Dynamic Control of Detectors, Polarized Beams, and Polarized Targets Using AI/ML

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

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

JLab hosts the CEBAF electron accelerator which has a focus on understanding the structure of protons, neutrons, and nuclei in terms of quarks and gluons.

- 12 GeV electrons
- CW beam (bunch every 4ns)
- Fixed Target NP Experiments

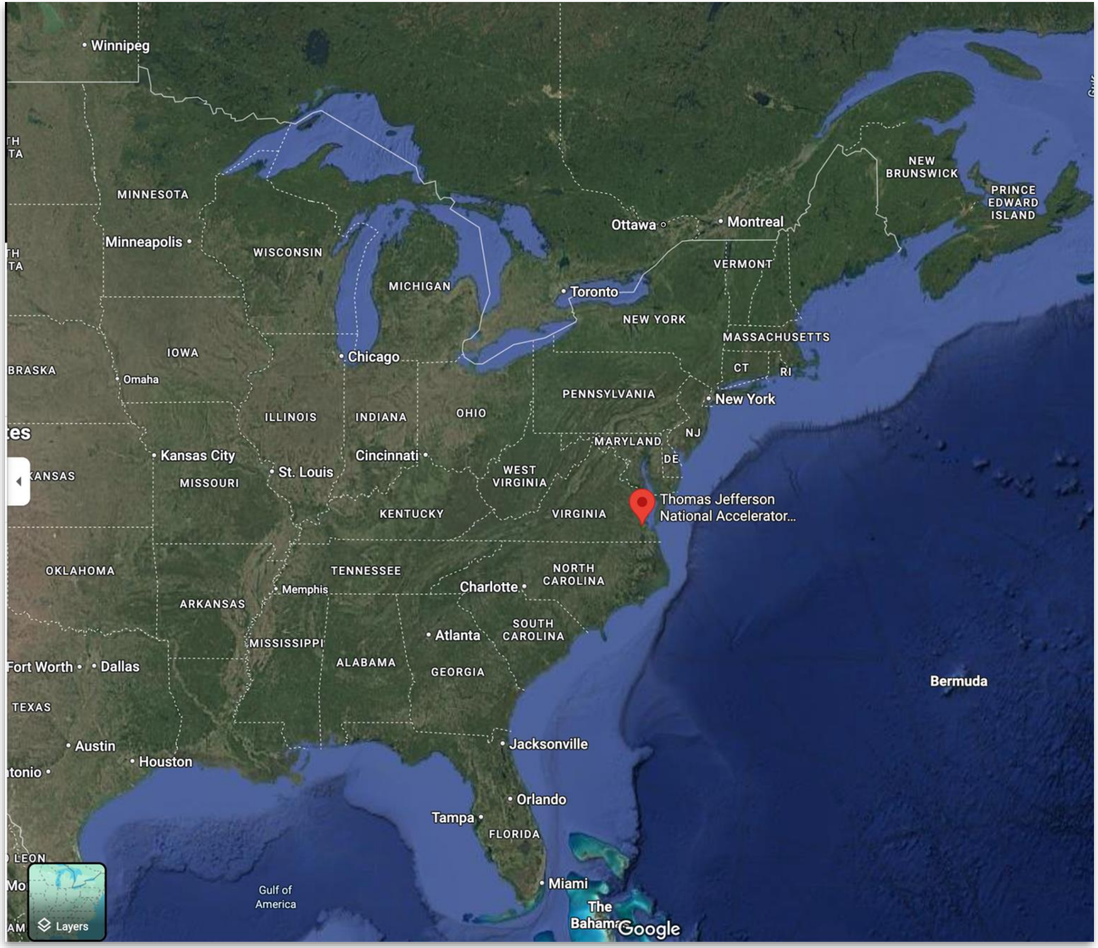

The GlueX Detector

[GlueX detector]() located in Hall D at Jefferson Lab

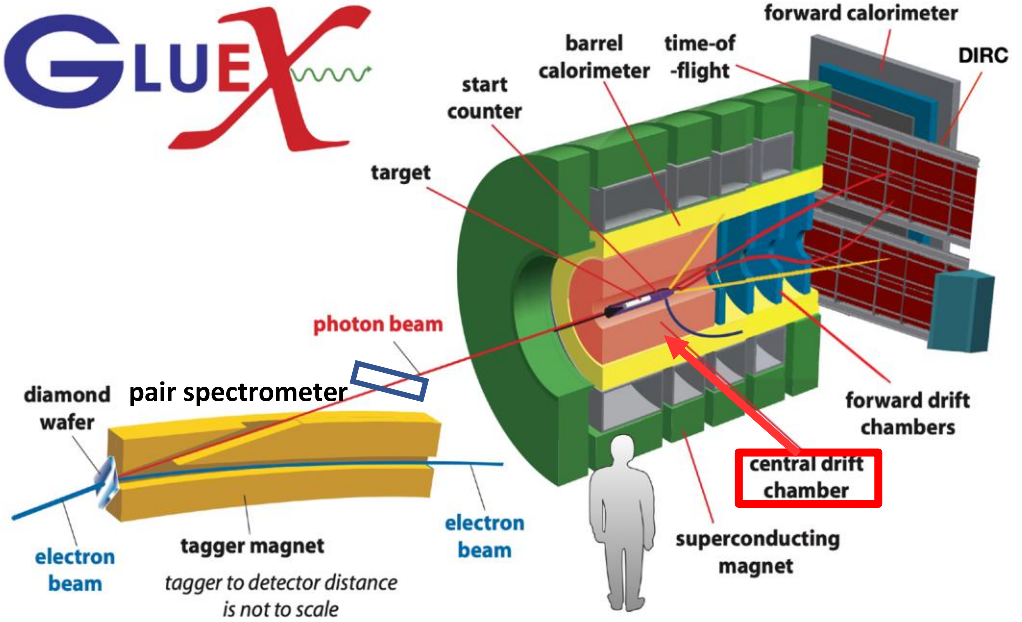

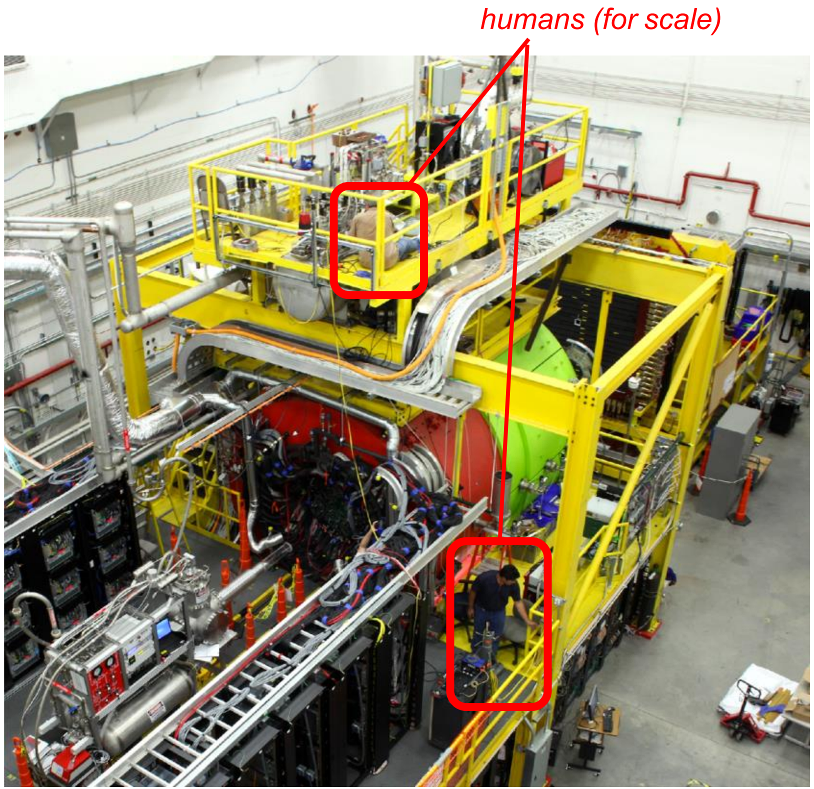

System is part of standard production data taking

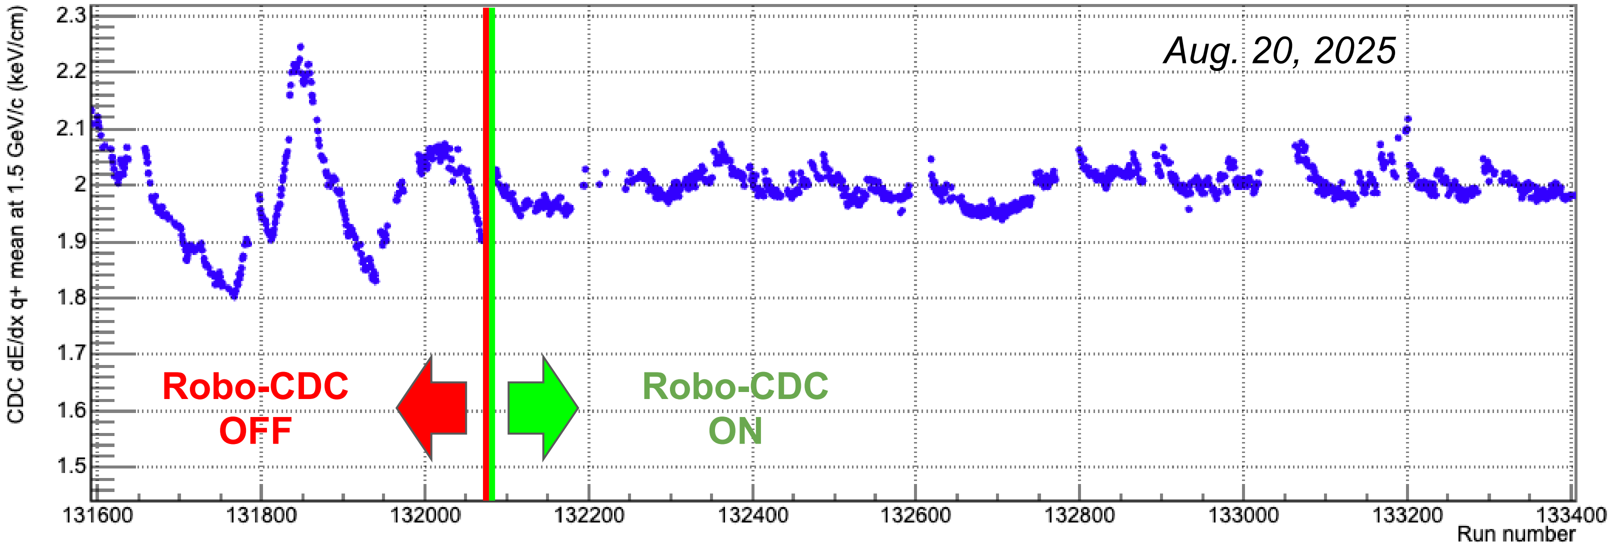

Polarization at Jefferson Lab

Polarized cryotargets are used throughout Jefferson Lab to study nuclear spin structure

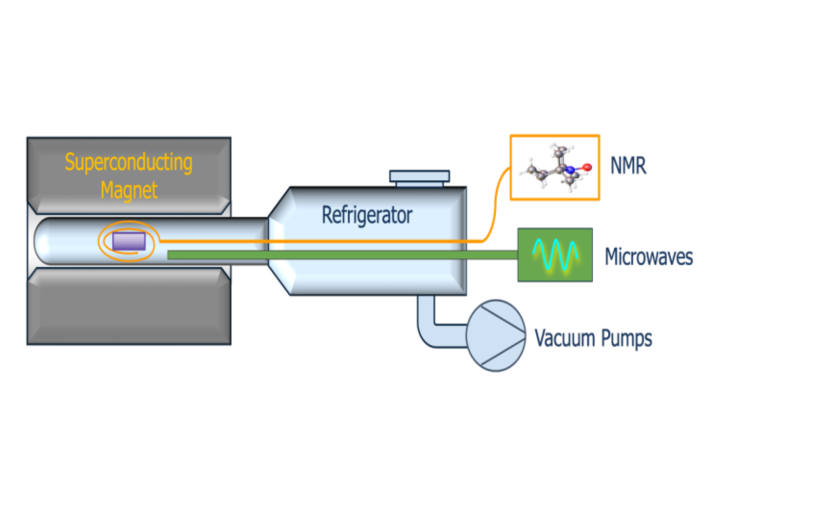

Polarized targets

Hall-D uses a polarized photon beam to search for and measure exotic hybrid mesons

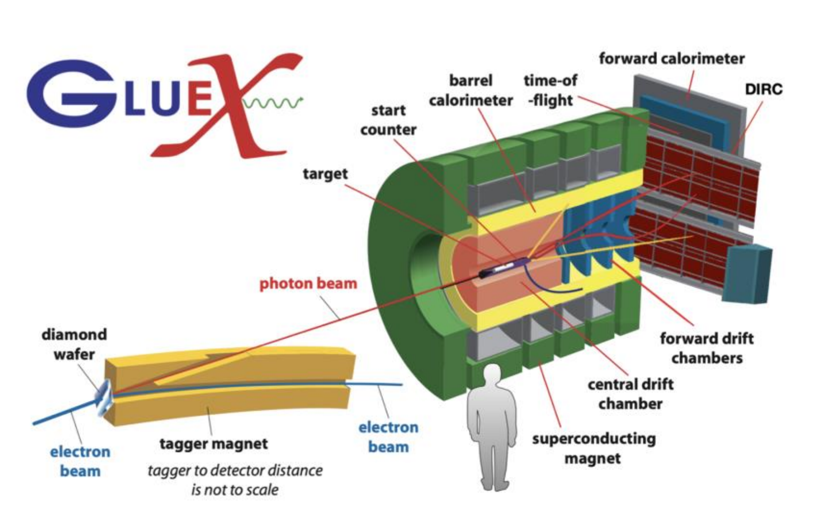

Polarized photon beam

AI OPTIMIZED POLARIZATION

Polarized Photon Beam

Coherent Bremsstrahlung

Polarization indicates fraction of high energy photons with electric field pointing in the same direction

Linearly Polarized Photons are generated by passing the electron beam through a diamond which has regular molecular structure

Enhancements in the bremsstrahlung energy spectrum are correlated with the angle of the produced photon

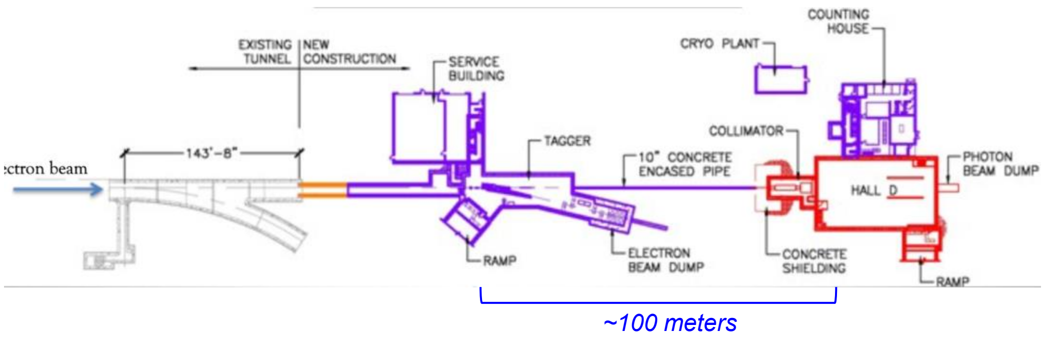

Polarized Photon Beam

Diamond Radiator

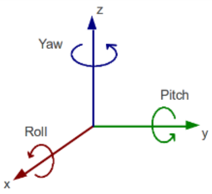

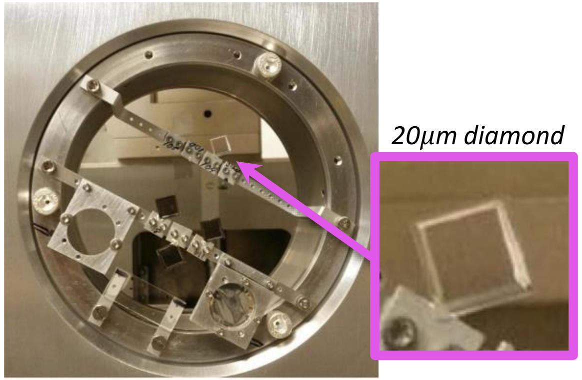

Radiators on goniometer

	e^- beam	Coherent peak
GlueX	11.7 GeV	9 GeV
CPP	11.7 GeV	6 GeV

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

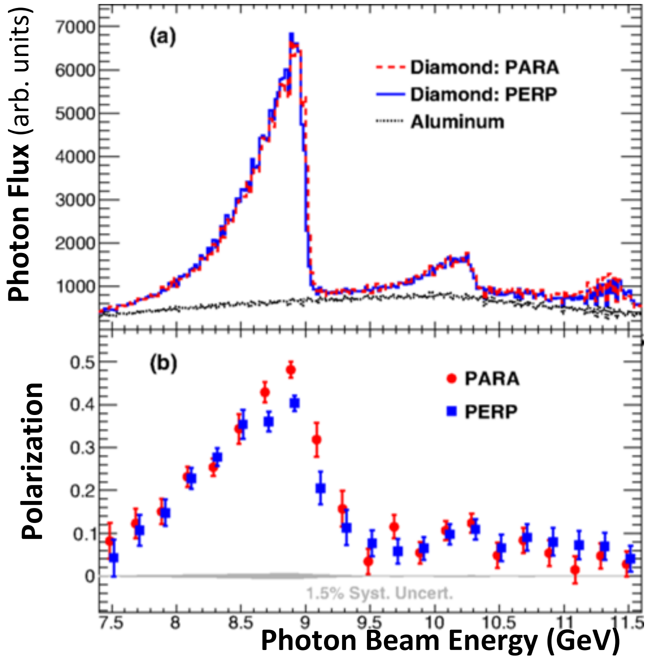

(a) Photon beam intensity versus energy as measured by the pair spectrometer (not corrected for instrumental acceptance)

(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.

AIOP
AI OPTIMIZED POLARIZATION

Jefferson Lab

Polarized Photon Beam

Standard operation

Experts determine the ideal coherent edge position for each run period.

The shift crew is responsible for maintaining this position throughout data taking by “nudging” the orientation of the diamond via pitch and yaw angles.

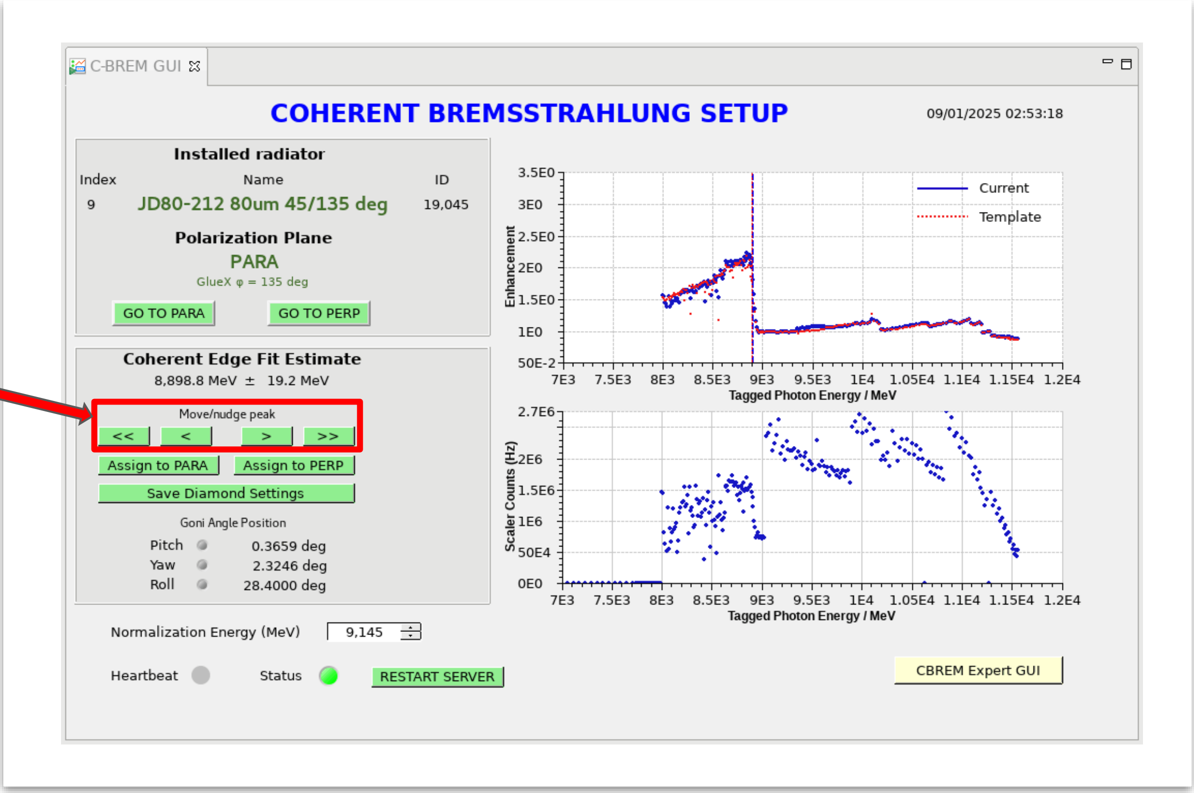

Polarized Photon Beam

Effect 1: 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.

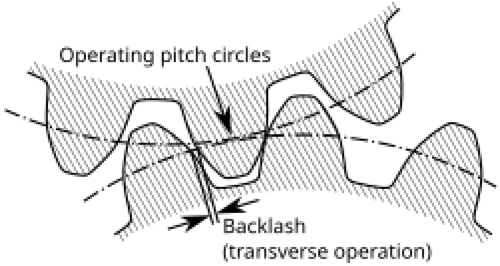

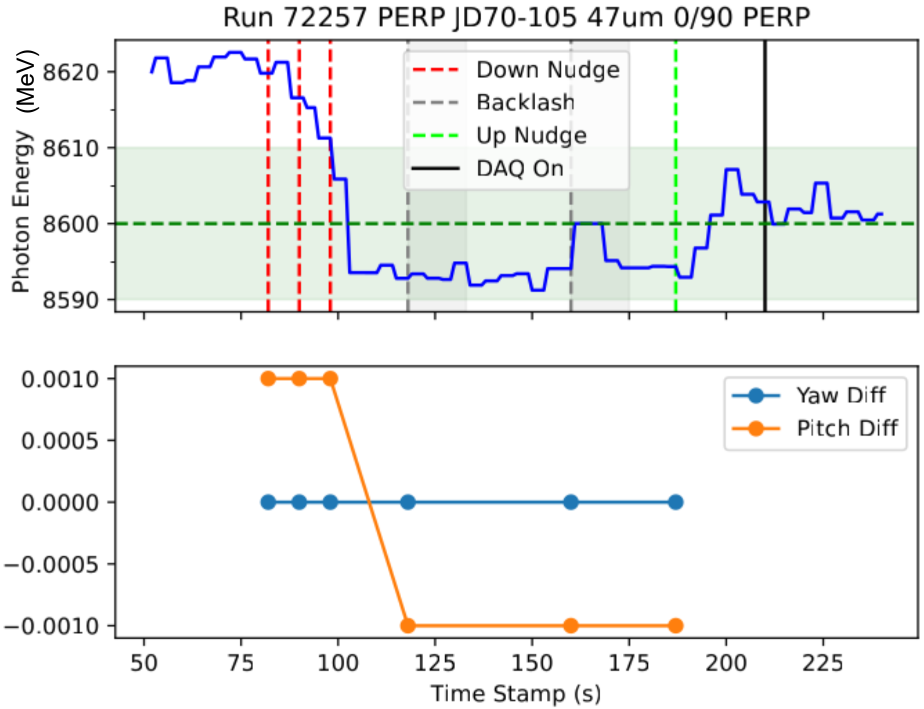

Polarized Photon Beam

Effect 2: 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.

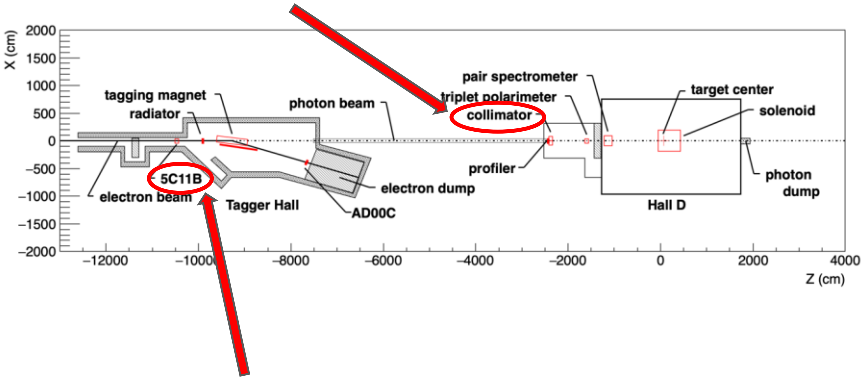

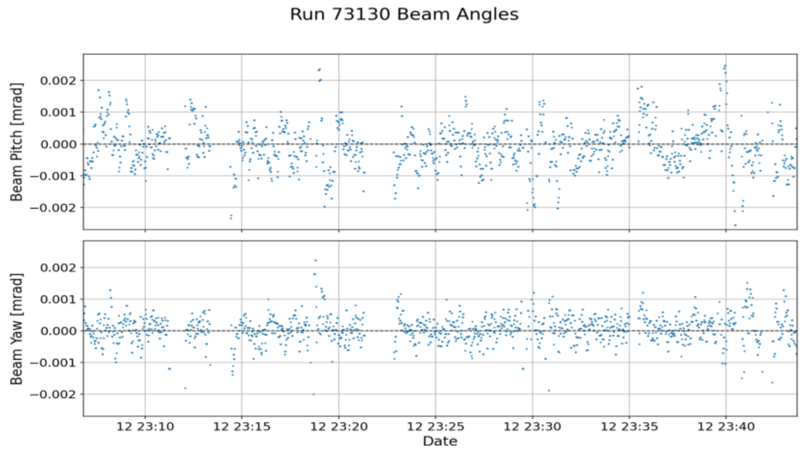

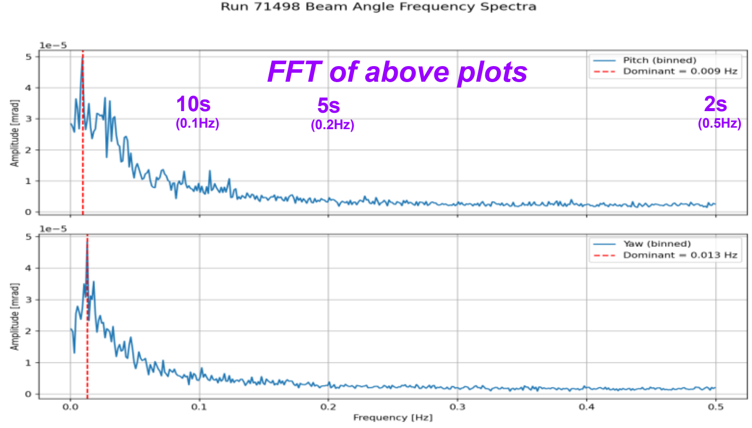

AIOP
AI OPTIMIZED POLARIZATION

Jefferson Lab

Polarized Photon Beam

Effect 3: 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** for **damaged** parts of the crystal for a dose of ~10^{19} electrons.

Edge widths for undamaged and damaged regions of diamond

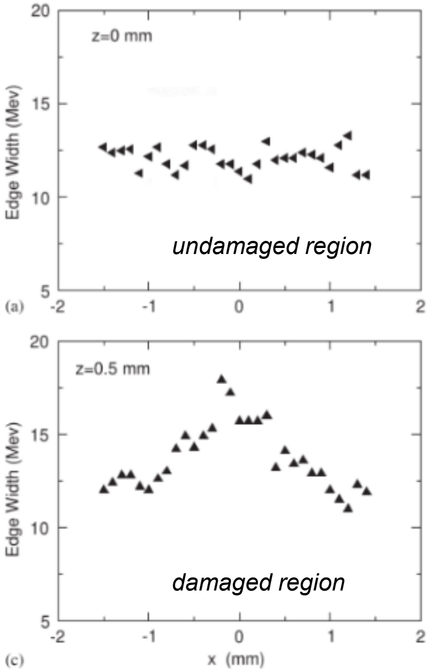

Shifting of coherent edge

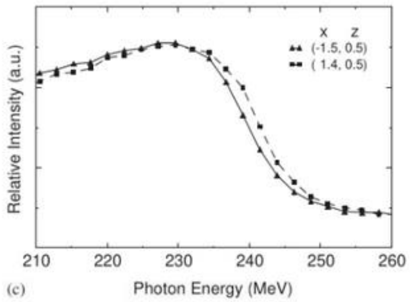

Source: J.D. Kellie et al. The selection and performance of diamond radiators used in coherent bremsstrahlung experiments. NIM-A 545 (2005) 164-180.

AIOP
AI OPTIMIZED POLARIZATION

Polarized Photon Beam

Simulated Beam Environment

Reinforcement Learning requires a simulation and environment on which to train the actor model.

Coherent Edge calculated from Monte Carlo driven beamline effects

AI OPTIMIZED POLARIZATION

Polarized Photon Beam

Beam Spot Finder with Bayesian Optimization

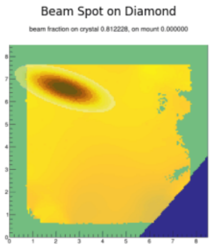

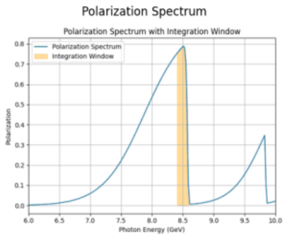

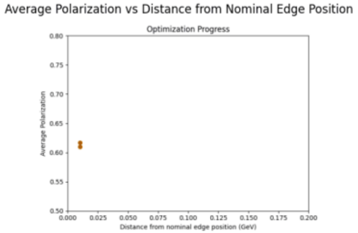

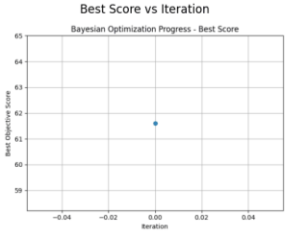

A separate, related project is to use BO to identify the optimal positions to center the beam in order to maximize the overall polarization for data taken over the lifetime of a diamond*.

$$score = \frac{P_{avg}}{|E_c - E_{nom}| + \epsilon}$$

**This specific part of the project was started, but is currently on hold while we focus on the automated goniometer tuning.*

Polarized Target

Dynamic Nuclear Polarization

The optimal polarization decreases due to the electron beam creating additional radicals, requiring further adjustments from the shift crew.

Target samples are warmed up (annealed) to ~100k to remove unwanted impurities, with eventual replacement after 5-10 anneals

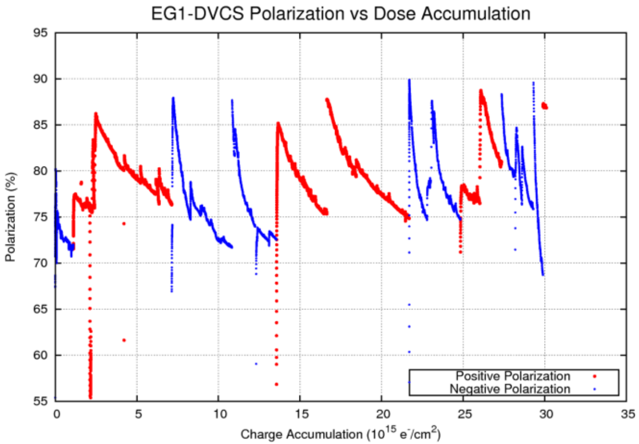

Sample container filled with ND3

Beam spot on target

AI OPTIMIZED POLARIZATION

Polarized Target

Standard Operation

Shift workers adjust the microwave frequency as needed to maintain the target polarization.

The experience of the shift workers significantly influences the average target polarization maintained throughout an experiment.

Polarization extraction

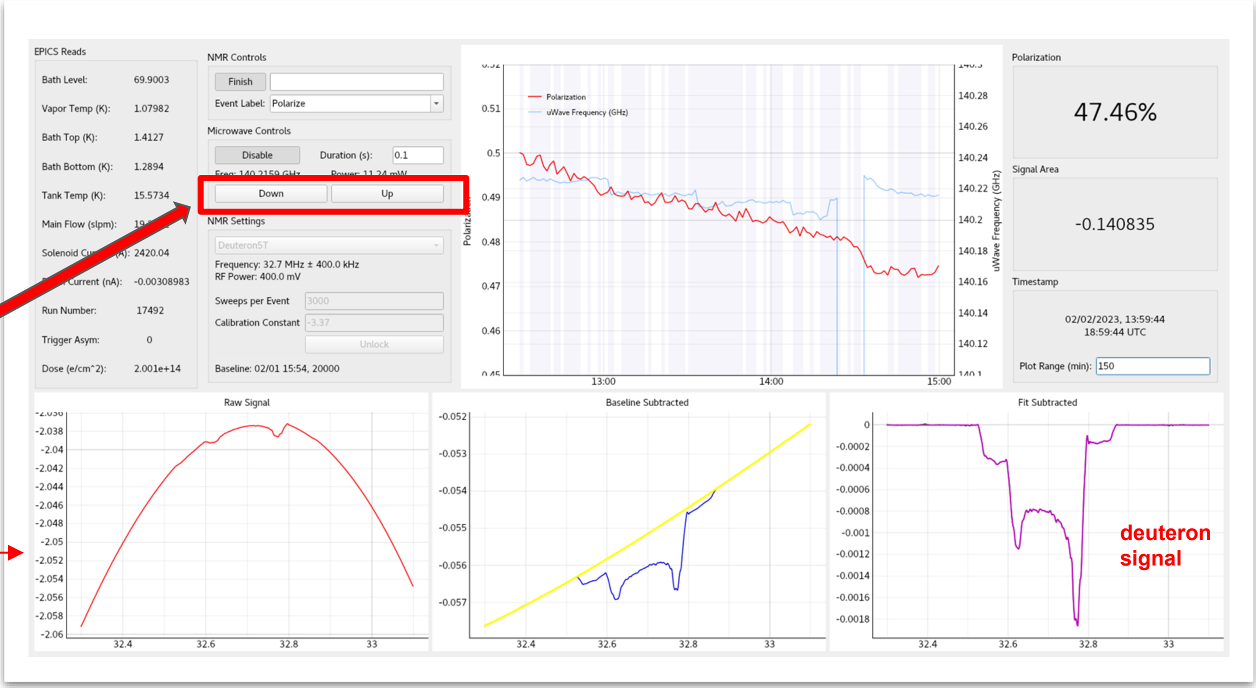

Polarized Target

Polarization extraction

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

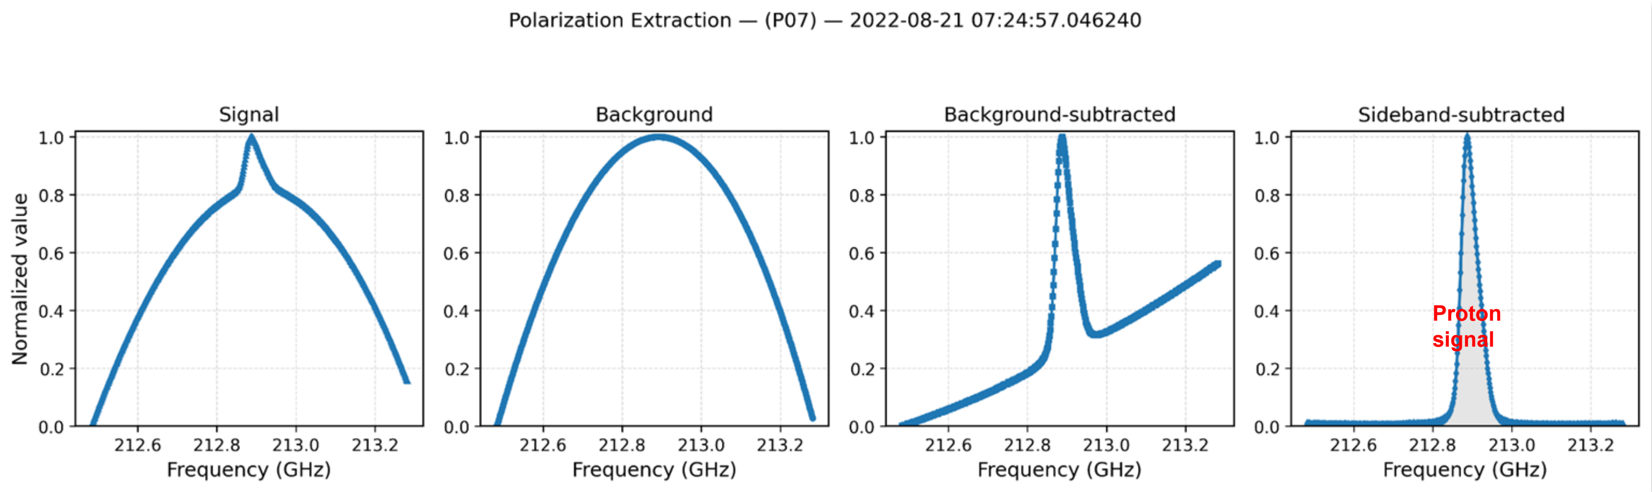

Polarized Target

Measured Polarization and Relevant Features

Input Features:

- Beam Current
- Target Dose
- Microwave Frequency

Target Value:

- Optimal Polarization
  (now and future)

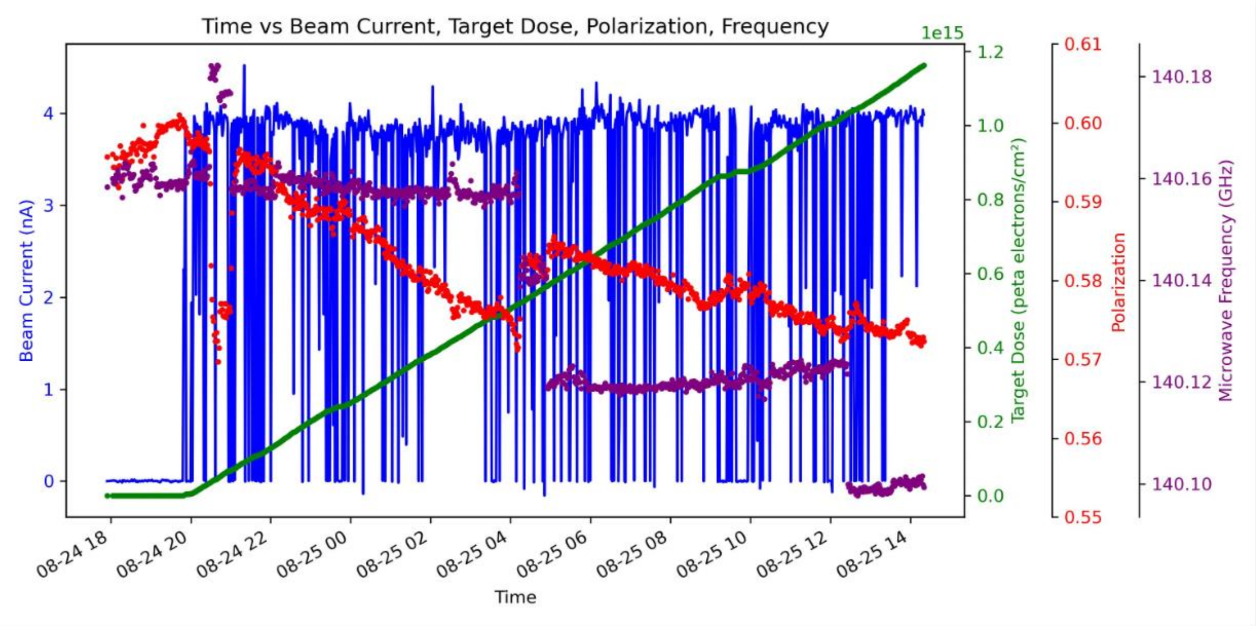

Polarized Target

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

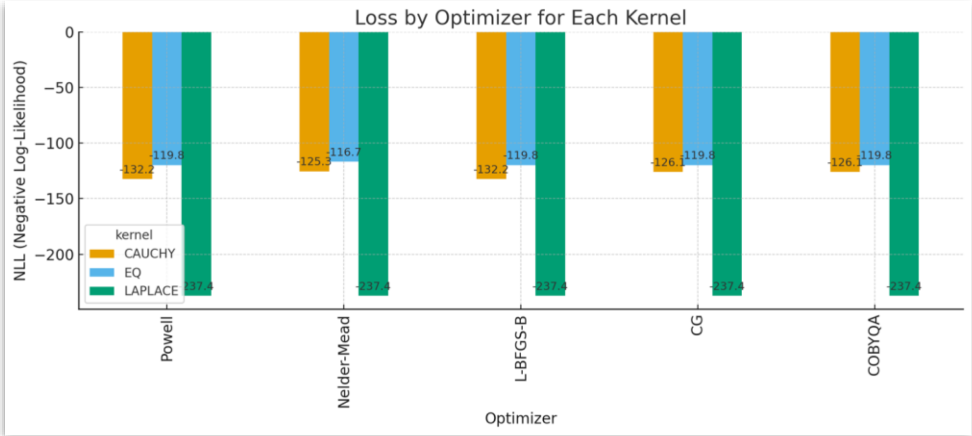

Laplace kernel consistently achieved lower NLL loss than RBF and Cauchy.

AI OPTIMIZED POLARIZATION

Polarized Target

Gaussian Process Approximation

Input data is converted into **Random Fourier Features (RFF)** for simplification.

Random frequencies are generated based on the selected kernel type.

RBF, **Laplace**, and **Cauchy kernels** are used to sample frequencies according to their respective distributions in the RFF method.

Hyperparameters (length scale, coefficient, and noise) are **optimized** based on the RFF for the corresponding kernel.

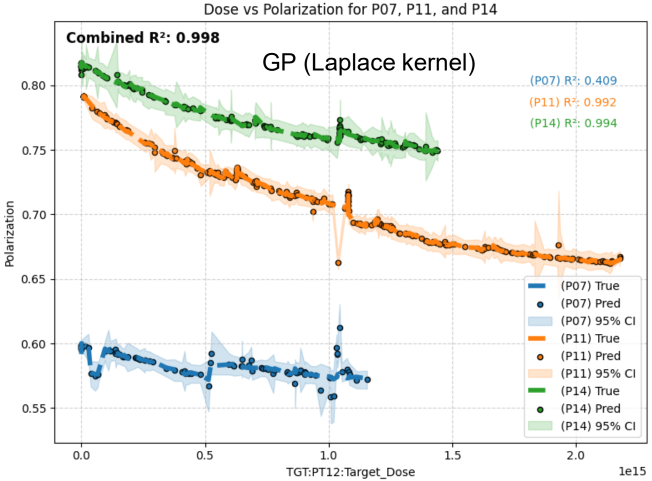

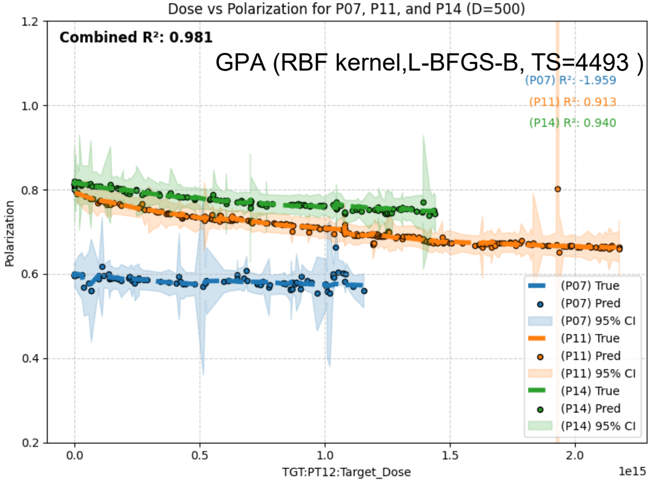

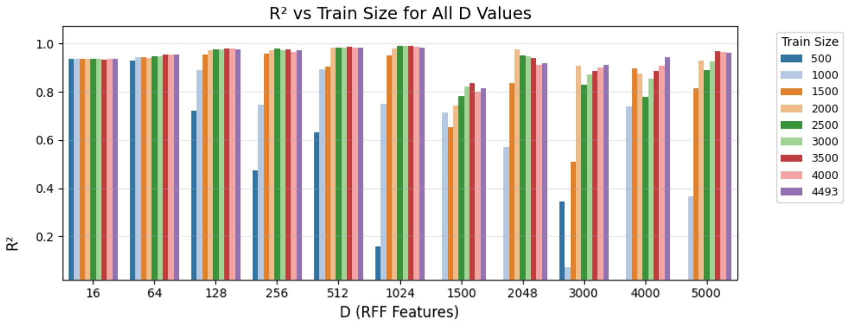

UQ Surrogate Model Manual Tuning Demo

- 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.

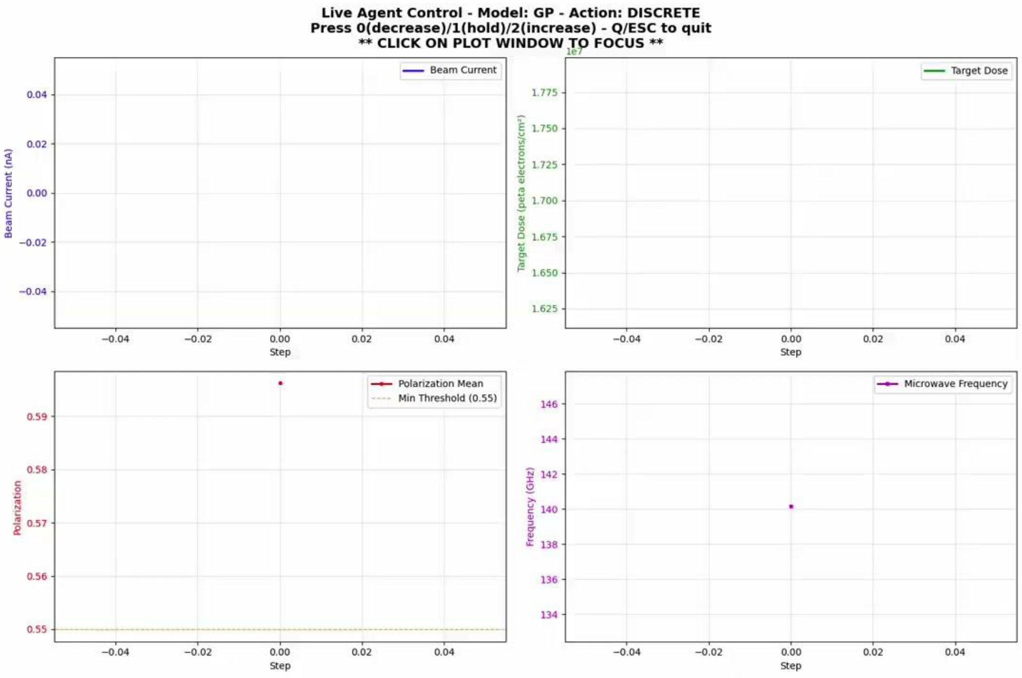

Top Left: Beam Current (nA), Top Right: Target Dose
Bottom Left: Polarization, Bottom Right: Microwave Frequency

Polarized Target

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

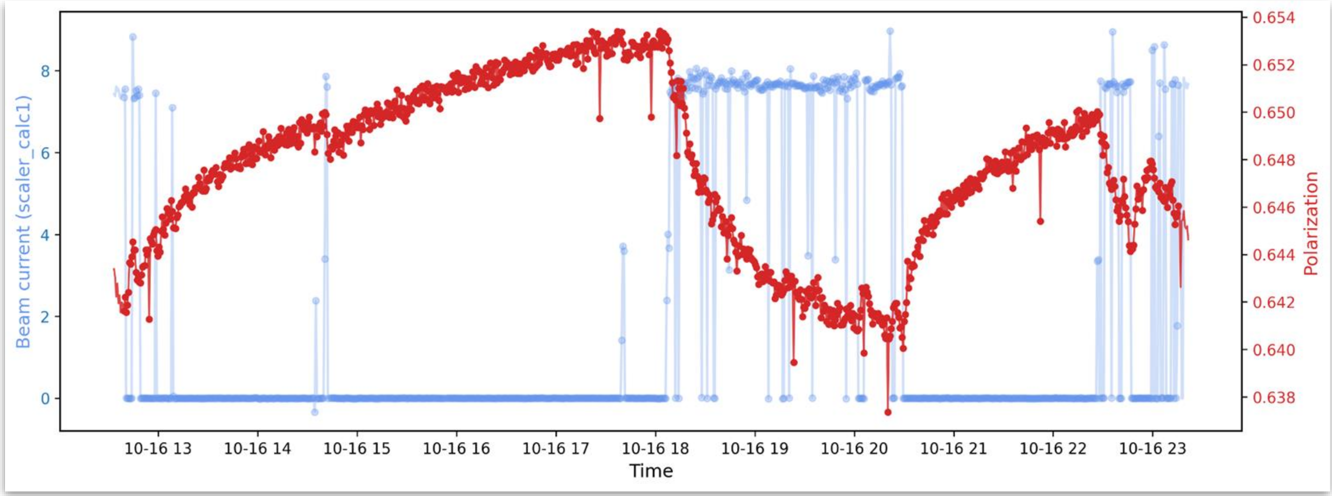

Incorporating surrogate model with Gymnasium and SciOpt Control Toolkit

- *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.

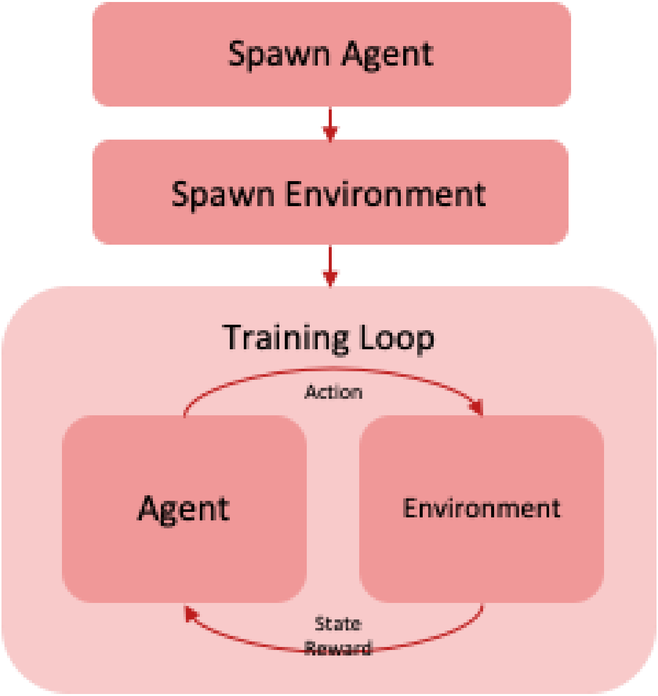

SOCT Schematic

https://github.com/JeffersonLab/SciOptControlToolkit

Budget, Deliverables, and Schedule

	Year 1 (FY24)	Year 2 (FY25)	Year 3 (FY26)	Totals
a) Funds allocated	$218.7k	$546.3k	$535.1k	$1,300.0k
b) Actual costs to date	$218.7k	$546.3k	$153.8k	$918.7k

Total available at end of Oct: $381.3k

Current spending and schedule projection has us about 1 FYQ behind. This may be affected by potential government shutdown at end of Jan. 2026. May need to submit NCE to July 2026, but will know more in Jan.

AIOP
AI OPTIMIZED POLARIZATION

Major Deliverables and Schedule

FY26Q1	Finalize environment simulation for Polarized Target/Beam System
FY26Q2	Integrate simulation with SOCTS system to train RL model
FY26Q2	Model training and validation
FY26Q3	Deploy AI/ML control system in experimental hall

POLARIZED TARGET MILESTONES	
Previous Milestones	
4/30/2024 (2 months)	Identify and curate appropriate historical data sets of measured polarization. This will include the CLAS12 Run Group C archives.
9/30/2024 (5 months)	Clean data. Map correlations. Investigate UQ models for the RGC data set.
2/28/2025 (5 months)	Separate data into anneal groups and fit Gaussian Process model to positive and negative polarization sets for RGC. Investigate use of single model for positive and negative sets.
4/30/2025 (2 months)	Develop simulation of target polarization behavior based on historical archives that can be used for AI/ML model development.
5/31/2025 (1 month)	Collect historical waveform data for NMR signal from polarized target and prepare for use in model training.
Future Target Milestones	
6/30/2025 (2 months)	Implement signal extraction technique for accurate extraction of NMR signal from electronics background.
8/30/2025 (2 months)	Identify and train an appropriate model for controlling microwave frequency based on historical data and direct NMR feedback. This should start with a Deep Reinforcement Learning model.
10/31/2025 (2 months)	Test model against simulation and adjust to optimize performance.
12/31/2025 (2 months)	Utilize the Polarized Target Group's test facility to test the model and further refine it.
2/28/2026 (2 months)	Integrate models and appropriate codes into the AI/ML controls ecosystem and deploy in Hall-B.

POLARIZED SOURCE MILESTONES	
Previous Milestones	
3/31/2024 (1 month)	Identify all potentially relevant parameters (e.g. beam positions, energy, collimator, etc..) and gather historical data. Curate into form suitable for processing with modern data science tools.
7/31/2024 (4 months)	Identify "nudge" events and responses to build data set for training.
12/31/2024 (5 months)	Investigate "nudge" sequences in GlueX 2020 and 2023 data sets
3/31/2025 (3 months)	Investigate correlations of measured beam position/angle drift with coherent peak position
Beam Spot Milestones	
5/31/2025 (2 months)	Port existing UConn beam spot finder tool webpage into format that can be used via command line interface.
7/15/2025 (1.5 months)	Determine degradation parameter(s) for coherent peak as function of dose
9/30/2025 (1.5 months)	Formula for calculating degradation map as function of dose for a given diamond map and beam profile. Report FOM for any given diamond and configuration of N beam spots.
12/31/2025 (3 months)	Create a Bayesian optimization tool for identifying optimal configuration of N beam spots on a given diamond
2/28/2026 (2 months)	Document tool and deploy alongside existing Cobbrems tools
Automated Coherent Peak Positioning Milestones	
6/30/2025 (1 month)	Identify "nudge" events and responses to build data set for training.
7/31/2025 (1 month)	Train model to provide response of coherent peak position to nudges w/ backlash included
8/31/2025 (1 month)	Wrap model in simulation layer to include drifts in coherent peak
12/31/2025 (4 months)	Train RL model to "push" nudge buttons using simulation
2/28/2026 (2 months)	Connect AI/ML model from the larger lab DS ecosystem to the control system for the goniometer. Include appropriate elements into the standard control system GUIs.

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).

https://wiki.jlab.org/epsciwiki/index.php/AI_Optimized_Polarization
https://wiki.jlab.org/epsciwiki/index.php/AI_For_Experimental_Controls

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 Numbers: DE-FOA-0002875 - GRANT13779668

BACKUPS

Measured Polarization vs. Target dose

Polarization of cryo target as a function of accumulated dose

The polarization tends to drop linearly with dose at a fixed u-wave frequency

Temporary beam trips cause brief spikes in polarization as the temperature drops slightly

Occasional adjustments to the u-wave frequency can improve polarization

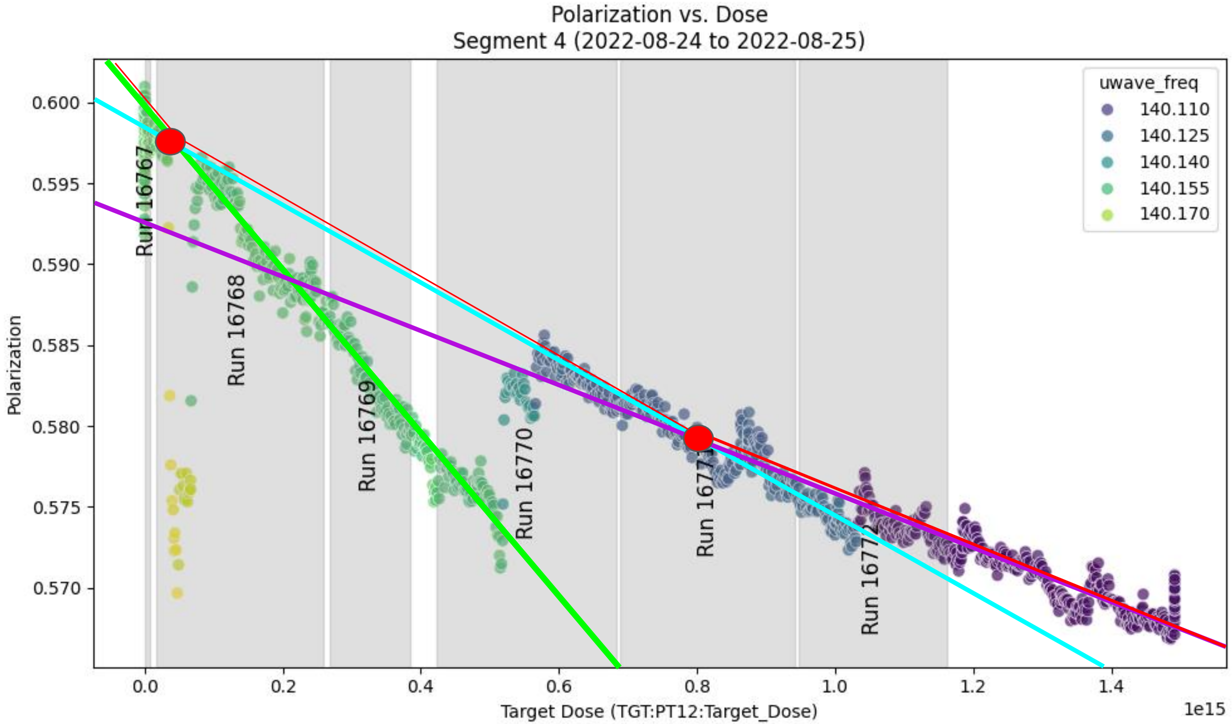

Surrogate Model

- 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.

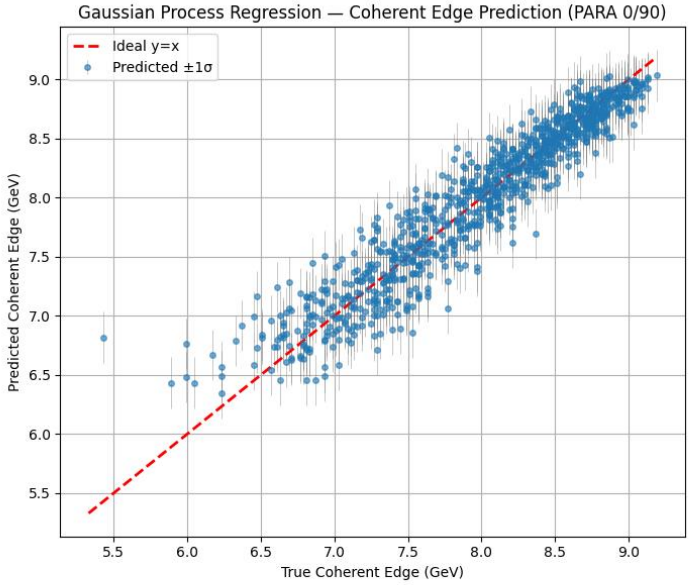

Surrogate Model (MLP)

- **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**

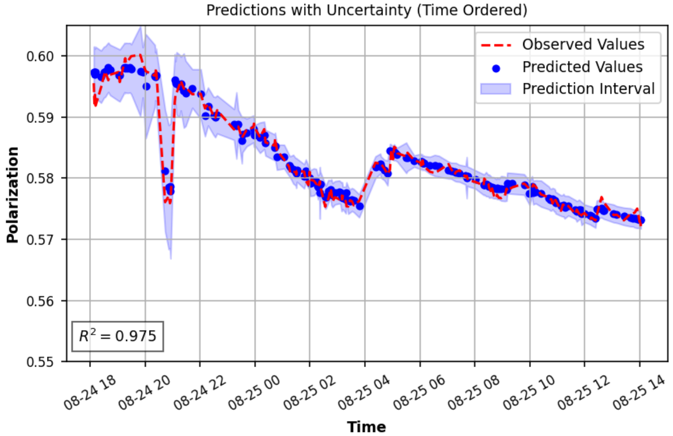

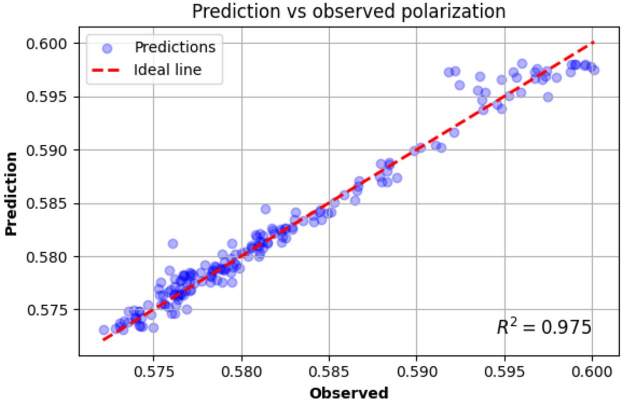

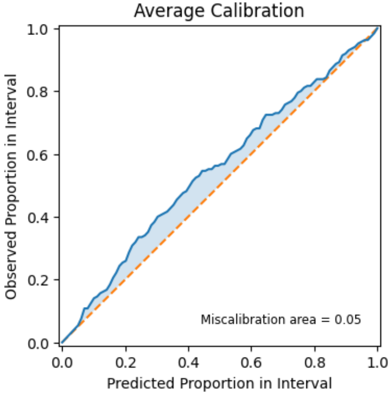

Surrogate Model (GP)

- GP 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

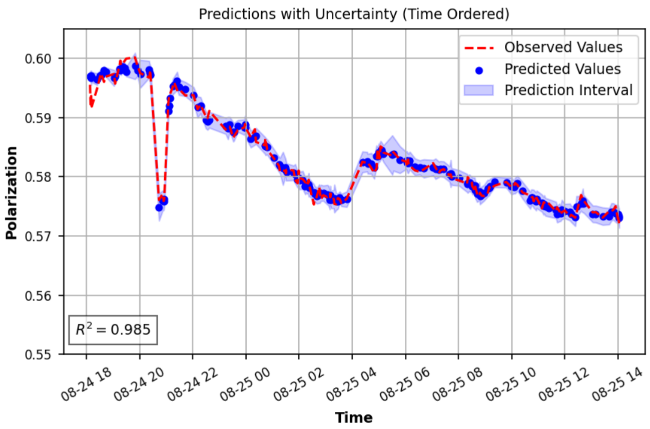

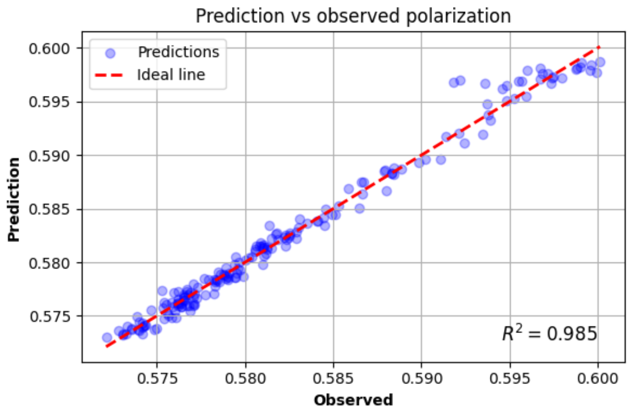

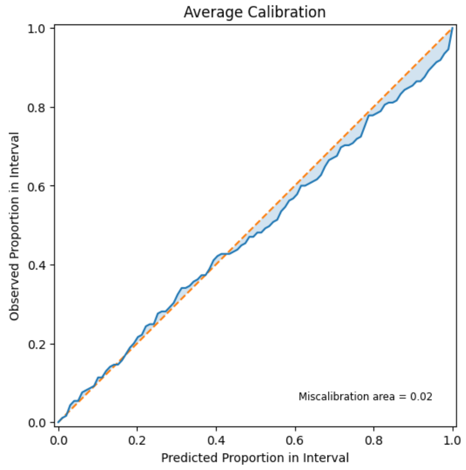

Polarization vs. Dose

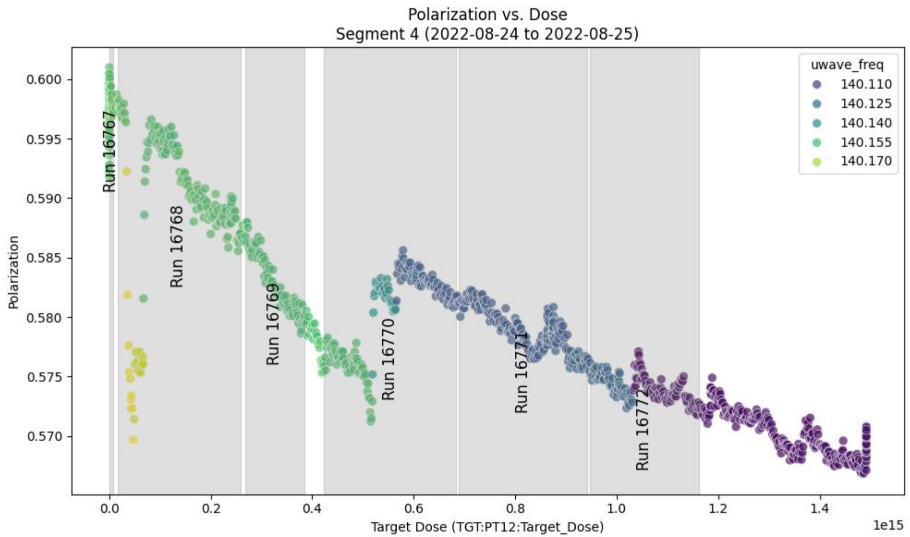

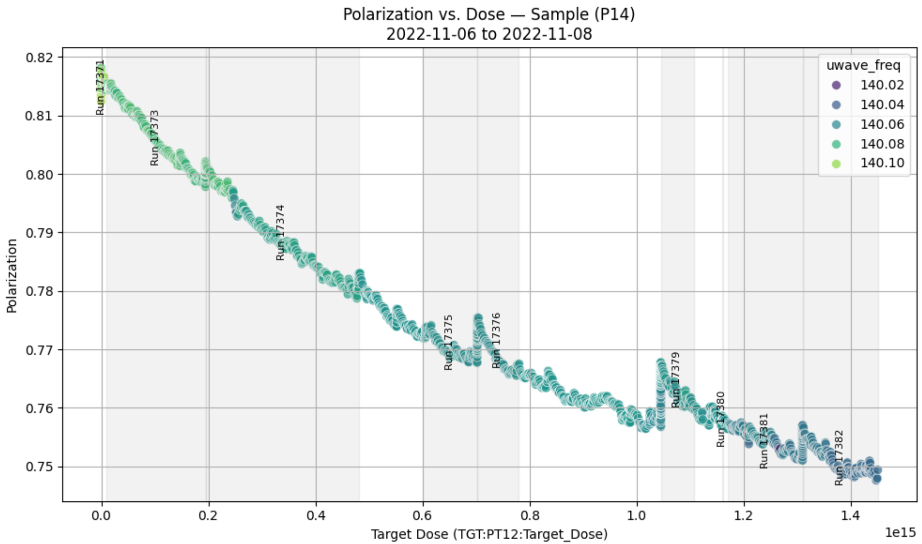

Environment Setup - V0

- Goal: Maximize total polarization over run period
- Environment setup
  - Polarization starts at fixed point (0.600) configurable
  - Environment resets when polarization reaches (0.55
  - Environment dynamics determined by setting on fre
    - Current Polarization based of Y = Mx+b
      - x is dosage
      - M is determined by what frequency
      - b is determined by what frequency
  - Simple environment only has 3 deterministic actions
    - Setting 1 (140.155) Highest slope
    - Setting 2 (140.125) Middle slope
    - Setting 3 (140.110) Lowest slope
  - State:
    - Current Polarization
    - Current Dosage
    - Current Frequency Setting
  - Reward:
    - Current polarization

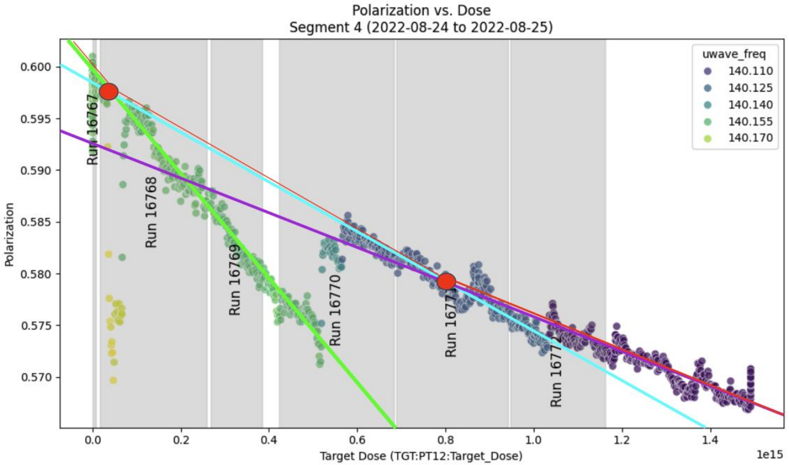

V0 Visualization - Recreating Human Behaviours

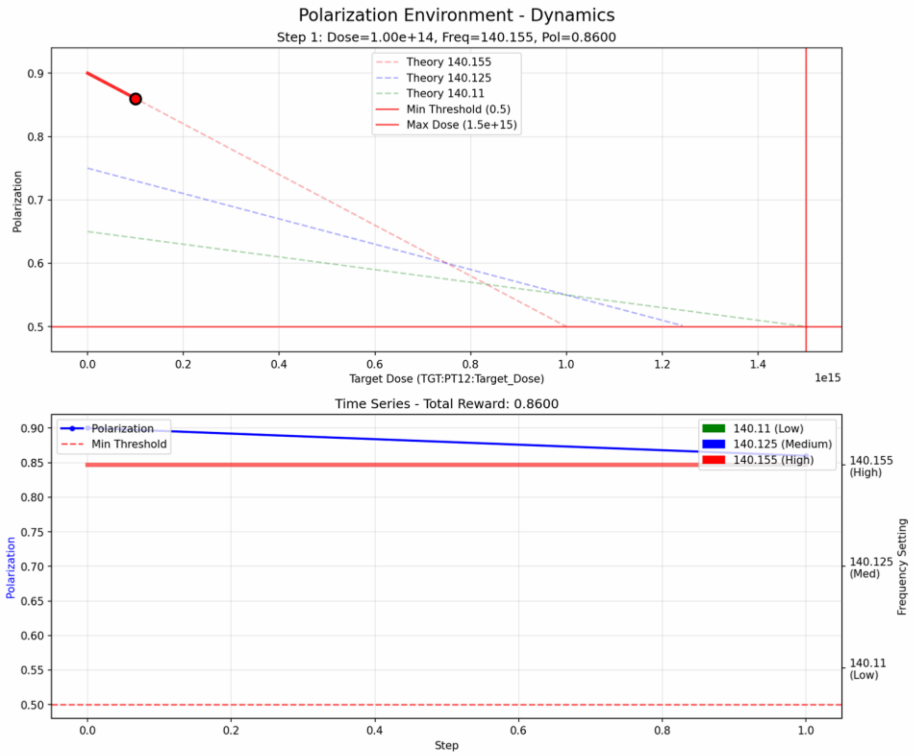

- Find estimated equations of each frequency
  - Not exact but following the same principle
- Y=Mx+b
  - 140.155: (-4.0e-16, 0.900)
  - 140.125: (-2.0e-16, 0.750)
  - 140.110: (-1.0e-16, 0.650)

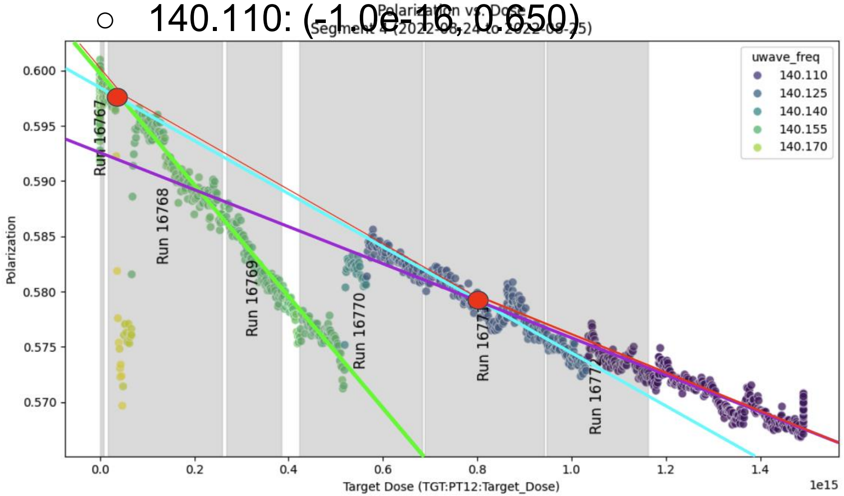

V0 Visualization - Human vs Optimal Behaviours

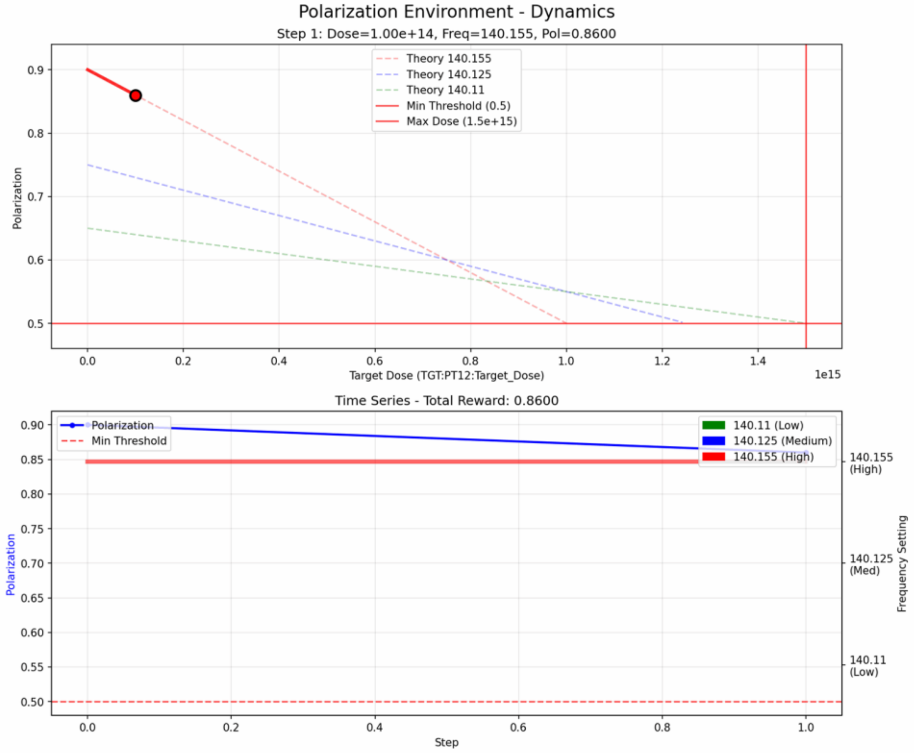

Human Behaviour

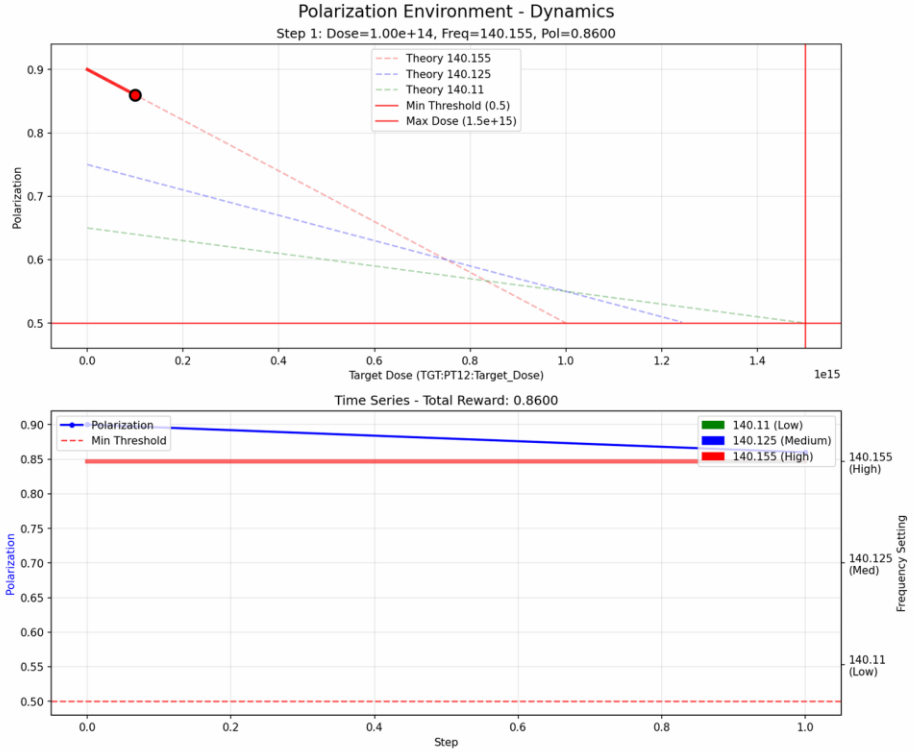

Optimal Behaviour

Environment Setup - V1

- Goal: Maximize total polarization over run period
- Environment setup
    - Polarization starts at fixed point (0.900) configurable
    - Environment resets when polarization reaches (0.550) or dosage gets too high (both configurable)
    - More advanced environment only has **continuous action space** to simulate button presses
    - Environment dynamics determined by setting on frequency
        - Current Polarization based of Y = Mx+b
            - x is dosage
            - M is determined by what frequency
            - b is determined by what frequency
        - M and b are determined by equations of **frequency** (through quadratic interpolation from previous equations and a linear fit)
            - These should be changed and are stand ins until we have a better understanding of the data!
                - Designed to be configurable!
            - M(frequency) = 1*10^-16*(-66.667*frequency +9339.65)
            - b(frequency) = -37.037*frequency^2 + 10385.728*frequency - 728077.315
    - Action:
        - Delta change of current frequency
    - State:
        - Current Polarization
        - Current Dosage
        - Current Frequency Setting
    - Reward:
        - Current polarization

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.

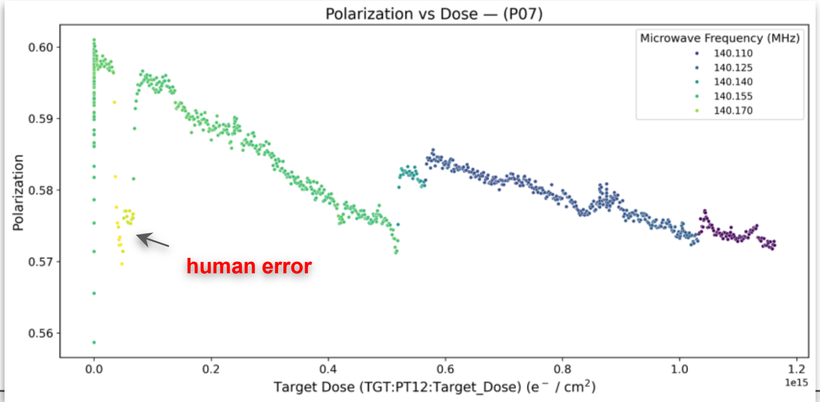

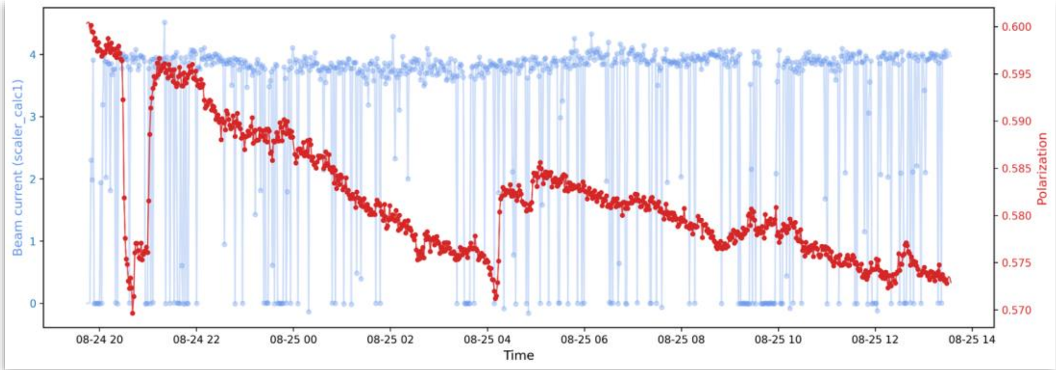

Surrogate Models

- 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.

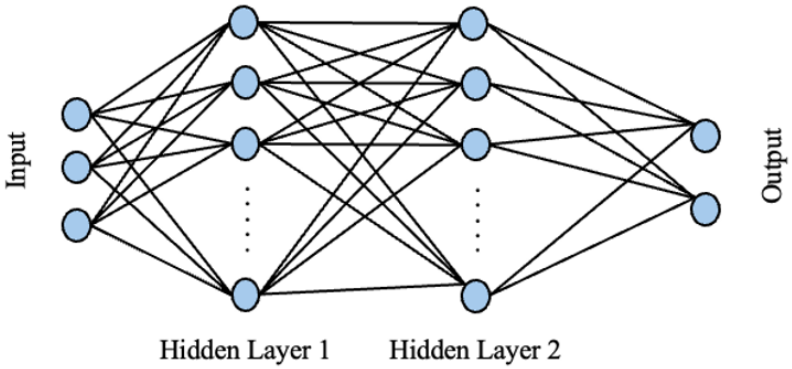

multi-layer perceptron

Goniometer Offsets

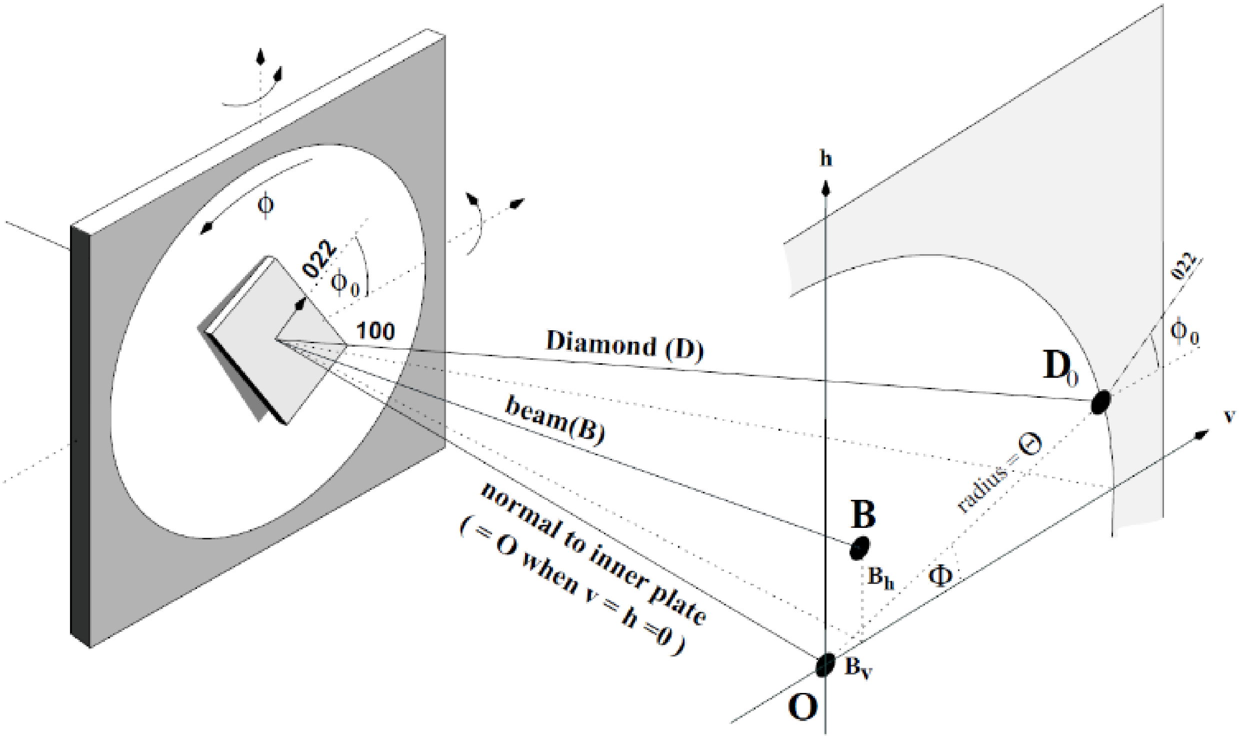

OOD frequency above maximum training frequency (140.1822 GHz)

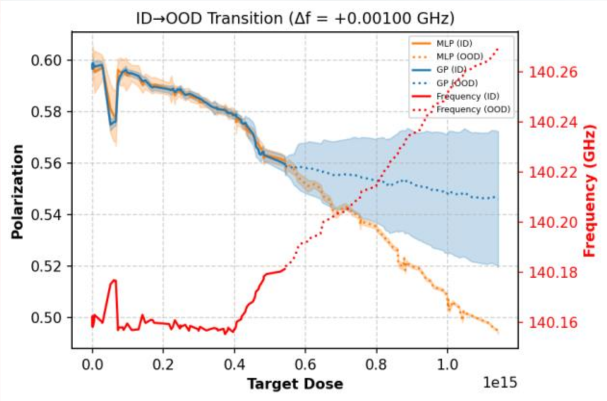

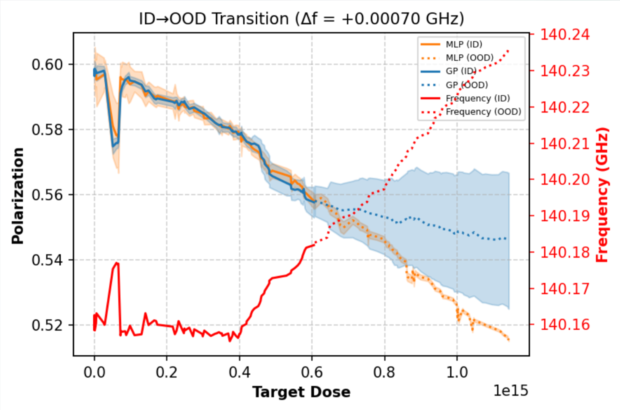

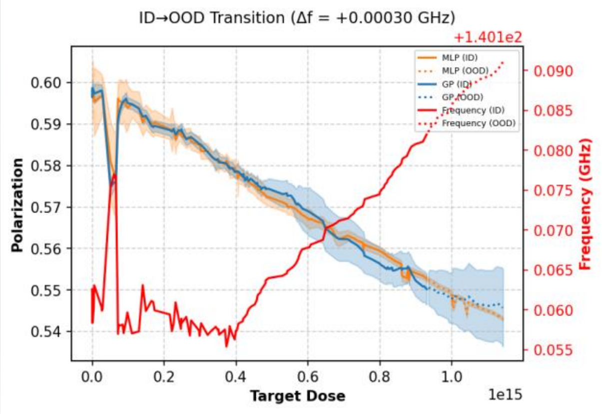

OOD frequency below minimum training frequency (140.0979 GHz)

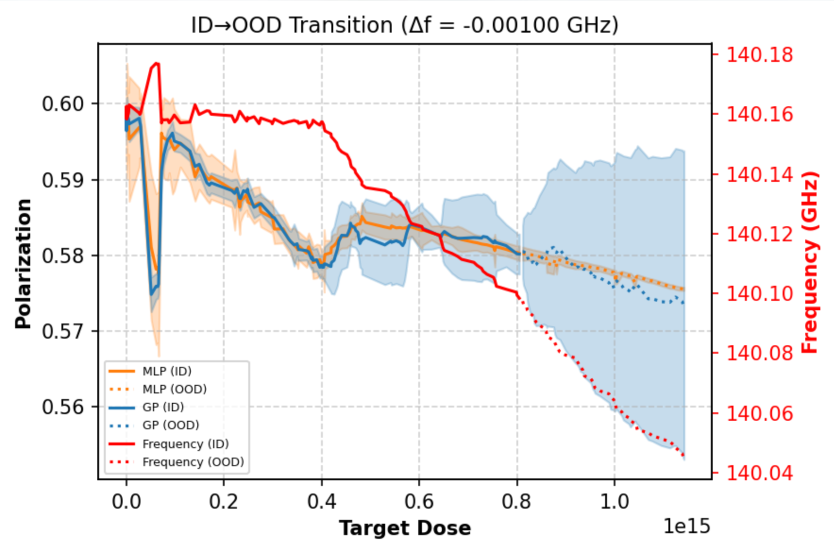

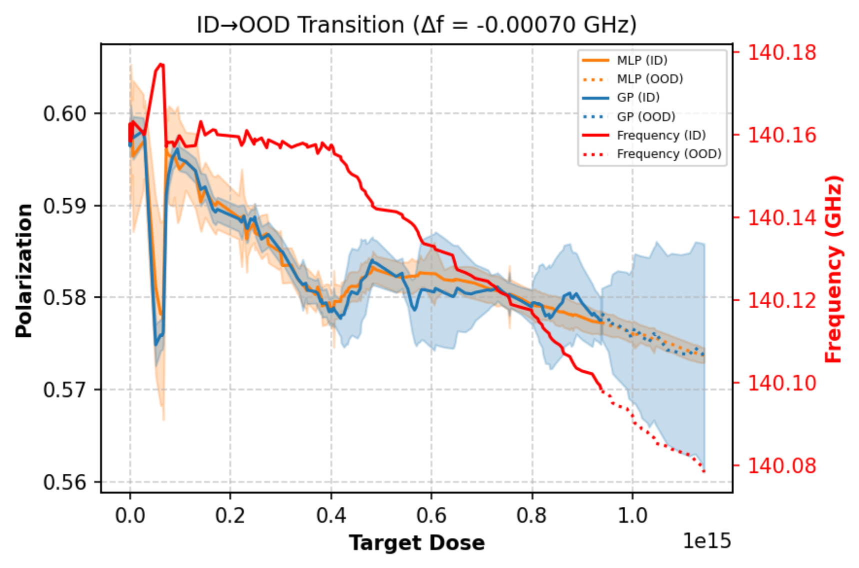

GPA Results [Laplace Kernel]

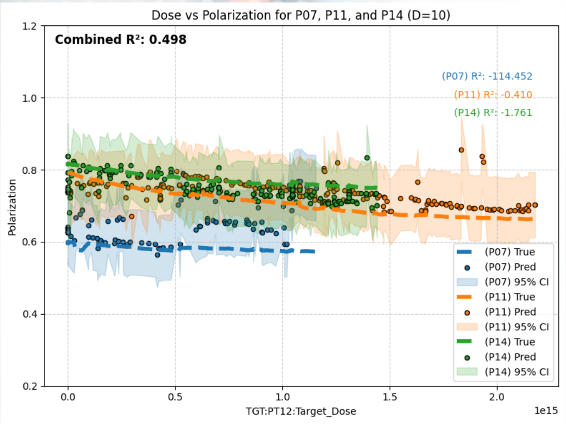

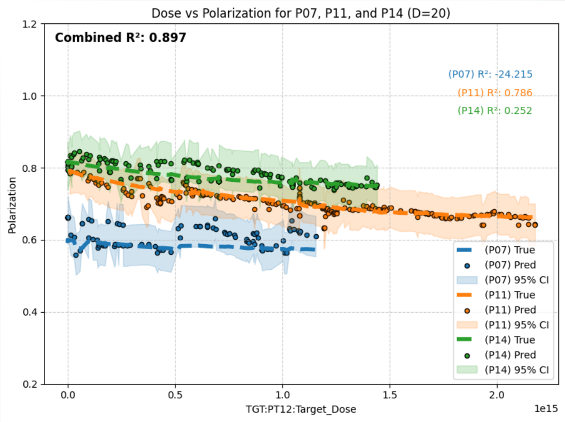

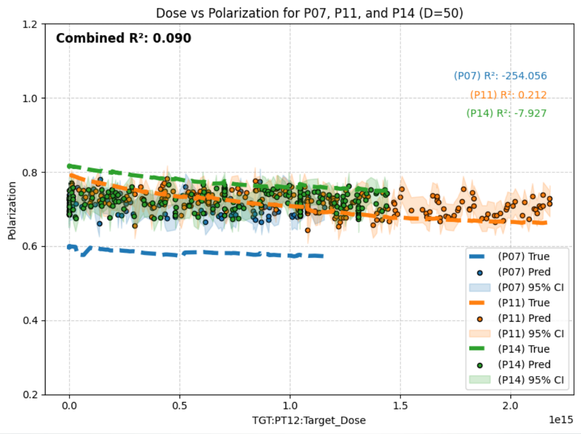

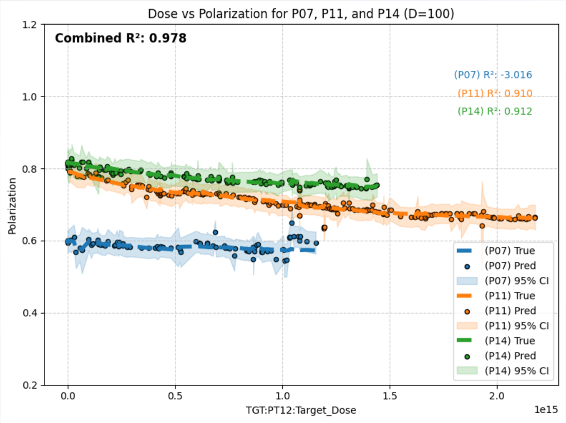

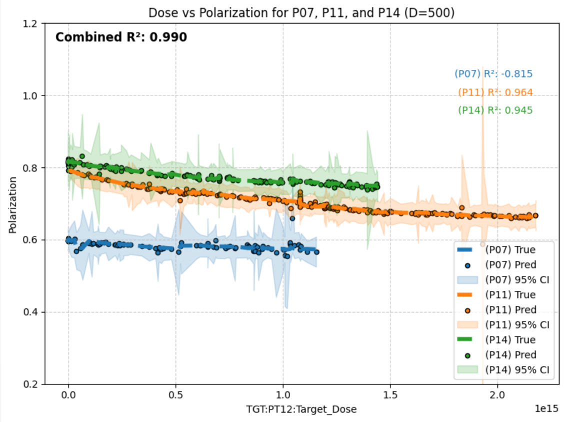

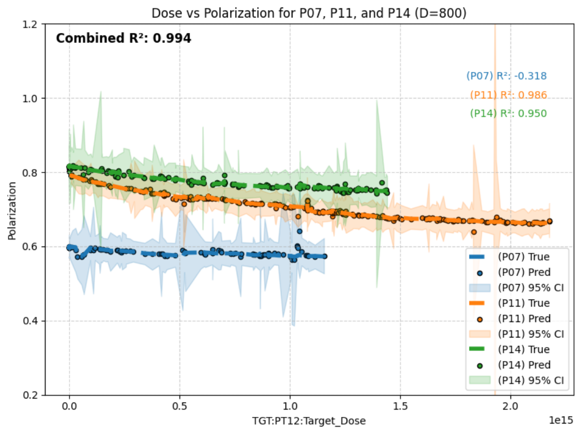

GPA Results[Laplace Kernel]

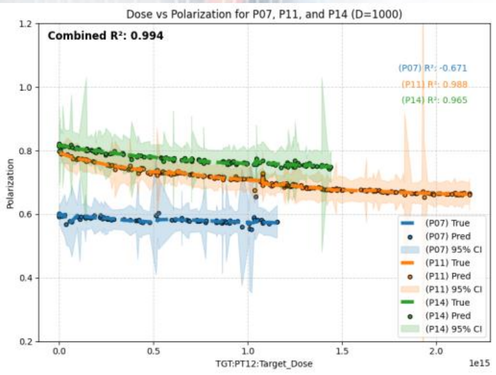

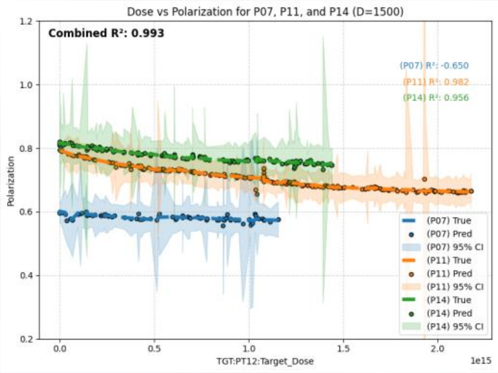

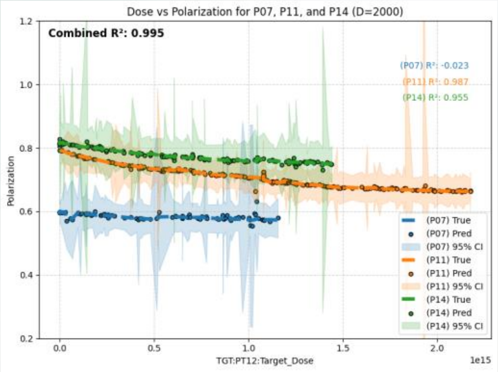

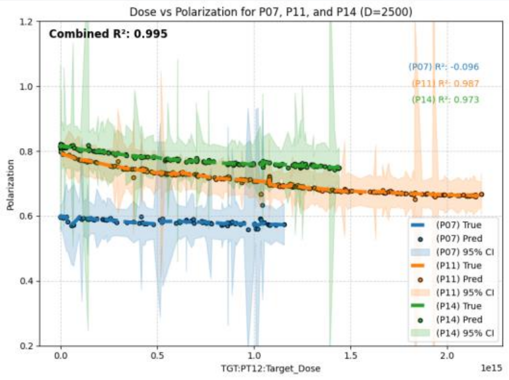

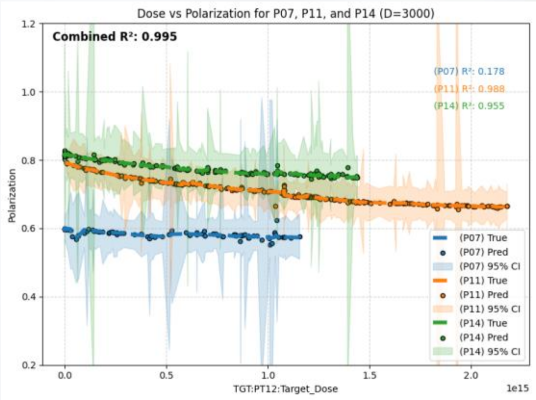

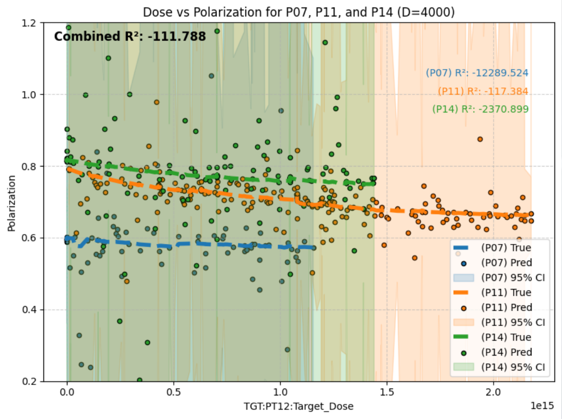

GPA Residuals [Laplace Kernel]

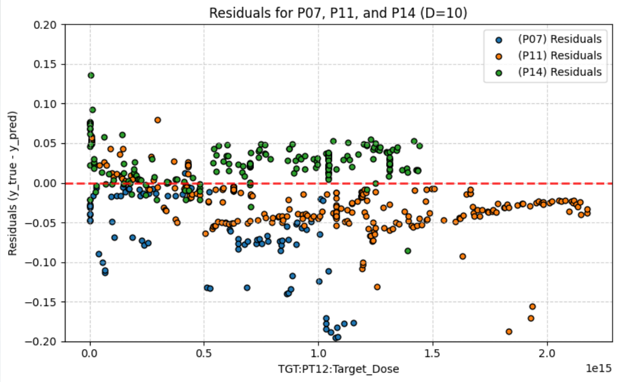

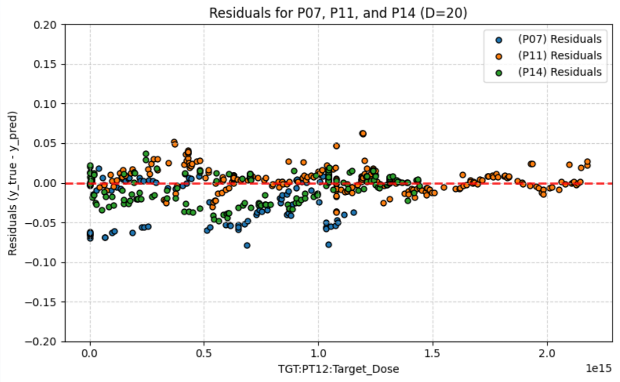

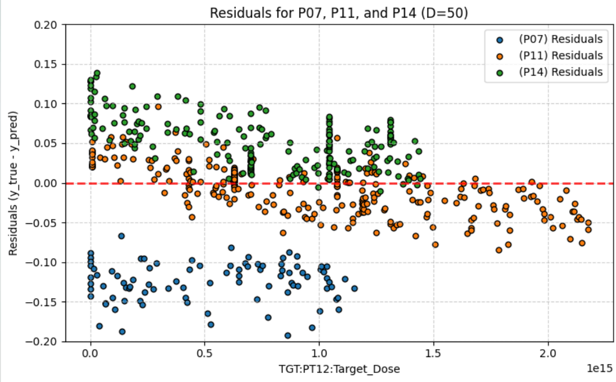

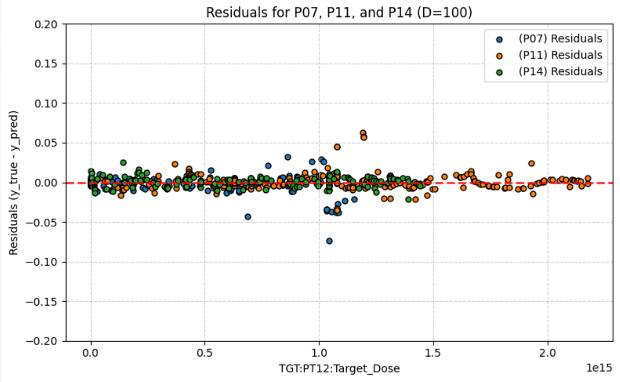

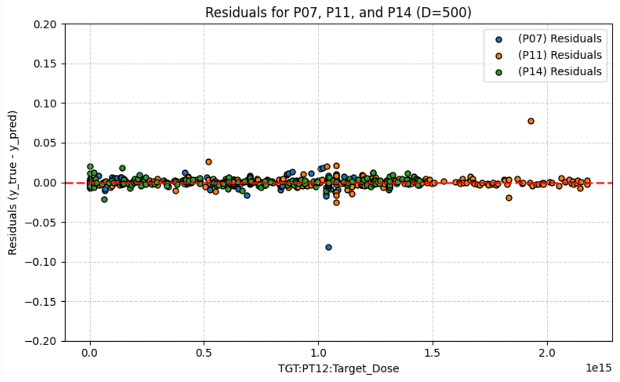

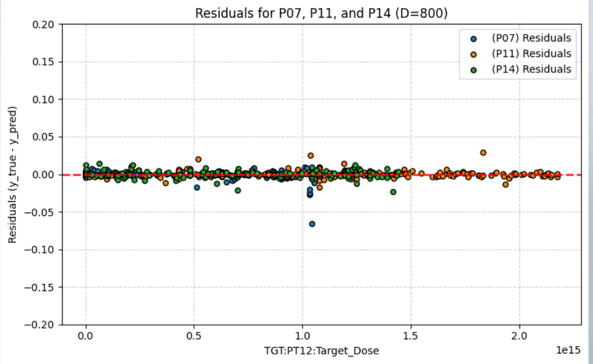

GPA Residuals [Laplace Kernel]

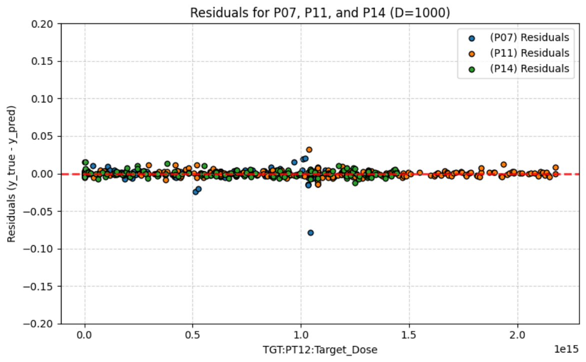

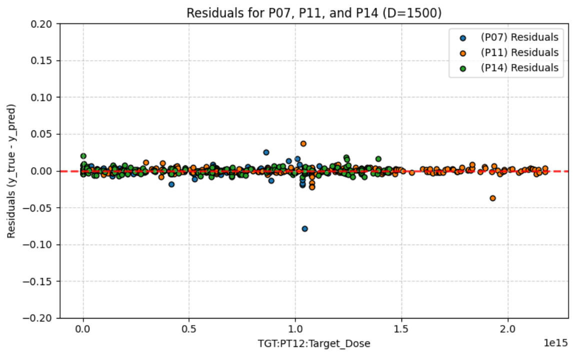

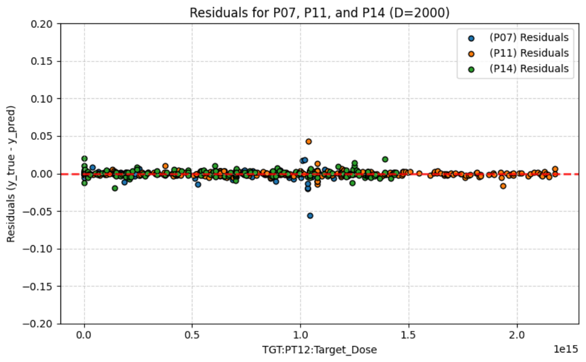

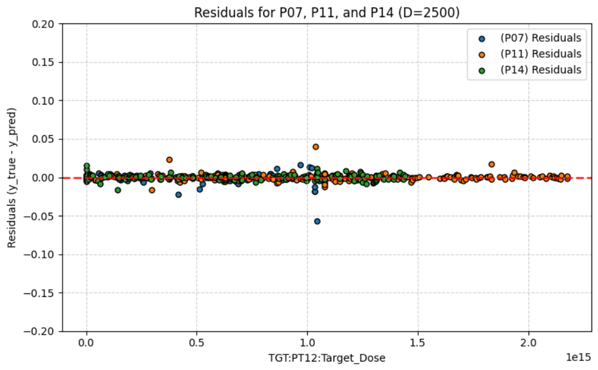

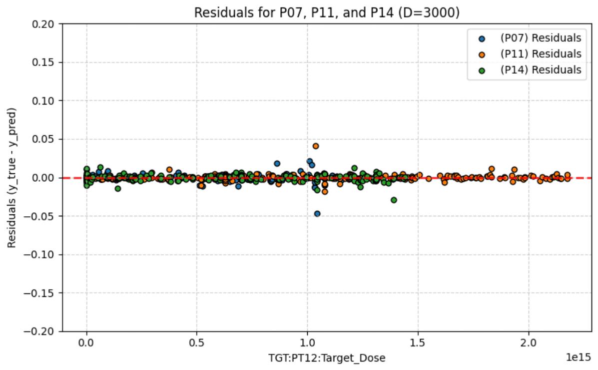

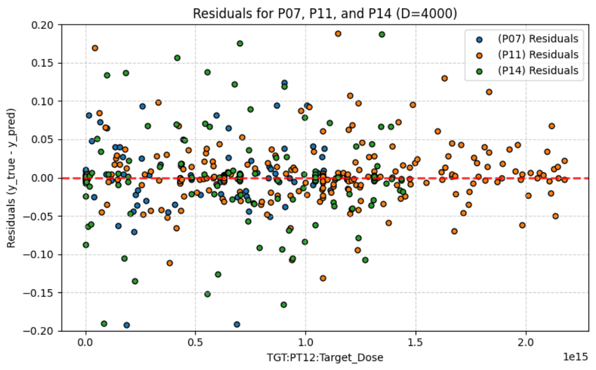

GP [Laplace Kernel]

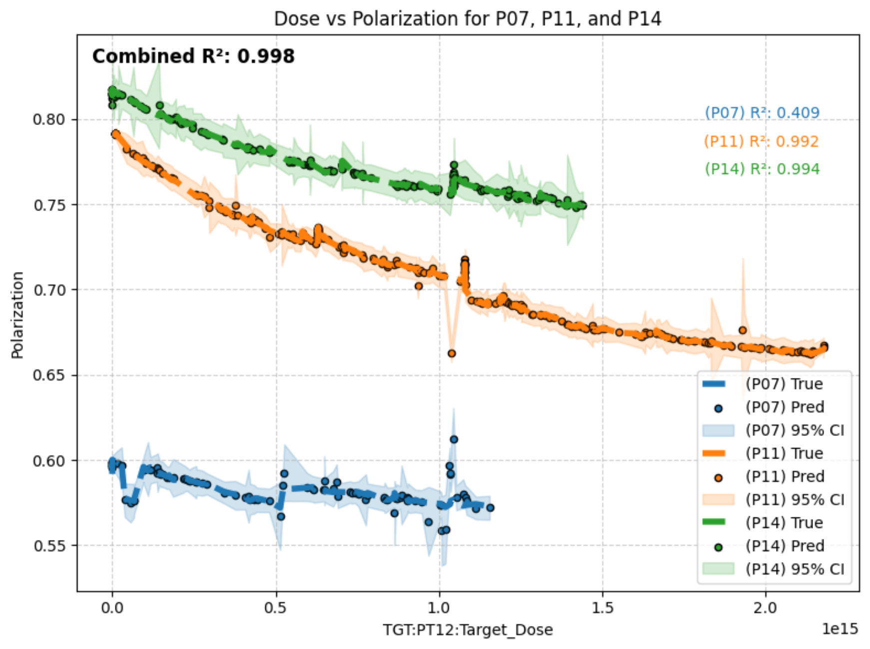

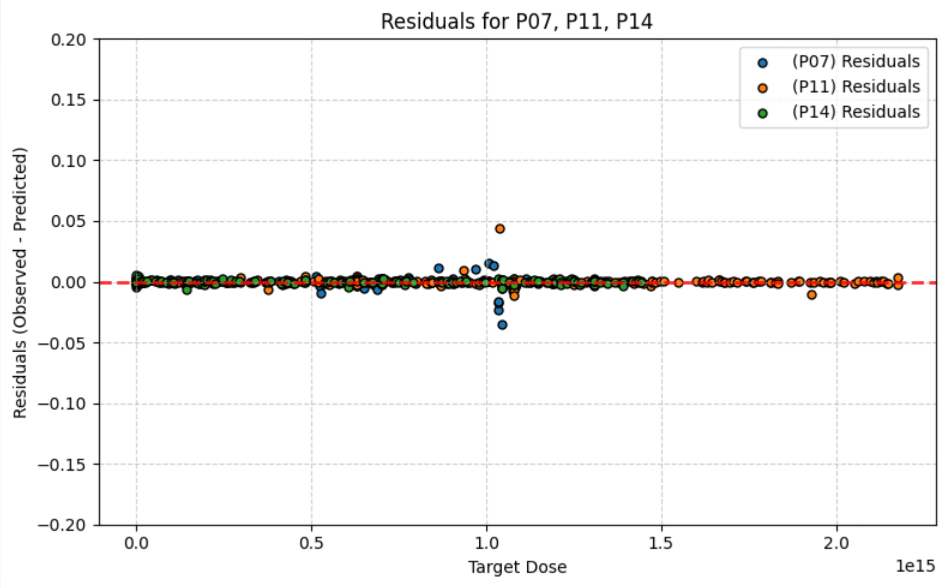

GP exact (RBF)

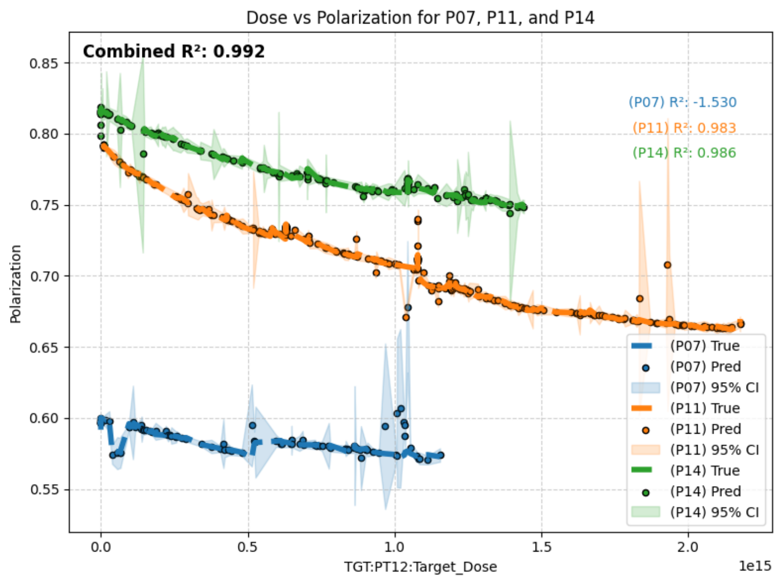
Dose vs Polarization for P07, P11, and P14
Combined R²: 0.992
(P07) R²: -1.530
(P11) R²: 0.983
(P14) R²: 0.986
Polarization
TGT:PT12:Target_Dose
1e15
(P07) True
(P07) Pred
(P07) 95% CI
(P11) True
(P11) Pred
(P11) 95% CI
(P14) True
(P14) Pred
(P14) 95% CI

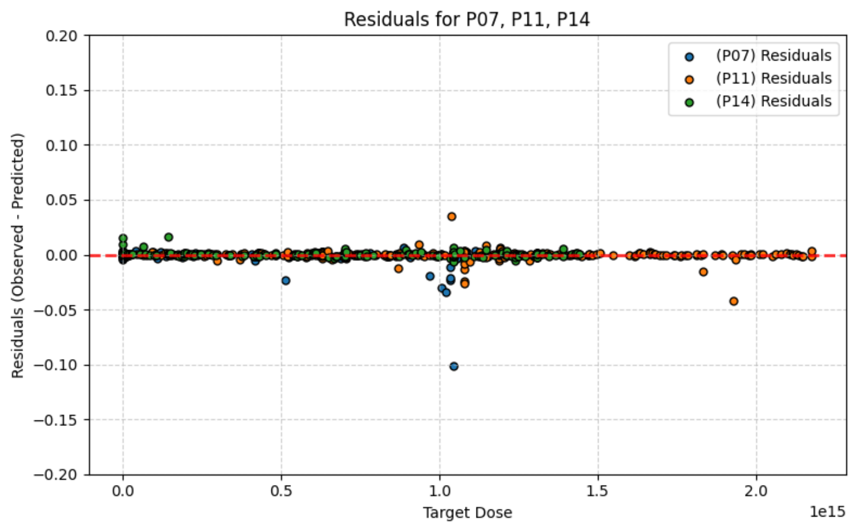
Residuals for P07, P11, P14
Residuals (Observed - Predicted)
Target Dose
1e15
(P07) Residuals
(P11) Residuals
(P14) Residuals