Automated Experiments, Machine Learning and Reinforcement Learning for tuning and control applications in microscopy

Rama Vasudevan
Center for Nanophase Materials Sciences,
Oak Ridge National Laboratory

With K. Kelley, S. Jesse, S. V. Kalinin, M. Ziatdinov

AI/ML for Particle Accelerator, X-Ray Beamlines and Electron Microscopy

November 1, 2021

Background

OAK RIDGE National Laboratory | CENTER FOR NANOPHASE MATERIALS SCIENCES

CNMS is a national user facility with a mission to advance nanoscience

About CNMS:

- **Unlike many user facilities, you don't need to have samples to apply for time**
- **Two calls per year for continuous access; anytime for short-term projects**
- **Simple 2-page proposal**
- **Free access to laboratories, equipment and expertise if you agree to publish**
- **Proposal deadlines: early May and mid-October**
- **Joint proposals with neutron sources (SNS, HFIR)**

Research areas:

- **Synthesis –** 2D, precision synthesis, selective deuteration
- **Nanofabrication –** direct-write, microfluidics, cleanroom
- **Advanced Microscopy –** AFM, STM, aberration-corrected TEM/STEM, atom-probe tomography
- **Functional Characterization –** laser spectroscopy, transport, magnetism, electromechanics
- **Theory and Modelling –** including gateway to leadership-class high performance computing

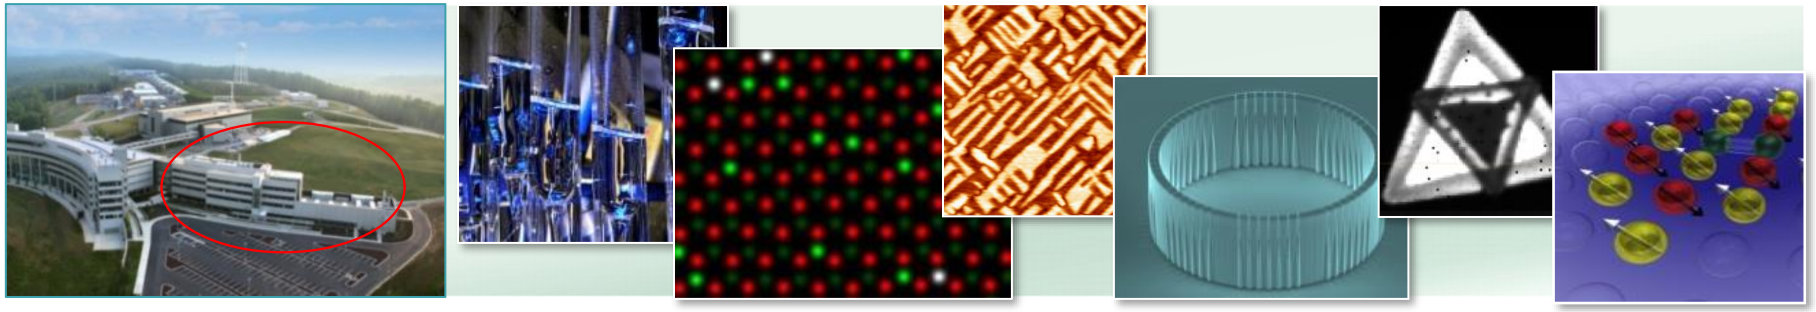

CNMS is a Nanoscale Science Research Center supported by the U.S. Department of Energy, Office of Science, Scientific User Facilities Division

Introduction: ‘Smart’ experiments

Automated chemical synthesis with feedback

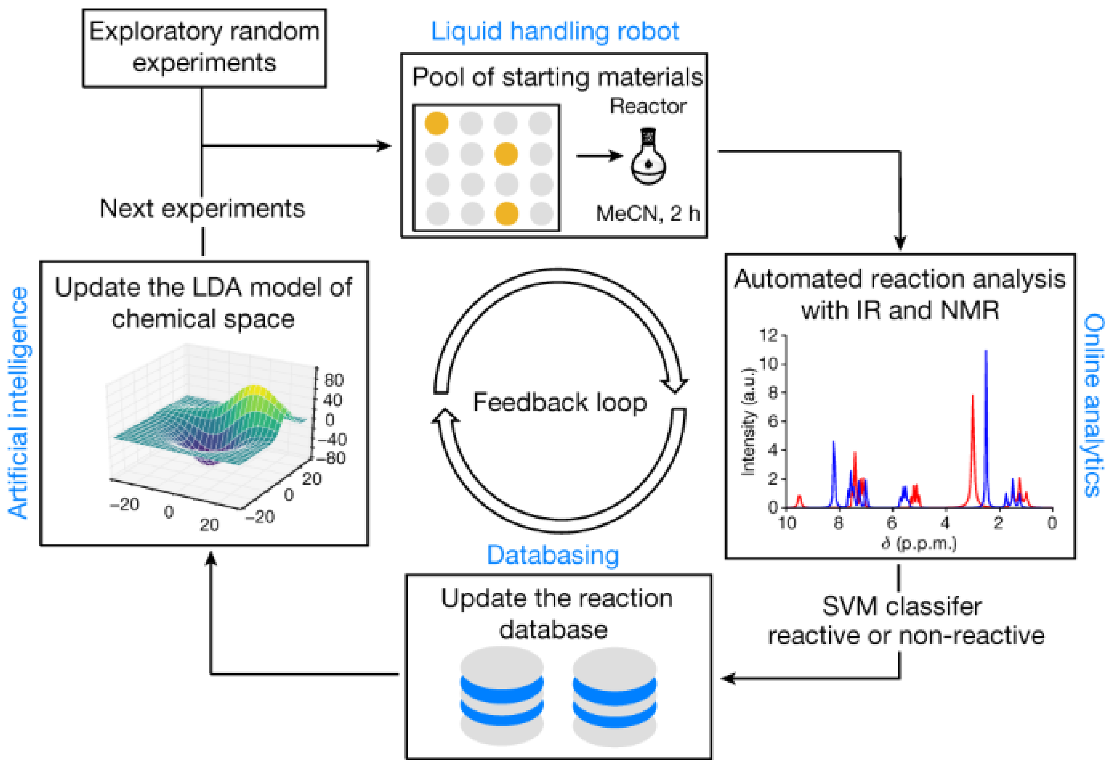

Granda et al. Nature **559**, 377 (2018)

‘Phase Mapping’ at a synchrotron

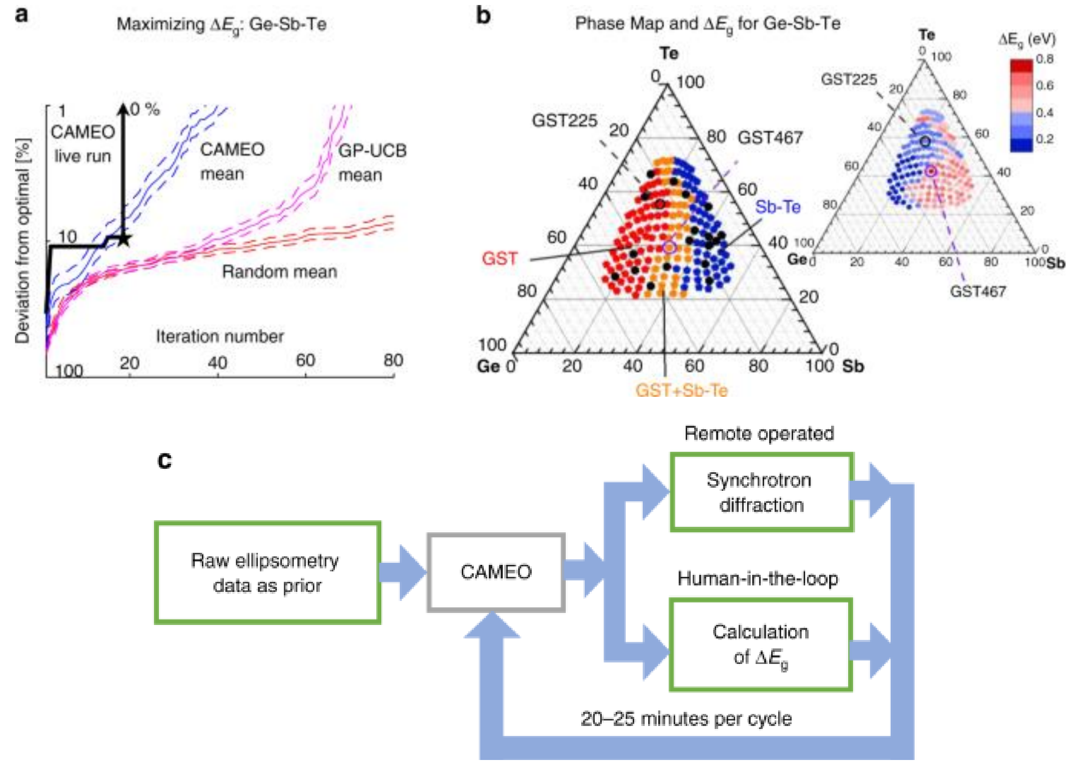

Kusne et al. Nat. Comm. **11**, 5966 (2020)

Oak Ridge National Laboratory

Outlines

RL: Defect engineering

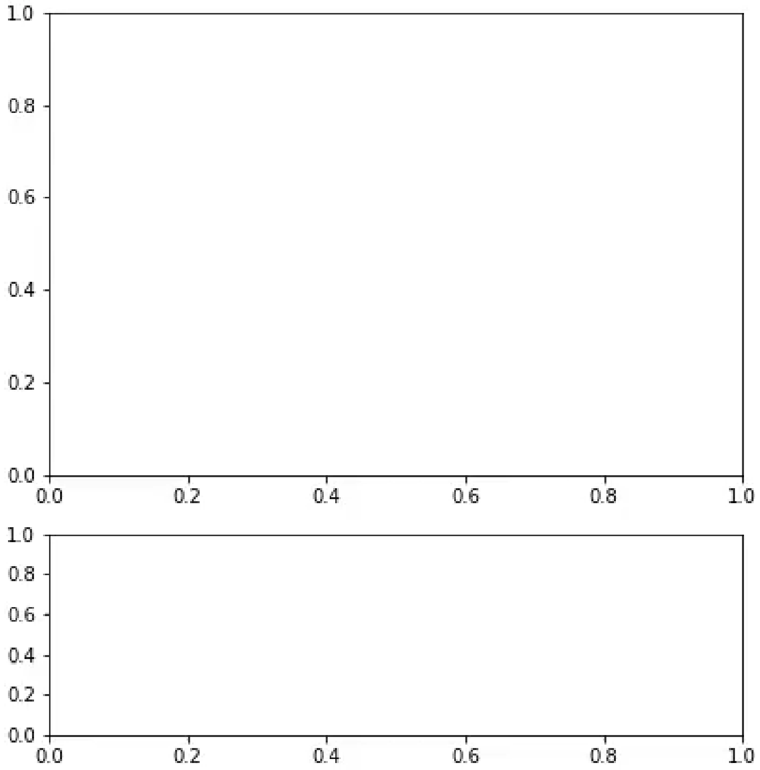

Automated Exp.: Manipulations

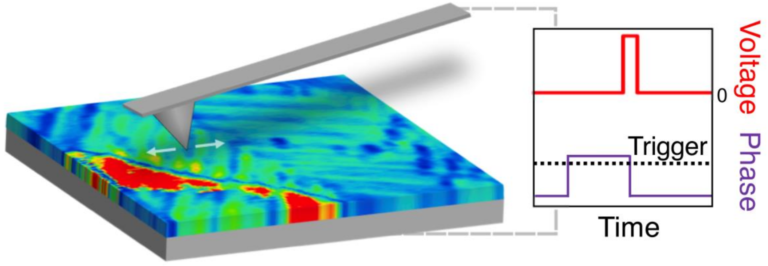

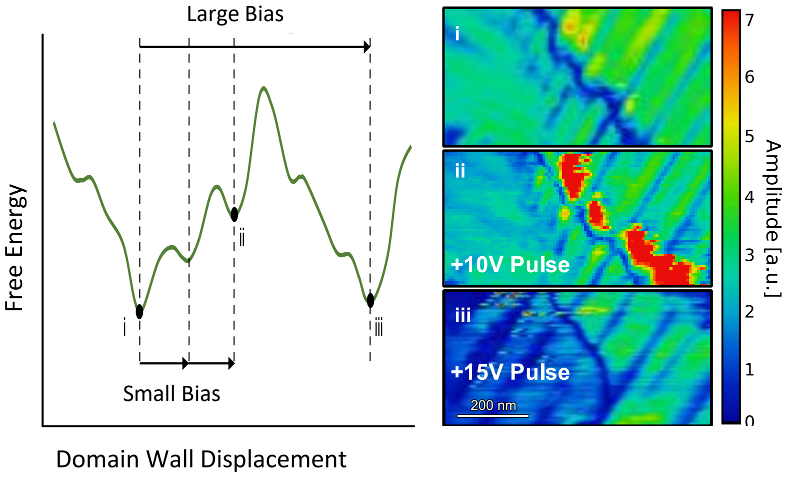

Adaptive Sampling

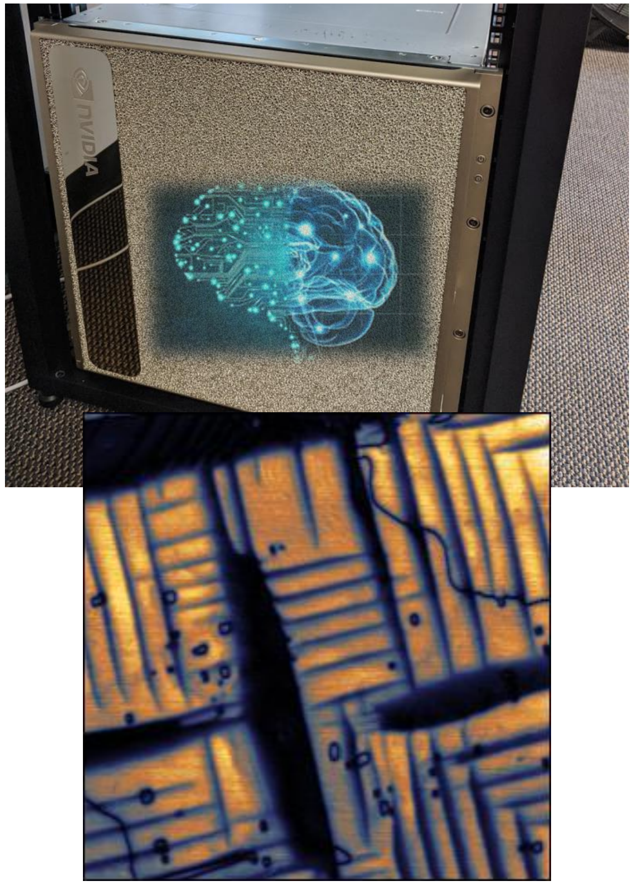

Outlines

RL: Defect engineering

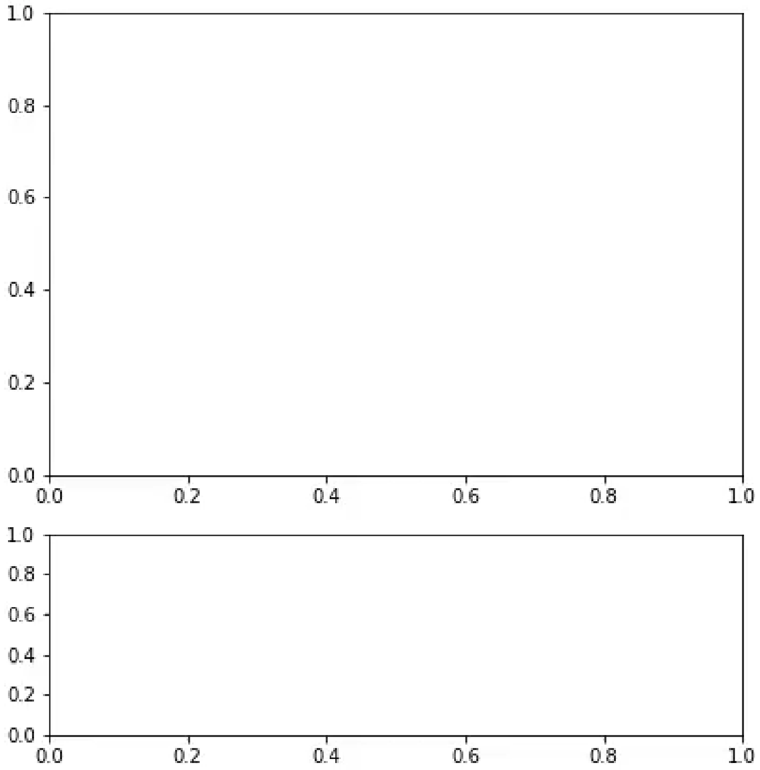

Automated Exp.: Manipulations

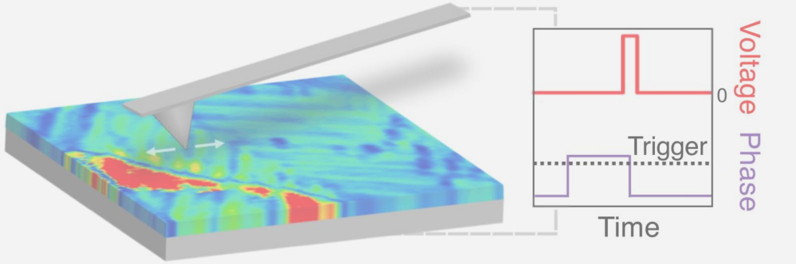

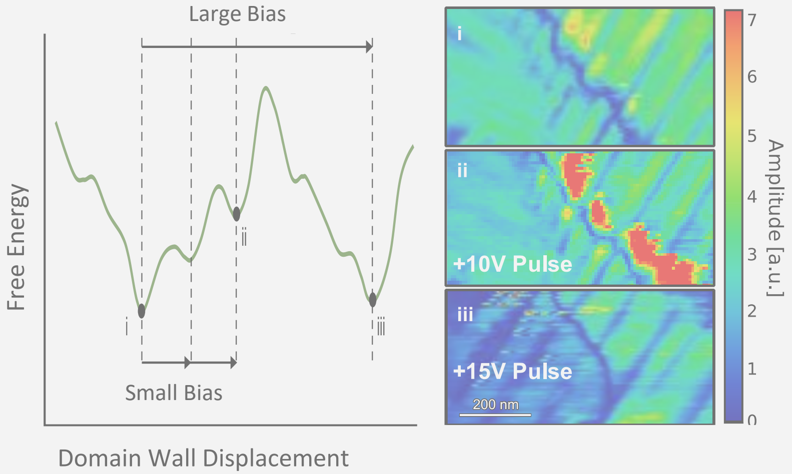

Adaptive Sampling

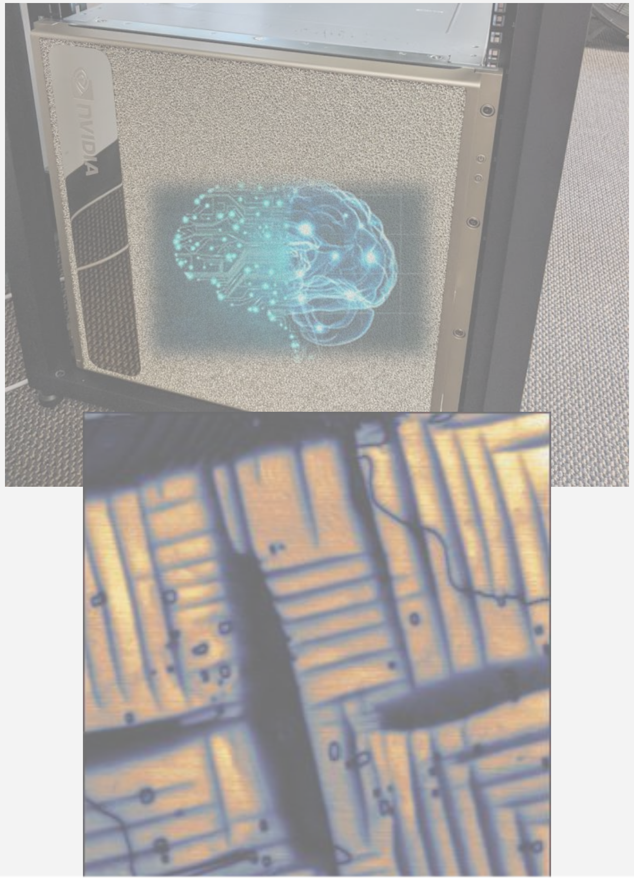

Reinforcement Learning Basics

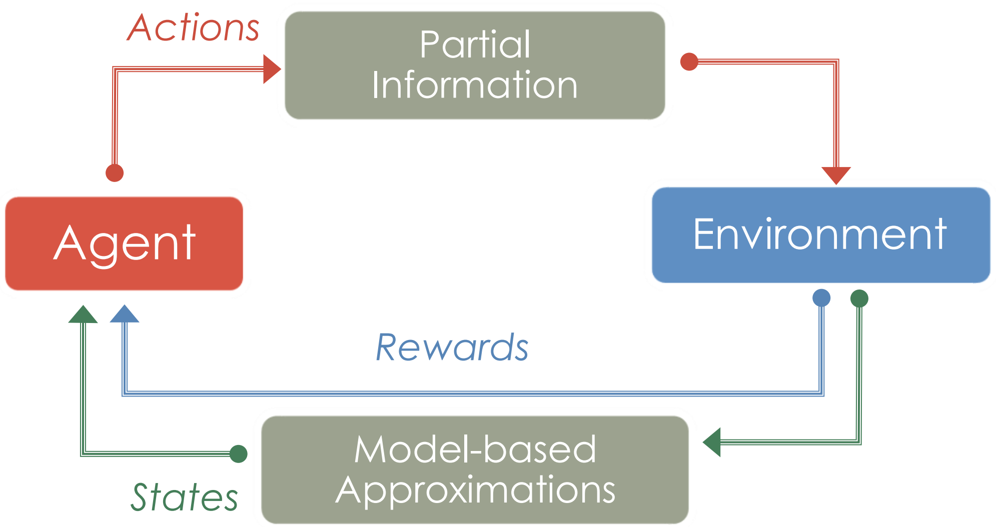

- RL is neither supervised nor unsupervised – it deals with optimal decision making in uncertain environments
- Variants of RL include on-policy learning and off-policy learning, and fully offline learning.

- We wish to learn stochastic policies that map states to actions to maximize some reward
- Two main types of RL: model-based, and model-free
- We can deal with continuous and discrete action spaces
- Policies are generated that aim to maximize expected future rewards emitted from the environment

Main ideas of Reinforcement Learning

- A policy defines how an actor behaves in a Markov Decision Process (MDP), and is defined as a distribution of actions over states:

$$\pi(a|s) = \mathbb{P}(A_t = a\,|S_t = s)$$

- We can sample the policy to obtain trajectories τ through the MDP

$\tau_k = (S_1, A_1, R_1, \ldots,)$

- The goal in RL is to solve the MDP to maximize the cumulative rewards. The policy is parametrized by parameters θ. So we can write the objective function as

$$J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta}\,[R(\tau)]$$

Where R() is the reward of the trajectory

Reinforcement Learning Intuition

- In RL, we are not using supervised or unsupervised machine learning. We don't know the 'correct' answer through supervision. So where to start? Answer: Trial random actions

Pong

Up-Down-Up-Down-Up-Up **Bad**

Down—Down-Up-Up-Down **Good**

Down—Down-Up-Up-Down **Good**

Up—Down-Up-Down-Down **Bad**

Supervised ML

maximize

$$\sum_i \log p(y_i|x_i)$$

Reinforcement Learning

maximize $\log p(y_i|x_i)$

maximize $-1*\log p(y_i|x_i)$

Discrete Landau Model

2D Discrete Landau Model

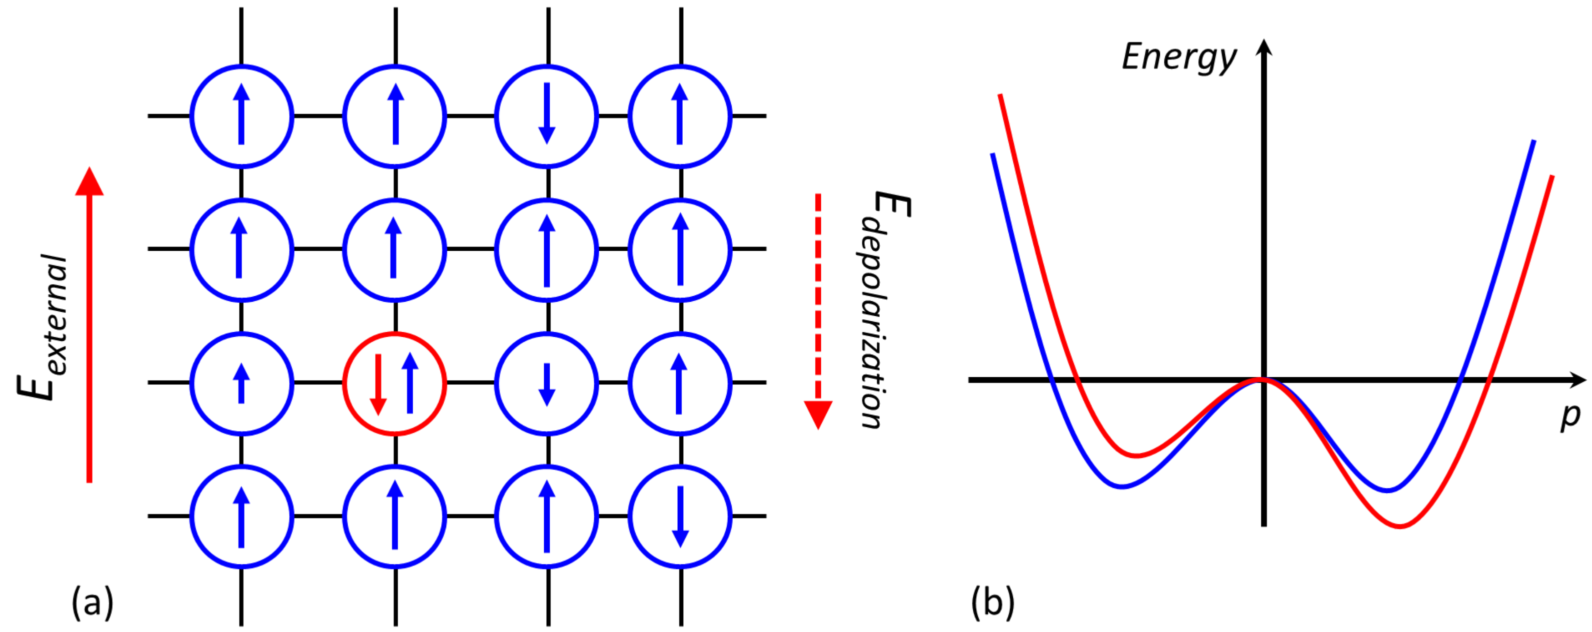

$$F = \sum_{i,j}^{N} \left((\alpha/2)p^2 + (\beta/4)p^4 - E_{loc}p + K \sum_{k,l} \left(p_{i,j} - p_{i+k,j+l}\right)^2 \right)$$

$$\frac{dp_{ij}}{dt} = -\gamma^{-1} \left(\beta p_{ij}^3 + \alpha p_{ij} + K \sum_{k,l} \left(p_{ij} - p_{kl}\right) - E_{loc} \right)$$

- Simple discrete time-dependent Landau formulation for ferroelectrics
- Code is available at github.com/ramav87/FerroSim

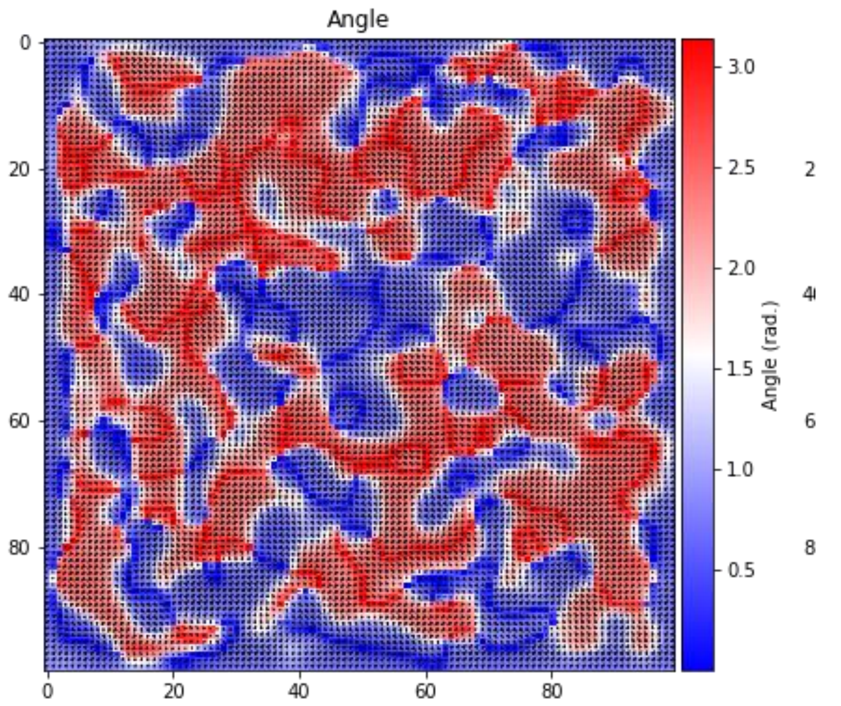

Kalinin, Ziatdinov, Vasudevan (JAP) (2020)

FerroSIM: Simulator for defects in ferroelectrics

- 2D Discrete Landau Model

Uniaxial

$$F_{unixial} = \sum_{i,j}^{N} \left(\frac{\alpha_1}{2}\right) p_{i,j}^2 + \left(\frac{\alpha_2}{4}\right) p_{i,j}^4 + K \sum_{k,l} (p_{i,j} - p_{i+k,j+l})^2 \; - E_{loc} p_n$$

In all cases,

$E_{loc} = E_{ext} + E_{dep} + E_d(i,j)$

Note also that E_{dep} is calculated as αP_{avg}

Tetragonal or Rhombohedral

$$F = \sum_{i,j}^{N} \alpha_1 \left(p_{x_{ij}}^2 + p_{y_{ij}}^2\right) + \alpha_2 \left(p_{x_{ij}}^4 + p_{y_{ij}}^4\right) + \alpha_3 p_{x_{ij}}^2 p_{y_{ij}}^2 + K \sum_{k,l} (p_{ij} - p_{i+k,j+l})^2 \; - E_{loc_x} p_{x_{ij}} - E_{loc_y} p_{y_{ij}}$$

Squareelectric

In the 'Square electric' we essentially have decoupled polarization

$$F_{square} = \sum_{i,j}^{N} \left[\left(\frac{\alpha_1}{2}\right) p_{x_{ij}}^2 + \left(\frac{\alpha_2}{4}\right) p_{x_{ij}}^4 + K \sum_{k,l} \left(p_{x_{ij}} - p_{x_{i+k,j+l}}\right)^2 + \left(\frac{\alpha_1}{2}\right) p_{y_{ij}}^2 + \left(\frac{\alpha_2}{4}\right) p_{y_{ij}}^4 + K \sum_{k,l} \left(p_{y_{ij}} - p_{y_{i+k,j+l}}\right)^2 - E_{loc_y} p_{y_{ij}} \right]$$

FerroSIM: Add defects to lattice and observe P map

- And easily add defects (changes to local E)

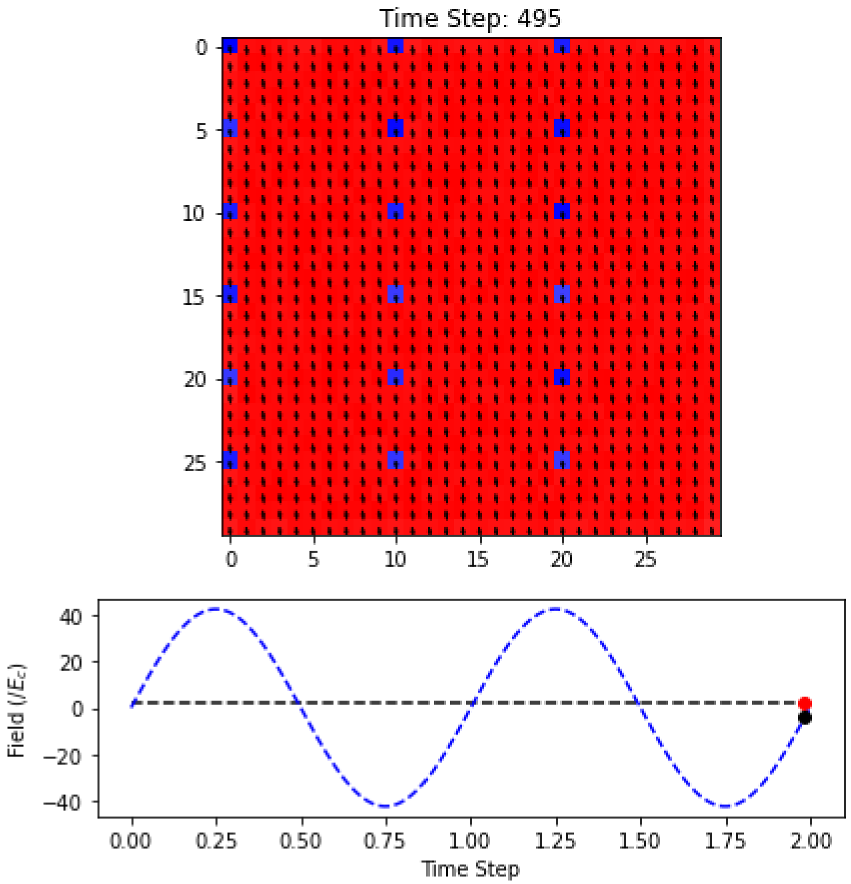

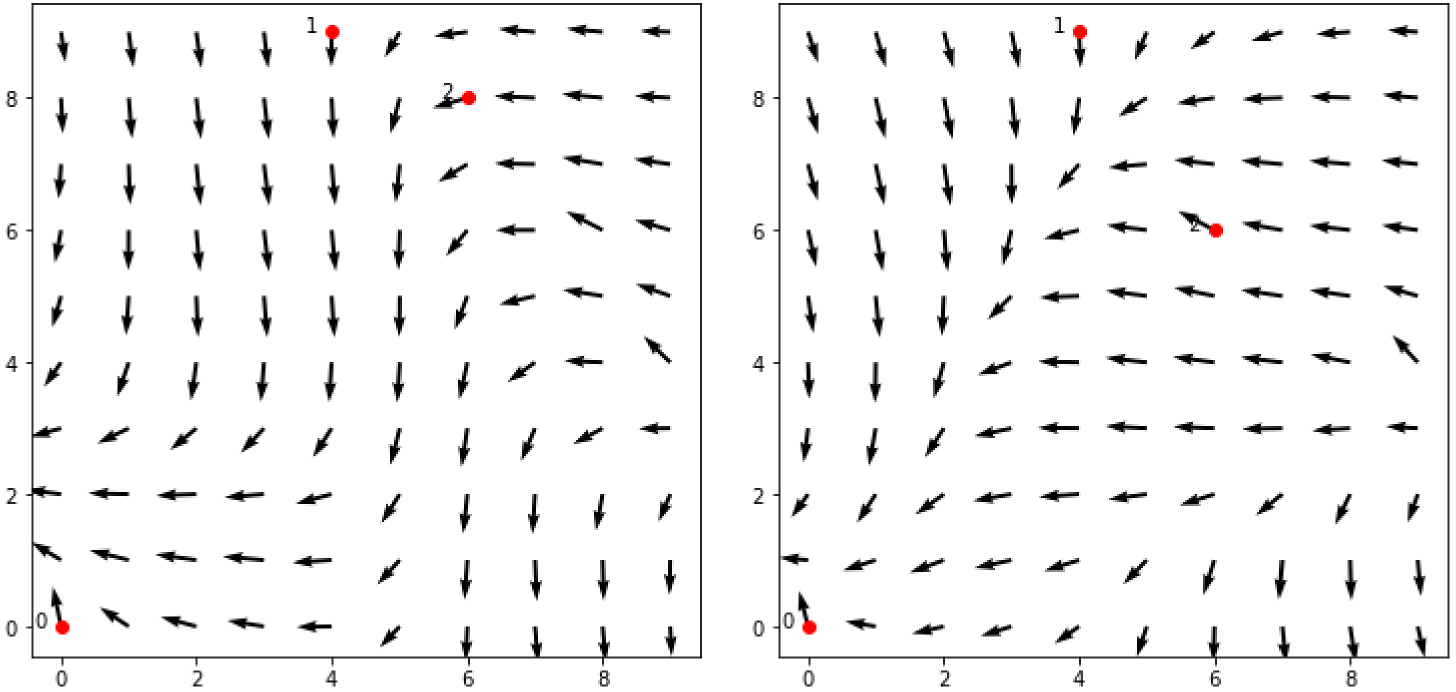

Moving defects to see how they affect the ground polarization state

OAK RIDGE National Laboratory

DQN for defect manipulation -> maximize curl

Initial

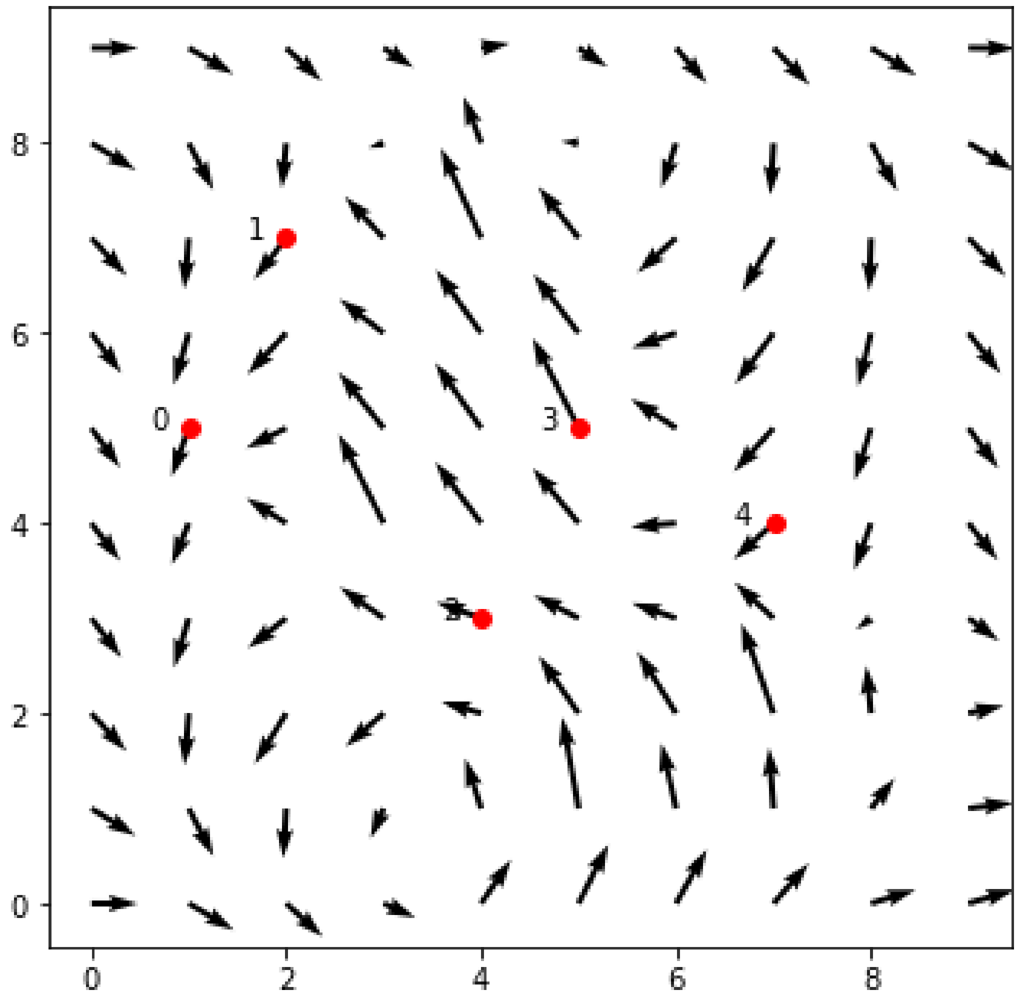

Final

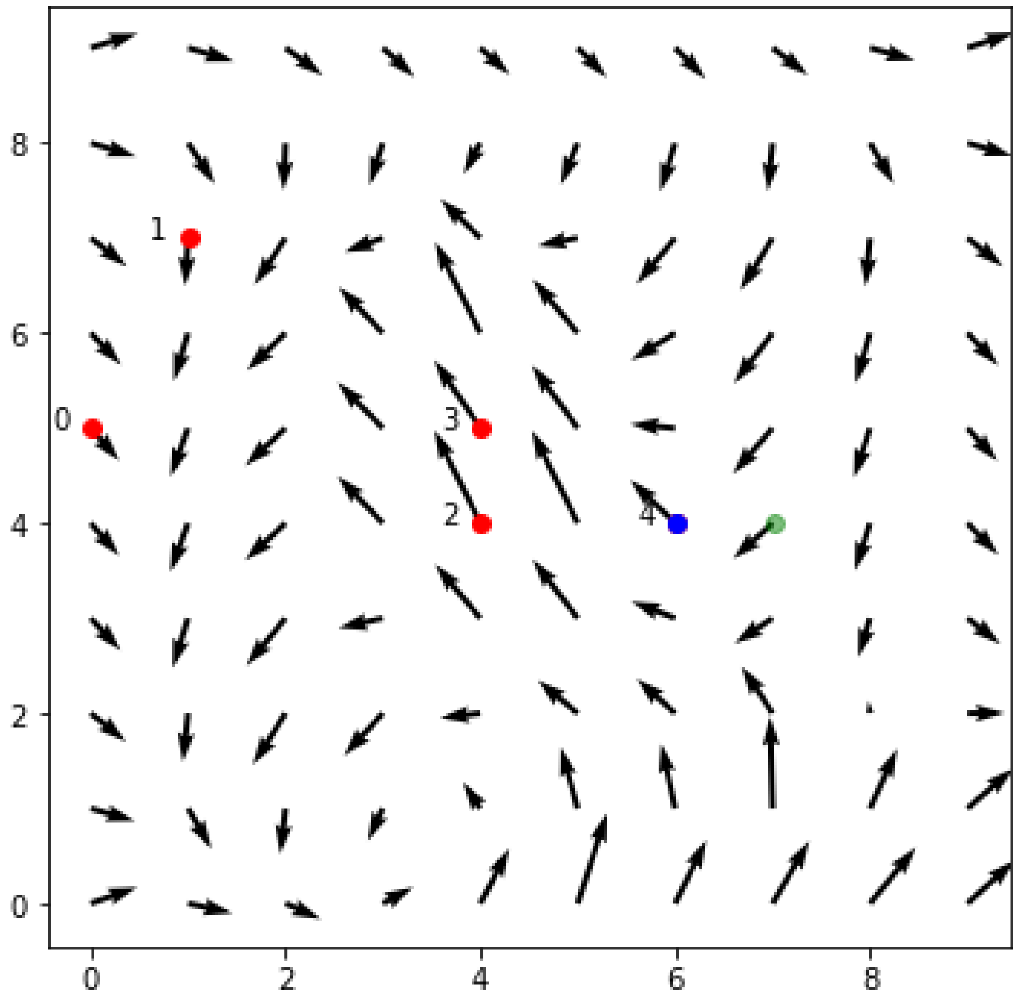

Agent trying to cluster defects

Another app: Molecular Dynamics Environment

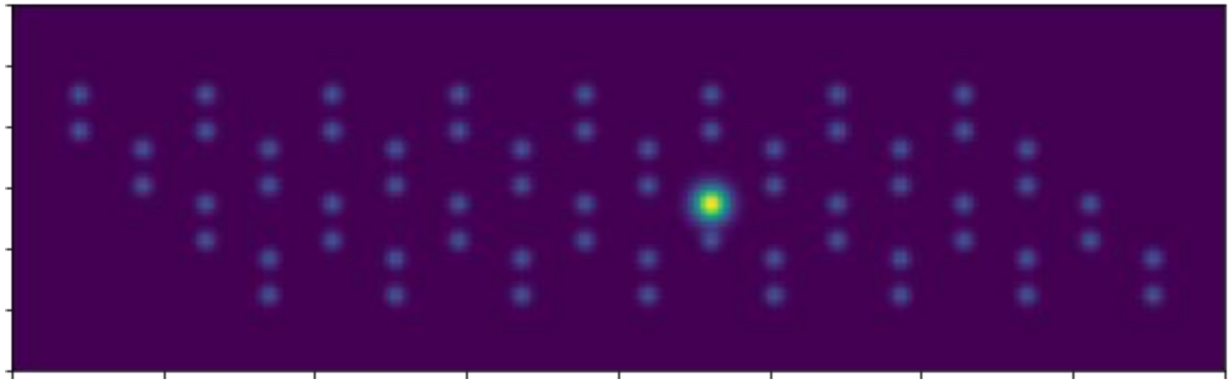

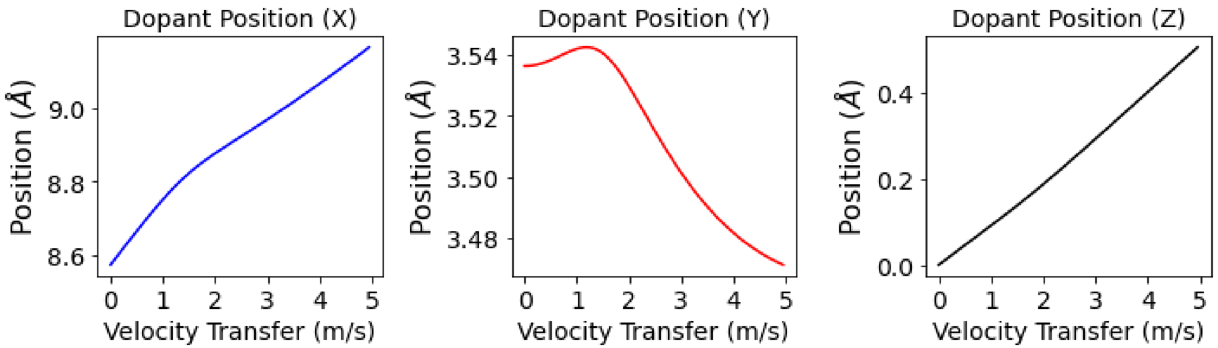

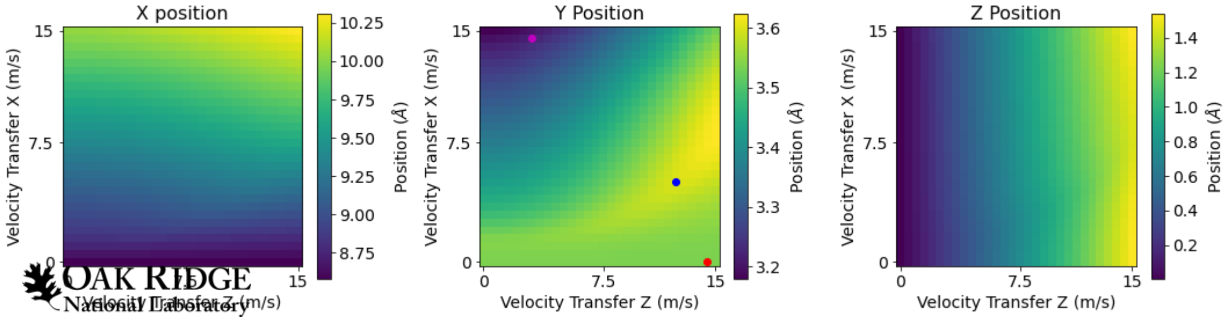

X_v = 1E-4 , Z_v = 14.5

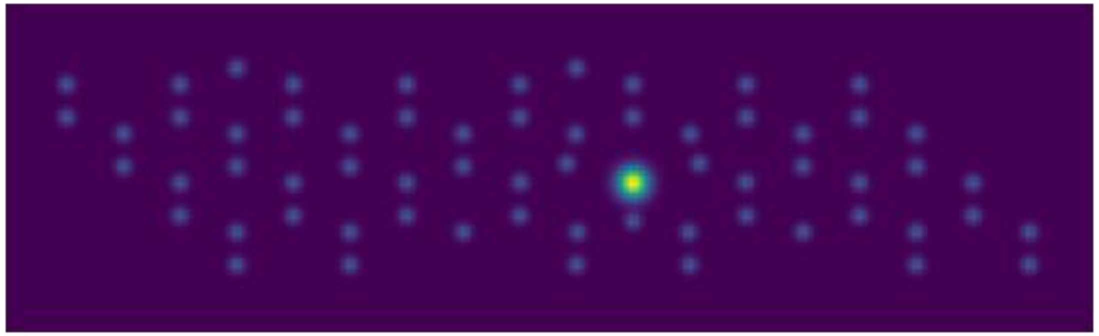

X_v = 5.2, Z_v = 12.4

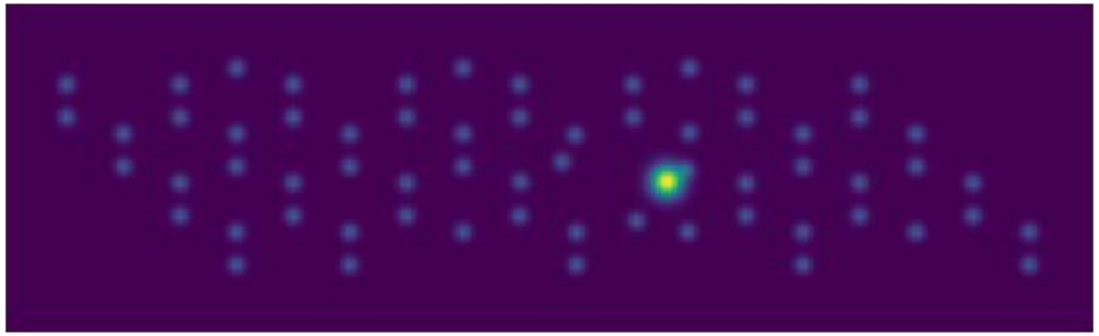

X_v = 14.5 , Z_v = 3.1

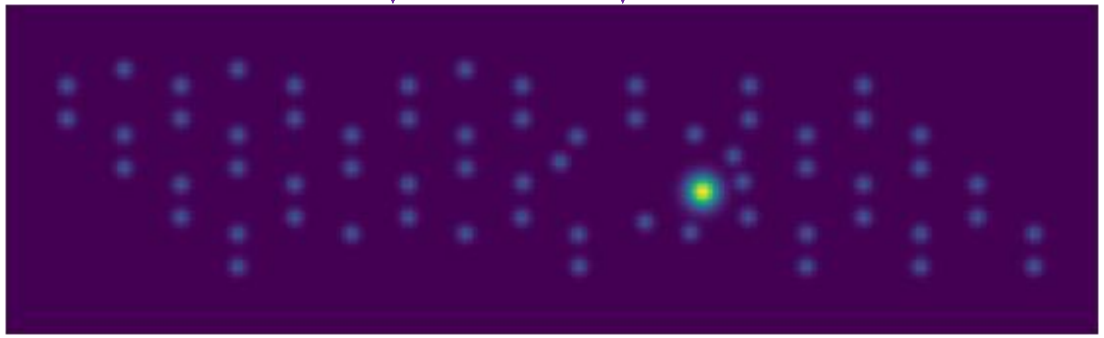

RL Environment for Atomic Fabrication: MD exploration

State: Si dopants in graphene

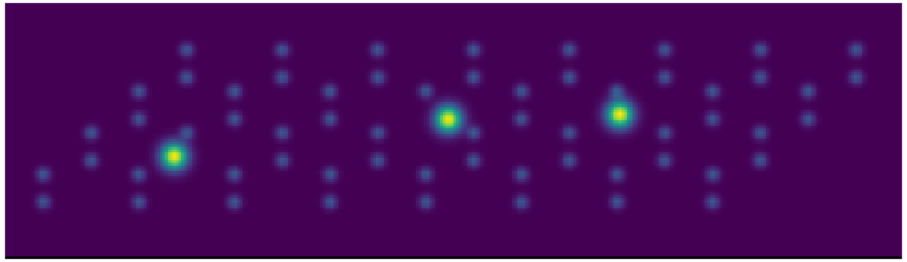

Goal: Move dopants together by changing momentum in MD

Action

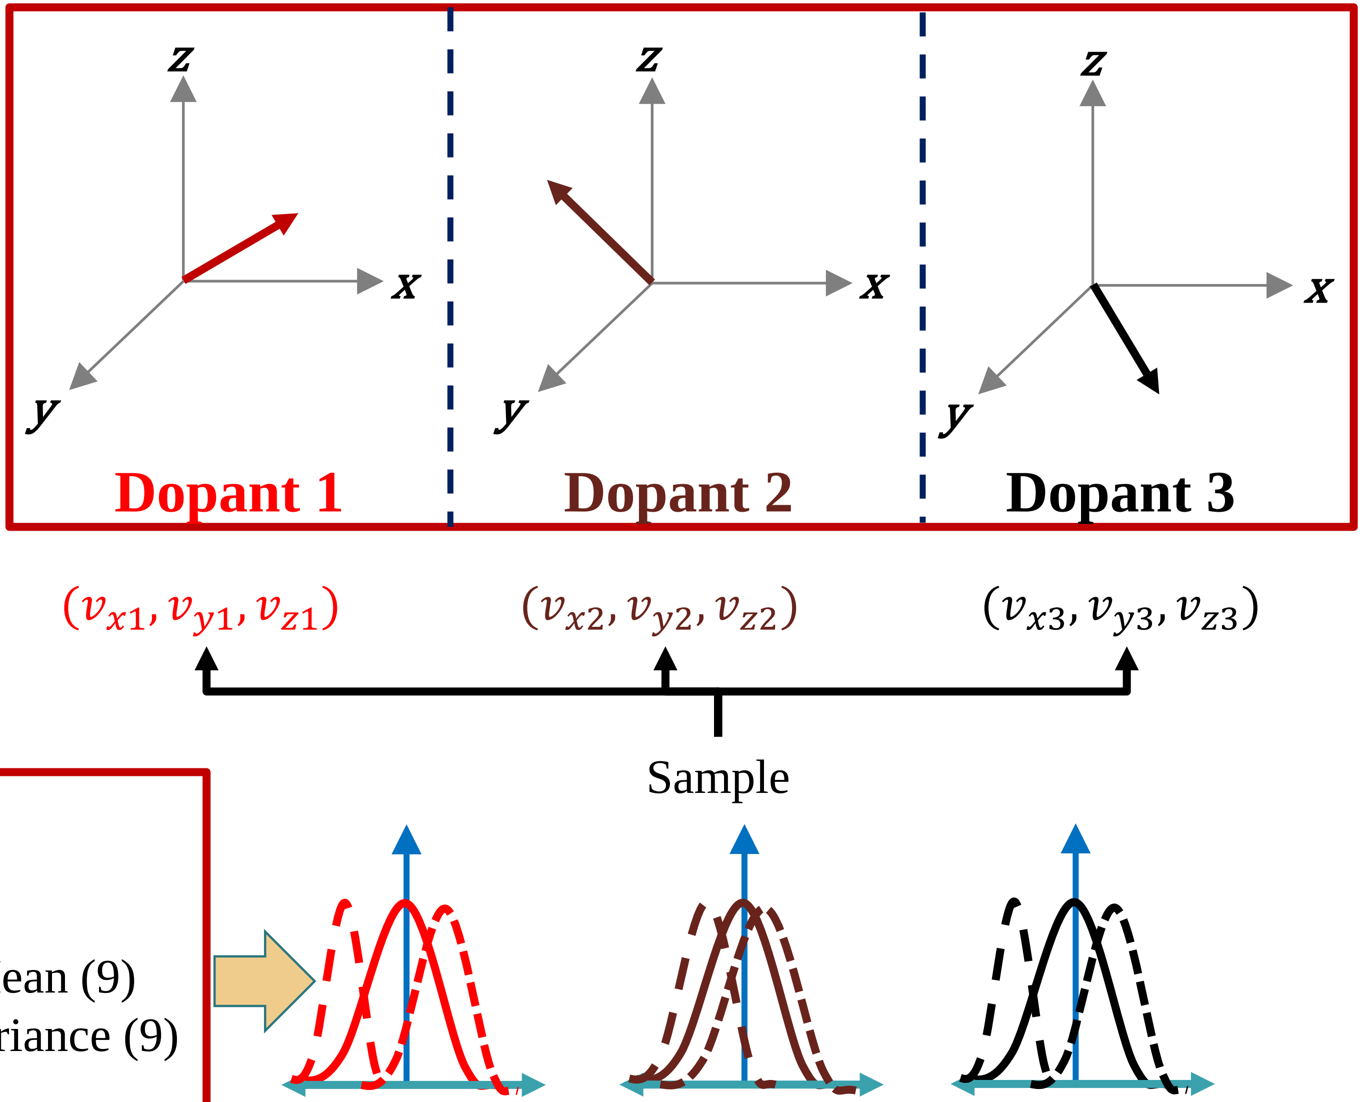

Multivariate Gaussian Policy

Actor/Critic Networks

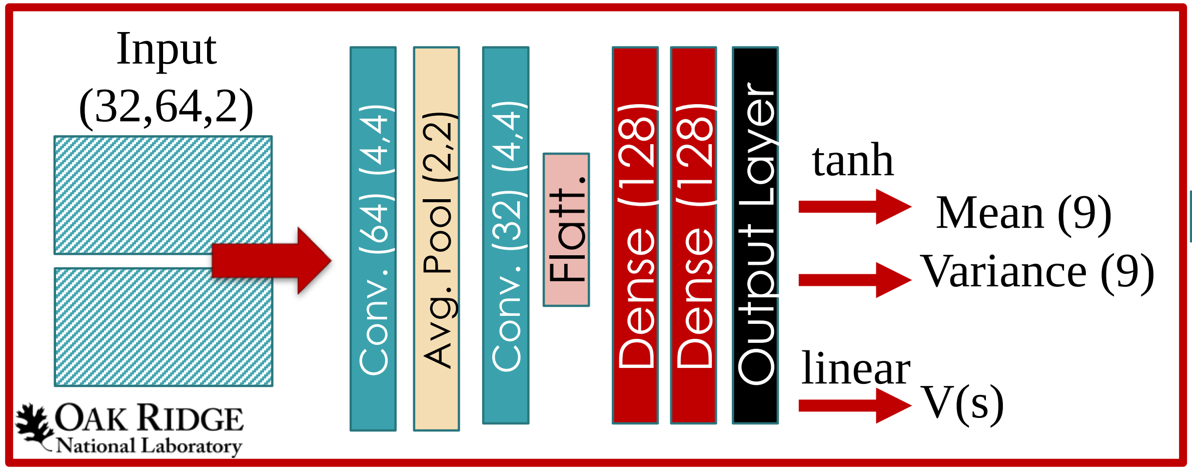

OAK RIDGE National Laboratory

Vasudevan et al. (unpublished)

Results: SVPG

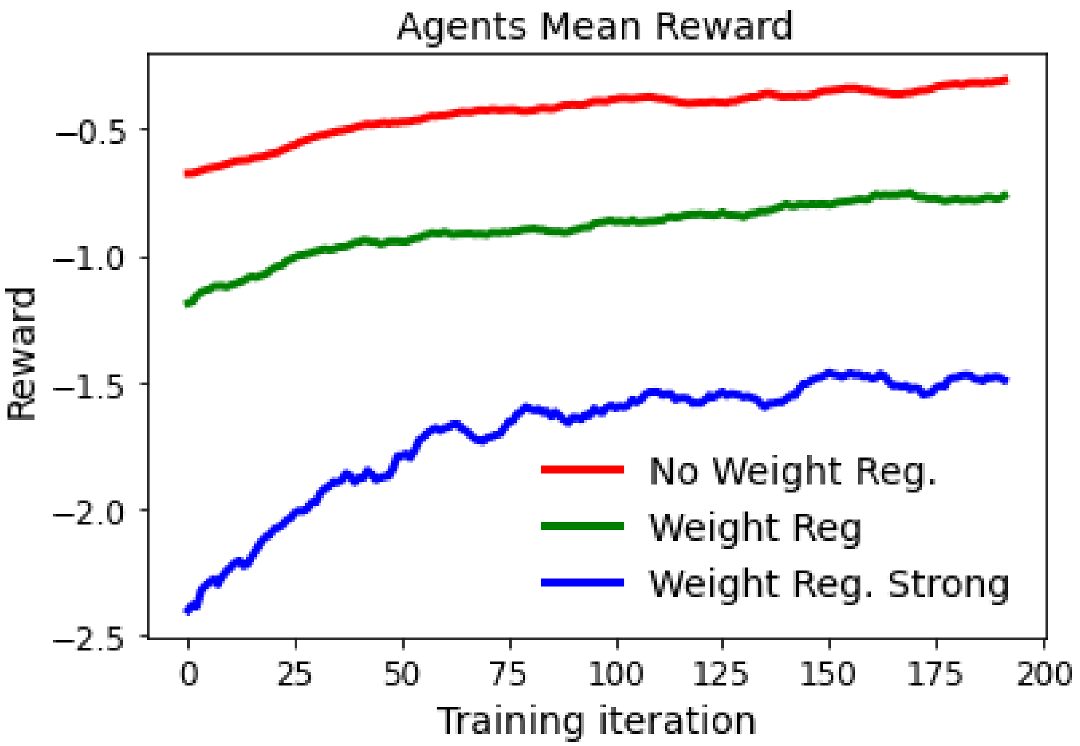

Vasudevan et al. (unpublished)

Runs of trained agents

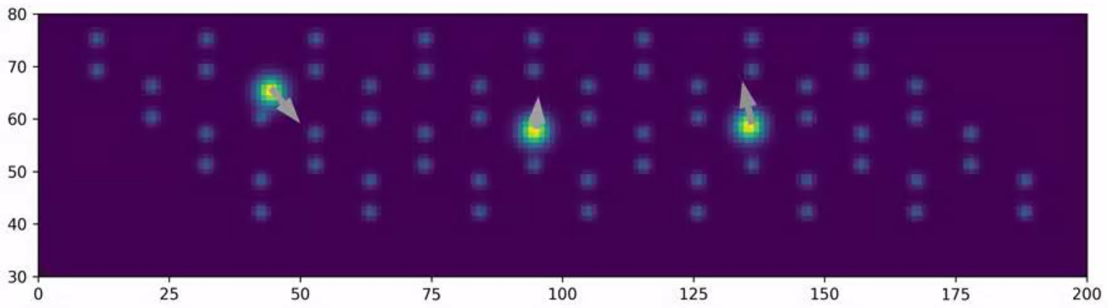

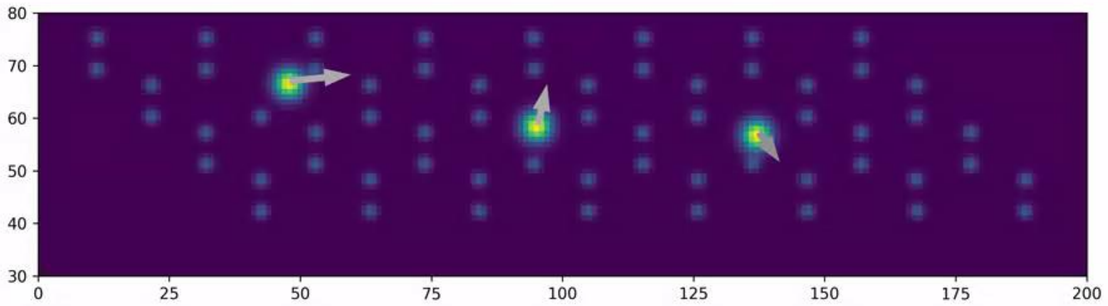

Results: Policy Inspection

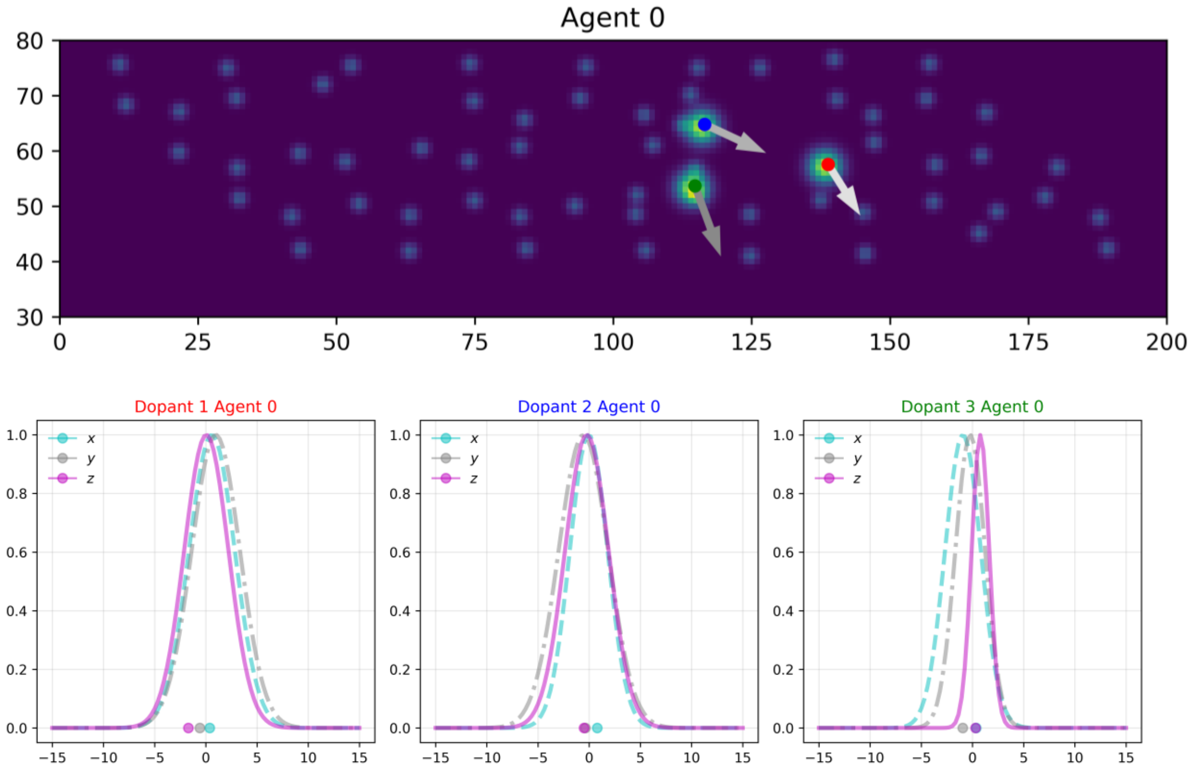

Vasudevan et al. (under review)

- Highly stochastic environment leads to conservative policies
- Z-component is not a delta function around zero- implies small z component is necessary to move dopant (also backed up by theoretical work)
- Policy inspection may become a useful tool to understanding the dynamics of the system -> relevant dynamics are learned, somewhat simplifying the problem.

OAK RIDGE National Laboratory

Outlines

RL: Defect engineering

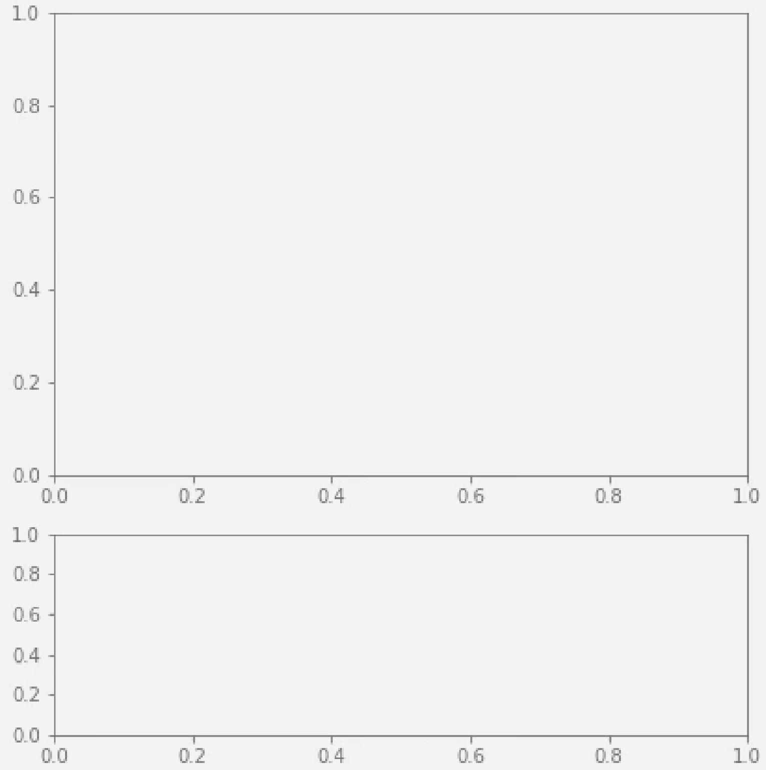

Automated Exp.: Manipulations

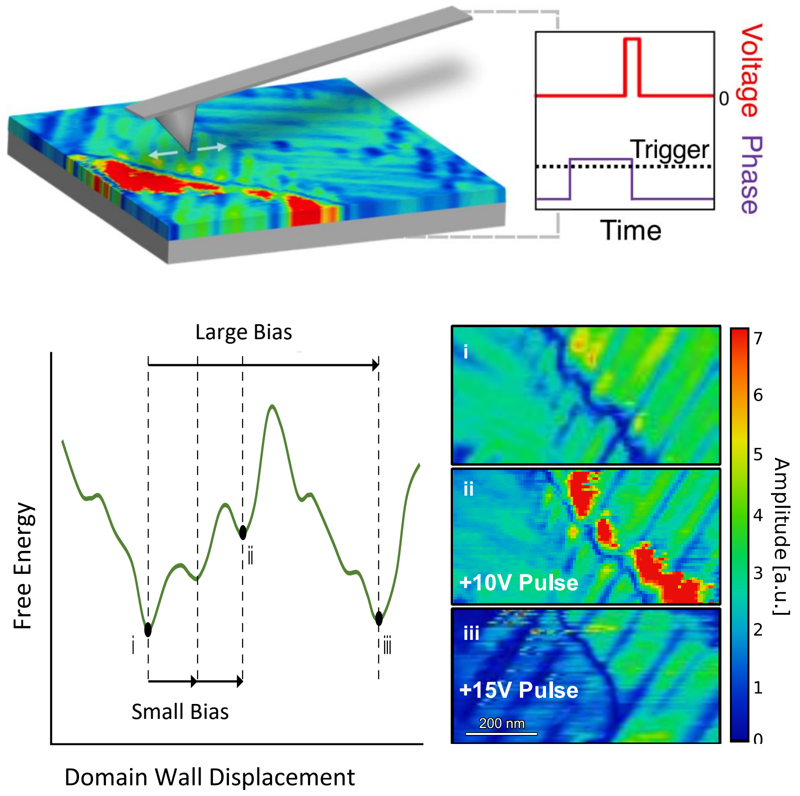

Adaptive Sampling

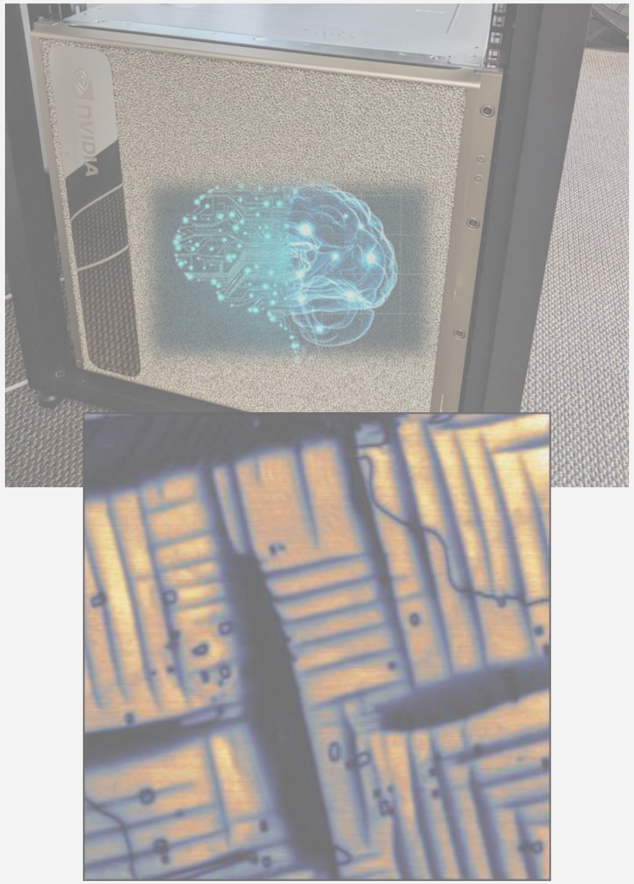

OAK RIDGE National Laboratory

FerroBOT: Automated manipulation of domain walls

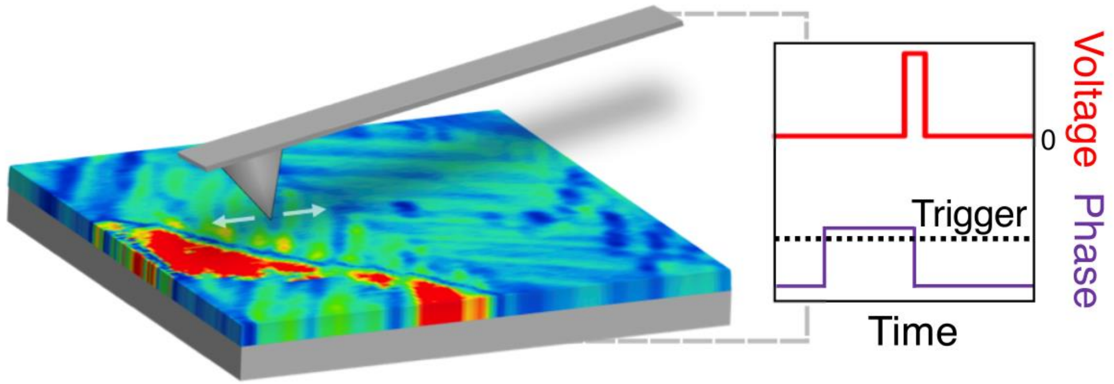

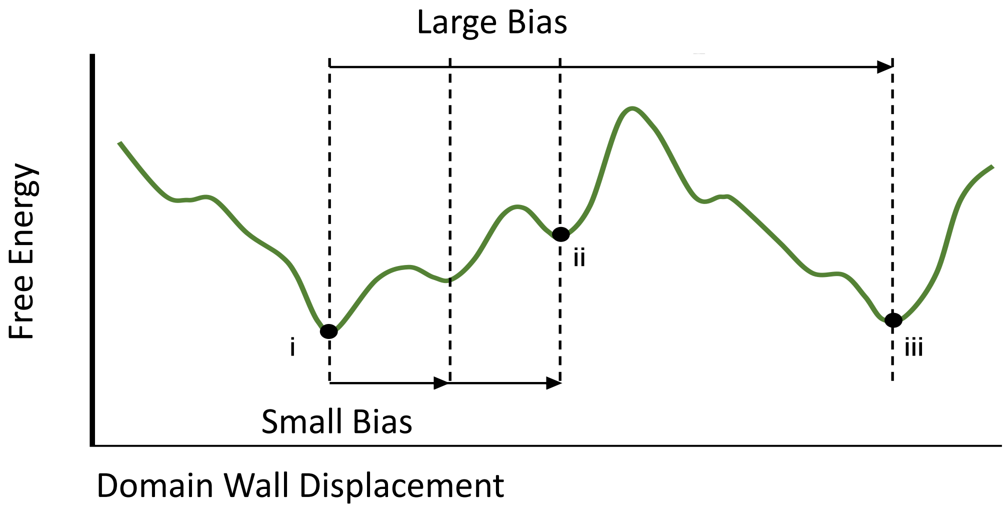

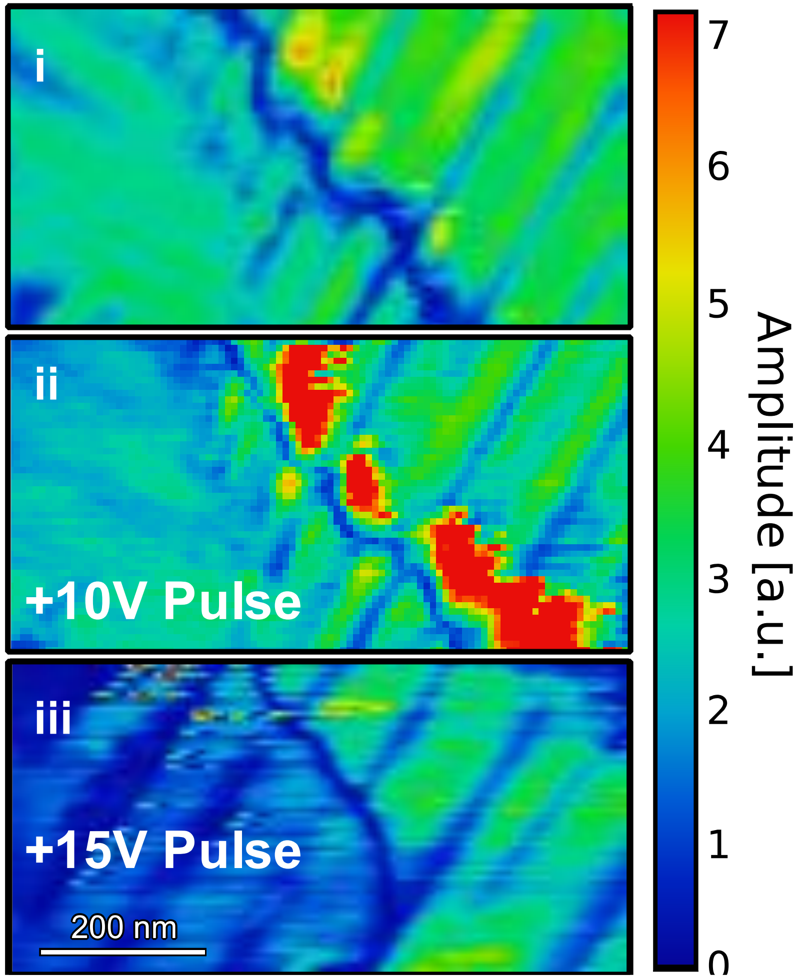

Kelley et al., ACS Nano **14**, 8, 10569–10577 (2020).

FerroBOT: Recent Extensions

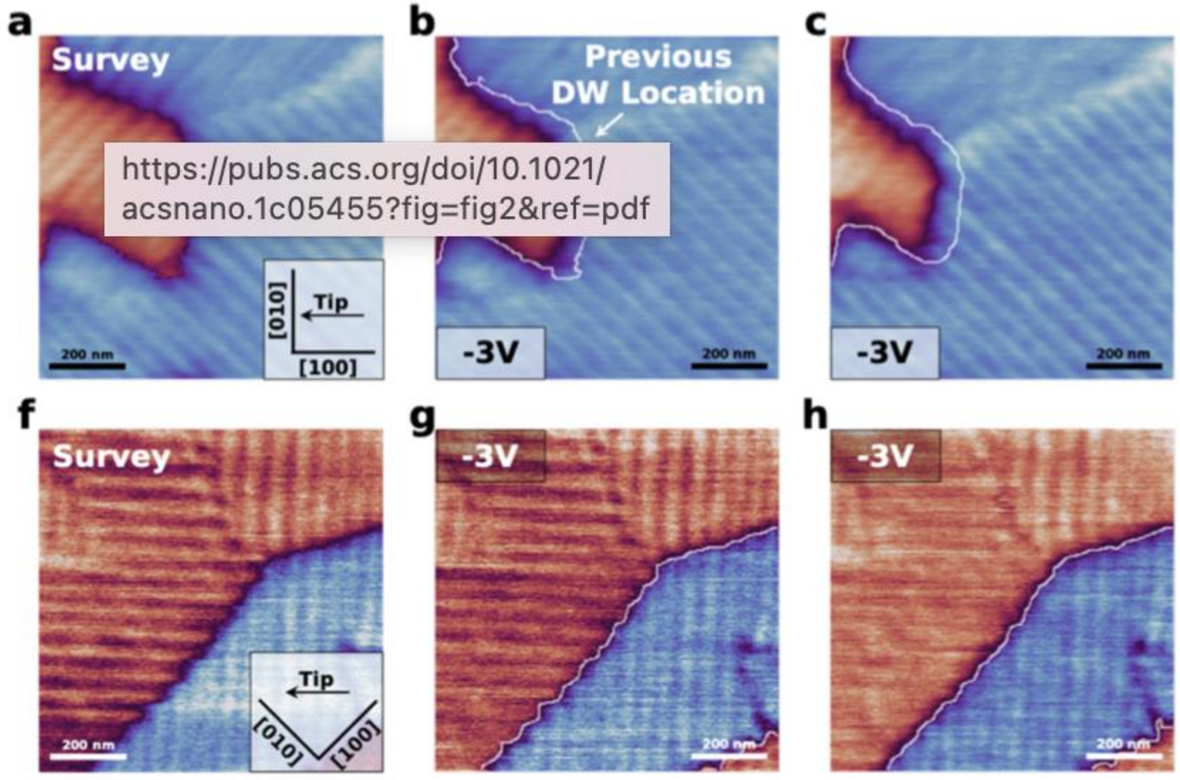

Kelley et al., (in press)

Triggered stimulus to manipulate domain walls automatically

Image-based feedback modes

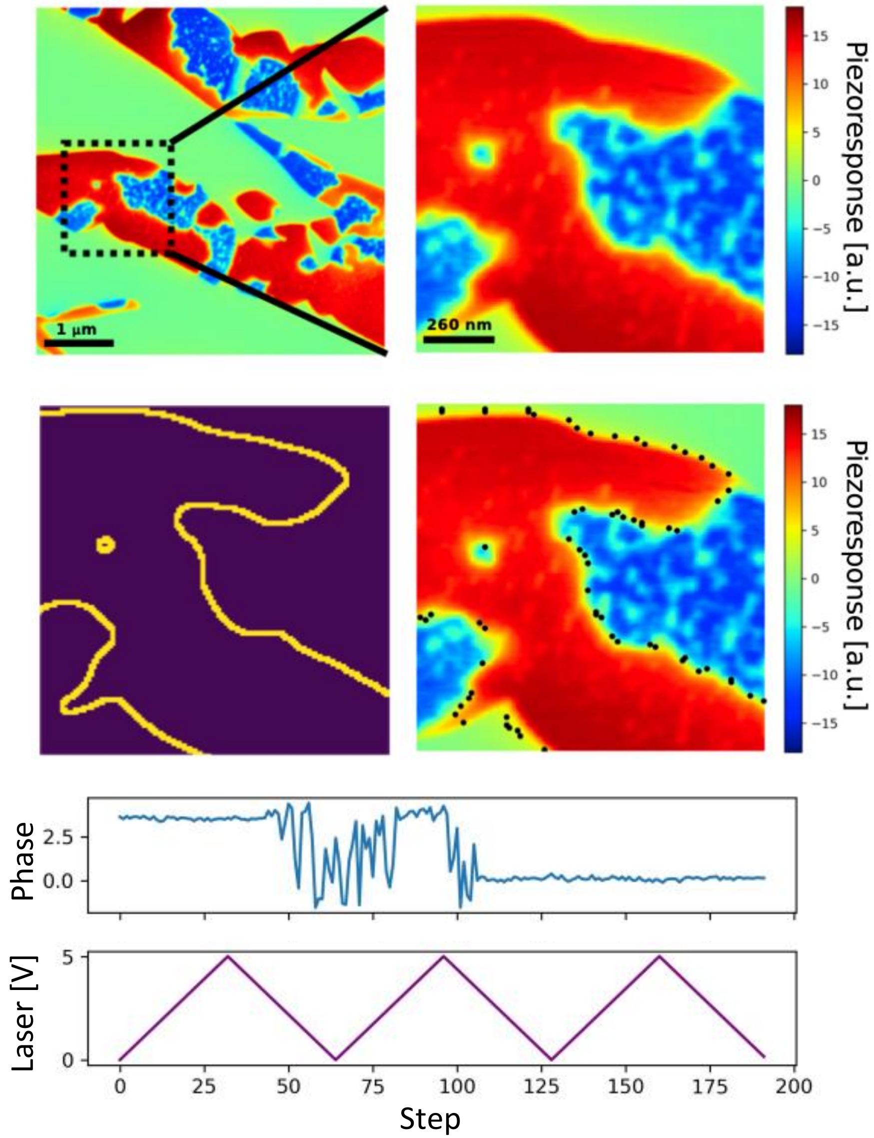

Kelley et al., (in press)

Oak Ridge National Laboratory

Outlines

RL: Defect engineering

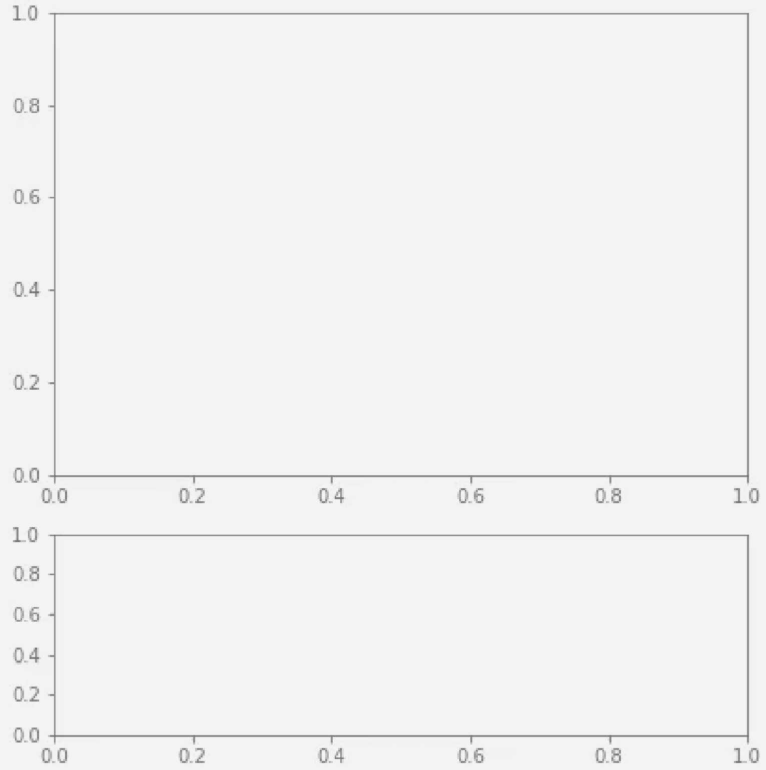

Automated Exp.: Manipulations

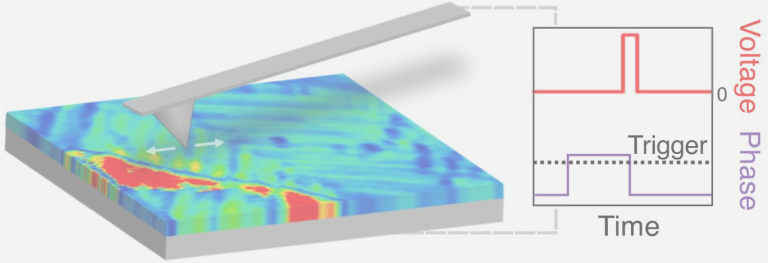

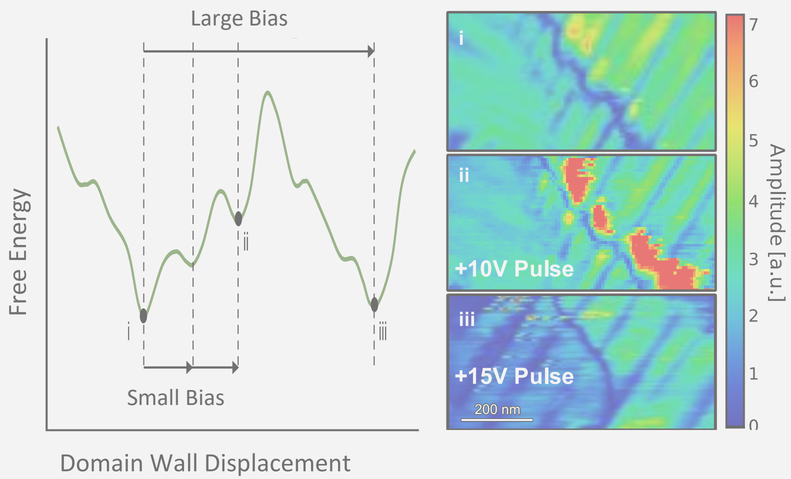

Adaptive Sampling

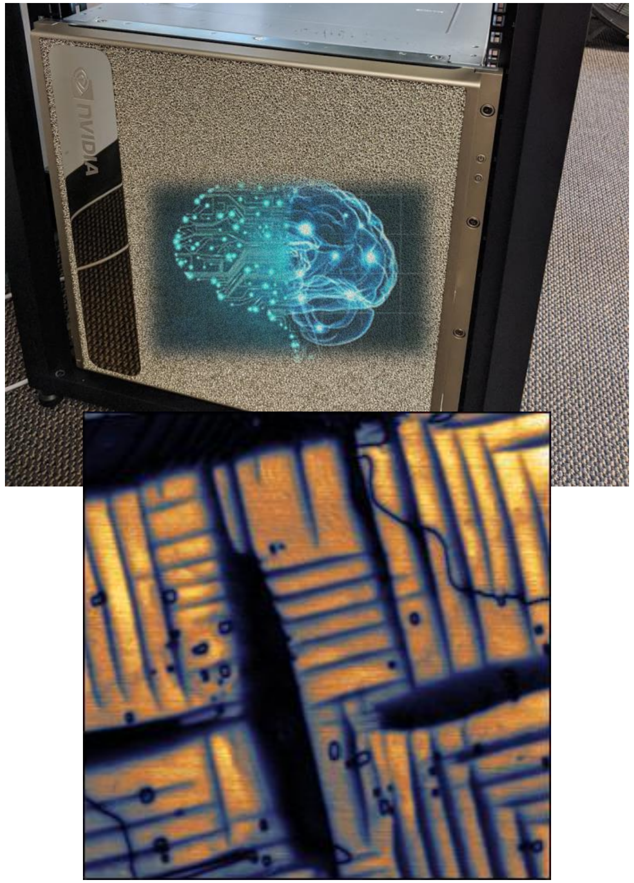

OAK RIDGE National Laboratory

Bayesian Optimization

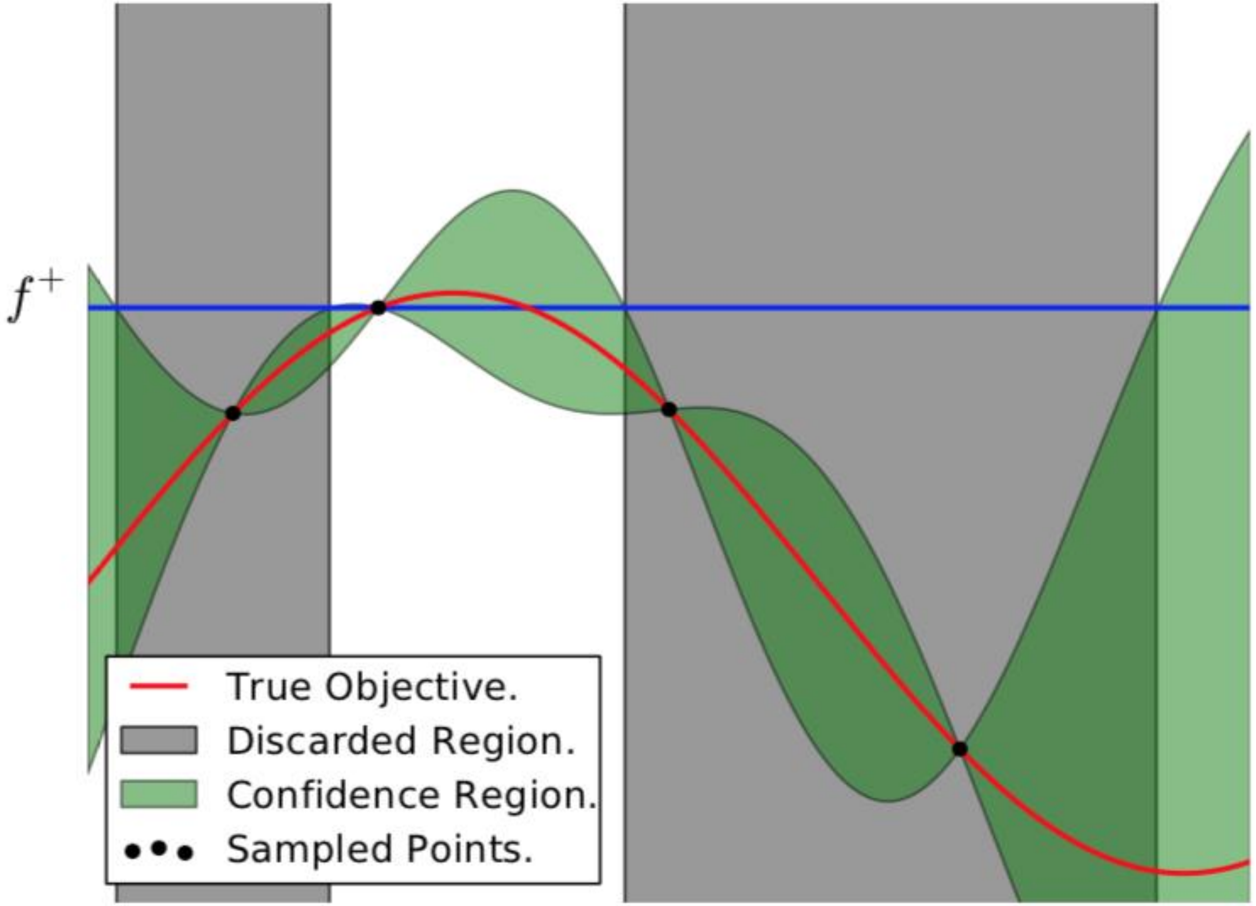

N. de Freitas et al., Taking the Human Out of the Loop: A Review of Bayesian Optimization , *Proceedings of the IEEE* **104**, 148 (2015)

- We have some measurements in space X, and we want to maximize some property f(X).
- Generalizes to higher dimensions
- Recently has become computationally tractable

Acquisition Functions

Probability of Improvement Acquisition Function

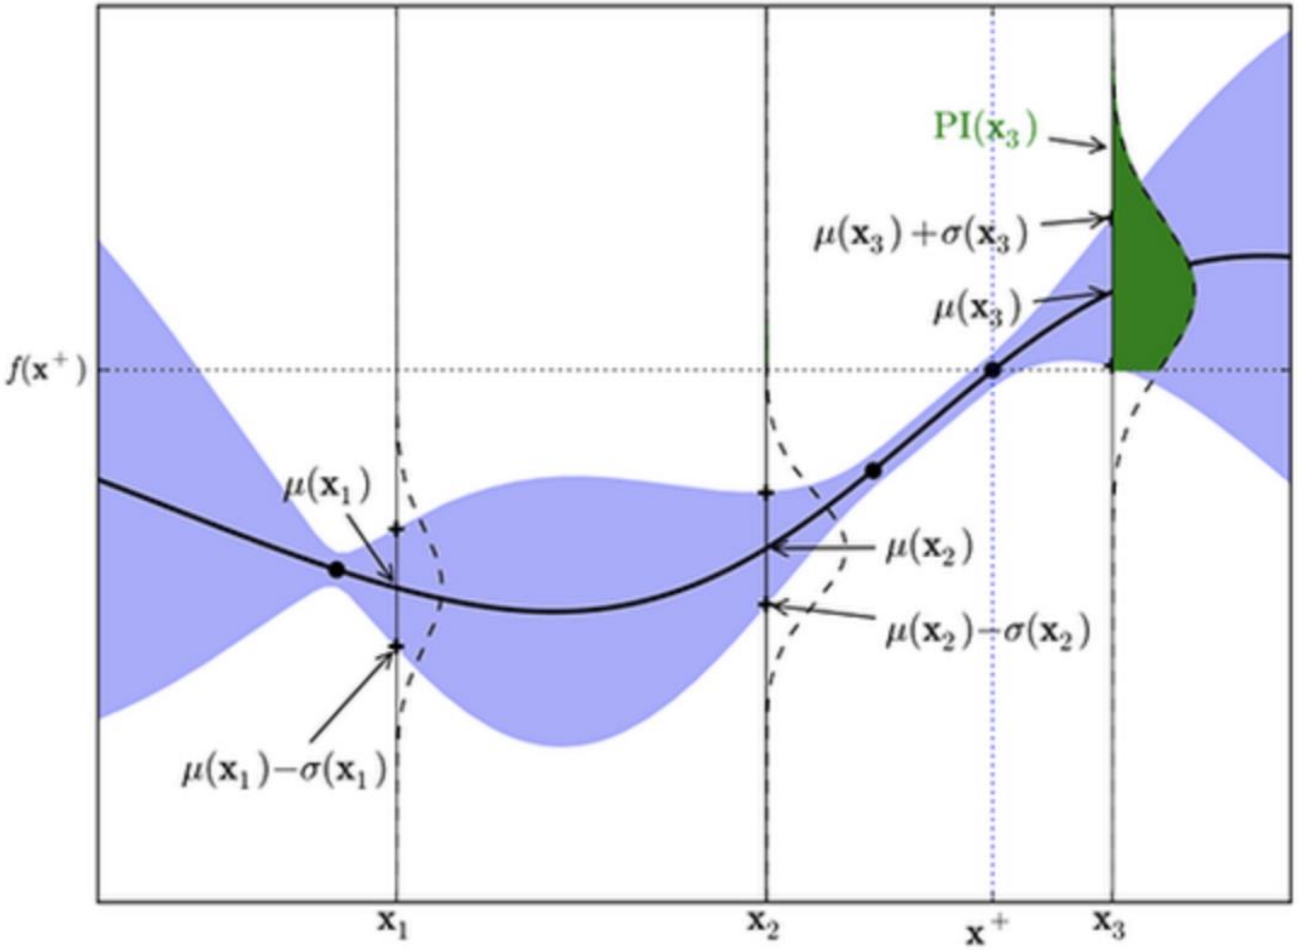

1. **Confidence bound:** simplest possible - just take the upper confidence bound from the prediction
2. **Probability of Improvement:** Integral from current functional maximum to upper limit of distribution as test point
3. **Expected Improvement:** Instead of probability of improvement, we want to maximize the expected increase in the function value
4. **There are (always) more…**

Efficient sampling required!

- Large spectroscopic datasets take too long to capture: efficiency in sampling required. Can be done via Bayesian optimization.

Provide list of new measurement conditions

Measure a batch of points

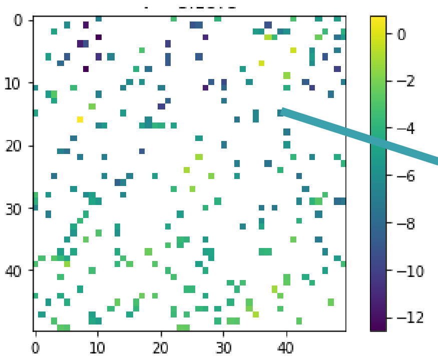

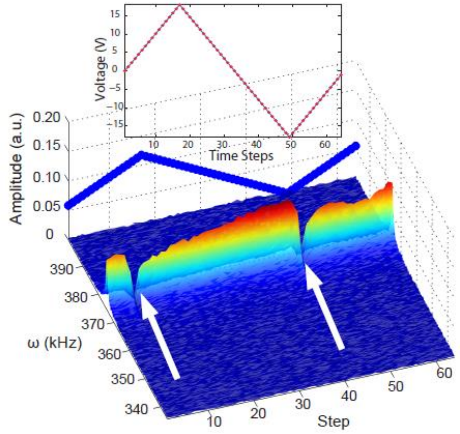

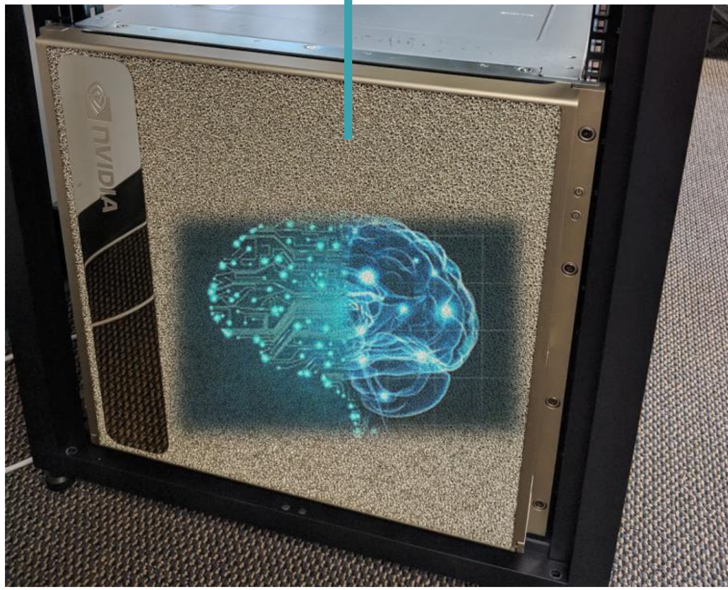

Upload to dgx for Bayesian Optimization

Data preprocessing

Spectroscopy on a ferroelectric film

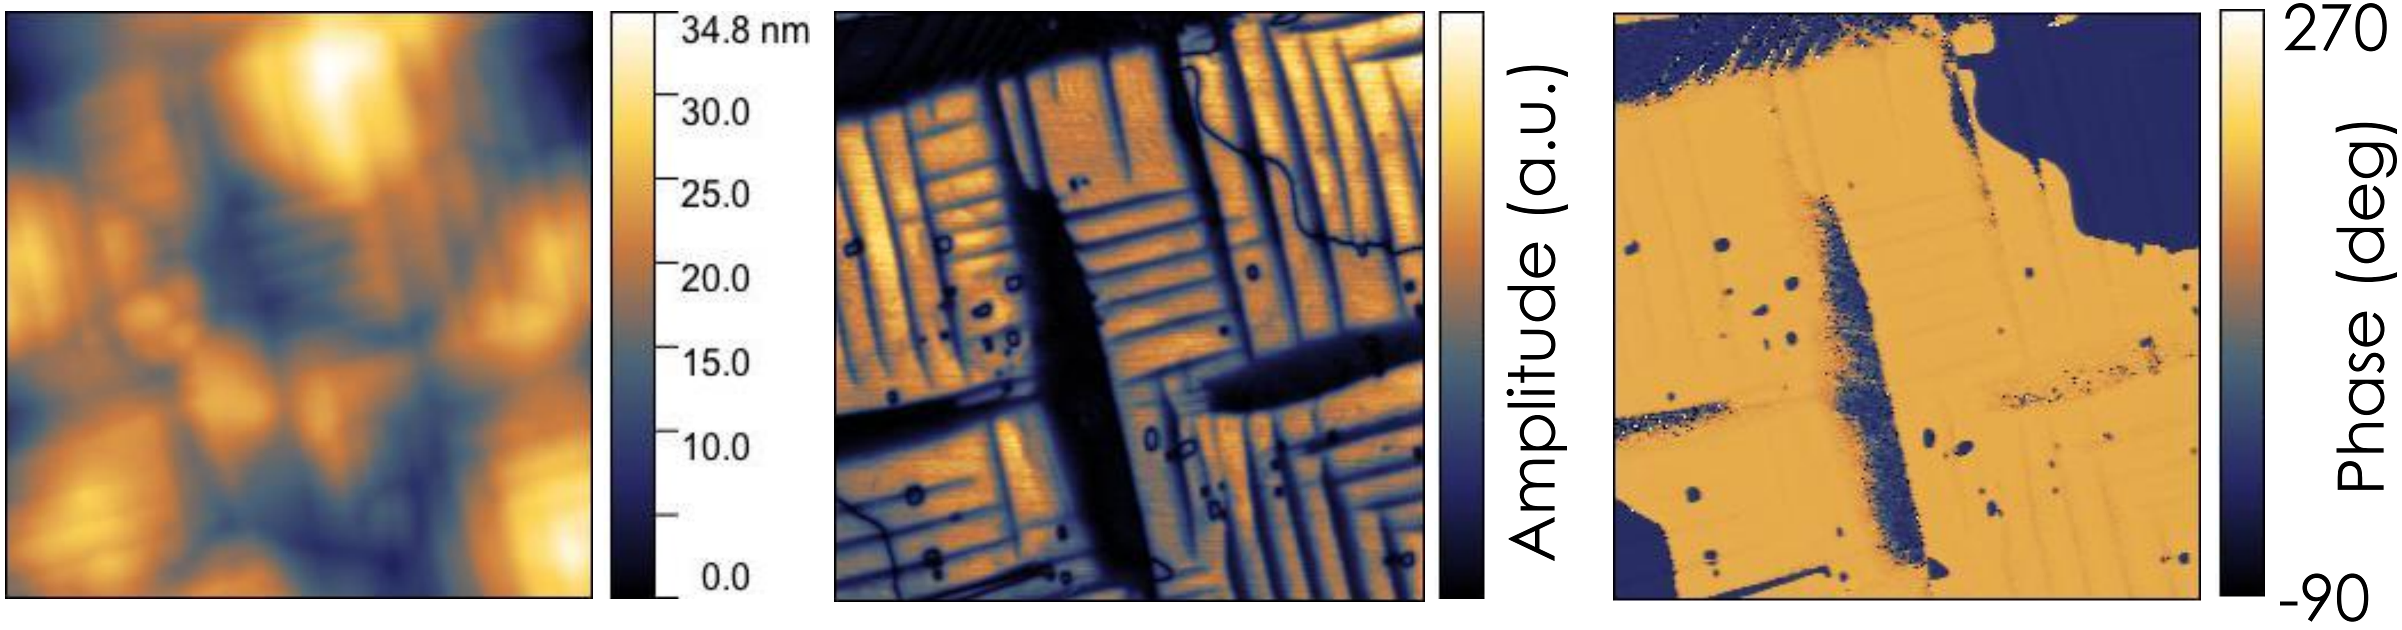

Can image the domain structure with the microscope -> 4 minutes
Spectroscopy – obtaining spectra pixel by pixel – can take 2-24 hours depending on type of measurement.

Automated Experiment example

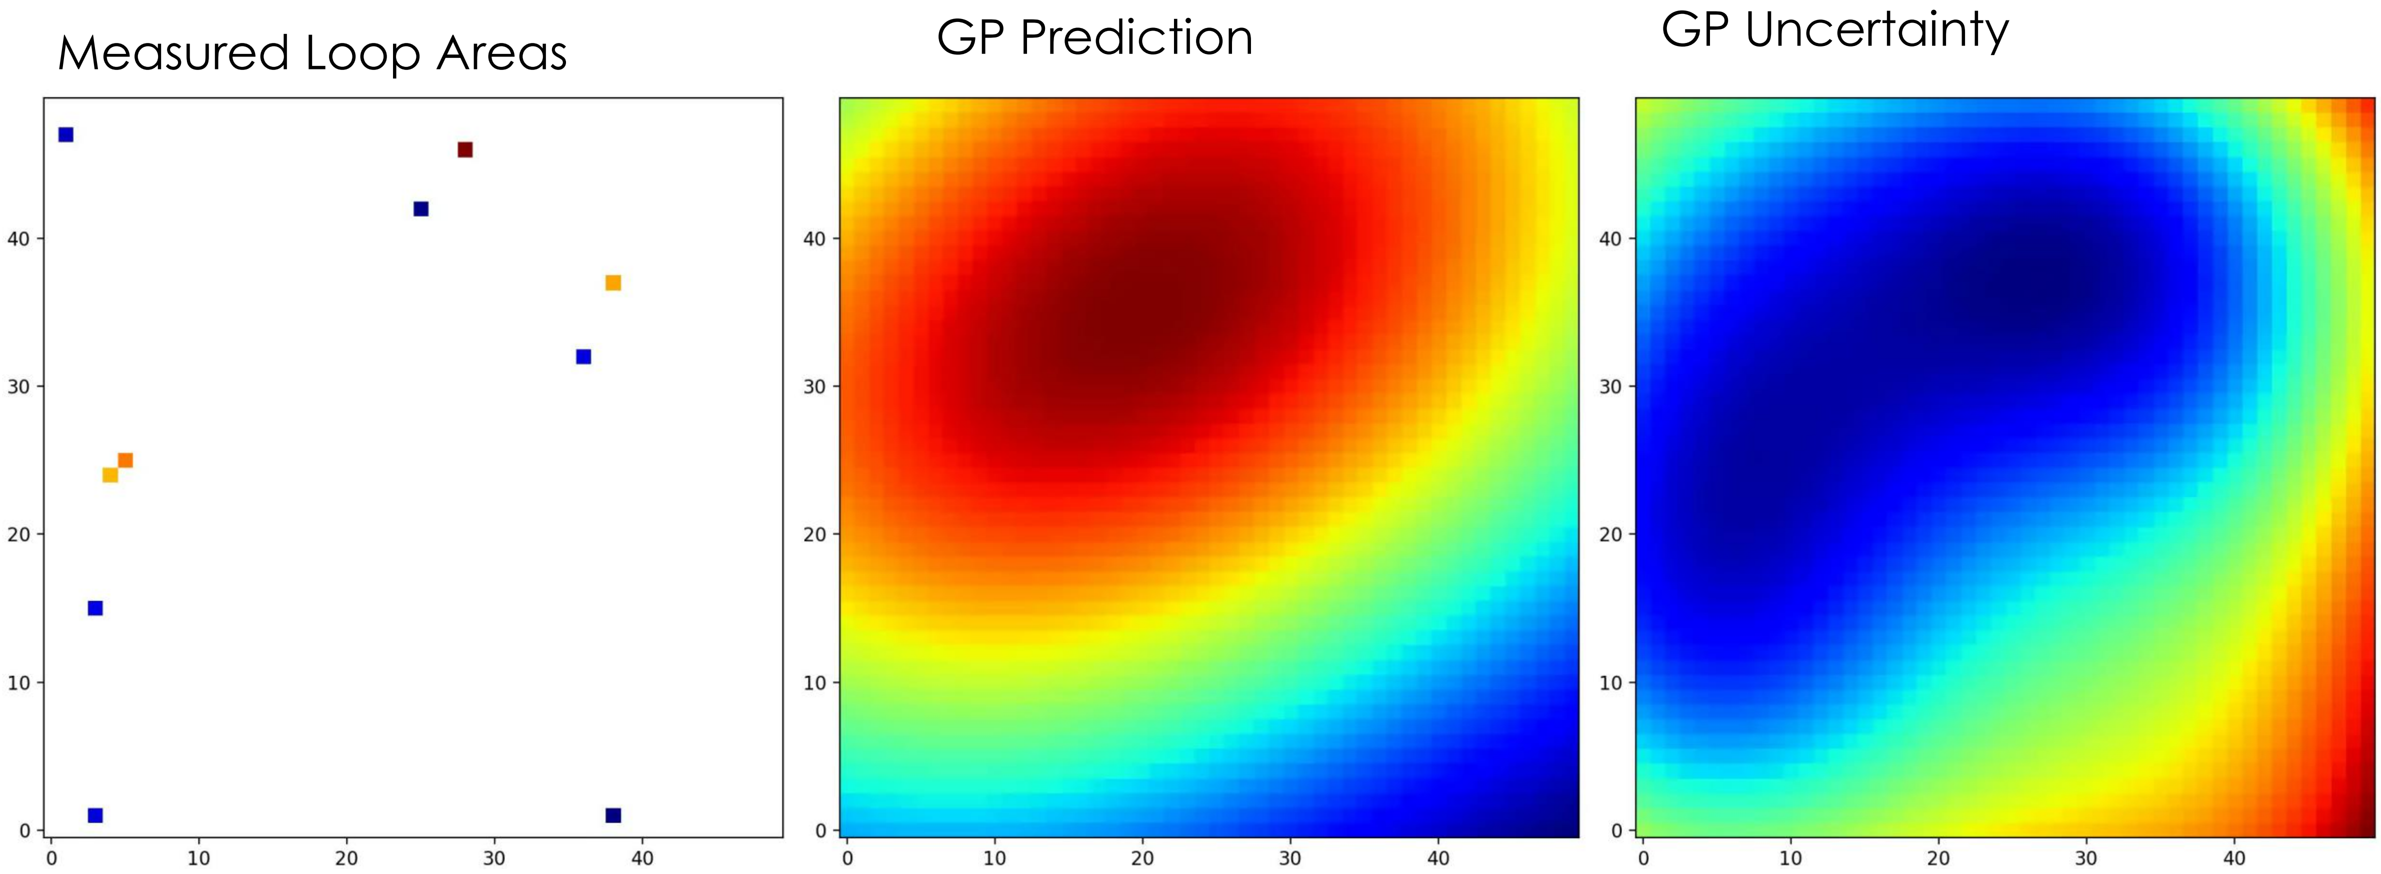

R. Vasudevan et al., arXiv:2011.13050 (under review)

OAK RIDGE National Laboratory

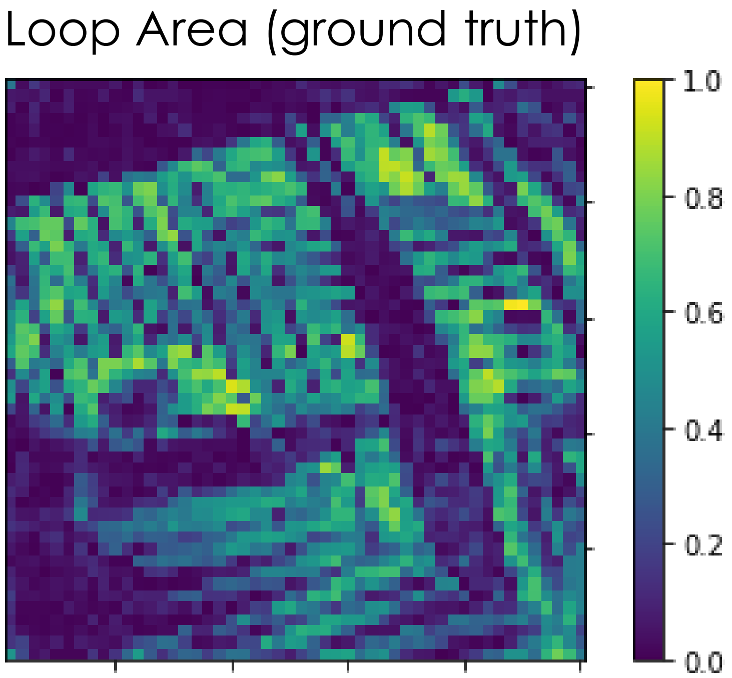
Loop Area (ground truth)
1.0
0.8
0.6
0.4
0.2
0.0

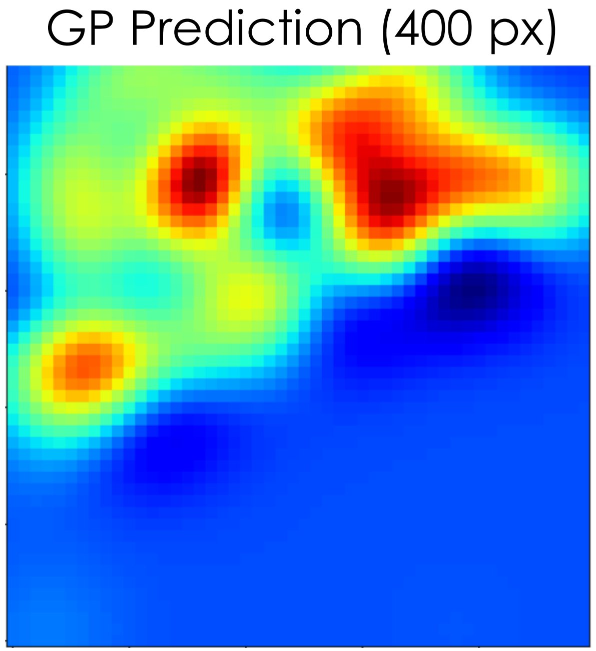
GP Prediction (400 px)

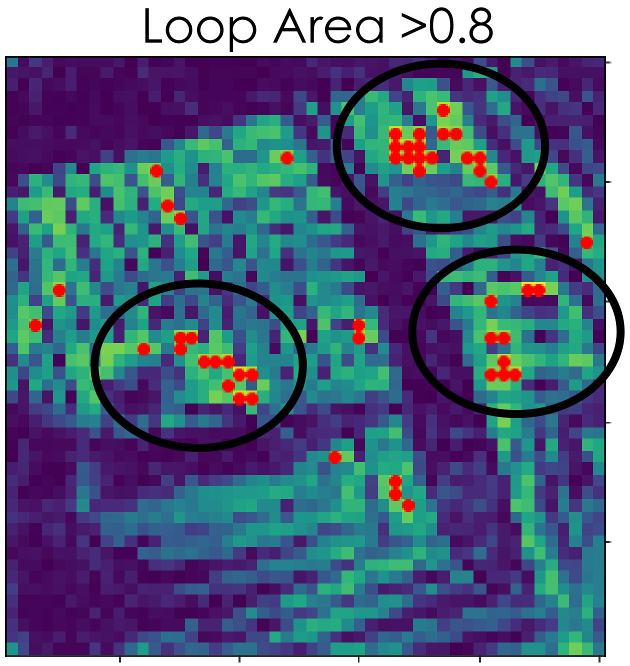
Loop Area >0.8

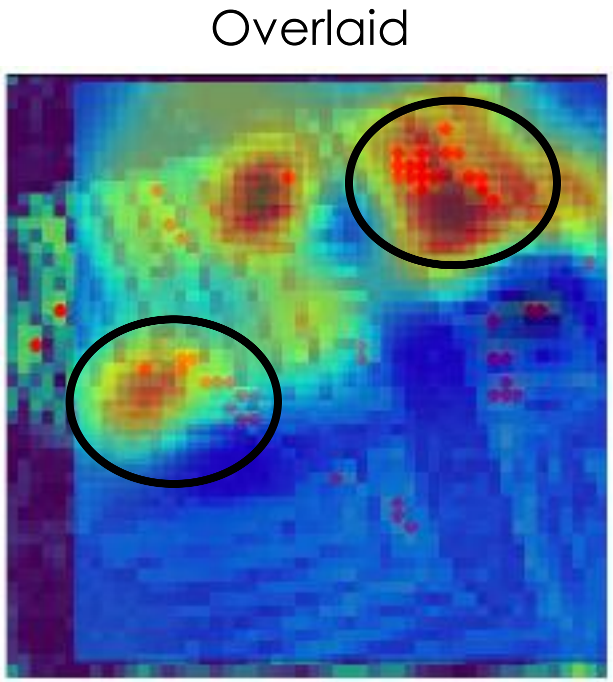
Overlaid

Deep Kernel Learning: Better priors means better results

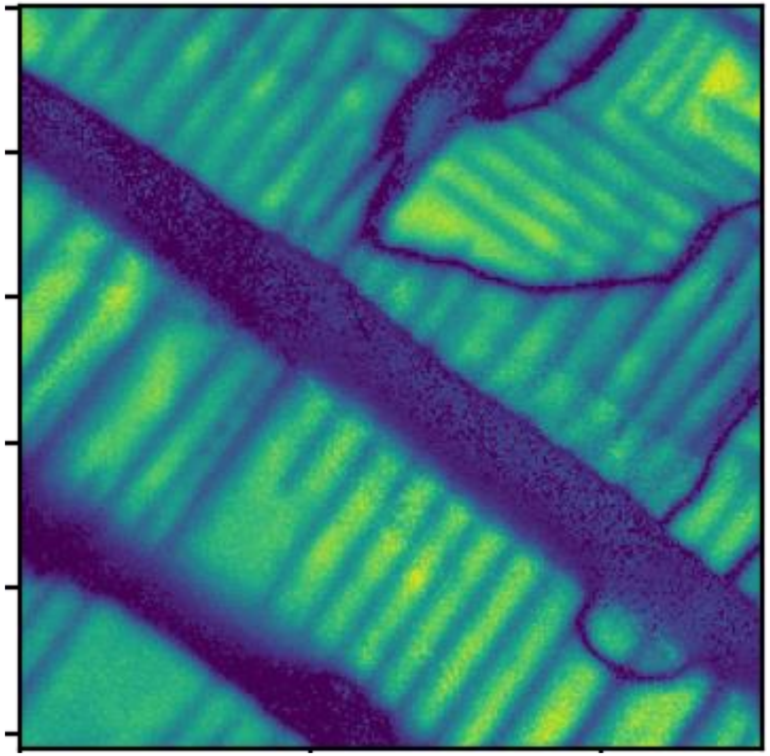

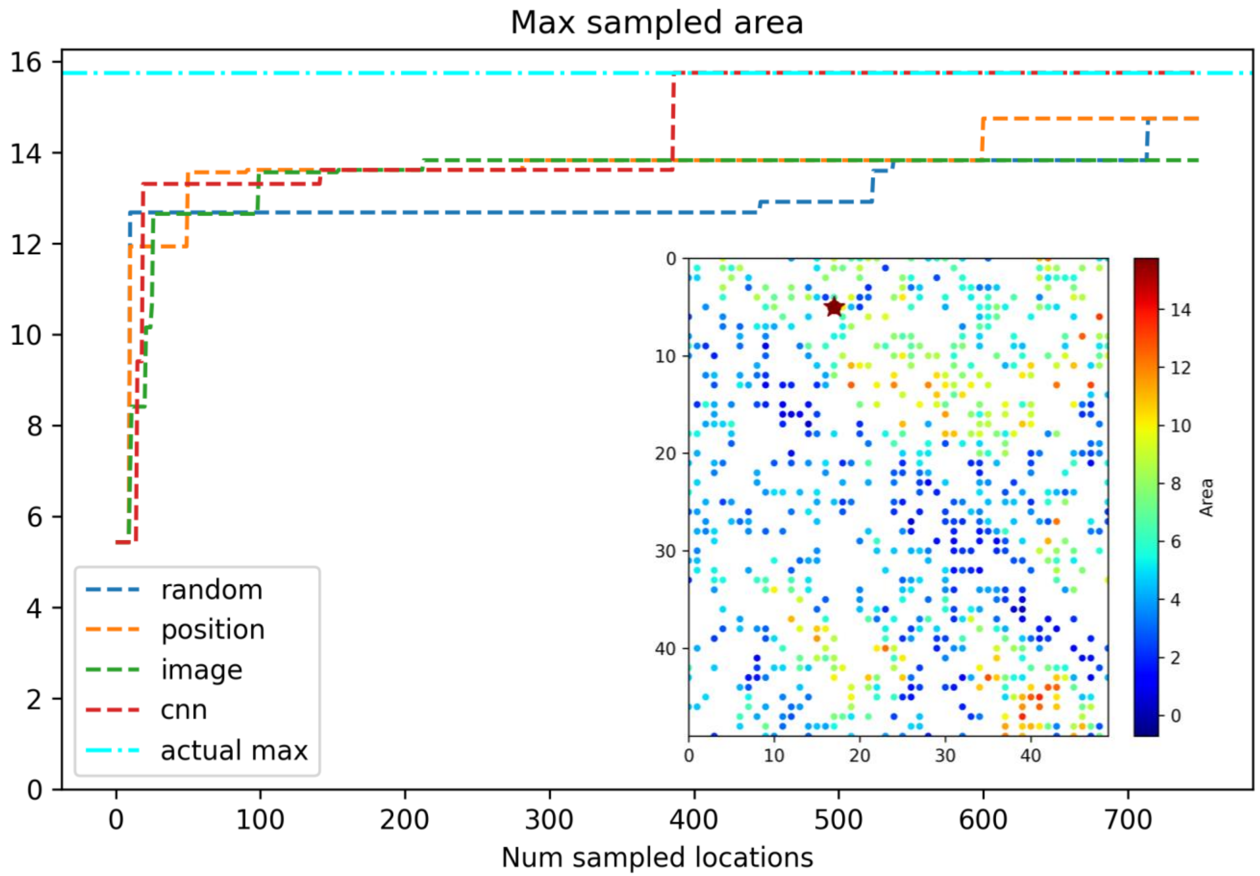

Utilize a CNN or just directly image pixels to better determine next sampling locations

Vasudevan et al. ACS Nano, 2021

Computational Needs: Streaming, Near Edge and HPC

~64 MB/s

~ MB-GB/run

~ 10+GB/job

Instrument

“Streaming” Edge

FPGA*

Jetson
Nano

GPU “Far” Edge

DGX-2
16 GPUs

Leadership Class

Summit
~27000 GPUs

Simulations and Model Refinements

Feedback for control

OAK RIDGE
National Laboratory

OAK RIDGE National Laboratory | CENTER FOR NANOPHASE MATERIALS SCIENCES

Thank you