ComptonINR: Implicit Neural Representations for Fast Modeling of Compton Telescope Point Spread Functions

Anirudh Kotamraju¹² Andreas Zoglauer²

Abstract

Compton telescopes enable the observation of the MeV universe; however, detection of gamma-ray sources requires the generation of the instrument response function, typically computed by intensive simulations on high-performance computers. We introduce ComptonINR, a small, coordinatebased neural network that learns the mapping between a point source location in the image space and the Compton camera's measurements in the data space. ComptonINR's continuous structure enables training on a small, simulated set of coarse resolution point spread functions (PSFs) instead of the whole response. Moreover, ComptonINR can interpolate PSFs during inference at multiple times higher resolution-this significantly reduces the response simulation requirements. The model accurately generalizes to untrained source locations, and source detection is achieved with 0.987 precision, 0.964 recall, and ~0.63° median angular error. ComptonINR reduces the required simulation time by roughly a factor of 70 and trains in \sim 53 minutes on a MacBook. Furthermore, ComptonINR scales favorably with event count, charting a realistic path towards high-resolution gamma-ray imaging on consumer hardware.

1. Introduction

Gamma rays in the energy range from hundreds of keV to tens of MeV offer a powerful lens into the study of nucleosynthesis, black holes, neutron stars, pulsars, galactic nuclei, supernovae, and gamma-ray bursts (Grenier & Harding, 2015; De Angelis et al., 2021). However, building telescopes with sufficient sensitivity to detect emissions in this energy band, known as the MeV gap (Kierans et al., 2022), is a challenging task. Specialized telescopes known as Compton telescopes overcome this challenge by relying on a process called Compton scattering. During this process, a gamma ray enters the detector and transfers a portion of its momentum to an electron in the detector. This deflects the gamma ray by a Compton scatter angle ψ determined by energy-momentum conservation (see Appendix A.5). This photon, now with reduced energy, may Compton scatter again or undergo a final photo-absorption. The initial scattering direction of the gamma ray, defined as the unit vector from the first photon-electron interaction to the second interaction, is represented by (θ, ϕ) . By measuring the location of each interaction and the energy transferred to each electron, the Compton-scatter angle ψ of the first interaction and the total measured energy $E_{measured}$ can be calculated (Kierans et al., 2022).

The point spread function (PSF) is a structure that encodes physical imperfections of the telescope. The set of PSFs for all possible gamma-ray source locations is called the instrument response. This response is the key component to creating images of the MeV sky from measured data (Kierans et al., 2022). Currently, the best way to create the response of a Compton telescope is to simulate $O(10^{13})$ gamma rays using highly realistic Monte-Carlo simulations. However, these simulations are extremely time-consuming and memory-intensive, requiring significant super-computer usage (Zoglauer, 2005).

A machine learning model that learns the response from a small set of simulated PSFs could interpolate all untrained PSFs, avoiding the compute-demanding simulation of the full response. Moreover, it could also significantly reduce the memory footprint of the 7-dimensional response. We present ComptonINR, a novel deep-learning approach to modeling the response function of Compton telescopes with minimal training data, compute resources, training time, and inference time. Our approach builds upon Implicit Neural Representations (INRs), which are fully connected neural networks that map coordinates to the value(s) of a signal. For example, an INR for an image could learn $f(x, y) \mapsto (\text{Red, Green, Blue})$ (Sitzmann et al., 2020).

¹Department of Electrical Engineering and Computer Sciences, University of California - Berkeley, Berkeley, USA ²Space Sciences Laboratory, University of California - Berkeley, Berkeley, USA. Anirudh Kotamraju <akotamraju@berkeley.edu>.

ML4Astro 2025, Vancouver, CA. Copyright 2025 by the author(s).



Figure 1. An ideal PSF in the Cartesian data space for one energy value and one source location. Events from a point source at $(\theta_{src}, \phi_{src})$ produce a cone (the PSF for this source) with fixed 90° opening angle in the (θ, ϕ, ψ) data space.

2. Methods

2.1. Image Space vs. Data Space Formulation

Two spaces are crucial to Compton imaging. First, $(\theta_{src}, \phi_{src}, E_{src})$ span the image space. This space contains the true locations of the gamma-ray sources (Kierans et al., 2022).

Four dimensions (θ, ϕ, ψ, E) span what is known as the data space of a Compton telescope. In this space, each measured gamma ray corresponds to one single point. For an ideal, noise-free detector, the amalgamation of many photons from a single source forms the surface of a cone, centered at $(\theta_{src}, \phi_{src})$ — the origin of the gamma rays. Figure 1 shows an ideal data space. This cone is the Compton telescope's PSF for one source location (Kierans et al., 2022).

The response function is the probabilistic mapping between the image space and the data space. Formally, the response can be defined as $P(\theta, \phi, \psi, E \mid \theta_{src}, \phi_{src}, E_{src})$: how likely is it that a photon emitted from a point source at $(\theta_{src}, \phi_{src})$ with energy E_{src} is measured in data space bin $(\theta, \phi, \psi, E_{measured})$ (Kierans et al., 2022)? For this work, we focus on a single energy, the 511 keV energy line —our approach models $P(\theta, \phi, \psi \mid \theta_{src}, \phi_{src}, E \sim 511$ keV).

2.2. Data Generation

Using MEGAlib (Zoglauer et al., 2006), we simulated 228 gamma-ray point-sources with \sim 1,000,000 events per source. In other words, we have 228 PSFs. These events were simulated with a spherical Compton telescope on an AMD Threadripper 3970X (32 cores/64 threads) with 256 GiB RAM and an NVIDIA RTX 3080 GPU (12 GB VRAM). 170 PSF files were utilized for training, while 58 were utilized for validation.

2.3. Data Binning

Although continuous, the data space must be discretized into bins to perform processing on a computer. We placed events in the simulation data space into one of the following ψ bins: $[10^\circ, 20^\circ), [20^\circ, 30^\circ), \dots, [80^\circ, 90^\circ)$. Backscattered events with a ψ outside of these bins were excluded.

Within each bin, we performed a change of coordinates. Each event's (θ, ϕ) was treated as spherical coordinates on an arbitrary unit sphere, rather than Cartesian coordinates. After this change of coordinates, the spherical data space still contains a cone, centered at the spherical location $(\theta_{src}, \phi_{src})$ (see Appendix A.1).

Using HEALPix's fast hierarchical representation of spherical data, we discretized the spherical data space into equal area pixels. The NSIDE parameter determines the resolution of a HEALPix sphere. We chose a coarse NSIDE of 16 for fast training, which corresponds to 3,072 pixels for one cross section of the spherical data space (Gorski et al., 2005). Each pixel contains the number of events with that $(\theta_{src}, \phi_{src}, \theta_{pix}, \phi_{pix}, \psi)$ combination, where θ_{pix} and ϕ_{pix} are the spherical coordinates for the center of that pixel.

2.4. Network Architecture

1

$$f\left(\underbrace{\sin\theta_{src},\cos\theta_{src},\sin\phi_{src},\cos\phi_{src}}_{\text{point source location}}, \underbrace{\sin\theta_{pix},\cos\theta_{pix},\sin\phi_{pix},\cos\phi_{pix},\psi}_{\text{spherical data space pixel being queried for}}\right) \mapsto \hat{I}$$
(1)

ComptonINR is a fully-connected neural network that learns a continuous mapping between coordinates to event intensity (Equation 1). This category of network is known as an implicit neural representation (INR) (Sitzmann et al., 2020). For a given point source location, rather than try to output the entire data space at once, we query for a specific pixel in the data space. This significantly reduces the complexity of the network. With this approach, we can repeatedly perform inference on ComptonINR for every pixel to generate the full data space.

To prevent discontinuities around the $-180^{\circ}/180^{\circ} \phi$ border, we added sin / cos periodic encodings, which uniquely identify each angle. A sample output when the encodings are removed is presented in Appendix Figure A.2.

ComptonINR has 4 hidden layers, each with 1024 hidden units. We chose ReLU for our intermediate layer activations, Softplus as our final layer activation, and Poisson Negative Log Likelihood Loss (PoissonNLL) as our loss function. We chose PoissonNLL as gamma-ray events typically follow Poisson distributions (Schmitt et al., 2009). The addition of



Example Measured Data

Figure 2. End-to-end deconvolution of one example measured data space into an NSIDE=32 sky map, running Richardson-Lucy for 50 iterations with ComptonINR as the PSF. In the sky map, red \times 's mark detected sources; green dots mark ground truth sources.

a final Softplus activation ensures strictly positive intensities. The output of ComptonINR is a single intensity value \hat{I} : the number of events in that specific scatter angle bin and pixel bin. We can transform ComptonINR's predicted data space from raw counts into a probability distribution through dividing by the total number of events across all ψ bins.

2.5. Training Details

ComptonINR was trained with a batch size of 1024, learning rate of 0.001, and 40 maximum epochs. On a 2021 MacBook M1 Pro (10-core CPU, 32 GB RAM), training converged in **53 min (79 s/epoch)**. On an A40 GPU (48 GB VRAM, 10-core CPU), it converged in **30 min (45 s/epoch**). See Appendix A.3 for further details.

2.6. Data Space Deconvolution

Raw measurements in the data space must be deconvolved into a sky map, a map in the image space of the point sources. We chose the Richardson-Lucy (RL) algorithm. RL assumes that a ground truth signal has been blurred by some inherent device response function. Using a model of the response, RL iteratively deconvolves the blurred signal towards the ground truth signal (Richardson, 1972; Lucy, 1974). Richardson-Lucy can be performed by precomputing the response cube, which is a stored map from every possible point source to the PSF. As a sidenote, we

Source Detection Metrics	Mean	Std. Dev.
Precision	0.9867	0.0361
Recall (TPR)	0.9640	0.0686
F1-Score	0.9747	0.0518
False Positive Rate (FPR)	0.0000	0.0000
Flux Absolute Error (%)	7.80	12.13
Angular Error Metrics		
Median Angular Error	0.6295°	_
25th Percentile AE	0.4998°	_
75th Percentile AE	0.9067°	-
Runtime Metrics		
Run on MacBook M1 Pro, 10-c	ore CPU, 3	2GB RAM
Response Cube Generation	34 min	_
\Rightarrow One-time, then cached		
Single Sky Image Creation (s)	88.4s	3.1s

Table 1. Sky Map Creation with ComptonINR (NSIDE=32)

omit a background model for RL but are currently integrating one. Our response cube consumed ~4.8 GB of storage and had dimensions 12, $288^2 \times 8$ (with NSIDE=32: 12,288 possible source locations, 12,288 pixels per ψ bin, 8 ψ bins).

3. Results

3.1. Sky Image Creation

ComptonINR's continuous structure excels during deconvolution. While trained on images with a coarse NSIDE resolution of 16, the model has no input dependence on NSIDE as it interpolates for any $(\theta_{src}, \phi_{src})$. Accordingly, we chose a larger NSIDE of 32 (12,288 pixels per ψ bin) during RL to produce a sky map with higher resolution. Figure 2 shows an example deconvolved sky map, and Table 1 summarizes metrics on source detection.

A key figure is that for a full response, we would have to simulate 12,288 source locations, not just 170 — Compton-INR reduces simulation time by a factor of ~70. Consider increasing sky map resolution to NSIDE=256 and the number of events to 10,000,000. This would require simulating $O(10^{13})$ gamma rays. In contrast, ComptonINR only scales with a small number of training file events $\ll O(10^{13})$ and is largely independent of inference time NSIDE resolution.

Each simulation file contains the data space for a single point source (θ_{src}, ϕ_{src}). We drew 50 random combinations of 10 point sources along with the expected data space for each combination. We enforced that point sources are separated by 8° pairwise to ensure viable deconvolution. To compute the expected data space if all point sources in a combination existed simultaneously, we summed the data spaces of the



Figure 3. ComptonINR generalizes to unseen scatter angles. The figure shows (a) the predicted data space for an untrained 55°-65° ψ bin (b) the ground truth data space for this bin (c) the difference between the ground truth data space and the predicted data space (d) the natural logarithm of the predicted data space, highlighting the PSF wings—correctly learned imperfections in the imaging device.

individual simulation files.

3.2. ComptonINR Predicted Data Space



Figure 4. Per-pixel Poisson deviances (see Appendix A.4 for formula) for all eight ψ bins across all validation files. Medians stay close to the ideal value of 1, indicating statistical consistency between ComptonINR and an ideal PSF model.

Since gamma-ray events can be modeled using Poisson distributions (Schmitt et al., 2009), we treated each HEALPix pixel as the average of a Poisson distribution. Figure 4 shows box plots of the Poisson deviances per pixel for each of the ψ bins. For noisy systems, the expected Poisson deviance when normalized by the degrees of freedom—in our case, the number of pixels in one ψ cross section—is 1 (Mohr et al., 2022). Moreover, one powerful feature of the continuous, coordinate-based capabilities of ComptonINR is the ability to interpolate outputs for $(\theta_{src}, \phi_{src}, \psi)$ input data that the model was not trained on, shown in Figure 3.

4. Conclusion

We introduce ComptonINR-a fast, continuous implicit neural network to represent the response of Compton telescopes. In our example, ComptonINR reduced the required simulation time by a factor of 70. With our model, the telescope's response function can be produced in less than 90 minutes of MacBook compute time (including training), and the model can be run for sky map production in under 90 seconds. For this work, we focus on the \sim 511 keV line and choose point sources with a similar number of events. However, future work can naturally extend ComptonINR to accept energy as an input for other energy bins and modify the simulation script to create point sources with significantly differing flux. Moreover, the NSIDE resolution is easily configurable in the codebase for higher resolution images. The code and dataset will be released shortly. We are currently training on response data that will be used for an upcoming Compton telescope mission, and preliminary tests suggest response generation with our method could achieve speedups of over 100. Fast, continuous PSF models similar to ComptonINR are expected to significantly accelerate telescope response generation, paving the path towards fast gamma-ray imaging on consumer hardware.

References

- De Angelis, A. et al. Gamma-ray astrophysics in the MeV range: The ASTROGAM concept and beyond. *Exper. Astron.*, 51(3):1225–1254, 2021. doi: 10.1007/s10686-021-09706-y.
- Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005.
- Grenier, I. A. and Harding, A. K. Gamma-ray pulsars: A gold mine. *Comptes Rendus. Physique*, 16(6–7):641–660, August 2015. ISSN 1878-1535. doi: 10.1016/j.crhy.2015. 08.013. URL http://dx.doi.org/10.1016/j.crhy.2015.08.013.
- Kierans, C., Takahashi, T., and Kanbach, G. Compton Telescopes for Gamma-Ray Astrophysics, pp. 1–72. Springer Nature Singapore, Singapore, 2022. ISBN 978-981-16-4544-0. doi: 10.1007/978-981-16-4544-0_46-1. URL https: //doi.org/10.1007/978-981-16-4544-0_ 46-1.
- Lucy, L. B. An iterative technique for the rectification of observed distributions. , 79:745, June 1974. doi: 10.1086/111605.
- Mohr, D. L., Wilson, W. J., and Freund, R. J. Chapter 13 special types of regression. In Mohr, D. L., Wilson, W. J., and Freund, R. J. (eds.), *Statistical Methods (Fourth Edition)*, pp. 623–649. Academic Press, fourth edition edition, 2022. ISBN 978-0-12-823043-5. doi: https://doi.org/10.1016/B978-0-12-823043-5.00013-8. URL https://www.sciencedirect.com/science/article/pii/B9780128230435000138.
- Oliveira, N. L., Lei, J., and Tibshirani, R. J. Unbiased test error estimation in the poisson means problem via coupled bootstrap techniques. *Electronic Journal of Statistics*, 19 (1):361–396, 2025.
- Richardson, W. H. Bayesian-based iterative method of image restoration*. J. Opt. Soc. Am., 62(1):55–59, Jan 1972. doi: 10.1364/JOSA.62.000055. URL https://opg.optica.org/abstract.cfm? URI=josa-62-1-55.
- Schmitt, J., Starck, J.-L., Fadili, J., Grenier, I., and Casandjian, J.-M. Poisson denoising on the sphere. In *Wavelets XIII*, volume 7446, pp. 397–408. SPIE, 2009.

- Sitzmann, V., Martel, J., Bergman, A., Lindell, D., and Wetzstein, G. Implicit neural representations with periodic activation functions. *Advances in neural information* processing systems, 33:7462–7473, 2020.
- Zoglauer, A., Andritschke, R., and Schopper, F. Megalib the medium energy gamma-ray astronomy library. *New Astronomy Reviews*, 50(7):629–632, 2006. ISSN 1387-6473. doi: https://doi.org/10.1016/j.newar.2006.06.049. URL https://www.sciencedirect.com/ science/article/pii/S1387647306000972. Astronomy with Radioactivities. V.
- Zoglauer, A. C. *First Light for the next Generation of Compton and Pair telescopes*. PhD thesis, Technische Universität München, 2005.

A. Appendix

A.1. Cartesian Data Space vs. Spherical Data Space



Figure 5. Cartesian Data Space vs. Spherical Data Space, for the same simulation file. (a) Approximated Cartesian view of the data space, with the PSF cone for one source location. For clear illustration of sample bins, (θ, ϕ) were rounded down to the nearest 3° bin, and ψ was rounded down to the nearest 10° bin. (b) Spherical re-parameterization used for HEALPix binning; the same events now lie on concentric rings about $(\theta_{src}, \phi_{src})$. The cross sections for the [10°, 20) ψ bin and the [80, 90)° ψ bin are shown.

A.2. Removal of Periodic Encodings



Figure 6. Sample data space with ablation of periodic encodings. Removing \sin / \cos angle encodings introduces artifacts at the -180°/180° ϕ boundary, highlighting the need for periodic input encodings.

A.3. Isolated Hyperparameter Sweeps



Figure 7. Isolated hyperparameter sweeps. Top-left: the best Poisson-NLL validation loss across 40 epochs does not vary with a consistent trend as the number of hidden layers changes. Top-right: the optimal batch size in this sweep appears to be 1024. Bottom-left and bottom-right: as hidden dimension size and the number of training files increase, the PoissonNLL validation loss decreases toward an asymptote.

We performed isolated hyperparameter sweeps across the size of the hidden dimension, number of simulation files used to train the model, batch size, and the number of hidden layers. It is important to highlight that each simulation file contained 8 ψ bins, each of which contains 3,072 HEALPix pixels, creating a total of 24,576 training points per file. As

the hidden dimension size and number of training files increased, the best validation loss achieved by the model decreased asymptotically.

A.4. Poisson Deviance Formula

The following formula is defined as the Poisson deviance, normalized by the degrees of freedom ν . The Poisson deviance measures how well a model's predictions \hat{y}_i match ground truth counts y_i from a Poisson distribution. The equation is adapted from (Oliveira et al., 2025) to add the normalization component and modify notation for clarity.

$$D(y,\hat{y}) = \frac{2}{\nu} \sum_{i=1}^{n} \left(y_i \log \frac{y_i}{\hat{y}_i} + \hat{y}_i - y_i \right)$$
(2)

A.5. Scatter Angle Formula

The following formula defines the Compton scatter angle ψ for a gamma ray. The energy of the scattered gamma ray is denoted by E_2 , while E_1 represents the energy of the energized electron with initial rest mass m_e . The total energy $E_1 + E_2$ corresponds to the initial photon energy. This equation is adapted with modifications from (Kierans et al., 2022) to reflect the notation for the scatter angle chosen in our work.

$$\cos\psi = 1 - \frac{m_e c^2}{E_2} + \frac{m_e c^2}{E_1 + E_2},\tag{3}$$

A.6. Richardson-Lucy Results for Coarse NSIDE=16

We performed Richardson-Lucy as in Section 3.1 but with a coarser sky map resolution of NSIDE=16. While source detection metrics of true positives generally improve, angular error metrics are worsened. This is expected due to the lower resolution.

Source Detection Metrics	Mean	Std. Dev.
Precision	0.9829	0.0441
Recall (TPR)	0.9700	0.0671
F1-Score	0.9761	0.0551
False Positive Rate (FPR)	0.0001	0.0001
Angular Error Metrics		
Median Angular Error	1.5570°	_
25th Percentile AE	1.1341°	_
75th Percentile AE	1.9318°	-
Inference Time Metrics		
Run on MacBook M1 Pro, 10-c	ore CPU, 3	2GB RAM
Response Cube Generation	407s	_
\Rightarrow One-time, then cached		
Single Sky Image Creation (s)	21.8s	1.3s

Table 2. Sky Map Creation Metrics