








ComptonUNet: A Deep Learning Model for Gamma-Ray Burst Localization with Compton Cameras under Noisy and Low-statistic Conditions

Shogo Sato^{1,2} , Kazuo Tanaka¹ , Shojun Ogasawara¹, Kazuki Yamamoto¹, Kazuhiko Murasaki² , Ryuichi Tanida² , and Jun Kataoka¹ 

¹ Faculty of Science and Engineering, Waseda University, Japan; shg.sato@ntt.com

² NTT, Inc.

Received 2025 July 3; revised 2026 January 30; accepted 2026 February 3; published 2026 March 6

Abstract

Gamma-ray bursts (GRBs) are among the most energetic transient phenomena in the Universe and serve as powerful probes for high-energy astrophysical processes. In particular, faint GRBs originating from a distant Universe may provide unique insights into the early stages of star formation. However, detecting and localizing such weak sources remains challenging owing to low photon statistics and substantial background noise. Although recent machine learning models address individual aspects of these challenges, they often struggle to balance the trade-off between statistical robustness and noise suppression. Consequently, we propose ComptonUNet, a hybrid deep learning framework that jointly processes raw data and reconstructs images for robust GRB localization. ComptonUNet was designed to operate effectively under conditions of limited photon statistics and strong background contamination by combining the statistical efficiency of direct reconstruction models with the denoising capabilities of image-based architectures. We perform realistic simulations of GRB-like events embedded in background environments representative of low-Earth orbit missions to evaluate the performance of ComptonUNet. Our results demonstrate that ComptonUNet significantly outperforms existing approaches, achieving improved localization accuracy across a wide range of low-statistic and high-background scenarios.

Unified Astronomy Thesaurus concepts: [Convolutional neural networks \(1938\)](#); [Gamma-ray bursters \(1878\)](#); [Gamma-ray astronomy \(628\)](#); [Astronomy data analysis \(1858\)](#); [Astronomical simulations \(1857\)](#); [Astronomical reference materials \(90\)](#)

1. Introduction

Gamma-ray bursts (GRBs) are among the most energetic transient phenomena in the Universe and provide critical insights into the deaths of massive stars, neutron star mergers, and the early evolution of the cosmos (T. Piran 2004; S. Woosley & J. Bloom 2006). Since the 1990s, missions such as the Burst and Transient Source Experiment (BATSE) aboard the Compton Gamma Ray Observatory (CGRO) satellite have enabled the detection of thousands of GRBs, revealing their isotropic distribution and establishing their extragalactic origins (W. S. Paciesas et al. 1999). More recently, the Swift satellite, launched in 2004, provided rapid GRB localization and multiwavelength follow-up with its Burst Alert Telescope, X-ray Telescope, and Ultraviolet/Optical Telescope, leading to major breakthroughs in understanding GRB progenitors and their afterglows (N. Gehrels et al. 2004). However, the detection and accurate localization of weak and short GRBs are still challenging primarily because of limited photon statistics under background contamination. These faint events, which often originate from the most distant Universe, provide rare opportunities for probing the early stages of star formation. Consequently, the development of advanced techniques for accurately estimating gamma-ray directions under these challenging conditions is highly desired.

Although traditional GRB detectors such as BATSE employed large scintillator arrays with wide fields of view (FoV) and high sensitivity, recent advances in miniaturized satellite platforms, particularly small satellites, have enabled the low-cost and high-frequency deployment of high-energy detectors in space (N. Werner et al. 2018; S. Serjeant et al. 2020). One such mission is the INnovative Space Probe for Imaging R-process Emission (INSPIRE) with a Compton camera, developed by Waseda University and the Institute of Science Tokyo, and scheduled for launch in 2027 (J. Kataoka et al. 2024; K. Tanaka et al. 2024). The wide FoV of the INSPIRE CC-Box, which covers approximately one-fourth of the entire sky, makes it highly suitable for observing not only steady sources but also unpredictable transient events such as GRBs, black hole binaries, and active galactic nuclei flares. In particular, the accurate GRB localization is expected to play a crucial role in multimessenger astronomy by coordinating with gravitational wave observations from LIGO and Virgo (B. P. Abbott et al. 2017, 2020). However, compared to large-scale predecessors like BATSE, the microsatellite-based INSPIRE mission is constrained by a limited detector size, making it difficult to accumulate sufficient photon statistics within such short timeframes. To address these limitations, INSPIRE aims to detect weak GRBs through the incorporation of deep learning techniques tailored to the platform's limited detection area and the short duration of these events.

Recently, deep learning has been increasingly applied to Compton-camera data analysis to enable accurate image reconstruction under challenging conditions (G. Daniel et al.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

2022; Z. Yao et al. 2022). A common approach involves convolutional neural networks such as UNet (O. Ronneberger et al. 2015) to estimate the gamma-ray directions from the reconstructed Compton image (S. Sato et al. 2020). This approach is inherently robust to background noise because event selection during the reconstruction process removes most of the background events, and the model is trained to reduce background noise. However, this event selection may discard events that implicitly contain information regarding gamma-ray directions, resulting in degraded performance under limited photon statistics. More recently, ComptonNet (S. Sato et al. 2024) was proposed, which directly processes raw data using a neural network to estimate gamma-ray directions. This approach is particularly effective with limited statistics, making it suitable for transient detection. However, because ComptonNet treats all measured events, including background and noise, as valid inputs, its performance may degrade significantly in high-noise environments dominated by the cosmic X-ray background (CXB; E. Boldt 1987) or atmospheric albedo (G. L. Stephens et al. 2015).

Consequently, we propose ComptonUNet, a hybrid deep learning framework that jointly processes raw data in a ComptonNet manner and reconstructed images in a UNet manner. By leveraging the complementary strengths of both models, ComptonUNet enables accurate GRB localization even under limited photon statistics and significant background contamination. We simulated GRB-like events under realistic space radiation backgrounds and compared their performance with that of existing models to evaluate ComptonUNet’s performance.

The remainder of this paper is organized as follows. Section 2 describes the configuration of the Compton camera, the architectures of the proposed models with their baselines, and the simulation setup used for dataset generation. Section 3 presents the comparative results of the reconstruction performance. Section 4 discusses the noise robustness of the model, the contribution of each input component via an ablation study, and the astrophysical implications of the results. Finally, Section 5 concludes the paper. Detailed implementation parameters of the neural networks are provided in the Appendix.

2. Materials and Methods

2.1. Compton-camera Configuration and Reconstruction Technique

The primary gamma-ray detector aboard the INSPIRE satellite is the Compton-camera box (CC-Box), which is a compact, lightweight instrument designed for MeV gamma-ray observations (J. Kataoka et al. 2024; K. Tanaka et al. 2024), as depicted in Figure 1. The CC-Box integrates both pinhole and Compton imaging capabilities, covering the energy range from 30 keV to 3 MeV. The effective FoV of the CC-Box is approximately 3.1 sr, enabling the detection of transient phenomena such as GRBs. Its design was optimized for high sensitivity and angular resolution within the constraints of a 50 kg class satellite platform.

The CC-Box comprises a 2×2 array of sensor modules, each comprising a pixelated cerium-doped gadolinium aluminum gallium garnet (Ce:GAGG) scintillator optically coupled to a 16×16 -channel multipixel photon counter (MPPC) array. The front layer is a 5 mm thick Ce:GAGG array with

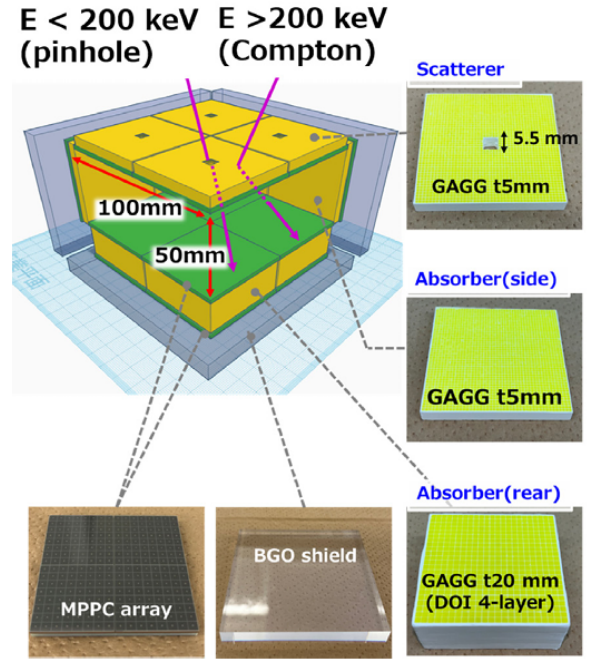


Figure 1. Detailed configuration of the CC-Box onboard the INSPIRE satellite (J. Kataoka et al. 2024). The figure illustrates the front Ce:GAGG pixel array, side and bottom BGO scintillator active shields, pinhole structure, and multilayer detector arrangement. This design enables high-sensitivity and high-angular-resolution observations over a wide energy range (30 keV–3 MeV).

$1 \times 1 \text{ mm}^2$ pixels, featuring a central $5.5 \times 5.5 \text{ mm}^2$ pinhole for low-energy imaging. The rear section includes four layers of 20 mm thick Ce:GAGG arrays with $2 \times 2 \text{ mm}^2$ pixels, forming a depth-of-interaction structure to improve spatial resolution. Additional Ce:GAGG elements are positioned on the sides to enhance off-axis sensitivity. The CC-Box is equipped with bismuth germanate (BGO) scintillator shields surrounding its sides and bottom to suppress background radiation and identify incomplete energy deposition events. Each BGO shield is read out by a monolithic MPPC, providing active shielding and aiding in background discrimination.

For image reconstruction in Compton mode, the CC-Box utilizes Compton kinematics to determine the gamma-ray directions. The scattering angle θ can be calculated using the Compton formula by measuring the energy deposits and interaction positions in the two layers,

$$\cos \theta = 1 - \frac{m_e c^2}{E_2} + \frac{m_e c^2}{E_1 + E_2}, \quad (1)$$

where E_1 and E_2 are the energy deposits in the scatter and absorber layers, respectively, and $m_e c^2$ is the electron rest mass energy. Each event defines a “Compton cone” representing the possible gamma-ray directions, and we reconstruct images by superimposing Compton cones. In pinhole mode, gamma-ray directions are calculated using a pinhole principle that relates the detection position to the incoming direction. Conventional reconstruction methods, such as back projection (BP) and maximum likelihood expectation maximization, integrate multiple Compton cones or pinhole events to visualize the spatial distribution of gamma-ray incidences. However, these methods require a large number of events for accurate reconstructions, making them less effective under low-statistic

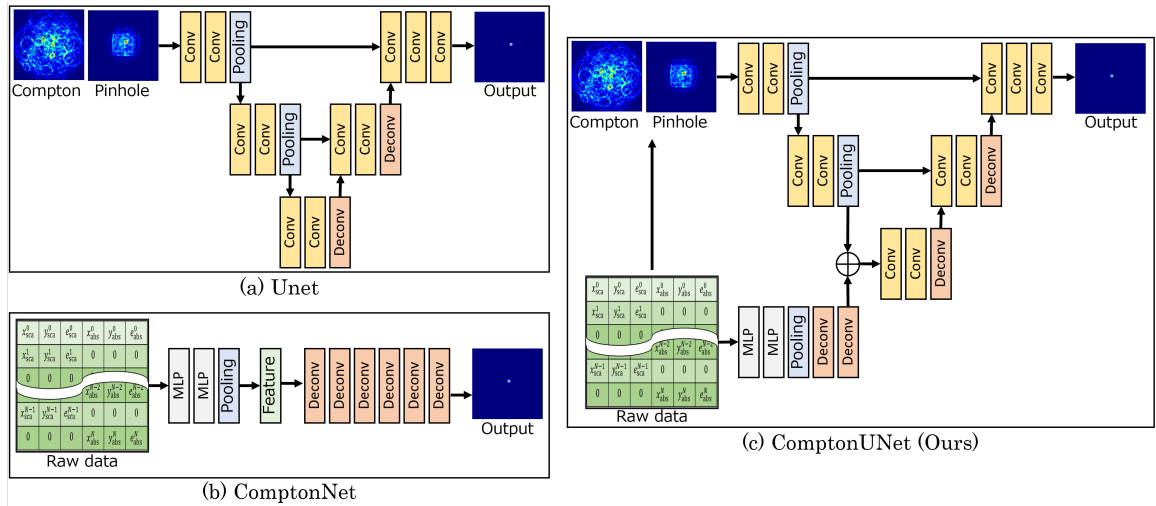


Figure 2. Overview of (a) UNet, (b) ComptonNet, and (c) ComptonUNet architectures. UNet processes reconstructed images, while ComptonNet directly estimates gamma-ray directions from raw data. ComptonUNet combines both approaches by transforming raw data into feature maps using a ComptonNet-inspired encoder, followed by a UNet-style decoder to estimate the gamma-ray directions. This design improves robustness under low-statistics and high-background conditions, making it suitable for transient gamma-ray observations.

conditions, which is typical for short-duration or weak-flux GRBs. Consequently, we investigated deep learning-based reconstruction methods capable of operating under low-statistics and background conditions, thereby enhancing the performance of compact Compton cameras likesuch as the CC-Box, for weak-GRB observations.

2.2. Machine Learning Models

We adopted deep learning models for GRB localization to overcome the limitations of traditional image reconstruction under low-statistics and high-background conditions. In this section, we introduce two established architectures, UNet (O. Ronneberger et al. 2015) and ComptonNet (S. Sato et al. 2024) as baseline models. We then propose a novel hybrid framework, ComptonUNet, which integrates the strengths of both models.

UNet, originally developed for biomedical image segmentation, has been widely adopted for image reconstruction. In the context of GRB localization with Compton cameras, UNet-based models process reconstructed images using Compton and pinhole modes to estimate gamma-ray directions. The reconstructed images were calculated through event selection based on the energy deposits and respective BP algorithms. The architecture comprised an encoder-decoder structure with skip connections, which allows it to learn hierarchical features and perform effective denoising, as shown in Figure 2(a). This approach provides robust performance in noisy environments by implicitly learning to denoise the input data. However, UNet-based models rely on reconstructed images from sufficient events and tend to degrade their performance when the number of input events is severely limited.

In contrast, ComptonNet is a specialized neural network designed to directly estimate gamma-ray directions from raw data without requiring intermediate image reconstruction. Here, the “raw data” consists of normalized feature vectors with 16 channels containing energy depositions and interaction coordinates (E , x , y , z) from each detector segment (front, rear, side, and BGO shields), hence ComptonNet can effectively utilize the full statistical information available from the

detector. The ComptonNet comprises shared-weight multilayer perceptrons (MLPs), max pooling, and deconvolution layers, as shown in Figure 2(b). This approach provides robust performance under low-statistic conditions by directly mapping raw data to the gamma-ray directions. However, ComptonNet is highly sensitive to background contamination, because it treats all measured events as valid inputs. In high-background environments, such as those dominated by CXB and albedo, the model can struggle to distinguish between signal and noise, leading to degraded performance.

To address these complementary limitations, we propose ComptonUNet, a hybrid deep learning framework that jointly processes raw data and reconstructed images for robust GRB localization. As shown in Figure 2(c), this model first transforms raw data into feature maps with a ComptonNet-inspired encoder and then applies the UNet-style convolutional model to estimate the gamma-ray directions. This design allows the model to tolerate both limited statistics and significant noise contamination, making it well-suited for space-based transient gamma-ray observations such as GRBs. In this study, following ComptonNet, we selected the mean squared error (MSE) for the loss function during training. We trained and evaluated all models using simulated data that incorporated both signal and realistic background sources, and compared their performance across a range of GRB durations. Detailed specifications of these input formats and the reconstruction logic are provided in the Appendix.

2.3. Dataset Preparation

We generated a synthetic dataset using Geant4 (S. Agostinelli et al. 2003) simulations in which the full geometry and material properties of the CC-Box were faithfully reproduced to train and evaluate the proposed machine learning models under realistic observational conditions. The detector configuration, including the Ce:GAGG scintillators, MPPC, and shielding components, was modeled based on the engineering specifications described in (J. Kataoka et al. 2024).

GRB events were simulated as point sources toward the CC-Box with a fixed photon flux of $1.0 \text{ photons cm}^{-2} \text{ s}^{-1}$, which




Figure 3. Summary of quantitative performance for the three models: (a) MSE, (b) SSIM, and (c) peak offset. ComptonUNet outperforms both UNet and ComptonNet across all metrics by combining the strengths of both architectures.

corresponds to the typical value reported in GRB population studies (W. S. Paciesas et al. 1999). Simulations were performed with burst durations of 1, 3, 10, 30, and 100 s to evaluate model performance across a range of photon statistics. For each duration, we generated 1000 independent simulation runs, each with a GRB source placed at a random position within the 30° half-angle FoV of the CC-Box. To train and evaluate the models, we split the dataset into training (640 runs), validation (160 runs), and test (200 runs) sets. The energy spectra of the GRB sources were assumed to follow a power-law distribution with a photon index Γ randomly sampled from the range of $[-2.5, -1.3]$, which is consistent with the observed GRB spectral slopes in the 50 keV–10 MeV range. Each simulation included background contributions from both the CXB and the albedo. Following the empirical fit (D. Gruber et al. 1999), we adopted a simple power-law approximation of the CXB and albedo spectrum with a photon index of $\Gamma = -2.88$ and $\Gamma = -1.35$, respectively.

All photons were tracked through the full detector geometry to produce event-level data, including energy deposits, reaction positions, and timing information. Energy and position resolutions were considered in this simulation, following the CC-Box. No preselection or event filtering was applied prior to model training, to preserve the statistical and noise characteristics of realistic GRB observations.

2.4. Quantitative Evaluation Metrics

We employed a set of quantitative metrics commonly used in image analysis, as well as specialized criteria tailored to GRB localization, to assess the performance of each reconstruction model.

The overall reconstruction accuracy was evaluated using two standard metrics: MSE and structural similarity index measure (SSIM; Z. Wang et al. 2004). The MSE quantifies the pixelwise differences between the reconstructed image and the corresponding ground truth, providing a measure of the global deviation. SSIM, in contrast, evaluates the perceived structural similarity between the predicted images and ground truth by considering luminance, contrast, and texture information. In addition to general performance metrics, we utilized a task-specific criterion for GRB localization: the angular offset between the predicted and ground-truth gamma-ray directions. For each image, the predicted direction was estimated by computing the intensity-weighted centroid of the reconstructed image. The angular separation between the centroid and ground-truth directions was then calculated. This “peak offset” directly quantifies localization accuracy and serves as a key

indicator of a model’s suitability for GRB localization. Following the ComptonNet paper, we trained each model 10 times independently for each experimental condition (i.e., for each GRB duration). In these runs, the dataset split was fixed, but the weight initialization and data shuffling were varied using different random seeds. For quantitative assessment, we selected the top five runs that achieved the lowest MSE on the validation dataset and computed the mean and variance of their evaluation metrics using the test dataset. This approach ensures that the reported results are robust against random initialization and training fluctuations, thereby providing a reliable estimate of the model performance under each condition.

3. Results

3.1. Quantitative and Visual Performance Comparison

We first compared the performances of the three models, UNet, ComptonNet, and ComptonUNet (ours), using the quantitative metrics described in Section 2.4. Figure 3 summarizes the results in terms of the MSE, SSIM, and peak offset across all test cases. For MSE, all models exhibited improved estimation accuracy with increasing GRB duration. In the 1–10 s range, accurate localization remains difficult because of severely limited photon statistics, resulting in low signal amplitudes and a high MSE across all models. For durations beyond 30 s, the photon statistics are still insufficient for UNet, and the background effects are still critical for ComptonNet. In contrast, ComptonUNet consistently achieves lower MSE values, demonstrating its ability to balance statistical efficiency with noise robustness. The SSIM values exhibit 0.88–0.90 among the models for short durations (1–10 s), as all models produce low pixel values close to zero due to the severely limited photon statistics. For durations exceeding 30 s, ComptonUNet demonstrates superior performance, mirroring the trend observed in the MSE results. Similarly, the peak offset follows a trend comparable to that of the MSE, with ComptonUNet consistently providing more accurate localization across most conditions. Consequently, even with the limited detector size of a 50kg-class satellite, ComptonUNet achieves a localization accuracy of 7.5° at 30 s and 2.5° at 100 s, assuming a fixed flux of $1.0 \text{ photons cm}^{-2} \text{ s}^{-1}$. Note that the geometrical area of INSPIRE is only $\sim 1/20$ of one of eight NaI(Tl) scintillators constituting the BATSE detector.

Because ComptonNet directly estimates gamma-ray directions from raw data without using reconstructed images and includes background events, its training process is unstable,




Figure 4. Visual comparison of reconstructed images from the three models. From top to bottom: Ground truth, BP image of Compton mode (BP Compton), BP image of pinhole mode (BP Pinhole), UNet output, ComptonNet output, and ComptonUNet output. Additionally, from left to right, the images represent different GRB durations (1, 3, 10, 30, and 100 s). The ComptonUNet model demonstrates superior performance in accurately reconstructing the source morphology and peak location compared to the other models.

and the inference performance varies significantly depending on the initialization parameters. In contrast, ComptonUNet incorporates reconstructed images as additional inputs, which provide analytically calculated gamma-ray directions. This enables ComptonUNet to achieve greater stability and reduced variance in its predictions compared with ComptonNet.

For comparison, we also evaluated the localization accuracy of the traditional BP method. The peak offsets for BP were 20.7° , 19.7° , 16.3° , 14.8° , and 13.9° for durations of 1, 3, 10, 30, and 100 s, respectively. While the BP method shows some improvement with duration, the accuracy is significantly inferior to that of ComptonUNet, which achieves 2.2° at 100 s. Note that we do not compare MSE or SSIM for BP, as these pixelwise metrics are heavily biased by image sparsity in low-statistic regimes and do not correctly reflect localization performance in this context.

Figure 4 presents representative output images from each model under identical simulation conditions. The ground truth, BP images, and model predictions are shown for the GRB durations of 1, 3, 10, 30, and 100 s. In Compton-mode BP, only a few cones are visible for 1–3 s, which makes

localization difficult. At 100 s, the GRB directions are roughly distinguishable. In pinhole-mode BP, the GRB source is located outside the FoV; consequently, it does not appear in the reconstructed images. The UNet performance improves with increasing duration, particularly from 30 to 100 s, but still exhibits residual noise owing to the limited photon counts. ComptonNet performs worse than UNet under background environments; the localized position remains deviated from the ground truth even at 30 s, and the source position becomes distinguishable only at 100 s. In contrast, ComptonUNet recovers the GRB peak direction with less noise, outperforming both baseline models. While UNet tends to oversmooth the signal and ComptonNet introduces noise artifacts, ComptonUNet provides a robust compromise, confirming the trends observed in the quantitative evaluations.

4. Discussion

4.1. Noise Robustness

In this section, we evaluate the noise robustness of ComptonUNet by comparing its performance with those of




Figure 5. Noise robustness evaluation of the three models under eliminated background conditions. The dashed lines and simple lines represent the performance of each model under background conditions (as shown in Figure 3) and under eliminated background conditions, respectively. ComptonUNet maintains its performance even in the presence of background noise, while ComptonNet shows significant degradation.

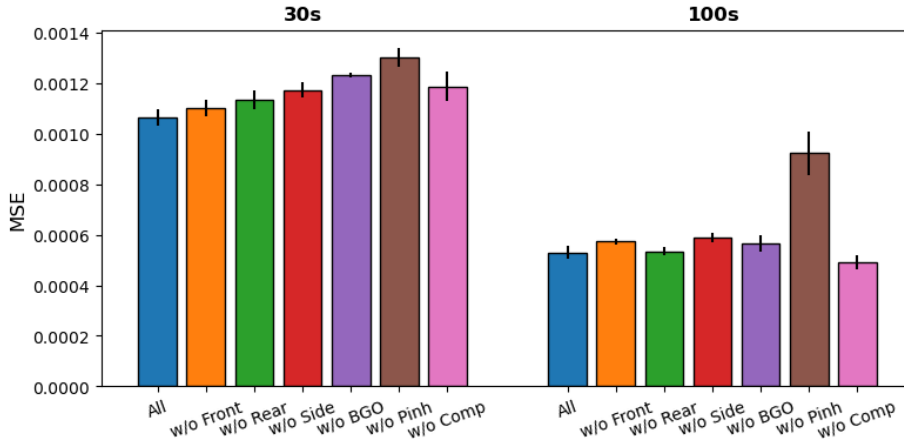


Figure 6. Ablation study of ComptonUNet. The ablation study demonstrates the importance of each component in the ComptonUNet inputs for MSE evaluation. “All” represents the full input. Additionally, w/o Front, Rear, Side, BGO, Pinh, and Comp ignored the front, rear, side, BGO, pinhole image, and Compton image components, respectively. The results indicate that the pinhole image plays a crucial role in improving the performance of ComptonUNet, since it provides stable hints of gamma-ray direction through analytical reconstruction based on the pinhole principle.

UNet and ComptonNet under different background conditions. In particular, we assess the model performance with and without background contributions from the CXB and atmospheric albedo, allowing us to isolate the impact of noise on the measured events.

Figure 5 shows the results in terms of MSE, SSIM, and peak offset for all models under both background-present and background-removed scenarios. As expected, the performance of all models improves when background noise is eliminated. UNet exhibits relatively stable performance under both conditions, as much of the background is removed during the event selection processes. In contrast, ComptonNet exhibits a substantial degradation in performance under background-present conditions, primarily because it processes all measured events, including noise, without explicit background filtering. ComptonUNet, however, demonstrates a considerably smaller performance gap between background and no-background conditions. This indicates that its hybrid architecture effectively mitigates the influence of noise, balancing statistical sensitivity with denoising capability.

4.2. Ablation Study

We conducted an ablation study in which each component was removed in turn to assess the contribution of the individual input components to the performance of

ComptonUNet. ComptonUNet receives raw data from multiple detector segments (front, rear, side, and BGO), and reconstructs the images (pinhole mode and Compton mode). We evaluated the impact on the reconstruction performance by systematically omitting each component individually.

Figure 6 presents the results of this ablation analysis, showing the mean MSE and standard deviation (error bars) derived from the top five independent training runs for each configuration. The figure indicates that the full ComptonUNet model achieves the best performance at both 30 s and 100 s. In the 30 s analysis, removing the pinhole image component results in a significant performance degradation, highlighting its critical role in providing stable hints of the gamma-ray direction through analytical reconstruction based on the pinhole principle. Additionally, the BGO component plays an important role by providing gamma-ray direction constraints from count ratios and removing escape events via veto signals. At 100 s, the pinhole image remains significant for source localization.³

³ Although the “w/o Comp” configuration appears to have a slightly lower mean MSE at 100 s compared to the full model, this difference falls within the standard deviation, suggesting that it is attributable to statistical fluctuations rather than a physical advantage.

4.3. Comparison with BATSE Observations

BATSE aboard the CGRO achieved remarkable success in detecting various GRBs with different fluxes and durations. BATSE comprised eight large-area NaI(Tl) scintillation detectors, each with an effective area of approximately 2025 cm^2 , mounted on the corners of the satellite to provide nearly all-sky coverage. The system employed simple count-rate threshold triggers and triangulation techniques to localize GRBs, with typical positional uncertainties of several degrees depending on the photon statistics and incident angle. Owing to its large detection area and wide FoV, BATSE could detect thousands of GRBs, including several short and faint events.

We compared the localization accuracy of BATSE with that of the CC-Box on INSPIRE combined with ComptonUNet to contextualize the performance of our proposed system. We selected BATSE as a benchmark because it shares a similar observation energy band and wide-field survey objective, serving as a suitable reference for degree-level localization, unlike coded-mask missions (e.g., Swift).

In particular, we focused on GRBs with peak fluxes in the range of $0.8\text{--}1.2 \text{ photons cm}^{-2} \text{ s}^{-1}$, which is consistent with our simulation setup. We compared the positional uncertainty reported in the 4B BATSE GRB catalog (W. S. Paciesas et al. 1999) with the peak offset distribution derived from our simulations, using the interquartile range (IQR; 25th–75th percentile) as a metric. For GRBs with $T_{90} \sim 30 \text{ s}$, BATSE exhibited a positional uncertainty of $3.04\text{--}5.57$, whereas INSPIRE with ComptonUNet yielded a peak offset range of $3.08\text{--}9.03$. At $T_{90} \sim 100 \text{ s}$, BATSE achieved $1.80\text{--}2.26$, whereas ComptonUNet produced a comparable IQR of $1.02\text{--}2.32$. It should be noted that these ComptonUNet performances differ from the mean performance shown in Figure 3. While Figure 3 presents the average of the top five training runs to demonstrate model stability, the values reported here are derived from the single best-performing model among them, assuming a realistic deployment scenario where the optimal model is selected for operation.

Notably, this comparison is not strictly equivalent. The BATSE results were derived from real observational data, whereas the INSPIRE+ComptonUNet results are based on detailed Geant4 simulations. Additionally, the BATSE values represent estimated uncertainties rather than measured offsets from true positions, as ground truth is unavailable for most BATSE events. Moreover, BATSE was specifically designed for transient detection, whereas the CC-Box on INSPIRE was optimized for steady gamma-ray sources and did not employ a geometry specialized for GRB localization. Furthermore, the CC-Box has a significantly smaller effective detection area ($\sim 100 \text{ cm}^2$) than BATSE's $\sim 2000 \text{ cm}^2$, which inherently limits its photon collection capabilities. Nonetheless, the comparison provides a useful benchmark, suggesting that advanced deep learning–based reconstruction enables competitive localization performance, even with compact instruments operating under constrained photon statistics.

4.4. Limitations

Despite these advantages, several limitations of this study must be acknowledged.

First, it relied entirely on simulated data. Although the simulations incorporate realistic background sources, such as CXB and albedo, the generalization of ComptonUNet to real-

world conditions, including detector response nonuniformities, calibration uncertainties, and hardware-related artifacts, remains to be evaluated. The incorporation of experimental data from laboratory setups or orbital missions is essential for validating the practical performance of a model.

Second, the GRB-like sources in this study were restricted to within 30° of the CC-Box FoV center, which is consistent with the FoV limits of 30° for the pinhole mode and 60° for the Compton mode. However, the actual GRB sky distribution is isotropic, and the model performance may vary significantly in off-axis regions. Future work should explore the robustness of the model across a wider range of gamma-ray directions and angles, particularly using the BGO count rates such as BATSE.

Third, the size of the training dataset (1000 samples per duration) was selected to verify the proposed method while maintaining reasonable computational costs. We confirmed the sufficiency of this size for stable training by conducting 10 independent runs, which yielded consistent results with small variance as shown in Figure 3. However, it is possible that further improvements in localization accuracy could be achieved by significantly increasing both the dataset size and the model capacity (i.e., the number of parameters). Exploring such large-scale models, potentially at the cost of higher computational requirements, remains an important direction for future research.

4.5. Scientific Implications

In the context of multimessenger astronomy, a localization accuracy of approximately $3^\circ\text{--}5^\circ$ is the practical threshold for coordinating effective follow-up observations. Our results with ComptonUNet demonstrate that effective localization achieving this accuracy is feasible for long GRBs (e.g., 100 s) even under typical fluxes ($1.0 \text{ photons cm}^{-2} \text{ s}^{-1}$) using a micro-satellite-scale detector. Regarding short GRBs, while detection remains challenging under typical flux conditions due to limited photon statistics, the proposed method could potentially be applicable to exceptionally bright events where the integrated photon counts are comparable to those of long GRBs. Precise position determination further enables the application of angular resolution measure cuts (J. Kataoka et al. 2025). By filtering events based on the derived source direction, the signal-to-noise ratio can be significantly improved, thereby facilitating accurate spectral extraction even for faint sources.

4.6. Strategy for Application to Real Observations

In this study, we evaluated models trained on discrete durations and a fixed flux to demonstrate the feasibility of the proposed method. For application to real observational data, where GRBs exhibit continuous variations in duration and flux, we propose a ‘‘Model Bank’’ strategy. In this approach, a library of network weights is pretrained on a grid of conditions covering various durations and average fluxes. In the analysis pipeline, the system first detects a GRB trigger and estimates its approximate duration (T_{90}) and average flux from the light curve. Subsequently, the model weights trained on the condition closest to the observed parameters are selected for inference. Furthermore, since the inputs to our model are time-integrated images and count vectors, the system is robust to the

temporal variability of the flux as long as the integrated photon statistics align with the selected model’s training domain.

5. Conclusion

We presented ComptonUNet, a hybrid deep learning architecture for estimating gamma-ray direction using Compton-camera data under low-statistic and high-background conditions. ComptonUNet demonstrated robust performance across varying GRB durations and background environments by combining the event-level inference capability of ComptonNet with the denoising strengths of UNet. Through detailed Geant4-based simulations of the CC-Box instrument aboard the upcoming INSPIRE mission, we demonstrated that ComptonUNet consistently outperformed baseline models in terms of reconstruction accuracy, noise robustness, and localization precision. These results highlight the potential of hybrid deep learning approaches for enhancing transient gamma-ray detection using compact space-based instruments. Future work will focus on validating the model performance with experimental and in-orbit data to further advance the applicability of deep learning techniques in high-energy astrophysics.

Acknowledgments

This research was supported by the Japan Science and Technology Agency (JST) ERATO grant No. JPMJER2102.

Author Contributions

S.S. conducted the core research activities, including the development of the ComptonUNet model, implementation of simulations, data analysis, and manuscript preparation. K.T., K.M., and R.T. contributed to the conceptual development of the study through technical discussions and provided interpretation of results, particularly in the Discussion section. S.O. and K.Y. assisted in the construction and preprocessing of simulation datasets used in the model training and evaluation phases. J.K. served as the principal investigator of the project, securing research funding, guiding the overall scientific direction, and providing continuous feedback on the methodology and manuscript content. All authors reviewed and approved the final manuscript.

Appendix Implementation Details

This paper presents a novel hybrid deep learning architecture, ComptonUNet, for gamma-ray direction estimation using Compton-camera data. This Appendix provides additional implementation details for ComptonUNet and the baseline models to facilitate reproducibility and further understanding of the model.

UNet. The UNet baseline is implemented as an encoder-decoder network with skip connections. The encoder consists of two stages (depth = 2) of convolutional blocks. Each block includes two 3×3 convolutional layers, each followed by batch normalization and a LeakyReLU activation. The input consists of two channels (256×256 pixels). For this study, the UNet model was configured as a lightweight version with 18,817 trainable parameters. During training, a batch size of 16 was used.

ComptonNet. ComptonNet processes raw event data directly. The encoder comprises shared-weight MLPs, transforming the 16-channel input features into a 128-dimensional latent space. The decoder consists of five transposed convolutional layers. This model contains 3,855,120 trainable parameters and was trained with a batch size of 4.

ComptonUNet. ComptonUNet integrates the encoders of both architectures. The raw data path uses two NonLinear layers (64 and 128 units), while the image path processes reconstructed images through three convolutional stages (32, 64, and 128 filters). These paths are concatenated at the bottleneck (256 channels). The total number of trainable parameters is 5,550,593. Similar to ComptonNet, it was trained with a batch size of 4.

Shared Training Details and Computational Cost. All models were implemented using PyTorch and trained on an NVIDIA GeForce RTX 2080 Ti GPU. The training process employed the Adam optimizer with the MSE loss function used as the primary objective for all architectures. To ensure optimal generalization and prevent overfitting, an early stopping strategy was implemented; in particular, training was terminated if the validation loss failed to improve for 30 consecutive epochs after an initial warm-up period of 50 epochs. During this process, the model weights corresponding to the minimum validation loss were preserved for final evaluation. The computational efficiency was assessed by monitoring peak GPU memory consumption and processing time during training with the 100 s duration dataset. Under these conditions, the UNet model required 4.37 GB of GPU memory, while ComptonNet and the proposed ComptonUNet utilized 1.28 GB and 2.50 GB, respectively. Regarding computational time, UNet was the fastest, with an average training time of 8.54 s per epoch (21.6 minutes total for 152 epochs). ComptonNet required 30.4 s per epoch (41.5 minutes total for 82 epochs). As expected, the proposed ComptonUNet was the most demanding, requiring 31.9 s per epoch (50.0 minutes total for 94 epochs). The relatively small difference in per-epoch time between ComptonNet and ComptonUNet implies that the processing of raw event data dominates the computational load. For reference, the inference time for processing 200 test samples was 1.13 s (5.65 ms/sample) for UNet, 7.49 s (37.5 ms/sample) for ComptonNet, and 7.72 s (38.6 ms/sample) for ComptonUNet.

Input Data (Raw Data). The input to the ComptonNet-based encoder consists of raw event-level information derived from the CC-Box. For each detected event, we construct a 16-channel feature vector that encapsulates the energy deposition recorded in the Front, Rear, and Side GAGG scintillators, as well as the BGO active shield. Furthermore, the vector includes the three-dimensional interaction coordinates (x, y, z) and its deposit energy E for each event. These features are normalized to a [0, 1] range before being processed by the shared-weight MLP blocks. In cases where multiple interactions occur within the same detector type, the position corresponding to the highest energy deposition is selected as the interaction location, and the total energy deposition from all interactions is summed to represent the event’s energy.

Input Data (Reconstructed Images). The image-based path of ComptonUNet and the UNet baseline receives two types of 256×256 pixel images representing different detection modes. The first type is the pinhole image, generated by utilizing the front detector layer as a pinhole camera; hit

positions on the rear detector are back-projected through the physical aperture of the CC-Box to the sky plane, effectively capturing low-energy photons. The second type is the Compton image, used for multihit events identified as Compton scattering. A simple Compton BP algorithm calculates the scattering angle θ from energy depositions using the Compton formula, projecting possible source directions as cones onto the sky map. The final image is formed by accumulating these cones, which provides spatial constraints for higher-energy photons. To enhance the quality of the reconstructed images, we use front, rear, and side GAGG detector data for image reconstruction, while employing the BGO active shielding to veto incomplete energy deposition events, thereby reducing background noise.

ORCID iDs

Shogo Sato  <https://orcid.org/0000-0003-3070-5880>
 Kazuo Tanaka  <https://orcid.org/0000-0001-6916-9654>
 Kazuhiko Murasaki  <https://orcid.org/0000-0001-7697-9575>
 Ryuichi Tanida  <https://orcid.org/0000-0002-5379-3150>
 Jun Kataoka  <https://orcid.org/0000-0003-2819-6415>

References

- Abbott, B. P., Abbott, R., Abbott, T., et al. 2020, *LRR*, **23**, 3
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **119**, 161101
 Agostinelli, S., Allison, J., Amako, K., et al. 2003, *NIMPA*, **506**, 250
 Boldt, E. 1987, *PhR*, **146**, 215
 Daniel, G., Gutierrez, Y., & Limousin, O. 2022, *NuEnT*, **54**, 1747
 Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *ApJ*, **611**, 1005
 Gruber, D., Matteson, J., Peterson, L., & Jung, G. 1999, *ApJ*, **520**, 124
 Kataoka, J., Iwashita, R., Tanaka, K., et al. 2024, *NIMPA*, **1065**, 169518
 Kataoka, J., Ogasawara, S., Mori, R., et al. 2025, *JInst*, **20**, P10009
 Paciesas, W. S., Meegan, C. A., Pendleton, G. N., et al. 1999, *ApJS*, **122**, 465
 Piran, T. 2004, *RvMP*, **76**, 1143
 Ronneberger, O., Fischer, P., & Brox, T. 2015, in *Medical Image Computing and Computer-assisted Intervention-MICCAI 2015: 18th Int. Conf., Munich, Germany, October 5-9, 2015, Proc., part III (Springer)*, 234, 18
 Sato, S., Kataoka, J., Kotoku, J., et al. 2020, *JInst*, **15**, P10026
 Sato, S., Tanaka, K., & Kataoka, J. 2024, *ApPhL*, **124**, 253702
 Serjeant, S., Elvis, M., & Tinetti, G. 2020, *NatAs*, **4**, 1031
 Stephens, G. L., O'Brien, D., Webster, P. J., et al. 2015, *RvGeo*, **53**, 141
 Tanaka, K., Kataoka, J., Iwashita, R., et al. 2024, *SPIE*, **13093**, 130932O
 Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. 2004, *ITIP*, **13**, 600
 Werner, N., R^ípa, J., Pál, A., et al. 2018, *SPIE*, **10699**, 106992P
 Woosley, S., & Bloom, J. 2006, *ARA&A*, **44**, 507
 Yao, Z., Shi, C., Tian, F., et al. 2022, *MedPh*, **49**, 7336