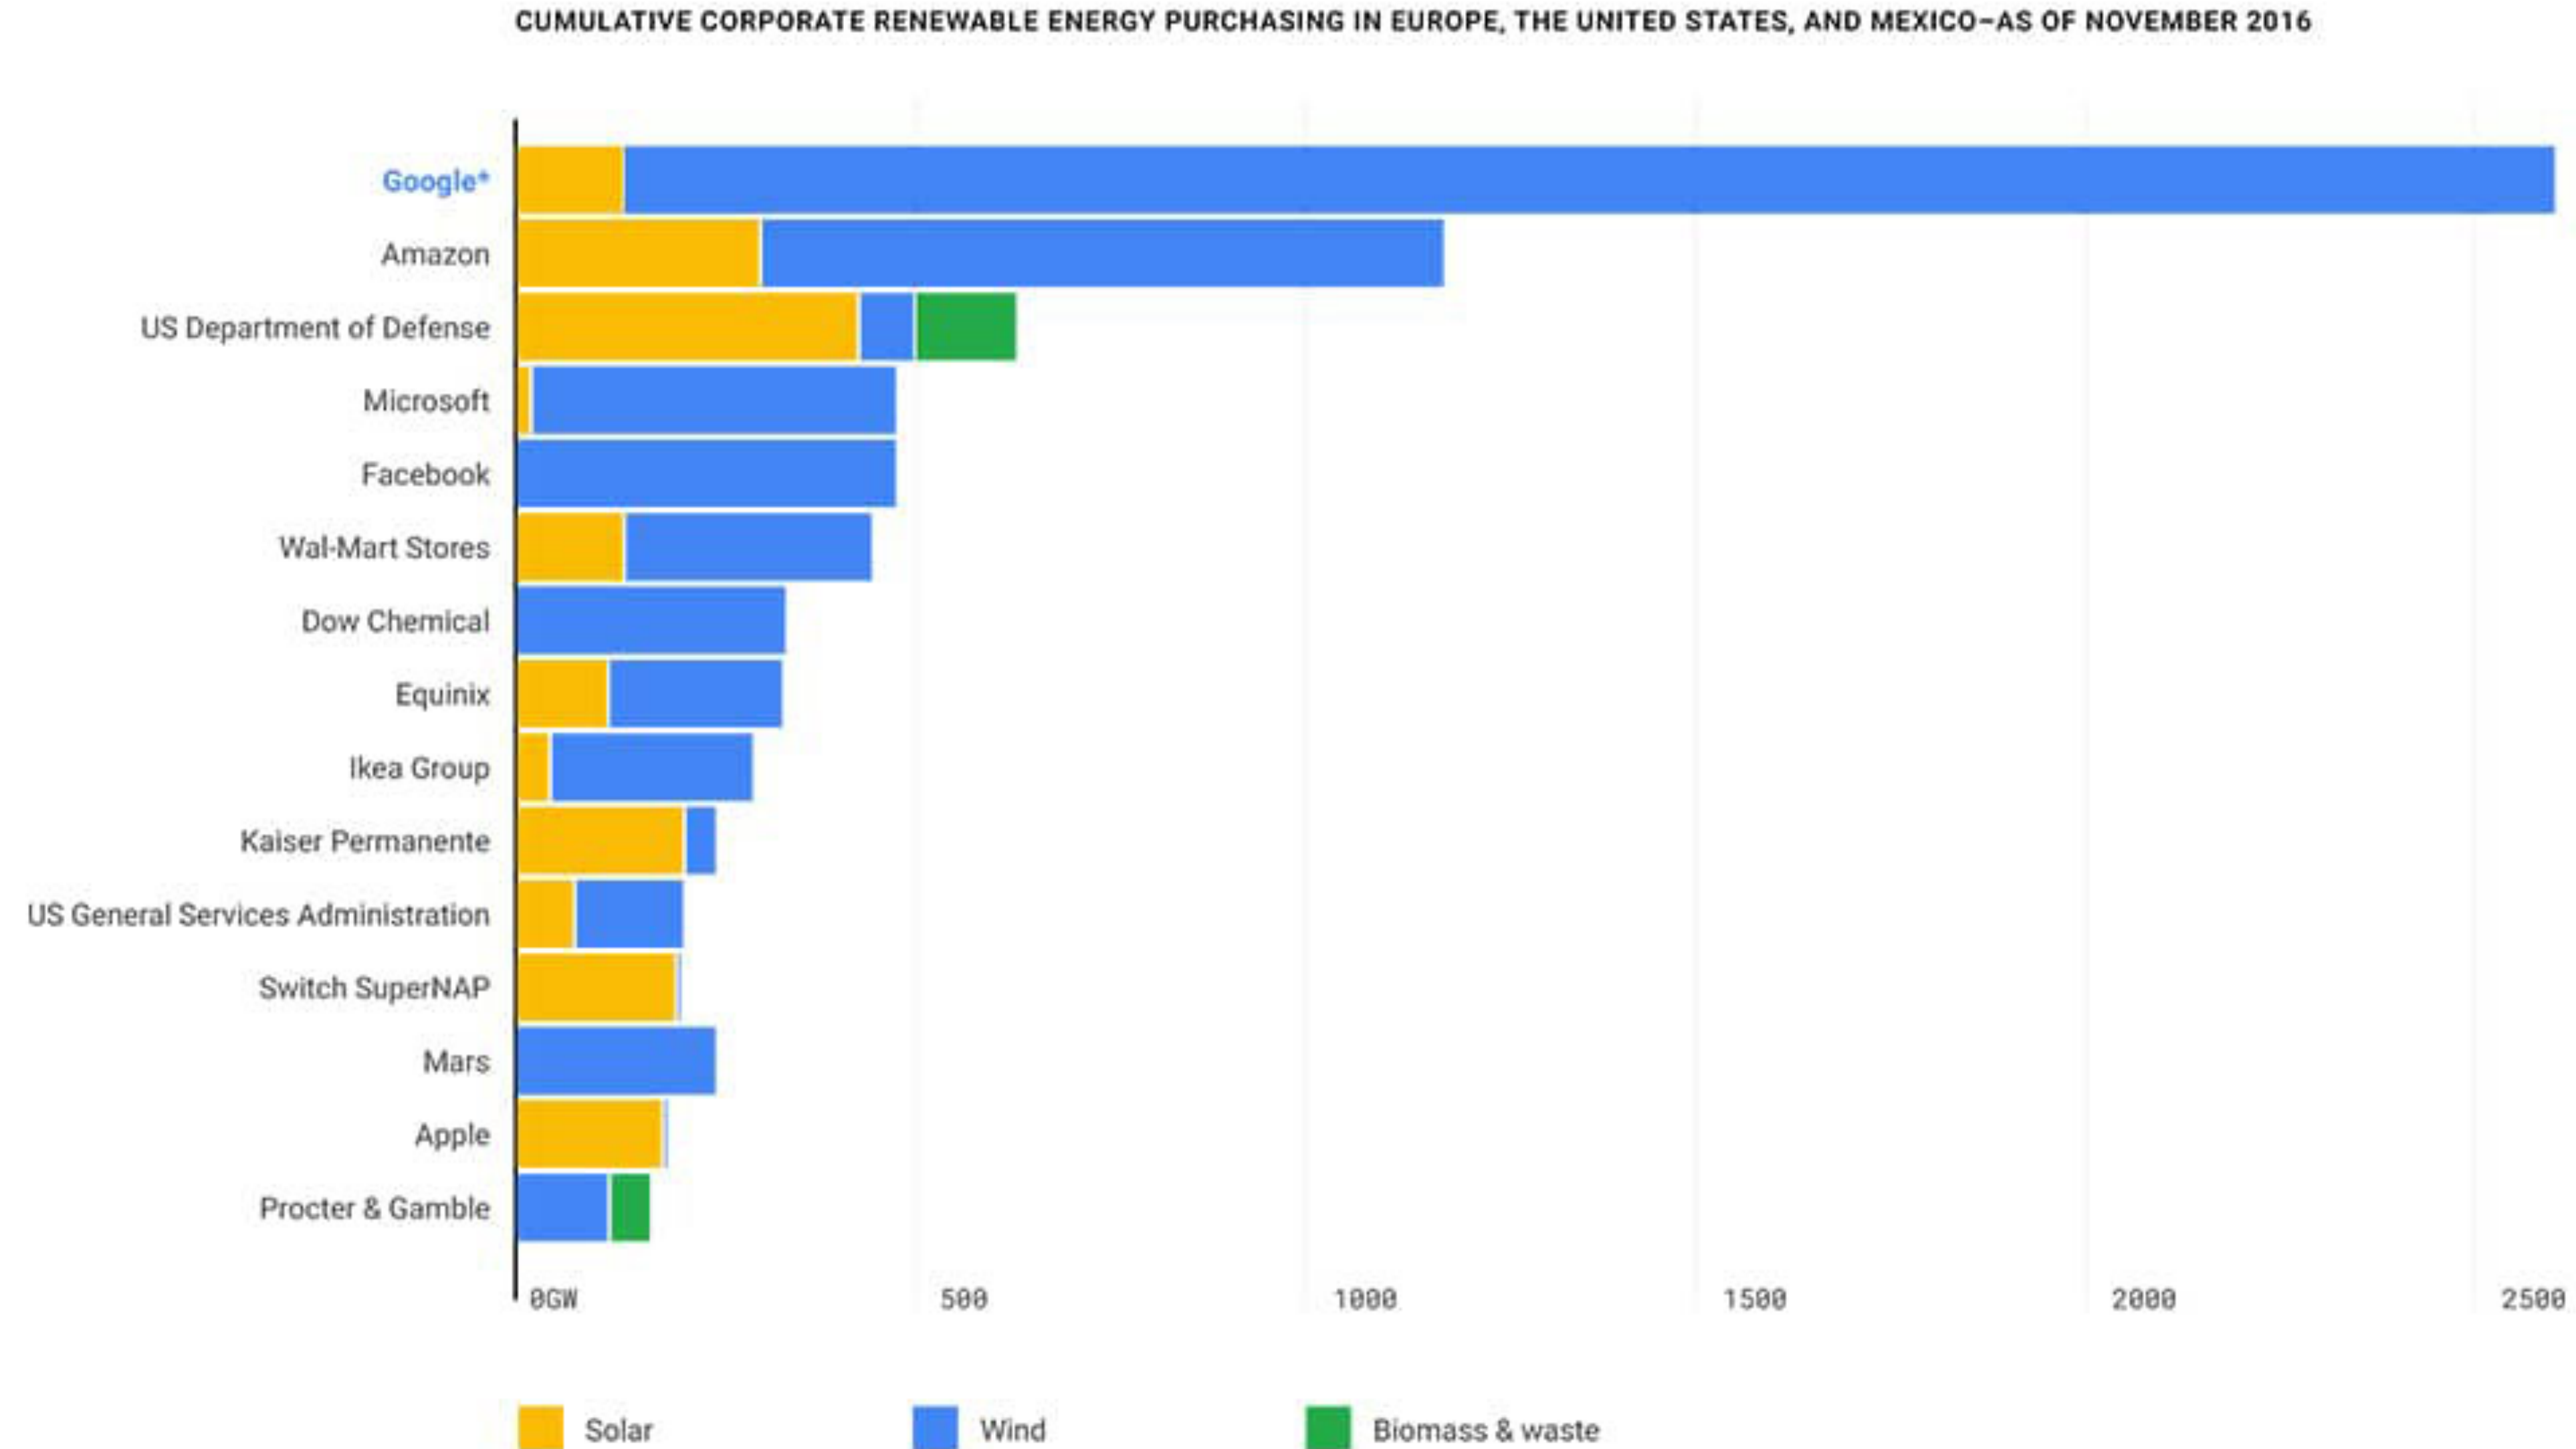
Improving the economics of wind energy to **accelerate adoption**

The cost of turbines has plummeted, but wind is **unpredictable** and **intermittent**

The unpredictability of renewable energy makes it **less valuable** than fossil fuel energy

One strategy: train a system for **predicting** and **scheduling** wind energy

Applying ML algorithms to **700MW** of Google's wind farm portfolio.

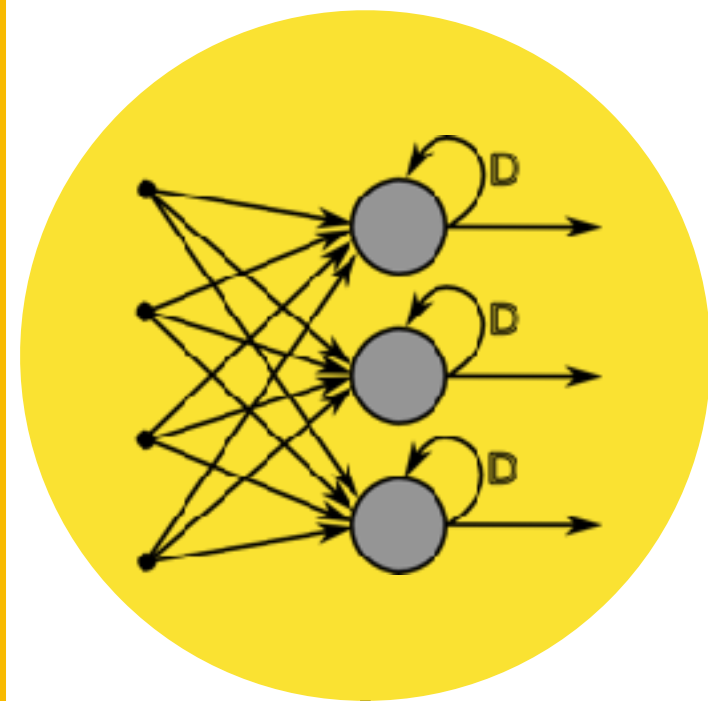


Source: Bloomberg New Energy Finance

*Google total also includes one project in Chile for 80 MW

Inputs

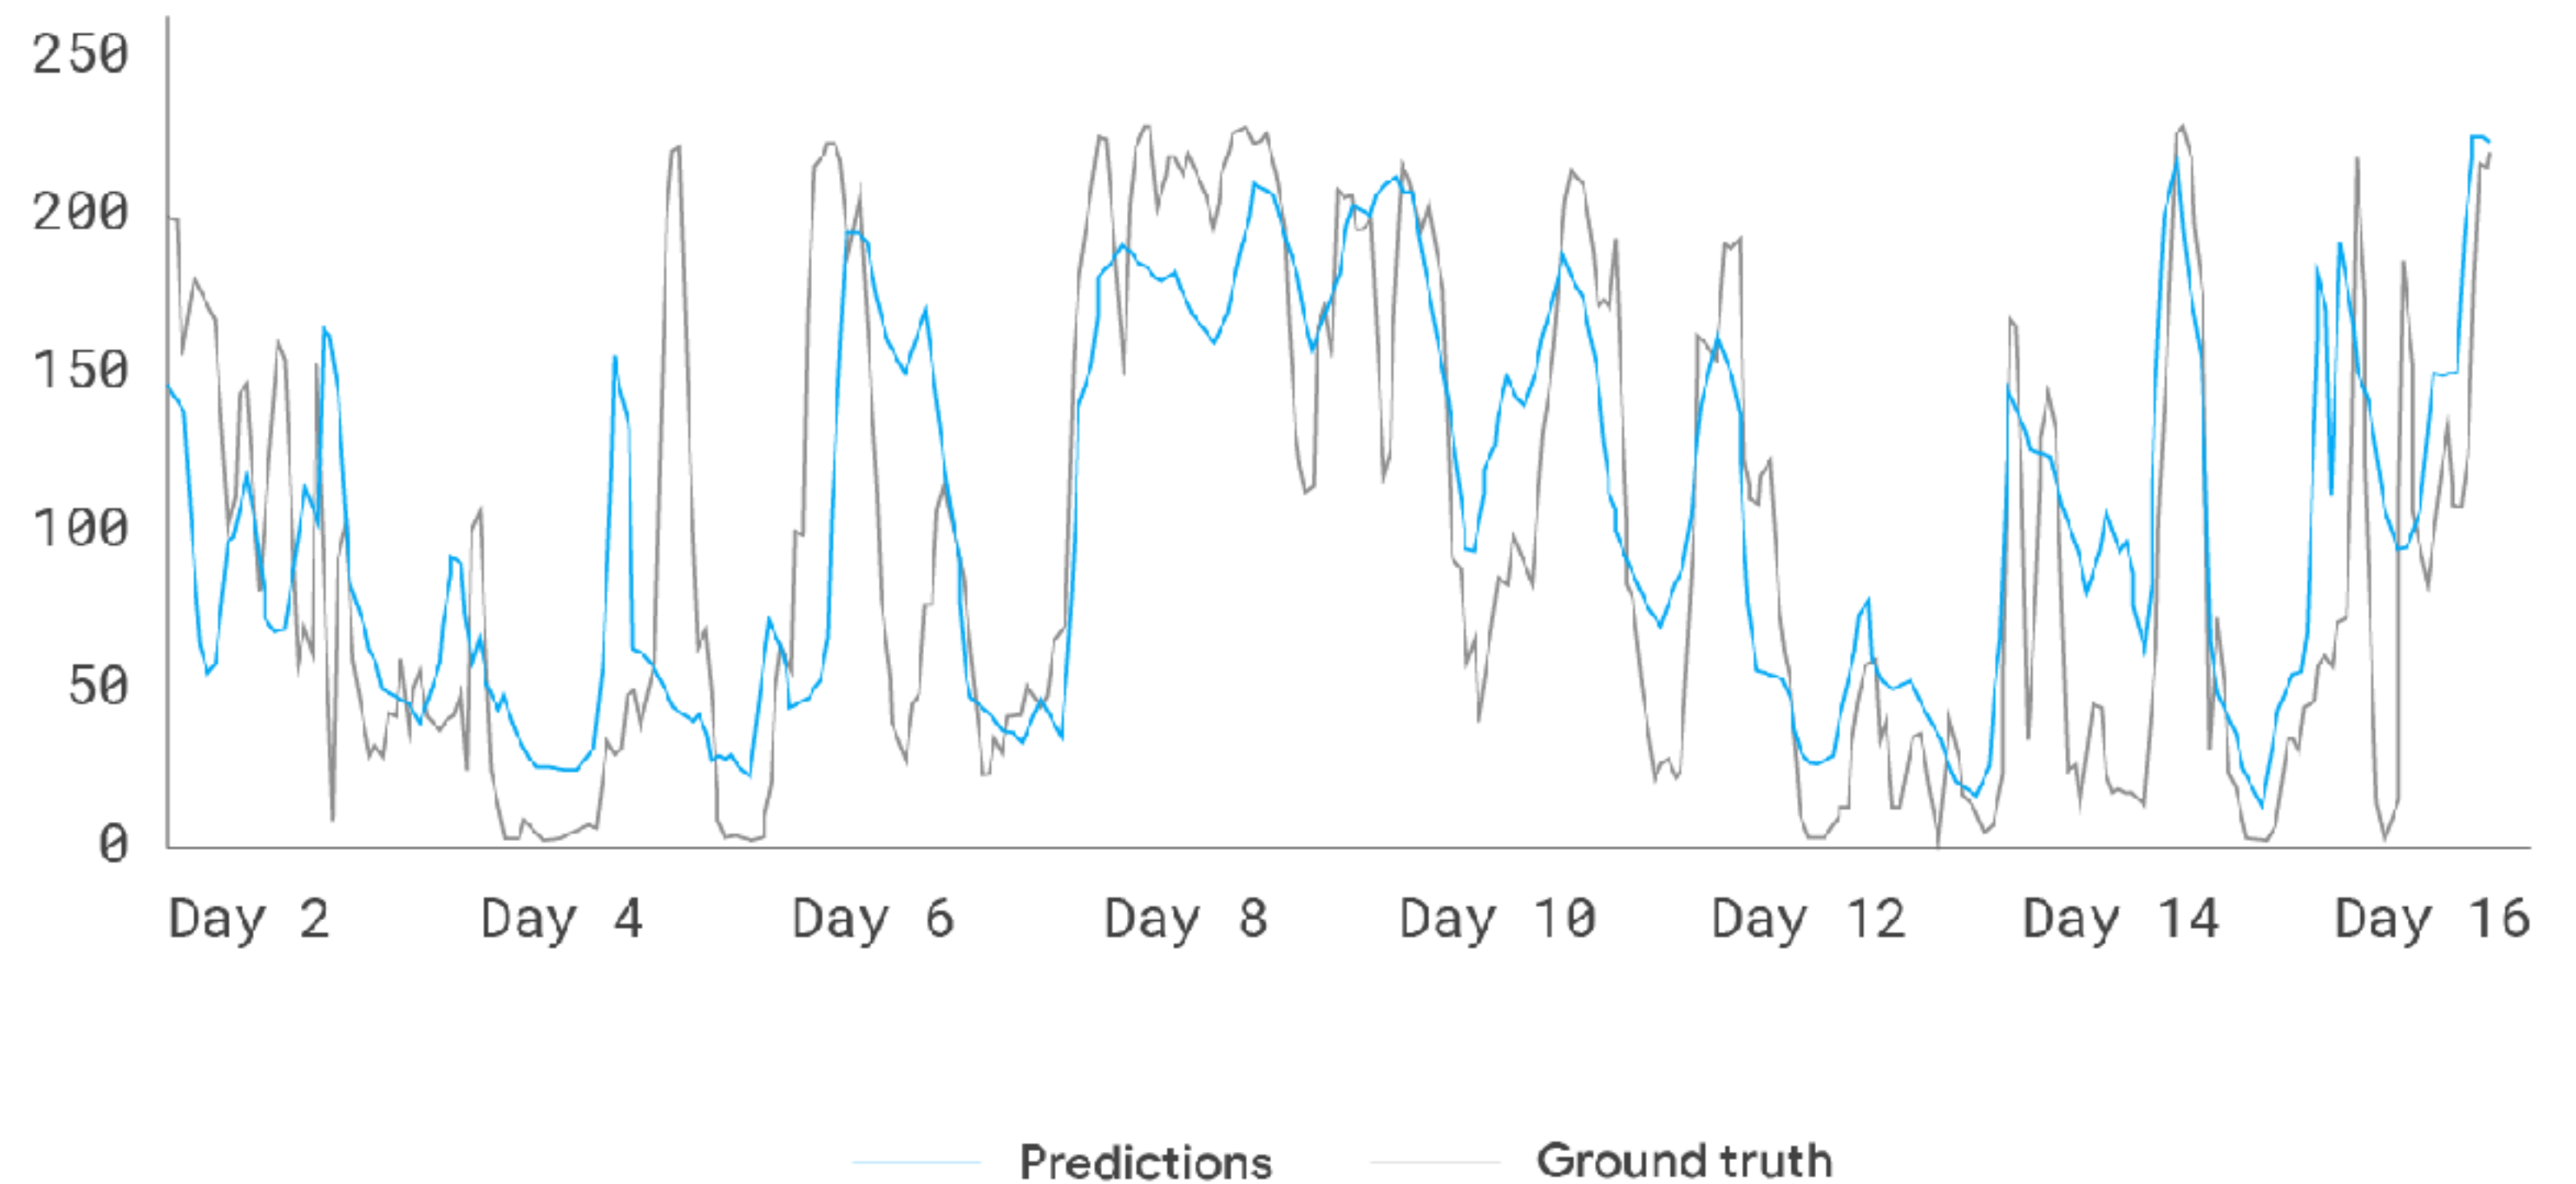
- Global numerical weather forecasts
- Local weather observations



Outputs

Future wind power output
(36 hours in advance)

Wind Power: Predicted Output v Ground Truth



20%

increase in economic
value, compared to
baseline of no time-based
commitments to grid



Tackling Climate Change with Machine Learning

David Rolnick^{1*}, Priya L. Donti², Lynn H. Kaack³, Kelly Kochanski⁴, Alexandre Lacoste⁵, Kris Sankaran^{6,7}, Andrew Slavin Ross⁸, Nikola Milojevic-Dupont^{9,10}, Natasha Jaques¹¹, Anna Waldman-Brown¹¹, Alexandra Luccioni^{6,7}, Tegan Maharaj^{6,7}, Evan D. Sherwin², S. Karthik Mukkavilli^{6,7}, Konrad P. Kording¹, Carla Gomes¹², Andrew Y. Ng¹³, Demis Hassabis¹⁴, John C. Platt¹⁵, Felix Creutzig^{9,10}, Jennifer Chayes¹⁶, Yoshua Bengio^{6,7}

¹University of Pennsylvania, ²Carnegie Mellon University, ³ETH Zürich, ⁴University of Colorado Boulder,

⁵Element AI, ⁶Mila, ⁷Université de Montréal, ⁸Harvard University,

⁹Mercator Research Institute on Global Commons and Climate Change, ¹⁰Technische Universität Berlin,

¹¹Massachusetts Institute of Technology, ¹²Cornell University, ¹³Stanford University,

¹⁴DeepMind, ¹⁵Google AI, ¹⁶Microsoft Research

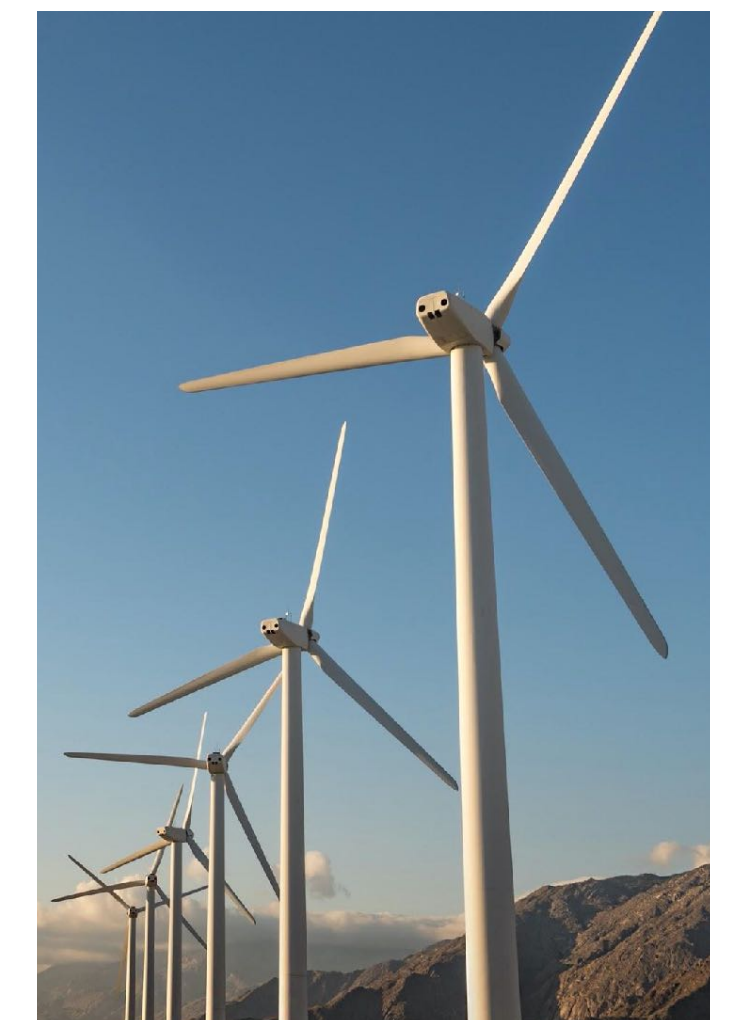
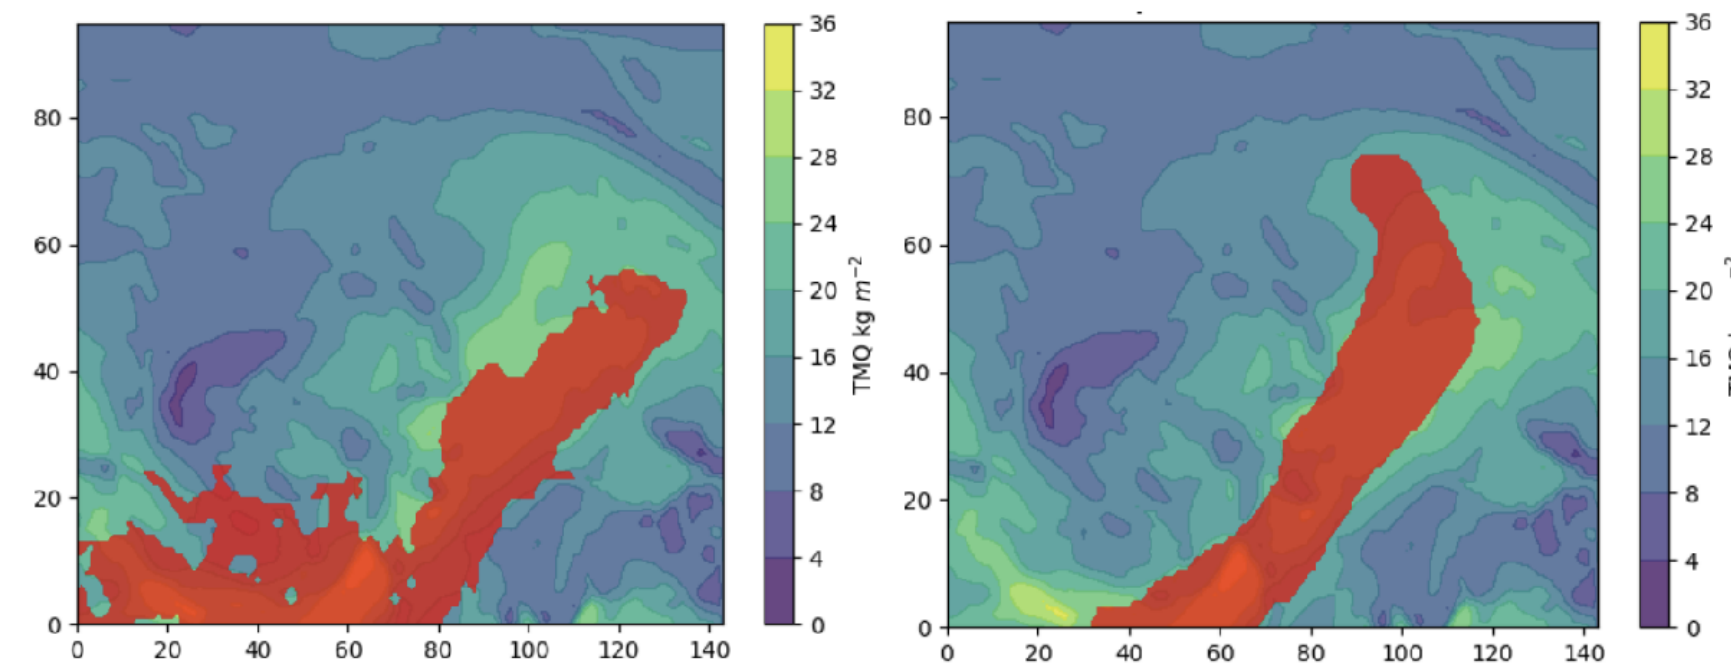
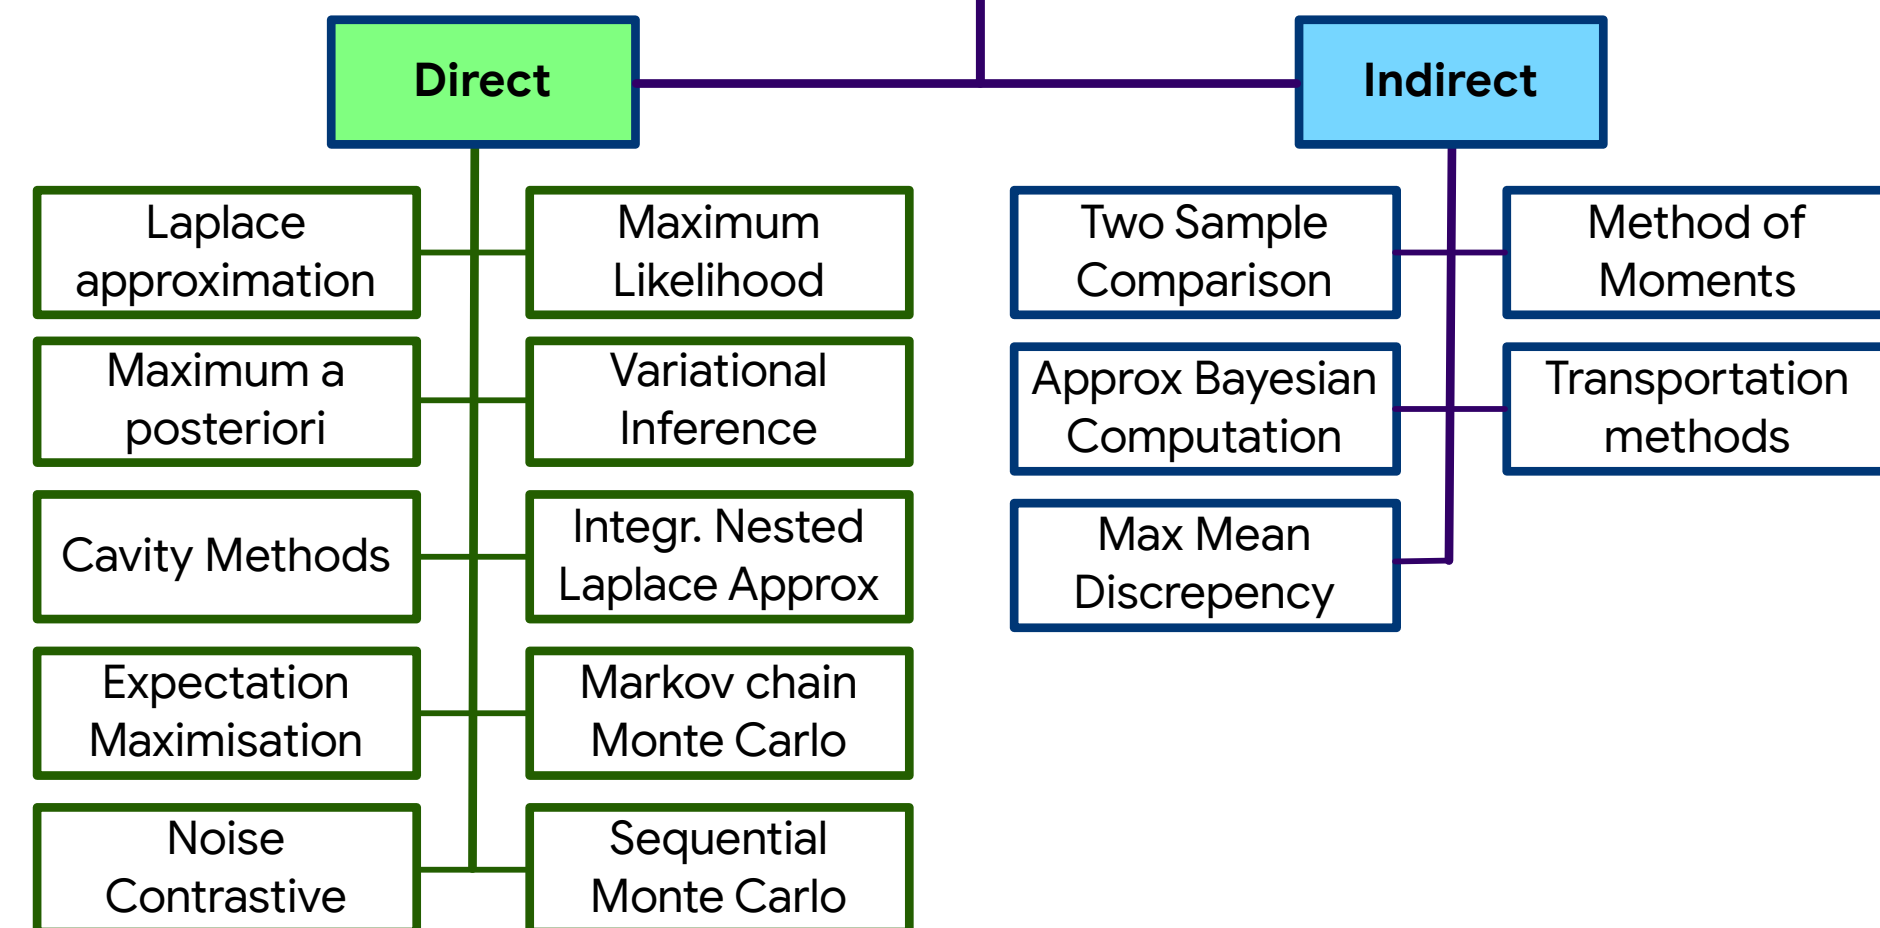
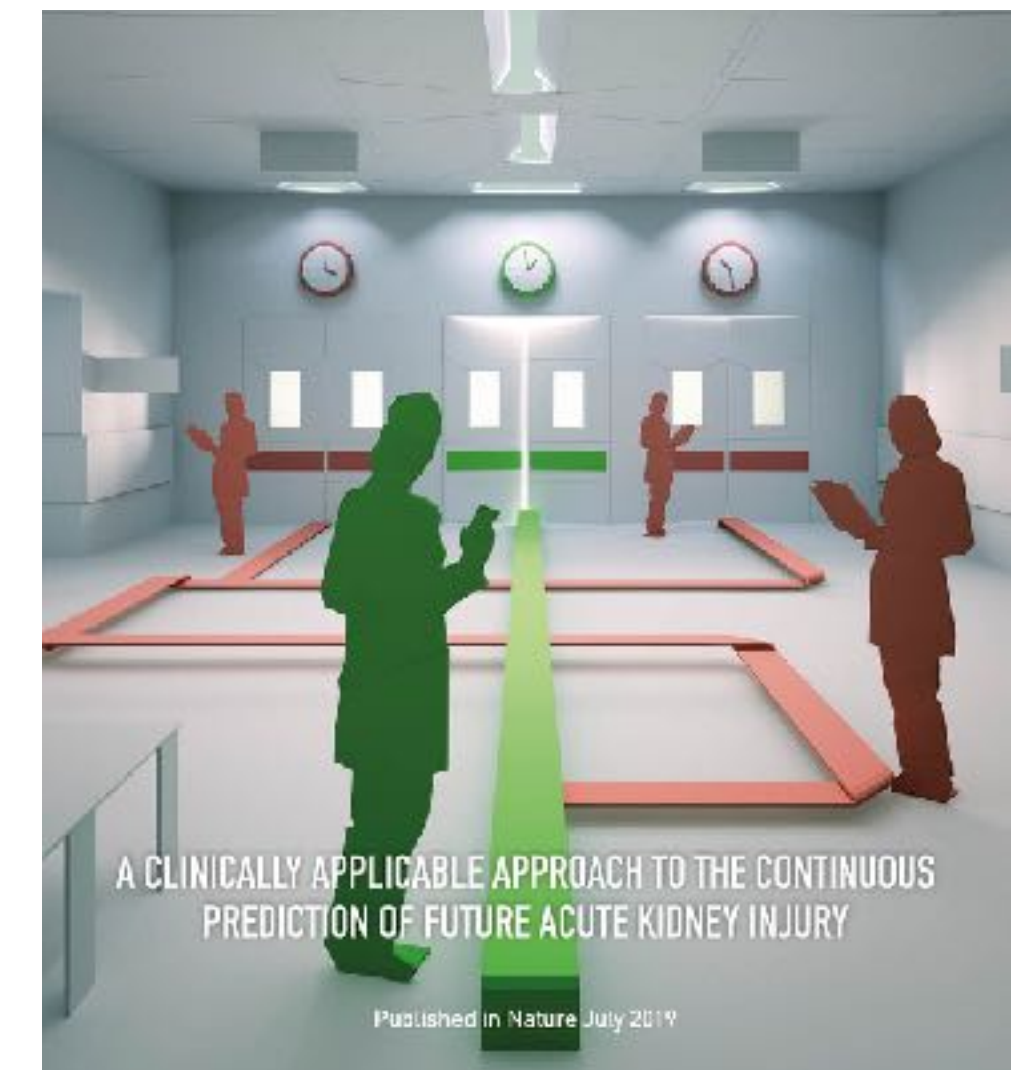
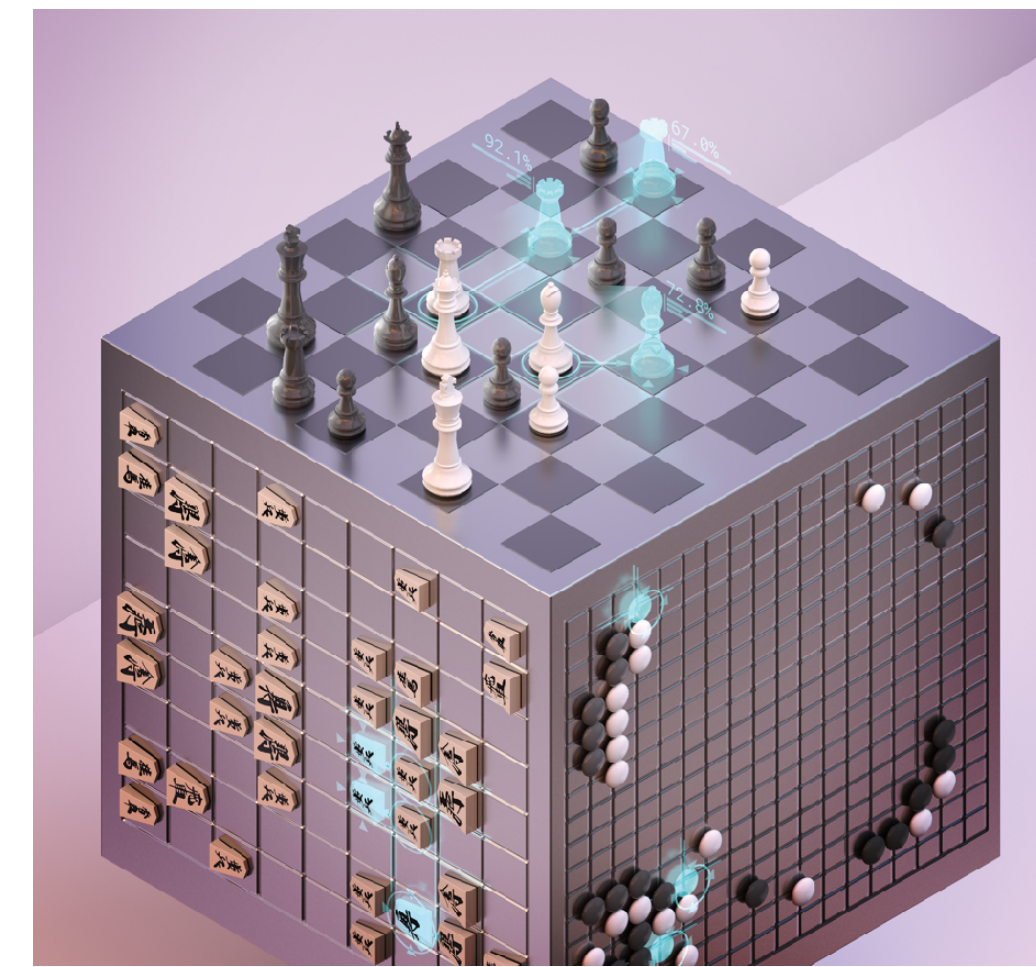
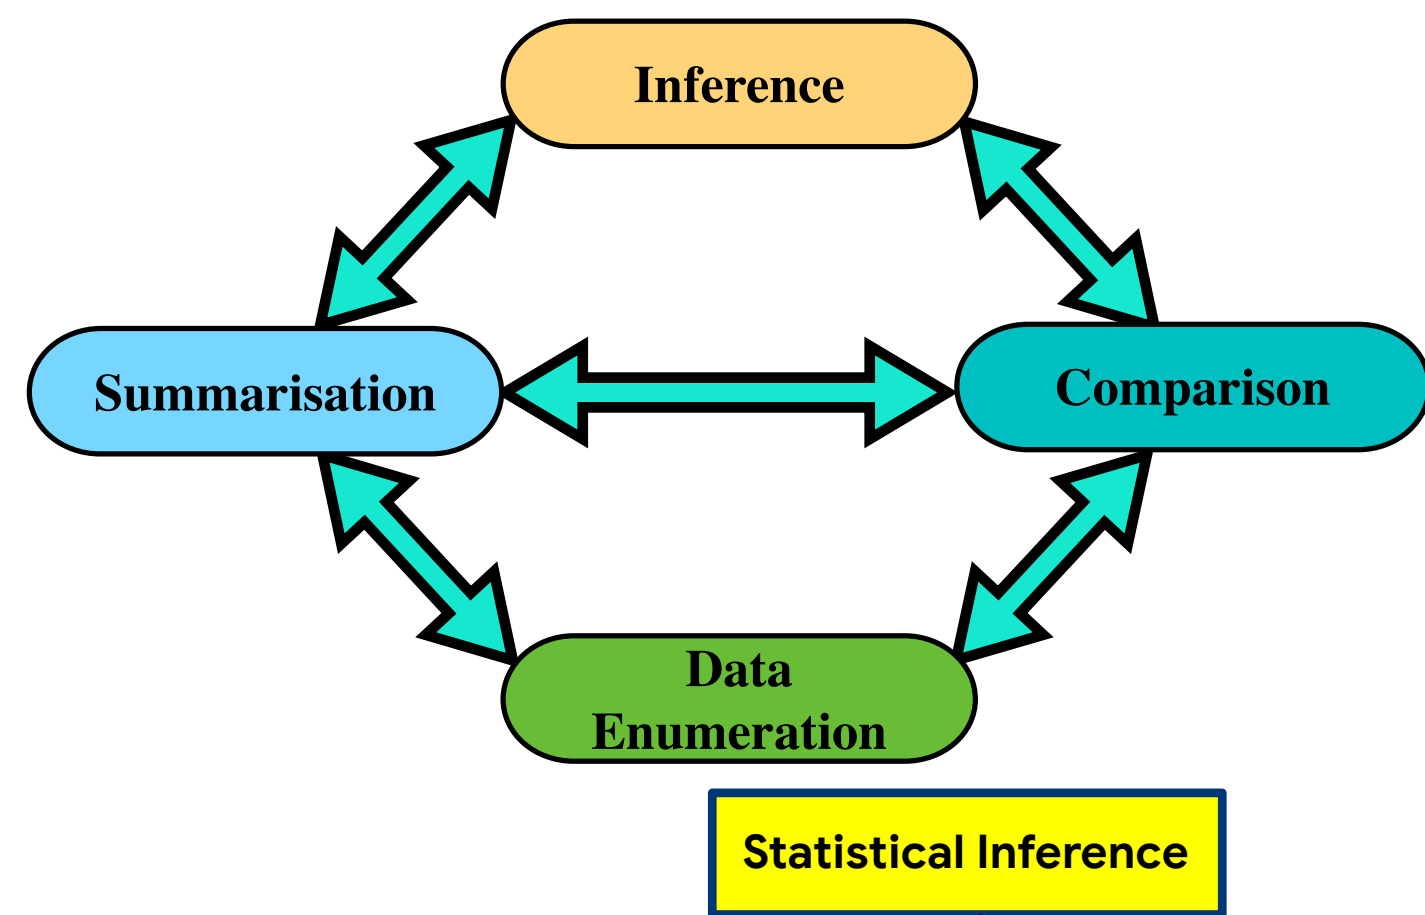
Abstract

Climate change is one of the greatest challenges facing humanity, and we, as machine learning experts, may wonder how we can help. Here we describe how machine learning can be a powerful tool in reducing greenhouse gas emissions and helping society adapt to a changing climate. From smart grids to disaster management, we identify high impact problems where existing gaps can be filled by machine learning, in collaboration with other fields. Our recommendations encompass exciting research ques-

CHAPTER 4

Climate Informatics

Claire Monteleoni, Gavin A. Schmidt,
Francis Alexander, Alexandru Niculescu-Mizil,
Karsten Steinhaeuser, Michael Tippett,
Arindam Banerjee, M. Benno Blumenthal,
Auroop R. Ganguly, Jason E. Smerdon,
and Marco Tedesco



References

- Mohamed, S., and Balaji L.. "Learning in implicit generative models." arXiv preprint arXiv:1610.03483(2016).
- Mohamed, S., and Rezende D. J. . "Variational information maximisation for intrinsically motivated reinforcement learning." Advances in neural information processing systems. 2015.
- Mohamed, S., Rosca, M., Figurnov, M., & Mnih, A. (2019). Monte Carlo Gradient Estimation in Machine Learning. arXiv preprint arXiv:1906.10652.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Petersen, S. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... & Lillicrap, T. (2017). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. arXiv preprint arXiv:1712.01815.
- Tomašev, Nenad, et al. "A clinically applicable approach to continuous prediction of future acute kidney injury." Nature 572.7767 (2019): 116.
- Bansal, A., Ma, S., Ramanan, D., & Sheikh, Y. (2018). Recycle-gan: Unsupervised video retargeting. In Proceedings of the European Conference on Computer Vision (ECCV) (pp. 119-135).
- Mudigonda, M., Kim, S., Mahesh, A., Kahou, S., Kashinath, K., Williams, D., ... & Prabhat, M. (2017). Segmenting and tracking extreme climate events using neural networks. In Deep Learning for Physical Sciences (DLPS) Workshop, held with NIPS Conference.
- de Bézenac, E., Pajot, A., & Gallinari, P. Towards a Hybrid Approach to Physical Process Modeling.
- J. Kelly (2019). Open Climate Fix. <https://openclimatefix.github.io/>
- C. Gamble and J. Gao. Safety-first AI for autonomous data centre cooling and industrial control
- <https://deepmind.com/blog/article/safety-first-ai-autonomous-data-centre-cooling-and-industrial-control>
- C. Elkin and S. Witherspoon (2019). Machine learning can boost the value of wind energy. <https://deepmind.com/blog/article/machine-learning-can-boost-value-wind-energy>

Machine Learning for Environmental Grand Challenges

With thanks to colleagues and the
work of many others referenced here
from our ML community.

Shakir Mohamed

Research Scientist, DeepMind