

This is the accepted manuscript made available via CHORUS. The article has been published as:

Measuring the neutron star equation of state with gravitational waves: The first forty binary neutron star merger observations

Francisco Hernandez Vivanco, Rory Smith, Eric Thrane, Paul D. Lasky, Colm Talbot, and Vivien Raymond

Phys. Rev. D **100**, 103009 — Published 13 November 2019

DOI: [10.1103/PhysRevD.100.103009](https://doi.org/10.1103/PhysRevD.100.103009)

Measuring the neutron star equation of state with gravitational waves: the first forty binary neutron star mergers

Francisco Hernandez Vivanco,^{1,2,*} Rory Smith,^{1,2,†} Eric Thrane,^{1,2}
Paul D. Lasky,^{1,2} Colm Talbot,^{1,2} and Vivien Raymond³

¹*School of Physics and Astronomy, Monash University, Vic 3800, Australia*

²*OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia*

³*School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, United Kingdom*

Gravitational waves from binary neutron star coalescences contain rich information about matter at supranuclear densities encoded by the neutron star equation of state. We can measure the equation of state by analyzing the tidal interactions between neutron stars, which is quantified by the tidal deformability. Multiple merger events are required to probe the equation of state over a range of neutron star masses. The more events included in the analysis, the stronger the constraints on the equation of state. In this paper, we build on previous work to explore the constraints that LIGO and Virgo are likely to place on the neutron star equation of state by combining the first forty binary neutron star detections, a milestone we project to be reached during the first year of accumulated design-sensitivity data. We carry out Bayesian inference on a realistic mock dataset of binaries to obtain posterior distributions for neutron star tidal parameters. In order to combine posterior samples from multiple observations, we employ a random forest regressor, which allows us to efficiently interpolate the likelihood distribution. Assuming a merger rate of $1540 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and a LIGO-Virgo detector network operating for one year at the sensitivity of the third-observation run, plus an additional eight months of design sensitivity, we find that the radius of a $1.4 M_{\odot}$ neutron star can be constrained to better than $\sim 10\%$ at 90% confidence. The pressure at twice the nuclear saturation density can be constrained to better than $\sim 45\%$ at 90% confidence. We show that these results are robust to our choice of parametrization for the neutron star equation of state.

I. INTRODUCTION

With the improving sensitivity of LIGO [1, 2] and Virgo [3], we expect to increase the number of binary neutron star detections using gravitational waves [4]. Binary neutron star coalescences provide rich information about the neutron star equation of state, which can be constrained by measuring the neutron star’s tidal interactions. These tidal interactions are determined by the neutron star’s dimensionless tidal deformability Λ , which determines how the neutron star’s quadrupole moment changes in response to the companion’s tidal field (e.g. [5, 6]). The gravitational-wave detection of a binary neutron star coalescence, GW170817 [7], placed the first constraints on Λ , ruling out some of the stiffest proposed equations of state [8].

In order to obtain a better measurement of the neutron star equation of state, it is necessary to combine information from an ensemble of binary neutron stars. The first fully Bayesian study that combined information to constrain the equation of state using tidal interactions was proposed by Del Pozzo et al. [9]. They simulated an astrophysically distributed gravitational-wave dataset and showed that a few tens of observations were enough to constrain the tidal deformability $\lambda = \Lambda/G(Gm/c^2)^5$ to about 10% at a reference mass of $1.4 M_{\odot}$. The method in Ref. [9] fits the tidal deformability λ by Taylor expand-

ing $\lambda(m)$ around a fiducial mass of $m = 1.4 M_{\odot}$. Then, Lackey and Wade [10] showed that λ can be measured at 10-50% in the range between $1-2 M_{\odot}$, where most of the information comes from the loudest ~ 5 events. Their method used Markov chain Monte Carlo simulations to fit a model that follows the piecewise polytrope parametrization of the equation of state. In order to reduce the computational cost of their study, they used Fisher matrix approximations in some of the events contained in their mock data study. Agathos et al. [11] improved on [10] by including waveforms with tidal effects up to 1 post-Newtonian (1PN) order, neutron-star spins, and termination of the waveform at the contact frequency. While the results from Agathos et al. [11] are broadly consistent with the findings of [10], Agathos et al. find that a poorly chosen neutron-star mass prior can bias inferences about the equation of state.

In this paper, we study how well we can constrain the equation of state with gravitational-wave observations using the first forty expected binary neutron star detections from the LIGO-Virgo network. We perform a realistic mock data study, which introduces a few innovations: (1) we use a reduced order quadrature model waveform based on the IMRPHENOMD_NRTIDAL approximant [12–14] which significantly reduces the cost of parameter estimation; (2) we run full parameter estimation on all the events in our dataset that contain significant information about the equation of state and (3) we introduce a new interpolation method for combining events, yielding an interpolation error of $\approx 0.3\%$.

The problem of combining equation of state constraints from multiple events is an example of an inter-

*Electronic address: francisco.hernandezvivanco@monash.edu

†Electronic address: rory.smith@monash.edu

esting general class of problems in which (1) the prior for an individual measurement is conditional on some hyper-parameter; (2) the conditional prior spans some d -dimensional subspace with fewer dimensions than the full D -dimensional parameter space and (3) the equation of state curves are deterministic, i.e., they live exactly on a curve rather than have a probability defined over the whole parameter space. As we discuss in greater detail below, standard techniques from hierarchical modelling fail when applied to such problems. In this paper, we discuss the subtleties associated with combining events. We document the challenges arising from this unique class of problem and propose a new solution, which employs a random forest regressor to interpolate the likelihood distribution.

This paper is organized as follows. In Section II, we describe the difficulties of combining a population of gravitational-wave measurements to constrain the equation of state. In order to solve this problem, we describe a new method that is based on interpolating the likelihood distribution using a random forest regressor. In Section III, we present constraints that the LIGO-Virgo detector network will place on the equation of state after the first forty gravitational-wave observations. We discuss our results in Section IV and conclude in Section V.

II. METHOD

A. Equation of state parametrization

Gravitational-wave observations can be used to constrain the equation of state by measuring the tidal deformability of neutron stars. If we can measure these tidal interactions, we can place constraints on relations that model the equation of state such as pressure p and density ρ (see e.g. [15]). While we cannot measure pressure and density directly, we can map $p(\rho)$ onto the measurable astrophysical parameters $\Lambda(m)$ using the Tolman-Oppenheimer-Volkoff equations and the relationship between the tidal deformability, second Love number, and radius [16, 17].

Read et al. [18] showed that $p(\rho)$ can be modeled assuming a piecewise polytropic relation between pressure and density, where each polytrope is defined between three different sections, $i = 1, 2, 3$. Each polytrope is modeled by

$$p = K_i \rho^{\Gamma_i}. \quad (1)$$

Here, K_i is a constant and Γ_i is the adiabatic index. The density transitions occur at $\rho_1 = 10^{14.7} \text{ g cm}^{-3}$ and $\rho_2 = 10^{15} \text{ g cm}^{-3}$ [5, 17–19]. Demanding that $p(\rho)$ is continuous implies the equation of state can be fully determined with only four numbers: Γ_i for $i = 1, 2, 3$ and p_1 , the value of pressure at ρ_1 .

The mapping of $p(\rho)$ onto $\Lambda(m)$ allows us to cast the neutron star equation of state measurement as a hierarchical model, where the neutron star parameters Λ and

Hyper-parameter	Prior	Minimum	Maximum
$\log(p_1/[\text{dyne/cm}^2])$	Uniform	33.6	34.8
Γ_1	Uniform	2	4.5
Γ_2	Uniform	1.1	4.5
Γ_3	Uniform	1.1	4.5

TABLE I: Hyper-prior distributions used in our study. We model the equation of state with three polytropes and assume that we know the equation of state below the pressure p_1 . The hyper-priors are chosen based on the discussions from Refs. [10, 20] and are set to cover a wide range of equations of state.

m are conditional on hyper-parameters Υ , which describe the shape of the piecewise polytrope.

In this study, we follow the piecewise polytrope model [18] to sample over the equation of state hyper-parameters. We choose the hyper-priors following the discussions from [10, 20]. The hyper-priors are shown in Tab. I. Additionally, when we sample over the piecewise polytrope hyper-parameters, we impose three astrophysically-motivated constraints:

1. The maximum speed of sound cannot exceed the speed of light. However, since the piecewise polytrope only fits the equations of state to an accuracy of $\sim 4\%$, we set the causality constraint only when the speed of sound exceeds 1.1 times the speed of light.
2. We only allow equations of state with maximum masses above $m = 1.97 M_\odot$, which is consistent with the maximum TOV mass observed from pulsar PSR J0348+0432 [21] and PSR J0740+6620 [22].
3. We restrict the tidal deformability to $\Lambda \leq 5000$, conservatively consistent with the gravitational-wave measurement GW170817 [7].

B. Formalism

Our objective is to calculate the posterior distribution on the hyper-parameters $\Upsilon = \{\Gamma_1, \Gamma_2, \Gamma_3, p_1\}$ that describe the equation of state. The posterior is given by Bayes theorem,

$$p(\Upsilon|\vec{d}) = \frac{\mathcal{L}(\vec{d}|\Upsilon)\pi(\Upsilon)}{\mathcal{Z}_\Upsilon}. \quad (2)$$

Here, we introduce the likelihood $\mathcal{L}(\vec{d}|\Upsilon)$, the hyper-evidence $\mathcal{Z}_\Upsilon$ and the hyper-prior $\pi(\Upsilon)$. The hyper-prior $\pi(\Upsilon)$, in our case, describes our prior belief about the hyper-parameters Υ that model the equation of state.

We can think of this problem as follows. We have some total likelihood function, which is the product of N




FIG. 1: Example of posterior samples from parameter estimation of two binary neutron star observations. Each gravitational-wave measurement results in posterior samples on the $\Lambda(m)$ curve, represented by blue and orange dots. The equation of state is defined by an infinitesimally thin curve, represented by the black solid curve. None of the posterior samples fall exactly on the black curve. This is the reason why the usual hierarchical Bayesian inference is not directly applicable to this problem.

single-event likelihoods

$$\mathcal{L}_{\text{tot}}(\vec{d}|\Upsilon) = \prod_i^N \int d\theta_i \mathcal{L}(d_i|\theta_i) \pi(\theta_i|\Upsilon). \quad (3)$$

Here, θ_i are the waveform parameters for the i th event, $\mathcal{L}(d_i|\theta_i)$ is the likelihood function assuming the signal hypothesis and $\pi(\theta_i|\Upsilon)$ is the conditional prior for parameters θ_i given hyper-parameters Υ . In our specific problem, θ_i refers to the 17 binary neutron star parameters and the hyper-parameters in Υ depend on the parameters that determine the equation of state, which we denote $\omega = (m, \Lambda)$. We introduce κ to denote the parameters in θ that are not in ω . Equation 3 is challenging to compute. We would normally analyze each event individually and then “recycle” the posterior samples (see e.g. [23]):

$$\mathcal{L}_{\text{tot}}(\vec{d}|\Upsilon) = \prod_i^N \frac{\mathcal{Z}_{\mathcal{O}}^i}{n_i} \sum_k^{n_i} \frac{\pi(\theta_i^k|\Upsilon)}{\pi(\theta_i^k|\mathcal{O})} \quad (4)$$

$$= \prod_i^N \frac{\mathcal{Z}_{\mathcal{O}}^i}{n_i} \sum_k^{n_i} \frac{\pi(\omega_i^k|\Upsilon)}{\pi(\omega_i^k|\mathcal{O})}. \quad (5)$$

Here, we sum over n_i posterior samples, where $\pi(\theta_i^k|\mathcal{O})$ is the fiducial prior used for the initial analysis and $\mathcal{Z}_{\mathcal{O}}$ is the corresponding signal evidence. Since $\pi(\theta_i^k|\omega)$ is described by a curve in the (Λ, m) plane, no posterior sample will fall exactly on this curve (see Fig. 1), and so the recycling procedure fails because $\pi(\theta_i^k|\Upsilon)$ always evaluates to zero. In the parlance of importance sampling, our proposal distribution fails to provide samples for our target distribution.

In order to solve this problem, we devise a method to interpolate the likelihood marginalized over nuisance parameters. The method is as follows.

1. Run parameter estimation on each event using the fiducial prior $\pi(\theta|\mathcal{O})$. The priors used in our injection set are shown in Tab. II. For simplicity, we consider non-spinning binary neutron stars.
2. The fiducial parameter estimation yields samples from the posterior distribution in ω . The density of these samples is proportional to the height of the marginalized likelihood function $\mathcal{L}_{\kappa}(d|\omega)$. Using the posterior sample density as a rough guide, we chose “interpolation points” $\{\omega_j\}$ for which we calculate the κ -marginalized likelihood in the next step:

$$\mathcal{L}_{\kappa}(d|\omega) \equiv \int d\kappa \mathcal{L}(d|\omega, \kappa) \pi(\kappa). \quad (6)$$

The values of $\{\omega_j\}$ are selected to give us the best possible representation of the marginalized likelihood so we can interpolate between different values of $\{\omega_j\}$. Additionally, we draw random samples from the prior to better sample the tails of the distribution.

3. We launch a new inference job for each interpolation point $\{\omega_j\}$. These inference jobs are identical to the original inference job except we employ a delta-function prior on ω so that these parameters are fixed to the values of the interpolation point. All of these inference jobs can be run in parallel so this step takes less time than the previous step, though it requires more compute cores. Since the metric perturbation depends only on the intrinsic parameters, it can be calculated just once at the beginning of the job and cached for later use. Calculating the strain given the metric perturbation is an extremely fast calculation, enabling these new jobs to run comparatively quickly. The Bayesian evidence given by each job represents a marginalized likelihood $\mathcal{L}_{\kappa}(d|\omega)$.
4. We create a function which interpolates between interpolation points $\mathcal{L}_{\kappa}^{\text{int}}(d|\omega)$ with a random forest regressor [24]. Once the random forest is trained, we can give it a new value of $\omega \notin \{\omega_j\}$, and the random forest provides an estimate of the marginalized likelihood $\mathcal{L}_{\kappa}(d|\omega)$.
5. Using our new interpolated marginal likelihood, the total likelihood for N events is

$$\mathcal{L}_{\text{tot}}(d|\Upsilon) = \prod_i^N \int d\omega_i \mathcal{L}_{\kappa}^{\text{int}}(d_i|\omega_i) \pi(\omega_i|\Upsilon). \quad (7)$$

As we change Υ , we obtain different curves in ω , which do not pass through any of our original

Parameter	Unit	Prior	Minimum	Maximum
$\mathcal{M}$	$M_\odot$	Uniform	1.05	1.45
q	-	Uniform	0.7	1
RA	rad.	Uniform	0	2π
DEC	rad.	Cos	$-\pi/2$	$\pi/2$
$\cos(\theta_{jn})$	-	Uniform	-1	1
ψ	rad.	Uniform	0	π
ϕ	rad.	Uniform	0	2π
d_L	Mpc	Comoving	10	400

TABLE II: Prior distributions used in our injection set. The intrinsic parameters are the chirp mass $\mathcal{M}$, and mass ratio q . The dimensionless tidal deformabilities Λ are determined using the SLy equation of state. The extrinsic parameters are the right ascension (RA) and declination (DEC), inclination angle θ_{jn} , phase of coalescence ϕ , polarization angle ψ , and luminosity distance d_L . The comoving prior in the luminosity distance means that we assume a uniform prior in comoving volume. We assume non-spinning binaries in all the injections.

samples, but we are able to evaluate the hyper-likelihood anyway because we are able to interpolate throughout ω . In practice, the integral in Eq. 7 is carried out on a discrete grid so that the integral is replaced by a sum over mass bins.

$$\mathcal{L}_{\text{tot}}(d|\Upsilon) = \prod_i^N \int d\omega_i \mathcal{L}_\kappa^{\text{int}}(d_i|\omega_i) \pi(\omega_i|\Upsilon) \quad (8)$$

$$\begin{aligned}
\mathcal{L}_{\text{tot}}(d|\Upsilon) &= \prod_i^N \int d\omega_i \mathcal{L}_\kappa^{\text{int}}(d_i|\omega_i) \pi(m_1^i, m_2^i) \\
&\quad \delta(\Lambda_1 - \Lambda_1'(m_1|\Upsilon)) \delta(\Lambda_2 - \Lambda_2'(m_2|\Upsilon)) \\
&= \prod_i^N \int dm_1^i \int dm_2^i \pi(m_1^i, m_2^i) \\
&\quad \mathcal{L}_\kappa^{\text{int}}\left(d_i \middle| m_1^i, m_2^i, \Lambda_1^i(m_1^i), \Lambda_2^i(m_2^i)\right) \\
&= \prod_i^N \sum_{\alpha, \beta} \pi(m_1^{i, \alpha}, m_2^{i, \beta}) \mathcal{L}_\kappa^{\text{int}}\left(d_i \middle| m_1^{i, \alpha}, m_2^{i, \beta}, \right. \\
&\quad \left. \Lambda_1^i(m_1^{i, \alpha}|\Upsilon), \Lambda_2^i(m_2^{i, \beta}|\Upsilon)\right) \Delta m^2. \quad (9)
\end{aligned}$$

Here, the sum over α, β run over 200 primary and secondary mass bins with width Δm . For each bin, Λ is determined by the equation of state parameters Υ .

C. Interpolation with random forest

In order to interpolate between samples, we employ the random forest regressor from the PYTHON package SCIKIT-LEARN [25]. Broadly speaking, the random forest algorithm combines information from different decision




FIG. 2: Posterior distributions of the intrinsic parameters $\omega = (\mathcal{M}, q, \Lambda_1, \Lambda_2)$. The orange posteriors are obtained by sampling the interpolated likelihood distribution. Our results are compared to the original posteriors shown in blue. This plot indicates that an interpolation error of $\sim 0.3\%$ reproduces the original posteriors.

trees. A decision tree can be understood as a tree-like model of branches where all outcomes are made based on conditional control statements. The results of random forest regressors are obtained by averaging over different decision trees [24, 26].

The evaluation time and accuracy of the interpolation depends on the number of decision trees and the number of interpolation points. We use 9×10^3 interpolation points, where 85% of our dataset is used for training and 15% is used for testing. We find that with 50 decision trees, the interpolation error is $\sim 0.3\%$ and the interpolated-likelihood evaluation time for one value of $\omega = (\mathcal{M}, q, \Lambda_1, \Lambda_2)$ is ≈ 0.05 ms. Increasing the number of random trees may increase the accuracy, but this level of accuracy is sufficient for our present purposes and the interpolated likelihood evaluation time increases proportionally to the number of decision trees.

In order to vet the interpolation, we sample the new interpolated-likelihood to obtain posterior distributions of the intrinsic parameters. The posterior distributions from the exact likelihood and the interpolated likelihood are shown in Fig. 2. The interpolated likelihood distributions shown in Fig. 2 are consistent with the exact likelihood, which indicates that the random forest interpolation is effective.

There are other machine-learning algorithms that could be as effective as random forests, for example, neural network regressors [27, 28]. While these algorithms have been used in recent studies, exploring the advantages/disadvantages of these algorithms is left for future study. All our simulations are done within the Bayesian inference library BILBY [29] and the LIGO algorithm library LALSUITE [30]. We use the DYNESTY nested-sampler [31].

III. MOCK DATA STUDY

We demonstrate our method with a mock data study. We assume a detector network consisting of the two Advanced LIGO detectors and Advanced Virgo. We assume that LIGO and Virgo operate at the expected sensitivity of the third observation run (O3) for one year, plus eight months of design sensitivity. We assume that the binary neutron star range of LIGO Hanford and LIGO Livingston at design sensitivity is 190 Mpc, while the range of Virgo at design sensitivity is 140 Mpc. During O3, we additionally assume that the binary neutron star range of LIGO Hanford, LIGO Livingston and Virgo are 100 Mpc, 130 Mpc and 50 Mpc respectively [32].

We assume a merger rate of $1540 \text{ Gpc}^{-3} \text{ yr}^{-1}$, which is consistent with a realistic binary neutron star merger rate [7]. Based on this rate, we expect to detect $N = 40_{-11}^{+10}$ binary neutron star mergers (90% credible interval) with a network matched filter signal-to-noise ratio exceeding the usual threshold of 12. In Tab. II we show the prior distributions of our injection set. This results in binaries with individual masses between $1 M_{\odot}$ and $2 M_{\odot}$. The tidal deformabilities are determined using the SLy equation of state [33].

Due to the fact that binary neutron stars can be in-band for long periods of time—for example GW170817 was in band for ≈ 2 minutes [7]—we use a reduced-order-model (ROM) gravitational waveform to accelerate Bayesian parameter estimation [34]. A reduced-order model removes redundant waveform calculations by reduced-order quadrature (ROQ) integration. In this study, we use an ROQ implementation of the IMRPHENOMD_NRTIDAL waveform approximant [12–14], which is a spin-aligned waveform that includes tidal interactions and which models the inspiral, merger and ring-down. This allows us to run full parameter estimation for each gravitational-wave measurement in ≈ 4 hours. This is an improvement compared to Ref. [10], where they used the TAYLORF2, TAYLORT1 and TAYLORT4 approximants without an ROQ implementation of these waveforms.

We investigate the result from Lackey and Wade [10], which states that most of the equation of state information comes from the loudest $\sim 12\%$ of events. We calculate the Bayes factor comparing the non-black hole evidence $\mathcal{Z}_{\Lambda \neq 0}$ to the black hole evidence $\mathcal{Z}_{\Lambda = 0}$,

$$\text{BF} = \frac{\mathcal{Z}_{\Lambda \neq 0}}{\mathcal{Z}_{\Lambda = 0}}. \quad (10)$$

Then, we evaluate (10) for events with a range of different signal-to-noise ratios. We rank events from loud to quiet, where “event rank”=1 is defined as the loudest event. Our results are shown in Fig. 3, where we plot the cumulative $\ln \text{BF}$ as a function of “event rank”. We see that the $\ln \text{BF}$ plateaus at $\text{SNR} \approx 20$. This means that events with $\text{SNR} < 20$ do not add significant information about the equation of state, compared to the loudest events in our dataset. We find that most of the




FIG. 3: Cumulative $\ln(\text{BF})$ as a function of event rank. We rank events from loud to quiet, where event rank=1 is defined as the loudest event. The $\ln \text{BF}$ is calculated between the hypothesis that $\Lambda \neq 0$ and $\Lambda = 0$. The event rank is shown in the lower horizontal axis and its corresponding SNR is shown in the upper horizontal axis. We find that after $\text{SNR} \approx 20$, the $\ln \text{BF}$ starts to plateau. This means that past $\text{SNR} \approx 20$, signals do not significantly add substantial information about the equation of state. Our analysis focuses only on signals where $\text{SNR} > 20$.

information comes from the loudest $\sim 15\%$ of events, which is consistent with Ref. [10]. The remainder of our study focuses only on analyzing the loudest eight events.

In Fig. 4 we show the joint posterior for (R, m) and (p, ρ) obtained using the loudest eight events in our mock data study. Assuming the piecewise polytrope parametrization of the equation of state, we find that the radius of a $1.4 M_{\odot}$ neutron star is constrained to better than 10% at 90% confidence corresponding to a precision of $\pm 1 \text{ km}$. At the same time, the pressure at twice the nuclear saturation density can be constrained to better than 45% at 90% confidence.

In order to explore how the constraint on the neutron star equation of state depends on the equation of state parametrization, we analyze the same dataset with the spectral decomposition parametrization [37]. Using the same priors outlined in Ref. [20], we find that the radius of a $1.4 M_{\odot}$ neutron star can be constrained to better than $\sim 10\%$, consistent with the piecewise polytrope parametrization.

IV. DISCUSSION

We show that after an observation time of one year of LIGO/Virgo operating at O3 sensitivity, plus eight months of design sensitivity, we can constrain the neutron star radius to $\sim 10\%$ at 90% confidence. Our results are consistent with Ref. [10], where they showed that the neutron star radius can be measured to $\sim 10\%$ after one year of LIGO/Virgo design sensitivity. Our results are consistent with Ref. [10] because the number of detections used in their study correspond roughly to the




FIG. 4: Top panel: 90% credible interval of the posterior distribution of radius as a function of mass. The shaded grey and green posteriors are obtained from analyzing respectively the loudest four and eight events from our dataset with the piecewise polytrope parametrization. The dark-blue dashed posteriors are obtained from analyzing the loudest eight events of our dataset with the spectral decomposition parametrization. Bottom panel: 90% credible interval of the pressure as a function of radius and the ratio between the pressure and the pressure given by the SLy equation of state p_{SLy} . The gray vertical lines represent the dividing densities of the piecewise polytrope. We assume that binary neutron stars follow the SLy equation of state and the merger rate is fixed to $1540 \text{ Gpc}^{-3} \text{ yr}^{-1}$.

number of detections used in this study.

The work in [10] showed that the vast majority of information comes from the loudest five events. We confirm

this statement by calculating the Bayes factor between $\mathcal{Z}_{\Lambda \neq 0} / \mathcal{Z}_{\Lambda = 0}$. From Fig 3, we can see that the cumulative Bayes factor significantly grows for loud events, but plateaus for events with $\text{SNR} \approx 20$, which confirms the statement that the vast majority of information comes from the loudest few events. Our Bayes factor estimation is possible because we use a nested sampler algorithm, which returns a Bayesian evidence for every parameter estimation run.

In this study, we use an ROQ implementation of the IMRPHEMOMD_NRTIDAL waveform approximant [12–14], which includes inspiral, merger and ringdown. This is in contrast to the waveforms used in Ref. [10], which used inspiral-only waveforms containing post-Newtonian (PN) terms up to 3.5PN, in the point-particle terms. Their study shows that not including 4PN terms adds systematic biases which can sometimes be larger than the 95% credible interval of their mass and pressure constraints. In this study, we use the IMRPHEMOMD_NRTIDAL approximant, which is calibrated with effective-one-body and numerical relativity waveforms.

Agathos et al. [11] showed that a poorly chosen prior for neutron-star mass can bias inferences about the neutron-star equation of state. While we do not account for prior mismatch in our study, this effect should be taken into account when data from many binary neutron star detections are combined. The long-term solution, we believe, is to employ a population model, which can be fit using hierarchical modeling [23]. Such models are already used to infer the shape of the binary black hole mass spectrum [35]. In this way, the neutron-star mass prior for each event is informed by the distribution inferred by the other events.

In this study, we explore in greater detail a different method to interpolate the likelihood distribution using a random forest regressor which leads to an interpolation error of $\sim 0.3\%$. Our method is different from the Gaussian kernel density estimation algorithm used in [10]. The interpolation method is important because a poor interpolation accuracy can bias the hyper-posterior $p(\Upsilon|d)$.

One factor that could change our results is the choice of priors for the piecewise polytrope parametrization. Particularly, [20] shows that different low-density starting points of the piecewise polytrope parametrization result in different mass-radius constraints. In this study, we followed the parametrization used in Refs. [10, 18].

In our analysis, we require that the equation of state allow for a maximum mass in excess of $m = 1.97 M_\odot$ in order to be consistent with the pulsars PSR J0348+0432 and PSR J0740+6620. However, Miller et al. [36] argue that a smooth prior on maximum mass (informed by radio observations) is preferable to a hard cut-off. In future work, it will be useful to revisit prior constraints on the maximum neutron star mass.

We compare our results using two different parametrizations of the equation of state: the piecewise polytrope and the spectral decomposition parametrization. Although these parametrizations are consistent

between $\rho \approx 10^{14.2} \text{ g cm}^{-3}$ and $\rho \approx 10^{15} \text{ g cm}^{-3}$, we note differences at densities near $10^{14.1} \text{ g cm}^{-3}$ and $\gtrsim 10^{15} \text{ g cm}^{-3}$. The differences near $10^{14.1} \text{ g cm}^{-3}$ are explained by the choice of reference pressure p_0 , where we use the value of $p_0 = 5.36 \times 10^{32} \text{ dyne cm}^{-2}$ to reflect that the equation of state is well constrained below the nuclear density [20].

We find that the differences above $\approx 10^{15} \text{ g cm}^{-3}$ are explained by the fact that this region of parameter space is mostly dominated by the prior on causality and high mass neutron star observations. Since the prior that we use in the piecewise polytrope is narrower above $\approx 10^{15} \text{ g cm}^{-3}$, we find that the piecewise polytrope is more tightly constrained above this density. In particular, we find that the pressure at six times the nuclear saturation density is constrained to $\sim 30\%$ and $\sim 50\%$ assuming the piecewise polytrope and spectral decomposition parametrization, respectively.

Another algorithm that can be used to infer equation of state hyper-parameters is RIFT [38]. This method interpolates the likelihood distribution over a grid of source parameters using a Gaussian process regressor. In order to recover the parameters of astrophysical sources, the algorithm obtains an interpolated continuous representation of the likelihood, which can also be used to overcome the “stacking problem” that we address here in our study.

Finally, in some circumstances, it is possible that we will observe bias in our hyper-posterior $p(\Upsilon|d)$ due to small errors in our interpolation. If this happens, we can fix the problem with an iterative procedure. We can carry out steps 1-5 of our interpolation method described in Sec. IIB in order to obtain a posterior predictive distribution in ω . We can use the posterior predictive distribution to choose additional interpolation points ω_j . This amounts to an adaptive mesh refinement where the interpolated likelihood $\mathcal{L}_k$ is calculated with greater density in ω depending on the data, helping to control systematic error as we combine more events. This ability to adaptively refine the interpolation is a potential advantage

over Gaussian kernel density estimation.

V. CONCLUSION

We present a new method to combine information from an ensemble of gravitational-wave observations to constrain the neutron-star equation of state. Our method interpolates the likelihood distribution with an error of 0.3% using a random forest regressor. Our interpolation algorithm is an alternative to the Gaussian kernel density estimation method used in [10]. We demonstrate how LIGO and Virgo are likely to constrain the radius and pressure of binary neutron stars with the first forty detections. We show that the radius of a $1.4 M_\odot$ neutron star can be constrained to $\sim 10\%$ at 90% confidence and the density at twice the nuclear saturation density can be constrained to $\sim 45\%$ at 90% confidence. Finally, we combine gravitational-wave measurements by running full parameter estimation on events with $\text{SNR} \geq 20$ using an ROQ implementation of the IMRPHENOMD_NRTIDAL approximant, which significantly reduces the cost of parameter estimation.

Acknowledgments

This work is supported through Australian Research Council Grant No. CE170100004, No. FT150100281, No. FT160100112, and No. DP180103155. F.H.V. is supported through the Monash Graduate Scholarship (MGS). The results presented in this manuscript were calculated using the computer clusters at California Institute of Technology and Swinburne University of Technology (OzSTAR). The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This is LIGO Document No. P1900262.

-
- [1] J. Aasi et al., *Classical and Quantum Gravity* **32**, 115012 (2015), URL <https://doi.org/10.1088/2F0264-9381/2F32/2F11/2F115012>.
 - [2] B. P. Abbott et al. (The LIGO Scientific Collaboration and the Virgo Collaboration), *Classical and Quantum Gravity* **32**, 074001 (2015), URL <http://stacks.iop.org/0264-9381/32/i=7/a=074001>.
 - [3] F. Acernese et al., *Classical and Quantum Gravity* **32**, 024001 (2014), URL <https://doi.org/10.1088/2F0264-9381/2F32/2F2/2F024001>.
 - [4] B. P. Abbott et al., *Living Reviews in Relativity* **21**, 3 (2018), ISSN 1433-8351, URL <https://doi.org/10.1007/s41114-018-0012-9>.
 - [5] E. E. Flanagan and T. Hinderer, *Phys. Rev. D* **77**, 021502 (2008), URL <https://link.aps.org/doi/10.1103/PhysRevD.77.021502>.
 - [6] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read, *Phys. Rev. D* **81**, 123016 (2010), URL <https://link.aps.org/doi/10.1103/PhysRevD.81.123016>.
 - [7] B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **119**, 161101 (2017), URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101>.
 - [8] B. P. Abbott et al. (The LIGO Scientific Collaboration and the Virgo Collaboration), *Phys. Rev. Lett.* **121**, 161101 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevLett.121.161101>.
 - [9] W. Del Pozzo, T. G. F. Li, M. Agathos, C. Van Den Broeck, and S. Vitale, *Phys. Rev. Lett.* **111**, 071101 (2013), URL <https://link.aps.org/doi/10.1103/PhysRevLett.111.071101>.
 - [10] B. D. Lackey and L. Wade, *Phys. Rev. D* **91**,

- 043002 (2015), URL <https://link.aps.org/doi/10.1103/PhysRevD.91.043002>.
- [11] M. Agathos, J. Meidam, W. Del Pozzo, T. G. F. Li, M. Tompitak, J. Veitch, S. Vitale, and C. Van Den Broeck, *Phys. Rev. D* **92**, 023012 (2015), URL <https://link.aps.org/doi/10.1103/PhysRevD.92.023012>.
- [12] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevD.93.044006>.
- [13] T. Dietrich, S. Bernuzzi, and W. Tichy, *Phys. Rev. D* **96**, 121501 (2017), URL <https://link.aps.org/doi/10.1103/PhysRevD.96.121501>.
- [14] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevD.93.044007>.
- [15] F. zel and P. Freire, *Annual Review of Astronomy and Astrophysics* **54**, 401 (2016), <https://doi.org/10.1146/annurev-astro-081915-023322>, URL <https://doi.org/10.1146/annurev-astro-081915-023322>.
- [16] L. Lindblom, *The Astrophysical Journal* **398** (1992).
- [17] T. Hinderer, *The Astrophysical Journal* **677**, 1216 (2008), URL <https://doi.org/10.1086%2F533487>.
- [18] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009), URL <https://link.aps.org/doi/10.1103/PhysRevD.79.124032>.
- [19] J. M. Lattimer and M. Prakash, *Phys. Rept.* **442**, 109 (2007), astro-ph/0612440.
- [20] M. F. Carney, L. E. Wade, and B. S. Irwin, *Phys. Rev. D* **98**, 063004 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevD.98.063004>.
- [21] J. Antoniadis, P. C. C. Freire, N. Wex, T. M. Tauris, R. S. Lynch, M. H. van Kerkwijk, M. Kramer, C. Bassa, V. S. Dhillon, T. Driebe, et al., *Science* **340** (2013), ISSN 0036-8075, URL <https://science.sciencemag.org/content/340/6131/1233232>.
- [22] H. T. Cromartie, E. Fonseca, S. M. Ransom, P. B. Demorest, Z. Arzoumanian, H. Blumer, P. R. Brook, M. E. DeCesar, T. Dolch, J. A. Ellis, et al., arXiv e-prints arXiv:1904.06759 (2019), 1904.06759.
- [23] E. Thrane and C. Talbot, *Publ. Astron. Soc. Austral.* **36**, 10 (2019), 1809.02293.
- [24] L. Breiman, *Machine Learning* **45**, 5 (2001), ISSN 1573-0565, URL <https://doi.org/10.1023/A:1010933404324>.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., *Journal of Machine Learning Research* **12**, 2825 (2011).
- [26] G. Biau and E. Scornet, *TEST* **25**, 197 (2016), ISSN 1863-8260, URL <https://doi.org/10.1007/s11749-016-0481-7>.
- [27] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al., *TensorFlow: Large-scale machine learning on heterogeneous systems* (2015), software available from tensorflow.org, URL <https://www.tensorflow.org/>.
- [28] Y. LeCun, Y. Bengio, and G. Hinton, *Nature* **521**, 436 (2015), ISSN 0028-0836.
- [29] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov, et al., *The Astrophysical Journal Supplement Series* **241**, 27 (2019), 1811.02042.
- [30] LIGO Scientific Collaboration, *LIGO Algorithm Library - LALSuite*, free software (GPL) (2018).
- [31] J. S. Speagle, arXiv e-prints arXiv:1904.02180 (2019), 1904.02180.
- [32] B. P. Abbott et al., *Living Reviews in Relativity* **21**, 3 (2018), ISSN 1433-8351, URL <https://doi.org/10.1007/s41114-018-0012-9>.
- [33] F. Douchin and P. Haensel, *Astron. Astrophys.* **380**, 151 (2001), astro-ph/0111092.
- [34] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürrer, V. Raymond, and P. Schmidt, *Phys. Rev. D* **94**, 044031 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevD.94.044031>.
- [35] B. P. Abbott et al., Accepted in *Astrophys. J. Lett.* (2018), arxiv/1811.12940.
- [36] M. C. Miller, C. Chirenti, and F. K. Lamb (2019), 1904.08907.
- [37] L. Lindblom, *Phys. Rev. D* **82**, 103011 (2010), URL <https://link.aps.org/doi/10.1103/PhysRevD.82.103011>.
- [38] J. Lange, R. O'Shaughnessy, and M. Rizzo (2018), 1805.10457.