Project title:
AI/ML Optimized Polarization

Institution: Thomas Jefferson National Accelerator Facility (TJNAF)

Street address: 12000 Jefferson Avenue, Newport News, VA
Postal Address: Suite 15, Newport News, VA 23606

Administrative Point of Contact: Deborah Dowd, (757)269-7180, dowd@jlab.org

PI: David Lawrence, (757)269-5567, davidl@jlab.org

DOE National Laboratory Announcement Number: **DE-FOA-0002875**
DOE/Office of Science Program Office: Nuclear Physics (NP)
DOE/Office of Science Program Office Technical Contact: Dr. Manouchehr Farkhondeh

Duration of project: 2 years

Contents

A Background/Introduction

This project seeks to develop AI/ML based control applications for optimizing polarized targets and beams for use in all 4 experimental halls throughout Jefferson Lab (JLab). The focus will be on two applications: cryogenic polarized targets and the linearly polarized photon source. Dynamically polarized proton and deuteron targets [1] are crucial tools for the study of nuclear spin structure at JLab (see [2], [3], [4]) and other facilities [5]. Throughout its 6 GeV era, these systems were used in each of the lab's three experimental halls to enable nine experimental programs, and the first of six programs at 12 GeV is currently underway in Hall B. The linearly polarized photon beam is an essential component to carry out the GlueX physics program, namely beam asymmetry studies and searches for long-sought exotic hybrid mesons via partial wave analysis that would confirm the role of gluonic excitations in a mass region that is largely unexplored and favorable for their observation. In addition to GlueX, there is a rich experimental program in Hall-D that relies on this polarized photon source including the Charged Pion Polarizability (CPP) [6] and Short Range Correlations [7] experiments, and the future K-Long Facility[8]. Currently, both of these applications require constant attention from shift takers and experts. Implementing uncertainty aware, robust AI/ML based automation and control systems can improve the stability of the polarization, effectively increase statistics at no cost, and reduce inconsistencies inherent in human control.

This work will lay additional foundational stones that can eventually be used to build a fully self-driven experiment. A future goal for NP and HEP experiments is the ability to control multiple detector and beam systems in order to optimize on a set of physics goals. This will need to be a layered approach in which individual systems already under autonomous AI/ML control can be coupled by a larger AI/ML system that optimizes on global goals. For example, this work, which focuses on aspects of targets and beam controlled by the experimental halls, could be coupled with AI/ML controlled detector systems such as the one that the successful AI for Experimental Controls (AIEC) project started at JLab under LAB-20-2261[9]. While coupling independent systems is beyond the scope of this proposal, this work will add to the number of AI/ML controlled systems that could be brought together under a global AI/ML as part of a future FOA.

A detailed description of the two polarization control applications is provided in the following sections.

A.a Polarized Target

Cryogenic polarized targets are the most complex target systems utilized at JLab, and their performance is sensitive to numerous factors including temperature, the instantaneous and integrated beam currents, and fluctuations in both the microwave apparatus that drive the polarization and the NMR electronics that measure it. The most commonly used materials for dynamic polarization are irradiated solid ammonia (NH_3) and deuterated ammonia (ND_3). To prepare these materials for experiments, the gases are frozen, crushed into millimeter-sized granules, and then irradiated with a low-energy electron beam to create paramagnetic radicals within the solid lattice. After irradiation, they can be stored in liquid nitrogen until needed. During the scattering experiments, the radicals' electron spins are polarized to nearly 100% at ultra low temperatures and high magnetic fields. This high polarization is then transferred to nuclear spins using microwave-induced transitions that simultaneously flip both electronic and nuclear spins. Depending on the microwave frequency, either positive or negative nuclear polarization can be produced. In a well optimized system, proton (deuteron) polarizations exceeding 90% (40%) can be achieved under the "standard" JLab conditions of 1 K and 5 T.

Unfortunately, the average polarization that can be maintained during the electron-scattering experiment is lower because the beam creates additional radicals in the lattice that are deleterious to the polarization process. To reduce the rate at which polarization is lost, the microwave frequency must be adjusted as the integrated beam charge (dose) on the sample accumulates. Figure 1 at left shows the polarization achieved in Hall B over the charge accumulated lifetime of one target sample's lifetime. This shows large, regular spikes in the polarization as the material is annealed to 80-100 K to partially remove the unwanted radicals and recover polarization. It also displays several smaller jumps that are due to beam trips and microwave adjustments. Traditionally, these adjustments have been performed by shift workers periodically during the experiment as they feel it is warranted. Figure 1 at right shows the manually selected microwave frequencies as a function of dose since the last anneal event for all target samples used in a Hall C experiment. This plot illustrates that the optimal frequency and its behavior with beam is not straightforward to predict. It can vary from one sample to the next and usually shifts after a given sample is annealed. As a result, the experience of the shift workers plays a large role in determining the average target polarization that can be maintained over the course of the experiment. Machine Learning and Artificial Intelligence methods can be utilized to reduce this effect. Moreover, because the target polarization enters the experiment Figure-of-Merit (FOM) as a square, even modest improvements in the average polarization can noticeably improve the statistical precision of a given experiment, or reduce the time required for the experiment to reach its desired precision.

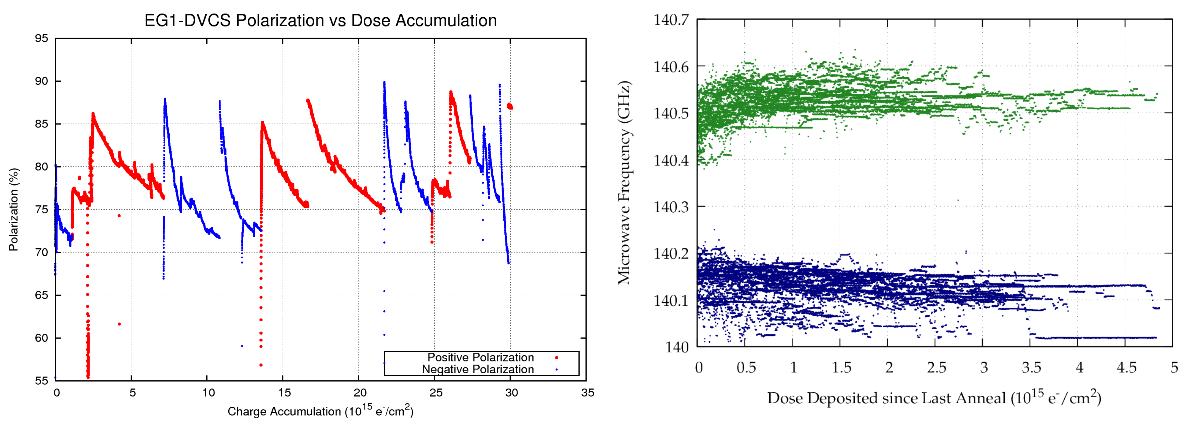

Figure 1: **Left**: Target polarization versus beam dose over the lifetime of one target material sample in Jefferson Lab's Hall B. **Right**: Chosen microwave frequency versus beam dose since the last anneal in Hall C, showing the trend of the frequency for negative polarization (blue, near 140.1 GHz) and positive polarization (green, near 140.5 GHz).

In addition to improving the average polarization of the target during the experiment, AI/ML techniques may improve the accuracy to which this polarization is measured using standard continuous-wave NMR techniques. Figure 2 illustrates the process of extracting the target polarization from the NMR signal. The raw signal ("a" in the figure) is a plot of the output voltage as the circuit is swept through the nuclear resonance frequency. The background under the resonance signal is inherent in the circuit and is too complex to fit with great accuracy. Instead, the background is measured by adjusting the magnetic field so the target sample is no longer on resonance and subtracted from subsequent on-resonance sweeps. Next, a polynomial fit is performed to compensate for any drifts in the background due to small temperature changes both inside and outside the target cryostat ("b").

Even small changes in temperature, both inside and outside the target cryostat, produce drifts in

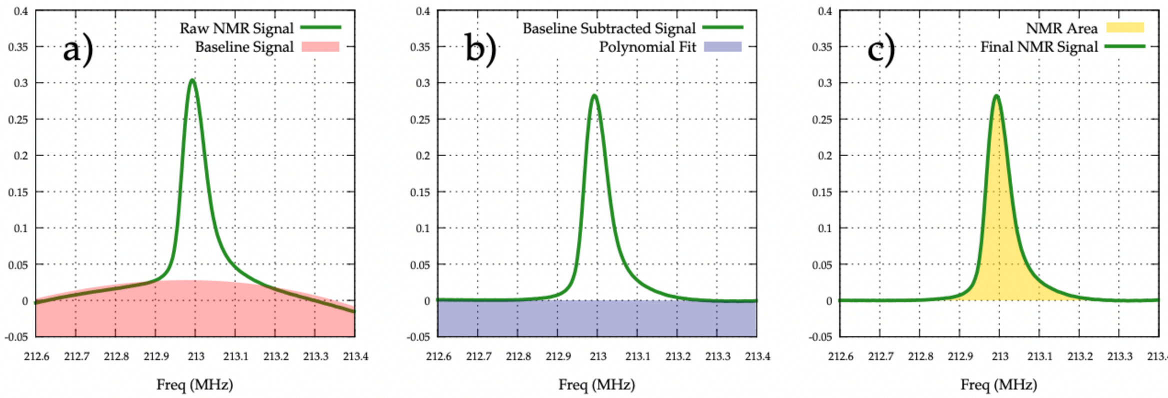

Figure 2: NMR signal background extraction process for a proton signal, showing background subtraction, polynomial residual subtraction and final integration [10].

the circuit response, so it has been important to remeasure the electronic background periodically. First, a previously measured background signal is subtracted ("a" in the figure), then a polynomial fit is done to attempt to remove any drift of the current background ("b" in the figure). Finally, the area under the subtracted curves is taken to determine the total circuit response from the polarization ("c"). In the case of deuteron targets, this becomes even more complicated due to the double peak structure. Figure 3 shows an example of this from real data. The polynomial fit fails to completely remove the background drift, resulting in significant residual background and an inaccurate polarization extraction.

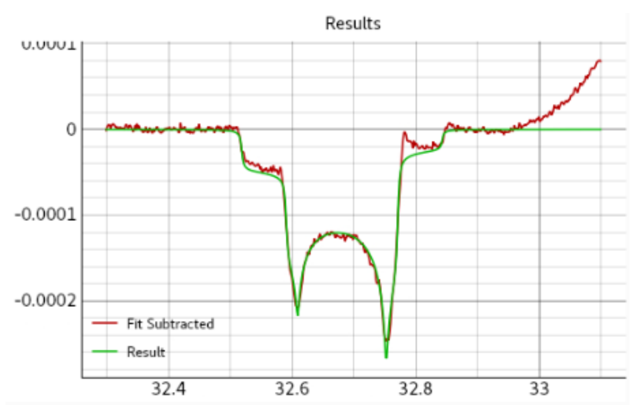

Figure 3: A deuteron polarization signal, showing a characteristic double NMR peak, after the extraction process. This illustrates the failure of the polynomial fit to completely remove the background drift for a small NMR signal, causing inaccuracy in the measurement.

Figure 4 shows the microwave frequency(left) and corresponding polarization(right) for an example 500 minute time period during the CLAS12 Run Group C run period in 2022. A change in the microwave frequency and the corresponding change in polarization can be seen at around the 75th minute. The red line near the top of the polarization plot indicates the average polarization of the top 15% of points in the graph. Taking the ratio of the area under the red curve to the area under the blue curve gives a rough estimate of the improved polarization for this period of time. The statistical equivalent of beam time goes like the polarization squared leading to an additional

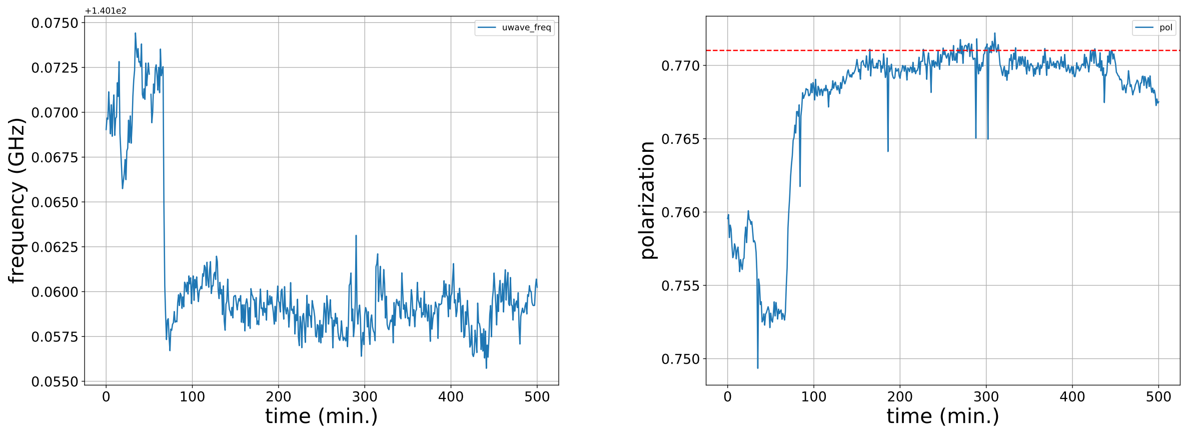

Figure 4: Data from a single 500 minute data file from CLAS12 RGC. **Left**: microwave frequency showing a change at around 75 min. when it was adjusted. **Right**: Measured polarization for the same time period. The red line indicates the average polarization of the top 15% of measurements. The ratio of the integral under red line to that of the blue can be used to estimate a benefit of automating control of the polarization (see text).

1% in this region. Scaling this up to a 36 week year of running, controlling the polarization would gain the equivalent of 2.5 days of beam time per year. This is summarized in Table 1. Given the operating cost of the accelerator (estimated at $\sim \$10k$/hr-$\$15k$/hr) this is significant.

Estimated Benefit Polarized Target	
FOM	0.99
gained statistics	1%
equivalent beam time gained	2.5 days/year

Table 1: Estimated benefit per year of successful implementation of proposed work. (Polarized cryo-target part only).

While the Jefferson Lab Target Group has decades of experience designing, building, and operating these systems, in this proposal we identify two key areas where new AI/ML techniques can be utilized to increase both the average polarization of the target sample and the accuracy to which the polarization is known.

A.b Polarized Photon Source

In the case of the GlueX experiment, a linearly polarized photon beam is incident on a proton target to search for and measure exotic hybrid mesons [11]. The photon beam is produced via coherent bremsstrahlung radiation from an electron passing through a thin diamond wafer (radiator). The diamond radiators are mounted on a goniometer which can change the orientation of the diamonds via three axes of rotation. This enables photoproduction experiments with different orientations of linear polarization direction. GlueX uses four orientations of the polarization vector to have a better handle on the systematic uncertainties, namely at 0, 45, 90, and 135 degrees with respect to the y-axis in the lab frame. The linearly polarized photons have an energy distribution that includes a primary and smaller, secondary and tertiary enhancements that can be adjusted via the angles of the diamond. The polarization is strongly correlated with the photons in the enhanced energy region. Figure 5 shows the polarized photon source control GUI, where the enhancements are shown in the top plot. The vertical line identifies the "coherent edge" as determined by a fit to the photon energy spectrum. The position of the coherent peak edge is dependent on the electron beam position, which fluctuates during the course of an experiment. Ensuring the stability of the primary coherent peak edge is important to control possible systematic effects during data taking.

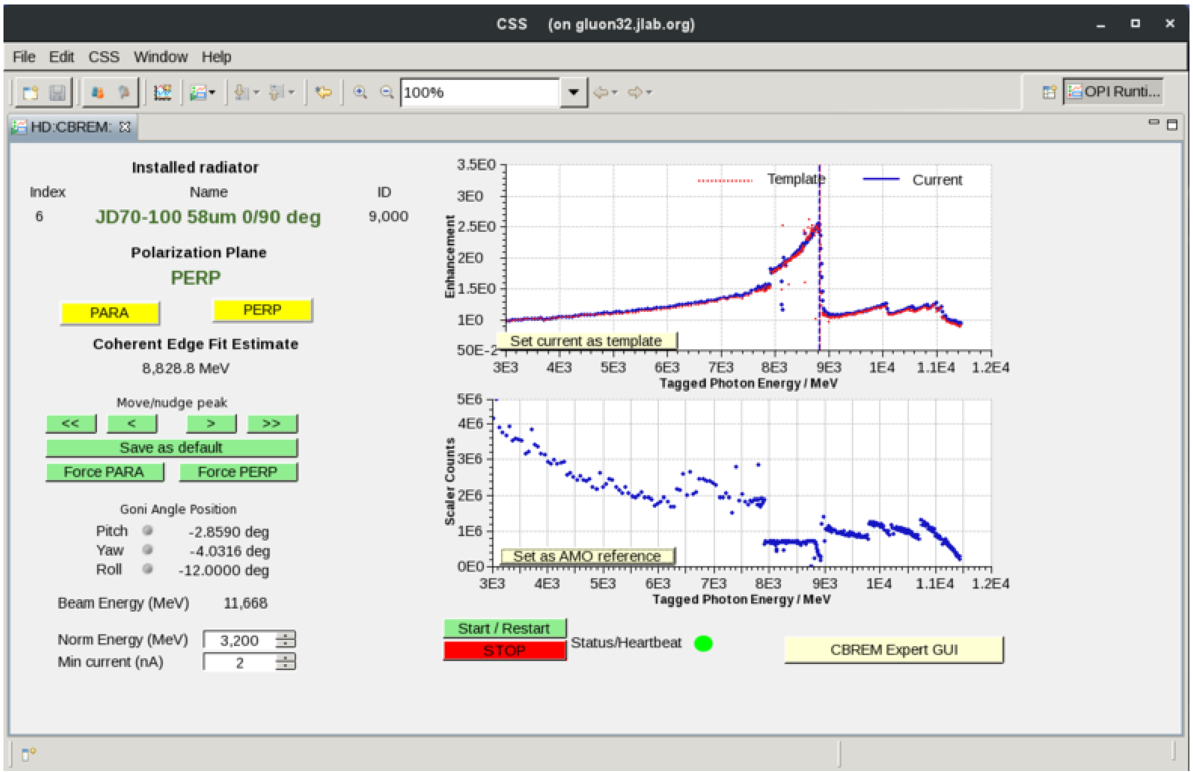

Figure 5: The Hall-D polarized photon source control GUI. The top plot shows the enhancements in the photon energy spectrum with the vertical line identifying the edge of the primary coherent peak. The polarization is strongly correlated with the enhancement.

In Figure 6, the coherent edge position (yellow), beam current (blue), and beam positions as measured by beam position monitors (other colors) are shown. An anti-correlation can be seen between the coherent edge position and the beam position as it drifts. If the coherent peak drifts too high, then the polarization decreases. If the peak drifts too low, then the high energy photon statistics are decreased. An example of the instability in the coherent edge fit parameters as a function of time from the CPP experiment is shown in Figure 7.

Currently, adjustment of the coherent peak position is done manually by shift takers through the GUI shown in Figure 5. This involves moving the coherent edge to within ± 10 MeV of the desired edge value at the start of each run (approximately every 2 hours during data taking). "Nudging" the peak technically corresponds to tuning the orientation of the diamond via the pitch and yaw angles. An alignment procedure is performed on the goniometer at the beginning of each run period that leverages the Stonehenge fitting procedure described in [12]. This procedure allows us to keep one of the diamond orientation angles fixed while permitting changes in the other orthogonal angle. Some difficulties in nudging the coherent peak arise from the fact that the fitter only works in the near neighborhood of the energy peak. In addition, there is a significant time lag on the goniometer

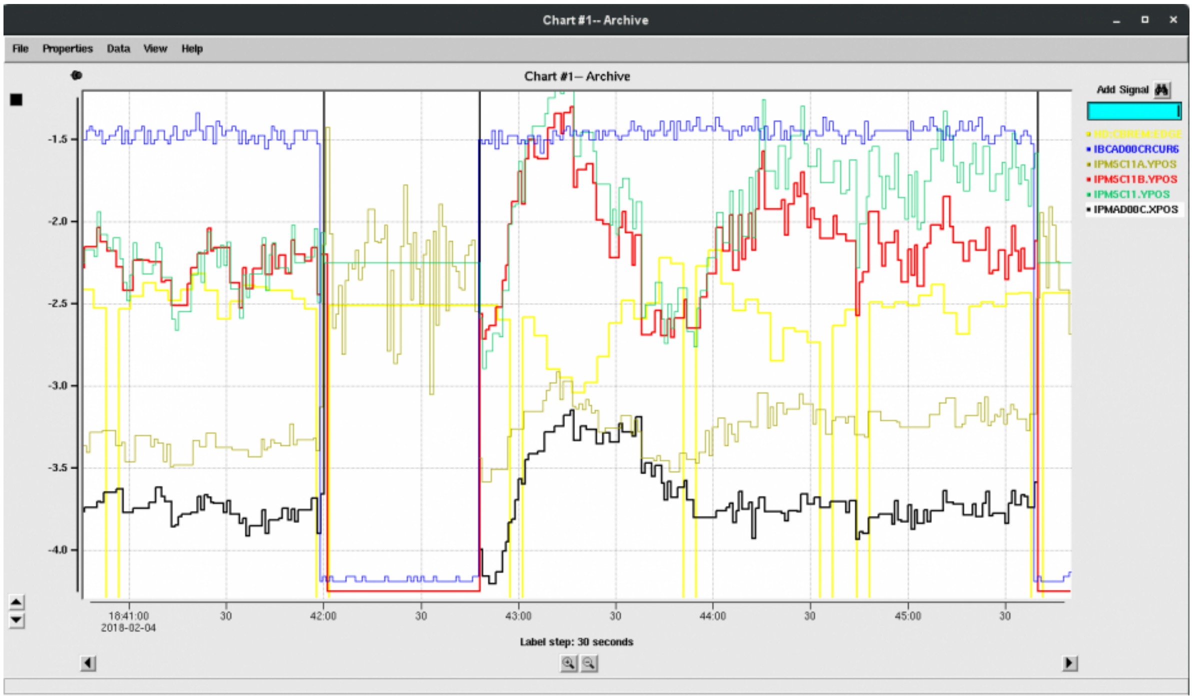

Figure 6: Example of coherent bremsstrahlung edge drifts as a function of beam position drifts. (See text.)

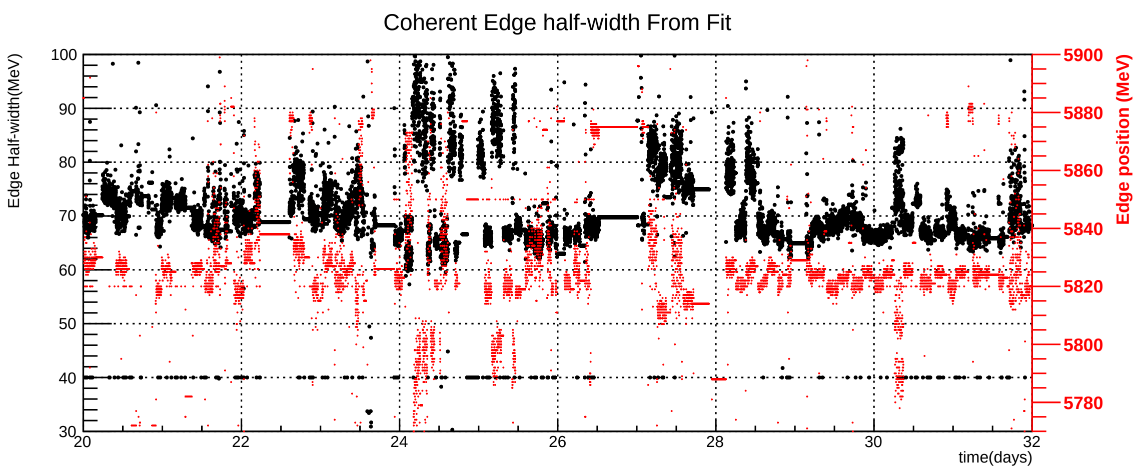

Figure 7: Coherent edge fit parameters for half-width (black plotted against left axis) and position (red plotted against right axis). This is from data taken during a portion of the CPP experiment in the Summer of 2022.

motor motion as well as the response of the readout spectrum. This is typically around 30 seconds each time the peak is "nudged" by the shift worker.

This project aims to replace the manual tuning of the coherent peak done at the beginning of each run with an AI/ML based system that would perform continuous, real-time tuning. Figure 8 shows a calculation used to estimate the potential benefit of the proposed work. The left plot is a histogram of the measured edge positions for the peak relative to the nominal position of 5820 MeV for times when the beam was on and the Data Acquisition system was active. The vertical lines indicate the ± 10 MeV window the proposed work will confine the coherent edge to via AI/ML based control. The curve on the right plots the (normalized) FOM as a function of the coherent edge position integrated over a fixed window of 500 MeV. The FOM is calculated from:

$$FOM = \wp^2 \times dN/dt \times E^2, \tag{1}$$

where $\wp$ is the polarization, dN/dt is the coherent bremsstrahlung rate as a function of energy, and E is the photon energy (squared due to the Primakoff cross section). The 500 MeV integration window was fixed based on the optimal window position for a coherent edge position of 5820 MeV. The vertical axis of the right side of Figure 8 therefore represents the fraction of signal statistics as a function of the coherent edge position. The lines on the plot indicate the points for statistical losses of 0.5%, 1%, and 2% correspond to coherent edge drifts of 21 MeV, 29.5 MeV, and 42 MeV respectively.

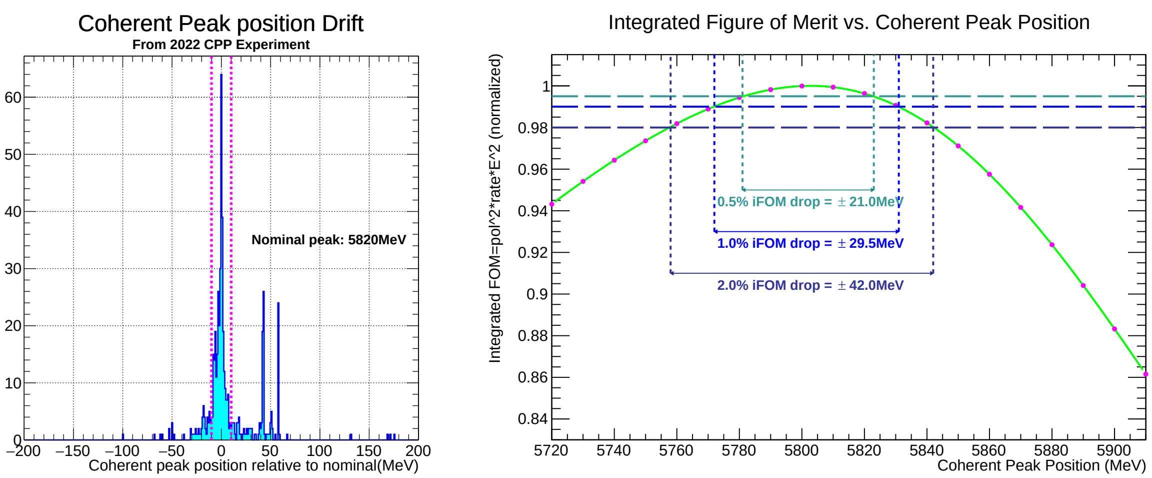

Figure 8: Figure of merit estimation of potential gain based on 2022 Charged Pion Polarizability (CPP) experiment. **Left**: Distribution of measured peak positions relative to the nominal position of 5820MeV. **Right**: Figure Of Merit as a function of measured peak position. Y-axis is fraction of statistics lost due to coherent peak drifting away from nominal (see text).

	All	$\pm 10 MeV window$
Avg. iFOM	0.9909	0.9998
lost statistics	0.906%	0.019%
equivalent beam time lost	2.3 days/year	0.05 days/year

Table 2: Estimated benefit per year of successful implementation of proposed work. (Hall-D Polarized Photon Source part only).

Table 2 indicates the estimated equivalent value of statistics gained if the proposed work is successful. The values come from the convolution of the two plots in Figure 8 with the left side plot centered on the peak position of the right side plot. The “All” column includes the entire spectrum while the “$\pm 10 MeV$ window” column includes only the data within the limits shown on the left side plot. The “$\pm 10 MeV$ window” column therefore represents the practical limit of statistical loss compared to optimal. As noted in the previous section, the accelerator operating cost is estimated at $\sim \$10k/\text{hr}$-$\$15k/\text{hr}$ making this a significant gain.

B Project Objectives

> The primary objective is to improve polarization for fixed target experiments at Jefferson Lab through the use of automated AI/ML controls.

There are multiple parts to the project and within those, multiple objectives. For the purposes of this section, the objectives are grouped relevant to the domains of Nuclear Physics (NP) and Data Science (AI/ML).

B.a NP Objectives

NP1. **Polarized Cryogenic Target Optimization**: The problem is that the microwave frequency corresponding to optimal polarization at any point in time is a function of several parameters: 1) thermal fluctuations, 2) beam trips, 3) target degradation and 4) time since target annealing. It is difficult for a human to predict the optimal polarization when manually adjusting the microwave frequency occasionally during the run.

> NP1: The objective is to learn the optimal microwave frequency control policy to maximize the target polarization continuously throughout an experiment such that it is 0.5% better relative to the current technique.

NP2. **Hall D Polarized Photon Source**: The problem is that the angles of the diamond radiator corresponding to the maximum photon beam polarization are a function of several parameters including: (1) changes in the electron beam position/energy, (2) thermal effects, (3) diamond degradation and (4) vibrational effects in the goniometer on which the diamond is mounted. Human adjustment of these is done occasionally and via trial and error every few hours.

> NP2: The objective is to continuously learn and apply autonomous, real-time angular shifts to the goniometer that nudge the coherent peak of polarized bremsstrahlung photons to its nominal position. This will keep the coherent peak to within 10MeV.

B.b AI/ML Objectives

The project objectives are organized into two parallel applications of uncertainty-aware surrogate models to experimental control and two independent *applications* of that AI/ML research to real NP scientific facilities. Details of each of these objectives are discussed in Section C.

Model Development:

> ML1. For NP1, develop a model for NMR signal extraction.
>
> ML2. Uncertainty-aware surrogate models
>
> ML2.1. For NP1, develop and evaluate a surrogate model that captures uncertainties for microwave frequency control that optimizes polarization of the cryogenic target.
>
> ML2.2. For NP2, develop and evaluate a physics-informed ML-based surrogate model.

Control Applications:

CA1. ML-based controllers

- CA1.1. ML-based control system for Hall B polarized cryogenic target microwave frequency control.
- CA1.2. ML-based for continuous real-time tuning of the orientation of the Hall D diamond target.

C Proposed Research and Methods

C.a Polarized Cryogenic Target Optimization

The current proposal is to apply AI/ML in two parts to address the issue of continual control of the microwave frequency such that the polarization is optimized. The first will be to develop a model to extract polarization signal directly from the NMR system response. Identifying the current background from each signal dynamically will allow a clean extraction of the polarization without the need for frequent background measurements or polynomial fits to correct for signal drift. This will result in a higher quality signal value since it will effectively be subtracting the current background as opposed to one measured in the past. Many thousands of NMR signals from past JLab experiments are available to train models on both background electronic responses and polarization signals. We plan to start by leveraging a technique developed at the University of North Carolina using Cyclic positional U-Net[13]. The technique was developed to extract signals from germanium based detectors under this same funding program.

The second part of the effort will be the use of extracted NMR signals along with the historical data stream consisting of beam current (including trips), annealing events, and temperature to identify and set the optimal microwave frequency. Automating the necessary microwave frequency changes will remove the need for the effort of trained target operators and experts who must constantly monitor the polarization during running. It will also improve the final polarization by reducing the opportunities for user error and delays in user action.

A reach goal for this part of the project would be for the system to estimate the integrated $beam \times polarization$ at some point in the future to determine the optimal beam current. This would rely on effects such as a slightly lower beam current causing a slightly lower target temperature resulting in a higher polarization. Such a self-driving experiment could result in a data set with a similar number of signal events but a smaller background, improving the systematic uncertainties.

C.a.1 Timetable of Activities - Polarized Target

Duration	Activity
1 month	Identify and curate appropriate historical data sets of measured polarization. This will include the CLAS12 Run Group C archives.
3 months	Develop simulation of target polarization behavior based on historical archives that can be used for AI/ML model development.
1 month	Collect historical waveform data for NMR signal from polarized target and prepare for use in model training.
3 months	Implement signal extraction technique for accurate extraction of NMR signal from electronics background.
3 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.
3 months	Test model against simulation and adjust to optimize performance.
3 months	Utilize the Polarized Target Group's test facility to test the model and further refine it.
2 months	Integrate models and appropriate codes into the AI/ML controls ecosystem and deploy in Hall-B.
3 months	Improve the simulation accurately simulate high luminosity beam conditions in Hall-A for SoLID.
2 months	Duplicate and refine model developed for CLAS12 for use in SoLID.
24 months	**TOTAL**

C.b Hall D Polarized Photon Source

This project proposes to replace the manual tuning of the coherent peak done at the beginning of each run, with a continuous real-time tuning. The bremsstrahlung coherent peak will be adjusted through a control loop that leverages AI to minimize the number of steps needed to tune the peak within a precision of 10 MeV from the nominal peak. The proposed procedure will take into account the peak position, electron beam energy, electron beam positions and photon beam positions. This information will be converted into two motor motions based on the polarization plane, and six input parameters. The input data can be taken from the MYA EPICS archive for recent time periods when the beam current was non-zero using an average over a time interval.

In summary:

- This new automated procedure will result in a larger figure of merit for the experiment due to higher degree of linear polarization of the photon beam when integrated over time as well as in larger statistics, as the polarization and the coherent peak enhancement both have sharp drop offs on both sides of the coherent edge. Thus, we will avoid smearing the edge by not maintaining a constant the peak position. With a smeared edge, one would need to select the energy cuts for coherent peak range within the area that excludes the nominal position of the coherent edge itself, thus reducing the average polarization and the data sample. Therefore, this procedure will improve the figure of merit for all experiments in Hall D involving linearly polarized photon beam including GlueX.

- The fact that the peak position is smeared because of the lack of control of the goniometer angles also results in a larger systematic uncertainty for the measured polarization using the triplet polarimetry, which will directly translate into larger systematic uncertainties for the measured asymmetry observable.

A reach goal of the project will be to track and predict when the running experiment will benefit from moving the beam spot to a new location on the diamond. This will involve using the original x-ray scans of the diamond and a digital-twin like simulation to estimate the localized degradation of the diamond quality as a function of the beam dose, position, and measured coherent spectrum. The system would advise experimenters when such a change would benefit the experiment overall and the position of the diamond that should be moved into the beam.

C.b.1 Timetable of Activities - Polarized Source

Duration	Activity
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.
2 months	Identify "nudge" events and responses to build data set for training.
1 month	Perform Shapley analysis based on polarization FOM ($= pol^2 \times photonenergy$) to determine most relevant parameter set.
4 months	Develop and train model to predict polarization FOM based on available inputs.
3 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.
1 month	Develop safety policies for operation of the system. Interface with the EPICS alarm system.
1 month	Create outward facing monitoring pages for the system using Grafana or similar.
3 months	Create simulation of realistic operating conditions that includes regular beam trips, DAQ transitions and configuration changes.
1 month	Refine model and deployment to operate in continuous mode.
1 month	Port existing UConn beam spot finder tool webpage into format that can be used via command line interface.
4 months	Build digital twin like model of diamond based on existing x-ray scans. Map should be adjusted based on integrated beam spot size, and conditions during operation.
2 months	Implement automated monitoring system that uses input from digital twin to determine optimal location on diamond. Notify experimenters when new position should be used.
24 months	**TOTAL**

C.c Data Science Research and Methods

C.c.1 ML1: Develop A Model For NMR Signal Extraction.

ML1 requires the development of an ML model to extract a clean, near-real-time polarization signal from the raw NMR signal. The resonant circuit of the continuous-wave NMR system is coupled to the magnetization, and thus polarization, of the target material when the circuit is swept in frequency through the larmor frequency of the species of interest. The circuit itself contributes a background response versus frequency that must be removed to isolate the response from the material's polarization[14]. Historically, these backgrounds have required performing special measurements regularly throughout the experiment, and drifts in the circuit response over time required further correction, which have introduced error, particularly during real-time data-taking. Near-real-time NMR signal cleaning will directly extract the polarization for each signal, providing real value to Hall B and future polarized target experiments by reducing polarization error from background fits. The accurate extraction of the polarization in near-real-time is required for the control of the microwave frequency, and is thus the development of ML1 is critical to the implementation of ML2.1 and CA1.

We will leverage a signal extraction technique, the Cyclic Positional U-Net (CPU-Net)[13]. CPU-Net has demonstrated the ability to learn a detector pulse transformation without explicit programming of detector physics. CPU-Net is built on a 1-D U-Net[15] architecture, and we expect to employ CPU-Net to subtract the drifting electronic baseline with low latency.

C.c.2 ML2: Uncertainty-Aware Surrogate Models

The research team will use ML methods with well-calibrated[16], out-of-distribution (OOD)-aware uncertainties to develop surrogate models to explore experimental control policies. Previous work under LAB-20-2261[9] demonstrated successful control of drift chamber gain throughout multiple run periods using a Gaussian Process (GP)[17]. The GP provided well-calibrated uncertainty of predictions. High voltage control of the drift chamber was performed only when predictions were within an uncertainty threshold; otherwise, the drift chamber was returned to the traditional high voltage, and additional data was collected that was used for offline training of the GP to increase the in-distribution coverage of a re-trained GP. GP's are applicable in control scenarios with a small number of observations due to $\mathcal{O}(n^3)$ and a smaller number of features.

When the number of observations, features, or both exceeds the feasibility of a GP, techniques such as Spectral-normalized Neural Gaussian Process [18] provide the expressiveness of a Deep Neural Network (DNN) and the OOD uncertainty awareness required for our applications. This approach has been used in recent work for anomaly detection for Spallation Neutron Source at ORNL [19] and for a regression model for the Fermilab Booster Accelerator Complex [20].

Monte-Carlo (MC) Dropout [21] is an additional technique to estimate DNN prediction uncertainty and may provide a valuable comparison to GP-based uncertainty estimates. The MC dropout approach relies on introducing a tunable uncertainty parameter into a network training process by adding a dropout layer that randomly removes nodes from the following layer with a set probability at each forward pass. The MC dropout technique can overcome the computational challenges of estimating uncertainty using Bayesian models and can be used as a second method to compare predictions and estimated prediction uncertainties. The review by Abdar et.al [22] provides comprehensive details on estimating uncertainty in deep learning.

ML2.1: Polarized Target

We propose creating a data-driven surrogate model of the cryogenic target polarization using historical data, including thermal fluctuations, beam trips, target degradation, and time since target annealing. This surrogate model will be our offline, ML-based system to simulate the behavior of the target over time and as conditions change. Previously, the development of a Monte-Carlo simulation was in-progress[23]; however, the MC simulation has not yet been validated, is not expected to be differentiable, and the use of this simulation is expected to be relatively slow for training an AI-based controller.

We propose to develop the data-driven surrogate model of the target's polarization using the many thousands of available historical data files from Hall B experiments from 2003 to 2022. Even with the thousands of experimental data files, we do not expect complete coverage of the input feature space. We desire an ML method that recognizes data that is OOD of training data, and the OOD distance must be well-calibrated. Uncertainty awareness in an ML surrogate model is critical for safe and interpretable machine learning and informs the control policy when the predictions are no longer reliable, i.e., predictions with high uncertainty. Traditional ML methods poorly estimate uncertainties, especially OOD uncertainties.

ML2.2: Polarized Source

The primary application for this use case is to create a surrogate model of the polarization FOM system. We will use historical data to develop the uncertainty-aware surrogate model. This will provide a fast inference model that is needed to train the ML-based control models. Additionally, using an ML-based surrogate model ensure that it's differentiable since this is required to train the DMPC control model described in Section C.c.5.

We know that the diamond target degrades overtime, which needs to be modeled in order to properly capture the time-dependent performance of the system. In order to extract that from data, we propose to use data-driven discovery methods, such as the sparse identification of nonlinear dynamics (SINDy) [24] approach, to extract the dominate terms.

C.c.3 CA1. AI-based Controllers

Utilizing the ML-based surrogate models in Section C.c.2, we will exploit the uncertainty estimation techniques to enable uncertainty-aware AI-based controllers for both the polarized target (NP1) and polarized source (NP2). The uncertainty estimates provided in conjunction with the policy model's predictions determine how reliable the predictions are. Each application can determine its reliability threshold and the control behavior based on prediction uncertainty. Both NP1 and NP2 propose replacing manual adjustments by shift takers with AI-based control; however, with well-calibrated, distance-aware-UQ control policies, if experimental conditions stray into feature space that is OOD from training data, shift takers could be alerted and traditional, manual methods could resume until experimental conditions return to feature space with higher control policy certainty. Ultimately the data gathered during traditional/manual adjustment could be used as additional training data control policy, resulting in less OOD feature space.

Both NP1 and NP2 will learn control policies that may include physics constraints, take advantage of prediction uncertainties from the surrogate model, and provide control predictions with the associated estimated uncertainties.

For this proposal, we will explore the use of Deep Model Predictive Control (DMPC) and Deep Reinforcement Learning (DRL) which are two actively evolving and powerful control techniques. Each technique has strengths and weakness which can be leveraged to solve the technical problems in this proposal and will be discussed below.

Model predictive control (MPC) is a control method that uses a system model to predict the future actions of a system time horizon. MPC performs very well when the system model is accurate and there is an existing reference control target. By introducing a deep learning policy model (π) in the MPC workflow we can provide fast inference and the ability for continuous learning. In order to use DNN model for the policy model the system model must be differentiable to be able to back-propagate when training the policy model. It should be stressed that the policy developed by the DMPC controller is only as good as the system model and will need to be continuously refined using new data when the experimental configuration changes.

Reinforcement Learning (RL) is aimed at optimizing the control or planning of complex tasks by interacting and learning from the environment. The main components of RL are an environment and an agent, as shown in Figure 9. In this proposal, the environment encapsulates the experimental information (our surrogate models) that is relevant to the optimization task. The agent is composed of the control policy along with other critical tools that help guide the exploration of the environment and eventually train an optimal control policy. The learning process for the control policy is iterative and involves interacting between the agent and the environment. During the interaction at time step t, the agent receives the current observable state S_t from the environment and applies an action A_t. Based on this action, the environment updates its internal state and observation state S_{t+1}, and provides the actor with a quantitative value for how "good" the action was (aptly called the reward) R_{t+1}. The goal of the agent is to maximize the total reward over a

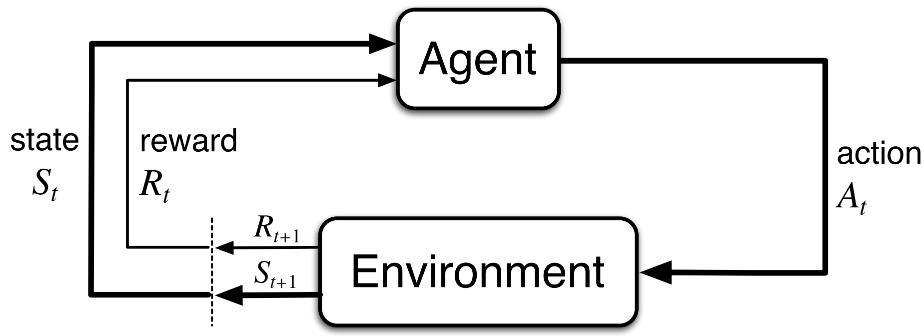

Figure 9: The agent-environment interaction in a reinforcement learning workflow [25]. The agent executes a policy that selects an action A_t given the current S_t, which results in a reward R_{t+1} and a new state S_{t+1} of the environment.

defined number of steps.

Within the RL algorithm, the total reward can be calculated using Bellman optimality equation as stated in Equation 2. Where $Q(S, A)$ is a long term expected reward (also called Q-value) for taking action A at state S, R_{t+1} is the instantaneous reward, γ is the discount factor that is used to weight long term reward, and $Q(S', A')$ is the maximum Q value that can be achieved at next state S'.

$$Q(S, A) = E\left[R_{t+1} + \gamma \max_{A'} Q(S', A')\right] \tag{2}$$

In Deep Reinforcement Learning (DRL), the agent is composed of ML-based (deep learning) model(s).

C.c.4 CA1.1 ML-based Control System For Hall-B Polarized Cryogenic Target Microwave Frequency Control.

The DMPC method requires a fully differentiable system model, and while [23] began the development of a simulation of the polarization over time as a function of microwave frequency, it is unclear if the system was developed to support backpropagation for the ML training workflow. Consequently, to address CA1.1, we propose to develop an offline DRL model to learn the microwave frequency control policy to optimize target polarization.

C.c.5 CA1.2. ML-based Continuous Real-time Tuning Of The Orientation Of The Hall-D Diamond Radiator.

For this use case there is a physics system model [26], [27] and known reference targets; therefore, we will initially explore the use of DMPC method. DMPC provides a simpler solution to the DRL since it has additional information.

C.c.6 MLOps, Monitoring, and Controls

An accurate and robust monitoring system is required for an ML system to sense its environment and take appropriate actions in given situations. It requires access to all the quantities needed for decision-making for control and to know when control has been successfully executed. We will exploit the Experimental Physics and Industrial Control System (EPICS), the MLFlow centralized model repository, and Granfana dashboards; all are already in use at Jefferson Laboratory. EPICS provide access to the parameters that models would need and access to controls via well-established APIs. MLFlow provides API access to published ML models, and Grafana is currently used for real-time monitoring of ML implementations. Examples of Grafana's use at Jefferson Lab for ML monitoring are for the drift chamber control system completed under LAB-20-2261[9] and for *Hydra*[28], an ML system for data quality monitoring, currently used in three of Jefferson Lab's experimental halls.

C.c.7 Novelty and Impact

Although we expect the number of features to be few, ten to twenty, the number of samples is large, on the order of thousands, making a GP computationally impractical. A GP with well-calibrated uncertainty estimates would be ideally suited for real-time control for both the Polarized Cryogenic Target Optimization and Hall D Polarized Photon Source. An uncertainty-aware surrogate model of Fermilab's Booster Accelerator Complex [29] demonstrated an out-of-distribution aware deep learning model by adding a radial basis function kernel to the final layer of a deep neural network (DNN). Distance awareness between input and DNN output was preserved using a bi-Lipshitz constraint as part of the training loss function. Model uncertainty calibration was confirmed using the Uncertainty Toolkit [16], which can also be utilized for after-the-fact calibration if required. Methods, such as DGPA, that address the importance of uncertainty in real-time experimental controls and consider the requirement for a large number of examples to learn control policies will impact both data science and the nuclear physics program.

Appendix 1: BIBLIOGRAPHY & REFERENCES CITED

References

[1] D. G. Crabb and W. Meyer, "Solid polarized targets for nuclear and particle physics experiments," *Annual Review of Nuclear and Particle Science*, vol. 47, no. 1, pp. 67–109, 1997. DOI: 10.1146/annurev.nucl.47.1.67.

[2] C. Keith, M. Anghinolfi, M. Battaglieri, *et al.*, "A polarized target for the clas detector," *Nuclear Instruments and Methods in Physics Research A*, vol. 501, pp. 327–339, 2003. DOI: https://doi.org/10.1016/S0168-9002(03)00429-7.

[3] A. Liyanage, W. Armstrong, H. Kang, *et al.*, "Proton form factor ratio $\mu_p G_E^p / G_M^p$ from double spin asymmetry," *Phys. Rev. C*, vol. 101, p. 035 206, 3 Mar. 2020. DOI: 10.1103/PhysRevC.101.035206. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevC.101.035206.

[4] J. Pierce, J. Maxwell, and C. Keith, "Dynamically polarized target for the g_2^p and G_E^p experiments at jefferson lab," *Physics of Particles and Nuclei*, pp. 303–304, 2014. DOI: https://doi.org/10.1134/S1063779614010808.

[5] P. Anthony, R. Arnold, T. Averett, *et al.*, "Precision measurement of the proton and deuteron spin structure functions g2 and asymmetries a2," *Physics Letters B*, vol. 553, no. 1, pp. 18–24, 2003, ISSN: 0370-2693. DOI: https://doi.org/10.1016/S0370-2693(02)03015-0. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0370269302030150.

[6] D. Lawrence, R. Miskimen, E. S. Smith, and A. Muskarenkov, "Measuring the charged pion polarizability in the $\gamma\gamma \to \pi^+\pi^-$ reaction," *PoS*, vol. CD12, p. 040, 2013. DOI: 10.22323/1.172.0040.

[7] O. Hen, M. Patsyuk, E. Piasetzky, *et al.*, *Studying short-range correlations with real photon beams at gluex*, 2020. DOI: 10.48550/ARXIV.2009.09617. [Online]. Available: https://arxiv.org/abs/2009.09617.

[8] KLF Collaboration, M. Amaryan, M. Bashkanov, *et al.*, *Strange hadron spectroscopy with secondary kl beam in hall d*, 2020. DOI: 10.48550/ARXIV.2008.08215. [Online]. Available: https://arxiv.org/abs/2008.08215.

[9] T. Jeske, D. McSpadden, N. Kalra, T. Britton, N. Jarvis, and D. Lawrence, "AI for experimental controls at jefferson lab," *Journal of Instrumentation*, vol. 17, no. 03, p. C03043, Mar. 2022. DOI: 10.1088/1748-0221/17/03/c03043. [Online]. Available: https://doi.org/10.1088%2F1748-0221%2F17%2F03%2Fc03043.

[10] J. Maxwell, "Probing proton spin structure: A measurement of g_2 at four-momentum transfer of 2 to 6 gev," https://doi.org/10.18130/V3526T, PhD, University of Virginia, Dec. 2011.

[11] H. Al Ghoul, E. G. Anassontzis, F. Barbosa, *et al.*, "First results from the gluex experiment," *AIP Conference Proceedings*, vol. 1735, no. 1, p. 020 001, 2016. DOI: 10.1063/1.4949369. eprint: https://aip.scitation.org/doi/pdf/10.1063/1.4949369. [Online]. Available: https://aip.scitation.org/doi/abs/10.1063/1.4949369.

[12] K. Livingston, "The Stonehenge technique. A method for aligning coherent bremsstrahlung radiators," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 603, no. 3, pp. 205–213, 2009.

[13] A. Li, J. Gruszko, B. Bos, T. Caldwell, E. León, and J. Wilkerson, *Ad-hoc pulse shape simulation using cyclic positional u-net*, 2022. DOI: 10.48550/ARXIV.2212.04950. [Online]. Available: https://arxiv.org/abs/2212.04950.

[14] J. Maxwell, W. Armstrong, S. Choi, *et al.*, "Design and performance of the spin asymmetries of the nucleon experiment," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 885, pp. 145–159, 2018, ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2017.12.008. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168900217313621.

[15] O. Ronneberger, P. Fischer, and T. Brox, *U-net: Convolutional networks for biomedical image segmentation*, 2015. DOI: 10.48550/ARXIV.1505.04597. [Online]. Available: https://arxiv.org/abs/1505.04597.

[16] Y. Chung, I. Char, H. Guo, J. Schneider, and W. Neiswanger, "Uncertainty toolbox: An open-source library for assessing, visualizing, and improving uncertainty quantification," *arXiv preprint arXiv:2109.10254*, 2021.

[17] C. E. Rasmussen and W. C. K. I., *Gaussian processes for machine learning.* MIT Press, 2008.

[18] J. Liu, Z. Lin, S. Padhy, D. Tran, T. Bedrax Weiss, and B. Lakshminarayanan, "Simple and principled uncertainty estimation with deterministic deep learning via distance awareness," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33, Curran Associates, Inc., 2020, pp. 7498–7512. [Online]. Available: %7Bhttps://proceedings.neurips.cc/paper/2020/file/543e83748234f7cbab21aa0ade66565f-Paper.pdf%7D.

[19] W. Blokland, K. Rajput, M. Schram, T. Jeske, P. Ramuhalli, C. Peters, Y. Yucesan, and A. Zhukov, "Uncertainty aware anomaly detection to predict errant beam pulses in the oak ridge spallation neutron source accelerator," *Physical Review Accelerators and Beams*, vol. 25, no. 12, Dec. 2022. DOI: 10.1103/physrevaccelbeams.25.122802. [Online]. Available: https://doi.org/10.1103%2Fphysrevaccelbeams.25.122802.

[20] M. Schram, K. Rajput, K. S. P. Li, J. S. John, and H. Sharma, *Uncertainty aware ml-based surrogate models for particle accelerators: A study at the fermilab booster accelerator complex*, 2022. DOI: 10.48550/ARXIV.2209.07458. [Online]. Available: https://arxiv.org/abs/2209.07458.

[21] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *Proceedings of The 33rd International Conference on Machine Learning*, M. F. Balcan and K. Q. Weinberger, Eds., ser. Proceedings of Machine Learning Research, vol. 48, New York, New York, USA: PMLR, Jun. 2016, pp. 1050–1059. [Online]. Available: https://proceedings.mlr.press/v48/gal16.html.

[22] M. Abdar, F. Pourpanah, S. Hussain, *et al.*, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," *Information Fusion*, vol. 76, pp. 243–297, 2021, ISSN: 1566-2535. DOI: https://doi.org/10.1016/j.inffus.2021.05.008. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1566253521001081.

[23] G. Darshana-Perera, "Automated microwave frequency control of dynamically polarized targets," 22nd International Spin Symposium, Sep. 2016.

[24] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Sparse identification of nonlinear dynamics with control (sindyc)**slb acknowledges support from the u.s. air force center of excellence on nature inspired flight technologies and ideas (fa9550-14-1-0398). jlp thanks bill and melinda gates for their active support of the institute of disease modeling and their sponsorship through the global good fund. jnk acknowledges support from the u.s. air force office of scientific research (fa9550-09-0174).," *IFAC-PapersOnLine*, vol. 49, no. 18, pp. 710–715, 2016, 10th IFAC Symposium on Nonlinear Control Systems NOLCOS 2016, ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2016.10.249. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896316318298.

[25] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[26] R. Jones, *Hall D Coherent Bremsstrahlung Rate Calculator*. [Online]. Available: http://zeus.phys.uconn.edu/halld/cobrems/ratetool.cgi.

[27] R. Jones, *Diamond radiator spot finder tool*. [Online]. Available: https://cn410.storrs.hpc.uconn.edu/spotfinder.

[28] Britton, Thomas, Lawrence, David, and Rajput, Kishansingh, "AI enabled data quality monitoring with hydra," *EPJ Web Conf.*, vol. 251, p. 04 010, 2021. DOI: 10.1051/epjconf/202125104010. [Online]. Available: https://doi.org/10.1051/epjconf/202125104010.

[29] M. Schram, K. Rajput, K. S. NS, P. Li, J. St. John, and H. Sharma, "Uncertainty aware ml-based surrogate models for particle accelerators: A study at the fermilab booster accelerator complex," *TBD*, Sep. 2022. [Online]. Available: https://www.osti.gov/biblio/1892259.

Appendix 2: FACILITIES

0.1 Thomas Jefferson National Accelerator Facility

The proposed research will be carried out at Jefferson Lab and the partner institutions: The College of William and Mary and Carnegie Mellon University. The major equipment on which this work will be based is part of the standard Jefferson Lab facilities. These are listed in the tables below. The fraction of the listed resources needed for this project are modest (<1%) and will not disrupt current operations at the lab.

The research team will have access to the major equipment when it is not actively in use by a running experiment. All members of the research team currently have or will have access to a portion of the computational resources of the IT division of Jefferson Lab. For the needs of the proposed research the research team will have ready access to the dedicated machine learning nodes, ample CPU compute power, and sufficient storage (several TeraBytes).

Additionally, the research team will have access to the existing historical data archives of all four experimental halls for the purposes of this project.

Compute Resources

- GPU-enabled Machine Learning worker nodes
- 8 NVIDIA A100 GPUs (80GB)
- 56 NVIDIA T4 GPUs
- Equivalent of 9000 cores of AMD Milan CPUs drawing from several generations of Intel and AMD CPUs
- LQCD compute clusters including
- 400 KNL nodes cluster (OPA fabric)
- 256 NVIDIA 2080 GPU cluster (OPA fabric)
- 64 AMD MI100s GPU cluster (InfiniBand EDR fabric)

Test Stand Hardware Resources

- R&D computing resources in the TJNAF data center and associated lab spaces in CEBAF Center with fiber optic connectivity between them
- Access to TJNAF compute cluster for testing at scale
- 6 Compute Servers with FPGA and NVIDIA dual-port ConnectX-6Dx and 100Gbit Ethernet NIC
- File server with multiple NVMe drives, FPGA and NVIDIA dual-port ConnectX-6 DX and 100Gbit Ethernet NICs
- Arista 400Gbit capable Ethernet switches
- Access to TJNAF streaming data sources and NP experimental data as test data source
- The NVIDIA ARM HPC Development Kit
- High speed networking including EDR InfiniBand and 400Gbit Ethernet

Storage capabilities for varying IO profiles

- 6PB in two Lustre Parallel filesystems
- 5PB NFS and XRootD storage on ZFS
- A tape library capable of over 100PB of long term storage
- Burst disk storage for tape buffering at ¿10 gbps

Data center and Wide area Networking tuned for scientific workflows

- Internet-facing Science DMZ network for fast off-site data transfer including Globus and XRootD Data Transfer Nodes; jumbo frame support
- Dual 10Gbit internet connections. An upgrade to dual 100Gbit connections is in process with ESNet
- Data Center EDR Infiniband and OPA fabrics
- Data center 100Gbit Ethernet redundant routed core network with support for virtual machines on an ESX cluster

Software Infrastructure

- Jupyter Notebook environment with GPU capabilities
- Slurm Fairshare job scheduling for all clusters
- MPI and parallel job support
- Complex batch workflow orchestration tooling (SWIF)
- Full Open Science Grid Support including HTCondor
- Federated identity support for remote users including InCommon, SciTokens, Shibboleth, COManage, and CILOGON.

Data Center Capabilities

- Rack Space, power, and cooling for hardware installs including
- Hot aisle containment for air cooling
- 3 Phase 208VAC power to racks (typical: 2x50A @ 208V)
- UPS, generator, automatic power transfer to two substations

Appendix 3: EQUIPMENT

0.2 Thomas Jefferson National Accelerator Facility

Major equipment items such as the polarized targets and photon sources already exist. The only equipment needed that is not already available will be personal computers for the new hires which will be purchased using funds from this proposal.

Appendix 4: DATA MANAGEMENT PLAN

The proposed work involves no sensitive or personally identifiable information. As such, this work will follow the DOE Office of Science Statement on Digital Data Management, as well as the guidelines for data management set forth by Jefferson Lab's Scientific Computing Division.

0.3 Data Sources

Most of the data necessary for this work will come directly from or be derived from the quantities stored in the EPICS archive system at Jefferson Lab. Additionally, data from feedback systems, such as those required to measure target polarization, will be required. These data will be stored in appropriate forms in keeping with best-practice and established systems. For example, data stored in a database in the case of the EPICS archive and histograms derived from root files in the case of measured quantities as produced via an analysis.

0.4 Types of data

Raw data

Raw data for use in the proposed work will be primarily stored in human readable formats as well as containerized formats such as ROOT histogram files, h5 formatted files, etc.

Meta data

Meta data involved in analysis or scientific record keeping may involve text files, root files, or other saved file formats (eg JSON) and will be sufficient to provide repeatability of analyses.

0.5 Data Sharing

Believing that the sharing of data fosters innovation and advancement all data used in the course of this research will be made public. Produced codes will be preserved in an open GitHub repository. Any models produced along with their metadata will be permanently archived on magnetic tape at Jefferson Lab. Items stored as "production" in the archive are automatically copied to a second backup tape that is stored in a protected vault ensuring preservation in the event of catastrophe. Other documents generated and classified in the course of this project will be shared among collaborators from the start of the project, and will be continually upgraded and updated. E-mail will be the most common method to share information. Prior to making them available for wider distribution, such as through journal publications or conference presentations, applications for the protection of Intellectual Property would be filed, as agreed upon by all collaborators.

0.6 Data Preservation

This research has no need for any personally identifiable information (PII), nor does it have the need for any sensitive information related to national security interests or any other interests outside of the advancement of the scientific objectives of the Department of Energy (DOE). Any data gathered or produced in the course of this research will thus be stored in an unencrypted, human readable (where appropriate), format suitable for the training and/or testing of the models produced in the course of this research. Any codes developed during this research will be version controlled by and stored on github.com. Upon completion of the research the code will be made open source and released, with suitable documentation to the public. In this way future projects may benefit from the models and code developed in pursuit of the proposed research.

All the electronic files generated for this project will be preserved according to Jefferson Lab's record management policy (https://www.jlab.org/div_dept/cio/IR/records/index.html) and DOE's Research and Development Records Schedule (http://www.archives.gov/recordsmgmt/rcs/schedules/departments/department-of-energy/rg-0434/n1-434-08-002_sf115.pdf).

0.7 Data Validation

All data used in analysis, model development, or other related activities will be validated before use or publication. This validation may include expert review as well as checks for consistency. These consistency checks may include redundant backups with checksums as well as checks like those performed in continuous integration.

Appendix 5: PROMOTING INCLUSIVE AND EQUITABLE RESEARCH (PIER) PLAN

The proposal team includes staff from Jefferson Lab, Carnegie Mellon University and William & Mary. William & Mary is classified as a R2 university. Over 30% of the team identify as women, and one of these is a member of the GlueX collaboration's Diversity and Inclusion subcommittee. The graduate students for this project will be hired by the universities. The postdocs will be hired by Jefferson Lab and William & Mary.

0.8 Outreach Strategies

The proposal team has integrated several students from regional universities, such as Old Dominion University, William & Mary, and the University of Virginia (UVA), into data science and accelerator projects at Jefferson Lab. In addition, the proposal team has provided opportunities to UVA's Master of Data Science students in the form of real world research problems for their Capstone projects, and is mentoring the student researchers.

We propose to organize a data science workshop for local middle schools. Adolescent years are a critical stage when role models, peer and societal pressures form an influence which can be critical to the student's eventual career path (see `https://doi.org/10.1002/sce.21492` and `https://journals.aps.org/prper/pdf/10.1103/PhysRevSTPER.5.010101`). Jefferson Lab's Activities Group already has a strong relationship with local schools. We will form a partnership between this network and our project team to organise an afternoon 'hackathon' workshop, where the students write their own code to explore data science techniques and solve real life problems. The project team will lead the workshop in person.

Hampton University Graduate Studies (HUGS), a summer session at Jefferson Lab, is for experimental and theoretical nuclear and particle physics graduate students who have completed their studies and have at least one year of research experience. Jefferson Lab's data science department has presented talks at HUGS on distributed computing and an overview of data science techniques. We propose to expand our contribution by providing a mini-workshop within HUGS focused on teaching introductory techniques used in this proposal. The curriculum will include hands-on tutorials on building surrogate models and applications of control techniques.

0.9 Development & Mentoring

The postdocs and graduate students joining the project team will be mentored by project staff. Mentors are required to complete the ethics and mentorship course 'Mentoring at Jefferson Lab' that addresses trust, respect, and open communication. Students are required to complete the course 'Fostering a Positive and Respectful Workplace' that teaches cultural norms, community standards, and ethics. Students are also invited to use their voice, participating in networking events promoted across the campus to foster inclusion.

Opportunities for professional growth include presenting at workshops and conferences and publishing our work. Junior staff within the project team receive priority for conference attendance, and the publications are prepared collaboratively. Conference presentation statistics are noted and reviewed to ensure parity. The team also has a regular series of meetings where staff present new works from the literature and lead an informal discussion; the leadership of each meeting rotates throughout the group.

0.10 Diversity, Equity, Inclusion and Accessibility (DEIA) at Jefferson Lab

Everyone at Jefferson Lab has a responsibility to foster an environment where all employees, users, students, guests, visitors, and subcontractors feel safe, welcomed and supported in advancing the Lab's mission. Expectations are outlined in Jefferson Lab's Community Standards as well as the Lab's Standards of Conduct & Code of Ethics. Jefferson Lab actively promotes a diverse and harassment-free experience for all. Jefferson Lab contracts with an independent reporting service to provide an Ethics Hotline - a risk-free way to anonymously and confidentially report concerns that may involve unsafe, fraudulent, unethical, or otherwise inappropriate behavior.

Jefferson Lab partners with Circa, a diversity recruiting network, to target job boards and organizations for women, minorities, and individuals with disabilities. DEIA is integrated in all recruiting cycle policies, procedures and practices.

0.11 Carnegie Mellon University

Naomi Jarvis is a Research Scientist in the Department of Physics at Carnegie Mellon University.

Carnegie Mellon University has multiple college and departmental initiatives that the senior investigators use to promote inclusive and equitable research within their research groups.

Our graduate admissions committee receives instructions from our College of Science's Associate Dean for Diversity, Equity, and Inclusion (DEI) on selection practices to promote a diverse graduate student population. The College and the Department participate in conferences and other targeted opportunities to advertise our graduate program to members of underrepresented groups, to expand the applicant pool. To serve on faculty search committees, members of the faculty must attend training on equitable search practices offered by the College. Similarly, faculty members hiring postdocs may be advised on those practices by the College on a voluntary basis.

There are also numerous groups at CMU dedicated to providing a supportive environment to members of our community from a variety of underrepresented groups; these include not only support groups, but also training services such as Tartan Allies, which provides regular training to university members committed to inclusion of LGBTQ+ members of the community.

Our department participates in the American Physical Society Inclusion, Diversity, and Equity Alliance (IDEA), which facilitates sharing of experiences and best practices between departments that seek to improve DEI. This year the Department's focus in APS IDEA is on developing spaces where all members, including those from underrepresented groups, feel a sense of belonging and inclusion.

This and other DEI initiatives are driven by the department's DEI committee, which engages actively with the Department including an anonymous feedback form used to identify issues for resolution. Department trust in and engagement with this committee is essential especially for early career department members, who may not feel empowered to openly speak up in situations with tricky power dynamics. Our Department also has a group, 'Constructive Interference', dedicated to supporting women and minorities in physics.

The Department has adopted a shared leadership model for some committees, providing opportunities for graduate students and postdocs to serve, which increases their sense of belonging while also providing them with training opportunities for their future careers. Graduate program processes are documented in a handbook that is used to communicate expectations clearly to students and faculty. The Department also has a graduate program committee that works to ensure departmental awareness of those processes meant to support the training and professional development of the students. These processes include the early formation of a thesis committee and annual reviews to ensure progress towards a degree and provide students with opportunities to form mentoring relationships with faculty besides their direct supervisors.

0.12 William & Mary

Cristiano Fanelli is an Associate Professor in the recently formed Data Science Department and also affiliated with the Department of Applied Science at William & Mary (W&M).

The Department of Applied Science recognises gender parity as an important aspiration and is committed to recruiting highly qualified women with interdisciplinary scientific foci. For all new hires, the department has set up multiple mentoring relationships within and outside of the department for its diverse faculty to have multiple support structures. Furthermore, the department has advocated for maximally advantageous start-up packages for our new hires, particularly women. The Department's DEI workplan has three outcome-driven specific aims: (1) achieve and maintain gender parity in our tenured core faculty, (2) increase the number of tenured Black core faculty members, and (3) inculcate inclusive ideals among our students, postdocs and research staff.

Part of W&M's overall mission is to cultivate creative thinkers, principled leaders, and compassionate global citizens equipped for lives of meaning and distinction. W&M's statement of values includes creating a welcoming and caring community that embraces diverse people and perspectives, commitment to the highest ethical standards, and treating one another with mutual respect, recognizing and upholding each person's inherent dignity and worth.