# AI-Optimized Polarization at Jefferson Lab

Patrick Moran[2] David Lawrence[1] Thomas Britton[1] Hovanes Egiyan[1] Cristiano Fanelli[2] Jiawei Guo[3] Will Imoehl[1,3] Naomi Jarvis[3] Torri Jeske[1] Armen Kasparian[1] Chris Keith[1] James Maxwell[1] Malachi Schram[1]

[1]Thomas Jefferson National Accelerator Facility

[2]The College of William & Mary

[3]Carnegie Mellon University

Digital Twins for Nuclear and Particle Physics 2026

# GlueX Experiment

- Hall D experiment at Jefferson Lab
- Linearly polarized photon beams for photoproduction
- Physics goal: exotic hybrid mesons, nonperturbative QCD
- Requires precise control of photon spectrum and polarization
- Beam quality directly impacts statistical precision of measurements

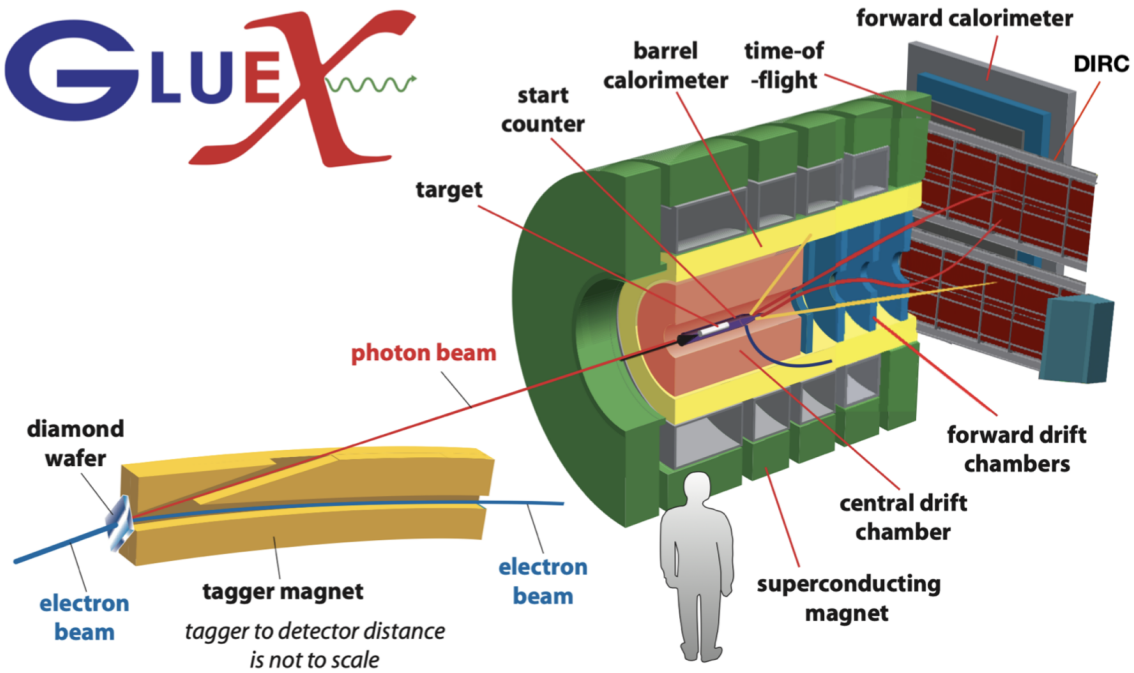

# Coherent Bremsstrahlung

- Electron beam interacts with periodic diamond lattice
- Constructive interference enhances photon yield at specific energies
- Produces a sharp coherent edge in the spectrum
- Edge position determines usable polarization region
- Sensitive to sub-milliradian changes in crystal orientation

PHYSICAL REVIEW C **95**, 042201(R) (2017)

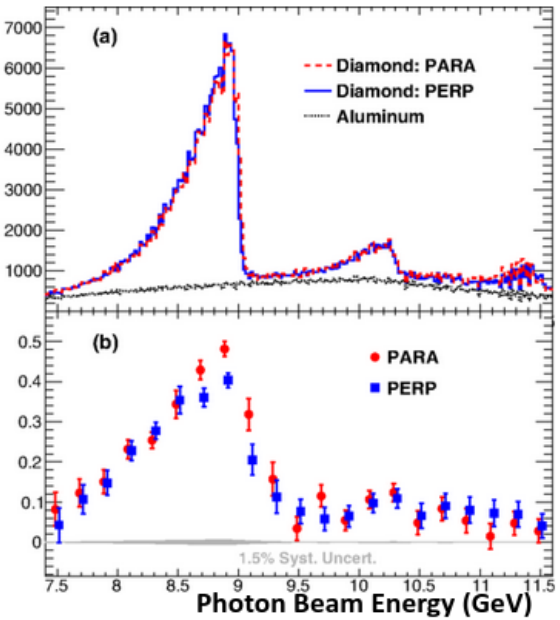


*(a) Photon beam intensity versus energy as measured by the pair spectrometer (not corrected for instrumental acceptance)*
*(b) Photon beam polarization as a function of beam energy, as measured by the triplet polarimeter, with data points offset horizontally by ±0.015 GeV for clarity.*

# Diamond + Goniometer

- Thin diamond radiator mounted on precision multi-axis goniometer
- Degrees of freedom:
  - Pitch: vertical tilt
  - Yaw: horizontal rotation
  - Roll: sets polarization orientation (fixed during run)
- Sub-milliradian control resolution required

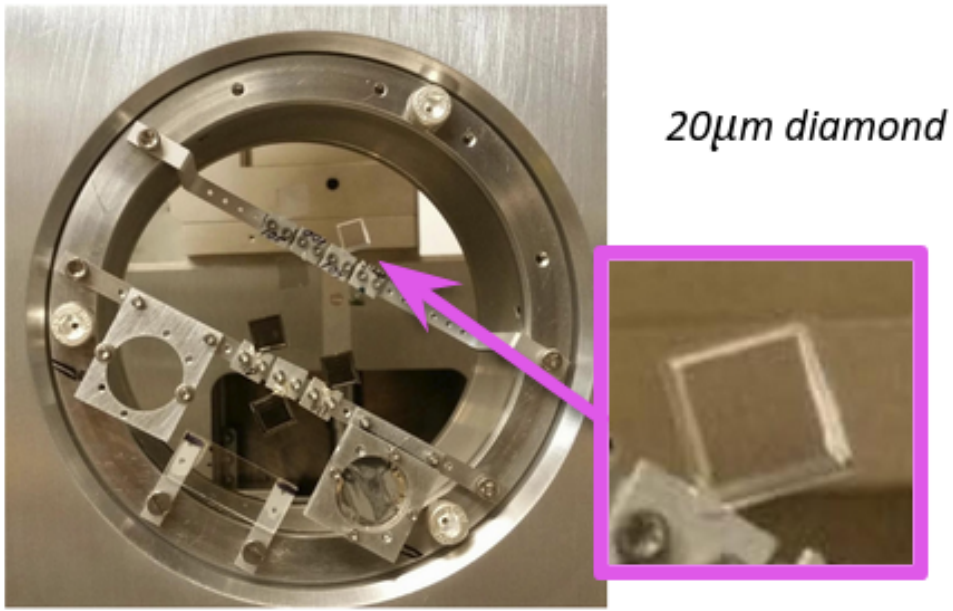

# Goniometer System

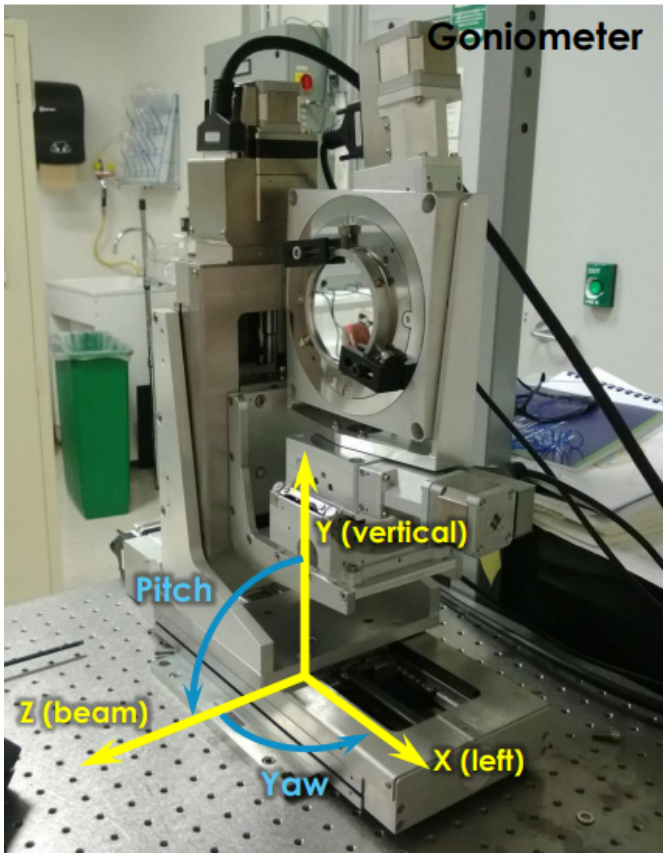

# Operational Workflow

- Initial setup:
  - Select beam spot on diamond
  - Set polarization orientation (roll)
- Online tuning:
  - Adjust pitch/yaw to align coherent edge
  - Monitor spectrum and apply corrections
- Continuous maintenance during data taking

**Limitations:**

- Requires expert operators
- Time-consuming and subjective
- Sensitive to drift and fluctuations

# Current Goniometer Control GUI

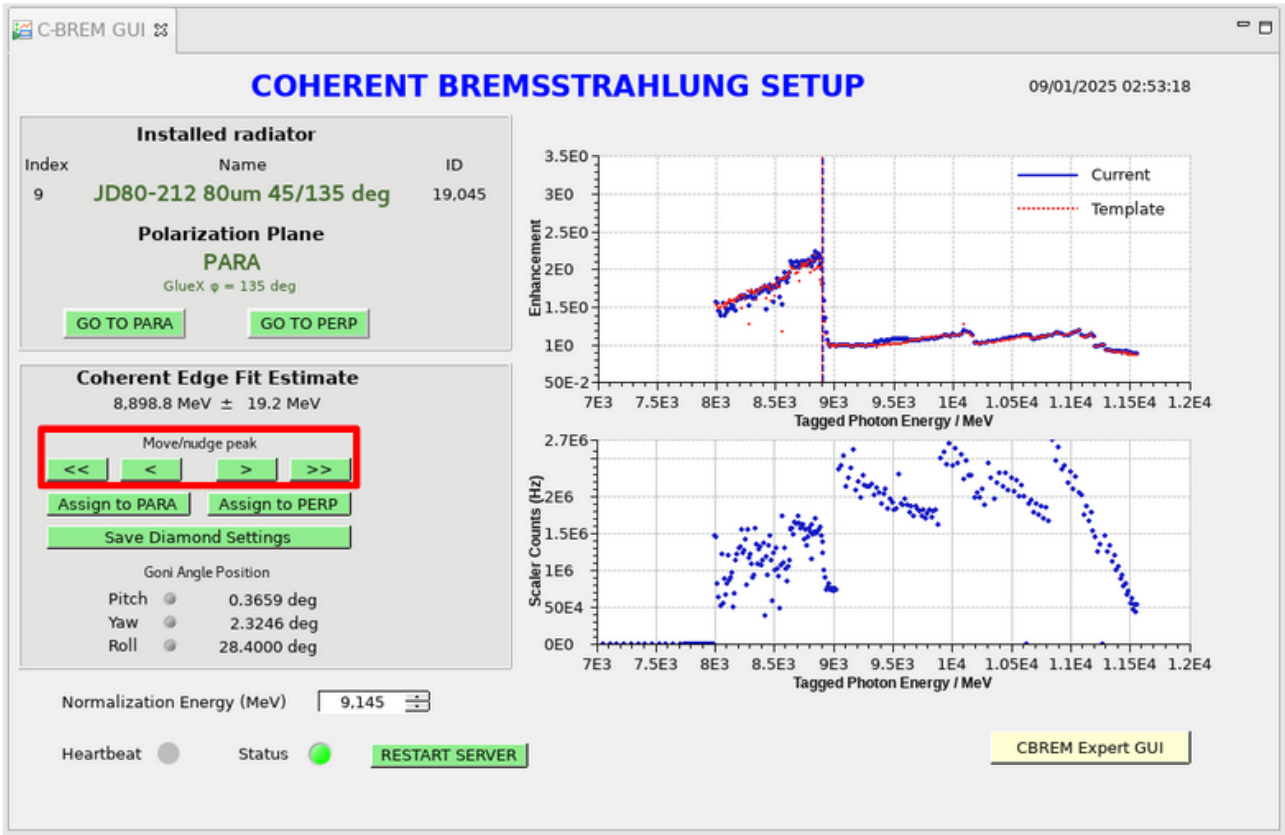

# Control Problem Formulation

- Objective: maintain $E_{peak} \approx E_{target}$
- Control inputs: discrete pitch and yaw adjustments
- Observations include:
  - Current peak position and target
  - Relative error and sign
  - Recent action history
  - Beam position and angle
  - Accumulated beam dose
- System evolves over both fast (noise) and slow (degradation) timescales

# Coherent Edge Physics

$$c \approx \frac{k}{gE_{beam}^2}\left(\frac{1}{E_{edge}} - \frac{1}{E_{beam}}\right)$$

- $c$: angle between beam and crystal plane
- Determines coherent edge position
- Small angular changes lead to MeV-scale energy shifts

# Non-Ideal Effects

- **Backlash:**
  - Mechanical slack when reversing direction
  - Leads to ineffective control steps
- **Beam fluctuations:**
  - Variations in beam pitch/yaw
  - Introduces stochastic energy shifts
- **Radiation damage:**
  - Accumulated dose degrades crystal quality
  - Alters system response over time

# Backlash Illustration

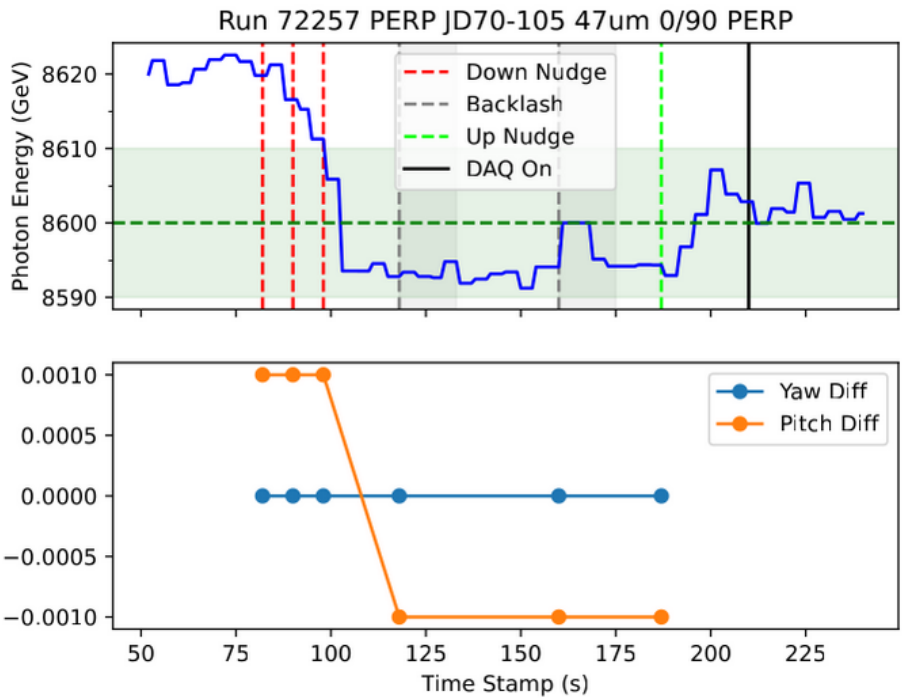


- Not all control actions produce motion
- Effect depends on recent direction history
- Creates non-Markovian dynamics

# Why Classical Control Struggles

- Nonlinear coupling between angles and energy
- Hidden state (backlash, beam motion)
- Discrete and delayed actuation
- Time-varying system due to degradation

**Result:** difficult to design stable model-based controller

# Simulation (Control Environment)

- Combines physics model with empirical effects
- Includes:
  - Coherent edge kinematics
  - Beam pitch/yaw noise
  - Backlash dynamics
  - Dose-dependent drift
- Provides realistic training environment for control policies

*Can be viewed as a control-oriented digital twin*

# Reinforcement Learning Setup

- PPO algorithm (policy gradient method)
- Discrete action space: $a \in [-10, 10]$
- State vector includes:
  - Peak position and target
  - Error metrics
  - Orientation index
  - Action history (handles hysteresis)
  - Beam pitch and yaw
  - Accumulated dose

# Reward Function

$$R_t = -\lambda_{err}\epsilon_{rel} + \lambda_{bonus}e^{-\epsilon_{rel}^2/2\sigma^2} - \lambda_{act}\|a_t\|$$

**Relative error:** $\epsilon_{rel} = \frac{|E_{peak} - E_{target}|}{E_{target}}$

- Tracking penalty: drives global convergence
- Gaussian bonus: stabilizes near target
- Action penalty: discourages unnecessary motion

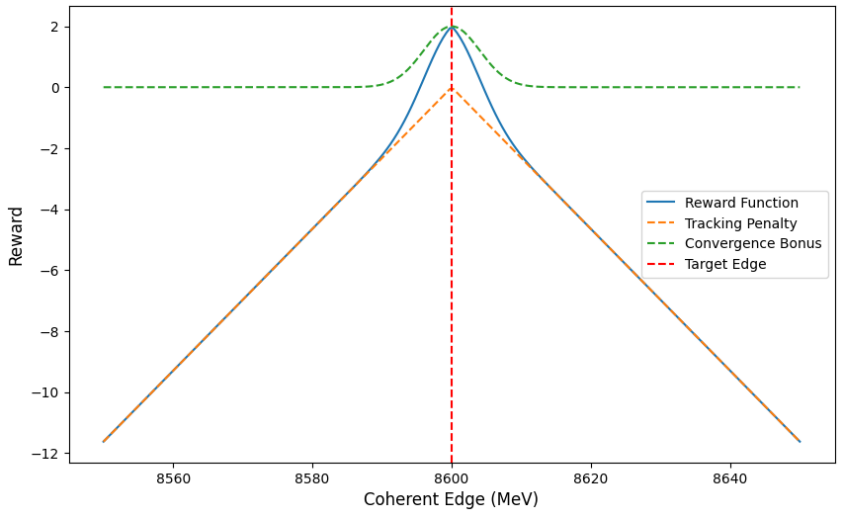

# Training Strategy

- Randomized initial conditions and targets
- Train separate agents for all four possible diamond orientations
- Deterministic evaluation for performance tracking
- Learning rate scheduling for stable convergence

| | |
|---|---|
| Algorithm | PPO (MLP policy) |
| Total timesteps | 300,000 |
| Max episode steps | 500 |
| Learning rate | $1 \times 10^{-4}$ (linear decay) |
| Discount factor $\gamma$ | 0.99 |
| Rollout steps | 2048 |
| Batch size | 64 |
| Epochs/update | 10 |
| Target edge range | 8.4–8.8 GeV |
| Action resolution | $2 \times 10^{-4}$ deg $\times$ [-10,10] |

# Learning Performance

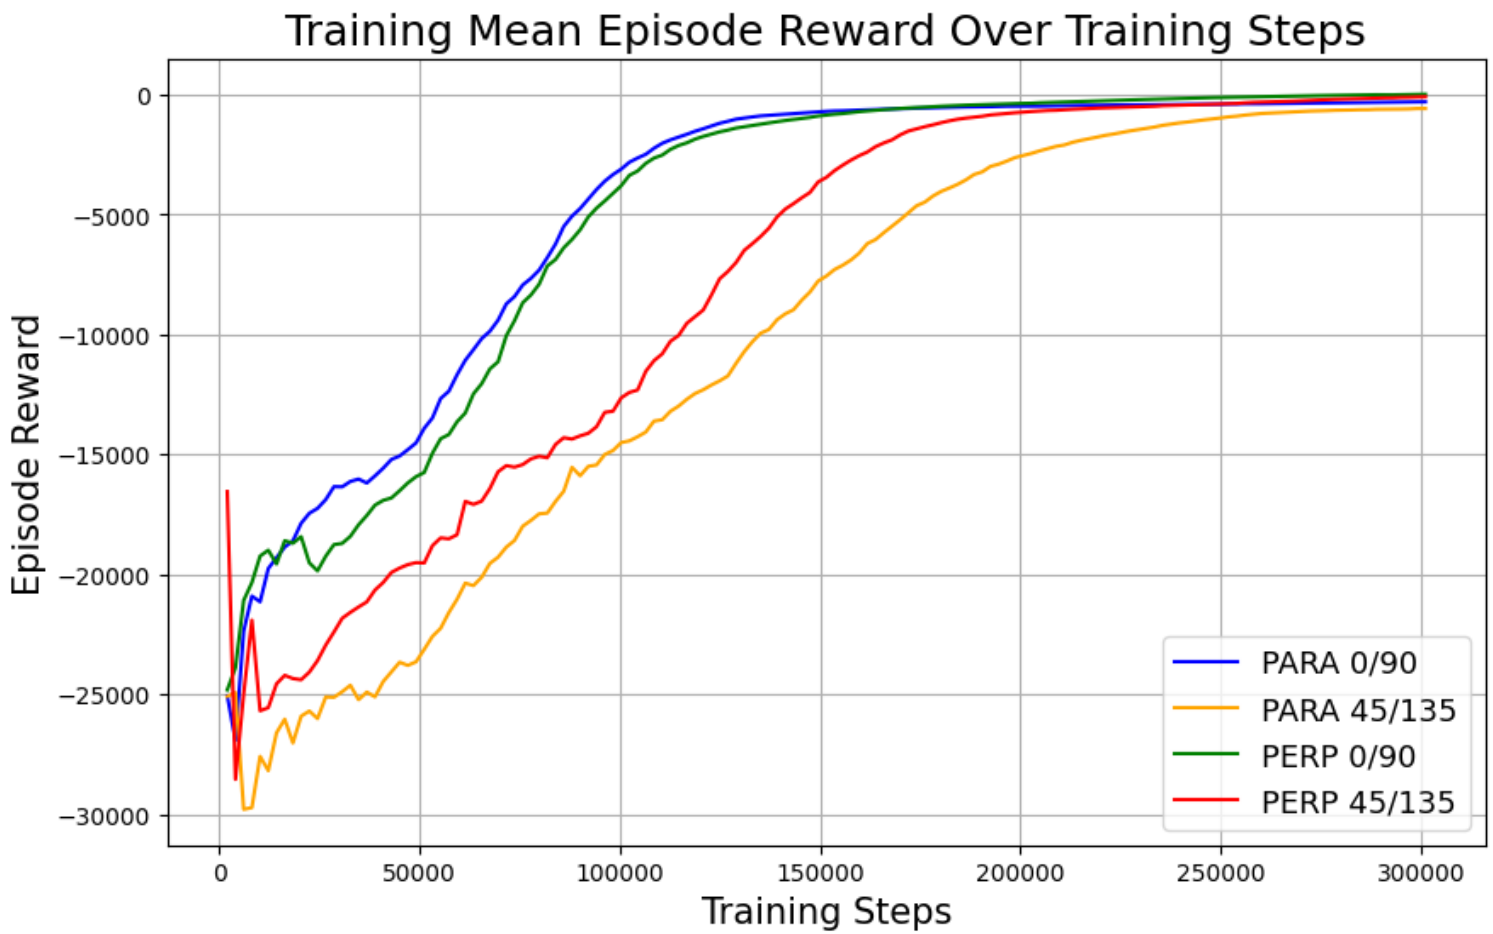


- Reward improves steadily during training
- Evaluation performance stabilizes

# Control Performance

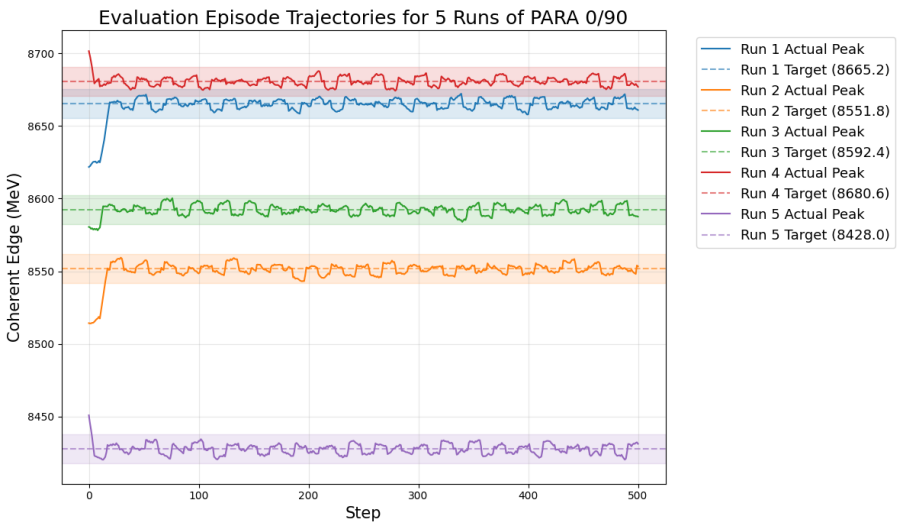


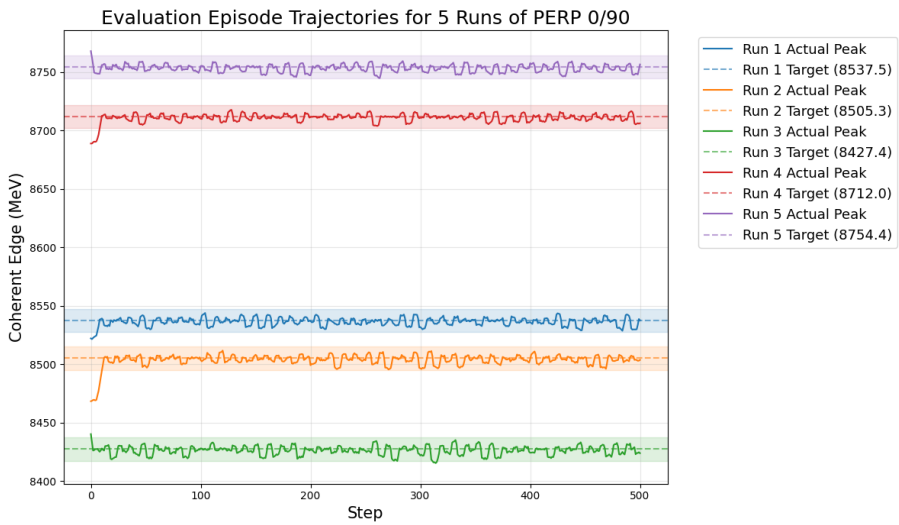


- Rapid convergence to target (10–20 steps) over multiple target edge values
- Stable tracking within tolerance band
- Consistent behavior across different initial conditions

# Behavior and Interpretation

- Policy learns implicit model of system dynamics
- Uses history to overcome backlash
- Adapts to beam fluctuations and slow degradation
- Balances accuracy with minimal actuation

# Integration at JLab

- RL controller being integrated into control system
- Guardrails ensure safe operation:
  - Limits on motion and step size
  - Monitoring of beam conditions
  - Safety interlocks
- Operational flexibility:
  - AI-driven control
  - Human-in-the-loop supervision
  - Seamless switching between modes

# Future Work

- Deployment and validation on real beamline
- Incorporate uncertainty-aware control
- Automate the optimization of choosing the beam spot
- Validate/compare results using offline RL with real data rather than online RL with simulation

# Conclusion

- Autonomous beam tuning using RL
- Physics-informed simulation captures real effects
- Promising path toward reliable beamline automation
- A parallel AIOP effort to optimize target polarizations at JLab is wrapping up as well