Automating Particle Accelerator Operations with Machine Learning

Jonathan Edelen, Kevin Bruhwiler, Nathan Cook, Christopher Hall, and Stephen Webb

Kevin Brown, Vincent Schoefer, Philip Dyer, Auralee Edelen, and Dean Edstrom

*jedelen@radiasoft.net

October 29th 2021

2021 New York Scientific Data Summit

This material is based upon work supported by the U.S. Department of Energy, Office of Science,
Office of Nuclear Physics Award Number DE-SC0019682
and Office of High Energy Physics Award Number DE-SC0019979

Boulder, Colorado USA | radiasoft.net

Particle accelerators and their applications

Electron Accelerators

E-Beam sources
(Electron sources for a range of applications)
- Cancer treatment
- Polymer linking
- Wastewater treatment

Free electron lasers
(Highly tunable coherent photon sources ranging from IR to Hard X-Ray)
- Basic research
- Directed energy
- Industrial research

Synchrotron light sources
(Tunable photon sources usually in the x-ray)
- Basic research (NSLS-II, APS)
- Materials science, biology, etc.

Hadron accelerators

Proton accelerators
- High flux neutron sources (SNS)
- Accelerator driven subcritical reactors
- Basic research (fixed target and colliders)
- Cancer treatment

Ion accelerators
- Basic research (nuclear physics)
- Cancer treatment
- Isotope production

Particle accelerators and their applications

Electron Accelerators

E-Beam sources
(Electron sources for a range of applications)

- Cancer treatment
- Polymer linking
- Wastewater treatment

Free electron lasers
(Highly tunable coherent photon sources ranging from IR to Hard X-Ray)

- Basic research
- Directed energy
- Industrial research

Synchrotron light sources
(Tunable photon sources usually in the x-ray)

- Basic research (NSLS-II, APS)
- Materials science, biology, etc.

Hadron accelerators

Proton accelerators

- High flux neutron sources (SNS)
- Accelerator driven subcritical reactors
- Basic research (fixed target and colliders)
- Cancer treatment

Ion accelerators

- Basic research (nuclear physics)
- Cancer treatment
- Isotope production

See DOE Report "Accelerators for America's Future"
(https://science.energy.gov/~/media/hep/pdf/accelerator-rd-stewardship/Report.pdf)

Overview of accelerator operations

Accelerator R+D

Machine Development Time

Specialized R+D Facilities

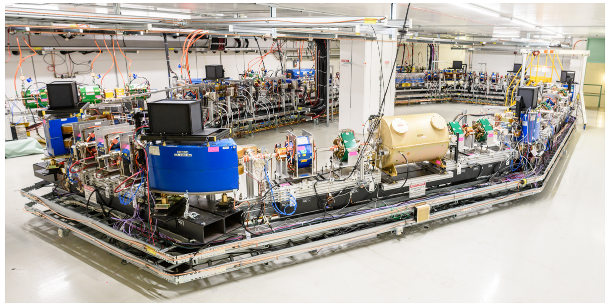

Beam for Experimentalists

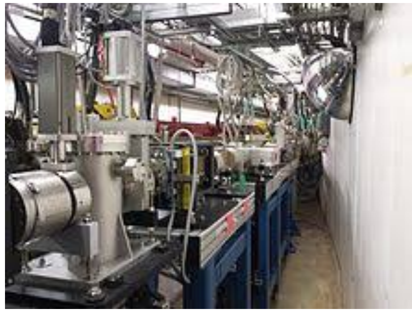

Small single user end stations

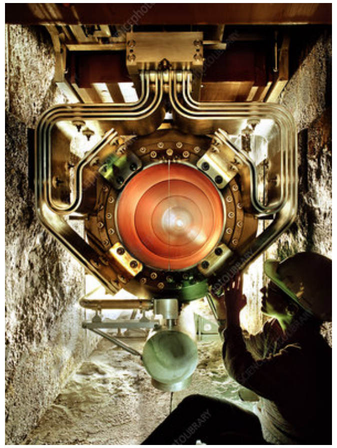

Large experimental collaborations

Down Time

Scheduled Maintenance

Unscheduled Maintenance

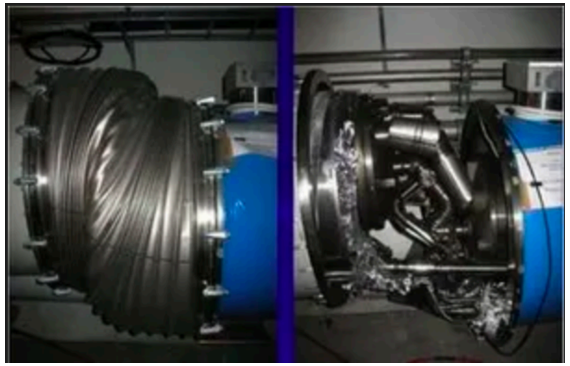

Machine learning applications for accelerators

Accelerator R+D

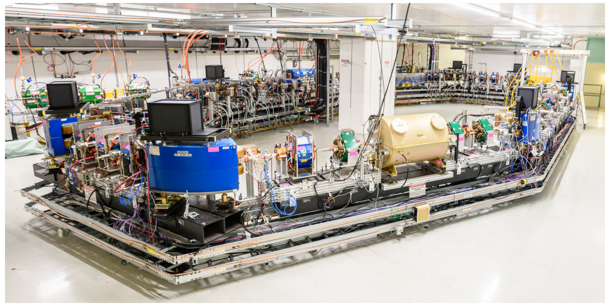

Beam for Experimentalists

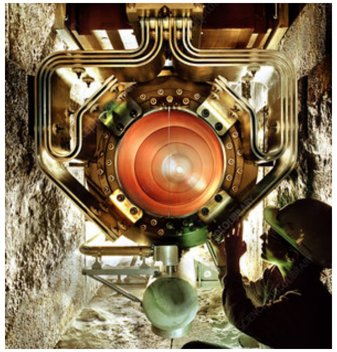

Down Time

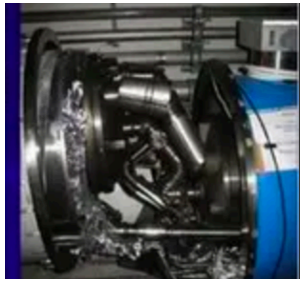

Tuning

- Inverse Models
- Neural Network + Optimization

Surrogate Modeling

- Neural Network
- Neural Network + Optimization

Anomaly Detection

- Inverse Models
- Autoencoders

Machine learning applications for accelerators

Accelerator R+D

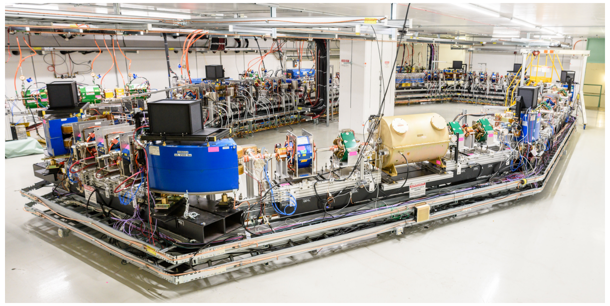

Beam for Experimentalists

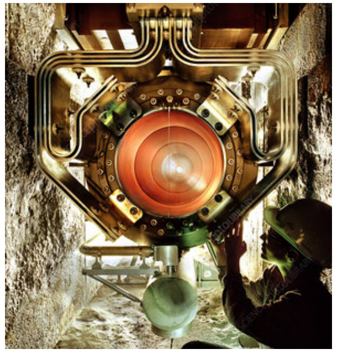

Down Time

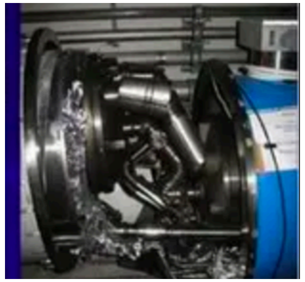

Tuning
- Inverse Models
- Neural Network + Optimization

Surrogate Modeling
- Neural Network
- Neural Network + Optimization

Anomaly Detection
- Inverse Models
- Autoencoders

Inverse models for tuning

- Direct use of inverse models for tuning
  - Train a model to predict settings from desired diagnostic output
  - A common application is beam steering
    - Inputs are the requested BPM readings
    - Outputs: Suggested corrector settings

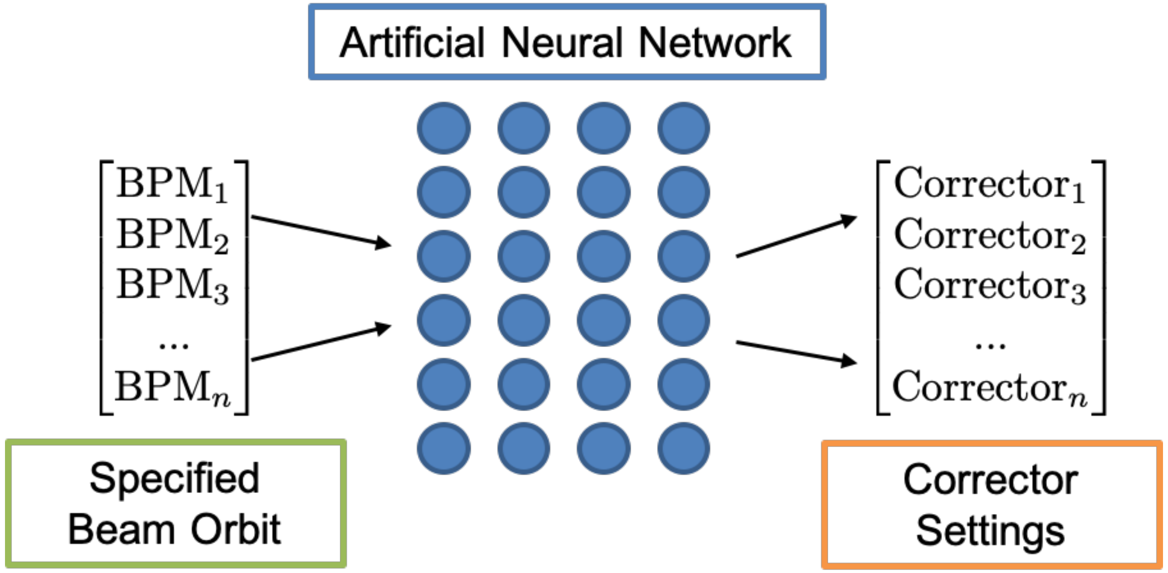

Inverse models for tuning

- Direct use of inverse models for tuning
  - Train a model to predict settings from desired diagnostic output
  - A common application is beam steering
    - Inputs are the requested BPM readings
    - Outputs: Suggested corrector settings

- Use inverse model as a starting point for optimization
  - Speeds up switching between beamline configurations

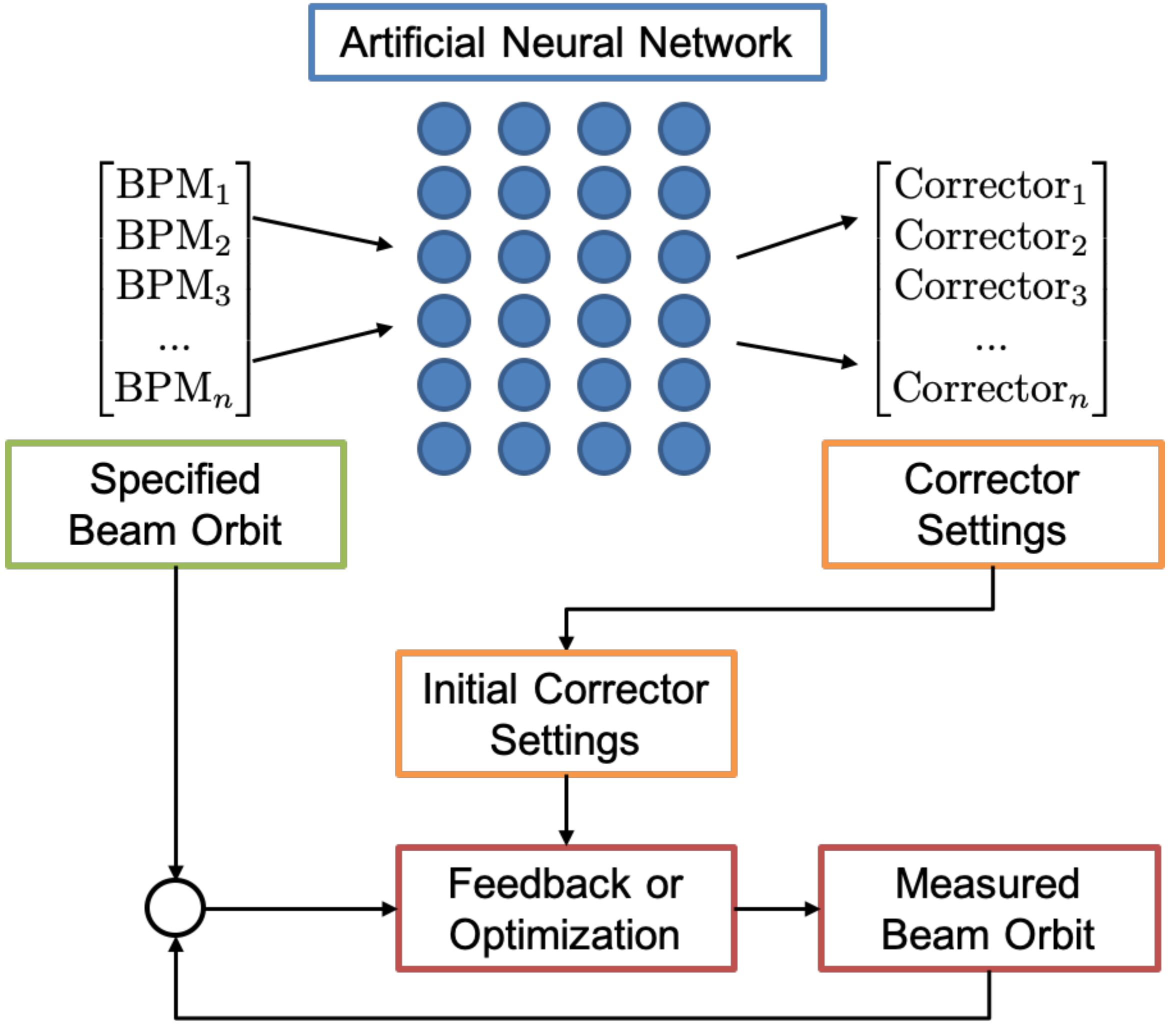

AGS to RHIC transfer line

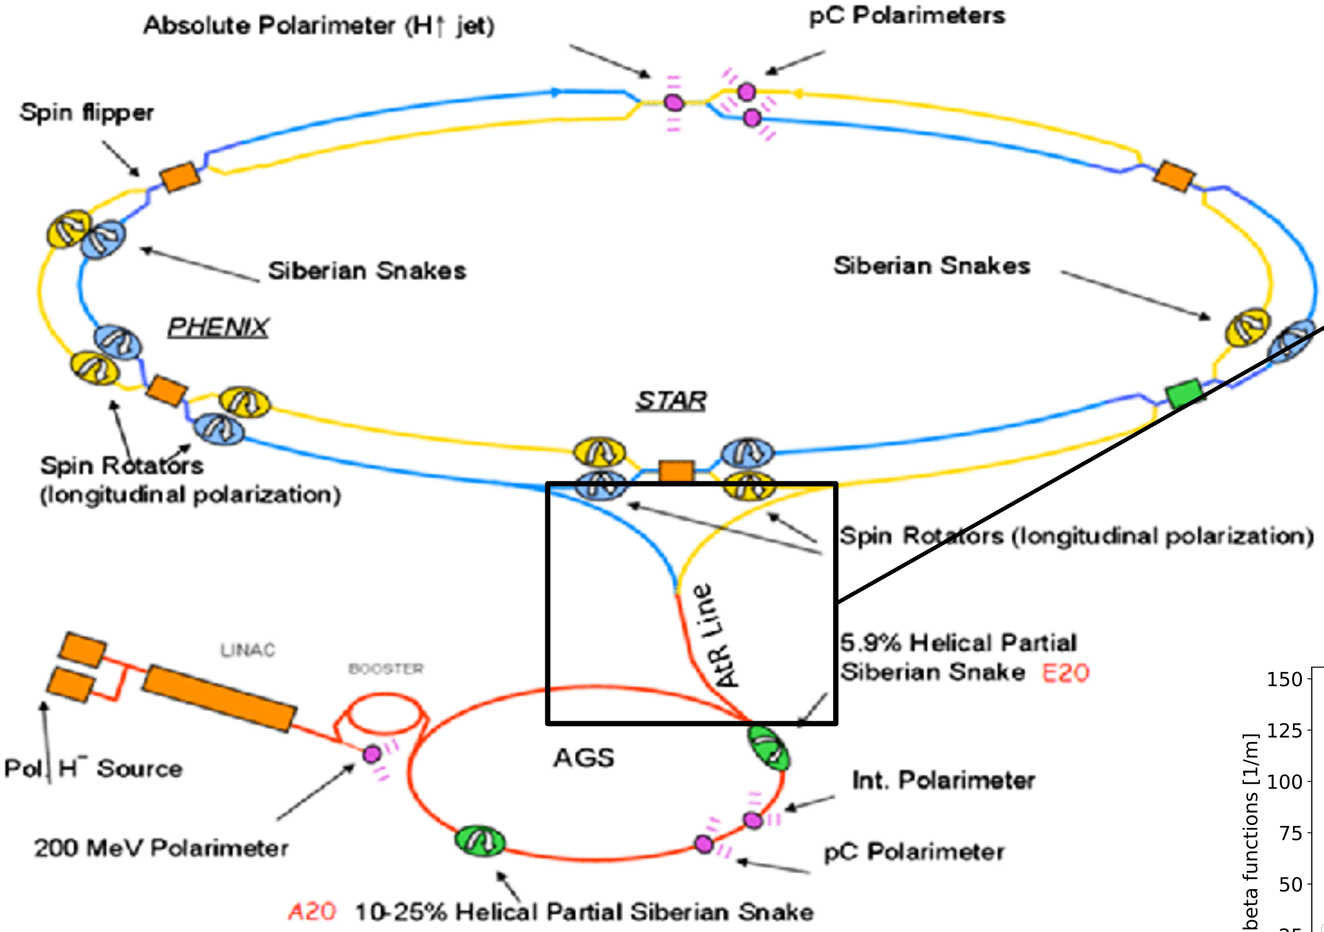

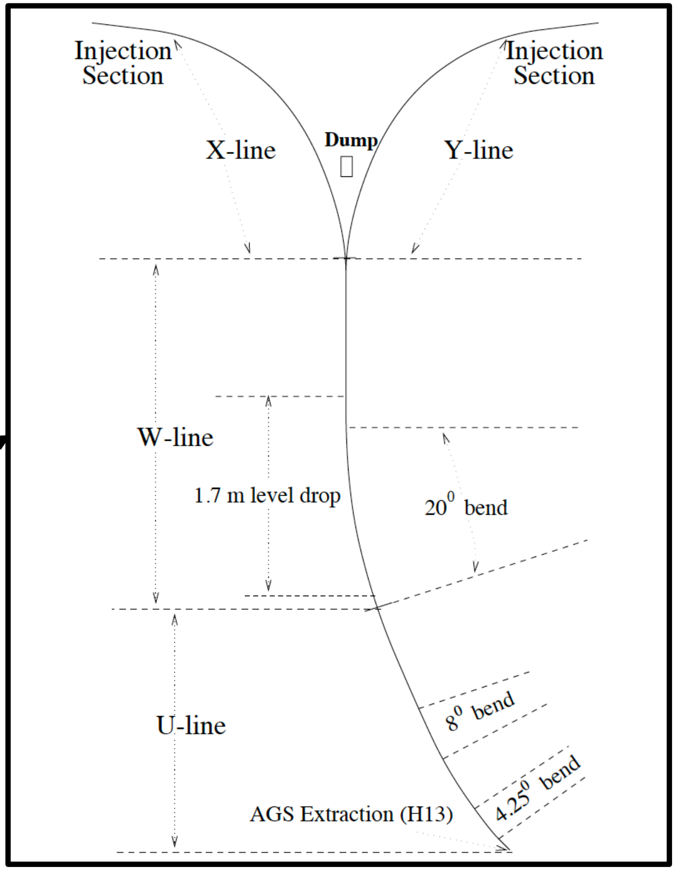

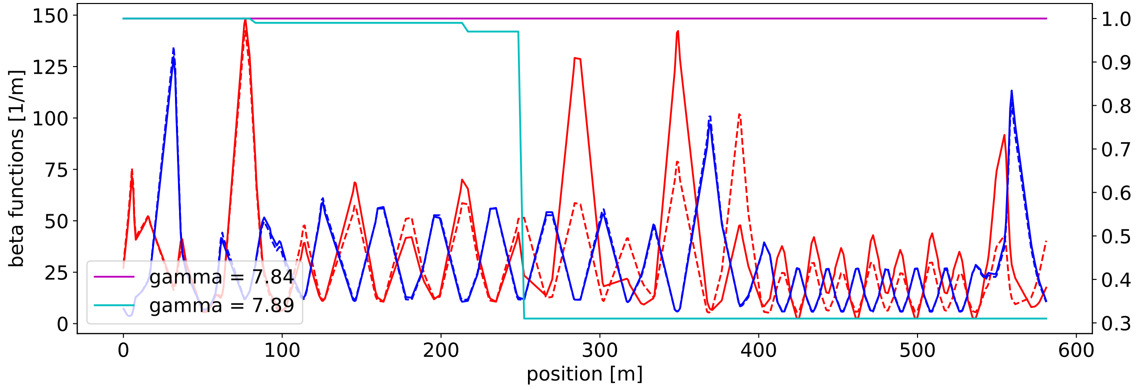

Beam steering with the AGS to RHIC transfer line

- Machine Learning (top)
  - Build inverse model of bpm-readings to corrector settings
  - Make feed-forward correction
  - Inverse models are fast and effective

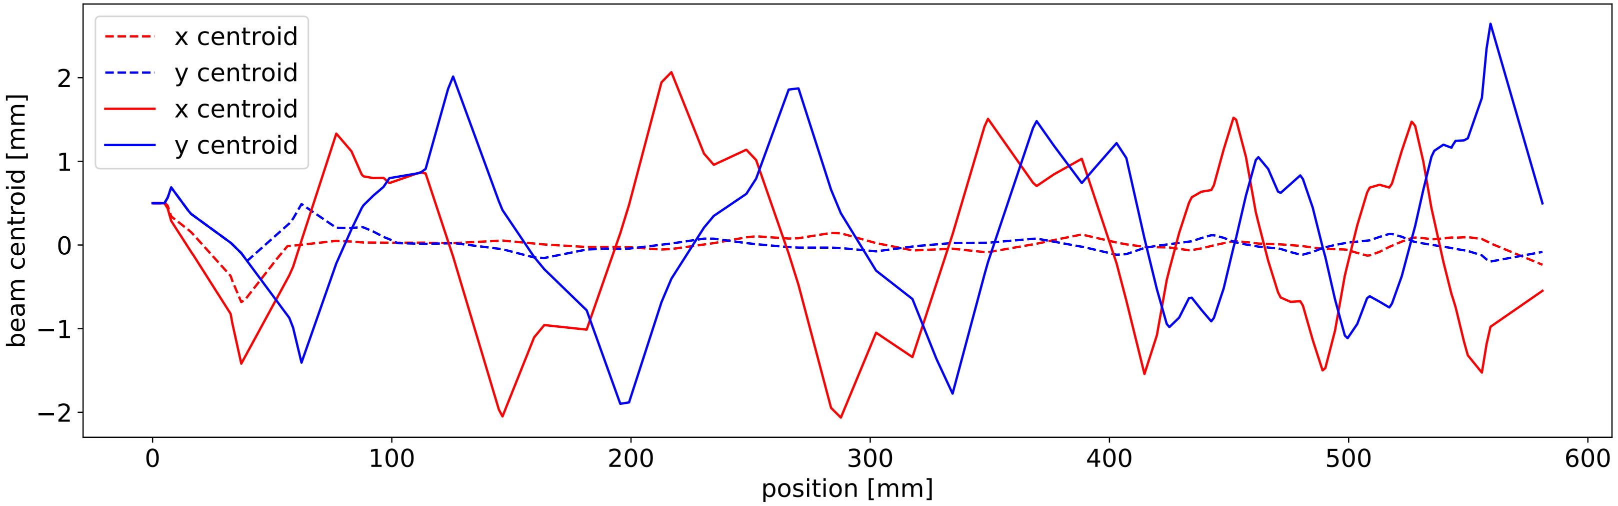

Beam steering with the AGS to RHIC transfer line

- Machine Learning (top)
  - Build inverse model of bpm-readings to corrector settings
  - Make feed-forward correction
  - Inverse models are fast and effective

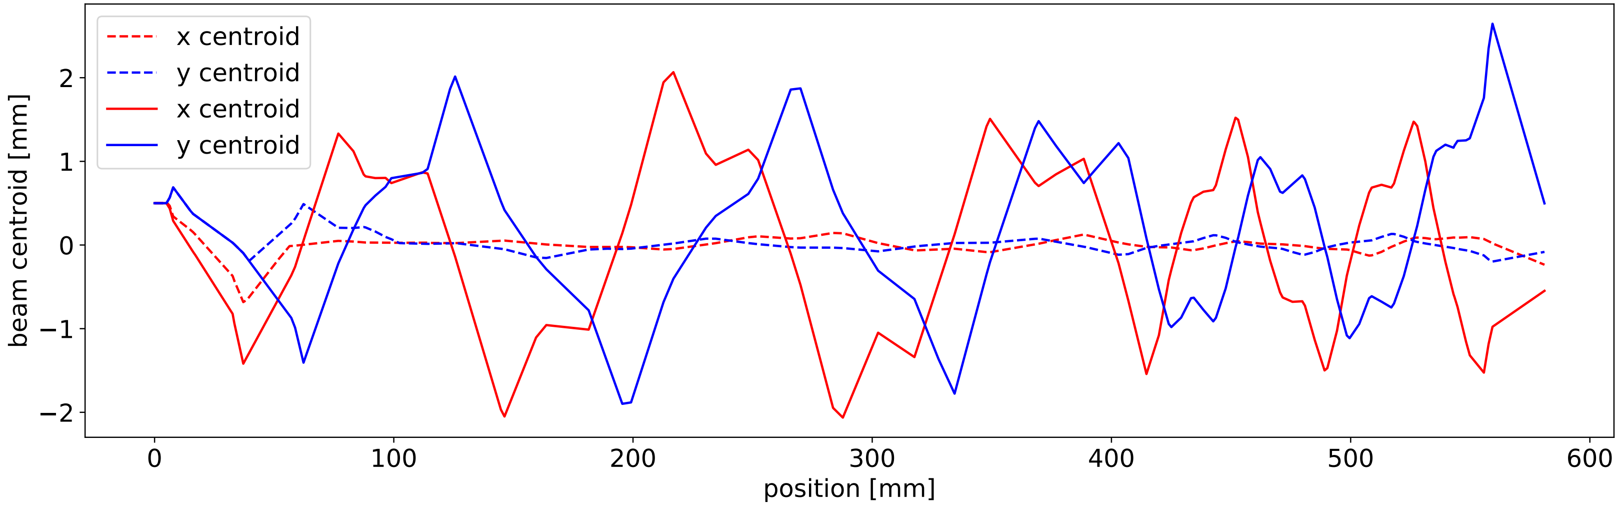

- Machine learning + optimization (bottom)
  - Connect accelerator simulation simulation to python optimization tools using our middle layer
  - Use output of the neural network as a starting point for a Nelder-Mead optimization

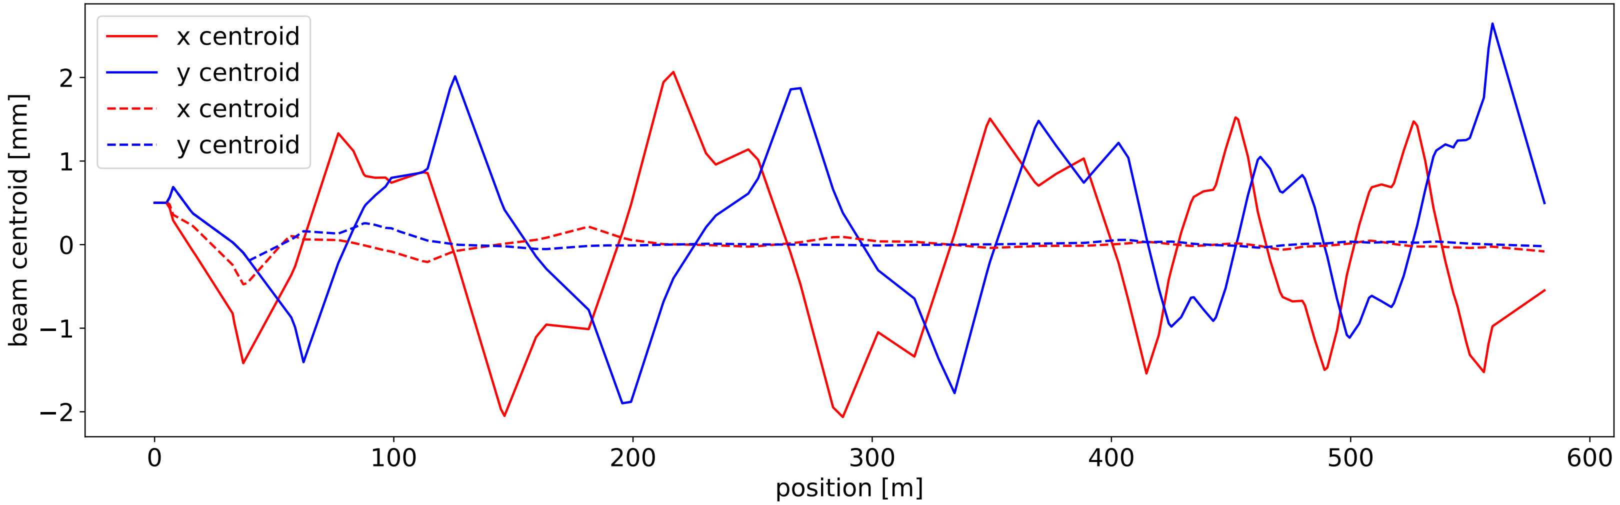

Machine learning for accelerator control

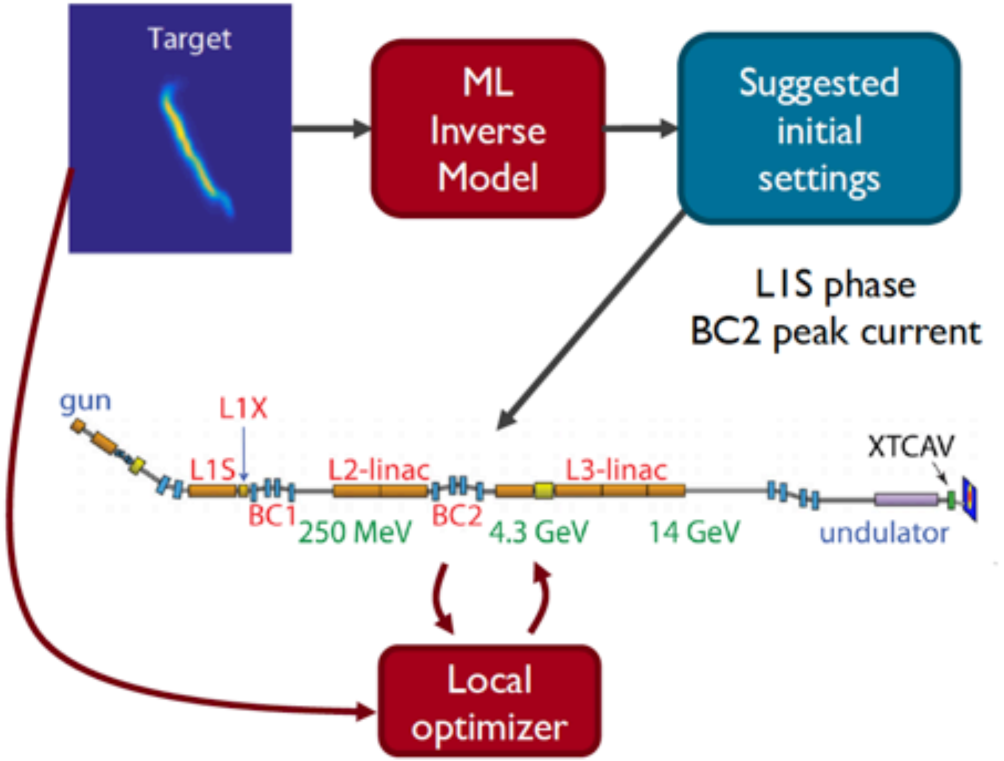

A. Scheinker, A. Edelen, et al., PRL 121, 044801 (2018)

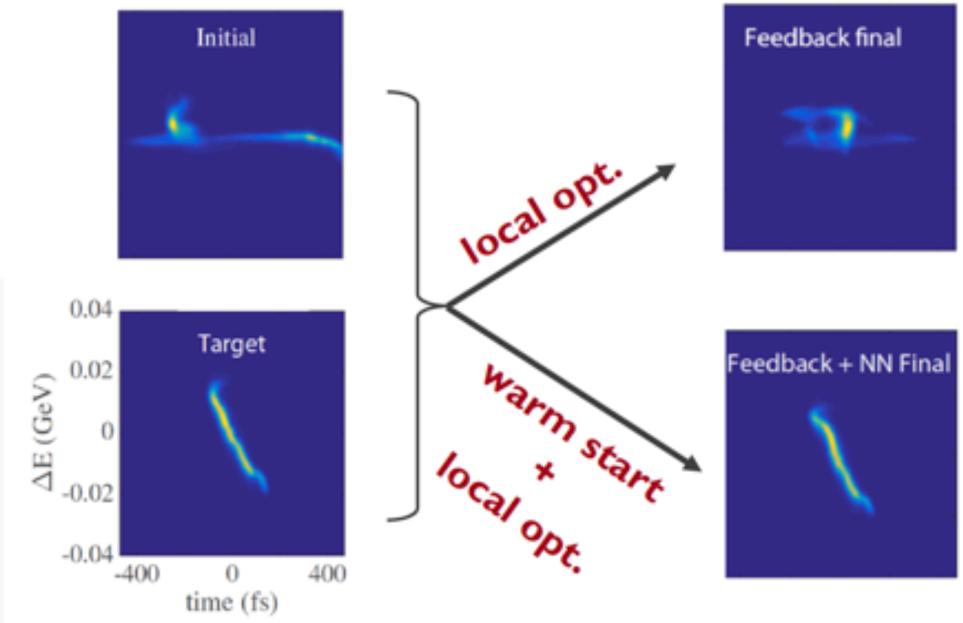

Local optimizer alone was unable to converge → **able to converge after initial settings from neural network**

Machine learning applications for accelerators

Accelerator R+D

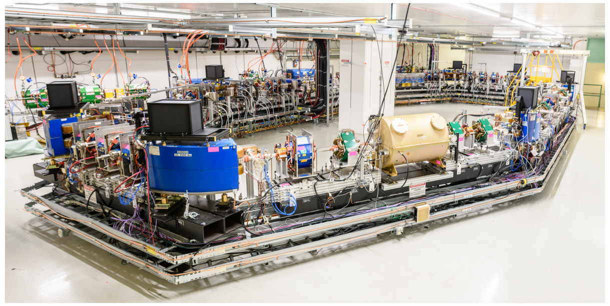

Beam for Experimentalists

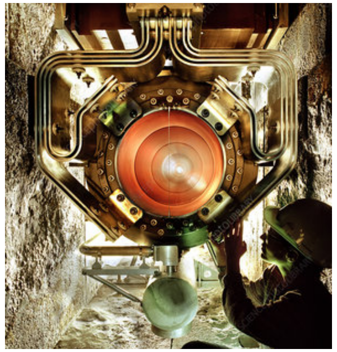

Down Time

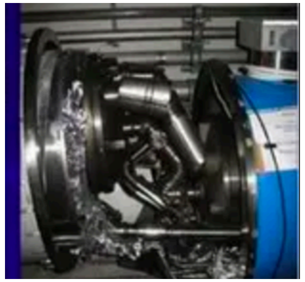

Tuning

- Inverse Models
- Neural Network + Optimization

Surrogate Modeling

- Neural Network
- Neural Network + Optimization

Anomaly Detection

- Inverse Models
- Autoencoders

Neural networks for surrogate models

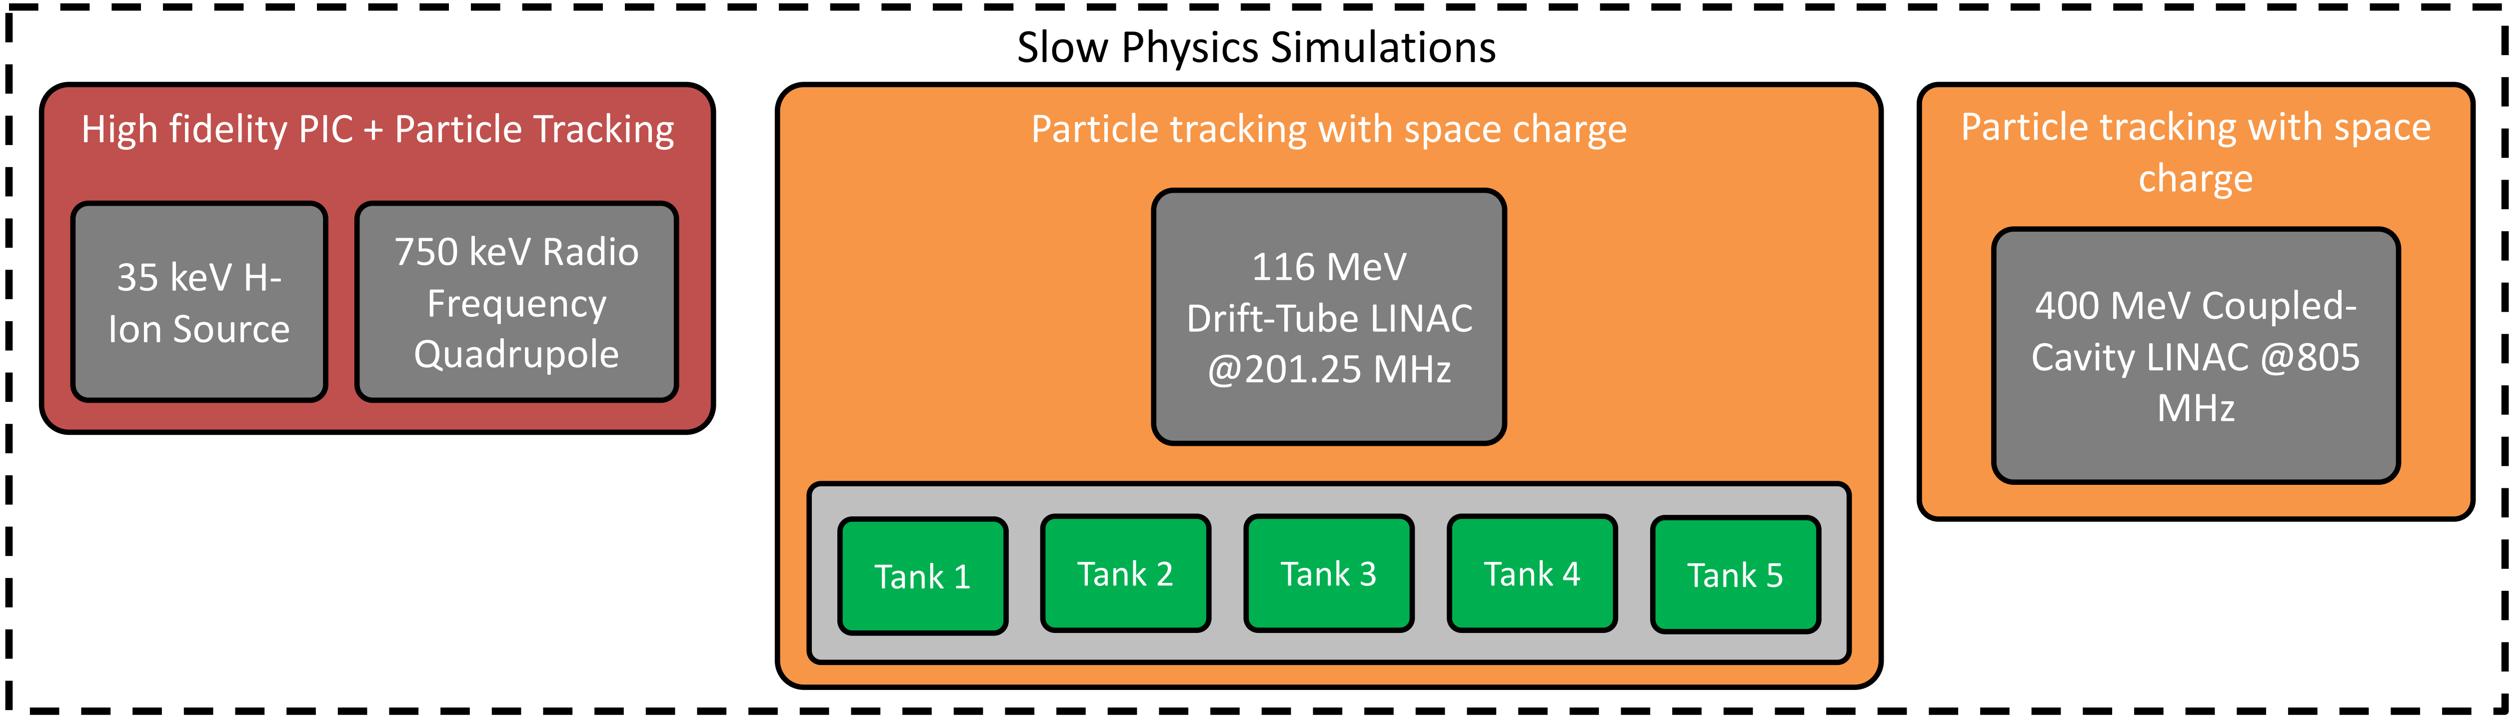

Neural networks for surrogate models

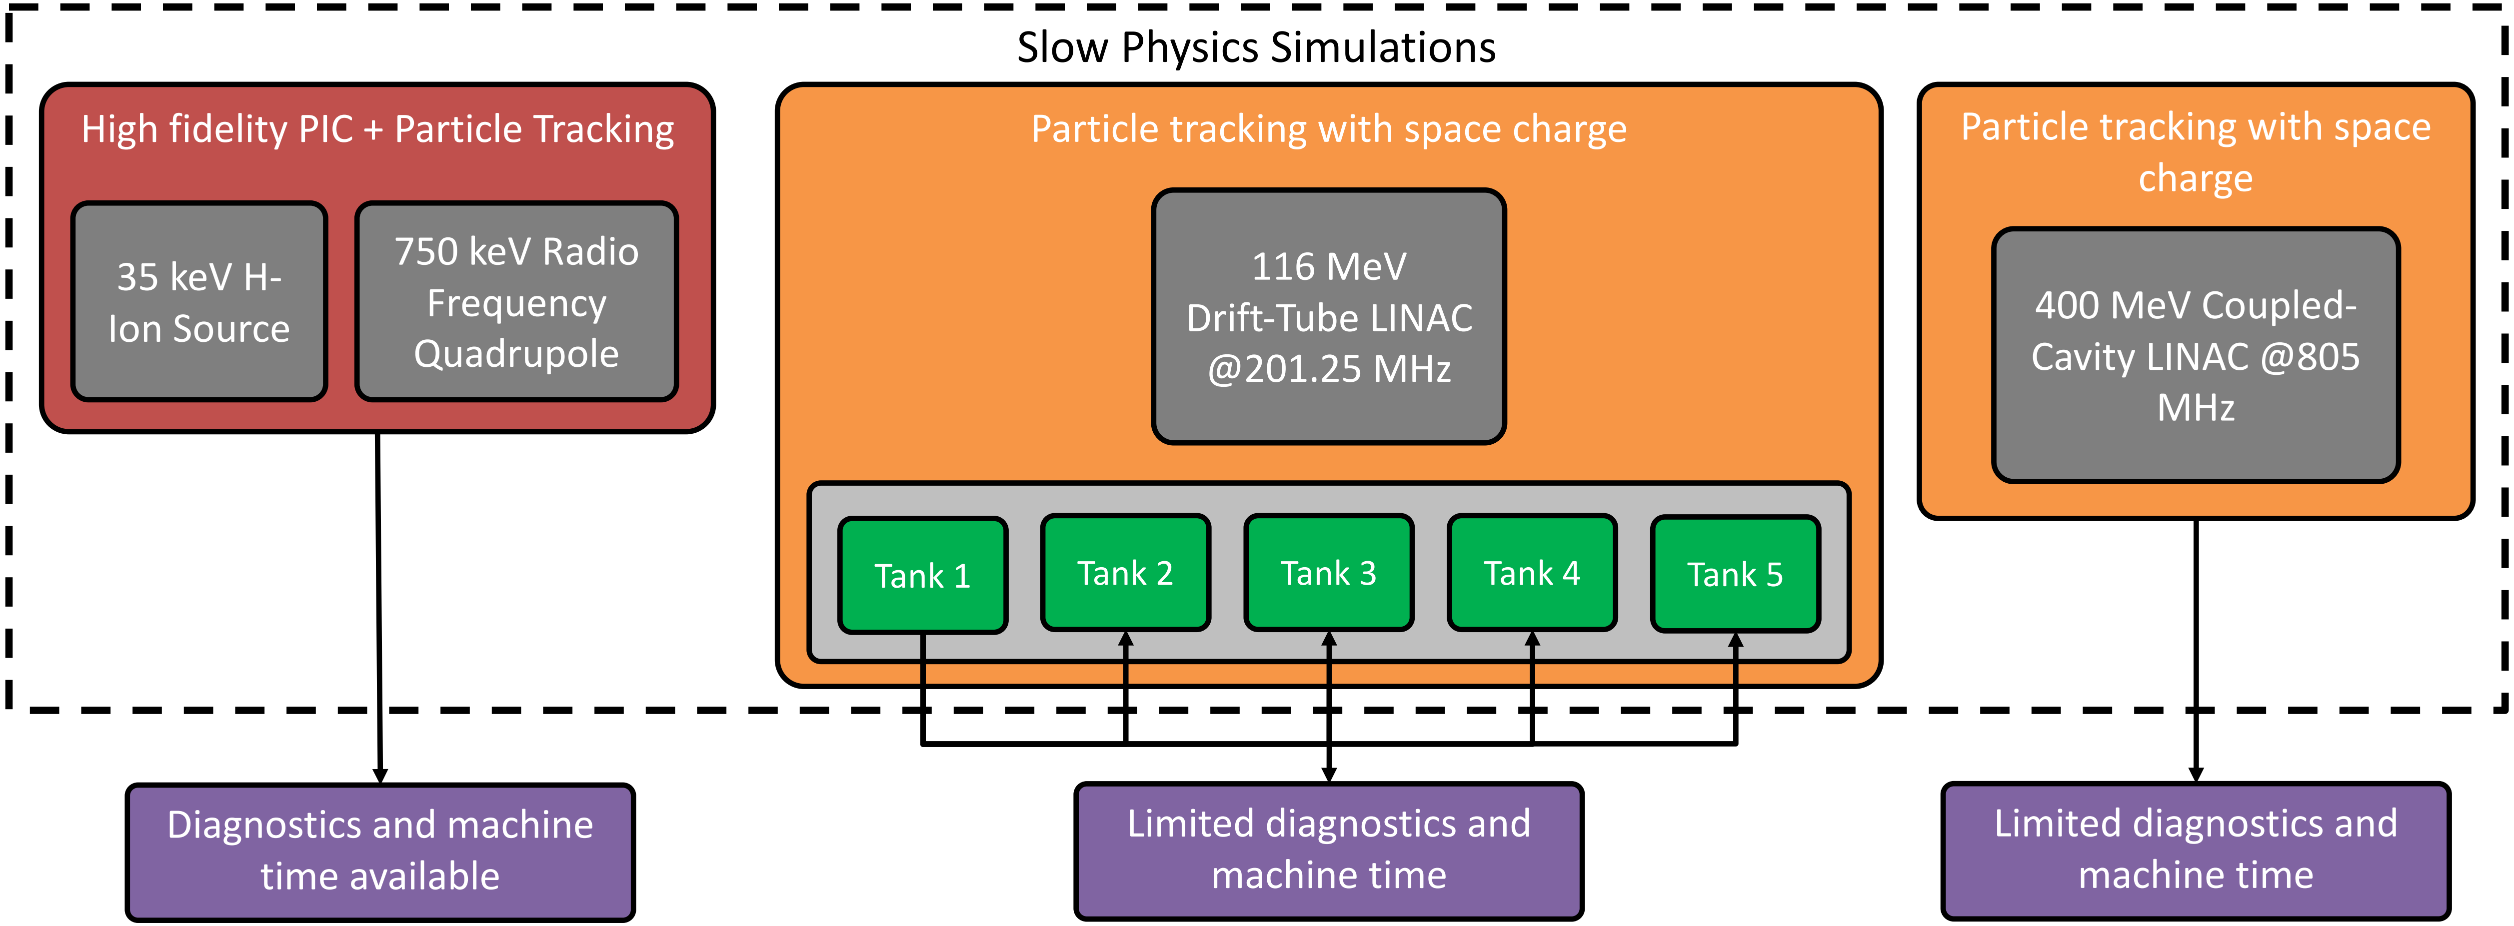

Neural networks for surrogate models

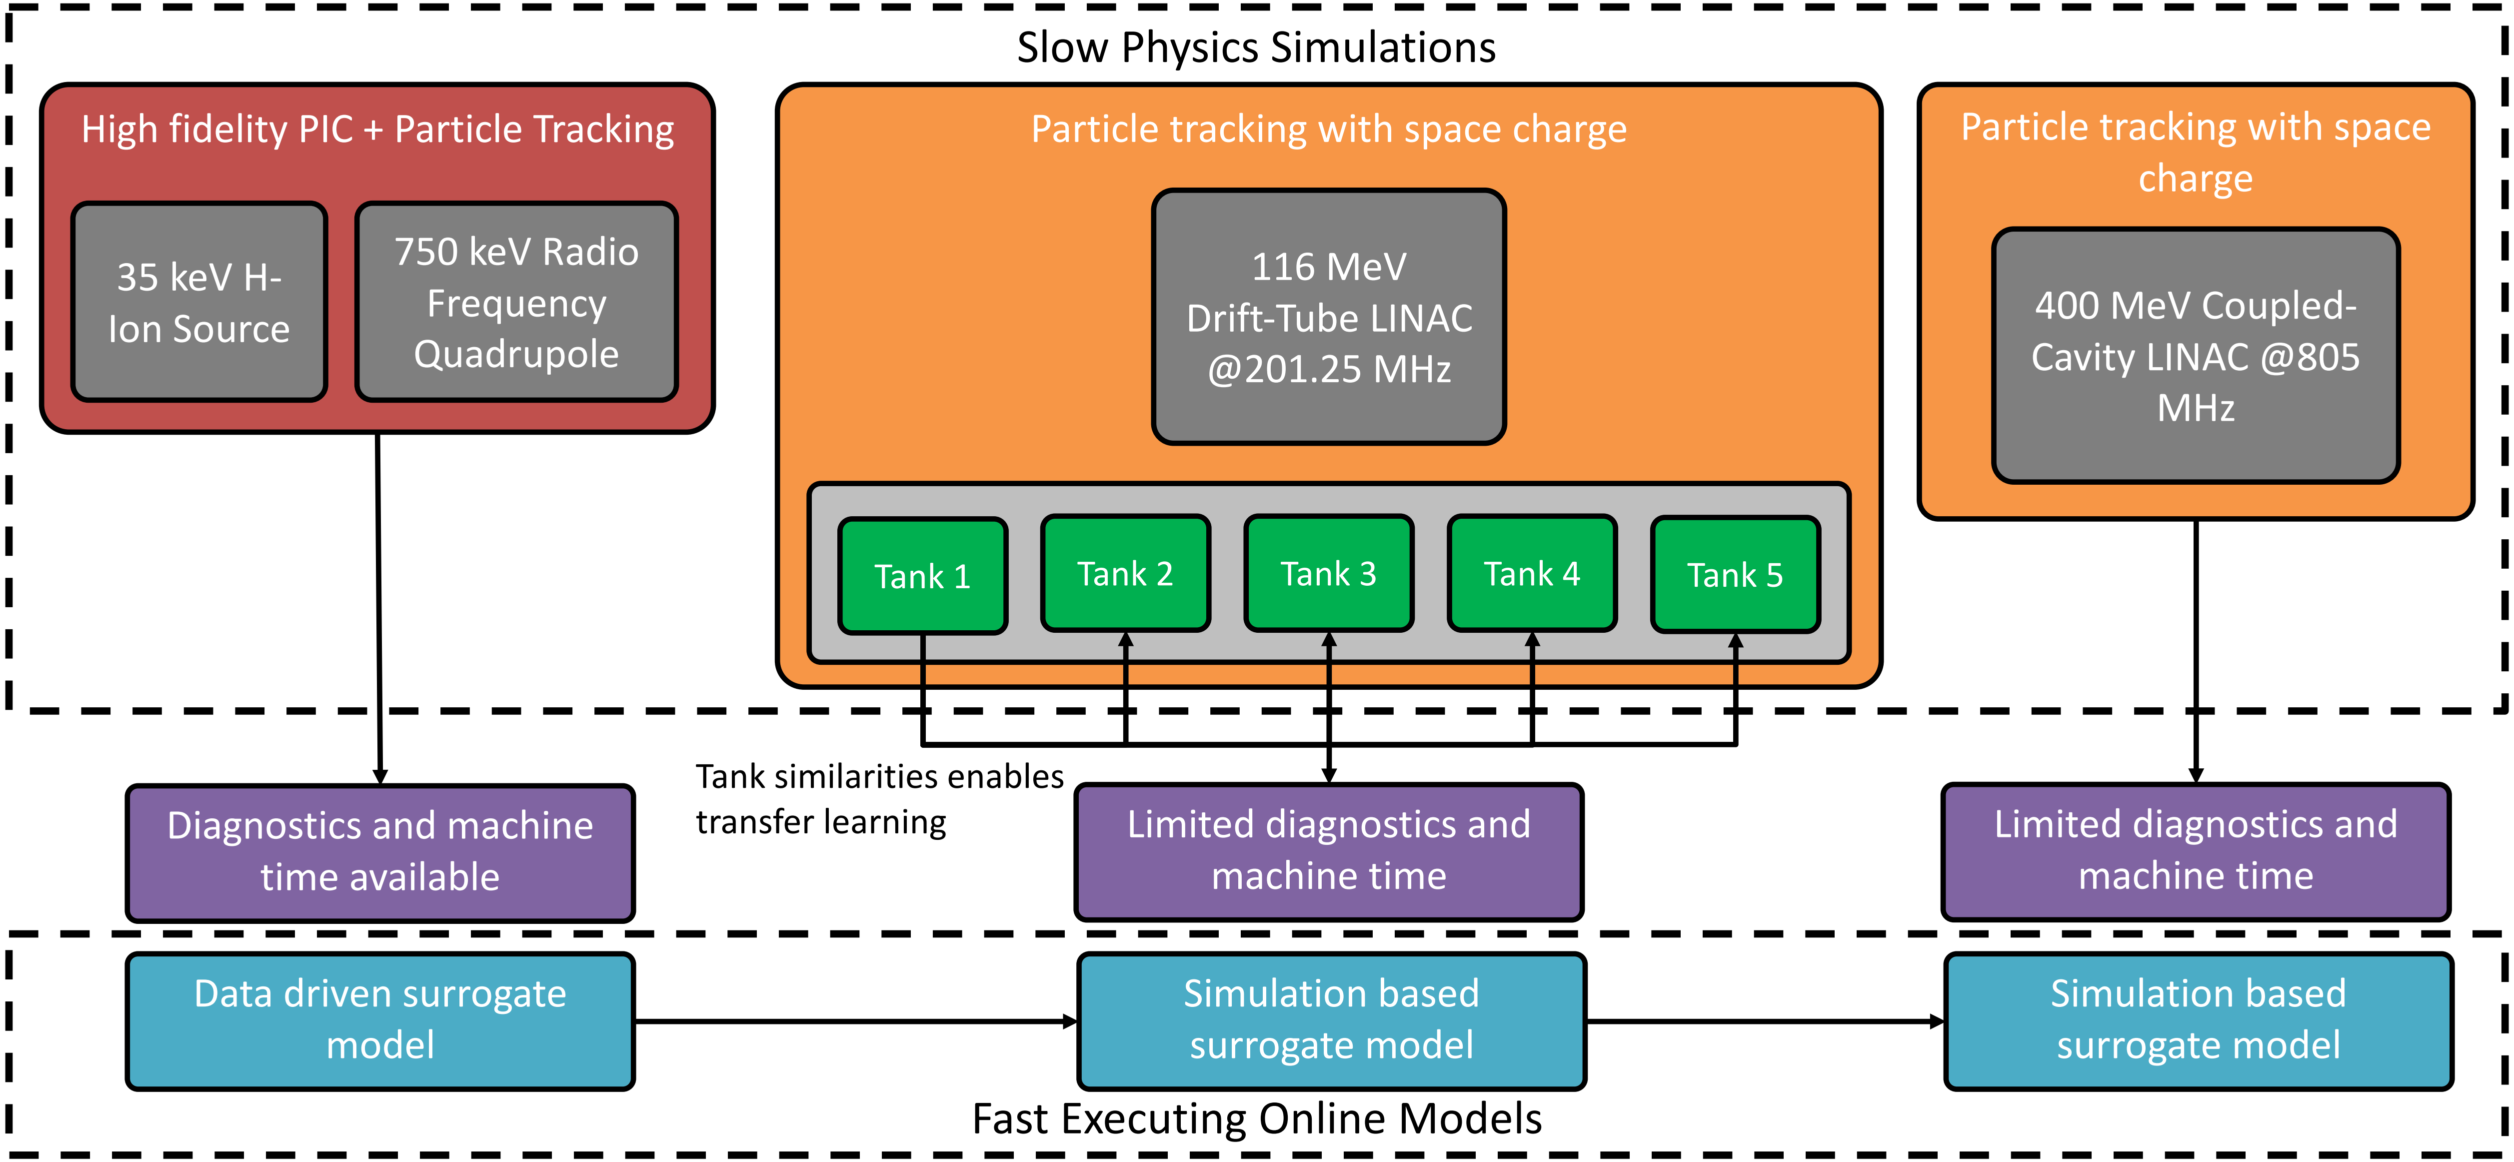

Transfer learning enables portable solutions between accelerators

- Case Study: The Fermilab linac
- Neural networks trained on data from DTL Tanks 2, 3, and 4 for 1k epochs
    - Model from tank 2 is trained on data from tanks 3 and 4 for 1k epochs
    - Transfer learning trains faster and reaches a better overall solution

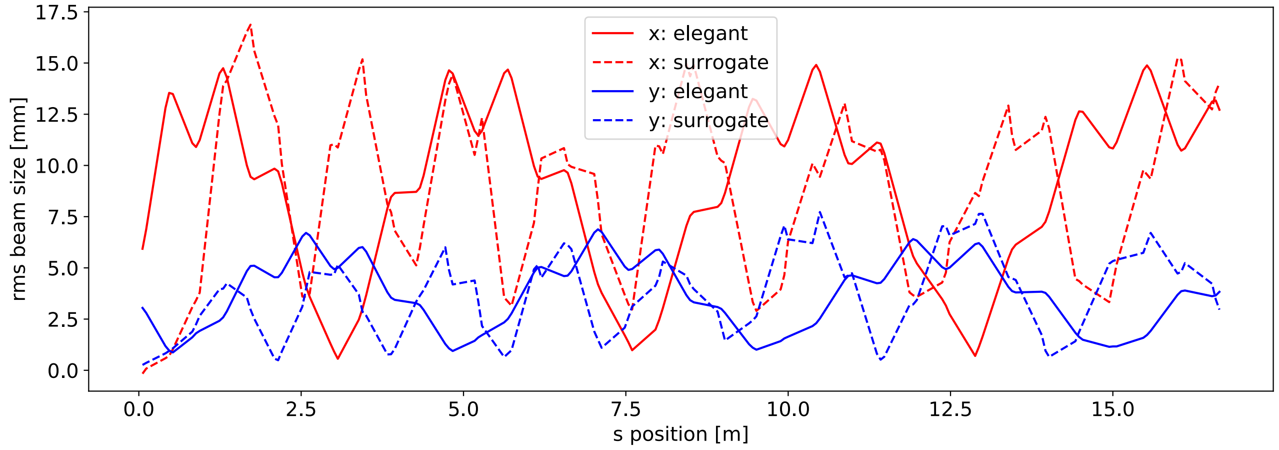

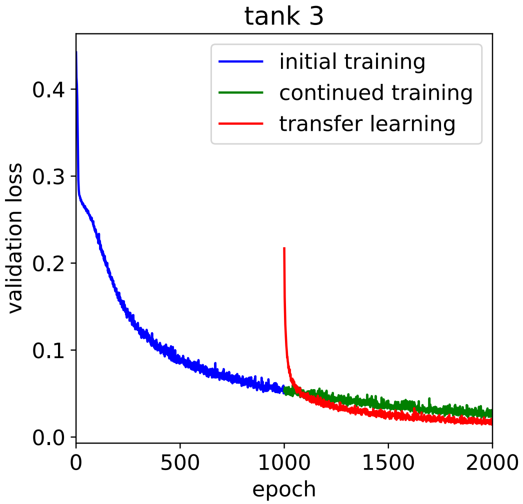

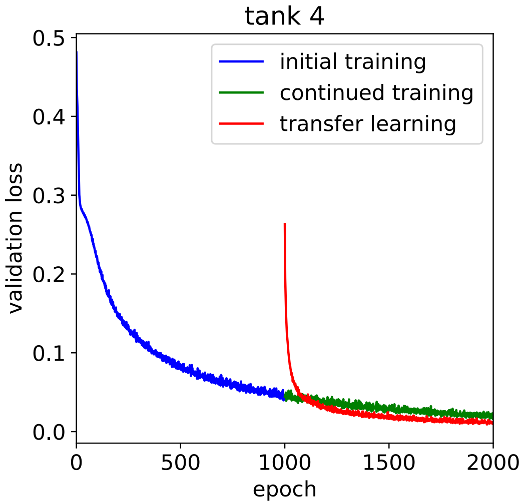

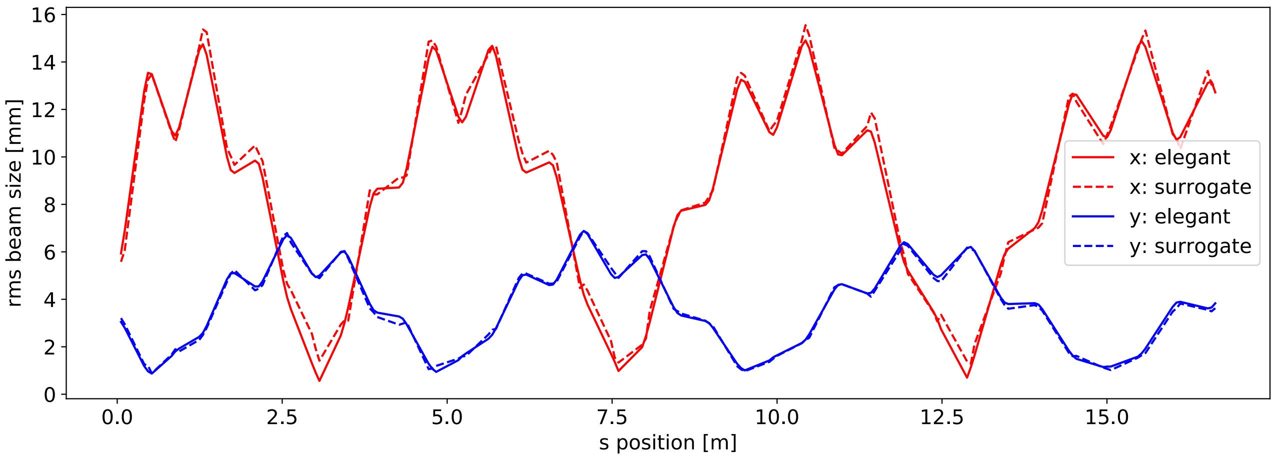

Machine learning for accelerator optimization

Test Case: Argonne Wakefield Accelerator Injector

OPAL simulation (PIC) :
3D space charge
3D field maps

NSGA-II for optimization:
200 generations
~350 individuals

100 - 5K random
points for training

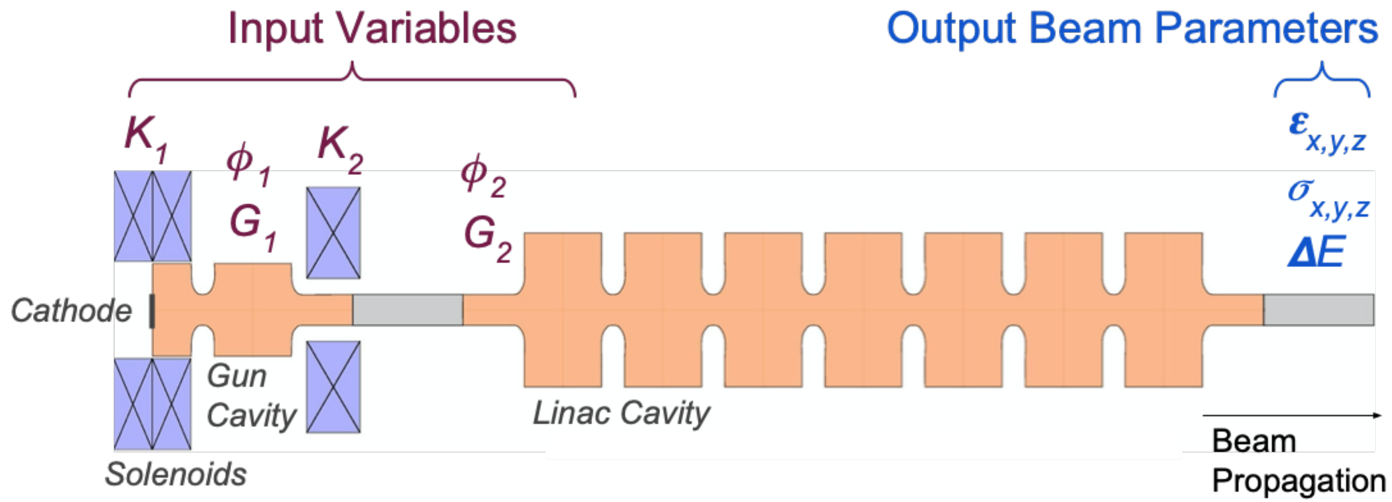

Generate ML Model using Sparse Random Sample

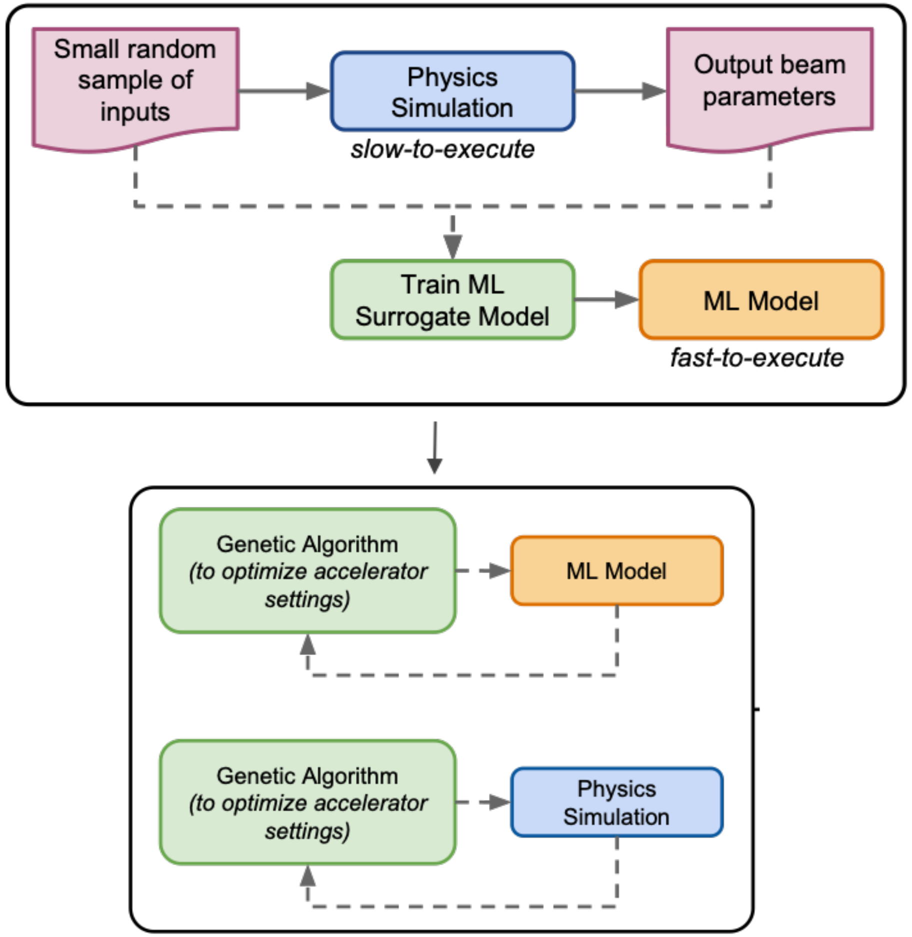

A. Edelen, et al., PRAB 23, 044601 (2020)

Machine learning for accelerator optimization

Test Case: Argonne Wakefield Accelerator Injector

OPAL simulation (PIC) : 3D space charge 3D field maps	NSGA-II for optimization: 200 generations ~350 individuals	100 - 5K random points for training

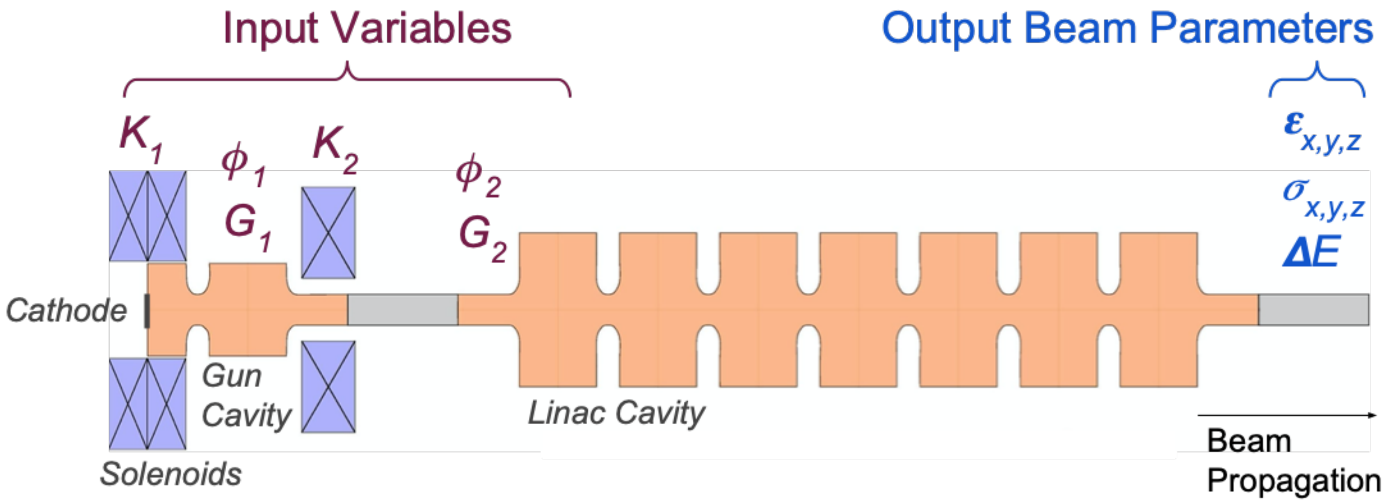

A. Edelen, et al., PRAB 23, 044601 (2020)

Generate ML Model using Sparse Random Sample

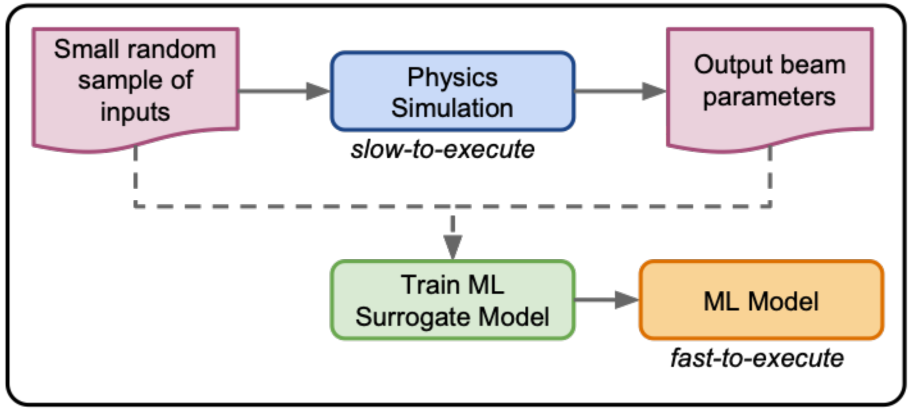

Compare Resulting Pareto Fronts

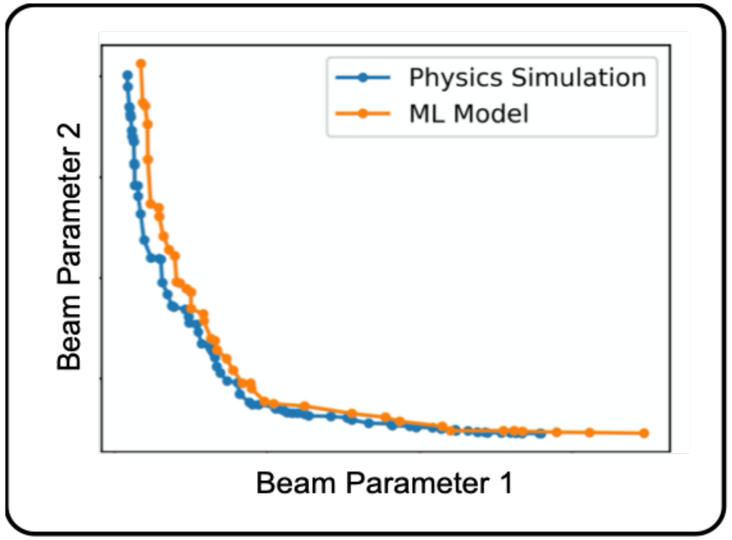

Machine learning for accelerator optimization

In some cases, optimization over simulation takes too long to converge

→ *validate Pareto front from neural network more directly*

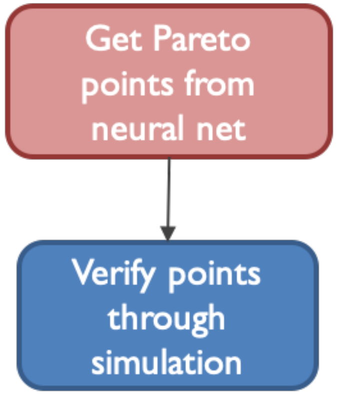

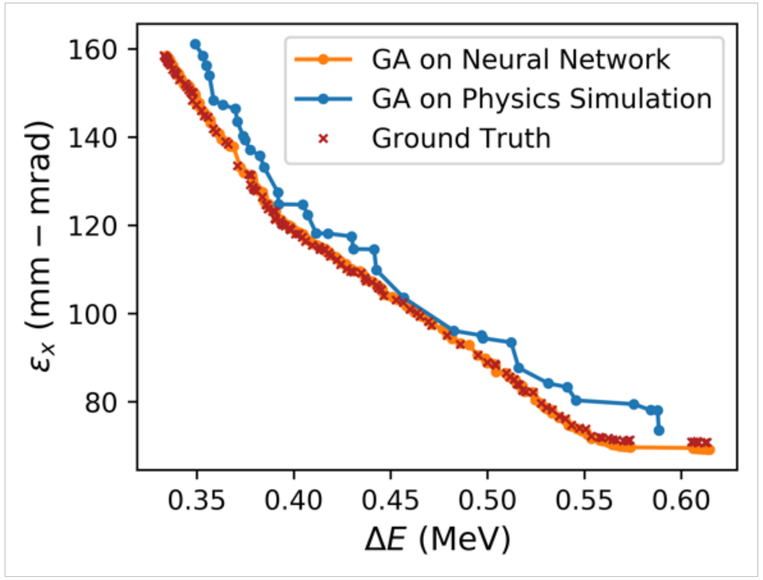

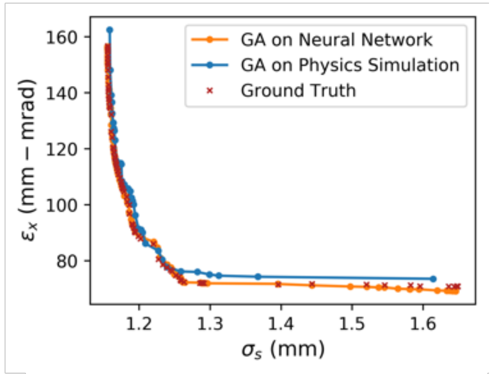

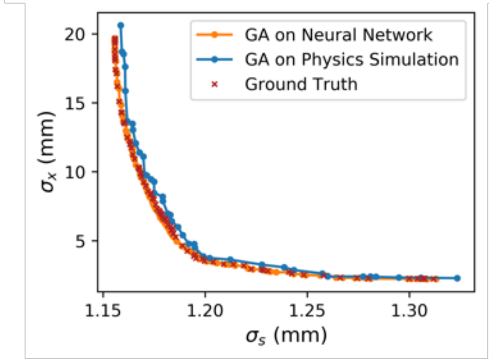

Required **130x fewer simulations** and had **10^6 times faster execution** in the optimization

A. Edelen, et al., PRAB 23, 044601 (2020)

Machine learning applications for accelerators

Accelerator R+D

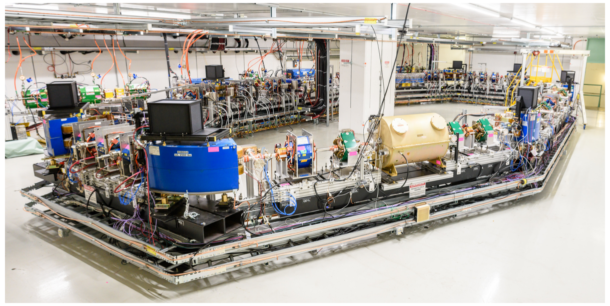

Beam for Experimentalists

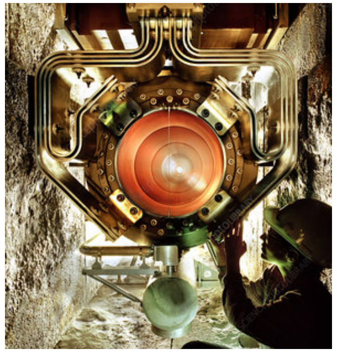

Down Time

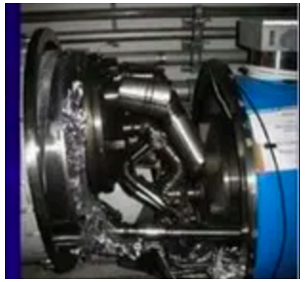

Tuning

- Inverse Models
- Neural Network + Optimization

Surrogate Modeling

- Neural Network
- Neural Network + Optimization

Anomaly Detection

- Inverse Models
- Autoencoders

Inverse models for diagnostics

- Inverse models as a diagnostic in a supervised fashion
  - Direct comparison between predicted settings and actual settings informs operations of a potential anomaly with that magnet

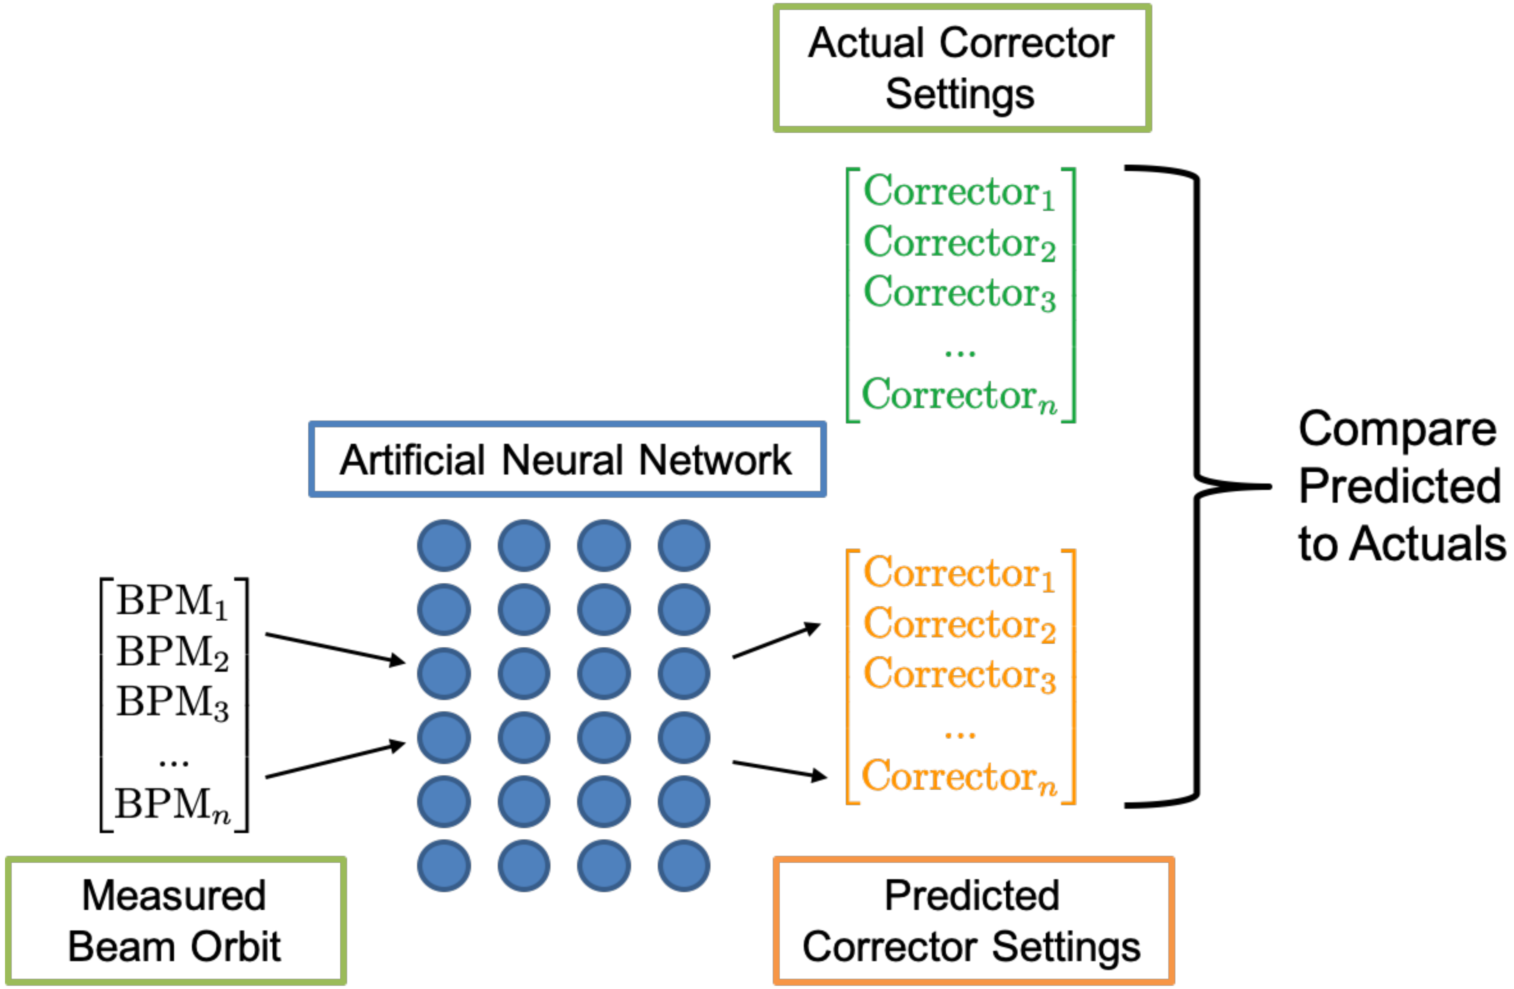

Inverse models for diagnostics

- Inverse models as a diagnostic in a supervised fashion
  - Direct comparison between predicted settings and actual settings informs operations of a potential anomaly with that magnet
- Inverse models as a diagnostic in an unsupervised fashion
  - Assumptions
    - model errors are caused by other beamline elements
    - each beam-line element will have a unique error signature
  - Use this for tuning

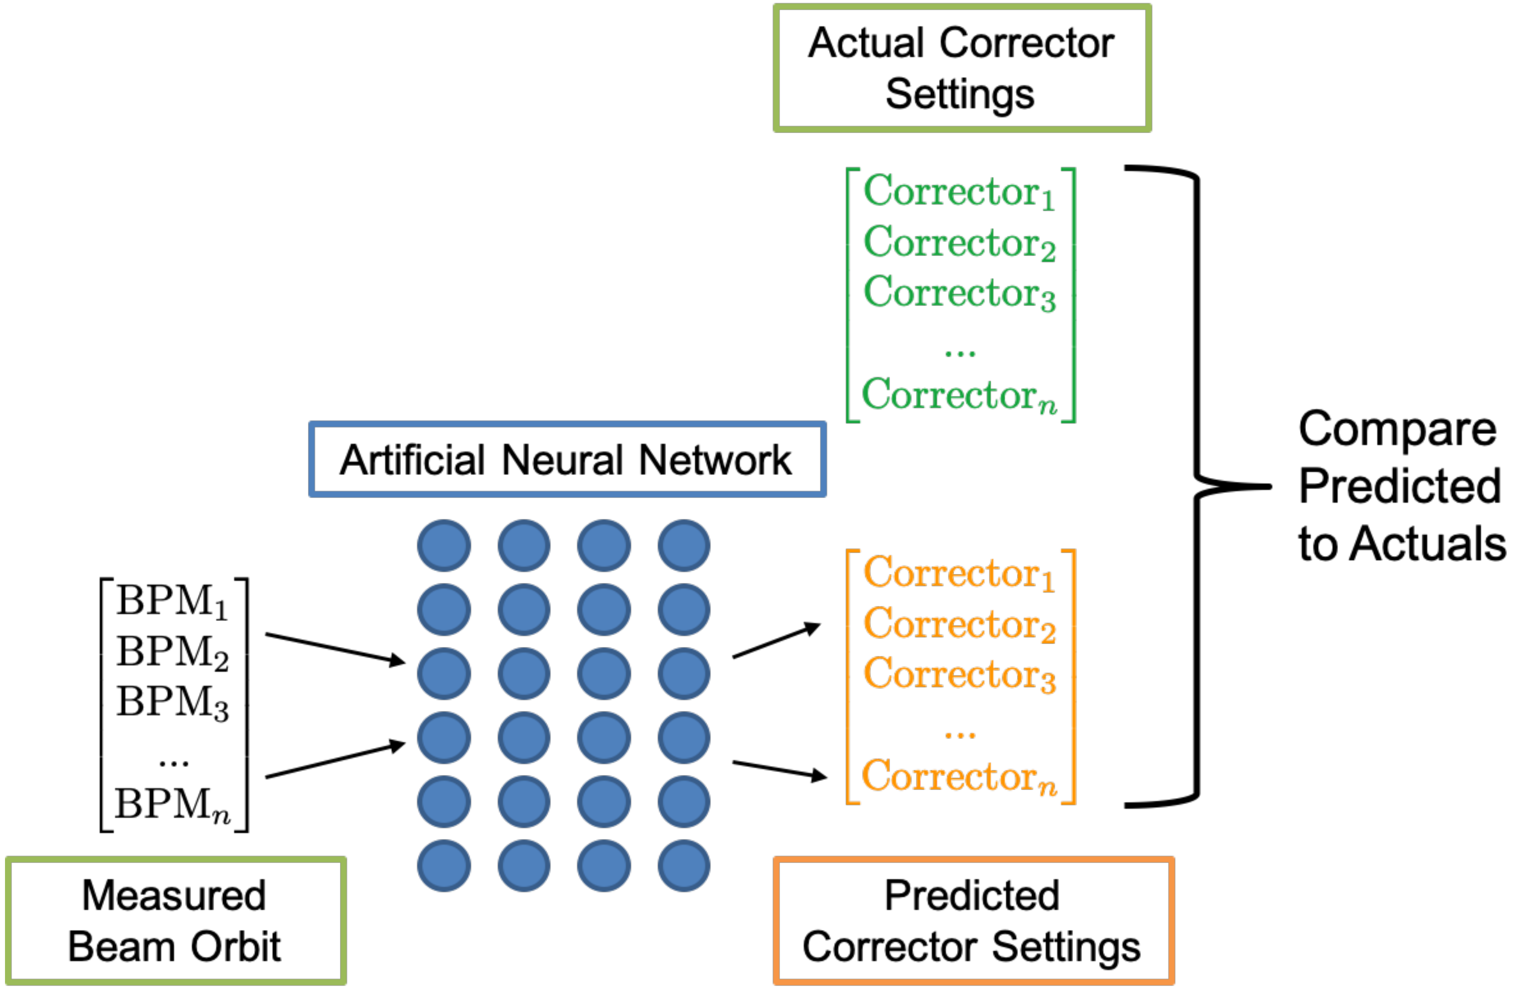

AGS to RHIC transfer line

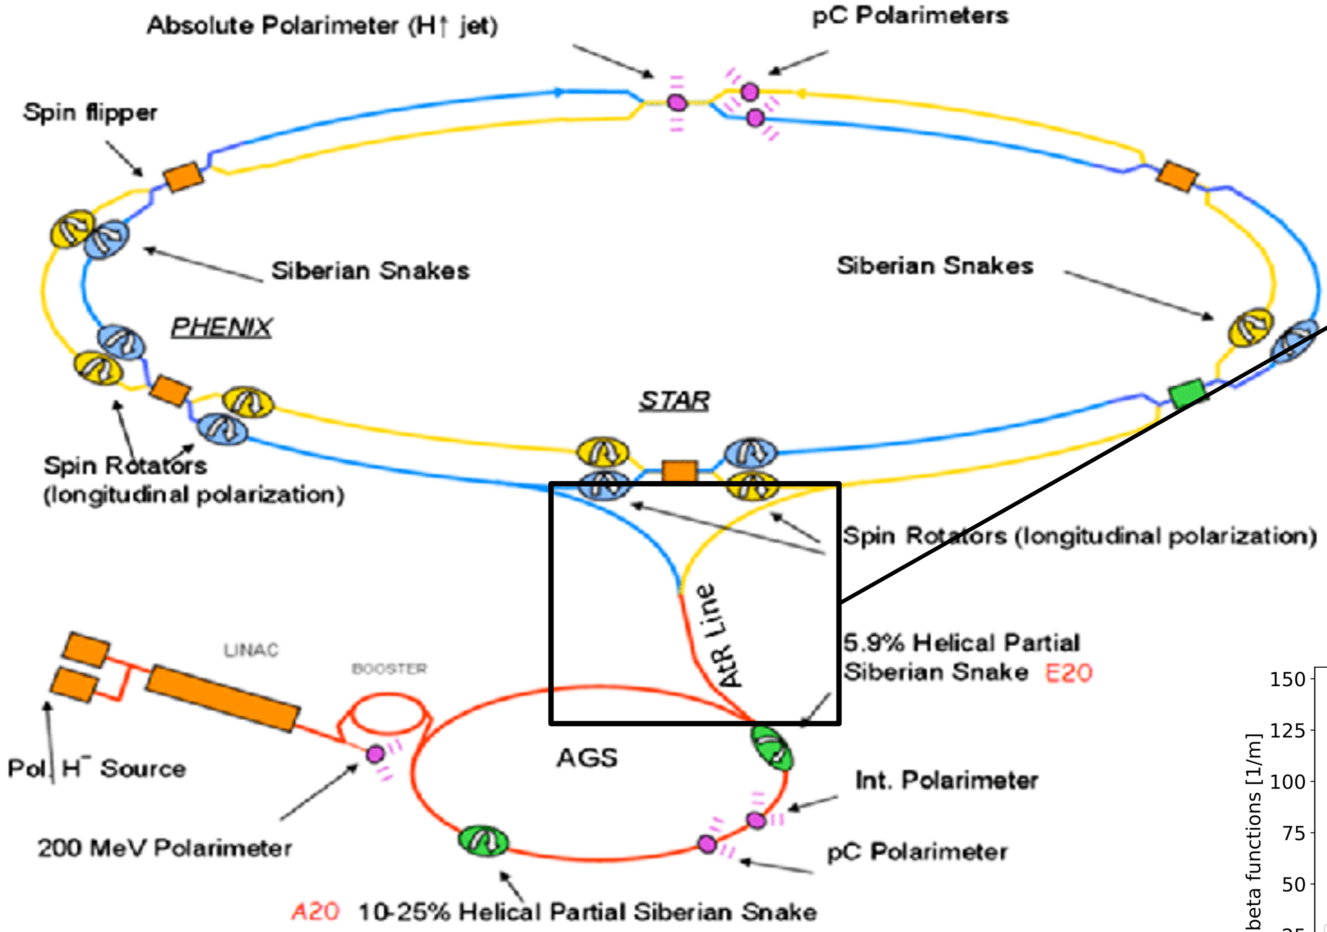

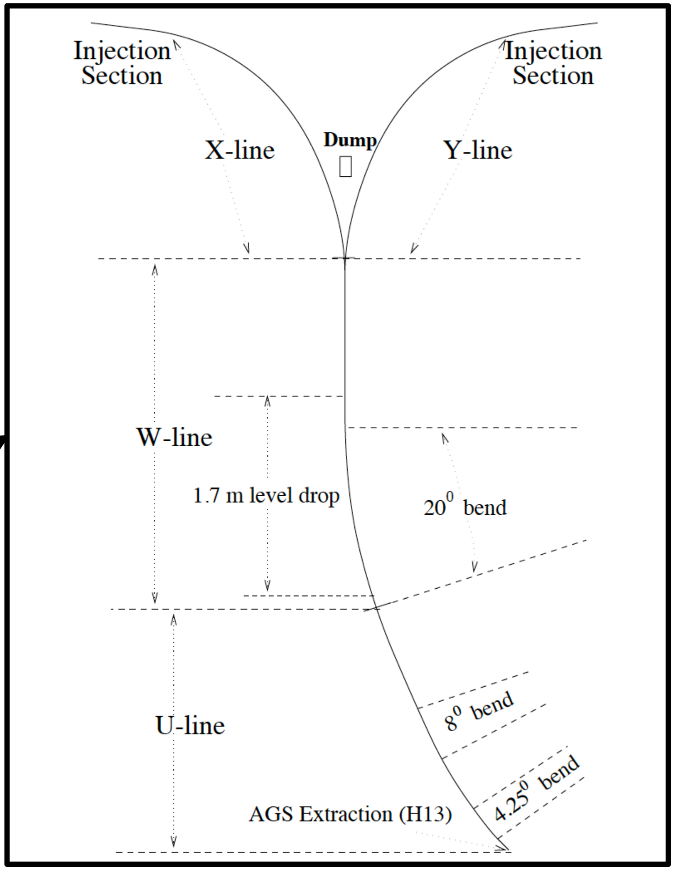

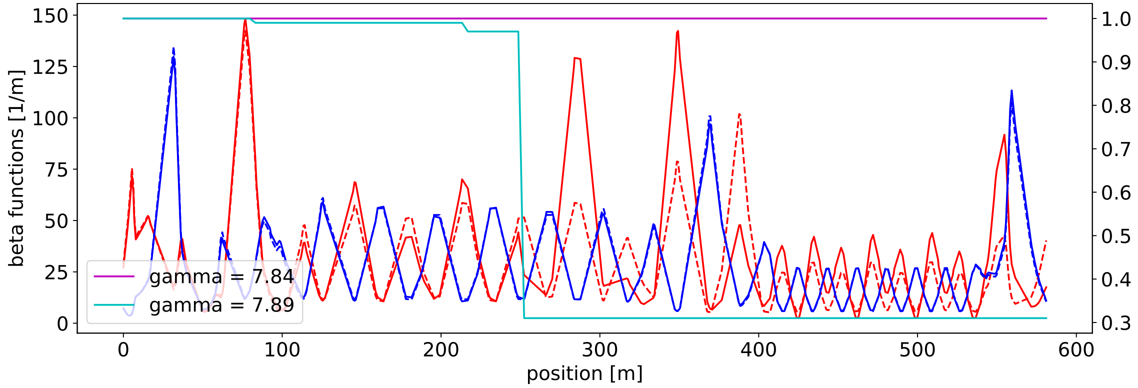

AGS to RHIC transfer line study

Training Data Generation

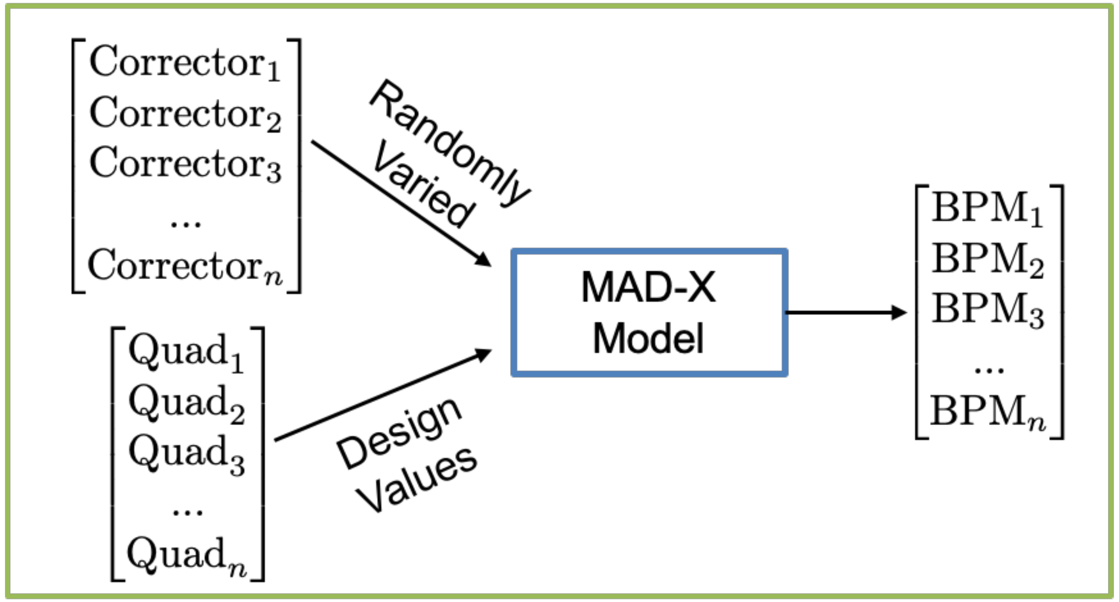

Test Data Generation

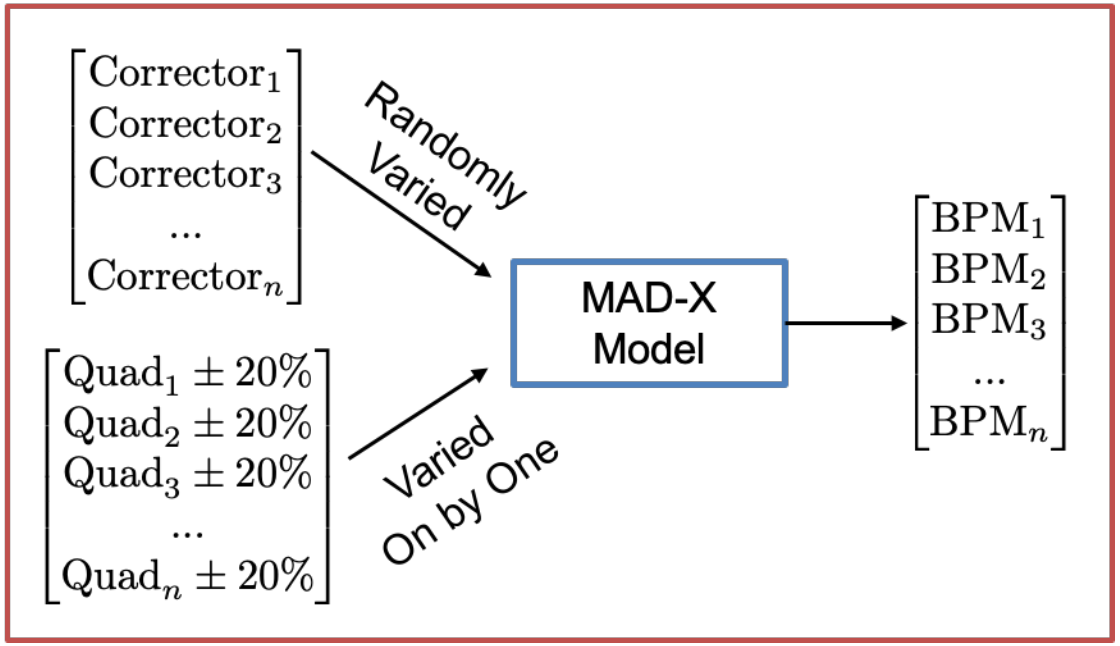

AGS to RHIC transfer line study

- Two configurations were used: one where the initial positions were also varied randomly and one where the initial positions were not varied.
- Right: Predicted corrector settings vs the ground truth for the validation set
  - Black: without quadrupole errors
  - Red: a single quadrupole error and random initial position errors
  - Blue: a single quadrupole error without initial position errors

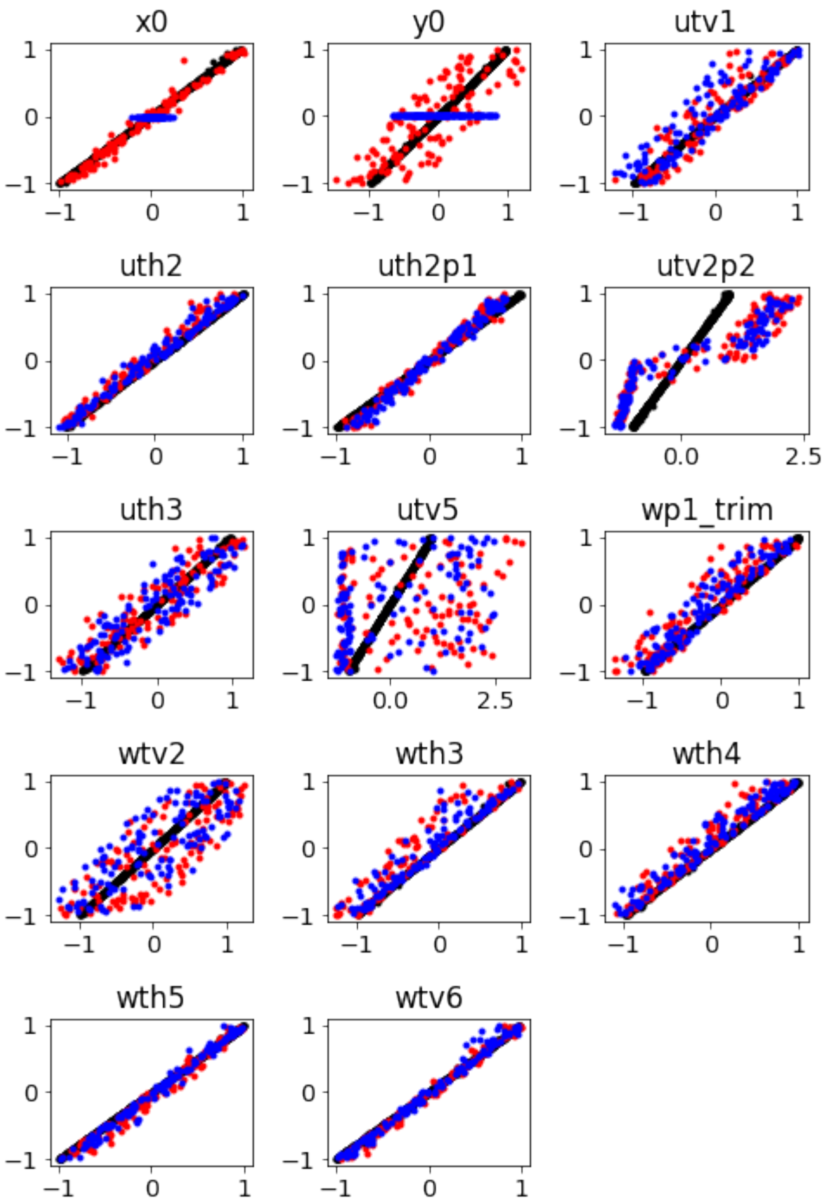

AGS to RHIC transfer line study

- Sensitivity of each corrector prediction to a particular quadrupole
  - Unique signatures for each quadrupole
  - The model clearly identifies errors in these magnets without any explicit knowledge of their existence

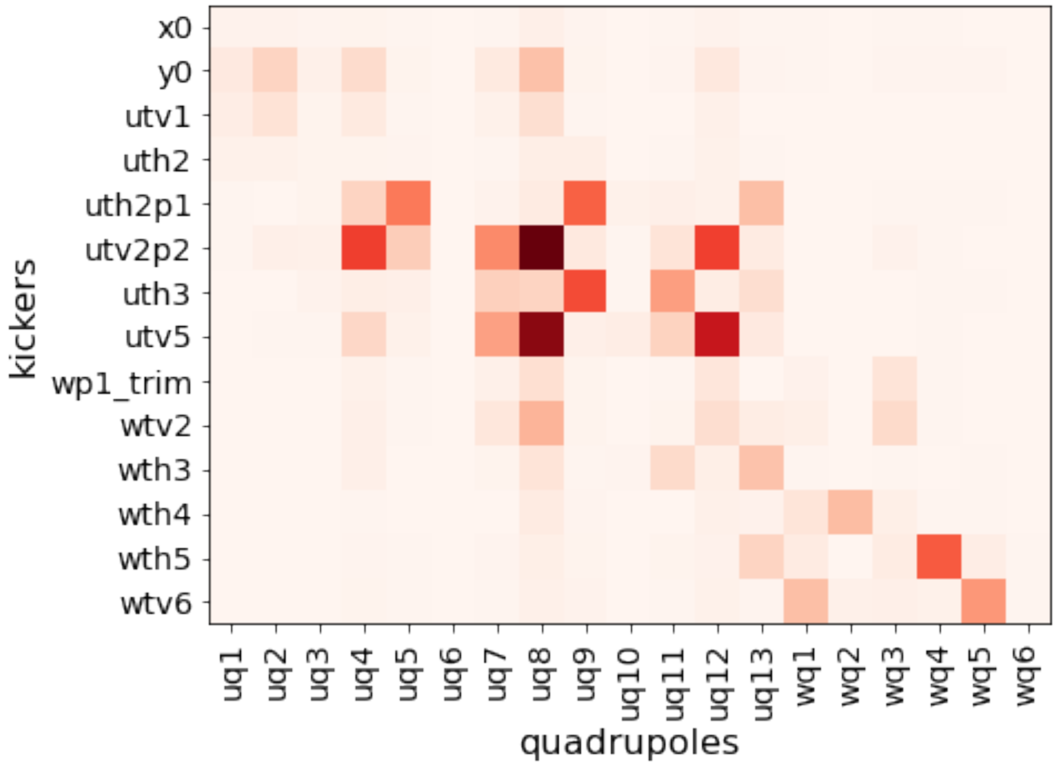

- Future work
  - Use signatures to predict unknown quadrupole errors
  - Use model errors to tune out quadrupole errors

Detecting faulty magnet power supplies in the APS

- Can we predict if a fault will occur?
  - If yes, can we predict which magnet will fault

- Components of interest
  - 1320 magnet power supplies / 40 sectors (each has A (green) and B (blue) sections)

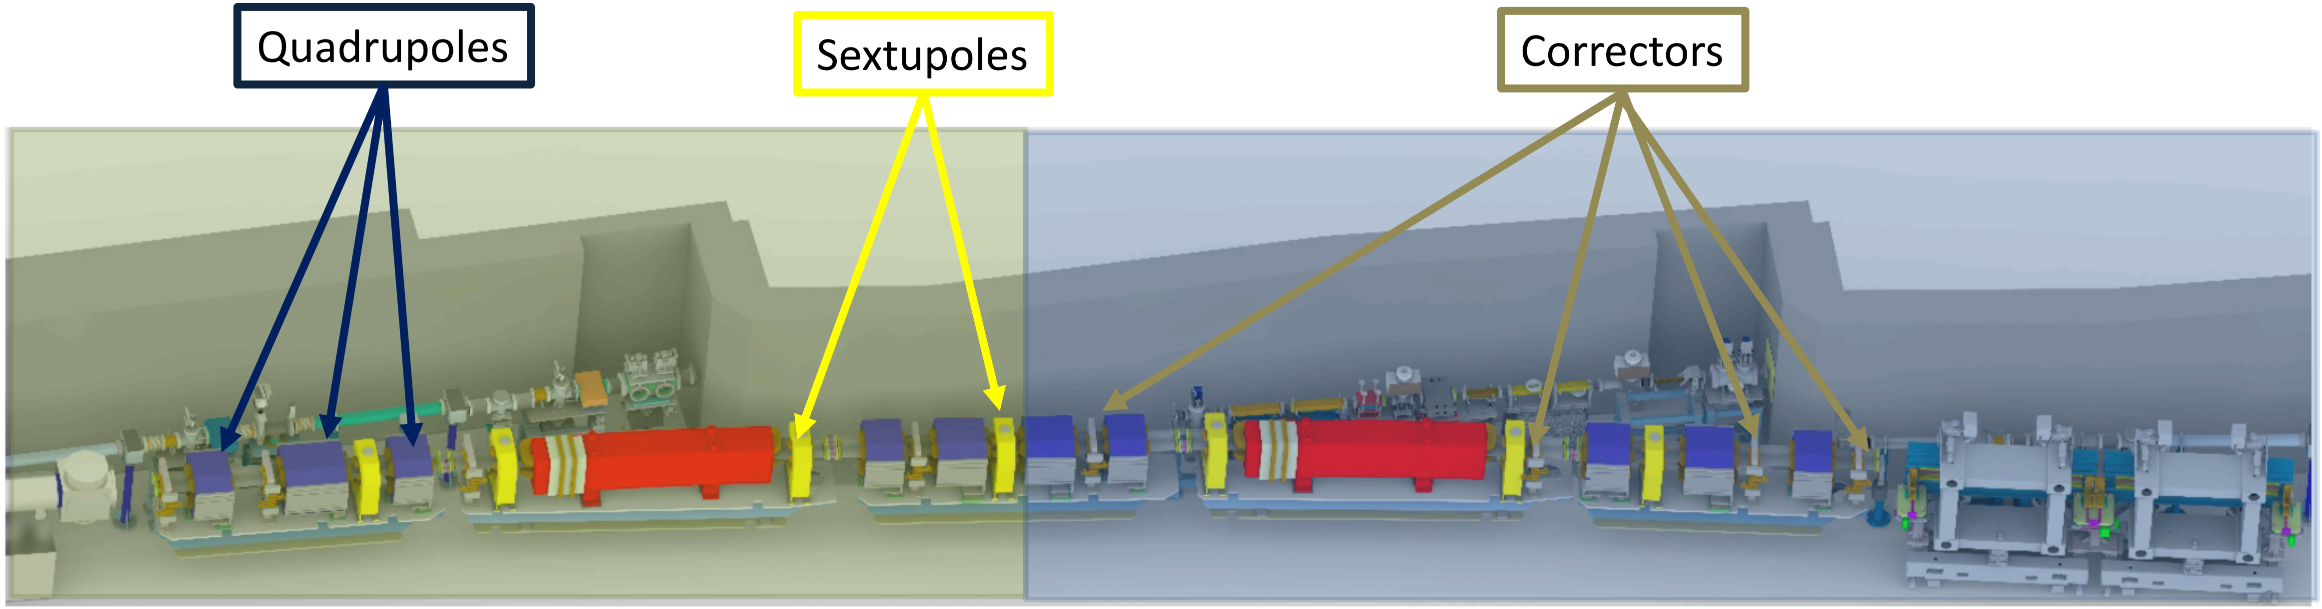

https://www.energy.gov/sites/prod/files/2019/04/f62/Advanced-Photon-Source-Upgrade-Project.pdf

Reconstruction tests

- Reconstruct unknown data using an autoencoder
  - Train and validate the autoencoder on known good datasets
  - Test on unknown data (may be good or bad)
  - Measure the degree to which the autoencoder successfully reconstructs the unknown data

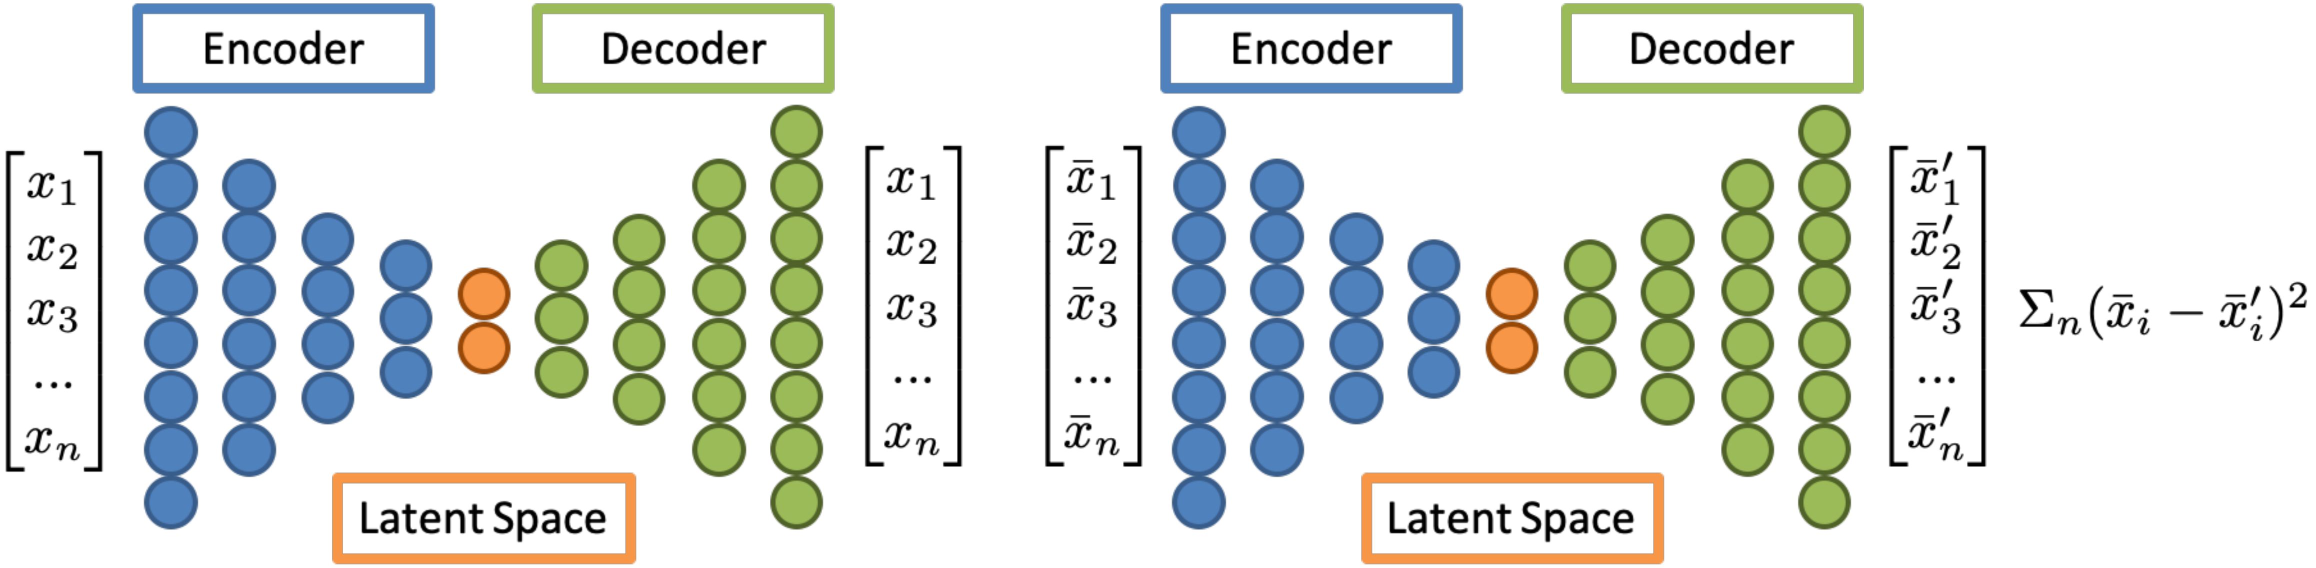

Data collected by the APS over three years of operation

- Time series data for 1320 magnets
  - Power supply cap temperature
  - Current
  - Magnet temperature
- Data is aggregated by sector
  - Reference data (left) used for training and validation
  - Test data (middle) with known anomalies
  - Histogram difference (right)

Reference data (blue): no fault occurs in vicinity, normal operations. Test data (orange): magnet failure occurs; data is clipped and does not include final minutes prior to magnet fault.

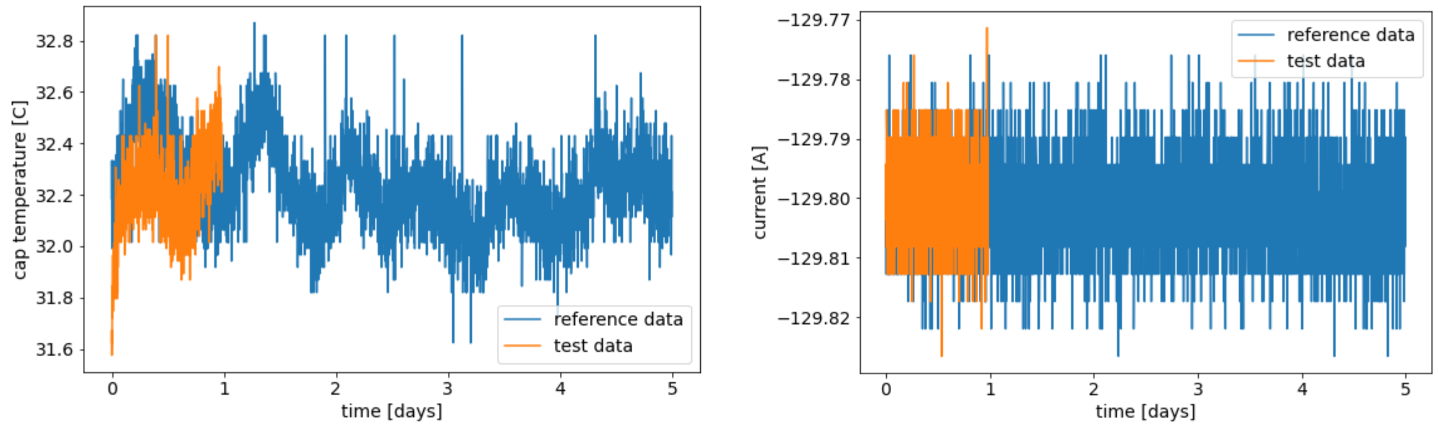

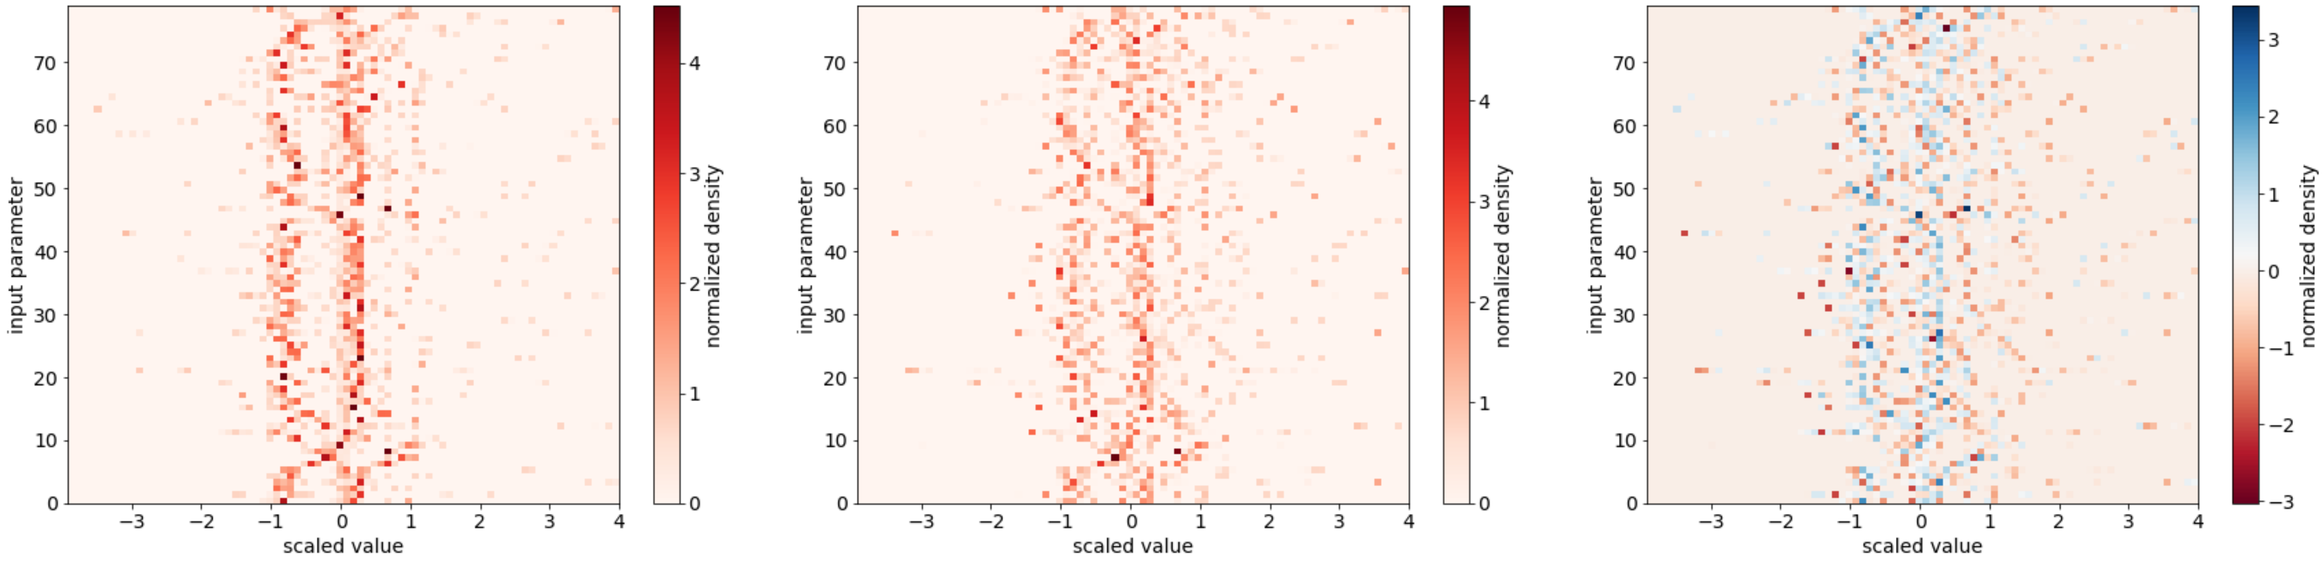

Machine learning for anomaly detection

- Reconstruct unknown data using an autoencoder
  - Train and validate the autoencoder on known good datasets
  - Test on unknown data (may be good or bad)
  - Measure the degree to which the autoencoder successfully reconstructs the unknown data

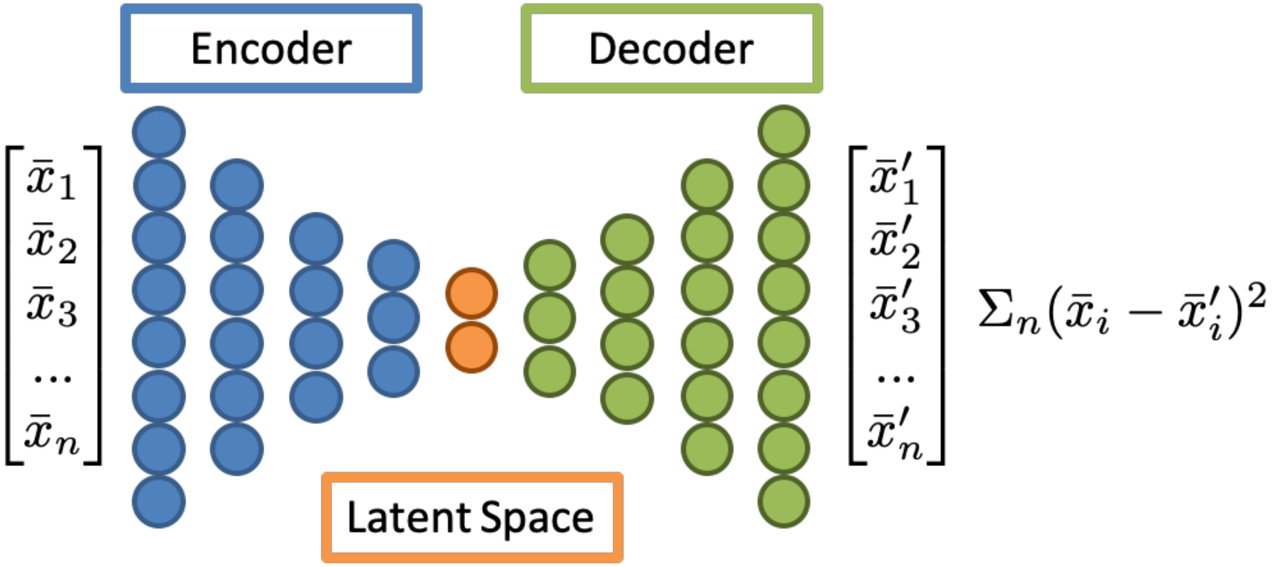

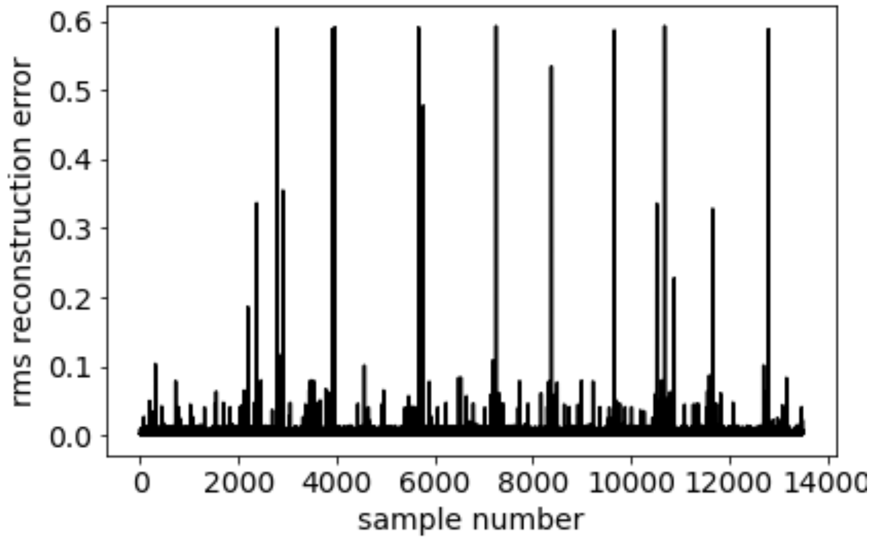

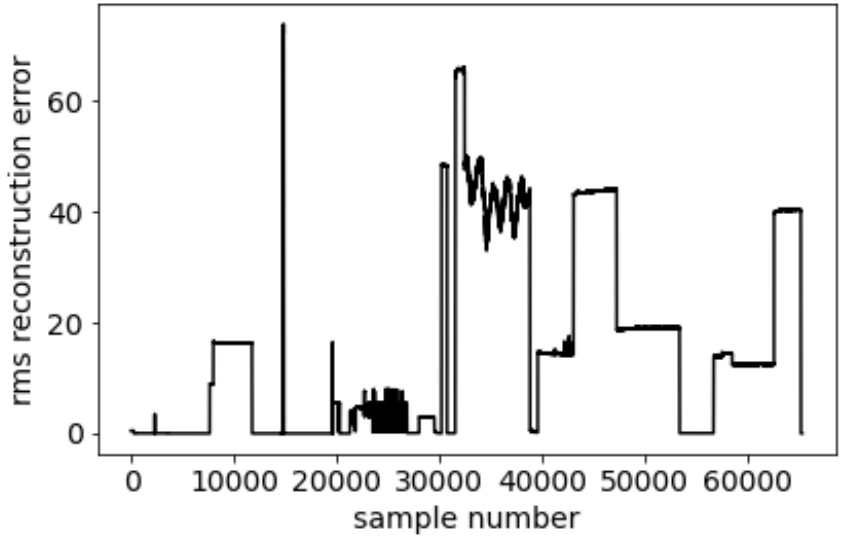

Reconstruction of reference data and test data (by sector)

Region of convergence plot for the RMS error and squared error evaluation metrics. The main plot shows the true positive rate vs the false positive rate as a function of anomaly threshold. Inset a) shows the true positive rate as a function of the error threshold and inset b) shows the false positive rate as function of the error threshold. Note that the threshold is normalized to the peak value of the reconstruction error computed on the reference data.

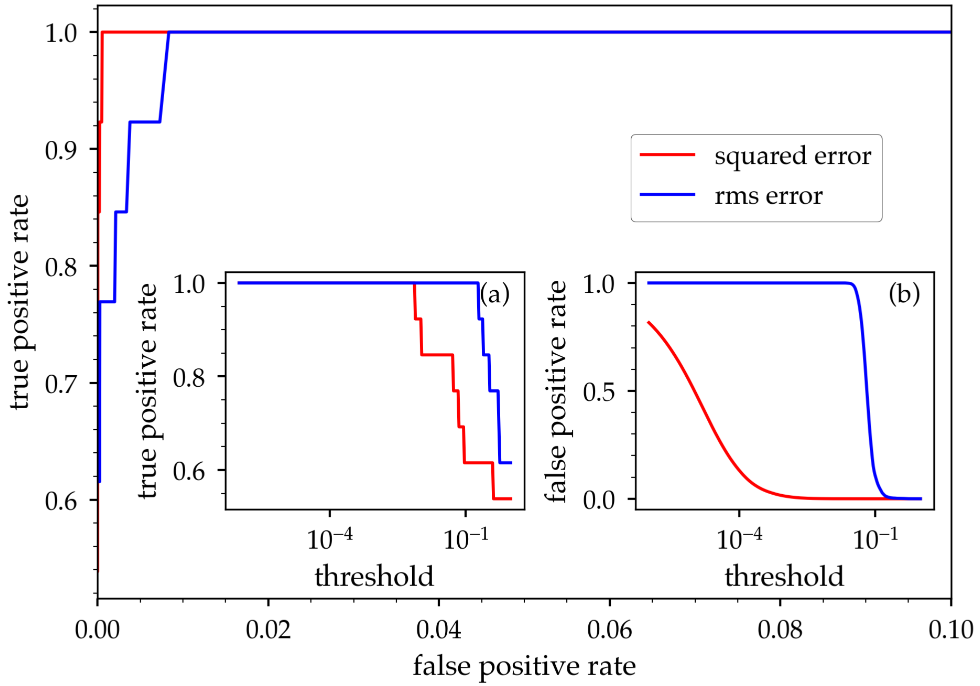

Number of faulty sectors for a given fault run. The data are sorted by the number of faulty sectors identified in the semisupervised case.

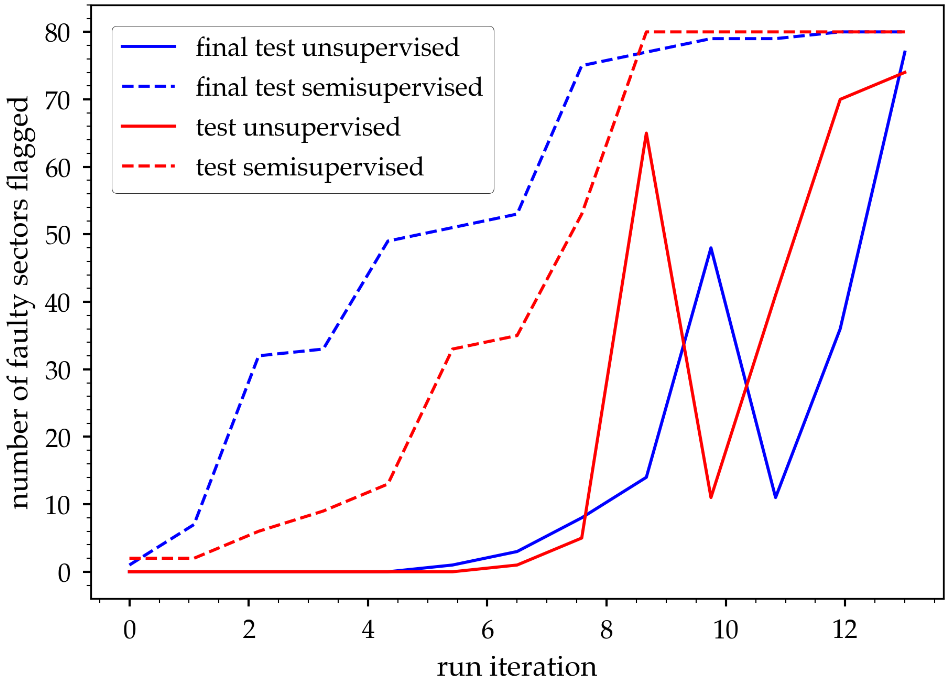

Forecasting faults using unsupervised and semi-supervised learning

First indication of an anomaly as a function of the run time for the fault data using the RMS error metric. Red is the data used to tune the detection threshold while blue is the final test data that is not used in any of the training or parameter tuning. The dashed lines represent the unsupervised case while the solid line is the semisupervised case.

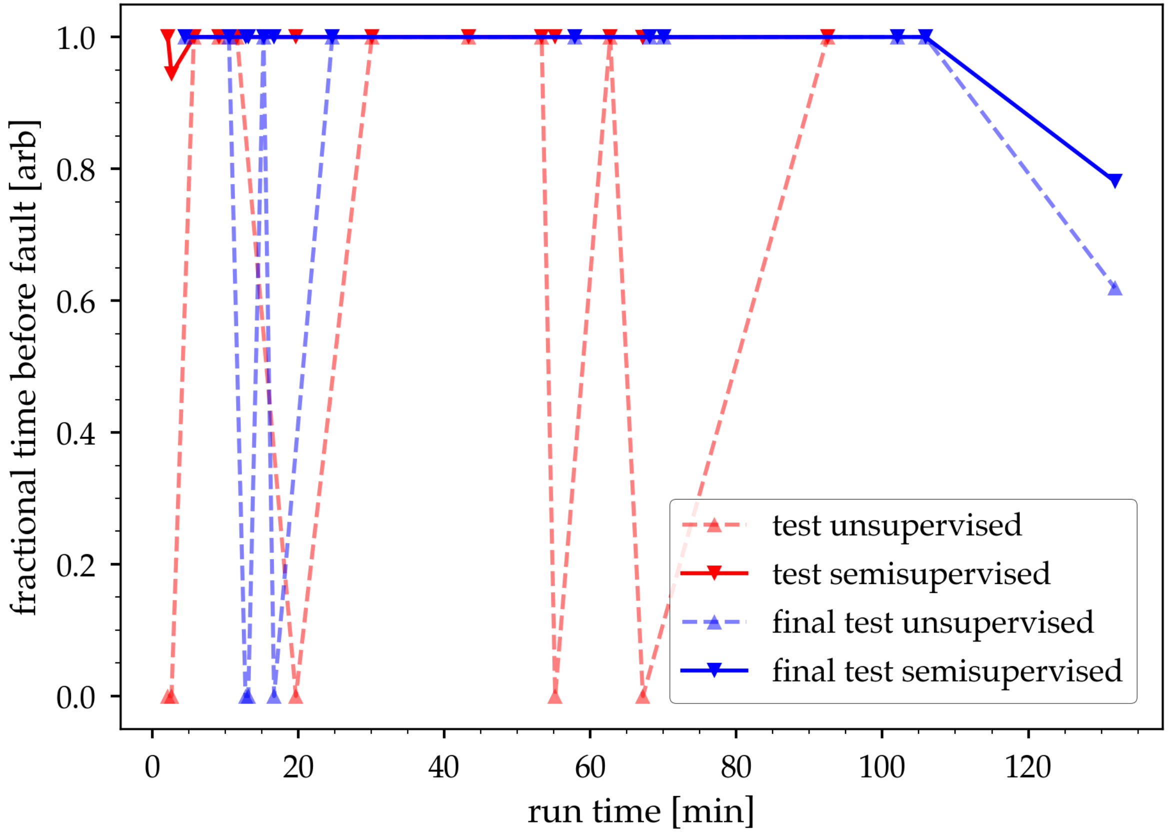

First indication of an anomaly as a function of the run time for the fault data using the squared error metric. Red is the data used to tune the detection threshold while blue is the final test data that is not used in any of the training or parameter tuning. The dashed lines represent the unsupervised case while the solid line is the semisupervised case

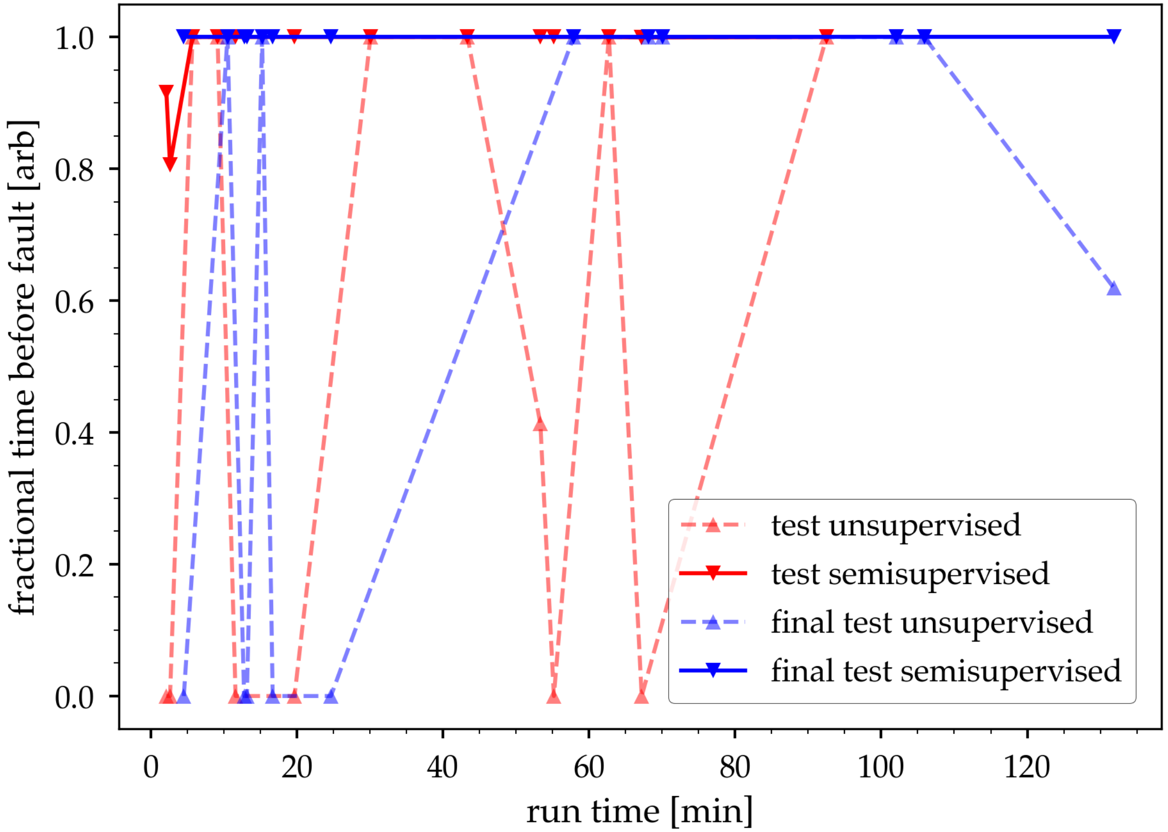

Conclusions

- Accelerator facilities rely heavily on human operators for tuning/control

- Modeling and control of these machines is challenging
  - Nonlinear systems with large parameter spaces
  - Variety of diagnostics (e.g. beam images), but these are limited and number, and some are not continuously available for use
  - Time-varying/ non-stationary behavior

- Strong incentives for improving control (and understanding system)
  - High user demand → want to switch between custom user requests quickly
  - High cost for unintended down-time → personnel time, user time, scientific output
  - Achieve challenging beam setups for new science goals

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.