Draft version January 3, 2022
Typeset using LaTeX **twocolumn** style in AASTeX62

arXiv:2112.15470v1 [astro-ph.HE] 31 Dec 2021

Using Neural Networks to Perform Rapid High-Dimensional Kilonova Parameter Inference

Mouza Almualla,[1] Yuhong Ning,[2] Mattia Bulla,[3] Tim Dietrich,[4,5] Michael W. Coughlin,[6] and Nidhal Guessoum[1]

[1] *American University of Sharjah, Physics Department, PO Box 26666, Sharjah, UAE*
[2] *Zhiyuan College (School of Mathematical Sciences), Shanghai Jiao Tong University, Shanghai, 200240, China*
[3] *The Oskar Klein Centre, Department of Astronomy, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden*
[4] *Institute of Physics and Astronomy, University of Potsdam, Karl-Liebknecht-Str. 24/25, 14476, Potsdam, Germany*
[5] *Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, D-14476 Potsdam, Germany*
[6] *School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

ABSTRACT

On the 17th of August, 2017 came the simultaneous detections of GW170817, a gravitational wave that originated from the coalescence of two neutron stars, along with the gamma-ray burst GRB170817A, and the kilonova counterpart AT2017gfo. Since then, there has been much excitement surrounding the study of neutron star mergers, both observationally, using a variety of tools, and theoretically, with the development of complex models describing the gravitational-wave and electromagnetic signals. In this work, we improve upon our pipeline to infer kilonova properties from observed light-curves by employing a Neural-Network framework that reduces execution time and handles much larger simulation sets than previously possible. In particular, we use the radiative transfer code `POSSIS` to construct 5-dimensional kilonova grids where we employ different functional forms for the angular dependency for the dynamical ejecta component. We find that incorporating an angular dependence improves the fit to the AT2017gfo light-curves by up to ∼50% when quantified in terms of the weighted Mean Square Error.

1. INTRODUCTION

The driving idea of multi-messenger astronomy is that a single event observed through different messengers improves our understanding of the physical processes. This concept has gotten additional momentum through the Binary Neutron Star (BNS) detection GW170817 (Abbott et al. 2017a) connected with the observation of electromagnetic radiation from gamma-rays to radio waves (Abbott et al. 2017b). This event followed the birth of GW astronomy in 2015 (Abbott, B. P. et al. 2016), brought about with the success of the Advanced Virgo (Acernese et al 2015) and the Advanced LIGO (Aasi et al 2015) interferometers in detecting GWs emitted by compact object mergers. In general, BNS and Neutron Star – Black Hole (NSBH) mergers may be accompanied by an optical/infrared counterpart, referred to as a *kilonova*, which will give us important insights into what is possibly one of the most dominant sites of rapid neutron capture (r-process) elements in the Universe; see Metzger (2017) for a review. Our first and as yet only confirmed discovery of a kilonova from a BNS merger, AT2017gfo in NGC 4993 (D ∼ 40 Mpc) (Coulter et al. 2017), was achieved through follow-up of GW170817 by multiple telescope facilities around the globe, and is an exceptional demonstration of the potential of multi-messenger astronomy in contributing to scientific knowledge (Abbott et al. 2017b). These detections were also accompanied by the discovery of short gamma-ray burst GRB 170817A (Abbott et al. 2017c). Through this event, we learned more about the aforementioned r-process nucleosynthesis (e.g., Chornock et al. 2017; Coulter et al. 2017; Cowperthwaite et al. 2017; Pian et al. 2017; Smartt et al. 2017; Watson et al. 2019; Kasliwal et al. 2019) and also placed constraints on the expansion rate of the Universe (Abbott et al. 2017; Guidorzi et al. 2017; Hotokezaka et al. 2019; Coughlin et al. 2020; Dhawan et al. 2019; Dietrich et al. 2020) and the neutron star equation of state (e.g., Bauswein et al. 2017; Margalit & Metzger 2017; Coughlin et al. 2019a,b; Annala et al. 2018; Most et al. 2018; Radice et al. 2018; Abbott et al. 2018; Lai et al. 2019; Dietrich et al. 2020; Huth et al. 2021).

Inferring kilonova model parameters from observational data requires a thorough understanding of the ejecta composition, behaviour, and morphology (e.g., Dietrich et al. 2020; Nedora et al. 2022; Nicholl et al. 2021; Breschi et al. 2021); in addition, there are various parts of the parameter inference pipeline that must be improved upon and optimized such that we are able to efficiently utilize the available data and learn more

about the progenitors. Radiative transfer simulations have been most commonly used to generate expected supernova and kilonova light-curves, (e.g., Tanaka 2016; Kasen et al. 2017; Wollaeger et al. 2017; Bulla 2019; Kawaguchi et al. 2020).

There are two major sources of ejecta from the progenitor system. The first is known as the *dynamical ejecta*, ejected 'dynamically' around the moment of merger through torque or shocks (at merger or in the early-postmerger phase) (e.g., Metzger et al. 2008; Bauswein et al. 2013). The second source of ejecta comes from the *post-merger wind*, i.e., outflows from an accretion disk formed around the central remnant, consisting of the unbound debris resulting from the merger (e.g., Siegel & Metzger 2018; Metzger 2019). The dynamical source is ejected on a timescale of milliseconds, while the latter is formed on a timescale of up to seconds.

There are two major dynamical ejection processes in BNS mergers. The first is through shocks that forms at the contact interface of the merging stars or at core bounces in the early postmerger phase (Hotokezaka & Piran 2015). The second component is the tidal ejecta, and is due to tidal interactions that arise from the non-axisymmetric gravitational forces at play in the binary system (Hotokezaka & Piran 2015; Metzger 2019). Recent simulations (Radice et al. 2018; Nedora et al. 2022) show that the more neutron-rich, lower Y_e ejecta ($Y_e \lesssim 0.1$) arises from the aforementioned tidal component, and is located closer to the equatorial plane, while the ejecta arising from the shock component tends to have a relatively higher electron fraction (going up to $Y_e \sim 0.4$) and is approximately isotropic.

The disk mass is $\sim 0.01 - 0.3 M_\odot$ (Oechslin & Janka 2006), and is usually much lower when the remnant immediately collapses into a black hole; this is due to the lack of time for the remnant to redistribute its angular momentum and mass as it transforms from a differentially rotating to a solid rotating body, thus preventing the formation of a more massive disk (Metzger 2019). A fraction of the disk is ejected in the form of a post-merger disk wind, with the exact value being uncertain and ranging from 20 to 40 per cent (e.g. Just et al. 2015; Siegel & Metzger 2018; Miller et al. 2019; Fernández et al. 2019; Fujibayashi et al. 2020). The post-merger wind usually dominates the dynamical ejecta (Wu et al. 2016). Works such as Fernández & Metzger (2013) show the disk wind ejecta Y_e ranging within $\sim 0.2 - 0.4$ when the remnant collapses into a black hole. The electron fractions present in the case of both dynamical and post merger wind cases are suitable for the production of heavy r-process elements (Rosswog et al. 2014).

With the gravitational-wave detector networks' fourth observing run nearing, one needs to continue to improve the instruments and software tools that we use to analyze the forthcoming multi-messenger detections. These mergers bring rich knowledge to various fields of physics and astronomy, from the understanding of such events and the big, powerful jets of relativistic particles that they produce, to the physical structure of both the neutron stars (internally) and the matter that is ejected in the merger, not to mention the detailed characteristics of the gravitational waves that are emitted. Much work is being done on the theoretical front to prepare us for the 'treasure trove' of data from upcoming detections. In particular, models of kilonovae are being constructed with increased complexity and fine-tuning (e.g., Korobkin et al. 2021; Zhu et al. 2021; Wollaeger et al. 2021).

In this work, we use `POSSIS` (Bulla 2019), a Monte-Carlo radiative transfer code, in order to produce "model grids" that sample the parameter space for any given ejecta morphology. In particular, we use an updated version of `POSSIS` including improved heating rates and wavelength- and time-dependent opacities from Tanaka et al. (2020) that depend on local properties of the ejecta (density, temperature and Y_e). Details about these new implementations will be discussed elsewhere (Bulla et al., in prep). Previous works such as Coughlin et al. (2018); Dietrich et al. (2020); Heinzel et al. (2021) used Gaussian Process Regression (GPR) to interpolate within kilonova grids to perform parameter estimation. Here, we introduce a Neural Network (NN) framework that significantly enhances the speed of the interpolation step and scales well for much larger grids. In addition, due to the importance of understanding the effects of different ejecta morphologies on the computed light-curves, we explore different functional forms for the angular dependence in the density profile of the dynamical ejecta.

Our paper is structured as follows, in Sec. 2, we describe the NN architecture. In Sec. 3, we discuss and evaluate the different density profiles for the dynamical ejecta in the context of GW170817/AT2017gfo. Finally, we summarize our conclusions and outline ways of moving forward in Sec. 4.

2. A NEW SURROGATE GENERATION FRAMEWORK

2.1. *Preprocessing*

Instead of directly inputting our bolometric and photometric light-curves into our Neural Network for interpolation, we follow the procedure of Doctor et al. (2017) and Coughlin et al. (2018), using Singular Value Decom-

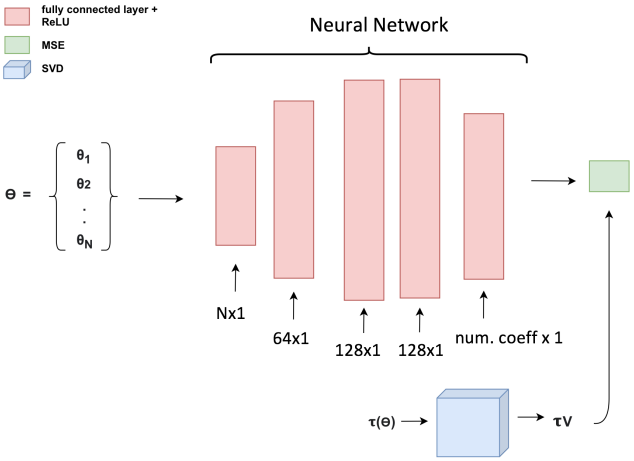

Figure 1. A diagram of the interpolation framework used in this work. The input to our network is the model parameters, and the network is trained on a PCA-like projection of our photometric and bolometric data onto the right-singular vector. Five fully-connected layers are shown, and an MSE loss function is employed.

position (SVD) to perform a dimensionality reduction on our data and produce the principle components of our light-curve vectors, i.e., a new basis for our data in which different dimensions of the data are most uncorrelated. We then train our neural network on this output in order to produce a so-called "surrogate model" that interpolates our original model grid. Denoting our photometric and bolometric light-curves as $\mathcal{T}(\mathbf{\Theta})$, where $\mathbf{\Theta}$ represents the set of parameters for a given model, SVD decomposes $\mathcal{T}(\mathbf{\Theta})$ as follows:

$$\mathcal{T} = U\Sigma V^T \tag{1}$$

where Σ is a diagonal matrix containing the singular values of $\mathcal{T}(\mathbf{\Theta})$, and the columns of U and V are orthonormal bases called the left-singular and right-singular vectors, respectively. We can then obtain the principle component analysis (PCA) output by projecting $\mathcal{T}(\mathbf{\Theta})$ onto the right singular matrix:

$$\mathcal{T}V = PCA \tag{2}$$

Note that we first normalized our light curves before performing PCA, and the right singular vector V was truncated to 100 basis vectors. A train/validation split of 90/10 was used.

2.2. *Neural Network Architecture*

Our NN consists of five ReLU-activated (Rectified Linear Units, i.e., $R(z) = \max(0, z)$ for each neuron) fully-connected layers, as shown in Figure 1, with the first layer containing N weights corresponding to the number of parameters for the model, three hidden layers, and the output layer having a number of weights equal to the number of columns in our PCA output. We used a batch size of 32 and trained for 200 epochs, using a mean-square error (MSE) loss function.

The use of a NN is much more efficient for higher-dimensional kilonova models than GPR, which usually performs poorly when a model has more than four parameters. To demonstrate the performance of our interpolation framework, we use a model (to be described in-depth in the following sections) that has 5 parameters and 1,260 total parameter combinations comprising our simulation set. Training 10 different networks (one for each photometric band u, g, r, i, z, y, J, H, K, and in addition one for the bolometric luminosity) took on the order of $\sim$ 40 mins. The MSE loss is shown as a function of epoch in Figure 2. Our model does not over- or under-fit the train data, since we can see very similar performance on our validation set in the right subplot as compared to the train set. In addition, our network converges very quickly, so even 50 epochs would be sufficient in constructing our surrogate.

We also compare our NN framework to state-of-the-art interpolation methods, such as the GPR framework used in Coughlin et al. (2018); Heinzel et al. (2021), as well as that in Ristic et al. (2021), in terms of its computational efficiency. It took upwards of 13 hours for the former GPR implementation from Coughlin et al. (2018) (which uses the `scikit-learn` library) to construct our model in the u-band, and so we only use the framework from Ristic et al. (2021) in this comparison. The performance for each model is shown in Figure 3. In addition to the benefits in terms of convergence time, the output of the NN interpolation scheme takes up a lot less memory – on the order of tens of MB, in comparison to a few GBs for the `scikit-learn` GPR scheme. Even more importantly, the total time it takes to perform parameter inference is significantly reduced due to the much simpler loading procedure for NNs; for the 5-dimensional model used here, loading for all of the bands and the bolometric luminosity takes a total of $\sim$700 ms for the NN, while for the GPR approach from Ristic et al. (2021), it takes $\sim$ 80 min (since some components need to be recalculated[1]). This is a significant improvement when considering that the parameter inference takes $\approx$20 min excluding the loading step.

Future work in improving the performance of our model includes incorporating batch normalization, drop

[1] There may be some room for improvement in loading for the GPR method, so consider this an upper limit on the execution time.

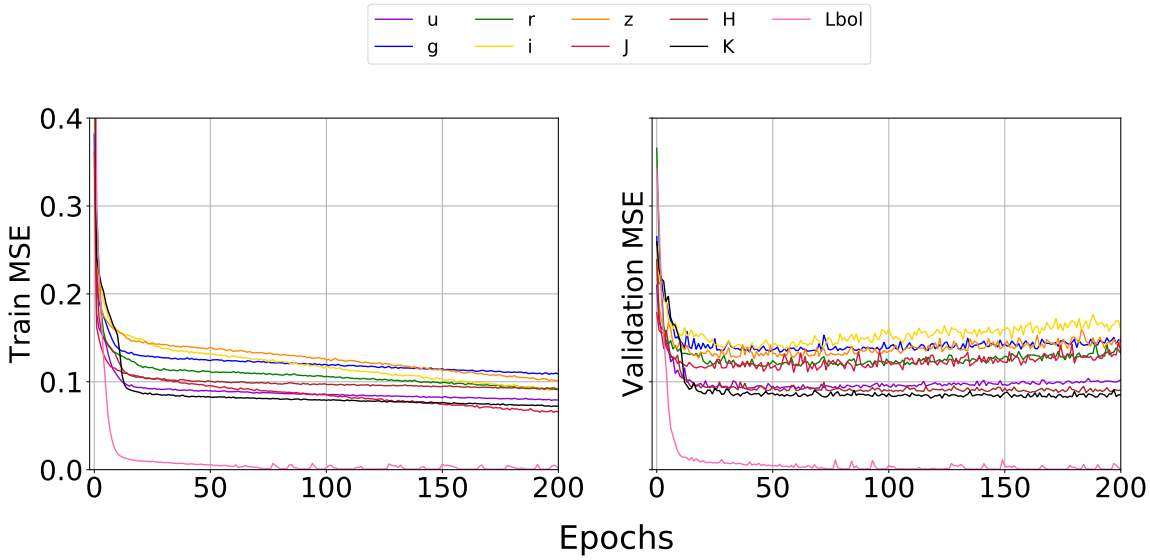

Figure 2. The MSE loss is shown for the u, g, r, i, z, y, J, H, and K bands, as well as for the bolometric luminosity. The loss for the train set is shown in the left subplot, and for the validation set in the right subplot. After $\sim$ 200 epochs, the network tends to over-fit the train data, so we stop training at this point.

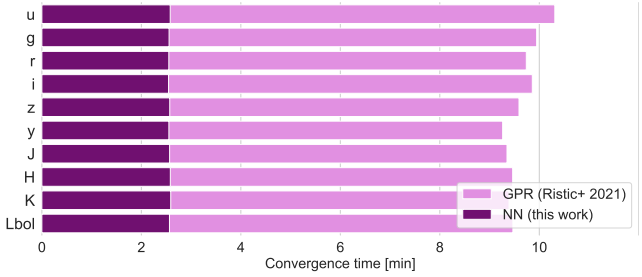

Figure 3. Convergence time for all of the photometric bands used in this work, as well as for the bolometric luminosity. The NN framework is shown in dark purple, while the interpolation scheme from Ristic et al. (2021) is shown in light purple.

out layers, and various other easily implementable techniques. The wealth of developments in deep learning allows for a significant improvement in computational efficiency whilst still maintaining high accuracies and developing surrogate models with high fidelity.

3. ANGULAR DEPENDENCIES IN THE DYNAMICAL EJECTA

Numerical relativity simulations provide valuable information about what we can expect from the ejected matter in the merger of BNSs. As the field develops, these simulations have become more and more complex, incorporating neutrino transport and more realistic neutron star equations of state. We know that there are two primary ejection processes, described in Section 1, but the exact geometry of each of these components are now known; note that also the ejecta opacities are another point of uncertainty in kilonova simulations and we refer to Appendix A for a preliminary exploration of this. Generally, the density profile of the post-merger wind is thought to be relatively spherically symmetric, while the density profile of the dynamical ejecta will have an angular dependence such that the mass is more concentrated in equatorial regions as compared to polar regions (e.g., Dietrich & Ujevic 2017; Kawaguchi et al. 2020; Nedora et al. 2022).

There are various functional forms that have been used in kilonova modelling to incorporate this angular dependency into the dynamical ejecta. We will explore a few of them, and then use data from AT2017gfo to infer which provides the best fit to the observed light-curves. We remark that in our simulations, we will only consider the case in which the post-merger wind is slower than the dynamical ejecta. We use the so-called *Spherical segment-Spherical Cap* geometry (shortened as SSCr, with r signifying that re-processing between the different ejecta components is taken into account); see the left-most geometry in Figure 1 of Heinzel et al. (2021). The density profiles for the different ejecta components are represented as follows:

$$\rho(r,t) \propto \begin{cases} r^{-\alpha}t^{-3} & v_{\rm pm,min} \leq r/t \leq v_{\rm pm,max} \\ \eta(\theta) r^{-\alpha}t^{-3} & v_{\rm pm,max} \leq r/t \leq v_{\rm dyn,max} \end{cases}$$

where we set α=3, $v_{\rm pm,min} = 0.025c$, $v_{\rm pm,max} = 0.08c$, $v_{\rm dyn,max} = 0.3c$. $\eta(\theta)$ is the term introducing the angular dependence with respect to θ, the angle measured from

the polar axis, into the dynamical ejecta. The forms that we will first explore are:

1) a sinusoidal relationship, i.e., $\eta(\theta) = \sin^2(\theta)$, also used by Perego et al. (2017) based on numerical-relativity simulations from Radice et al. (2018); and

2) a function defined in Kawaguchi et al. (2020):

$$\eta(\theta) = (1 - \Theta(\theta))f_{\rm d} + \Theta(\theta) \tag{3}$$

$$\Theta(\theta) = \frac{1}{1 + \exp\left[-10(\theta - \pi/4])\right]} \tag{4}$$

where f_d determines the strength of the angular dependence from equatorial to polar angles. We will call the latter the "Kawaguchi Model." As can be seen in Figure 4, with the sinusoidal functions, we have a very gradual increase, whereas with the dependency from Kawaguchi et al. (2020) the density profile will be much more concentrated for $\pi/4 \leq \theta \leq 3\pi/4$. For our

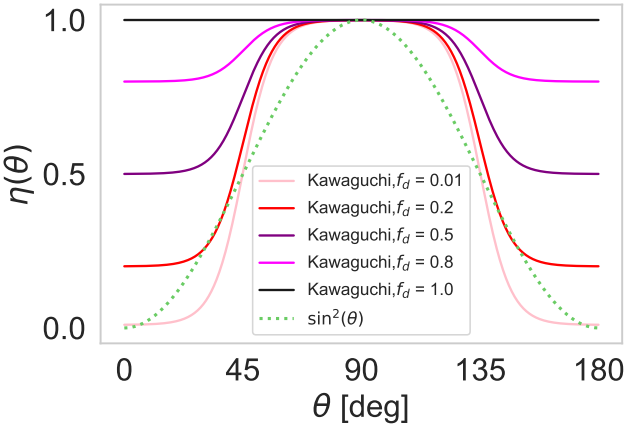

Figure 4. Different functional forms for the angular dependency in the dynamical ejecta. This is shown for $\sin^2(\theta)$ and Equations (3) and (4). θ represents the angle from the polar axis.

first analysis here, we will be pinning f_d at 0.01 as done in Kawaguchi et al. (2020), since this produces the strongest angular dependency and will be a useful comparison. Our model has the following four free parameters: the dynamical ejecta mass $M_{\rm ej,dyn}$, post-merger wind ejecta mass $M_{\rm ej,dyn}$, half-opening angle of the lanthanide-rich component of the dynamical ejecta ϕ, and the observing angle $\theta_{\rm obs}$. For the dynamical and post-merger wind ejecta masses, we perform logarithmic sampling across the interval [0.001, 0.1]. Our entire sample set used to generate our model grids is shown in Table 1.

We generate 3 model grids: 1) no angular dependence, 2) angular dependence $\eta = \sin^2(\theta)$, and 3) the Kawaguchi model with f_d=0.01. We thus run `POSSIS` over the aforementioned parameter space, generating $6 \times 6 \times 7 = 252$ simulations for each model, and then used the NN interpolation described in Section 2 to interpolate the grid. We then perform parameter inference and find the maximum likelihood estimates for AT2017gfo using `PyMultiNest`, which is capable of performing both model selection and parameter inference. Using `PyMultiNest`, we extract posterior distributions for each of the aforementioned parameters. We assume uniform priors that extend over the limits of the grid in the parameter space. The results of the inference (i.e., the best-fit light-curves), assuming ± 1 mag systematic error, for each of these models are shown in Figure 5.

Table 1. The parameters of our model and the sample values for which we run `POSSIS`.

Parameters	Samples
$M_{\rm ej,dyn}/M_\odot$	[0.001, 0.00251, 0.00631, 0.0158, 0.0398, 0.1]
$M_{\rm ej,win}/M_\odot$	[0.001, 0.00251, 0.00631, 0.0158, 0.0398, 0.1]
ϕ [deg]	[0, 15, 30, 45, 60, 75, 90]
$\cos(\theta_{\rm obs})$	[0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0]

Inspecting by eye, we can see in Figure 5 that across all bands the $\sin^2(\theta)$ and Kawaguchi models result in residuals closer to 0. We can further quantify this by obtaining the MSE for each model. However, since the Local Thermodynamic Equilibrium (LTE) assumption in `POSSIS` is likely to fail at late times when the ejecta become optically thin, and the brighter early-time emission is more important for observational purposes, it is reasonable to perform a Weighted MSE (WMSE) prioritizing the early observations. We thus use the following equation:

$$\text{WMSE} = \frac{1}{n}\frac{\sum_{i=0}^{n-1}\omega_i(M_{\rm obs} - M_{\rm pred})^2}{\sum_{i=0}^{n-1}\omega_i} \tag{5}$$

where $\omega_i = \ln(|M_{\rm obs}| + 1) + 1$. Having done so, we obtain the WMSE values shown in Table 2, and verify our conclusion that incorporating some form of an angular dependence improves the fit to AT2017gfo across almost all bands.

The parameter f_d allows us also to assess which dependence provides the best fit to the measured light-curves. In Figure 6, we show the 2D density profile for different f_d's (0.01, 0.5, 1.0) for comparison.

For this analysis, we include f_d as an additional parameter to those listed in Table 1, sampling $f_d = [0.01, 0.2, 0.5, 0.8, 1.0]$. This increases our simulation set to 1260 simulations. The resulting corner plot of our posteriors is shown in Figure 7. We also include the

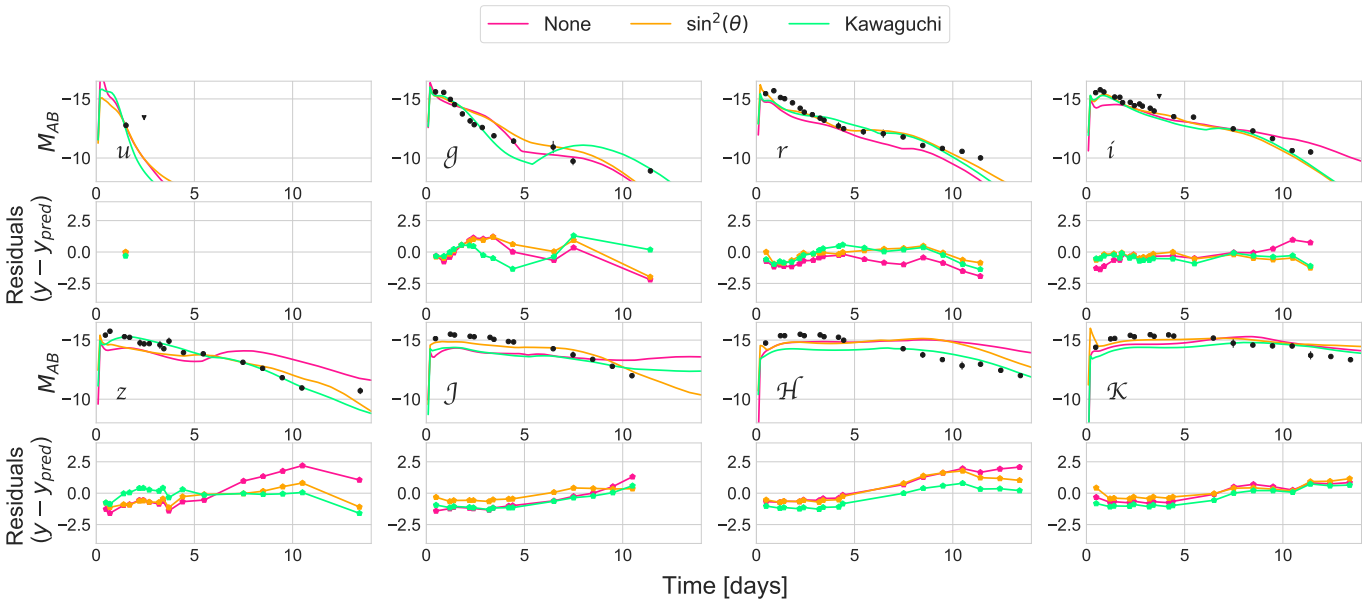

Figure 5. Best-fit light-curves in the u, g, r, i, z, J, H, and K bands are plotted for three different models in which the angular dependency in the dynamical ejecta is varied. The AT2017gfo observations are plotted as black dots, the Kawaguchi best-fit is plotted in green, the $\sin^2(\theta)$ best-fit is plotted in orange, and the model with no angular dependency is plotted in pink. We also plot the residuals in each band.

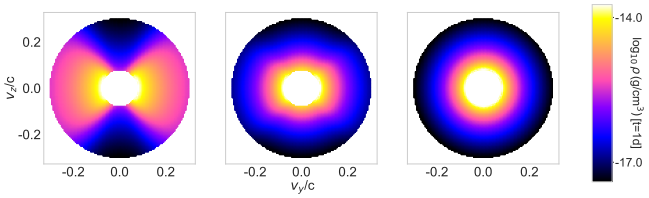

Figure 6. 2D dynamical ejecta density profiles for f_d=0.01, f_d=0.5, f_d=1.0.

WMSE values with our maximum-likelihood f_d in Table 2.

We can see that $f_d \sim 0.50^{+0.28}_{-0.32}$ provides the most optimal fit. When fitting for f_d, the previous WMSE of 0.0333 decreases to 0.0311. The fit is significantly better in redder bands with f_d=0.50, while it would be better using a stronger dependency $f_d = 0.01$ with respect to bluer g and r bands. In order to confirm this behavior, we also run the inference with an assumed systematic error of 0.5 mag (instead of 1.0), and find that the distribution is even more clearly distributed around $f_d \sim 0.6$, and the WMSE further decreases to 0.0252. A moderate angular dependence is thus preferred.

To finalize our exploration of different angular dependencies, we summarize the inferred constraints on $M_{\rm ej,pm}$, $M_{\rm ej,dyn}$, ϕ, and $\theta_{\rm obs}$ for each model in Table 3.

4. CONCLUSION

In this work, we have modified and improved upon the present parameter inference pipeline used to infer ejecta parameters from observed kilonova light-curves (Coughlin et al. 2018; Dietrich et al. 2020). We constructed an NN framework that sufficiently models the sampled lightcurves and produces a surrogate model with a much improved computational efficiency as compared to previous surrogate construction methods.

We explored the possible nature of AT2071gfo's ejecta geometry. Specifically, we looked at the angular dependence in the density profile of the dynamical ejecta, showing that some form of angular dependence is beneficial and provides better fits to GW710817/AT2017gfo,

Table 2. WMSE values for each band and model, for when no angular dependence is incoporated, versus a $\sin^2(\theta)$ dependency, versus the Kawaguchi equation when pinning f_d at 0.01, and finally when allowing f_d to vary (i.e., $f_d = 0.50$).

Band	**Models**			
	None	$\sin^2(\theta)$	$f_d = 0.01$	best-fit f_d
g	0.0642	0.0571	**0.0291**	0.0738
r	0.0461	**0.0127**	0.0204	0.0227
i	0.0242	0.0119	0.0154	**0.0034**
z	0.0771	0.0313	0.0160	**0.0109**
J	0.0781	**0.0164**	0.0649	0.0317
H	0.0846	0.0593	0.0506	**0.0412**
K	0.0249	**0.0168**	0.0370	0.0340
average	0.0570	0.0294	0.0333	0.0311

Table 3. Resultant posteriors for $M_{\rm ej,pm}$, $M_{\rm ej,dyn}$, ϕ, and $\theta_{\rm obs}$ for each of the models explored in this section. Uncertainties are shown at the 15th and 85th percentiles.

Model	**Parameters**			
$\eta(\theta)$	$M_{\rm ej,dyn}$ $[\times 10^{-3} M_\odot]$	$M_{\rm ej,pm}$ $[M_\odot]$	ϕ [deg]	$\theta_{\rm obs}$ [deg]
1	$4.57^{+21.88}_{-2.48}$	$0.110^{+0.0488}_{-0.0302}$	$44.29^{+11.85}_{-13.69}$	$32.01^{+19.38}_{-19.00}$
$\sin^2(\theta)$	$3.09^{+18.29}_{-1.51}$	$0.0912^{+0.209}_{-0.032}$	$36.03^{+14.04}_{-12.77}$	$24.86^{+24.89}_{-17.43}$
kawaguchi ($f_d = 0.01$)	$19.05^{+35.90}_{-11.98}$	$0.0832^{+0.0456}_{-0.0405}$	$26.95^{+12.28}_{-8.03}$	$37.87^{+15.77}_{-19.50}$
kawaguchi (varied f_d)	$5.89^{+15.49}_{-3.43}$	$0.0977^{+0.0645}_{-0.0236}$	$29.32^{+17.25}_{-10.16}$	$25.84^{+18.18}_{-17.72}$

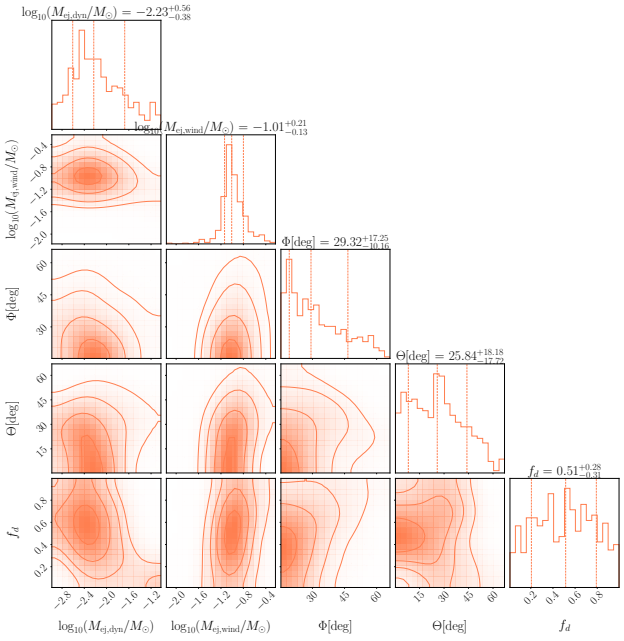

Figure 7. Corner plot showing the resulting posteriors for our model with f_d as a free parameter. We can see that all of our inferred parameters are within reasonable ranges and that the distribution for f_d is centered around 0.50.

with the $\sin^2(\theta)$ and the Kawaguchi functional forms both outperforming the model with no angular dependence. The $\sin^2(\theta)$ dependence improves the fit across all bands by $\sim$50%.

When using the Kawaguchi equation – parameterized by f_d, which is approximately the ratio of the density in the polar region ($0 \leq \theta \leq \pi/4$) to that in the equatorial region ($\pi/4 \leq \theta \leq \pi/2$) – to study the angular dependence, we found that $f_d \sim 0.50$ is most suitable. Fitting for f_d especially improves performance in the redder bands, decreasing the WMSE by 34% when considering lightcurves in the i-band and onwards.

Although we explored one specific aspect of the ejecta structure in this work, and there have been other recent explorations of different kilonova ejecta morphologies (Darbha & Kasen 2020; Heinzel et al. 2021; Korobkin et al. 2021), there is still much to explore. There are countless other geometries that could be considered for the kilonova ejecta, so an interesting path forward would be performing inference on a "super-geometry" that samples a much wider range of possible ejecta morphologies.

M.B. acknowledges support from the Swedish Research Council (Reg. no. 2020-03330). M.W.C. acknowledges support from the National Science Foundation with grant Nos. PHY-2010970 and OAC-2117997. T.D. acknowledges financial support through the Max Planck Society.

We thank Stephan Rosswog, Oleg Korobkin, and Masaomi Tanaka for sharing the heating rate libraries and opacities used in this work. We are grateful for computational resources provided by the Leonard E Parker Center for Gravitation, Cosmology and Astrophysics at the University of Wisconsin-Milwaukee, as well as by the Supercomputing Laboratory at King Abdullah University of Science and Technology (KAUST) in Thuwal, Saudi Arabia.

APPENDIX

A. INTRODUCING AN OPACITY-SCALING PARAMETER

Our simulations use state-of-the-art opacities for r-process elements (Tanaka et al. 2020). However, uncertainties in the opacities are likely to be present due to the challenges in modelling the myriad of line transitions expected for these heavy elements.

Given these inherent uncertainties, we introduce the free parameter α_κ that scales the opacities, and in turn, fine-tunes the emitted flux and color of the kilonova. We assume, for simplicity, that the relative uncertainties are constant across time and wavelength, although this is likely to be more complicated in reality. In exploring this, we use the

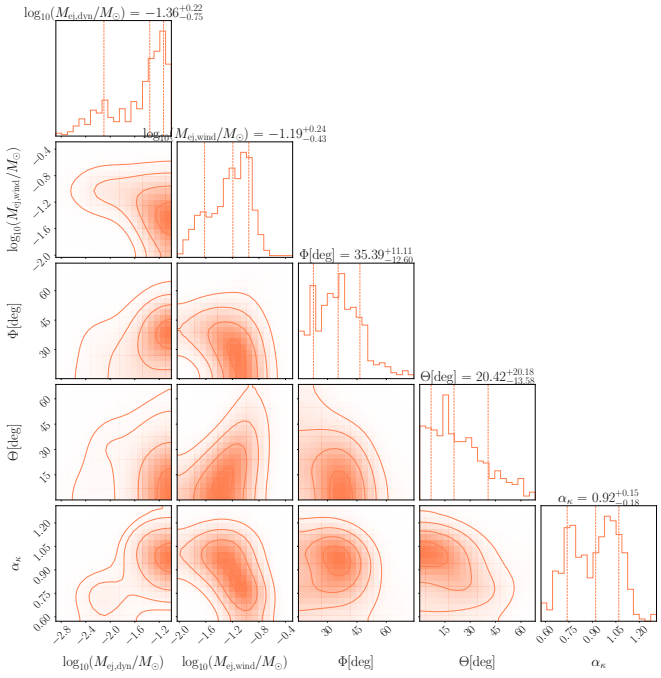

Figure 8. Posteriors showing that α_κ has a bimodal distribution when fit to GW170817/AT2017gfo lightcurves and a $\sin^2(\theta)$ dependency is assumed.

$\sin^2(\theta)$ model, given its good performance in Section 3. We use the same sample set as in Table 1, including our new parameter α_κ=[0.75,1.0,1.25]. The resulting corner plot is shown in Figure 8, and we can clearly see a bimodal distribution, with one peak centered at $\alpha_\kappa \sim 1.04$, and one centered at ~ 0.75.

Hence, while at first glance it seems that introducing an opacity rescaling might improve the performance of the model, more work is required for a reliable estimation of the opacities. In particular, the strong bimodal nature of the distribution indicates that a simple time-independent rescaling of the opacities, as employed in this work, might not be applicable. And so, further work exploring either the temporal or the wavelength dependence of the ejecta opacities would be interesting, and may serve to limit the associated systematic errors.

REFERENCES

Aasi et al. 2015, Classical and Quantum Gravity, 32, 074001

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, Phys. Rev. Lett., 121, 161101. https://link.aps.org/doi/10.1103/PhysRevLett.121.161101

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, Nature, 551, 85

Abbott, B. P. et al. 2016, Phys. Rev. Lett., 116, 061102. https://link.aps.org/doi/10.1103/PhysRevLett.116.061102

Abbott et al. 2017a, Phys. Rev. Lett., 119, 161101. https://link.aps.org/doi/10.1103/PhysRevLett.119.161101

—. 2017b, Astrophysical Journal Letters, 848, L12

—. 2017c, The Astrophysical Journal Letters, 848, L13. http://stacks.iop.org/2041-8205/848/i=2/a=L13

Acernese et al. 2015, Classical and Quantum Gravity, 32, 024001

Annala, E., Gorda, T., Kurkela, A., & Vuorinen, A. 2018, Phys. Rev. Lett., 120, 172703. https://link.aps.org/doi/10.1103/PhysRevLett.120.172703

Bauswein, A., Goriely, S., & Janka, H.-T. 2013, The Astrophysical Journal, 773, 78. http://stacks.iop.org/0004-637X/773/i=1/a=78

Bauswein et al. 2017, The Astrophysical Journal Letters, 850, L34. http://stacks.iop.org/2041-8205/850/i=2/a=L34

Breschi, M., Perego, A., Bernuzzi, S., et al. 2021, Mon. Not. Roy. Astron. Soc., 505, 1661

Bulla, M. 2019, Monthly Notices of the Royal Astronomical Society, 489, 5037–5045. http://dx.doi.org/10.1093/mnras/stz2495

Chornock et al. 2017, The Astrophysical Journal Letters, 848, L19. http://stacks.iop.org/2041-8205/848/i=2/a=L19

Coughlin, M. W., Dietrich, T., Margalit, B., & Metzger, B. D. 2019a, Monthly Notices of the Royal Astronomical Society: Letters, 489, L91. https://doi.org/10.1093/mnrasl/slz133

Coughlin, M. W., Dietrich, T., Doctor, Z., et al. 2018, Monthly Notices of the Royal Astronomical Society, 480, 3871. http://dx.doi.org/10.1093/mnras/sty2174

Coughlin, M. W., Dietrich, T., Antier, S., et al. 2019b, Monthly Notices of the Royal Astronomical Society, 492, 863. https://doi.org/10.1093/mnras/stz3457

Coughlin, M. W., Dietrich, T., Heinzel, J., et al. 2020, Physical Review Research, 2, doi:10.1103/physrevresearch.2.022006. http://dx.doi.org/10.1103/PhysRevResearch.2.022006

Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, Science, 358, 1556

Cowperthwaite, P. S., Berger, E., Villar, V. A., et al. 2017, ApJL, 848, L17

Darbha, S., & Kasen, D. 2020, ApJ, 897, 150

Dhawan, S., Bulla, M., Goobar, A., Carracedo, A. S., & Setzer, C. N. 2019, arXiv:1909.13810

Dietrich, T., Coughlin, M. W., Pang, P. T. H., et al. 2020, Science, 370, 1450. https://science.sciencemag.org/content/370/6523/1450

Dietrich, T., & Ujevic, M. 2017, Classical and Quantum Gravity, 34, 105014. http://dx.doi.org/10.1088/1361-6382/aa6bb0

Doctor, Z., Farr, B., Holz, D. E., & Pürrer, M. 2017, ArXiv e-prints, arXiv:1706.05408

Fernández, R., & Metzger, B. D. 2013, Monthly Notices of the Royal Astronomical Society, 435, 502. +http://dx.doi.org/10.1093/mnras/stt1312

Fernández, R., Tchekhovskoy, A., Quataert, E., Foucart, F., & Kasen, D. 2019, MNRAS, 482, 3373

Fujibayashi, S., Shibata, M., Wanajo, S., et al. 2020, PhRvD, 101, 083029

Guidorzi, C., et al. 2017, Astrophys. J. Lett., 851, L36

Heinzel, J., Coughlin, M. W., Dietrich, T., et al. 2021, Monthly Notices of the Royal Astronomical Society, 502, 3057–3065. http://dx.doi.org/10.1093/mnras/stab221

Hotokezaka, K., Nakar, E., Gottlieb, O., et al. 2019, Nature Astronomy, doi:10.1038/s41550-019-0820-1. https://doi.org/10.1038/s41550-019-0820-1

Hotokezaka, K., & Piran, T. 2015, Monthly Notices of the Royal Astronomical Society, 450, 1430–1440. http://dx.doi.org/10.1093/mnras/stv620

Huth, S., Pang, P. T. H., Tews, I., et al. 2021, Constraining Neutron-Star Matter with Microscopic and Macroscopic Collisions, , , arXiv:2107.06229

Just, O., Bauswein, A., Pulpillo, R. A., Goriely, S., & Janka, H.-T. 2015, Monthly Notices of the Royal Astronomical Society, 448, 541. +http://dx.doi.org/10.1093/mnras/stv009

Kasen, D., Metzger, B., Barnes, J., Quataert, E., & Ramirez-Ruiz, E. 2017, Nature, 551, 80 EP . http://dx.doi.org/10.1038/nature24453

Kasliwal, M. M., Kasen, D., Lau, R. M., et al. 2019, Monthly Notices of the Royal Astronomical Society: Letters, http://oup.prod.sis.lan/mnrasl/advance-article-pdf/doi/10.1093/mnrasl/slz007/27503647/slz007.pdf. https://doi.org/10.1093/mnrasl/slz007

Kawaguchi, K., Shibata, M., & Tanaka, M. 2020, The Astrophysical Journal, 889, 171. http://dx.doi.org/10.3847/1538-4357/ab61f6

Korobkin, O., Wollaeger, R. T., Fryer, C. L., et al. 2021, The Astrophysical Journal, 910, 116. http://dx.doi.org/10.3847/1538-4357/abe1b5

Lai, X., Zhou, E., & Xu, R. 2019, The European Physical Journal A, 55, 60. https://doi.org/10.1140/epja/i2019-12720-8

Margalit, B., & Metzger, B. 2017, The Astrophysical Journal Letters, 850, doi:10.3847/2041-8213/aa991c

Metzger, B. D. 2017, Living Rev. Rel., 20, 3

—. 2019, Living Reviews in Relativity, 23, 1. https://doi.org/10.1007/s41114-019-0024-0

Metzger, B. D., Piro, A. L., & Quataert, E. 2008, Monthly Notices of the Royal Astronomical Society, 390, 781. +http://dx.doi.org/10.1111/j.1365-2966.2008.13789.x

Miller, J. M., Ryan, B. R., Dolence, J. C., et al. 2019, PhRvD, 100, 023008

Most, E. R., Weih, L. R., Rezzolla, L., & Schaffner-Bielich, J. 2018, Phys. Rev. Lett., 120, 261103. https://link.aps.org/doi/10.1103/PhysRevLett.120.261103

Nedora, V., Schianchi, F., Bernuzzi, S., et al. 2022, Class. Quant. Grav., 39, 015008

Nicholl, M., Margalit, B., Schmidt, P., et al. 2021, Monthly Notices of the Royal Astronomical Society, 505, 3016–3032. http://dx.doi.org/10.1093/mnras/stab1523

Oechslin, R., & Janka, H.-T. 2006, Monthly Notices of the Royal Astronomical Society, 368, 1489–1499. http://dx.doi.org/10.1111/j.1365-2966.2006.10238.x

Perego, A., Radice, D., & Bernuzzi, S. 2017, ApJL, 850, L37

Pian et al. 2017, Nature, 551, 67 EP . http://dx.doi.org/10.1038/nature24298

Radice, D., Perego, A., Hotokezaka, K., et al. 2018, Astrophys. J., 869, 130

Radice, D., Perego, A., Hotokezaka, K., et al. 2018, ApJ, 869, 130

Radice, D., Perego, A., Zappa, F., & Bernuzzi, S. 2018, The Astrophysical Journal Letters, 852, L29. http://stacks.iop.org/2041-8205/852/i=2/a=L29

Ristic, M., Champion, E., O'Shaughnessy, R., et al. 2021, Interpolating Detailed Simulations of Kilonovae: Adaptive Learning and Parameter Inference Applications, , , arXiv:2105.07013

Rosswog, S., Korobkin, O., Arcones, A., Thielemann, F. K., & Piran, T. 2014, Mon. Not. Roy. Astron. Soc., 439, 744

Siegel, D. M., & Metzger, B. D. 2018, ApJ, 858, 52

Smartt et al. 2017, Nature, 551, 75 EP . http://dx.doi.org/10.1038/nature24303

Tanaka, M. 2016, Advances in Astronomy, 2016, 1. http://dx.doi.org/10.1155/2016/6341974

Tanaka, M., Kato, D., Gaigalas, G., & Kawaguchi, K. 2020, MNRAS, 496, 1369

Watson, D., Hansen, C. J., Selsing, J., et al. 2019, Nature, 574, 497

Wollaeger, R. T., Korobkin, O., Fontes, C. J., et al. 2017, ArXiv e-prints, arXiv:1705.07084

Wollaeger, R. T., Fryer, C. L., Chase, E. A., et al. 2021, The Astrophysical Journal, 918, 10. http://dx.doi.org/10.3847/1538-4357/ac0d03

Wu, M.-R., Fernández, R., Martínez-Pinedo, G., & Metzger, B. D. 2016, MNRAS, 463, 2323

Zhu, Y. L., Lund, K. A., Barnes, J., et al. 2021, The Astrophysical Journal, 906, 94. http://dx.doi.org/10.3847/1538-4357/abc69e