LHC

Scattering and Neutrino Detector
at the LHC

Hadronic shower reconstruction for the SND@LHC experiment

Stefan Katsarov
Supervisors: Martina Ferrillo, Antonio Iuliano, Cristovao Vilela

September 22, 2023

Abstract

The SND@LHC experiment aims to measure the cross-section of high-energy neutrinos produced in LHC collisions to study charmed hadron production and extract the gluon PDF at low momentum fractions. In this project, the first procedure for hadronic shower energy and vertex position reconstruction is developed for the SND detector. Shower propagation is profiled using PDFs and a log-likelihood function is built to reconstruct energy and vertex. Energy reconstruction is successful for showers initiated early in the target. For showers initiated further downstream low energies are overestimated while high energies are underestimated. Vertex reconstruction consistently overestimates the initiation depth, with the performance improving for later showering and higher energies. A discrepancy is identified with the true shower vertex depth. Results suggest vertex reconstruction would perform significantly better if the true vertex depth is corrected.

Contents

1 Introduction

1.1 Overview

The SND@LHC is a compact experiment which exploits the large flux of high-energy neutrinos, comprised of all three flavours, produced in LHC collisions at CERN, Geneva [1]. The experiment is located in the far-forward region from the ATLAS collision point (IP1). The detector is placed 480m downstream of IP1, in the TI18 tunnel, and covers the pseudo-rapidity (η) range from 7.2 to 8.4. The detector acceptance range and orientation can be seen in Figure 1 below. The highest energy and flux of neutrinos is produced along the collision axis, which corresponds to the highest η range. To increase the fraction of neutrinos originating from charmed hadrons the detector is placed slightly off-axis with respect to the collision axis.

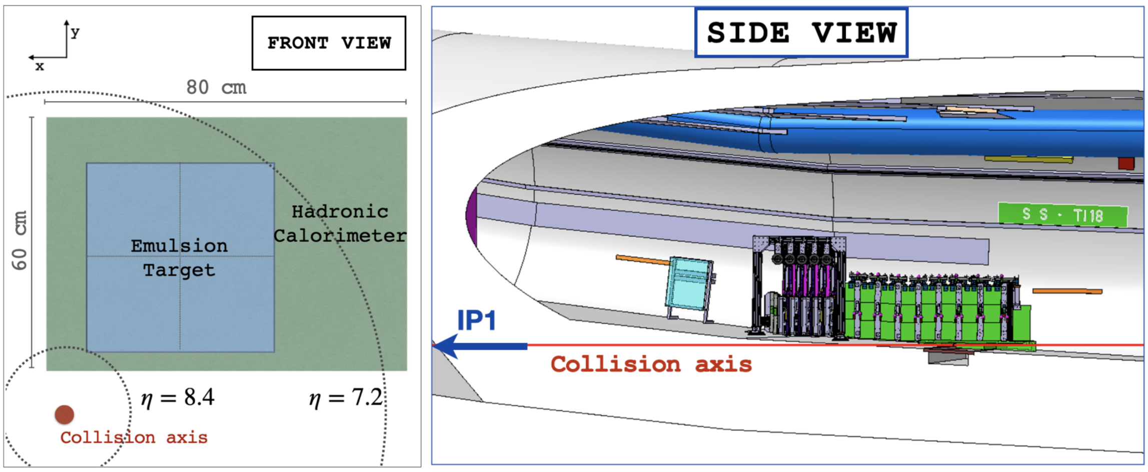

Figure 1: The experimental range of η acceptance can be seen from the front view in (a) [2] and shown from the side view is the relative placement of the detector within the TI18 tunnel in (b) [1].

Prior to the SND experiment, a long-standing interest in the LHC was to make the first direct collider neutrino observations [3]. The major experiments such as ATLAS and CMS consist of cylindrical detectors that study high p_T physics at $\eta < 5$ and are therefore blind to the far-forward region of neutrino flux. The remaining LHC experiments are otherwise prohibited from making direct neutrino observations due to having large, uncontrollable backgrounds and lacking a suitable target to increase the probability of neutrino interactions. Therefore, any small-scale experiment with the correct acceptance coverage and target region has a large physics potential in observing the large flux of high-energy neutrinos in the forward region [4]. The SND experiment was designed to fill this gap of detection capability and use the neutrino flux to probe heavy-flavour production in a previously unexplored range of energy and η [1]. In its location, the experiment is separated from IP1 collisions by 100m of rock and concrete which shields the detector from all hadrons and most charged leptons that are produced. In addition, LHC magnets along the collision axis deflect many of the charged leptons from IP1 that would otherwise hit the detector. These factors significantly reduce the neutrino background.

The SND experiment will take data throughout the whole of Run 3 and is expected to accumulate a total of 250 fb^{-1} by the end, corresponding to approximately 2000 high-energy neutrino interactions [1].

1.2 Physics goals

Neutrino measurements can allow us to make a precise test of the SM and hence can be used as a powerful probe for new physics [1]. Currently, the neutrino-nucleon cross-section in the region between 350 GeV and 10 TeV remains unexplored. With its novel capacity, the SND experiment looks to study this energy range corresponding to the highest energy collider-produced neutrinos that are below the energy of cosmic neutrinos. The exact contribution made by SND in the neutrino energy cross-section measurements can be seen in Figure 2.

In the η range which SND uniquely explores, electron neutrinos and anti-electron neutrinos are predominantly produced in the decays of charmed hadrons [1]. Therefore these electron neutrinos can be used to probe the production of charm. The SND experiment can probe an average momentum fraction (x) of interacting gluons as low as 10^{-6}. The extraction of the gluon PDF at these low values is currently unknown and could provide new insights for future experiments that will probe this x range, such as the FCC. Furthermore, using SND's efficient identification of all three neutrino flavours, lepton flavour universality can be tested in the neutrino sector. This can be achieved by measuring the ratio of $\frac{\nu_e}{\nu_\tau}$ and $\frac{\nu_e}{\nu_\mu}$.

Through its design, the SND detector also has sensitivity to extend the parameter range which can be accessed to search for Feebly Interacting Particles (FIPs) that can be detected via scattering off nuclei and electrons in the detector target [1].

1.3 SND detector

The design of the SND detector was focused on allowing the identification of the three neutrino flavours and the direct search for FIPs [1]. A representation of the SND detector without the neutron shield can be seen in Figure 3.

The detector is comprised of two main sections corresponding to a target region followed downstream by a hadronic calorimeter and muon identification system [1]. In front of the target region are placed two planes of scintillator bars that act as a veto for charged particles, with the flux largely made up of muons from IP1. The target region, weighing about 830 kg, consists of five walls of Emulsion Cloud Chambers (ECC), each wall followed by a Scintillating Fibre (SciFi) plane.

The emulsion walls are comprised of layers of alternating tungsten and emulsion plates [1]. The tungsten plates act as the passive material, which is the neutrino target, while the emulsion plates act as a tracking medium with micrometre accuracy. Tungsten is chosen as the passive material to maximise the density of the target and hence increase the probability of neutrino interactions.

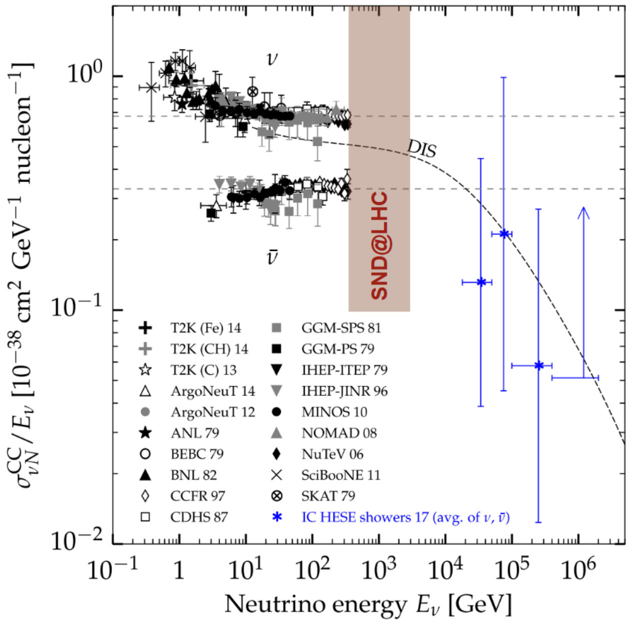

Figure 2: Current charged-current inclusive neutrino-nucleon cross-section measurements [5]. The thick dashed curve represents a standard prediction of deep inelastic scattering (DIS), averaged between neutrinos and anti-neutrinos. The shaded region represents the neutrino energy range that is being studied by SND@LHC.

The SciFi planes provide a timestamp and spatial information about the hits detected in the scintillating fibres [1]. Fibres in the SciFi planes are arranged in alternating layers of vertical and horizontal orientation. SciFi hits can then be combined across multiple planes and be used to produce the shower profile to reconstruct energy and vertex position.

The emulsion walls and SciFi planes within the target region form a coarse sampling EM calorimeter with a total of 85 radiation lengths, X_0 [1] [1]. This allows to effectively contain the full EM showers within the target. The veto and target region are both contained within a neutron shield.

The hadronic calorimeter and muon system are located downstream of the target and are comprised of eight 20 cm-thick iron slabs [1]. These slabs correspond to 9.5 interaction lengths, λ [2], in total. Each iron slab is followed by either one or two planes of scintillating bars. In this configuration, hadronic showers begin developing within the target region, which corresponds to 1.5 radiation lengths, giving 11 radiation lengths in total for the whole detector. This provides effective containment of hadronic showers within the detector volume.

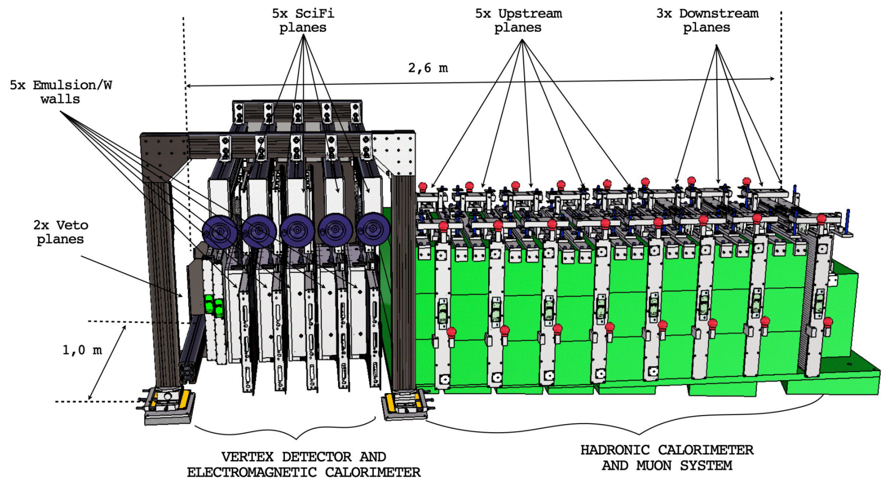

Figure 3: The layout of the SND@LHC detector [1].

[1] X_0 is defined as the mean length (in cm) into the material at which the energy of an electron is reduced by the factor 1/e ($\sim$ 0.368) [6]

[2] λ is defined as the mean distance travelled by a hadronic particle before undergoing an inelastic nuclear interaction [7]

2 Methodology

The goal of this project was to develop the first procedure for hadronic shower energy and vertex position[3] reconstruction for the SND detector. The most effective shower reconstruction would involve combining information from both the electromagnetic and the hadronic components of the detector.

For this analysis, the shower reconstruction is initially demonstrated using only the target (electromagnetic) region of the detector to serve as a proof of concept prior to adding the hadronic information of the detector to the reconstruction procedure. Since at the time of this project the tracks from the emulsion walls have not yet been processed and digitised, only the SciFi data from the target region is used for the reconstruction.

2.1 Simulating pion interactions within the SND detector

GEANT4 [8] was used to simulate charged pions being fired into the SND detector to obtain data that is representative of the hadronic showers initiated by neutrino interactions. The detector geometry is implemented using the TGeo package of ROOT. The GEANT4 simulation stops with the deposition of energy in the sensitive detectors. A digitisation step takes this energy and simulates an electronic signal.

To generate the simulations, charged pions were fired perpendicularly to the plane of the target and in the centre. In the GEANT4 geometry, the centre of the walls corresponds to -28 cm and 35.5 cm in the x and y planes respectively. To effectively capture the possible neutrino interactions with the detector, pions were fired with a range of energies and depths into the detector.

The energies used were: E = {50,100,200,300,400,500,750,1000,1250,1500,2000} GeV

Whilst the depths used were: Z = {291,295,304,308,317,321,330,334,343,347} cm

The set of energy points was chosen to realistically represent the range of the shower energy and the frequency at each energy point. Lower energies have smaller spacing due to being more frequent. Ten depths are used in total, where two depths are used for pions fired in each emulsion wall. The first depth corresponds to the beginning while the second corresponds to the middle of each emulsion wall respectively. In this case, the beginning of the first wall lies at 291 cm and the middle of the fifth wall lies at 347 cm within the GEANT4 geometry.

A simulated event display of a charged pion showering from the first SciFi station can

[3]Vertex position and shower initiation depth will be used interchangeably.

be seen in Figure 4. The pion shower extends into the hadronic calorimeter where it fully saturates the scintillating bars in the first two hadronic scintillating planes.

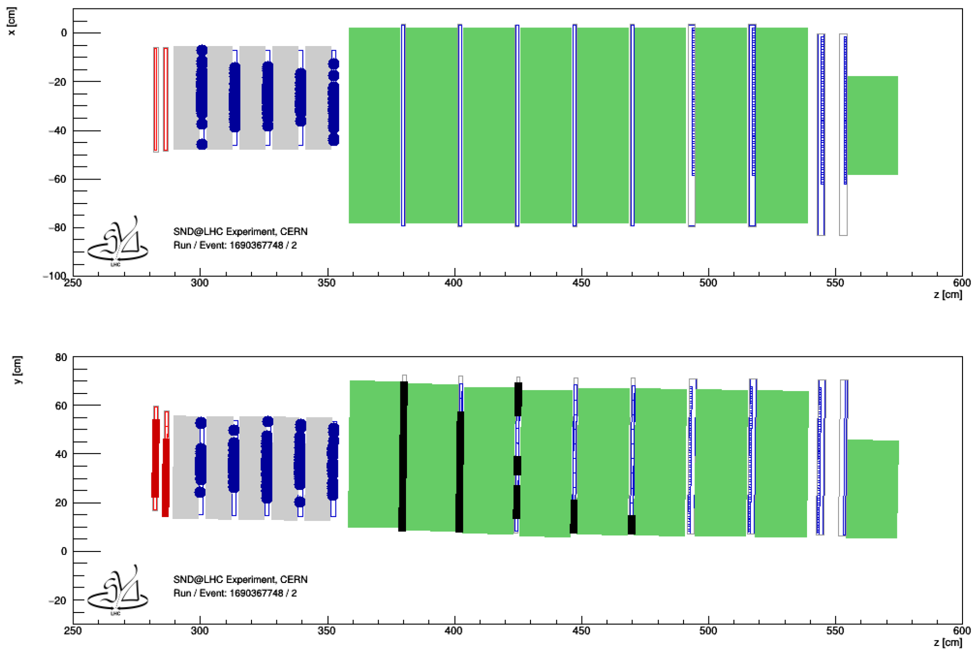

Figure 4: A simulated event display of a charged pion shower that is initiated from the first SciFi station. The upper plot shows the top view while the lower plot shows the side view of the detector.

The GEANT4 simulations were preliminarily run on the SND collaboration software framework called SNDSW [9]. The final GEANT4 simulations and the later processing stages were executed through HTCondor which is a software for high-throughput computing run over the CERN lxplus computing system.

2.2 Profiling showers using PDFs

To characterise the shower shapes for different energies of pions and different depths of shower initiation, a probability density function (PDF) can be created to represent the average activity within each of the SciFi stations for each configuration.

First, the total number of hits per unit length in a given station are obtained for a given energy and depth configuration. The total hits are scaled by the number of pions that were fired to obtain the average number of hits per event per unit length in each of the SciFi stations. The scaled hits are then divided by the number of silicon photomultipliers (SiPMs) in each of those (position) bins. The resulting distribution then corresponds to the SciFi hit probability as a function of the SiPM position, corresponding to a PDF. A different PDF is generated for each station and configuration.

A schematic representation of the PDF generation process can be seen in Figure 5. The PDF procedure is demonstrated for the first SciFi station in the x plane for the configuration of energy 300 GeV and depth 291 cm.

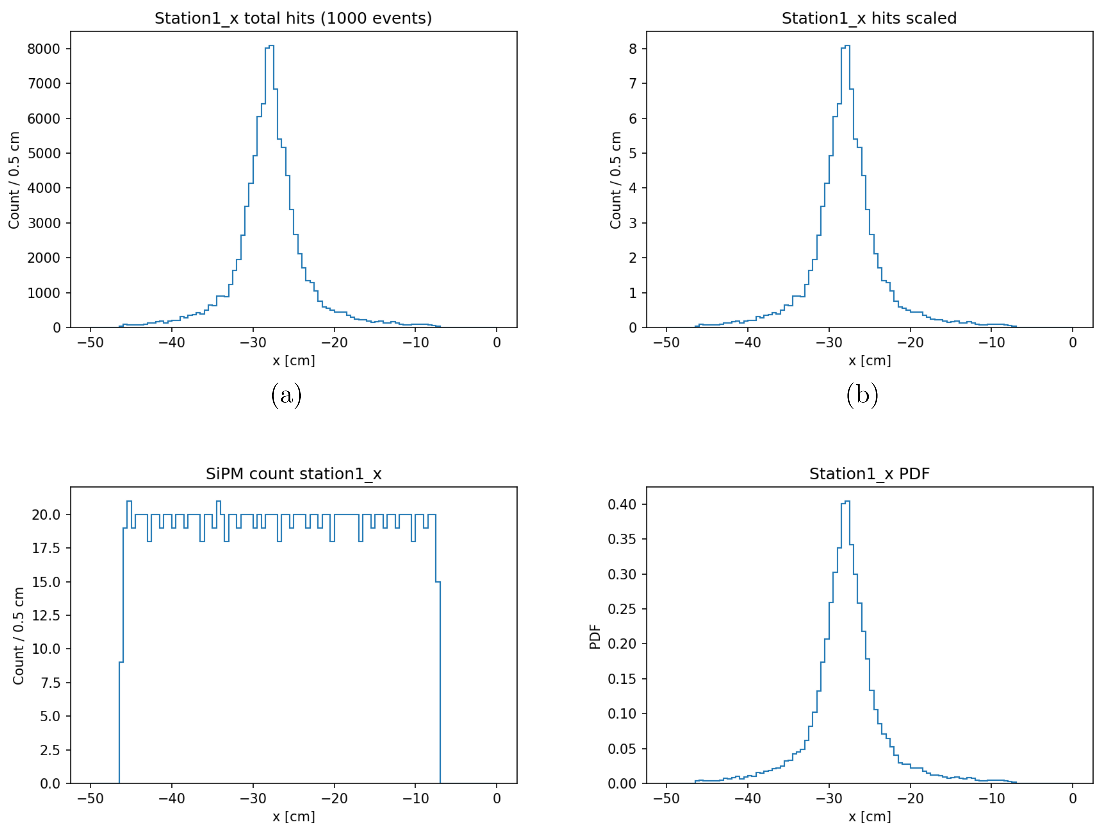

Figure 5: The components of the PDF generation process for pions with energy 300 GeV, fired at a depth of 291cm. The total hit distribution from 1000 events (a) is scaled down to the average hits per event (b). The scaled distribution is divided by the SiPM bin counts in (c) to obtain the PDF distribution in (d).

2.3 Building a log-likelihood function for reconstruction

Once the showers have been profiled with the PDFs, individual pion shower hit patterns can be compared with the PDFs to determine how likely it is that the shower belongs to each configuration. This procedure is achieved by building a log-likelihood function to scan over the different configurations to find the likeliest fit. The log-likelihood formula is reported in Equation 1:

$$\log L = \sum_{i}^{Nbins} (N_{SiPMs_i} - N_{hits_i}) \log (1 - p_{hit_i} + \epsilon) + N_{hits_i} \log (p_{hit_i} + \epsilon) \tag{1}$$

Where the $SiPMs$ are the SciFi bin counts (max possible hits), N_{hits} are the number of recorded SciFi hit bin counts, p_{hit} is the PDF bin counts for a given energy and depth and ϵ is a small constant value added to all PDF bin counts to make sure there are no empty PDF bins when taking the log.

The two terms in the log-likelihood formula allow to make the function symmetric. For an individual SciFi station, the function is at a maximum where there is an identical match between the recorded hit bin counts for a given event and the PDF profile for the given energy and depth.

3 Results and discussion

3.1 Diagnostic plots of the simulations

To confirm that the simulations of the hadronic showers were successful we can look at the diagnostic plots of the event hits. As an example of some simulated events, we can look at pions fired at 291 cm (beginning of the target) and with an energy of 100 GeV. Figure 6 shows the 2D, x and y plane, distribution of the shower profile for all the hits in 1000 events.

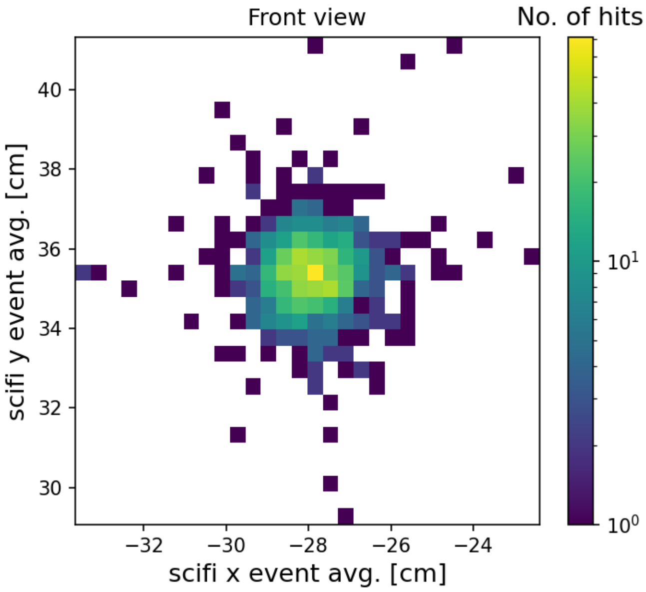

Figure 6: The 2D (x and y plane) distribution of the shower profile for all the hits recorded in 1000 events. Pions of 100 GeV energy and 291cm showering depth are used.

The x-plane hits correspond only to vertical fibre hits, while the y-plane hits correspond only to horizontal fibre hits. In this case, the x and y positions of the hits for each event are averaged to obtain a single point for each event. This plot shows the combined hits in all the five planes overlayed into a single plot. The distributions in both the x and y planes are Gaussian and have the same shape. The centre of the distribution is on the firing coordinates and the distribution away from the centre is due to the shower spreading out in width for the later stations as it develops.

To obtain a view of the shower development with each of the SciFi stations we can look at the top and side view representations of the SciFi hits. Figure 7 shows both the side (horizontal hits) and top (vertical hits) views.

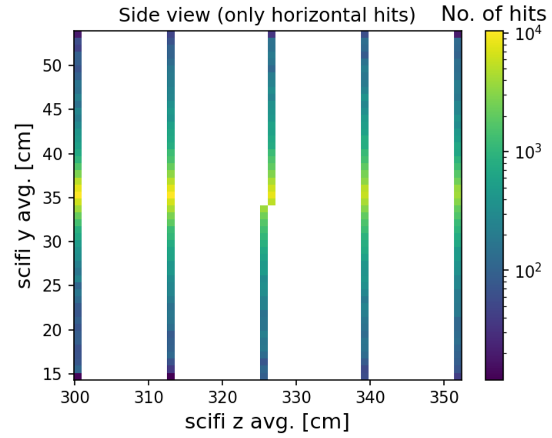

(a)

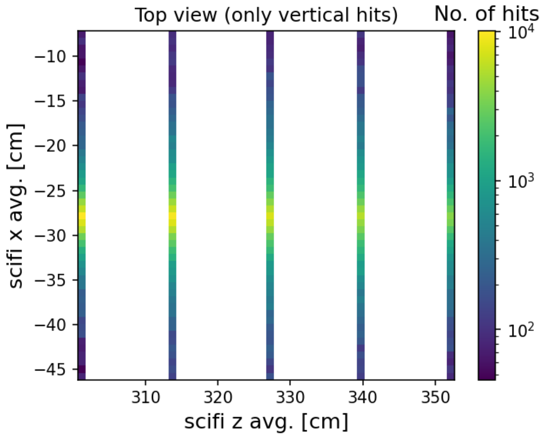

(b)

Figure 7: The side and top view distributions of the shower profile for all hits recorded in 1000 events of pions with energy 100 GeV and showering depth 291cm.

In these plots the shower begins from the first station, where it is very narrow and concentrated at the firing coordinates. For each successive station, the shower spreads out and loses energy, which can be seen through the more diluted distributions further downstream. Both the side and top views show the same shower development, as we expect, since the shower is on average symmetrical about the central firing axis.

3.2 Generating the PDFs

Given that there are 11 energies, 10 depths, 5 x planes and 5 y planes, a total of $(11 \times 10) \times (5 + 5) = 1100$ PDFs need to be generated to fully profile the showers in all configurations. 1000 events were used for each configuration of energy and depth to generate the individual PDFs. To generate the PDFs 100 bins were used for the histograms with a 0.5cm bin width. The GEANT4 geometrical ranges used were between 0cm and -50cm and 10cm and 60cm for the x and y planes respectively.

A selection of the PDFs at different energies and at a fixed depth of 291 cm can be seen in Figure 8 for stations 1, 2 and 3.

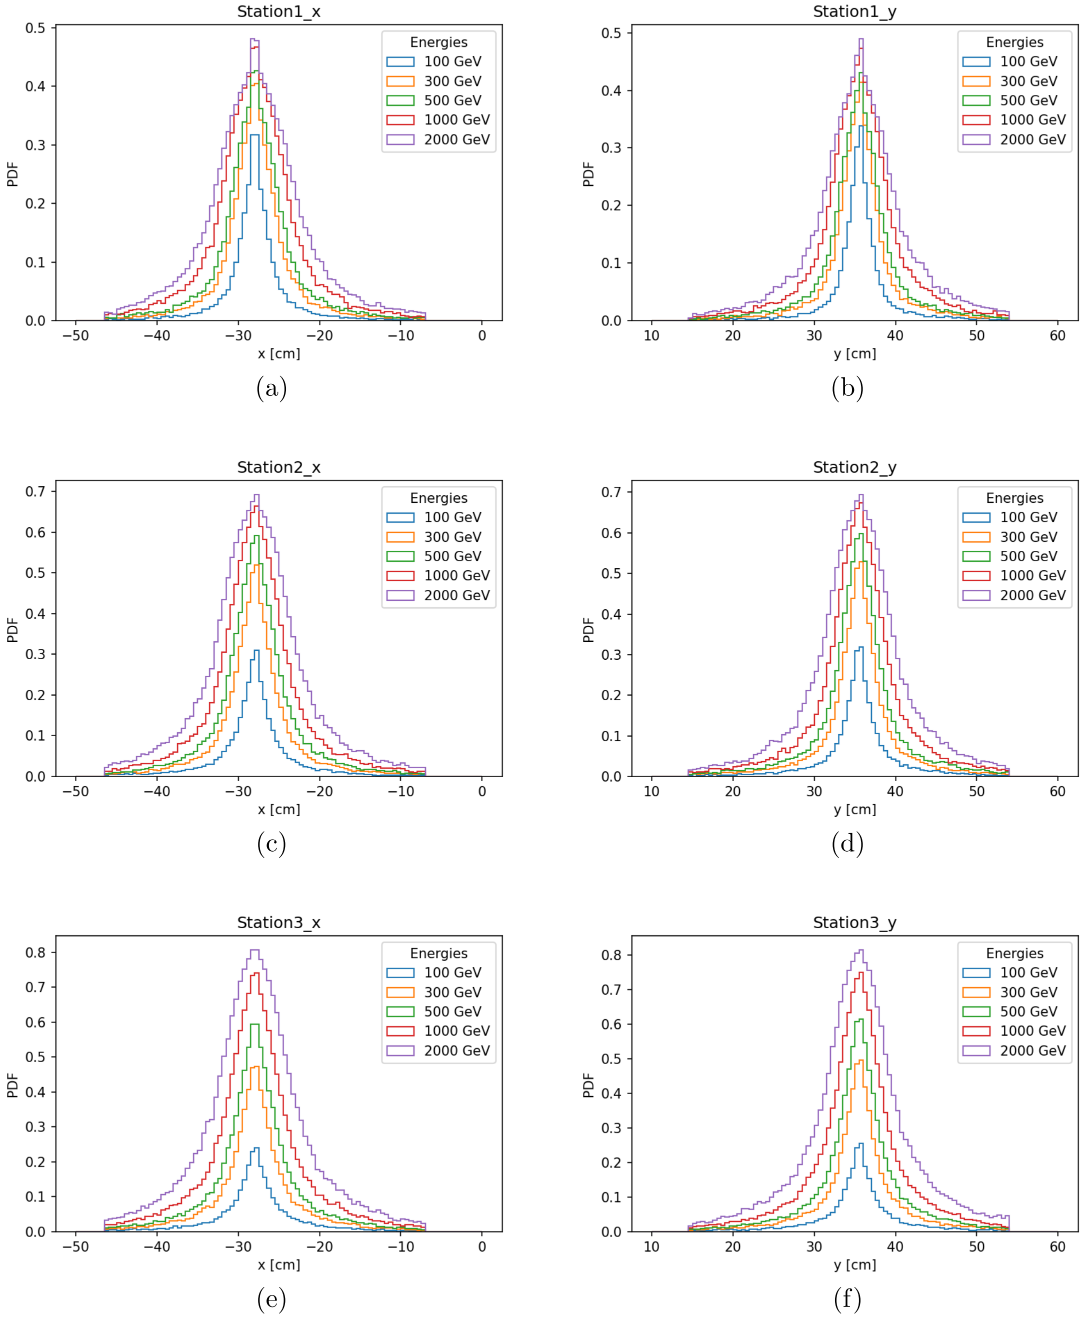

Figure 8: The PDFs of stations 1, 2 and 3 at energies of 100, 300, 500, 1000, and 2000 GeV at a fixed depth of 291cm.

The PDFs all show a consistent trend where the activity in the SciFi stations is correlated with the energy of the pion which initiated the shower. Pions with higher energy leave higher activity in the SciFi station compared to lower energy pions. The activity is highest for the central region where the core of the showering particles are contained. For the lower energies, the shower hits peak in the initiating station and the activity further downstream only gets weaker. On the other hand, for higher-energy showers, the average shower activity is relatively low in the initiating station and increases consecutively in stations 2 and 3. Therefore, the peak of higher-energy showers happens further downstream from the station of initiation.

In addition, a selection of PDFs at different depths for a fixed energy of 300 GeV can be seen in Figure 9 for stations 1, 2 and 3.

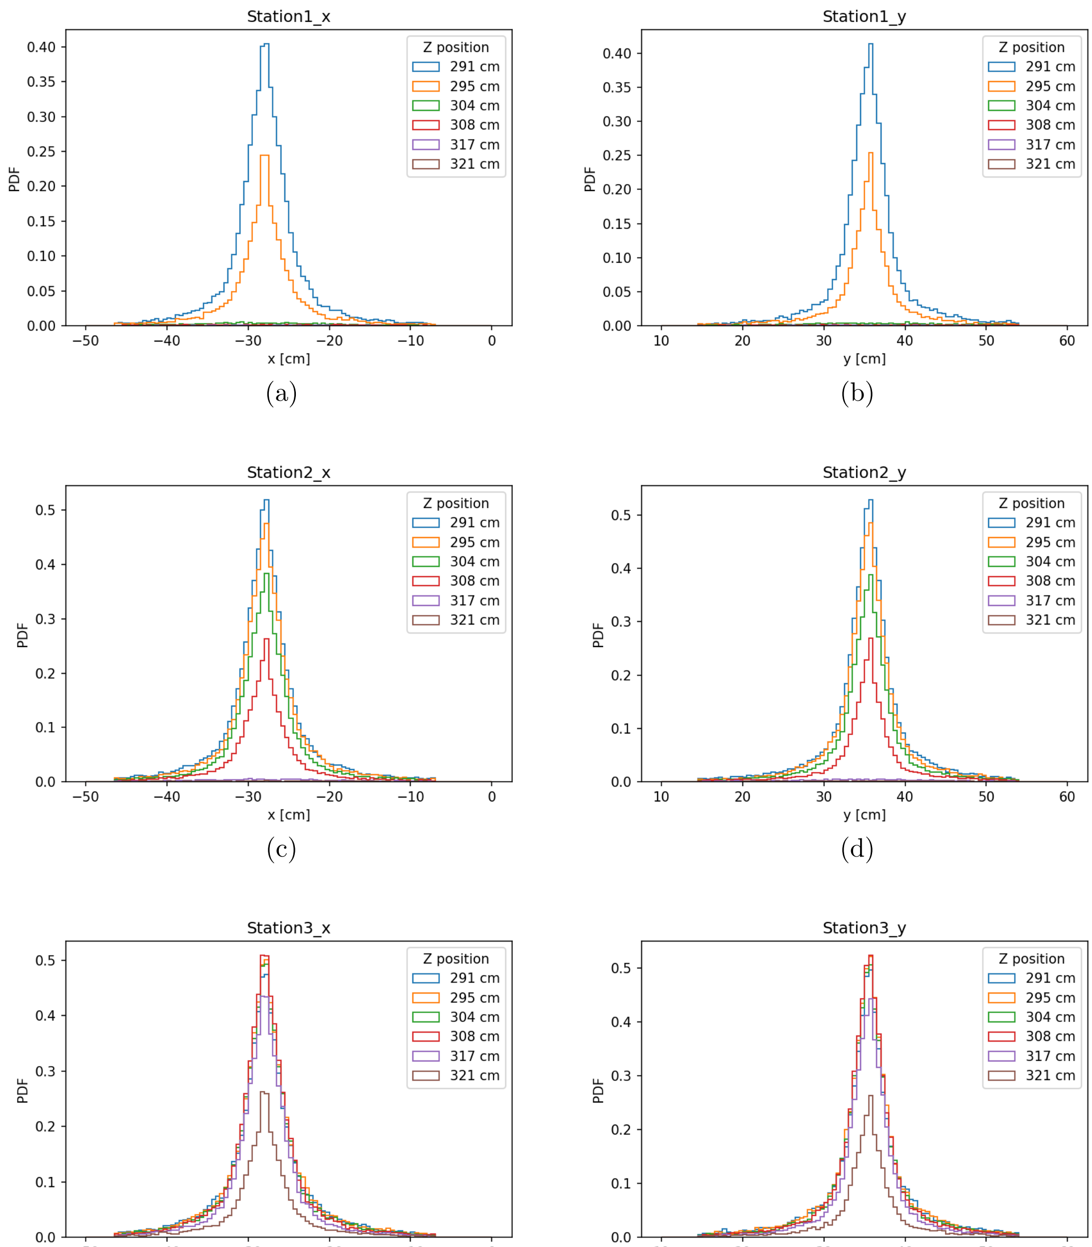

Figure 9: The PDFs of stations 1, 2 and 3 at the depths of the first three target walls and fixed energy of 300 GeV

The PDFs are shown for showers initiated at the beginning and middle of the first three target walls. In these PDFs, it can be seen that showers which are initiated deeper into the target walls leave close to no activity in the previous stations. Any activity which is recorded in the previous stations, prior to the shower initiation, is residual activity of the shower which has travelled backwards and has interacted with the SciFi station. For the energy of 300 GeV it can be seen that the activity between the initiating station and the consecutive station is approximately the same. The shower profile is relatively uniform as the shower propagates downstream.

3.3 Performing a log-likelihood scan

To reconstruct a given energy or vertex position for each set of hit counts, a scan has to be carried out over all the PDFs. The scan is performed over each SciFi station separately and the log-likelihood values for all five stations are summed for each energy or depth.

During the scanning procedure, a constant ϵ value of 5×10^{-5} is added to each PDF bin to ensure that there are no empty PDF bins when the log is taken. The value of 5×10^{-5} was chosen as it corresponds to the lowest non-zero activation in the PDF bins that is possible. This is equivalent to 1 hit for all 1000 events in all 20 SciFi fibres per bin: $\frac{1}{1000 \times 20} = 5 \times 10^{-5}$. Smaller values of the correction term ϵ were tested but the reconstruction results did not improve.

The reconstruction is currently only performed using the x-plane PDFs and hit counts. In sections 3.3.1 and 3.3.2, separate examples of the procedures are shown for energy and vertex reconstruction.

3.3.1 Energy reconstruction procedure

For example, we can look at the successful energy reconstruction of an event with 300 GeV energy and a depth of 291 cm. In Figure 10, the hit distributions[4] from a single event can be seen for the 5 SciFi stations.

[4]Hit distributions in all plots are shown with half the number of bins that were used for the analysis to improve clarity. The maximum possible hits per bin in this binning configuration is 40.

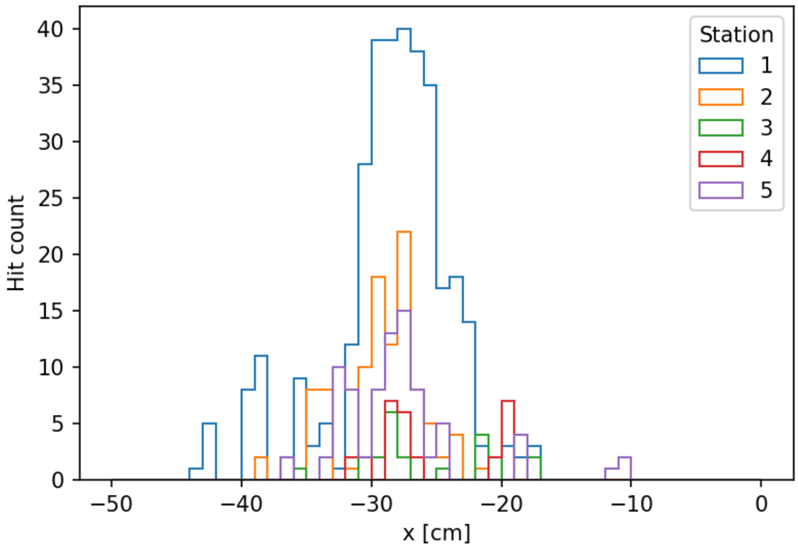

Figure 10: The SciFi hit distributions from a single pion shower of energy 300 GeV and depth of 291 cm, for all the 5 SciFi stations. This set of hit distributions allows for accurate energy reconstruction.

The hit distributions show that the pion shower fully develops directly from the initiating depth, where the fibres in the first station become fully activated with a maximum of 20 hits possible per bin. The next highest activity left by the shower is in stations 2 and 5, whereas the shower leaves minimal activity in stations 3 and 4.

This shows that a shower's activity throughout the SciFi stations is not constant and continuous, the shower can die down, like in stations 3 and 4, and become active again further downstream, like in station 5. This is in contradiction to the trend that was observed from the PDF plots which showed that the average activity of the shower for 300 GeV to be constant through at least the first three SciFi stations. This indicates that the behaviour of each shower is quite different and that shower development is not consistent between individual events.

The log-likelihood scan results based on this set of hit distributions can be seen below, separately for each SciFi station in Figure 11a and the combined results between all stations in Figure 11b.

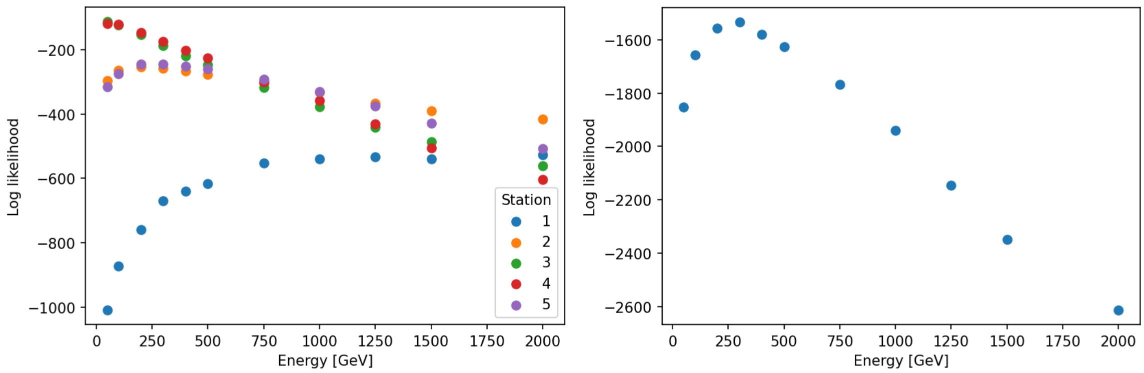

(a) (b)

Figure 11: The results of a successful log-likelihood scan in energy shown separately for each SciFi station in (a) and summed between all the stations in (b).

For station 1 the maximal activity in the central region looks closest to the PDF of the highest energy pions showered and this is reflected in the log-likelihood curve where the highest energies are at the maximum of the distribution. On the other extreme, the minimal activity in stations 3 and 4 looks the most like the PDFs of the lowest energy pions that were showered and the maximum of the log-likelihood is located at the lowest energies. For stations 2 and 5 the activity that is left is closer to the average activity of 300 GeV pions and hence the log-likelihood is maximal in the region of energy around 300 GeV.

The strong correlation observed between the activity in the stations and the shape of the log-likelihood curves confirms that the scanning procedure was well implemented and works as expected.

In the summed results, the above-average activity in station 1 is cancelled out by the below-average activity in stations 3 and 4 to produce a peak in the summed log-likelihood values at 300 GeV, which is the correct reconstruction energy.

3.3.2 Vertex position reconstruction procedure

The equivalent approach is shown for the vertex reconstruction of an event with 300 GeV energy and 304 cm depth.

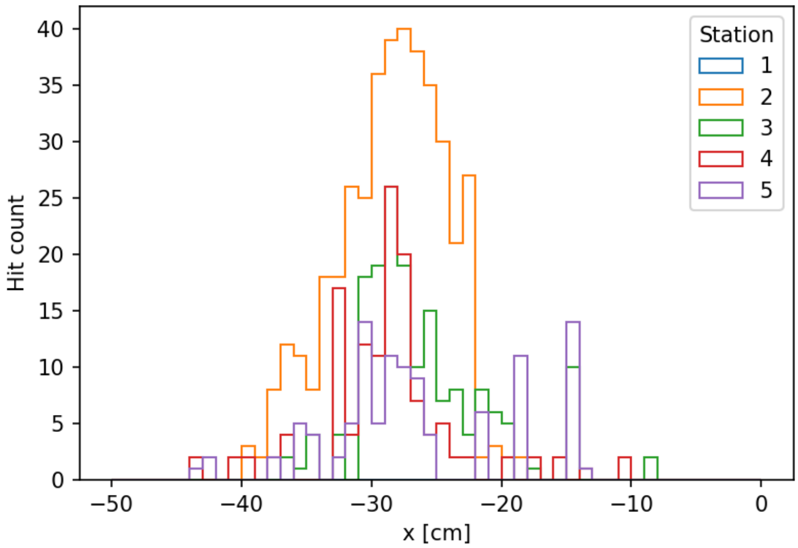

Figure 12: The SciFi hit distributions from a single pion shower of energy 300 GeV and depth of 304 cm, for all the 5 SciFi stations.

In the hit distribution in Figure 12, there are no hits in the first station as the shower is only initiated from the second station and there have not been any backscattered shower particles that originated from further downstream. Station 2 has maximal activation in the central region as the shower develops from its initiation depth. Stations 3, 4 and 5 show consecutively decreasing moderate activity.

These hit distributions can be compared directly with those shown in the energy reconstruction section (section 3.3.1) but shifted down one station and with the possibility of only populating station 1 with backscattered particles. The log-likelihood scan results can be seen in Figure 13.

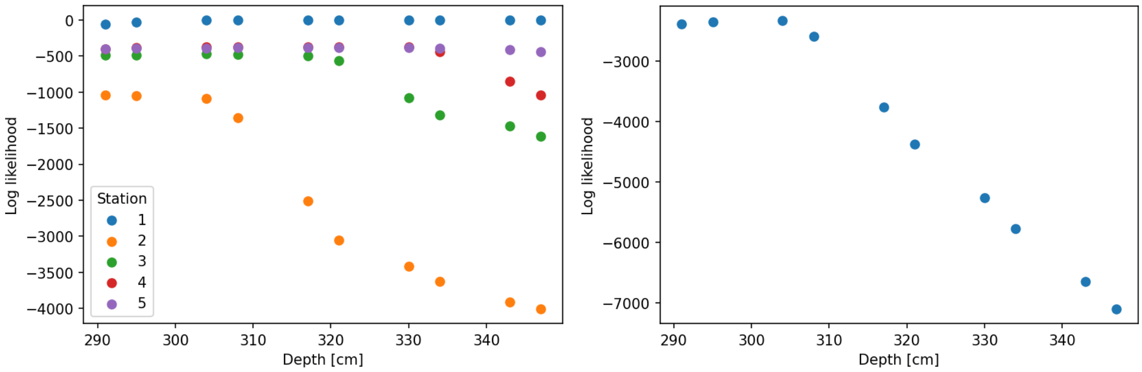

(a)

(b)

Figure 13: The results of a successful log-likelihood scan in depth shown separately for each SciFi station in (a) and summed between all the stations in (b).

For station 1, since there are no hits in this station only the first term of the log-likelihood formula contributes:

$$\log L = \sum_{i}^{Nbins} N_{SiPMs_i} \log\left(1 - p_{hit_i} + \epsilon\right) \tag{2}$$

PDFs for depths 291cm and 295cm have active shower development so they contribute the most. PDFs for 304cm, and 308cm and onwards have only backscattered particles and hence close to empty PDF bins. So for these PDFs we get i.e. $ln(1-x) < -1 \times 10^{-4}$ for $x < 1 \times 10^{-4}$. Therefore the log-likelihood for further depths is even closer to zero.

At station 2 all shower depths have close to equal PDF distributions for showering starting before or at station 2. Station 2 PDFs have only backscattered activity for showering starting at a depth after station 2. Therefore, the log-likelihood for depths after 308cm becomes significantly more negative. The same happens for stations 3, 4 and 5 relative to the depths of those stations. The individual distributions of the hits only have an influence on the plateau of each log-likelihood curve.

In Figure 13b we can see that when all the log-likelihood values are added up between the stations the maximum occurs at 304 cm, which is the correct reconstructed depth.

When comparing the combined log-likelihood distributions for energy and depth it can be seen that the energy distribution has a much more defined peak compared to the depth distribution. Therefore, for these examples, the reconstruction of the energy is much more reliable as there is a bigger margin between neighbouring reconstruction values.

3.4 Energy reconstruction in the first three SciFi stations

The log-likelihood scanning procedure can be applied to all energies at a given depth to obtain the full reconstruction performance. Figure 14 shows the energy reconstruction plots for the first three SciFi stations corresponding to the depths of 291cm, 304cm and 317cm.

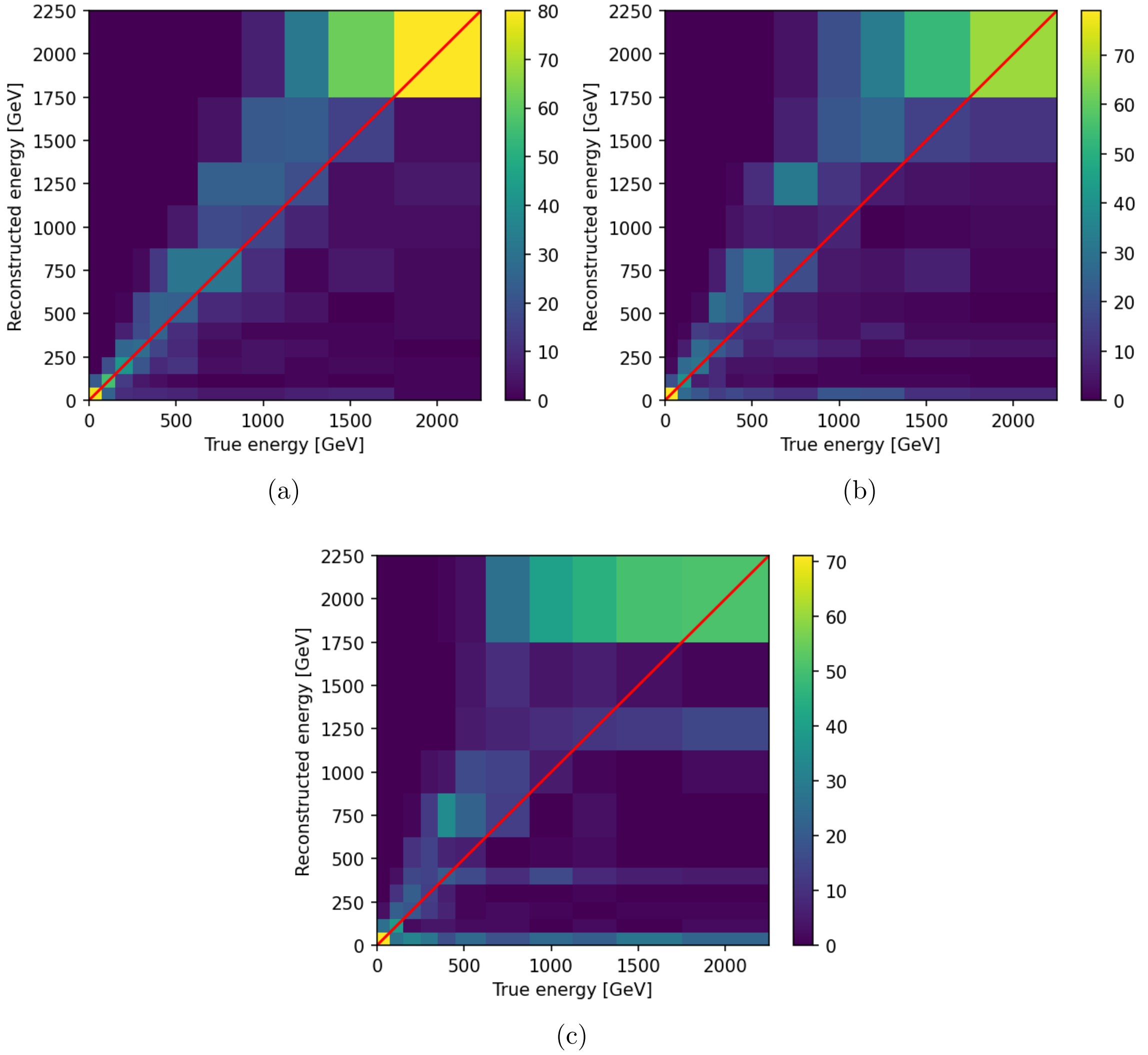

Figure 14: The energy reconstruction plots for the first three stations at depths (a) 291cm, (b) 304cm, and (c) 317cm. The true energy corresponds to the pion firing energy used in GEANT4.

These figures are generated by reconstructing the energy for individual events and then comparing the reconstructed energy with the true energy. For each true energy, 100 events are tested. In these figures, it is desired for all events to lie on the diagonal which would correspond to perfect reconstruction.

For the first station, the energy reconstruction is carried out quite accurately where it can be seen that events lie on the diagonal or very close to it. Lower energies get closer predictions to the diagonal but for higher energies, the predictions begin to overestimate the energy, where events lie above the diagonal.

For stations 2 and 3 the same trends can be seen as station 1 but with increasingly more energies being overestimated. Furthermore, there is a gradual increase in the events that are reconstructed at either the highest or the lowest energies, corresponding to 2000 GeV and 50 GeV.

This trend is likely due to having less information for later initiating showers as less of the shower is contained in the detector and also due to there being more fluctuations in the shower development for the shorter shower distance considered.

3.5 Vertex position reconstruction for 100 GeV, 300 GeV and 1000 GeV

The vertex reconstruction procedure can likewise be applied to all of the depths at a given energy. Figure 15 shows the depth reconstruction plots for the energies of 100 GeV, 300 GeV and 1000 GeV. The reconstruction is again performed for each event individually and 100 events are used to test each true depth value.

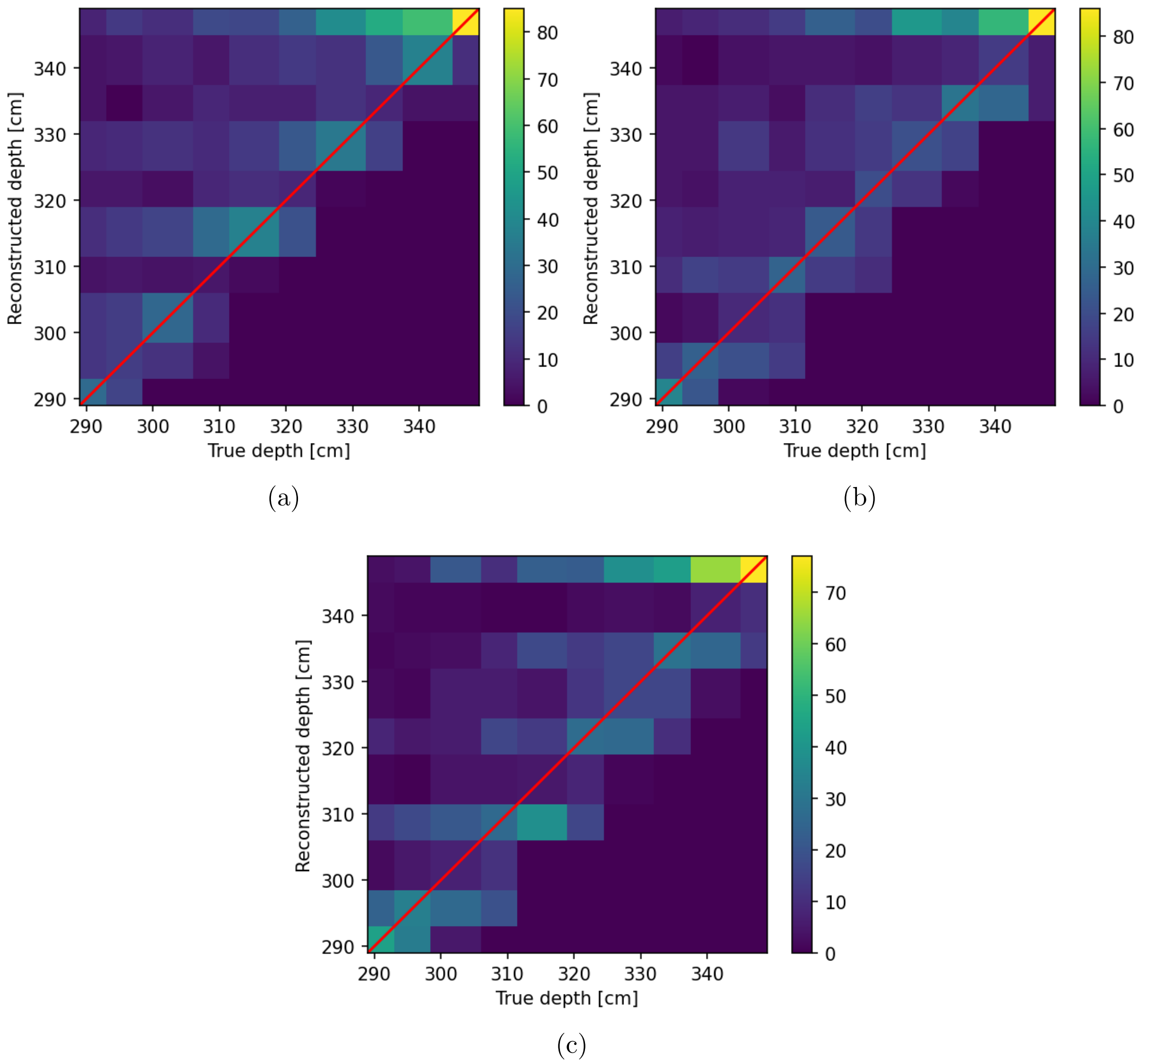

Figure 15: The depth reconstruction plots for the first at the energies of (a) 100 GeV, (b) 300 GeV and (c) 1000 GeV. The true depth corresponds to the pion firing depth used in GEANT4.

For the energy of 100 GeV a pattern is formed in the reconstruction performance. Depths in the middle of each station have enhanced reconstruction accuracy. Depths at the beginning of each station are reconstructed as the middle of either the current or the previous stations.

In the plot for 300 GeV the reconstruction performance pattern is more smoothed out and does not favour the middle station. However, in the reconstruction for 1000 GeV the opposite is observed, where the middle of each station is often poorly reconstructed but the beginning of the stations is relatively well reconstructed.

For all three sets of energies, there is a bias to overestimating the depth reconstruction, where the last depth is disproportionately favoured. In general, for 100 GeV and 300 GeV few events are underestimated, but for 1000 GeV the rate of underestimation significantly increases.

3.6 Summary reconstruction performance plots

The overall performance of energy reconstruction is summarised below. In Figure 16a is the bias for a selection of energies and in Figure 16b is the resolution for all of the energies.

The bias is given by:

$$R = \langle \frac{E_{reco} - E_{true}}{E_{true}} \rangle, \tag{3}$$

whereas the resolution is given by:

$$R = std(\frac{E_{reco} - E_{true}}{E_{true}}), \tag{4}$$

where std stands for taking the standard deviation.

In this case the smaller the bias or resolution, the better.

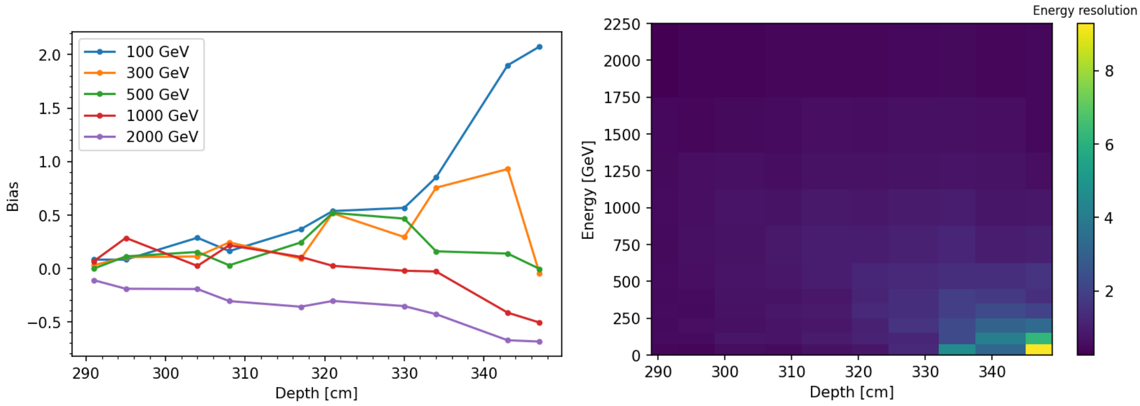

(a) (b)

Figure 16: The (a) bias and (b) resolution plots for energy reconstruction.

In Figure 16a it is apparent that the bias increases with further depth into the SciFi stations. This can be expected since showers which are initiated further downstream have less of the shower development contained within the target and there is less information to base the reconstruction on. Furthermore, activity from a smaller section of the shower is also likely to fluctuate more which biases the results of the log-likelihood scan.

For the lower energies there is almost always a positive bias whereas for the higher energies the bias gradually becomes negative. This supports the trends observed in section 3.4. The bias is the largest for the lowest energy range.

In Figure 16b a similar trend can be seen. The energy resolution for all energies at low depths is very low, meaning the reconstruction is more consistent for early shower initiation. A trend of worse resolution can be seen along the anti-diagonal section, with the resolution being the worst at the lowest energies and highest depths. This indicates that as energies decrease and the pion firing depth increases, the reconstruction fluctuates the most.

For the depth reconstruction summary plots, the absolute bias and resolution are used instead. The absolute bias is given by:

$$R = \langle E_{reco} - E_{true} \rangle, \tag{5}$$

whereas the absolute resolution is given by:

$$R = std(E_{reco} - E_{true}), \tag{6}$$

where std stands for taking the standard deviation.

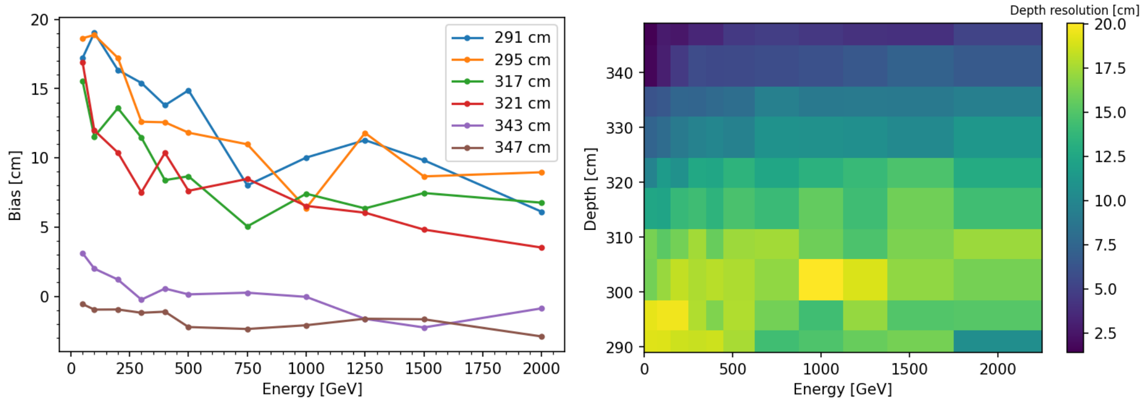

(a) (b)

Figure 17: The (a) bias and (b) resolution plots for depth reconstruction.

In Figure 17a there is a clear separation between the bias in the three different stations. Station one shows the highest reconstruction bias where the depth is significantly overestimated. The reconstruction consistently predicts that the pion shower is initiated on average one and a half stations downstream from the true depth value. This trend improves with the energy of the shower where the bias decreases for higher energies.

For station 3 a very similar trend is observed. The bias is lower throughout the energy range compared to station 1 likely as there are fewer possible stations further downstream that can be predicted.

In the last station, the bias measured is the lowest but this result is unreliable since there is only one station that can be favoured in the reconstruction. All the previous stations have only backscattered hits whereas the last station is the only one with direct shower activity.

Figure 17b shows the resolution of the depth reconstruction. The results show the highest level of fluctuation in the reconstruction to be centred around the lower depth and energy quadrant. As the depth increases the resolution of the reconstruction improves. The resolution is the best for showers starting at the end of the target.

3.7 Diagnosing failed reconstruction attempts

3.7.1 Energy reconstruction example

Below we can look at another instance where this time the energy reconstruction fails. In Figures 18 and 19 below the hit distributions in the SciFi stations can be seen for this event, including the results of the log-likelihood scan.

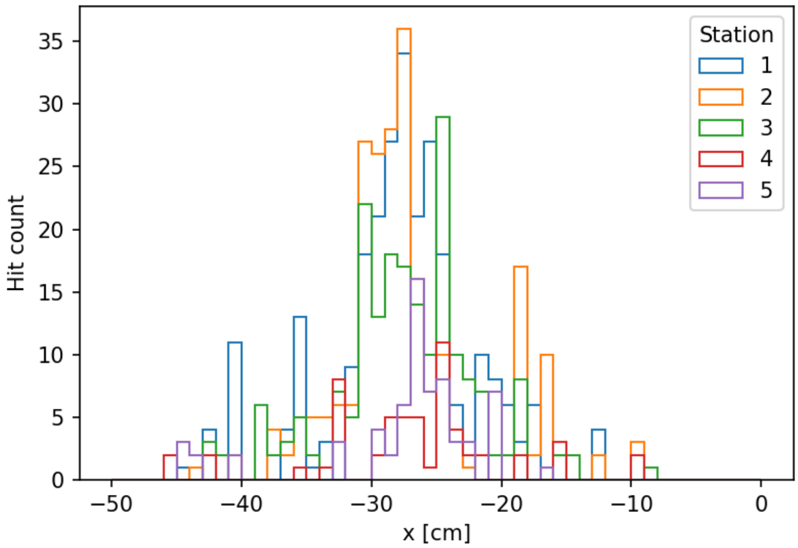

Figure 18: The SciFi hit distributions from a single pion shower of energy 300 GeV and depth of 291 cm, for all the 5 SciFi stations. For this set of hit distributions energy reconstruction fails.

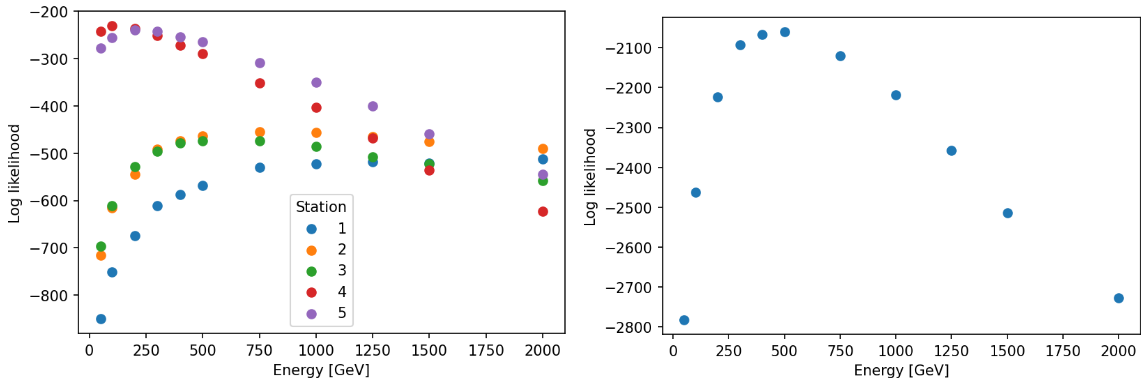

(a) (b)

Figure 19: The results of a failed log-likelihood scan in energy shown separately for each SciFi station in (a) and summed between all the stations in (b).

In these hit counts the shower activity is much more extended through the SciFi stations and the shower has a significantly different behaviour compared to the first example in Figure 10, section 3.3.1. The first two stations contain a high level of activity with the activity of station 3 being moderately high and less localised. Stations 4 and 5 have a consecutively lower level of minimal activity. This activity is reflected in the log-likelihood curves for the separate stations. The first three stations have a log-likelihood maximum at the high energy range whereas for the last two stations the maximum lies at the lowest energies. Since more of the stations were skewed towards a higher energy reconstruction the overall summed log-likelihood does not balance out and the energy is overestimated. In this case, an incorrect reconstruction of 500 GeV is made.

Therefore the reconstruction in this case fails when the hit distributions become too different from the PDF distribution of the activity for the given energy. This is the case when shower propagation is inconsistent between separate events.

3.7.2 Vertex position reconstruction example

In Figures 20 and 21 we can look at an example where the vertex reconstruction does not give the truth depth value.

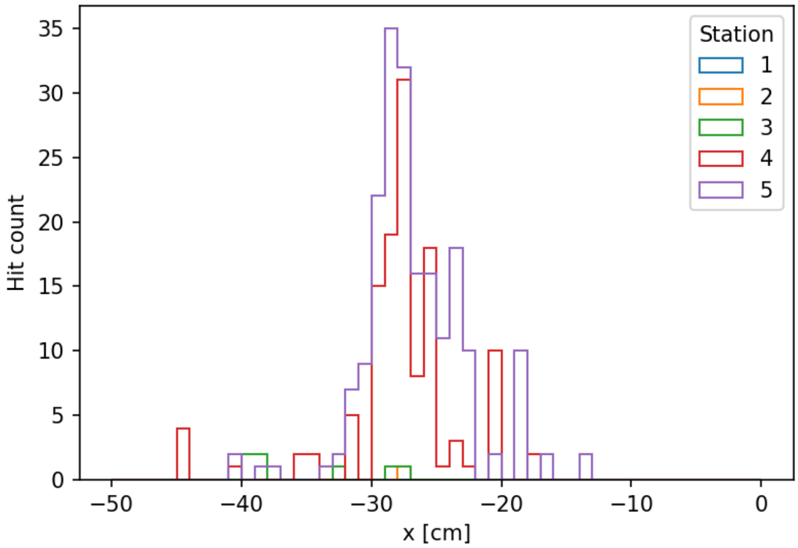

Figure 20: The SciFi hit distributions from a single pion shower of energy 300 GeV and depth of 304 cm, for all the 5 SciFi stations. For this set of hit distributions depth reconstruction fails.

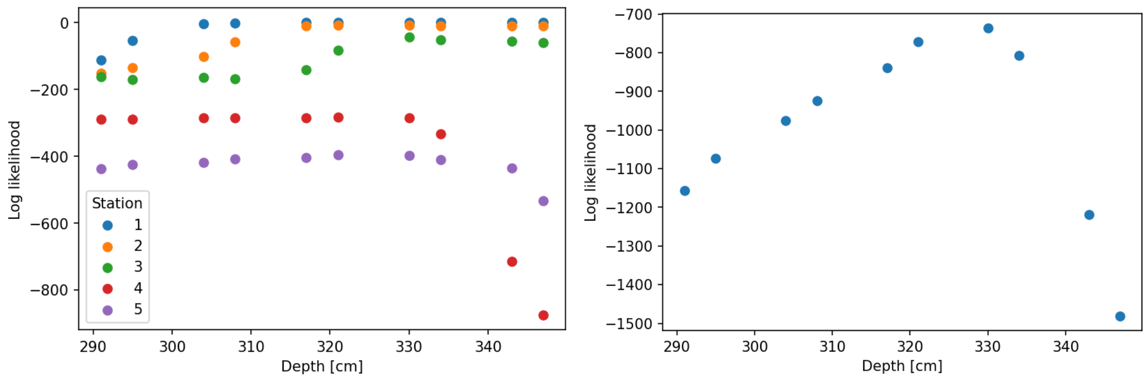

(a) (b)

Figure 21: The results of a failed log-likelihood scan in depth shown separately for each SciFi station in (a) and summed between all the stations in (b).

In this case, although the pion is fired from the beginning of the 2nd station there is no shower development until the 4th and 5th stations. The per-station log-likelihood shows that depths before and after the shower starts are penalised while the showering depth has close to a maximal value for each station. When combining the per station results in Figure 21b it can be seen that the maximum log-likelihood occurs at the 4th station where the shower activity really begins. Therefore in this case the reconstruction was accurate for the actual showering depth but since this showering depth does not correspond to the pion firing depth the reconstruction is deemed as failed.

This example brings to attention the discrepancy between the firing depth used as the truth label and the actual shower initiation depth. The firing depth of pions in GEANT4 is not a reliable method to determine the shower initiation depth as pions often travel some distance before showering.

For this instance, the depth reconstruction is actually successful but the truth label is incorrect. As showers are more likely to develop further downstream from the firing depth this would lead to a tendency to overestimate the depth during reconstruction.

3.7.3 Discussion

The energy and vertex position reconstruction method via sampling of the shower profile relies on consistent behaviour between individual showers where the average PDF can accurately represent individual events. However, as the activity between different events becomes less consistent the average PDF created from these events becomes less representative of each individual event and the log-likelihood becomes less reliable.

Energy reconstruction is less affected by the reliability of the true shower initiation depth and is more dependent on the shower development over the stations.

Vertex position reconstruction is less affected by the shower development over the stations and more by the reliability of the true shower initiation depth.

4 Conclusion

The SND@LHC experiment aims to measure the cross-section of high-energy neutrinos produced in LHC collisions to study charmed hadron production and extract the gluon PDF at low momentum fractions. In this project, the first procedure for hadronic shower energy and vertex reconstruction was developed for the SND detector.

The reconstruction procedure was applied on simulated charged pions initiating hadronic showers with different energies and depths into the target region of the detector. Only the SciFi planes were used for reconstruction. Shower propagation was profiled by creating PDFs for each energy and depth configuration for each SciFi station. A log-likelihood function was built to scan over the PDFs to reconstruct the energy and vertex position of the showers.

Energy reconstruction was very accurate for showers initiated in the first station, with the energy being increasingly overestimated for the low energies and underestimated for high energies for stations further downstream. It was determined that energy reconstruction begins to fail when the shower activity in the stations becomes too different from the average activity in the PDF distributions. For later station showering, a smaller section of the shower is sampled which is more prone to fluctuations that bias the log-likelihood scan.

Vertex reconstruction consistently overestimated the shower initiation depth, with the effect becoming weaker for later stations as there were fewer stations further downstream to bias the reconstruction. Overestimation of the depth improved at higher energies. A discrepancy was identified between the true vertex depth label and the actual shower initiation depth. It was found that showers often developed further downstream than the firing depth. In these cases shower vertex depths were correctly reconstructed but look like overestimates when the firing depth is used as the true vertex depth. Results suggested vertex reconstruction would perform significantly better if the true vertex depth is corrected.

References

[1] SND Collaboration, *SND@LHC: The Scattering and Neutrino Detector at the LHC*, (2023), arXiv: 2210.02784 [hep-ex].

[2] SND Collaboration, *Observation of Collider Muon Neutrinos with the SND@LHC Experiment*, Phys. Rev. Lett. **131** (2023) 031802, arXiv: 2305.09383 [hep-ex].

[3] FASER Collaboration, *First Direct Observation of Collider Neutrinos with FASER at the LHC*, Phys. Rev. Lett. **131** (2023) 031801, arXiv: 2303.14185 [hep-ex].

[4] N. Beni et al., *Further studies on the physics potential of an experiment using LHC neutrinos*, J. Phys. G: Nucl. Part. Phys. **47** (2020) 125004, arXiv: 2004.07828 [hep-ex].

[5] M. Bustamante and A. Connolly, *Extracting the Energy-Dependent Neutrino-Nucleon Cross Section above 10 TeV Using IceCube Showers*, Phys. Rev. Lett. **122** (2019) 041101, arXiv: 1711.11043 [astro-ph.HE].

[6] M. Gupta et al., *Calculation of radiation length in materials*, PH-EP-Tech-Note-2010-013, 2010, URL: https://cds.cern.ch/record/1279627?ln=en.

[7] J. V. Tilburg, *Information on the absorption of hadrons*, URL: https://twiki.cern.ch/twiki/bin/view/LHCb/TrackingEffAbsLength (visited on 16/09/2023).

[8] GEANT4 Collaboration, S. Agostinelli et al., *Geant4 – a simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250.

[9] SND Collaboration, *SNDSW*, URL: https://github.com/SND-LHC/sndsw (visited on 16/09/2023).