

Particle identification in the GlueX detector using a neural network

E. Habjan^{a,b} R. Dube^a R. Jones^a

^a*University of Connecticut, Department of Physics,
196A Auditorium Road, Unit 3046, Storrs, CT 06269, USA*

^b*Northeastern University, Department of Physics,
100 Forsyth Street 111, Boston, MA 02115, USA*

E-mail: habjan.e@northeastern.edu

ABSTRACT: Accurate Particle Identification (PID) is a crucial element for successful reconstruction of interactions measured in particle physics experiments. In the GlueX experiment at Jefferson Laboratory, PID is achieved by making cuts on the kinematic properties of tracks and showers reconstructed from hits in the detector. However, in this work we seek to improve upon these simple cuts-based procedures using machine learning with neural networks. The promise of this approach is the ability to exploit hidden correlations between PID variables in the reconstructed kinematics data. We demonstrate that both charged and neutral particles can be identified in simulated GlueX events with significantly improved accuracy using a neural network.

KEYWORDS: Particle identification methods, Analysis and statistical methods

Contents

1	Introduction	1
2	Monte Carlo Simulation Dataset	2
3	DNN model description	4
3.1	Cross Entropy Loss Function	4
3.2	Adam Optimizer	4
3.3	Structure and Activation Functions	4
3.4	Hyperband	5
4	Methods	5
4.1	Manual PID	5
4.2	Neural Network PID	8
5	Results	9
5.1	Comparing PID techniques	9
5.2	Feature Importance with Shapley Values	10
6	Conclusion	10

1 Introduction

A long-standing question in modern physics is concerning the nature of the quark confinement mechanism of quarks and gluons in quantum chromodynamics (QCD). To begin unraveling this mystery, the GlueX experiment in Hall D at Jefferson Laboratory aims to better understand the quark and gluonic degrees of freedom that are predicted by QCD through gluonic excitations of hybrid mesons originating from $\gamma - p$ collisions [1, 2]. A comprehensive study of this spectrum of hybrid mesons requires accurate determination of the final hadronic states through good particle identification (PID). The PID capabilities of the GlueX detector is enabled by four detector systems: a forward time-of-flight wall (TOF), a barrel calorimeter with a cylindrical geometry (BCAL), a forward lead-glass calorimeter with a planar geometry (FCAL), a thin scintillator start counter (SC), and a Čerenkov detector [3, 4]. GlueX uses time-of-flight (TOF and BCAL), forward going tracks (Čerenkov detector), and dE/ds (CDC) information to classify particles using a likelihood-based PID system in combination with kinematic fitting.

In this work, the aim is to demonstrate that PID accuracies for GlueX analyses can be improved through the use of machine learning (ML) [5]. The advent of ML has led to innumerable studies in a wide range of fields, including a GlueX study [6] that used a discriminating neural network (NN) to reduce background in the FCAL. Motivated by such studies and the demand for high accuracy PID,

NNs are used to classify particles using GlueX simulation data and directly compared to 'manual' PID methods using the same simulation dataset. This work exhibits that ML models can be a more effective alternative to the contemporary methods of PID.

2 Monte Carlo Simulation Dataset

Both the training and test data are extracted from the low-momentum GlueX particle gun simulations. In these simulations, a particle is spawned at a random location within the target and fired in a random direction with a random magnitude of momentum under 1 GeV/c. The interactions between the particle and the detector (along with any decays that may occur) are handled by Geant4 [7, 8], with the simulated detector hits being stored in a Hall D Data Model (HDDM) format. These data are then reconstructed using the `hald_recon` package to identify showers and tracks, which significantly decreases the number of features needed to describe each event. Finally, these data are converted from the hierarchical HDDM format to a tabular format that can be used in either manual and NN PID. The labels of each quantity in our final dataset are shown in Table 1.

Although the particle gun simulations allow for easy event labeling, decays and other interactions in the detector may produce tracks or showers that are not produced directly by the generated particle. To exclude events in which the generated particle decayed before interacting with the detector, we remove any event that has more than one vertex in the first 500 seconds of the simulation, as indicated by the truth information of the Monte Carlo simulation. Note that the initial spawning of the generated particle is counted as a vertex. To eliminate events where interactions with the detector produced secondary tracks or showers, cuts were placed on the number of tracks and showers per event. For events with a charged generated particle, the event was only included in the training and test samples if the reconstructed event contained exactly one track and one shower, and the shower must be associated with the track. Events with neutral generated particles must have exactly one shower and no tracks. These cuts are necessary to ensure the event label matches the particle that produced the shower or track that is included in the training or test dataset, though it may inflate the accuracy of PID techniques due to the exclusion of complicated interactions with the detector.

The training dataset consists of 80,000 events per particle type. For events with a charged generated particle, only the track hypothesis that matches the generated particle type is included in the training dataset. For events with muons, only the pion track hypothesis was added to the training dataset, as there is no muon hypothesis in the default reconstruction. This results in a training dataset in which each row of the dataset represents a different event. In contrast, a row corresponding to each hypothesis is added to the test dataset for events with charged generated particles; the event number is identified by the `group` label, which only appears in the test dataset. For neutral particles, there is only one row in the test dataset per event, as no hypotheses are used in the shower reconstruction process. The test dataset contains 40,000 events per particle type, though the number of rows is substantially larger due to the inclusion of multiple hypotheses per event for charged generated particles.

Table 1: Feature labels of the particle gun dataset.

Column	Unit	Description	Overflow Value
true_ptype		The true generated particle type (Geant3 coding)	
ptype		Particle hypothesis (Geant3 coding)	
group		Event number	
E	<i>GeV</i>	Particle total energy	-5
px	<i>GeV/c</i>	Particle momentum X-component	-500
py	<i>GeV/c</i>	Particle momentum Y-component	-500
pz	<i>GeV/c</i>	Particle momentum Z-component	-500
q	<i>e</i>	Particle charge	-10
E1E9		E1/E9 ratio for the matched FCAL cluster	-5
E9E25		E9/E25 ratio for the matched FCAL cluster	-5
docaTrack	<i>cm</i>	Impact parameter of track to FCAL cluster	-5
preshowerE	<i>GeV</i>	Shower energy in the 1st layer of the BCAL	-5
sigLong	<i>cm</i>	RMS of BCAL shower along depth	-5
sigTrans	<i>cm</i>	RMS of BCAL shower along azimuth	-5
sigTheta	<i>rad</i>	RMS of BCAL shower along Z	-5
E_L2	<i>GeV</i>	Shower energy in the 2nd layer of the BCAL	-5
E_L3	<i>GeV</i>	Shower energy in the 3rd layer of the BCAL	-5
E_L4	<i>GeV</i>	Shower energy in the 4th layer of the BCAL	-5
dEdxCDC	<i>keV/cm</i>	Average dE/ds of track in the CDC	-5
dEdxFDC	<i>keV/cm</i>	Average dE/ds of track in the FDC	-5
tShower	<i>ns</i>	Mean shower time in the BCAL or FCAL	-10
thetac	<i>rad</i>	Track Cerenkov angle measured by DIRC	-5
bCalPathLength	<i>cm</i>	Track distance from vertex to BCAL entry	-5
fCalPathLength	<i>cm</i>	Track distance from vertex to FCAL entry	-5
dEdxTOF	<i>keV/cm</i>	Average track dE/ds in the TOF	-5
tofTOF	<i>ns</i>	Time from track vertex to impact on the TOF	-5
pathLengthTOF	<i>cm</i>	Distance from track vertex to impact on the TOF	-5
dEdxSc	<i>keV/cm</i>	dE/ds of track in the SC	-5
pathLengthSc	<i>cm</i>	Distance from track vertex to impact on the SC	-100
tofSc	<i>ns</i>	Time from track vertex to impact on the SC	-100
xShower		Shower X-component	-500
yShower		Shower Y-component	-500
zShower		Shower Z-component	-500
xTrack		Track X-component	-500
yTrack		Track Y-component	-500
zTrack		Track Z-component	-500
CDChits		Number of straws in the CDC producing hits	-5
FDChits		Number of anode wires in the FDC producing hits	-5
DOCA	<i>cm</i>	Impact parameter of track at the BCAL cluster	-5
deltaz	<i>cm</i>	Impact parameter of track at the BCAL along Z	-100
deltaphi	<i>rad</i>	Impact parameter of track at the BCAL along azimuth	-10
tFlightSc	<i>ns</i>	Calculated time from vertex to SC	
tFlightBCAL	<i>ns</i>	Calculated time from vertex to BCAL	
tFlightTOF	<i>ns</i>	Calculated time from vertex to TOF	
tFlightFCAL	<i>ns</i>	Calculated time from vertex to FCAL	

3 DNN model description

In this section we explain and justify the TensorFlow implementations of the Adam optimizer, the cross entropy loss function, the activation functions and Hyperband optimization.

3.1 Cross Entropy Loss Function

The advent of logistic regression by [9] and the creation of the idea of cross-entropy in the early years of information theory has evolved into a loss function that ubiquitous in machine learning: the cross entropy loss function. In general, the minimization of cross entropy between two distributions is equivalent to the maximization of the log likelihood [10]. The log likelihood can be defined as:

$$l(\theta) = \frac{1}{N} \sum_{i=1}^N \log(P(x_i|\theta)). \quad (3.1)$$

Where x_i is a given detection (or in the case of this article, a particle) and θ defines our parameter space (e.g., energy loss, momenta). By maximizing the log likelihood, we can best predict the probability of detecting a given x_i when provided with θ . Also defined in terms of the probability of x_i and θ , the cross entropy $H(P_D(x), P_\theta(x))$ is defined as:

$$H(P_D(x), P_\theta(x)) = -\frac{1}{N} \sum_{i=1}^N \log(P(x_i|\theta)) = -l(\theta). \quad (3.2)$$

Thus, we see that maximizing the log likelihood is equivalent to minimizing the cross entropy; otherwise known as the cross entropy and maximum likelihood principle. In our DNN, we utilize the TensorFlow implementation of the cross entropy loss function, which calculates the cross entropy loss between our dataset θ and the generated particle x_i . The minimization of the cross entropy loss is made to be the objective of our DNNs, and the optimization process is described in Section 3.4.

3.2 Adam Optimizer

As stated in Section 3.1, we make minimizing the cross entropy loss function the objective of our machine learning problem. To do this, we implement the Adam optimizer [11]. The Adam optimizer is a method to efficiently optimize the subfunctions that make up the entire objective function by taking gradient steps with respect to each individual subfunction. This process also describes Stochastic Gradient Descent, however the Adam optimizer is particularly designed to optimize parameters for stochastic subfunctions in high dimensional space, while only requiring first-order gradients. The large level of stochasticity in measured quantities within particle colliders makes the Adam optimizer an ideal choice for minimizing the cross entropy loss function.

3.3 Structure and Activation Functions

The structure of our neural network is comprised of an input layer, one or multiple hidden layers (determined via optimization in Section 3.4), and an output layer. The input layer of our models is

comprised of 38 nodes, which is equal to the number of feature labels shown in Table 1. In each of the hidden layers of our neural network we add the Rectified Linear unit (ReLU) activation function [12, 13]. The ReLu activation function is described as

$$f(x) = \max(0, x), \quad (3.3)$$

so that for any input x from a previous neuron, a non-negative output $f(x)$ will be produced from that neuron. The non-linearity of ReLu introduces sparsity and avoids saturation at large values, while also remaining simple. These advantages allow for efficient training and for meaningful connections to be drawn between complex relationships in the data. In our output layers, we use the sigmoid activation function, shown in Equation 3.4.

$$S(x) = \frac{1}{1 + e^{-x}}. \quad (3.4)$$

With an input x , the output $S(x)$ will always be between 0 and 1, which makes the sigmoid function effective for PID.

3.4 Hyperband

In order to minimize the Cross Entropy Loss function, we must find optimized values for the number of hidden layers, the number of neurons in each hidden layer, and the learning rate of the Adam optimizer. In this work we find the optimized values of each of these hyperparameters using Hyperband [14]. We chose to employ Hyperband as the optimization algorithm since it is more computationally efficient and performs better than Bayesian optimization [14]. Hyperband selects a different set of hyperparameters and trains the neural network for a fixed number of epochs. A method known as *successive halving* is then utilized, which removes half of the models with the largest Cross Entropy loss. This procedure is repeated until only a single set of hyperparameters remains; these hyperparameters are then used for training.

4 Methods

In this section we describe our manual PID cuts and the training process for our neural network models used for PID.

4.1 Manual PID

In this work, we identify pions ($\sim 139.6 \text{ MeV}/c^2$) and muons ($\sim 105.7 \text{ MeV}/c^2$) as the same particle, denoted as $\pi^+ | \mu^+$ or $\pi^- | \mu^-$ for the positive and negative counterparts, respectively. This simplification is made for our PID methods in this paper since pions and muons have similar masses and the GlueX detector does not have a hadronic calorimeter, which makes discerning these particles particularly difficult. Muons and Pions can be distinguished by investigating the momentum distributions of a given event in the FCAL, however this must be done prior to reconstruction, which is outside the scope of this work.

The timing cuts we implement in this make use of the Spring 2017 Analysis Launch Cuts [1]; each of these cuts are shown in Table 2. The measured BCAL and FCAL times are recorded as a single variable in our dataset: `tShower`. If an event has a detection for E_L2 then we label `tShower` as the mean shower time in the BCAL and if there is a detection for $E1E9$ then we label `tShower` as the mean shower time in the FCAL. We find the difference between the mean shower times in each detector with the calculated time from the vertex to the respective detectors (`tFlightBCAL` or `tFlightFCAL`). In order to assess the quality of a given hypothesis, we calculate a chi-squared value between the mean shower time and calculated shower times. Only hypotheses with a chi-squared value of less than 0.075 are considered robust; any hypotheses that are above this threshold are labeled as no identification (no ID). We can only perform timing cuts on charged particles as there is no calculated timing information in our simulation dataset.

In addition to timing cuts, we also implement track energy loss cuts using the `dEdxCDC` variable and the particle momentum vectors added in quadrature. To create a decision boundary between each particle, we use the same functional form of the equations used in the Spring 2017 Analysis Launch Cuts:

$$dE/ds = e^{a \cdot p+b} + c, \quad (4.1)$$

where p is the momentum of a particle in units of GeV/c , dE/ds is the energy loss in the CDC in units of KeV/cm and e is Euler's number. a , b and c are constants that are varied in order to best classify each particle. Using the training dataset, we minimize the number of incorrectly identified particles by making each of these constants as free parameters and utilizing the `minimize` method from the `scipy.optimize` [15] module. Similarly to the timing cuts, we only derive $dE/ds - p$ decision boundaries for charged particles. We show the optimized decision boundaries in Equations 4.2 – 4.4.

$$dE/ds_1 = e^{-5.095 \cdot p - 10.205} + (2.080 \cdot 10^{-6}), \quad (4.2)$$

$$dE/ds_2 = e^{-3.947 \cdot p - 12.284} + (1.936 \cdot 10^{-6}), \quad (4.3)$$

$$dE/ds_3 = e^{-0.185 \cdot p + -19.215} + (2.190 \cdot 10^{-6}), \quad (4.4)$$

Each decision boundary is overlaid on the test dataset in Figure 1. Additionally, only for electrons and muons/pions, we use the particle's total energy E divided by the total momentum; a decision boundary of $0.83 c$ is chosen. Lastly, we only consider a hypothesis if the particle hypothesis matches the predicted hypothesis from our manual PID. Each of the conditions for our manual PID are shown in Table 2. A PID is made for every hypothesis in our test dataset that passes all of the cuts shown in Table 2. If a given even meets none of the cuts made, then no ID is designated; if an event has two or more PIDs that match the hypotheses in our test dataset, then the particle type with the highest chi-squared value is designated. The confusion matrix presenting the results of our manual PID on charged particles is shown in Figure 2a.




Figure 1: A 2-Dimensional histogram of the average track energy loss in the CDC plotted against the total momentum from our test dataset. The manual PID cuts described in Section 4.1 are overlaid to show the classification boundaries; the functional form of each decision boundary is shown in Equations 4.2 – 4.4. Regions of the plot shaded in red are classified as p or \bar{p} , purple as $K^{+,-}$, yellow as $e^{-,+}$, and blue as $\pi^{+,-}$ or $\mu^{+,-}$.

Table 2: Manual PID cuts. If an entry is missing, there is no cut for that particle. All manual PIDs must match the given hypothesis.

Particle	Δt BCAL [ns]	Δt FCAL [ns]	dE/ds [keV/cm]	E/p [c]
e^+	± 1.0	± 2.0	$dE/ds_3 < dE/ds$ & $dE/ds_2 > dE/ds$	$E/p > 0.83$
e^-	± 1.0	± 2.0	$dE/ds_3 < dE/ds$ & $dE/ds_2 > dE/ds$	$E/p > 0.83$
$\mu^+ \pi^+$	± 1.0	± 2.0	$dE/ds_3 > dE/ds$	$E/p < 0.83$
$\mu^- \pi^-$	± 1.0	± 2.0	$dE/ds_3 > dE/ds$	$E/p < 0.83$
K^+	± 0.75	± 2.5	$dE/ds_2 < dE/ds$ & $dE/ds_1 > dE/ds$	
K^-	± 0.75	± 2.5	$dE/ds_2 < dE/ds$ & $dE/ds_1 > dE/ds$	
p	± 1.0	± 2.0	$dE/ds_1 < dE/ds$	
\bar{p}	± 1.0	± 2.0	$dE/ds_1 < dE/ds$	




Figure 2: The confusion matrix for our manual PID on charged particles shown in Figure 2a and the confusion matrix for our NN PID on charged particles shown in Figure 2b. The generated particle is shown on the y-axis and the identified particle is shown on the x-axis. For events in our manual PID scheme that do not meet our chi-squared criteria described in Section 4.1, a no identification (no ID) classification is given. Similarly, for our NN PID method, a no ID classification is given when the confidence criteria described in Section 4.2 is not achieved.

4.2 Neural Network PID

We split our datasets into charged and neutral datasets. Furthermore, to ensure that the input feature space of each event/hypothesis is uniform, we replace any value in our datasets that do not have a detection with the 'Overflow Value' seen in Table 1. Additionally, only the features in Table 1 that have an Overflow Value are used to train our two models.

Our neural network models we make use of the TensorFlow implementations of the Cross Entropy Loss Function, Adam Optimizer and ReLu activation function, all of which are described in Section 3. The number of neurons, the number of hidden layers and the learning rate of the Adam optimizer are optimized to have the maximum validation accuracy by Hyperband. We allow variation between 1 and 6 hidden layers, between 100 and 600 neurons per hidden layer and between 10^{-4} and 10^{-2} for the learning rate. We use the optimized hyperparameters to train a NN model for a maximum of 50 epochs and stop the training if the TensorFlow implementation of Early Stopping if the Cross Entropy Loss changes by less than 0.01 after 5 epochs.

Our trained models are used to make a classification on every hypothesis in the test dataset. Each prediction made by the `predict` method from the TensorFlow models yields a confidence value for each possible classification (i.e. particle). We take the highest confidence value across all hypotheses in an event and label that event with the corresponding particle type. A no ID label is given for any PID that has a confidence of less than 0.4. In the same way as Section 4.1, we consider positive pions and muons and negative pions and muons as the same particles. The confusion




Figure 3: The confusion matrix for our NN PID for neutral particles; the generated particle is shown on the y-axis and the particle classified by our neutral NN model is shown on the x-axis. Particles that do not meet the confidence criteria discussed in Section 4.2 are classified with a no identification (no ID) label.

matrix for our charged NN model is shown in Figure 2b and the confusion matrix for our neutral NN model is shown in Figure 3.

5 Results

In this section we present the results of our traditional PID cuts and NN PID. We make a direct comparison between these two methods and demonstrate advantages of NNs in PID. We also determine the importance of each feature in the simulation dataset to understand what features the NN models suggest are the most crucial for PID.

5.1 Comparing PID techniques

In Figure 2a, the p sample is correctly identified with an accuracy of 0.942 for the manual PID method, with only 0.052 of samples being allocated to the no ID column of the confusion matrix. The p is the particle with the highest accuracy, while the K^- , e^- and e^+ have PID accuracies in the range 0.82 – 0.84. The lower end of classification accuracies are \bar{p} , K^+ , $\pi^+ | \mu^+$ and $\pi^- | \mu^-$ in the range 0.57 – 0.64. We see that many of the samples that are not correctly identified by our manual PID method are designated as no ID. The most prominent exceptions of this are 0.124 of the K^+ sample misidentified as p while 0.092 and 0.107 of the $\pi^+ | \mu^+$ sample identified as K^+ and e^+ , respectively. For $\pi^- | \mu^-$, 0.087 and 0.107 were identified as K^- and e^- , respectively. In addition to these common misclassifications, there are several cases in which our manual PID method misclassified samples more infrequently (< 0.05). The large number of events that pass the chi-squared timing cut and which are incorrectly identified highlight the areas of improvement of the classification accuracies of the manual PID method.

In Figure 2b, the confusion matrix of our charged NN model is presented. The p , e^- and e^+ are correctly identified in the range 0.94 – 0.98 and the five other charged particles have accuracies

in the range 0.83 – 0.91. It is found that for all particles the PID accuracy is better for our NN PID method in comparison to our manual PID method. p and K^- have an improvement of ~ 0.03 , e^- and e^+ have an improvement of 0.1 – 0.13, and the four other particles have substantial improvements of more than 0.2. In addition to substantial improvements in PID accuracy, we see decreases in the misclassification of K^+ as p (0.075), $\pi^+ | \mu^+$ as K^+ (0.06) and e^+ (0.041), and $\pi^- | \mu^-$ as K^- (0.052) and e^- (0.033). Despite the decreases in these particular misidentifications, we see increases in \bar{p} as K^- (0.064), K^+ as $\pi^+ | \mu^+$ (0.087), and K^- as $\pi^- | \mu^-$ (0.098). The increase in misidentification of these particles is the lone shortcoming of the NN PID method in contrast to our manual PID method.

Our classification scheme discussed in 4.2 indicates that a particle type is assigned only when the confidence interval from our model is 0.4. Even when this criteria is increased from 0.4 to 0.95, these misclassifications from our NN model still persist, which indicates that the only way to remove these misidentifications is through adopting a different ML architecture or making more stringent cuts.

5.2 Feature Importance with Shapley Values

We need to make a distinction between Shapley values and SHapley Additive exPlanations (SHAP) [16].

To analyze the PIDs made by our neural network, we use Shapley values to assess the importance of each feature. For a given PID made by our models, a Shapley value for a given feature measures the average contribution of that feature across the entire feature space. A Shapley value is computed for each feature for a given classification by considering possible permutations of features and then taking the average of all marginal contributions by a feature to the resultant prediction. In our case, this process is very computationally expensive, so we use `random.choice` from NumPy [17] to randomly sample 1,000 classifications from our test sample. We take the median Shapley value for each particle and for each feature label used in our neural network models.

The absolute value of our SHAP values for positively charged particles is shown in Figure 4, negatively charged particles in Figure 5, and neutral particles in Figure 6.

6 Conclusion

Discuss briefly:

1. Sample selection bias
2. Advantages of Neural Networks in PID

Acknowledgments

This is the most common positions for acknowledgments. A macro is available to maintain the same layout and spelling of the heading.




Figure 4: We denote the SHAP value magnitude as the absolute value of our SHAP values. We randomly sample 10^3 hypotheses from our test sample for each particle type and calculate the SHAP value [16] for each feature. Here we show only the $q = +1$ particles from our charged particle NN model; p in pink, K^+ in red, e^+ in light blue and $\pi^+ | \mu^+$ in green. The black line in each box plot represents the mean SHAP value in a given box plot. The upper and lower end of a given box plot represents the 75th and 25th percentile of the data, respectively. Each of the feature labels used to train our charged NN model from Table 1 are shown on the x-axis unless a SHAP value of zero is calculated, then these features are omitted.




Figure 5: The same as Figure 4, except we show SHAP value magnitudes for particles with $q = -1$. We show \bar{p} in pink, K^- in red, e^- in light blue and $\pi^- | \mu^-$ in green.




Figure 6: The same as Figure 4 and 5, except we show the SHAP value magnitudes for particles included in our neutral particle NN model. We show γ in pink, K_L^0 in red and n in light blue.

References

- [1] S. Adhikari, C.S. Akondi, H. Al Ghouli, A. Ali, M. Amaryan, E.G. Anassontzis et al., *The GLUEX beamline and detector*, *Nuclear Instruments and Methods in Physics Research A* **987** (2021) 164807 [2005.14272].
- [2] D.S. Carman, *GlueX: The Search for Gluonic Excitations at Jefferson Laboratory*, in *Hadron Spectroscopy*, A. Reis, C. Göbel, J.D.S. Borges and J. Magnin, eds., vol. 814 of *American Institute of Physics Conference Series*, pp. 173–182, AIP, Feb., 2006, DOI [hep-ex/0511030].
- [3] The GlueX Collaboration, A. AlekSejevs, S. Barkanova, M. Dugger, B. Ritchie, I. Senderovich et al., *An initial study of mesons and baryons containing strange quarks with GlueX*, *arXiv e-prints* (2013) arXiv:1305.1523 [1305.1523].
- [4] The GlueX Collaboration, *Mapping the spectrum of light quark mesons and gluonic excitations with linearly polarized photons presentation to pac30 – the glueX collaboration*, https://www.jlab.org/exp_prog/proposals/06/PR12-06-102.pdf.
- [5] L. Alzubaidi, J. Zhang, A. Humaidi, A. Al-Dujaili, Y. Duan, O. Al-Shamma et al., *Review of deep learning: concepts, cnn architectures, challenges, applications, future directions*, *Journal of Big Data* **8** (2021) .
- [6] R. Barsotti and M.R. Shepherd, *Using machine learning to separate hadronic and electromagnetic interactions in the GlueX forward calorimeter*, *Journal of Instrumentation* **15** (2020) P05021 [2002.09530].
- [7] Geant4 Collaboration, “Geant4: A Simulation Toolkit for the Passage of Particles through Matter.” Astrophysics Source Code Library, record ascl:1010.079, Oct., 2010.
- [8] J. Allison, K. Amako, J. Apostolakis, P. Arce, M. Asai, T. Aso et al., *Recent developments in GEANT4*, *Nuclear Instruments and Methods in Physics Research A* **835** (2016) 186.
- [9] D.R. Cox, *The regression analysis of binary sequences*, *Journal of the Royal Statistical Society: Series B (Methodological)* **20** (1958) 215 [https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1958.tb00292.x].

- [10] Z. Shangnan and Y. Wang, *Quantum Cross Entropy and Maximum Likelihood Principle*, *arXiv e-prints* (2021) [arXiv:2102.11887](https://arxiv.org/abs/2102.11887) [[2102.11887](https://arxiv.org/abs/2102.11887)].
- [11] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, *arXiv e-prints* (2014) [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [[1412.6980](https://arxiv.org/abs/1412.6980)].
- [12] R.H.R. Hahnloser, R. Sarpeshkar, M.A. Mahowald, R.J. Douglas and H.S. Seung, *Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit*, **405** (2000) 947.
- [13] A.F. Agarap, *Deep Learning using Rectified Linear Units (ReLU)*, *arXiv e-prints* (2018) [arXiv:1803.08375](https://arxiv.org/abs/1803.08375) [[1803.08375](https://arxiv.org/abs/1803.08375)].
- [14] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh and A. Talwalkar, *Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization*, *arXiv e-prints* (2016) [arXiv:1603.06560](https://arxiv.org/abs/1603.06560) [[1603.06560](https://arxiv.org/abs/1603.06560)].
- [15] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau et al., *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature Methods* **17** (2020) 261.
- [16] S. Lundberg and S.-I. Lee, *A Unified Approach to Interpreting Model Predictions*, *arXiv e-prints* (2017) [arXiv:1705.07874](https://arxiv.org/abs/1705.07874) [[1705.07874](https://arxiv.org/abs/1705.07874)].
- [17] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau et al., *Array programming with NumPy*, *Nature* **585** (2020) 357.