Fermilab

July 15, 2026

Bayesian Exploration and Surrogate Emulation of Nonlinear Beam-Response Geometry in the LBNF Beamline

Sudeshna Ganguly, Katsuya Yonehara, Pavel Snopok
Dedicated AI/ML Accelerator Test Facility

U.S. DEPARTMENT of ENERGY
Fermi National Accelerator Laboratory is managed by FermiForward for the U.S. Department of Energy Office of Science

FERMILAB-SLIDES-26-0125-AD

Precision Beam Control for Future High-Power Neutrino Beams

- Precise measurements of CP violation and neutrino oscillations
- Physics sensitivity depends on reducing beam-related systematic uncertainties

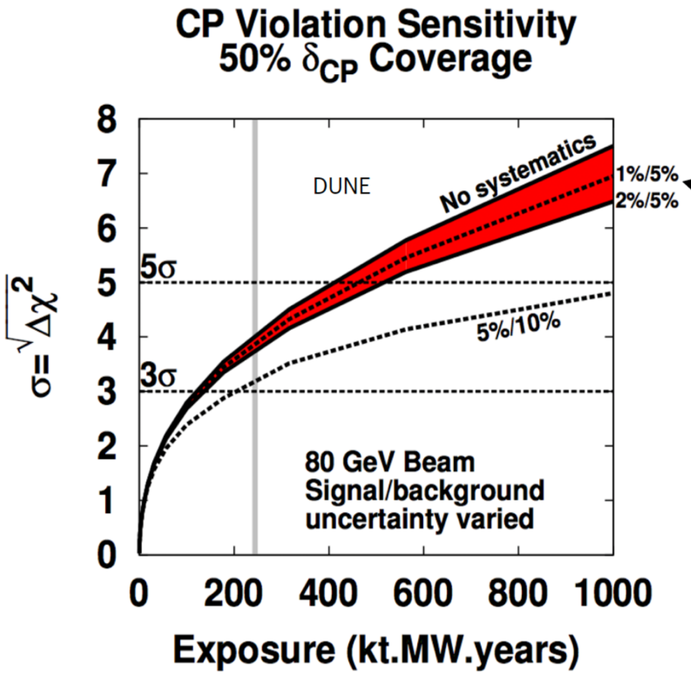

This is **DUNE's goal**

Taken from Laura Fields, University of Notre Dame, Neutrino Flux Uncertainties USPAS Neutrino Beams Course talk , 2024

Beam-related uncertainties remain a significant limitation on DUNE sensitivity

Total focusing uncertainty at 3% level with uncertainties piling up at the falling edge of focusing peak up to 6-7%

POT counting uncertainty at 2% level

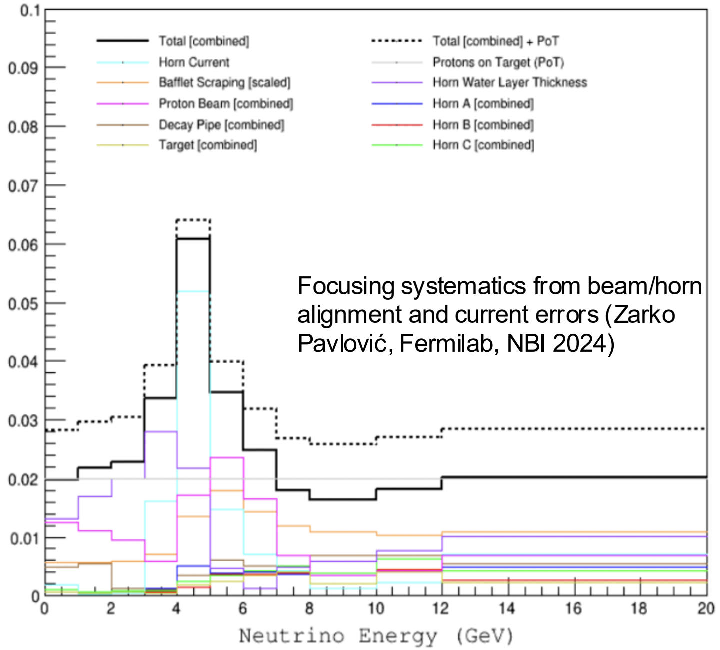

Focusing systematics from beam/horn alignment and current errors (Zarko Pavlović, Fermilab, NBI 2024)

- Beam focusing uncertainty ≈ 3%
- Larger (6–7%) near the focusing edge
- Beam steering, horn current and target effects dominate

FERMI NATIONAL ACCELERATOR LABORATORY

Muon Monitors as a Beam Operations Tool

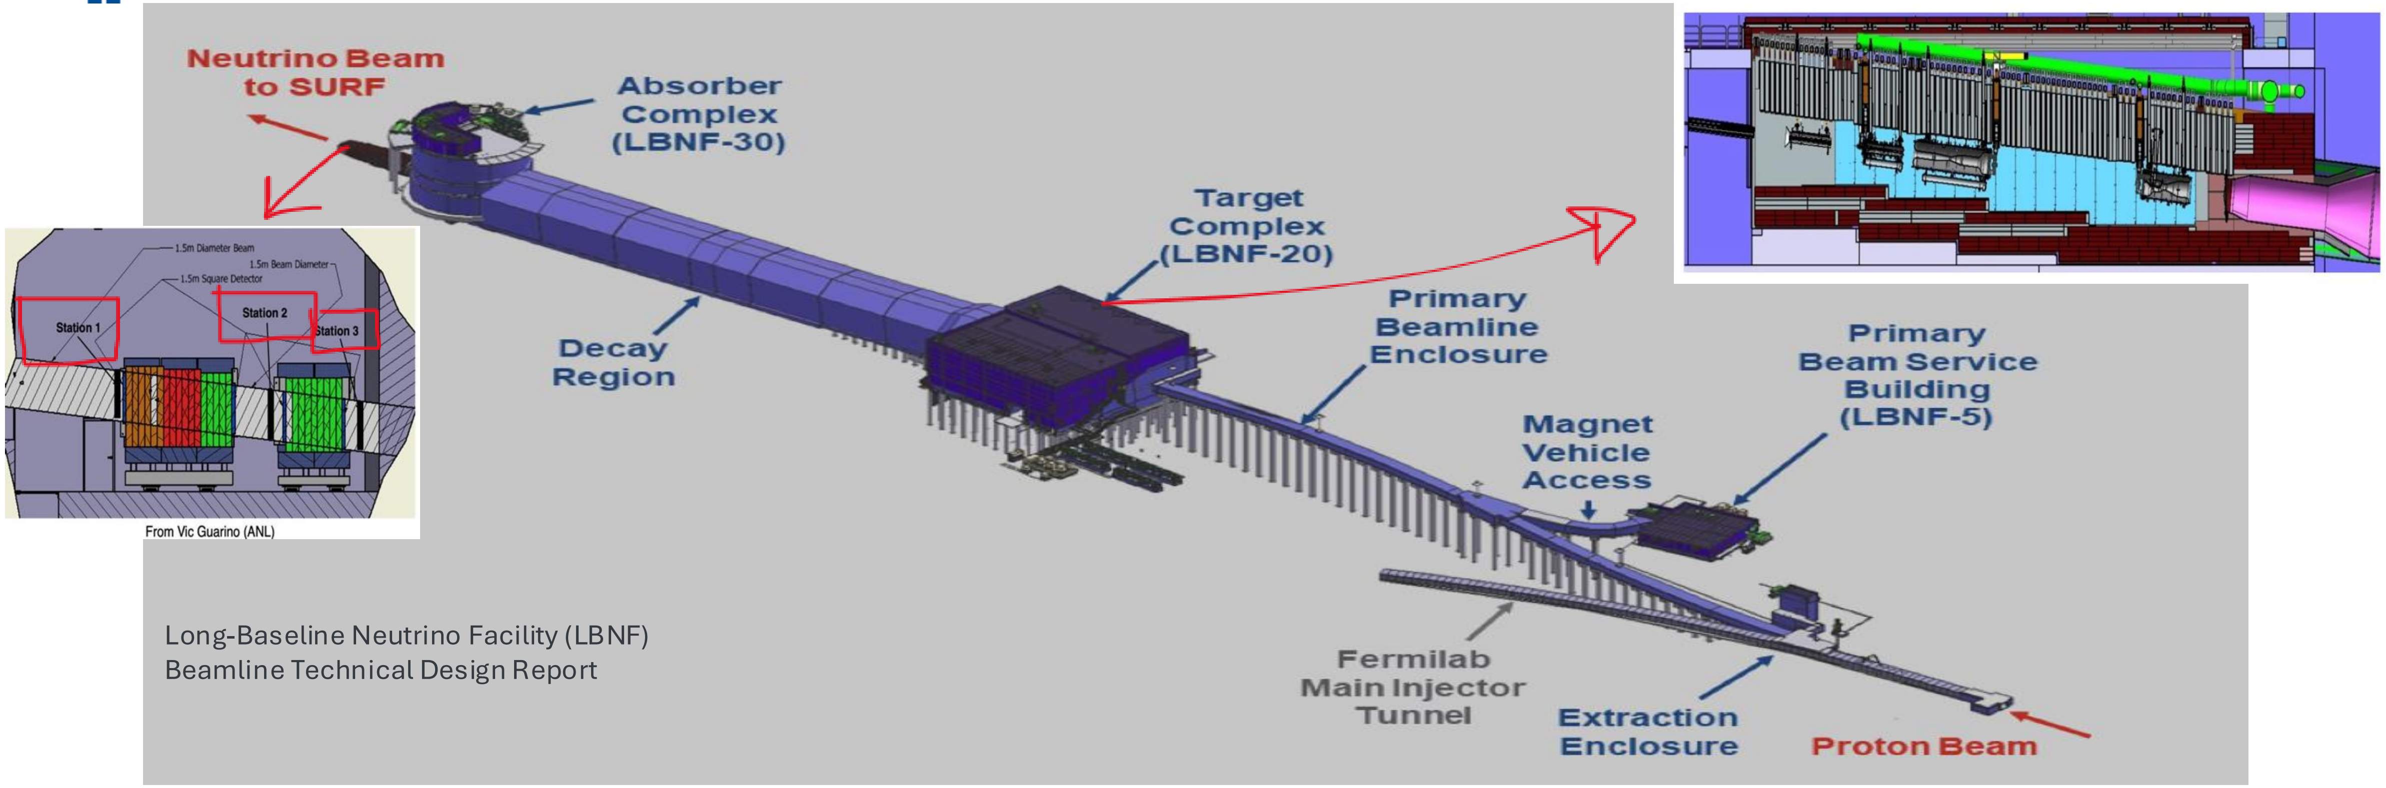

Proton beam

↓

Target + Horns

↓

Muon monitors

↓

Infer proton beam parameters

- **During operation**
- 120 GeV, 1.2 MW proton beam (upgradable to 2.4 MW)
- Muon monitors provide only continuous downstream beam diagnostic
- Beam steering, horn current and focusing change muon distribution
- Goal: infer beam parameters in real time from muon monitor images
- **Muon monitors observe consequences of beam changes—not beam parameters themselves**

Muon Monitors Encode Beam Focusing

- Muon distributions change predictably with proton beam position, beam size, and horn current
- Each station samples a different pion energy range
- Together they provide a fingerprint of beam focusing

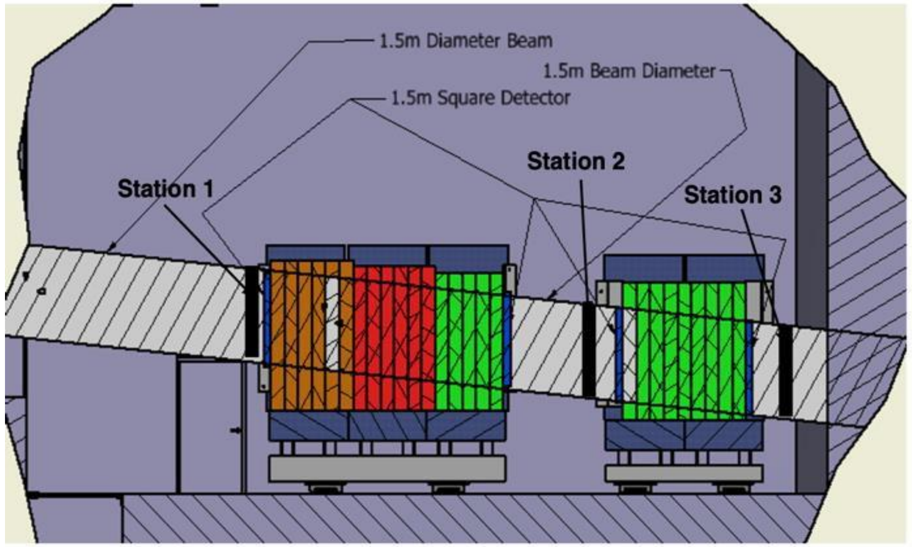

From Vic Guarino (ANL)

- Each station is effectively a different projection of the pion phase space
- Three muon monitor images jointly contain enough information to reconstruct beam parameters

Station 1 → low-energy pions

Station 2 → intermediate-energy pions

Station 3 → high-energy pions

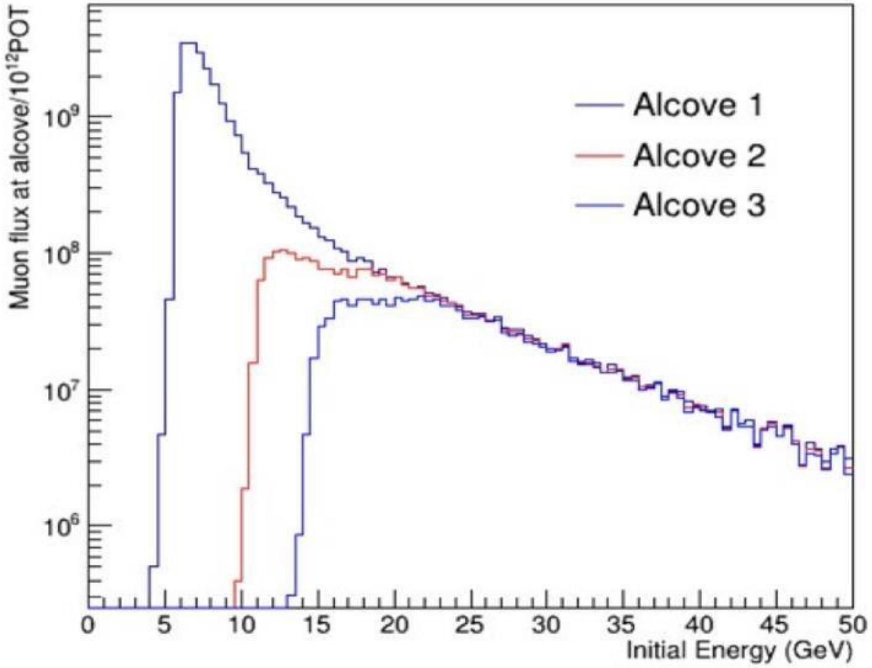

Muon Monitor Response to Proton Beam Position

- Proton beam scanned from −2 to +2 mm
- Muon charge distributions recorded at all three stations
- All stations respond systematically to beam motion
- Each station has different sensitivity because it samples a different pion energy range
- Small proton beam shifts produce systematic and measurable changes in muon distributions
- **Can these changes be inverted to recover original beam parameters?**

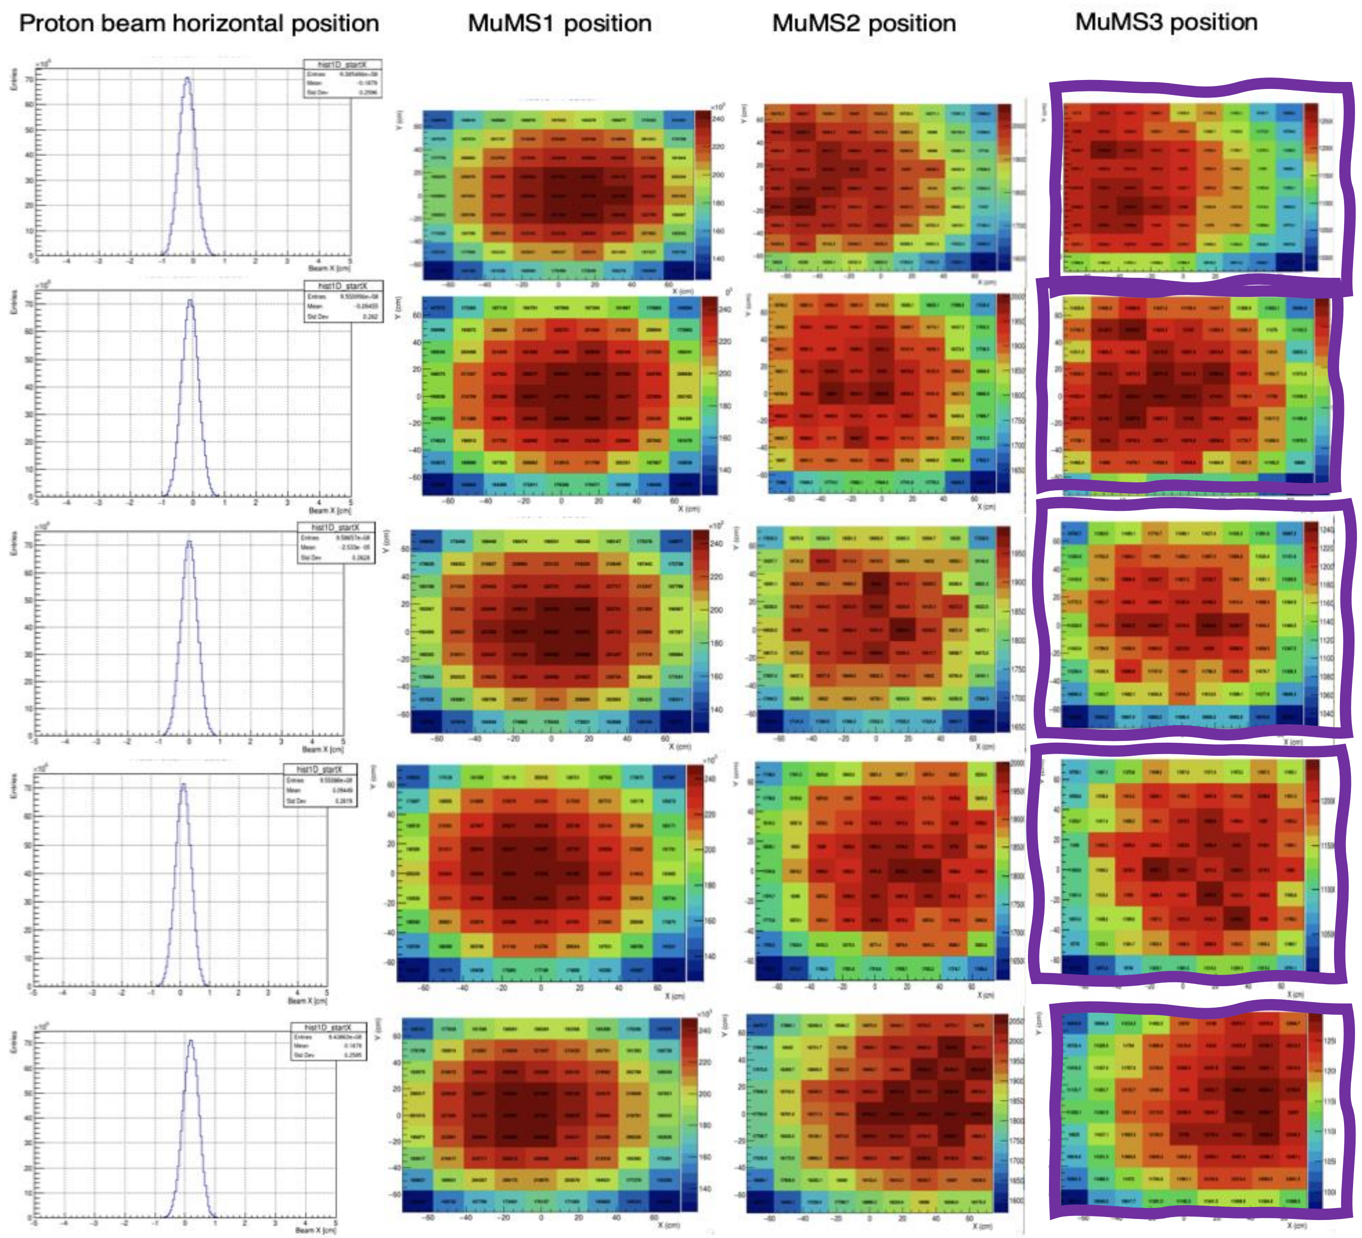

Muon Monitor Response is Linear and Predictable

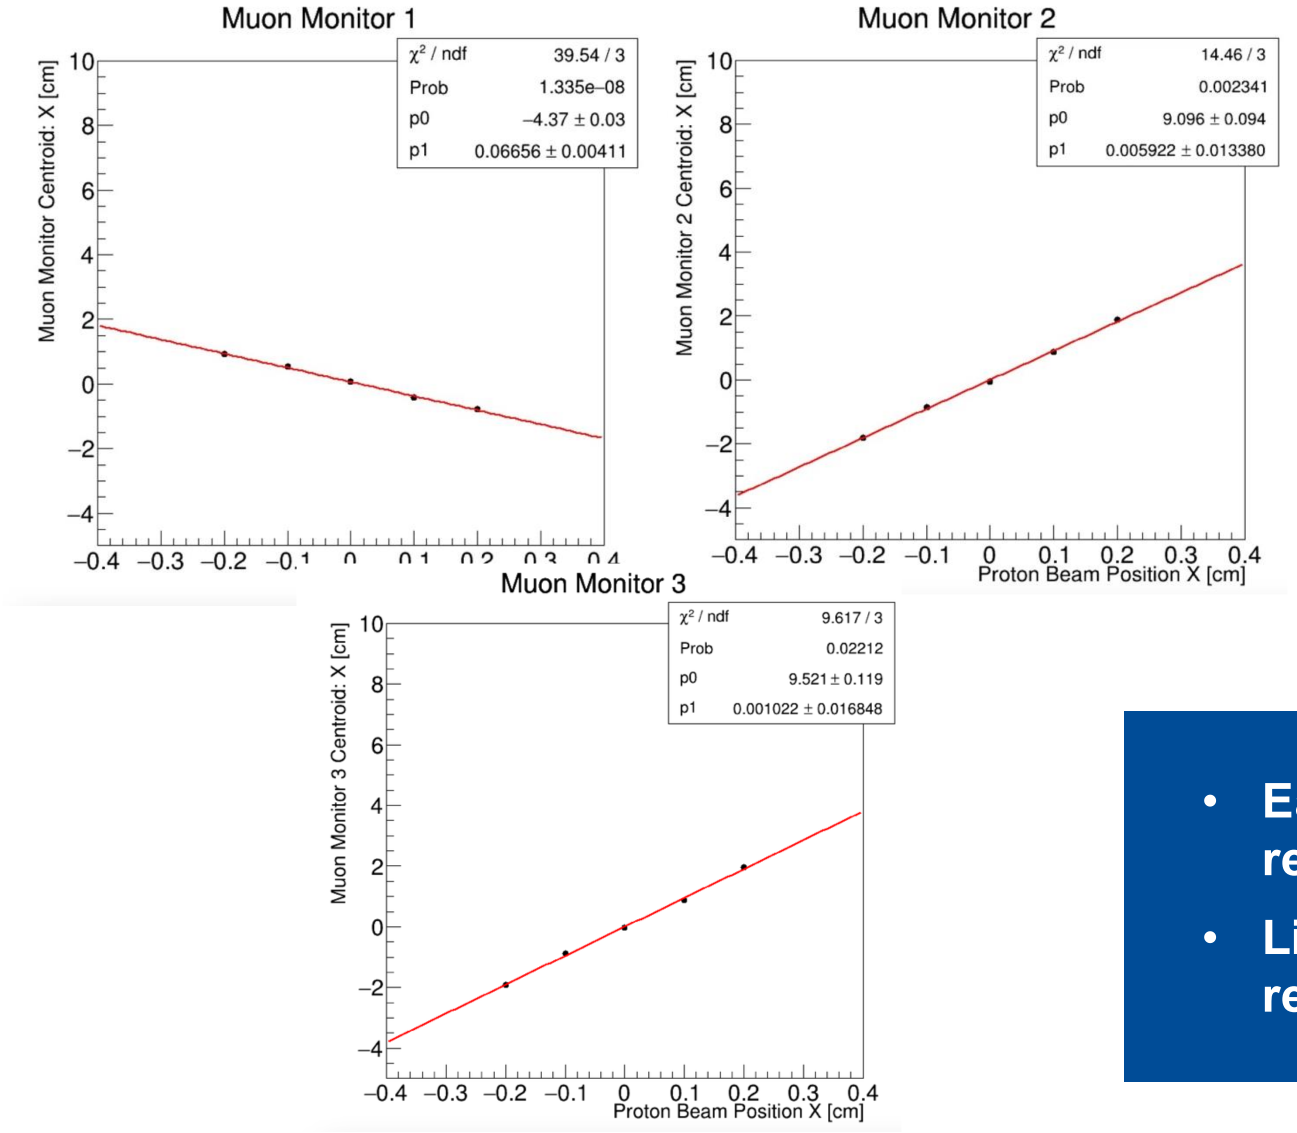

Station	Response	Physics
MuMS1	Negative slope	Low-energy muons
MuMS2	Positive slope	Medium-energy muons
MuMS3	Positive slope	High-energy muons

- **Each monitor has a unique but deterministic response to proton beam position**
- **Linear, monotonic response ⇒ suitable for real-time beam parameter inference**

Why Machine Learning is Needed

Muon monitor images depend simultaneously on

- Proton beam position
- Beam size
- Horn current
- Beamline alignment
- Hadron production

Each pixel provides a local measurement of the muon flux

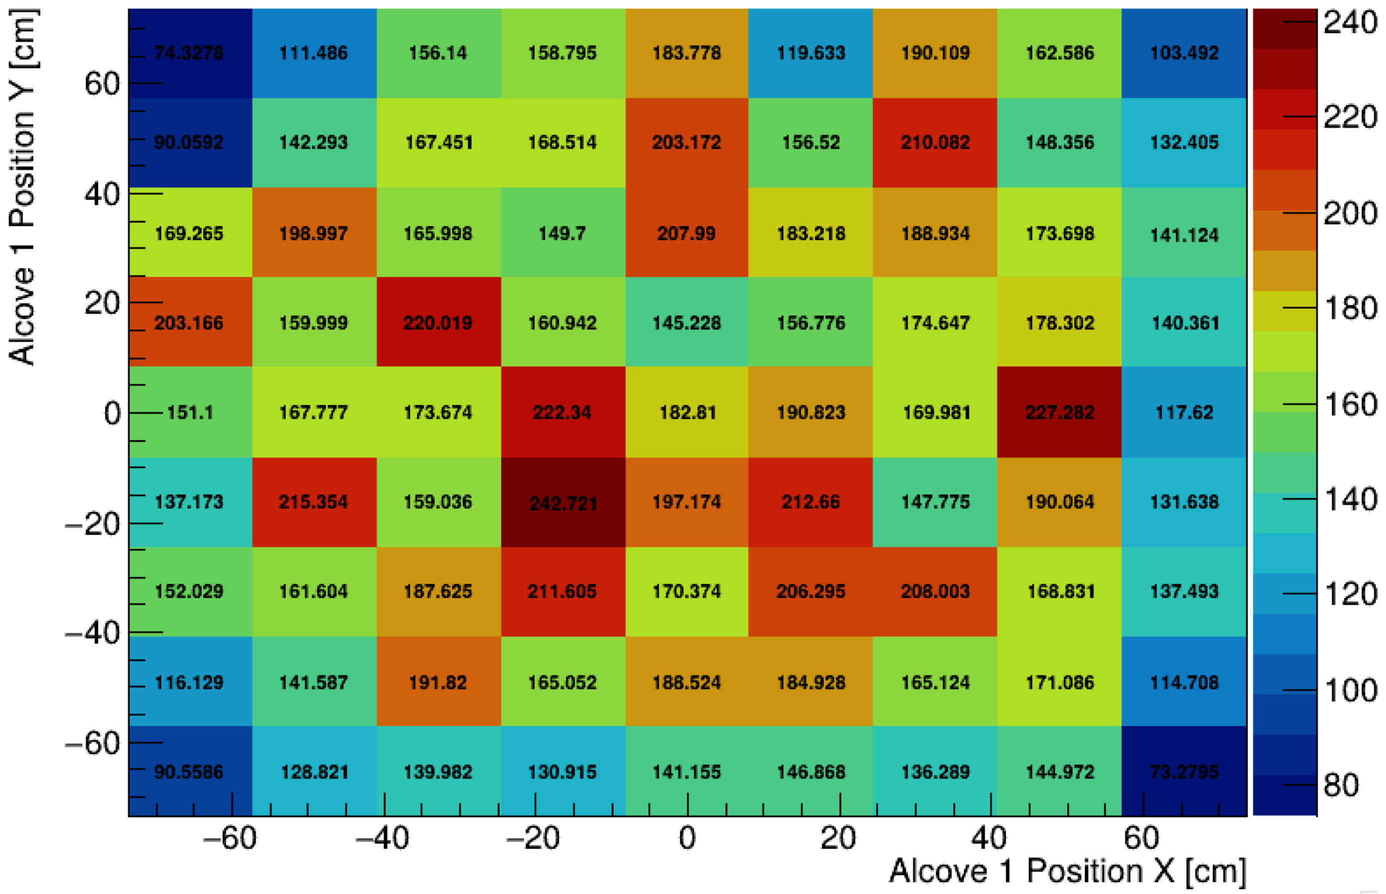

Each parameter leaves a predictable signature, but multiple parameters are superimposed in detector image
{MuMS pixel intensities}⟶{beam position, beam size, horn current, ...}

Beam Monitoring Framework

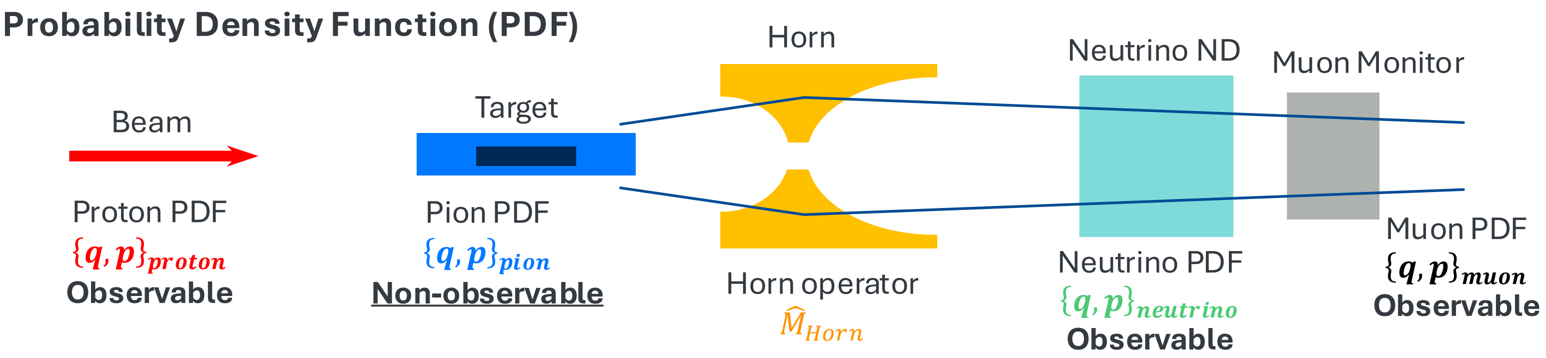

$$\begin{pmatrix} q \\ p \end{pmatrix}_{muon} \oplus \begin{pmatrix} q \\ p \end{pmatrix}_{neutrino} = \widehat{M}_{Horn} \cdot \begin{pmatrix} q \\ p \end{pmatrix}_{pion}$$
$$= \widehat{M}_{Horn} \cdot \widehat{M}_{Hadron\ prod.} \cdot \begin{pmatrix} q \\ p \end{pmatrix}_{proton}$$

- **Objective**
  - Recover beam parameters from MuMS images in real time
- **Inputs**
  - Fast beam instrumentation (WCT, BPM, BLM)
  - MuMS detector images

Outputs

Beam position
Beam size
Horn current
Neutrino beam prediction…

- In forward simulation we go from proton beam to muon monitors
- Our neural network learns opposite direction—from muon monitors back to proton beam conditions

Previous Demonstration at NuMI Beamline

- [Precision beam diagnostics at the NuMI facility using muon monitor observations]

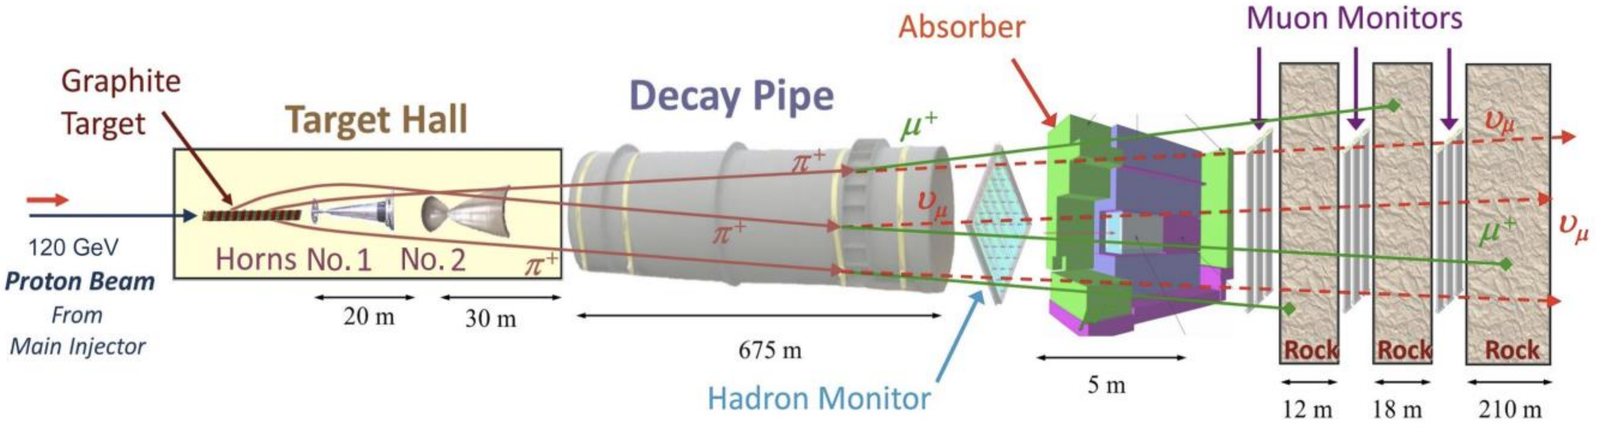

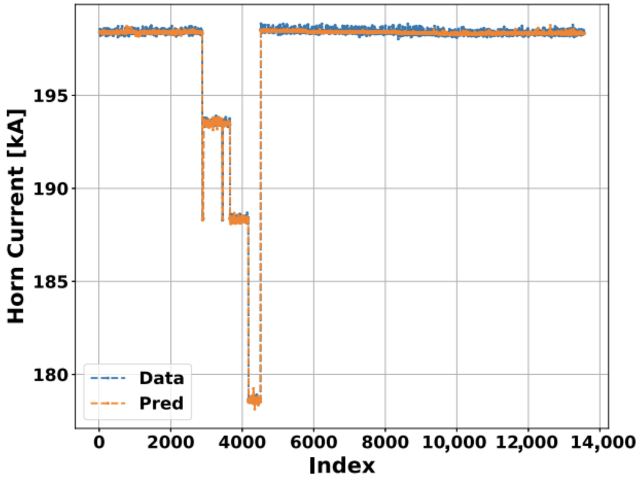

FIG. 30. Comparison between the horn current measured by the horn CT and that predicted by the trained ANN using muon monitor signals.

•Neural networks successfully reconstructed horn current and beam position directly from muon monitor signals

•Achieved sub-percent precision using operational NuMI data

•Our work extends this concept toward multi-parameter, simulation-trained beam inference for DUNE/LBNF beamline

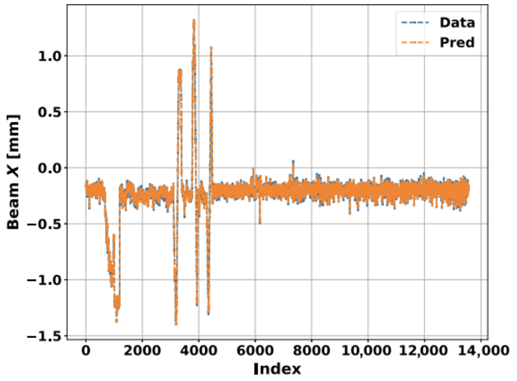

FIG. 32. Comparison between the horizontal beam position at the target, as measured by BPMs, and that predicted by the trained ANN model using muon monitor signals.

Demonstration using Operational J-PARC Muon Monitor Data

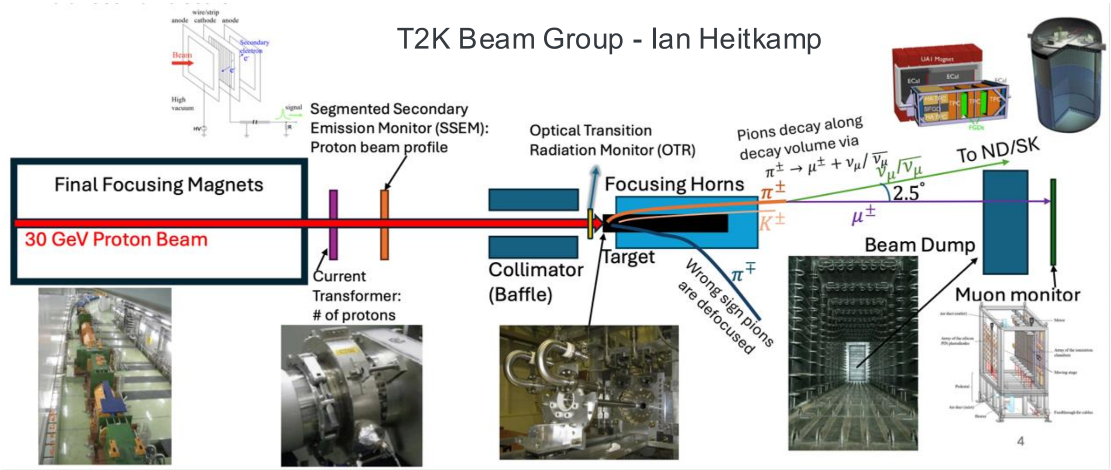

- Neural networks reconstruct multiple beam parameters simultaneously from muon monitor signals

- Proton beam position generalizes across independent physics runs (trained on commissioning + previous physics run, evaluated on a later physics run)

- Horn current is accurately reconstructed within a physics run

- These results show that muon monitor images contain rich information about coupled beam parameters and motivate more robust models for run-to-run beam operations

Within-run validation

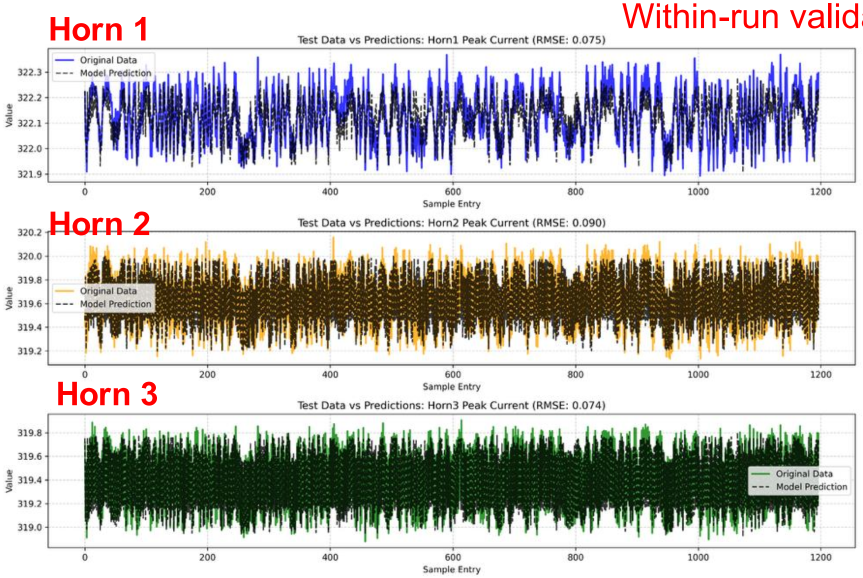

Independent-run validation

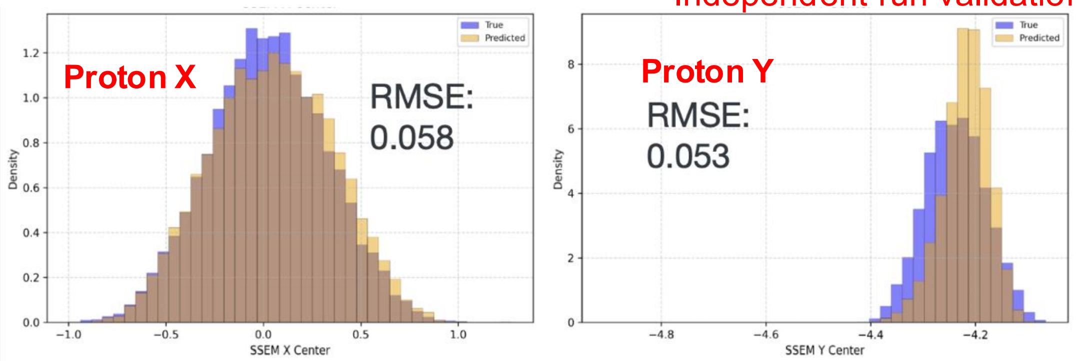

Work done by Metcalf 2026 Interns Maxwell Chen and Noah Aney

FERMI NATIONAL ACCELERATOR LABORATORY

Motivation for Physics-Informed Beam Emulation

NuMI /J-PARC demonstrations

↓

MuMS contains beam information

↓

Challenge:
Generalize to LBNF
Changing optics
Changing horn settings
Changing beam conditions

↓

How do we address challenge of adapting to LBNF and changing beam conditions?

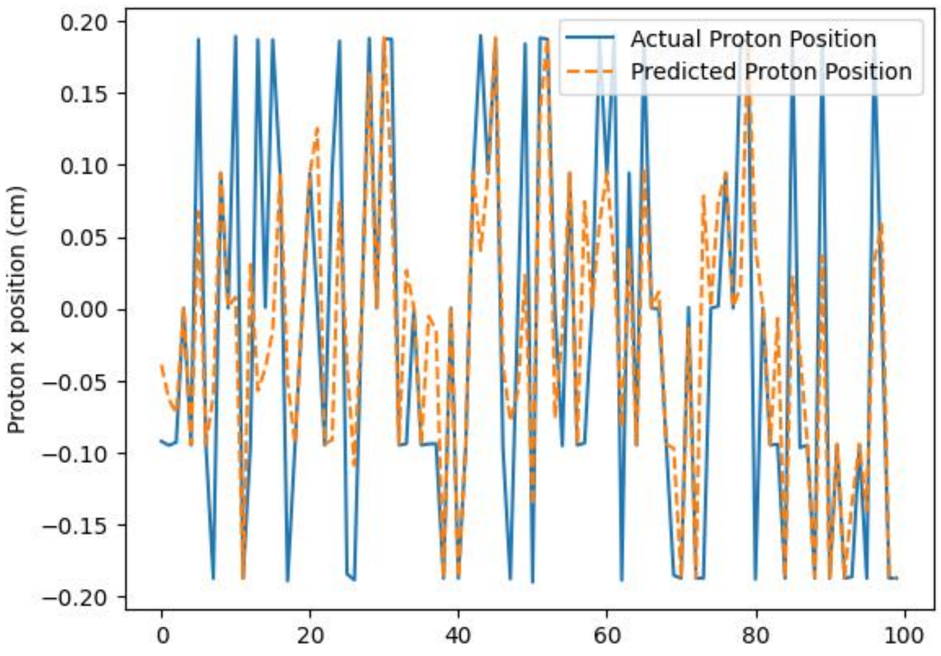

10^{10} POT simulation, No Bootstrapping

Bridget O'Brien
CCI Summer Intern, 2025
[CCI report]

Toward Real-Time AI-Assisted Beam Operations for LBNF

Online
- Measured MuMS observables are passed to the inference model
- Inference model reconstructs underlying beam parameters
- Physics emulator provides fast predictions and physics consistency checks

Continuous improvement
Periodically retrain using validated operational data

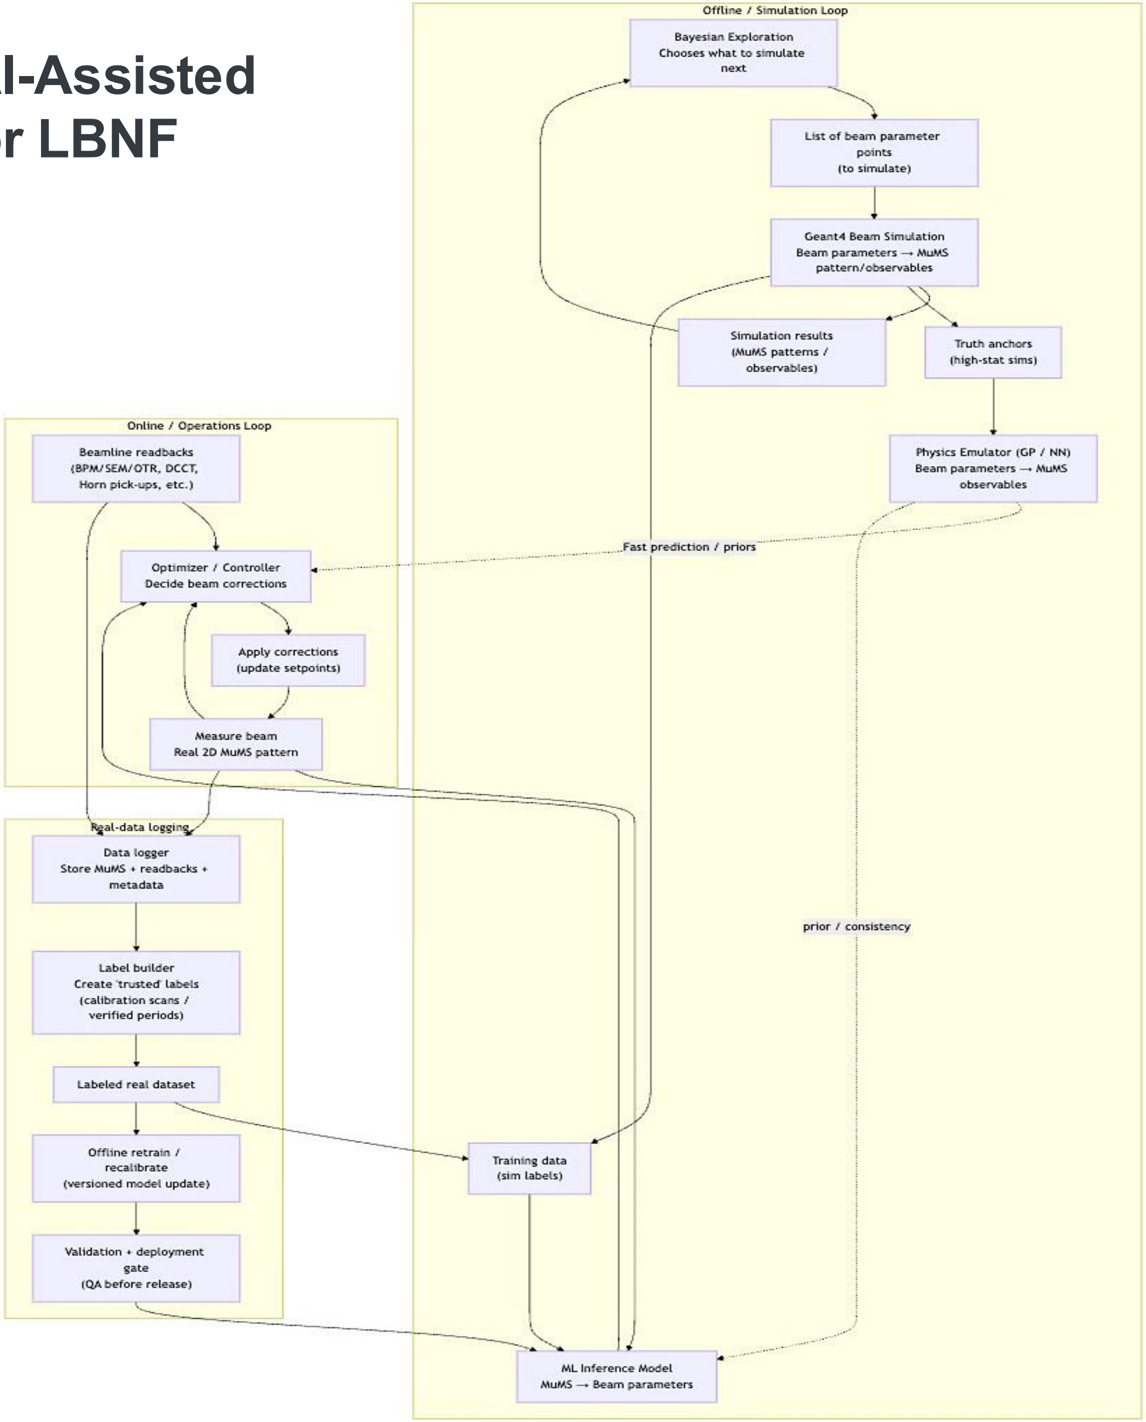

Offline
- Bayesian Exploration identifies most informative beam conditions to simulate
- High-stat Geant4 simulations become **physics truth anchors**
- A physics-informed emulator learns forward mapping:
  Beam parameters → MuMS observables

Toward Real-Time AI-Assisted Beam Operations for LBNF

Online
- Measured MuMS observables are passed to the inference model
- Inference model reconstructs underlying beam parameters
- Physics emulator provides fast predictions and physics consistency checks

Continuous improvement
Periodically retrain using validated operational data

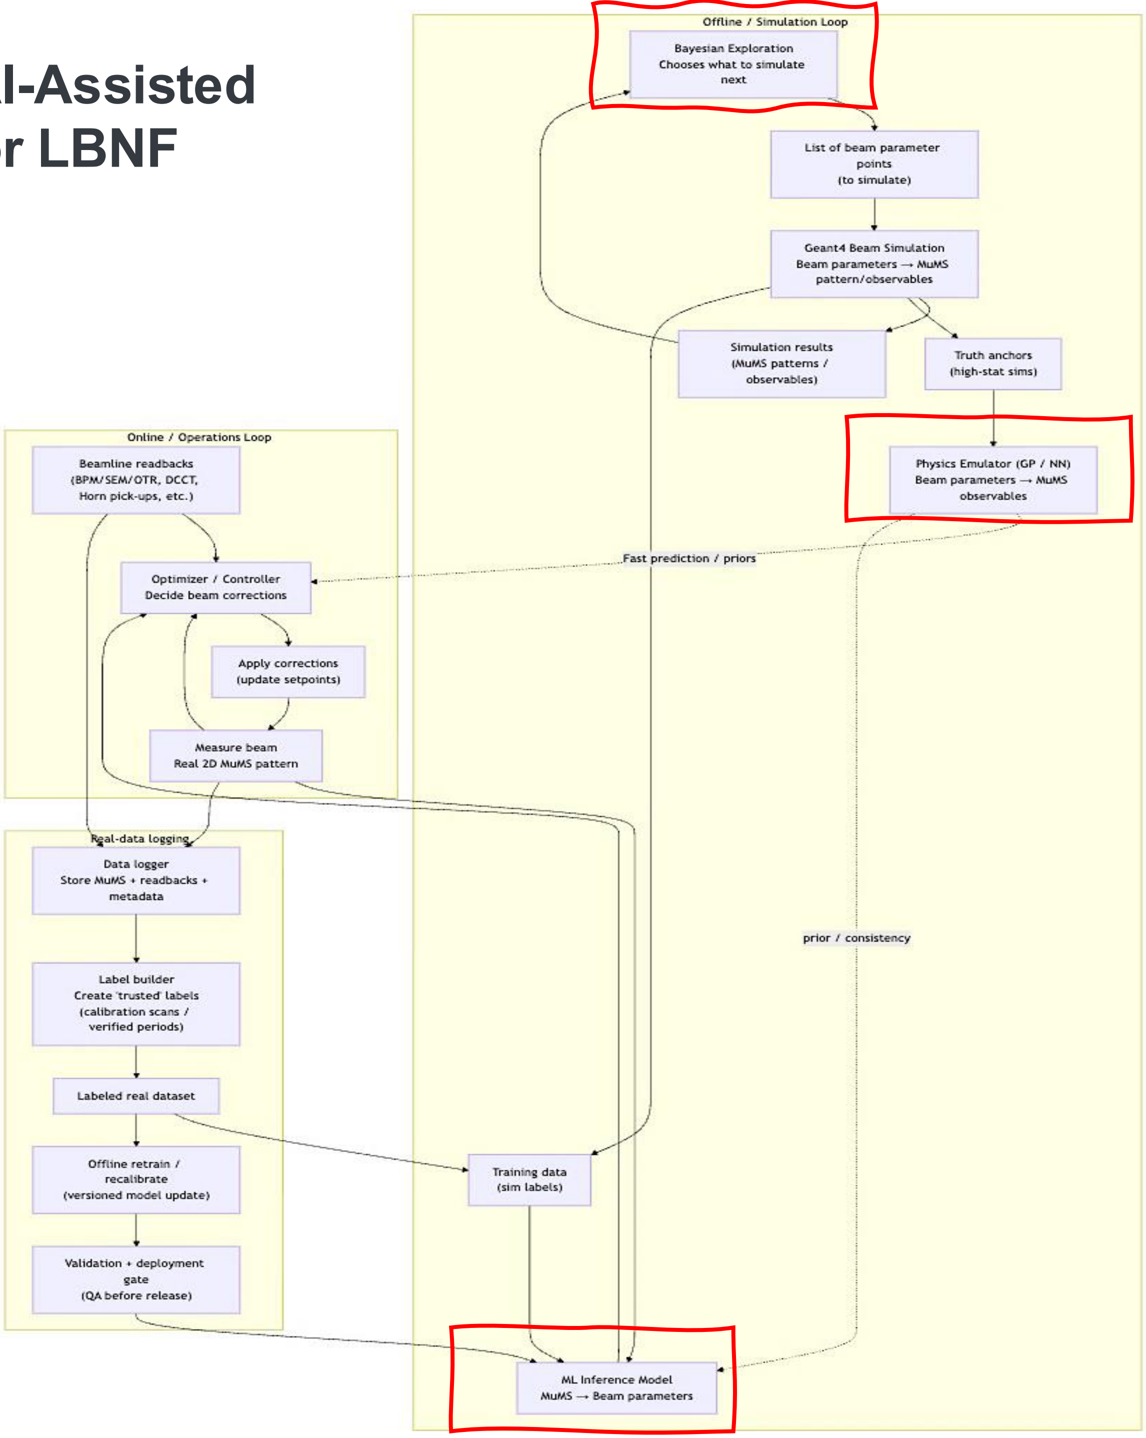

Offline
- Bayesian Exploration identifies most informative beam conditions to simulate
- High-stat Geant4 simulations become **physics truth anchors**
- A physics-informed emulator learns forward mapping: **Beam parameters → MuMS observables**

Building Physics Emulator Efficiently

Beam parameters
↓
Geant4 simulation
↓
MuMS observables
↓
Gaussian Process emulator
↓
Prediction uncertainty
↓
Bayesian Exploration chooses next simulation

Why?

- Geant4 takes hours per simulation
- Parameter space is continuous
- Exhaustive scans are impossible

Result

Bayesian exploration reduces required simulations from hundreds or thousands to ~40

Physics-Guided Bayesian Exploration

Each Muon monitor pattern gives three physics-motivated metrics, computed **relative to a nominal reference beam**

Δshape
Measures change in transverse MuMS intensity pattern
→ sensitive to focusing distortions

Δgradient
Measures change in spatial slope of the MuMS profile
→ sensitive to beam steering and asymmetry

Δcentroid
Measures shift of MuMS center of mass
→ sensitive to beam pointing

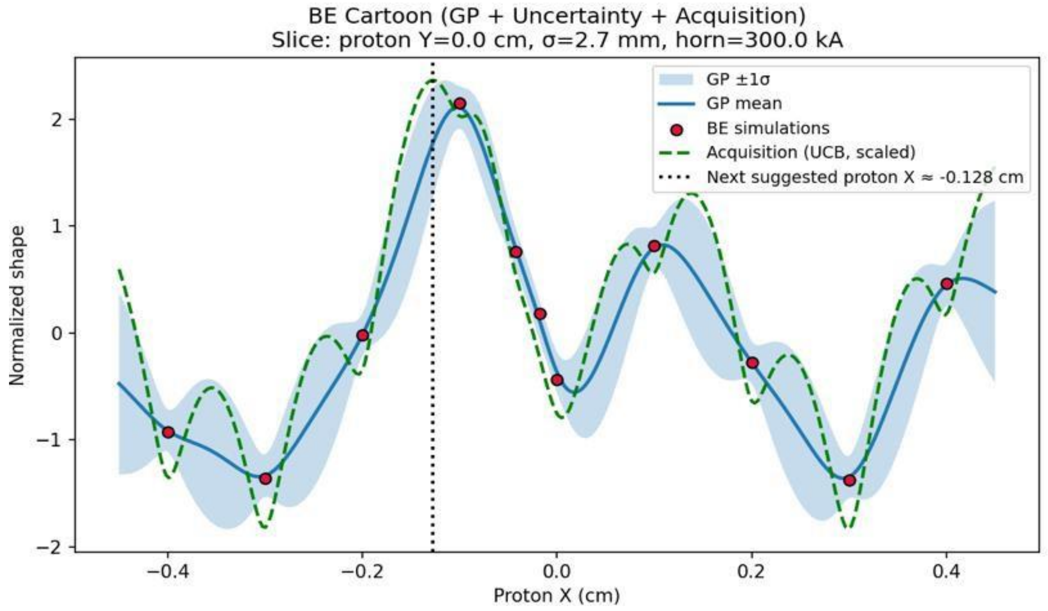

•Physics observables compress full MuMS images into meaningful features

•Gaussian Process learns response + uncertainty

•Bayesian Exploration places new simulations where information gain is larges

FERMI NATIONAL ACCELERATOR LABORATORY

Bayesian Exploration Learns Beamline Response

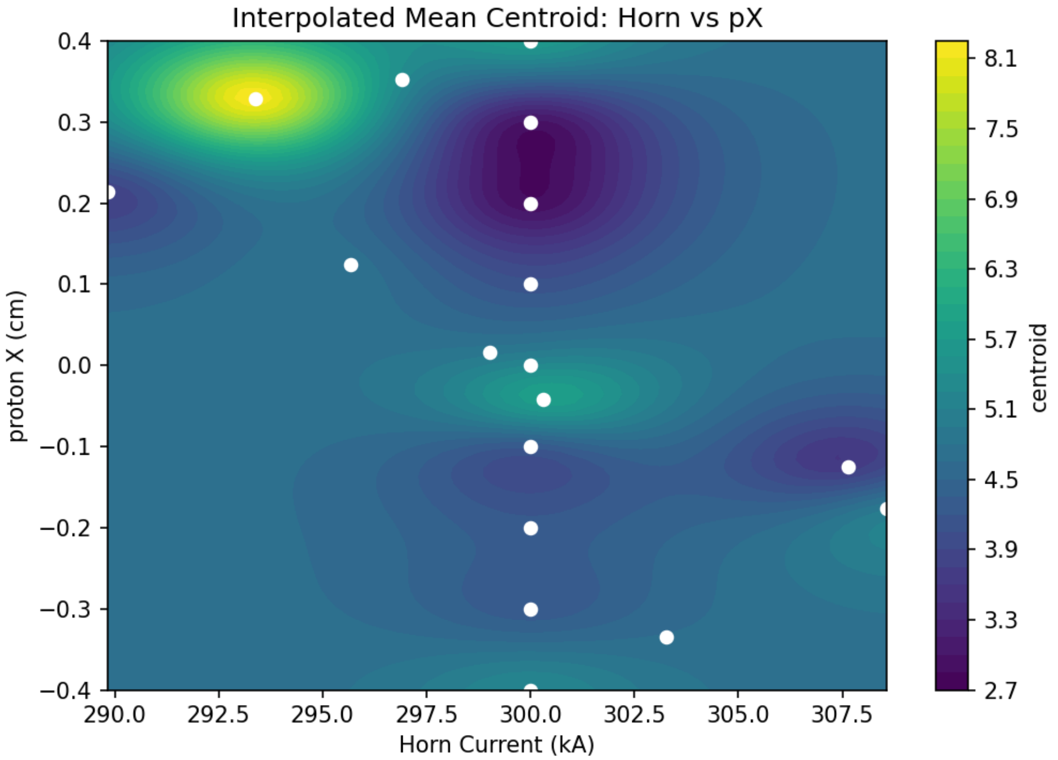

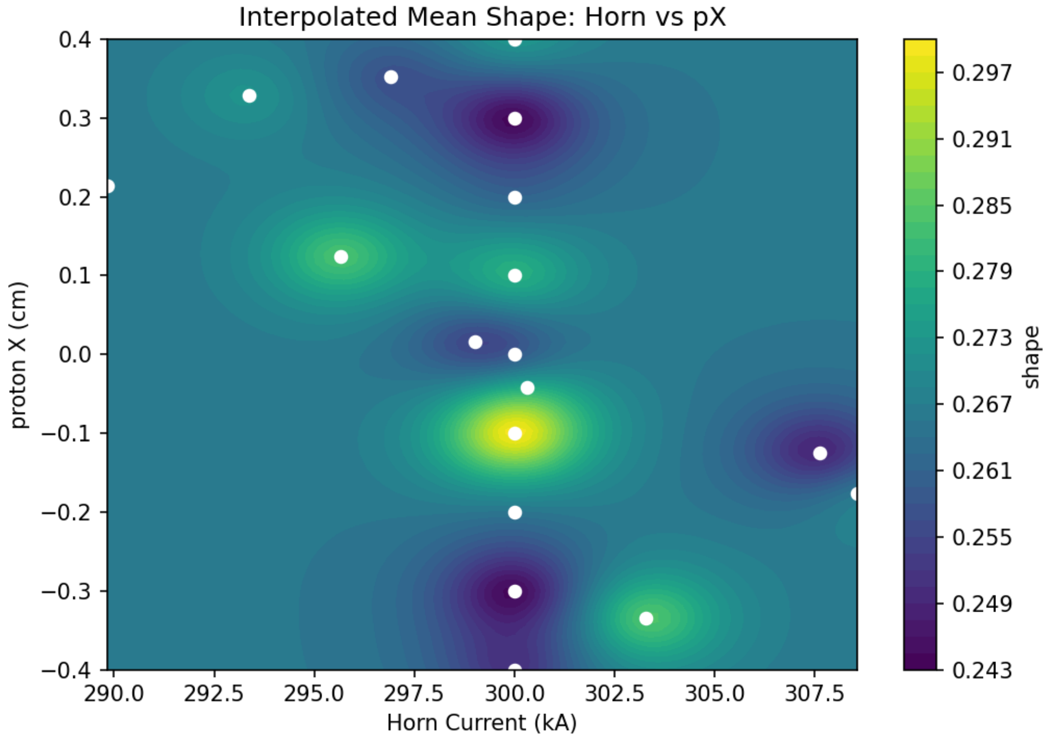

Centroid varies smoothly with horn current and proton X. The structure is mostly horn-dominated

Shape shows localized extrema and strong coupling between proton X and horn current

- Bayesian Exploration learns multidimensional beam-response surface
- Muon monitor observables depend jointly on proton position and horn current
- Learned response identifies where the beamline is most sensitive

FERMI NATIONAL ACCELERATOR LABORATORY

What Bayesian Exploration Provides

•Identifies which beam parameters most strongly affect the muon-monitor response

•Identifies regions where beam parameters produce similar detector signatures

•Identifies regions where additional Geant4 simulation is most informative

Bayesian Exploration converts an expensive Geant4 simulation campaign into a compact set of informative truth anchors that capture beam-response physics while minimizing number of simulations required

↓

These truth anchors become training and validation dataset for physics-informed surrogate emulator

Proof-of-Principle Operational Surrogate Emulator

Using an initial set of Geant4 beam simulations, we trained a surrogate emulator and evaluated its ability to interpolate to unseen beam conditions

- **Training data**
  - Geant4 proton-X scan
  - Training points at selected beam positions
  - Independent intermediate positions reserved for validation
- **Validation**
  - Validation beam positions not used during training
  - Tests interpolation between simulated beam conditions
- **Results**
  - Predicts centroid and gradient observables
  - Validation points lie within predicted uncertainty
  - Captures smooth beam-response trends
- This provides initial surrogate that will seed Bayesian Exploration in future iterations

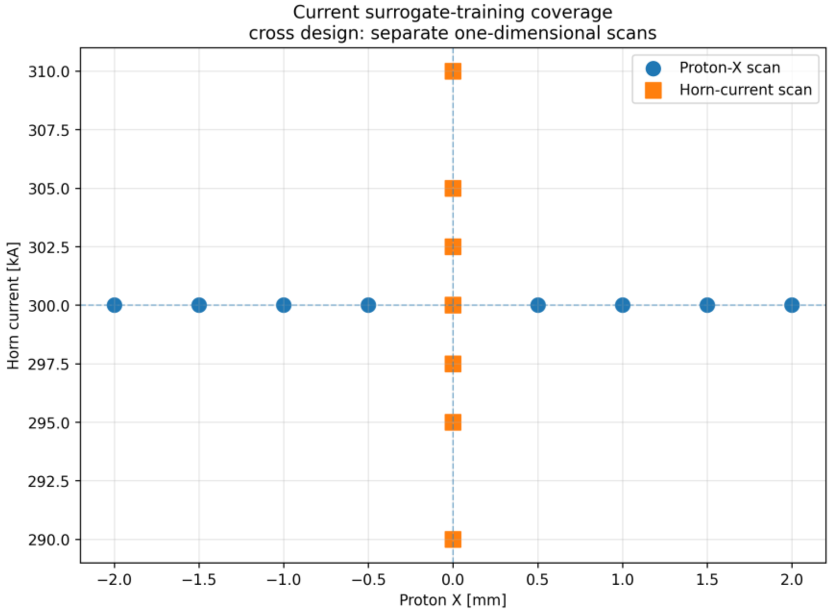

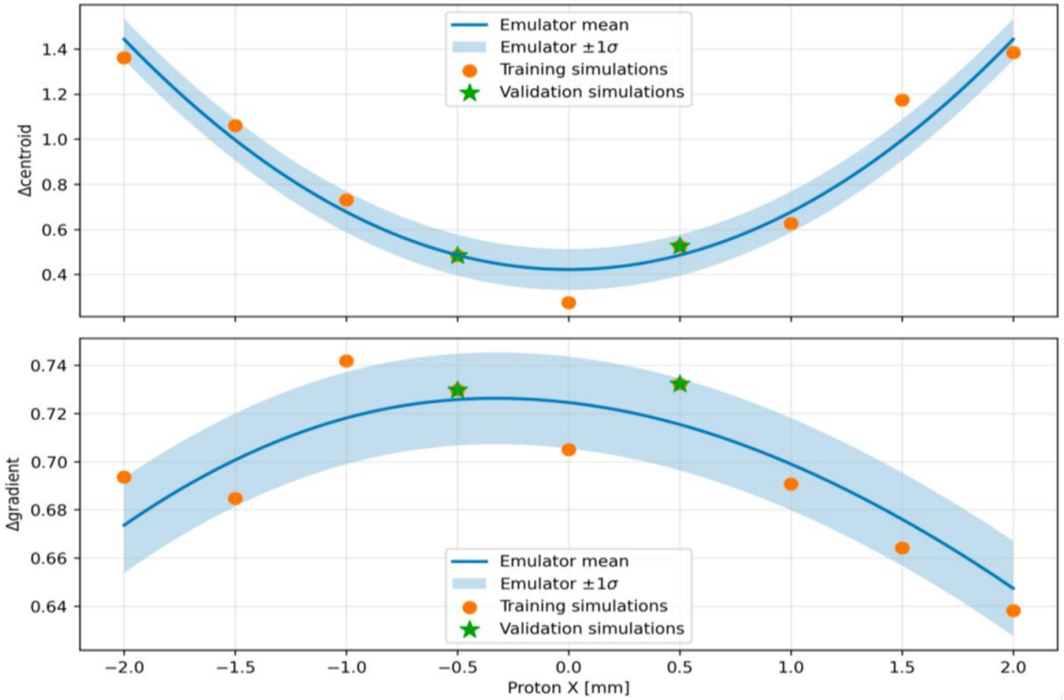

Distinct Muon-Monitor Signatures from Proton Position and Horn Current

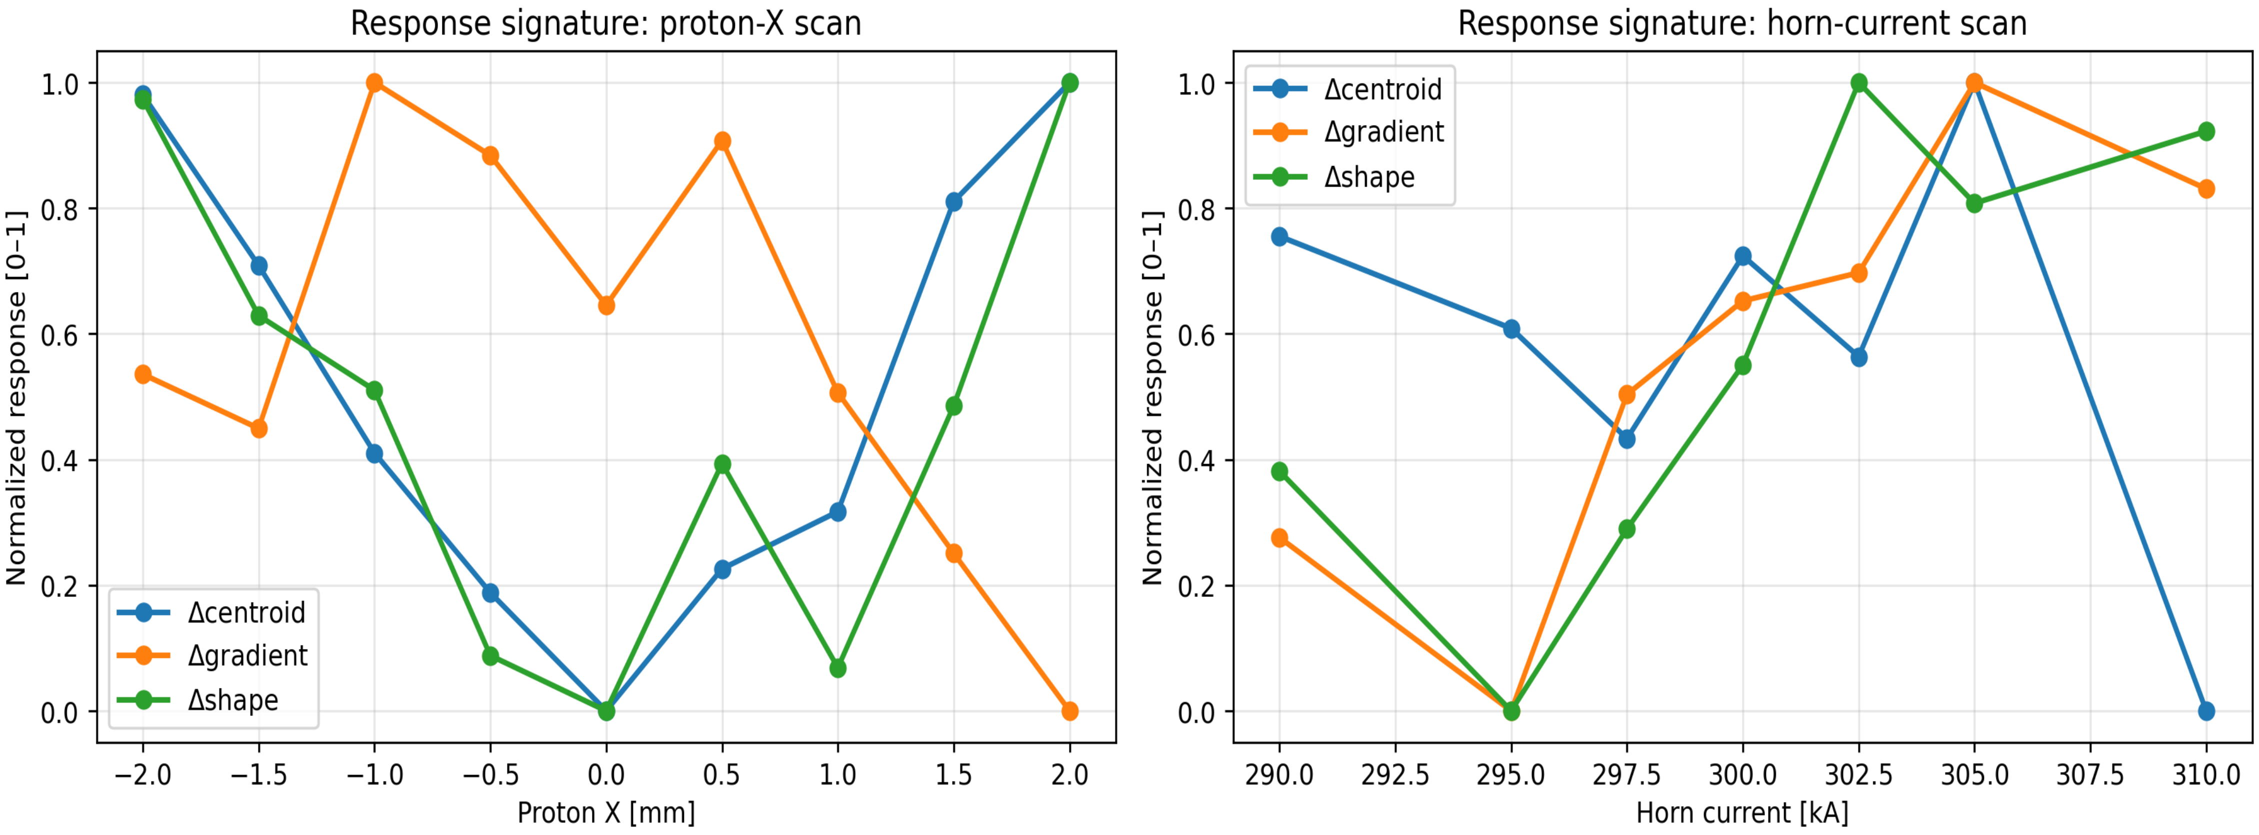

- **Detector does not respond to every beam parameter in same way**
- Proton steering and horn focusing imprint different patterns across set of observables

Can Different Beam Perturbations Be Distinguished?

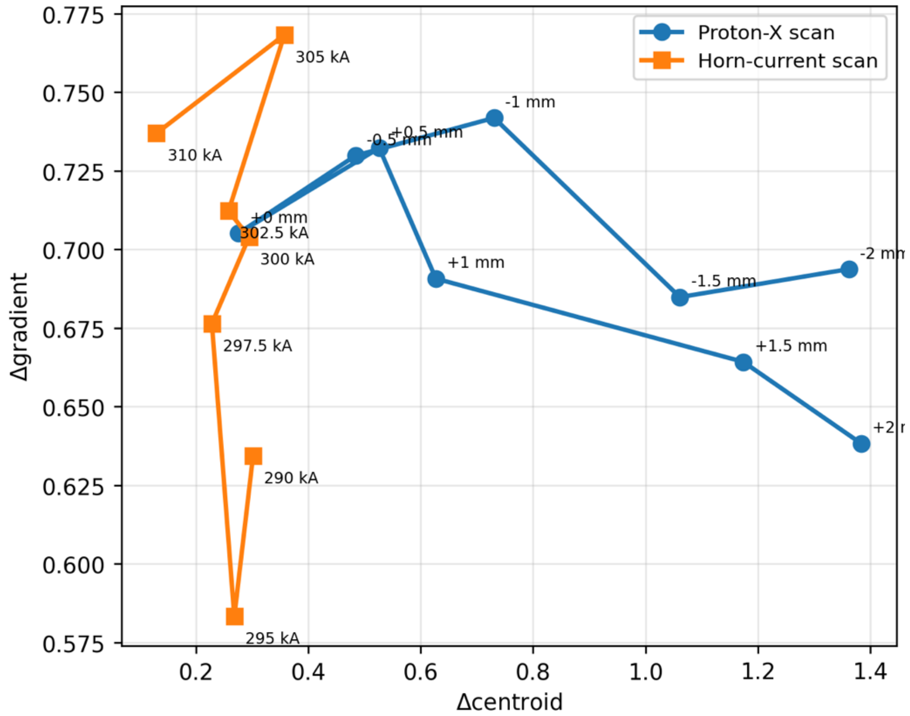

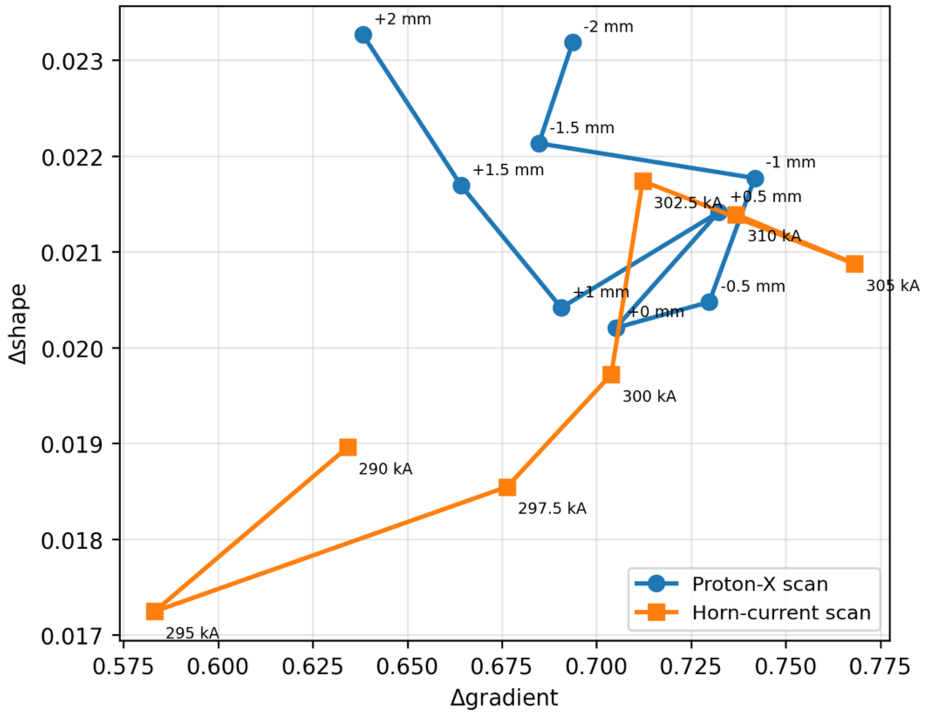

Example
Suppose muon monitor measures
•ΔCentroid = 0.70
•ΔGradient = 0.72
Is this due to a **0.5 mm proton-X shift** or a **302.5 kA horn current**?
Include shape observable, trajectories separate

•Some observables respond similarly to different beam perturbations

•Full set of observables separates proton steering from horn focusing

•This multidimensional response enables beam-state reconstruction

FERMI NATIONAL ACCELERATOR LABORATORY

How are these signatures used?

- Distinct muon-monitor signatures encode beam state
- Bayesian Exploration selects most informative Geant4 simulations
- Selected Geant4 simulations train a physics-informed forward emulator

Beam parameters → Muon-monitor observables

- Inference model solves the inverse problem
  Muon-monitor observables → Beam parameters
- These learned detector signatures provide information required for spill-by-spill beam-state reconstruction

Key idea:
Forward model: learns detector response to changes in beam state
Inference model: jointly reconstructs beam state by resolving degeneracies in detector response.

Motivation from J-PARC
- **Multiple beam parameters can produce similar MuMS signatures**
- Including beam angle in inference model breaks this degeneracy & significantly improves proton-position reconstruction

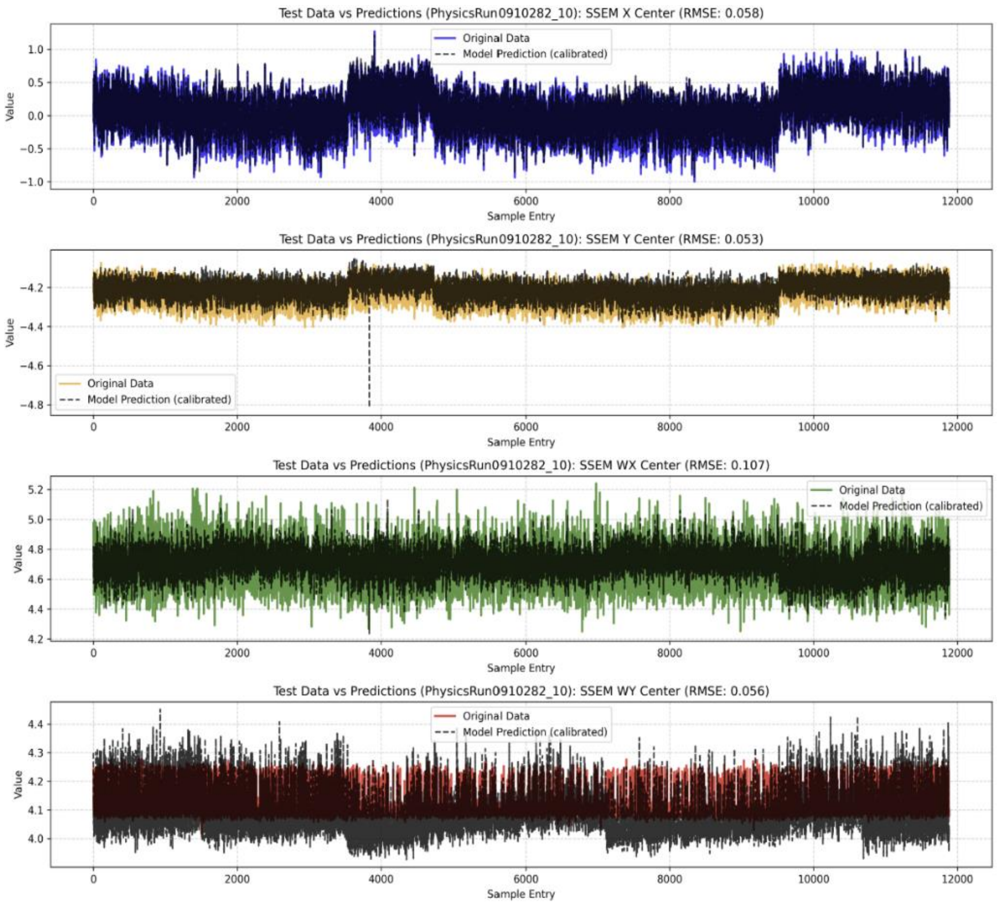

Work done by Metcalf 2026 Interns Maxwell Chen and Noah Aney

Toward AI-Assisted Beam Operations for LBNF

- Bayesian Exploration reduces expensive Geant4 simulations
- Physics-informed emulator predicts unseen beam conditions
- Existing beam data improves inference robustness
- Framework extends naturally toward spill-by-spill beam-state reconstruction

Simulation → Emulator → Inference → AI-Assisted Beam Operations

Toward AI-Assisted Beam Operations for LBNF

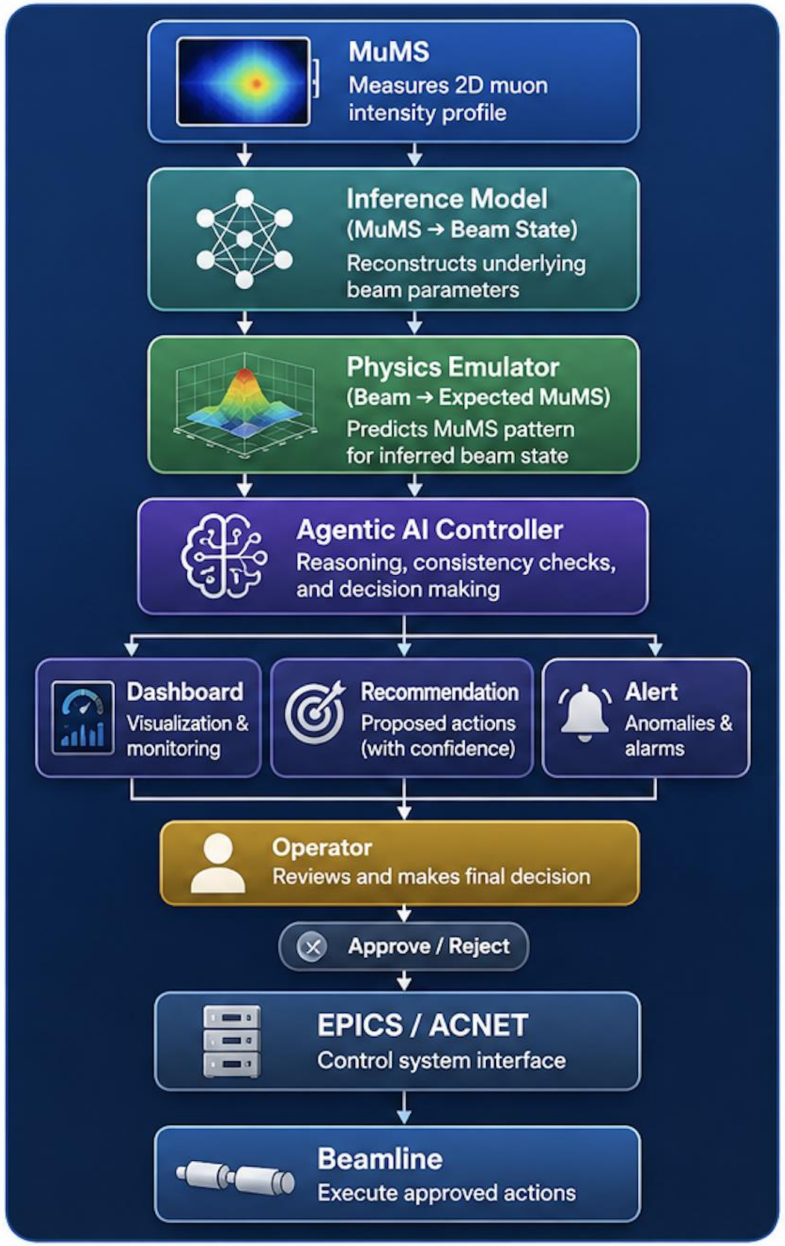

Agent has access to

- inferred beam parameters
- emulator predictions
- uncertainties
- machine readbacks
- operator constraints
- beam history

Inference model answers

- What beam do I think produced this MuMS image?

Emulator answers

- If beam really had those parameters, what MuMS image should I see?

Agent asks

- Given all this information, what should I do next?

Thank you!

Fermilab

FermiFORWARD

U.S. DEPARTMENT of ENERGY

FERMI NATIONAL ACCELERATOR LABORATORY