Point-cloud based machine learning for classifying rare events in the Active-Target Time Projection Chamber

Poulomi Dey[a,b], Adam K. Anthony[a,c,d], Curtis Hunt[a], Michelle P. Kuchera[e,f], Raghuram Ramanujan[f], Jessica W. Ajongbah[a,b], Kyle Brown[a,b], Zbigniew Chajecki[h], Skyler Gangestad[d], William G. Lynch[a,c], ManYee Betty Tsang[a,c], Joseph M. Wieske[a,c], HoTing Wong[g,h]

[a] *Facility for Rare Isotope Beams, Michigan State University, East Lansing, MI, 48824, USA*
[b] *Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI, 48824, USA*
[c] *Department of Physics and Astronomy, Michigan State University, East Lansing, MI, 48824, USA*
[d] *Department of Physics and Astronomy, High Point University, High Point, NC, 27268, USA*
[e] *Department of Physics, Davidson College, Davidson, NC, 28035, USA*
[f] *Department of Mathematics and Computer Science, Davidson College, Davidson, NC, 28035, USA*
[g] *Department of Physics, The Chinese University of Hong Kong, Shatin, Hong Kong, China*
[h] *Department of Physics, Western Michigan University, Kalamazoo, MI, 49008, USA*

Abstract

In this work, we assess the use of machine learning to classify fission events in the Active Target Time Projection Chamber (AT-TPC) using data from an experiment performed at the National Superconducting Cyclotron Laboratory at Michigan State University. The experiment produces an extremely large quantity of data, less than 3% of which are fission events. Therefore, separating fission events from the background beam events is critical to an efficient analysis. A heuristic method was developed to classify events as `Fission` or `Non-Fission` based on hand-tuned parameters. However, this heuristic method places 5% of all events into an `Unlabeled` category, including 15% of all fission events. We present a PointNet model trained on the data labeled by the heuristic method. This model is then used to generate labels for the events in the `Unlabeled` category. Using the heuristic and machine learning methods together, we can successfully identify 99% of fission events.

1. Introduction

In many physics applications, isolating rare events of interest is essential. Maximizing the number of these rare events that can be identified is critical for performing robust downstream analyses. In some applications, traditional approaches for identifying such reactions have relied on carefully handcrafted heuristic approaches, but recent advances in machine learning (ML) point to a promising alternative. In this work, we present the effectiveness of a deep learning method in identifying rare fission events by bootstrapping from annotations generated by a heuristic algorithm.

The HiRA group at the National Superconducting Cyclotron Laboratory (NSCL) conducted an experiment to study the fission properties of isotopes around ^{196}Pb using the Active Target Time Projection Chamber (AT-TPC). The AT-TPC is a detector capable of tracking charged particles moving through a gas volume which acts as the reaction target [1]. Measurements of the fission cross section as a function of energy and their charge distributions are of particular interest. The cross sections are used to model the fission barrier, allowing comparisons to theory. In the experiment, less than 3% of the measured events are fission. A heuristic method for fission event identification was developed which recovers 85% of the fission events [2] with minimal computing time. It is important, both to accurately determine the cross section and increase the efficiency of the analysis, to precisely identify fission events at all energies and quantify any systematic limitations in the identification algorithm.

Various studies have explored practical applications of ML in nuclear physics [3, 4]. ML for anomaly detection—flagging events that are atypical or rare—has been used successfully in particle physics applications [5–7]. ML was applied to similar classification tasks in the AT-TPC with promising outcomes. Most prior ML work with the AT-TPC utilizes a two-dimensional projection of the data for tractable computation with the chosen architectures. For example, Kuchera et al. used Convolutional Neural Networks (CNNs) to classify proton scattering events in the AT-TPC [8]. Solli et al. used the latent space of deep

learning models to cluster AT-TPC events with unsupervised learning methods [9]. Wu et al. employed two CNNs, ResNet and VGG-16, to identify ^{12}C events from background events using simulated data for a similar active target time projection chamber [10]. We build off this work to classify fission events in the AT-TPC using ML methods.

Specifically, we apply point cloud based ML methods to identify fission events in the AT-TPC that are missed by our traditional analysis. We use events that have been labeled as either `Fission` or `Non-Fission` by a heuristic method with high confidence to train a PointNet model [11]. We then use this model to predict labels for events the heuristic method could not classify in order to gain additional statistics. This also compensates for shortcomings in the heuristic method when classifying events with a small folding angle or those close to the pad plane.

We start with a description of the experimental setup, detectors, and traditional analysis techniques used in the experiment in Section 2. Section 3 discusses the machine learning methods used. This includes both the data sets used to train and evaluate the model along with a description of the model itself and its training. Section 4 details the results of applying the machine learning model to our data, its ability to recover additional fission events over the traditional analysis methods employed, and the implications for our ability to extract physics from the data.

2. Background

2.1. The Active Target Time Projection Chamber (AT-TPC)

The AT-TPC has a cylindrical active volume that is one meter long and about half a meter in diameter filled with a target gas. For this experiment, the target gas was He at one atmosphere. Within this volume, we can record the trajectories of charged particles and their energy loss through the gas over time [1]. Figure 1 shows a schematic representation of the detector.

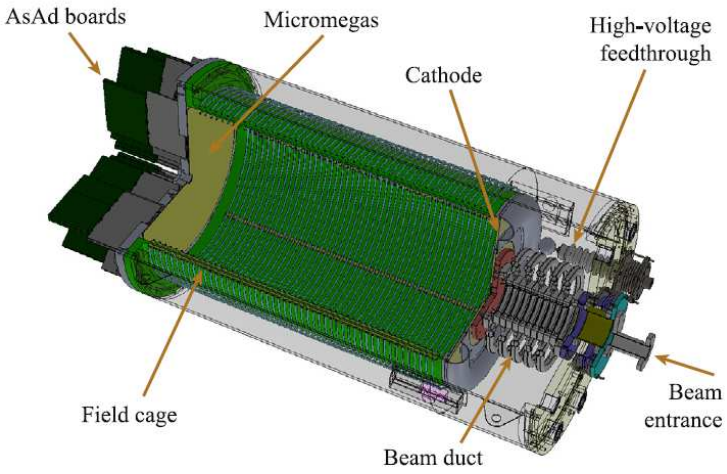

Figure 1: Schematic representation of the AT-TPC. Reprinted with permission from [1].

As a charged particle passes through the active volume and loses energy, it will ionize the gas molecules, producing electron/ion pairs. The number of electrons produced is proportional to the energy lost by the charged particle. A constant electric field in the chamber causes the electrons to drift at a constant velocity to a highly segmented pad plane (with more than 10,000 pads) opposite the entrance to the TPC. As the electrons reach the pad plane, they are amplified through a Micromegas [12] and Thick-GEM [13]. This amplified charge is collected on the pads in the pad plane. The resulting voltage on each pad is individually read out over time using the General Electronics for TPCs (GET) system [14]. The digitized signal results from two stages of amplification, a charge sensitive pre-amplifier and a shaping amplifier. By measuring the time the electrons arrive relative to an external timing reference (see Sec. 2.2) and their location on the pad plane, the initial 3D position where the electron was ionized can be reconstructed. In addition, by measuring the magnitude of the charge collected, the energy loss of the charged particles as they pass through the gas can be inferred. In order to reconstruct the charge from the measured signal, we use a standard deconvolution method for GET electronics [2, 15]

These two pieces of information, charge and 3D-position, are used to generate

the four-dimensional (three spatial, one charge) hit patterns that are used in the analysis of the experiment. These hits are grouped into tracks using a modified RANSAC [16] algorithm. The algorithm, part of the ATTPCROOT analysis package [2, 17], is tuned for TPC data and can incorporate information on the expected event geometry into its sampling [2, 18, 19]. The resulting tracks describe the path and energy loss of a charged particle through the active volume of the detector.

2.2. Experimental Details

A recent experiment was conducted at the NSCL using the AT-TPC to study the fission of neutron deficient isotopes around ^{196}Pb, near the transition between symmetric and asymmetric fission [2]. In the experiment, a stable beam of $^{208}Pb^{+63}$ was accelerated to 85 MeV/A using the NSCL coupled cyclotron facility [20, 21]. This stable beam was fragmented on a ^{9}Be target and the isotopes of interest were isolated using the A1900 fragment separator. The beam, which contained about a dozen different isotopes, was identified on an event-by-event basis using the HEIST particle identification system [22] and then directed into the AT-TPC. The HEIST system uses the dE-Bρ-ToF method for particle identification and is partially composed of two micro-channel plate detectors for time of flight (ToF) and a MUSIC detector for energy loss (dE). The downstream MCP detector serves as the beam arrival time in the TPC. This external timing reference is crucial for establishing the absolute Z-position of events in the AT-TPC. A detailed description of the detectors and experimental setup can be found in [2, 22].

As the beam travels through the AT-TPC, it can fuse with a He nucleus in the gas. The resulting compound nucleus will be in an excited state and quickly decay. The two primary decay mechanisms are fission (our reaction of interest) and neutron decay, a channel we have negligible sensitivity to. Since the nucleus of an atom is about 25,000 times smaller than its atomic size, the probability of a projectile reacting with a He nucleus in the gas is extremely small. Similarly, the probability of the excited compound nucleus decaying though fission is given,

roughly, by the ratio of the fission to the neutron emission decay widths. For the nuclei and energies involved in this experiment, this ratio is expected to be on the order of 10^{-4} to 10^{-6}.

To reduce the relative amount of non-fission events we store to disk, we only record data when there is a signal on both the Micromegas mesh at the pad plane and at the downstream MCP detector in HEIST. To ensure any signal in the pad plane is primarily due to fission fragments, the gas gain of the pads in the center of the AT-TPC was lowered by biasing the pads in the path of the beam. This ensures that there is a large signal on the mesh only when a charged particle leaves the central beam region of the detector. Even with these precautions, fission events still only comprise an estimated 2-3% of the over 6 million events that were collected. This estimation is based on hand labeling a small subset of the data. The remaining 97-98% of the data is attributed to other kinds of events. The `Non-Fission` category is comprised primarily of beam events, i.e., events where the isotope beam passes through the AT-TPC without reacting. An accurate and efficient method is therefore necessary to sift through this large amount of data to isolate events of interest. Furthermore, since fission is a rare event, it is important that we capture as many fission events in our analysis as possible. In this paper, we will focus on maximizing fission recall, which is the percentage of true fission events that our algorithm identifies as fission.

When fission occurs, the AT-TPC records the tracks of both the beam before the vertex and the fission fragments. Fission events have a characteristic "Y"-shaped track, where the arms are the fission fragments and the base stem is the beam. Typical beam events look like a straight line through the central axis of the AT-TPC. Figure 2 presents examples of typical fission and beam events, using the standard coordinate convention for the AT-TPC: the beam direction is the negative z-direction, with the pad plane at zero and the window at 1000 mm. The y-direction is the vertical axis, while the x-direction is the horizontal axis, with the center of the pad plane serving as the origin. Charge is measured in arbitrary units by the electronics, ranging from 0 to 4095 channels.

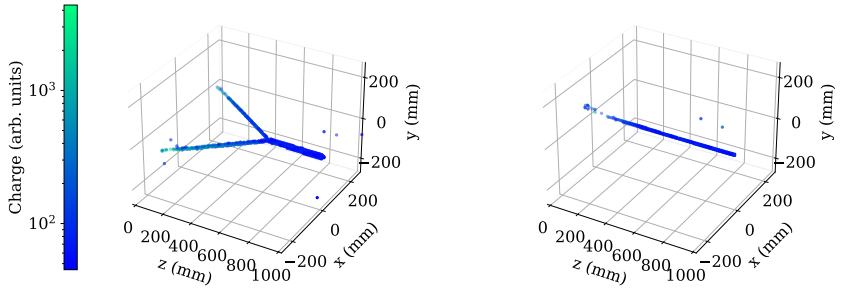

(a) `Fission` (b) `Non-Fission`

Figure 2: Examples of typical fission and beam events from this experiment. Beam events make up the vast majority of the `Non-Fission` category. These events were labeled by the heuristic method as `Fission` and `Non-Fission`, respectively, and come from the training set. To produce this figure, each pad in the pad plane was processed to allow more than one "hit," or point, in the point cloud. This allows us to visualize the beam prior to fission to clearly show the "Y"-shape of the fission event.

2.3. Heuristic Analysis Method

The HiRA group has developed a heuristic method to separate out fission events from all the events detected by the TPC [2]. The method quickly tags events as either `Fission`, `Non-Fission`, or `Unlabeled` based on various characteristics of the tracks. The specific parameters used are given in Table 1. The heuristic method was developed using a set of 2,500 hand labeled events. Independent hand-labeling of a sub-sample of `Unlabeled` events, discussed in Sec. 3.1, showed an agreement between experts of more than 99%.

The heuristic method consists of various cuts on the parameters. Each cut is a range of values for a given parameter, such as the angle between tracks, that ensures a clean event identification. For the angle between tracks, if no tracks have an angle above a certain threshold, that event is classified as beam. The cuts are conservative, designed to exclude incorrect identifications, even if it means forgoing an identification for some events. A series of these cuts are made using different combinations of values for these parameters to classify an event. When the first cut fails to make a classification, the next is used and so

Individual Cut Parameters	units
Number of RANSAC tracks found	#
Largest opening angle between any pair of tracks in the event	°
Smallest polar angle (w.r.t. beam axis) of each track in the event	°
z_{min}	mm
z_{max}	mm
Maximum $\sqrt{\sigma_x^2 + \sigma_y^2}$	mm

Table 1: A list of the parameters used by the heuristic algorithm. A sequence of cuts based on these parameters were applied to the data to classify the data into `Fission` and `Non-Fission`. Beam conditions are typified by small track angles and low spread of the point cloud hits. Fission conditions involve multiple tracks, larger angles, and high point cloud spread. The `Unlabeled` category contains events which did not satisfy any condition for beam or fission.

Category	Count	Percentage
`Fission`	135,959	2.2%
`Non-Fission`	5,827,822	92.8%
`Unlabeled`	316,581	5.0%
Total Events	6,280,362	

Table 2: Results from using the heuristic method to classify all events in all runs used in the data analysis.

Category	Count	Percentage
`Fission`	147	84.5%
`Unlabeled`	27	15.5%
Hand-Labeled Fission Events	174	

Table 3: A total of 174 fission events were identified by processing a set of 5,000 events with the heuristic method. Of these, 147 were classified as `Fission` and 138 were classified as `Unlabeled`. Hand-labeling the `Unlabeled` events identified 27 fission events. Extrapolating this finding to the rest of the dataset, we estimate that 15% of all fission events are sorted into the `Unlabeled` category.

on until all 50 cuts are exhausted. An event is classified by the first cut in the sequence that it fits, which means that subsequent cuts do not need to account for that type of event and some cuts could otherwise provide contradictory classifications. The method is only capable of definitively classifying about 95% of events (Table 2); events that fail all of the cuts are classified as `Unlabeled`. By hand-labeling a subset of these `Unlabeled` events, it is estimated that about 15% of all fission events were sorted into the `Unlabeled` category (Table 3).

There are two primary factors that contribute to difficulty in classifying some events. The first is that some events contain multiple beam particles, leading to non-fission events where more than one track is identified by RANSAC. This is further complicated by the relatively wide beam profile and the effects of space charge warping the tracks, leading to pileup events having two or more tracks with a measurable angle between them. The second is the existence of a region down the centerline of the detector that is low gain. Fission fragments are undetectable in that region. For events with a small angle between the fission fragments or where the reaction occurs close to the pad plane, the fission fragment tracks may be too short for the RANSAC algorithm to detect. The consequence of this is that a disproportionate number of fission events with a small folding angle or those close to the pad plane are not successfully identified by the heuristic method.

The experiment lasted several days and the data was split up into roughly

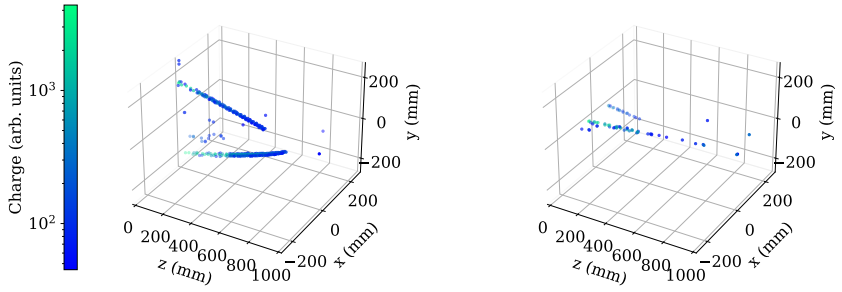

(a) Labeled as `Fission` by heuristic method

(b) `Unlabeled` by heuristic method

Figure 3: Fission events that were `Unlabeled` by the heuristic method often have visual differences from the ones labeled as `Fission`. These events were randomly selected from the training set and application set, respectively.

hour-long "runs". Across all runs, the percentage of all fission events sorted into the `Unlabeled` category remained consistent. The number of `Non-Fission` events sorted into the `Unlabeled` category did vary over time, possibly due to a varying beam rate. If the properties of the beam were changing we would expect to see a difference in the number of recorded fission events, which was not observed.

Events close to the pad plane are of particular interest to fission cross section studies. They correspond to lower energy events and the cross section is expected to be the most sensitive to changes in the fission barrier at such energies. Figure 3 shows the visual difference between typical fission events that were labeled as `Fission` and `Unlabeled` events by the heuristic method. Further improvement to the heuristic method would not only be time consuming to develop, but also likely to rely on data further along the analysis pipeline. This would negate one of the primary benefits of classifying events—reducing the time needed to analyze the data.

Dataset	`Fission`	`Non-Fission`	Labeling Method
Training	1,691	3,946	Heuristic
Validation	440	1,027	Heuristic
Test	2,056	104,113	Heuristic
Application	28	905	Manual

Table 4: Count of `Fission` and `Non-Fission` events in the data sets. The training and validation sets were rebalanced to a 30-70 `Fission`-to-`Non-Fission` ratio. The events in the training, validation, and test sets were labeled as `Fission` or `Non-Fission` by the heuristic method. The events in the application set were those `Unlabeled` by the heuristic method that were subsequently tagged by hand as `Fission` or `Non-Fission`.

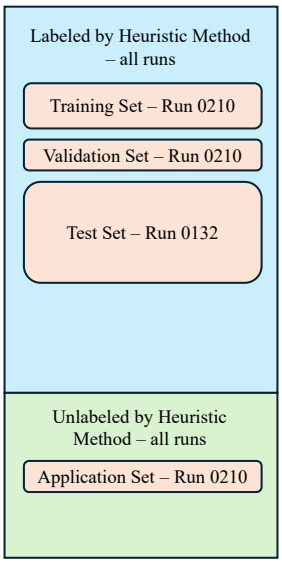

Figure 4: The source of all datasets used. The training, validation, and application sets were subsets of one run. The test set was composed of all `Fission` and `Non-Fission` events from another run. The specific run does not affect the properties of the events. This diagram is not to scale. Exact event distributions are presented in Tables 2, 3, and 4.

3. Methods

3.1. Machine Learning Data Processing

We prepared our data for training the ML model by first labeling a set of $65,583$ events from a single run as `Fission`, `Non-Fission`, or `Unlabeled` using the heuristic method. The properties of fission events remain unchanged across the course of the experiment. This was verified in multiple ways, including analysis of the beam profile in HEIST and the measurements of the electron drift velocity in the AT-TPC across runs. In addition, this was further confirmed by the uniformity in the percentage of events labeled as fission by the trained ML model across runs.

In this initial stage of data preparation, any events that were `Unlabeled` by the heuristic method were disregarded. This yielded a dataset with $2,131$ `Fission` events and $60,891$ `Non-Fission` events. Training statistical models in regimes with such highly imbalanced class distributions can be challenging, and is a problem that has been extensively studied by the ML community [23, 24]. Indeed, our initial attempts to fit models directly using this data yielded underwhelming results, predicting nearly all events as `Non-Fission`. We thus resorted to a simple and standard remedy—undersampling the majority class [23, 24]—to address this issue. Specifically, we rebalanced the dataset by retaining all `Fission` events and discarding a subset of the `Non-Fission` events chosen at random, so that the final dataset contained a 30-70 distribution of `Fission` and `Non-Fission` events, respectively. Following standard ML procedures, this dataset was then randomly partitioned using an 80-20 split to create a *training set* (for fitting model parameters) and a *validation set* (for determining ideal hyperparameter settings).

Simultaneously, an independent set of events were heuristically labeled as `Fission` and `Non-Fission` to form a *test set*. The events in this set did not participate in the model fitting process and were used to obtain an unbiased estimate of model performance, as well as to quantify the degree of agreement between the heuristic and ML approaches to event classification. No rebalancing

was performed on the events in this set.

Finally, we compiled a third dataset consisting solely of events that were `Unlabeled` by the heuristic method. The events in this *application set* were then hand-labeled `Fission` and `Non-Fission`, and used to evaluate the ML model's performance on events on which the heuristic approach failed. These four datasets, their class distributions, and data sources are summarized in Table 4 and illustrated schematically in Figure 4.

To assess the accuracy of hand-labeling the data, the application set was independently labeled by three individuals: two experts familiar with the data and one non-expert who was trained in the labeling process. All three sets of labels agreed 95.5% of the time, with the two experts agreeing more than 99% of the time. For each event, the mode of the labels was taken as the ground truth. It is important to note that the events in this sample are those that are expected to be hardest to classify. The strong agreement among experts, even with this challenging sample, highlights the reliability of the hand-labeling process and represents a conservative estimate of consistency.

Each event was processed into a point cloud, with each point along a particle track represented by its three-dimensional position in the chamber and charge amplitude information, to produce a four-dimensional (x, y, z, Q) representation. The (x, y) coordinates are determined by the centroid of the pad where charge was deposited, Q is the maximum charge per unit time collected on a given pad, and z is the distance from the pad plane to the point where the maximum charge occurs along the beam axis. Each pad on the pad plane was processed to contain only one "hit". Thus, the beam track prior to fission has fewer points, while the fission fragments have many points because they cross multiple pads.

Since the ML methods we utilize expect every point cloud to have the same size (i.e., the same number of points), we standardize each event to contain 64 points. Events containing more points were downsampled via uniform random sampling; events containing fewer than 64 points were upsampled by duplicating randomly selected existing points. Further, the spatial coordinates of every point in the events were linearly rescaled so that the pad plane corresponds to

$z = 0$ and the beam entrance corresponds to $(0, 0, 1)$, with the positive x and y directions chosen to maintain the geometry of the detector. Since the recorded charge values Q ranged over several orders of magnitude, we first performed a "safe" log transformation on these, given by $Q' = \log_{10}(Q + 10^{-10})$. The resulting Q' values were then linearly min-max scaled, using the minimum and maximum Q' values encountered in the training set[1].

3.2. Model Training and Evaluation

We modify the PointNet architecture for our application, which is a machine learning architecture that directly consumes point cloud data as input and can be used to classify or segment objects [11]. We trained a PointNet model on resampled and rescaled (x, y, z, Q) point clouds on a binary classification task: that of predicting whether a given point cloud depicted a `Fission` or `Non-Fission` event. Hyperparameters were found via a grid search, and the final, tuned values are listed in Table A.9. We found that 30 epochs of training was typically sufficient to train the model to convergence. We determined the optimal early stopping point—where the model's loss on the validation set was minimized—by visual inspection of the learning curves generated by the training procedure. We report this model's performance on both the test set and the application set.

4. Results

4.1. Test Results

When directly comparing the ML model to the heuristic method, we were able to achieve 0.997 fission recall, a measure of the amount of fission events we are able to capture from the test set with the ML model. Table 5 compares how the heuristic method and the ML model labeled the same events. The left

[1]The events presented in plots in this paper, such as those in Figure 2, are depicted prior to resampling and rescaling.

	Machine Learning	
Heuristic	`Fission`	`Non-Fission`
`Fission`	2,051	5
`Non-Fission`	464	103,649

Table 5: A confusion matrix tabulating the `Fission` and `Non-Fission` classifications made by the heuristic and ML methods on the examples comprising the test set. The numbers along the main diagonal show where the two methods agree and the off-diagonal values are where the methods disagree.

	Test Results	Application Results
Precision	0.816	0.963
Recall	0.997	0.929
Accuracy	0.996	0.997
F1 Score	0.897	0.945

Table 6: Evaluation of Test Results and Application Results. Precision is the proportion of events labeled as `Fission` by both methods over all events labeled as `Fission` by the machine learning method. Recall is the proportion of events labeled as `Fission` by both methods over all events labeled as `Fission` by the heuristic method (Test Results) or by hand-labeling (Application Results). Accuracy is the proportion of events where both methods agree on a label over all events. F1 score is the harmonic mean of precision and recall.

	Machine Learning	
Hand-Labeled	`Fission`	`Non-Fission`
`Fission`	26	2
`Non-Fission`	1	904

Table 7: A confusion matrix tabulating the `Fission` and `Non-Fission` classifications made by a human annotator and the ML method on the examples comprising the application set. These are events that were `Unlabeled` by the heuristic method that were then manually labeled as `Fission` or `Non-Fission`. The numbers along the main diagonal show where the two methods agree and the other values are where the methods disagree. The rarity of fission events meant that their count was relatively small in the application set.

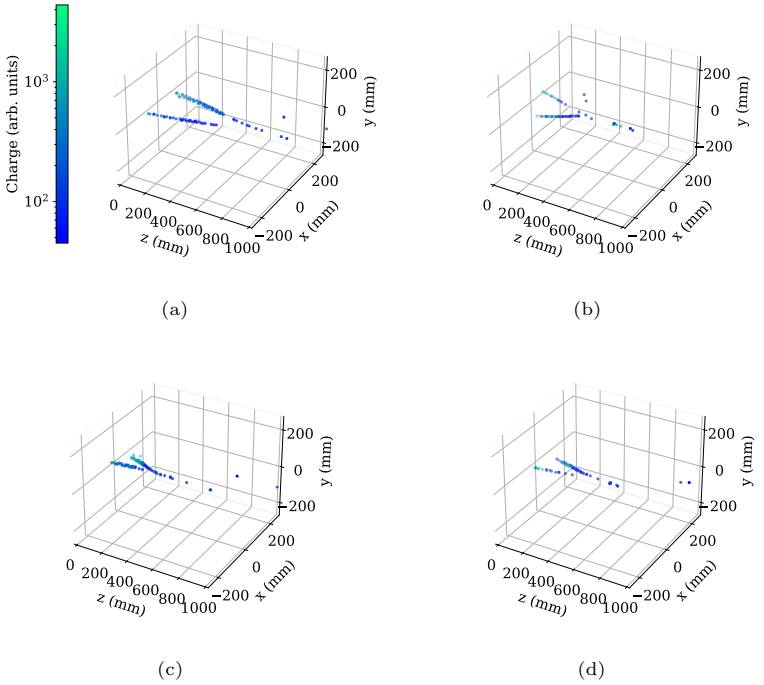

Figure 5: A sampling of fission events from the application set that were correctly labeled as `Fission` by the ML model. These were among the 26 events in the top-left cell of Table 7 that were originally `Unlabeled` by the heuristic method.

column of Table 6 shows various metrics measuring the performance of the ML model compared to the heuristic method.

Upon closer inspection, all of the five events that the heuristic method labeled as `Fission` but that the machine learning model labeled as `Non-Fission` were in fact fission events, albeit ones triggered by a different beam particle entering the TPC within the same measurement time window. These cannot be used for analysis purposes and therefore should actually be in the `Non-Fission` category.

4.2. Application Results

We then used the model to make predictions on the application set. The model achieved a recall of 0.929 for fission events that the heuristic method was

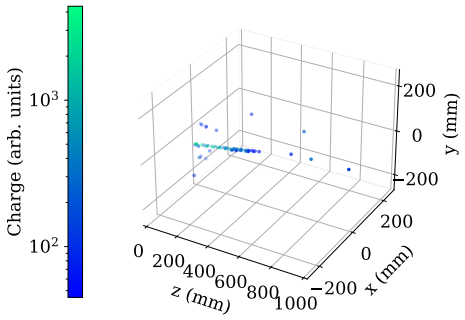

(a) A `Non-Fission` event that was mis-classified as `Fission`.

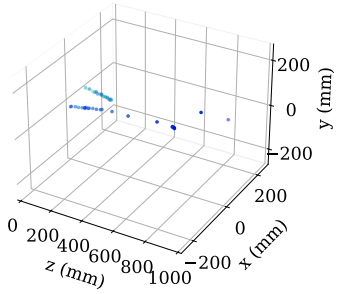

(b) A `Fission` event that was classified as `Non-Fission`.

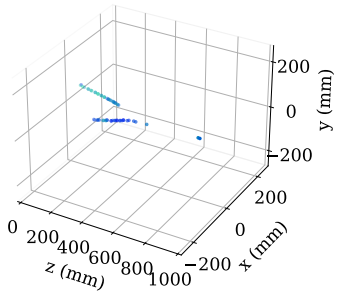

(c) A `Fission` event that was classified as `Non-Fission`.

Figure 6: Events from the application set that were incorrectly labeled by the ML model. These are the off-diagonal entries in Table 7; these events were originally `Unlabeled` by the heuristic method.

Event Type	Count	Percentage
`Fission`	23,720	7.5%
`Non-Fission`	292,861	92.5%
Unlabeled Category	316,581	

Table 8: ML classification of all events that the heuristic method could not label (`Unlabeled` in Table 2).

unable to label. Table 7 compares how the ML model labeled events against how a human annotator labeled them. The human labels were determined through a majority voting process, with each event independently labeled three times, as described in Sec. 3.1. The right column of Table 6 shows metrics measuring the performance of the ML model on `Unlabeled` data. The table also compares the performance of the ML model on the application set to the performance on the test set as a baseline. The model was tuned to optimize performance on the `Unlabeled` data in the application set, particularly the fission recall. A sample of events from the `Unlabeled` data that the ML model correctly labeled as `Fission` can be seen in Figure 5. The most likely reason for these events being classified as `Unlabeled` is the lack of sufficient hits to create a RANSAC track for the beam. All of the events the ML model misclassified from the application set are shown in Figure 6. The fission events have a characteristic "V"-shape, with both tracks leading to a common vertex. In contrast, the upper non-fission event in 6a only has one clear track, with any possible second track being at such a small angle as to be kinematically impossible for fission.

4.3. Physics Analysis

We used the ML model to generate labels for the `Unlabeled` data across all runs of the experiment, the results of which can be seen in Table 8. The gain in statistics from these additional events is not uniformly distributed. The counts as a function of position in the TPC are shown in Figure 7, along with a ratio of the additional events to the total events. The ML model recovered a significantly higher percentage of events closer to the pad plane than those further from

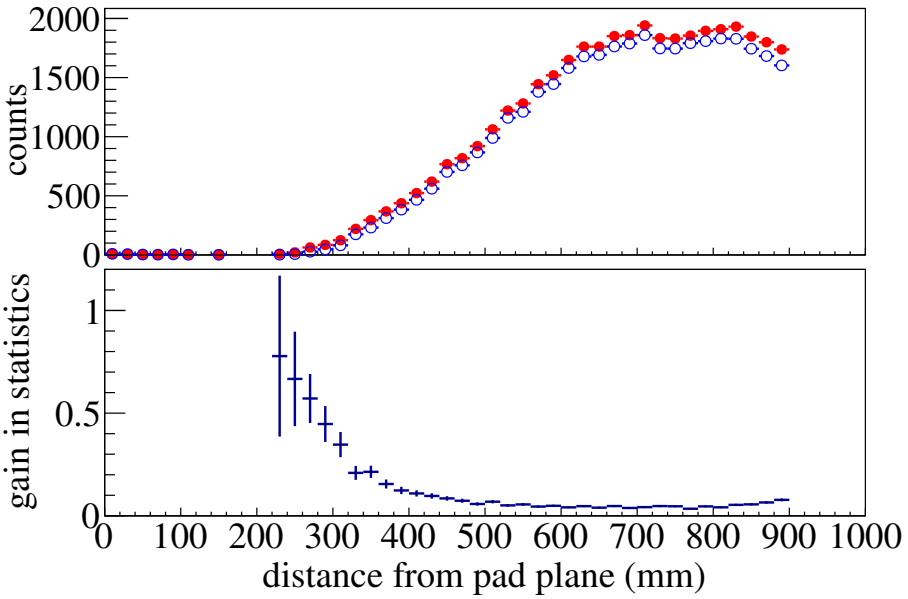

Figure 7: Top: The total number of fission events as a function of the position of the reaction vertex in the TPC. The blue empty circles represent the counts for events labeled `Fission` by the heuristic method; the red filled circles are the sum of the events labeled `Fission` by the heuristic method and the `Unlabeled` events labeled `Fission` by the ML model. Bottom: The ratio of the events labeled by the ML model to the total number of events from both methods.

the pad plane. A similar plot looking at the folding angle between the fission fragments is shown in Figure 8. The ML model recovered a significantly higher percentage of events with a smaller angle. There is additional gain at larger angles as well, due to the gain in events closer to the pad plane which can only be detected at larger folding angles.

The non-uniform distribution in the ratio of recovered fission events to total fission events as a function of the reaction vertex position and folding angle suggests the heuristic method had difficulty definitively classifying fission events close to the pad plane and fission events with a smaller angle between the fission fragments. As mentioned in Section 2.3, this was a suspected limitation of the heuristic method which could not be previously quantified. This result is important for a few reasons. The model was trained on events that were identified

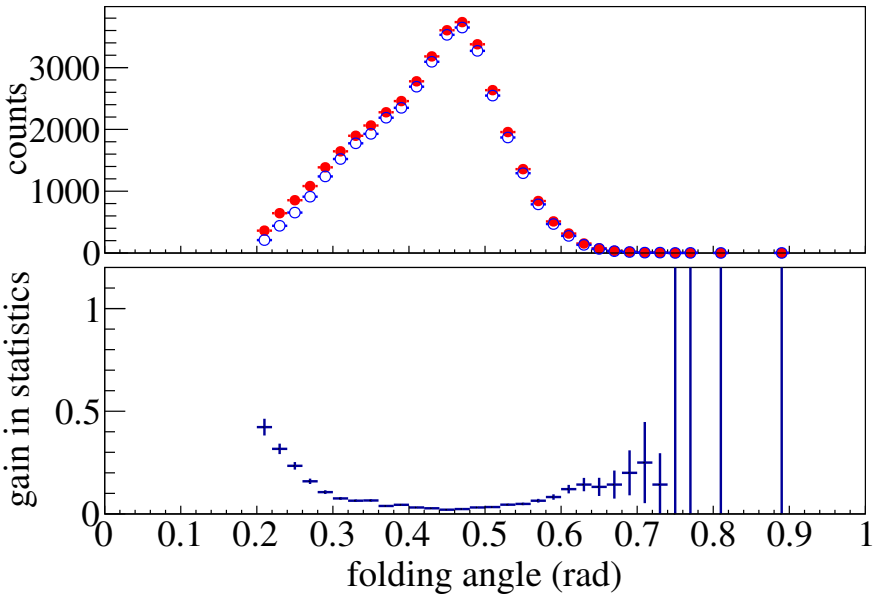

Figure 8: Top: The total number of fission events as a function of the angle between the fission fragments. The blue empty circles represent the counts for events labeled `Fission` by the heuristic method, the red filled circles are the sum of the events labeled `Fission` by the heuristic method and `Unlabeled` events labeled `Fission` by the ML model. Bottom: The ratio of the events labeled by the ML model to the total number of events from both methods.

by the heuristic model, now known to struggle with identifying certain types of events. This means that certain types of fission events were under-represented in the training set. Despite this under-representation, the trained model was still able to correctly identify events missed by the heuristic method, almost doubling the number of events in the regions the heuristic method was suspected to have problems making a classification. This has important implications for the physics analysis. As mentioned earlier, the extraction of certain physical properties, especially the fission barrier, can be very sensitive to the number of counts in the region nearest the pad plane.

5. Conclusions

In this paper, we presented an ML model that classified the 5% of events that the heuristic method was unable to classify. By combining both methods, we were able to identify 98.9% of all fission events. The ML model identified a systematic limitation in the heuristic method, showing that the latter had difficulty labeling events close to the pad plane of the AT-TPC, as well as events with a small angle between the fission fragments. Despite being trained on a data set which underrepresented certain fission events, the machine learning model was able to accurately identify these events in the unlabeled data.

As previously noted, the fission properties measured from the fission cross section can be sensitive to the shape at low energies, corresponding to events close to the pad plane with the AT-TPC. The characterization of the systematic limitations in our event sorting and then recovering these events was important for confidently obtaining an accurate cross section at all energies.

6. Acknowledgments

This work is supported in part by the National Science Foundation under Grant No. PHY-2209145 (AKA, CH, MBT, WGL, JW), Grant No. PHY-2012865 (MPK, RR), Grant No. PHY-2309923 (KB), and Grant Nos. PHY-2110218 and PHY-1712832 (ZC). Additionally, this work is in part supported by

the U.S. Department of Energy, Office of Science, Nuclear Physics under Award No. DE-SC0021938 (KB). PD acknowledges the financial support from the Director's Research Scholar Program of the Facility for Rare Isotope Beams. HTW acknowledges the financial support from the Summer Undergraduate Research Experience. SG acknowledges financial support from the Summer Research Program in the Sciences.

References

[1] J. Bradt, et al., Commissioning of the Active-Target Time Projection Chamber, Nucl. Instrum. Meth. A 875 (2017) 65–79. doi:10.1016/j.nima.2017.09.013.

[2] A. K. Anthony, Fission in the Lead Region, Ph.D. thesis, Michigan State University, 2023.

[3] A. Boehnlein, M. Diefenthaler, N. Sato, M. Schram, V. Ziegler, C. Fanelli, M. Hjorth-Jensen, T. Horn, M. P. Kuchera, D. Lee, W. Nazarewicz, P. Ostroumov, K. Orginos, A. Poon, X.-N. Wang, A. Scheinker, M. S. Smith, L.-G. Pang, Colloquium: Machine learning in nuclear physics, Rev. Mod. Phys. 94 (2022) 031003. URL: https://link.aps.org/doi/10.1103/RevModPhys.94.031003. doi:10.1103/RevModPhys.94.031003.

[4] W.-B. He, Y.-G. Ma, L.-G. Pang, H.-C. Song, K. Zhou, High-energy nuclear physics meets machine learning, Nuclear Science and Techniques 34 (2023) 88. URL: https://link.springer.com/10.1007/s41365-023-01233-z. doi:10.1007/s41365-023-01233-z.

[5] M. Crispim Romão, N. F. Castro, R. Pedro, Finding new physics without learning about it: anomaly detection as a tool for searches at colliders, The European Physical Journal C 81 (2021) 27. URL: https://link.springer.com/10.1140/epjc/s10052-020-08807-w. doi:10.1140/epjc/s10052-020-08807-w.

[6] V. S. Ngairangbam, M. Spannowsky, M. Takeuchi, Anomaly detection in high-energy physics using a quantum autoencoder, Phys. Rev. D 105 (2022) 095004. URL: https://link.aps.org/doi/10.1103/PhysRevD.105.095004. doi:10.1103/PhysRevD.105.095004.

[7] J. Collins, K. Howe, B. Nachman, Anomaly detection for resonant new physics with machine learning, Phys. Rev. Lett. 121 (2018) 241803. URL:

https://link.aps.org/doi/10.1103/PhysRevLett.121.241803. doi:10.1103/PhysRevLett.121.241803.

[8] M. Kuchera, R. Ramanujan, J. Taylor, R. Strauss, D. Bazin, J. Bradt, R. Chen, Machine learning methods for track classification in the at-tpc, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 940 (2019) 156–167. URL: https://www.sciencedirect.com/science/article/pii/S0168900219308046. doi:https://doi.org/10.1016/j.nima.2019.05.097.

[9] R. Solli, D. Bazin, M. Hjorth-Jensen, M. Kuchera, R. Strauss, Unsupervised learning for identifying events in active target experiments, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1010 (2021) 165461. URL: https://www.sciencedirect.com/science/article/pii/S0168900221004460. doi:https://doi.org/10.1016/j.nima.2021.165461.

[10] H. Wu, Y. Wang, Y. Wang, X. Deng, X. Cao, D. Fang, W. Ma, W. He, C. Fu, Y. Ma, Machine learning method for 12c event classification and reconstruction in the active target time-projection chamber, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1055 (2023) 168528. URL: https://www.sciencedirect.com/science/article/pii/S0168900223005181. doi:https://doi.org/10.1016/j.nima.2023.168528.

[11] C. R. Qi, H. Su, K. Mo, L. J. Guibas, Pointnet: Deep learning on point sets for 3d classification and segmentation, CoRR abs/1612.00593 (2016). URL: http://arxiv.org/abs/1612.00593. arXiv:1612.00593.

[12] Y. Giomataris, P. Rebourgeard, J. Robert, G. Charpak, Micromegas: a high-granularity position-sensitive gaseous detector for high particle-flux environments, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated

Equipment 376 (1996) 29–35. URL: https://www.sciencedirect.com/science/article/pii/0168900296001751. doi:https://doi.org/10.1016/0168-9002(96)00175-1.

[13] M. Cortesi, S. Rost, W. Mittig, Y. Ayyad-Limonge, D. Bazin, J. Yurkon, A. Stolz, Multi-layer thick gas electron multiplier (m-thgem): A new mpgd structure for high-gain operation at low-pressure, Review of Scientific Instruments 88 (2017) 013303. doi:10.1063/1.4974333.

[14] E. Pollacco, S. Anvar, H. Baba, P. Baron, D. Bazin, C. Belkhiria, B. Blank, J. Chavas, P. Chomaz, E. Delagnes, F. Druillole, P. Hellmuth, C. Huss, E. Galyaev, B. Lynch, W. Mittig, T. Murakami, L. Nalpas, J.-L. Pedroza, R. Raabe, J. Pibernat, B. Raine, A. Rebii, A. Taketani, F. Saillant, D. Suzuki, N. Usher, G. Wittwer, Get: A generic electronic system for tpcs for nuclear physics experiments, Physics Procedia 37 (2012) 1799–1804. URL: https://www.sciencedirect.com/science/article/pii/S1875389212018974. doi:https://doi.org/10.1016/j.phpro.2012.02.506, proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011).

[15] J. Giovinazzo, T. Goigoux, S. Anvar, P. Baron, B. Blank, E. Delagnes, G. Grinyer, J. Pancin, J. Pedroza, J. Pibernat, E. Pollacco, A. Rebii, T. Roger, P. Sizun, Get electronics samples data analysis, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 840 (2016) 15–27. URL: https://www.sciencedirect.com/science/article/pii/S0168900216309408. doi:https://doi.org/10.1016/j.nima.2016.09.018.

[16] M. A. Fischler, R. C. Bolles, Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography, Commun. ACM 24 (1981) 381–395. URL: https://doi.org/10.1145/358669.358692. doi:10.1145/358669.358692.

[17] A. Anthony, Y. Ayyad, ACeulemans, freund17, lisacarpenter, cnhunt,

H. A. Pol, J. Zamora, N. Rijal, A. M. Ramos, bolaizol, AriAtari, Ruchi-GADGETII, Simon, Wieskejo, diazcort, Attpcroot, 2023. URL: https://doi.org/10.5281/zenodo.10113377. doi:10.5281/zenodo.10113377.

[18] J. Zamora, G. Fortino, Tracking algorithms for tpcs using consensus-based robust estimators, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 988 (2021) 164899. doi:10.1016/j.nima.2020.164899.

[19] Y. Ayyad, W. Mittig, D. Bazin, S. Beceiro-Novo, M. Cortesi, Novel particle tracking algorithm based on the random sample consensus model for the active target time projection chamber (at-tpc), Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 880 (2018) 166–173. URL: https://www.sciencedirect.com/science/article/pii/S0168900217311798. doi:https://doi.org/10.1016/j.nima.2017.10.090.

[20] A. Gade, B. Sherrill, NSCL and FRIB at Michigan State University: Nuclear science at the limits of stability, Physica Scripta 91 (2016) 053003.

[21] D. . J. Morrissey, B. Sherrill, M. Steiner, A. Stolz, I. Wiedenhoever, Commissioning the A1900 projectile fragment separator, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 204 (2003) 90–96.

[22] A. K. Anthony, C. Y. Niu, R. S. Wang, J. Wieske, K. W. Brown, Z. Chajecki, W. G. Lynch, Y. Ayyad, J. Barney, T. Baumann, D. Bazin, S. Beceiro-Novo, J. Boza, J. Chen, K. J. Cook, M. Cortesi, T. Ginter, W. Mittig, A. Pype, M. K. Smith, C. Soto, C. Sumithrarachchi, J. Swaim, S. Sweany, F. C. E. Teh, C. Y. Tsang, M. B. Tsang, N. Watwood, A. H. Wuosmaa, Beam particle identification and tagging of incompletely stripped heavy beams with HEIST, Review of Scientific Instruments 93 (2022) 013306. URL: https://doi.org/10.1063/5.0068180.

490 doi:10.1063/5.0068180. arXiv:https://pubs.aip.org/aip/rsi/
491 article-pdf/doi/10.1063/5.0068180/ 16576488/013306_1_online.pdf.

492 [23] H. He, E. A. Garcia, Learning from imbalanced data, IEEE Transactions
493 on Knowledge and Data Engineering 21 (2009) 1263–1284. doi:10.1109/
494 TKDE.2008.239.

495 [24] Y. Sun, A. K. C. Wong, M. S. Kamel, Classification of imbalanced data:
496 a review, Int. J. Pattern Recognit. Artif. Intell. 23 (2009) 687–719. URL:
497 https://api.semanticscholar.org/CorpusID:27118324.

Appendix A. Architecture and hyperparameter details

Table A.9 contains hyperparameters used in the model whose results are reported in this paper. The architectural details remained the same as specified in the PointNet paper [11] except where noted in Table A.9. Hyperparameters were found through a grid search and analysis of loss curves and validation results.

Hyperparameter	Value
input layer dim	64×4
output layer dim	2
batch size	32
epochs	30
optimizer	adam
learning rate	5×10^{-4}

Table A.9: Hyperparameters used in training the final model which correspond to all results presented in this paper.