

Interpolating detailed simulations of kilonovae: Adaptive learning and parameter inference applications

M. Ristic¹, E. Champion¹, R. O’Shaughnessy¹, R. Wollaeger^{2,3}, O. Korobkin^{2,3}, E. A. Chase^{2,4,5,6},
C. L. Fryer^{2,3,7,8}, A. L. Hungerford^{2,4} and C. J. Fontes^{2,4}

¹Center for Computational Relativity and Gravitation, Rochester Institute of Technology, Rochester, New York 14623, USA

²Center for Theoretical Astrophysics, Los Alamos National Laboratory, Los Alamos, New Mexico, 87545, USA

³Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, New Mexico, 87545, USA

⁴Computational Physics Division, Los Alamos National Laboratory, Los Alamos, New Mexico, 87545, USA

⁵Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, Evanston, Illinois, 60201, USA

⁶Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208, USA

⁷The University of Arizona, Tucson, Arizona 85721, USA

⁸Department of Physics and Astronomy, The University of New Mexico, Albuquerque, New Mexico 87131, USA



(Received 15 June 2021; accepted 2 November 2021; published 21 January 2022)

Detailed radiative transfer simulations of kilonovae are difficult to apply directly to observations; they only sparsely cover simulation parameters, such as the mass, velocity, morphology, and composition of the ejecta. On the other hand, semianalytic models for kilonovae can be evaluated continuously over model parameters, but neglect important physical details which are not incorporated in the simulations, thus introducing systematic bias. Starting with a grid of two-dimensional anisotropic simulations of kilonova light curves covering a wide range of ejecta properties, we apply adaptive learning techniques to iteratively choose new simulations and produce high-fidelity surrogate models for those simulations. These surrogate models allow for continuous evaluation across model parameters while retaining the microphysical details about the ejecta. Using a code for multimessenger inference developed by our group, we demonstrate how to use our interpolated models to infer kilonova parameters. Comparing to inferences using simplified analytic models, we recover different ejecta properties. We discuss the implications of this analysis which is qualitatively consistent with similar previous work using detailed ejecta opacity calculations and which illustrates systematic challenges for kilonova modeling. An associated data and code release provides our interpolated light-curve models, interpolation implementation which can be applied to reproduce our work or extend to new models, and our multimessenger parameter inference engine.

DOI: [10.1103/PhysRevResearch.4.013046](https://doi.org/10.1103/PhysRevResearch.4.013046)

I. INTRODUCTION

As exemplified by GW170817, neutron star mergers are empirically known to produce a rich array of multimessenger emission [1,2]. The presence of matter is most unambiguously indicated by electromagnetic emission from nuclear matter ejected during the merger itself, which produces distinctive “kilonova” emission [3–6] via radioactive heating of this expanding material. Kilonova observations can provide insight into uncertain nuclear physics [7–11] and help constrain the expansion rate of the universe [12–15], particularly in conjunction with gravitational wave observations [1,2,16–21].

In principle, kilonova observations encode the amount and properties of the ejected material in their complex multiwavelength light curves (and spectra) [5,22,23]. For example, several studies of GW170817 attempted to infer the amount of

material ejected [22–31]. In practice, these observations have historically been interpreted with semianalytic models, as they can be evaluated quickly and continuously over the parameters which characterize potential merger ejecta. However, it is well known that these semianalytic models contain oversimplified physics of already simplified anisotropic radiative transfer calculations [32–34] that neglect detailed anisotropy, radiative transfer, opacity, sophisticated nuclear reaction networks, and composition differences.

To circumvent these biases, some groups have attempted to construct surrogate kilonova light-curve models, calibrated to detailed radiative transfer simulations [22,25,35]. For example, Coughlin *et al.* [22] used Gaussian process (GP) regression of principal components to construct a multiwavelength surrogate calibrated to a fixed three-dimensional grid of simulations [36], describing flux F_k from a single component of ejected material. This study generated a “two-component” ejecta model by adding the fluxes of two independent calculations ($F = F_1 + F_2$), ignoring any photon reprocessing effects. More recently, Heinzel *et al.* [25] applied this method to construct an anisotropic surrogate depending on two components M_1, M_2 and viewing angle, calibrating to their own anisotropic radiative transfer calculations. They also included

reprocessing effects, showing that their previous simplified approach which treats the radiation from each of the two components of the outflow independently introduces biases in inference for the components' parameters. These strong reprocessing or morphology-dependent effects are expected in kilonova light curves [37–40]. Finally, a recent study by Breschi *et al.* [24] favored an anisotropic multicomponent model.

In this work, extending [35], we apply an adaptive learning technique to generate surrogate light curves from simulations of anisotropic kilonovae. Starting with a subset of 36 simulations reported in [34], we use these adaptive learning methods to identify new simulations to perform, refining our model with 448 simulations so far. We apply our surrogate light curves to reassess the parameters of GW170817. We distribute the updated simulation archive, our current best surrogate models, and our training algorithms in Ref. [41].

This paper is organized as follows. In Sec. II we describe the kilonova simulation family we explore in this study and the active learning methods we employ to target new simulations to perform. We also briefly comment on our model's physical completeness. In Sec. III we describe the specific procedures we employed to interpolate between our simulations to construct surrogate light curves. In Sec. IV we describe how we compare observations to our surrogate light curves to deduce the (distribution of) best-fitting two-component kilonova model parameters for a given event. We specifically compare our model to GW170817. In Sec. V we describe how our surrogate models and active learning fit into the broader challenges of interpreting kilonova observations. We conclude in Sec. VI.

II. KILONOVA SIMULATION PLACEMENT

A. Kilonova simulations

The kilonova simulations described in this work adopt a similar setup as and expand on the work of [34]. The simulations discussed throughout were generated using the SUPERNU [41] time-dependent radiative transfer code, using tabulated binned opacities generated with the Los Alamos suite of atomic physics codes [42,43]. We use results from the WINNET code [44] to determine radioactive heating and composition effects. We employ the thermalization model of [45], but use a gray Monte Carlo transport scheme for gamma ray energy deposition [32].

The ejecta model is based on a symmetrically shaped ideal fluid expanding in vacuum described by the Euler equations of ideal hydrodynamics. The assumption of a radiation-dominated polytropic equation of state allows for an analytic representation of the ejected mass M and average velocity $\bar{v}$ as a function of initial central density ρ_0 , initial time t_0 , and the velocity of the expansion front $v_{\max}$ [Eqs. (11) and (12) in [32]]. When combined with Monte Carlo based radiative transfer and a specified elemental composition for the ejecta, the code produces time- and orientation-dependent spectra. Convolution of these spectra with standard observational filters produces light curves such as the ones in Figs. 1 and 4.

Real neutron star mergers have (at least) two mechanisms for ejecting material, denoted as dynamical and wind ejecta




FIG. 1. Bolometric luminosities of initial and adaptively placed simulations: The top panel shows the $\log_{10}$ bolometric luminosity in CGS units versus time in days for the simulations we initially used to train our grid. These simulations all extend out to roughly 8 days. The bottom panel shows the bolometric light curves for our adaptively placed simulations overlaid on top of the initial grid light curves. Most of these simulations extend past 32 days. Both panels exhibit significant diversity in behavior and timescale.

[46]. Due to the difference in formation mechanisms of dynamical and wind ejecta [5], a multicomponent approach is necessary for accurate modeling. Each of the two types of ejecta, dynamical and wind, is modeled by a separate component with a specified morphology, elemental composition, ejecta mass, and ejecta velocity. The components are modeled together as one radiative outflow [32]. The thermal decay energy is treated by mass weighting between the components where they overlap. The end product represents a time-dependent spectral energy distribution contained in 54 angular bins, equally spaced in $\cos\theta$ from 1 to -1 . For the purposes of this study, the spectra are convolved with broadband filters to produce a series of broadband light curves. Specifically, we use the Vera C. Rubin Observatory *grizy* filters for optical and near-infrared bands, 2MASS *JHK* filters for longer wavelength near infrared, and the mid-infrared *S* filter for the Spitzer 4.5- μm band. For each band and emission direction, we estimate the AB magnitude for that filter, defined for a source at 10 pc in terms of energy flux in $\text{g cm}^2/\text{s}^2 F_\nu$ per unit frequency via $m_{X,AB} = -2.5 \log_{10} \langle F_\nu \rangle - 48.6$. All observations used in this work are provided or are translated into this AB -magnitude system [47–49]. Because our simulations tend toward reflection-symmetric behavior across the $z = 0$ plane, we only consider the independent information

TABLE I. Kilonova simulation parameters: Within the framework of models explored in [34], parameters of the initial kilonova simulations used to initialize our adaptive learning process in this work. All simulations used in this work adopt a two-component model where the morphology and composition of each component are fixed.

Ejecta	Morphology	Y_e	M_{ej} $M_\odot$	$\bar{v}$ c
Dynamical	Torus	0.04	0.001, 0.01, 0.1	0.05, 0.3
Wind	Peanut	0.27	0.001, 0.01, 0.1	0.05, 0.3

contained in the upper half ($z > 0$) of these angular bins. To reduce the acquisition cost of each simulation, we evolved each kilonova simulation in our initial grid out to 7.66 days. To minimize data handling and training cost, unless otherwise noted, we manipulate a subset of our simulation output based on a log-uniform grid. For the initial simulations, this log-uniform grid consists of 191 time points ranging from 1/8 to 7.66 days. For the remaining simulations, this grid is extended in log-time to cover their available duration, up to a maximum of 64 days. Because of several systematics associated with modeling emission at early times (e.g., in the ionization states of the medium and in the contribution from and interaction with any strong jet), we do not report on behavior prior to 3 hours postmerger. In this work, we use the orientation-averaged luminosity for simulation placement, but reconstruct the luminosity continuously in angle and time.

The original simulation hypercubes discussed in [32,34] consider multiple wind ejecta morphologies and compositions. To simplify the dimensionality of the problem, this work only considers simulations from the initial grid with a peanut-shaped morphology [37] and lower $Y_e = 0.27$ composition describing the wind ejecta. Table I summarizes the parameters for the 36 simulations in our four-dimensional hypercube and highlights variation in only ejected mass M and average velocity $\bar{v}$ for each of the two components: the mass and velocity of the dynamical and wind ejecta, denoted henceforth as M_d , v_d , M_w , v_w . Every simulation in our hypercube adopts the same morphologies for the dynamical and wind ejecta, respectively. This initial simulation hypercube thus consists of only two of the three velocities and three of the five masses explored in the companion study [34].

As expected and discussed elsewhere [34], these simulations exhibit significant viewing-angle dependence on the relative speed of the components. The obscuration of the wind by the dynamical ejecta becomes less significant closer to the symmetry axis and the peanut morphology itself also produces orientation dependence. The two-component model shows “blanketing” of slow blue components by fast red components [50]. Also expected and observed are qualitative trends versus the component masses and velocities: more wind ejecta mass increases the g-band luminosity along the symmetry axis.

B. Illustrating systematics of kilonova simulations

Before extensively discussing our ability to reproduce this specific family of simulations, we first comment on their systematic limitations. Our simulation archive explores only a




FIG. 2. Impact of removing zirconium: Solid and dashed lines show simulations with otherwise identical assumptions about composition, morphology, and velocity structure, differing only by the presence (solid) and elimination (dashed) of Zr. The selected simulation parameters $M_d = 0.01M_\odot$, $v_d = 0.3c$, $M_w = 0.01M_\odot$, and $v_w = 0.15c$ are our closest-matching representations of the simulation parameters considered during the Zr-omitting study in [56].

limited range of initial conditions for the ejecta, with specific assumptions about the composition, morphology, and velocity profiles; with specific assumptions about nucleosynthetic heating; and with specific assumptions about (the absence of) additional power and components, such as a jet or a central source to provide additional power or light [51–54].

Several previous studies have indicated that these and other aspects of kilonova simulations can noticeably impact the outcome [8,32,33,37,38,55,56]. Where possible, we very briefly comment on how current and previous SUPERNU simulations’ results change when making similar changes in assumptions.

Prior work with SUPERNU has explored the impact of composition [33]. However, recently, Kawaguchi *et al.* [56] (henceforth K20) demonstrated that Zr makes a substantial contribution to the final light curve. Figure 2 shows how our simulations depend on a similar change in composition, noting substantial change in the late-time optical light curves when we remove Zr.

As demonstrated by many previous studies using SUPERNU, the morphology and velocity structure also have a notable impact on the post-day light-curve behavior [32,37,38]. Several other groups have demonstrated similar strong morphology and orientation dependence in their work [24,25,56,57]. For example, in their Fig. 8, K20 demonstrate how the light curve

changes when a specific polar component of the ejecta is removed.

Uncertain nuclear physics inputs also propagate into notable uncertainties about the expected light curve; see, e.g., [7,8]. Even for the same morphology and amount of ejecta, nuclear physics uncertainties can modify the effective heating rate, particularly for material with low Y_e which has the greatest prospect for producing r -process elements.

Given limited exploration of possible kilonova initial conditions and physics, we can only at present quantify the uncertainties of the type listed above. In future work, we will employ our parametrized models to assess the impact of these uncertainties on inferences about kilonova parameters. Future work could require kilonova models which include equation of state (EOS) parameters to enable joint inference which also simultaneously constrains the equation of state.

As discussed elsewhere [34], at late times some light curves show a modest deficit of blue light (g -band) relative to observations of GW170817 (unless the dynamical ejecta mass is large). Notably, our g -band light curves fall off significantly more rapidly after their peak in all viewing directions and for most parameters considered here. Previous work with other morphologies also recovers similar falloff in these bands (see, e.g., [28]), though additional components could conceivably contribute. Similar g -band behavior has been seen in other detailed kilonova simulations; see, e.g., Fig. 12 in [58]. As noted above, this behavior depends on the assumed composition, notably Zr.

C. Interpolation methodology

In this work, we principally interpolate using Gaussian process (GP) regression. In GP regression, given training data pairs (x_a, y_a) , the estimated function $\hat{y}(x)$ and its variance $s(x)^2$ are approximated by

$$\hat{y}(x) = \sum_{a,a'} k(x, x_a) K_{aa'}^{-1} y_{a'}, \quad (1a)$$

$$s(x)^2 = k(x, x) - k(x, x_a) K_{aa'}^{-1} k(x_{a'}, x), \quad (1b)$$

where the matrix $K_{aa'} = k(x_a, x_{a'})$ and where the function $k(x, x')$ is called the kernel of the Gaussian process. In this work, unless otherwise noted, we used a squared-exponential kernel and a white-noise (diagonal) kernel

$$k(x, x') = \sigma_o^2 e^{-(x-x')Q(x-x')/2} + \sigma_n^2 \delta_{x,x'}, \quad (2)$$

where Q is a diagonal matrix of possible length scales and σ_o, σ_n are hyperparameters that characterize the amount of noise allowed in the problem. The other interpolation method considered in this work was random forest (RF) regression [59]. Unlike the GP, the RF output had no error quantification and was used primarily as a consistency check on the Gaussian process prediction. Unless otherwise noted, we performed all GP and RF regression with SCIKIT-LEARN [60].

Because of the substantial dynamic range of our many outputs, we interpolate the $\log_{10}$ luminosity (for placement) or AB magnitudes (for all other results). Unless otherwise noted, we quantify the performance of our interpolation with the root-mean-square (RMS) difference between our prediction

and the true value

$$\ell^2 = \frac{1}{n} \sum_{j=1}^n [y_j - \log_{10}(L_{bol})_j]^2. \quad (3)$$

(This expression overweights the importance of large errors when the source is not detectable at late times; see the Appendix.)

We employ GP interpolation in two standard use cases. In the first case, used for our exported production results, we interpolate the AB magnitude $m_\alpha(t_*|\Lambda)$ at some fixed reference time t_* and band α versus our four simulation hyperparameters (and, in the end, also across the extrinsic parameters of angle and wavelength) contained in Λ . In this case, the prediction $y(x_a)$ has a single scalar value at each point; the x_a refer to model hyperparameters; and the interpolation provides us with a scalar function of four or more variables. GP regression [Eq. (1)] provides an error estimate for m_α at this specific time t_* , which in general will depend on time.

In the second case, used for simulation placement, we interpolate the logarithm bolometric luminosity *light curve* $\log_{10} L_{bol}(t|\Lambda)$ versus *all time*. (In terms of each simulation's spectrum, the bolometric luminosity is $L_{bol} = 4\pi R^2 \int_0^\infty F_\nu d\nu$ where $R = 10$ pc.) In this case, the prediction $\vec{y}(x_a)$ is vector valued at each point; the x_a refer to model hyperparameters; and the interpolation provides us with a vector-valued function of four or more variables. For simplicity and given our use case, we reduce our error estimate to a single overall value for the entire light curve, reflecting the overall uncertainty in $\vec{y}(x_a)$.

D. Active learning scheme

Gaussian processes have long been used for active learning because they provide an error estimate: followup simulations can be targeted in regions with the largest expected error (and thus improvement) [61]. We follow this approach in our active learning scheme; see [62–69] for a broader discussion of active learning methods and their tradeoffs. To reduce the data volume needed for targeting followup simulations, we used vector-valued interpolation as described above, applied to *orientation-averaged outputs* of our simulations. This approach has the substantial advantage of providing a single error estimate per light curve (both in training and off sample), which we can immediately use as an objective function in a minimization algorithm.

We pursued an active learning simulation placement approach in order to maximally explore the parameter space and reduce the amount of redundant information obtained from each new simulation. The subset of 36 light curves discussed in Sec. II A was used as the initial training set. Thousands of parameter combinations were subsequently drawn from uniform distributions with maxima and minima matching those of the varied parameters in Table I. Each of these parameter combinations was evaluated by an interpolator to produce an initial light-curve prediction as well as an error on the entire light-curve output. The prediction with the largest error across all the tested parameter combinations was selected as the next placed simulation.




FIG. 3. Impact of adaptive placement on interpolation: Example of interpolation output at a point with large predicted fitting error, both before and after placing the simulation. In both panels, the solid black curve shows the true simulated bolometric light curve versus time. The red band shows the GP-predicted one-sigma error bar about the expected value. Top panel: predictions from our RF and GP interpolations versus time. The large error and low practical utility of the GP fit are apparent. Bottom panel: after including this simulation in the training set, the revised RF and GP predictions much more closely conform with this specific simulation as expected.

E. Prediction improvement and interpolation results

We verified our active learning strategy for simulation placement by randomly sampling combinations of parameters and creating two light-curve predictions based on those parameters. The first prediction was trained solely on our initial grid of simulations from Sec. II A, while the second prediction was trained on the same initial grid, but with an added simulation output characterized by the aforementioned random combination of parameters. Figure 3 shows these before- and after-inclusion predictions which show that, as expected, the GP interpolation capability is improved. This pair of figures anecdotally illustrates the degree to which new training data improve our surrogate light-curve models.

With over 400 placed simulations since the start of the active learning process, the training library is built up enough to allow for physically meaningful interpolation of off-sample events. The performance of our adaptive learning is best illustrated with our production-quality interpolation scheme, illustrated in Fig. 4 and described in the next section.




FIG. 4. Off-sample interpolation with original and refined grid: Example of an interpolated stitched fixed-time prediction compared to a simulation output created from the same corresponding input parameters. The top panel shows our estimate based on the initial 36 simulations; the bottom panel shows the result after adaptive learning. Different colors denote different filter bands, described in the legend. The dashed lines show full simulation output for each band. The colored points show our interpolated bolometric magnitude predictions at the 191 evaluation times. The solid lines show our final interpolated light curves, interpolating between the points shown. The largest error in this example occurs for the g band at late times. The simulated parameters and viewing angle for this configuration are $M_d = 0.097\,050M_\odot$, $M_w = 0.083\,748M_\odot$, $v_d = 0.197\,642c$, and $v_w = 0.297\,978c$, viewed on axis ($\theta = 0$). The exaggerated modulations in the top panel's solid lines and dotted curves illustrate interpolation failures, arising from adopting an initially insufficient training set.

Despite producing many followup simulations, we achieve success with a very sparse coverage of our parameter space. To illustrate the sparsity of our parameter-space coverage, and how slowly our added simulations increase coverage, we evaluated the median “intersimulation” distance, using a simple Euclidean (L^2) norm over $\log_{10} L_{bol}(t_k)$ for several reference times t_k . As expected given the high apparent dimension of our output, this median distance changes very slowly with n , owing to the large effective dimension of the output light curves. The median distance is also larger than the residual error in our fit, as reported below. The success of our interpolation relies not on an overwhelmingly large training sample, but on the smoothness and predictability of our physics-based light curves.

III. LIGHT-CURVE INTERPOLATION

A. Stitched fixed-time interpolation

To efficiently interpolate across the whole model space, we follow a strategy illustrated in Fig. 1 of [70]: we pick several fiducial reference times t_q (and angles); use GP interpolation to produce an estimate $m_\alpha(t_q|\Lambda)$ versus Λ ; interpolate in time to construct a continuous light curve at the model hyperparameters Λ at each reference angle; and then interpolate in angle to construct a light curve for an arbitrary orientation. For an error estimate, we stitch together the error estimates in each band to produce a continuous function of time. Figure 4 shows the output of our interpolation (smooth lines), compared to a validation simulation at the same parameters (dashed lines). Our predictions generally agree, though less so for the shortest wavelengths at the latest times. Subsequent figures also illustrate the typical GP error estimate, which is usually $O(0.1)$ in $\log_{10} L$ for most bands and times considered.

B. Trends identified with interpolated light curves

In Fig. 5 we show the results of our fit evaluated at a fixed viewing angle ($\theta = 0$), varying one parameter at a time continuously, relative to a fiducial configuration with $M_d = M_w = 0.01M_\odot$, $v_w/c = v_d/c = 0.05$. The fixed value for the ejected mass of $M = 0.01M_\odot$ was chosen as the middle ground of the initial grid's sampled mass space, which does not introduce any biases toward lighter or heavier masses. Since no similar central value was initially available for the velocity parameters, the lower value was selected in the case of both components. The slower velocity resulted in the ejecta not dissipating as quickly and allowed for more variation in the light curves as the nonstatic parameter was varied. For this viewing angle, changes in the amount and velocity of the dynamical ejecta have relatively modest effect, in large part because that ejecta is concentrated in the equatorial plane. By contrast, changes in the mostly polar wind ejecta have a much more substantial impact on the polar light curve ($\theta = 0$). Specifically, increasing the amount of wind ejecta brightens and broadens the light curve, as expected from classic analytic arguments pertaining to how much material the light must diffuse through [5,71–73]. Similarly, increasing the velocity of wind ejecta causes the peak to occur at earlier times (diffusion is easier) and be brighter.

C. Interpolation in viewing angle

All of the interpolated light curves discussed thus far have been trained at some fixed viewing angle. In Fig. 6, we explore the interpolation of several families of models, each of which was trained using simulation data at a different viewing angle. The symmetry of the ejecta across the orbital plane allows for the assumption that any angular variation between 0 and $\pi/2$ can simply be mirrored across the symmetry axis.

Figure 6 indicates that the first day postmerger does not introduce much angular variation and, as such, is quite well predicted even when interpolating across only 11 angles. After 1 day, the luminosity across different angles begins to change considerably as the peanut-shaped wind ejecta becomes more dominant. *Particularly at late times, there is a strong angular variability which manifests near the orbital plane, most*




FIG. 5. Interpolated and simulated g-band light curves: In this figure, we generate $\log L_g(t|\Lambda)$ for a one-parameter family of simulations Λ where either one of the M parameters vary from $0.001M_\odot$ to $0.1M_\odot$ or one of the v parameters vary from $0.05c$ to $0.3c$, and the viewing angle is $\theta = 0$. The remaining model parameters are fixed to $(M/M_\odot, v/c) = (0.01, 0.05)$. Contours in M are uniform in $\log M$, while those for v are linearly uniform. For comparison, the heavy dashed lines show the initial training simulation results for the two parameter end points. The g-band light curve has the largest dynamic range and is the most sensitive to interpolation errors; notably, the interpolation does not always conform tightly to the underlying simulation data at late times.




FIG. 6. Interpolation of g -, y -, and K -band luminosity at different viewing angles: This figure compares the g -, y -, and K -band luminosity at select times as a function of viewing angle. The solid points represent fixed angles at which the different families of models were trained. The solid lines connecting the points indicate the interpolated prediction of the angular variation at some given time in the light curve. The dashed lines represent the simulation data and show the true angular variation. The shaded regions denote the 1σ error estimate derived from our Gaussian process fit versus time, extended in angle.

strongly apparent in the blue (g) and near-infrared (K) bands. In the blue bands, the angular variation reflects lanthanide curtaining; in the red bands, the angular variation reflects red emission from the late-peaking red dynamical ejecta. (At




FIG. 7. Light curve versus time for selected angles and bands: comparison to Fig. 7 of [58] indicating angular dependence of light-curve predictions across the g , z , and K bands.

the latest times and faintest luminosities along the equatorial plane, numerical uncertainty in our Monte Carlo simulations is apparent in the light-curve results.) In all panels, the solid band denotes an estimated error bar from our GP fit in time, extended in angle. For comparison, Fig. 7 shows the light curve behavior for increasing viewing angle for short-, mid-, and long-wavelength filters.

D. Predictive accuracy versus time, angle grid sizes

To better understand the systematic limitations and computational inefficiencies introduced by our stitched-time interpolation grid, we investigated the accuracy of our fits when only using a subset of the time or angular grid.

First, we consider a simple analysis of loss of predictive accuracy as the number of GP interpolators used to make a surrogate light curve is decreased. We denote $t \in T$ as the subset of times represented by the GP interpolators used to make a prediction, T as the total available number of time points, and thus interpolators, which can be used to make a light curve, and $\bar{t}$ as all the other times in T which are not represented by t such that $t \cap \bar{t} = \emptyset$ and $t \cup \bar{t} = T$.

Thus, when using any number of interpolators at times $t \in T$ which is less than the total number of possible time points T , we first generate predictions $y(t)$ with the chosen subset of interpolators. These predictions $y(t)$, along with the times t at which the predictions were made, are then used as inputs for SCIPY's UnivariateSpline method from which the remainder of the light curve $z(t) = f(t, y(t))$ is constructed, where in this last case the function can be evaluated $\forall t \in T$.

Figure 8 shows how the average residual between on-sample light-curve predictions and the respective simulation data changes as a function of the number of time points used as the base for constructing the time-interpolated light curve. For the current scheme, we can remove up to roughly 75% of the initial set of time points without substantially diminishing our overall accuracy. Future work will explore smarter selection of representative time points in an effort to further reduce the number of interpolators which can be removed without significant loss of accuracy.




FIG. 8. Average residual as a function of number of considered time points: A plot of the average residuals between on-sample time-interpolated light curves and the respective simulation data as a function of how many time points are used to generate the light curves. In each case, we drew the respective number of samples from a log-uniform distribution between the start and end time of our light curves.

IV. PARAMETER INFERENCE OF RADIOACTIVELY POWERED KILONOVAE

In this section, we describe and demonstrate the algorithm we use to infer kilonova parameters given observations, using the interpolated light-curve model above. Unless otherwise noted, for simplicity all calculations in this section assume the kilonova event time and distance are known parameters. We likewise assume observational errors are understood and well characterized by independent Gaussian magnitude errors in each observation, and that our model families include the underlying properties of the source (i.e., we neglect systematic modeling errors due to the parameters held constant in our simulation grid: morphology, initial composition, etc.).

A. Framework and validation

As in many previous applications of Bayesian inference to infer parameters of kilonovae [22,23,25–27], we seek to compare the observed magnitudes x_i at evaluation points i (denoting a combination of band and time) to a continuous model that makes predictions $m(i|\theta)$ [henceforth denoted by $m_i(\theta)$ for brevity] which depend on some model parameters θ . Bayes theorem expresses the posterior probability $p(\theta)$ in terms of a prior probability $p_{\text{prior}}(\theta)$ for the model parameters θ and a likelihood $\mathcal{L}(\theta)$ of all observations, given the model parameters, as

$$p(\theta) = \frac{\mathcal{L}(\theta)p_{\text{prior}}(\theta)}{\int d\theta \mathcal{L}(\theta)p_{\text{prior}}(\theta)}. \quad (4)$$

Unless otherwise noted, for simplicity we assume the source sky location, distance, and merger time are known. We adopt a uniform prior on the ejecta velocity $v/c \in [0.05, 0.3]$ and a log-uniform prior on the ejecta masses $m/M_{\odot} \in [10^{-3}, 0.1]$. We assume the observations have Gaussian-distributed *magnitude* errors with presumed known observational (statistical) uncertainties σ_i , convolved with some additional unknown

systematic uncertainty σ , so that

$$\ln \mathcal{L}(\theta) = -0.5 \sum_{i=1}^n \left[\frac{[x_i - m_i(\theta)]^2}{\sigma_i^2 + \sigma^2} + \ln [2\pi(\sigma_i^2 + \sigma^2)] \right], \quad (5)$$

where the sum is taken over every data point in every band used in the analysis. In tests, we treat σ as an uncertain model parameter, de facto allowing for additional systematic observational uncertainty (or for some systematic theoretical uncertainty). For our GP surrogate models, we set σ to the estimated GP model error.

Unlike prior work, we eschew Markov-chain Monte Carlo, instead constructing the posterior distribution by direct Monte Carlo integration as in [74,75]. To efficiently capture correlations, we employ a custom adaptive Monte Carlo integrator; see Wofford *et al.* [76] for implementation details. In the Appendix, we describe several tests we performed to validate this inference technique using synthetic kilonova data drawn from a previously published semianalytic kilonova model. Our tests include recovering the parameters of 100 synthetic kilonova sources. In future work, we will demonstrate how our parameter inference method can be incorporated efficiently and simultaneously with gravitational wave (GW) parameter inference with the rapid iterative fitting (RIFT) parameter estimation pipeline [75].

B. Inference with surrogate kilonova model

Figure 9 demonstrates parameter inference using our surrogate light curves, for a synthetic source generated using our own model. As expected, we can recover a known source, including constraining the viewing angle θ . Figure 10 performs a similar test, but now using a specific simulation, without interpolation. As expected given our adopted systematic error, we recover the simulation parameters. Finally, Fig. 11 repeats the test above, using the semianalytic model described in the Appendix. This comparison emphatically demonstrates large systematic differences between this semianalytic model and our detailed simulations.

C. Example: GW170817

SUPERNU-based kilonova models have already been successfully used to interpret GW170817, though as noted previously these models have a rapid falloff in the late-time optical magnitudes that is not present in the observations; see [28]. Because of the close proximity of GW170817, only distance modulus (but not redshift) corrections are needed to translate our predictions to apparent magnitudes which can be directly compared to electromagnetic observations. Observational results are taken from a Ref. [23] compilation of photometry reported in [27,28,31,77–84]. Figure 12 shows the results of directly comparing our extended simulation archive directly to observations of GW170817, selecting for simulations (parameters and angles) with the highest overall likelihood. The solid black curves in these figures show the 50 highest-likelihood configurations, where the likelihood requires simultaneously reproducing all observed bands. Except for reddest three bands (*JHK*), many simulations compare extremely favorably to the observations. The parameters of these




FIG. 9. Synthetic source recovery with surrogate model: Recovery of a parameter of a known two-component surrogate kilonova model, using inference based on our interpolated model. Solid black curves show results adopting a strong angular prior motivated by radio observations of GW170817. Top panel: synthetic light-curve data in several bands. Bottom panel: inferred distribution of the four model parameters, and viewing angle. The blue cross denotes the injected values. Red contours show results without adopting a prior on observing angle; black contours show results inferred when adopting a prior on viewing angle consistent with observations of GW170817.

simulations, however, do not represent the optimal parameters of this model family: because our placement algorithm minimizes interpolation error, the selected points preferentially occur at the edges of our domain. Finally, for the reddest band (K), our fits exhibit notable systematic uncertainty relative to the underlying simulation grid.

We have performed parametric inference on GW170817 using our surrogate light-curve model to the underlying SUPERNU results. Motivated by the direct comparisons above, we perform two analyses. In the first, we use all observing data at all times. In the second, we omit the reddest (K) band.




FIG. 10. Simulation parameter recovery with surrogate model: Recovery of a parameter of a known two-component kilonova simulation, using inference based on our interpolated model. The parameters corresponding to the relevant simulation are $M_d = 0.052780M_\odot$, $v_d = 0.164316c$, $M_w = 0.026494M_\odot$, and $v_w = 0.174017c$. Top panel: synthetic light-curve data in several bands. Bottom panel: inferred distribution of the four model parameters.

Figure 13 shows the results of these comparisons. Because of the systematic fitting uncertainties at late times, we highlight the analysis omitting K -band observations as our preferred result. Although previously reported inferences about ejecta masses cover a considerable dynamic range (see, e.g., Fig. 1 in [85]), our inferred masses are qualitatively consistent with selected previous estimates including previous inferences with similar SUPERNU models [28] and recent surrogate models adapted to simplified multidimensional radiative transfer [21]. Notably, however, we infer a large amount of “dynamical” (red, lanthanide-rich) ejecta mass [i.e., $M_{ej} \simeq O(1/30)M_\odot$], more dynamical ejecta than wind, and the velocities for the dynamical and wind components are inverted relative to customary expectations (i.e., $v_d < v_w$). Our dynamical and wind




FIG. 11. Simulation parameter recovery with analytic model: Recovery of a parameter of a known two-component kilonova *simulation*, using inference based on the simplified analytic model described in the Appendix. The analytic model cannot fit our simulation data. While only a one-component fit is shown, similar results arise when employing multiple components. The parameters corresponding to the relevant simulation are $M_d = 0.01M_{\odot}$, $v_d = 0.3c$, $M_w = 0.01M_{\odot}$, and $v_w = 0.3c$. Top panel: synthetic light-curve data in several bands, including error bars on both the synthetic data and posterior light-curve predictions. The analytic model cannot fit our simulated data well. Bottom panel: inferred distribution of the *four* model parameters. The blue cross denotes the injected values.

component masses differ from typical semianalytical treatments (see, e.g., [86]).

We also weakly constrain the misalignment angle between the outflow and the line of sight to be consistent with independent late-time radio observations, which (weakly supplemented with gravitational wave constraints) constrain the opening angle between the jet and the line of sight to be roughly 20° [20,28,78,87–89]; see also [21] for previous, weaker constraints on binary alignment from kilonova observations. Reanalyzing the light curves using this prior information, imposed as a Gaussian prior on θ with mean 20°

with $\sigma_\theta = 5^\circ$, we find modestly improved overall constraints on the ejecta; see also, e.g., Fig. 2 of [21] for previous, weaker constraints derived using joint radio and kilonova observations.

V. DISCUSSION

We have demonstrated that our surrogate models can be operationally compared to real kilonova observations, allowing us to deduce what the range of parameters for the original simulation family best fits the observations. In this section, we emphasize several systematic limitations of our approach, to more clearly distinguish the ways in which the answers so obtained could differ from a description of physical reality. When possible, we comment on ways in which these systematic limitations could be mitigated with future work.

First and foremost, our surrogate models introduce some modest bias, being an imperfect representation of the simulations they mimic. We have demonstrated that these errors are relatively small (see Figs. 4 and 6). In this work, we principally employed two standard interpolation methods (GP and RF interpolations) to construct synthetic light curves. Recent substantial advances in machine learning have led to many new algorithms and architectures for adaptive learning and interpolation. Our prior work suggests that neural networks can also usefully interpolate kilonova light curves [35], which we will describe at greater length in future work. Other groups have also successfully produced surrogate light curves with modest error. Previously, surrogate light curves have been produced by interpolating the coefficients $c_g(\Lambda)$ of a basis-function expansion $\log L_\alpha(t|\Lambda) = \sum_g c_g(\Lambda)\phi_g(t)$ [22,25], with the appropriate basis functions identified by principal component analysis of the raw simulation output. Because of the prohibitive cost of GP on large data sets, this analysis had to decimate input data to enable interpolation. Our reference-time method offers several notable advantages. The most important advantage is that our method is embarrassingly parallel (interpolations at every reference time can be performed independently, without need to select suitable basis functions in advance) and completely decoupled between time samples. Our interpolation is also inherently local in time, so artifacts inherited from late-time simulation data of low-photon-count light curves cannot contaminate our estimates of early-time behavior. Finally, our method can in principle be applied to all available data, without decimation, particularly when we employ other interpolation techniques.

Second, we adopt simulations with an imperfect model of the relevant opacities and nuclear physics. For example, we have adopted a conventional nuclear mass and decay model to predict nuclear heating and element abundances [32,90]. We also find that although our detailed multifrequency opacities yield a more realistic representation of the physics in the system, the assumption of thermalization breaks down much sooner than anticipated. Uncertainties in nuclear physics can play a substantial role in kilonova light curves [7,8]. Given sufficient simulations, surrogate light curves can be constructed for a wide range of nuclear inputs. We defer a systematic treatment of nuclear physics uncertainties and non-local thermodynamic equilibrium opacities to later work.




FIG. 12. Comparison to GW170817: Simulations only: this figure shows the results of direct comparison between all our simulations and our observations. Each simulation and angle is assigned a likelihood; the top 50 highest likelihood simulations are shown (black), compared with the observational data (red). The 50 black curves are drawn from a total number of potential candidate angles and simulations of 54×448 . For comparison, the blue curves show our continuously interpolated model, evaluated at the same parameters as the underlying simulations. We emphasize the identified simulations are at the edges of our simulation domain, where interpolation error is most severe.

Third, we employ simulations with phenomenological initial conditions, that are not initialized with the appropriate orientation-dependent distribution of mass, velocity, and composition versus time. More suitable initial conditions could be provided by detailed disk simulations [91,92].

Fourth, we do not (and cannot) initialize our simulations with initial data that are set by physics of the merger. Unlike previous work [22,26], we have not adopted a relationship between our two-component ejecta parameters and the progenitor masses m_1, m_2 , motivated by substantial uncertainty in the nuclear equation of state and remnant lifetime [18,93–95]. Even the best-available fits have considerable systematic uncertainty [96]. Similarly, we have not adopted assumptions about the lifetime of any hypermassive remnant and the duration of neutrino illumination [97], nor have we incorporated radiation from any associated jet [51]. Instead, given substantial systematic uncertainty in merger simulations (relative to the small amount of ejecta), we treat the ejecta purely phenomenologically, implicitly allowing for many potential nuisance parameters to characterize the outflow.

Fourth, our models were trained on a subset of simulations with fixed ejecta morphologies and mass fractions for both components. The predictive capability of our interpolations is restricted to two-component models represented by the

parameters in Table I. Our parameter estimation inherits these limitations and should thus also be considered for bias which stems from the selected subset of models in our interpolation training library.

Finally, we note that GW170817 was likely an exceptional case which contributed sufficient quantities of observational data for extremely informative parameter inference. Our methodology is still applicable in cases with sparser light-curve data; however, the level of detail in the inference results will vary based on the aforementioned sparsity.

VI. CONCLUSIONS

We have adaptively constructed detailed anisotropic models for kilonovae that cover a four-dimensional space describing two components' masses and velocities. From these models, we have constructed surrogate multiband light curves which can be evaluated continuously over this space. We have demonstrated how our model can be used for kilonova source parameter inference, including the kilonova associated with GW170817.

Although we limited our study to a specific set of assumptions, this analysis is an important stepping stone towards a better understanding of kilonova systematics. Recently,




FIG. 13. Comparison to GW170817: The left panel shows the results of interpreting observations of GW170817 using our surrogate light-curve model and adopting a strong angular prior $\theta \simeq 20^\circ$. In the left panel, the solid black shows inferences using all observing data, while the red curve omits K -band observations. The right panel shows inferred light curves corresponding to the full data set analysis (i.e., the light curves correspond to the black contours in the left panel).

several studies have demonstrated that several physical assumptions can notably impact the deduced light curve. However, these impacts could have effects that are partially degenerate with modest shifts in ejecta properties. To understand the practical impact of these uncertainties, in future work we will employ our parametrized models with these sources of error.

In this work we emphasized inference on only phenomenological kilonova parameters. Several studies have demonstrated the value in using multimessenger information to more tightly constrain parameters like source inclination (see, e.g., [13,20,21,87]), even without adopting strong assumptions about the relationship between ejecta and progenitor masses. With such assumptions, even stronger constraints have been widely explored. In future work we will show how the electromagnetic inference strategy applied here can be tightly and efficiently integrated with the RIFT parameter inference engine, enabling concordance inference about multimessenger sources.

All of our input data products, fitted light curves, and the code we used to produce them are available in Ref. [41]. The underlying full simulations are available in Ref. [98].

ACKNOWLEDGMENTS

R.O.S. and M.R. acknowledge support from NSF Grant No. AST 1909534. E.A.C. acknowledges financial support from the IDEAS Fellowship, a research traineeship program funded by the National Science Foundation under Grant No. DGE-1450006. C.L.F., C.J.F., A.H., O.K., and R.W. were supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for

the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). Research presented in this paper was supported by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Project No. 20190021DR. This research used resources provided by the Los Alamos National Laboratory Institutional Computing Program, which is supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. 89233218CNA000001. This document has been assigned LA-UR-21-24289.

APPENDIX: VALIDATION OF PARAMETER INFERENCE METHOD

1. Simple analytic kilonova model

To validate our parameter inference codes, we implemented a standard semianalytic kilonova model previously presented in [23]. This model consists of single-component, two-component, and three-component models, combined by flux addition and not allowing for anisotropy. In the following equations, M is the r -process ejecta mass (in $M_\odot$) and v is the ejecta velocity. Note that for now we assume the ejecta consists entirely of r -process material, so M is the full ejecta mass. The radioactive heating rate at time t is given by [90]

$$L_{\text{in}}(t) = 4 \times 10^{18} M \times \left[0.5 - \pi^{-1} \arctan \left(\frac{t - t_0}{\sigma} \right) \right]^{1.3} \times \text{erg s}^{-1}, \quad (\text{A1})$$

where $t_0 = 1.3$ s and $\sigma = 0.11$ s are constants.

Only a fraction of L_{in} powers the kilonova, given by the thermalization efficiency ϵ_{th} . This is approximated analyti-




FIG. 14. Analytic kilonova model injection. Top panel: posterior distribution for the injection/recovery test (the blue lines show the true parameter values). For this example, T_c was fixed for each component and the distance was fixed to 40.0 Mpc. Bottom panel: fake photometry data and light-curve models for the injection/recovery test. The solid line shows the light-curve model evaluated at the maximum-likelihood parameters. Both the data and light curves include a confidence interval.

cally in [45]

$$\epsilon_{\text{th}}(t) = 0.36 \left[e^{-at} + \frac{\ln(1 + 2bt^d)}{2bt^d} \right]. \quad (\text{A2})$$

The parameters a , b , and d are constants that depend on the ejecta mass and velocity; an interpolation of Table 1 in [45] is used in the model. The bolometric luminosity is calculated via [73]

$$L_{\text{bol}}(t) = \frac{2}{t_d} \exp\left(\frac{-t^2}{t_d^2}\right) \int_0^t L_{\text{in}} \epsilon_{\text{th}} \exp\left(\frac{t'^2}{t_d^2}\right) \frac{t'}{t_d} dt', \quad (\text{A3})$$

where t_d is the diffusion timescale, $t_d = \sqrt{2\kappa M/\beta v c}$, κ is the opacity, and $\beta = 13.7$ is a dimensionless constant related to the ejecta's geometry.




FIG. 15. Probability-probability (P-P) tests: P-P plot for 100 synthetic injections generated with one random component.

Light curves are calculated by assuming the kilonova behaves as a blackbody photosphere that expands at a velocity v . The blackbody temperature is generally defined by its bolometric luminosity; however, once it cools to a critical temperature T_c , the photosphere recedes into the ejecta and the temperature remains fixed. The photosphere temperature is

$$T_{\text{phot}}(t) = \max \left[\left(\frac{L_{\text{bol}}(t)}{4\pi\sigma_{\text{SB}}v^2t^2} \right)^{1/4}, T_c \right], \quad (\text{A4})$$

where σ_{SB} is the Stefan-Boltzmann constant. When $T_{\text{phot}} > T_c$, the photosphere radius is simply $R_{\text{phot}} = vt$. When $T_{\text{phot}} = T_c$ (i.e., the photosphere has receded into the ejecta), the photosphere radius is

$$R_{\text{phot}}(t) = \left(\frac{L_{\text{bol}}(t)}{4\pi\sigma_{\text{SB}}T_c^4} \right)^{1/2}. \quad (\text{A5})$$

The flux density at frequency ν is given in [5]

$$F_{\nu}(t) = \frac{2\pi h\nu^3}{c^2} \frac{1}{\exp[h\nu/kT_{\text{photo}}(t)] - 1} \frac{R_{\text{photo}}^2(t)}{D^2}, \quad (\text{A6})$$

where D is the source distance. We use a fixed fiducial distance of $D = 10$ pc to calculate $F_{\nu}(t)$, then calculate AB magnitude with a distance modulus if necessary.

To compute multicomponent light curves we assume each component has a photosphere that evolves independently of the others. The total flux density is the sum of the flux densities of the individual components. The version of the model implemented in the code uses three components with fixed opacities (a blue component with $\kappa = 0.5 \text{ cm}^2 \text{ g}^{-1}$, a purple

component with $\kappa = 3 \text{ cm}^2 \text{ g}^{-1}$, and a red component with $\kappa = 10 \text{ cm}^2 \text{ g}^{-1}$.

2. Validation test: Recovery of synthetic kilonova

Figure 14 illustrates our implementation of this synthetic model and our parameter inference code. The error bars on the bottom panel show a densely sampled synthetic multiband light curve with plausible kilonova parameters for all three components. The top panel shows a standard corner plot representing the parameters of our synthetic kilonova with thin blue lines; one-dimensional marginal posterior distributions on the diagonal for each parameter; and two-dimensional marginal posterior distributions in the bottom corner.

3. Validation test: Random synthetic kilonova and probability-probability plot

We also demonstrated our inference technique using 100 randomly generated light curves, drawn uniformly from the same priors we use for inference and incorporating noise consistent with our Gaussian noise model. Using these 100 synthetic events and inferences, we can construct a probability-probability (P-P) plot [99], which corroborates that the one-dimensional marginal distributions are consistent. Figure 15 shows the results of our analysis. More extensive tests of the underlying integration algorithm, including P-P plots using more complex and higher-dimensional models, are reported elsewhere [76].

-
- [1] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.*, Multi-messenger observations of a binary neutron star merger, *Astrophys. J. Lett.* **848**, L12 (2017).
 - [2] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.* **119**, 161101 (2017).
 - [3] J. M. Lattimer and D. N. Schramm, Black-hole-neutron-star collisions, *Astrophys. J. Lett.* **192**, L145 (1974).
 - [4] L.-X. Li and B. Paczyński, Transient events from neutron star mergers, *Astrophys. J. Lett.* **507**, L59 (1998).
 - [5] B. D. Metzger, Kilonovae, *Living Rev. Relativity* **23**, 1 (2019).
 - [6] S. Bernuzzi, Neutron star merger remnants, *Gen. Relativ. Gravit.* **52**, 108 (2020).
 - [7] J. Barnes, Y. L. Zhu, K. A. Lund, T. M. Sprouse, N. Vassh, G. C. McLaughlin, M. R. Mumpower, and R. Surman, Kilonovae across the nuclear physics landscape: The impact of nuclear physics uncertainties on r-process-powered emission, *Astrophys. J.* **918**, 44 (2021).
 - [8] Y. L. Zhu, K. Lund, J. Barnes, T. M. Sprouse, N. Vassh, G. C. McLaughlin, M. R. Mumpower, and R. Surman, Modeling kilonova light curves: Dependence on nuclear inputs, *Astrophys. J.* **906**, 94 (2021).
 - [9] N. Vassh, G. C. McLaughlin, M. R. Mumpower, and R. Surman, Markov chain monte carlo predictions of neutron-rich lanthanide properties as a probe of r-process dynamics, *Astrophys. J.* **907**, 98 (2021).
 - [10] C. J. Horowitz, Neutron rich matter in the laboratory and in the heavens after GW170817, *Ann. Phys.* **411**, 167992 (2019).
 - [11] K. Chatziioannou, Neutron-star tidal deformability and equation-of-state constraints, *Gen. Relativ. Gravit.* **52**, 109 (2020).
 - [12] H.-Y. Chen, P. S. Cowperthwaite, B. D. Metzger, and E. Berger, A program for multimessenger standard siren cosmology in the era of LIGO a, rubin observatory, and beyond, *Astrophys. J.* **908**, L4 (2021).
 - [13] M. W. Coughlin, S. Antier, T. Dietrich, R. J. Foley, J. Heinzel, M. Bulla, N. Christensen, D. A. Coulter, L. Issa, and N. Khetan, Measuring the Hubble constant with a sample of kilonovae, *Nat. Commun.* **11**, 4129 (2020).
 - [14] M. W. Coughlin, T. Dietrich, J. Heinzel, N. Khetan, S. Antier, M. Bulla, N. Christensen, D. A. Coulter, and R. J. Foley, Standardizing kilonovae and their use as standard candles to measure the Hubble constant, *Phys. Rev. Research* **2**, 022006(R) (2020).
 - [15] Z. Doctor, Thunder and lightning: Using neutron-star mergers as simultaneous standard candles and sirens to measure cosmological parameters, *Astrophys. J. Lett.* **892**, L16 (2020).
 - [16] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, A standard siren measurement of the Hubble constant, *Nature (London)* **551**, 85 (2017).
 - [17] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.*, GW170817: Measurements of Neutron Star Radii and Equation of State, *Phys. Rev. Lett.* **121**, 161101 (2018).
 - [18] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott *et al.*, Model comparison from LIGO-Virgo data on GW170817's binary components and consequences for the merger remnant, *Class. Quantum Grav.* **37**, 045006 (2020).
 - [19] A. Bauswein, S. Blacker, V. Vijayan, N. Stergioulas, K. Chatziioannou, J. A. Clark, N.-U. F. Bastian, D. B. Blaschke, M. Cierniak, and T. Fischer, Equation of State Constraints from the Threshold Binary Mass for Prompt Collapse of Neutron Star Mergers, *Phys. Rev. Lett.* **125**, 141103 (2020).
 - [20] K. Hotokezaka, E. Nakar, O. Gottlieb, S. Nissanke, K. Masuda, G. Hallinan, K. P. Mooley, and A. T. Deller, A Hubble constant measurement from superluminal motion of the jet in GW170817, *Nat. Astron.* **3**, 940 (2019).
 - [21] T. Dietrich, M. W. Coughlin, P. T. H. Pang, M. Bulla, J. Heinzel, L. Issa, I. Tews, and S. Antier, Multimessenger constraints on the neutron-star equation of state and the Hubble constant, *Science* **370**, 1450 (2020).
 - [22] M. W. Coughlin, T. Dietrich, Z. Doctor, D. Kasen, S. Coughlin, A. Jerkstrand, G. Leoudas, O. McBrien, B. D. Metzger, R. O'Shaughnessy *et al.*, Constraints on the neutron star equation of state from AT2017gfo using radiative transfer simulations, *Mon. Not. R. Astron. Soc.* **480**, 3871 (2018).

- [23] V. A. Villar, J. Guillochon, E. Berger, B. D. Metzger, P. S. Cowperthwaite, M. Nicholl, K. D. Alexander, P. K. Blanchard, R. Chornock, T. Eftekhari *et al.*, The combined ultraviolet, optical, and near-infrared light curves of the kilonova associated with the binary neutron star merger GW170817: Unified data set, analytic models, and physical implications, *Astrophys. J. Lett.* **851**, L21 (2017).
- [24] M. Breschi, A. Perego, S. Bernuzzi, W. Del Pozzo, V. Nedora, D. Radice, and D. Vescovi, At2017gfo: Bayesian inference and model selection of multicomponent kilonovae and constraints on the neutron star equation of state, *Mon. Not. R. Astron. Soc.* **505**, 1661 (2021).
- [25] J. Heinzel, M. W. Coughlin, T. Dietrich, M. Bulla, S. Antier, N. Christensen, D. A. Coulter, R. J. Foley, L. Issa, and N. Khean, Comparing inclination-dependent analyses of kilonova transients, *Mon. Not. R. Astron. Soc.* **502**, 3057 (2021).
- [26] M. W. Coughlin, T. Dietrich, B. Margalit, and B. D. Metzger, Multimessenger Bayesian parameter inference of a binary neutron star merger, *Mon. Not. R. Astron. Soc.* **489**, L91 (2019).
- [27] S. J. Smartt, T. W. Chen, A. Jerkstrand, M. Coughlin, E. Kankare, S. A. Sim, M. Fraser, C. Inserra, K. Maguire, K. C. Chambers *et al.*, A kilonova as the electromagnetic counterpart to a gravitational-wave source, *Nature (London)* **551**, 75 (2017).
- [28] N. R. Tanvir, A. J. Levan, C. González-Fernández, O. Korobkin, I. Mandel, S. Rosswog, J. Hjorth, P. D'Avanzo, A. S. Fruchter, C. L. Fryer *et al.*, The emergence of a lanthanide-rich kilonova following the merger of two neutron stars, *Astrophys. J. Lett.* **848**, L27 (2017).
- [29] K. D. Alexander, E. Berger, W. Fong, P. K. G. Williams, C. Guidorzi, R. Margutti, B. D. Metzger, J. Annis, P. K. Blanchard, D. Brout *et al.*, The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo GW170817. VI. Radio constraints on a relativistic jet and predictions for late-time emission from the kilonova ejecta, *Astrophys. J. Lett.* **848**, L21 (2017).
- [30] R. Chornock, E. Berger, D. Kasen, P. S. Cowperthwaite, M. Nicholl, V. A. Villar, K. D. Alexander, P. K. Blanchard, T. Eftekhari, W. Fong *et al.*, The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo GW170817. IV. Detection of near-infrared signatures of r-process nucleosynthesis with gemini-south, *Astrophys. J. Lett.* **848**, L19 (2017).
- [31] P. S. Cowperthwaite, E. Berger, V. A. Villar, B. D. Metzger, M. Nicholl, R. Chornock, P. K. Blanchard, W. Fong, R. Margutti, M. Soares-Santos *et al.*, The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. II. UV, Optical, and Near-infrared Light Curves and Comparison to Kilonova Models, *Astrophys. J. Lett.* **848**, L17 (2017).
- [32] R. T. Wollaeger, O. Korobkin, C. J. Fontes, S. K. Rosswog, W. P. Even, C. L. Fryer, J. Sollerman, A. L. Hungerford, D. R. van Rossum, and A. B. Wollaber, Impact of ejecta morphology and composition on the electromagnetic signatures of neutron star mergers, *Mon. Not. R. Astron. Soc.* **478**, 3298 (2018).
- [33] W. Even, O. Korobkin, C. L. Fryer, C. J. Fontes, R. T. Wollaeger, A. Hungerford, J. Lippuner, J. Miller, M. R. Mumpower, and G. W. Misch, Composition effects on kilonova spectra and light curves. I, *Astrophys. J.* **899**, 24 (2020).
- [34] R. T. Wollaeger, C. L. Fryer, E. A. Chase, C. J. Fontes, M. Ristic, A. L. Hungerford, O. Korobkin, R. O'Shaughnessy, and A. M. Herring, A broad grid of 2d kilonova emission models, *Astrophys. J.* **918**, 10 (2021).
- [35] M. Ristic, Smarter simulation placement of kilonova light curve models for computationally inexpensive surrogate model creation, Master's thesis, Rochester Institute of Technology, 2020.
- [36] D. Kasen, B. Metzger, J. Barnes, E. Quataert, and E. Ramirez-Ruiz, Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event, *Nature (London)* **551**, 80 (2017).
- [37] O. Korobkin, R. Wollaeger, C. Fryer, A. L. Hungerford, S. Rosswog, C. Fontes, M. Mumpower, E. Chase, W. Even, J. M. Miller *et al.*, Axisymmetric radiative transfer models of kilonovae, *Astrophys. J.* **910**, 116 (2021).
- [38] A. Perego, D. Radice, and S. Bernuzzi, AT 2017gfo: An Anisotropic and Three-component Kilonova Counterpart of GW170817, *Astrophys. J. Lett.* **850**, L37 (2017).
- [39] B. O'Connor, P. Beniamini, and C. Kouveliotou, Constraints on the circumburst environments of short gamma-ray bursts, *Mon. Not. R. Astron. Soc.* **495**, 4782 (2020).
- [40] B. O'Connor, E. Troja, S. Dichiara, E. A. Chase, G. Ryan, S. B. Cenko, C. L. Fryer, R. Ricci, F. Marshall, C. Kouveliotou *et al.*, A tale of two mergers: constraints on kilonova detection in two short GRBs at $z \sim 0.5$, *Mon. Not. R. Astron. Soc.* **502**, 1279 (2021).
- [41] R. T. Wollaeger and D. R. van Rossum, Radiation transport for explosive outflows: Opacity regrouping, *Astrophys. J. S.* **214**, 28 (2014).
- [42] C. J. Fontes, H. L. Zhang, J. Abdallah, J., R. E. H. Clark, D. P. Kilcrease, J. Colgan, R. T. Cunningham, P. Hakel, N. H. Magee, and M. E. Sherrill, The los alamos suite of relativistic atomic physics codes, *J. Phys. B: At. Mol. Phys.* **48**, 144014 (2015).
- [43] C. J. Fontes, C. L. Fryer, A. L. Hungerford, R. T. Wollaeger, and O. Korobkin, A line-binned treatment of opacities for the spectra and light curves from neutron star mergers, *Mon. Not. R. Astron. Soc.* **493**, 4143 (2020).
- [44] C. Winteler, R. Käppeli, A. Perego, A. Arcones, N. Vasset, N. Nishimura, M. Liebendörfer, and F. K. Thielemann, Magnetorotationally driven supernovae as the origin of early galaxy r-process elements? *Astrophys. J. Lett.* **750**, L22 (2012).
- [45] J. Barnes, D. Kasen, M.-R. Wu, and G. Martínez-Pinedo, Radioactivity and thermalization in the ejecta of compact object mergers and their impact on kilonova light curves, *Astrophys. J.* **829**, 110 (2016).
- [46] D. Radice, S. Bernuzzi, and A. Perego, The dynamics of binary neutron star mergers and GW170817, *Annu. Rev. Nucl. Part. Sci.* **70**, 95 (2020).
- [47] M. S. Bessell, F. Castelli, and B. Plez, Model atmospheres broad-band colors, bolometric corrections and temperature calibrations for O - M stars, *Astron. Astrophys.* **333**, 231 (1998).
- [48] M. R. Blanton and S. Roweis, K-corrections and filter transformations in the ultraviolet, optical, and near-infrared, *Astron. J.* **133**, 734 (2007).
- [49] P. C. Hewett, S. J. Warren, S. K. Leggett, and S. T. Hodgkin, The UKIRT infrared deep sky survey ZY JHK photometric system: passbands and synthetic colours, *Mon. Not. R. Astron. Soc.* **367**, 454 (2006).
- [50] D. Kasen, R. Fernández, and B. D. Metzger, Kilonova light curves from the disc wind outflows of compact object mergers, *Mon. Not. R. Astron. Soc.* **450**, 1777 (2015).
- [51] L. Nativi, M. Bulla, S. Rosswog, C. Lundman, G. Kowal, D. Gizzi, G. P. Lamb, and A. Perego, Can jets make the

- radioactively powered emission from neutron star mergers bluer? *Mon. Not. R. Astron. Soc.* **500**, 1772 (2021).
- [52] H. Klion, P. C. Duffell, D. Kasen, and E. Quataert, The effect of jet-ejecta interaction on the viewing angle dependence of kilonova light curves, *Mon. Not. R. Astron. Soc.* **502**, 865 (2021).
- [53] L. Piro, E. Troja, B. Zhang, G. Ryan, H. van Eerten, R. Ricci, M. H. Wieringa, A. Tiengo, N. R. Butler, S. B. Cenko *et al.*, A long-lived neutron star merger remnant in GW170817: constraints and clues from X-ray observations, *Mon. Not. R. Astron. Soc.* **483**, 1912 (2018).
- [54] S. Ai, H. Gao, Z.-G. Dai, X.-F. Wu, A. Li, B. Zhang, and M.-Z. Li, The allowed parameter space of a long-lived neutron star as the merger remnant of GW170817, *Astrophys. J.* **860**, 57 (2018).
- [55] R. T. Wollaeger, C. L. Fryer, C. J. Fontes, J. Lippuner, W. T. Vestrand, M. R. Mumpower, O. Korobkin, A. L. Hungerford, and W. P. Even, Impact of pulsar and fallback sources on multi-frequency kilonova models, *Astrophys. J.* **880**, 22 (2019).
- [56] K. Kawaguchi, S. Fujibayashi, M. Shibata, M. Tanaka, and S. Wanajo, A low-mass binary neutron star: Long-term ejecta evolution and kilonovae with weak blue emission, *Astrophys. J.* **913**, 100 (2021).
- [57] S. Darbha and D. Kasen, Inclination dependence of kilonova light curves from globally aspherical geometries, *Astrophys. J.* **897**, 150 (2020).
- [58] K. Kawaguchi, M. Shibata, and M. Tanaka, Diversity of kilonova light curves, *Astrophys. J.* **889**, 171 (2020).
- [59] L. Breiman, Random forests, *Mach. Learn.* **45**, 5 (2001).
- [60] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [61] K. P. Murphy, *Machine Learning: A Probabilistic Perspective* (The MIT Press, Cambridge, MA, 2012).
- [62] H. Solomon and S. Zacks, Optimal design of sampling from finite populations: A critical review and indication of new research areas, *J. Am. Stat. Assoc.* **65**, 653 (1970).
- [63] A. Krause and C. Guestrin, Nonmyopic active learning of gaussian processes: an exploration/exploitation approach, in *International Conference on Machine Learning (ICML)* (Corvallis, Oregon, 2007).
- [64] D. Cohn, Z. Ghahramani, and M. Jordan, Active learning with statistical models, *J. Artif. Intell. Res.* **4**, 129 (1996).
- [65] Y. Gal, R. Islam, and Z. Ghahramani, Deep Bayesian active learning with image data, *arXiv:1703.02910* (2017).
- [66] D. J. MacKay, *Bayesian Methods for Adaptive Models* (California Institute of Technology, Pasadena, 1992).
- [67] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, Information-theoretic regret bounds for gaussian process optimization in the bandit setting, *IEEE Trans. Infor. Theor.* **58**, 3250 (2012).
- [68] J. Mockus, V. Tiesis, and A. Zilinskas, in *The Application of Bayesian Methods for Seeking the Extremum* (Springer, Berlin, 1978), pp. 117–128.
- [69] J. Wu and P. I. Frazier, in *Neural Information Processing Systems* (NeurIPS, 2016).
- [70] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, Fast Prediction and Evaluation of Gravitational Waveforms Using Surrogate Models, *Phys. Rev. X* **4**, 031006 (2014).
- [71] W. D. Arnett, Analytic solutions for light curves of supernovae of Type II, *Astrophys. J.* **237**, 541 (1980).
- [72] W. D. Arnett, Type I supernovae. I - Analytic solutions for the early part of the light curve, *Astrophys. J.* **253**, 785 (1982).
- [73] E. Chatzopoulos, J. C. Wheeler, and J. Vinko, Generalized semi-analytical models of supernova light curves, *Astrophys. J.* **746**, 121 (2012).
- [74] C. Pankow, P. Brady, E. Ochsner, and R. O'Shaughnessy, Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences, *Phys. Rev. D* **92**, 023002 (2015).
- [75] J. Lange, R. O'Shaughnessy, and M. Rizzo, Rapid and accurate parameter inference for coalescing, precessing compact binaries, *arXiv:1805.10457*.
- [76] E. Champion, V. Delfavero, H. Gallagher, J. Lange, O'Shaughnessy R., J. Wofford, D. Wysocki, and A. Yelikar (unpublished).
- [77] I. Arcavi, G. Hosseinzadeh, D. A. Howell, C. McCully, D. Poznanski, D. Kasen, J. Barnes, M. Zaltzman, S. Vasylyev, D. Maoz *et al.*, Optical emission from a kilonova following a gravitational-wave-detected neutronstar merger, *Nature (London)* **551**, 64 (2017).
- [78] E. Troja, L. Piro, H. van Eerten, R. T. Wollaeger, M. Im, O. D. Fox, N. R. Butler, S. B. Cenko, T. Sakamoto, C. L. Fryer *et al.*, The X-ray counterpart to the gravitational-wave event GW170817, *Nature (London)* **551**, 71 (2017).
- [79] M. R. Drout, A. L. Piro, B. J. Shappee, C. D. Kilpatrick, J. D. Simon, C. Contreras, D. A. Coulter, R. J. Foley, M. R. Siebert, N. Morrell *et al.*, Light curves of the neutron star merger GW170817/SSS17a: Implications for r-process nucleosynthesis, *Science* **358**, 1570 (2017).
- [80] S. Valenti, D. J. Sand, S. Yang, E. Cappellaro, L. Tartaglia, A. Corsi, S. W. Jha, D. E. Reichart, J. Haislip, and V. Kouprianov, The discovery of the electromagnetic counterpart of GW170817: Kilonova at 2017gfo/DLT17ck, *Astrophys. J. Lett.* **848**, L24 (2017).
- [81] E. Pian, P. D'Avanzo, S. Benetti, M. Branchesi, E. Brocato, S. Campana, E. Cappellaro, S. Covino, V. D'Elia, J. P. U. Fynbo *et al.*, Spectroscopic identification of r-process nucleosynthesis in a double neutron-starmerger, *Nature (London)* **551**, 67 (2017).
- [82] M. M. Kasliwal, E. Nakar, L. P. Singer, D. L. Kaplan, D. O. Cook, A. Van Sistine, R. M. Lau, C. Fremling, O. Gottlieb, J. E. Jencson *et al.*, Illuminating gravitational waves: A concordant picture of photons from a neutron star merger, *Science* **358**, 1559 (2017).
- [83] B. J. Shappee, J. D. Simon, M. R. Drout, A. L. Piro, N. Morrell, J. L. Prieto, D. Kasen, T. W. S. Holoi, J. A. Kollmeier, D. D. Kelson *et al.*, Early spectra of the gravitational wave source GW170817: Evolution of a neutron star merger, *Science* **358**, 1574 (2017).
- [84] Y. Utsumi, M. Tanaka, N. Tominaga, M. Yoshida, S. Barway, T. Nagayama, T. Zenko, K. Aoki, T. Fujiyoshi, H. Furusawa *et al.*, J-GEM observations of an electromagnetic counterpart to the neutron star merger GW170817, *PASJ* **69**, 101 (2017).
- [85] D. M. Siegel, GW170817 -the first observed neutron star merger and its kilonova: Implications for the astrophysical site of the r-process, *Eur. Phys. J. A* **55**, 203 (2019).
- [86] M. Nicholl, B. Margalit, P. Schmidt, G. P. Smith, E. J. Ridley, and J. Nuttall, Tight multimessenger constraints on the neutron star equation of state from GW170817 and a forward model for

- kilonova light-curve synthesis, *Mon. Not. R. Astron. Soc.* **505**, 3016 (2021).
- [87] K. P. Mooley, A. T. Deller, O. Gottlieb, E. Nakar, G. Hallinan, S. Bourke, D. A. Frail, A. Horesh, A. Corsi, and K. Hotokezaka, Superluminal motion of a relativistic jet in the neutron-star merger GW170817, *Nature (London)* **561**, 355 (2018).
- [88] P. A. Evans, S. B. Cenko, J. A. Kennea, S. W. K. Emery, N. P. M. Kuin, O. Korobkin, R. T. Wollaeger, C. L. Fryer, K. K. Madsen, F. A. Harrison *et al.*, Swift and NuSTAR observations of GW170817: Detection of a blue kilonova, *Science* **358**, 1565 (2017).
- [89] E. Troja, H. van Eerten, B. Zhang, G. Ryan, L. Piro, R. Ricci, B. O'Connor, M. H. Wieringa, S. B. Cenko, and T. Sakamoto, A thousand days after the merger: Continued X-ray emission from GW170817, *Mon. Not. R. Astron. Soc.* **498**, 5643 (2020).
- [90] O. Korobkin, S. Rosswog, A. Arcones, and C. Winteler, On the astrophysical robustness of the neutron star merger r-process, *Mon. Not. R. Astron. Soc.* **426**, 1940 (2012).
- [91] J. M. Miller, B. R. Ryan, J. C. Dolence, A. Burrows, C. J. Fontes, C. L. Fryer, O. Korobkin, J. Lippuner, M. R. Mumpower, and R. T. Wollaeger, Full transport model of GW170817-like disk produces a blue kilonova, *Phys. Rev. D* **100**, 023008 (2019).
- [92] S. De and D. Siegel, Igniting weak interactions in neutron star postmerger accretion disks, *Astrophys. J.* **921**, 94 (2021).
- [93] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, and M. Shibata, Dynamical mass ejection from binary neutron star mergers: Radiation-hydrodynamics study in general relativity, *Phys. Rev. D* **91**, 064059 (2015).
- [94] T. Dietrich, S. Bernuzzi, and W. Tichy, Closed-form tidal approximants for binary neutron star gravitational waveforms constructed from high-resolution numerical relativity simulations, *Phys. Rev. D* **96**, 121501(R) (2017).
- [95] M. Lucca and L. Sagunski, The lifetime of binary neutron star merger remnants, *J. High Energy Astrophys.* **27**, 33 (2020).
- [96] C. J. Krüger and F. Foucart, Estimates for disk and ejecta masses produced in compact binary mergers, *Phys. Rev. D* **101**, 103002 (2020).
- [97] V. Decoeene, C. Guépin, K. Fang, K. Kotera, and B. D. Metzger, High-energy neutrinos from fallback accretion of binary neutron star merger remnants, *JCAP* **04** (2020) 045.
- [98] https://ccsweb.lanl.gov/astro/transient/transients_astro.html.
- [99] S. Cook, A. Gelman, and D. Rubin, Validation of Software for Bayesian Models Using Posterior Quantiles, *J. Comput. Graph. Stat.* **15**, 675 (2012).