Active Feature Acquisition with Generative Surrogate Models

Yang Li, Junier Oliva

UNC DEPARTMENT OF COMPUTER SCIENCE

LUPA LAB

lupalab.com

Motivation

typical machine learning

Patient

collect features

Age: 47
Weight: 160
Heart Rate: 102

Blood Test:

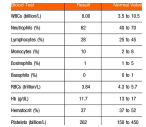

X-Ray:

MRI:

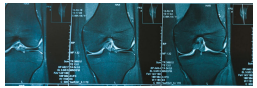

Biopsy:

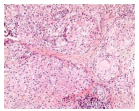

...

predict

what disease?

Motivation

human doctor

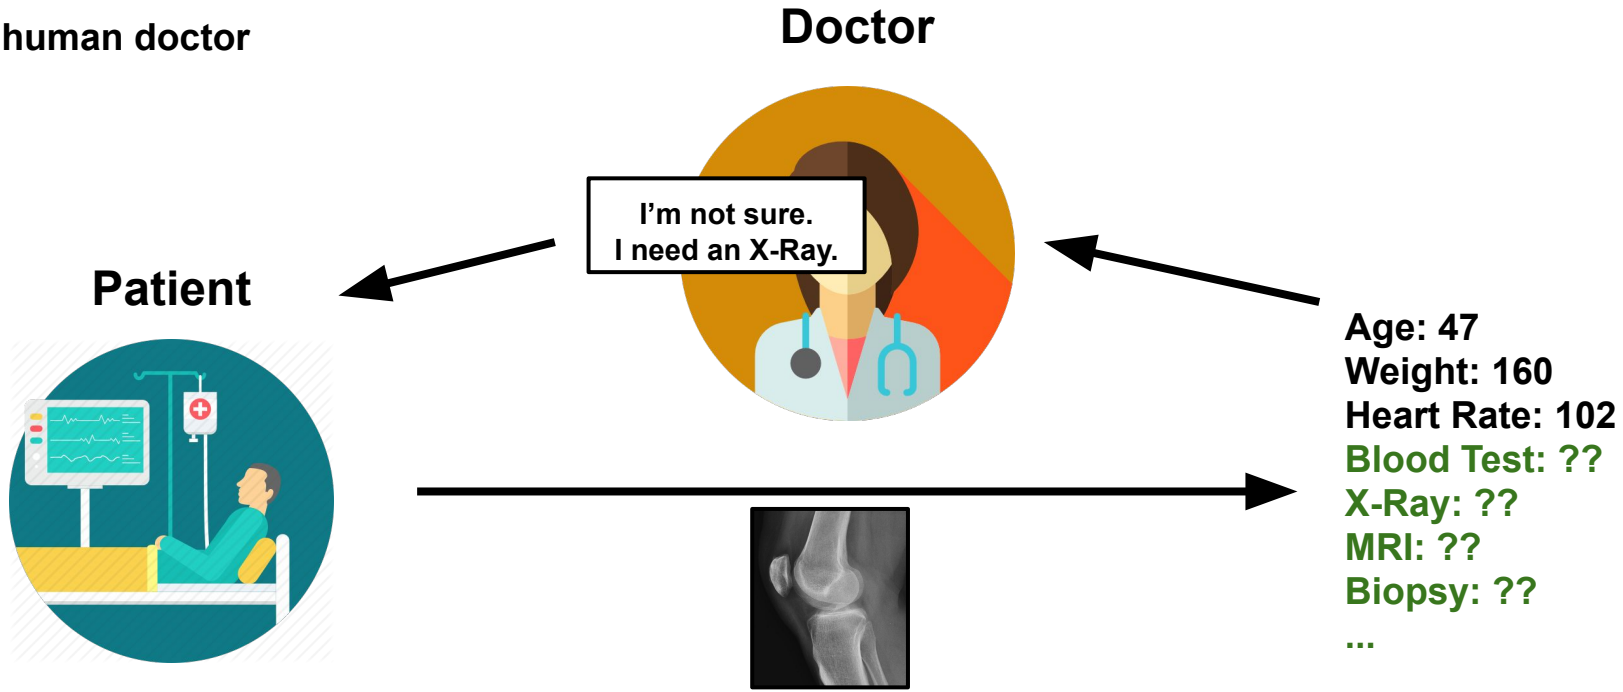

Motivation

human doctor

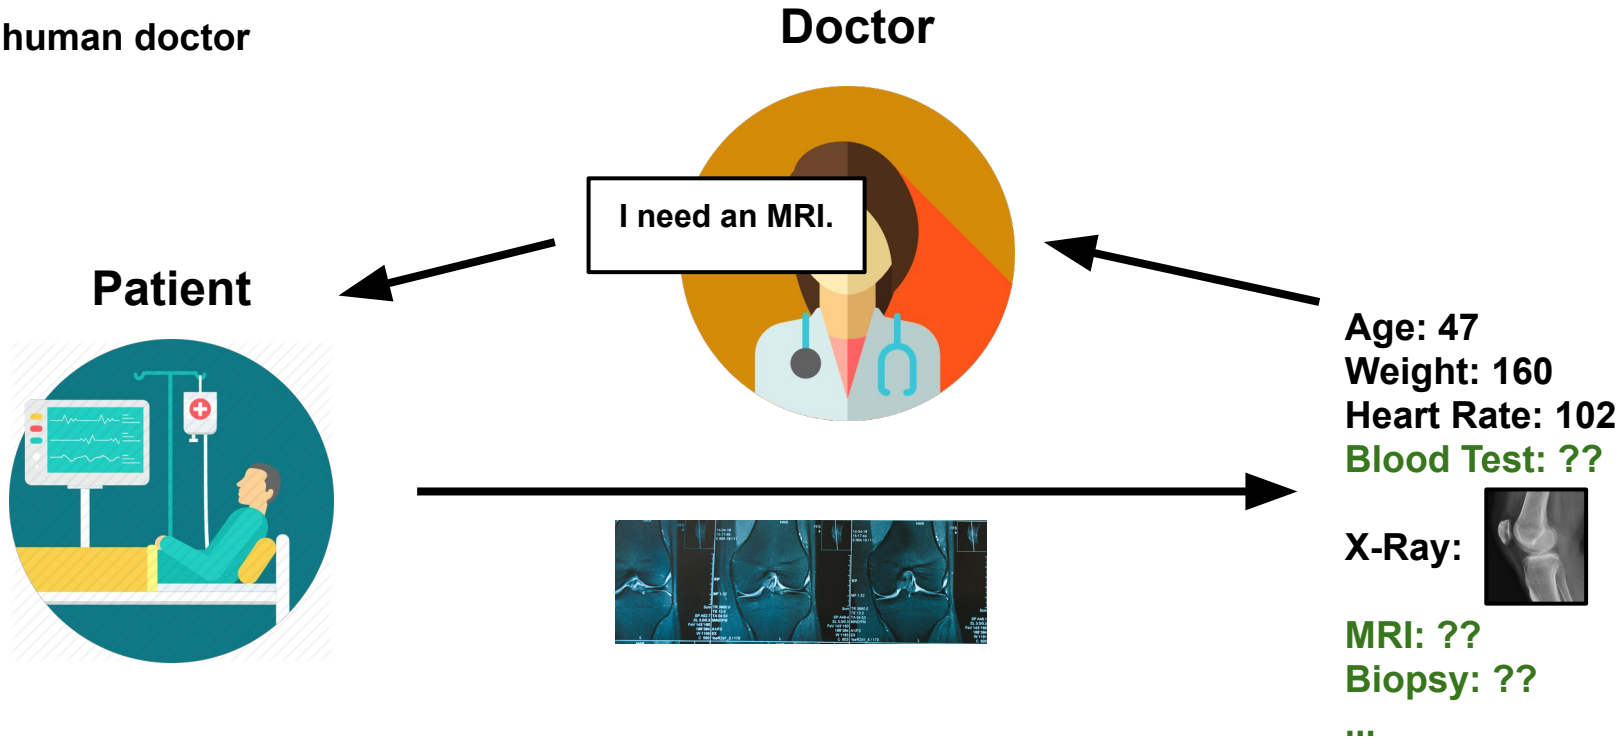

Motivation

human doctor

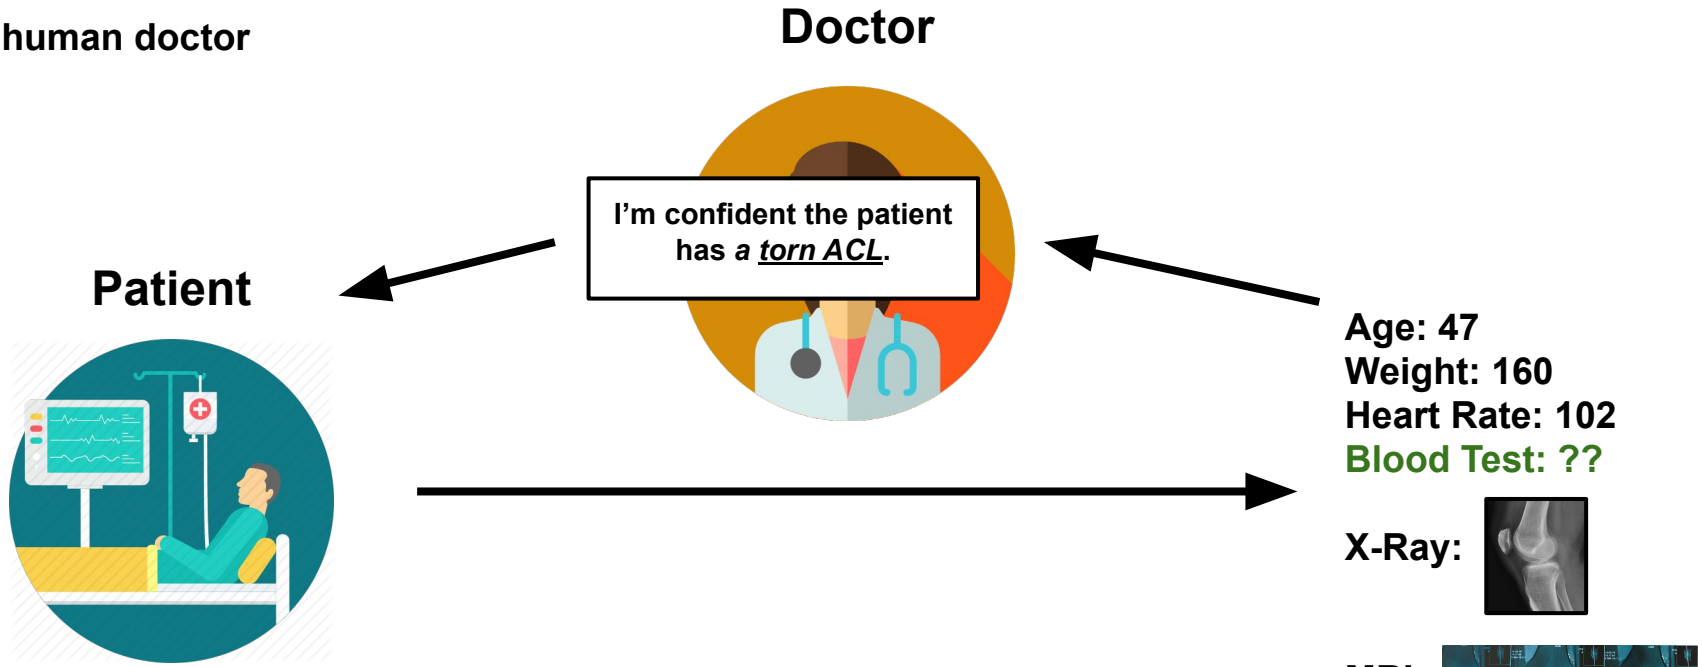

Active Feature Acquisition (AFA)

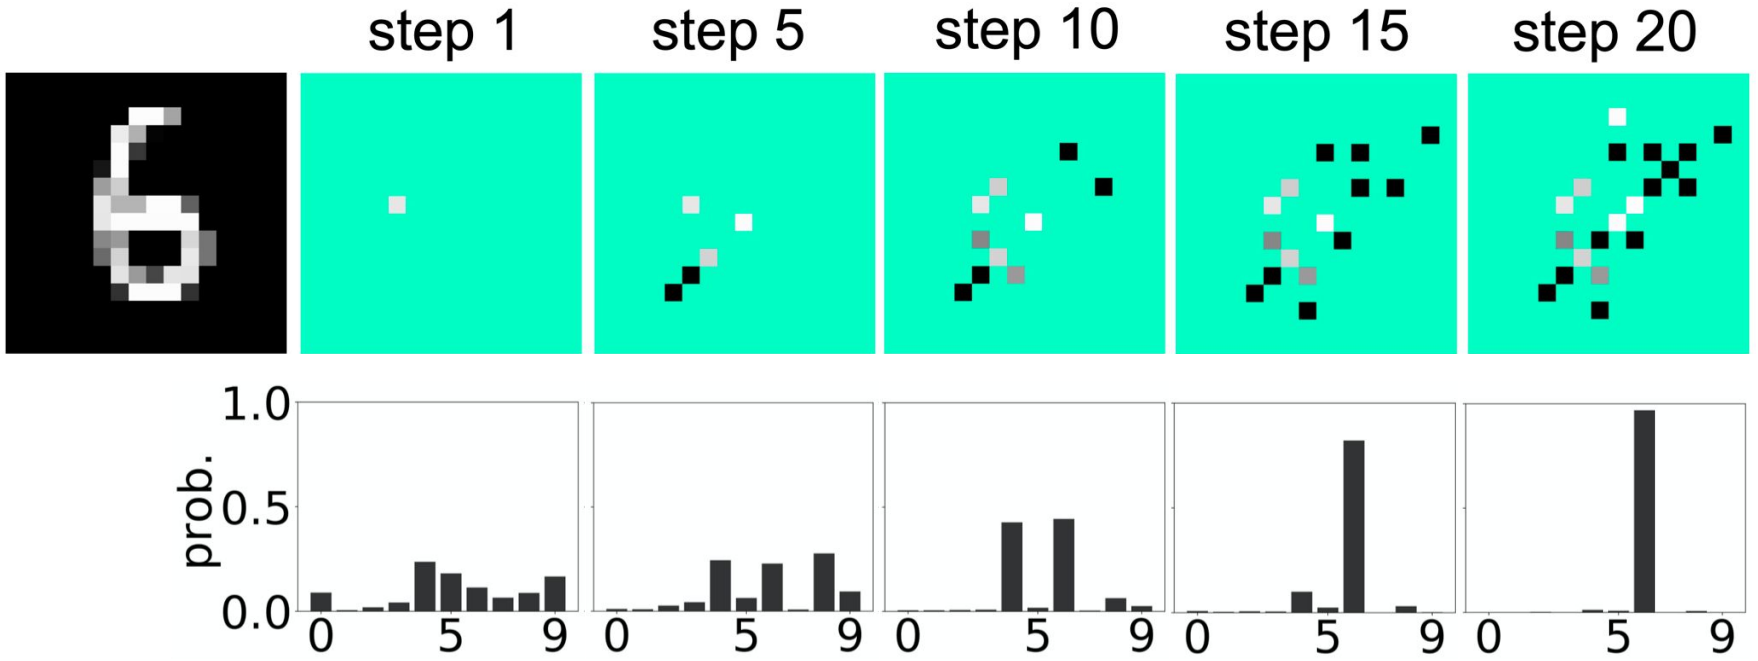

Active Feature Acquisition (AFA)

Informing Feature Acquisition with Conditionals

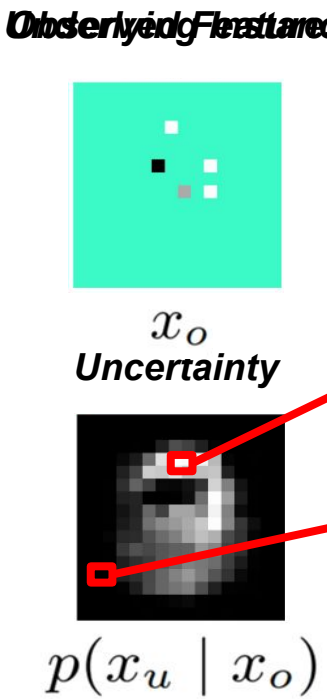

We can use conditionals to guide the acquisition of features by avoiding features that:

we are very certain about

If we are sure of the value of a missing feature, then it's redundant to observe it.

will not help the prediction regardless of the true value

If the feature will not inform our prediction then there is no point in observing it.

Active Feature Acquisition (AFA)

Arbitrary Conditionals

$$p(x_u \mid x_o) \qquad u, o \subseteq \{1, \ldots, d\}$$

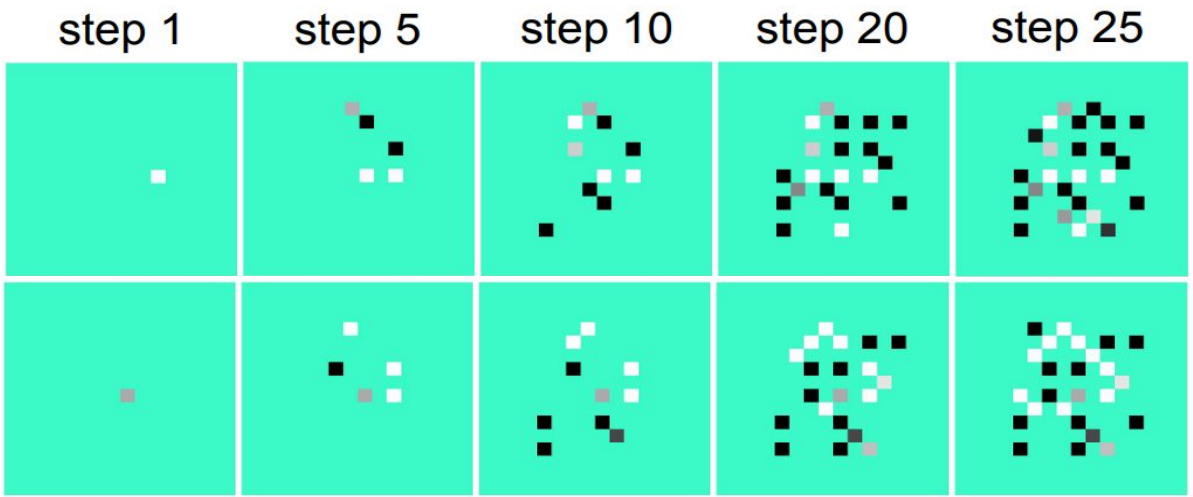

AFA as MDP

State: x_o

Actions:

Acquire $j \in \{1, \ldots, d\} \setminus o$

or Predict $Y = y$

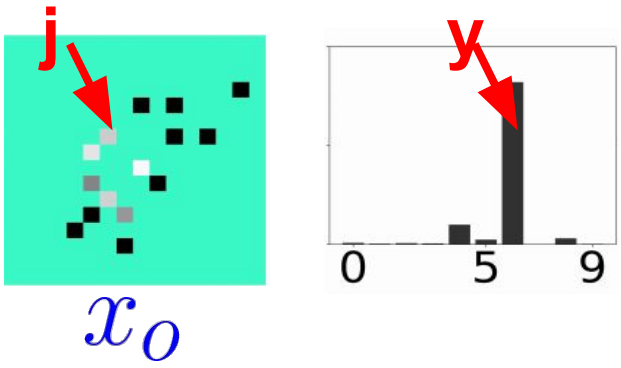

Rewards:

Cost of obtaining feature *j*

or cost of incorrect prediction (reward for correct)

AFA with Model-based RL

GSM $p(y, x_u \mid x_o)$

Agent

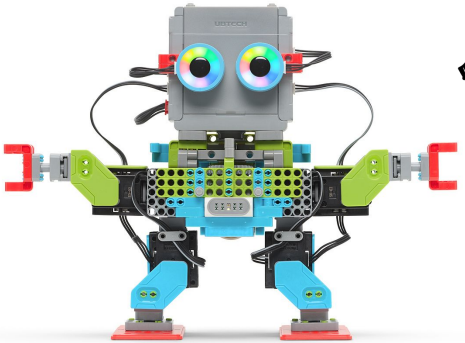

Auxiliary Information

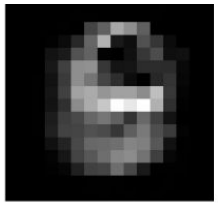

Current Imputation & Uncertainties

Expected **Information Gain**

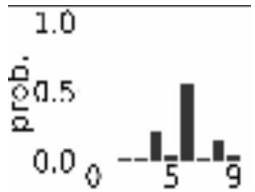

Auxiliary Classification

Action

acquire feature *i*

cost of acquiring feature *i*

Reward

State

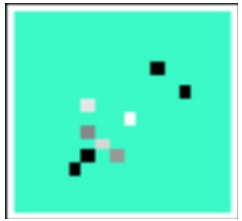

x_o

Auxiliary Reward

\+ $H(y \mid x_o) - H(y \mid x_o, x_i)$. **Information Gain**

AFA for MNIST Classification

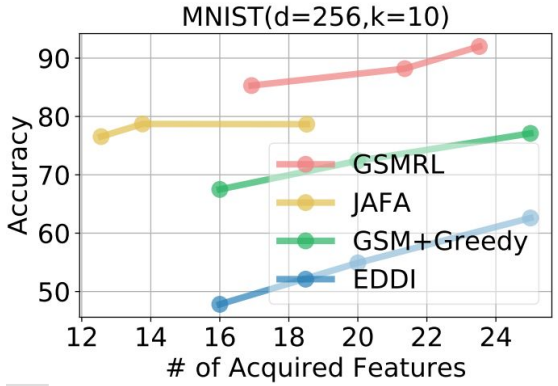

- **EDDI**: Efficient Dynamic Discovery of High-Value Information with Partial VAE (ICML 2019)
  - Greedy acquisition based on VAEs
- **JAFA**: Joint Active Feature Acquisition and Classification with Variable-Size Set Encoding (NeurIPS 2018)
  - Plain reinforcement learning optimization with Q-learning

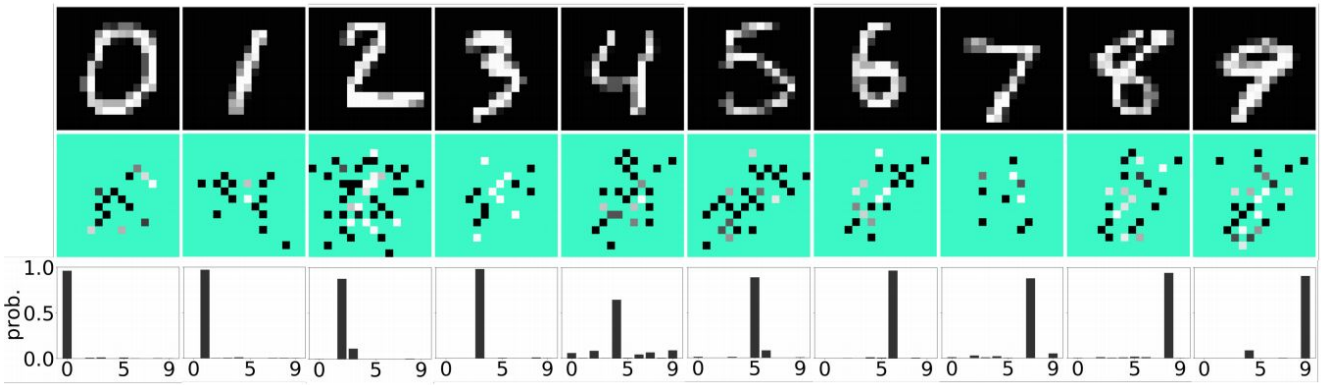

AFA for MNIST Classification

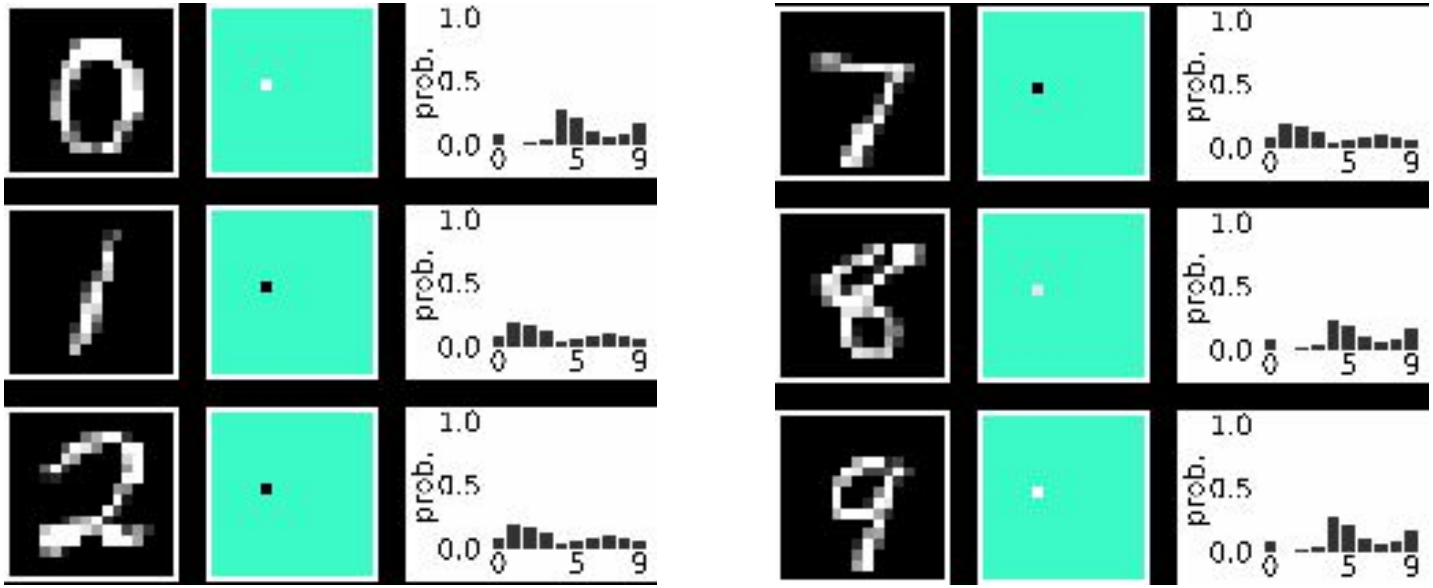

Active Instance Recognition (AIR)

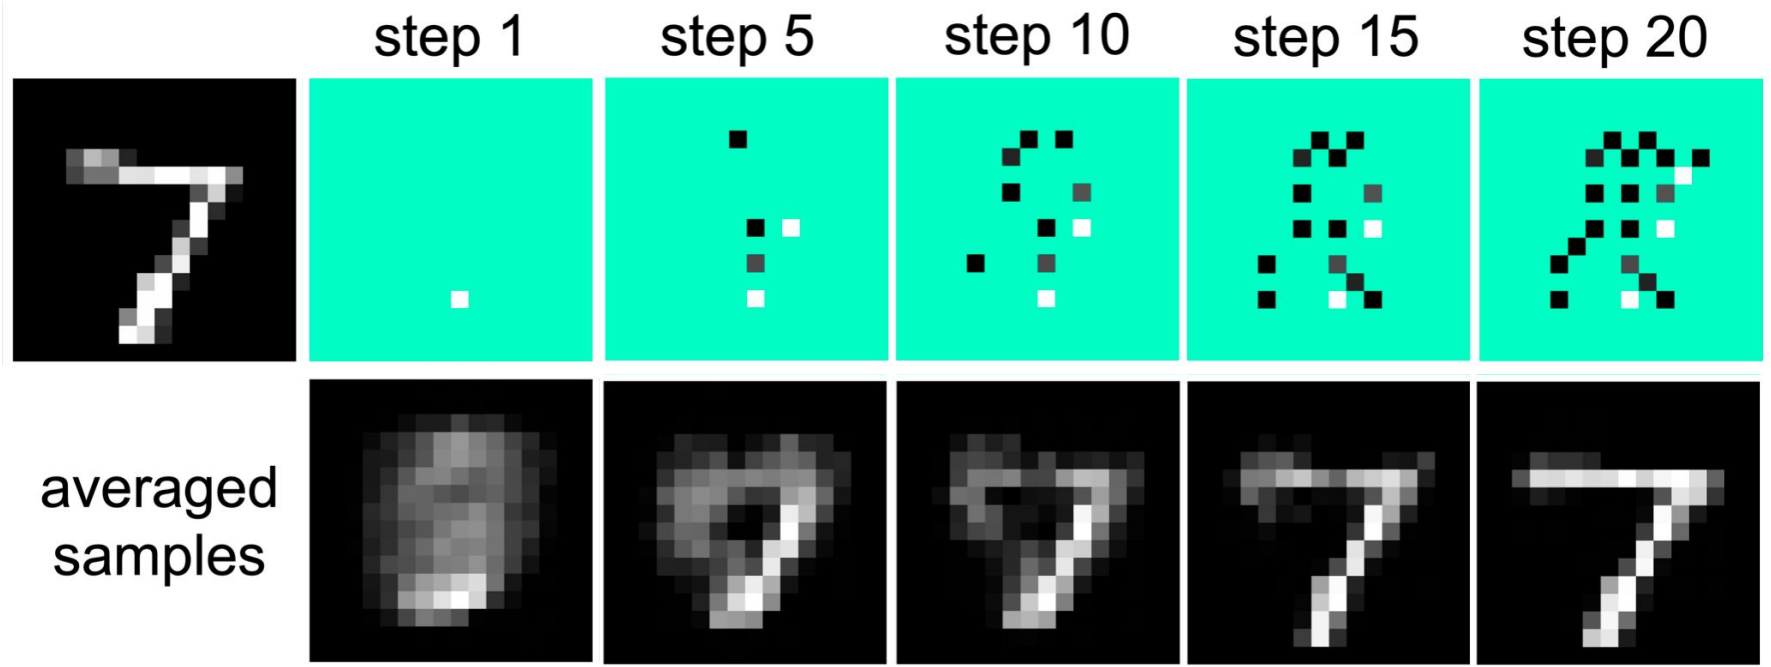

AIR for MNIST Reconstruction

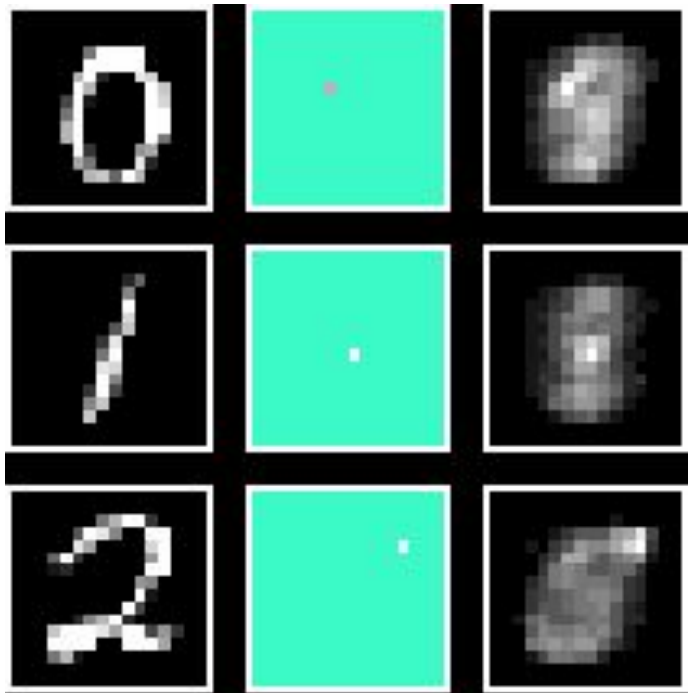

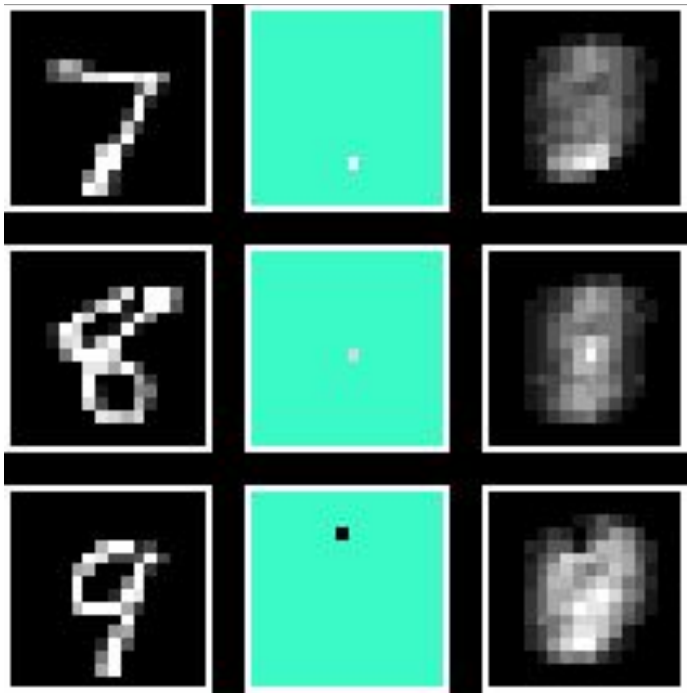

Summary

- We reformulates the MDP that underlies the AFA problem as a generative modeling task and optimizes a policy via a model-based approach.

- The GSM is leveraged to provide intermediate rewards and auxiliary information to aid the agent navigate a complicated high-dimensional action space and sparse rewards.

- We extend AFA in a task we coin active instance recognition (AIR) for the unsupervised case where the target variables are the unobserved features themselves and the goal is to collect information for a particular instance in a cost-efficient way.

Thanks!

UNC
DEPARTMENT OF
COMPUTER SCIENCE

LUPA LAB

lupalab.com