15TH INTERNATIONAL CONFERENCE ON HEAVY ION ACCELERATOR TECHNOLOGY

Argonne NATIONAL LABORATORY | 75 1946–2021

REINFORCEMENT LEARNING AND BAYESIAN OPTIMIZATION FOR ION LINAC OPERATIONS AT ATLAS

JOSE L. MARTINEZ-MARIN
Postdoc
Physics Division
Argonne National Laboratory

HIAT 2022 | Darmstadt

U.S. DEPARTMENT OF ENERGY
Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

June 30th, 2022
HIAT 2022 (Virtual)

OUTLINE

✓ATLAS AI/ML Project
- Brief Description
- Main Objectives & Approach

✓ Progress
- **Data Collection**
- **Bayesian Optimization with Gaussian Processes** to support online tuning
- **Deep Reinforcement Learning** to support online tuning
- Surrogate models for speeding simulations

✓Conclusions and Next Steps

ARGONNE TANDEM LINEAR ACCELERATOR SYSTEM

- ✓ 1st Superconducting heavy-ion linac in the world
- ✓ It has been operating for over 35 years
- ✓ National user facility serving ~ 400 users per year

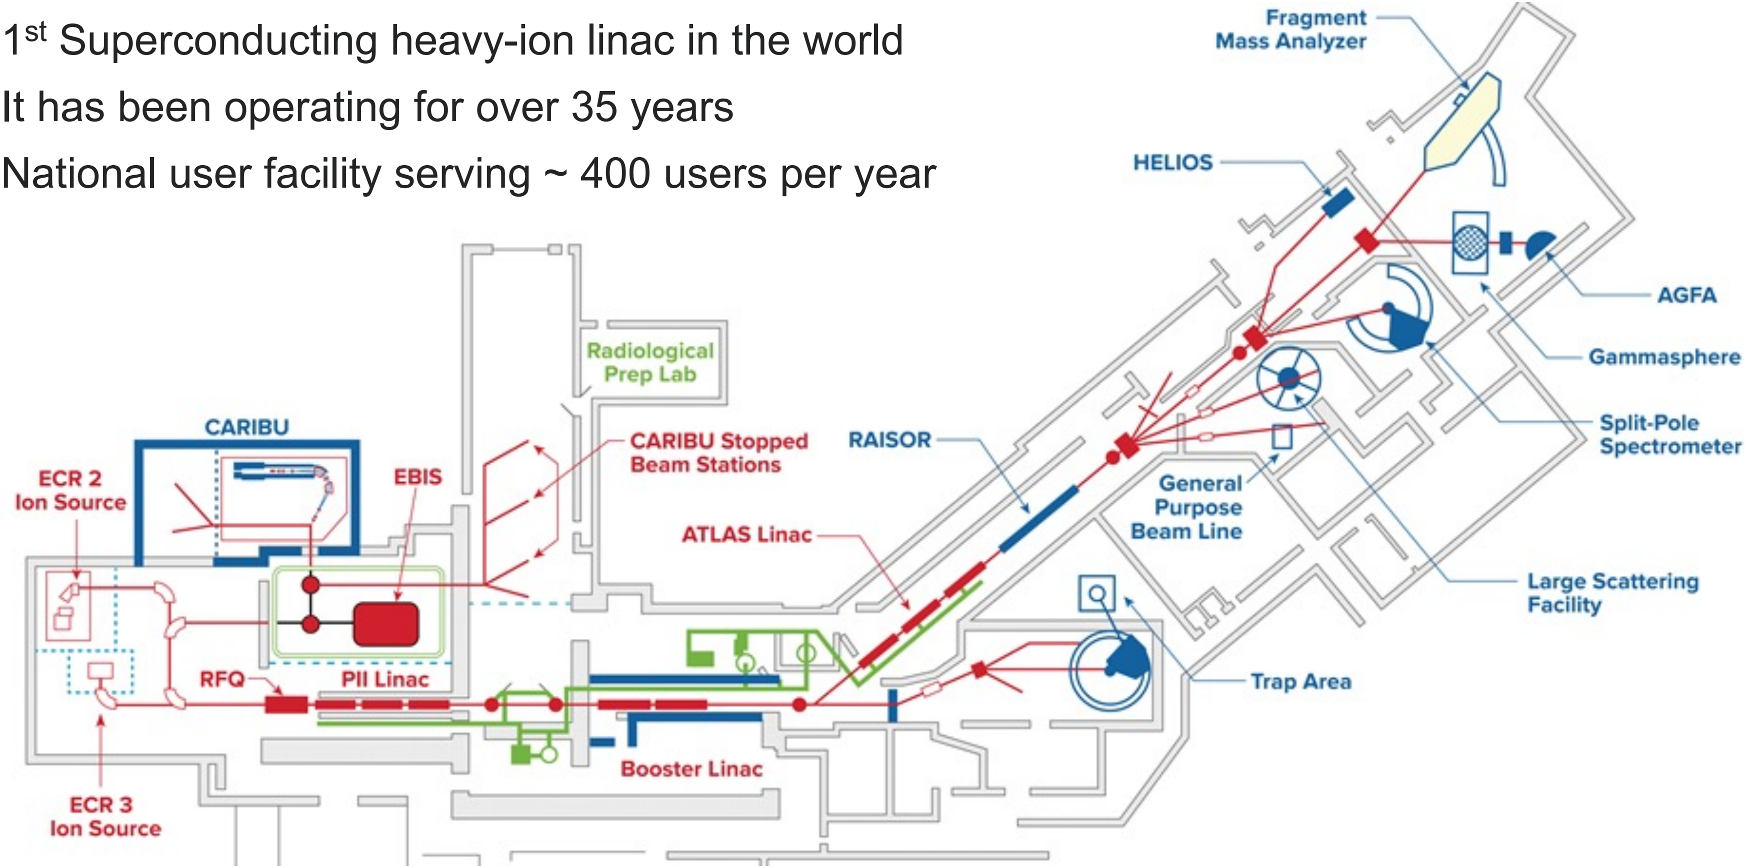

THE ATLAS AI / ML PROJECT

Use of artificial intelligence to optimize accelerator operations and improve machine performance

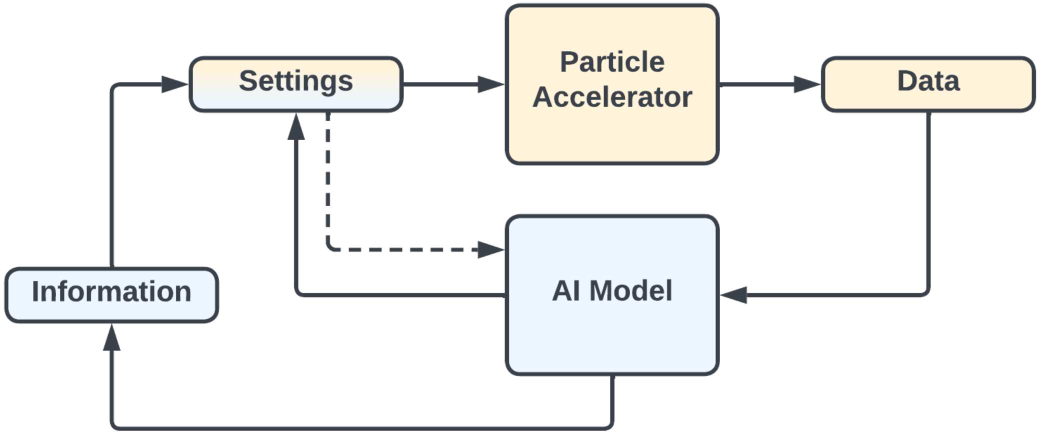

- ✓Surrogate Models
- ✓Virtual Diagnostics
- ✓Tuning Control Model
- ✓…

THE ATLAS AI / ML PROJECT

Use of artificial intelligence to optimize accelerator operations and improve machine performance

- ✓ At ATLAS, ion beam species are switched every 3-4 days … → Using AI could streamline beam tuning & help improve machine performance
- ✓ The main project goals are:
  - Data collection, organization and classification, towards a **fully automatic and electronic data collection** for both machine and beam data
  - **Online tuning model to optimize operations and shorten beam tuning time** in order to make more beam time available for the experimental program
  - <u>Virtual model to enhance our understanding of the machine behavior</u> in order to improve performance and optimize particular and new operating modes

Project Started in 2021

ATLAS – FIRST STEPS IN DATA COLLECTION

~80% time of a Data Scientist is Collecting Data, Cleaning and Organizing Data

- ✓ Kind of data?
- ✓ How much data?
- ✓ Accessible?
- ✓ Automated?

ATLAS - DATA AT ATLAS

~80% time of a Data Scientist is Collecting Data, Cleaning and Organizing Data

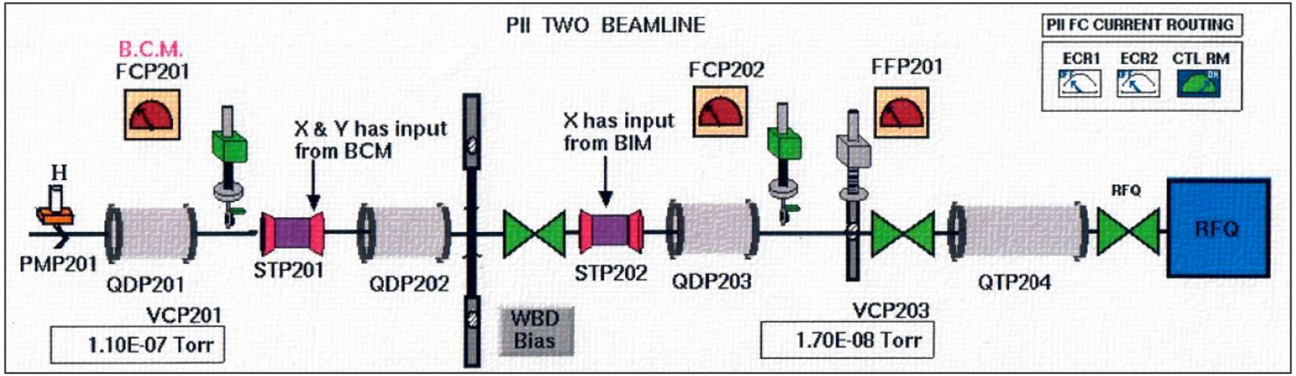

ATLAS sub-section

ATLAS - DATA AT ATLAS

~80% time of a Data Scientist is Collecting Data, Cleaning and Organizing Data

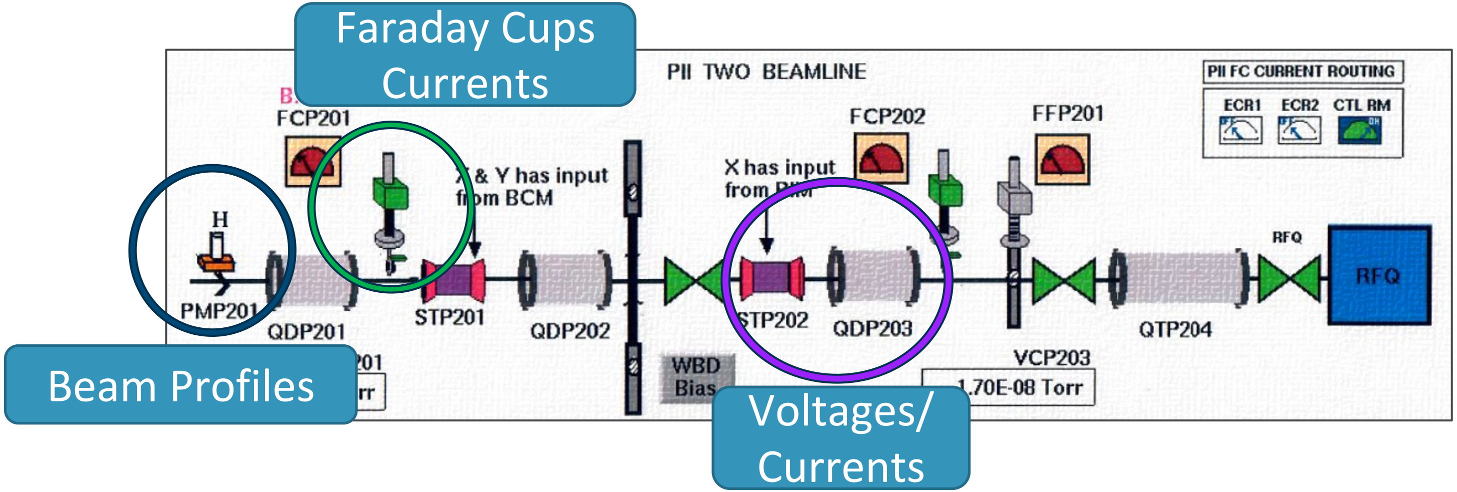

ATLAS - DATA AT ATLAS

~80% time of a Data Scientist is Collecting Data, Cleaning and Organizing Data

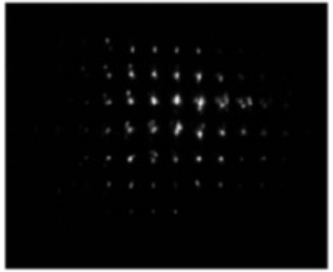

Pepper Pot Images

Faraday Cups Currents

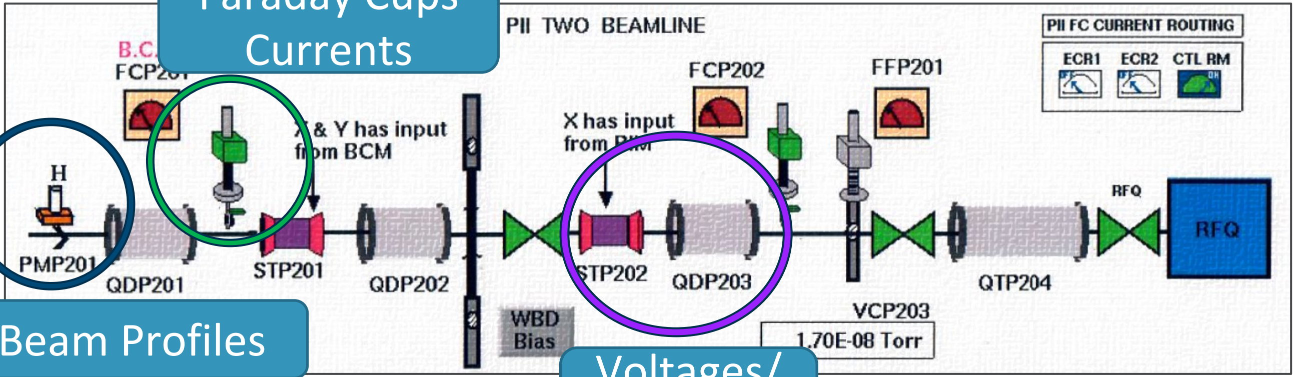

Beam Profiles

Voltages/ Currents

ATLAS - DATA AT ATLAS

~80% time of a Data Scientist is Collecting Data, Cleaning and Organizing Data

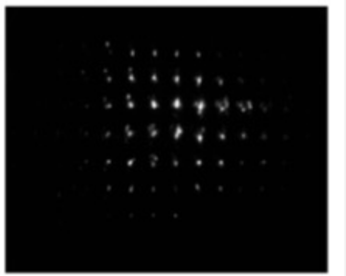

Pepper Pot Images

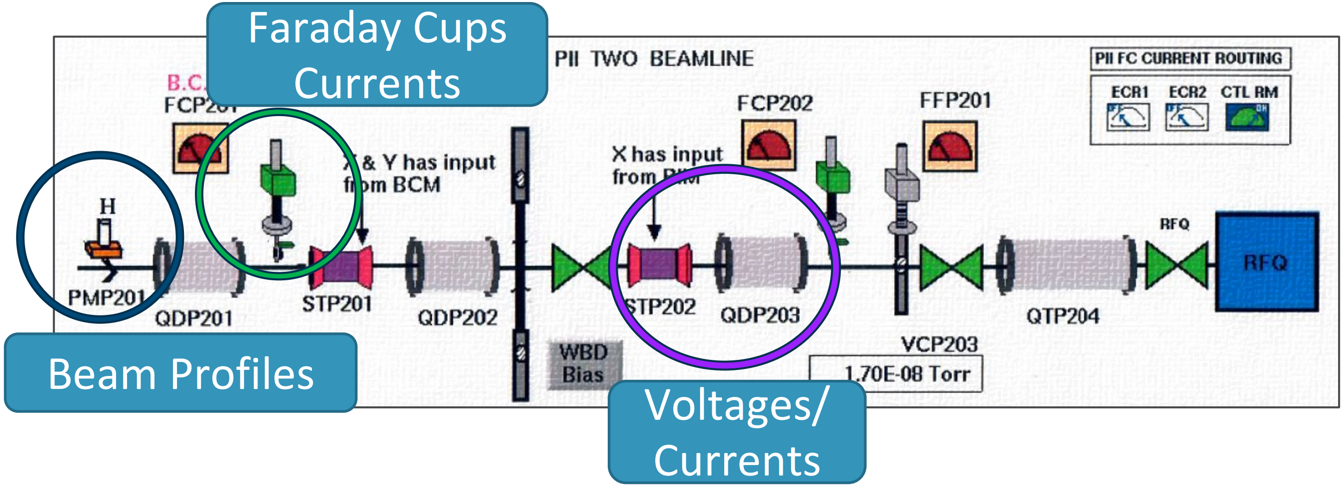

- ✓ Only settings could be saved automatically using the Control System (vsystem)
- ✓ Faraday Cups and Beam Profile Monitor in Control System but not automated

ATLAS - DATA COLLECTION NOW

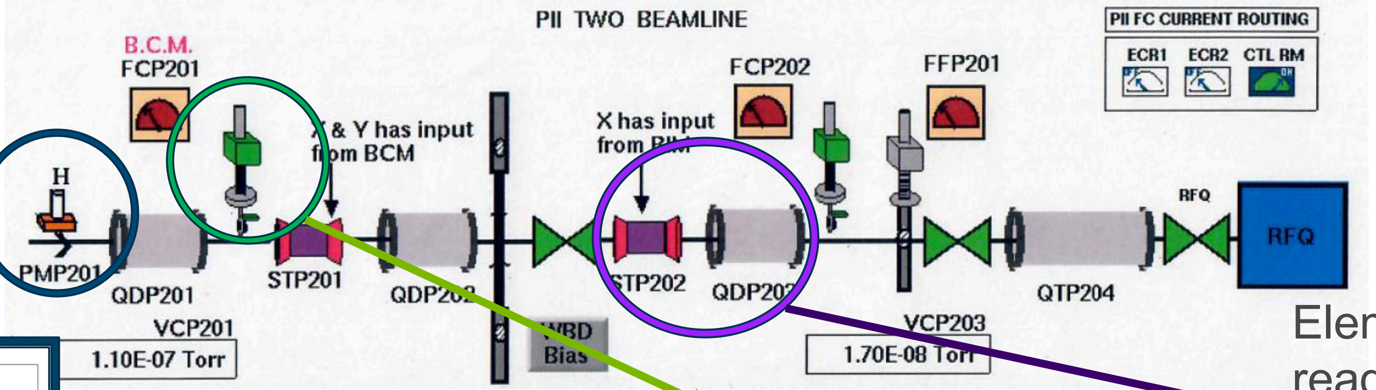

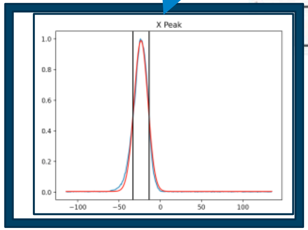

BPM: digitized, read

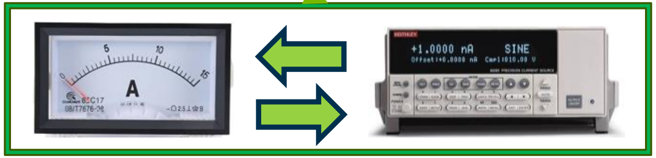

FC: digitized, insert, read

Elements: read/set

*Pepper pot setup in progress

ATLAS - DATA COLLECTION

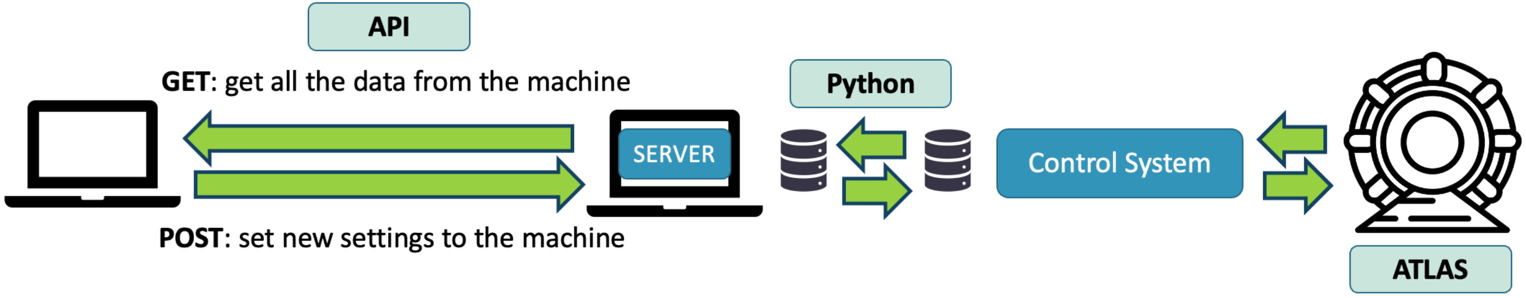

SIMULATION - DATA COLLECTION

- ✓ Python wrapper for TRACK (Simulation Code)
- ✓ Generation of data easily
- ✓ Different conditions and inputs
- ✓ Integration with modeling

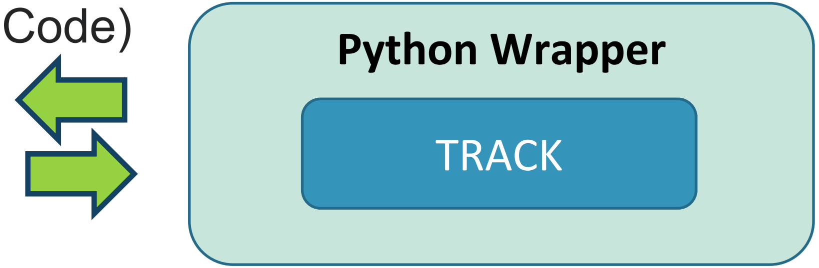

TUNING/CONTROL OF ATLAS

Online tuning model to optimize operations and shorten beam tuning time

- ✓ Develop a baseline model to tune/control a small section of ATLAS Linac using Simulation Data (from TRACK simulation code)
- ✓ Followed Approaches: **Bayesian Optimization with Gaussian Processes** and **Deep Reinforcement Learning**
- ✓ Test models on real machine
- ✓ Improve models
- ✓ Expand to other parts of the Linac

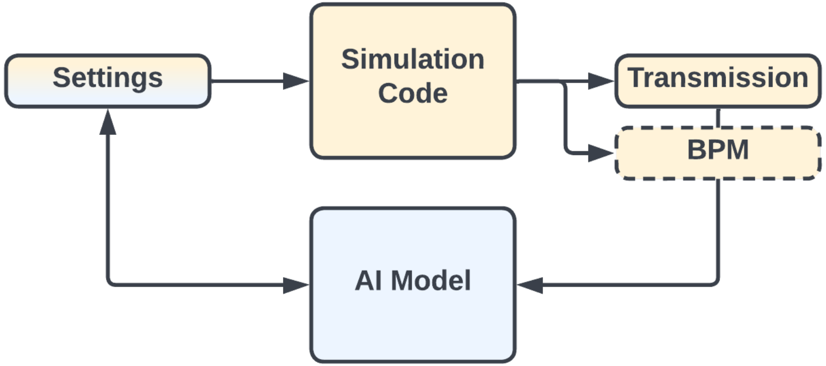

TUNING/CONTROL OF ATLAS

Online tuning model to optimize operations and shorten beam tuning time

- ✓ Develop a baseline model to tune/control a small section of ATLAS Linac using Simulation Data (from TRACK simulation code)
- ✓ Followed Approaches: **Bayesian Optimization with Gaussian Processes** and **Deep Reinforcement Learning**
- ✓ **Test models on real machine**
- ✓ Improve models
- ✓ Expand to other parts of the Linac

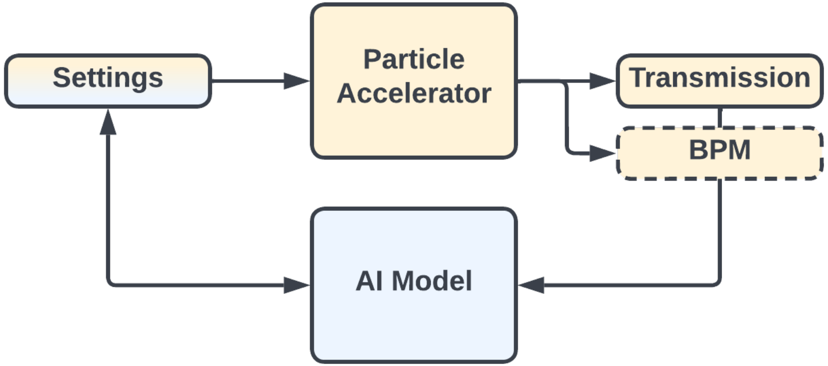

BAYESIAN OPTIMIZATION

- Model-based Bayesian optimization **combines the complementary strengths of human and numerical optimization**.
  - Life-long learning, learns by experience / Juggle many things at once, Fast decisions + estimate of their own uncertainty + global optimum in a minimum number of steps.
- Bayesian optimization incorporates prior belief about f(x) and updates the prior with samples drawn from f(x) to get a posterior that better approximates f(x)

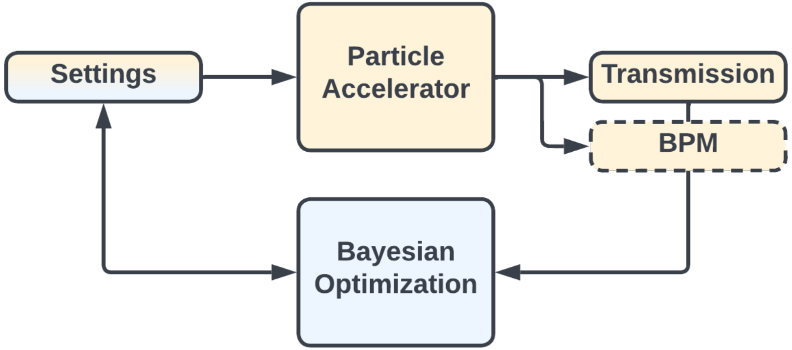

BAYESIAN OPTIMIZATION

- **Probabilistic surrogate model** for approximating the objective function.
  - Gaussian Process (GP): give a reliable estimate of their own uncertainty and shape our prior belief via the choice of kernel.
- **Acquisition function** that tells where to query the system next for a more likely improvement

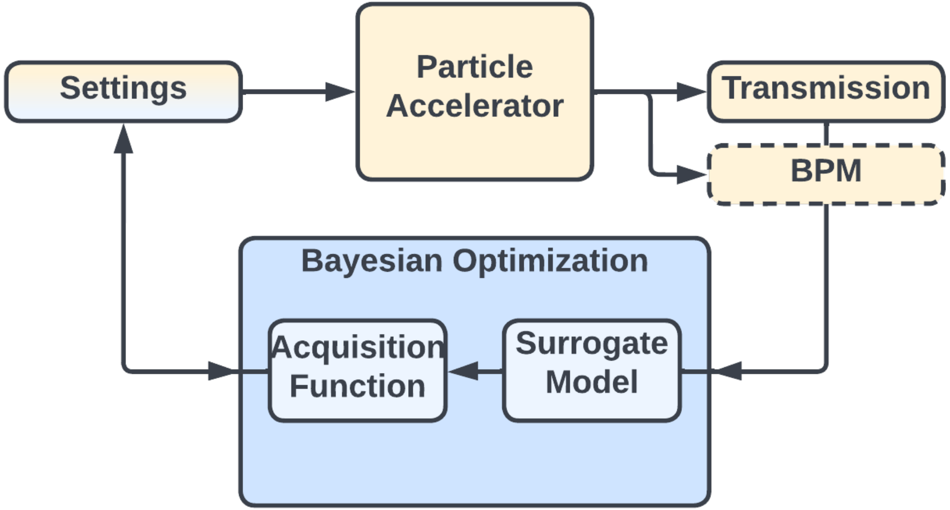

GAUSSIAN PROCESS

- Non-parametric approach: <u>learning model and hyperparameters from data</u>.
- It finds a distribution over the possible functions f(x) that are consistent with the observed data
- Begins with a prior distribution, which can be converted into a posterior over functions by observing more data *Bayes' rule*.
- Example using a Gaussian Kernel and assuming a mean of 0 for prior:

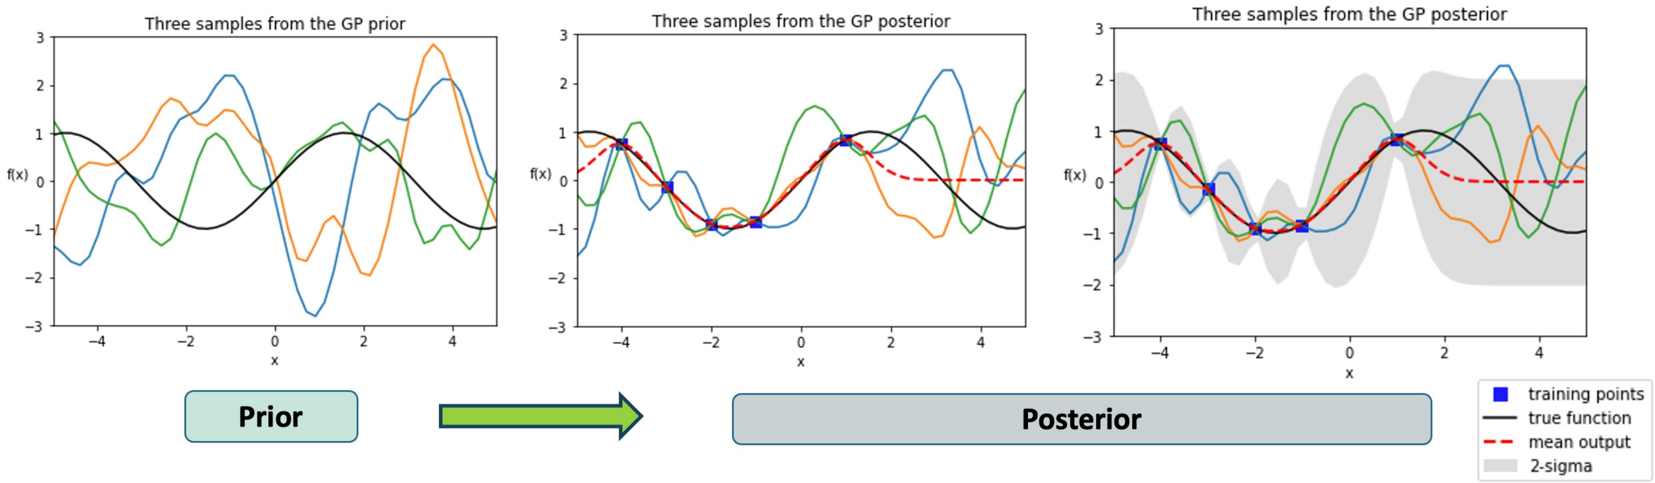

BAYESIAN OPTIMIZATION WITH GP

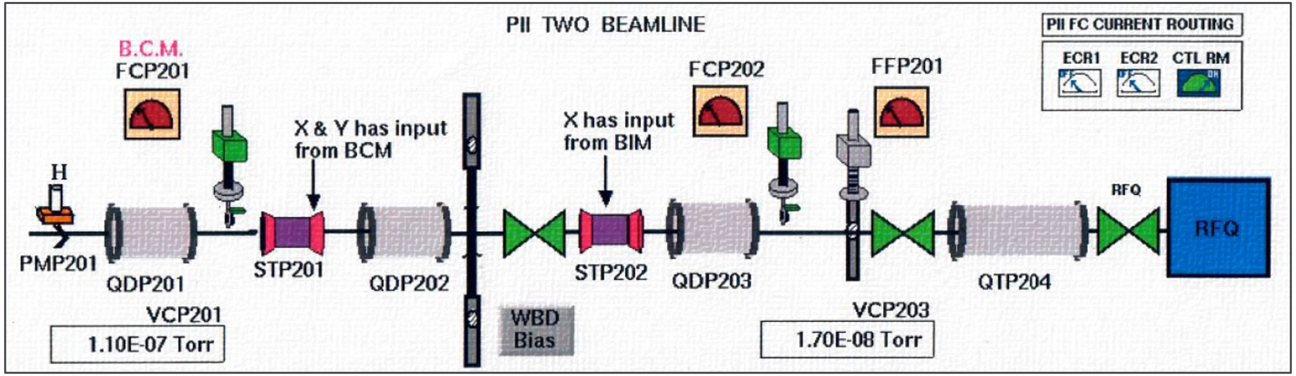

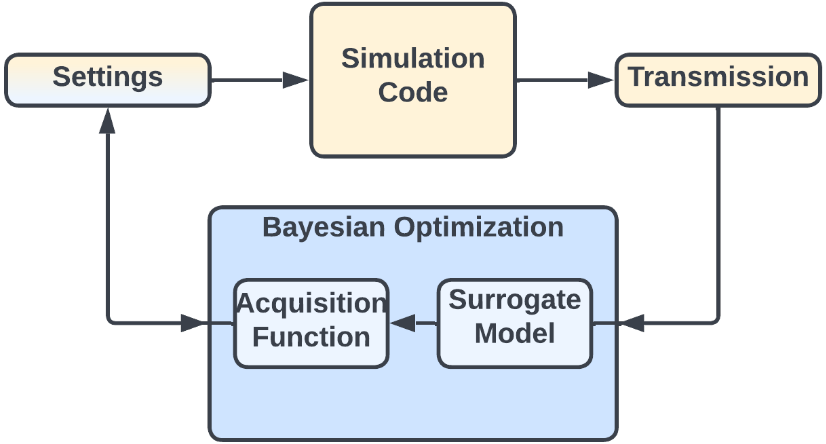

BO WITH SIMULATION DATA CASES

No RFQ

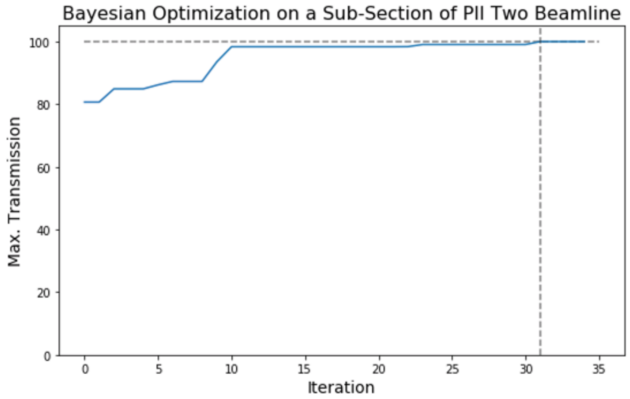

Random Initial Data

BO WITH SIMULATION DATA CASES

No RFQ

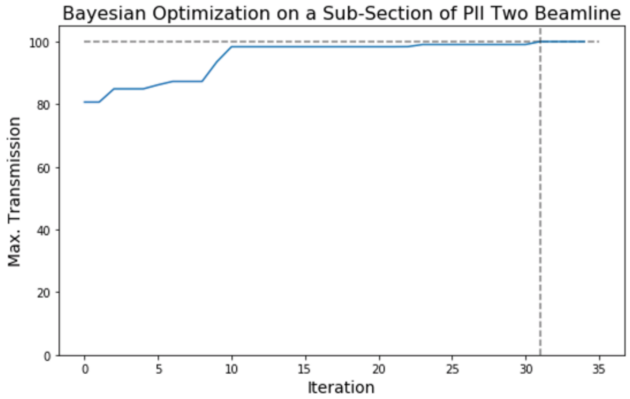

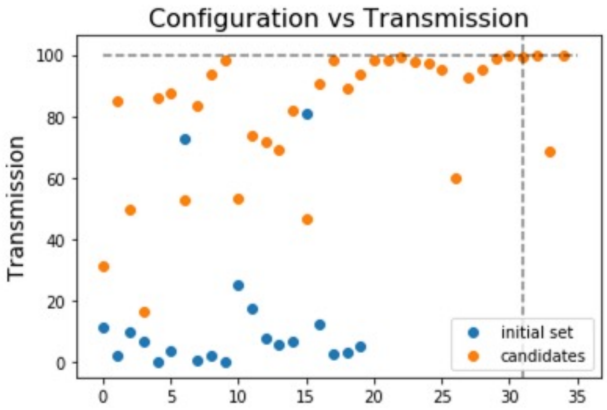

Random Initial Data

BO WITH SIMULATION DATA CASES

No RFQ

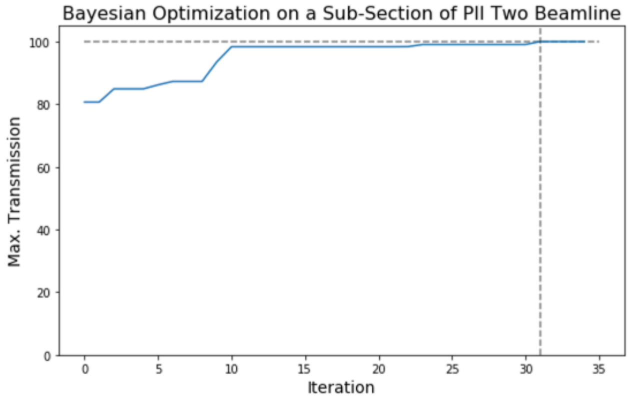

Prior knowledge/conditions helps a lot

RFQ

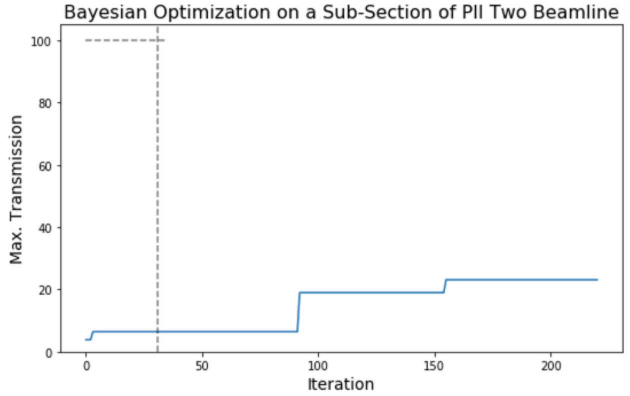

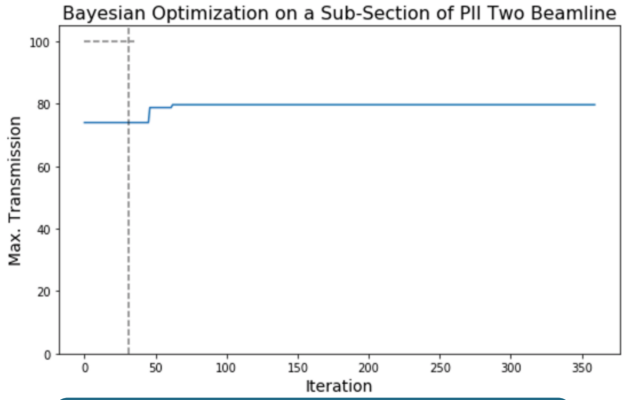

Random Initial Data

Historical Initial Data

BO WITH SIMULATION DATA – RFQ AND HISTORICAL INITIAL DATA CASE

- Bayesian Optimization with Gaussian Processes
- Transmission = f(9D-configuration)
- Quadrupoles limited based on historical data
- Surrogate Model: Gaussian Process with Matern Kernel and Gaussian likelihood.
- Acquisition function: Expected Improvement
- GPyTorch + BoTorch
- TRACK simulating the real machine.

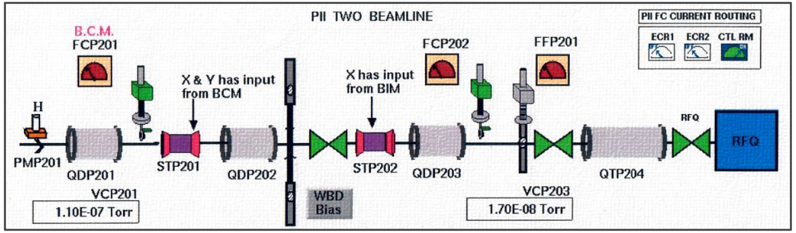

ATLAS sub-section from MHB to RFQ (incuded). 29 configurations with constrained settings.

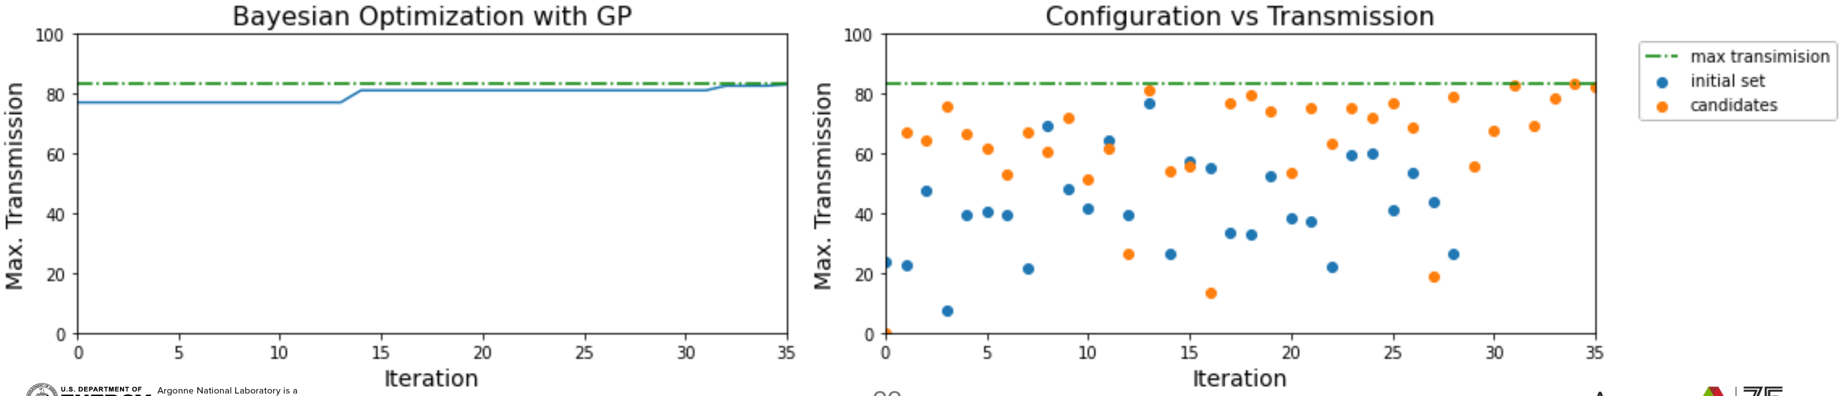

BO WITH GP - ATLAS

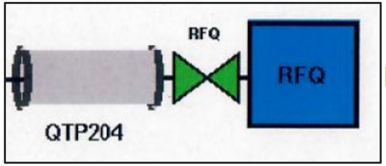

2xSteerers

- ✓ 7 input parameters (3 quadrupoles + 2 steerers)
- ✓ Optimization of the transmission
- ✓ Case of $^{14}N^{3+}$:
  - ✓ 29 historical tuned beams + 33 random configurations.
- ✓ Case $^{40}Ar^{9+}$:
  - ✓ 29 historical tuned beams

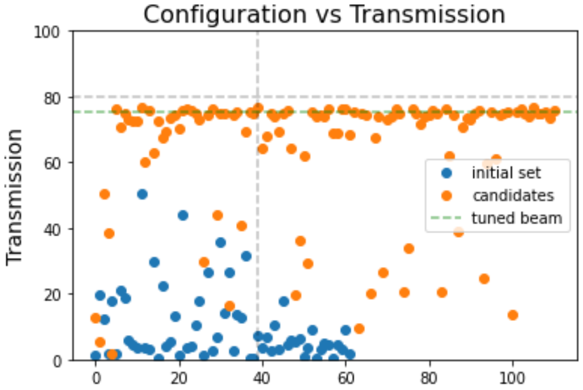

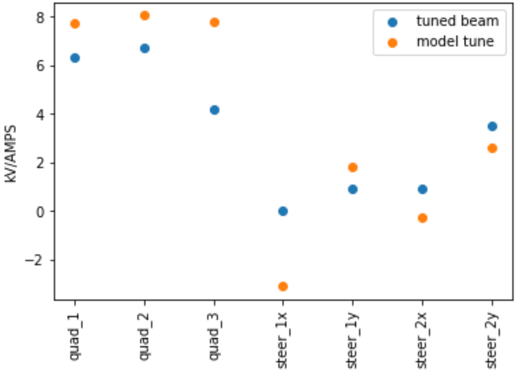

$^{14}N^{3+}$

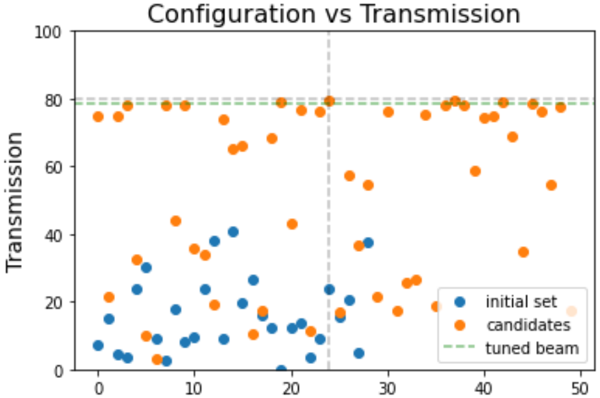

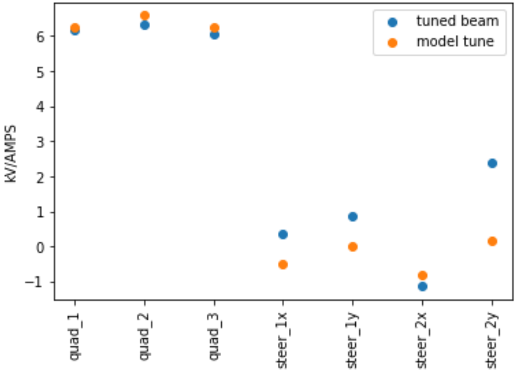

$^{40}Ar^{9+}$

REINFORCEMENT LEARNING

- ✓Reinforcement learning is learning what to do - how to map situations to actions in order to maximize a numerical reward.
- ✓Classic Control

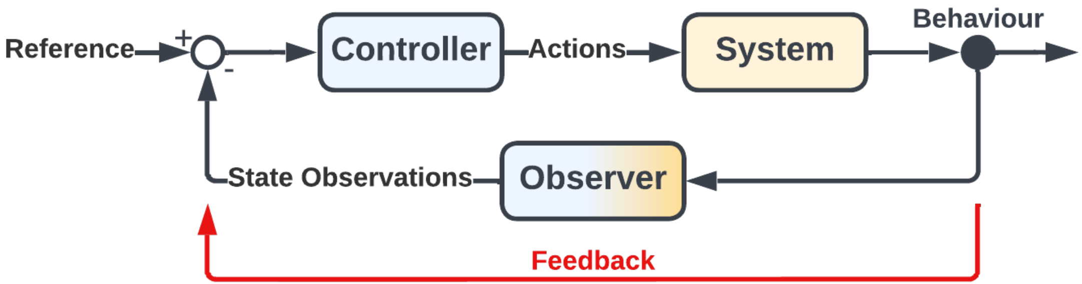

- ✓Particle Accelerators are among the most complex machines
- ✓Goal:

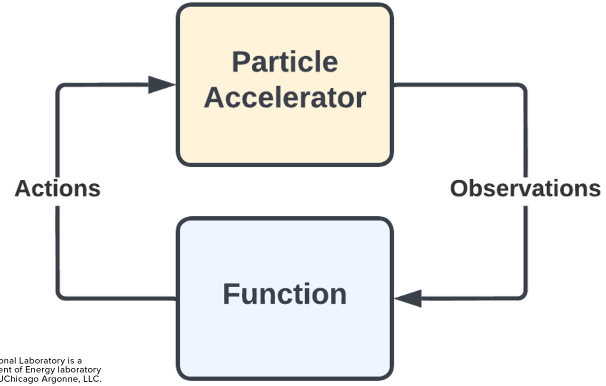

- ✓What does this function look like?
- ✓How do you design it?

REINFORCEMENT LEARNING

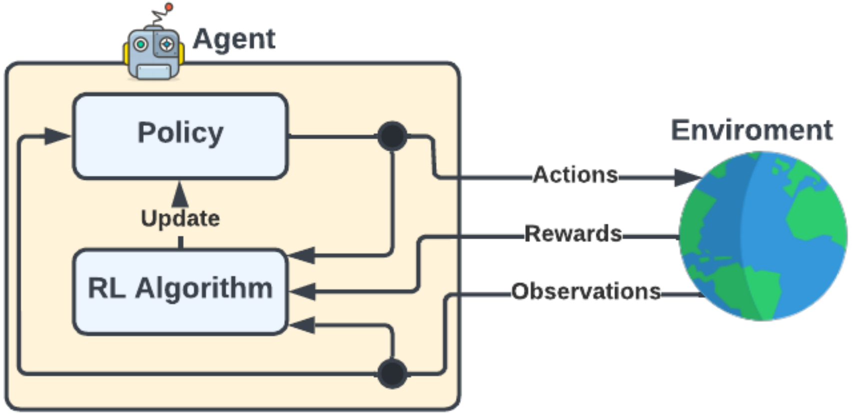

REINFORCEMENT LEARNING

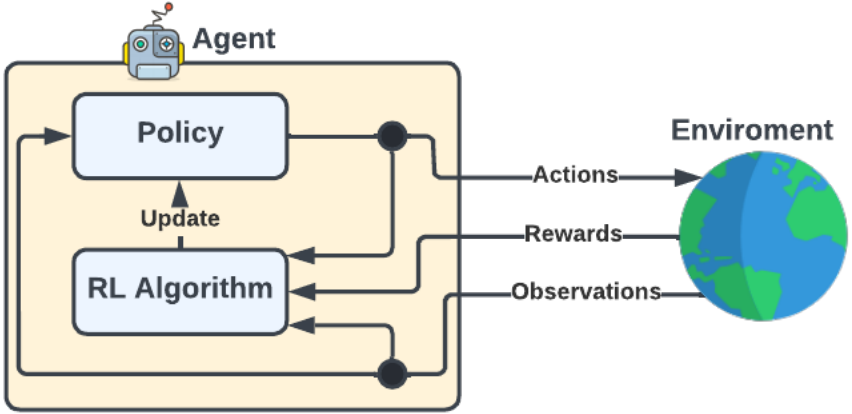

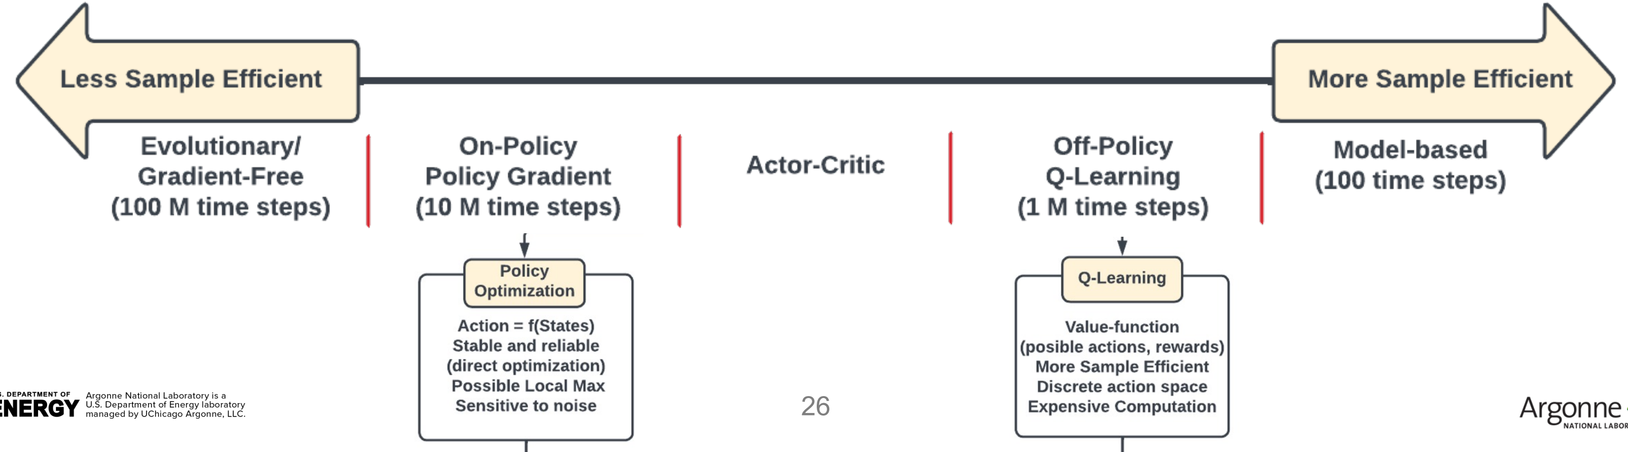

REINFORCEMENT LEARNING

Actor-Critic

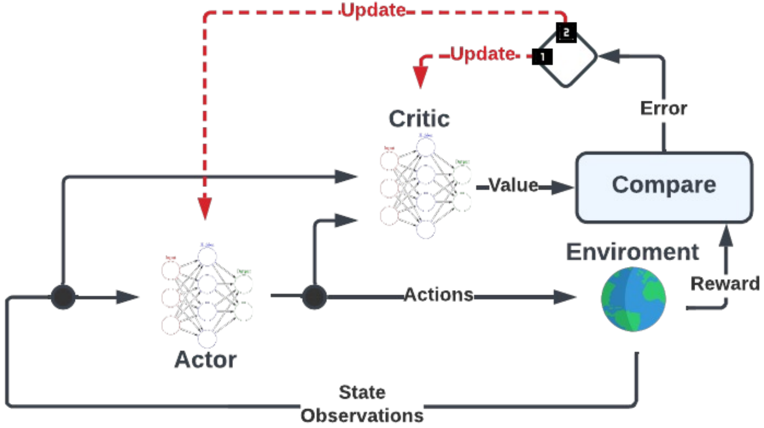

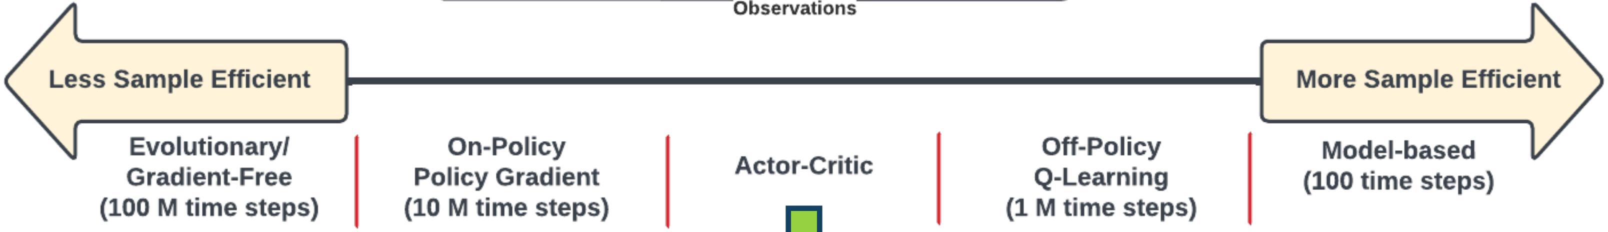

Deep Deterministic Policy Gradient (DDPG)

DDPG – 3 QUADRUPOLE AND BEAM SIZE – SIMULATION DATA

Training

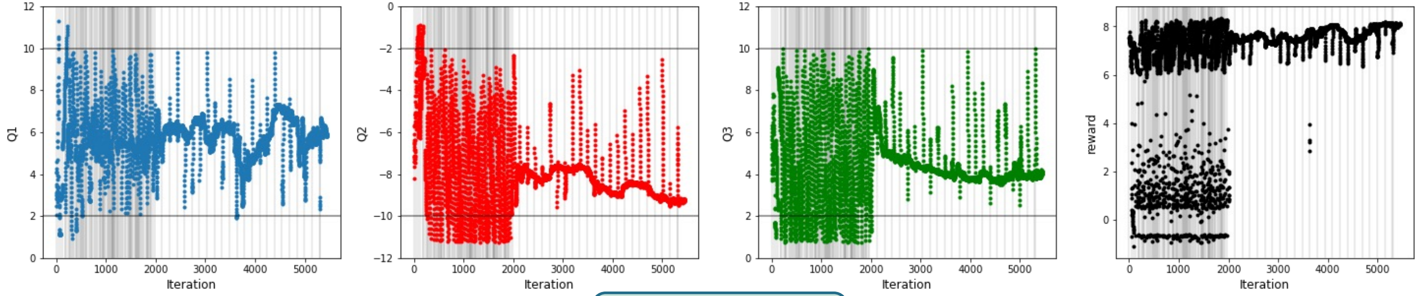

Prediction

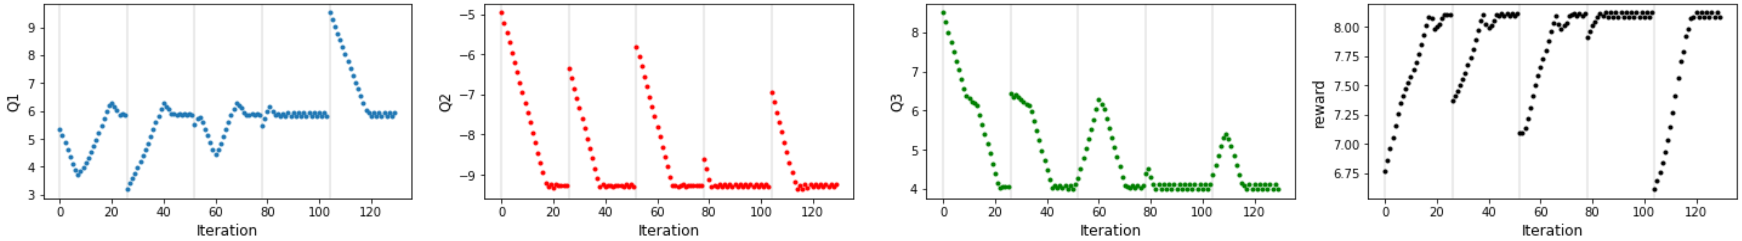

✓ Electrostatic Quadrupoles:
- 2 kV to 10 kV
- Max action +- 0.25 kV

DDPG – 4 QUADRUPOLE + 2 MAGNETS AND TRANSMISSION – REAL DATA

Training

*In progress

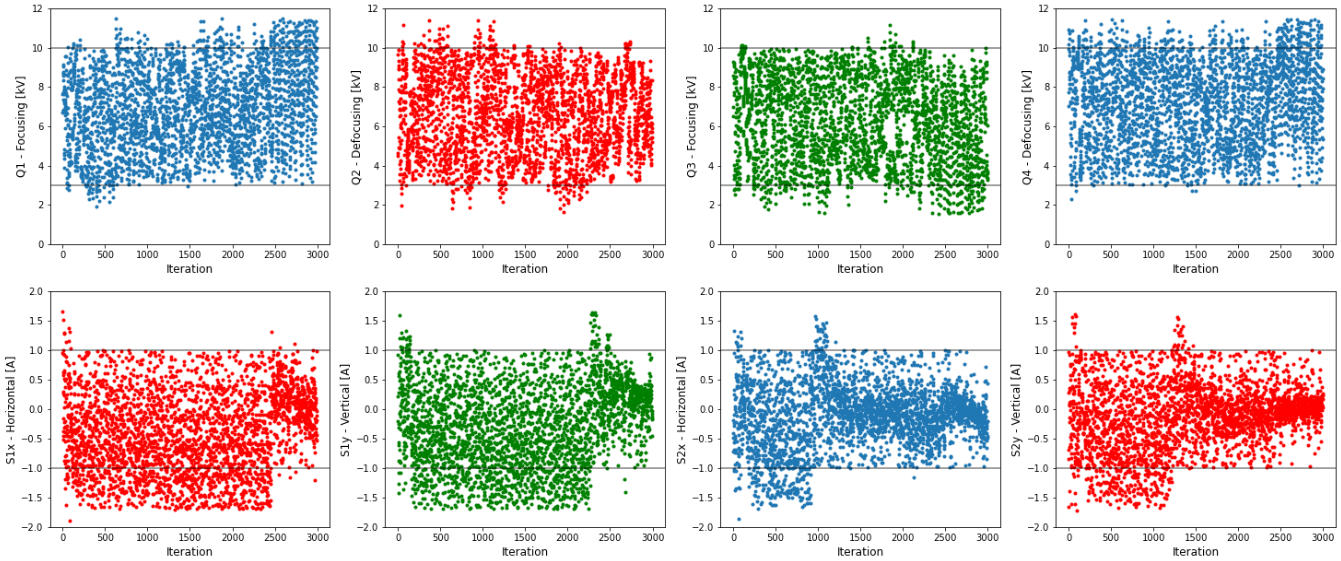

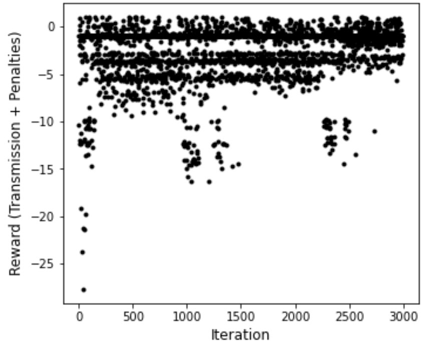

✓ Electrostatic Quadrupoles :

- 3 kV to 10 kV
- Max action +- 0.25 kV

✓ Steering Magnets:

- -1 A to 1 A
- Max action +- 0.25 A

DDPG – 3 QUADRUPOLE CASE + RFQ AND TRANSMISSION

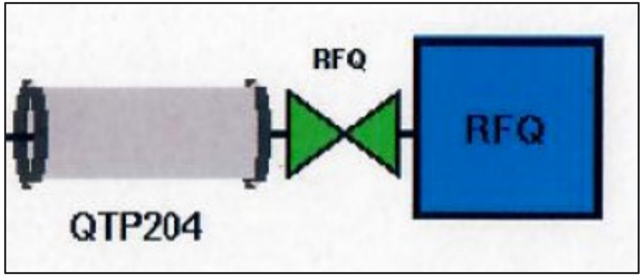

✓Long simulation times because of RFQ.
✓10^4 particles -> +45 seconds per simulation.
✓RL requires a lot of iterations.

Offline training and online fine tune

SURROGATE MODEL

✓ Physics Simulation Codes → nonlinear/collective effects/3D fields + **slow**

- Impedes:
  - Start-to-end optimization
  - Use as an online model / virtual diagnostic
  - Use in control
- Cannot always replicate the real machine behavior

SURROGATE MODEL

✓ <u>Physics Simulation Codes</u> → nonlinear/collective effects/3D fields + **slow**

- Impedes:
  - Start-to-end optimization
  - Use as an online model / virtual diagnostic
  - Use in control
- Cannot always replicate the real machine behavior

✓ Faster Codes:

- Simpler (↓ Accuracy)
- Parallelization
- Faster Algorithms

✓ <u>ML surrogate model</u>

- Once trained, fast execution
- Be able to optimize multiple objectives
- Fulfill multiple constraints
- **Be fast and accurate enough**
- Handle noise
- Learn from past experiments and simulations

SURROGATE MODEL

✓ML Surrogate Model can be used for **virtual diagnostics, offline experiment planning, design of new setups, control and tuning**.

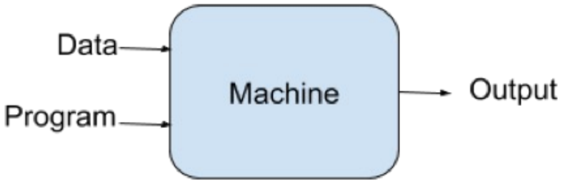

Traditional Programming

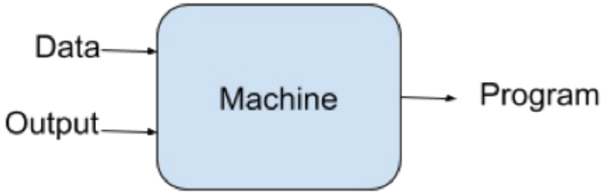

Machine Learning

SURROGATE MODEL

✓ML Surrogate Model can be used for **virtual diagnostics, offline experiment planning, design of new setups, control and tuning**.

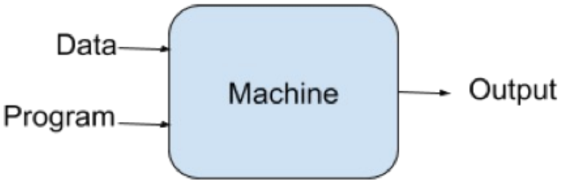

Traditional Programming

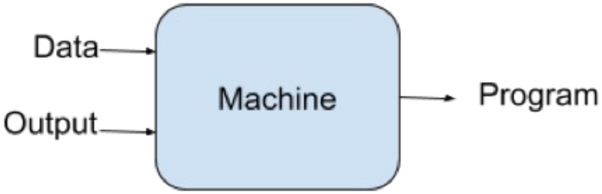

Machine Learning

- A. Edelen *et al 2019*: Multi-objective optimization using surrogate model based on neural networks → beam parameters = f(settings)

  ~ O(10^6) - O(10^7) more efficient to execute
- Lipi Gupta *et al* 2021: surrogate model to predict scalar beam parameters and the transverse beam distribution downstream for the LCLS-II injector taking into account the impact of time-varying non-uniformities in the initial transverse distribution

  ~ MAPE <= 9%

EXPLORING ML SURROGATE MODELS

- ✓ Preliminary studies
- ✓ Radio-frequency quadrupole (RFQ)
- ✓ Data from TRACK simulations
- ✓ Neural network architectures
- ✓ TensorFlow

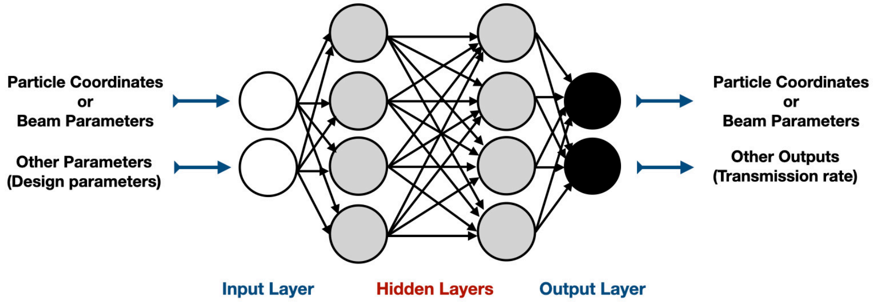

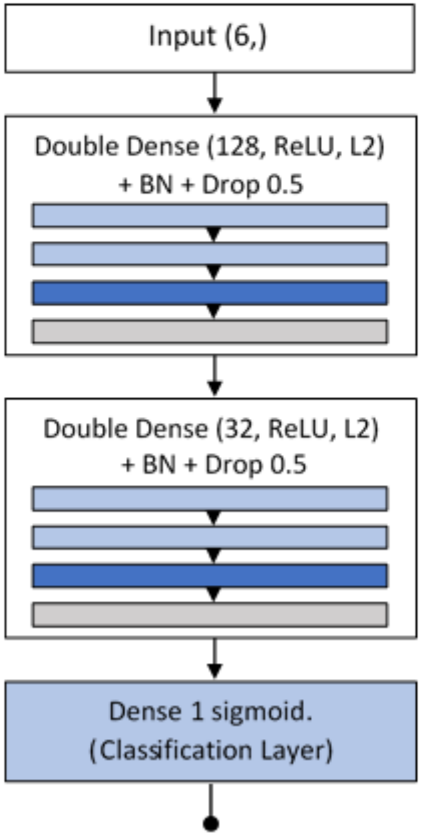

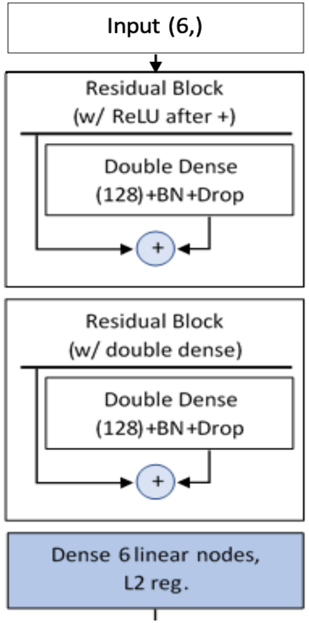

SURROGATE MODELS FOR BEAM TRANSPORT

Goal: Predict acceptance + Twiss parameters based on input distribution

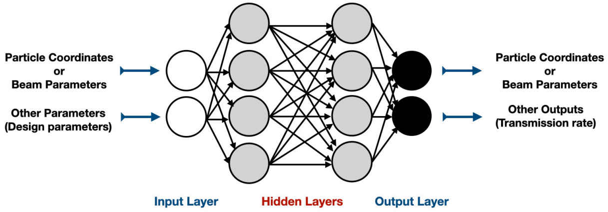

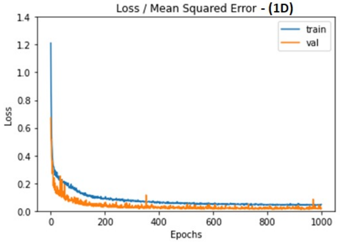

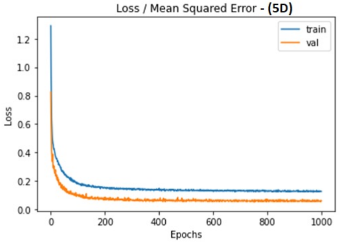

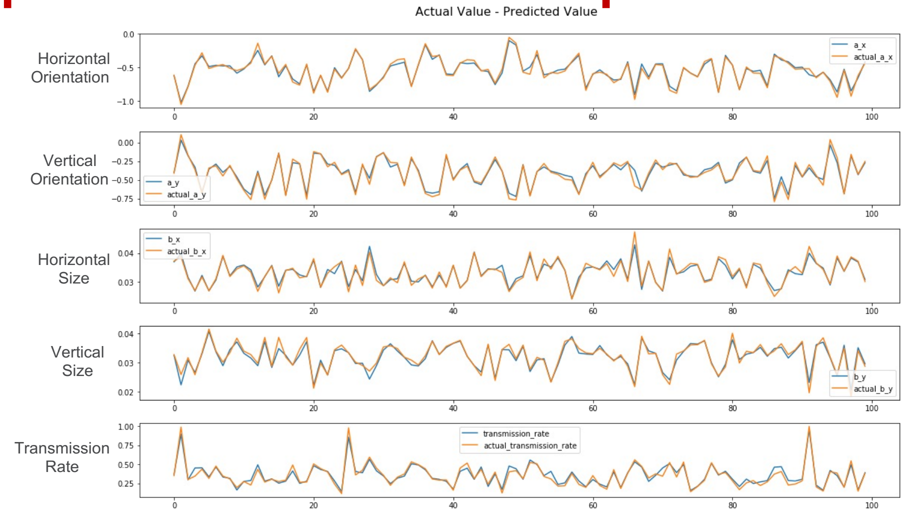

- ✓ Data generated using TRACK 3D simulations, 7000 x 10^4 particles each, with different transverse emittances, phase width and energy spread.
- ✓ Excellent agreement with TRACK 3D beam simulations
- ✓ Much faster than TRACK, **speed-up factor ~ 30,000**.

SURROGATE MODELS FOR BEAM TRANSPORT

Goal: Predict acceptance + Twiss parameters based on input distribution

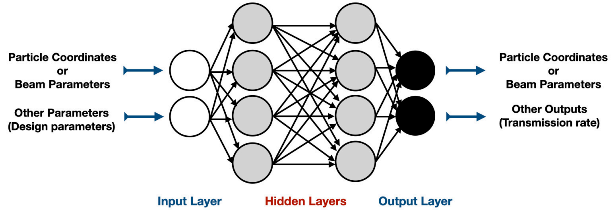

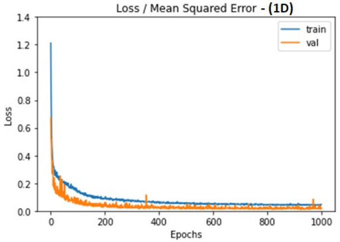

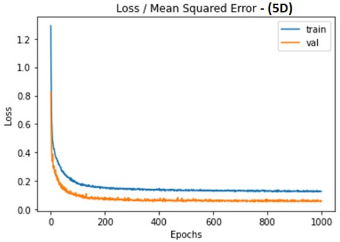

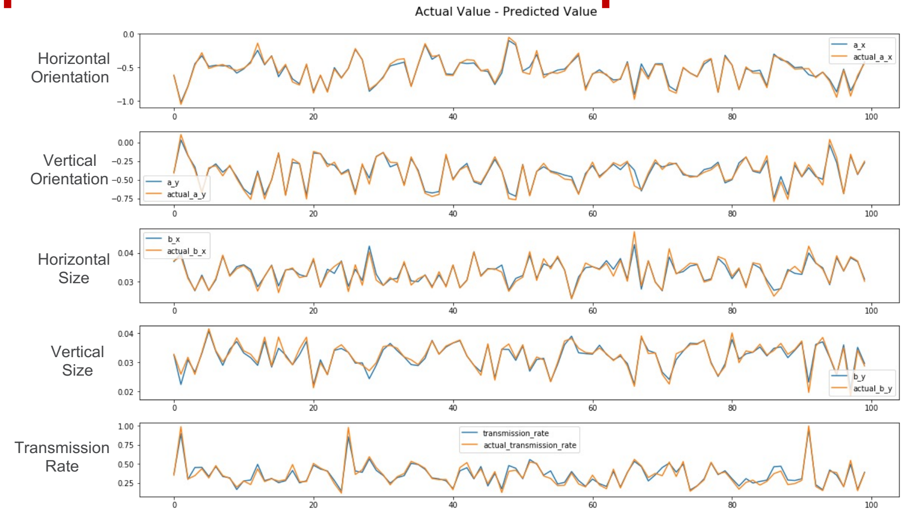

- ✓ Data generated using TRACK 3D simulations, 7000 x 10^4 particles each, with different transverse emittances, phase width and energy spread.
- ✓ Excellent agreement with TRACK 3D beam simulations
- ✓ Much faster than TRACK, **speed-up factor ~ 30,000**.

Next Step: Use RFQ surrogate model in RL for offline training

BO VS RL

Analogous concepts, different terminology and usually different settings.

Objective → Reward
Surrogate Model → Value Function
Acquisition Function → Policy
Acquire new sample → Take an Action

- ✓ BO and RL both are useful for high-level tuning and control but excel in different regimes.
  - ✓ BO: exploratory/optimization new setups + low data regime/slow measurements
  - ✓ RL: high data regime, continuous control

- ✓ BO would be more suitable for new tuning configurations and RL for continuous control after pre-trained offline.

CONCLUSIONS AND NEXT STEPS

✓**Automated data collection and integration of new devices as the pepper pot**.

✓**Successfully trained and deploy a BO with GP on real machine for a subsection of ATLAS.**

✓**Integration of RL model with the real machine (preliminary results).**

✓Next Steps

- Test Pepper Pot, get more useful data and test RL on machine.
- Improve existing models (other architectures, new type of data (adding beam profilers or pepper pot images, incorporate more Physics information, use of surrogate models, etc.).
- TRACK lattice including misalignments.

✓Current Challenges:

- Possible damage to devices when beam is lost during model training

ACKNOWLEDGMENTS

Brahim Mustapha, Ben Ryan Blomberg, Eric Letcher, Daniel Stanton, Clayton Dickerson, Kenneth Bunnell, Daniel Santiago, Megan McIntyre, Alexander F Grabenhofer, Gavin Matthew Dunn, Henry Brito, Samantha Burtwistle, Tony Krupa, Leland Luecke, etc.


Argonne
NATIONAL LABORATORY
1946
75
2021