

Particle Tracking Classification in the CONNIE Experiment

Sara Mirthis Dantas dos Santos¹, Irina Nasteva², Paula Dornhofer Paro Costa¹
on behalf of the CONNIE collaboration

¹ Faculdade de Engenharia Elétrica e de Computação – Universidade
Estadual de Campinas (UNICAMP)
Campinas – SP – 13083-852 – Brazil

²Instituto de Física – Universidade Federal do Rio de Janeiro (UFRJ)
Rio de Janeiro – RJ – 21941-909 – Brazil

s224018@dac.unicamp.br, irina@if.ufrj.br, paulad@unicamp.br

Abstract. *The Coherent Neutrino-Nucleus Interaction Experiment (CONNIE) is a particle physics experiment located at the Angra 2 Nuclear Reactor. Using silicon Skipper-CCDs, CONNIE aims to identify and investigate coherent elastic neutrino-nucleus scattering (CE ν NS). However, background particles from various sources, combined with the large volume of unlabeled data, present significant challenges, making statistical analysis difficult. To address these challenges, this work investigates and implements machine learning and image processing methodologies to classify distinct CONNIE events. A labeled dataset of real experimental events will also be created using the Annotation Redundancy with Targeted Quality Assurance method. This approach involves multiple annotators labeling the same data, enabling comparison to identify discrepancies, refine labeling accuracy, and enhance dataset reliability.*

1. Introduction

Using silicon detectors, the CONNIE experiment aims to identify and investigate Coherent Elastic Neutrino-Nucleus Scattering (CE ν NS), an extremely rare process in which a neutrino, or antineutrino, weakly interacts with an entire atomic nucleus, causing the nucleus to recoil [Aguilar-Arevalo et al. 2024]. In CONNIE, this interaction occurs between reactor-produced antineutrinos and the silicon nuclei within the detectors, causing a small recoil that can be detected by the experiment, as shown in Figure 1.




Figure 1. A neutrino (ν) hits a silicon (Si) nucleus, producing a recoil that creates electron-hole pairs in a CCD pixel for readout. Extracted from [Bonifazi 2017].

Located 30 meters from the Angra 2 reactor core of the Amirante Álvaro Al-
 berto Nuclear Power Plant in the state of Rio de Janeiro, Brazil, the experiment has been
 operating since 2016 [Aguilar-Arevalo et al. 2019]. Initially, it was equipped with 14
 charge-coupled devices, or CCDs, each with 4000 x 4000 pixels. Then, in 2021, it was
 upgraded with 2 Skipper-CCDs, making CONNIE the first experiment to use this type of
 sensor for detecting neutrinos from a nuclear reactor. This upgrade was a significant step,
 as Skipper-CCDs are part of a new generation of CCDs with a lower detection threshold
 and greater sensitivity, enabling the counting of the exact number of electrons collected
 in each detector pixel [Nasteva 2021].

While the CONNIE experiment is designed to detect $\text{CE}\nu\text{NS}$, the presence of
 background particles from various sources creates significant challenges. Differentiating
 between these background events and neutrino interactions is a key aspect to the experi-
 ment's success, as neutrinos are among the most abundant yet least understood particles
 in the universe due to their weak interaction with other matter [Kayser and Parke 2009].
 The main cause of background noise is cosmic rays, but additional sources include envi-
 ronmental radiation and radioactive contamination of the materials in the detector itself,
 as well as spontaneous charge creation in the sensor.

Each particle interacting with the detector sensors leaves a signature that depends
 on variables such as type, energy, and angle of incidence [Bartovsky et al. 2011]. For
 example, muons produce straight tracks, electrons produce curly or irregular tracks, al-
 pha particles produce high-energy blobs, and low-energy electrons or photons produce
 dots or blobs [Ventura 2024]. However, identifying particles solely by shape is chal-
 lenging, as the same particle can produce different trajectories depending on its crossing
 angle [Čelko 2021], and particles such as low-energy X-rays and neutrinos may leave
 similar signature tracks, further complicating classification. Given these challenges and
 the large volume of events, defined as particle interactions or collisions observed by the
 detector [Haas 2007], statistical analysis is particularly difficult.

In this context, a few works have explored classifying particles and interactions
 detected in the sensors based on the track shape, often using image processing techniques
 and deep learning. This approach improves classification accuracy, allowing researchers
 to better track particle interactions and properties, furthering their understanding of how
 particles interact with detector materials. Additionally, it accelerates data analysis by
 making it easier to extract meaningful information from the datasets.

Therefore, the objective of this work is to investigate and implement machine
 learning and image processing methodologies to classify distinct CONNIE events that
 are detected by the Skipper-CCD sensors. In parallel, a comprehensive database of real
 events from the experiment will be created to support this research.

2. Concepts and Related Work

2.1. The CONNIE Experiment

The CONNIE setup (Figure 2a) includes 14 standard CCDs and was later upgraded with 2
 Skipper-CCDs installed in a copper box (Figure 2b), surrounded by shielding to minimize
 background noise. Each silicon Skipper-CCDs has a thickness of 675 μm , a resolution
 of 1022×682 pixels, and a pixel size of $15 \times 15 \mu\text{m}^2$, resulting in an area of 1.57 cm^2

(Figure 2c). They were placed in the lowest slots of the copper box, while the standard CCDs from the previous CONNIE setup remained in their original positions.




Figure 2. CONNIE experiment. Adapted from [Aguilar-Arevalo et al. 2024] and [Aguilar-Arevalo et al. 2019].

Skipper-CCDs are advanced scientific CCDs with a modified readout stage that allows multiple nondestructive measurements per pixel [Tiffenberg et al. 2017]. Averaging these measurements suppresses electronic noise to sub-electron levels. In this work, $N_{skp} = 400$ samples per pixel are used, with data collected in runs under consistent detector conditions. The output consists of 2D FITS (Flexible Image Transport System) files, widely used in astronomy for storing scientific data [Pence, W. D. et al. 2010]. Each set of N_{skp} consecutive pixels in a row corresponds to repeated measurements of the same physical pixel. Processing follows the steps detailed in [Aguilar-Arevalo et al. 2024]:

- Average of N_{skp} : the bidimensional image is created by averaging each pixel's N_{skp} measurements;
- Baseline correction: a horizontal baseline for each row based on the overscan (unexposed pixels) and a vertical baseline for each column (based on the first 100 pixels). These baselines are subtracted from the respective rows and columns;
- Energy calibration: convert the charge from analog-to-digital units (ADU) to electrons (e^-) and consequently to energy in eV ($1e^-$ corresponds to 3.752 eV, which is the mean energy necessary for ionizing one silicon atom).

A raw image with the average of N_{skp} and a processed image are shown in Figures 3a and 3b respectively. In the current processing pipeline, after obtaining the processed image, a mask is created to remove Serial Register Events (SRE) and hot pixels, as seen in Figure 3c. SREs appear as horizontal charge deposits across a single row of pixels, whereas hot pixels are defective pixels represented by single bright pixels, columns, or rows.

The next step in the data processing pipeline is creating a catalog from the processed energy-calibrated images and masks. Stored in ROOT format developed by CERN [Brun and Rademakers 1997], the catalog organizes events as a TTree, which can store large quantities of same-class objects as branches, in this case the variables for each individual event such as total pixels or position. The leaves contain the actual data for that variable.




Figure 3. CONNIE images.

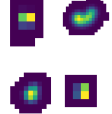


Events are identified by locating a seed pixel with charge above 1.6 e^- (about ten times the readout noise), then reconstructed by adding neighboring pixels with charge above 0.64 e^- . Information such as X and Y coordinates, pixel energy, total charge, number of pixels, barycenter position, and sensor number are stored as branches in the ROOT catalog. Events with a barycenter within a masked pixel are assigned a catalog flag corresponding to that masked pixel: 128 for SRE, and 64 for hot pixels. For events without any associated issues, the flag is set to 0.

Despite shielding to reduce background contamination mainly from cosmic rays (particles that travel through space and interact with the earth's atmosphere) and environmental radiation, some background particles still reach the detector. Expected event categories include Muon, Electron, Blob, Diffusion Hit, Alpha, and Others. Alpha particles are rare and usually come from external contamination. Blobs come from low-energy electrons or photons, while Diffusion Hits may include X-rays, low-energy gamma photons, and $\text{CE}\nu\text{NS}$ interactions, the main focus of the experiment.

Table 1 shows the classification criteria and example events, plotted by X , Y coordinates and pixel energy.

Table 1. Event types and selection criteria

Type	Criteria	Example
Muon	Straight, uninterrupted track	
Electron	Not straight track, worm shape	
Blob	1 or more pixels in the event, energy above 600 eV	

Diffusion hit	1 or more pixels in the event, energy below 600 eV	
Alpha	Round shape, high energy (usually from 3 to 9 MeV)	
Others	Unidentified type, usually more than one event with crossed trajectory	

2.2. Related Works

Experiments such as SENSEI [Adari et al. 2025] and DAMIC-M [Arnquist et al. 2023] use Skipper-CCD sensors for dark matter searches but do not report applying particle classification algorithms, likely due to their low-background underground environments. In contrast, above-ground experiments like CONNIE face higher background levels, making classification essential.

Related works explore particle classification with other detector technologies, including Complementary Metal Oxide Semiconductor (CMOS) sensors and Timepix hybrid detectors. For example, [Čelko 2021] developed a set of tools that includes a 2D and 3D cluster viewer, cluster property computation, cluster filtering based on properties and particle classifiers for Timepix data, achieving up to 87.5% accuracy using a feature-based Multi-Layer Perceptron (MLP) with multi-level classifier, which combines the outputs from multiple classifiers in a tree-like structure. Similarly, [Furnell et al. 2019] analyzed data from a Timepix-based detector on the TechDemoSat-1 satellite, applying clustering techniques followed by neural network classifiers to identify particle tracks as alpha, beta (electron), X-ray, proton, muon, and others. Using a manually labeled dataset, both a Metric-Based Network (MBN) and a Convolutional Neural Network (CNN) were trained, achieving classification accuracies above 80% on the training and test datasets, although the MBN demonstrated superior performance.

Using smartphone camera images from the Distributed Electronic Cosmic-ray Observatory (DECO) dataset, [Winter et al. 2019] developed a binary classifier using geometric features to categorize particles as either tracks or non-track, but the model struggled to separate worm-like shapes from straight tracks. To improve classification, a DECO-specific CNN was trained on manually labeled data to identify tracks, worms, spots, and various types of noise (e.g., hot pixels, thermal fluctuations, and sensor artifacts), achieving over 90% accuracy across all classes, suggesting that it may be suitable for other experiments using CCD and CMOS sensors for particle detection, such as DAMIC-M. Similarly, [Hachaj et al. 2021] used the CREDO dataset to classify cosmic-ray event types into spots, tracks, worms, and artefacts, with five pre-trained CNNs – VGG16, NASNetLarge, MobileNetV2, Xception, and DenseNet201 – via transfer learning and a threshold-based scheme. VGG16 achieved the highest accuracy of 85.8% at a 0.9 threshold, though 62.6% of events remained unclassified due to low confidence, with most misclassifications occurring between tracks and worms due to their visual similarities.

In contrast, [Bar et al. 2021] extracted Zernike moments as shape descriptors from CREDO images and trained 19 classifiers: 11 individual models (e.g., K-Nearest Neighbors, Decision Tree, MLP) and 8 ensemble methods (e.g., Gradient Boosting, Extra Trees, One-vs-Rest). Ensemble classifiers performed best, likely because they better handle particles with similar visual morphology. The highest performance was achieved by an ensemble MLP with a One-vs-Rest strategy, reaching 91.0% accuracy in cross-validation and 89.7% on the test set. The comparative outcomes are summarized in Table 2.

Table 2. Accuracy comparison of different studies on cosmic ray event classification. Adapted from [Bar et al. 2021].

Class	[Hachaj et al. 2021]	[Winter et al. 2019]	[Bar et al. 2021]
Spots	98.71%	98.90%	98.13%
Tracks	88.89%	95.40%	93.67%
Worms	89.65%	92.90%	88.52%
Artefacts	97.70%	98.20%	91.96%

The ensemble MLP in [Bar et al. 2021] improved track accuracy compared to [Hachaj et al. 2021] but slightly underperformed on worms. Unlike [Hachaj et al. 2021], which discarded low-confidence events via thresholding, [Bar et al. 2021] classified all events, resulting in slightly lower overall accuracy. Both methods fell short of [Winter et al. 2019], which achieved the highest accuracy but only for events exceeding an 80% probability threshold. Despite these differences, the ensemble MLP showed comparable performance to CNN models.

The review shows that feature-based methods and CNNs are the most commonly employed, often achieving comparable results depending on implementation. As both methods rely on supervised learning, established models were first applied to labeled data to create a baseline and evaluate performance, guiding future improvements.

Given that reviewed models achieve over 80% accuracy, training a classifier specifically on CONNIE data can significantly improve the identification of relevant events. This is especially important because CONNIE is the first experiment to search for CE ν NS using a pixel-based detector, and it features unique event classes, such as diffusion hits which may correspond to these rare interactions. Adapting established models to these characteristics enables more targeted classification, helping prioritize key events and streamline the analysis process.

3. Proposed Approach

As outlined in Section 2.2, the majority of related works rely on either feature-based or convolutional models. Furthermore, [Winter et al. 2019] demonstrated that a CNN trained on DECO data could generalize to other experiments with similar sensor configurations. As the CONNIE dataset is currently unlabeled, it is therefore proposed to evaluate whether a CNN model trained on a different labeled dataset can effectively classify CONNIE events as well, before focusing on manual labeling.

The CREDO database was chosen due to its use in previous studies and contains cosmic-ray event images captured by smartphones and other CMOS-equipped devices [Homola et al. 2020]. The dataset contains 2356 RGB images, each showing a

cropped detected hit with a resolution of 60×60 pixels, stored in 8-bit per channel (uint8) format. Figure 4 displays images representing the four classes introduced in [Piekarczyk et al. 2021]: worms, dots, tracks, and artefacts.



Figure 4. CREDO dataset classes. Extracted from [Piekarczyk et al. 2021].

For this analysis, the dataset was converted to grayscale and split into training, validation, and testing sets, representing 70%, 15%, and 15% of the data, respectively. The number of images in each set is detailed in Table 3.

Table 3. Number of images in each class of the Credo dataset.

Class	Train	Validation	Test
Worms	212	46	46
Dots	374	80	81
Lines	275	59	59
Artefacts	785	168	169

After processing the CREDO data, the CONNIE data were prepared similarly by cropping detected events from processed images rather than using the existing catalog, which is a more general approach that can be suitable for other datasets. Run 125, containing 1004 calibrated images, was used. To avoid losing relevant events, image masks were not applied except for masking 5-pixel borders. Pixels below 0.64 e^- , considered irrelevant or negative pixels, were set to zero.

Using Python’s OpenCV, each image underwent Gaussian blur and a global threshold of 10 eV, since the sensor’s detection threshold is 15 eV, and lower values added noise. Various morphological transformations were tested for optimal contour detection, with the best sequence being opening, followed by dilation and erosion. These steps remove noise, fill gaps, and refine contours [Gonzalez and Woods 2008]. Contours were cropped and padded to match training data, yielding 30704 PNG images. A sample is shown in Figure 5.



Figure 5. Example of cropped events from the CONNIE images.

To optimize the classifier, several architectures and hyperparameters were tested, including batch size, learning rate, epochs, and regularization. ResNet-18 was selected for its balanced architecture and residual connections, which improve training stability

and feature extraction. Data augmentation was applied, including horizontal and vertical flips and random rotations. Hyperparameters were tuned using grid search to identify the most accurate combination. The Adam optimizer was used with weight decay of $1e-3$ and an initial learning rate of $5e-4$, reduced by a factor of 0.9 every 10 epochs.

3.1. Data Labeling

Concurrently, a labeled dataset from the CONNIE experiment is being developed for future analysis and model training, using catalogs from runs 118 and 125. After filtering for valid events (flag 0) and applying a fiducial cut to exclude problematic regions, 13390 images were obtained from run 118 and 13336 from run 125.

A custom GUI was created to allow physicists to visually inspect each event and assign a label. Alongside the event plot, relevant information such as the total number of pixels, energy, and X/Y barycenters is displayed to assist classification (Figure 6a). To reduce human error, the Annotation Redundancy with Targeted Quality Assurance (QA) method [Tseng et al. 2020] is applied, in which 10% of the data are labeled by three independent annotators. When the GUI loads a ROOT file, these events are automatically flagged for redundant labeling. Each event is shown with its associated data, and users may skip or assign a label. Upon submission, the label and annotator identification are stored unless the event has already been labeled, in which case the next image is loaded.

For reviewing discrepancies, a review mode within the annotation interface was developed (Figure 6b). This mode is limited to a default user and secured with a password to prevent unintentional or unauthorized access. It displays only events where annotators disagreed, along with the assigned labels and how many times each was selected. A button is also available to automatically resolve conflicts when there is a majority agreement.

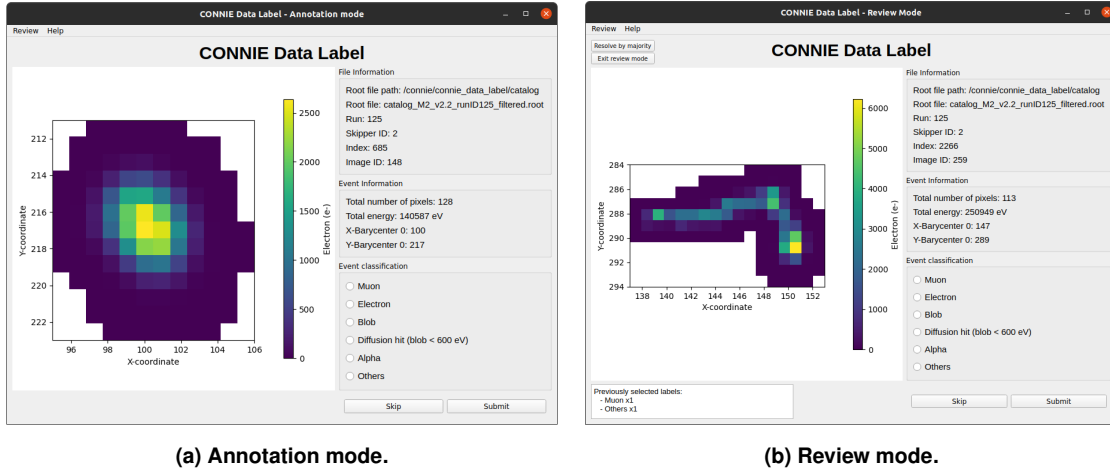


Figure 6. Data label GUI.

4. Results

4.1. Training with CREDO Data

The Resnet-18 was trained with 150 epochs, where the training loss and accuracy were 7.80% and 96.90%, respectively, while the validation loss was 11.19% and the accuracy was 96.60%. After the training phase, the model was tested and evaluated on the test

subdataset. The confusion matrix in Table 4 shows that only a few errors occurred in each class.

Table 4. Confusion matrix on the test dataset

		Predicted			
		Artefacts	Dot	Line	Worm
True	Artefacts	168	0	0	1
	Dot	0	81	0	0
	Line	1	0	50	8
	Worm	0	0	6	40

Based on these results, the accuracy, precision, recall, and F1-score metrics were calculated. The metrics results for each class are shown in Table 5, where Worms performed the worst and Dots the best. This outcome is expected, as Worms has the least amount of data and Dots has the most distinct shape compared to the other event types. Nevertheless, all classes achieved metrics above 80%, indicating strong overall performance.

Table 5. Evaluation metrics for each class.

Class	Precision	Recall	F1-score
Worms	81.63	86.96	84.21
Dots	100.00	100.00	100.00
Lines	89.28	84.74	86.95
Artefacts	99.41	99.41	99.41

As the classification is multiclass, to get the average precision, recall, and F1-score were calculated as the weighted averaged across all labels, resulting in 92.59%, 95.49%, 95.50%, respectively, and an accuracy of 95.49%.

4.1.1. CONNIE Data Process and Classification

The previously trained model was applied to the newly labeled CONNIE dataset (Section 3.1). Figure 7 shows examples of images classified into each category.




Figure 7. Examples of event classification in the CONNIE dataset.

Feedback from CONNIE collaborators indicated that artifacts in the CREDO dataset differ from those in CONNIE, leading the model to misclassify valid events as artifacts. Also, worms are the most frequently misclassified class, often mistaken for either an elongated Dot or with Lines, whereas dots are classified most accurately due to their distinct shape.

Overall, the CONNIE collaboration noted that classification performance was limited, likely due to differences in the sensor technologies used by CONNIE and CREDO, as well as disparities in the types of events collected.

4.2. Data Labeling and Training with CONNIE Data

Given the unsatisfactory previous results, using labeled CONNIE events in future analyses was deemed necessary. The dataset, currently being developed with CONNIE collaboration members, includes 2076 annotated events saved in PNG and ROOT formats. Among these, 227 have been labeled by multiple users, with 45 showing disagreement. The distribution by label is as follows:

- Alpha: 5, Blob: 184, Diffusion Hit: 9, Electron: 166, Muon: 1469, Others: 243.

Due to the high number of muon events caused by the abundance of cosmic-ray muons at surface level and the resulting class imbalance, the One-vs-Rest (OvR), or One-vs-All (OvA), classification strategy was adopted as an initial approach. This method divides the task into several binary classification problems, beginning with distinguishing muons from non-muons.

For preliminary analysis, PNG files were used to train a new ResNet-18. The labels with disagreement were removed from the dataset, retaining a total of 1741 events. A grid search was performed to tune hyperparameters, setting weight decay to $5e - 5$ and initial learning rate to $5e - 4$, which decayed by 0.5 every 10 epochs. Stratified K-fold cross-validation with $k = 5$ subsets was applied, training the model iteratively on $k - 1$ folds while validating on the remaining one. The same data augmentation techniques that were previously used were applied to the new training and validation datasets, including horizontal and vertical flips and random rotations. ResNet-18 was trained for up to 100 epochs, with early stopping if the validation accuracy did not improve after 10 epochs. The model had an average training accuracy of 0.9359 and a loss of 0.1685, with validation accuracy of 0.9076 ± 0.0177 and loss of 0.1580 ± 0.0485 .

To compare CNN results and further explore the problem, feature-based classifiers were trained on the same 1741 events in ROOT format. Analyzing the event features, constant value features were removed, and the mean of features such as pixel energy was calculated. Features with correlations above 0.99 were eliminated, reducing the total from 46 to 22. Random Forest and XGBoost were chosen due to their strong performance and robustness across diverse datasets [Fatima et al. 2023]. XGBoost used 100 estimators, a learning rate of 0.05, and a regularization parameter gamma of 0.98 to reduce overfitting, using log loss as the evaluation metric. Random Forest was configured with 300 trees, a maximum of 30% of features considered at each split, and a minimum of 10 samples required to split a node. As before, stratified K-fold cross-validation with $k = 5$ was used to ensure balanced evaluation across classes. XGBoost and Random Forest achieved training accuracies of 0.9880 and 0.9692, and validation accuracies of 0.8803 ± 0.0098 and 0.8844 ± 0.0113 , respectively. Both classifiers performed comparably to ResNet-18, suggesting that feature-based methods may be an effective approach for this task.

5. Conclusion

This work explores the application of machine learning and image processing techniques to classify particle interaction signatures and the creation of a labeled dataset. These

tools aim to assist the CONNIE experiment in its goals of detecting neutrino interactions and advancing our understanding of neutrino-nucleus interactions. Additionally, it opens doors to parallel investigations, such as the analysis of muon flux, which can contribute to better detector calibration and shielding strategies.

The results indicate that the CNN model trained on the CREDO data does not perform well when applied to the CONNIE data. While the image processing used for event cropping was effective, the event extraction method used for creating the CONNIE catalog is more robust. It not only flags potential issues such as hot pixels and SRE but also provides valuable features for each event, such as pixel count, energy, and barycenter, that can support the development of a stronger classifier.

On the other hand, the CNN trained directly on CONNIE events delivered promising results. Moving forward, a benchmark will compare this approach with the feature-based model, which achieved accuracy comparable to ResNet-18 with lower computational cost. Both approaches will be improved, with the feature-based model refined to better exploit catalog information. Due to class imbalance, only the muon class has been trained so far, but remaining classes will be incorporated as annotation progresses. As OvA classifiers are developed, they may be combined into an ensemble model to improve robustness, potentially using confidence-based voting or similar techniques.

6. Acknowledgments

We thank the CONNIE collaboration for their invaluable support and contributions to this research.

References

- Adari, P. et al. (2025). First Direct-Detection Results on Sub-GeV Dark Matter Using the SENSEI Detector at SNOLAB. *Phys. Rev. Lett.*, 134(1):011804. (SENSEI Collab.).
- Aguilar-Arevalo, A. et al. (2019). Exploring low-energy neutrino physics with the Coherent Neutrino Nucleus Interaction Experiment. *Physical Review D*, 100(9):092005. (CONNIE Collab.).
- Aguilar-Arevalo, A. A. et al. (2024). Searches for $\text{CE}\nu\text{NS}$ and Physics beyond the Standard Model using Skipper-CCDs at CONNIE. arXiv:2403.15976 [hep-ex]. (CONNIE Collab.).
- Arnquist, I. et al. (2023). The DAMIC-M Experiment: Status and First Results. *SciPost Phys. Proc.*, 12:014. (DAMIC-M Collab.).
- Bar, O., Bibrzycki, L., Niedźwiecki, M., et al. (2021). Zernike Moment Based Classification of Cosmic Ray Candidate Hits from CMOS Sensors. *Sensors*, 21(22).
- Bartovsky, J., Schneider, D., Dokladalova, E., et al. (2011). Morphological classification of particles recorded by the Timepix detector. In *2011 7th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pages 343–348.
- Bonifazi, C. (2017). Experimento CONNIE para detecção de neutrinos com CCDs na usina nuclear de Angra II. <https://indico.cern.ch/event/647995/contributions/2652498/attachments/1497846/2349509/CONNIE-Bonifazi-Renafae.pdf>. (CONNIE Collab.). Instrumentação Científica em Física de Altas Energias no Brasil.

- Brun, R. and Rademakers, F. (1997). ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth. A*, 389:81–86.
- Fatima, S., Hussain, A., Amir, S., Ahmed, S. H., and Aslam, S. (2023). XGBoost and Random Forest Algorithms: An in Depth Analysis. *Pakistan Journal of Scientific Research*, 3:26–31.
- Furnell, W., Shenoy, A., Fox, E., and Hatfield, P. (2019). First results from the LUCID-Timepix spacecraft payload onboard the TechDemoSat-1 satellite in Low Earth Orbit. *Adv. Space Res.*, 63:1523–1540.
- Gonzalez, R. C. and Woods, R. E. (2008). *Digital Image Processing*. Prentice Hall, Upper Saddle River, NJ, 3rd edition.
- Haas, T. (2007). Particle event. Symmetry Magazine. Available at: <https://www.symmetrymagazine.org/article/junejuly-2007/particle-event>. Accessed: 2024-10-20.
- Hachaj, T., Bibrzycki, L., and Piekarczyk, M. (2021). Recognition of Cosmic Ray Images Obtained from CMOS Sensors Used in Mobile Phones by Approximation of Uncertain Class Assignment with Deep Convolutional Neural Network. *Sensors*, 21(6).
- Homola, P. et al. (2020). Cosmic-Ray Extremely Distributed Observatory. *Symmetry*, 12(11):1835. (CREDO Collab.).
- Kayser, B. and Parke, S. (2009). The neutrinos. Fermi National Accelerator Laboratory. Available at: <https://physics.ucsc.edu/~joel/Phys129/Kayser&Parke-Neutrino-overview.pdf>. Accessed: 2024-10-20.
- Nasteva, I. (2021). Low-energy reactor neutrino physics with the CONNIE experiment. *J. Phys. Conf. Ser.*, 2156(1):012115. (CONNIE Collab.).
- Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R. A., and Stobie, E. (2010). Definition of the Flexible Image Transport System (FITS), version 3.0. AA, 524:A42.
- Piekarczyk, M. et al. (2021). CNN-Based Classifier as an Offline Trigger for the CREDO Experiment. *Sensors*, 21(14):4804.
- Tiffenberg, J., Sofo-Haro, M., Drlica-Wagner, et al. (2017). Single-Electron and Single-Photon Sensitivity with a Silicon Skipper CCD. *Physical Review Letters*, 119(13). (SENSEI Collab.).
- Tseng, T., Stent, A., and Maida, D. (2020). Best Practices for Managing Data Annotation Projects. *CoRR*, abs/2009.11654.
- Ventura, P. Z. M. (2024). Skipper-CCDs para detecção de neutrinos: primeiros resultados com o experimento CONNIE. Master’s thesis, Universidade Federal do Rio de Janeiro.
- Winter, M., Bourbeau, J., Bravo, S., Campos, F., Meehan, M., et al. (2019). Particle Identification In Camera Image Sensors Using Computer Vision. *Astroparticle Physics*, 104:42–53.
- Čelko, T. (2021). Support for annotating and classifying particles detected by Timepix3. Bachelor’s thesis, Charles University.