# PHY379 Final Report

Cary Kan, Lyssa Dale-Davies

**Department:** University of Toronto, Department of Physics

**Supervisor:** Dr. Pekka K. Sinervo

**Date:** April 2026

**Abstract**

Dark matter is a theoretical form of matter which interacts gravitationally but does not interact with light. One of the most promising dark matter candidates is the weakly interactive massive particle (WIMP), which is the focus of the Super Cryogenic Dark Matter Search (SuperCDMS) experiment at SNOLAB, Ontario. The experiment aims to detect the nuclear recoils of WIMPs colliding with the target nuclei of its detector crystals. This report summarized our work training and using a machine learning (ML) model to reconstruct real position and energy test data from the Cryogenic Underground Tower Test (CUTE) facility using Germanium high voltage (HV) detectors. The model, a histogram gradient boosted regressor (HGBR), was trained using Detector Monte Carlo (DMC) samples. To account for more energy sharing observed in the CUTE data than in the DMC sample, we tested two forms of energy rescaling: a nonlinear rescale and a linear rescale applied to the DMC sample before training. We plotted RadialPar, a parameter used to see the correlation between energy sharing and radius and found the linear rescale model caused discontinuities in the plot, while the nonlinear rescale model appeared to accurately reconstruct RadialPar. We used the nonlinear rescale model for the rest of the analysis. We examined the effects of the rescale on the model resolution and found that the radial and X resolution worsened slightly, the Y resolution improved, and the Z and energy resolutions remained the same. The radial reconstruction for the K-shell improved after energy scaling, although did not reconstruct events all the way down to the center of the detector. There was also a large concentration of events reconstructed at the surface of the detector. We found the same for the L-shell radial reconstruction. The X and Y distributions for the K- and L-shells were uniform following energy rescale. We found that plotting $\sqrt{X^2+Y^2}$ yielded a smoother distribution than the original radial distribution for the K-shell, although there was still a slight peak at large radii. The energy distributions were also reconstructed for the K- and L-shell. The expected energy peaks were 10.37 keV and 1.30 keV while our reconstructed energy peaks were at 12 keV and 3.5 keV for the K- and L-shell, respectively. These results indicate there are still issues with the energy rescaling and calibration of the DMC.

# Introduction

## Dark Matter

There are a multitude of scientific observations which support the existence of dark matter. One prominent example was the observation of the rotational velocity curve of the Andromeda galaxy (M31) observed by Rubin and Ford in 1970 [2]. If all of M31's mass were baryonic, the rotational velocity curve shown in Figure 1 would decrease as $r^{-1/2}$ at large radii [8]. However, the figure shows that the actual curve appears constant for large radii, which was only possible if the galaxy contained a large amount of non visible matter at large distances from its center .

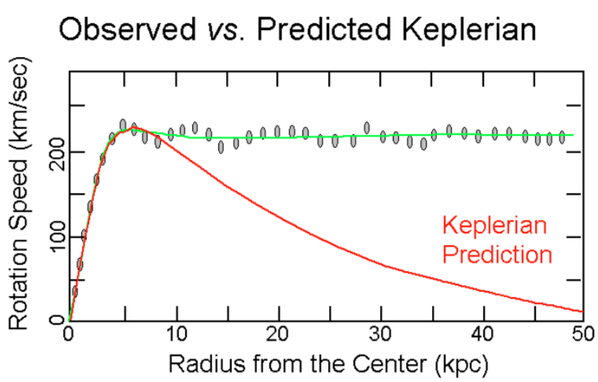


Figure 1: The figure shows the extented rotation curve of the Andromeda Galaxy. In red is the expected decrease and in green is the actual observed trajectory. Courtesy of Daniel Murnane [5].

This evidence, along with temperature fluctuations in the cosmic microwave background radiation, gravitational lensing, dynamics of colliding galaxy clusters, and other discoveries that cannot be explained by baryonic matter, imply the existence of dark matter.

## SuperCDMS Experiment

The Super Cryogenic Dark Matter Search (SuperCDMS) is a direct particle detection experiment which aims to detect the nuclear recoils of Weakly Interacting Massive Particles (WIMPs) colliding with target nuclei. WIMPs are one of the most promising candidates for dark matter. WIMPs are theorized to interact only with a weak force, which makes their detection a challenge. The sensitivity required to detect these nuclear recoils therefore necessitates detectors with excellent background discrimination and shielding from outside radiation and cosmic rays. SuperCDMS is one such direct detection experiment, building off of previous CDMS experiments to probe masses below 10 GeV/c$^2$ with greater sensitivity than previous experiments [11]. The detectors are shielded, located underground, and employ two types of detectors to achieve the desired background discrimination.

The design of the SuperCDMS detectors makes it possible to reconstruct the position and energy of events, which is crucial for background rejection of electron recoils and of events occurring on the surface or outside the detectors [3]. However, the resulting data files are large and complicated, making it difficult to extract useful information. Machine learning (ML) techniques are thus needed to extract the position and energy information from the raw signal data [3]. The ML techniques used in this report build on Matthew Penner's work detailed in his thesis "Position Reconstruction for the SuperCDMS Experiment" [7]. To train the models, we used a sample from the Detector Monte Carlo (DMC) simulation. To test the trained model on real data, we used data taken from Run 37 for Tower 3 testing at the SuperCDMS collaboration's Cryogenic Underground Tower TEst (CUTE).

# Background

## Detector Design

The detectors are housed in towers kept in a copper cryostat connected to a common dilution refrigerator capable of keeping the system at temperatures as low as 15 mK [11]. The cryostat is capable of housing 31 detector towers and houses four as a baseline [11]. The detectors themselves are cylindrical crystals of either germanium or silicon, with diameters of 100 mm and thicknesses of 33.3 mm [11]. SuperCDMS uses two different detector designs to improve performance. High Voltage (CDMS-HV, or just HV) detectors have excellent phonon sensor resolution in order to measure the lowest dark matter masses [11]. Interleaved Z-dependent Ionization and Phonon (iZIP) detectors measure both ionization and phonon signals to discriminate between nuclear recoil signals and electron recoil backgrounds for masses as low as 3 GeV/$c^2$ [11]. The combination of these detectors allows excellent background discrimination while searching for the smallest dark matter masses.

Figure 2 shows a simplified schematic of the interaction between a WIMP and a target nucleus in the detector. The collision generates an electron-hole pair and phonons which propagate through the crystal to be collected by the sensors. The bias voltage amplifies the production of phonons via the Neganov-Trofimov-Luke (NTL) effect to improve the phonon sensor resolution [11].

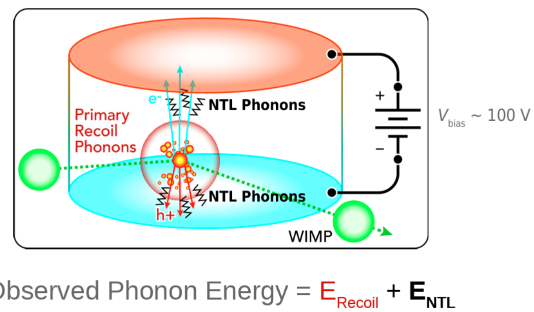


Figure 2: Figure shows the detector response to a WIMP particle detection. Image courtesy of Noah Kurinsky and Paul Brink [4].

The detectors use Quasiparticle-trapping-assisted electrothermal-feedback Transition-Edge Sensors (QETs) to measure the athermal phonon energy produced from particle interactions [11]. Superconducting aluminum fins are attached to the detectors in order to absorb the athermal phonons. The energy from this absorption breaks Cooper pairs in the Al fins, generating quasiparticles. A subset are then absorbed by tungsten Transition Edge Sensors (TESs), which causes a change in the electrical resistance of the tungsten. The electrothermal feedback required to return the TES to its nominal operating resistance provides a calibration for the received energy and generates the phonon signal as a smooth trace. Because the athermal phonons propagate quasi-ballistically through the crystal, the phonon signal carries event-position information needed for a “fiducialization” cut which removes events at the detector surface [11]. To access this information, the phonon sensors are partitioned into multiple channels. The phonon channel layout for a HV detector is shown in Figure 3. Both detectors have 12 such phonon channels, with 6 on each side [11]. For the HV detectors, the phonon channels are named P*S1 and P*S2 for either side of the detector, with * ranging from A-F. We will be focusing on the Germanium HV detectors for this report.

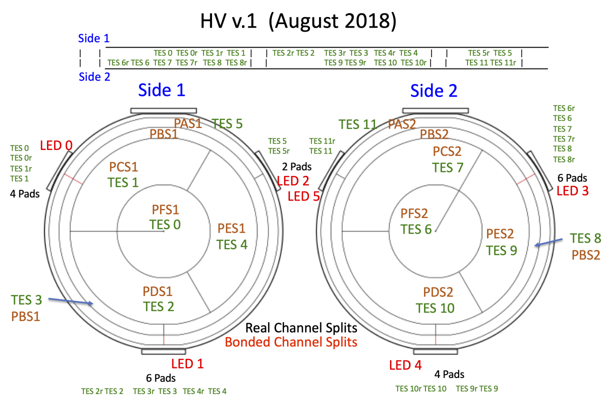


Figure 3: Phonon channel layout on either side of a HV detector. TESs are also labelled. Image courtesy of Noah Kurinsky and Paul Brink [4].

## Simulation

The simulation is split into two main stages. The first stage of the simulation, called SourceSim, simulates incident particles, calibration sources, and background sources. The second stage DetectorSim starts with simulating the detector response to interactions inside the crystal and is comprised of three components: CrystalSim, TESSim, and FETSim. CrystalSim simulates the propagation of phonons and charges in the crystal, TESSim simulates the TESs, and FETSim simulates the charge readout [3]. Next, noise is added to the pulse using DAQSim, which uses Power Spectral Densities (PSDs) of real noise recorded from previous CDMS experiments [3]. The pulses are then processed using the reconstruction package CDMSBats, which shifts the data to align the signal with a pulse template derived from real data. The outputs include Reduced Quantities (RQs) describing the shifts made, with the phonon channel pulse amplitudes (P*OFamps) and time delays (P*OFdelay) as well as the total phonon amplitude (PTOFamps) being the most relevant to this report [1]. We refer to the samples as DMC samples.

DMC samples undergo two types of calibration before being used to train ML models: standard calibration and relative calibration. Standard calibration converts values with units µA to units of energy (keV). Relative calibration is used to make the channel responses as identical as possible in order to reconstruct the energy properly. To understand, consider if some channel P*S1 was compared to another channel P**S1, and it was observed that P*S1 was degraded and only recorded 60% of the total energy. This would cause the same event to be reconstructed with different amplitudes depending on its proximity to either P*S1 or P**S1. Relative calibration is therefore necessary for events to be reconstructed accurately.

[1]The ”*” represents the phonon channel A-F.

## Machine Learning

The basis for the code used comes from Matt Penner's SuperCDMS Event Reconstruction, which in turn builds upon Getin Qin's research in the paper "Ensemble Learning Model for Improving Position and Energy Resolution for SuperCDMS" [7]. This code was further cleaned up and edited by Dr. Madeleine Zurowski and Warren Perry. The training starts by initializing an Estimator class, which is done by defining an ML model, a list of observable parameters the model will use to make predictions, a list of parameters to reconstruct, and a function which reformats the data. The observable parameters used in this report are P*OFamps, PTOFamps, and P*OFdelay, while the reconstructed parameters are the energy, X, Y, Z, and R positions. After matching the observables with the DMC positions, the events are saved in a CSV file. The reconstructed radial parameter is the most relevant here, and will be the main parameter studied in this report. Because the crystal is a uniform cylinder, we expect the number of events to increase uniformly with radius.

We will also use the parameter "RadialPar", defined by

$$\text{RadialPar} = \frac{PF - PA}{PF + PA}, \tag{1}$$

where $PF = PFS1 + PFS2$ and $PA = PAS1 + PAS2$, to check the quality of the reconstruction. Because we can expect events occurring near the center of the radius to be detected more by the central PF channels and vice versa for events occurring near the edge for PA channels, we will expect an anti-correlation between RadialPar and radius.

We used a Histogram Gradient Boosting Regressor (HGBR) model from the Python library scikit-learn. A HGBR model is an ensemble model which builds a sequence of decision trees, with each new tree trained to correct the errors of the predictions of all previous trees [7].

To test the model's ability to accurately reconstruct energy data, we can study how the model reconstructs the energy peaks of the CUTE data. The $^{70}$Ge in the detector produces the isotope $^{71}$Ge via radiation from a $^{252}$Cf source [9]. $^{71}$Ge then decays spontaneously through electron capture, a process in which an unstable nucleus captures one of its atomic electrons. The bound electron will subsequently combine with a proton in the nucleus to form a neutron and an electron neutrino:

$$p + e^- \rightarrow n + \nu_e \tag{2}$$

This leaves the atom in an excited ionized state, causing it to emit X-rays or Augur electrons[1]. Figure 4 shows electron capture occurring in the K-shell of an arbitrary atom.

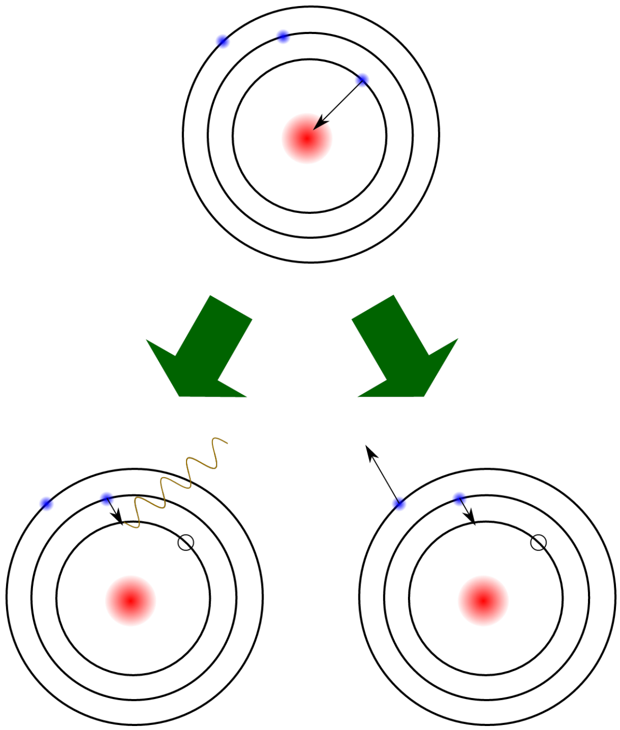

Figure 4: The figure depicts two types of electron capture decay occurring in the K-shell of an arbitrary atom. *Top*: The nucleus absorbs an electron in the k-shell. *Lower left*: An L-shell electron replaces the vacancy, and electromagnetic radiation equal in energy to the difference between the two shells is emitted. *Lower right*: Instead of emission as gamma rays, the energy absorbed through the process is instead transferred to an outer shell electron. That outer shell electron is ejected from the atom in the Auger effect, leaving a positive ion [6].

The capture of either a K, L, or M shell electron generates the characteristic emission lines corresponding to the K-, L-, and M-shell peaks. In this report, we will be be looking at only the K-shell and L-shell energies. In $^{71}$Ge, these correspond to energy peaks of 10.37 keV and 1.30 keV. These peaks provide a way to verify that the model is reconstructing CUTE energy data correctly and to identify mismatches between the DMC and CUTE.

# Results

## CUTE K-Shell Radial Reconstruction

We trained the Histogram Gradient Boosting Regressor (HGBR) model wrapped with a Multi Output Regressor on a K-shell 100k event 50V DMC sample. The CUTE sample we used in the reconstruction was the 50V low background sample from Run 37 of the Tower 3 testing at CUTE. The hyperparameters were found using GridSearchCV. The validation plots are shown in Figure 5 and Figure 6, while the CUTE reconstruction is plotted in Figure 7.

Need a paragraph describing the definition of the CUTE sample in terms of at least energy requirements.

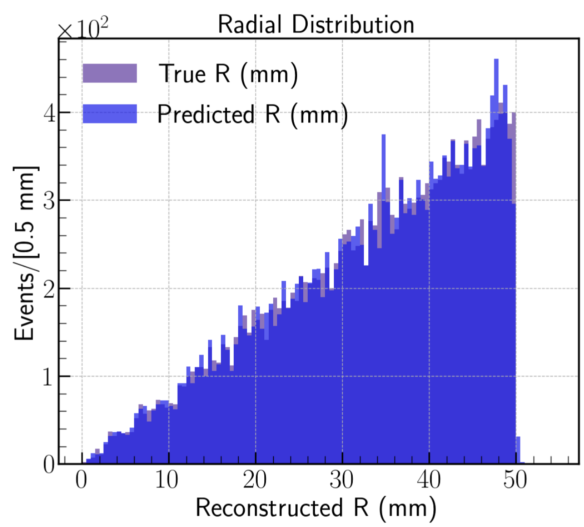


Figure 5: Radial validation plot as a histogram for the HGBR model trained on the 100k events DMC K-shell sample.

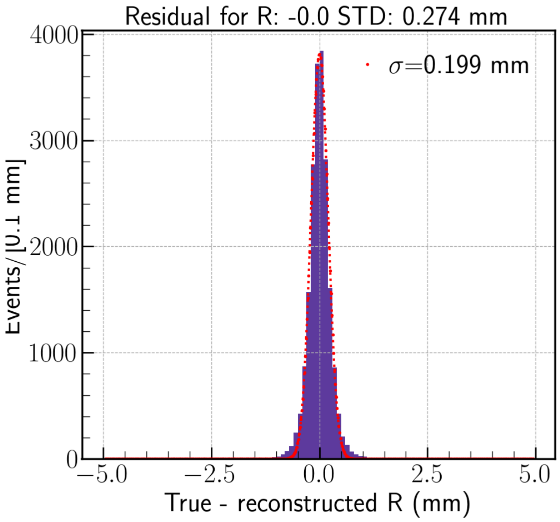


Figure 6: Validation residuals for the HGBR model trained on the 100k events DMC K-shell sample. The mean is 0.000 mm and the standard deviation is 0.274 mm.

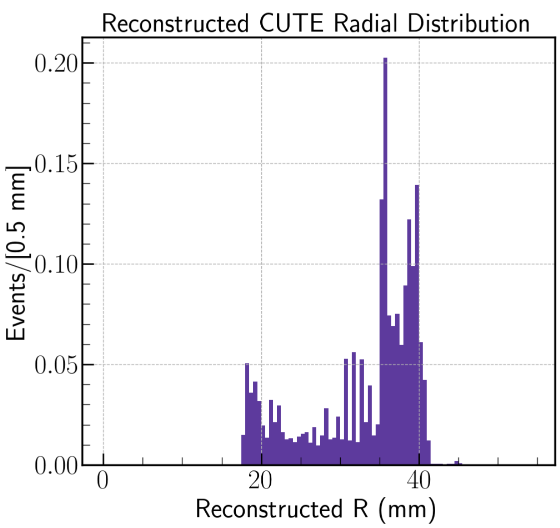


Figure 7: The figure shows the CUTE radial data plotted as a histogram after reconstruction by the HGBR model trained on the DMC sample.

The plot in Figure 5 increases linearly with the radius, which was expected. When used to reconstruct the CUTE data, however, the model does not yield the expected linear trend and only spans between 17 mm and 43 mm.

## Data-driven Rescale of DMC Response

To further investigate these behaviours, in Figure 8 we plotted PASOFamps/PTOFamps vs PFSOFamps/PTOFamps for the raw CUTE data and the DMC sample.

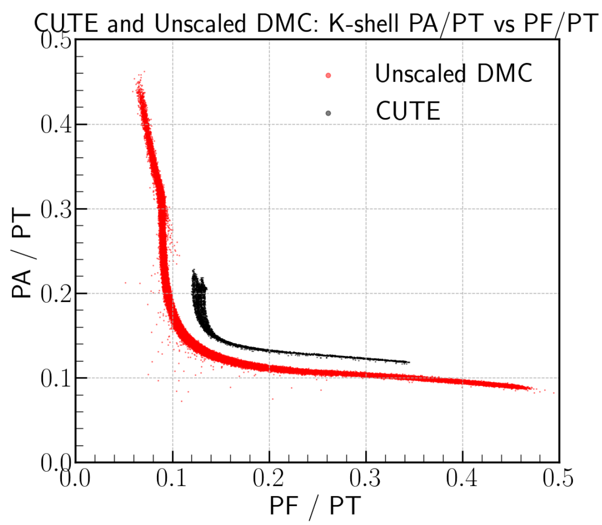


Figure 8: PA/PT vs PF/PT is plotted for the DMC sample and the CUTE data. No model has been applied to either the data or the sample.

Figure 8 shows that there are noticeable differences between DMC and CUTE data, with the most apparent difference being that the DMC predicts a stronger response in the outer channel for events occurring near the edge of the detector and a stronger response in the inner channel for events occuring near the center of the detector. This suggests that more energy is being shared between channels in the CUTE data than in the DMC. In order for DMC response to be compatible with CUTE, we must account for the difference in phonon channel response. Warren Perry found a method to rescale each channel's response by scaling each channel's amplitude. The total energy $E_t$ is given by

$$E_t = E_r + E_{eh} = \sum P^*OFamps \quad (3)$$

Here, $E_r$ is the recoil energy, and $E_{eh}$ is the energy generated by the electron hole pairs produced in each event. The P*OFamps are the separate channel amplitudes from both faces of the detector ranging from A to F, which we will denote as $E_i$. We can write $E_i$ as

$$E_i = \frac{1}{12}E_r + \beta_i E_{eh} \quad (4)$$

This assumes that recoil energy is distributed uniformly across all 12 channels. However, since the electron hole contribution is not expected to be uniform, $\beta_i$ acts as a channel specific coefficient. The relative calibration is then applied to $E_i$ as a fraction of the total energy, where the correction can be applied using either a linear:

$$E_i' = (a_i(\frac{E_i}{E_t}) + b_i) \times E_t \quad (5)$$

or nonlinear function:

$$E_{new} = (a_i(\frac{E_i}{E_t})^2 + b_i(\frac{E_i}{E_t}) + c_i) \times E_t \quad (6)$$

The factors $a_i, b_i, \&\ c_i$ for each channel were computed by Warren Perry 'by eye' and are shown in Table 1 and table 2.

| Channel | a | b | c |
|---|---|---|---|
| PAS1 | -2 | 0.77 | 0.03 |
| PBS1 | -2.8 | 1 | 0.023 |
| PCS1 | 0 | 0.6 | 0.038 |
| PDS1 | 0 | 0.6 | 0.03 |
| PES1 | 0 | 0.62 | 0.032 |
| PFS1 | 0 | 0.6 | 0.035 |
| PAS2 | -2 | 0.8 | 0.028 |
| PBS2 | -2.8 | 1 | 0.023 |
| PCS2 | 0 | 0.6 | 0.035 |
| PDS2 | 0 | 0.68 | 0.025 |
| PES2 | 0 | 0.65 | 0.032 |
| PFS2 | 0 | 0.65 | 0.035 |

Table 1: Nonlinear relative scaling factors for each channel found by Warren Perry [10].

| Channel | a | b | c |
|---|---|---|---|
| PAS1 | 0 | 0.35 | 0.045 |
| PBS1 | 0 | 0.53 | 0.04 |
| PCS1 | 0 | 0.6 | 0.038 |
| PDS1 | 0 | 0.6 | 0.03 |
| PES1 | 0 | 0.62 | 0.032 |
| PFS1 | 0 | 0.6 | 0.035 |
| PAS2 | 0 | 0.44 | 0.04 |
| PBS2 | 0 | 0.53 | 0.04 |
| PCS2 | 0 | 0.6 | 0.035 |
| PDS2 | 0 | 0.68 | 0.025 |
| PES2 | 0 | 0.65 | 0.032 |
| PFS2 | 0 | 0.65 | 0.035 |

Table 2: Linear relative scaling factors for each channel found by Warren Perry [10].

## Linear Versus Nonlinear Rescale

To evaluate the linear rescale against the nonlinear rescale, we will compare RadialPar CUTE plots produced using nonlinear and linear rescaling. Figure 9 plots RadialPar vs radius for the nonlinear rescaled DMC and produces the expected anti-correlation as a smooth curve.

In Figure 10, the linear rescale causes a discontinuity at a radius of approximately 28 mm to 34 mm. This discontinuity could be caused by surface edge effects at the edge of different phonon channels, but a full investigation is beyond the scope of this report. For our purposes, this indicates the nonlinear rescaling more accurately reconstructs the CUTE data. We will use the nonlinear rescaling for all following training and reconstruction.

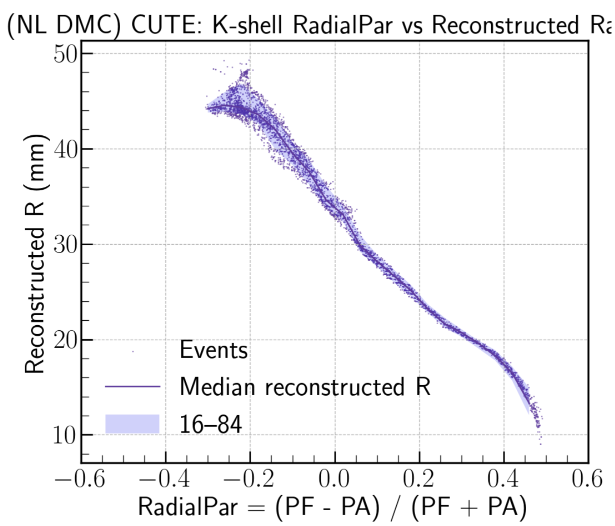


Figure 9: The figure shows reconstructed radius as a function of RadialPar for CUTE K-shell using HBGR trained on a nonlinearly rescaled 1 million event 50V sample in the 0.1–12 keV range. The shaded band is $\approx \pm 1\sigma$.

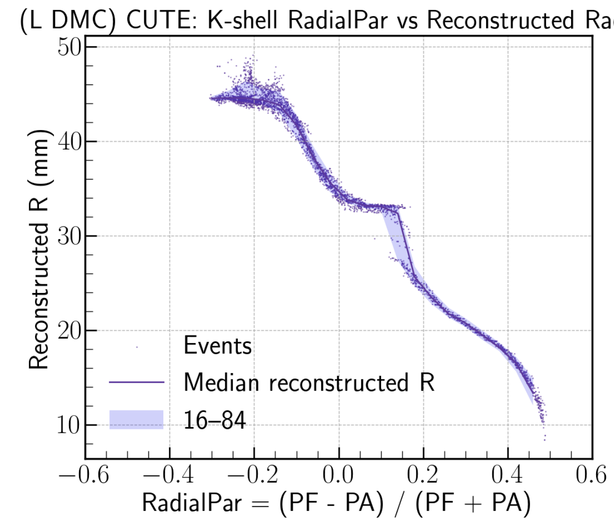


Figure 10: The figure shoes reconstructed radius as a function of RadialPar for CUTE K-shell using a model trained on a linearly rescaled 1 million event 50V sample in the 0.1–12 keV range.

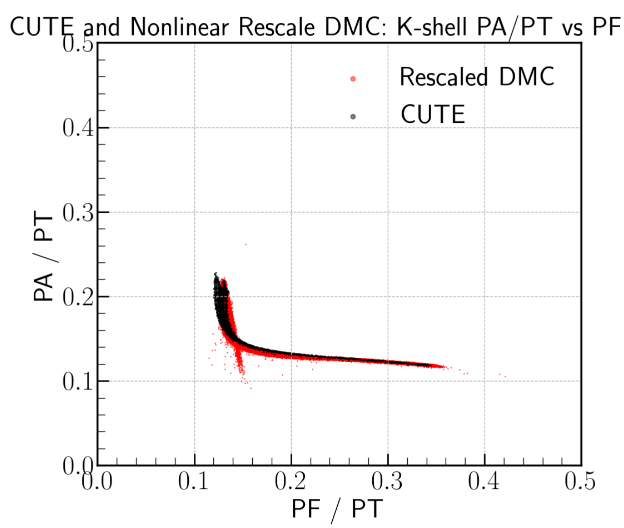


Figure 11: PA/PT vs PF/PT is plotted for the DMC sample with nonlinear rescaling and the CUTE data. The plots show only K-shell energies.

Figure 11 shows PA/PT vs PF/PT results with nonlinear scaling applied to the DMC are shown with the CUTE data plotted on top.

The DMC overlaps much more with CUTE than Figure 8, which is a promising sign for the nonlinear rescaling. There is some nonlinear behaviour in the DMC for lower PA/PT and PF/PT, indicating there are still some issues with the DMC and energy rescaling.

## Effects of Data-driven Rescale on DMC Reconstruction

Now using the nonlinear rescaling, we can examine the effects of the rescaling on the model's resolution through the validation residuals. The sample used is a 1 million event 50V event sample in the 0.1–12 keV range.

The resulting distributions, shown in Figures 12, 13, 14 and 27 in Appendix , were evaluated using a Gaussian fit to determine bias and resolution. Although the rescale may improve the discrepancy between the DMC and CUTE, we expect the rescale to worsen the model's position resolution. This is because the rescaling of channel response to spread the energy more evenly across all channels should result in less position sensitivity. We plotted the before and after rescaling model residuals for radius, X, and Y in Figures 12, 13, and 14, respectively.

In Figure 12 (a), before the rescale, the residual is centered very close to zero. In Figure 12 (b), after the rescale, the distribution shifts negative and widens. The bias suggests that the model now tends to under-predict the radius and the resolution is worse by about 11%.

We see the same effects in Figure 13 for X, though the changes are less drastic since radial reconstruction depends more on channel response differences than X or Y reconstruction.

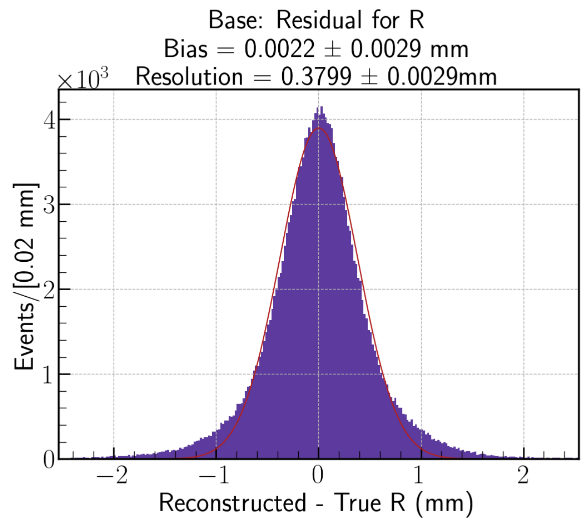


(a)

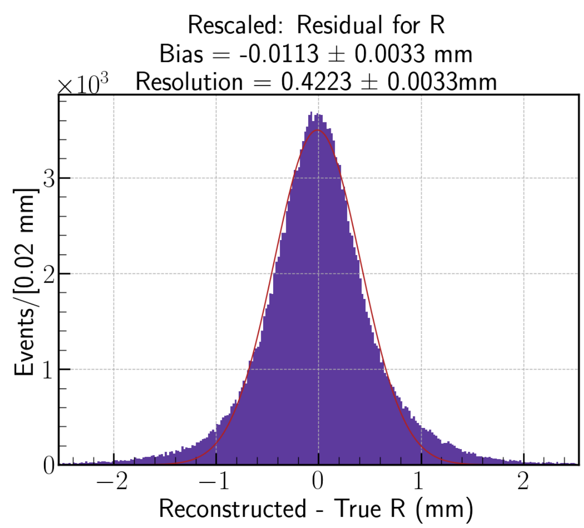


(b)

Figure 12: The figure shows a comparison of R residuals from validation. Figure (a) shows the model without scaling and figure (b) shows the model with nonlinear rescaling. The absolute difference in bias is 0.0135 mm and the absolute difference in resolution is 0.0424 mm.

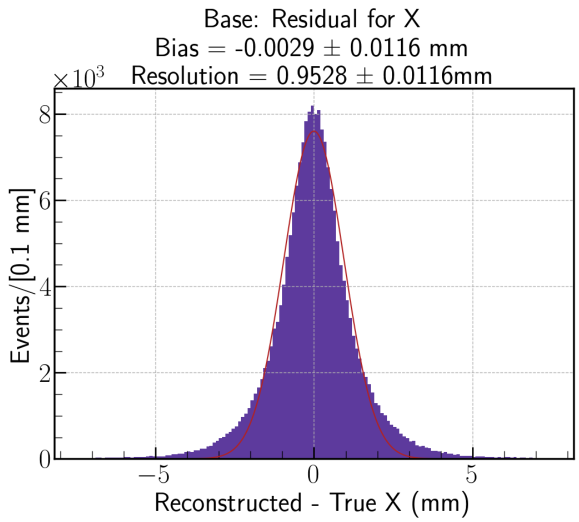


(a)

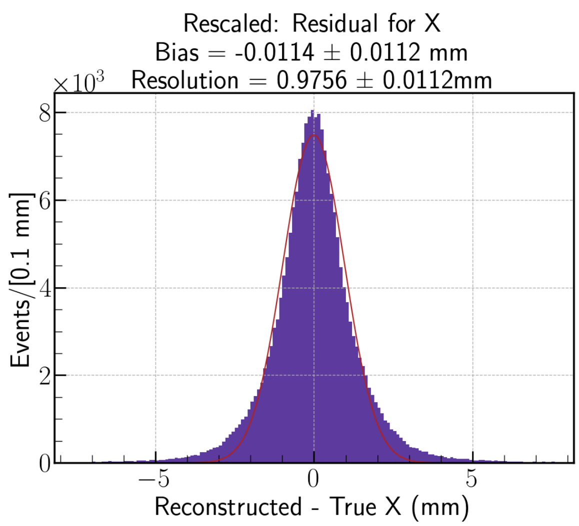


(b)

Figure 13: The figure shows the comparison of X residuals. In (a) the reconstruction was predicted using the DMC sample, while in (b) a nonlinear rescale was applied to the sample. For details see App.1. The absolute difference in bias is 0.0085 mm and the absolute difference in resolution is 0.0028 mm.

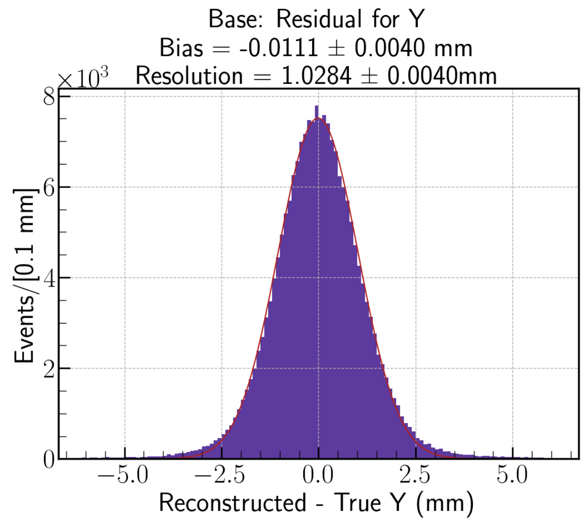


(a)

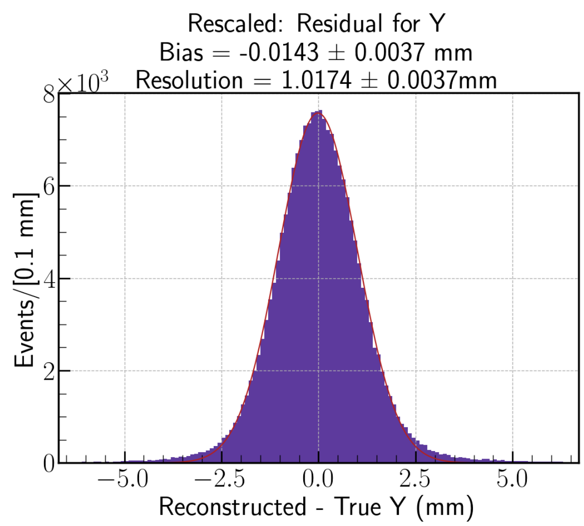


(b)

Figure 14: The figure shows the comparison of Y residuals. Figure (a) plots the residuals for the model trained on the unscaled DMC sample. Figure (b) plots the residuals for the model trained on the DMC sample after nonlinear rescaling. The absolute difference in bias is 0.0032 mm and the absolute difference in resolution is 0.011 mm.

We see an exclusion to this trend in the Y resolution in Figure 14.

Figure 14 has a small shift in bias of 0.0032 mm, and resolution actually improves from 1.0284 mm to 1.0174 mm. The same behavior can be seen in the Z and energy residuals, shown in Figure 27 in Appendix . However, in the case of energy and Z, the lack of change is expected as neither heavily on channel response differences. Z relies on timing information from P*OFdelays, while energy reconstruction depends mostly on total amplitude, which is conserved in the rescale.

## Data-driven Rescale on CUTE Reconstruction

### CUTE K-Shell Reconstruction

To further test the effects of rescaling, we retrained the model on the full DMC energy sample (validation plots shown in Figure 15 and Figure 15) and reconstructed the K-shell CUTE radial in Figure 17.

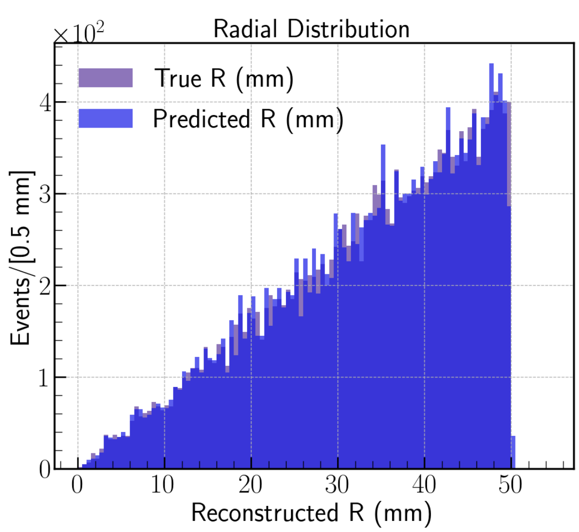


Figure 15: Figure shows the radial validation plot as a histogram for the HGBR model trained on the DMC sample after the data driven rescales were applied.

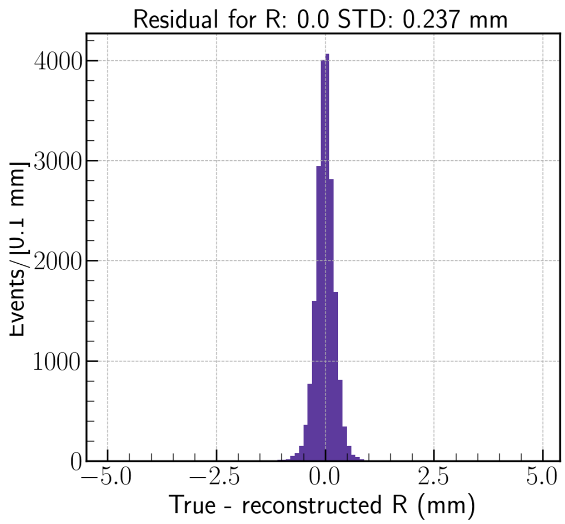


Figure 16: Validation residuals for the HGBR model trained on the DMC sample after the data driven rescales were applied. The mean is 0.001 mm and the standard deviation is 0.237 mm.

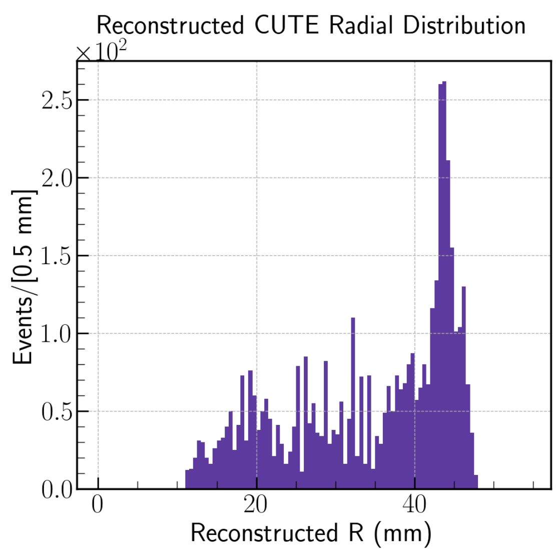


Figure 17: The figure shows the CUTE radial data as a histogram. The reconstruction was done by the HGBR model trained on the rescaled DMC sample.

The validation plots show the model is still working as expected. The CUTE reconstruction has improved from Figure 7, with a max around 50mm and a min around 7mm. The plot appears a bit more linear, though it is still not smooth and has a very large peak at large radii. This peak could be due to surface background events, but it may not be the only cause.

Motivated by the improved or unchanged model resolution Y shown in Figure 14, we next plotted the reconstructed Y and X positions as a scatterplot in Figure 18 and then used them to plot R as $\sqrt{X^2+Y^2}$ in Figure 19. This was done to see if reconstruction R from X and Y would be more accurate than directly reconstructing using the model.

Figure 18 appears uniformly distributed and Figure 19 is smooth, with a minimum close to 0 mm and a maximum a bit above 45 mm. There is a larger peak at large radii, but it is less pronounced than the one seen in Figure 17.

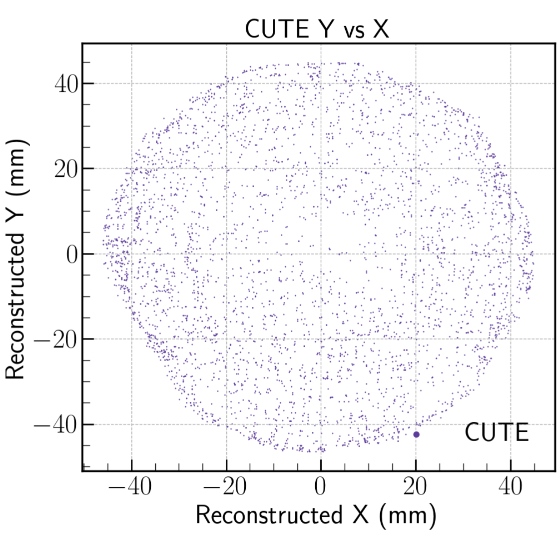


Figure 18: CUTE Y vs X reconstruction plotted as a scatterplot. The events appear mostly uniformly distributed.

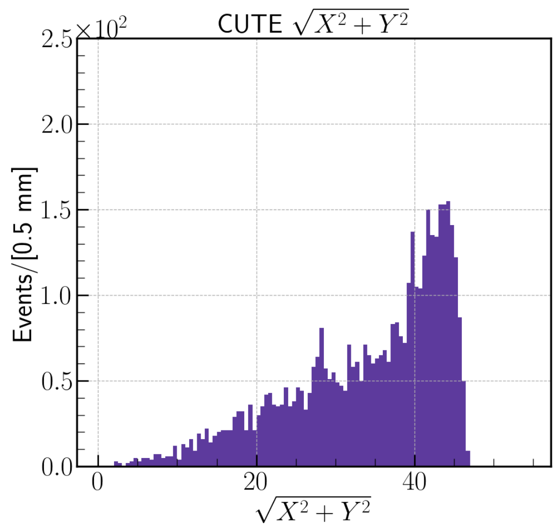


Figure 19: CUTE $\sqrt{X^2+Y^2}$ plotted as a histogram.

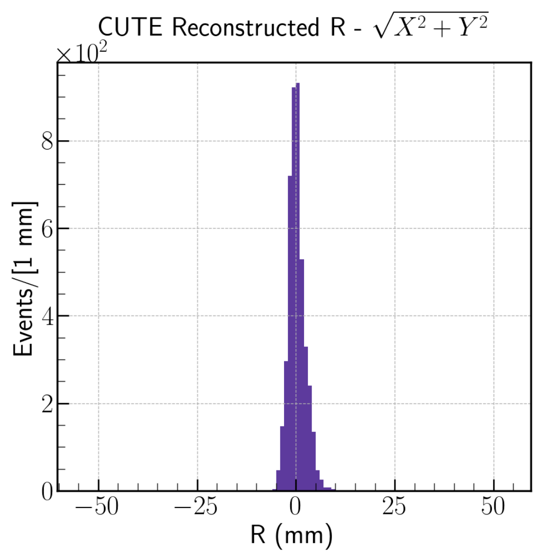


Figure 20: CUTE R-$\sqrt{X^2+Y^2}$ plotted against reconstructed R as a histogram. The plot is centered at 0 mm, with a standard deviation of 1.69 mm.

To better see the difference between the radial reconstruction methods, Figure 20 shows the difference between reconstructed radius and $\sqrt{X^2+Y^2}$ against reconstructed radius.

It has a mean centered at 0 mm, but a resolution of 1.69 mm. This suggests the differences between the two methods varies in extremity. This agrees with the visual difference between the two plots, where the direct reconstruction is less smooth than and has a larger peak than the $\sqrt{X^2 + Y^2}$ reconstruction.

Next, Figure 21 plots the reconstructed energy distribution for the CUTE K-shell.

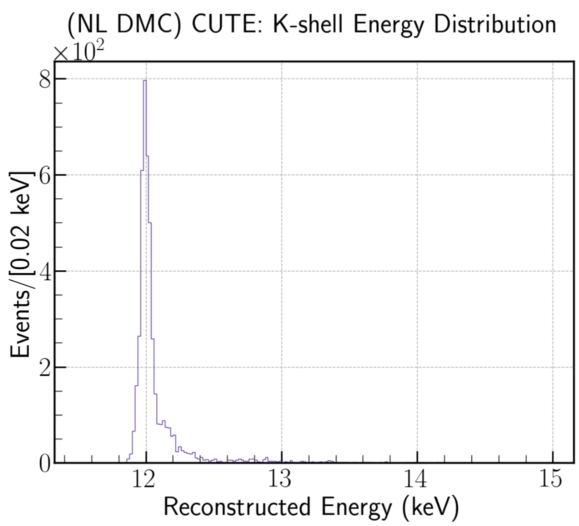


Figure 21: The figure shows reconstructed energy distribution for CUTE K-shell using a model trained on a nonlinearly rescaled 1 million event 50V sample in the 0.1–12 keV range. The K-shell has 4392 events. The peak is concentrated at $\approx$ 12 keV.

In Figure 21, the distribution is narrow and dominant, which indicates that reconstruction resolution is relatively high. However, there is a clear offset from the expected K-shell energy, with the peak at around 12 keV instead of the expected 10.37 keV. Along with this offset, the high energy tail is non-negligible, suggesting a problem with the CUTE energy reconstruction. This issue, in addition to the remaining problems with the radial reconstruction, indicates a mismatch between the K-shell DMC and CUTE data.

### CUTE L-shell Reconstruction

RadialPar analysis of L-shell in Fig.22 shows the anti correlation we expect, although appears fainter than the K-shell distribution in Figure 9 due to fewer L-shell events. Notably, the trend if also entirely smooth, with no discontinuity as seen in Fig.10, which further suggests that the nonlinear rescaling better reconstructs CUTE.

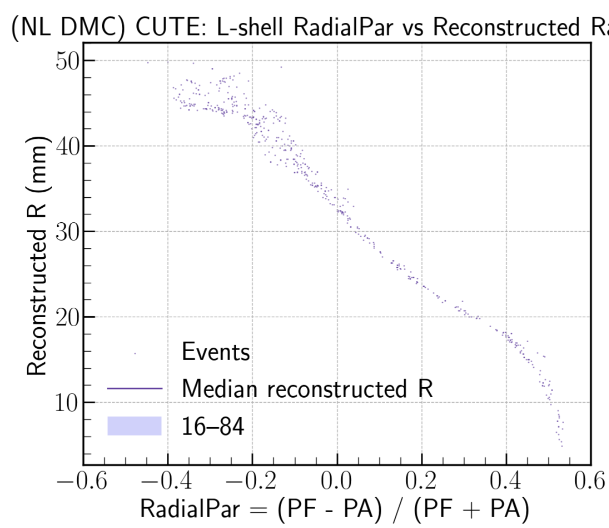


Figure 22: The figure shows reconstructed radius as a function of RadialPar for CUTE L-shell. Each point corresponds to an individual event, the solid curve is the median reconstructed radius in bins of RadialPar, and the shaded band is $\approx \pm 1\sigma$.

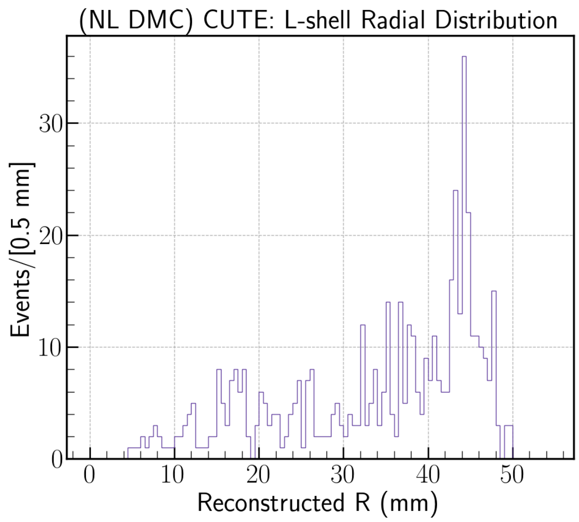


Figure 23: The figure shows reconstructed radial distribution for CUTE L-shell using the HGBR model trained on the nonlinearly rescaled 1 million event 50V sample in the 0.1–12 keV range. The L-shell has 516 events.

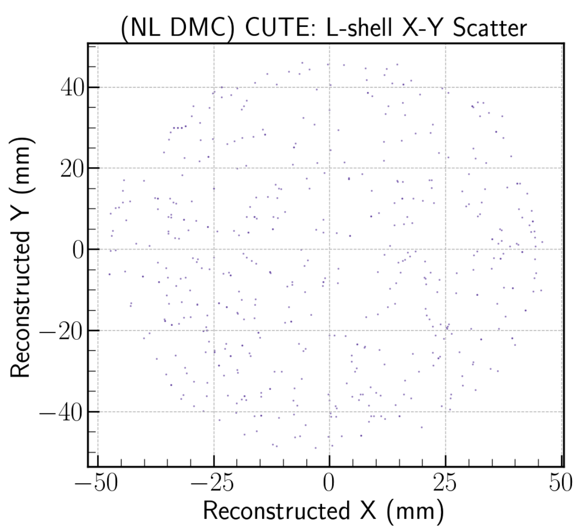


Figure 24: The figure shows reconstructed X-Y distribution for CUTE L-shell.

Figure 23 shows that the L-shell events are not uniformly distributed in radius, but the increase roughly follows a linear trend. Similar to the K-shell radial distribution, the highest concentration of events occurs near the edge of the detector, around 43 mm to 46 mm. However, there is also significant noise in the mid-section of the distribution, and a sharp drop at the very edge of the detector. The results suggest that, similar to the results for the K-shell, there are still low energy background and detector edge effects that are distorting the results.

Fig.24 shows that the X-Y scatter for L-shell is indeed circular and extends nearly to the detector edges. We can also confirm what we saw in the radial distribution in Fig.23, which is a 'ring' of events at higher radius. The consistency is reassuring, but this further suggests presence of edge effects.

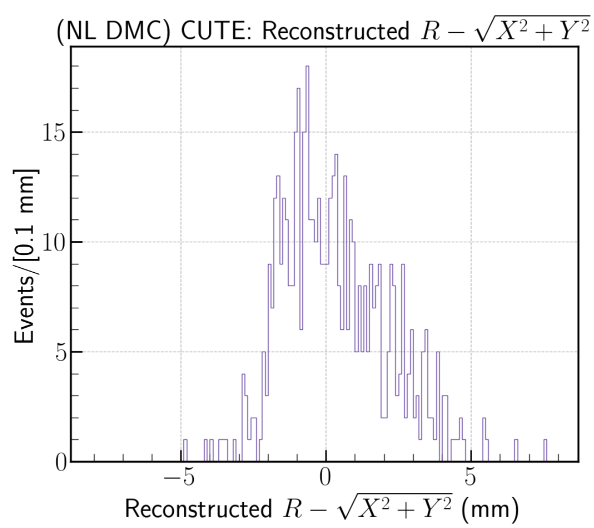


Figure 25: L-shell CUTE R-$\sqrt{X^2+Y^2}$ plotted against reconstructed R as a histogram. The plot is centered slightly left of 0 mm, with a standard deviation of about 1.3 mm.

Figure 25 shows the difference between reconstructed R and $\sqrt{X^2+Y^2}$ for the L-shell.

The mean is slightly left of 0 mm and the standard deviation is 1.3 mm. This implies the two methods produce noticeable differences, which is consistent with what we found for the K-shell.

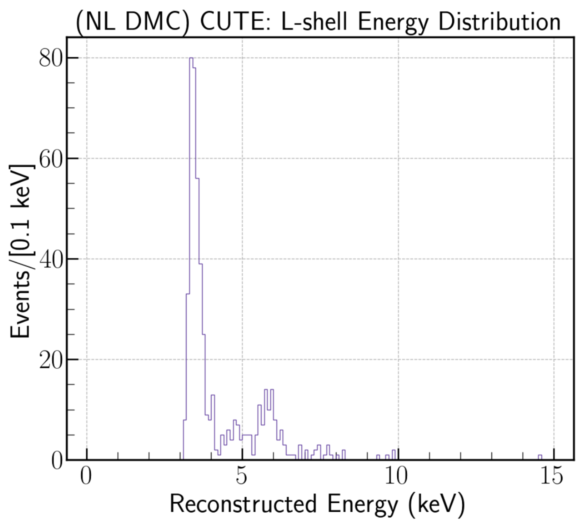


Figure 26: The figure shows reconstructed energy distribution for CUTE L-shell. The peak is concentrated at $\approx 3.5\,\text{keV}$.

Fig.26 shows the same kind of behavior we observed in the energy distribution of the K-shell in Fig.21. The entire peak is shifted to a substantially higher energy at $\approx 3.5\,\text{keV}$ instead of occurring at the expected L-shell peak of $1.30\,\text{keV}$. There is also a broader and noisier tail, but this is to be expected since the L-peak occurs at a much lower energy than the K-peak, and lower energy events are intrinsically harder to reconstruct. Moreover, the L-shell region is more susceptible to backgrounds since it is lower in energy. Regardless, the large discrepancy between expected energy and the reconstructed energy means there is a mismatch between the L-shell DMC sample and the CUTE data even after applying nonlinear rescaling.

# Discussion

The shell study energy and radial distributions yielded promising results. RadialPar's response to nonlinear and linear scaling indicated that the linear scaling was not sufficient to accurately model the energy sharing in the crystal. From this we can conclude that there are nonlinear effects in how the real energy is distributed across channels.

Beyond RadialPar, the nonlinear scaling improved the reconstructed radial distribution significantly. Using the rescaled training, the radial reconstruction reached a minimum of about 7 mm and a maximum around 50 mm for the K-shell, while the L-shell radial distribution appeared to reach a minimum around 5 mm and a maximum of 50 mm. Both distributions had the highest event concentration located near the surface at the detector, manifesting visually as large peaks.

Plotting the K-shell radial distribution using reconstructed X and Y in $\sqrt{X^2+Y^2}$ showed an improvement over the model's reconstruction. The reason for this could tie back to the change in the model resolutions in , where the radial resolution got worse while the X and Y resolutions remained relatively unchanged post nonlinear rescale. Figure 20 and Figure 25 demonstrate that for both the K-shell and L-shell, the two methods have noticeable differences. Further investigation will be needed to fully determine the reason and decide which method produces better results.

The reconstructed energy peaks are much too high, with a 2.5-3 keV discrepancy between the results and what we expect the energies to be (this discrepancy is shown visually in units of $\mu A$ for the K-shell and L-shell in Figure 29 and Figure 30, respectively). The reasons behind this are currently under investigation.

# Conclusion

The results show that nonlinear energy rescaling of the DMC improves the model's reconstruction of CUTE position data. However, the non-uniformity of the reconstructed radial distribution and the large peak in events at large radii are issues which could be caused by background surface events or more problems with the energy scaling. We also saw these issues in the reconstructed X, Y, and $\sqrt{X^2+Y^2}$ distributions, but to a lesser extent. The $\sqrt{X^2+Y^2}$ distribution, in particular was much smoother and more uniform than the direct radial reconstruction. This could be caused by the model's radial resolution worsening with the nonlinear rescales while the X and Y changed very little, if at all. The radial L-shell reconstruction had similar issues with a large peak at large radii and a lack of uniformity. Exploring the reason behind the large peak and investigating the different radial reconstruction methods are the next steps to improving the radial distribution.

Another issue we have is the discrepancy between CUTE's expected total energy for the K- and L-shells and the reconstructed energies. This is possibly due to errors with the calibration of the energy that can be fixed by changing certain calibration factors and ensuring we are working with correct units. More work needs to be done to understand the effects of the energy rescaling and relative calibration on the reconstructed energy if we want to accurately reconstruct energy and position information. Regardless, the improvements made by the energy rescaling show promise towards training models that accurately reconstruct position and energy data.

# References


[1] W. Bambynek et al. "Orbital electron capture by the nucleus". In: *Rev. Mod. Phys.* 49 (1 Jan. 1977), pp. 77–221. DOI: 10.1103/RevModPhys.49.77. URL: https://link.aps.org/doi/10.1103/RevModPhys.49.77.

[2] Gianfranco Bertone and Dan Hooper. "History of dark matter". In: *Reviews of Modern Physics* 90.4 (Oct. 2018). ISSN: 1539-0756. DOI: 10.1103/revmodphys.90.045002. URL: http://dx.doi.org/10.1103/RevModPhys.90.045002.

[3] Cui, Stephanie. *Optimizing Machine Learning Techniques for Dark Matter Detection: Enhancing Position Resolution in SuperCDMS Detectors*. Tech. rep. Department of Physics, University of Toronto, 2024. URL: https://confluence.slac.stanford.edu/spaces/CDMS/pages/424774503/SuperCDMS+Detector+Signal+ML?preview=%2F424774503%2F532905711%2F2024_Summer_NSERC_Project_Report.pdf.

[4] N. Kurinsky and P. Brink. *100mm HV SNOLAB*. SuperCDMS Detector Design. URL: https://confluence.slac.stanford.edu/x/SNZ0Cw.

[5] Daniel Murnane. "Supersymmetric Dark Matter Searches with CMS Data". PhD thesis. Oct. 2014. DOI: 10.13140/RG.2.2.13923.50724.

[6] Pamputt. *Scheme of two types of electron capture*. Own work, licensed under CC BY-SA 4.0, Wikimedia Commons. 2015. URL: https://commons.wikimedia.org/w/index.php?curid=38001613 (visited on 04/20/2026).

[7] Matthew A Penner. "Position Reconstruction for the SuperCDMS Experimen". MA thesis. University of Toronto, 2025. URL: https://confluence.slac.stanford.edu/spaces/CDMS/pages/550702530/UofT+Event+Reconstruction+Summary+and+tutorial?preview=%2F550702530%2F630207426%2FPosition_Reconstruction_For_SuperCDMS_Experiment.pdf.

[8] M. Pepin. "Low-Mass Dark Matter Search Results and Radiogenic Backgrounds for the Cryogenic Dark Matter Search". PhD thesis. University of Minnesota, 2016.

[9] W. Perry. "Comparing Ge activation rates in data and simulation". SuperCDMS Collaboration. URL: https://confluence.slac.stanford.edu/x/LyndIw.

[10] W. Perry. "Comparing RadialPar in data and simulation". SuperCDMS Toronto Machine Learning Meetings. URL: https://confluence.slac.stanford.edu/x/xrMoJQ.

[11] SuperCDMS Collaboration. *SuperCDMS SNOLAB Technical Design Report*. URL: https://confluence.

slac . stanford . edu / display / CDMS / SuperCDMS + SNOLAB + Project ? preview = /180554723 / 445874563 / supercdms-snolab-fdr.pdf.

# Appendix

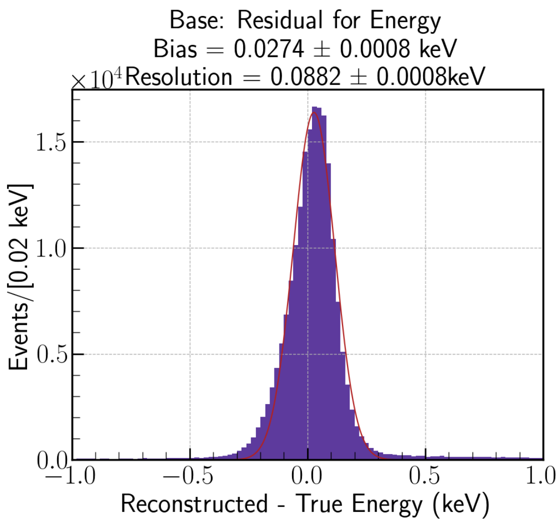


(a)

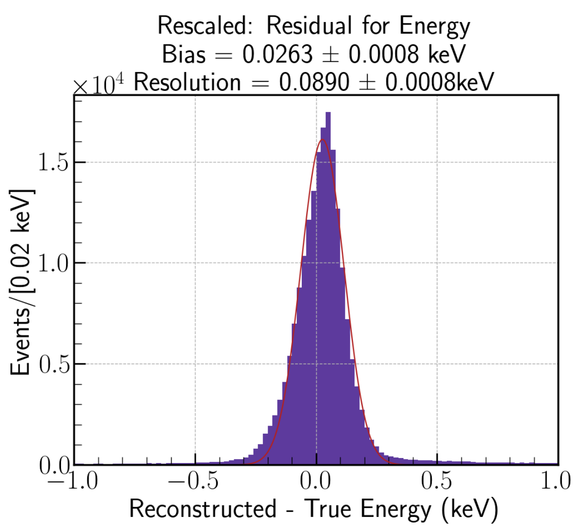


(b)

Figure 27: The figure shows a comparison of energy residuals from validation. In (a) the reconstruction was predicted using the DMC sample, while in (b) a nonlinear rescale was applied to the sample before training. The absolute difference in bias is 0.0011 mm and the absolute difference in resolution is 0.0008 mm.

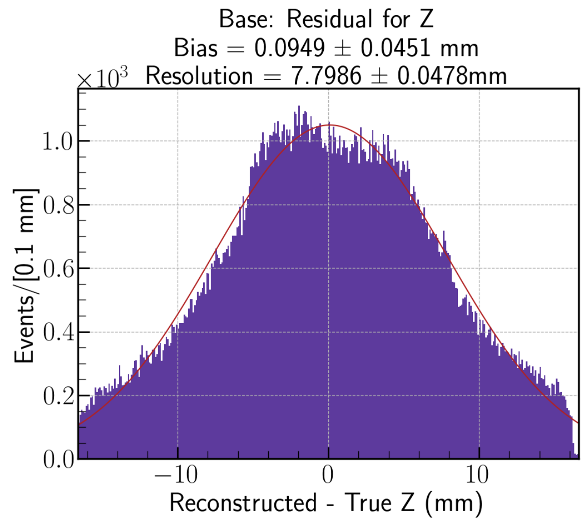


(a)

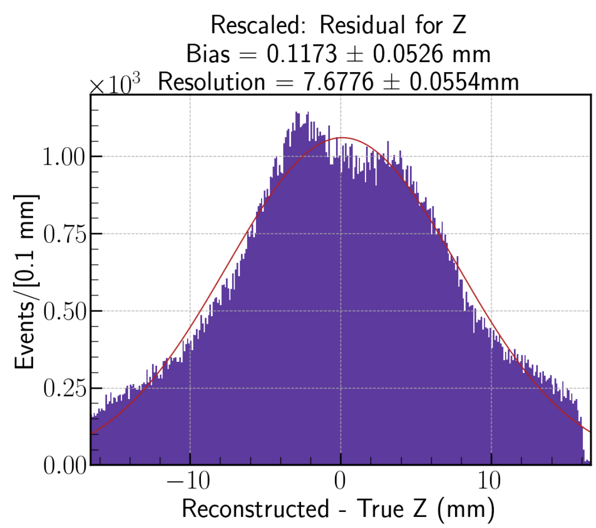


(b)

Figure 28: The figure shows a comparison of Z residuals from validation. In (a) the reconstruction was predicted using the DMC sample, while in (b) a nonlinear rescale was applied to the sample before training. The absolute difference in bias is 0.0224 mm and the absolute difference in resolution is 0.121 mm.

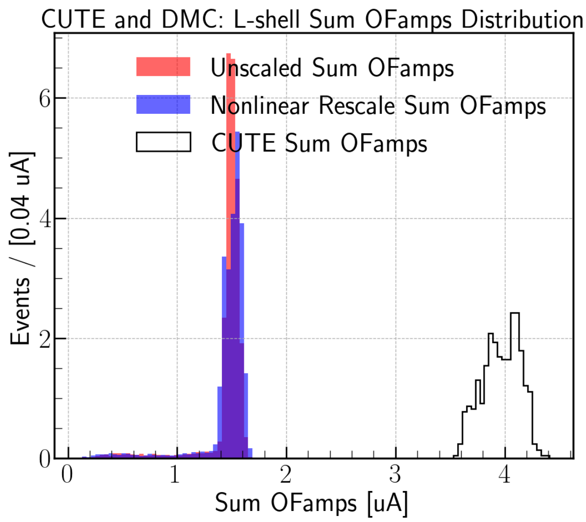


Figure 30: L-Shell energy peak distribution in $\mu$A plotted for unscaled DMC, nonlinearly scaled DMC, and reconstructed CUTE. The model is the HGBR model trained on the full energy DMC sample with nonlinear rescaling.

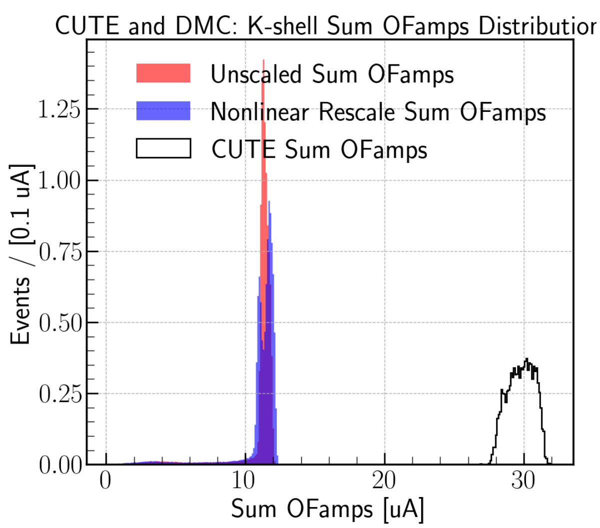


Figure 29: K-Shell energy peak distribution in $\mu$A plotted for unscaled DMC, nonlinearly scaled DMC, and reconstructed CUTE. The model is the HGBR model trained on the full energy DMC sample with nonlinear rescaling.