

---

# Multi-Wavelength Analysis of Kilonova Associated with GRB 230307A: Accelerated Parameter Estimation and Model Selection Through Likelihood-Free Inference

---

**Phelipe Darc\***

Centro Brasileiro de Pesquisas Físicas (CBPF)  
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.

**Clecio R. Bom**

Centro Brasileiro de Pesquisas Físicas (CBPF)  
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.  
Centro Federal de Educação Tecnológica Celso Suckow da Fonseca (CEFET-RJ),  
Av. Maracanã, 229, Maracanã, Rio de Janeiro, Brazil.

**Gabriel S. M. Teixeira**

Centro Brasileiro de Pesquisas Físicas (CBPF)  
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.

**Charles Kilpatrick**

Center for Interdisciplinary Exploration and Research in Astro- physics (CIERA)  
and Department of Physics and Astronomy  
North-western University, Evanston, IL 60208, USA.

**Nora F. Sherman**

Boston University  
750 Commonwealth Avenue Boston, MA, USA.

**Marcelo P. Albuquerque**

Centro Brasileiro de Pesquisas Físicas (CBPF)  
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.

**Paulo Russano**

Centro Brasileiro de Pesquisas Físicas (CBPF)  
Rua Xavier Sigaud, 150, Urca, Rio de Janeiro, Brazil.

## Abstract

The mergers of binary compact objects are of central interest to several areas of astrophysics, including as the progenitors of short gamma-ray bursts (GRBs). Dozens of GRBs have been confidently associated with rapidly-decaying optical transients (“afterglows”) and more recently with late-time emission, such as from a kilonovae. Traditional methods for modeling these phenomena are computationally expensive for comparison and parameter inference because of the high-dimension parameter space. We propose using Simulation-Based Inference (SBI) as a fast

and scalable alternative to conventional likelihood-based approaches. Preliminary results using SBI to fit multi-wavelength light curves of GRB 230307A across different emission models, highlight its efficiency in managing high-dimensional parameter space and show that SBI yields posterior distributions consistent with those from likelihood-based methods, while significantly reducing computational time from 110 hours to just 5 hours. This is a significant saving that can allow for more exhaustive model comparisons in future works.

## 1 Introduction

A gamma-ray burst (GRB) is a highly energetic ( $> 10^{50}$  erg) explosion characterized by a brief (seconds-long) but intense burst of gamma rays. GRBs are classified into two main categories: long-duration GRBs (LGRBs), which generally result from the collapse of massive stars, and short-duration GRBs (SGRBs), which are associated with the mergers of compact objects ([1]). Although the traditional classification typically divides the GRBs population at a gamma-ray duration of 2 s, recent observations have revealed intriguing exceptions to this. For instance, the LGRBs 211211A and 230307A have been linked to a kilonovae, indicating that such bursts may arise from compact object mergers despite their long duration ([2, 3]). Therefore, the merger of two compact objects, including neutron star/black hole, binary neutron star, and white dwarf/neutron star merger can give rise to a GRB and a short-lived thermal transient ([4, 5]). In this context, Bayesian inference and model selection applied to observational data can offer valuable insights into the GRB emission mechanisms and progenitor system.

Traditional Bayesian inference methods, such as Markov Chain Monte Carlo (MCMC) and nested sampling, have been employed to fit physical models to GRB observations. These methods, however, can, depending on the complexity of the model, take between tens of minutes to several days on a single CPU to sample the posterior, particularly when dealing with high-dimensional parameter spaces ([6]). The rapid advancements in computational resources have significantly expanded the availability of models for kilonovae, supernovae, and other thermal emission mechanisms ([7, 8]), which are used to model the additional optical excess beyond standard GRB afterglows models. As these models become more detailed, the increasing number of parameters creates challenges for efficient inference and model selection due to high computational costs to compute the likelihood. To address these challenges, we investigate the application of Simulation-Based Inference (SBI; [9]), a likelihood-free approach, to rapidly estimate the posterior distribution  $p(\theta|x)$  of GRB physical properties in a high-dimensional parameter space. Here,  $\theta$  denotes the model parameters for the combined afterglow and kilonova models and  $x$  represents the observed light curves (LCs).

SBI has gained traction in various fields, such as galaxy evolution ([10]), gravitational waves ([11]), and cosmology ([12]), but its application to GRB science remains underexplored. While recent work has applied amortized inference to kilonova spectral data ([13]) and likelihood-free inference to kilonova light curves ([14]), these studies have been limited to optical kilonova observations with low number of free parameters ( $N < 10$ ).

Here, we present the first application of SBI to model both afterglow and kilonova emissions, incorporating multi-wavelength data from both infrared and optical bands. We further compute model evidence within the SBI framework ([15]), leveraging a harmonic mean estimator ([16]) to enable efficient Bayesian model comparison across multiple emission models. We validate our approach using data from the GRB 230307A event, comparing it to traditional inference techniques.

This paper is structured as follows: Section 2 describes the kilonova and GRB forward shock models; Section 3 outlines the SBI methodology used for parameter estimation and Bayesian model selection; Section 4 provides details on the training procedure and compares the results with traditional methods for the GRB230307 event; Finally, Section 5 discusses the implications of the study and potential directions for future work.

## 2 Multi-wavelength afterglow and kilonovae modelling

To model the multi-wavelength counterpart of GRB230307A, we employ a combined afterglow and kilonova model. The first component describes the non-thermal synchrotron emission from the GRB's

forward shock interacting (FS) with the surrounding medium ([17]). To model the FS emission, we used the Python package `afterglowpy` ([18]), which includes the following free parameters: the isotropic-equivalent kinetic energy  $E_0$ , the circumburst density  $n_0$ , the fraction of burst kinetic energy in magnetic fields ( $\epsilon_B$ ) and in electrons ( $\epsilon_e$ ), the power-law slope  $p$  of the electron energy distribution, the jet core’s opening angle  $\theta_c$ , and the electron participation fraction  $\xi_N$ . The second component is an isotropic kilonova model from [19], which assumes a gray opacity and models the spectrum using a simple blackbody function. The free parameters for this model are the ejecta mass ( $M_{ej}$ ), the minimum velocity of the ejecta ( $v_{ej}$ ), the velocity index ( $\beta_v$ ), opacity ( $\kappa$ ), and electron fraction ( $Y_e$ ).

We fit the multi-wavelength data using three distinct models: (1) a forward shock model (FS), (2) a forward shock combined with a one-component kilonova (FS+1KNE), and (3) a forward shock plus a two-component kilonova (FS+2KNE). These analytical models serve as "simulators" in our framework, with uniform priors chosen for all parameters following the values reported in [5].

### 3 Methodology and Data




Figure 1: **Afterglow and kilonova models fit to observations of the GRB230307A counterpart:** The Light curves for each model were generated using the median values of the inferred posterior distributions. The solid thin lines are generated from 100 parameter sets randomly sampled from the FS+2KNE inferred posterior distribution. It was adopted an offset of -10.

early-phase and late-phase (>7 days) segments. Each segment was passed through a distinct fully connected neural network with 3 layers, with outputs from both phases combined in a final layer with 128 neurons. We evaluated various embedding network architectures and density estimations, finding that this configuration provided the most precise constraints.

**Bayesian Model Selection:** As shown by [15], Bayesian model selection is a natural extension of parameter inference in the context of SBI. In addition to inferring model parameters, SBI can be

**Parameter Inference:** For a rapid parameter inference, we employ Truncated Sequential Neural Posterior Estimation (TSNPE; [20]) with a Mixture Density Network (MDN; [21]) as the density estimator and an embedding network as the feature extractor. TSNPE is a modified version of SNPE ([22]), which is part of the family of SBI techniques. SNPE employs a multi-round adaptive inference strategy, where model parameters are sampled from an updated proposal distribution—an approximate posterior—generated in previous inference rounds (illustrated by Fig.2). This adaptive sampling improves posterior accuracy for a specific observation, which in this case corresponds to the GRB230307A event. Truncated SNPE introduces an additional truncation step that rejects parameter samples outside the current posterior approximation. This ensures only parameters within the posterior’s support are simulated, increasing inference efficiency by concentrating computational resources on regions more likely to yield an accurate posterior.

The MDN is used to model the posterior distribution as a mixture of Gaussian distributions. Although we tested other density estimators such as Normalizing Spline Flow and Masked Autoregressive Flow ([23]), MDN demonstrates superior precision. To further improve the inference accuracy, an embedding network is used to extract features from the observational data. The main purpose of the embedding network is to transform the raw observational data into a representation that is more informative for the inference task. Our embedding network processed the input observations by separating them into

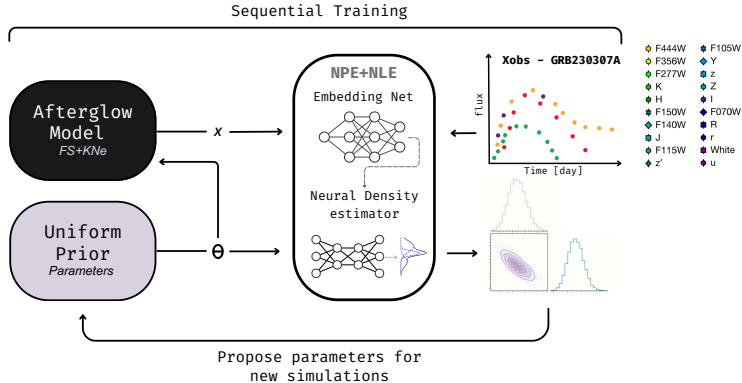


Figure 2: Inference pipeline using SNPE and SNLE — Sequential inference enables rapid convergence by focusing on a more informative region of the parameter space. In the training phase, a neural-network-based density estimator (MDN) learns the probabilistic relationship between the model’s parameters ( $\theta$ ) and “simulated” light curve ( $X$ ). Then, in the inference phase, the trained density estimator takes the observed real data ( $X_{obs}$ ; GRB 230307A) as input and infers the posterior distribution of the parameters, which is used as prior to the next round.

used to compare different models by computing the model evidence (also known as the marginal likelihood). The model evidence allows one to compute the Bayes factor, a ratio of the evidence of two competing models, which provides a way to assign relative probabilities to models given the observed data. Traditionally, fast and accurate estimation of Bayesian model evidence in a likelihood-based method poses significant computational challenges. This process requires solving a high-dimensional integral, which can become infeasible for parameter spaces of even moderate dimensionality.

In order to compute the model evidence, we adopt the learned harmonic mean estimator (LHME; [16]), which uses samples drawn directly from the surrogate posterior produced by the TSNPE. The application of LHME, however, requires evaluating the likelihood function at the sampled parameter positions. In many SBI settings, the likelihood function is either intractable or computationally prohibitive to evaluate directly. To overcome this, we trained a Neural Likelihood Estimation (NLE; [24]) to provide a surrogate likelihood function. Therefore, this approach involves training two density estimators at the same time, NLE and NPE, under the identical sequential training setup. For the LHME, we adopted the Rational quadratic spline normalizing flow model, with 80% of the samples being used to train the NF for 200 epochs.

**Data:** In this analysis, we utilize the dataset detailed in [5]. GRB230307A was detected by the Fermi Gamma-ray Burst Monitor on 7 March 2023. Its dataset includes extensive follow-up observations across multiple wavelengths: optical and near-infrared data from the Gemini South telescope and the Very Large Telescope, X-ray data from the Swift X-ray Telescope and Chandra X-ray Observatory, and radio data from the Australia Telescope Compact Array and MeerKAT. The James Webb Space Telescope (JWST) observations were initiated on 5 April 2023, with six-color imaging performed with the Near Infrared Camera (NIRCam) at +28.4 days after the GRB. For our current analysis, we focused exclusively on the optical and near-infrared observations. Future work will extend our fit to incorporate radio and X-ray observations as well.

## 4 Results and Discussion

We compare the performance of the FS+2KNE model against traditional MCMC method using both qualitative and quantitative analyses. The FS+2KNE model is chosen as the fiducial model due to its capacity to incorporate additional kilonova emission components and the complexity introduced by its 17 free parameters. In addition, we calculate the Bayesian evidence for the three models (FS, FS+1KNE, and FS+2KNE) using the learned harmonic mean estimator. We anticipate that FS+2KNE will yield the highest evidence, indicating a better fit to the observational data due to its ability to capture more complex physical processes ([5]).

**Training Setup:** We trained the TSNPE and SNLE models over 10 rounds with a minimum of 600 epochs per round, each involving 10,000 simulations. We used the posterior distribution from the NPE method as the proposal distribution because it's faster to sample directly from the NPE than from the NLE. The total training time for the FS+2KNE model was under 5 hours using 80 CPUs. In contrast, an MCMC approach was estimated to require over 110 hours with comparable computational resources. For comparison, we initialized an MCMC run using the most probable parameter values obtained from the TSNPE method and executed it for 5,000 steps with 100 walkers (referred to as MCMC FS+2KNE), which took 52 hours to complete. Additionally, we trained a TSNPE model under a low simulation budget of 100 simulations for each round, achieving comparable precision for most parameters in under 30 minutes. This improvement in performance could allow for an exhaustive search of the parameter space and extensive model comparisons within a feasible timeframe, even with a limited CPU budget.

**Parameter Retrieval:** Fig.3(a) shows that the inferred posterior distributions for the FS+2KNE model align well with both the likelihood-based method from [5] and MCMC FS+2KNE approach, with median values falling within the one sigma interval. Similar consistency is observed in the FS+1KNE model, with median values agreeing within one sigma. Notably, we observe larger uncertainties in parameters such as  $E_0$  and  $p$ , which may stem from our use of uniform priors, in contrast to the truncated-Gaussian prior adopted for the electron index  $p$  in [5].

As shown in Fig. 3(b), SBI effectively constrains key parameters like  $\epsilon_B$ ,  $\epsilon_e$ ,  $\beta_v$  and  $Y_e$ , performing similarly or better than traditional methods. Despite some larger uncertainties in parameters such as  $E_0$  and  $M_{ej}$ , SBI remains accurate while significantly reducing computational time – 28 minutes, for the low budget inference, compared to over 110 hours with MCMC.

**Bayesian Model Comparison:** Using samples from TSNPE and the surrogate likelihood from SNLE, we estimate model evidences with the LHME methodology [15]. The evidence increases with the number of model components, aligning with theoretical expectations due to additional kilonova emission in FS+2KNE. Fig. 1 visually confirms these results. It shows LCs computed from the most probable parameters for each model, compared with GRB230703A data across ten photometric bands. The FS+2KNE model provides the best fit to the observational data, outperforming the FS and FS+1KNE models.

**Discussion:** Despite being a preliminary result, this work showcases the efficiency and scalability of SBI for modeling GRB afterglows and kilonovae in high-dimensional parameter spaces. By generating posterior distributions and enabling model comparisons with observational data in less than 5 hours, SBI offers a scalable alternative to traditional inference techniques, allowing for broader model coverage and a more comprehensive characterization of the progenitor system, even with a limited CPU budget. Future directions include expanding the analysis to consider more complex kilonova models, such as those involving three-component emission profiles ([8]), and applying this approach to additional GRB events [25]. Moreover, we plan to utilize observations from X-ray and radio to further enhance model precision and to explore alternative inference approaches, including Bayesian neural networks (BNN), to improve performance and scalability. The code used to perform the experiments presented in this paper is openly available in our GitHub repository: NeurIPS-2024-GRB-SBI

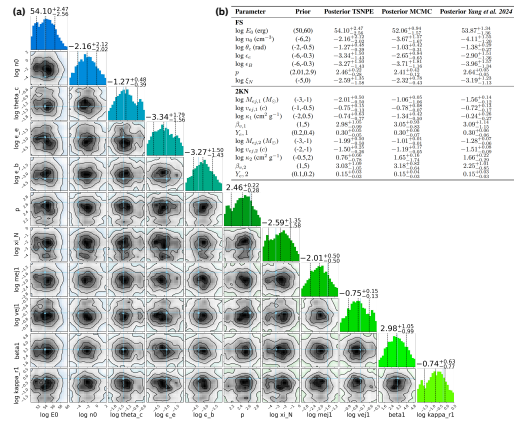


Figure 3: **Results for a forward shock plus two-component kilonova model:** (a) Corner plot of the inferred GRB parameters for the GRB230703A at 10%, 32%, 68%, and 95% confidence intervals. The median and 68% confidence are shown by vertical solid and dashed lines, respectively, and reported above each column. The results from [5] (blue solid line) are also shown for comparison. Due to the size of the parameter space, we only display the first two components of the FS+2KNE model, the full image is displayed in the Appendix section. (b) The prior bounds and inferred posterior medians for the FS+2KNE model.

## References

- [1] Eleonora Troja. “Eighteen Years of Kilonova Discoveries with Swift”. In: *Universe* 9.6, 245 (May 2023), p. 245. DOI: 10 . 3390 / universe9060245. arXiv: 2305 . 18531 [astro-ph.HE].
- [2] Jillian C. Rastinejad et al. “A kilonova following a long-duration gamma-ray burst at 350 Mpc”. In: 612.7939 (Dec. 2022), pp. 223–227. DOI: 10 . 1038/s41586-022-05390-w. arXiv: 2204.10864 [astro-ph.HE].
- [3] Andrew J. Levan et al. “Heavy-element production in a compact object merger observed by JWST”. In: 626.8000 (Feb. 2024), pp. 737–741. DOI: 10 . 1038/s41586-023-06759-1. arXiv: 2307.02098 [astro-ph.HE].
- [4] Shu-Qing Zhong, Long Li, and Zi-Gao Dai. “GRB 211211A: A Neutron Star-White Dwarf Merger?” In: 947.2, L21 (Apr. 2023), p. L21. DOI: 10 . 3847/2041-8213/acca83. arXiv: 2304.04009 [astro-ph.HE].
- [5] Yu-Han Yang et al. “A lanthanide-rich kilonova in the aftermath of a long gamma-ray burst”. In: 626.8000 (Feb. 2024), pp. 742–745. DOI: 10 . 1038/s41586-023-06979-5. arXiv: 2308.00638 [astro-ph.HE].
- [6] Noemi Anau Montel, James Alvey, and Christoph Weniger. *Scalable inference with Autoregressive Neural Ratio Estimation*. 2023. arXiv: 2308.08597 [astro-ph.IM].
- [7] Tim Dietrich et al. “Multimessenger constraints on the neutron-star equation of state and the Hubble constant”. In: *Science* 370.6523 (Dec. 2020), pp. 1450–1453. DOI: 10 . 1126/science.abb4317. URL: <https://doi.org/10.1126%2Fscience.abb4317>.
- [8] Shreya Anand et al. “Chemical Distribution of the Dynamical Ejecta in the Neutron Star Merger GW170817”. In: *arXiv e-prints*, arXiv:2307.11080 (July 2023), arXiv:2307.11080. DOI: 10 . 48550/arXiv.2307.11080. arXiv: 2307 . 11080 [astro-ph.HE].
- [9] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. “The frontier of simulation-based inference”. In: *Proceedings of the National Academy of Sciences* 117.48 (May 2020), pp. 30055–30062. DOI: 10 . 1073/pnas . 1912789117. URL: <https://doi.org/10.1073%2Fpnas.1912789117>.
- [10] ChangHoon Hahn and Peter Melchior. “Accelerated Bayesian SED Modeling Using Amortized Neural Posterior Estimation”. In: 938.1, 11 (Oct. 2022), p. 11. DOI: 10 . 3847/1538-4357/ac7b84. arXiv: 2203.07391 [astro-ph.GA].
- [11] Maximilian Dax et al. “Real-Time Gravitational Wave Science with Neural Posterior Estimation”. In: 127.24, 241103 (Dec. 2021), p. 241103. DOI: 10 . 1103/PhysRevLett . 127 . 241103. arXiv: 2106.12594 [gr-qc].
- [12] Pablo Lemos et al. “Robust simulation-based inference in cosmology with Bayesian neural networks”. In: *Machine Learning: Science and Technology* 4.1, 01LT01 (Mar. 2023), 01LT01. DOI: 10 . 1088/2632-2153/acbb53. arXiv: 2207.08435 [astro-ph.CO].
- [13] P. Darc et al. “Kilonova Spectral Inverse Modelling with Simulation-based Inference: An Amortized Neural Posterior Estimation Analysis”. In: 971.1, 82 (Aug. 2024), p. 82. DOI: 10 . 3847/1538-4357/ad53c7. arXiv: 2311.09471 [astro-ph.HE].
- [14] Malina Desai et al. “Likelihood-free Inference for Kilonova Light Curve Parameter Estimation”. In: *American Astronomical Society Meeting Abstracts*. Vol. 243. American Astronomical Society Meeting Abstracts. Feb. 2024, 426.06, p. 426.06.
- [15] A Spurio Mancini et al. “Bayesian model comparison for simulation-based inference”. In: *RAS Techniques and Instruments* 2.1 (Jan. 2023), pp. 710–722. DOI: 10 . 1093/rasti/rzad051. arXiv: 2207.04037 [astro-ph.CO].
- [16] Jason D. McEwen et al. “Machine learning assisted Bayesian model comparison: learnt harmonic mean estimator”. In: *arXiv e-prints*, arXiv:2111.12720 (Nov. 2021), arXiv:2111.12720. DOI: 10 . 48550/arXiv.2111.12720. arXiv: 2111 . 12720 [stat.ME].
- [17] He Gao et al. “A complete reference of the analytical synchrotron external shock models of gamma-ray bursts”. In: 57.6 (Dec. 2013), pp. 141–190. DOI: 10 . 1016/j . newar . 2013 . 10 . 001. arXiv: 1310.2181 [astro-ph.HE].
- [18] Geoffrey Ryan et al. “Gamma-Ray Burst Afterglows in the Multimessenger Era: Numerical Models and Closure Relations”. In: 896.2, 166 (June 2020), p. 166. DOI: 10 . 3847/1538-4357/ab93cf. arXiv: 1909.11691 [astro-ph.HE].

- [19] Brian D. Metzger. “Kilonovae”. In: *Living Reviews in Relativity* 23.1 (Dec. 2019). DOI: 10.1007/s41114-019-0024-0. URL: <https://doi.org/10.1007%2Fs41114-019-0024-0>.
- [20] Michael Deistler, Pedro J Goncalves, and Jakob H Macke. “Truncated proposals for scalable and hassle-free simulation-based inference”. In: *arXiv e-prints*, arXiv:2210.04815 (Oct. 2022), arXiv:2210.04815. DOI: 10.48550/arXiv.2210.04815. arXiv: 2210.04815 [stat.ML].
- [21] George Papamakarios and Iain Murray. “Fast  $\epsilon$ -free Inference of Simulation Models with Bayesian Conditional Density Estimation”. In: *Advances in Neural Information Processing Systems*. Ed. by D. Lee et al. Vol. 29. Curran Associates, Inc., 2016. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2016/file/6aca97005c68f1206823815f66102863-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/6aca97005c68f1206823815f66102863-Paper.pdf).
- [22] David Greenberg, Marcel Nonnenmacher, and Jakob Macke. “Automatic Posterior Transformation for Likelihood-Free Inference”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri and Ruslan Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. Long Beach, California, USA: PMLR, Sept. 2019, pp. 2404–2414. URL: <http://proceedings.mlr.press/v97/greenberg19a.html>.
- [23] George Papamakarios et al. “Normalizing flows for probabilistic modeling and inference”. In: *The Journal of Machine Learning Research* 22.1 (2021), pp. 2617–2680.
- [24] George Papamakarios, David C. Sterratt, and Iain Murray. “Sequential Neural Likelihood: Fast Likelihood-free Inference with Autoregressive Flows”. In: *arXiv e-prints*, arXiv:1805.07226 (May 2018), arXiv:1805.07226. DOI: 10.48550/arXiv.1805.07226. arXiv: 1805.07226 [stat.ML].
- [25] J. C. Rastinejad et al. *Uniform Modeling of Observed Kilonovae: Implications for Diversity and the Progenitors of Merger-Driven Long Gamma-Ray Bursts*. 2024. arXiv: 2409.02158 [astro-ph.HE]. URL: <https://arxiv.org/abs/2409.02158>.

## A Appendix

Table 1: Parameter Priors and Posteriors for the FS+2KNE model

Parameter	Prior	Posterior TSNPE	Posterior MCMC	Posterior Yang et al. 2024
<b>FS</b>				
$\log E_0$ (erg)	(50,60)	$54.10^{+2.47}_{-2.56}$	$52.06^{+0.94}_{-1.57}$	$53.87^{+1.34}_{-1.36}$
$\log n_0$ (cm <sup>-3</sup> )	(-6,2)	$-2.16^{+2.12}_{-2.02}$	$-3.67^{+1.57}_{-1.67}$	$-4.11^{+1.55}_{-1.20}$
$\log \theta_c$ (rad)	(-2,-0.5)	$-1.27^{+0.48}_{-0.39}$	$-1.03^{+0.42}_{-0.31}$	$-1.38^{+0.29}_{-0.27}$
$\log \epsilon_e$	(-6,-0.3)	$-3.34^{+1.50}_{-1.43}$	$-2.65^{+0.84}_{-0.65}$	$-2.90^{+1.51}_{-1.36}$
$\log \epsilon_B$	(-6,-0.3)	$-3.27^{+1.50}_{-1.43}$	$-3.71^{+1.81}_{-1.16}$	$-3.96^{+1.57}_{-1.34}$
$p$	(2.01,2.9)	$2.46^{+0.22}_{-0.28}$	$2.41^{+0.42}_{-0.12}$	$2.64^{+0.05}_{-0.05}$
$\log \xi_N$	(-5,0)	$-2.59^{+1.35}_{-1.58}$	$-2.32^{+0.78}_{-0.43}$	$-3.19^{+1.23}_{-1.13}$
<b>2KN</b>				
$\log M_{ej,1}$ ( $M_\odot$ )	(-3,-1)	$-2.01^{+0.50}_{-0.50}$	$-1.06^{+0.05}_{-1.06}$	$-1.56^{+0.14}_{-0.12}$
$\log v_{ej,1}$ (c)	(-1,-0.5)	$-0.75^{+0.15}_{-0.13}$	$-0.78^{+0.08}_{-0.07}$	$-0.72^{+0.13}_{-0.17}$
$\log \kappa_1$ (cm <sup>2</sup> g <sup>-1</sup> )	(-2,0.5)	$-0.74^{+0.63}_{-0.77}$	$-1.34^{+0.42}_{-0.50}$	$-0.24^{+0.26}_{-0.27}$
$\beta_{v,1}$	(1,5)	$2.98^{+1.05}_{-0.99}$	$3.05^{+0.93}_{-0.83}$	$3.09^{+1.14}_{-1.15}$
$Y_{e,1}$	(0.2,0.4)	$0.30^{+0.05}_{-0.05}$	$0.30^{+0.06}_{-0.07}$	$0.30^{+0.06}_{-0.06}$
$\log M_{ej,2}$ ( $M_\odot$ )	(-3,-1)	$-1.99^{+0.50}_{-0.50}$	$-1.01^{+0.01}_{-0.01}$	$-1.28^{+0.07}_{-0.06}$
$\log v_{ej,2}$ (c)	(-2,-1)	$-1.50^{+0.25}_{-0.26}$	$-1.19^{+0.17}_{-0.05}$	$-1.51^{+0.08}_{-0.09}$
$\log \kappa_2$ (cm <sup>2</sup> g <sup>-1</sup> )	(-0.5,2)	$0.76^{+0.66}_{-0.78}$	$1.65^{+0.16}_{-1.74}$	$1.66^{+0.22}_{-0.29}$
$\beta_{v,2}$	(1,5)	$3.03^{+1.09}_{-1.05}$	$3.18^{+0.82}_{-0.64}$	$2.25^{+1.01}_{-0.85}$
$Y_{e,2}$	(0.1,0.2)	$0.15^{+0.03}_{-0.03}$	$0.15^{+0.04}_{-0.03}$	$0.15^{+0.03}_{-0.03}$

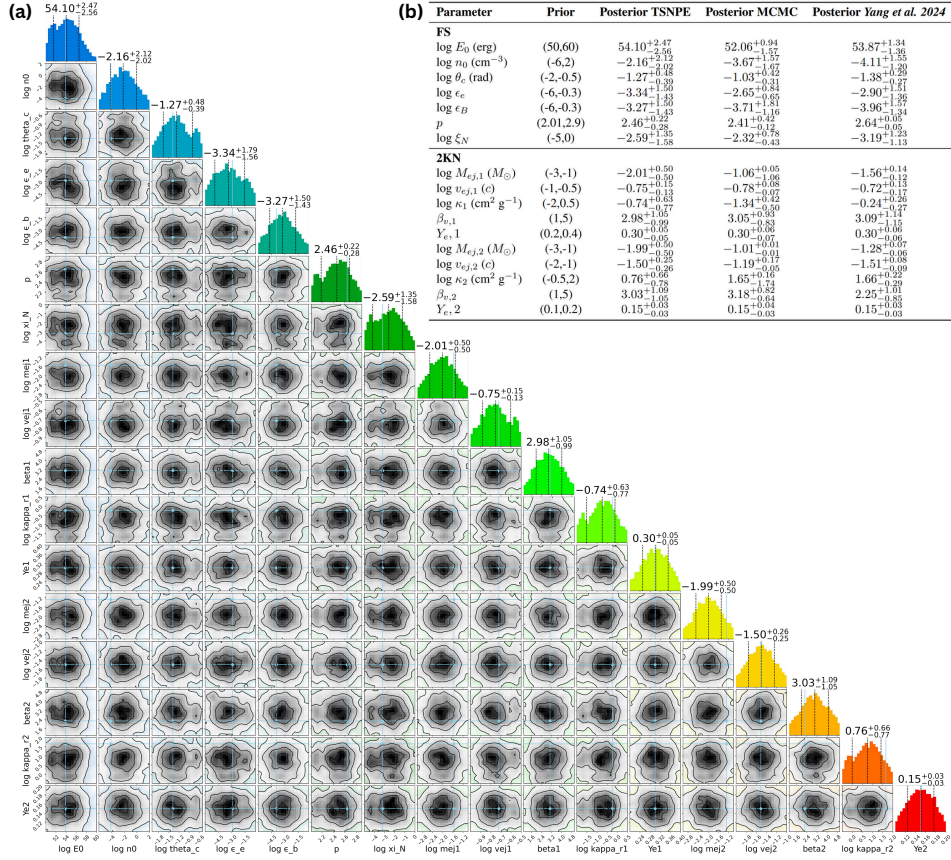


Figure 4: Enhanced version of Figure 2.




Figure 5: Results for a forward shock plus two-component kilonova model in a low-budget simulation regime.