Optimal Control of Polarized Sources and Targets

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

[1]Thomas Jefferson National Accelerator Facility
[2]The College of William & Mary
[3]Carnegie Mellon University

AI4EIC Workshop 2025
October 27, 2025

Jefferson Lab EPSCI WILLIAM & MARY CHARTERED 1693 Carnegie Mellon University

Motivation

- Polarized sources and targets are complex, dynamical systems.
- Experiments at Jefferson Lab rely on human operators to control and optimize the performance of these systems.
- AI/ML control algorithms could exceed the performance of human operators.
- Potential benefits both in terms of improved statistics and cost savings.

AI-Optimized Polarization (AIOP)

- AIOP is a 2-year, DOE-funded project that began in March 2024.
- An initiative of the Experimental Physics Software and Computing Infrastructure (EPSCI) group at Jefferson Lab.
- Consists of two sub-projects:
  a. Polarized Photon Beam;
  b. Polarized Cryogenic Target.

The GlueX Experiment

- The GlueX detector is located in Jefferson Lab’s Hall D.
- GlueX uses a polarized photon beam to search for and measure exotic hybrid mesons predicted from lattice QCD.

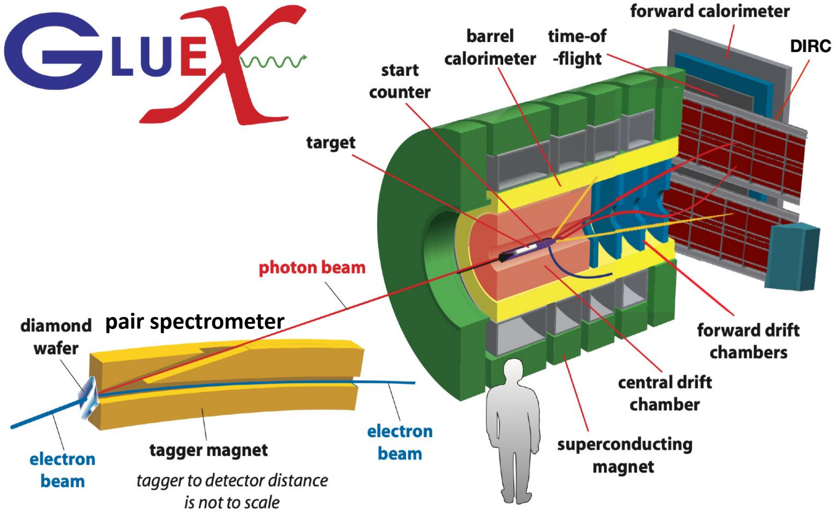

Surrogate Model

- Can we train a **data-driven surrogate model**—a reduced-order representation built directly from measured or collected data—to emulate the behavior of the underlying physical system?
- Simulation was developed based on the GlueX HDGeant4 simulation (R. Jones, D. Lawrence, et al).
  - AIOP Photon simulation allows for changing of the goniometer pitch and yaw angles, and the electron beam pitch and yaw angles.
  - Included degradation model that shifts and broadens the cobrem peak as a function of dose.
- Gaussian Process Regression is a flexible, non-parametric approach to regression, which allows for uncertainty quantification in predictions.
  - GP kernel used a Radial Basis Function + White Noise.
- Inputs for GP surrogate model include both the goniometer pitch and yaw and beam pitch and yaw.

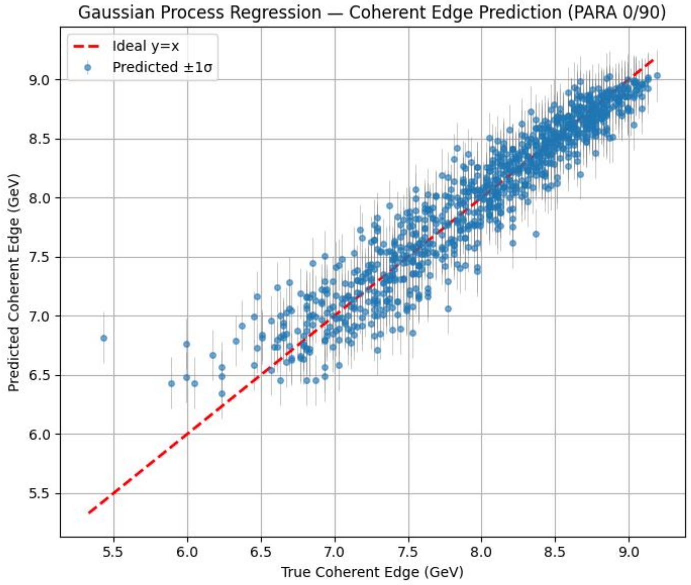

GlueX Photon Beam

- The interaction of the CEBAF electron beam with a thin (20-60μm) diamond radiator produces a linearly polarized photon beam via coherent bremsstrahlung radiation.
- The position of the primary peak (E_γ) is determined by the orientation of the diamond with respect to the electron beam.
- Data-taking runs alternate between PARA and PERP configurations, referring to the orientation of the polarization.

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

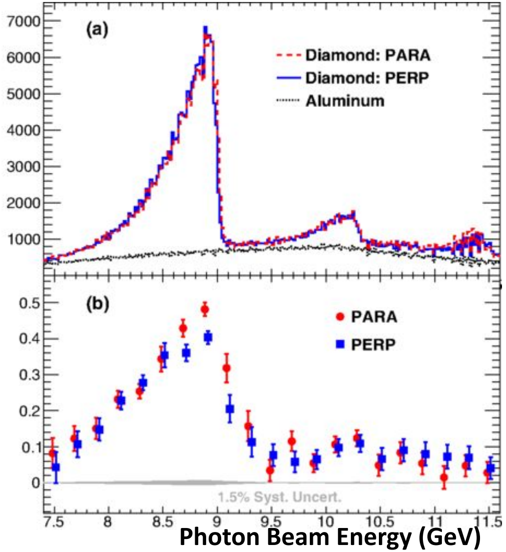

(a) Photon beam intensity versus energy as measured by the pair spectrometer (not corrected for instrumental acceptance)

(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.

GlueX Goniometer

- The diamond is mounted on a goniometer which can be rotated with respect to the x, y, and z axes of the lab frame (called pitch, yaw, and roll angles, respectively).
- The roll angle determines the polarization plane and is held constant.
- The pitch and yaw angles determine the location of the coherent bremsstrahlung peak and, if necessary, are adjusted at the start of a run.

20μm diamond

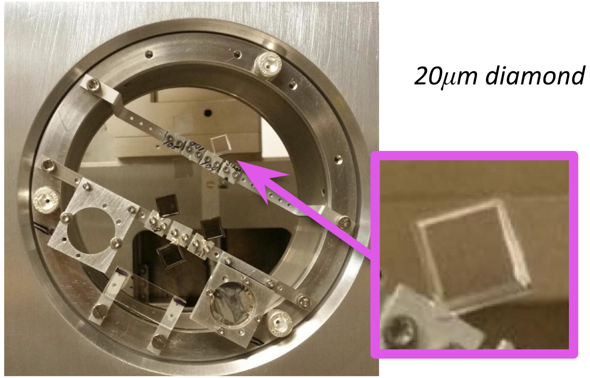

Radiators on goniometer

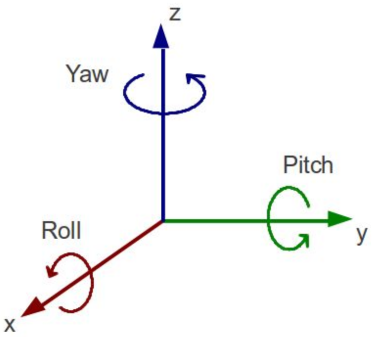

Standard Goniometer Operation

- Shift workers use a control GUI to align the coherent peak with respect to a reference.
- These buttons nudge the pitch and yaw value to shift the coherent edge to a higher or lower energy.

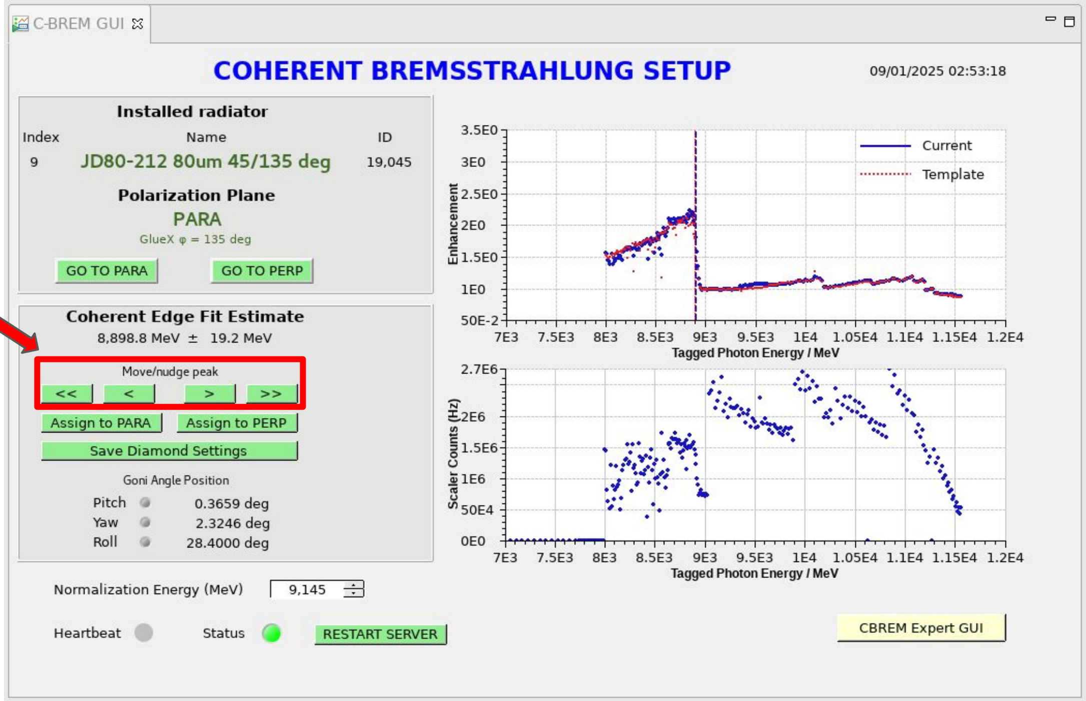

Goniometer Backlash

- Changing nudge directions causes lost motion due to mechanical backlash.
- Figure to the right shows a clear example of backlash from the Spring 2020 data: 3 down nudges + 3 up nudges leads to change in the energy.
- Control system must account by learning this backlash.

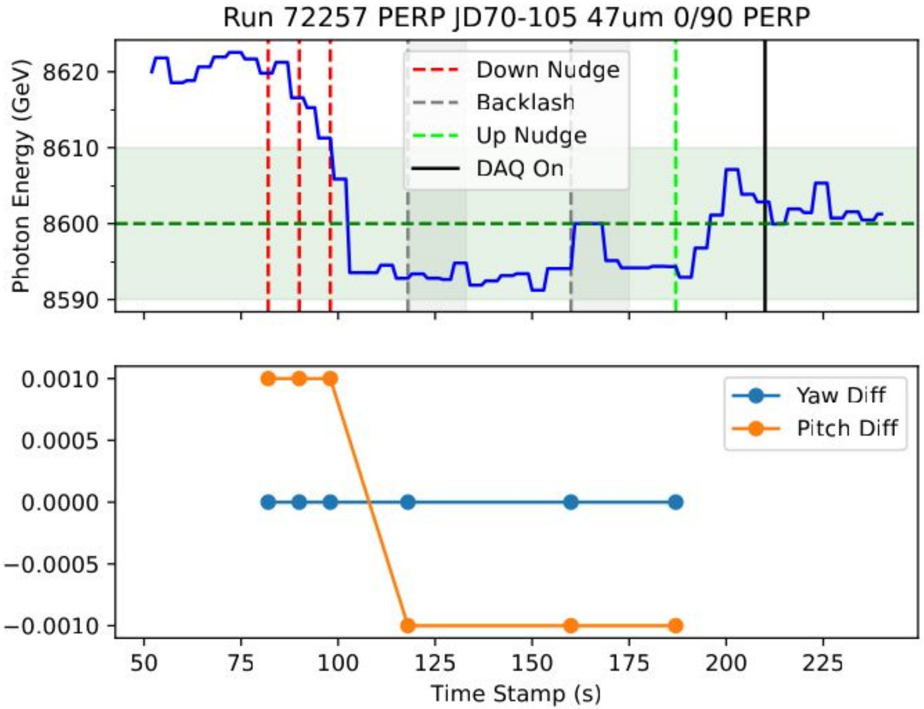

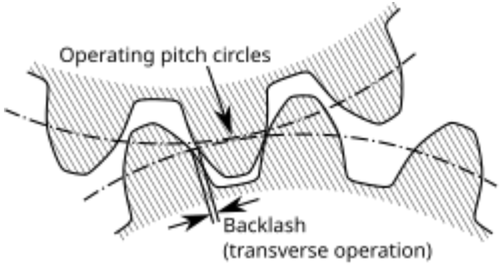

Beam Angle

- Fluctuations of the electron beam angle shift the coherent edge by changing the orientation of the crystal lattice with respect to the beam.
- Calculate the beam angle using beam position monitors upstream (5C11B) and downstream (Active Collimator) of the diamond.

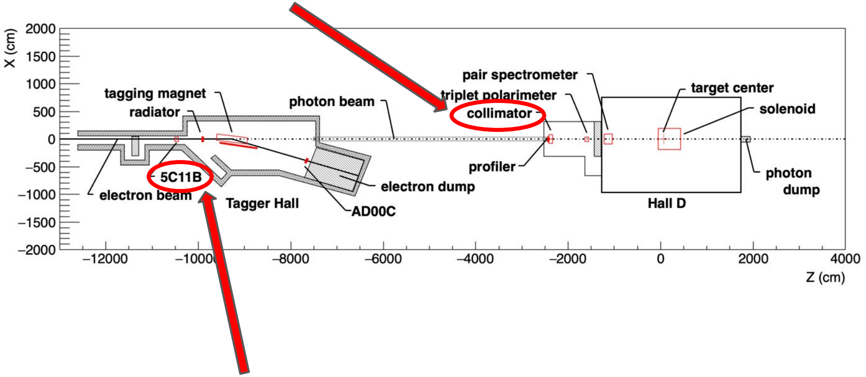

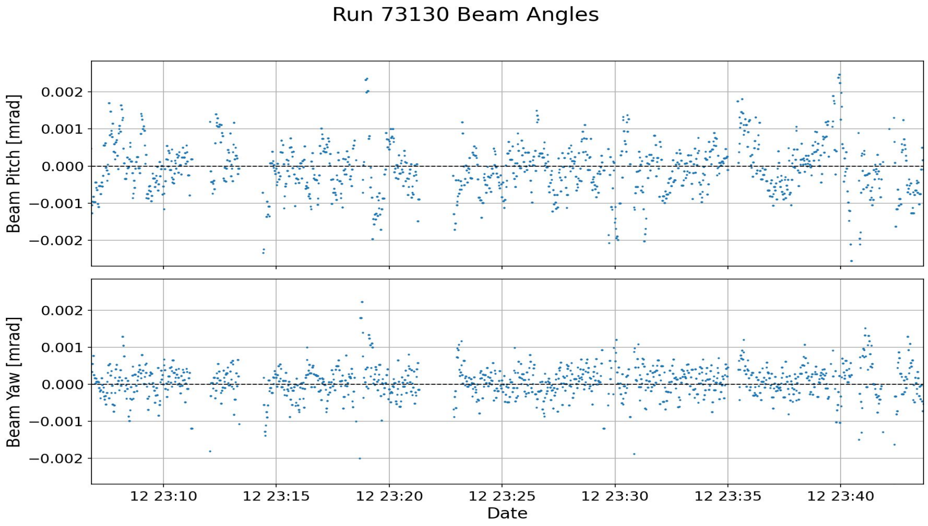

Diamond Degradation

- Radiation dose will lead to lattice orientation spread (or “mosaic spread”).
- This causes the coherent edge to shift and broaden.
- At GlueX, the width of the coherent edge increases from an estimated value of 46 ± 6 MeV for undamaged parts of the crystal to 270 ± 36 MeV damaged parts of the crystal for a dose of ~10^{19} electrons.
- Coherent edge shifts by a few MeV.

Edge widths for undamaged and damaged regions of diamond

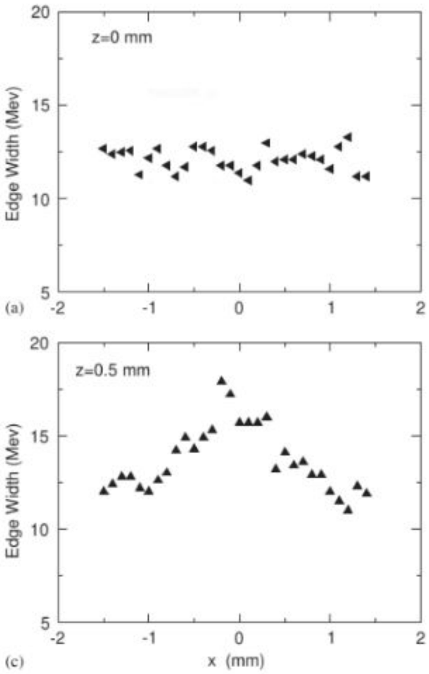

Shifting of coherent edge

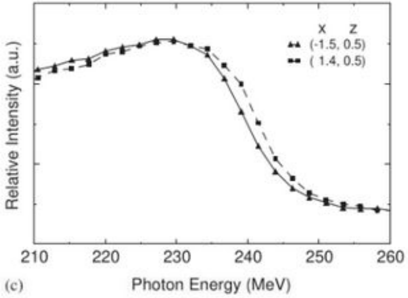

Source: J.D. Kellie et al. The selection and performance of diamond radiators used in coherent bremsstrahlung experiments. NIM-A 545 (2005) 164-180.

Ongoing Work

- Improve on diamond degradation model and incorporate it as input to the surrogate model.
- Incorporate backlash into simulation to be learned by control algorithm.
- Test control algorithms using surrogate model environment.
- Integrate with Hall D controls system.

Polarized Targets

- Polarized cryotargets are used at Jefferson Lab to study nuclear structure.
- Dynamic Nuclear Polarization (DNP) is used to transfer the high electron polarization to nearby nuclei.
- This polarization degrades with the accumulation of free radicals produced when the target is irradiated by the electron beam.

Sample container filled with ND_3

Beam spot on target material

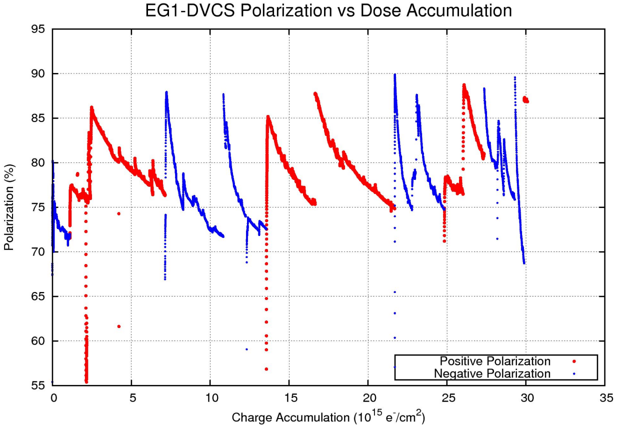

Shift crew operation

Polarization decay over time requires the shift crew to continuously adjust the microwave frequency. The average target polarization that can be maintained over the course of an experiment heavily depends on the shift crew's experience.

Microwave frequency controls

Polarization extraction

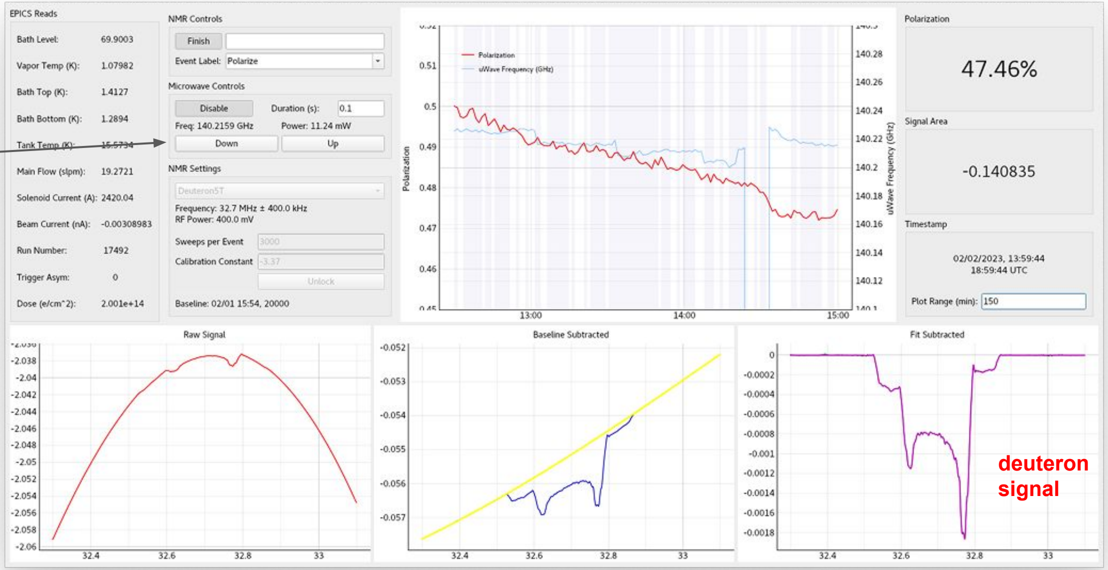

Python-based NMR monitoring and control GUI

Polarization extraction

- The target polarization is obtained by integrating the proton signal after the background has been sufficiently removed.

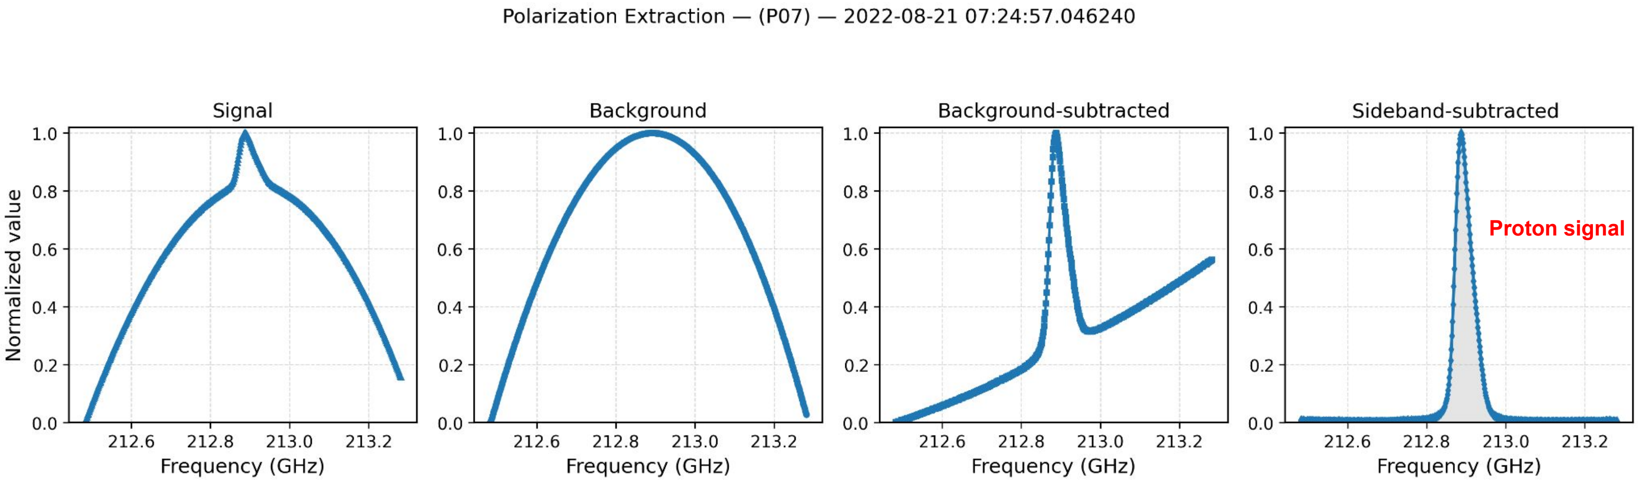

Polarization dynamics

The polarization increases when beam is off because the target cools down, and then degrades when beam is on due to a combination of beam heating and increased concentration of free radicals (e-, atomic H) in the sample

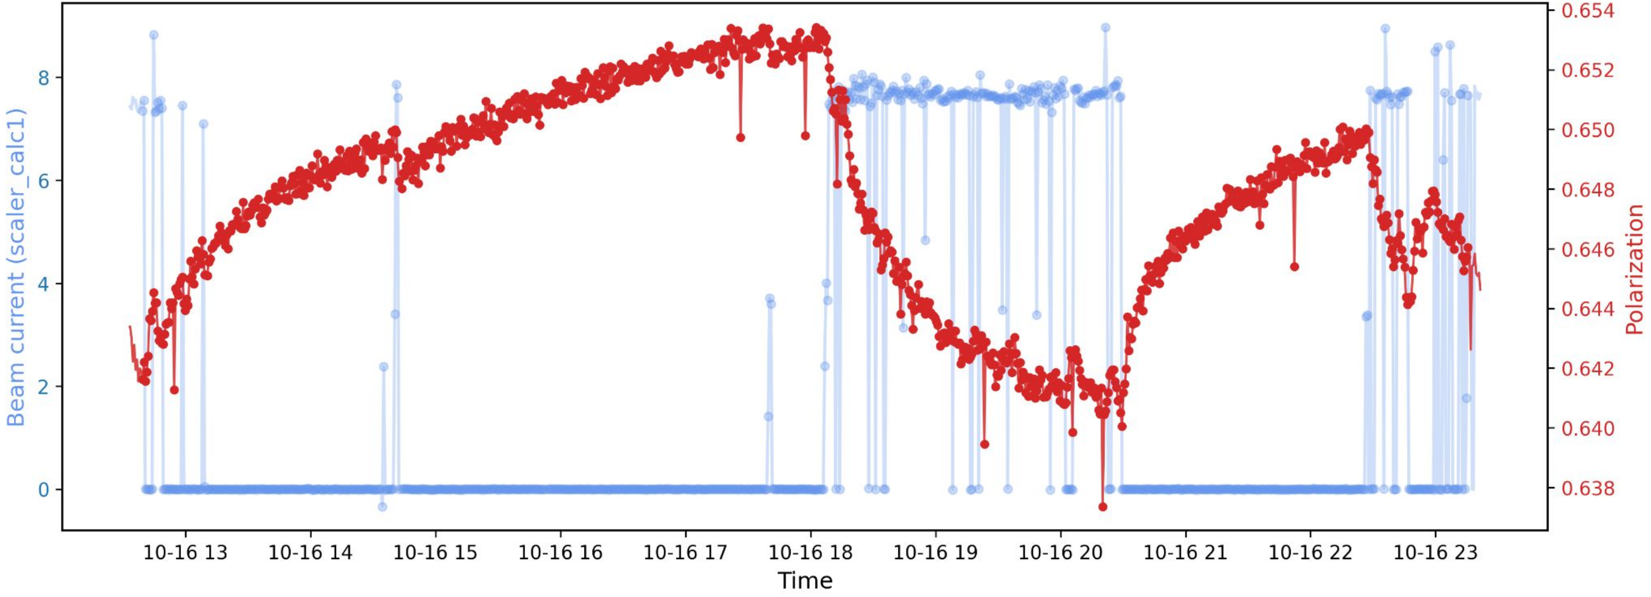

Dataset selection

- Sample “P07” is NH_3 material used during the Run Group C experiments in Hall B
- It includes a notable episode where raising the microwave frequency reduced the polarization, and lowering it afterward improved polarization.
- Microwave frequency, target dose per unit area, and beam current are used as input features to predict the target polarization.

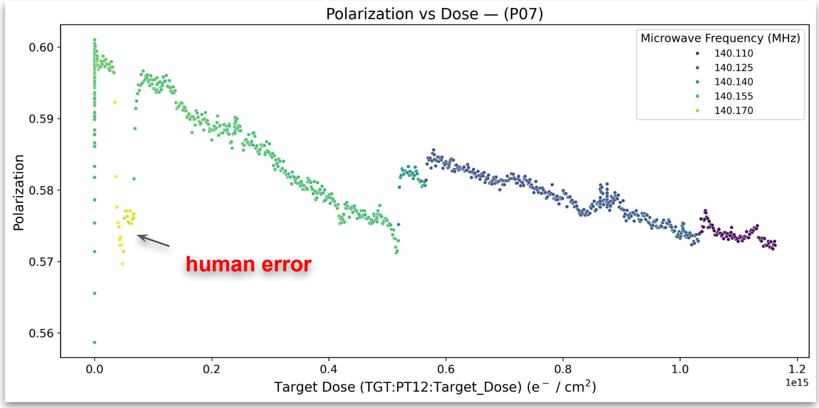

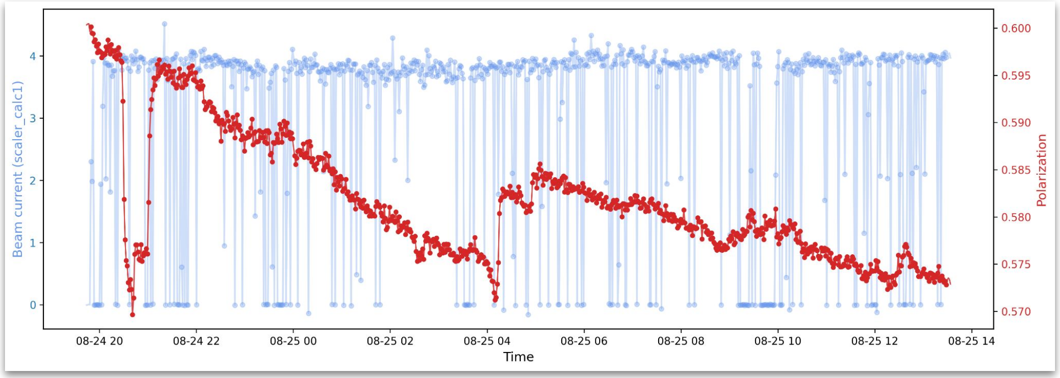

Surrogate Models

- Since real polarized target experiments are not currently running, building a surrogate model provides a practical way to simulate and analyze system behavior virtually.
  - The surrogate enables data-driven analysis and prediction without requiring physical runs.
- This work uses two surrogate models: Gaussian Processes (GP) and Multi-Layer Perceptron (MLP).
- Both models use Negative Log-Likelihood (NLL) loss to predict values and uncertainties.

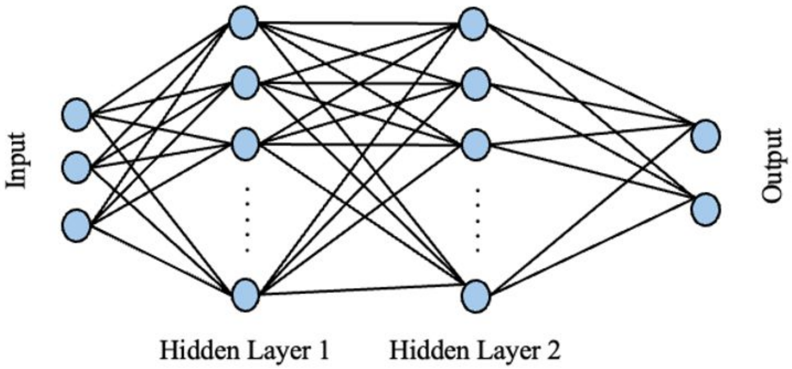

multi-layer perceptron

Surrogate Model (MLP)

MLP architecture with 3 input features, 2 hidden layers, and two outputs

- **Learning rate: 0.0001**
- **Optimizer: Adam**
- **Loss: Negative Log-likelihood**
- **Early stopping: 25**
- **Batch size: 32**

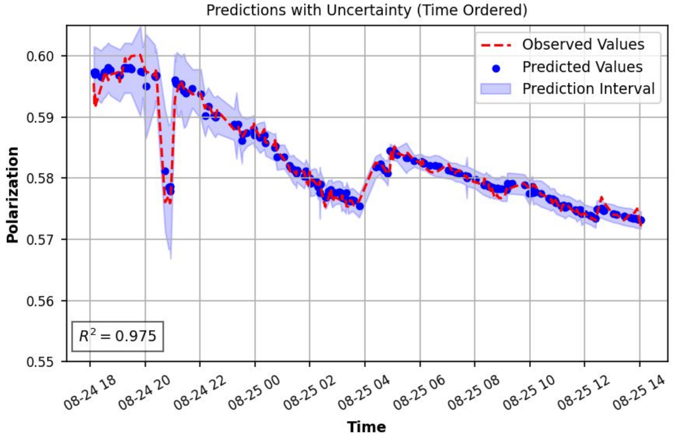

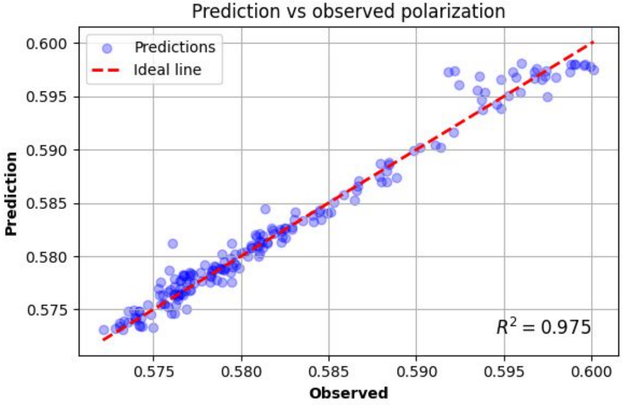

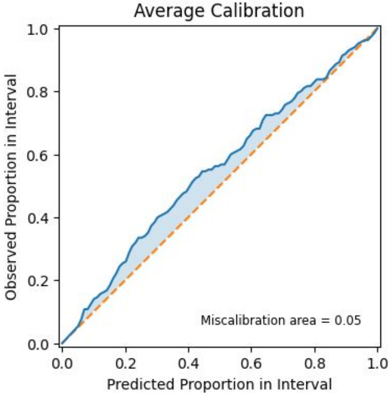

Surrogate Model (GP)

- GP hyperparameters (length scale, coefficient, and noise) are optimized by minimizing NLL loss.
- Optimization is performed for 3 kernels (Laplace, RBF, Cauchy) using 5 optimization methods:
  - Powell, Nelder-Mead, L-BFGS-B, CG, and COBYQA
- Laplace kernel consistently achieved lower NLL loss than RBF and Cauchy.

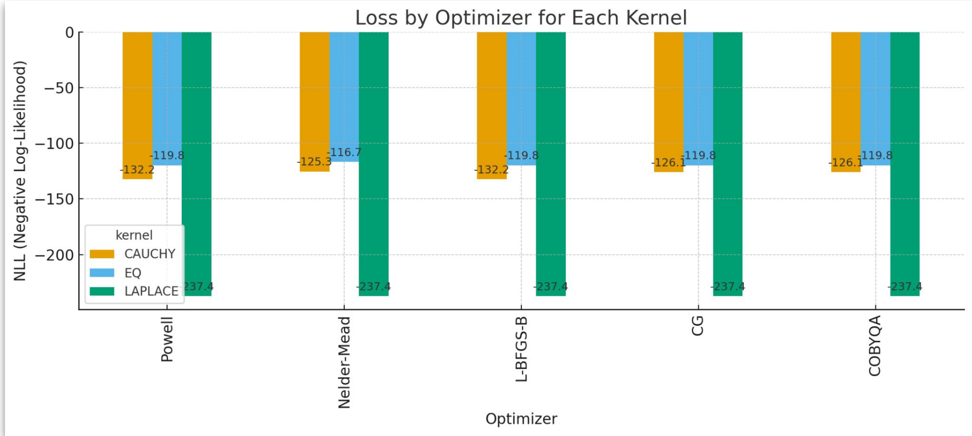

Surrogate Model (GP)

- GP with Laplace kernel and optimized with L-BFGS-B algorithm gives
  - Higher R^2 values compared to MLP → predictions are closer to the actual observed values.
  - Lower miscalibration area → predicted uncertainties aligning more closely to observed outcomes
  - Narrower uncertainty ranges → more precise and reliable estimates

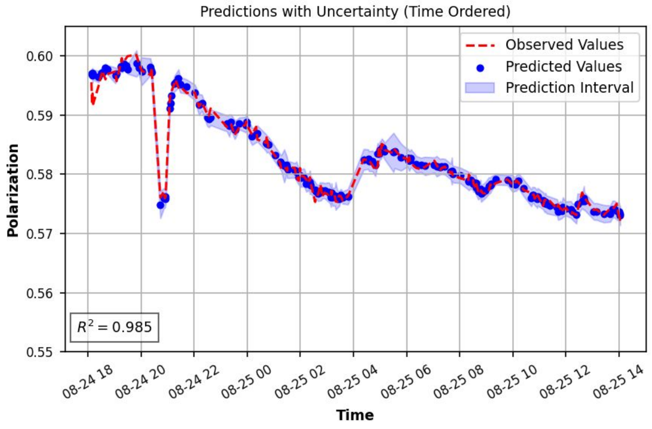

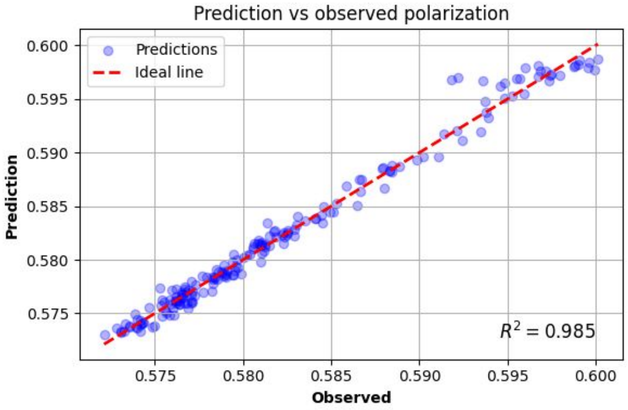

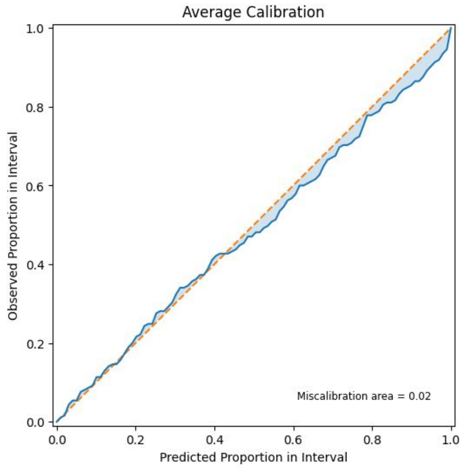

UQ Surrogate Model Manual Tuning Demo

- Interactive demonstration of GP-predicted polarization response to microwave frequency adjustments.
  - Here, we are adjusting the microwave frequency manually as the shift crew would, except the polarization and uncertainty is determined from the GP

- Polarization uncertainty increases when the microwave frequency is out-of-domain

- We have an environment to find the optimal frequency selection given the conditions.

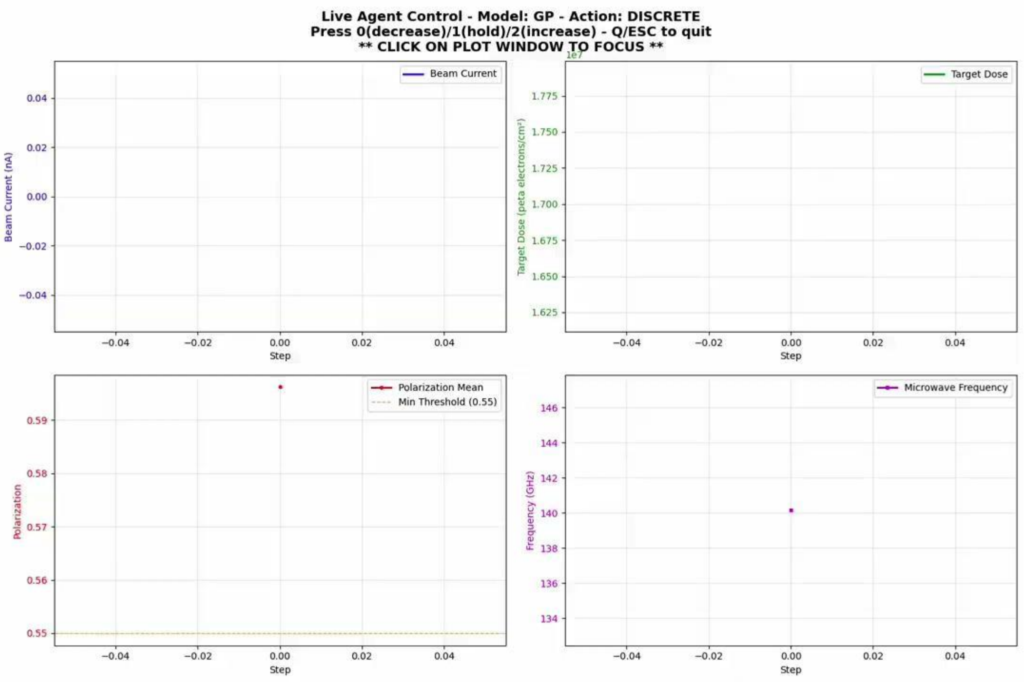

Top Left: Beam Current (nA), Top Right: Target Dose
Bottom Left: Polarization, Bottom Right: Microwave Frequency

Incorporating surrogate model with Gymnasium and SciOpt Control Toolkit

- *Scientific Optimization Control Toolkit (SOCT)* is a modular, Gymnasium-compatible framework for building, training, and deploying control agents with pluggable environments, live TensorBoard monitoring, and reproducible configs/models.

- Gymnasium Wrapper:
  - UQ surrogate models are “wrapped” using the Gymnasium package to integrate into the SOCT workflow
  - With both simulations integrated into the environment, we can train and test multiple control algorithms and compare their learned policies under different simulations and uncertainties.

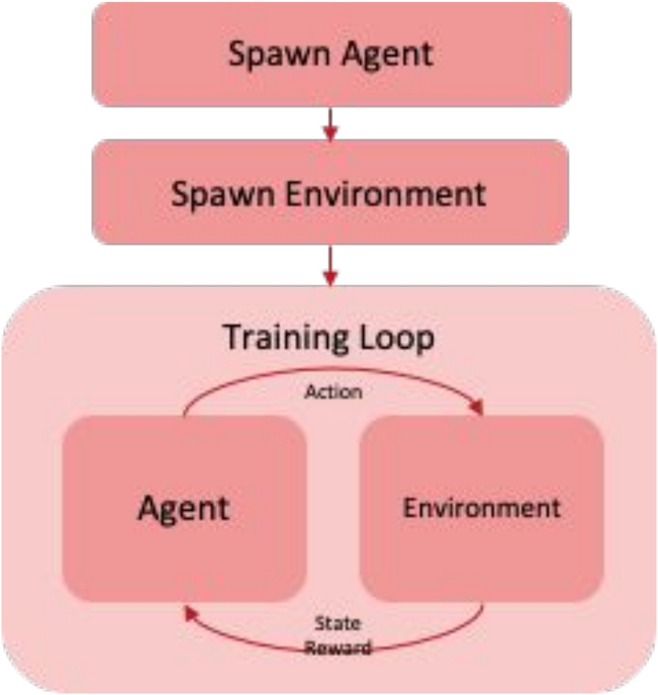

SOCT Schematic

Ongoing work

- Use existing SOCT algorithms for polarization tuning: Bayesian Optimization (BO), Twin Delayed DDPG (TD3), and Soft Actor-Critic (SAC).
- Incorporating more sample in surrogate model training
- Implementation with existing NMR software
  - Dual view of polarization obtained from traditional method and predicted polarization from surrogate model
  - Ability to run both methods simultaneously
  - Toggle human-in-the-loop vs automated control
- Use statistical and ML methods to improve background subtraction

Summary

- AIOP seeks to optimize nuclear physics measurements using AI/ML for experimental control.
- Working to integrate surrogate models into control algorithms for both polarized source and target subprojects.
- Plan to integrate with the JLab experimental hall controls system in 2025-26.
- Results could help lay the foundation for future autonomous experiments at other facilities (e.g. EIC).

U.S. DEPARTMENT of ENERGY

Awardee

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Award Number: DE-FOA-0002875 - GRANT13779668

Backup Slides

Goniometer Offsets