



CNNCat: categorizing high-energy photons in a Compton/Pair telescope with convolutional neural networks

Jan Peter Lommler¹ · Uwe Gerd Oberlack¹

Received: 22 May 2024 / Accepted: 23 October 2024 / Published online: 4 November 2024
© The Author(s) 2024

Abstract

A Compton/Pair telescope, designed to provide spectral resolved images of cosmic photons from sub-MeV to GeV energies, records a wealth of data in a combination of tracking detector and calorimeter. Onboard event classification can be required to decide on which data to down-link with priority, given limited data-transfer bandwidth. Event classification is also the first and one of the most crucial steps in reconstructing data. Its outcome determines the further handling of the event, i.e., the type of reconstruction (Compton, pair) or, possibly, the decision to discard it. Errors at this stage result in misreconstruction and loss of source information. We present a classification algorithm driven by a Convolutional Neural Network. It provides classification of the type of electromagnetic interaction, based solely on low-level detector data. We introduce the task, describe the architecture and the dataset used, and present the performance of this method in the context of the proposed (e-)ASTROGAM and similar telescopes.

Keywords γ -Astronomy · Machine learning · Event classification · Compton telescope · Pair telescope

1 Introduction

Gamma-ray astronomy in the energy range of sub-MeV to several 10s of MeV remains little explored. At these energies, the Compton effect dominates the total cross-section of photons with matter, changing to dominant pair production in the upper part of this range. The intransparency of the atmosphere requires Compton and pair telescopes to be operated from space. Currently, INTEGRAL [1] provides varying sensitivity from 15 keV to 10 MeV, while Fermi-LAT [2] covers the range from 30 MeV to 300 GeV. Higher energy photons are observed by ground-based observatories like H.E.S.S. [3],

✉ Jan Peter Lommler
lommler@uni-mainz.de

¹ Institute of Physics & PRISMA⁺ Excellence Cluster, Johannes Gutenberg University Mainz, 55099 Mainz, Germany

Veritas [4] and MAGIC [5] or, in the future, CTA [6]. More than 20-year old all-sky maps in 0.7 – 30 MeV are available from the Compton telescope COMPTEL [7] on the Compton Gamma-Ray Observatory (CGRO), yet at the limited sensitivity of a pioneering instrument. Comparatively poor sensitivity in the energy range from about 0.1 MeV to 100 MeV leaves a window open for discovery, often referred to as the *MeV Gap*.

INTEGRAL as well as Fermi-LAT are successful observatories with a long service record, but at some point they will need to be replaced. In fact, INTEGRAL will be placed in standby at the end of 2024. The likelihood for a large-scale mission like CGRO with four gamma-ray instruments to happen again soon is low. This motivates the development of detector concepts that cover the MeV gap and beyond with one instrument at the cost scale of a medium-size mission. A possible solution to this challenge is a Compton/Pair telescope. Example design proposals are the ASTROGAM detector family (e.g., e-ASTROGAM [8], AllSky-ASTROGAM [9]) and AMEGO [10].

1.1 Imaging with a Compton/Pair telescope

For energies between several hundred keV and several GeV the interaction of γ -radiation with matter is dominated by Compton scattering and pair production. A telescope that is able to image both Compton and pair events, and to determine the energy of the source photon is called a Compton/Pair telescope. It can cover up to five orders of magnitude of the electromagnetic spectrum. This study uses the e-ASTROGAM design proposal [8] as a reference detector but the concept could be applied to other Compton/Pair telescopes accordingly. Figure 1 shows a schematic of the e-ASTROGAM concept. The main components are a tracker (D1) consisting of stacked Double-Sided Silicon Strip Detectors (DSSD), a calorimeter unit (D2) made of pixelated Tl-doped cesium iodide (CsI(Tl)), which allows for a depth and x/y-position-sensitive energy measurement, and an active veto system (AC) made of segmented plastic scintillator to prevent triggering on charged primaries. This design




Fig. 1 Left panel: Schematic view of e-ASTROGAM, consisting of 56 layers of DSSD-tracker (D1) on top of a pixelated CsI(Tl) calorimeter (D2). D1 and D2 are surrounded by an active veto made of plastic scintillator [8]. Right panel: Monte Carlo mass model generated with MEGALib [11]

principle enables the tracking of photons through their multiple hits in a Compton-scattering sequence, as well as of charged particles as they traverse the detector layers. The ability to accurately measure the energy deposit of charged particles allows the instrument to distinguish between tracked and untracked Compton events, pair events, and charged-particle tracks originating from, e.g., cosmic rays or albedo secondaries, as depicted in Fig. 2.

1.2 Misclassification and data handling in space

Bandwidth in Low Earth Orbit is a limited resource. It has to be shared with house-keeping data, satellite telemetry, and secondary instruments on the respective satellite bus. One important goal of this study therefore consists in reliably identifying and discarding events that will not be usable in a later stage of event reconstruction. Event classification is crucial in reaching this goal. A small and computationally lean model that divides events into potentially reconstructable "good" events and non-reconstructable or background events can help to reduce the amount of data that has to be transmitted. It can also improve event-type-specific preprocessing onboard. Application of an incorrect reconstruction algorithm to an otherwise valid event leads to a false reconstruction of source photons into the background. The purpose of this study is to investigate whether it is possible to classify events with high efficiency and speed




Fig. 2 Event topologies encountered in a Compton/Pair telescope. Compton scatters can occur without and with a trackable recoil electron (top left and center). To determine the cone of origin of the source photon, the correct sequence of first and second hit has to be identified. For tracked Compton events, the recoil electron track must be isolated from the photon hit pattern and reconstructed separately. This defines a second cone of origin, whose intersection with the photon cone further constrains the origin of the photon. Pair events are reconstructed using the opening angle between electron and positron, providing the photon's origin and energy if the pair-creation vertex occurs in the tracker (top right). Possible background events (bottom row) include charged particle tracks (bottom left) and pair events created in the calorimeter (bottom center and right): Bremsstrahlung, annihilation photons, and electron-positron tracks can mimic legitimate Compton or pair events with no correlation to the original source photon's energy and origin

into *Compton events*, reconstructable *Pair events*, and *non-reconstructable events* (3-class classification), while restricting the evaluation of the pursuing improvement to energies up to 20 MeV.

1.3 Deep feedforward convolutional neural networks

A deep feedforward neural net defines a mapping

$$\mathbf{y} = f(\mathbf{x}, \Theta) \quad (1)$$

where \mathbf{x} is an input vector and Θ are the learned model parameters [12]. In the case of a classification task, \mathbf{y} is identified with a category. Feedforward means that data is processed inside the model only in one direction towards the output nodes.

Convolutional Neural Nets (CNNs) [13] emerged as the go-to processing approach when dealing with images and gridded data with local dependencies. The main idea is to restrict the search for relevant information to small regions of the input by applying a moving filter with *learned* kernel values on the pixel values. The number of free (trainable) parameters of a pure CNN is independent of the input size and lower than in any conventional Multi Layered Perceptron (MLP) approach since kernel sizes in general do not exceed $7 \times 7 = 49$ entries. This is much smaller than the actual number of pixels in an image, which would form the input for an MLP. The number of calculations necessary in a CNN depends on the size of the input. To scale down the processing overhead, a common strategy is to sample down the intermediate input by applying either maximum or average pooling. Here a moving window is applied to the input, propagating either the maximum or the average value of the pixels inside the window.

CNN-driven classification models are a synthesis of CNN and MLP. The task of the CNN is to extract key features of the image and deliver them to the classifier. The classifier has to comprehend the data presented by the CNN part of the model and decides to which class the sample belongs. In multi-class scenarios, the final output function is a classification vector whose components are generated by the softmax-function:

$$\text{softmax}(\mathbf{z})_i = \frac{\exp(z_i)}{\sum_j^N \exp(z_j)} = y_{i,\text{pred}} \quad (2)$$

Per definition, the sum over all entries equals to 1 and allows a probabilistic interpretation of the model's output. $y_{i,\text{pred}}$ is therefore the probability ($0 \leq y_{i,\text{pred}} \leq 1$) of a given sample to belong to the i -th class. The weight-adaptation is driven by the minimization of the categorical cross-entropy:

$$CE = - \sum_{i=1}^N y_{i,\text{true}} \log(y_{i,\text{pred}}) \quad (3)$$

where $y_{i,\text{true}}$ is 1 if the sample belongs to class i and 0 else.

2 Model architecture and training

2.1 Architecture

Convolutions can be applied either on 2D images or 3D volumes, which could be stacks of images forming a video or a binned 3D representation of a detector. A full 3D representation of the tracker module of e-ASTROGAM using only one 8-bit color channel for the deposited energy results in 3D images with a size of around one GB. This would lead to a very large training dataset and intense VRAM usage during training and inference. Therefore a multi-view approach was chosen. For optimal performance, each view is processed in a separate stack of convolutional modules. The result of each stack is then merged and further processed. The design of the tracker makes the use of X-Z, Y-Z and X-Y-projections of an event a natural choice. In addition to the tracker information (one row of pixels per layer), discretized calorimeter information (one row of pixels per centimeter of crystal height) is included at the bottom of the tracker projection. In order to avoid large memory and processing overhead, the X-Y projection was down-sampled by a factor 10.

This reduces the memory footprint of the input data from around 1 GB down to 0.71 MB per event using the same precision for the deposited energy.

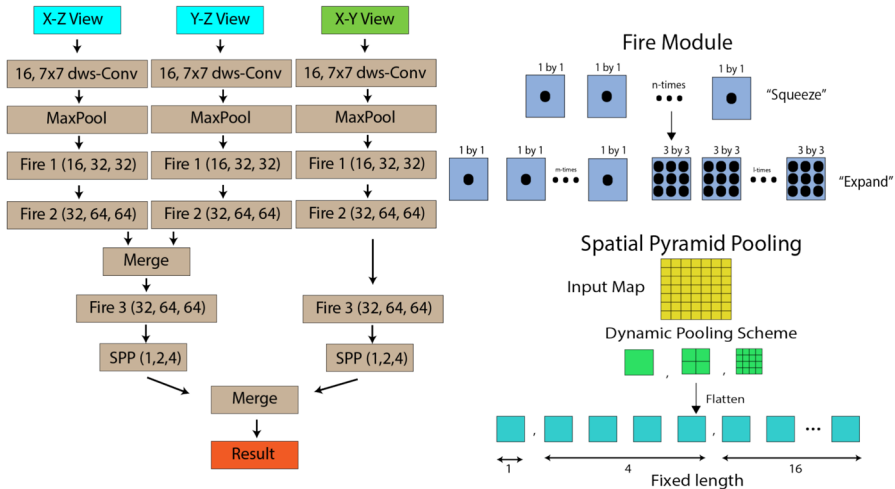


Fig. 3 Left: Sketch of the CNNCat model architecture. Each view is processed in a separate stack of convolutional (dws-Conv: depth-wise separable convolution) and pooling layers with identical layout. The results are merged and processed again. The final output is a classification vector whose components give the probability for the shown sample to be either a Compton, Pair or non-reconstructable event. Top Right: Sketch of the micro architecture of a SqueezeNet Fire Module. The number of input maps is reduced to n by the 1×1 "squeeze" filters that act as information bottle necks. The resulting data is processed by m 1×1 and 3×3 "expand" filters. The stacked output is handed to the next layer. Adapted from: [14]. Bottom right: Sketch of spatial pyramid pooling. The incoming feature maps are binned via pooling into 1×1 , 2×2 , 4×4 , ... representations, which are then flattened into the input of the dense layer. Here the length of the outgoing vector is fixed to $n \cdot (1 + 4 + 16)$, where n is the number of feature maps to process. Adapted from: [16]

The SqueezeNet architecture [14], enables the design of small, yet powerful models, which have been demonstrated to run on small-scale hardware like FPGA-accelerated System-On-Chips (SoCs) [15] and was chosen as design pattern of the convolutional modules. The presented model makes heavy use of SqueezeNet's Fire modules shown in the top right of Fig. 3 and is sketched in Fig. 3 left.

A common restriction in Feedforward Neural Nets in general lies in the fixed input size. In the case of CNNs, this restriction is self-imposed by using a dense layer to perform the classification. If used conventionally this layer takes the stack of final feature maps, flattens them, and uses the resulting 1D tensor as input. The number of free parameters introduced by the dense layer for processing is

$$\text{parameters} = \text{input length} \times \text{output nodes} \times \text{maps}$$

and can quickly reach millions because of the size and number of the feature maps to flatten. Due to their nature of connecting every input node with every output node, dense networks correlate each pixel in a feature map with every other pixel, which is also a source of faulty correlations either by pure chance or simulation biases. A more robust approach called Spatial Pyramid Pooling was introduced by [16]. The final feature maps are binned by applying Pooling with a set of dynamically sized windows. When flattened, this results in a fixed input size for the dense layer, regardless of the original image size. Each value now corresponds to regions of the processed feature map, effectively binning the image into $n \times n$ equally sized rectangles. This facilitates decoupling of the input size of the classifier (the dense layer) from the output of the feature generator (convolutional neural net). In addition, the size of the input vector of the dense layer is generally smaller than with conventional flattening. This reduces the number of free parameters. He et al. [16] also demonstrated that contemporary image classification networks using this approach outperform the designs with conventional flattening. In this study, max-pooling was used to generate the binned representation. Another application of this pooling scheme is to resize intermediate input with respect to each other. The feature maps of the X-Z filter tree have a different size with respect to the X-Z and Y-Z maps. This can be ramified by applying the dynamical pooling scheme in order to generate compatible input for the dense layer.

One go-to method for regularizing neural networks (i.e., to prevent overfitting) is Dropout [17]. During training in layers with Dropout applied (usually the dense layer or the first few convolutional layers), randomly selected neurons are switched off for each training step. To compensate, the model has to learn to propagate vital information for the classification over various paths in the graph. During inference, no Dropout is applied, i.e., the full model is used. Using weight adaptation only during training results in a more robust model.

Table 1 Data binning used for dataset balancing and simulation management

Type	Low energy	Medium energy	High energy	Θ
Compton	0.1 - 1 MeV	1 - 5 MeV	5-10 , 10-15, >15	0-180
Pair	10-200 MeV	0.2 - 1 GeV	1 - 3.5 GeV	0-180
Electron	10 - 100 MeV	0.1 - 1 GeV	1 - 10 GeV	90-180

The CNNCat model was implemented using TensorFlow [18] Version 2 with TensorFlow's internal keras implementation. Spatial Pyramid Pooling was implemented with a modified version of [19].

2.2 Dataset generation and selection criteria

The training and test data generation is based on MEGALib [11] simulations with the mass model shown on the right in Fig. 1. The simulation output was then processed into X-Y, X-Z and Y-Z images using custom-written python scripts. Each Z-projection has a size of 66 rows by 4267 columns with 3 color channels. The X-Y projection has a size of 427 by 427 and also 3 channels. Channel 1 contains the binned deposited energy in the pixel saturating at 2048 keV, channel 2 contains the deposited energy relative to the total energy deposited in the tracker, while channel 3 contains the deposited energy relative to the total energy measured in the calorimeter. If the calorimeter has no deposited energy, the values of channel 3 is set to 0. Since the event topology changes with incident angle and energy, the simulations were split as shown in Table 1. This allowed for a controlled dataset composition decoupled as much as possible from detector efficiencies.

In addition to enforcing the absence of a veto trigger, we restrict the training and test data to events that can be reconstructed in principle, as shown in Fig. 4. In order to reconstruct a Compton sequence one needs at least two hits. Since the position and energy resolution in the calorimeter is lower in general, a coincidence of at least two individual tracker layers or one tracker layer and the calorimeter is required for Compton events included in the training set. Pair events require at least three resolvable hits.

In this study, we use the following definitions of event types:

- *Compton event*: The first interaction in the sensitive part of the detector is a Compton scatter and the veto was not triggered (Fig. 2, top left and center).
- *Pair event*: Electron and positron interacted inside the sensitive part of D1 first and the veto was not triggered (Fig. 2, top right)
- *Non-reconstructable*: Events that cannot be reconstructed meaningfully. Either the event is resulting from an electron primary traversing the geometry without triggering the veto, or a pair was created inside D2 and did not trigger the veto. D2-first pairs are a case of lost source photons, events triggered by electron primaries are a genuine background component (Fig. 2, bottom row).




Fig. 4 Enforced trigger pattern to ensure a minimum amount of reconstructable information in each event. Compton events have to fulfill at least 1+1 or 2+0, pair events have to fulfill at least 3+0 or 2+1

It was also deemed important to enforce clean samples from which the model is tasked to learn the characteristics of good Compton and Pair events. The very first interaction of a photon in each training set sample, independent of the type of event, has to occur in sensitive detector material. Pair events have to be created inside the tracker. Since vetoed events in space-borne detectors are usually discarded immediately, they are not considered in the datasets of this study. For image augmentation, we applied Gaussian noise with the estimated energy resolution of the proposed detector to the input images to further solidify the model against overfitting.¹

3 Results

We start the presentation of the results with an overview of the computational demands, introduce performance measures used throughout the evaluation, followed by a performance analysis on two test sets. Test set 1 comprises the full observational range from 100 keV to 3.5 GeV from all incidence angles. Test set 2 considers a simulated Crab observation.

3.1 Computational demands and architectural optimizations

Since one of the possible applications should be an online data classifier used on the satellite itself, all the design choices were driven toward optimizing the number of computations per event. The choice of Fire-Modules as micro-architecture for the convolutional modules allows us to control the size and hence the computational demands of the model, due to their built-in bottlenecks. Dropout as a regularization method introduces no computational overhead at run-time while effectively constraining the model and preventing overfitting of the training data.

To further optimize the number of computations per convolutional layer, depth-wise separable convolution (DWS-Conv) [20] was applied wherever feasible, i.e., on the 7×7 and 3×3 convolutions².

A conventional convolutional layer applies $m \times n$ convolutions on l incoming channels on an image with $j \times k$ pixels, where m denotes the number of filters to be applied and n is the filter size along one axis in pixels. Assuming $m = 16$, $n = 3$, $l = 3$ and $j = k = 100$, this setup would result in

$$\text{parameters} = n^2 \times l \times m = 432 \quad (4)$$

and

$$\text{ops} = j \times k \times n^2 \times l \times m = 7 \times 10^6 \quad (5)$$

operations used in this layer.

¹ Since the D1 detector is highly inhomogeneous due to its layered design, the common technique of rotating the image cannot be applied.

² This speed-up may not be realized with non-optimized TensorFlow models since the underlying libraries (e.g. CuDNN) may not be optimized for this convolution type. Running the model through an inference engine like tensorRT offered by NVIDIA will show the expected behaviour.

Depth-wise separable convolution, on the other hand, tries to simplify this step without sacrificing too much performance. In this scheme, only one full $n \times n$ convolution is applied to each incoming channel. The intermediate data is put back together and then stretched to the required dimension by applying m 1×1 convolutions. In the setup described above, this results in

$$\text{parameters} = n^2 \times l + l \times m = 75 \quad (6)$$

and

$$\text{ops} = n^2 \times l \times j \times k + l \times j \times k \times l \times m = 1.2 \times 10^6. \quad (7)$$

Applying this to the precursor model allowed us to reduce the number of trainable parameters to 64k and the number of FLOPS per event to 0.9 GFLOPS. ESA identified the necessity to explore systems for onboard processing and funded the GPU4S (GPUs for Space) project that surveyed use-cases and possible embedded systems [21]. Follow-up studies also investigated potential improvements of commercial off-the-shelf embedded systems (e.g. [22]). Ortiz et al. also studied the Jetson AGX system which delivers around 1.3 TFLOPS per Watt at FP16 or 2.6 TOPS at INT8, equivalent to a rate of 1.4 kHz or 2.9 kHz per Watt, respectively. Note that either a smaller detector or a coarser binning of x-y strips in the input data would increase the amount of rate per Watt significantly as the computational complexity for CNNs scales mostly with $j \times k$ of the input data. A more recent study of Rodriguez-Ferrandez et al in cooperation with ESA was able to show that this kind of devices can sustain the radiation environment in orbit [23]. They obtained acceptable error rates, making them a candidate system for embedded systems especially designed for space applications.

3.2 Performance measures

The outcome of each binary event classification falls into one of the following four cases:

- (tp) True positive: the sample belongs to the positive class and is classified correctly (positive).
- (tn) True negative: the sample belongs to the negative class and is classified correctly (negative).
- (fn) False negative: the sample belongs to the positive class and is classified as negative.
- (fp) False positive: the sample belongs to the negative class and is classified as positive.

From this, several metrics can be derived, each bound between 0 and 1, with 1 being a perfect performance under the given metric. The most common one is *accuracy*, defined as

$$\text{Accuracy} = \frac{tp + tn}{tp + tn + fn + fp}. \quad (8)$$

It gives the fraction of correct predictions relative to all predictions. Accuracy is a workable metric on class-wise balanced datasets but may give deceptive results when the data is skewed towards one class.

On unbalanced datasets, performance must be measured with more robust metrics. Two of those, from the context of binary classification and information retrieval, are precision and recall. *Precision* is defined as

$$\text{Precision} = \frac{tp}{tp + fp} \quad (9)$$

and is a measure of how reliable a classification into a certain class is, i.e., it is a measure of the purity of the events in that class. It can be influenced by the classifier's performance on the samples of a class itself or by misclassification of samples belonging to other classes.

The other metric is the *recall* on a given class. It is defined as:

$$\text{Recall} = \frac{tp}{tp + fn}. \quad (10)$$

Recall, also called true positive rate, is a measure of how well a given class is recognized by a classifier, i.e., the efficiency with which these events are properly classified. In general, it is impossible to optimize both precision and recall to arbitrarily high values unless the samples in the dataset are perfectly separable.

Both metrics can be extended from binary to multi-class scenarios by applying a one-against-all scheme. The class to be evaluated is considered the positive class, while the other classes are combined into the negative class. False positives then have to be defined as the combined misclassifications into the investigated class from all other classes.




Fig. 5 Training Loss as function of epochs. Training was started with Learn Rate I. After epoch 7 early stopping terminated training run I. The model was reset to the previous best weights. The training was then restarted with Learn Rate II for the remaining 34 epochs

Table 2 Best training parameters found after optimization

Learn Rate I	Learn Rate II	Dropout	Batch Size	SPP	Delta	Patience
10^{-4}	10^{-6}	0.3	1	[1,2,4]	0.001	5

Delta gives the minimum improvement from the previous best loss-value, patience is the number of epochs in which the improvement has to be achieved

3.3 Training and validation

The training set consisted of 400,000 events evenly divided over classes, incident energy and incident angle binning. 300k stratified events were used for training, 100k similar selected events were held back for validation for the purpose of monitoring the training progress. We used the Adam optimizer [24] for parameter adaptation. Early stopping was used to prevent overfitting due to excessive training. During trials we observed that the improvement stalled after only a couple of epochs. To mitigate this, early stopping was used to trigger a restart of the optimizer with a lower learn-rate on the best-performing model (denoted as Training Reset) in order to force the optimizer to a lower minimum. The progress during training is shown in Fig. 5. The relevant training and hyper parameters are summarized in Table 2

Training for the best-performing model took 39 epochs with a total duration of around 9 days on an NVIDIA GeForce RTX 2080Ti graphics card. The model was reset after early stopping to the best-performing weights after epoch 34, which were then used for evaluation against the test sets.

3.4 Classification performance

The test set contains 2,562,406 events containing photons and electrons ranging from 100 keV to 3.5 GeV simulated from a flat spectrum uniformly from all incident angles. The energy range was divided analogously to the training-set generation. This dataset is reserved to analyze the energy dependency of the classification result averaged over all incident angle bins and was not used during training or validation.

We evaluate the performance of CNNCat over the full observation range of the instrument using the average over all directions as a function of incident energy of recall, precision, and class-wise false negative rates. The resulting plots are energy-dependent representations of rows of the confusion matrix, shown in Fig. 6a – c. On the input side, we further distinguish Compton events of three types: fully absorbed (golden) events, non-fully absorbed events starting in sensitive material (“genuine” Compton events - potentially reconstructable), and non-fully absorbed events starting in passive material (“non-genuine” Compton events), which, effectively, are source photons turned into background, as the initial scatter angle and energy loss is unknown.

Figure 6a shows the model performance on Compton events. The recall of all Compton subsets is very high (above 90 %) for incident energies up to 4 MeV. The recall of non-fully absorbed genuine Compton events steadily declines over the remaining energy range, while it stabilizes at 60 % at incident energies between 7 to 20 MeV for fully absorbed Compton events. Non-fully non genuine event are less accurately

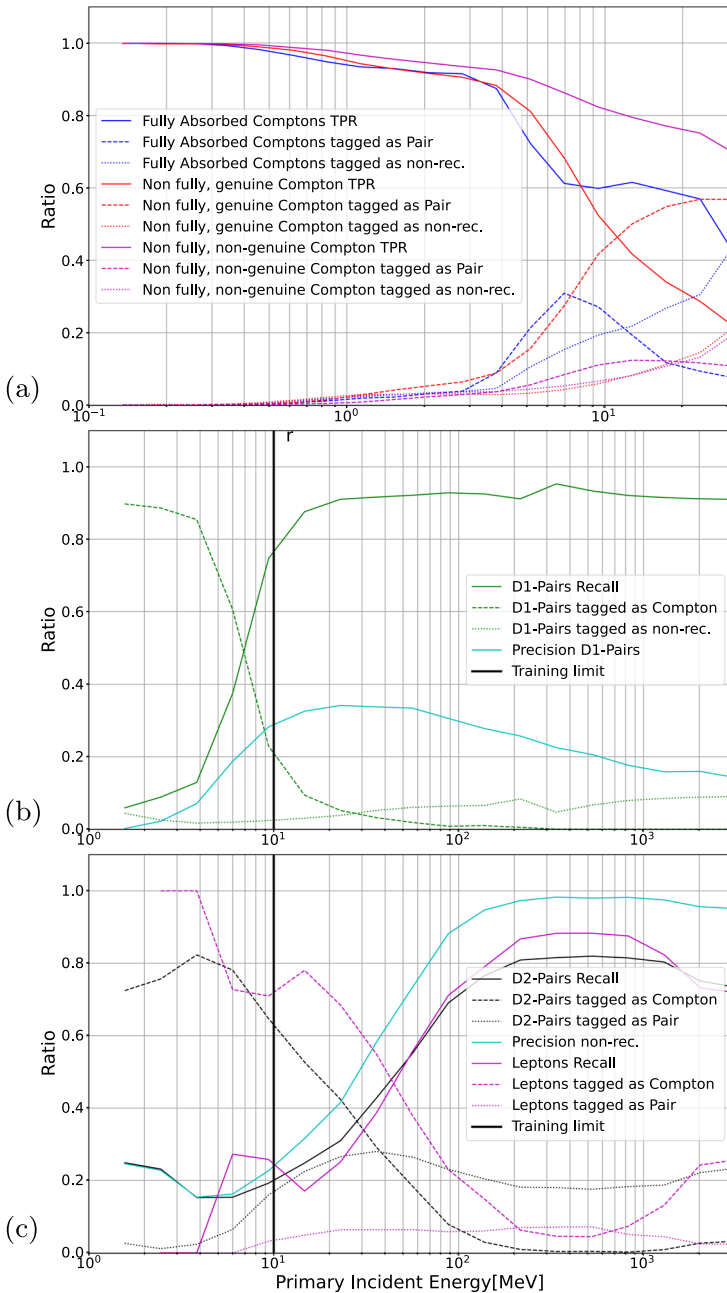


Fig. 6 Precision, recall and class-wise false negative rates vs. incident energy.

(a) Compton events. Recall for the three sub-categories of Compton events and corresponding false negative tags vs. incident energy up to 30 MeV.

(b) D1-first pair events. The model was trained on events with $E_{inc} > 10$ MeV, denoted by the bar.

(c) Non-reconstructable events. Misclassified events are mostly tagged as Compton for $E < 40$ MeV and as Pair above. The model was trained on events with $E_{inc} > 10$ MeV, denoted by the bar

tagged with increasing incident energy but stay above 70% over the shown energy range. Mistagged fully absorbed Compton events are mostly classified as Pair event at energies around 3 to 12 MeV. Beyond that, the majority of false negatives is tagged as non-reconstructable. The vast majority of false negative non-fully absorbed genuine Compton events are wrongly tagged as Pair. This can be explained by the very narrow scatter angle distribution at higher energies, which will give photon hits a topology similar to tracks which together with a recoil electron track can confuse the model. Another possibility is the generation of multiple recoil electron tracks before the photon leaving the detector. Non-genuine non-fully absorbed Compton events show a behaviour similar to fully absorbed Compton events, but shifted to higher energies. Here the turnover from Pair to non-rec. tagging occurs around 18 MeV.

Figure 6b shows the model performance on pair events with their first interaction measured in D1. The recall on pair events at incident energies below 7 MeV is very low. This is not surprising since the model only had pair events with incident energies larger than 10 MeV in the training data. Electrons and positrons in this energy range are very susceptible to multiple scattering in solid material and show a topology similar to recoil electrons in Compton events. Also, positrons are generating shorter tracks, leading to even greater similarities between Compton and pair events. At 10 MeV, however, the recall is already passing 78%. For energies higher than 100 MeV the recall saturates at 92%. Precision is in general low due to the large amount of false positives from Compton events at low energies and D2-first pair events over the whole energy range.

Figure 6c shows the model performance on non-reconstructable events. The recall on D2-first pair events is rising to around 80% at around 100 MeV leptonic events are recognized reliably from around 70 MeV and higher with a peak recall of 92% around 200 MeV. In the energy range up to 50 MeV, the majority of misclassified events is tagged as Compton. Over the remaining range, most of the misclassified events are tagged as pair. Overall precision for energies higher than 100 MeV peaks at 92%, leaving a relatively small number of potentially reconstructable events that were falsely rejected.

Figure 7 summarizes the recall vs. incident energy for all classes, distinguishing the three Compton event populations, D1- and D2-first pairs, and lepton events.

Figure 8 shows the ratio of D1-first ('good') to D2-first ('bad') pairs over the observed energy range. Since e-ASTROGAM was designed to measure both Compton and pair events, the detector design had to compromise between the demands of both detection channels. For instance, the tracker layers lack tungsten converter foil, which enhances the creation of pairs in Fermi-LAT, since this would severely hamper the measurement of (low-energy) Compton events. As a result, many high-energy source photons traverse the tracker without interaction and create pairs only in the calorimeter as indicated by the ratio before classification. This skews the ratio of D1-first to D2-pairs heavily towards non-reconstructable events that are potentially dangerous. While initially above 1, the ratio soon drops to values below 0.1, manifesting the large surplus of D2-first pairs. This ratio greatly improves after classification for incident energies greater than 15 MeV. Above 20 MeV, i.e., in the intended observation range in pair-imaging mode, the model greatly improves the 'good'/'bad' pair ratio by up to a factor 8, shown by the blue line. The weak performance at energies below 10 MeV is

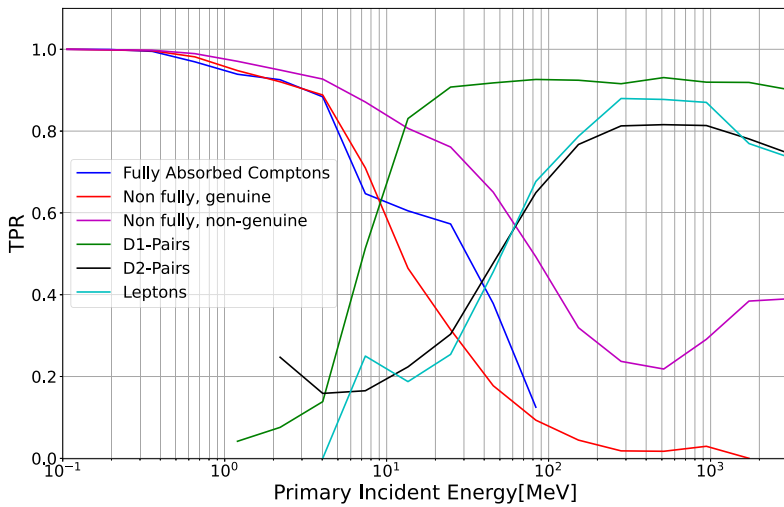


Fig. 7 Recall vs. incident energy for the three Compton event populations, D1- and D2-first pairs, and lepton events

tolerable, since, at these low energies, pair events are difficult to reconstruct and are dominated by Compton events.

Overall, one can conclude from this test set that a small-scale model is capable of distinguishing the presented events by type. As expected, the decision to classify an event as either Compton or pair is biased to some degree by the energy the model can




Fig. 8 Ratio of D1-first to D2-first pair events vs. incident photon energy. At lower energies (below 15 MeV) the effect of the classification on the ratio is negative. This is because of the lower recall of the model on pair events in this energy region. For higher energies, the model is able to improve the ratio of D1-first vs. D2-first pairs by a factor ranging up to 8 at 1000 MeV




Fig. 9 Directionally averaged ROC curves for each class for the test set. The performance of a random classifier corresponds to the dashed line. The Area Under the Curves are given in the legend. Classification of pair events performs best, with an AUC-score of 0.9. The other classes are still performing very well as indicated by scores larger than 0.8

see in an event. This can be seen from the fact that almost no events with an incident energy below 1 MeV are classified as pair event. The classification has a beneficial effect on the ratio of D1-first pairs to D2-first pairs for all but the lowest energies.

To assess the total performance of our classifier, we calculated the Receiver-Operator-Characteristic Area-under-the Curve (ROC-AUC) by determining the True Positive Rate and False Positive Rate for decreasing decision thresholds for the positive class. In multi-class scenarios this is done in a so called ‘one vs. all’ approach.

The resulting curves for the test dataset are shown in Fig. 9. All three classes are well separated from random guessing (indicated by the dashed line). There is no task-independent interpretation of the AUC-score, but one can employ a rule of thumb found in [25].

The scores for each class are around a value of 0.8, which indicates a good performance on the presented dataset. The best performance is achieved by the pair classification, which achieves an area of 0.9.

Following the scheme of [25], our model reaches the second highest performance category.

3.5 Crab nebula simulation

In this part, we compare the model classification to reconstructed data up to 20 MeV by applying the standard MEGALib reconstruction chain to a simulated Crab measurement. The Crab spectrum was modelled according to ref. [26] for energies up to 6 MeV, and extended to 20 MeV by a power-law with index 2.24, as shown in Fig. 10. The Crab simulation was performed in eight separate energy ranges. For each bin, a simulation




Fig. 10 Input spectra for the Crab-Nebula dataset. The Crab spectrum uses ref. [26] at $E < 6\text{ MeV}$ and a power-law with index 2.24 above. The simulation range is indicated by the solid line

with 500k pre-trigger events was performed, resulting in a total of 2.85 M events with a distribution shown in Table 3. The source was set up as a far-field point source with an incident angle $\theta = 30^\circ$. MEGAlib was tasked with reconstructing Compton and Pair events and eliminating events from minimum ionizing particle tracks (MIPs). It applies a cascade of reconstruction algorithms with increasing computational complexity to each event. If at any stage the algorithm is successful, the chain stops and the event is classified accordingly. The MIP-search tests the hit-pattern for linearity by applying a linear χ^2 fit. The next stage is the Pair-reconstruction. Here a pattern-recognition algorithm is searching for a creation vertex. If successful, the electron and positron tracks are reconstructed by applying a Pearson correlation test. The Compton reconstruction is separated into two stages. First, the recoil electron track is extracted and reconstructed based on a Spearman-Rank test. Then the Compton-sequence is reconstructed by minimizing the χ^2 of the difference between the kinematic scatter angle obtained via Compton's formula and the geometric scatter angle derived from the hits' locations. Details about the implementation in MEGAlib can be found in [27].

We use the result of the reconstruction as prediction of the event type of the baseline classification. We considered an event tagged as Compton when the reconstruction succeeded in Compton reconstruction. The event was considered tagged as pair, when

Table 3 Upper limit of the simulation bins, medians and contributing events per bin

E_{max} [MeV]	0.13	0.26	0.45	1.1	2.4	4.7	10	20
Median [MeV]	0.07	0.16	0.31	0.56	1.35	2.85	5.72	12.03
Events	106k	289k	378k	415k	422k	422k	413k	407k

The minimum photon energy was set to 60 keV. The total size of the dataset is 2.85 million events

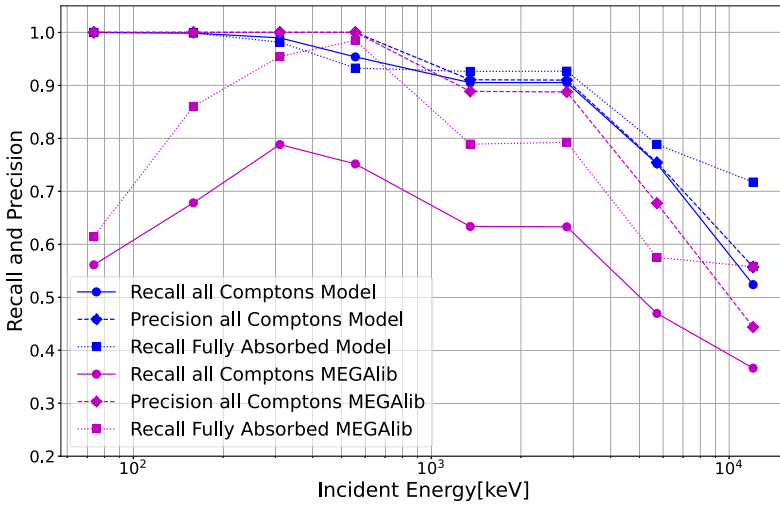


Fig. 11 Precision and Recall for Compton events as function of energy for the Crab-source dataset

the pair reconstruction accepted the event. An event was considered tagged as non-reconstructable when it was either tagged as MIP or all reconstruction algorithms rejected the event.

The performance on the Compton dataset is shown in Fig. 11. The default reconstruction tags up to 80% of all Compton events correctly at its peak around 550 keV and then drops off steadily at higher energies. The classification by our model outperforms the default classification, giving significantly higher recall values for each energy bin of the crab spectrum on the full Compton dataset. The precision on the




Fig. 12 Precision and Recall for Pair events as function of energy for the Crab-source dataset




Fig. 13 Precision and Recall for non-reconstructable events as function of energy for the Crab-source dataset

low-energy bins is nearly perfect in both classifications with a slight edge at higher energies for the model classification. An important subset constitute fully absorbed Compton events, which have the highest likelihood to be correctly reconstructed and therefore contain the proper directional information of the event. Here the model classification outperforms the default classification in every but one bin at 550 keV (93 % vs 98 %) by up to 38 percent-points.

The performance on D1 and D2 pairs is depicted in Figs. 12 and 13. At energies above the pair threshold and below 5 MeV, both the model and the default classification show very little discrimination power. Nonetheless, the model's Recall on D2-first pairs shows an improvement over the default case. The D1-pair Recall shows large improvements over the default classification, starting to recognize D1-pairs reliably at energies larger than 5 MeV and a very good class-coverage of 87 % at the highest simulated energies. The Precision of the model classifications on D1 and D2-pairs shows a slightly worse purity of the pair datasets after classification than in the default case, which can be explained by the relatively large amount of incorrectly tagged D2-pairs which have been tagged as D1-pair and vice-versa (see Fig. 6a, b and c).

Figure 14 shows the relative improvement of the model classification over the default case, given as the ratio of model performance relative to the default. Data points above 1 indicate improvement, below 1 indicate worse performance than the default. The recall on all Comptons, D1 and D2 pairs shows an overall improvement of around 25% to 60% over the default classification. In the low energy bins, the improvement ratio for D1-pairs reaches beyond a factor 10 while maintaining an improvement ratio of greater than 1.5 for high energies. Even though the recall on D2-pairs in this energy range is still sub-optimal, the model's performance is an improvement over the default classification over the whole inspected energy range, improving the class coverage by a factor 3 for lower energies to a factor of 1.3 at the higher energies. The observed

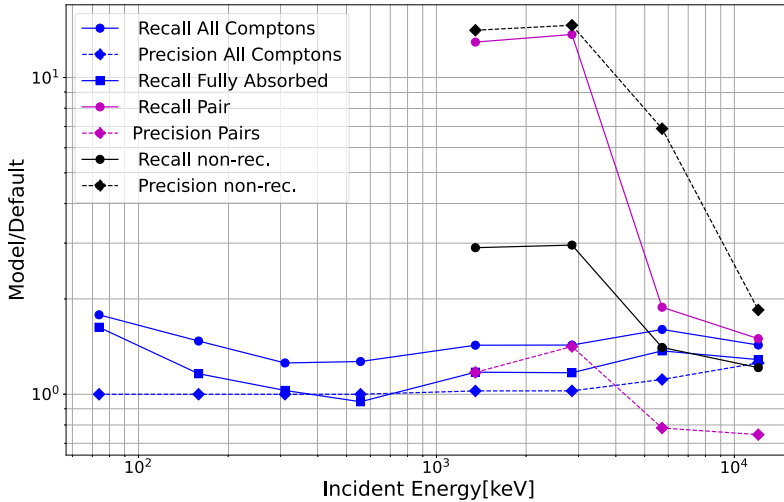


Fig. 14 Relative improvement of the model classification vs the default reconstruction as function of energy for the Crab-source only dataset

behaviour of the model classification is consistent with the expectations from the omni-directional test set.

4 Summary and outlook

The primary goal of this work was to explore the possibility to detect event types and event topologies with CNN-driven models. The successful model presented here is capable of distinguishing different event types and undesired event topologies, as shown on an omni-directional test set as well as on a simulated Crab measurement.

This model was trained to distinguish Compton, pair and non-reconstructable events. The latter consisted of signal photons with a non-reconstructable event topology and events caused by leptons, like cosmic and albedo electrons and positrons. Of course there are even more types of background events, like events resulting from delayed decays of activated materials, and events caused by background hadrons. These components will be the subject of a follow-up study.

The presented model has a computational overhead of about 0.9 GOPS per event, which fits well into the computational envelope of recently developed embedded systems. An architecture like the presented one is capable of processing even larger rates for smaller inputs. Studies are under way to test these embedded systems, making an online-approach viable, e.g., for CubeSat scenarios.

One particular aspect of this architecture remains unexplored. Since SPP enables the use of dynamic image sizes, it is possible to decrease the number of necessary operations per event even further by cropping the input images to a smaller size. For computational estimates, the average image size then becomes a variable of the

measurement environment. Such a study would make sense once a mission is funded and various orbit scenarios have to be explored.

Acknowledgements Jan Peter Lommler acknowledges support by DLR grant 500O2218.

Author Contributions J.L.: Main manuscript, review of the manuscript U.O.: Advisory, review of the manuscript

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Winkler, C., Courvoisier, T.J.-L., Di Cocco, G., Gehrels, N., Giménez, A., Grebenev, S., Hermesen, W., Mas-Hesse, J.M., Lebrun, F., Lund, N., Palumbo, G.G.C., Paul, J., Roques, J.-P., Schnopper, H., Schönfelder, V., Sunyaev, R., Teegarden, B., Ubertini, P., Vedrenne, G., Dean, A.J.: The INTEGRAL mission. *A&A* **411**, 1–6 (2003). <https://doi.org/10.1051/0004-6361:20031288>
2. Atwood, W.B., et al.: The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission. *Astrophys. Lett.* **697**(2), 1071–1102 (2009). <https://doi.org/10.1088/0004-637x/697/2/1071>
3. Aharonian, F., Akhperjanian, A.G., Bazer-Bachi, A.R., Beilicke, M., Benbow, W., Berge, D., Bernlöhr, K., Boisson, C., Bolz, O., Borrel, V., Braun, I., Breitling, F., Brown, A.M., Bühler, R., Büsching, I., Carrigan, S., Chadwick, P.M., Chounet, L.-M., Cornils, R., Costamante, L., Degrange, B., Dickinson, H.J., Djannati-Ataï, A., O’C. Drury, L., Dubus, G., Egberts, K., Emmanoulopoulos, D., Espigat, P., Feinstein, F., Ferrero, E., Fiasson, A., Fontaine, G., Funk, S., Funk, S., Gallant, Y.A., Giebels, B., Glicenstein, J.F., Goret, P., Hadjichristidis, C., Hauser, D., Hauser, M., Heinzlmann, G., Henri, G., Hermann, G., Hinton, J.A., Hofmann, W., Holleran, M., Horns, D., Jacholkowska, A., de Jager, O.C., Khélifi, B., Komin, N., Konopelko, A., Kosack, K., Latham, I.J., Le Gallou, R., Lemièrre, A., Lemoine-Goumard, M., Lohse, T., Martin, J.M., Martineau-Huynh, O., Marcowith, A., Masterson, C., McComb, T.J.L., de Naurois, M., Nedbal, D., Nolan, S.J., Noutsos, A., Orford, K.J., Osborne, J.L., Ouchrif, M., Panter, M., Pelletier, G., Pita, S., Pühlhofer, G., Punch, M., Raubenheimer, B.C., Raue, M., Rayner, S.M., Reimer, A., Reimer, O., Ripken, J., Rob, L., Rolland, L., Rowell, G., Sahakian, V., Saugé, L., Schlenker, S., Schlickeiser, R., Schwanke, U., Sol, H., Spangler, D., Spanier, F., Steenkamp, R., Stegmann, C., Superina, G., Tavernet, J.-P., Terrier, R., Théoret, C.G., Tluczykont, M., van Eldik, C., Vasileiadis, G., Venter, C., Vincent, P., Völk, H.J., Wagner, S.J., Ward, M.: Observations of the Crab nebula with HESS. *Astron. & Astrophys.* **457**(3), 899–915 (2006). [arXiv:astro-ph/0607333](https://doi.org/10.1051/0004-6361:20065351) [astro-ph]. <https://doi.org/10.1051/0004-6361:20065351>
4. Weekes, T.C., Badran, H., Biller, S.D., Bond, I., Bradbury, S., Buckley, J., Carter-Lewis, D., Catanese, M., Criswell, S., Cui, W., Dowkontt, P., Duke, C., Fegan, D.J., Finley, J., Fortson, L., Gaidos, J., Gillanders, G.H., Grindlay, J., Hall, T.A., Harris, K., Hillas, A.M., Kaaret, P., Kertzman, M., Kieda,

- D., Krennrich, F., Lang, M.J., LeBohec, S., Lessard, R., Lloyd-Evans, J., Knapp, J., McKernan, B., McEnery, J., Moriarty, P., Muller, D., Ogdén, P., Ong, R., Petry, D., Quinn, J., Reay, N.W., Reynolds, P.T., Rose, J., Salamon, M., Sembroski, G., Sidwell, R., Slane, P., Stanton, N., Swordy, S.P., Vassiliev, V.V., Wakely, S.P.: Veritas: the very energetic radiation imaging telescope array system. *Astropart. Phys.* **17**(2), 221–243 (2002). [https://doi.org/10.1016/S0927-6505\(01\)00152-9](https://doi.org/10.1016/S0927-6505(01)00152-9)
5. Ferenc, D.: The magic gamma-ray observatory. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Assoc. Equip.* **553**(1), 274–281 (2005). <https://doi.org/10.1016/j.nima.2005.08.085>. Proceedings of the fifth International Workshop on Ring Imaging Detectors
 6. Introducing the cta concept: CTA Consortium. *Astropart. Phys.* **43**, 3–18 (2013). <https://doi.org/10.1016/j.astropartphys.2013.01.007>. Seeing the High-Energy Universe with the Cherenkov Telescope Array - The Science Explored with the CTA
 7. Schoenfelder, V., Aarts, H., Bennett, K., de Boer, H., Clear, J., Collmar, W., Connors, A., Deerenberg, A., Diehl, R., von Dordrecht, A., den Herder, J.W., Hermsen, W., Kippen, M., Kuiper, L., Lichti, G., Lockwood, J., Macri, J., McConnell, M., Morris, D., Much, R., Ryan, J., Simpson, G., Snelling, M., Stacy, G., Steinle, H., Strong, A., Swanenburg, B.N., Taylor, B., de Vries, C., Winkler, C.: Instrument description and performance of the imaging Gamma-Ray telescope COMPTEL aboard the compton Gamma-Ray observatory. *ApJS* **86**, 657 (1993). <https://doi.org/10.1086/191794>
 8. De Angelis, A., Tatischeff, V., Tavani, M., et al: The e-ASTROGAM mission. Exploring the extreme Universe with gamma rays in the MeV - GeV range. *Exp. Astron.* **44**(1), 25–82 (2017) [arXiv:1611.02232](https://arxiv.org/abs/1611.02232) [astro-ph.HE]. <https://doi.org/10.1007/s10686-017-9533-6>
 9. Tatischeff, V., De Angelis, A., Tavani, M., et al: All-Sky-ASTROGAM: The MeV Gamma-Ray companion to multimessenger astronomy. *arXiv* (2019). <https://doi.org/10.48550/ARXIV.1905.07806>. ssss
 10. McEnery, J., et al: All-sky medium energy Gamma-ray observatory: exploring the extreme multimessenger universe. In: *Bulletin of the american astronomical society*, vol. 51, pp. 245 (2019)
 11. Zoglauer, A., Andritschke, R., Schopper, F.: MEGALib the medium energy Gamma-ray astronomy library. *New Astron. Rev.* **50**(7–8), 629–632 (2006). <https://doi.org/10.1016/j.newar.2006.06.049>
 12. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, (2016). <http://www.deeplearningbook.org>
 13. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998). <https://doi.org/10.1109/5.726791>
 14. Iandola, F.N., Moskewicz, M.W., Ashraf, K., Han, S., Dally, W.J., Keutzer, K.: Squeezenet: Alexnet-level vand <1mb model size. *CoRR* (2016). [arXiv:1602.07360](https://arxiv.org/abs/1602.07360)
 15. Gschwend, D.: Zynqnet: An fpga-accelerated embedded convolutional neural network. *ArXiv. arXiv:2005.06892* (2020)
 16. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. *CoRR* (2014). [arXiv:1406.4729](https://arxiv.org/abs/1406.4729)
 17. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014)
 18. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org (2015). <https://www.tensorflow.org/>
 19. Henon, Y.: keras-spp. <https://github.com/yhenon/keras-spp> (2017)
 20. Sifre, L.: Rigid-motion scattering for image classification. Phd thesis, Ecole polytechnique, CMAP (2014)
 21. Kosmidis, L., Rodriguez, I., Álvaro Jover, Alcaide, S., Lachaize, J., Abella, J., Notebaert, O., Cazorla, F.J., Steenari, D.: Gpu4s: Embedded gpus in space - latest project updates. *Microprocess. Microsyst.* **77**, 103143 (2020). <https://doi.org/10.1016/j.micpro.2020.103143>
 22. Ortiz, F., Monzon Baeza, V., Garcés-Socarras, L.M., Vázquez-Peralvo, J.A., Gonzalez, J.L., Fontanesi, G., Lagunas, E., Querol, J., Chatzinotas, S.: Onboard processing in satellite communications using ai accelerators. *Aerospace* **10**(2) (2023). 10.3390/aerospace10020101
 23. Rodriguez-Ferrandez, I., Tali, M., Kosmidis, L., Rovituro, M., Steenari, D.: Sources of single event effects in the nvidia xavier soc family under proton irradiation. In: 2022 IEEE 28th international

- symposium on on-line testing and robust system design (IOLTS), pp. 1–7 (2022). <https://doi.org/10.1109/IOLTS56730.2022.9897236>
24. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)
 25. Hosmer, D.W., Lemeshow, S.: Applied Logistic Regression, 3rd edn. Wiley-Interscience Publication, (2013)
 26. Jourdain, E., Roques, J.P.: The high-energy emission of the crab nebula from 20 keV to 6 MeV with integral SPI*. *Astrophys. J.* **704**(1), 17 (2009). <https://doi.org/10.1088/0004-637X/704/1/17>
 27. Zoglauer, A.C.: First light for the next generation of Compton and pair telescopes: development of new techniques for the data analysis of combined Compton and pair telescopes and their application to the MEGA prototype. PhD thesis, - (2006)
 28. Aggarwal, C.: Neural Networks and Deep Learning: A Textbook, 1st edn. Springer, (2018)
 29. Japkowicz, N., Shah, M.: Evaluating learning algorithms: a classification perspective. Cambridge University Press (2011). <https://doi.org/10.1017/CBO9780511921803>
 30. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR) (2017)
 31. De Angelis, A., Mallamaci, M.: Gamma-Ray Astrophysics. *Eur. Phys. J. Plus* **133**, 324 (2018). <https://doi.org/10.1140/epjp/i2018-12181-0>. [arXiv:1805.05642](https://arxiv.org/abs/1805.05642) [astro-ph.HE]
 32. Cumani, P., Hernanz, M., Kiener, J., Tatischeff, V., Zoglauer, A.: Background for a gamma-ray satellite on a low-earth orbit. *Exp. Astron.* **47**, 1–30 (2019). <https://doi.org/10.1007/s10686-019-09624-0>
 33. Zoglauer, A., Boggs, S.E., Galloway, M., Amman, M., Luke, P.N., Marc Kippen, R.: Design, implementation, and optimization of MeGaLib's image reconstruction tool mimrec. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. **652**(1), 568–571 (2011). <https://doi.org/10.1016/j.nima.2010.08.043>. Symposium on Radiation Measurements and Applications (SORMA) XII 2010
 34. Chollet, F.: keras. GitHub (2015)
 35. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR* (2017) [arXiv:1704.04861](https://arxiv.org/abs/1704.04861)
 36. Abbott, B.P., Abbott, R., et al, T.D.A.: Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817a. *Astrophys. J.* **848**(2), 13 (2017). <https://doi.org/10.3847/2041-8213/aa920c>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.