HUGS2021: Data Science (2)

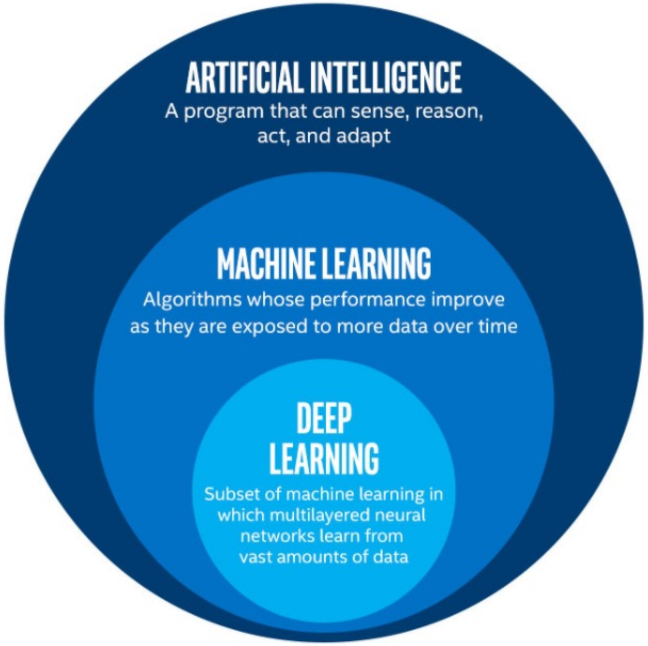

Malachi Schram, Ph.D.
Department of Data Science
Thomas Jefferson National Accelerator Facility

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

JSA

Goal of this talk

- Provide a high level overview on optimization for design and control
  - This is another massive research areas and there is no way to cover a fraction of these topics in a few lectures
- Provide some resources to get you started
  - Python centric ... sorry
- Cover some terminology
- Hopefully get you excited ☺

What is Data Science?

- Interdisciplinary field that leverages computer science, mathematics, and domain expertise to extract knowledge and insights from data
- Collaborative effort built on teams of experts

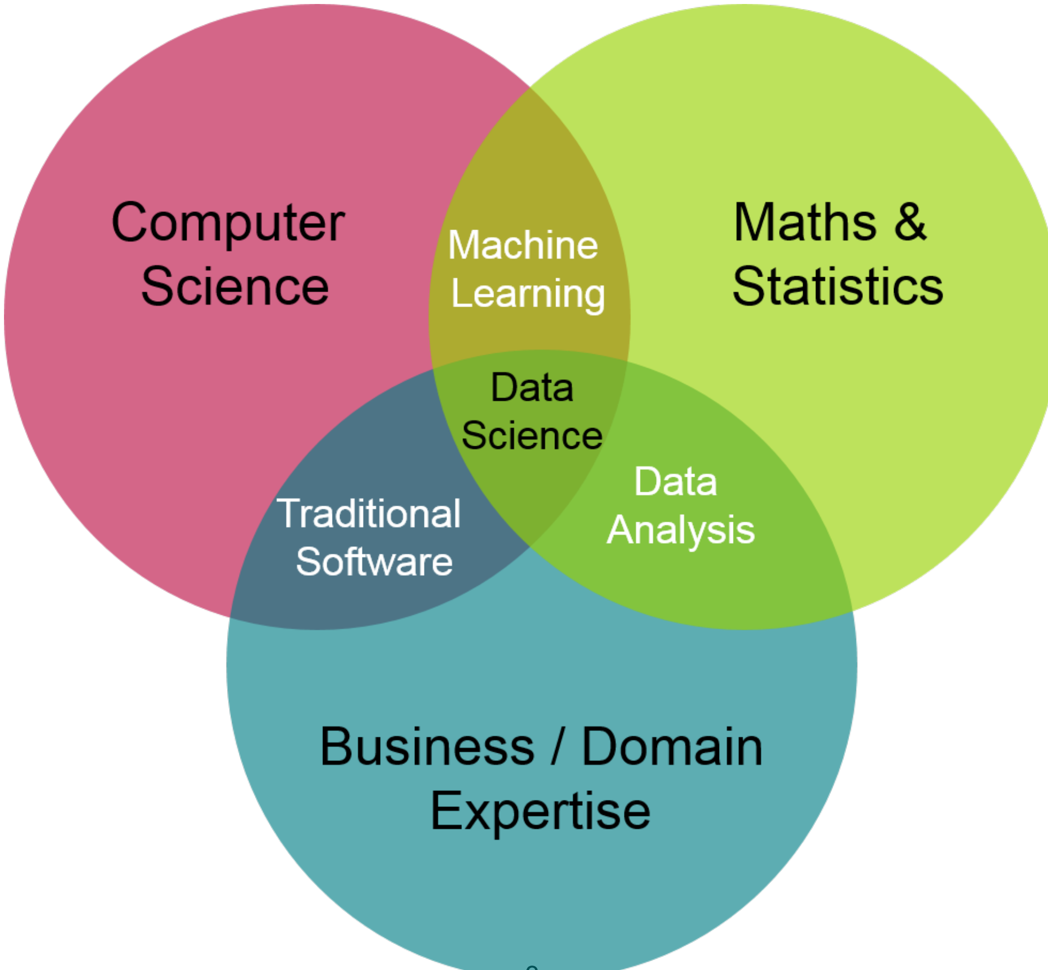

AI for Science Report

"New Deep Learning methods are required to ***detect anomalies*** and ***optimize operating parameters***..."

"... move from ***human-in-the-loop to AI-driven*** design, discovery, and evaluation also manifests across the ***design of scientific workflows***, *optimization of large-scale simulation codes*, and ***operation of next generation instruments***."

- Excerpts from the Executive Summary

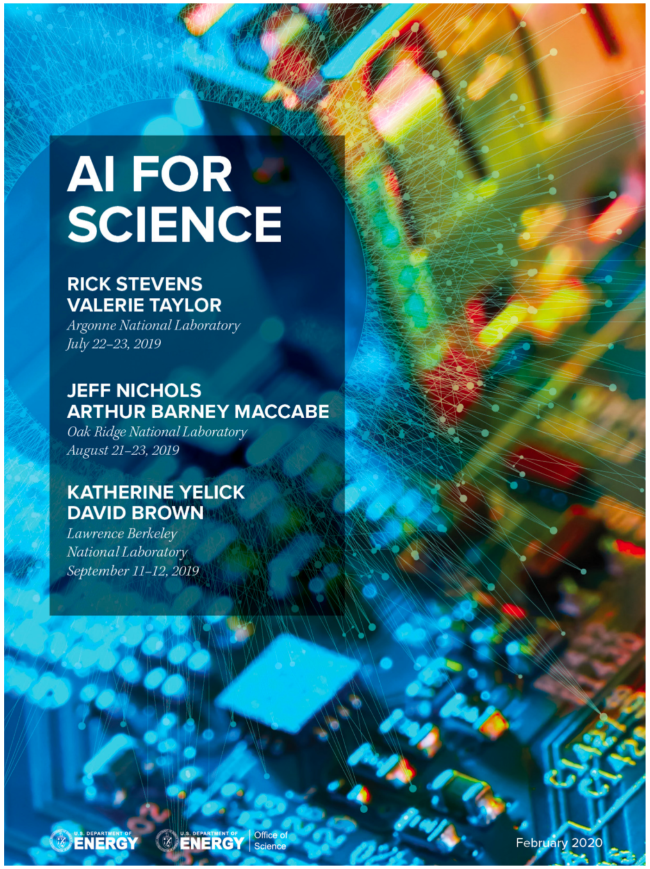

Definitions

Automated system is a sequence of instructions that executes a repeated set of processes.

Intelligent automated system is a sequence of processes which includes components of artificial intelligence and machine learning to improve the overall outcome.

Autonomous system is a sequence of processes dictated by an artificial intelligence.

- SAE defines levels of autonomous systems
- Strict requirements are put in place to avoid catastrophic results

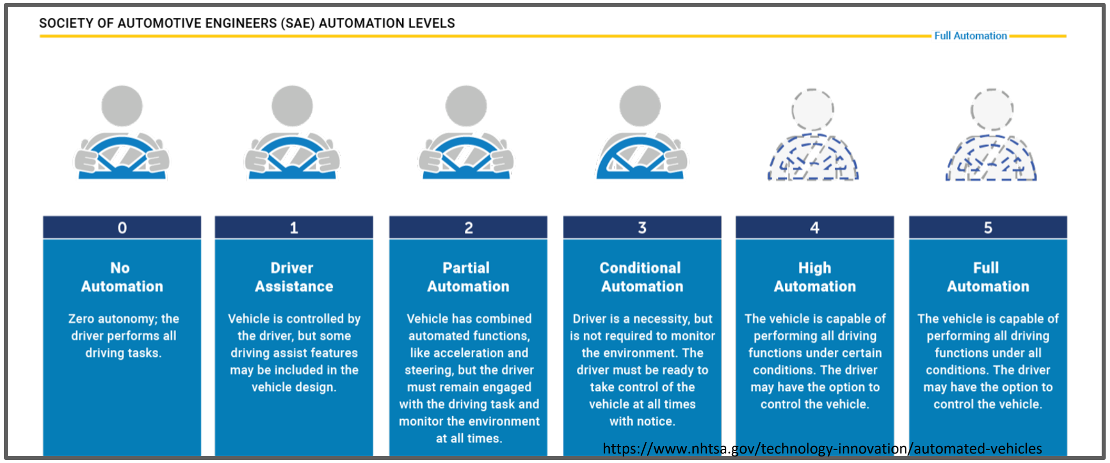

Data-driven automated workflows

- Automating a workflow is a critical component towards autonomous workflows
- Intrinsic value:
    - Increase efficiency, reduce tedious and laborious tasks, avoid mistakes, etc.
- However, there is nothing “intelligent” about it

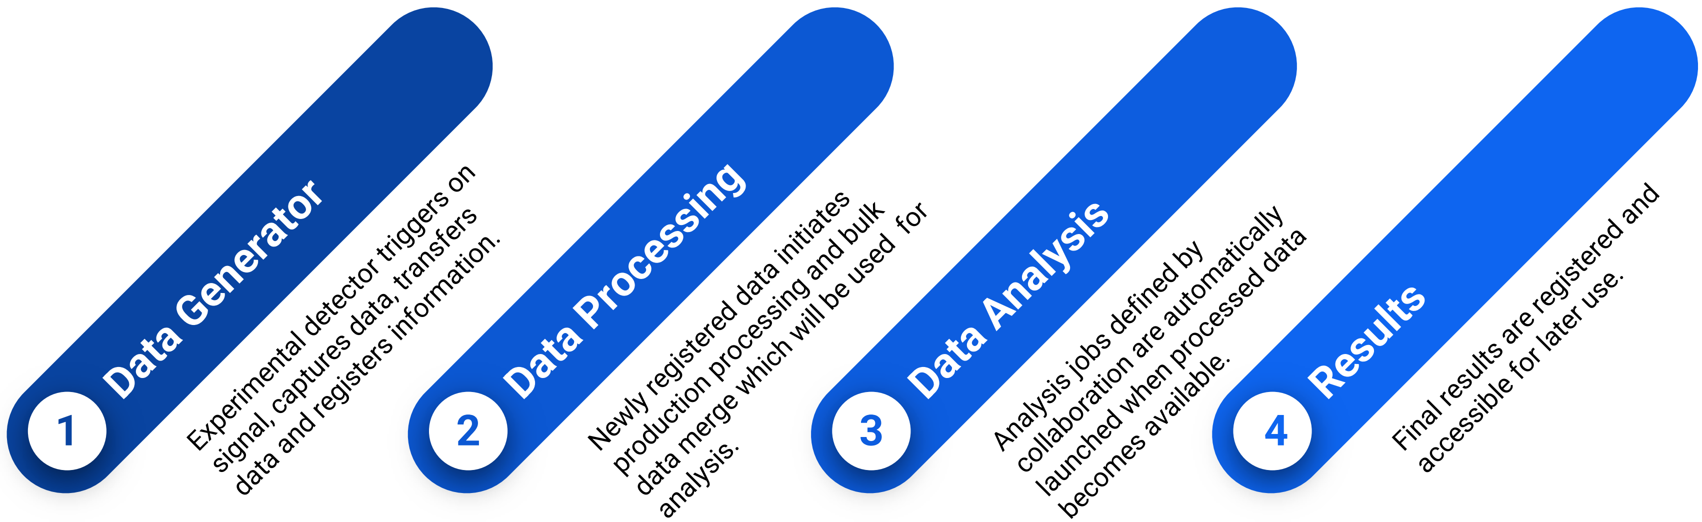

Jefferson Lab

Intelligent automated workflows

- Introducing an AI/ML component to improve automation workflow
- Requires a human-in-the-loop or rules to determine the appropriate action(s) to take
- Very well aligned with the current Scientific User Facilities AI/ML projects

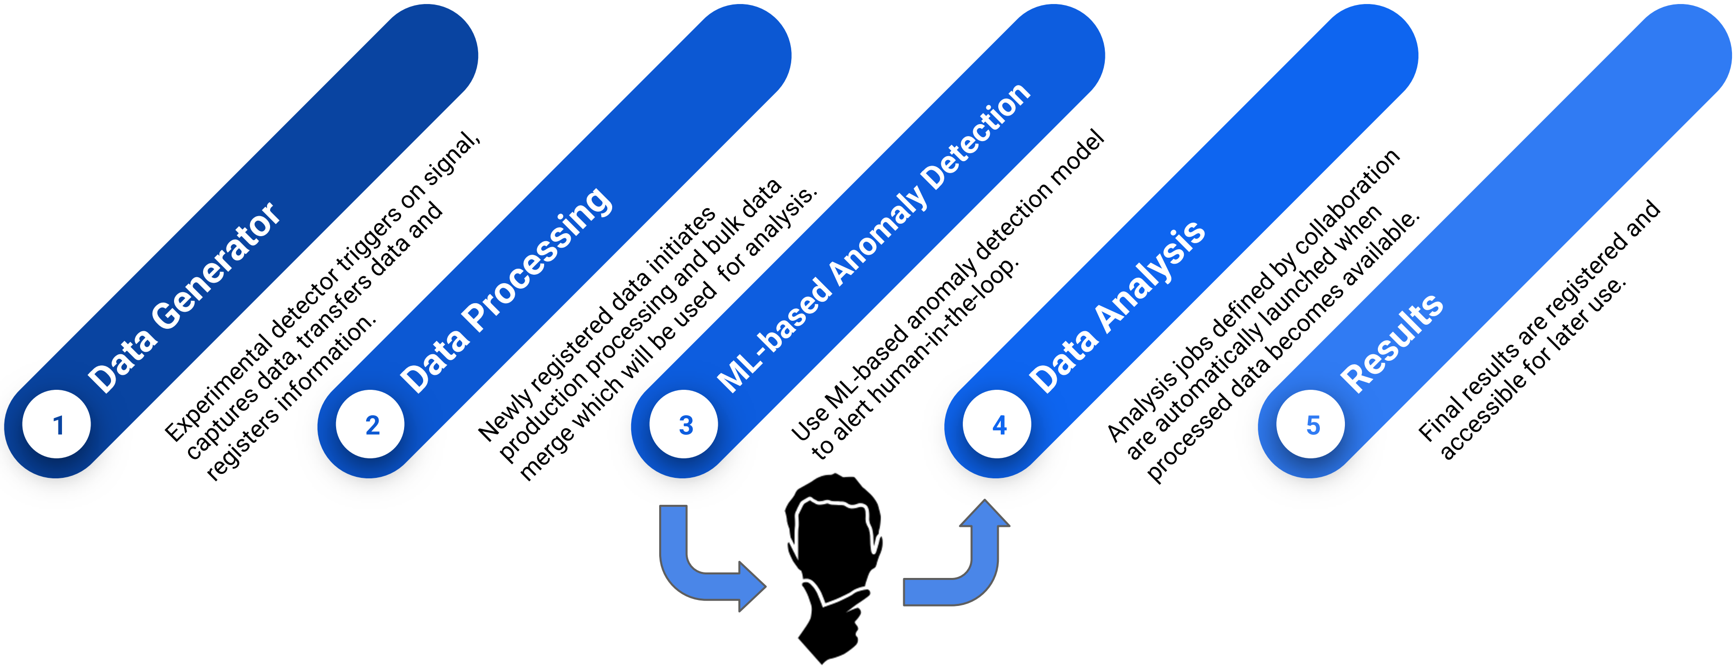

Autonomous control workflows

- Leverage AI/ML components to steer workflow
- No longer requires a human-in-the-loop or prescriptive rules, however, you might need multiple models

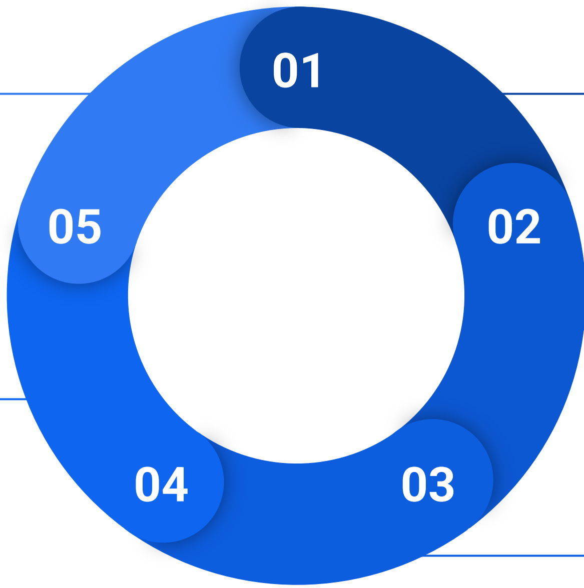

Data generation

Simulation and/or sensor data
Initial state or updated state based on new actions

Data processing

Data cleaning, re-structuring, etc.

ML-based augmented state representation (optional)

Enhance the state space representation using ML encoding, e.g. anomaly detection input, domain-aware models, etc.)

AI-based policy model

Continuous training/validating policy model

Control actions

Provide new actions for the system to perform

Optimizing the atomic inter-mixing properties

Problem definition:

Determine the optimal atomic inter-mixing between two materials to maximize interfacial thermal conductivity

1. What is the optimal intermixing fraction?
2. What is the optimal thickness of the intermixing region?

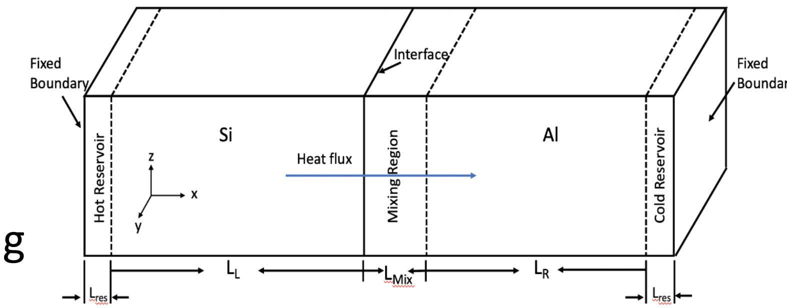

1 **Simulation input parameters** — Define a set of input parameters, appropriate physics model, material lattice configuration

2 **LAMMPS simulation** — Submit LAMMPS to cluster

3 **Thermal Conductance Analysis** — Process LAMMPS output to determine the thermal conductance for current setup

4 **Results** — Compare the new result with the previous results and determine next simulation (if needed)

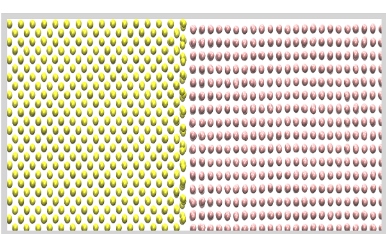

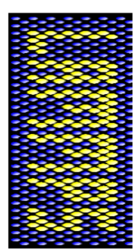

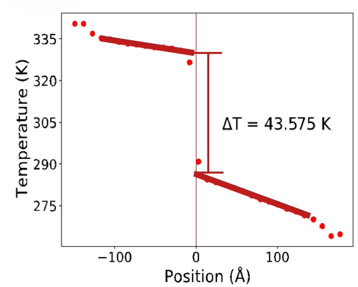

Jefferson Lab

Bayesian Optimization Workflow

Acquisition Function

Using quantile-based adaptive sampling

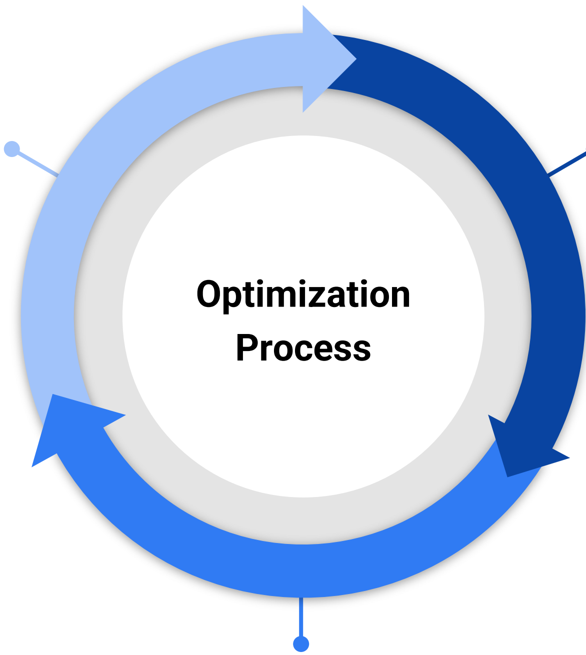

Data Generation

- **Initial Random Sampling**
- **Based on acquisition function**

Model Fitting

Gaussian Process Regressor

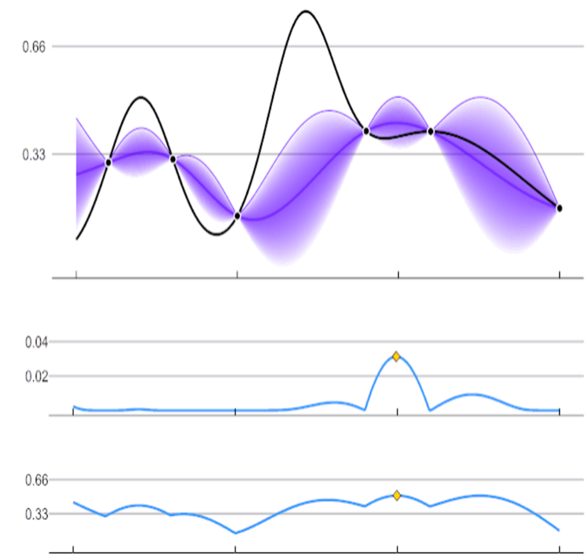

Scaling Bayesian Optimization on High Performance Computing system

- To accelerate the data generation we developed a MPI-based framework
- We created a single Bayesian Optimization agent that maps the action-reward for all simulations
- We created an OpenAI gym environment as a wrapper around the LAMMPS and the post simulation analysis
- A production job split the MPI communications between the BO agent and each LAMMPS environment

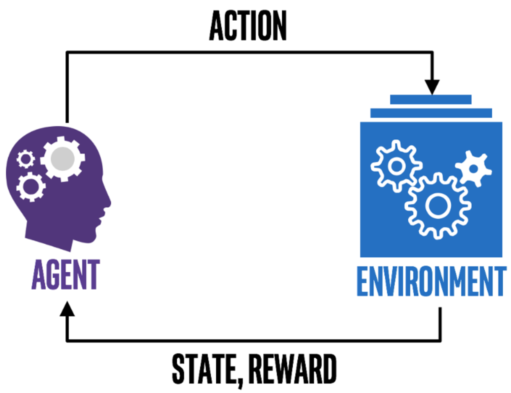

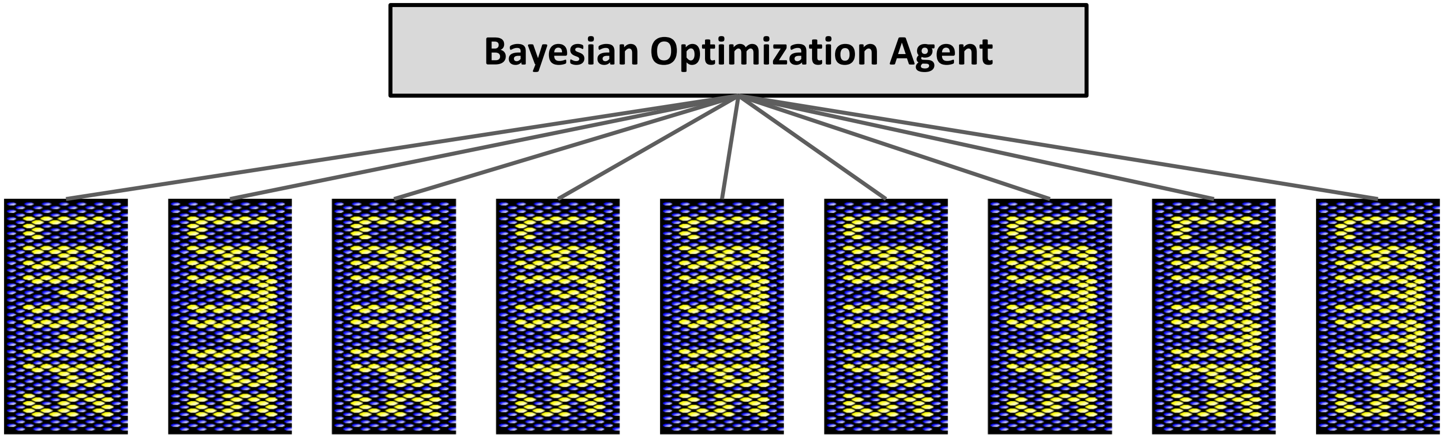

Quantile-based adaptive sampling

Algorithm 1: Quantile-based adaptive sampling

Initialize replay memory ($\mathcal{D}$) and action space ($\mathcal{A}$);

Initialize the quantile threshold value ($\mathcal{Q}$) and decay rate (ϵ);

for *trial* $= 1, T$ **do**

    if $t > m$ **then**

        Perform a gaussian process fit on $\mathcal{D}$;

        Calculate upper confidence bounds (UCB) using full action space;

        Calculate the UCB quantiles;

        Update $\mathcal{A}$ based on UCB quantiles satistying $\mathcal{Q}$;

    end

    Select random action a from $\mathcal{A}$;

    Execute a on environment and calculate the reward (r);

    Store transition (a,r) in $\mathcal{D}$;

    Update the $\mathcal{Q} = \mathcal{Q} \times \epsilon$;

end

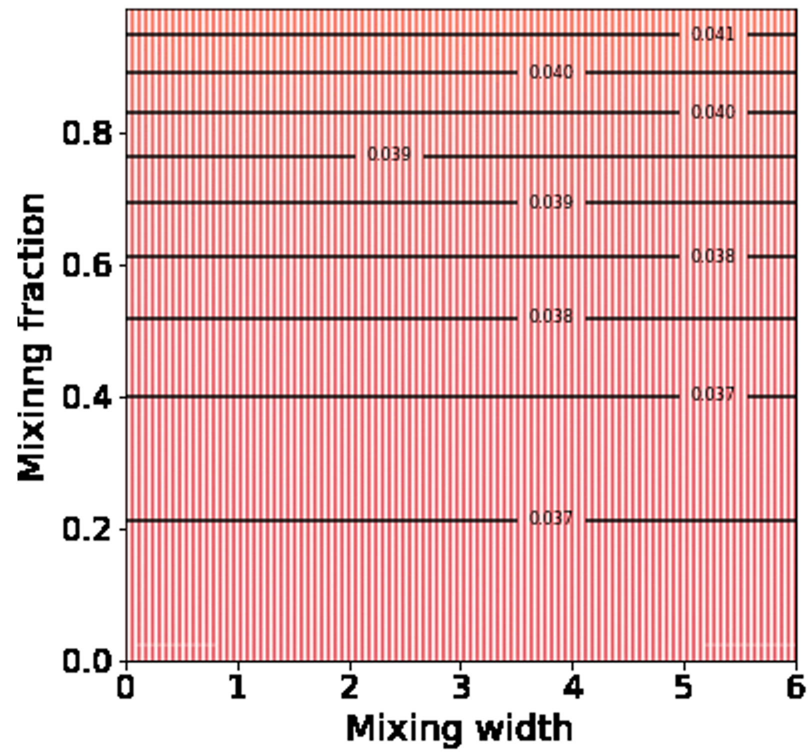

- Optimization improved the desired characteristics by nearly 50% relative to the pristine interface!

Example of an autonomous control workflow for FNAL Booster control policy

Problem definition:

Reduce beam losses in the FNAL Booster by developing a Machine Learning (ML) model that provides an optimal set of actions for accelerator controls

FNAL Accelerator Complex:

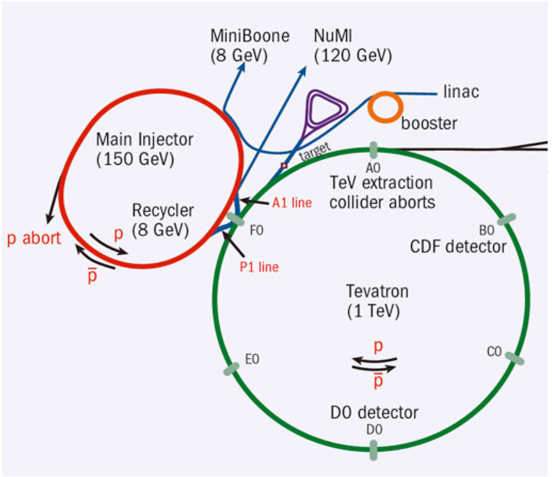

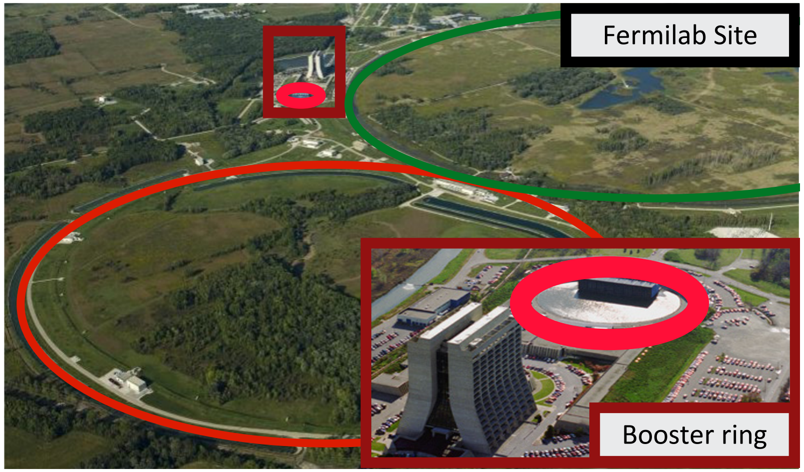

Courtesy: Christian Herwig

The Need for Improving Regulation

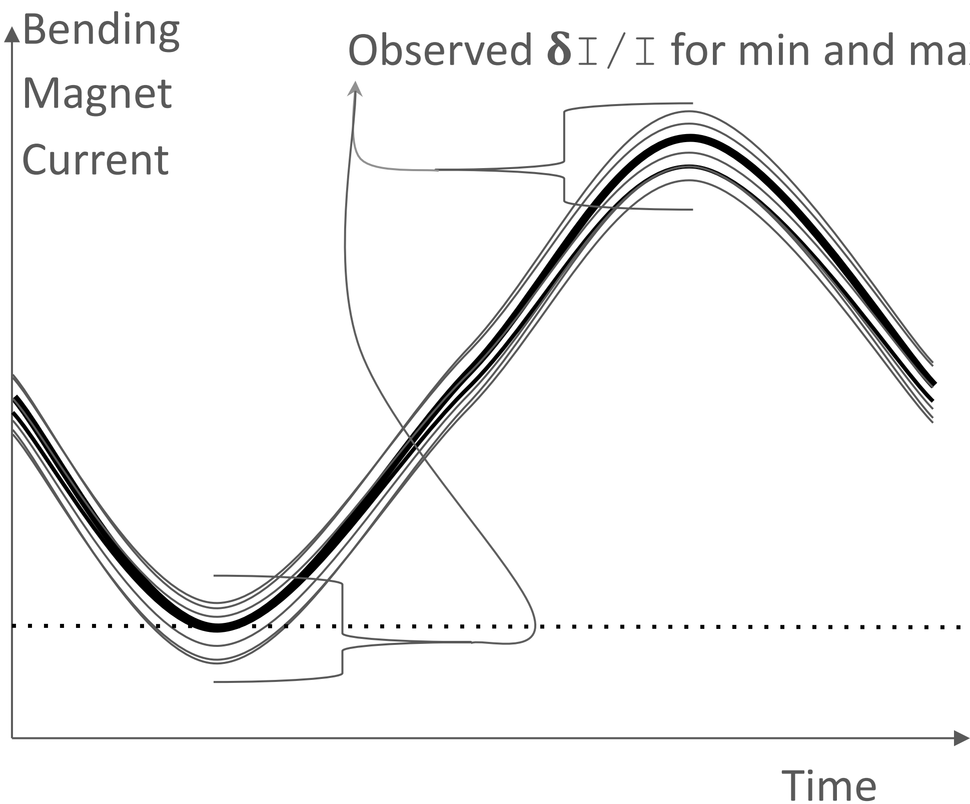

- Spread in B-field degrades beam quality and contributes to losses
- Focusing on min for now:
  B_VIMIN = Setting to achieve
  B:VIMIN = Prescribed remedy from PID regulator circuit
  B:IMINER = Error discrepancy
- Policy model is focused on controlling the regulator to reduce the error

Proof of Concept Workflow

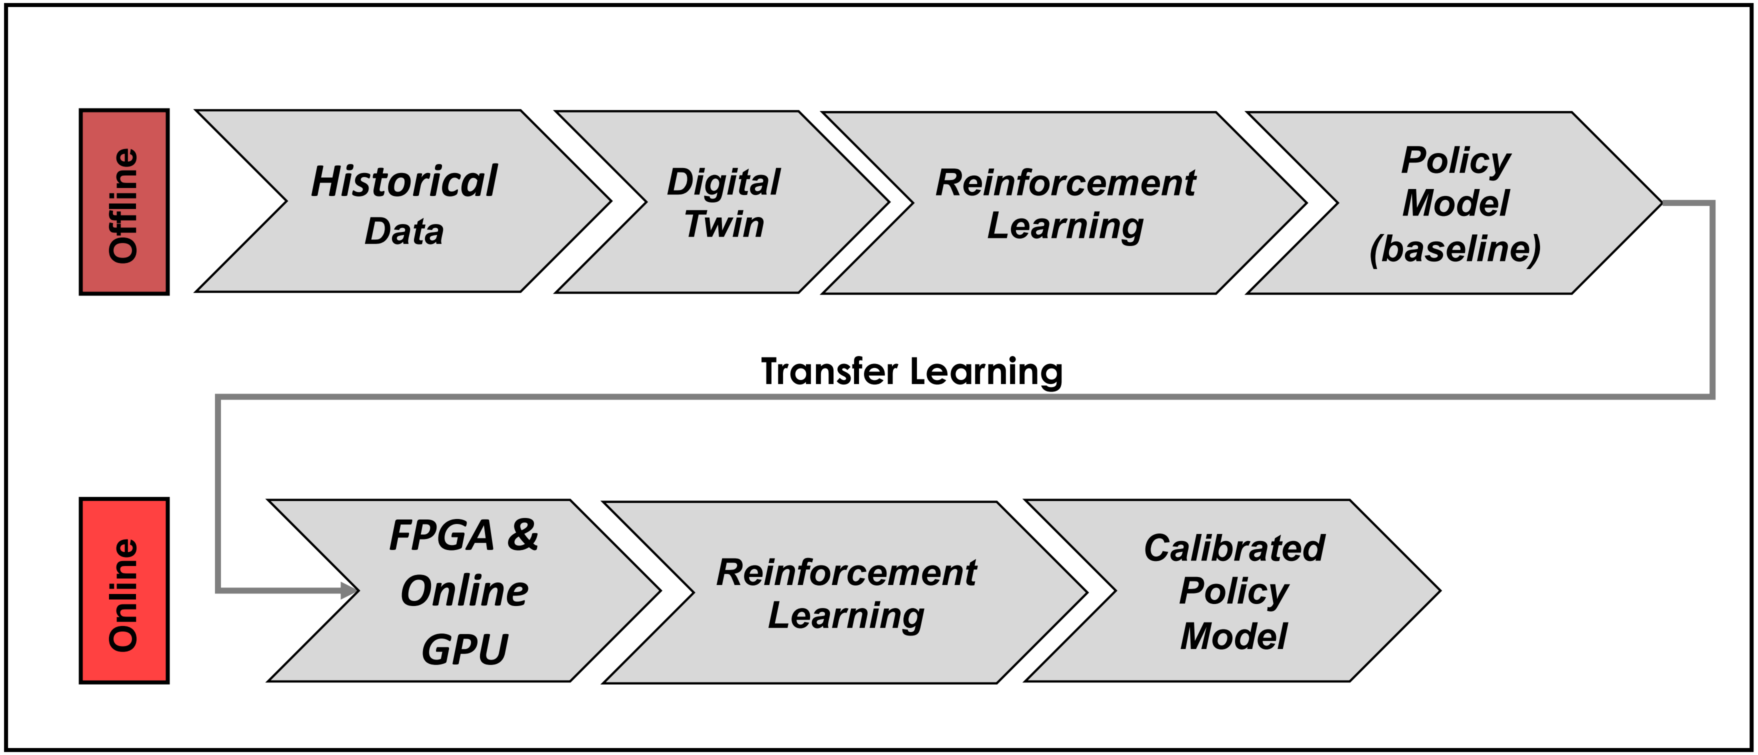

Machine learning-based “digital twin”

Scope and usage for digital twin:

- Provide accurate predictions of future time for key variables to be used by the reinforcement learning framework

Dataset provided:

- Historical temporal information from key variables was available based on subject matter expert input
- Caution:
  - Data did not include detailed history on commissioning, maintenance, etc.
  - Should conduct a full data inventory assessment

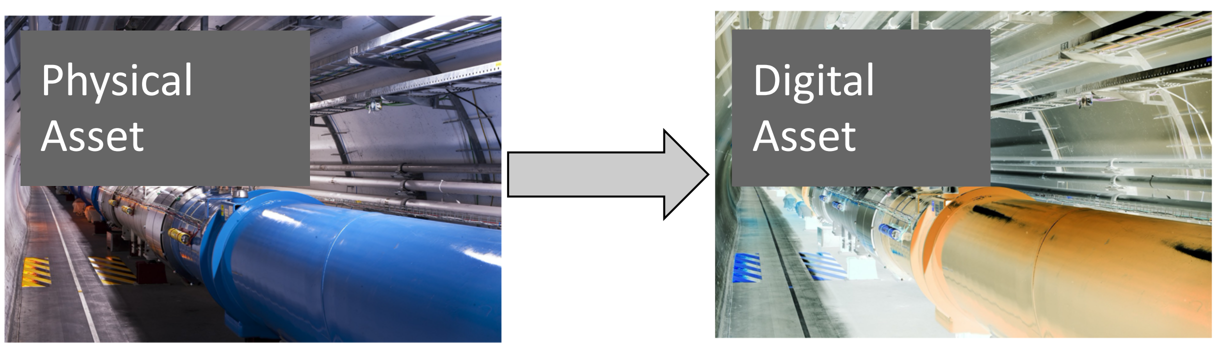

Digital Twin: Multivariate Stacked LSTM

Data Sources → Data Preparation Tools → AI/ML Applications → Robustness Tools → Results

Data Preparation Tools
- Cleaning
- Correlation studies
- Frequency analysis
- Dimensionality reduction

AI/ML Applications
- Recurrent Neural Networks
- Structural time series models
- Domain aware ML
- Bayesian regressors

Robustness Tools
- Validation
- Verification
- Integration
- Deployment

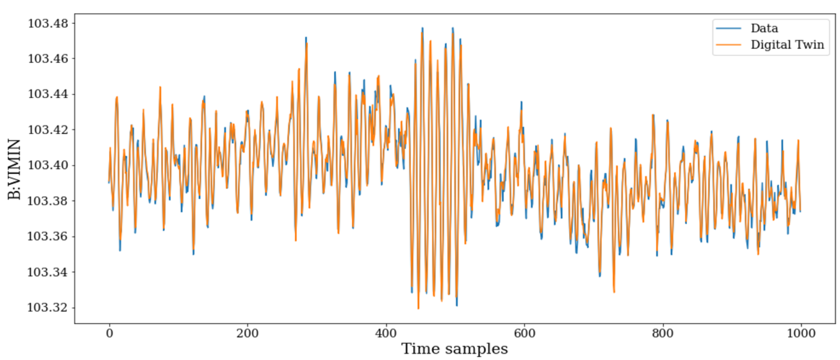

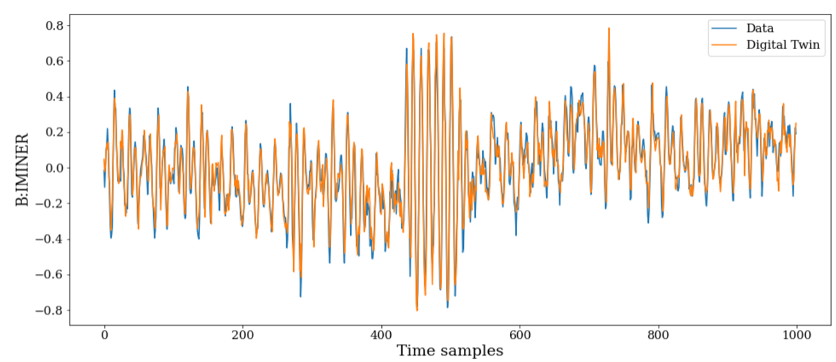

Autonomous System using Reinforcement Learning

“Reinforcement Learning is learning what to do — how to map situations to actions—so as to maximize a numerical reward signal. The learner is not told which actions to take, but instead must discover which actions yield the most reward by trying them.” - Barto & Sutton

Key concepts to Reinforcement Learning:

- Agent (controller – policy and sampling)
  - *Action* (control signal)
- Environment (controlled system)
  - *State* (representation of environment)
  - *Reward* (numerical consequence of action)
- Sequence of experience and agent forms trajectory:
  Example RL Trace: (S_0, A_0, R_0), (S_1, A_1, R_1), …

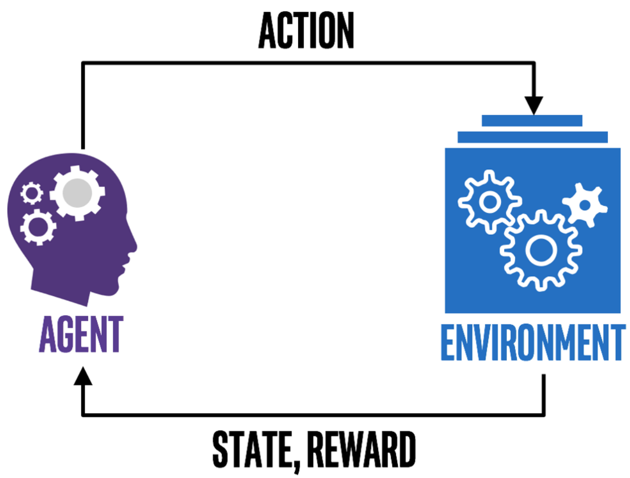

Deep Q-Networks Reinforcement Learning Algorithm

DQN uses a deep neural network to estimate the value of taking a specific action at a certain state, also called the state-action value or Q-value.

The DQN agent, once trained properly, suggests the action with the highest Q-value as its policy, and maximizes the total reward over the episode.

Five discrete actions were defined as possible control changes to the regulator.

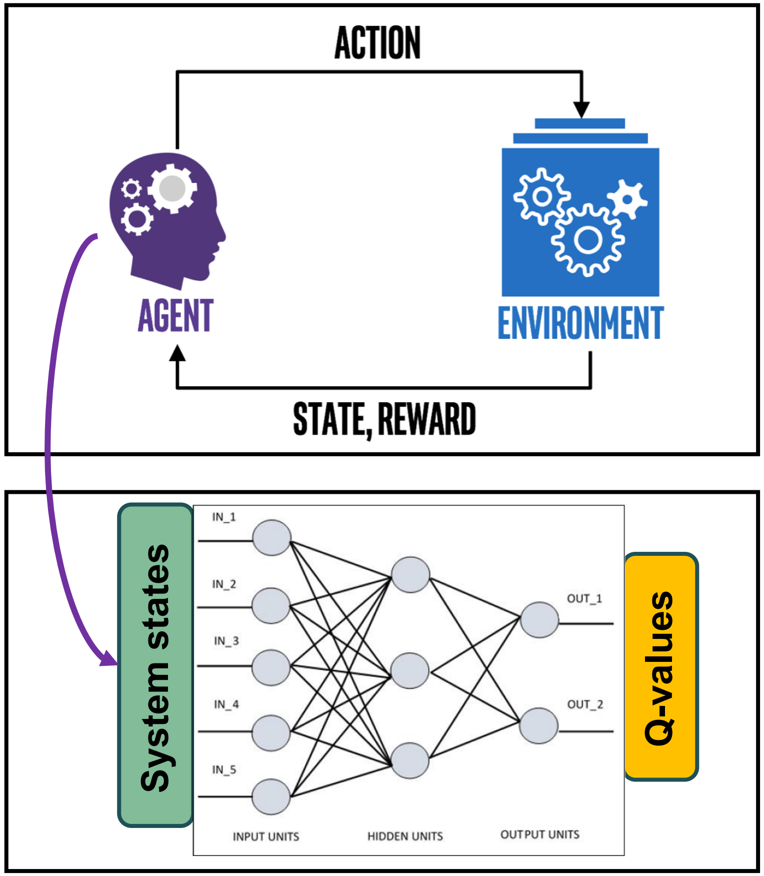

Reinforcement Learning FNAL Booster Workflow

The optimization was formulated as an episodic problem:

- An episode is composed of 50 sequential steps
- After each episode the environment was reset to the same initial state
- A batch size of 32 experiences were randomly sampled to train the active policy model
- A ϵ-greedy method was used to control the level of exploration/exploitation

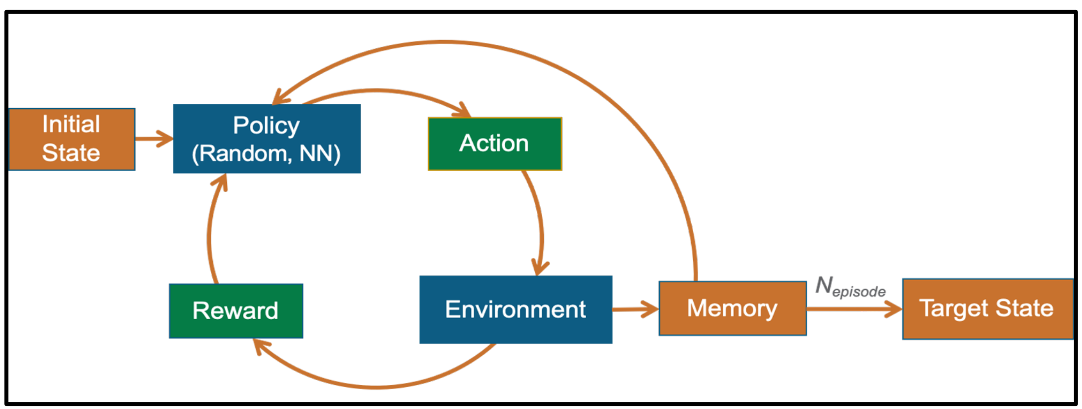

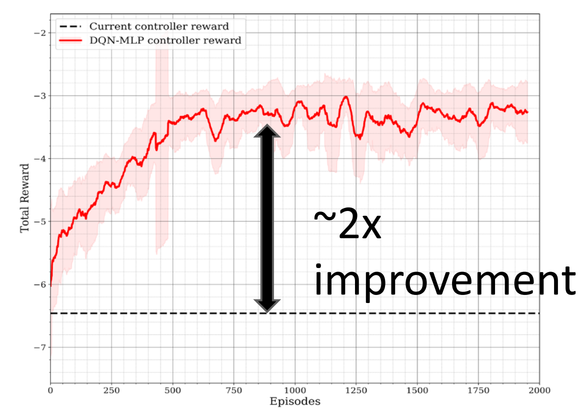

In summary

- DOE is making investments in AI/ML
- There are a lot of exciting opportunities to develop methods to accelerate science using AI/ML
- My current research ties into all components of the workflow

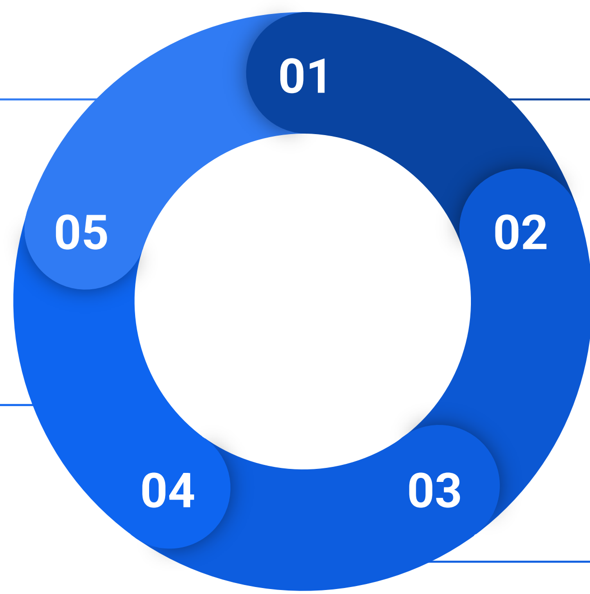

Data generation

- Chemistry/material data
- Accelerator data
- Experimental data
- Sensor data

Data processing

- Several common tools for all applications

ML-based augmented state representation (optional)

- Anomaly/prognostication and control
- Object identification
- Molecular descriptors via domain aware ML

AI-based policy model

- Scaling learning algorithms and workflows
- RL algorithm development

Control actions

- Scaling learning algorithms and workflows
- RL algorithm development

Malachi Schram, Ph.D.
Department of Data Science
Thomas Jefferson National Accelerator Facility
schram@jlab.org

- Next session we will go over some hands-on examples

- New data science position at JLab:
  https://careers.peopleclick.com/careerscp/client_jeffersonlab/external/jobDetails.do?functionName=getJobDetail&jobPostId=1910&localeCode=en-us