Machine Learning applications for digital twin development and polarization improvement at the BNL hadron injectors

Lucy Lin

Advisor: Georg Hoffstaetter de Torquat (Cornell), Kevin Brown (BNL)

AI4EIC 2023 Annual Workshop
Washington D.C., USA   Nov 29, 2023

Injector compound for RHIC and EIC

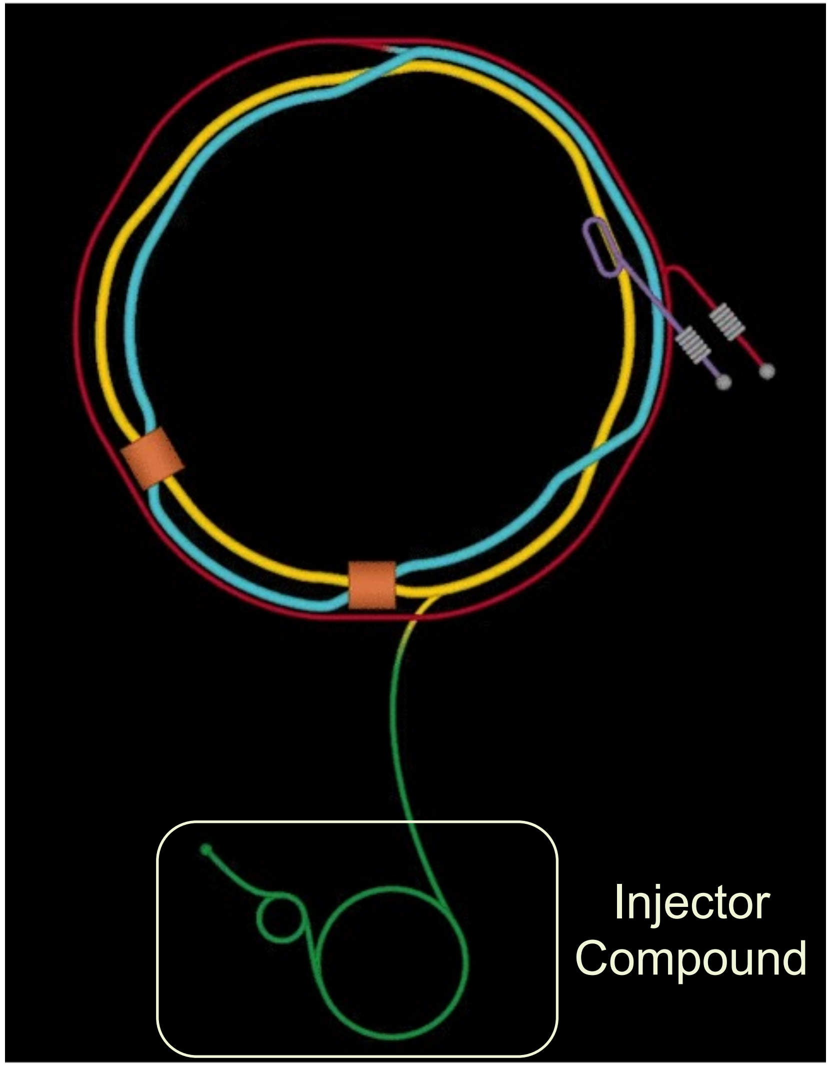

- **Relativistic Heavy Ion Collider (RHIC)**: largest operating accelerator in the US.
- **Electron Ion Collider (EIC)**: the nation's largest particle accelerator project.
- **Alternating Gradient Synchrotron** (AGS) and its **Booster** serve as part of the **injector compound** for RHIC and future EIC.
- **Bright ion beams** in the AGS and Booster are required for optimal luminosity and highest polarization in RHIC and EIC.
- Obtaining bright beam requires **more accurate beam control** in the injector compound, which is currently mostly hand tuned by operators.

RHIC Accelerator Complex

Accelerator Rings	Circumference [m]
Booster	201
AGS	807
RHIC	3833

Typical Top Energies [Total, GeV/N]

	Au	Pol. Protons
Linac (H^-)	--	1.1
Booster	1	2.3
AGS	10	23.8
RHIC	100	255

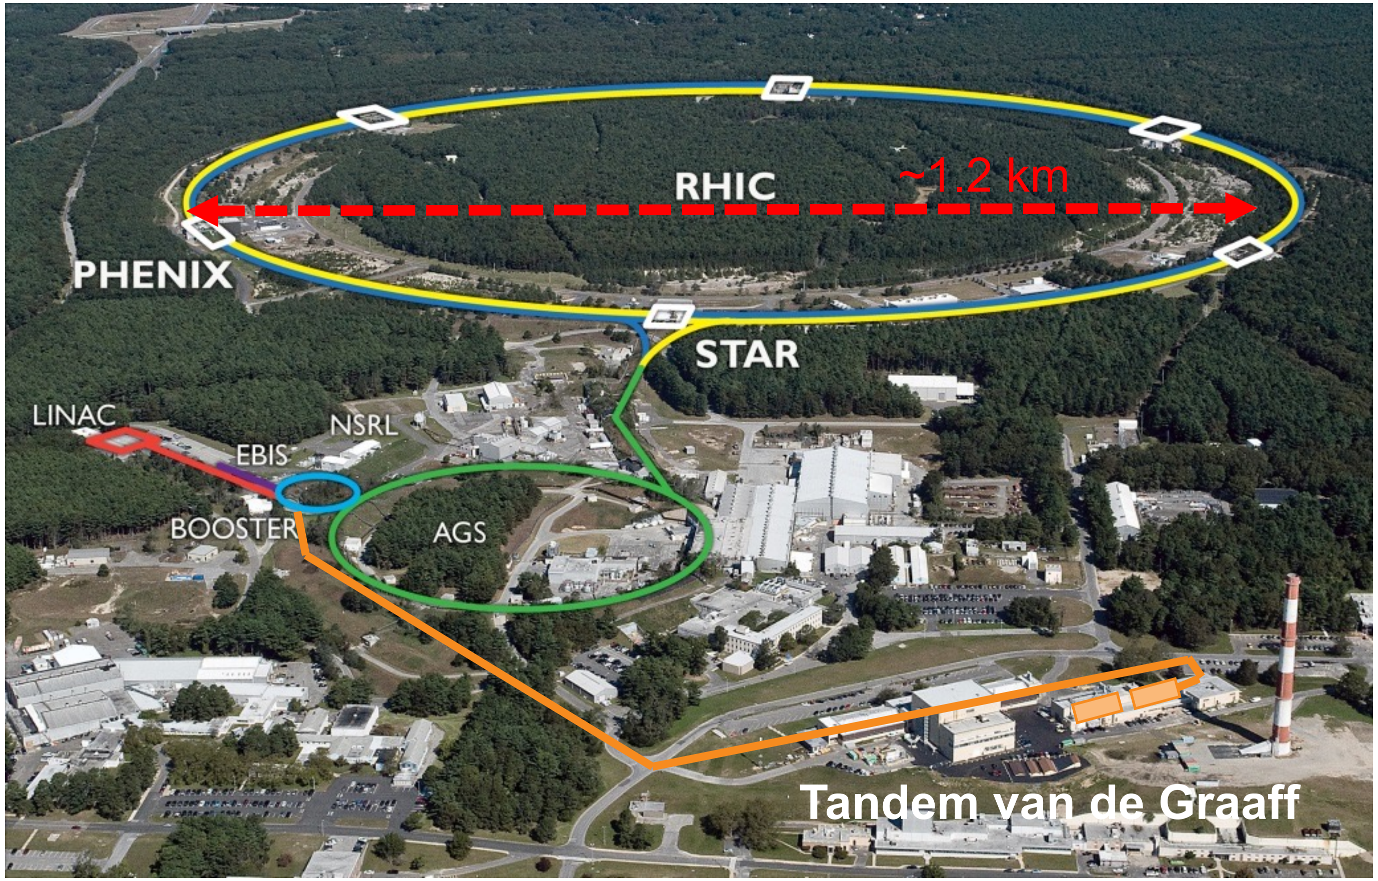

Heavy Ions	Protons
E-beam Ion Source (EBIS)	OPPIS (polarized)
Tandem Van de Graaf	High-intensity H^- (unpolarized)

Alternating Gradient Synchrotron (AGS)

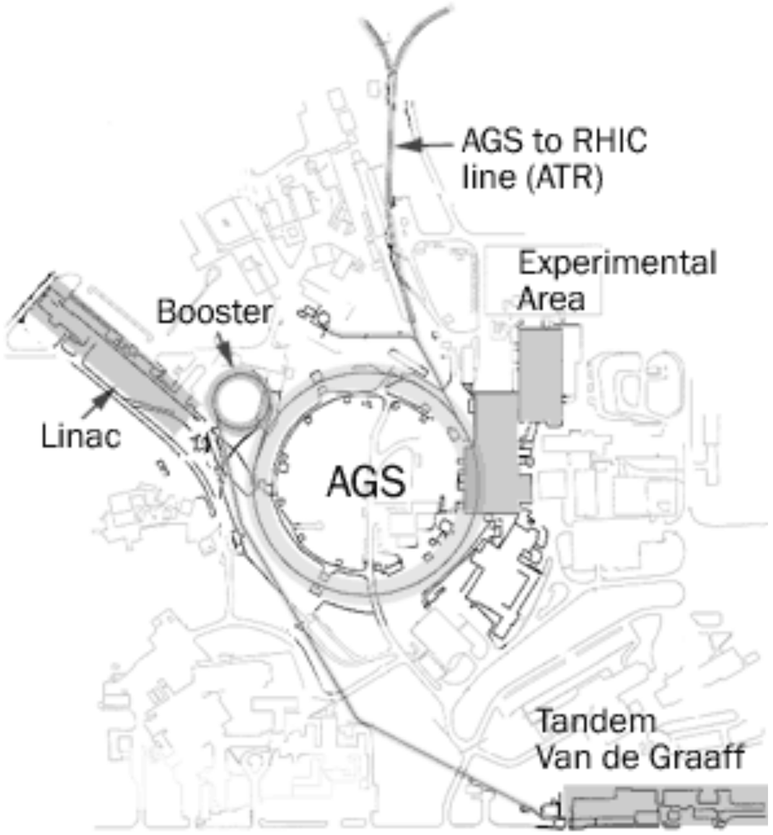

- Alternating gradient / strong focusing principle: achieve strong vertical and horizontal focusing of charged particle beam at the same time
- Accelerates proton to 33 GeV in 1960
- 12 super-periods (A to L), 240 main magnets, 810 m circumference
- Now serves as injector for Relativistic Heavy Ion Collider (RHIC)

Orbit Correction at the AGS

$$\begin{pmatrix} \Delta\vec{x} \\ \Delta\vec{y} \end{pmatrix} = \underline{R} \begin{pmatrix} \Delta\vec{\theta}_x \\ \Delta\vec{\theta}_y \end{pmatrix}$$

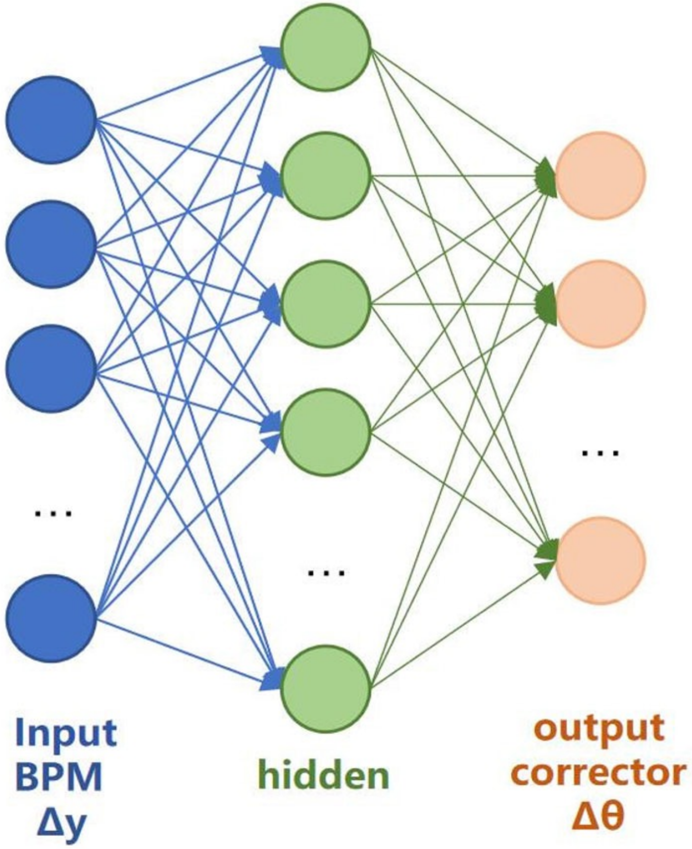

- Traditional orbit correction
  - obtain mapping $\underline{R}$ (orbit response matrix) from corrector settings $\vec{\theta}$ to orbit measurements $\vec{y}$
  - inverse mapping to get corrector settings $\Delta\vec{\theta}$ needed to cancel orbit deviations $\Delta\vec{y}$
- Orbit correction with NN
  - train directly to get inverse mapping, no need for extra calculation
  - easily update with new data and stay accurate

ML method: Neural Network (NN)

- Establish mapping between a given set of inputs $\vec{X}$ and corresponding outputs $\vec{Y}$
- Fully connected layers: output = activation(dot(input, weight) + bias)
- Activation function: Hyperbolic Tangent (Tanh) and Rectified Linear Unit (ReLU)
- Feed forward neural network (FFNN): most common, no feedback route

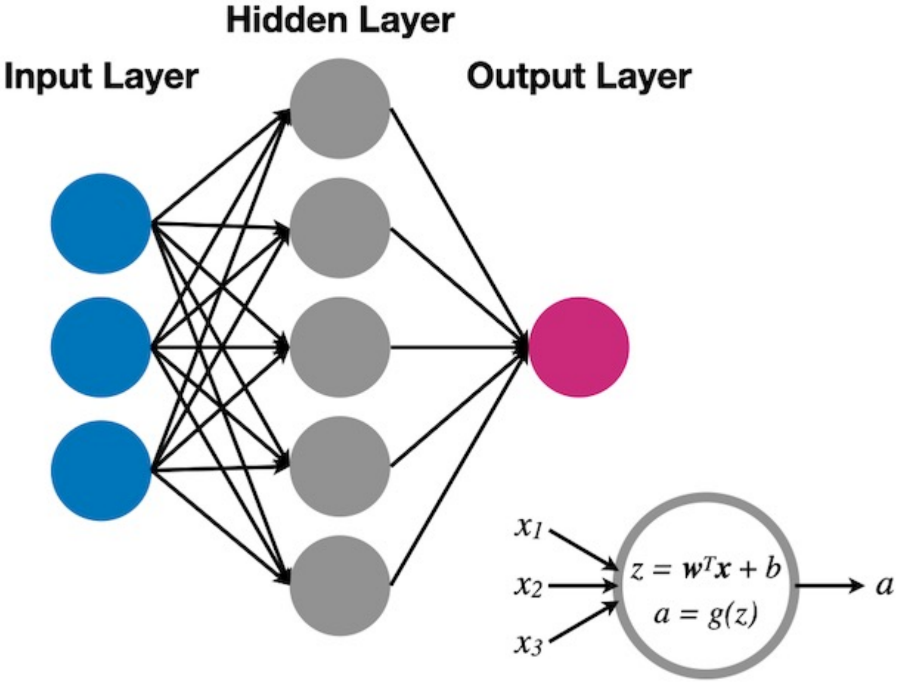

Hyperbolic tangent	$\phi(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$	Multi-layer Neural Networks	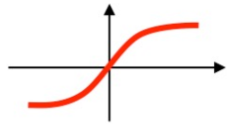
Rectifier, ReLU (Rectified Linear Unit)	$\phi(z) = max(0, z)$	Multi-layer Neural Networks	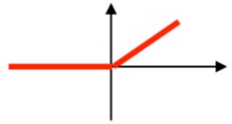

AGS ORM NN model: training results

- Input 48 vertical corrector kick → Output 72 y orbit measured at BPM
- Trained on 800 data pairs, tested on 200 data pairs: R^2 score = 0.998

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y}_i)^2}$$

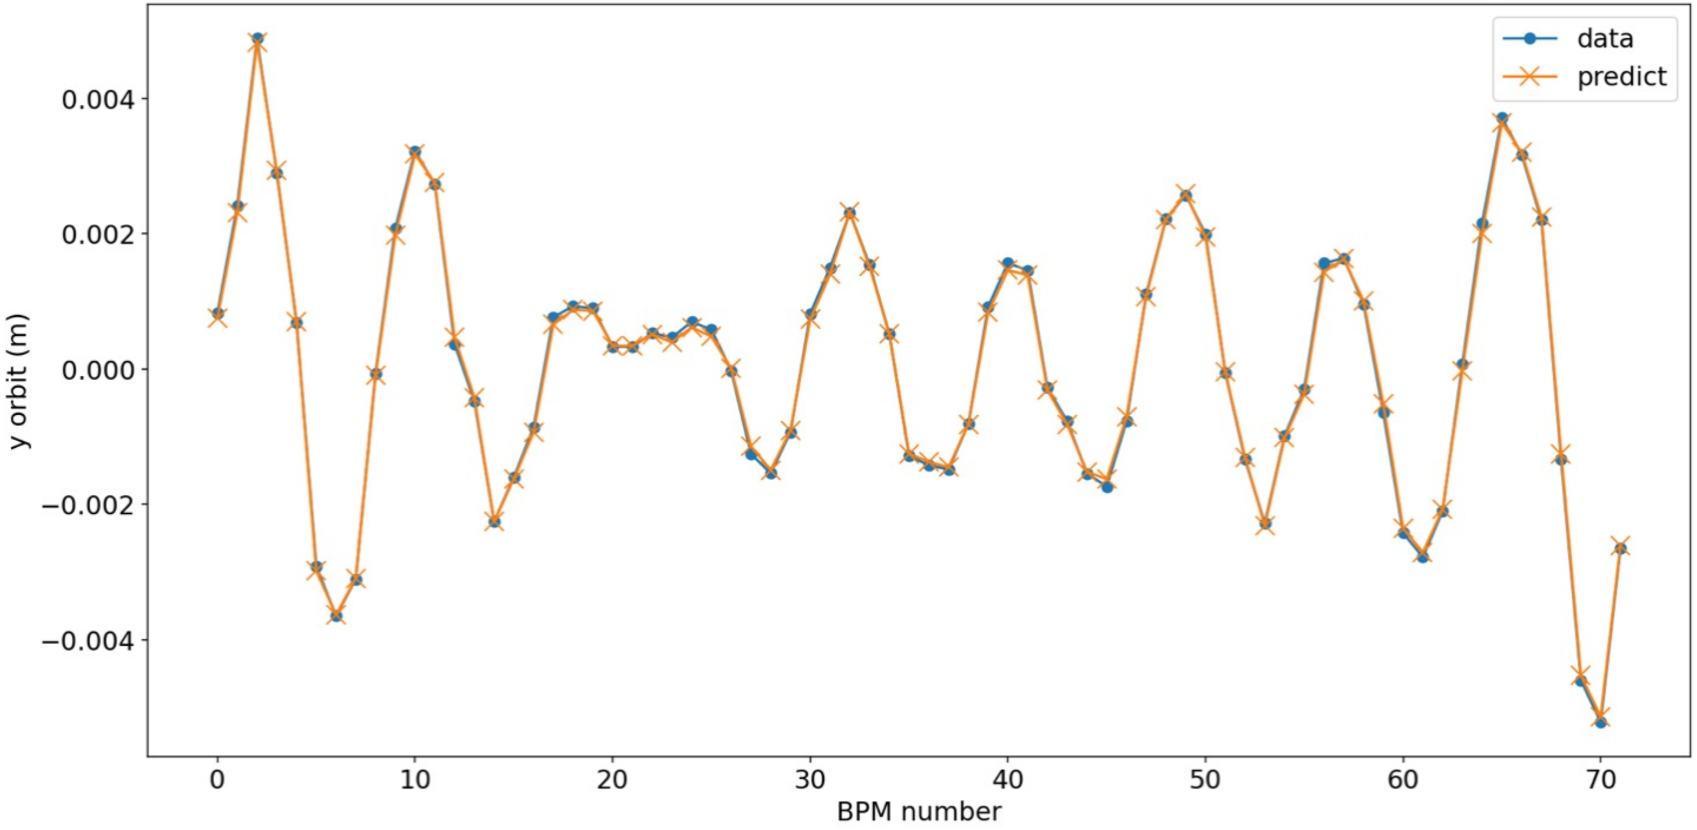

Inverse AGS ORM NN model: training results

- Input 72 y orbit measured at BPM → Output 48 vertical corrector kick
- Trained on 800 data pairs, tested on 200 data pairs: R^2 score = 0.993

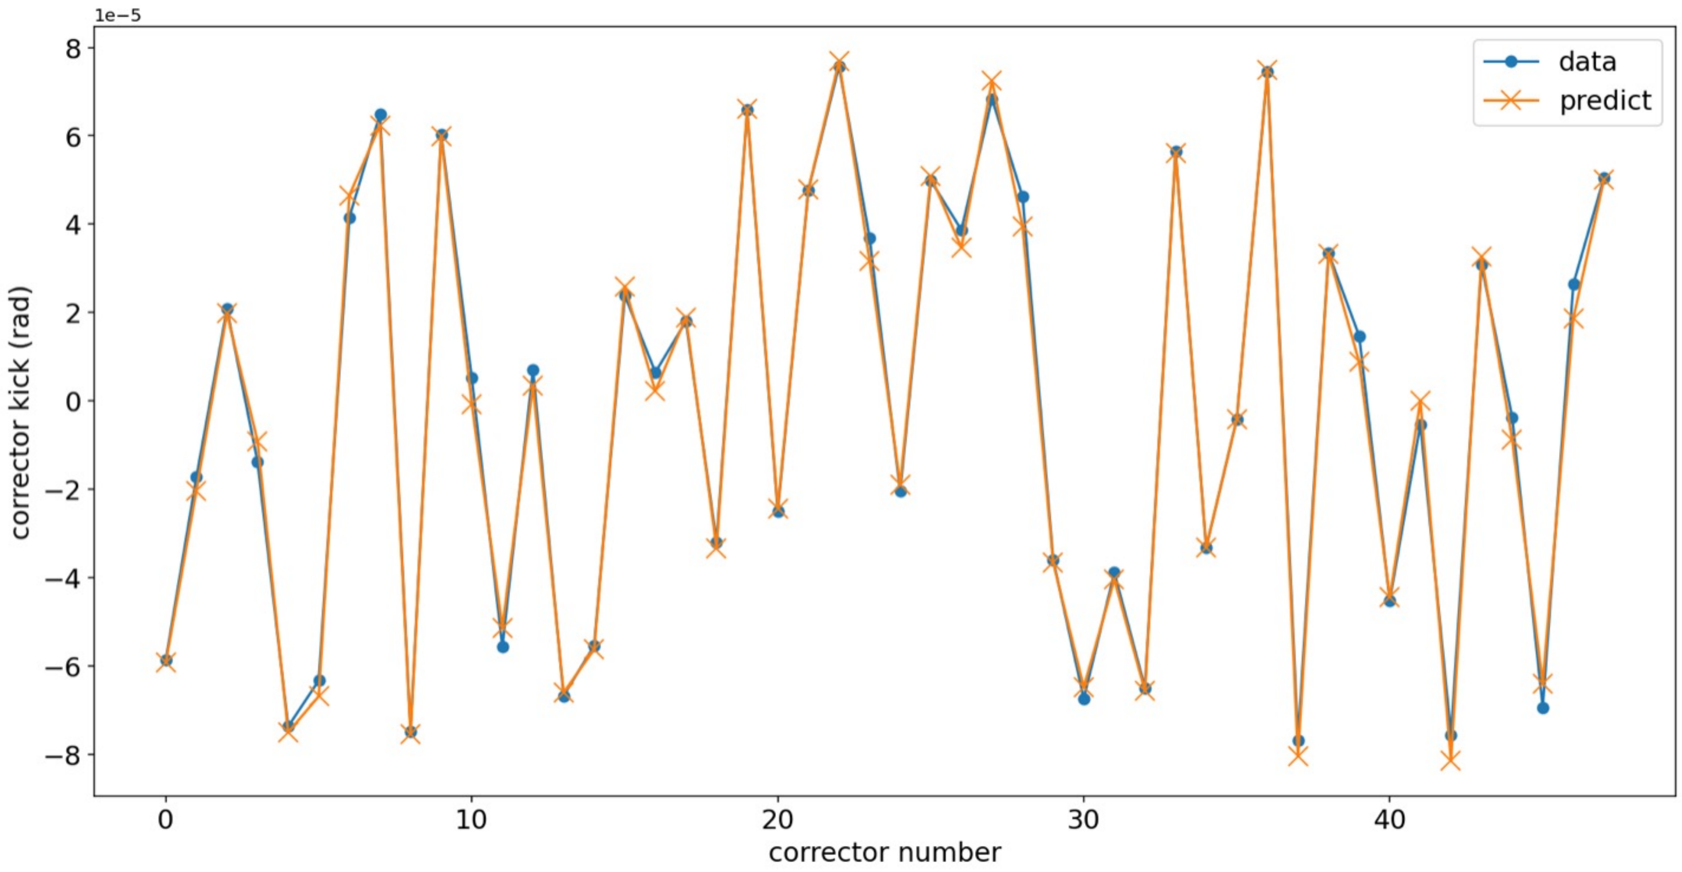

Alternating Gradient Synchrotron (AGS) Booster

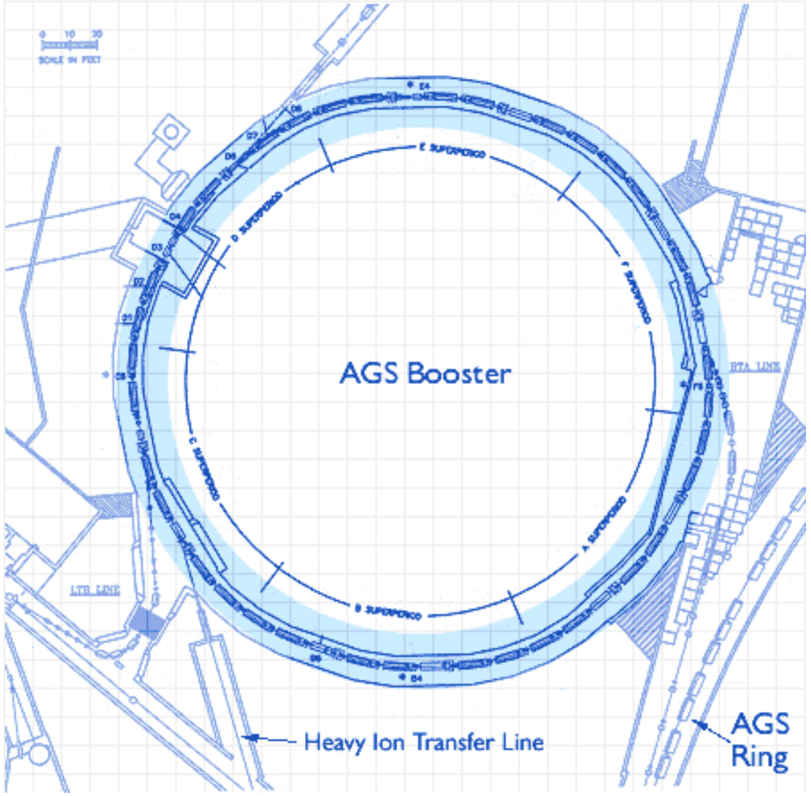

- Pre-accelerate particles entering the AGS ring
- Accepts heavy ions from EBIS or protons from 200 MeV Linac
- Serves as heavy ion source for NASA Space Radiation Laboratory (NSRL)
- 6 super-periods (A to F), 72 main magnets

Brookhaven National Laboratory
CLASSE Cornell Laboratory for Accelerator-based Sciences & Education

Booster magnet misalignment

- Magnet location in real machine from 2015 survey data
- Misalignment data for quadrupoles and dipoles
- Trouble with making physics simulation with misalignment agree with real orbit data

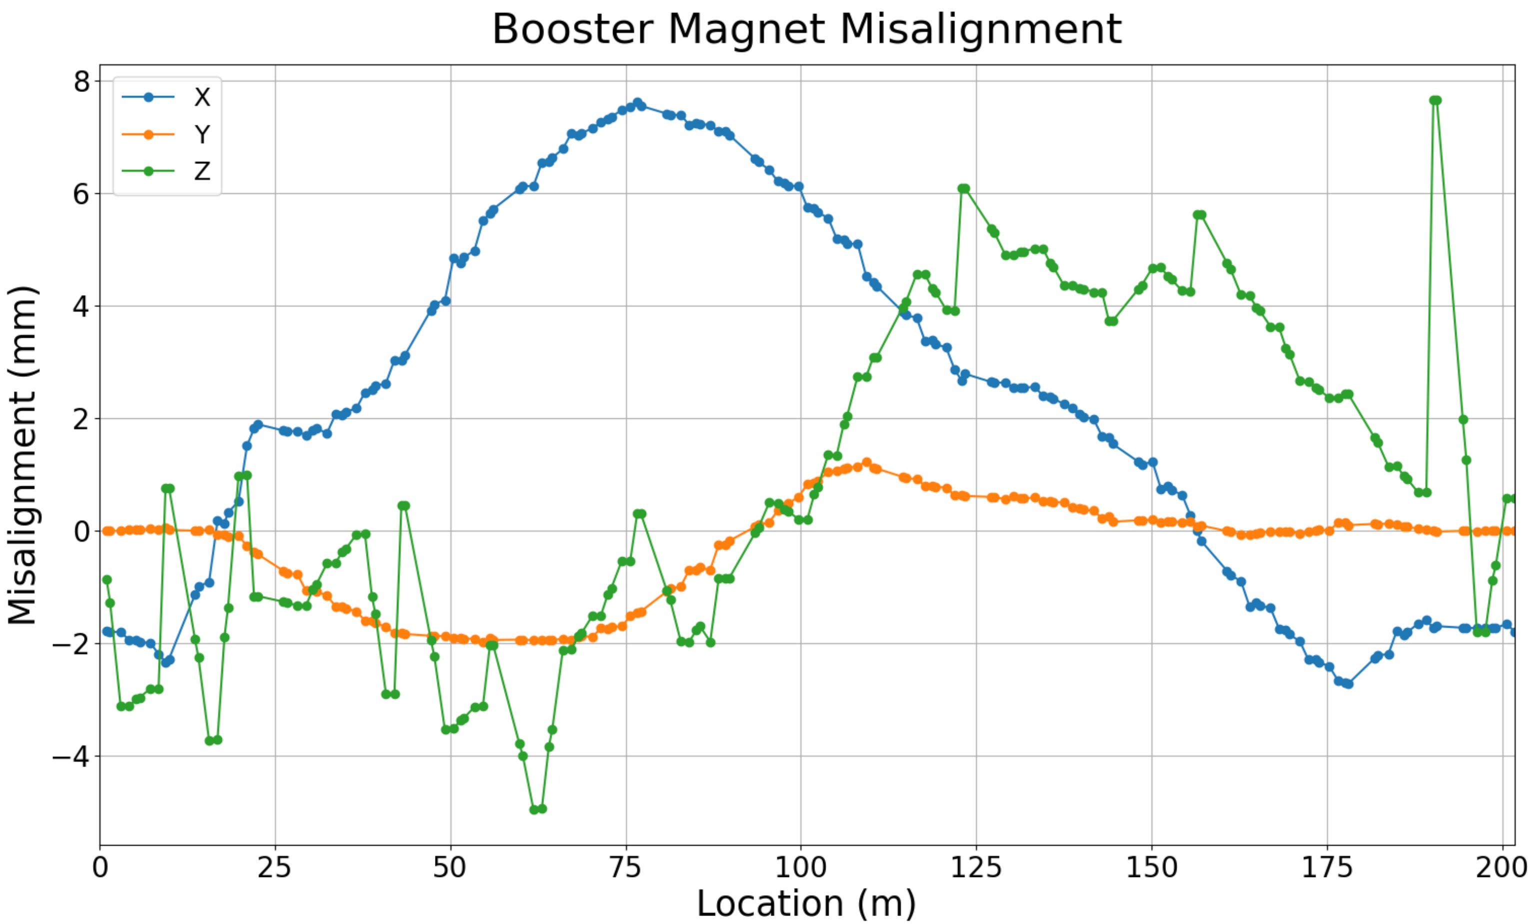

Script to get Booster orbit responses

- Script development with Collider Accelerator Department (CAD) Controls Group
- FunctionEditor: send trapezoid-like time-dependent function to corrector power supplies
- Script sets three corrector settings: positive, zero, negative; and save corresponding orbits

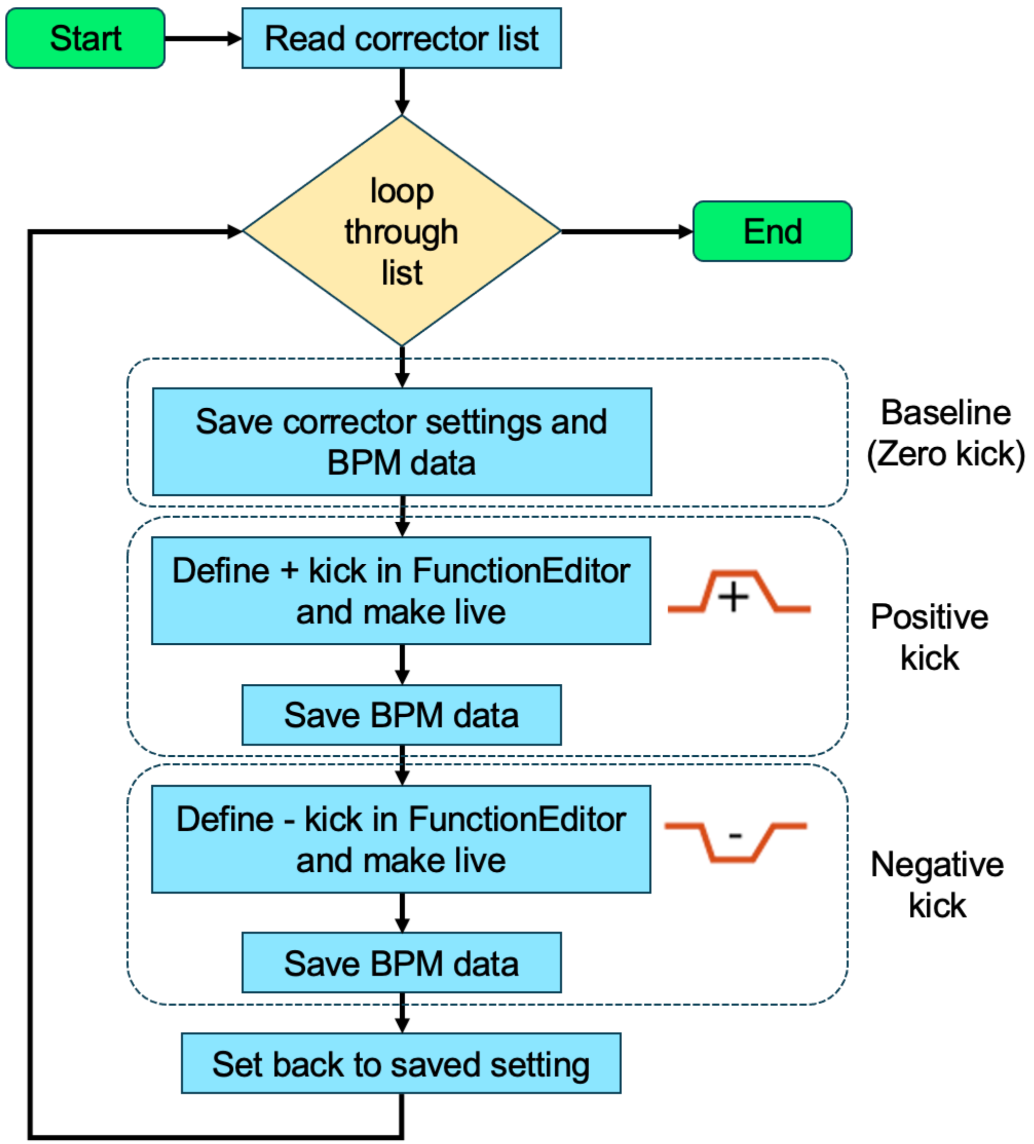

Sample data

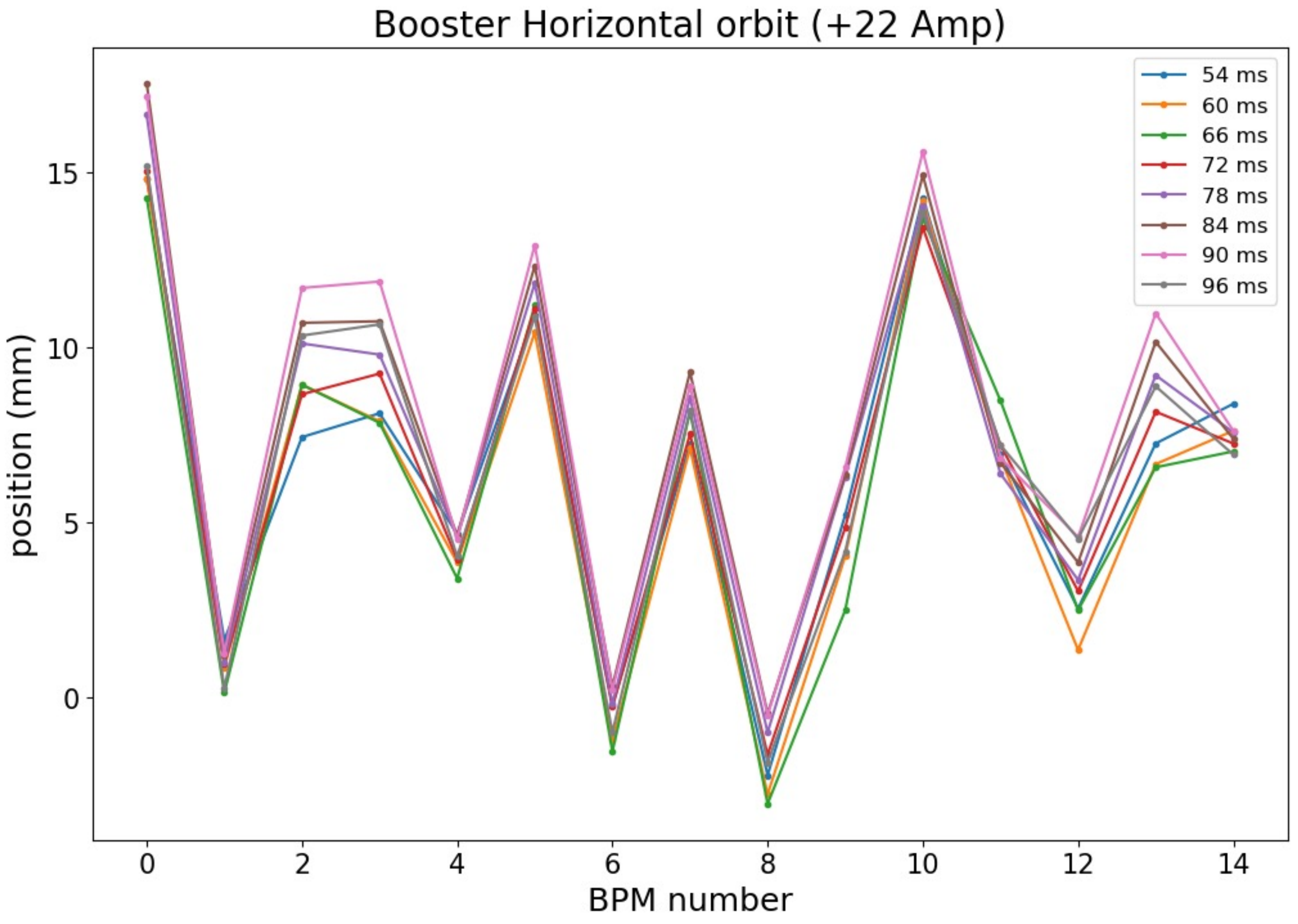

Booster Horizontal orbit (+22 Amp)
position (mm)
BPM number
0
5
10
15
0
2
4
6
8
10
12
14
54 ms
60 ms
66 ms
72 ms
78 ms
84 ms
90 ms
96 ms

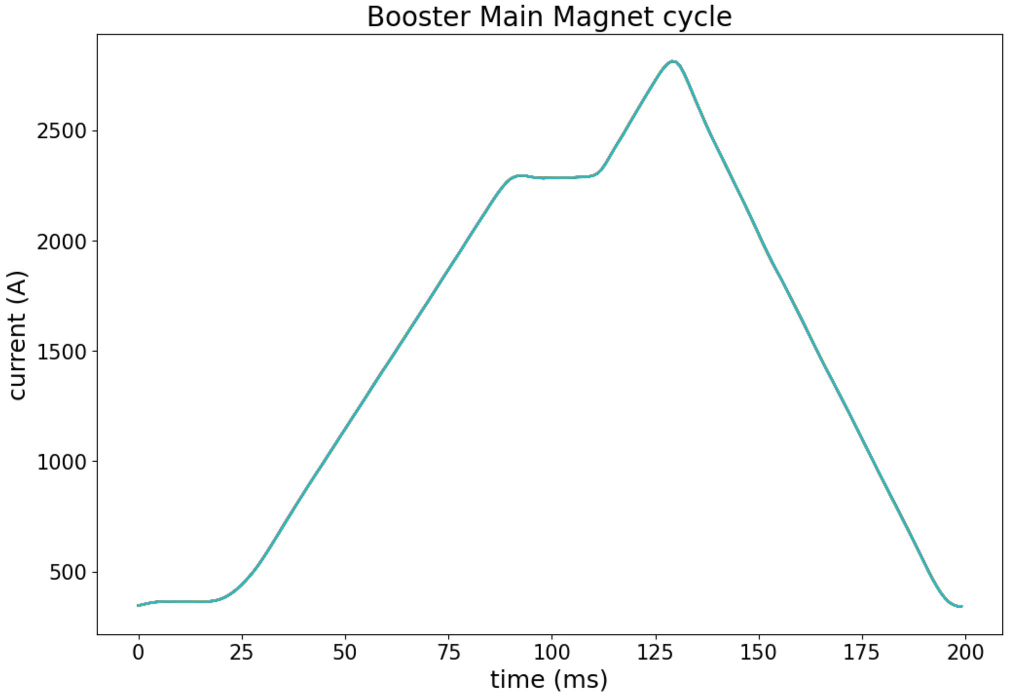

Booster Main Magnet cycle
current (A)
time (ms)
500
1000
1500
2000
2500
0
25
50
75
100
125
150
175
200

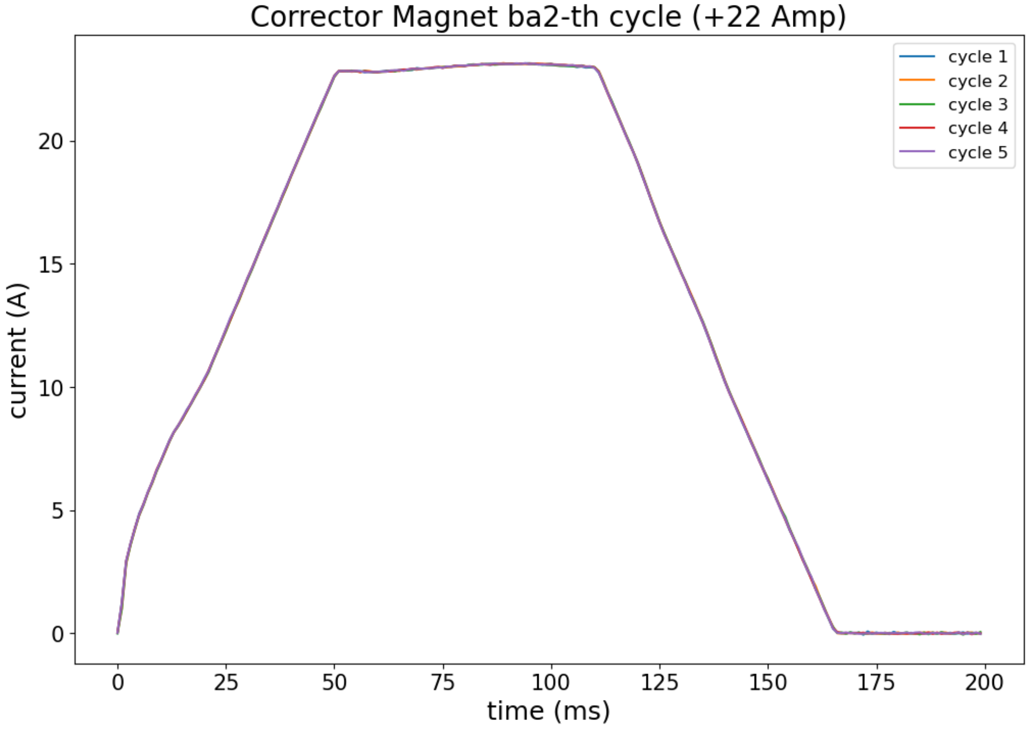

Corrector Magnet ba2-th cycle (+22 Amp)
current (A)
time (ms)
0
5
10
15
20
0
25
50
75
100
125
150
175
200
cycle 1
cycle 2
cycle 3
cycle 4
cycle 5

Booster model calibration

$$\left(I_{quad}, I_{corr}, \theta\right) \overset{model}{\Longleftarrow\!=\!=\!=\!=\!\Longrightarrow} \left(X_{BPM}^{(I_{quad}^{(1)}, I_{corr}^{(1)})}, X_{BPM}^{(I_{quad}^{(2)}, I_{corr}^{(2)})}, \dots\right)$$

- Control: power supply currents of quadrupoles and correctors
- Parameter θ: parameters that affect the orbit but not in our control → (magnet misalignments, magnet transfer functions, etc.)
- Output: orbit at the BPMs with certain current configuration
- Invert from measured BPM data to simulation model parameters
- Update beliefs on model parameters with real data → calibrated model m (“digital-twin”) can be used to optimize beam quality (objective F)

$$X_{BPM} = m\left(I_{quad}, I_{corr}; \theta\right) + \epsilon, \qquad \epsilon \sim N(0, \sigma)$$

$$I_{quad}^{*}, I_{corr}^{*} = \operatorname{argmax} F\left(m\left(I_{quad}, I_{corr}; \theta\right)\right)$$

Polarization at RHIC

	Max Energy [GeV]	Pol. At Max Energy [%]	Polarimeter
Source+Linac	1.1	82-84	
Booster	2.5	~80-84	
AGS	23.8	67-70	p-Carbon
RHIC	255	55-60	Jet, full store avg*

Loss in polarization along the chain

	Relative Ramp Polarization Loss (Run 17, full run avg)
AGS	17 %
RHIC	8 %

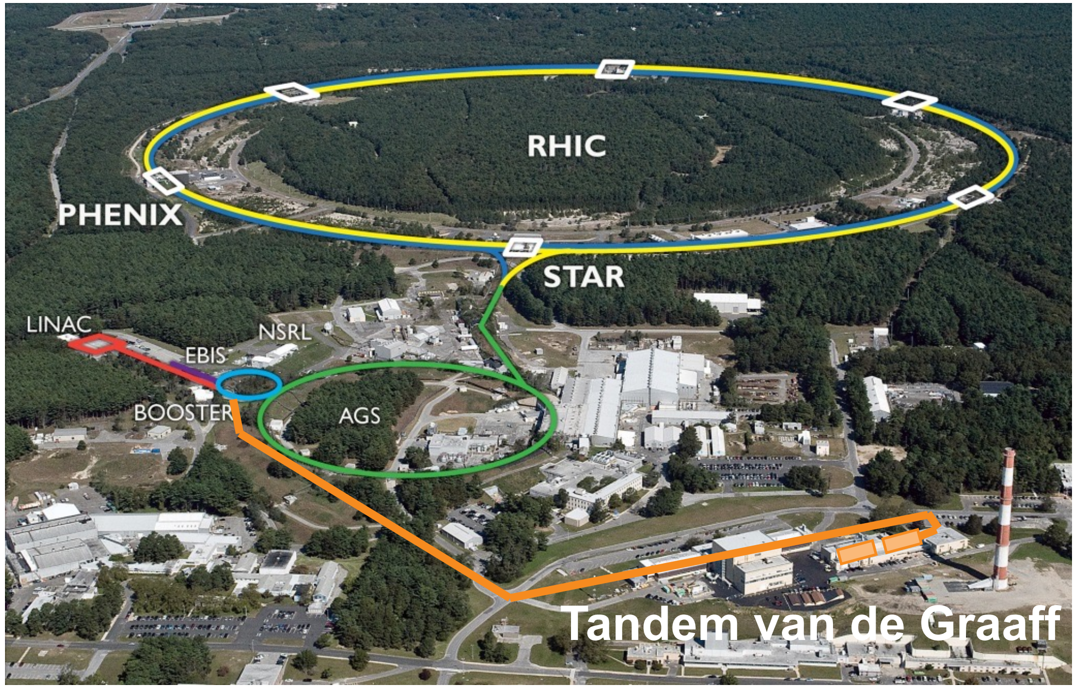

Polarimetry available at:
- Source
- End of Linac (200 MeV)
- AGS extraction
- RHIC injection energy
- RHIC flattop

<u>No Booster polarimeter</u>

Improve Polarization at RHIC

- Figure-of-merits (FOM) for the project ("experimental outputs"): emittance, beam intensity, polarization
- Trade-offs in optimizing **FOMs**:
  - Emittance ↓ Beam intensity ↑ Polarization ↑
- Trade-offs between **controls**:
  - Beam intensity ↑ → Emittance ↑
  - Emittance ↑ → Polarization ↓
- Main areas to optimize:
  - Booster injection / capture
  - AGS bunch splitting / merging scheme
  - AGS spin resonance compensation

AGS Polarization vs. Beam Intensity

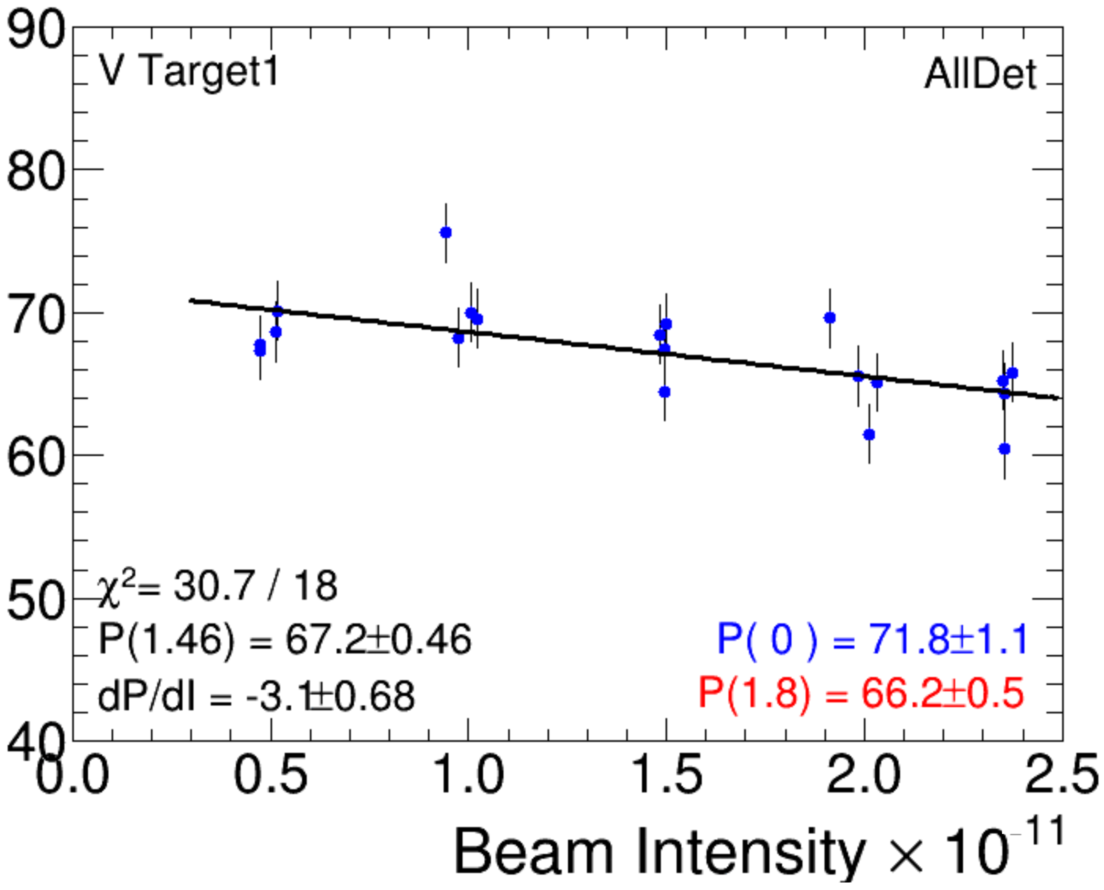

Polarization Improvement workflow

Data-model Integration

- Solve inverse problem for unknown model parameters
- Learn data-driven model for additional discrepancy

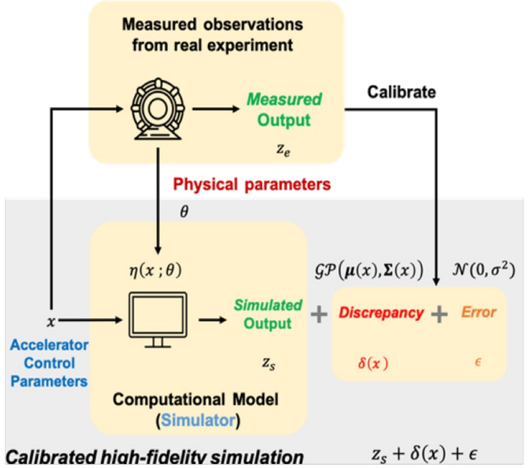

Scientific Machine Learning

- Include constraints for physics processes in surrogate model training
- "Soft" constraints as an objective penalty

Optimization with linear constraints

$$min_{\boldsymbol{x}\in D_{\boldsymbol{x}}} f(\boldsymbol{x}) \quad s.t. \;\; c_r(\boldsymbol{x}) \le 0 \quad \forall\, r \in [1,..,R]$$

Objective and constraints as GPs

$$f(\boldsymbol{x}) \sim \mathcal{GP}(\mu,\Sigma) \;\text{ and }\; c_r(\boldsymbol{x}) \sim \mathcal{GP}(\mu_r,\Sigma_r)$$

Integrate the feasibility through the CDF

$$cEI(\boldsymbol{x}) = EI(\boldsymbol{x}) \times \prod_{r=1}^{R} \Phi\left(\frac{\mu_r}{\sigma_r}\right)$$

Optimization under Uncertainty

- Bayesian optimization simultaneously trains a surrogate and identifies a maximum function evaluation

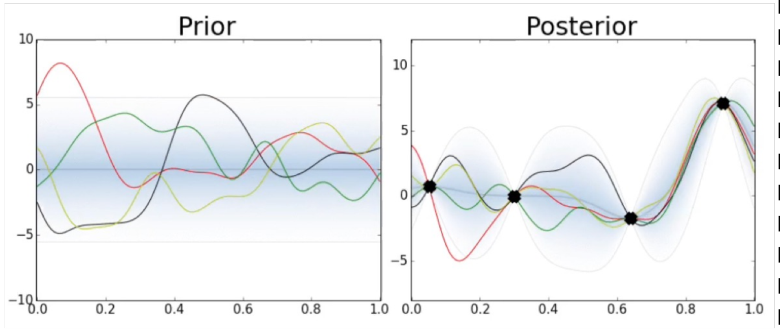

Planned project: Booter injection/capture

- Booster injection/early acceleration process sets maximum beam brightness for rest of acceleration though RHIC
- Linac pulse of 300 us, H^- beam ~6-9x10^{11} protons, strip through a carbon foil
- Intentional horizontal and vertical scraping reduce emittance (and intensity) to RHIC requirements ~2.5x10^{11} protons
- Goal: minimize beam loss at scraper
- Controls: Linac to Booster (LtB) transfer line optics
- Method: Bayesian Optimization

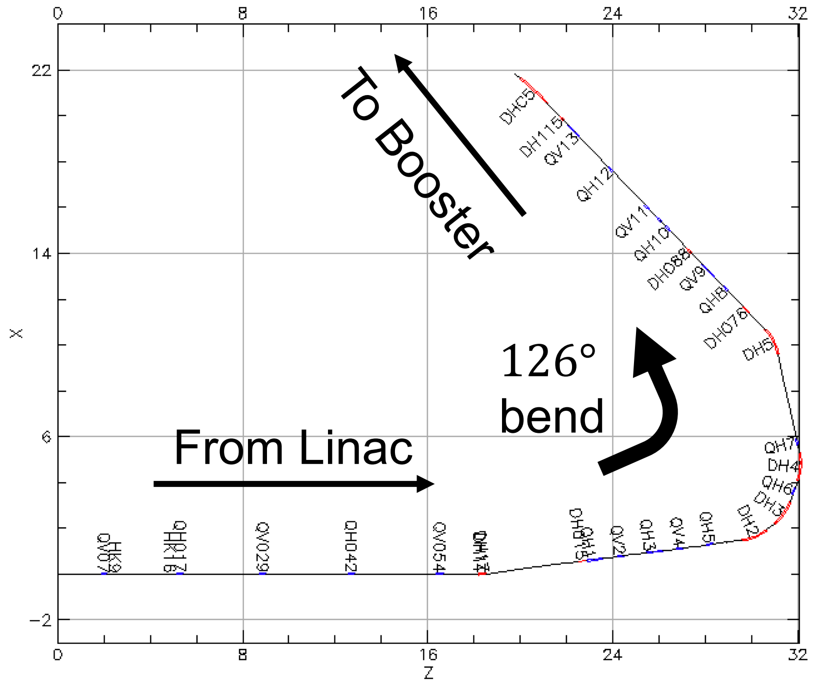

Planned project: AGS bunch splitting/merging

- Normal operation: One Linac pulse is captured as one bunch in the Booster and remains one bunch in AGS and RHIC
- Peak current (space charge) at AGS injection can be reduced by splitting the bunch into 2 longitudinally in Booster before transferring to AGS
- Bunches are later merged at AGS extraction
  - Requires expert tuning of many parameters, often done ‘by eye’
  - Prone to drift over time
- Goal: minimize longitudinal emittance
- Controls: RF voltages, phases
- Method: Reinforcement Learning

Merge of AGS proton bunches at flattop

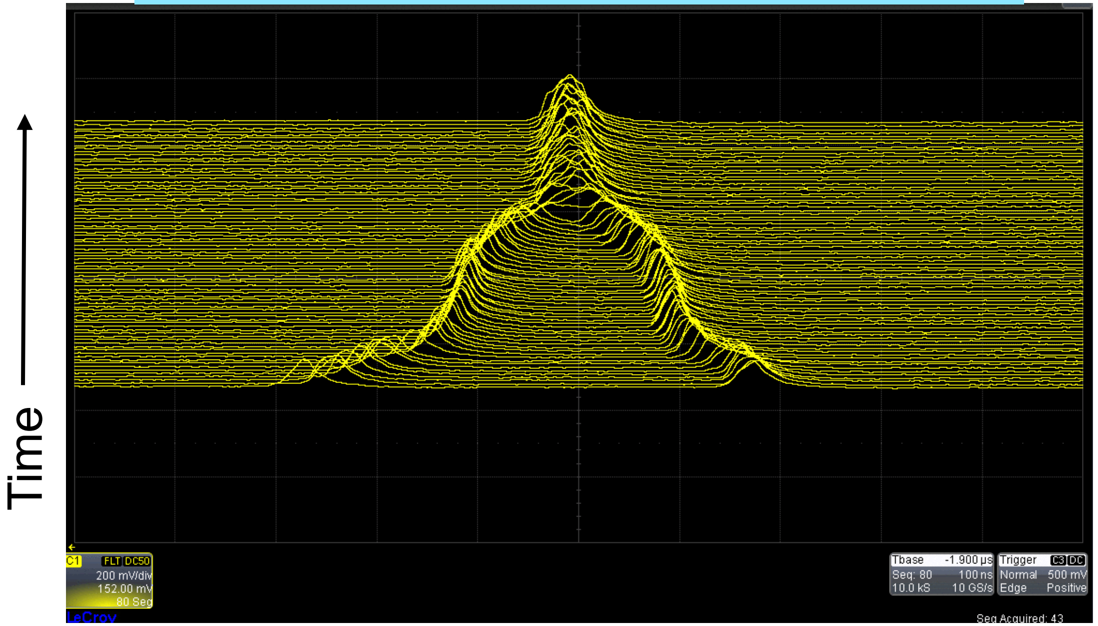

Planned project: AGS resonance compensation

- Partial snakes in the AGS keep the spin tune away from the integer (>0.96), avoiding vertical resonances
- Horizontal resonances remain, currently ‘jumped’ by moving the horizontal tune through the resonance
  - Each resonance is weak (~0.1% polarization loss)
  - But there are many of them (82)
- Proposal to use 15 pulsed skew quadrupoles to eliminate residual resonances
- Goal: minimize resonance strengths
- Controls: skew quadrupole currents
- Method: Reinforcement Learning / Bayesian Optimization (to be explored)

Betatron and spin tunes during AGS ramp

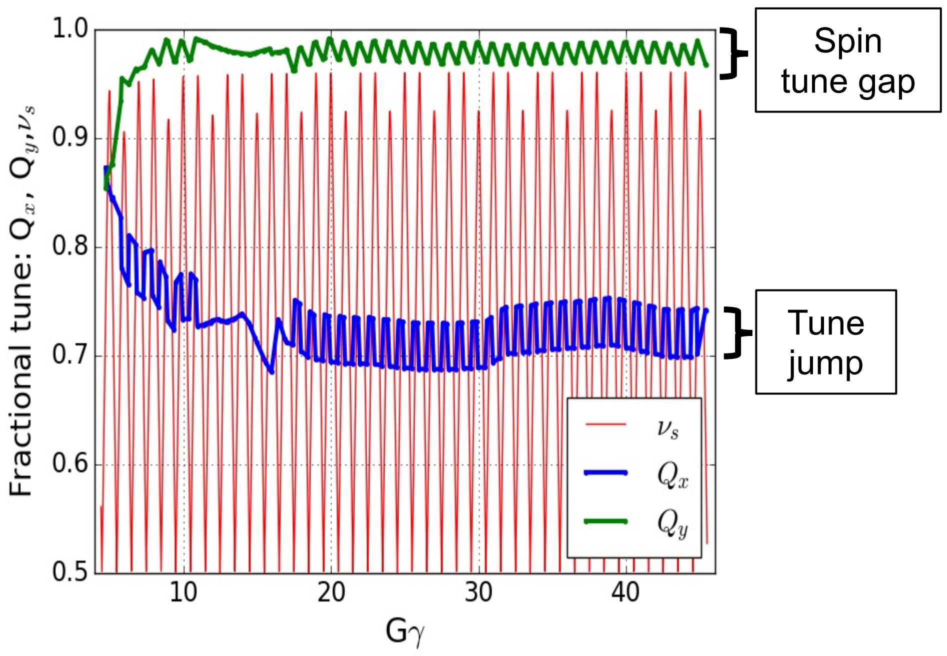

Spin resonance terms from skew quads in AGS

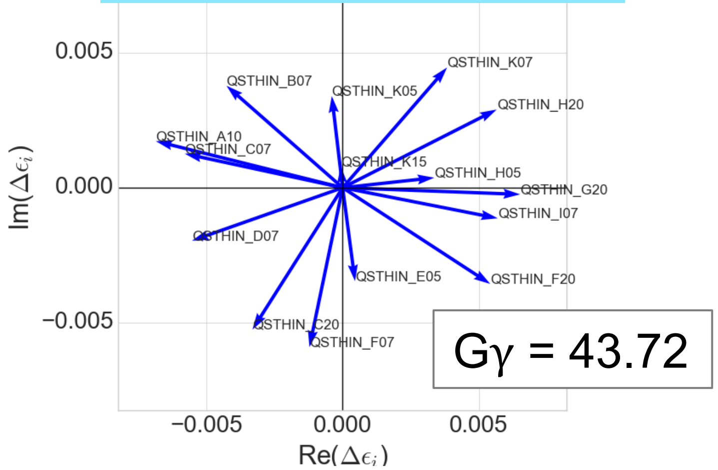

ML Method: Bayesian Optimization

- A powerful tool for finding the extrema of objective functions that are expensive to evaluate
- Bayes' theorem: probability of event based on previous knowledge of conditions

$$P(f|\mathcal{D}_{1:t}) \propto \boxed{P(\mathcal{D}_{1:t}|f)}P(f)$$

Tune hyperparameters of f to maximize likelihood of getting data $D_{1:t}$

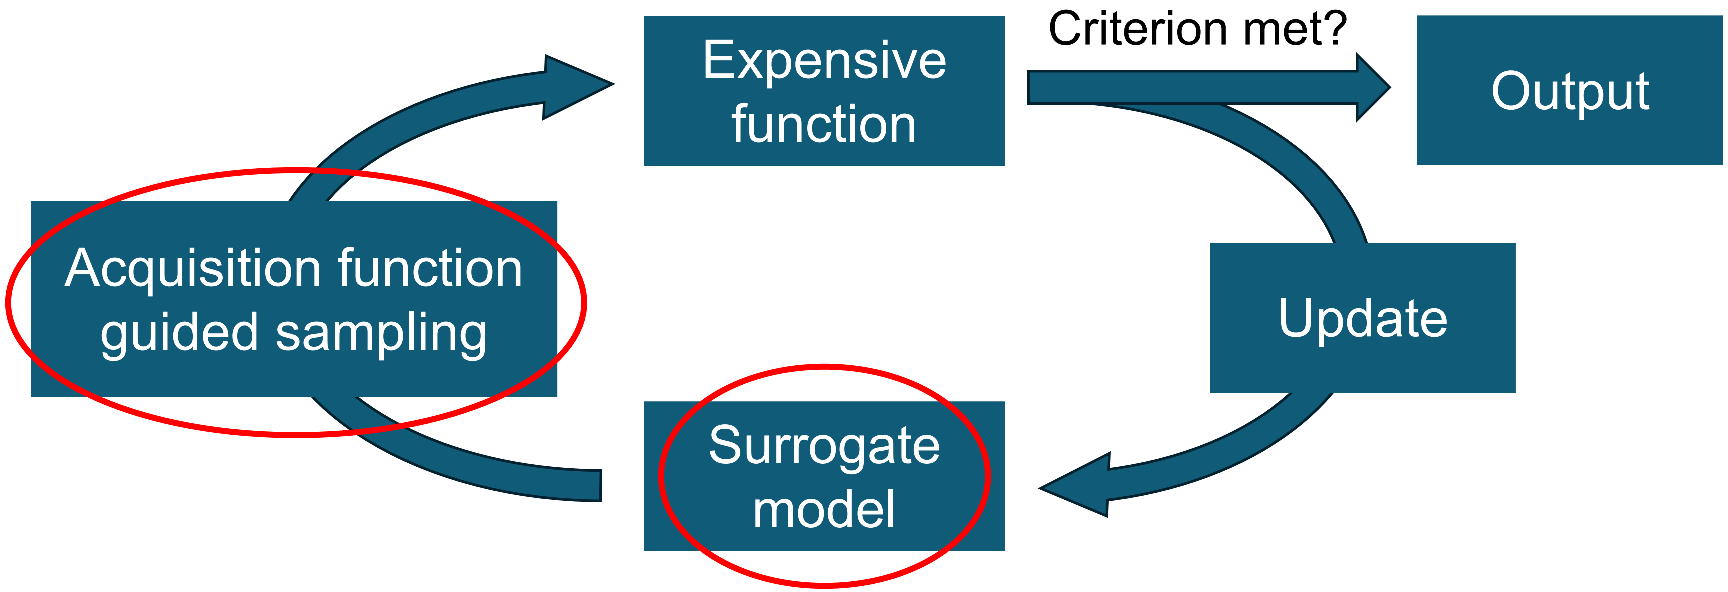

BO technique: Gaussian Process

- A probability distribution over possible functions that fit a set of points
- Mean function + Covariance function

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$$

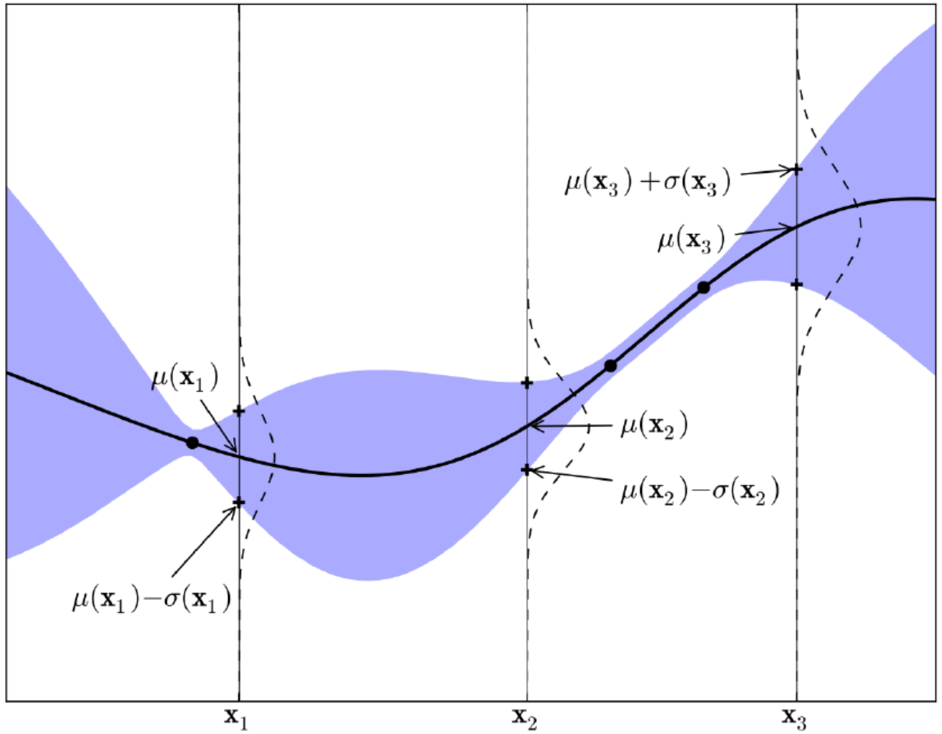

- Kernel: covariance function $k(x_i, x_j)$ of the input variables
- Covariance matrix $\mathrm{K} = k(X, X) = \begin{bmatrix} k(x_1, x_1) & \cdots & k(x_1, x_t) \\ \vdots & \ddots & \vdots \\ k(x_t, x_1) & \cdots & k(x_t, x_t) \end{bmatrix}$
- At a sample point x_i, Gaussian process returns mean $\mu(x_i|X) = m(x_i) + k(x_i, X)K^{-1}\big(f(X) - m(X)\big)$ and variance $\sigma^2(x_i|X) = k(x_i, x_i) - k(x_i, X)K^{-1}k(X, x_i)$

BO technique: Acquisition Function

- Guide how input space should be explored during optimization
- Combine predicted mean and variance from Gaussian Process model
  - Probability Improvement (PI)
  - Expected Improvement (EI)
  - **Upper Confidence Bound (UCB)**

$$\mathbf{UCB}(\boldsymbol{x}) = \boldsymbol{\mu}(\boldsymbol{x}) + \boldsymbol{\kappa\sigma}(\boldsymbol{x})$$

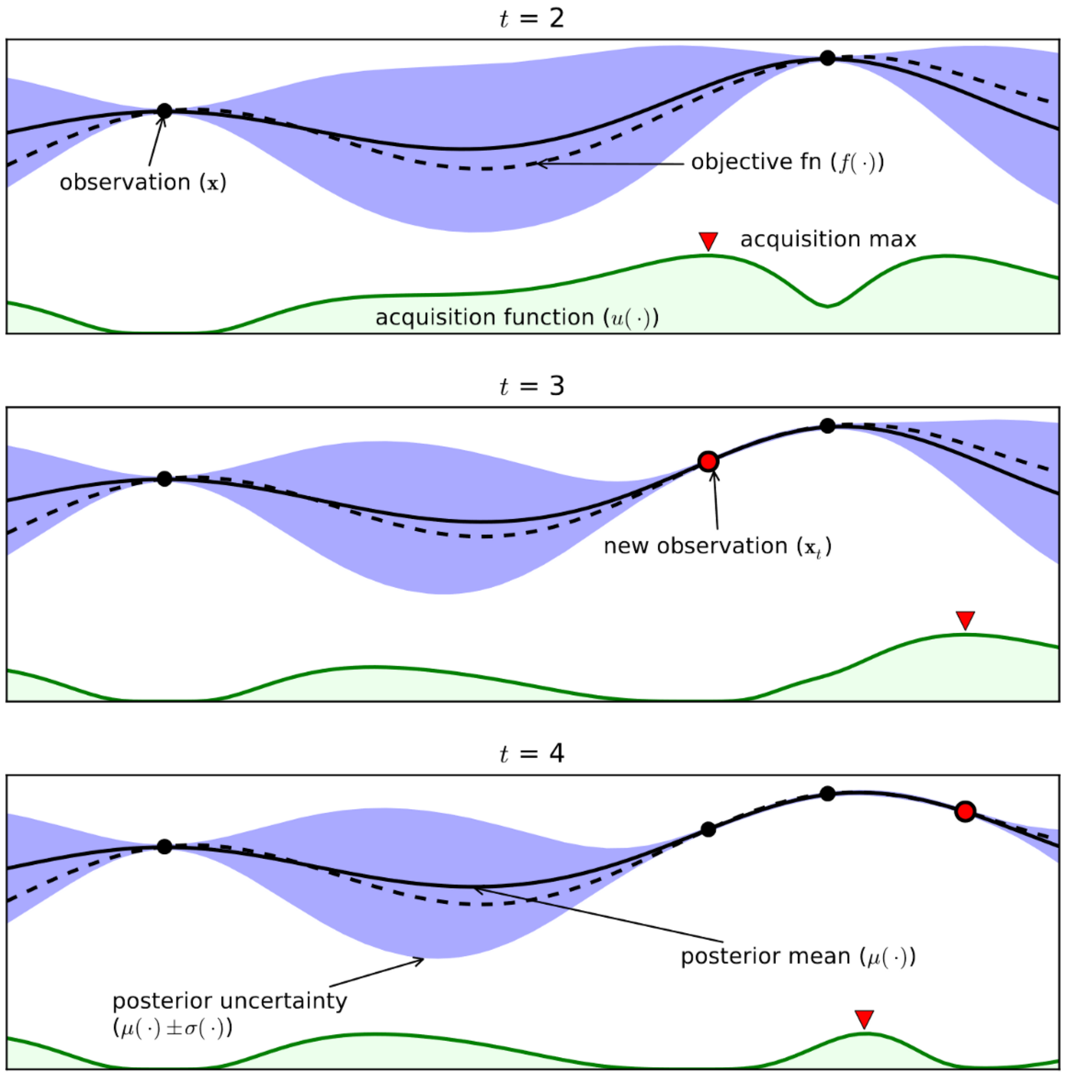

ML Method: Reinforcement Learning

- Learn optimal behavior in an environment to obtain maximum reward (e.g., highest polarization)
- Agent: controller, determine sampling policy
    - Action A: change control values
- Environment: controlled system
    - State S: representation of environment
    - Reward R: numerical evaluation of action
- Sequence of experience and agent forms trajectory $(S_0, A_0, R_0), (S_1, A_1, R_1), \cdots$

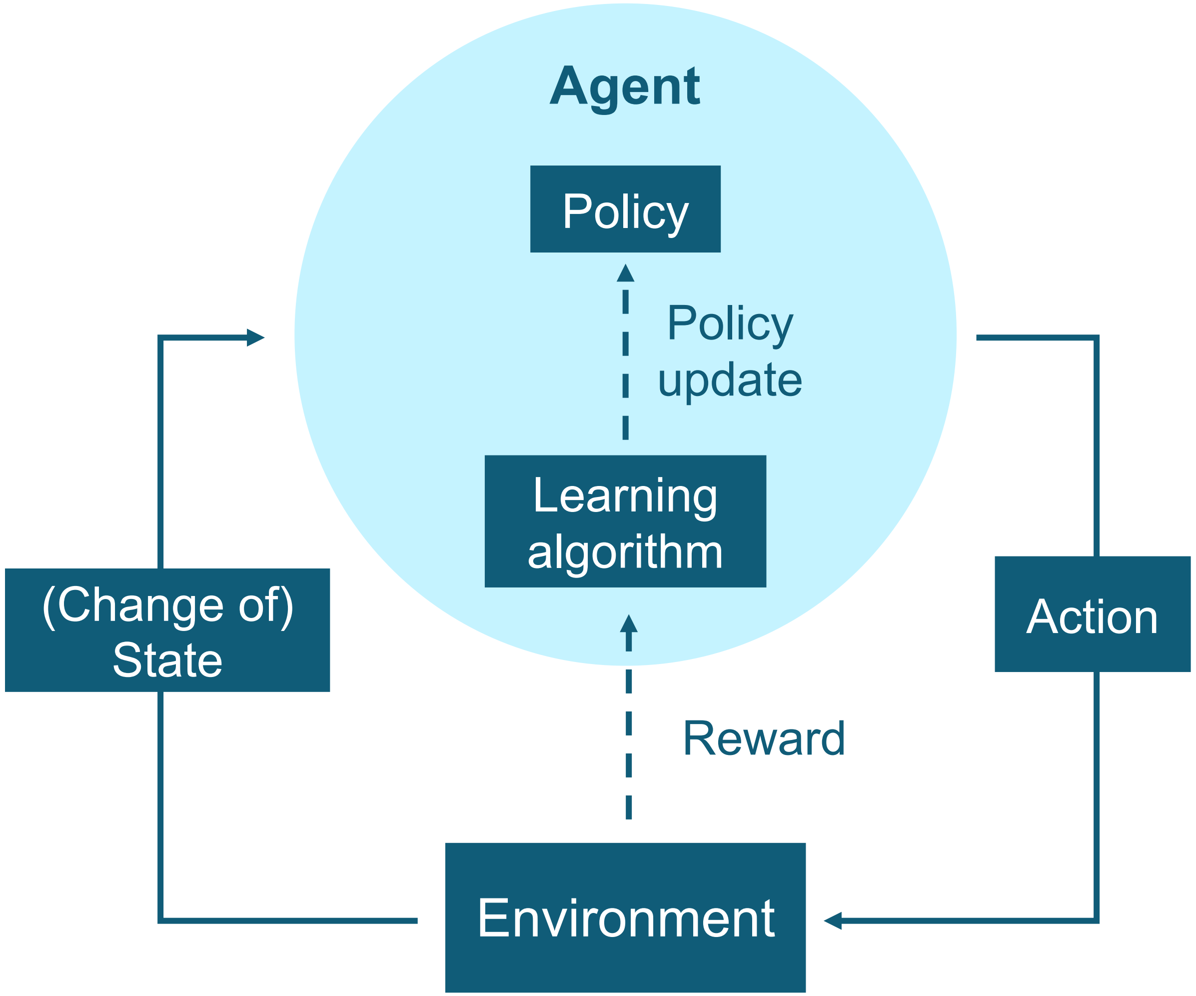

RL technique: Soft Actor-Critic (SAC)

- An entropy-based Reinforcement Learning (RL) aims to not only maximize total rewards, also to maximize the entropy of the policy

$$J(\pi) = \sum_{t=0}^{T} \mathbb{E}_{(\mathbf{s}_t,\mathbf{a}_t)\sim\rho_\pi} \left[r(\mathbf{s}_t, \mathbf{a}_t) + \alpha\mathcal{H}\big(\pi(\cdot\,|\mathbf{s}_t)\big)\right]$$

Final objective is weighted between a reward term r and an entropy term H by α

- SAC makes use of three networks: a state value function V parameterized by ψ, a soft Q-function Q parameterized by θ, and a policy function π parameterized by ϕ

- We can apply SAC to automatically tune RF phases and voltages so that a balanced beam profile can be achieved after bunch merge

Future: Digital twin and Optimal control

Control parameters + Environmental factors

Digital Twin

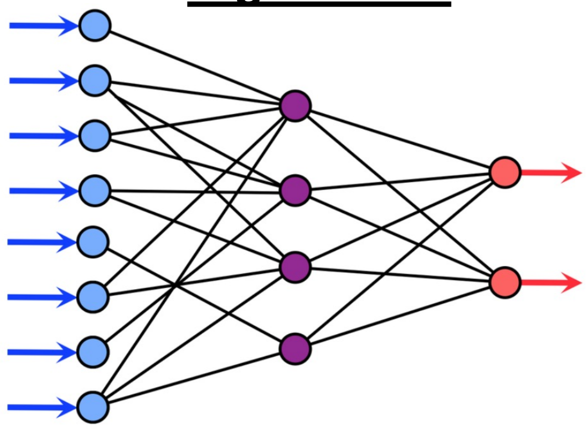

Evaluation / Analytic Metric (e.g., agreement with real measurements)

Find best settings for desired machine state

AI/ML Methods (RL, BO, etc.)

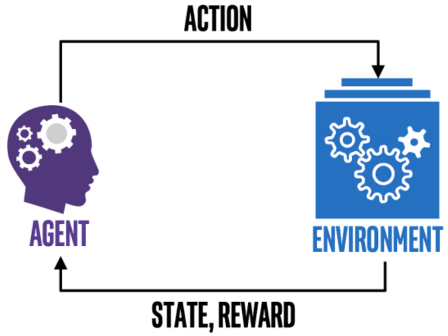

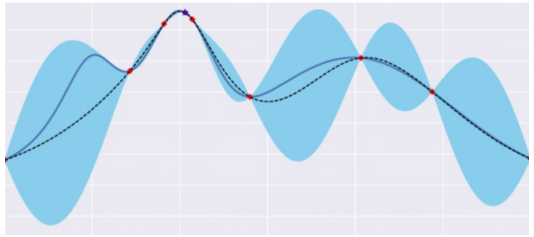

Thanks

- Kevin Brown, Yuan Gao, Levente Hajdu, Kiel Hock, Natalie Isenberg, Linh Nguyen, Vincent Schoefer, Nathan Urban, Keith Zeno
- Eiad Hamwi, Georg Hoffstaetter de Torquat, David Sagan
- Weining Dai, Bohong Huang, Thomas Robertazzi
- Yinan Wang
- Auralee Edelen
- Malachi Schram