Royal Astronomical Society

Downloaded from https://academic.oup.com/mnras/article/516/1/1137/6675824 by University of Michigan user on 01 September 2022

MNRAS **516,** 1137–1148 (2022) Advance Access publication 2022 August 25

KilonovaNet: Surrogate models of kilonova spectra with conditional variational autoencoders

K. Lukošiute [®], ¹★ G. Raaijmakers [®], ¹ Z. Doctor [®], ² M. Soares-Santos ^{®3} and B. Nord ^{®4,5,6}

Accepted 2022 August 10. Received 2022 August 3; in original form 2022 April 1

ABSTRACT

Detailed radiative transfer simulations of kilonova spectra play an essential role in multimessenger astrophysics. Using the simulation results in parameter inference studies requires building a surrogate model from the simulation outputs to use in algorithms requiring sampling. In this work, we present kilonovanet, an implementation of conditional variational autoencoders (cVAEs) for the construction of surrogate models of kilonova spectra. This method can be trained on spectra directly, removing overhead time of pre-processing spectra, and greatly speeds up parameter inference time. We build surrogate models of three state-of-the-art kilonova simulation data sets and present in-depth surrogate error evaluation methods, which can in general be applied to any surrogate construction method. By creating synthetic photometric observations from the spectral surrogate, we perform parameter inference for the observed light-curve data of GW170817 and compare the results with previous analyses. Given the speed with which kilonovanet performs during parameter inference, it will serve as a useful tool in future gravitational wave observing runs to quickly analyse potential kilonova candidates.

Key words: gravitational waves – methods: statistical.

1 INTRODUCTION

Mergers of compact stellar remnants, like neutron stars (NSs) and black holes (BHs), provide important testbeds for astrophysical processes, cosmological evolution, and matter under extreme conditions. In binary merger systems, as companions orbit each other, their physical separation decays due to energy loss through the emission of gravitational waves (GWs), which can be detected by modern ground-based interferometers, like those of the Laser Interferometer Gravitational-Wave Observatory (LIGO; Aasi et al. 2015), Virgo (Acernese et al. 2015), and Kamioka Gravitational Wave Detector (KAGRA; Akutsu et al. 2019) experiments. If the merging system contains one or two neutron stars, neutron star material can be ejected at high velocities and emit electromagnetic (EM) waves (Lattimer & Schramm 1974; Li & Paczyński 1998; Metzger & Berger 2012; Pian 2021). Using information jointly from GWs and EM emission is essential for maximizing our understanding of multiple phenomena: since both GW and EM measurements are sensitive to parameters of a merging binary, combining them can yield stronger constraints on properties of the binary (Coughlin et al. 2017), the neutron-star equation of state (Coughlin et al. 2018; Radice et al. 2018a), and the expansion rate of the Universe (Schutz 1986; Holz & Hughes 2005; Abbott et al. 2017b).

Amidst the neutron-rich material ejected during an NS-containing merger (Lattimer & Schramm 1974; Rosswog et al. 1998; Metzger et al. 2010), r-process nucleosynthesis generates heavy nuclei (Burbidge et al. 1957; Cameron 1957). The subsequent decay of these nuclei drives a kilonova – an ultraviolet-optical-infrared transient, whose brightness peaks two to three days after a merger (Li & Paczyński 1998). Most binary neutron star (BNS) systems and some black hole–neutron star (BHNS) systems are expected to produce a kilonova (Metzger 2019).

There are two primary channels for mass-ejection during an NS-containing merger (Metzger 2019) – dynamical ejection (during merger) and disc wind ejection (after merger). Dynamical ejecta are further subdivided into two categories – 'tidal' and 'polar.' Tidal dynamical ejecta outflows occur during the final stage of the inspiral in both BHNS and BNS systems when the neutron star is disrupted by the gravitational field of the companion. Polar dynamical ejecta result from shock heating caused by the direct collision of the neutron stars in BNS systems only (e.g. Radice et al. 2018b).

The detection of GW170817, the first observed BNS GW event, marks the onset of the era of multimessenger and GW astronomy (Abbott et al. 2017a). GW170817 was accompanied by counterparts across the EM spectrum (e.g. Abbott et al. 2017c, d; Arcavi et al. 2017; Coulter et al. 2017; Drout et al. 2017; Savchenko et al. 2017; Shappee et al. 2017; Smartt et al. 2017; Soares-Santos et al. 2017; Valenti et al. 2017; Goldstein et al. 2019). Then during the third LIGO and Virgo observing run (Run O3), tens of binary black hole mergers and multiple neutron star-containing mergers were

¹GRAPPA, University of Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands

²Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, 1800 Sherman Ave, Evanston, IL 60201, USA

³Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

⁴Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

⁵Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510, USA

⁶Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

^{*} E-mail: lukosiutekamile@gmail.com

detected, including BHNS events (Abbott et al. 2021a, b, c). Through multiple dedicated follow-up electromagnetic observing campaigns in O3 (Coughlin et al. 2019; Goldstein et al. 2019; Antier et al. 2020; Gompertz et al. 2020; Page et al. 2020; Anand et al. 2021), one candidate was identified (Graham et al. 2020), but zero were confirmed.

Accurate and fast models of kilonovae are critical for multiple tasks in multimessenger and GW astrophysics. First, we need fast models in order to quickly evaluate candidates during electromagnetic follow up campaigns (Soares-Santos et al. 2017). Identifying and observing a candidate in the first few hours is crucial, especially to understand the origin of early kilonova emission, such as the fast, blue emission seen in GW170817 (Arcavi 2018; Chase et al. 2022). After the data are all collected, analysis, inference can begin; we can use the data to infer parameters of the NS equation of state (e.g. Coughlin et al. 2018; Radice et al. 2018b; Raaijmakers et al. 2021b), the Hubble constant (e.g. Abbott et al. 2017b; Coughlin et al. 2020; Dietrich et al. 2020), and studying r-process nucleosynthesis (e.g. Drout et al. 2017), and the site of rapid neutron capture nucleosynthesis (e.g. Kasen et al. 2017). For these scientific studies, the most accurate models will lead to the least biased results.

Numerical hydrodynamic simulations provide the most detailed models of matter (i.e. velocity, mass, composition, and opacity) ejected during mergers (e.g. Dietrich & Ujevic 2017; Perego et al. 2017; Tanaka et al. 2020; Kawaguchi et al. 2021). The results of these simulations are then incorporated into radiative transfer calculations to compute observational properties of mergers (e.g. Kawaguchi et al. 2016; Kasen et al. 2017; Wollaeger et al. 2018; Bulla 2019; Korobkin et al. 2021; Wollaeger et al. 2021) to generate synthetic observations for comparison with real-sky observations (e.g. Raaijmakers et al. 2021a). Since the simulations are computationally complex, they take several hours to produce observables for one parameter set (Bulla 2019).

In order to use the outputs of the simulations in parameter inference studies, the outputs are trained to produce a 'surrogate' model that emulates the output of simulations but run considerably faster. For example, a commonly used surrogate-construction method is Gaussian process regression (GPR), a statistical method for interpolation; this method has been applied in a few GW and EM studies (Doctor et al. 2017; Coughlin et al. 2018; Dietrich et al. 2020; Ristic et al. 2022). One difficulty with using GPR for kilonova surrogate models for example is that it does not scale well with the number of training examples: the evaluation time goes as the square of the number of training points and the training time goes as the cube. State-of-the-art radiative transfer simulations, such as those published in Dietrich et al. (2020) and Anand et al. (2021), output spectra based on multiple input parameters in hundreds of frequency bins and with hundreds of time steps. The high-dimensionality of the inputs and outputs, combined with the grid over physical parameters means that the simulated sets become large quickly, making GPR a costly method. Recently, Almualla et al. (2021) presented a neural network based surrogate modelling framework which greatly increased inference time for the GW170817 light curve. We present a complementary approach to their method, instead training directly on spectra, thereby skipping the dimensional reduction step, and training one permanent network set, thereby allowing construction of photometric observations for any band from the surrogate as well as comparison to spectroscopic data.

kilonovanet is an algorithm for generating surrogate kilonova spectra on tens-of-millisecond time-scales via conditional variational autoencoders (cVAE) (Kingma & Welling 2014; Rezende, Mohamed & Wierstra 2014; Sohn, Yan & Lee 2015). We additionally

present a thorough error evaluation suite to understand the performance of kilonovanet across the input parameter space; these error evaluation methods can also be applied to any future surrogate construction methods. The code used to produce the results in our study is available at https://github.com/klukosiute/kilonovanet.

The work is organized as follows. In Section 2, we discuss the simulated data sets used for constructing the surrogate model. In Section 3, we discuss the cVAE, its limitations with respect to uncertainty estimation, and our approach for quantifying uncertainties. We discuss the performance of the cVAE surrogate model in Section 4 and conclude in Section 5.

2 DATA

In this work, we report the construction and testing of surrogate models for three publicly available data sets of simulated kilonova - in particular, spectra of the material ejected during a BNS or BHNS merger: the *Kasen BNS* simulations (Kasen et al. 2017), the Dietrich BNS simulations (Dietrich et al. 2020), and the Anand BHNS simulations (Anand et al. 2021). In the Kasen simulation, the key physical parameters of the single-component spherically symmetrically expanding ejecta are the ejecta mass $M_{\rm ej}$, the characteristic expansion velocity of the ejecta $v_{\rm ej}$, and the chemical composition as indicated by the lanthanide fraction χ . Each spectrum is computed in increments of 0.1 d, from 0.1 to 24.9 d post-merger and has 1629 logarithmically spaced wavelength bins from 149 Å to 99467 Å. There are two sets of publicly available parameter grids for these simulations available: a regularly spaced grid, whose parameter values are given in Table 1, and a narrower, irregularly spaced set. The values of the second set lie mostly within the ranges of the regularly spaced grid but also includes several models with $v_{\rm ei}/c = 0.4$. This set was was produced particularly for the study of GW170817. There are two sets of parameters. We combine the regular grid data (329 parameter sets¹) and the narrower, irregular grid data (22 parameter steps) to create our full training data set, resulting in 351 unique parameter combinations; with 249 time steps per parameter set, this leads to a total of 87 399 spectra in the *Kasen* simulated data set.

The Dietrich BNS simulations are generated with POSSIS, a multidimensional Monte Carlo radiative transfer code (Bulla 2019). The parameter sets consist of the mass of the dynamical ejecta $M_{\rm ej,\,dyn}$, the mass of the post-merger ejecta $M_{\rm ej,\,pm}$, the half-opening angle of the lanthanide-rich tidal dynamical ejecta Φ , and the cosine of the observer viewing angle $\cos \theta_{\rm obs}$. The half-opening angle Φ parametrizes the separation of the dynamical ejecta into two components: lanthanide-free matter ejected along the poles of the in-spiral plane above latitudes $\pm \Phi$, and lanthanide-rich ejecta tidally expelled equitorially below latitudes $\pm \Phi$. Each spectrum is computed at increments of 0.2 d, starting at 0.2 d up to 20 d post-merger and for 500 evenly spaced wavelength bins from 100 Å to 99 900 Å. There are in total 2156 combinations of the parameter values listed in Table 1, including 'redundant' simulations for different viewing angles at $\Phi = 0$, 90 deg. We keep these redundant simulations in the parameter set so that the surrogate models do not have to build in the spherical symmetry of the system. With 100 time steps for each parameter combination, this leads to 215 600 spectra for this simulation set.

¹There are 329 parameter sets rather than the $11 \times 5 \times 6 = 330$ expected from Table 1 because there exists no *Kasen* run for $(M_{ej}, v_{ej}, \log_{10} \chi) = (0.1, 0.3, -1)$.

Table 1. Parameter values for each of the three kilonova simulation data sets which we use for training our surrogate model.

Kasen et al. (2017) Parameters	Values			
$M_{ m ej}/M_{\odot}$	{0.001, 0.0025, 0.005, 0.01, 0.02, 0.025, 0.03, 0.04, 0.05, 0.075, 0.1}			
$v_{\rm ej}/c$	{0.03, 0.05, 0.1, 0.2, 0.3}			
$\log_{10}\chi$	$\{-9, -5, -4, -3, -2, -1\}$			
Dietrich et al. (2020)				
$M_{\rm ei,\ dyn}/M_{\odot}$	{0.001, 0.005, 0.01, 0.02}			
$M_{\rm ei,\ pm}/M_{\odot}$	{0.01, 0.03, 0.05, 0.07, 0.09, 0.11, 0.13}			
Φ (deg)	{0, 15, 30, 45, 60, 75, 90}			
$\theta_{ m obs}$	$\{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$			
Anand et al. (2021)				
$M_{\rm ej,dyn}/M_{\odot}$	$\{0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09\}$			
$M_{\rm ei, pm}/M_{\odot}$	$\{0.01, 0.02, 0.03, 0.040.05, 0.060.07, 0.08, 0.09\}$			
$\cos \theta_{\rm obs}$	$\{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$			

The Anand BHNS data set is also generated through POSSIS. The dynamical ejecta are concentrated within an angle $\phi=30$ deg above and below the equatorial plane, expanding with velocities ranging 0.1c to 0.3c, and with a lanthanide-rich composition. The post-merger ejecta are assumed to be spherical, expanding with velocities ranging 0.025c to 0.1c, and with an intermediate lanthanide composition. Table 1 lists all the values of these parameters as well as the final parameter $\cos\theta_{\rm obs}$. There are 891 unique parameter combinations and again with 100 time steps each.

3 METHODS

We use the conditional variational autoencoder (cVAE) to construct our surrogate models. We then highlight the subtle but critical challenge of deriving uncertainty estimates directly from Gaussian cVAE's, caused by variance shrinkage. Finally, we discuss data preprocessing, our training protocol, and our method for evaluating the trained cVAE-based model.

3.1 Conditional variational autoencoder

We start with a simplified derivation of the variational autoencoder (VAE; Kingma & Welling 2014), which is a generative model that relies on variational Bayesian methods for optimization. We note the following definitions:

- (i) x: vector of physical parameters, including those discussed in Section 2, like ejecta mass, velocity, and lanthanide fraction, as well as the time step t in the evolution a single kilonova.
 - (ii) y: vector of spectral data from a single kilonova.
 - (iii) z: vector of latent variables
 - (iv) p: generative model
 - (v) q: inference model
 - (vi) θ : parameters for generative neural model network²
 - (vii) ϕ : parameters for inference model neural network

Our goal is to obtain a model to efficiently approximate a generative process $p^*(y)$ (the radiative transfer model) to predict a spectrum y. We seek $p_{\theta}(y) = \int p_{\theta}(y|z)dz$ to approximate $p^*(y)$, where $p_{\theta}(y|z)$ is part of a deep latent variable model (DLVM) and is related to it via the definition of conditional probability.

To evaluate the integral, an inference model $q_{\phi}(\mathbf{z}|\mathbf{y})$ (the encoder) is used to approximate the true $p_{\theta}(\mathbf{z}|\mathbf{y})$, which, through Bayes' theorem, approximates $p_{\theta}(\mathbf{y}|\mathbf{z})$. A VAE is a deep neural network-based framework for co-optimizing a DLVM and an inference model.

A model *EncoderNeuralNet* is trained to return the mean μ and the variance σ of a multidimensional Gaussian for a particular data input y. The mean and variance are then used for sampling the generative model:

$$\begin{aligned} &(\boldsymbol{\mu}, \log \boldsymbol{\sigma}) = \text{EncoderNeuralNet}_{\boldsymbol{\phi}}(\mathbf{y}) \\ &q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{y}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}, \operatorname{diag}(\boldsymbol{\sigma})). \end{aligned}$$
 (1)

The optimization objective of the VAE is the Evidence Lower Bound (ELBO), which allows the simultaneous optimization of the parameters θ (generative model), and the variational parameters ϕ (inference model). The loss function is defined as

$$\mathcal{L}_{\theta,\phi}(\mathbf{y}) = -D_{KL}(q_{\phi}(\mathbf{z}|\mathbf{y})||p_{\theta}(\mathbf{z})) + \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{y})}[\log(p_{\theta}(\mathbf{y}|\mathbf{z}))], \tag{2}$$

where D_{KL} is the Kullback–Liebler Divergence, which measures the distance between the probability p and the inference model q, and \mathbb{E} is the reconstruction loss. Using Monte Carlo sampling and the 'reparametrization trick,' the networks p_{θ} and q_{ϕ} can be optimized using gradient descent methods (Kingma & Welling 2014, 2019; Rezende et al. 2014).

The variational autoencoder learns an approximation for the distribution $p^*(\mathbf{y})$, and the decoder returns the parameters of the distribution $p_{\theta}(\mathbf{y}|\mathbf{z})$. After training, the marginal likelihood $p(\mathbf{y})$ can be estimated through importance sampling, where random samples are drawn from $q_{\phi}(\mathbf{z}|\mathbf{y})$ (Rezende et al. 2014):

$$p_{\theta}(\mathbf{y}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{y})} \left[p_{\theta}(\mathbf{y}, \mathbf{z}) / q_{\phi}(\mathbf{z}|\mathbf{y}) \right] \approx \frac{1}{L} \sum_{l=1}^{L} \frac{p_{\theta}(\mathbf{y}|\mathbf{z}^{(l)}) p(\mathbf{z}^{(l)})}{q_{\phi}(\mathbf{z}|\mathbf{y})}.$$
(3)

Through this process, the VAE can learn a multidimensional distribution as a data-driven model. This distribution $p_{\theta}(\mathbf{y}|\mathbf{z})$ is then the approximation of the $p^*(\mathbf{y})$, the true generative model from the radiative transfer simulation.

In this work, we aim to learn a distribution over the spectral data \mathbf{y} . But, importantly, we must condition the VAE model on input physical parameters \mathbf{x} of the kilonova, transforming the model into a cVAE (Sohn et al. 2015; Kingma & Welling 2019): $p_{\theta}(\mathbf{y}|\mathbf{x})$. The

 $^{^2\}theta_{obs}$ in §2 is unrelated to θ .

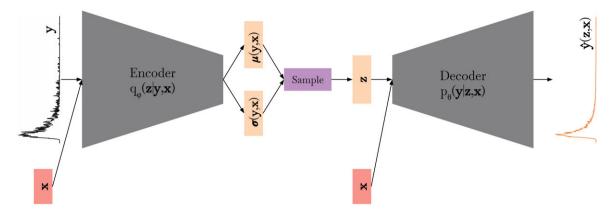


Figure 1. Diagram of the cVAE architecture. There are two neural networks – encoder (left-hand side, grey) and decoder (right-hand side, grey – and a latent space (\mathbf{z} , pink). The inputs to the encoder are the physical parameter \mathbf{x} and the spectrum \mathbf{y} at time step t for a single kilonova event. The outputs from the encoder are $\sigma(\mathbf{x}, \mathbf{y})$ and $\mu(\mathbf{x}, \mathbf{y})$, which are the parameters of the distribution $q_{\phi}(\mathbf{z}|\mathbf{y}, \mathbf{x})$ over the latent variables \mathbf{z} . The inputs to the decoder are \mathbf{z} and \mathbf{x} , which are used to draw a sample from the distribution $p_{\theta}(\mathbf{y}|\mathbf{z}, \mathbf{x})$ over the spectra. This produces an output $\hat{y}(z, x)$, a new surrogate spectrum. After training, the decoder is used as the generator of simulated data for the surrogate model: it takes as input a sample from $\mathcal{N}(0, \mathbf{I})$ and a set of physical parameters \mathbf{x} and predicts a spectrum $\hat{y}(\mathbf{z}, \mathbf{x})$.

optimization objective for the cVAE is the conditional likelihood

$$\mathcal{L}_{\theta,\phi}(\mathbf{y}|\mathbf{x}) = -D_{KL}(q_{\phi}(\mathbf{z}|\mathbf{y},\mathbf{x})||p_{\theta}(\mathbf{z}|\mathbf{x})) + \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{y},\mathbf{x})}[(\log p_{\theta}(\mathbf{y}|\mathbf{z},\mathbf{x}))].$$
(4)

A graphical representation of the model is shown in Fig. 1.

Since sampling from $p(\mathbf{z}) = \mathcal{N}(0, \mathbf{I})$ produces negligible differences in the prediction of new spectra \mathbf{y} , we choose to always use the same sample (the central value of $\mathcal{N}(0, \mathbf{I})$, a vector of zeroes) for \mathbf{z} when predicting spectra.

3.2 Variance shrinkage of gaussian VAEs

Using the prescription above, cVAEs can generate multidimensional data distributions conditioned on input variables when an appropriate likelihood function is chosen. For binary classification, the Bernoulli distribution is typically used to describe data that have binary outcomes. The log-likelihood is given by the binary cross-entropy (BCE) loss,

$$\log p_{\theta}(\mathbf{y}|\mathbf{z}) = -\sum_{i=1}^{D} y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i), \tag{5}$$

where D is the dimension of the vector to be predicted, y_i is a true label, and \hat{y}_i is a predicted label.

For real-valued data, like the kilonova spectra in this work, the multivariate Gaussian with a diagonal covariance offers flexibility and mathematical simplicity with a log-likelihood of the form

$$\log p_{\theta}(\mathbf{y}|\mathbf{z}) = \sum_{i=1}^{D} -\frac{1}{2\sigma_{\theta,i}^{2}(z)} ||y_{i} - \mu_{\theta,i}(z)||^{2} -\frac{1}{2} \log 2\pi \sigma_{\theta,i}^{2}(z),$$
(6)

where $\sigma_{\theta,i}^2(z)$ and $\mu_{\theta,i}(z)$ are the variance and mean, respectively, for the *i*-th dimension.

If we used the Gaussian likelihood for our continuous and real data outputs, we could consider interpreting the output distribution as an uncertainty learned by the model. However, it has recently been shown that the maximum-likelihood objective is ill-posed for continuous models, such as those employing Gaussian distributions (Mattei & Frellsen 2018). For models trained so that $y_i \sim \mu_{\theta,i}(z)$,

the $-\frac{1}{2}\log 2\pi\sigma_{\theta,i}^2(z)$ term will push the variance to zero before the $\frac{1}{2\sigma_{\theta,i}^2(z)}$ term can catch up. This can be seen in equation (6). Therefore, a Gaussian cVAE will produce a small variance that does not have useful physical interpretation. This 'variance shrinkage problem' prohibits use of the cVAE model as a probabilistic method (Mattei & Frellsen 2018; Skafte, Jørgensen & Hauberg 2019).

There are several standard procedures to avoid the variance shrinkage problem in practice (Skafte, Jørgensen & Hauberg 2019). For example, setting a globally constant variance – e.g. $\sigma^2 = 1$ – the log-likelihood becomes the mean squared error, and the cVAE loss function becomes

$$\log p_{\theta}(\mathbf{y}|\mathbf{z}) = \sum_{i=1}^{D} -||y_i - \mu_{\theta,i}(z)||^2.$$
(7)

This uncertainty is predetermined and not learned from the data.

Because the MSE loss does not provide statistical interpretability, the Bernoulli distribution is often used for non-binary data: optimizing a Bernoulli log-likelihood (the BCE loss in equation 5) is considerably simpler than optimizing the MSE (Skafte, Jørgensen & Hauberg 2019). This can be done when the data is scaled to the range [0,1]. However, the outputs will not have a meaningful statistical interpretation because the Bernoulli distribution is meant for discrete random variables with a binary outcome. Our initial experiments verify that training with BCE loss outperforms the MSE loss in terms of accuracy and efficiency.

The closed-form expression for our optimization objective is thus

$$\mathcal{L} = -\frac{1}{2} \sum_{i=1}^{D} \left[1 + \log \sigma_i^2 - \sigma_i^2 - \mu_i^2 \right] - \sum_{i=1}^{D} y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i),$$
 (8)

where we have used the BCE loss and $p(\mathbf{z}|\mathbf{x}) = \mathcal{N}(0, I)$, as per the definition of the VAE model, and where μ_i and σ_i^2 are the outputs of the encoder distribution for passing a data point pair (x_i, y_i) , with a resulting prediction \hat{y}_i . While variance shrinkage prevents learning of the variance from the data, the cVAE can still reproduce spectral data with high fidelity. We pursue estimates of the variance post-facto

through measurements of the surrogate model data with respect to the original simulated data.

3.3 Data pre-processing

The original *Kasen* spectra are in units of erg s⁻¹ Hz⁻¹, and the *Dietrich* and *Anand* spectra are in units of erg s⁻¹ cm⁻² Å⁻¹ at a distance of 10 pc. To use the data sets for network training and to evaluate the surrogate results, we first process all spectra such that they have units erg s⁻¹ Å⁻¹. We then re-scale the input physical parameters and the spectra to the range [0,1] so that the data lie in the supported range of the sigmoid activation function in the final layer of the decoder neural network. Finally, we separate each of the three simulation data sets into training, validation, and test sets with relative proportions 80:10:10, respectively. We use only the training and validation sets to optimize the training and hyperparameters.

3.4 Training

For each of the three simulated kilonova spectra data sets, we train a distinct cVAE model with the following procedure.

- (i) We use the training and validation data sets from the top-level split to perform hyperparameter selection for our final architecture. For each simulation data set, we perform a hyperparameter search using one of the train-validation data splits. We consider the hyperparameters of the cVAE model to be the dimensionality of the latent space z and the dimensionality of the hidden layers of the decoder and encoder. Because the encoder is considered the approximate inverse of the decoder, we use one hyperparameter for both the hidden layers. We perform hyperparameter selection only on these two values; all other hyperparameters of the model, such as the learning rate and batch size, are fixed across all of our tests and models. We trained 16 architectures with four values for each hyperparameter.
- (ii) We then split the training and validation sets into nine total sets, and train the architecture on these subsets.
- (iii) We then evaluate all nine models on the test set, which allows us to develop a statistical analysis of the model's predictive capability.
- (iv) We chose one of the nine models at random for the final production model that we use to show examples and make publicly available.

We train each model – both when performing hyperparameter selection and when training each of the nine data split models – for 200 epochs. The final training losses for the *Kasen*, *Dietrich*, and *Anand* data, averaged over the nine data split models, respectively, are 48.54, 11.23, and 8.20. The training during the hyperparameter search required ~24 h, and the final nine experiment models required approximately 26 h of training on the same GPU – both on an Nvidia GeForce 1080Ti GPU. The prediction of 100 spectra of the trained network from unique parameter sets requires approximately 10 ms on one Intel®CoreTM i7-7700HQ CPU. We use PyTorch for implementation (Paszke et al. 2019) and Adam as our gradient-based optimizer (Kingma & Ba 2015).

3.5 Model evaluation and quantifying prediction uncertainties

For each simulation data set, we have constructed a surrogate cVAE-based model with which we can generate data under the supported ranges of the physical input variables \mathbf{x} and output spectra \mathbf{y} .

We next seek to generate uncertainties associated with the predictions of the spectra. Variance shrinkage and the resulting training choices dictate that the cVAE cannot produce statistically meaningful

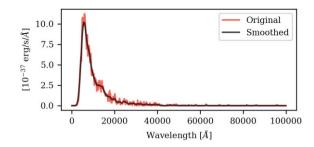


Figure 2. An example original spectrum (orange) and its smoothed version using a Gaussian kernel with $\sigma=2.5$ (blue). This spectrum is from the *Dietrich* data set and corresponds to the parameters $M_{\rm ej,\,dyn}/M_{\odot}=0.01$, $M_{\rm ej,\,pm}/M_{\odot}=0.07$, $\Phi=30.0^{\circ}$, $\cos\Theta_{\rm obs}=0.5$, t=1.8 d.

variances. Therefore the errors in the resulting surrogate models still need to be characterized. We use the cVAE only as a method to produce more samples of the original radiative transfer simulations with which the cVAE was trained: we do not seek to improve upon or extrapolate beyond those simulations. In the following sections, we discuss error sources and procedures for quantifying them.

3.5.1 Sources of error

We examined and estimated multiple sources of error with respect to the original simulations. First, for systematic error, we account for the bias of the predicted spectra as a function of input parameters by comparing predictions to the truth generated by the radiative transfer simulations. The surrogate model is generated to minimize bias in reconstructing spectra, so the model (with its chosen hyperparameter values) is designed to minimize this bias. Secondly, we account for statistical error of the predictions over the space of input parameters and different trained models, as we discuss in detail below.

3.5.2 Error metrics

The statistical error present in the training simulations will propagate to the surrogate data. Predicting synthetic observables from radiative transfer simulations results in Monte Carlo noise, leading to noisy spectra (Kasen et al. 2017; Bulla 2019): this represents the floor of the statistical noise. A surrogate model could recreate the noise present in the training data, but no surrogate model, even a hypothetical model with infinite capacity, will be able to perfectly predict the noise within the unseen test set. Therefore, the Monte Carlo noise from the radiative transfer simulations leads to a source of error in the final test set predictions. We would like to know the value of the error that is incurred because of the Monte Carlo noise present in the test data set. Estimating the true Monte Carlo noise would require simulating the observables for each parameter simulation several times and then computing residuals from the mean spectrum (Bulla, Sim & Kromer 2015) – a process that is computationally expensive. To estimate the error, we emulate the mean spectrum by smoothing each spectrum using a Gaussian smoothing method. We cannot know which value of the Gaussian kernel represents the true mean spectrum best, so we perform Gaussian smoothing for several values of the Gaussian kernel. We then visually inspect the spectra to check for an approximate best fit, which occurs at a value of $\sigma = 2.5$. Fig. 2 shows an original smoothed spectra and its smoothed counterpart. In Section 4, we compare our estimated fractional Monte Carlo noise with the errors of the kilonovanet predictions.

1142

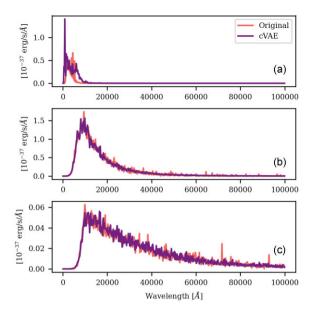


Figure 3. Three original spectra (orange) and corresponding cVAE predictions (blue) for physical parameters $M_{\rm ej,\ dyn}/M_{\odot}=0.02$, $M_{\rm ej,\ pm}/M_{\odot}=0.05$, $\Phi=45.0^{\circ}$, $\cos\Theta_{\rm obs}=0.8$, and times (a) 0.2 d, (b) 4.2 d, and (c) 14.2 d. The corresponding median spectral errors across the whole spectra (a) 12.91, (b) 0.37, and (c) 0.20.

We evaluate the model performance using three metrics: the spectral error, the bolometric luminosity error, and the band magnitude error.

The spectral error is the fractional error between the model prediction in a given wavelength bin $y_{\text{pred}}(\lambda)$ and the value of the original test data point $y_{\text{test}}(\lambda)$:

$$\epsilon_s(\lambda) = \frac{y_{\text{pred}}(\lambda) - y_{\text{test}}(\lambda)}{y_{\text{test}}(\lambda)}.$$
(9)

To connect with photometric survey observations, we quantify the bolometric luminosity error as

$$\epsilon_b = \frac{\int y_{\text{pred}}(\lambda) d\lambda - \int y_{\text{test}}(\lambda) d\lambda}{\int y_{\text{test}}(\lambda) d\lambda}.$$
 (10)

We also evaluate the performance of the surrogate model by constructing broad-band light curves at a distance of 40 Mpc for each of the unique parameter sets in our test data set. Broad-band AB magnitudes from simulated spectra are computed by convolving the flux with the broad-band filters at a chosen distance (40 Mpc in this work). The distance is chosen to correspond to the distance of GW170817 which allows us to compare with past analyses. These light curves are computed as AB magnitudes in each band,³ of which examples are shown in Fig. 7. We compute and evaluate the error in the AB magnitude in each band as

$$\Delta m = m_{\text{band,pred}} - m_{\text{band,test}},\tag{11}$$

where $m_{\rm band, \, pred}$ is the AB magnitude in a band at a given time step predicted by kilonovanet, and $m_{\rm band, \, test}$ is magnitude found from the corresponding spectrum in the test set.

Finally, we will evaluate the performance of the surrogate model by performing a representative inference task and comparing with previously published results. We perform parameter estimation using nested sampling with the dynesty (Speagle 2020) sampler on the GW170817 light-curve data using our surrogate models. We compare the differences in the best-fitting parameters between our fit using the same data set, which was first collated in Coughlin et al. (2018) and previously published fits for the same BNS kilonova models but using a different surrogate construction methods.

4 RESULTS

We present the results of training and evaluating the cVAE on the three simulated data sets. For each simulated data set, we have nine trained models (one for each data-splitting experiment as described in Section 3.3. We pass the test data sets – i.e. the pairs of input (\mathbf{x}) and output (\mathbf{y}) values that were not used in any aspect of training or model selection – through the decoders of each model to obtain a predicted spectrum. We compare the predictions and the true values of the spectra by computing the values presented in Section 3.5.2.

We report a detailed analysis of the surrogate model for the *Dietrich* BNS data set and then report only key values for the *Anand* and *Kasen* data sets; the analysis is the same for all three data sets. We focus on the *Dietrich* data set, because we have observational data for a single BNS event, and the *Dietrich* model is a newer BNS kilonova simulation set.

4.1 Quantifying uncertainties

4.1.1 Spectral error

We first compute $\mathtt{median}_{\mathbf{x},\lambda}\epsilon_s(\lambda)$ and $\mathtt{mean}_{\mathbf{x},\lambda}\epsilon_s(\lambda)$ for each of the nine data-split experiments. We calculate the median and mean over all sets of test input parameters (including the time parameter) and over all wavelengths in the spectra \mathbf{y} . We then compute the mean and standard deviation of the aforementioned means and medians over all nine data-split experiments.

We report the median and mean of the absolute value of $\epsilon(\lambda)_s$: 0.067 ± 0.20 and 6620 ± 3180 , where the error is given as the standard deviation over the aforementioned nine test splits. $\mathsf{mean}_{x,\lambda}\epsilon_s(\lambda)$ is heavily skewed by overpredictions when the true value is close to zero for a few outlier values and is therefore not a true representation of the of the typical error. Although it is not representative, we report the mean so that we may look for skew of the error distribution. Because the mean is greater than the median, the distribution of errors is skewed high towards overprediction or slightly brighter kilonovae. The corresponding values for the *Anand* and *Kasen* data sets are given in Table 2.

We then perform the same set of operations on the absolute value of the spectral error: the mean and variance of this value is 0.285 ± 0.004 , where the variance is again over all nine experiments. We estimate the median absolute fractional Monte Carlo noise using smoothing via a Gaussian kernel with $\sigma=2.5$ and find it to be 0.214. The size of the median spectral error median_{x,\(\lambda\)}|\epsilon_s(\(\lambda\))| is due to the cVAE surrogate learning the general shape of the spectra but not the Monte Carlo noise. We leave the interpretation and impact of the size of the error for when we discuss errors in broad-band filters.

In Fig. 3, we present the spectrum of a BNS kilonova and its corresponding surrogate model predictions for the input parameters $M_{\rm ej,\,dyn}/M_{\odot}=0.02$, $M_{\rm ej,\,pm}/M_{\odot}=0.05$, $\Phi=45.0^{\circ}$, and $\cos\Theta_{\rm obs}=0.8$ at three different time steps (i) 0.2 d, (ii) 4.2 d, and (iii) 14.2 d. The respective median_{λ} | ϵ_s | for these spectra are 12.91, 0.37, and 0.20. Fig. 3(a) shows significant differences between the original and the surrogate model where the spectrum is non-zero. In general, we find that all spectra at t=0.2 d are poorly predicted by the cVAE.

³We use the LSST filters provided at http://svo2.cab.inta-csic.es/svo/theory/fps3/

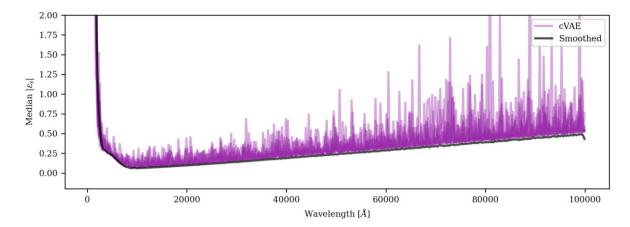


Figure 4. Median of $|\epsilon_s|$ over all the spectra in the *Dietrich* data set for the predictions of the cVAE (purple), along with the absolute spectral error from the Gaussian smoothed spectra for a Gaussian kernel with $\sigma = 2.5$ (black).

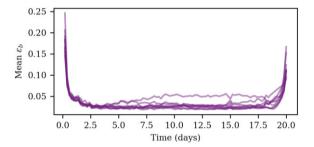


Figure 5. Mean of $|\epsilon_b|$ across all nine data-split models of the *Dietrich* data set. The two times at which the model is least accurate is at t = 0.2 and 20 d.

However, there is general simulation uncertainty at early times due to a lack of atomic data (Banerjee et al. 2020): due to the lack of detailed simulations of these light curves and the current lack of wide-field, rapid response UV telescopes to detect them, we can safely defer improvement of early-time simulation data to future work. In contrast, the spectra produced by the cVAE model at $t \gtrsim 1$ d are much more faithful to the original simulations, as evidenced by panels (b) and (c) in Fig. 3 and their lower median $|\epsilon_s|$.

Fig. 4 shows the median $median_x|\epsilon_s(\lambda)|$ i.e. the median absolute spectral error over all examples in the test set, including time. There are nine lines plotted for the medians, corresponding to the nine data split cVAE models. In addition, we show the estimate of the absolute fractional Monte Carlo noise using smoothing via Gaussian kernel with $\sigma=2.5$ as a function of wavelength. The computation of the fractional Monte Carlo noise is the same as the computation of $|\epsilon_s|$ and the values can thus be compared directly. The Monte Carlo noise estimate traces the bottom of the errors produced by the cVAE. This could imply that the Monte Carlo noise sets a lower limit for $|\epsilon_s|$ of the cVAE or perhaps any surrogate model.

4.1.2 Bolometric luminosity error

Next, we discuss the errors in the bolometric light curves. For Dietrich, mean_x $|\epsilon_b| = 0.033 \pm 0.005$, where the error is the standard deviation across all nine data split experiments.

An investigation into the outliers in the *Dietrich* data set shows that the spectra with the largest bolometric luminosity error occur at times t=0.2 and t=20. The spectra with the highest bolometric luminosity error (of 1.56) is shown in Fig. 3(a), which also corresponds to t=0.2. The spectra at t=20 are also often overpredicted, leading to ϵ_b values of approximately 0.7. The outliers at t=0.2 and t=20 are the

most prominent features across all nine data-split models in the plot of mean bolometric luminosity error $\max_{M_{\rm ej,dyn},M_{\rm ej,pm},\Phi,\cos\theta_{\rm obs}}|\epsilon_b|$ as a function time, which is shown in Fig. 5. The correlations with time are different for the other two data sets. The light curves of the *Kasen* set have a tendency to quickly fall to zero after ~ 14 d: the errors increase rapidly at about that time. The *Anand* data set also exhibits a large error peak at t=0.2, but since BHNS kilonovae dim more slowly, the errors remain consistently low even at t=20 d.

4.1.3 Broad-band magnitude error

Finally, we evaluate the performance of the surrogate model in broadband magnitudes, ugrizy. We construct light curves for each set of input parameters in the test data set and compute the mean magnitude differences. We report the mean of the magnitude differences $mean_x\Delta m$, where a mean is taken across all sets of input parameters in the test set, where time is also treated as an input parameter. We report the values for *Kasen* and *Anand* in Table 2.

For the Dietrich test set, we additionally show the distribution of $\mathsf{mean}_{M_{cj,dyn},M_{cj,pm},\Phi,\cos\theta_{obs}}\Delta m$ for selected time steps of the data set. Fig. 6 shows these distributions of Δm for all six bands. In general, for the time steps other than t=0.2, the width of the distributions of Δm is correlated with brightness. Overall, most of the predictions are within 0.1 mag of the test data set. This is well within the commonly assumed 1 mag uncertainty in other KN light-curve modelling fitting (Coughlin et al. 2018). The least accurate predictions occur for the Anand data set in u-band, with a median difference of 0.451 mag. Since the Anand data set is of a BHNS model, the model does not contain a lanthanide-free ('bluer') component and would thus be practically undetectable in the blue bands, leading to high errors in the surrogate model.

We present three examples for light curves produced by the cVAE, along with the original light curves and absolute differences between the prediction and original in Fig. 7. The three randomly chosen examples are representative of the remainder of the test set in that they show that the cVAE predictions in bands are smooth and that the predictions become unreliable when the light curve is too faint. Fig. 7(c) shows an example for a model free of lanthanide rich ejecta, since the half opening angle Φ of the equatorial component is zero. This particular example peaks quickly and dims quickly, and below mag~27, the cVAE fails to predict the original data. However, the kilonova at a distance of 40 Mpc would not be detectable at such a magnitude; for example, the design single-visit limiting magnitude for the LSST u-band is 23.9 (Ivezić et al. 2019). Fig. 7(c) shows

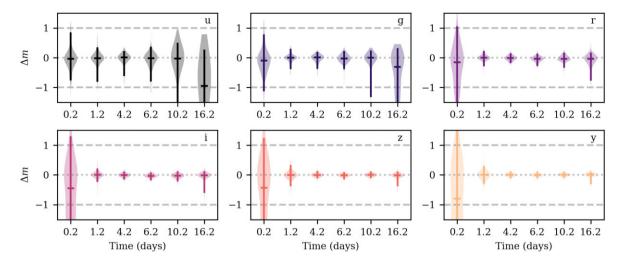


Figure 6. Distributions of magnitude errors Δm between the predictions and the test data set for five time steps in the *Dietrich* simulation set. The horizontal lines for each violin represents the median of the distribution, and the endpoints of the vertical lines represent the 95 per cent confidence interval.

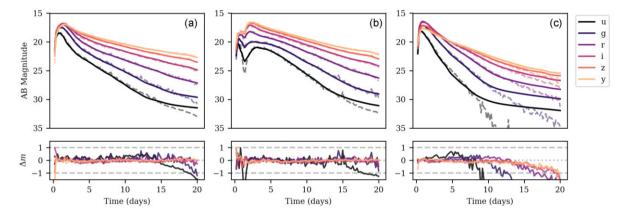


Figure 7. Three example light curves from the *Dietrich* BNS data sets for the original data set (dashed lines) and the cVAE predictions (solid lines) for the six ugrizy bands (colours black through yellow), as well as the residuals between original and prediction for each light curve and band. The three parameter sets are (a) $M_{ej, dyn}/M_{\odot} = 0.01$, $M_{ej, pm}/M_{\odot} = 0.01$, $M_{ej, pm}/$

Table 2. Summary of errors for all models.

	$ \epsilon_{s} $	MC noise estimate	$\epsilon_{s, ext{mean}}$	$\epsilon_{s,\mathrm{med}}$	$ \epsilon_b $	
Dietrich	0.285 ± 0.004	0.214	6620 ± 3180	0.067 ± 0.20	0.033 ± 0.005	
Anand	0.292 ± 0.005	0.241	1280000 ± 95900	0.114 ± 0.016	0.027 ± 0.004	
Kasen	0.202 ± 0.005	0.127	0.213 ± 0.044	0.009 ± 0.014	0.057 ± 0.008	
	$ \Delta m _u$	$ \Delta m _g$	$ \Delta m _r$	$ \Delta m _i$	$ \Delta m _z$	$ \Delta m _y$
Dietrich	0.281 ± 0.016	0.164 ± 0.014	0.099 ± 0.028	0.075 ± 0.019	0.065 ± 0.019	0.052 ± 0.006
Anand	0.462 ± 0.016	0.263 ± 0.016	0.145 ± 0.009	0.101 ± 0.013	0.088 ± 0.014	0.081 ± 0.009
Kasen	0.176 ± 0.007	0.136 ± 0.014	0.094 ± 0.011	0.090 ± 0.015	0.083 ± 0.010	0.092 ± 0.010

the limitations of the model, especially at the edges of the parameter space, but we emphasize that that the errors primarily occur when the kilonova has dimmed beyond detectability.

4.2 Fitting AT2017gfo

We test the performance of the surrogate model in the scientific context that it will be used: we use the *Dietrich*-based surrogate model to perform kilonova parameter estimation on the

object AT2017gfo (counterpart to the event GW170817) with the data collated in Coughlin et al. (2018) from the sources of Andreoni et al. (17ed), Arcavi et al. (2017), Chornock et al. (2017), Cowperthwaite et al. (2017), Drout et al. (2017), Evans et al. (2017), Kasliwal et al. (2017), Tanvir et al. (2017), Pian et al. (2017), Troja et al. (2017), Smartt et al. (2017), Utsumi et al. (2017). We use flat priors that extend over the published data range and therefore the range of supported parameter by the surrogate model $\log_{10}(0.001) \leq \log_{10}(M_{\rm ej,\,dyn}/M_{\odot}) \leq \log_{10}(0.02)$,

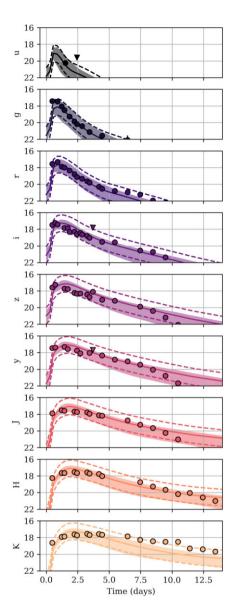


Figure 8. Light curves for AT2017gfo. Observed values (points and triangles, where the latter are lower bound observations) and the prediction based inferred parameters using the *Dietrich*-based cVAE surrogate model (solid lines). The shaded bands represent the 90 per cent confidence interval of light curves constructed from the posterior samples. The dashed lines represent the 1 mag tolerance typically used to represent modelling error of kilonova light curves.

 $\log_{10}(0.01) \le \log_{10}(M_{\rm ej,\,pm}/M_{\odot}) \le \log_{10}(0.13), 0 \le \cos{(\Phi)} \le 1$, and $0 \le \cos{(\theta)} \le 1$.

We use the dynesty sampler and the log-likelihood (up to a constant) of

$$\log \mathcal{L} = -\frac{1}{2} \sum_{i=1}^{N} \frac{(m_{i,\text{pred}} - m_{i,\text{obs}})^{2}}{\sigma_{i,\text{obs}}^{2} + \sigma_{\text{sys}}^{2}},$$
(12)

where *i* indexes observations, $m_{i, \text{pred}}$ is the magnitude of each proposed sample, $m_{i, \text{obs}}$ is the magnitude of each observation, and $\sigma_{i, \text{obs}}$ is the uncertainty of each observation. The likelihood is Gaussian with an additional systematic uncertainty of $\sigma_{\text{sys}} = 1$ mag added to account for the modelling uncertainty, as used in Coughlin et al. (2018). Fig. 8 shows the light curves derived from the fit,

and Fig. 9 shows the parameter posteriors, as well as the best-fitting median values on the same data set obtained by Dietrich et al. (2020).

There are some key differences to note for this comparison. The fit performed by Dietrich et al. (2020) on the light curve data uses as a prior gravitational wave and pulsar observations. In addition, they used different allowed ranges of the kilonova model parameters. With all these differences in mind, the agreement between our parameter recovery and the fit presented in their analysis shows remarkable consistency, with each of the medians of their recovered parameters lying within 1σ of our median recovered parameters, as shown in Fig. 9. The parameter inference performed with dynesty for this data set required 3.5 min on an Apple M1 Pro chip.

In the Appendix, we include a more directly comparable fit using the *Kasen* model as well. An additional test would involve generating light curves from injected parameters and seeing whether the same parameters are recovered through a sampling-based fit. This has been performed in Lukošiute (2021) using the surrogate constructions described in this work for the *Dietrich* and *Anand* data sets with success; see Lukošiute (2021) for details.

5 CONCLUSION

In this work, we presented an implementation of the conditional variational autoencoder (cVAE) as a method for surrogate model construction of kilonova spectra. We discussed the method's potential to produce complex, high-dimensional data, including the theoretical and practical limitations, especially with respect to uncertainty quantification. We applied the method to three commonly used kilonova spectra model data sets – those published in Kasen et al. (2017), Dietrich et al. (2020), and Anand et al. (2021). Because the model is unable to learn its own uncertainties, we delineated potential sources of error and developed a suite of metrics tailored to these specific data sets and their scientific uses.

We evaluated the surrogate model's performance in the context of directly predicting spectra, as well as on downstream products like the bolometric luminosities and *ugrizy* broad-band magnitudes. As a final test of the model, we applied it in a typical scientific use case: measuring the physical parameters of the kilonova associated with GW170817. While the cVAE method is limited by variance shrinkage, it can still quickly produce spectra with high fidelity to the original training simulations. Kilonova models that are published in the future can have a surrogate pre-trained and be used for fast inference for any set of photometric bands.

We leave for future work studying how to address the variance shrinkage problem so that the variance of the surrogate can be learned and the uncertainty incorporated into Bayesian analyses. Some potential avenues include using the CombVAE (Skafte, Jørgensen & Hauberg 2019) or the QR-VAE (Akrami et al. 2020), or other machine learning methods that allows for accurate density estimation, such as normalizing flows (Dinh, Krueger & Bengio 2015; Rezende & Mohamed 2015). We leave also for future work how to incorporate the errors incurred due to surrogate modelling into scientific analyses. Nevertheless, we have shown that the cVAE can already serve as a fast and accurate surrogate model for kilonova spectra and have evaluated uncertainty without modelling variance directly. We release the final trained models for the three data sets and user-friendly code to produce spectra at https://github.com/klukosiute/kilonovanet.

ACKNOWLEDGEMENTS

We acknowledge the contributions of the authors in detail.

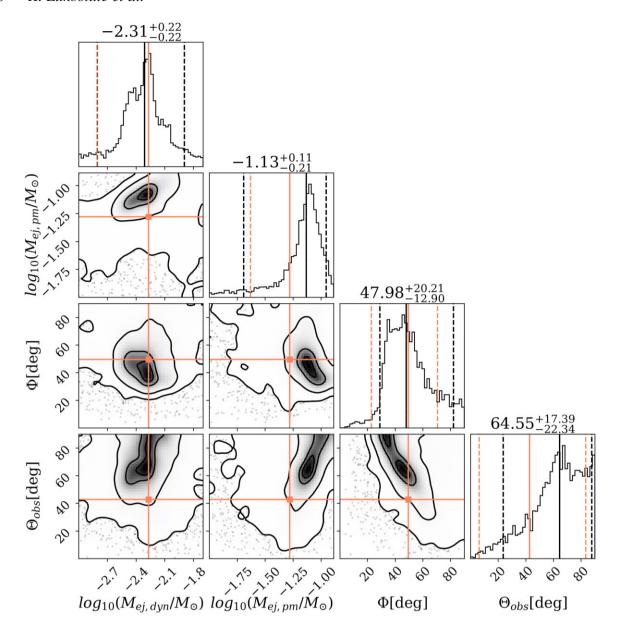


Figure 9. Inferred posteriors of model parameters $M_{\rm ej,\,dyn}$, $M_{\rm ej,\,pm}$, Φ , and $\cos\theta_{\rm obs}$ from the observations of AT2017gfo (as collated in Coughlin et al. 2018) at 10 per cent, 32 per cent, 68 per cent, and 95 per cent confidence intervals. The median values and 90 per cent confidence intervals are shown as vertical solid and dashed lines and above each column, where the orange lines indicate the 90 per cent confidence interval from Dietrich et al. (2020) using the same data set and kilonova model but a different surrogate construction and prior and the black lines indicate our 90 per cent confidence interval.

KL led all computations and writing. GR consulted on analysis of the observational data and co-wrote. ZD consulted on the statistical analysis and co-wrote. MSS posed the problem and edited the manuscript. BN helped develop the initial computational approach, co-designed the statistical analysis, and co-led writing of the paper.

We also thank Samaya Nissanke, Bas Dorsman, and James Annis for useful discussions and comments.

ZD is supported by the CIERA Board of Visitors Research Professorship. GR and KL are grateful for financial support from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) through the Projectruimte and VIDI grants (Nissanke). We acknowledge the Deep Skies Lab as a community of multidomain experts and collaborators who've facilitated an environment of open discussion, idea-generation, and collaboration. This community was important for the development of this project. We thank SURF-

sara (www.surfsara.nl) for the support in using the Lisa Compute Cluster.

Work supported by the Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

DATA AVAILABILITY

No previously unpublished data were generated nor analysed in this work. All data analysed in this article come from previously published articles and may be found by referencing the cited works. The models generated and analysed in this article may be accessed at https://github.com/klukosiute/kilonovanet.

REFERENCES

Aasi J. et al., 2015, Class. Quantum Gravity, 32, 074001

Abbott B. P. et al., 2017a, Phys. Rev. Lett., 119, 161101

Abbott B. P. et al., 2017b, Nature, 551, 85

Abbott B. P. et al., 2017c, ApJ, 848, L12

Abbott B. P. et al., 2017d, ApJ, 848, L13

Abbott R. et al., 2021a, Phys. Rev. X, 11, 021053

Abbott R. et al., 2021b, preprint (arXiv:2111.03606) [astro-ph, physics:gr-qc]

Abbott R. et al., 2021c, ApJ, 915, L5

Acernese F. et al., 2015, Class. Quantum Gravity, 32, 024001

Akrami H., Joshi A. A., Aydore S., Leahy R. M., 2020, preprint (arXiv:2010.09042) [cs, eess, stat]

Akutsu T. et al., 2019, Nat. Astron., 3, 35

Almualla M., Ning Y., Bulla M., Dietrich T., Coughlin M. W., Guessoum N., 2021, preprint (arXiv:2112.15470) [astro-ph]

Anand S. et al., 2021, Nat. Astron., 5, 46

Antier S. et al., 2020, MNRAS, 497, 5518

Arcavi I. et al., 2017, Nature, 551, 64

Arcavi I., 2018, ApJ, 855, L23

Banerjee S., Tanaka M., Kawaguchi K., Kato D., Gaigalas G., 2020, ApJ, 901, 29

Bulla M., 2019, MNRAS, 489, 5037

Bulla M., Sim S. A., Kromer M., 2015, MNRAS, 450, 967

Burbidge E. M., Burbidge G. R., Fowler W. A., Hoyle F., 1957, Rev. Mod. Phys., 29, 547

Cameron A. G. W., 1957, PASP, 69, 201

Chase E. A. et al., 2022, ApJ, 927, 163

Chornock R. et al., 2017, ApJ, 848, L19

Coughlin M. W. et al., 2018, MNRAS, 480, 3871

Coughlin M. W. et al., 2019, ApJ, 885, L19

Coughlin M. W. et al., 2020, Nat. Commun., 11, 4129

Coughlin M., Dietrich T., Kawaguchi K., Smartt S., Stubbs C., Ujevic M., 2017, ApJ, 849, 12

Coulter D. A. et al., 2017, Science, 358, 1556

Cowperthwaite P. S. et al., 2017, ApJ, 848, L17

Dietrich T., Coughlin M. W., Pang P. T. H., Bulla M., Heinzel J., Issa L., Tews I., Antier S., 2020, Science, 370, 1450

Dietrich T., Ujevic M., 2017, Class. Quantum Gravity, 34, 105014

Dinh L., Krueger D., Bengio Y., 2015, preprint (arXiv:1410.8516) [cs]

Doctor Z., Farr B., Holz D. E., Pürrer M., 2017, Phys. Rev. D, 96, 123011

Drout M. R. et al., 2017, Science, 358, 1570

Evans P. A. et al., 2017, Science, 358, 1565

Goldstein D. A. et al., 2019, ApJ, 881, L7

Gompertz B. P. et al., 2020, MNRAS, 497, 726

Graham M. J. et al., 2020, Phys. Rev. Lett., 124, 251102

Holz D. E., Hughes S. A., 2005, ApJ, 629, 15

Ivezić Ž. et al., 2019, ApJ, 873, 111

Kasen D., Badnell N. R., Barnes J., 2013, ApJ, 774, 25

Kasen D., Metzger B., Barnes J., Quataert E., Ramirez-Ruiz E., 2017, Nature, 551, 80

Kasliwal M. M. et al., 2017, Science, 358, 1559

Kawaguchi K., Fujibayashi S., Shibata M., Tanaka M., Wanajo S., 2021, ApJ, 913, 100

Kawaguchi K., Kyutoku K., Shibata M., Tanaka M., 2016, ApJ, 825, 52

Kingma D. P., Ba J., 2015, ICLR 2015, San Diego, CA

Kingma D. P., Welling M., 2014, ICLR 2014, Banff, Canada

Kingma D. P., Welling M., 2019, Foundations and Trends® in Machine Learning, 12, 307

Korobkin O. et al., 2021, ApJ, 910, 116

Lattimer J. M., Schramm D. N., 1974, ApJ, 192, L145

Li L.-X., Paczyński B., 1998, ApJ, 507, L59

Lukošiute K., 2021, Master's thesis, University of Amsterdam

Mattei P.-A., Frellsen J., 2018, in Bengio S., Wallach H., Larochelle H., Grauman K., Cesa-Bianchi N., Garnett R., eds, Proceedings of the 31st International Conference on Neural Information Processing Systems -NeurIPS 2018. Curran Associates, Inc., Montreal, Canada

Metzger B. D. et al., 2010, MNRAS, 406, 2650

Metzger B. D., 2019, Living Rev. Relat., 23, 1

Metzger B. D., Berger E., 2012, ApJ, 746, 48

Page K. L. et al., 2020, MNRAS, 499, 3459

Paszke A. et al., 2019, in Wallach H., Larochelle H., Beygelzimer A., eds, Proceedings of the 33rd International Conference on Neural Information Processing Systems - NeurIPS 2019. Curran Associates, Inc., Vancouver, Canada, p. 32

Perego A., Radice D., Bernuzzi S., 2017, ApJ, 850, L37

Pian E. et al., 2017, Nature, 551, 67

Pian E., 2021, Front. Astron. Space Sci., 7, 108

Raaijmakers G. et al., 2021b, ApJ, 918, L29

Raaijmakers G. et al., 2021a, ApJ, 922, 269

Radice D., Perego A., Hotokezaka K., Fromm S. A., Bernuzzi S., Roberts L. F., 2018b, ApJ, 869, 130

Radice D., Perego A., Zappa F., Bernuzzi S., 2018a, ApJ, 852, L29

Rezende D. J., Mohamed S., 2015, Proceedings of the 32nd International Conference on Machine Learning, PMLR, 37, 1530

Rezende D. J., Mohamed S., Wierstra D., 2014, Proceedings of the 31st International Conference on Machine Learning, PMLR, 32, 1278

Ristic M. et al., 2022, Phys. Rev. Research, 4, 013046

Rosswog S., Liebendoerfer M., Thielemann F.-K., Davies M. B., Benz W., Piran T., 1999, A&A, 341, 499

Savchenko V. et al., 2017, ApJ, 848, L15

Schutz B. F., 1986, Nature, 323, 310

Shappee B. J. et al., 2017, Science, 358, 1574

Skafte N., Jørgensen M., Hauberg S., 2019, in Wallach H., Larochelle H., Beygelzimer A., d'Alché-Buc F., Fox E., Garnett R., eds, Proceedings of the 33rd International Conference on Neural Information Processing Systems - NeurIPS 2019. Curran Associates, Inc., Vancouver, Canada, p. 32

Smartt S. J. et al., 2017, Nature, 551, 75

Soares-Santos M. et al., 2017, ApJ, 848, L16

Sohn K., Yan X., Lee H., 2015, in Cortes C., Lee D. D., Sugiyama M., Garnett R., eds, Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2. NIPS'15. MIT Press, Cambridge, MA, USA, p. 3483

Speagle J. S., 2020, MNRAS, 493, 3132

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

Tanvir N. R. et al., 2017, ApJ, 848, L27

Troja E. et al., 2017, Nature, 551, 71

Utsumi Y. et al., 2017, PASJ, 69

Valenti S. et al., 2017, ApJ, 848, L24

Wollaeger R. T. et al., 2018, MNRAS, 478, 3298 Wollaeger R. T. et al., 2021, ApJ, 918, 10

APPENDIX A:

In addition to performing parameter estimation for AT2017gfo for the *Dietrich* model, we also perform parameter estimation for the single component *Kasen* model. Most literature agrees that more than one component is present in kilonova outflows (Kasen, Badnell & Barnes 2013; Dietrich & Ujevic 2017; Coughlin et al. 2018; Raaijmakers et al. 2021a). We aim to compare the performance of our surrogate model directly to another published result. Therefore, we present this particular parameter estimation to compare with the parameter inference presented in Coughlin et al. (2018). We use the same likelihood (equation 12), sampling software (dynesty), and observational data set for the light curve of AT2017gfo as we used for the parameter inference for the *Dietrich* model (Speagle 2020). We use flat priors: $\log_{10}(0.001) \leq \log_{10}(M_{ei}/M_{\odot}) \leq \log_{10}(0.1)$, $0.03c \leq$

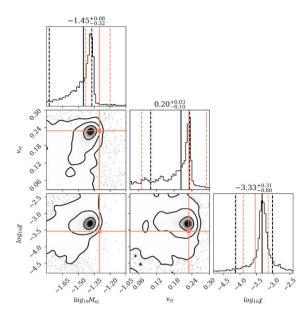


Figure A1. Corner plot for the inferred posterior distribution of the *Kasen* model parameters mass of the ejecta M_{ej} , velocity of the ejecta v_{ej} , and lanthanide fraction χ from the observations of AT2017gfo as collected in Coughlin et al. (2018) at 10 per cent, 32 per cent, 68 per cent, and 95 per cent confidence. The dashed lines in the 1D distributions, as well as the values given in the titles, represent 68 per cent confidence interval, with the median lying between the two dashed lines representing the median value. The orange solid lines indicate the median values from the results of the fit performed in Coughlin et al. (2018) using the same data set and kilonova model but a different surrogate construction and prior. The orange dashed lines indicate the ranges presented in Coughlin et al. (2018).

 $v_{ej} \le 0.3c$, and $-9 \le \log_{10}(\chi) \le -1$. The results for our fit are shown in Figs A1 and A2. We find relatively good agreement between our inferred parameters and those presented in Coughlin et al. (2018): all of their parameters fall within 1σ of our parameters. Some of the disagreement in fit is likely explained by the slight differences in the data that our surrogate was trained on and the prior ranges. We use both the systematic parameter survey and the narrower survey and combine them for our training set to construct the surrogate for the *Kasen* model, so our surrogate model uses more data than the surrogate model used in Coughlin et al. (2018). Additionally, Coughlin et al. (2018) use a wider prior that extends outside the ranges of the kilonova models published in Kasen et al. (2017).

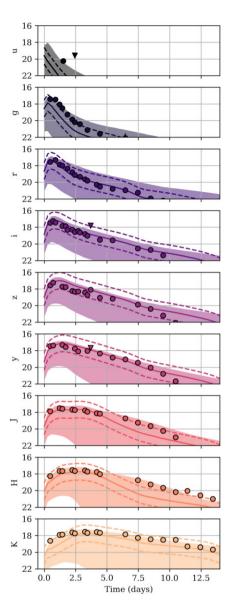


Figure A2. Comparison of the observed light curve for AT2017gfo (points) and the prediction from inferred parameters using the cVAE surrogate model (solid lines) using the *Kasen* model. The shaded bands represent the 90 per cent confidence interval of light curves constructed from the posterior samples. The dashed lines represent the 1 magnitude tolerance often used to represent modelling error of kilonova light curves.

This paper has been typeset from a TEX/LATEX file prepared by the author.