



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

Normalizing Flows as an Avenue to Studying Overlapping Gravitational Wave Signals

Jurriaan Langendorff, Alex Kolmus, Justin Janquart, and Chris Van Den Broeck

Phys. Rev. Lett. **130**, 171402 — Published 26 April 2023

DOI: [10.1103/PhysRevLett.130.171402](https://doi.org/10.1103/PhysRevLett.130.171402)

Normalizing flows as an avenue to study overlapping gravitational wave signals

Jurriaan Langendorff^{1,*}, Alex Kolmus^{2,†}, Justin Janquart^{1,3}, and Chris Van Den Broeck^{1,3}

¹ *Institute for Gravitational and Subatomic Physics (GRASP), Department of Physics, Utrecht University, Princetonplein 1, 3584 CC Utrecht, The Netherlands*

² *Institute for Computing and Information Sciences (ICIS),*

Radboud University Nijmegen, Toernooiveld 212, 6525 EC Nijmegen, The Netherlands and

³ *Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands*

(Dated: April 3, 2023)

Due to its speed after training, machine learning is often envisaged as a solution to a manifold of the issues faced in gravitational-wave astronomy. Demonstrations have been given for various applications in gravitational-wave data analysis. In this work, we focus on a challenging problem faced by third-generation detectors: parameter inference for overlapping signals. Due to the high detection rate and increased duration of the signals, they will start to overlap, possibly making traditional parameter inference techniques difficult to use. Here, we show a proof-of-concept application of normalizing flows to perform parameter estimation on overlapped binary black hole systems.

INTRODUCTION

Over the last few years, the improved sensitivity of the LIGO [1] and Virgo [2] detectors has made the detection of gravitational waves (GWs) originating from compact binary coalescences (CBCs) more and more common, with over 90 detections reported after the third observation run [3]. Soon, the upgrade of the current detectors and the addition of KAGRA [4–8] and LIGO India [9] to the network of ground-based interferometers will lead to even more detections. In addition, the passage from second-generation (2G) to third-generation (3G) detectors (Einstein Telescope (ET) [10, 11] and Cosmic Explorer (CE) [12–14]) will lead to an important increase in the number of observed CBCs. These detectors are also projected to have a reduced lower frequency cut-off [15], leading to longer signal durations. Therefore, CBC signals will overlap in 3G detectors [16–20].

Analyzing one of the overlapping signals without accounting for the presence of the other can lead to biases in the recovered posteriors, especially when the merger times of the two events are close [17–21]. These could impact any direct science case for CBCs (e.g. tests of general relativity [22]), but also indirectly related ones such as the hunt for primordial black holes [23–28]. In Ref. [29], the authors demonstrate on two overlapped binary black holes (BBHs) how adapted Bayesian inference can help reduce the biases. In particular, they perform joint parameter estimation, where the two signals are analyzed jointly. While accounting for all the noise characteristics, their analysis also suffers from some instabilities, and further upgrades are needed for it to be entirely reliable. An issue also mentioned in this work is the computational time. With hundreds of thousands of CBC mergers expected in the 3G era [17], analyses taking several weeks are not a realistic alternative.

Even if traditional methods can be sped-up [30–33], or quantum computing [34] could potentially be used in the future, the development of frameworks capable of doing

complete analyses in short timescales is crucial for the development of 3G detectors. Therefore, in this work, we propose the first step in that direction, showing how overlapping BBHs can be analyzed with a normalizing flow (NF) approach [35–37].

MACHINE LEARNING FOR OVERLAPPING GRAVITATIONAL WAVES

The use of machine learning (ML) in GW data analysis has been growing over the last years, having a wide range of applications [38]. A subset of these methods fall under the umbrella of simulation-based inference [39], and are being developed to perform parameter estimation for CBCs [40–47]. Refs.[43–45] use NFs to get posterior distributions for BBH parameters, obtaining results close to those from traditional Bayesian methods. Our approach is somewhat similar to theirs, with some notable differences explained below.

Our approach uses *continuous conditional NFs* [48, 49] (CCNFs), a variant of NFs suited for probabilistic modeling and Bayesian inference. Due to the recursive and continuous nature of these models, their memory footprint can be quite small [50], allowing for extensive training on home-grade GPUs while retaining the ability to capture complex distributions.

NFs are a method in ML through which a neural network can learn the mapping from some simple base distribution $p_u(\mathbf{u})$ to a more complex final distribution $q(\boldsymbol{\theta})$. This is done through a series of invertible and differentiable transformations, summarized by a function $g(\boldsymbol{\theta})$. However, in our case, the final distribution we seek depends on the GW data to analyze. Therefore, we use *conditional NFs* [51], where the transformation functions are dependent on the data \mathbf{d} (hence, $g = g(\boldsymbol{\theta}, \mathbf{d})$). A major difference with [51] is that our base distributions are kept static. Thus our model $g(\boldsymbol{\theta}, \mathbf{d})$ is a trainable conditional bijective function transforming a simple 30-D Gaussian

into a 30-D complex distribution. The bijectivity allows us to express and sample $q(\boldsymbol{\theta}|\mathbf{d})$ in terms of $g(\boldsymbol{\theta}, \mathbf{d})$ and $p_u(\mathbf{u})$ via:

$$q(\boldsymbol{\theta}|\mathbf{d}) = |\det(\mathbf{J}_{g^{-1}}(\boldsymbol{\theta}, \mathbf{d}))| p_u(g^{-1}(\boldsymbol{\theta}, \mathbf{d})), \quad (1)$$

where $\det(\mathbf{J}_{g^{-1}}(\boldsymbol{\theta}, \mathbf{d}))$ is the determinant of the Jacobian $\mathbf{J}_{g^{-1}}(\boldsymbol{\theta}, \mathbf{d})$ of the transformation. We train the model by minimizing the forward KL-divergence, which is equivalent to maximum likelihood estimation [37, 52]. As noted by [45], $q(\boldsymbol{\theta}|\mathbf{d})$ should cover the actual (Bayesian) posterior $p(\boldsymbol{\theta}|\mathbf{d})$, and asymptotically approach it as training progresses due to the mode-covering nature of the forward KL divergence.

A distinctive choice of our method is the continuous nature of the flow, which is linked to the transformation function itself. Neural ordinary differential equations (neural ODEs) [50] are the foundation of continuous NFs; they are not represented by a stack of discrete layers but by a hypernetwork [53]. Hypernetworks can be understood as regular networks where ‘external’ inputs such as a time or depth variable smoothly change the output of the network for identical inputs. They can thus represent multiple transformations. In [50], hypernetworks are used to represent ODEs and are trained by using ODE-solvers and clever use of the adjoint sensitivity method. A continuous NF uses neural ODEs as its transformations.

We will now explain the training of a continuous flow. For clarity, we will use h to refer to a continuous transformation and g for a discrete one. If $\boldsymbol{\theta}(t)$ represents the samples from the distribution at a given time t , when going from t_1 to t_2 , the continuous NF obeys

$$\frac{d\boldsymbol{\theta}(t)}{dt} = h(t, \boldsymbol{\theta}(t)). \quad (2)$$

The change in likelihood associated with this ‘step’ differs slightly from Eq. (1) due to the continuous nature of the flow:

$$\log(p(\boldsymbol{\theta}(t_1))) = \log(p(\boldsymbol{\theta}(t_0))) - \int_{t_0}^{t_1} \text{Tr}[\mathbf{J}_g(\boldsymbol{\theta}(t))] dt. \quad (3)$$

Assuming a non-stiff ODE the integration can be performed rapidly with state-of-the-art ODE-solvers, MALI [54] in our case. In addition, we have to calculate a trace instead of a determinant, speeding up the computation which reduces the complexity, going from $\mathcal{O}(D^3)$ to at most $\mathcal{O}(D^2)$ with D being the dimensionality of posterior space, speeding-up the computation [55, 56]. Moreover, using continuous NFs removes the need to use coupling layers between transformations, instead, all parameter dimensions can be dependent on each other throughout the flow. Combining the continuous and conditional flows leads to CCNFs, where the conditional consists of the GW \mathbf{d} and the time t .

We also need a better data representation than the raw strain to train and analyze the data. Therefore, we

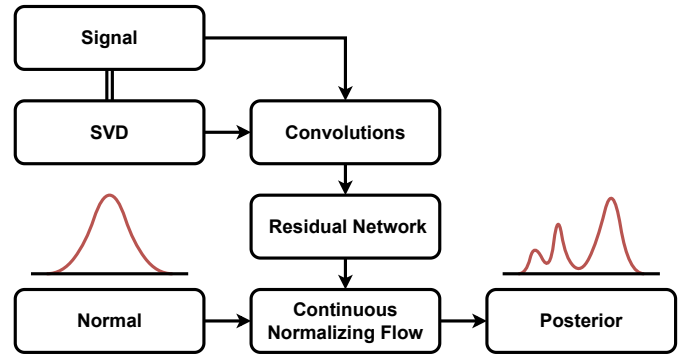


FIG. 1. Representation of our analysis framework. It is made of a pre-processing part where we build an SVD basis to filter the data, followed by a normalizing-flow-based neural network.

follow a similar approach as in [41–45], using a single value decomposition (SVD) [57] as summary statistics, reducing the dimension and the noise content of the data while retaining at least 99% of the original signal. Each of the 256 generated basis vectors is used as a kernel in 1D convolutions used as an initial layer in a ten-layer residual convolutional neural network (CNN), enabling one to capture the time variance of the signal. Therefore, we do not need to use a Gibbs sampler to estimate the time of the signal as done in [43–45], and can sample over time like any other variable. The CCNF itself is represented by two multi-layer perceptrons with 3 hidden layers of 512 units. Furthermore, we use a different representation for the angles. Instead of directly using their values, we project them onto a sphere for the sky location and onto a circle for the other angles. This makes for a better-posed domain for these angles, and plays on the strong interpolation capacities of the network, making the training step easier.

In the end, our framework combines data representation as a hybrid between SVD and CNN, followed by the CCNF network. A representation of our analysis framework is given in Fig. 1. Our entire framework is relatively small compared to the ones presented in [43], **both the residual network and CNF network**. Therefore, it can run on lower-end GPUs, but could also be limited in its capacity to model the problem.

DATA AND SETUP

To test our framework, we start with a simplified setup, considering a network made of the two LIGO detectors and the Virgo detector, at design sensitivity [2, 58], and with a lower sensitive frequency of 20Hz. We generate stationary Gaussian noise from their power spectral density (PSD) and inject two precessing BBH mergers using the IMRPhenomPv2 waveform [59]. Our data frames have an 8 seconds duration and are whitened af-

Parameter	Function
Chirp mass (\mathcal{M})	$\mathcal{U}(10, 100)M_{\odot}$
Mass ratio (q)	$\mathcal{U}(0.125, 1)$
Component masses ($m_{1,2}$)	Constrained in $[5, 100]M_{\odot}$
Luminosity distance (D_L)	Rescaled to follow SNR
SNR	$\mathcal{B}(10, 50)$
Coalescence time (t_c)	$\mathcal{U}(t_{\text{ref}} - 0.05, t_{\text{ref}} + 0.05)$
Spin Amplitudes ($a_{1,2}$)	$\mathcal{U}(0, 1)$
Spin tilt angles ($\theta_{1,2}$)	Uniform in sine
Spin vector azimuthal angle (ϕ_{jl})	$\mathcal{U}(0, 2\pi)$
Spin precession angle (ϕ_{12})	$\mathcal{U}(0, 2\pi)$
Inclination angle (θ_{jn})	Uniform in sine
Wave polarization (ψ)	$\mathcal{U}(0, \pi)$
Phase of coalescence (ϕ)	$\mathcal{U}(0, 2\pi)$
Right ascension (RA)	$\mathcal{U}(0, 2\pi)$
Declination (DEC)	Uniform in cosine

TABLE I. Summary of the parameters considered and the function used to generate the BBHs.

ter the signals are injected. The chirp mass ($\mathcal{M}_c = (M_1 + M_2)^{3/5}/(M_1 M_2)^{1/5}$) and mass ratio ($q = M_1/M_2$) are sampled from uniform distributions, between $10M_{\odot}$ and $100M_{\odot}$ and 0.125 and 1, respectively. The individual component masses are constrained between $5M_{\odot}$ and $100M_{\odot}$. During the data generation, the luminosity distance is kept fixed. It is then rescaled to result in a network signal-to-noise ratio value taken randomly between 10 and 50 from a beta distribution with a central value of 20. The coalescence time for the two events is set randomly around a time of reference, with $t_c \in [t_{\text{ref}} - 0.05, t_{\text{ref}} + 0.05]$ s, ensuring that the two BBH merge in the high bias regime [18]. The other parameters are drawn from their usual domain. Table I gives an overview of the parameters and the function from which they are sampled.

During the training, we continuously generate data by sampling the prior distributions for the events and making a new noise realization for each frame. The training is stopped when convergence is reached and before overfitting occurs. Our model was trained for about 12 days on a single *Nvidia GeForce GTX 1080*.

RESULTS

To demonstrate the method’s reliability, a P-P plot for the recovered parameters is shown in Fig. 2. It is constructed by sampling the posteriors of 1000 overlapped events¹ with parameters drawn from the distributions detailed in Table. I. Since the cumulative density aligns

¹ We refer the reader to Fig. 1 in Ref. [29] for an illustration of overlapping BBH signals.




FIG. 2. P-P plots for a subset of the recovered parameters for the two events in the data. The parameters shown are representative of all the BBH parameters for the two events. In both cases, the lines align along the diagonal, showing that our method can be trusted. The legend indicates which line corresponds to which parameters. The parameters for event 1 (resp. 2) are noted P_1 (resp. P_2), where P are the usual parameter symbols as presented in Table I. The values between the brackets are the KS test statistic.

along the diagonal, our network is reliable. Comparing this to the results given in [43] for single signals, there is a broadening of the shell around the diagonal, showing more variability in signal recovery, meaning our inference is less accurate than for single signals. Possible origins are the degenerate posteriors, increased complexity of the problem, and the reduced size of our network. This increased variability when going from single to joint parameter estimation has also been noted in Bayesian approaches [29].

While Bayesian methods have been developed in [29], they are not yet fully stable and take a long time to analyze a BBH system. Therefore, making a statistically significant study comparing the two approaches seems a bit premature at this stage. However, to have some sense of the performances of our network compared to traditional methods, we make 15 injections complying with our network’s setup and analyze them with the framework presented in [29]. Using these analyses, we can already identify some trends between the two pipelines. The first is that our ML pipeline typically has broader posteriors than the Bayesian approach. As mentioned in Ref. [29], the classical joint parameter estimation approach can sometimes get overconfident –see Ref. [29] for a discussion on the Bayesian algorithm–, where the recovered injected value lies outside of the 90% confidence interval. Our method is not confronted with this bottleneck as the broader posterior encapsulates the injected value. Fig. 3 illustrates the two representative situations: one

where the Bayesian approach finds the event correctly, and one where we see that our ML approach covers the injected values while it does not for the classical approach. Bias in the posterior, similar to the one noted in Ref. [29], can exist in our method and would not be seen because of the broad posteriors. However, because we are using the forward KL divergence, we expect the posteriors to have some support for the injected values. The origin of the larger posterior, which is not observed in the single parameter estimation machine learning-based methods, is probably due to the increased complexity of the problem combined with **the small residual and CNF network sizes**. One possible avenue is applying importance sampling after the normalizing flow [45, 60]. However, such methods can be tricky, and additional modifications to our network could be needed.

Finally, an important advantage of our method is its speed. After being trained, it can analyze two overlapping BBH signals in about a second, to compare with $\mathcal{O}(20\text{days})$ reported in [29]. While it is difficult to estimate the time gain for other CBC signals, we can expect the inference time after training not to be significantly larger than for BBHs. Since computational time is a crucial aspect of studies in the 3G era, ML approaches seem to be more suited to study realistic scenarios for these detectors.

CONCLUSIONS AND PERSPECTIVES

In this work, we have presented a proof-of-concept machine learning-based method to analyze overlapping BBH signals. We focused on a 2G detector scenario with the two LIGO, and the Virgo detectors at design sensitivity, with a lower frequency cutoff of 20Hz. Our approach is based on continuous normalizing flows.

While also using normalizing flows, as in [41–45], we bring extra modifications that seem to help in the inference task. We represent the data through a mixture of SVD and convolutions, enabling us to sample directly over the events’ arrival time, retaining the ability to access the likelihood of a sample. We also move to continuous conditional normalizing flows, reducing the computational cost of the method as we need to solve a trace instead of a determinant when going from one step to the other in the transformation. Finally, we also use a particular representation of the angles, projecting them onto circles (for the phase, the polarization, . . .) and spheres (for the sky location). We believe that these modifications make our network more flexible, enabling it to deal with overlapping signals even in a reduced form.

With this simplified setup, we have shown that our approach is reliable, with posteriors consistent with the injected values. Our method takes about one week to train on a single GPU. After that, it only takes about a second to analyze two overlapped BBHs. While, in reality,

other types of CBC mergers can happen, their inference after training should not be significantly longer than for BBHs. We also compared our machine learning method with classical Bayesian methods for overlapping signals. While our scheme leads to wider posteriors, it can correctly recover the injected values, even when the Bayesian approach gets overconfident and misses the injection. A possibility to correct for the widened posteriors is to use importance sampling.

Our method’s combined reliability and speed show that machine learning is a viable approach to analyzing CBC mergers in the 3G era. More interestingly, it would even be possible without needing to account for the development of more powerful computational means and could enable some science-case studies for ET and CE soon. For example, once trained for all possible BBH systems, it could help study the BBH mass function in the 3G era.

Still, one should note that extra improvements are needed before using our method in realistic 3G scenarios. One would first need to change our setup to the 3G detectors, where a lower frequency cutoff and extreme SNRs could be encountered. In addition, a wider range of objects should be accounted for. One should include higher-order modes and eccentricity as they could play a crucial role in the 3G era. Other modifications could also be implemented. Additionally, we need to account for the change in noise realization from one event to the other. Some of these steps, like changing the detector configuration, should be relatively easy. Others are more complex, as it is hard to perform parameter inference for long-lasting mergers due to the computational burden. So, extra developments in parameter estimation using machine learning would be required to get to the realistic 3G scenario. For overlapping signals, one would also benefit from developments in the classical study of the 3G scenario, such as how to deal with the noise characterization or the types of other events that could come into the data.

In the end, there is still work to be done before machine learning can be used in realistic 3G scenarios. However, we believe that this work shows it is an interesting avenue and could be practical on a relatively short time scale.

ACKNOWLEDGEMENTS

The authors are thankful to Tomasz Baka, Tom Hesses, and Twan van Laarhoven for discussions on related topics. The authors also thank Maximilian Dax for the careful re-reading of the manuscript. JL, JJ, and CVDB are supported by the research program of the Netherlands Organisation for Scientific Research (NWO). AK is supported by the NWO under the CORTEX project (NWA.1160.18.316). The authors are grateful for computational resources provided by the LIGO Laboratory and supported by the National Science Foundation Grants

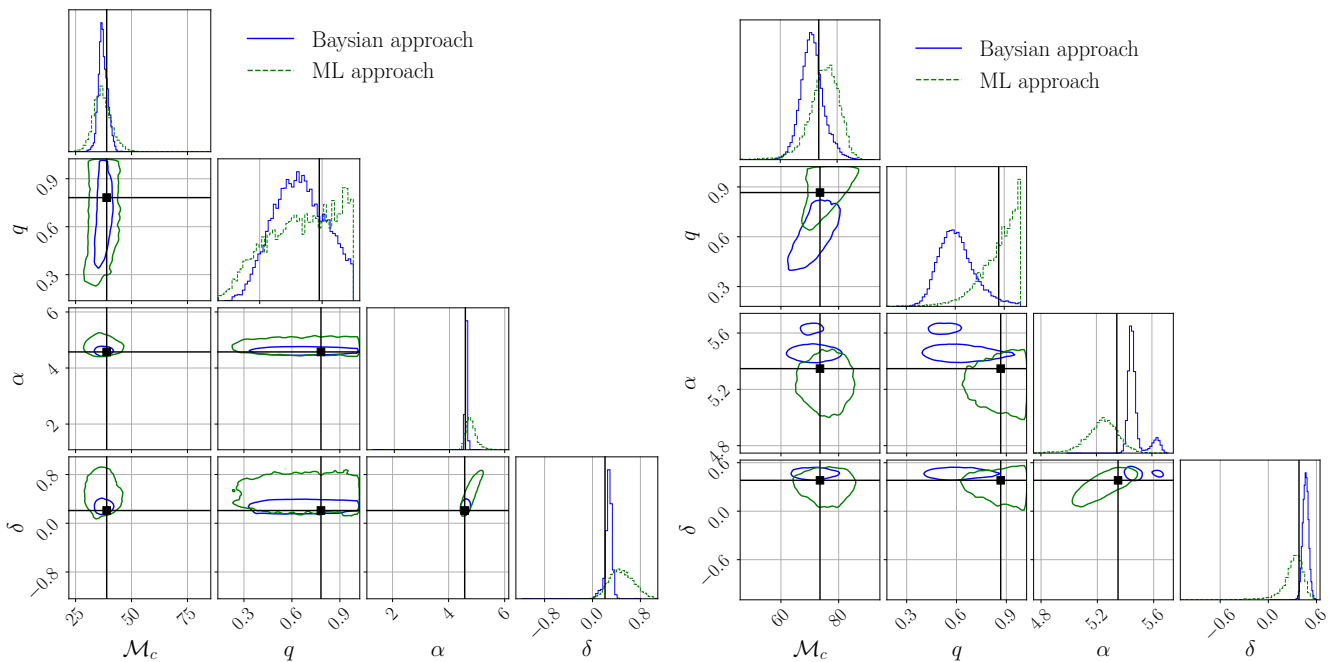


FIG. 3. Comparison between our approach and the one from [29] for two separate events and for the chirp mass, mass ratio, right ascension, and declination. The injected values are given by the black lines. For the left event the true value is encapsulated by the posteriors of both methods, for the right event this is only the case for our method. Our posteriors are generally broader but include the injected value within the 90% confidence interval. This could be corrected by applying importance sampling on the output samples.

No. PHY-0757058 and No. PHY-0823459. We are grateful for computational resources provided by Cardiff University, and funded by an STFC grant supporting UK Involvement in the Operation of Advanced LIGO.

* j.w.langendorff@uu.nl

† alex.kolmus@ru.nl

- [1] T. L. S. Collaboration, *Classical and Quantum Gravity* **32**, 074001 (2015).
- [2] F. Acernese *et al.* (VIRGO), *Class. Quant. Grav.* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [3] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), (2021), arXiv:2111.03606 [gr-qc].
- [4] K. Somiya (KAGRA), *Class. Quant. Grav.* **29**, 124007 (2012), arXiv:1111.7185 [gr-qc].
- [5] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (KAGRA), *Phys. Rev. D* **88**, 043007 (2013), arXiv:1306.6747 [gr-qc].
- [6] T. Akutsu *et al.* (KAGRA), *Nature Astron.* **3**, 35 (2019), arXiv:1811.08079 [gr-qc].
- [7] T. Akutsu, M. Ando, K. Arai, *et al.*, arXiv e-prints, arXiv:2005.05574 (2020), arXiv:2005.05574 [physics.ins-det].
- [8] T. Akutsu *et al.* (KAGRA), *PTEP* **2021**, 05A101 (2021), arXiv:2005.05574 [physics.ins-det].
- [9] B. Iyer, T. Souradeep, C. Unnikrishnan, S. Dhurandhar, S. Raja, and A. Sengupta, “LIGO-India Tech. rep.,” <https://dcc.ligo.org/LIGO-M1100296/public> (2011).
- [10] M. Punturo *et al.*, *Classical and Quantum Gravity* **27**, 194002 (2010).
- [11] S. Hild *et al.*, *Class. Quant. Grav.* **28**, 094013 (2011), arXiv:1012.0908 [gr-qc].
- [12] D. Reitze *et al.*, *Bull. Am. Astron. Soc.* **51**, 035 (2019), arXiv:1907.04833 [astro-ph.IM].
- [13] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, K. Ackley, and C. Adams, *Classical and Quantum Gravity* **34**, 044001 (2017).
- [14] T. Regimbau, M. Evans, N. Christensen, E. Katsavounidis, B. Sathyaprakash, and S. Vitale, *Phys. Rev. Lett.* **118**, 151105 (2017).
- [15] B. Sathyaprakash *et al.*, *Class. Quant. Grav.* **29**, 124013 (2012), [Erratum: *Class. Quant. Grav.* **30**, 079501 (2013)], arXiv:1206.0331 [gr-qc].
- [16] T. Regimbau and S. A. Hughes, *Phys. Rev. D* **79**, 062002 (2009), arXiv:0901.2958 [gr-qc].
- [17] A. Samajdar, J. Janquart, C. Van Den Broeck, and T. Dietrich, *Phys. Rev. D* **104**, 044003 (2021), arXiv:2102.07544 [gr-qc].
- [18] E. Pizzati, S. Sachdev, A. Gupta, and B. Sathyaprakash, *Phys. Rev. D* **105**, 104016 (2022), arXiv:2102.07692 [gr-qc].
- [19] P. Relton and V. Raymond, *Phys. Rev. D* **104**, 084039 (2021), arXiv:2103.16225 [gr-qc].
- [20] Y. Himemoto, A. Nishizawa, and A. Taruya, *Phys. Rev. D* **104**, 044010 (2021), arXiv:2103.14816 [gr-qc].
- [21] A. Antonelli, O. Burke, and J. R. Gair, *Mon. Not. Roy. Astron. Soc.* **507**, 5069 (2021), arXiv:2104.01897 [gr-qc].
- [22] S. Wu and A. H. Nitz, (2022), arXiv:2209.03135 [astro-ph.IM].

- [23] S. Sachdev, T. Regimbau, and B. S. Sathyaprakash, Phys. Rev. D **102**, 024051 (2020).
- [24] A. Sharma and J. Harms, Phys. Rev. D **102**, 063009 (2020).
- [25] S. Biscoveanu, C. Talbot, E. Thrane, and R. Smith, Phys. Rev. Lett. **125**, 241101 (2020).
- [26] B. Zhou, L. Reali, E. Berti, M. Çalışkan, C. Creque-Sarbinowski, M. Kamionkowski, and B. S. Sathyaprakash, (2022), arXiv:2209.01310 [gr-qc].
- [27] B. Zhou, L. Reali, E. Berti, M. Çalışkan, C. Creque-Sarbinowski, M. Kamionkowski, and B. S. Sathyaprakash, (2022), arXiv:2209.01221 [gr-qc].
- [28] L. Reali, A. Antonelli, R. Cotesta, S. Borhanian, M. Çalışkan, E. Berti, and B. S. Sathyaprakash, (2022), arXiv:2209.13452 [gr-qc].
- [29] J. Janquart, T. Baka, A. Samajdar, T. Dietrich, and C. Van Den Broeck, (2022), arXiv:2211.01304 [gr-qc].
- [30] B. Zackay, L. Dai, and T. Venumadhav, (2018), arXiv:1806.08792 [astro-ph.IM].
- [31] L. Dai, T. Venumadhav, and B. Zackay, (2018), arXiv:1806.08793 [gr-qc].
- [32] N. Leslie, L. Dai, and G. Pratten, Phys. Rev. D **104**, 123030 (2021), arXiv:2109.09872 [astro-ph.IM].
- [33] S. Morisaki, Phys. Rev. D **104**, 044062 (2021), arXiv:2104.07813 [gr-qc].
- [34] S. Gao, F. Hayes, S. Croke, C. Messenger, and J. Veitch, Phys. Rev. Res. **4**, 023006 (2022), arXiv:2109.01535 [quant-ph].
- [35] D. Jimenez Rezendé and S. Mohamed, arXiv e-prints, arXiv:1505.05770 (2015), arXiv:1505.05770 [stat.ML].
- [36] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, arXiv e-prints, arXiv:1606.04934 (2016), arXiv:1606.04934 [cs.LG].
- [37] G. Papamakarios, T. Pavlakou, and I. Murray, arXiv e-prints, arXiv:1705.07057 (2017), arXiv:1705.07057 [stat.ML].
- [38] E. Cuoco *et al.*, Mach. Learn. Sci. Tech. **2**, 011002 (2021), arXiv:2005.03745 [astro-ph.HE].
- [39] K. Cranmer, J. Brehmer, and G. Louppe, Proceedings of the National Academy of Sciences **117**, 30055 (2020).
- [40] A. Delaunoy, A. Wehenkel, T. Hinderer, S. Nissanke, C. Weniger, A. R. Williamson, and G. Louppe, (2020), arXiv:2010.12931 [astro-ph.IM].
- [41] S. R. Green, C. Simpson, and J. Gair, Phys. Rev. D **102**, 104057 (2020), arXiv:2002.07656 [astro-ph.IM].
- [42] S. R. Green and J. Gair, Mach. Learn. Sci. Tech. **2**, 03LT01 (2021), arXiv:2008.03312 [astro-ph.IM].
- [43] M. Dax, S. R. Green, J. Gair, M. Deistler, B. Schölkopf, and J. H. Macke, (2021), arXiv:2111.13139 [cs.LG].
- [44] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, Phys. Rev. Lett. **127**, 241103 (2021), arXiv:2106.12594 [gr-qc].
- [45] M. Dax, S. R. Green, J. Gair, M. Pürerer, J. Wildberger, J. H. Macke, A. Buonanno, and B. Schölkopf, (2022), arXiv:2210.05686 [gr-qc].
- [46] M. J. Williams, J. Veitch, and C. Messenger, Phys. Rev. D **103**, 103006 (2021), arXiv:2102.11056 [gr-qc].
- [47] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, Nature Physics **18**, 112 (2022).
- [48] I. Kobyzev, S. J. Prince, and M. A. Brubaker, IEEE Transactions on Pattern Analysis and Machine Intelligence **43**, 3964 (2021).
- [49] G. Papamakarios, E. Nalisnick, D. Jimenez Rezendé, S. Mohamed, and B. Lakshminarayanan, arXiv e-prints, arXiv:1912.02762 (2019), arXiv:1912.02762 [stat.ML].
- [50] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, Advances in neural information processing systems **31** (2018).
- [51] C. Winkler, D. Worrall, E. Hoogeboom, and M. Welling, arXiv e-prints, arXiv:1912.00042 (2019), arXiv:1912.00042 [cs.LG].
- [52] G. Papamakarios and I. Murray, arXiv e-prints, arXiv:1605.06376 (2016), arXiv:1605.06376 [stat.ML].
- [53] D. Ha, A. M. Dai, and Q. V. Le, in *International Conference on Learning Representations* (2017).
- [54] J. Zhuang, N. C. Dvornek, sekhar tatikonda, and J. s Duncan, in *International Conference on Learning Representations* (2021).
- [55] W. Grathwohl, R. T. Q. Chen, J. Bettencourt, I. Sutskever, and D. Duvenaud, arXiv e-prints, arXiv:1810.01367 (2018), arXiv:1810.01367 [cs.LG].
- [56] T. M. Nguyen, A. Garg, R. G. Baraniuk, and A. Anandkumar, arXiv e-prints, arXiv:1912.03978 (2019), arXiv:1912.03978 [cs.LG].
- [57] N. Halko, P.-G. Martinsson, and J. A. Tropp, arXiv e-prints, arXiv:0909.4061 (2009), arXiv:0909.4061 [math.NA].
- [58] L. Barsotti, P. Fritschel, M. Evans, and S. Gras, “Advanced LIGO anticipated sensitivity curves,” <https://dcc.ligo.org/LIGO-T1800044/public> (2021).
- [59] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürerer, X. Jiménez Forteza, and A. Bohé, Phys. Rev. D **93**, 044007 (2016), arXiv:1508.07253 [gr-qc].
- [60] A. Kolmus, G. Baltus, J. Janquart, T. van Laarhoven, S. Caudill, and T. Heskes, Phys. Rev. D **106**, 023032 (2022), arXiv:2111.00833 [gr-qc].