# Neural Likelihood Surrogates for Fast Gamma-Ray Burst Localization with Compton Telescopes

**Srinadh Reddy Bhavanam, Marco Ajello, Mark Leising, Dieter Hartmann, Yong Sheng, Christopher M Karwin**

Department of Physics and Astronomy, Clemson University, Clemson, SC 29634, USA
{srinadb, majello, lmark, hdieter, sheng2, ckarwin}@clemson.edu

## Abstract

Gamma-ray Compton telescopes do not directly image a source. Each detected event constrains the incident photon direction to a cone in the sky, and a gamma-ray burst (GRB) is localized from where many such noisy cones intersect. This makes localization a geometric inverse problem with a sharply peaked, computationally intensive all-sky likelihood, a bottleneck for the rapid, near-real-time response that multi-messenger follow-up demands. Targeting NASA's COSI mission, we introduce a novel approach that treats localization as fitting a per-burst neural likelihood surrogate. From the angular event triplets $(\chi, \psi, \phi)$, we define an unsupervised cone-residual log-likelihood on the sphere, fit a compact implicit neural representation (INR) to samples of this score, and select the final direction by exact local refinement seeded from the INR map, offering a new way to accelerate physically defined score surfaces. On 1,250 bright GRB simulations, the INR localizer reaches mean and median angular deviations of $0.255^\circ$ and $0.235^\circ$, versus $0.390^\circ$ and $0.380^\circ$ for the standard analysis algorithm.

## 1 Introduction

Fast and accurate MeV GRB localization is a bottleneck for multi-messenger astronomy. Inspiraling binary neutron stars emit gravitational waves; after merger, relativistic outflows can produce short GRBs when the jet geometry is favorable. GW170817/GRB 170817A provided the first unambiguous joint gravitational-wave and gamma-ray observation of a binary-neutron-star merger, establishing the observational link between compact-object coalescence and short GRBs [Abbott et al., 2017a,b, Goldstein et al., 2017]. It also exposed a practical limitation: the prompt gamma-ray detection arrived quickly, but the gamma-ray sky localization was broad, with Fermi-GBM credible regions of order hundreds to about a thousand square degrees [Goldstein et al., 2017]. Such localizations force optical, infrared, X-ray, and radio facilities to search large sky areas, delaying counterpart identification and consuming scarce telescope time. Future missions therefore need localization algorithms that are both accurate enough to focus follow-up and fast enough to support near-real-time, and potentially on-board, transient response.

The Compton Spectrometer and Imager (COSI) is a NASA Small Explorer soft gamma-ray mission planned for launch in 2027, designed to survey the sky over $0.2$–$5$ MeV with wide-field imaging, spectroscopy, and polarimetry [Tomsick et al., 2023, COSI Collaboration, 2026, NASA, 2025]. Among its transient-science goals, COSI will detect and localize short GRBs so that follow-up telescopes can respond to candidate multi-messenger events. This motivates localization methods that exploit the structure of Compton events without discarding their physics. Unlike conventional astronomical images, COSI observations arrive as sparse, irregular event lists: each photon contributes a noisy event circle rather than a pixel intensity. The measured interaction geometry and deposited energies determine a Compton scatter angle, constraining the unknown incoming direction to a circle on the sky centered on the reconstructed scattered-photon direction, with opening angle equal to

Conference on Physics and AI at Stanford University (PAI 2026).

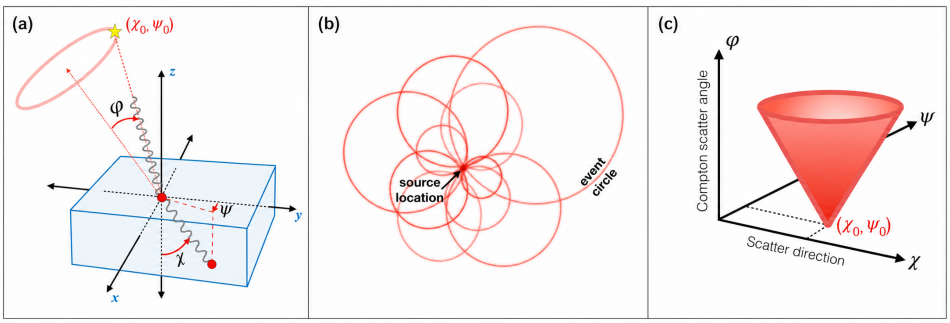


Figure 1: Compton-event geometry and Compton Data Space (CDS) representation. (a) Each event is described by an angular triplet $(\chi, \psi, \phi)$, where $\chi$ and $\psi$ are the polar and azimuthal angles of the scattered-photon direction, and $\phi$ is the Compton scatter angle. (b) Multiple event circles overlap near the source location, enabling gamma-ray source localization. (c) In Compton Data Space, point-source events populate a cone-like structure with apex at the source direction $(\chi_0, \psi_0)$.

that scatter angle. A GRB location is inferred from the common intersection region of many such noisy circles. More details can be illustrated from Fig 1. Brute-force likelihood evaluation or iterative imaging over the full sky can become computationally expensive, especially when rapid follow-up or on-board triggering is desired. A useful AI model should therefore not replace Compton geometry; instead, it should learn a fast, continuous surrogate of the physically defined localization score.

Classical Compton imaging methods, including backprojection and list-mode maximum-likelihood expectation maximization (LM-MLEM), exploit the fact that each valid Compton event constrains the unknown source direction to a cone, or equivalently to an event circle on the celestial sphere [Wilderman et al., 1998, Parajuli et al., 2022]. These methods are physically interpretable and do not require supervised labels, but dense all-sky reconstruction can be computationally expensive when many likelihood evaluations or iterative updates are required. Recent machine-learning approaches have explored complementary Compton-camera representations: ComptonINR learns a forward PSF/response surrogate, ComptonNet performs direct reconstruction from measured events, ComptonUNet combines raw-data and image-based information for noisy low-statistic GRB localization, and COSI shield-light-curve or HEALPix methods target satellite-level localization [Kotamraju and Zoglauer, 2025, Sato et al., 2024, 2026, Parmiggiani et al., 2026b,a]. Our setting is different: we operate directly on the raw reconstructed angular triplets of individual Compton events, as illustrated in Fig. 1, and use learning only as a per-burst inverse surrogate of an explicit Compton-cone likelihood rather than as a replacement for the physical imaging model.

Our contribution is a localization workflow that combines exact Compton-cone scoring with implicit neural representations (INRs), continuous neural functions popularized in computer vision and scientific machine learning [Sitzmann et al., 2020, Tancik et al., 2020, Mildenhall et al., 2020]. Instead of training a global network across many labelled GRBs, we train a small per-GRB INR on exact likelihood samples from that GRB. The approach is unsupervised with respect to source labels during inference, preserves a final exact refinement step, and produces dense probability maps suitable for localization diagnostics.

## 2 Dataset

The data are simulated with MEGAlib using the COSI Data Challenge 2 analysis context and the `massmodel-cosi-smex-v12` satellite mass model [Zoglauer et al., 2006, 2008, COSItools, 2024, COSI Collaboration, 2025]. The dataset contains 1,250 simulated GRB-like point-source realizations observed by the COSI Compton camera model. The injected prompt-emission spectrum is modelled with a Band function over 100–10,000 keV ($\alpha = -0.5$, $\beta = -2.3$, $E_{\text{break}} = 490$ keV). This study intentionally targets a favorable signal-rich regime: with 4,673 to 7,605 source photons, while the background is simulated for a one-second interval and contributes 15 to 53 photons per GRB, yielding detections with a signal-to-noise ratio of $\sim 70$. This isolates the raw localization problem and validates the INR surrogate before moving to realistic low-signal, high-background regimes. The method receives only unbinned event-level angular quantities; ground-truth coordinates are used only for evaluation, not for fitting the per-burst localizer.

## 3 Methodology

Our methodology proceeds in four stages: represent each GRB as unbinned Compton angular measurements in Compton Data Space (CDS), define a physics-based cone-consistency score, fit a per-burst INR surrogate over the sphere, and use the INR map to seed local refinement. This design keeps the Compton geometry explicit while using the neural network only as a fast continuous approximation to the localization surface.

**Event representation and Compton-event geometry:** For a burst with $N$ events, the raw input is $\mathcal{D} = \{(\chi_i, \psi_i, \phi_i)\}_{i=1}^{N}$, where $\phi_i$ is the Compton scatter angle of the first interaction (inferred from energy deposits via the Compton relation) and $(\chi_i, \psi_i)$ are the polar and azimuthal angles of the reconstructed scattered-photon direction. We focus on this angular triplet for localization.

Geometrically, event $i$ constrains the incoming direction to an event circle centered on $(\chi_i, \psi_i)$ with opening angle $\phi_i$ (Fig. 1). Let $v_i \in \mathbb{R}^3$ be the unit event-axis vector and $s \in \mathbb{R}^3$ a candidate source direction; for an ideal noiseless event, $\arccos(s^\top v_i) = \phi_i$. A GRB location is thus inferred from the common intersection of many noisy event circles, without rasterizing the event list into an image proxy.

**Compton-cone consistency score:** A candidate direction $s$ is consistent with event $i$ when its angle from $v_i$ matches the Compton scatter angle $\phi_i$. Since events are noisy, we define a signed cone residual and likelihood as

$$r_i(s) = \arccos\big(\mathrm{clip}(s^\top v_i, -1, 1)\big) - \phi_i, \qquad p_i(s \mid \phi_i, v_i) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{r_i(s)^2}{2\sigma^2}\right], \tag{1}$$

where $\sigma$ is a shared angular kernel width. The burst-level score is the mean log-likelihood

$$\mathcal{L}(s; \mathcal{D}) = \frac{1}{N} \sum_{i=1}^{N} \log p_i(s \mid \phi_i, v_i). \tag{2}$$

This score is label-free at inference time; the simulated truth is used only for evaluation. In this first controlled study, we do not include an explicit background term; instead, we isolate the raw angular Compton-cone localization problem in a favorable signal-rich regime.

**INR surrogate over the sphere:** Although Eq. (2) can be evaluated anywhere on the sky, dense evaluation is costly. We therefore fit a per-burst INR $f_\psi : \mathbb{S}^2 \to \mathbb{R}$ to approximate $\mathcal{L}(s; \mathcal{D})$. The input is a Fourier-feature positional encoding $\gamma(s)$ of the Cartesian sky vector, fed to a compact MLP with `tanh` activations [Tancik et al., 2020, Mildenhall et al., 2020]; this is a positional-encoding network with standard activations, not a sinusoidal-activation SIREN. The INR is optimized with Adam [Kingma and Ba, 2015]followed by L-BFGS refinement [Moritz et al., 2016], minimizing the unweighted squared error

$$\min_{\psi} \sum_{j=1}^{M} \left[f_\psi(\gamma(s_j)) - \mathcal{L}(s_j; \mathcal{D})\right]^2. \tag{3}$$

The exact-score samples $\{s_j\}_{j=1}^{M}$ combine a coarse sky grid, separated high-scoring coarse peaks, random global samples, and local patches around the top-$K$ peaks. This provides the INR with global context and high-resolution samples near plausible sources, reducing the risk of fitting a single narrow mode while missing competing peaks.

**Prediction, probability maps, and refinement:** After fitting, the INR is evaluated on a dense grid. Because $\mathcal{L}$ is a *mean* log-likelihood, naively exponentiating it gives $P(\mathcal{D} \mid s)^{1/N}$, a map over-broadened by an effective temperature $N$. We therefore form calibrated maps from the *total* log-likelihood, $P(\mathcal{D} \mid s) \propto \exp\big(N\, \mathcal{L}(s; \mathcal{D})\big)$, normalized over the grid; this leaves the localization argmax unchanged but recovers the correct probabilistic width. The highest-probability locations seed a final local search with Eq. (2), so the reported direction is the candidate with the best exact Compton-cone score in the refined neighbourhood, not the raw network maximum. The primary metric is angular separation $\Delta(s, s_\star) = \arccos(s^\top s_\star)$ to the simulated source $s_\star$, compared against the standard LM-MLEM separation.

## 4 Experimental Design and Results

We evaluate the proposed INR localization against the standard LM-MLEM baseline on 1,250 bright MEGAlib/COSI DC2 simulated GRBs. We also include a sampled-exact control that uses the same exact score samples and final refinement stage, but omits INR fitting. Ground truth is used only for evaluation. We use a fixed hyperparameter configuration: a $5^\circ$ coarse search grid, $1^\circ$ probability-map resolution, cone width $\sigma = 0.5^\circ$, hidden layers $(64, 64)$, four positional-encoding frequencies, 768 random global samples, top-$K = 3$ local seeds, a $12^\circ$ patch radius, 200 Adam iterations, learning rate $10^{-3}$, and random seed 12345. These values define our chosen operating point rather than a fully optimized configuration. The primary metric is angular separation between predicted and injected source directions. We report mean, median, best-case, and worst-case angular deviations, and inspect INR maps for compactness, entropy, and secondary modes.

Table 1 and Figure 2 summarize the results. The INR and the sampled-exact control both improve over LM-MLEM, and their near-identical errors show that the accuracy gain in this signal-rich regime comes from exact score sampling plus refinement, not from the surrogate.The INR's contribution is speed: it matches that accuracy while running 5.6$\times$ faster than the sampled-exact control.

Table 1: Quantitative and timing comparison of the INR localizer, the sampled-exact control, and LM-MLEM on 1,250 simulated GRBs. Angular deviations are in degrees, average time is per-burst wall-clock in seconds. LM-MLEM runs in a separate MEGAlib pipeline on different hardware, so a comparable time is not available.

| Method | Mean | Median | Best | Worst | Avg. Time (S) |
|---|---|---|---|---|---|
| INR | **0.255** | **0.235** | **0.002** | **0.899** | 2.98 |
| Sampled-exact control | **0.254** | **0.233** | **0.002** | **0.888** | 16.8 |
| List-mode MLEM baseline | 0.390 | 0.380 | 0.017 | 1.212 | - |

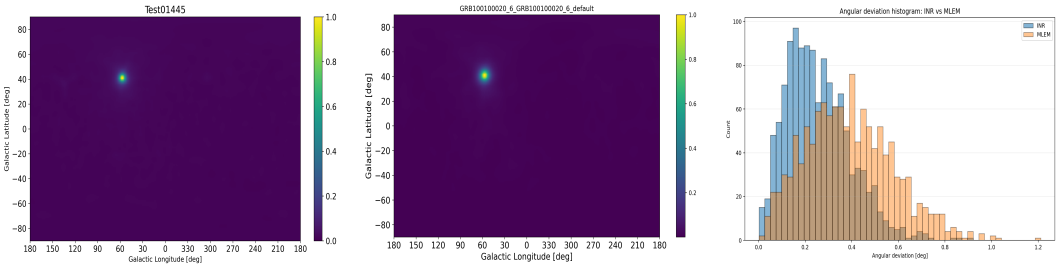


(a) INR map (b) LM-MLEM map (c) Angular-error histogram

Figure 2: **INR and LM-MLEM comparison.** Panels (a) and (b) show maps for the same burst with 6,403 signal photons and 35 background photons. The true source is $(57.373^\circ, 40.872^\circ)$. INR predicts $(57.375^\circ, 40.875^\circ)$ with AD $0.002^\circ$, while LM-MLEM predicts $(57.50^\circ, 40.5^\circ)$ with AD $0.385^\circ$. Panel (c) shows the angular-error distribution.

As shown in Fig. 2, both (a) INR and (b) LM-MLEM produce compact localization peaks near the injected source, and panel (c) shows the INR histogram shifted toward lower errors. The sampled-exact control matches INR's accuracy ($0.254/0.233^\circ$ vs. $0.255/0.235^\circ$ mean/median) but is slower (16.8 s vs. 2.98 s per burst): both share score sampling ($\sim 1.5$ s) and refinement ($\sim 0.07$ s), differing only in dense map evaluation ($\sim$ 12–15 s exact vs. $\sim 0.07$ s INR).

## 5 Critical Assessment and Conclusion

We present a neural likelihood surrogate for raw-event GRB localization in COSI-like Compton simulations: event-cone geometry defines an unsupervised likelihood, a per-burst INR learns a continuous score surface, and exact local refinement selects the final direction. This development reflects our collective effort to improve localization techniques, encouraging a sense of shared progress. The accuracy gain over LM-MLEM comes from exact sampling plus refinement, while the INR's benefit is computational, accelerating dense score-map evaluation. The method is best viewed as high-quality work in progress, not mission-level performance. Future validation should add confidence intervals, matched INR/LM-MLEM event selections, calibration-only kernel tuning, background-rich and extended-source tests, and calibrated containment regions.

## References


B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Physical Review Letters*, 119(16):161101, 2017a. doi: 10.1103/PhysRevLett.119.161101.

B. P. Abbott, R. Abbott, T. D. Abbott, et al. Multi-messenger observations of a binary neutron star merger. *The Astrophysical Journal Letters*, 848(2):L12, 2017b. doi: 10.3847/2041-8213/aa91c9.

COSI Collaboration. COSI data challenges. Zenodo, 2025. URL `https://zenodo.org/records/15126188`. doi: 10.5281/zenodo.15126188.

COSI Collaboration. COSI: The compton spectrometer and imager. Mission website, 2026. URL `https://cosi.ssl.berkeley.edu/`. Accessed May 2026.

COSItools. COSI data challenge 2. GitHub repository, 2024. URL `https://github.com/cositools/cosi-data-challenge-2`.

A. Goldstein, P. Veres, E. Burns, et al. An ordinary short gamma-ray burst with extraordinary implications: Fermi-GBM detection of GRB 170817a. *The Astrophysical Journal Letters*, 848(2):L14, 2017. doi: 10.3847/2041-8213/aa8f41.

D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

A. Kotamraju and A. Zoglauer. ComptonINR: Implicit neural representations for fast modeling of compton telescope point spread functions. In *ICML 2025 Workshop on Machine Learning for Astrophysics*, 2025. URL `https://ml4astro.github.io/icml2025/assets/camera_ready/46_ComptonINR_Implicit_Neural_.pdf`.

B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020.

P. Moritz, R. Nishihara, and M. Jordan. A linearly-convergent stochastic l-bfgs algorithm. In *Artificial intelligence and statistics*, pages 249–258. PMLR, 2016.

NASA. COSI: Compton spectrometer and imager. Mission webpage, 2025. URL `https://science.nasa.gov/mission/cosi/`. Accessed May 2026.

R. K. Parajuli, M. Sakai, R. Parajuli, and M. Tashiro. Development and applications of compton camera–a review. *Sensors*, 22(19):7374, 2022. doi: 10.3390/s22197374.

N. Parmiggiani, A. Bulgarelli, G. Panebianco, E. Burns, E. Neights, V. Fioretti, I. Martinez-Castellanos, L. Castaldini, A. Ciabattoni, A. Di Piano, R. Falco, S. Gallego, G. Mustafa, P. Patel, A. Rizzo, E. A. Wulf, D. H. Hartmann, C. A. Kierans, J. A. Tomsick, and A. Zoglauer. Localization and confidence region estimation of short GRBs with the COSI BGO shield using a HEALPix-based deep learning approach, 2026a.

N. Parmiggiani, A. Bulgarelli, G. Panebianco, et al. COSI short gamma-ray burst localization using BGO shield data. *The Astrophysical Journal*, 997(2):135, 2026b. doi: 10.3847/1538-4357/ae25fc.

S. Sato, K. S. Tanaka, and J. Kataoka. ComptonNet: A direct reconstruction model for compton camera. *Applied Physics Letters*, 124(25):253702, 2024. doi: 10.1063/5.0213950.

S. Sato, K. Tanaka, S. Ogasawara, K. Yamamoto, K. Murasaki, R. Tanida, and J. Kataoka. ComptonUNet: A deep learning model for GRB localization with compton cameras under noisy and low-statistic conditions. *The Astrophysical Journal*, 999(2):231, 2026. doi: 10.3847/1538-4357/ae422f.

V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, 2020.

M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, 2020.

J. A. Tomsick, S. E. Boggs, A. Zoglauer, et al. The compton spectrometer and imager. In *Proceedings of Science, 38th International Cosmic Ray Conference*, page 745, 2023. doi: 10.48550/arXiv.2308.12362.

S. J. Wilderman, N. H. Clinthorne, J. A. Fessler, and W. L. Rogers. List-mode maximum likelihood reconstruction of compton scatter camera images in nuclear medicine. In *1998 IEEE Nuclear Science Symposium and Medical Imaging Conference Record*, 1998.

A. Zoglauer, R. Andritschke, and F. Schopper. MEGAlib: The medium energy gamma-ray astronomy library. *New Astronomy Reviews*, 50(7–8):629–632, 2006.

A. Zoglauer, R. Andritschke, S. E. Boggs, F. Schopper, G. Weidenspointner, and C. B. Wunderer. MEGAlib: Simulation and data analysis for low-to-medium-energy gamma-ray telescopes. In *Proceedings of SPIE*, volume 7011, page 70113F, 2008.