

Transdimensional inference for gravitational-wave astronomy with BilbyHUI TONG,^{1,2} NIR GUTTMAN,^{1,2} TEAGAN A. CLARKE,^{1,2} PAUL D. LASKY,^{1,2} ERIC THRANE,^{1,2} ETHAN PAYNE,^{3,4}
ROWINA NATHAN,^{1,2} BEN FARR,⁵ MAYA FISHBACH,^{6,7,8} GREGORY ASHTON,⁹ AND VALENTINA DI MARCO^{1,2}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*Department of Physics, California Institute of Technology, Pasadena, California 91125, USA*⁴*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*⁵*Institute for Fundamental Science, Department of Physics, University of Oregon, Eugene, OR 97403, USA*⁶*Canadian Institute for Theoretical Astrophysics, 60 St George St, University of Toronto, Toronto, ON M5S 3H8, Canada*⁷*David A. Dunlap Department of Astronomy and Astrophysics, 50 St George St, University of Toronto, Toronto, ON M5S 3H8, Canada*⁸*Department of Physics, 60 St George St, University of Toronto, Toronto, ON M5S 3H8, Canada*⁹*Department of Physics, Royal Holloway, University of London, TW20 0EX, UK***ABSTRACT**

It has become increasingly useful to answer questions in gravitational-wave astronomy using *transdimensional* models where the number of free parameters can be varied depending on the complexity required to fit the data. Given the growing interest in transdimensional inference, we introduce a new package for the Bayesian inference Library (Bilby) called **tBilby**. The **tBilby** package allows users to set up transdimensional inference calculations using the existing **Bilby** architecture with off-the-shelf nested samplers and/or Markov Chain Monte Carlo algorithms. Transdimensional models are particularly helpful when we seek to test theoretically uncertain predictions described by phenomenological models. For example, bursts of gravitational waves can be modelled using a superposition of N wavelets where N is itself a free parameter. Short pulses are modelled with small values of N whereas longer, more complicated signals are represented with a large number of wavelets stitched together. Other transdimensional models have found use describing instrumental noise and the population properties of gravitational-wave sources. We provide a few demonstrations of **tBilby**, including fitting the gravitational-wave signal GW150914 with a superposition of N sine-Gaussian wavelets. We outline our plans to further develop the **tBilby** code suite for a broader range of transdimensional problems.

1. INTRODUCTION

Since the first detection of gravitational waves (Abbott et al. 2016a), Bayesian inference has been widely used to infer the astrophysical properties of merging binaries (Abbott et al. 2016b). Bayesian inference is used to search for physics beyond general relativity (Abbott et al. 2016c), to probe nuclear physics at extreme densities (Abbott et al. 2018), to measure the expansion of the Universe (Abbott et al. 2017; Hotokezaka et al. 2019), and to study the formation of merging binaries (Abbott et al. 2021, 2023).

In many applications, the framework underpinning these inferences is theoretically precise; that is, we have trustworthy, quantitative predictions for the data given

the model parameters. For example, when we infer the masses of merging black holes, we are able to leverage state-of-the-art gravitational waveforms, built from numerical-relativity simulations, to interpret data. In other cases, however, there is significant theoretical uncertainty and so we rely on phenomenological models. For example, following the detection of GW150914, the LIGO–Virgo Collaborations used the **BayesWave** algorithm (Cornish & Littenberg 2015) to perform a study to reconstruct the strain time series in the data with minimal assumptions using a superposition of N sine-Gaussian wavelets (Abbott 2016; Klimentenko et al. 2008).¹ If we treat N as a free parameter, then the total number of model parameters is itself variable. Such an analysis—where the number of free parameters is it-

Corresponding author: Hui Tong
hui.tong@monash.edu

¹ Sine-Gaussian functions are sometimes called Morlet or Gabor wavelets (Kronland-Martinet et al. 1987)

self a free parameter—is said to be *transdimensional*. The striking agreement between LIGO–Virgo’s minimal-assumption reconstruction and the waveform predicted by general relativity helped cement the interpretation of the signal as a binary black hole merger (Abbott 2016). It remains a powerful demonstration of the usefulness of transdimensional models.

There are other noteworthy applications of transdimensional inference in gravitational-wave astronomy. In the audio band where the LIGO–Virgo–KAGRA (LVK; Aasi et al. 2015; Acernese et al. 2015; Aso et al. 2013) observatories operate, the **BayesWave** package (Cornish & Littenberg 2015; Cornish et al. 2021) has been used for minimum-assumption model checking and waveform reconstruction (Millhouse et al. 2018; Pannarale et al. 2021; Dàya et al. 2021), improving the statistical significance of short and unmodeled “bursting” signals (Littenberg et al. 2016; Yi Shuen C. Lee & Melatos 2021), modelling astrophysically uncertain waveforms (e.g., from supernovae and hypermassive neutron stars) (Raza et al. 2022; Miravet-Tenés et al. 2023; Ashton & Dietrich 2022), modelling deviations from general relativity (Chatziioannou et al. 2021b; Johnson-McDaniel et al. 2022), and subtracting noise artifacts (glitches) (Littenberg & Cornish 2010; Pankow et al. 2018; Chatziioannou et al. 2021a; Davis et al. 2022; Hourihane et al. 2022). Meanwhile, the related **BayesLine** code is frequently used to estimate the noise power spectral density of gravitational-wave measurements (Littenberg & Cornish 2015; Gupta & Cornish 2023). Transdimensional analyses have also been demonstrated for use in the millihertz band by space-based observatories (Littenberg et al. 2020) and in the nanohertz band by pulsar timing arrays (Ellis & Cornish 2016). The code package **Eryn** (Karnesis et al. 2023) was recently introduced as a multi-purpose tool for transdimensional inference with special attention to problems relevant for the LVK and LISA.

In this work, we introduce **tBilby**, a package for transdimensional sampling with the Bayesian Inference Library **Bilby** (Ashton et al. 2019; Romero-Shaw et al. 2020). **Bilby** is widely used in gravitational-wave astronomy. It is designed and maintained with four guiding principles: modularity, consistency, generality, and usability. The mission for **Bilby** is to be intuitive enough to be used by new researchers, while still being applicable to a broad class of problems, and with the ability to easily swap samplers when needed. Our goal is to leverage these attributes, building on the existing **Bilby** infrastructure, in order to make it easier for **Bilby** users to carry out transdimensional analyses.

We envision the **tBilby** project as a long-term effort that will be developed gradually. With this in mind, we start here with a specific class of transdimensional problems: transient waveforms that can be modelled with a superposition of N component functions. In particular, we demonstrate a minimum-assumption reconstruction of GW150914 using a superposition of N sine-Gaussian functions. We use this demonstration to explain key concepts in transdimensional inference including the notion of ghost parameters and order statistics of a uniform distribution.² Readers can reproduce our calculation using the accompanying code.³ The remainder of this paper is organized as follows. In Section 2, we cover the basic principles of transdimensional inference and describe how they are implemented in **tBilby**. In Section 3, we demonstrate the **tBilby** package with two examples: a toy-model problem consisting of a superposition of Gaussian pulses and a minimum-assumption reconstruction of GW150914. We provide closing remarks in Section 4, briefly demonstrating another transdimensional example fitting LIGO’s noise amplitude spectral density with a sum of N power laws and M Lorentzian functions. We also sketch our priorities for future development.

2. METHOD

One of the goals of Bayesian inference is to determine the posterior distribution for model parameters $\vec{\theta}$ given a prior $\pi(\vec{\theta})$, data d , and likelihood $\mathcal{L}(d|\vec{\theta})$. In a transdimensional problem, the number of parameters N is itself a parameter:

$$\vec{\theta} \equiv \{\theta_1, \dots, \theta_N, N\}. \quad (1)$$

In some cases, this problem can be solved with brute-force parallelization: one can run multiple inference jobs, each with a different *fixed* number of parameters N , and then combine the resulting samples based on the Bayesian evidence for each fixed- N analysis $\mathcal{Z}_N$, as well as their prior preference. This approach works adequately when there is a relatively small range of values for N . However, it becomes inefficient when time is wasted exploring many values of N disfavoured by the likelihood function. The solution is to sample in N .⁴

The number of parameters N is treated similarly to any other discrete parameter in **Bilby**. In our demon-

² A dedicated appendix is provided for each of these topics: Appendix B covers ghost parameters while Appendix C discusses the mathematics of order statistics.

³ The code can be found at the git repository <https://github.com/tBilby/tBilby>.

⁴ In the context of Markov chain Monte Carlo samplers, this is essentially the same as the Reversible jump Markov chain Monte Carlo technique (Green 1995).

strations below, we take the prior $\pi(N)$ to be uniform on the interval $[N_{\min}, N_{\max}]$. At each step, the sampler draws a value of N along with values for all possible parameters in $\vec{\theta}$ —even parameters $\theta_{k>N}$ that are not used for the N -parameter model. We refer to the $\theta_{k>N}$ parameters as “ghost parameters” since they are not included in the likelihood evaluation, similar to the method in [Lodewyckx et al. \(2011\)](#); [Liu et al. \(2023\)](#). In App. B, we prove that when we marginalize over the ghost parameters, we obtain the same posterior as one would obtain without ghost parameters using either the brute-force method or the transdimensional sampler.⁵

From the perspective of the `tBilby` code, the transdimensional model behaves like a fixed-dimensional model in order to obtain the joint posterior:

$$p(\theta_1, \dots, \theta_{N_{\max}}, N | d) \propto \pi(\theta_1, \dots, \theta_{N_{\max}}) \pi(N) \mathcal{L}(d | \theta_1, \dots, \theta_N). \quad (2)$$

Since the likelihood does not depend on the ghost parameters, the marginal posterior distribution for the $k > N$ ghost parameters is equivalent to the prior for the ghost parameters

$$p(\theta_k | d, N, k > N, \theta_{j \leq N}) = \pi(\theta_k | \theta_{j \leq N}). \quad (3)$$

The $k > N$ ghost-parameter samples are removed in post-processing since they are not actually part of our model. The ghost-parameter framework is convenient since it allows us to perform transdimensional inference using the off-the-shelf samplers already available in `Bilby`.⁶

In this paper, we mainly focus on a specific set of transdimensional problems in which the data consists of a time series $d(t)$, and the signal model $s(t)$ is modelled with a sum of N components, each with parameters θ_k :

$$s(t | \vec{\theta}) = \sum_{k=1}^N s_k(t | \theta_k). \quad (4)$$

We refer to each $s_k(t | \theta_k)$ as a *component function*.

An interesting issue arises fitting signals with a superposition of identical component functions ([Buscicchio et al. 2019](#)). Since the component functions are identical, the labels of parameters can be swapped without changing the likelihood, leading to multiple likelihood

peaks all describing the same fit. Given N component functions, the number of these degenerate modes scales like $N!$. Thus, multimodality quickly becomes problematic if left unchecked. One must therefore intervene to avoid a multimodal likelihood surface that needlessly complicates our sampling efforts.

We address this issue by using the mathematics of *order statistics* ([David & Nagaraja 2004](#)). The basic idea of order statistics is to devise a method for ranking what would otherwise be indistinguishable component functions so that each one is uniquely identified. This eliminates the artificial multimodality arising from label swapping.

Of course, there are different ways in which component functions can be ordered: a set of Morlet wavelets, for example, can be ordered by frequency, time, amplitude, etc. The choice of a suitable ordering strategy is problem-dependent. However, for many transdimensional problems, it is useful to order component functions by descending signal-to-noise ratio (SNR).

As the sampler explores the likelihood surface, it is likely to find features in the order of each feature’s contribution to likelihood. Therefore, a natural way to order components is by descending SNR so that the first component is the one with the highest SNR. This way, the order of components are reasonably well matched with the order that the sampler finds the features, which reduces the chances that the sampler becomes stuck in a local likelihood maximum. The first component, with the largest SNR, is likely to be detected by the sampler first, followed by the second component with the second largest SNR component, and so on. In the next Section, we demonstrate the principles outlined in this Section with examples.

While we highlight the usefulness of SNR ordering, `tBilby` users are free to implement whichever ordering strategy they choose. In Appendix D, we discuss alternative ordering schemes, showing that they are mathematically equivalent, but that SNR ordering is more reliable for many problems of interest.

3. DEMONSTRATION

3.1. A superposition of Gaussian pulses

As a warm-up exercise, we consider a simple problem of fitting data with N Gaussian pulses in the presence of Gaussian white noise.⁷ Our data $d(t)$ is a time series

⁵ It is interesting to note that ghost parameters $\theta_{k>N}$ incur no Occam penalty. Since the ghost parameters do not appear in the likelihood, the flexibility of the model has not changed, so there is no penalty for adding unnecessary complexity.

⁶ The additional computational cost incurred by drawing prior samples that we do not use is (for most applications that we foresee) negligible compared the cost of the likelihood evaluation.

⁷ The calculations in this subsection are performed in the accompanying `jupyter` notebook, `pulse.ipynb`, in the git repository linked above.

consisting of signal $s(t)$ and noise $n(t)$:

$$d(t) = s(t) + n(t). \quad (5)$$

Our signal model is a superposition of component functions given by

$$s(t|\mu_k, A_k, \sigma_k) = \frac{A_k}{\sigma_k \sqrt{2\pi}} \exp\left(-\frac{(t - \mu_k)^2}{2\sigma_k^2}\right). \quad (6)$$

The mean μ_k , amplitude A_k and width σ_k are free to vary. We simulated data with a signal consisting of $N = 3$ pulses plus Gaussian white noise $n(t)$ with variance

$$\sigma_n = \sqrt{\langle n(t)n(t') \rangle} = 0.15 \delta_{tt'}, \quad (7)$$

where $\delta_{tt'}$ is a Kronecker delta function.

The corresponding SNR expression of each Gaussian can be written as

$$\text{SNR}_k \propto \frac{A_k}{\sqrt{\sigma_k \sigma_n}}. \quad (8)$$

The likelihood is

$$\mathcal{L}(d|s, \sigma_n) = \prod_i \frac{1}{\sqrt{2\pi\sigma_n^2}} \exp\left(-\frac{(d(t_i) - s(t_i))^2}{2\sigma_n^2}\right), \quad (9)$$

where d is the data, s is the signal template, t_i is a discrete time, and σ_n is the noise. We show simulated data in Fig. 2 created with parameters summarized in Table 1.

Our prior for the number of pulses N is a discrete uniform distribution $\mathcal{U}(0, 6)$. We order each pulse in descending order according to its SNR. We assume that the *unordered* SNR values are uniformly distributed on the interval $(0, \text{SNR}_{\max})$ where $\text{SNR}_{\max} = 10$. It follows that the prior on the first *ordered* SNR is given by a beta distribution with parameters $\alpha = N$ and $\beta = 1$ (David & Nagaraja 2004):

$$\pi(\rho_1) \propto \text{Beta}(\rho_1|N, 1). \quad (10)$$

Here, ρ_1 is the “relative signal-to-noise ratio” given by

$$\rho_i = \text{SNR}_i / \text{SNR}_{\max}. \quad (11)$$

The joint prior for subsequent ordered SNR values when $i < j$ are given by (David & Nagaraja 2004)

$$\pi(\rho_i, \rho_j) = N! \frac{(1 - \rho_i^{i-1})}{(i-1)!} \frac{(\rho_i - \rho_j)^{j-i-1}}{(j-i-1)!} \frac{\rho_j^{N-j}}{(N-j)!} \quad (12)$$

Using Eqs. 10-12, one can draw a full set of ordered SNRs. A corner plot showing an example prior for ordered SNRs is provided in Fig. 1. Additional details are




Figure 1. An example of descending SNR priors. Note here the variables are relative SNR ρ_i defined in Eq. 11




Figure 2. Simulated data (blue) for our Gaussian pulse model described in Subsection 3.1. The signal (orange) consists of three Gaussian pulses.

Pulse	Mean	Amplitude	Width
1	35	1.0	10
2	74	0.8	8
3	101	1.2	12

Table 1. Parameter values for the Gaussian pulses shown in Fig.2.

provided in C.

The prior for the central time of the pulse μ_k is uniform over the interval $(0, 150)$. For the widths, we employ a uniform priors $U[5, 20]$.

We run **tBilby** using two different samplers: the nested sampler **dynesty** (Speagle 2020) and the parallel-tempered Markov chain Monte Carlo sampler **ptemcee** (Vousden et al. 2015). For **ptemcee**, we update the number of pulses N and the parameters for each pulse θ_k separately every time a new move is proposed in the sampling process. Since we are using ghost parameters, the sampler behaves as though it is exploring a fixed-dimensional space. In each iteration, we randomly add a pulse, remove a pulse, or keep the number of pulses fixed with equal probability. Since **dynesty** draws samples from priors, jumps in N occur automatically by virtue of the discrete prior $\pi(N)$.

In Fig. 3, we plot the posterior odds

$$\mathcal{O}(N) = \frac{\mathcal{L}(d|N)}{\mathcal{L}(d|N'=3)}, \quad (13)$$

which compares the posterior support for different values of N to the best-fit $N = 3$ model.⁸⁹ The results obtained with **dynesty** are shown in orange while the results obtained with **ptemcee** are shown in green. As a sanity check, we also use **dynesty** to calculate the marginal likelihood for each value of N with separate fixed- N , which, combined with our prior of N , we use to estimate the ground-truth posterior obtained without transdimensional inference. All three methods produce a similar distribution. Some values of N are strongly disfavoured, and so the transdimensional sample records no posterior samples for that value of N . In such cases, we set an upper limit on $\ln \mathcal{O}$. Both **dynesty** and **ptemcee** produce $\ln \mathcal{O}$ values that are consistent with the fixed- N ground truth.

We compare the computational cost between the brute-force method of performing many fixed- N runs and using **tBilby**. The fixed- N runs for $N \in [0, 6]$ take roughly 5.2 times the sampling time of the **tBilby**

⁸ Astute readers may notice that the right-hand side of Eq. 13 does not include the prior odds. This is because the prior odds in this case are unity.

⁹ We estimate the uncertainty in our $\ln \mathcal{O}$ calculations as follows:

$$\sigma_{\ln \mathcal{O}}^2 = \frac{1}{n_N} + \frac{1}{n_\phi}. \quad (14)$$

Here, n_N is the number of posterior samples for the hypothesis that the data are described by N component functions while n_ϕ is the number of posterior samples describing the fiducial model—in this case, $N = 3$.




Figure 3. Natural log posterior odds obtained with different sampling techniques (see Eq. 13). The odds are measured relative to the favored $N = 3$ model. The uncertainties are one-sigma. The orange and green points are transdimensional sampling results using **dynesty** and **ptemcee**. The navy blue points labelled by “N fixed” where we calculate the evidence for each value of N with dedicated **dynesty** runs provide the ground truth.

dynesty run with the same sampler settings.¹⁰ As expected, when we run different N models separately, most of the computational time is spent exploring complicated models with large N , which may not be the models with the highest Bayes factor. Since transdimensional sampling accounts for the Occam factor during sampling process, it automatically prevents the sampler exploring disfavoured regions of the parameter space.

In Fig. 4, we present a corner plot showing the marginal posterior distribution of parameters $\text{SNR}_1, \text{SNR}_2, \text{SNR}_3, \mu_1, \mu_2, \mu_3$ given samples $N = 3$. As above, the fixed- N ground truth is shown in blue while the results obtained with **dynesty** and **ptemcee** are shown in orange and green, respectively. All three posteriors produce consistent credible intervals. The means of pulses inevitably show degeneracy to some extent because of the overlap of SNRs between different components. However, the sampling problem is still relatively simplified compared with the case of modelling without ordering when all the different modes in posterior distribution have exactly the same shape.

For the sake of interpretability, we reorder the posterior samples by ascending means of Gaussian pulses

¹⁰ Note this is not a rigorous apples-to-apples comparison. For example, we do not require the same number of effective samples between the brute-force calculation and **tBilby**. However, it does provide a rough estimate of the improvement in computational cost for this particular problem.




Figure 4. Posterior distribution for the SNRs and the means of $N = 3$ Gaussian pulses with descending SNR priors. The different shades show one, two, and three-sigma credible intervals. The ground truth, obtained with a fixed- N **dynesty** run, is shown in blue. In orange we plot the results obtained using a transdimensional implementation of **dynesty** while green shows a transdimensional implementation of **ptemcee**.

and present the corner plot in Fig. 5. It is clear for the means of three pulses detected which are matched with the true values.

3.2. GW150914

We now apply transdimensional inference to reconstruct the signal from the first gravitational-wave observation GW150914. Strain data for GW150914 is accessed via the Gravitational Wave Open Science Center (GWOSC; [LIGO Scientific Collaboration, Virgo Collaboration and KAGRA Collaboration 2018](#)). We char-

acterize the noise power spectral density using 128 s of data ahead of GW150914. We adopt sampling frequency as 2048 Hz with minimum and maximum cutoff frequencies of 20 Hz, 896 Hz. Following **BayesWave** ([Cornish & Littenberg 2015](#)), we assume that the source of gravitational waves is elliptically polarized so that the cross-polarized strain is completely determined by the plus-




Figure 5. Corner plot of posterior samples of $N = 3$ Gaussian pulses resorted by ascending means. The different shades show one, two, and three-sigma credible intervals. The ground truth, obtained with a fixed- N *dynesty* run, is shown in blue. In orange we plot the results obtained using a transdimensional implementation of *dynesty* while green shows a transdimensional implementation of *ptemcee*.

polarized strain:¹¹

$$h_{\times}(f) = \epsilon h_{+}(f) e^{i\pi/2}. \quad (15)$$

Here, $\epsilon \in [-1, 1]$ is the ellipticity, which characterizes the polarization. We fit the binary black hole signal GW150914 using a superposition of sine-Gaussian wavelets; see [Abbott \(2016\)](#):

$$\Psi(t|A, f_0, t_0, \phi) = A e^{-(t-t_0)^2/\tau^2} \cos(2\pi f_0(t-t_0) + \phi) \quad (16)$$

with $\tau = Q/(2\pi f_0)$. Here, A is the amplitude, τ is the damping time, Q is the quality factor, t_0 is the central time, f_0 is the central frequency, and ϕ is the phase offset. We approximate the SNR of a single wavelet using ([Cornish & Littenberg 2015](#))

$$\text{SNR} \simeq \frac{A\sqrt{Q}}{(8\pi)^{1/4} f_0 S_n(f_0)} \quad (17)$$

where S_n is the noise power spectral density. The plus-polarized strain h_{+} is the summation of several compo-

nents

$$h_{+}(t) = \sum_j \Psi(t|\text{SNR}_j, f_j, t_j, \tau_j, \phi_j). \quad (18)$$

We adopt the following priors. The distribution of non-ghost SNR $k \leq N$ are distributed according to Eqs. 10-12 for ordered statistics uniformly distributed on the interval $[0, 30]$.¹² Ghost parameter SNRs $k > N$ are uniformly distributed on the interval $[0, 30]$. The quality factor Q is taken from a uniform distribution on the interval $[0.1, 40]$, and ϕ follows a uniform distribution between 0 and 2π with periodic boundary conditions. We adopt a uniform prior for f_j between $[20 \text{ Hz}, 512 \text{ Hz}]$ and t_j between $[-0.3 \text{ s}, +0.2 \text{ s}]$ around the trigger time.

We analyze the LIGO–Virgo event GW150914 ([Abbott et al. 2016a](#)) using *dynesty* with the ghost parameter framework described above. We allow up to $N_{\text{max}} = 8$ wavelets (45 total parameters). We set up the analysis using *rwalk* method implemented in *Bilby* with $\text{nact}=80$ and $\text{nlive}=2000$. The sampling is converged in 4 days. The reconstructed waveform is shown in Fig. 7 (red) alongside the whitened data (peach), and the compact binary coalescence (CBC) template fit shown as the blue curve.¹³ The top panel is for the LIGO Hanford Observatory (LHO) while the bottom panel is for the LIGO Livingston Observatory (LLO). The wavelet fit produces a qualitatively similar reconstruction as the compact binary template fit. Both fits recover the morphology of key features in the whitened data.

In Fig. 6 we show the model selection result of GW150914 wavelet analysis. The posterior favours $N = 4$ over $N = 3$ and $N = 5$ with Bayes factor of 5.13 and 7.13 respectively. Models of $N < 3$ and $N > 5$ wavelets are strongly disfavored. The wavelet fit produces a higher maximum likelihood than the template fit ($\Delta \ln \mathcal{L} = 6$) at the sample of $N = 6$. Conditional posterior samples of the favored model $N = 4$ have maximum $\ln \mathcal{L} = 306.5$ which is comparable to the maximum $\ln \mathcal{L} = 306.7$ of the template fit. Since we expect the template derived from general relativity to fit the signal, we interpret this as evidence that the $N = 6$ wavelet fit is beginning to overfit features in the noise. In order to make further progress, it may be necessary to develop sampler settings that are better tuned for this transdimensional problem. We plan to adjust the im-

¹¹ For some bursting sources, it may be appropriate to adopt an unpolarized model so that $h_{\times}$ is modelled independently from h_{+} .

¹² Since we are only interested in relative SNR, we ignore additional normalization constants in Eq. 17, treating the proportionality symbol in that equation as an equality.

¹³ The CBC fit is obtained using the waveform approximant IMR-PHENOMXPHM ([Pratten et al. 2021](#)).




Figure 6. The model selection result for GW150914 wavelet fit. The ratio of number of samples at different N represents the Bayes factor of different models.




Figure 7. The reconstructed signal from GW150914. The red trace shows the whitened signal reconstructed using a transdimensional sine-Gaussian wavelet fit. The blue trace shows the maximum-likelihood template obtained with the IMRPHENOMXPHM approximant. The whitened data is shown in peach. The top panel is for the H1 observatory in Hanford, WA while the bottom panel is for the L1 observatory in Livingston, LA.

plementation of `dynesty` in `tBilby` as a focus of future work.

In Fig. 8, we show the posterior distributions for the frequencies of $N = 4$ wavelets. The blue posterior distributions are obtained with a fixed $N = 4$ analysis using `Bilby`, while the orange results are obtained allowing for any value of N using `tBilby`. We reorder the samples from both analyses by ascending frequen-

cies for the sake of interpretability. In Fig. 9 we show the sky localisation map for GW150914, where the blue curves are the 90% credible intervals obtained using the IMRPHENOMXPHM waveform approximant and the orange curves are obtained using our transdimensional sine-Gaussian wavelet fit.

As part of our validation, we also perform transdimensional inference on simulated data to validate our GW150914 reconstruction results. We inject 200 signals—drawn from our priors and consisting of different numbers of sine-Gaussians—to Gaussian noise, coloured to the LIGO O4 noise power spectral density. See Appendix E for complete details.

We find the sampler reconstructs 95% of injections correctly with the true number of wavelets N . In the remaining 5% of injections, the true value of N is incorrectly ruled out, but for an interesting reason. We determine that all of these injections include two wavelets, coincident in time, which the sampler fits using a single wavelet. Thus, for example, an injection with $N = 5$ is recovered with a strong preference for $N = 4$. In each case, the reconstructed waveform produces a good fit. However, the sampler has difficulty, for example, jumping from an adequate $N = 4$ fit to an also-adequate $N = 5$ fit with two wavelets on top of each other.

To the extent that we care only about the adequacy of the reconstructed waveform, these wrong- N recoveries are not especially concerning. After all, the wavelet model is entirely phenomenological, and so there is no true value for N in actual data. On the other hand, it is desirable for model interpretability to eliminate this corner case. A number of solutions may help including proximity priors (which prevent two wavelets from falling on top of each other in the first place) and custom jump proposals. We leave this for future work.

4. DISCUSSION AND CONCLUSIONS

We introduce the `tBilby` package that facilitates transdimensional inference calculations with `Bilby`. Focusing, to start with, on time-domain models with a superposition of component functions, we provide examples where users can employ off-the-shelf samplers in `Bilby` to reconstruct signals with minimal alterations. The package includes example implementations of ghost parameters and order statistics, useful ingredients for this class of transdimensional problems. We show how `tBilby` can be used to perform a minimum-assumption fit of GW150914 with sine-Gaussian wavelets as in Abbott (2016).

For future work, we propose to improve the efficiency of `tBilby` through the use of more finely tuned samplers, designed for specific classes of problems of interest




Figure 8. Posteriors for the frequencies of different sine-Gaussian wavelets fit to GW150914. The blue results are obtained with $N = 4$ fixed. The orange shows posteriors of samples with $N = 4$ from a transdimensional fit, which allows for any value of N .

in gravitational-wave astronomy. Thanks to the modular design of **Bilby**, it is relatively easy to experiment with different options. While we find that **dynesty** produces well-converged fits to GW150914 for $N_{\max} \leq 8$, we do not obtain reliable fits with **ptemcee** in a decent sampling time—at least using the default settings. Our work highlights the potential for carrying out transdimensional inference with nested sampling; see, e.g., [Brewer et al. \(2015\)](#).

We see this paper as the first step in a broader program to facilitate transdimensional inference with **Bilby**—in gravitational-wave astronomy and other con-

texts. We highlight a few priorities. First, as evidenced by work done with **BayesLine** ([Littenberg & Cornish 2015](#)), transdimensional inference is a powerful tool for modelling the noise in gravitational-wave observatories; see also [Gupta & Cornish \(2023\)](#). Noise modelling naturally lends itself to transdimensional models because the noise power spectral density can be characterised by some fiducial shape plus a variable number of spectral features superposed on top. Transdimensional models can be used to obtain smooth fits of the noise power spectral density while characterizing instrumental lines and other features, enabling us to study the evolution of




Figure 9. Sky map showing the 90 % credible intervals for GW150914 recovered by a compact binary coalescence template fit (blue) and a transdimensional sine-Gaussian wavelet fit (orange).

these features over the course of an observing run. We are also excited about the application of transdimensional sampling to model potential glitches simultaneously with compact binary signals (Chatziioannou et al. 2021a; Hourihane et al. 2022). This work may help astronomers to better interpret gravitational-wave events with potential data quality problems (Payne et al. 2022). A comprehensive study of noise modelling using **tBilby** will be detailed in a forthcoming publication.

Second, we envision extending **tBilby** to build more flexible models describing the population properties of binary black holes and neutron stars; see, e.g., Toubiana et al. (2023). For example, one may wish to model the distribution of primary black hole mass distribution with a variable number of peaks and troughs. Recent studies have highlighted the usefulness of flexible models to identify structure that might be missing from astrophysically inspired phenomenological models; see,

e.g., Tiwari & Fairhurst (2021); Edelman et al. (2022, 2023).

Finally, we propose to develop **tBilby** for applications beyond terrestrial gravitational-wave observatories. For example, pulsar timing measurements, which can be used to measure nanohertz gravitational-waves (Agazie et al. 2023; Antoniadis et al. 2023; Reardon et al. 2023; Xu et al. 2023), rely on measurements of the time of arrival of arbitrarily shaped radio pulses. By modelling these pulses using a superposition of component functions, it is sometimes possible to identify and account for aberrant behaviour in the pulsar evolution, ultimately improving sensitivity for gravitational-wave searches (Nathan et al. 2023). Transdimensional models may prove useful determining the number of component functions used in these fits. Of course, this is just one example. It is our hope that the **tBilby** package will facilitate the development of numerous transdimensional models for physics and astronomy.

- 1 This work is supported through Australian Research
- 2 Council (ARC) Centres of Excellence CE170100004,
- 3 CE230100016, Discovery Projects DP220101610 and
- 4 DP230103088, and LIEF Project LE210100002. T. A.
- 5 C. receives support from the Australian Government
- 6 Research Training Program. The authors are grateful
- 7 for for computational resources provided by the LIGO
- 8 Laboratory computing cluster at California Institute
- 9 of Technology supported by National Science Founda-
- 10 tion Grants PHY-0757058 and PHY-0823459, and the
- 11 Ngarrgu Tindebeek / OzSTAR Australian national fa-
- 12 cility at Swinburne University of Technology. LIGO was
- 13 constructed by the California Institute of Technology
- 14 and Massachusetts Institute of Technology with fund-
- 15 ing from the National Science Foundation and operates
- 16 under cooperative agreement PHY-1764464. This paper
- 17 carries LIGO Document Number LIGO-P2400105.

APPENDIX

A. CODE DESIGN

The objective of **tBilby** is to provide a comprehensive toolkit for handling transdimensional sampling. The **tBilby** package offers flexibility and automation. As outlined in this paper, the development of **tBilby** is part of a long-term project with multiple goals. At present, we have constrained the package to a set of essential tools and examples. **tBilby**'s design philosophy closely aligns with that of **Bilby**, emphasizing open-source code, modularity, generality, and usability (Ashton et al. 2019). Based on the ideas and infrastructure of **Bilby**, **tBilby** ensures a relatively smooth user experience, particularly for experienced users. Furthermore, we reinforce this philosophy by mandating that the sole requirement for **tBilby** is an installation of **Bilby**.

The structure of **tBilby** closely mirrors that of **Bilby**, with the core module including **base**, **prior**, and **sampler** modules, alongside an additional folder dedicated to examples. The **base** module contains fundamental functionality for

constructing transdimensional models and defining transdimensional priors. The **prior** folder houses priors intended for transdimensional sampling, while the **sampler** module facilitates support for transdimensional samplers.

The key building block of a transdimensional model in **tBilby** is the *transdimensional parameter*, which refers to a parameter of a component function that has multiple “orders” (in this language, each sine-Gaussian is a different order). Another fundamental concept is the *transdimensional prior*, which constitutes a set of priors related to a transdimensional parameter and which is attached to the parameter’s order. Transdimensional models with proximity priors employ conditional statements. These two elements serve as the basic building blocks.

For practical purposes, transdimensional priors in **tBilby** are categorized into four types: (i) transdimensional nested conditional priors, (ii) transdimensional conditional priors, (iii) conditional priors, and (iv) unconditional priors.¹⁴ Transdimensional nested conditional priors are defined by their dependence on previously sampled parameters of the same component function. If we assume that the current order being sampled is n , these priors depend on parameters of orders $n - 1, n - 2$, etc.

Transdimensional conditional priors, on the other hand, are dependent on parameters from all sampled orders of a component function, denoted by $k, k - 1$, etc. Conditional priors rely on a set of non-transdimensional parameters, whereas unconditional priors are independent of other parameters. The most general prior may combine elements of all these types, except for the last type, which by definition is an independent prior. In this framework, the most general form of a prior for transdimensional parameter ρ_n is:

$$\pi(\rho_n | \rho_{n-1}, \dots, \phi_{n-1}, \dots, \zeta_k, \zeta_{k-1}, \dots, \Lambda).$$

The variable ϕ represents another set of transdimensional parameters of the same order as the component function so that ρ_n does not depend on $\phi_{m \geq n}$. Meanwhile, ζ signifies another set of transdimensional parameters that depend on all available orders of the component function. Finally, Λ are parameters that may or may not be part of the component function parameters.

By allowing for the definition of conditional transdimensional priors, users can uniquely specify priors for each transdimensional parameter. Practically, this involves defining a class that inherits from a predefined transdimensional prior class and implementing an abstract function to define the mathematical relation between the conditional parameters and prior properties (this is a generalization of **Bilby**’s condition function, which is required when defining a conditional prior).

Facilitating such versatility and control over the priors allows users to gain flexibility in manipulating the prior distribution to suit their specific needs. The flexibility of **tBilby**’s extends further, enabling the construction of function superposition, each potentially comprising a different number of component functions. For instance, the LIGO noise power spectral density may be represented as a combination of several power law functions along with multiple Cauchy-like functions, addressing distinct spectral characteristics (Littenberg & Cornish 2015). Furthermore, **tBilby** offers supplementary tools for removing ghost parameters and generating relevant corner plots, thereby simplifying the analysis of component functions and individual transdimensional parameters.

B. GHOST PARAMETERS

The method outlined here is similar to Liu et al. (2023) who performed transdimensional inference using **Bilby** for gravitational-wave lensing study. In the ghost parameter framework, we introduce extra parameters that do not actually change the likelihood, and therefore do not change the posteriors for the original parameters—as long as the ghost-parameter prior is correctly normalized. For example, we consider the situation in Section 3.1 when $N = 3$. The signal is only determined by $\theta_{k \leq 3}$ while $\theta'_{k > 3}$ represents the ghost parameters. In this case, the posterior is

$$p(\theta_{k \leq 3}, \theta_{k > 3}, N = 3 | d) = p(\theta_{k \leq 3}, \theta_{k > 3} | d, N = 3) p(N = 3 | d). \quad (\text{B1})$$

The conditional posterior given N can be written as

$$p(\theta_{k \leq 3}, \theta_{k > 3} | d, N = 3) = \frac{\mathcal{L}(d | \theta_{k \leq 3}) \pi(\theta_{k \leq 3}) \pi(\theta_{k > 3} | \theta_{k \leq 3})}{p(N = 3 | d) / \pi(N = 3)}. \quad (\text{B2})$$

As the priors for the extra parameters are properly normalized by definition, i.e.,

$$\int p(\theta'_{k > 3} | \theta_{k \leq 3}) d\theta'_{k > 3} = 1, \quad (\text{B3})$$

¹⁴ Examples employing each of the prior types can be found at the [git repository](#).

the marginalized posterior for $\theta_{k \leq 3}$ is equivalent to the case where there are no extra parameters:

$$\begin{aligned} p(\theta_{k \leq 3}|d, N=3) &\propto \int \mathcal{L}(d|\theta_{k \leq 3})\pi(\theta_{k \leq 3})\pi(\theta_{k > 3}|\theta_{k \leq 3})d\theta_{k > 3} \\ &\propto \mathcal{L}(d|\theta_{k \leq 3})\pi(\theta_{k \leq 3}). \end{aligned} \quad (\text{B4})$$

Now we take a look at the denominator of Eq. B2. It is actually the marginal likelihood of N in transdimensional sampling:

$$\mathcal{L}(d|N=3) = p(N=3|d)/\pi(N=3). \quad (\text{B5})$$

Meanwhile, we note it is essentially a normalization factor, so the expression can be also written as

$$\begin{aligned} \mathcal{L}(d|N=3) &= \int \mathcal{L}(d|N=3, \theta_{k \leq 3}) \times \\ &\quad \pi(\theta_{k \leq 3})\pi(\theta'_{k > 3}|\theta_{k \leq 3})d\theta_{k \leq 3}d\theta_{k > 3} \\ &= \int \mathcal{L}(d|\theta_{k \leq 3})\pi(\theta_{k \leq 3})d\theta_{k \leq 3} \\ &= \mathcal{Z}_{N=3}. \end{aligned} \quad (\text{B6})$$

We make use of the fact that the priors for ghost parameters are properly normalized again.

So the model selection result of our transdimensional problem with ghost parameters is valid regardless of the inclusion of ghost parameters as the likelihood $\mathcal{L}(d|N=3)$ is correctly defined as the case without the implementation of ghost parameters.

As a comparison, the detailed balance equations of traditional reversible jump Markov chain Monte Carlo without ghost parameters is written as

$$p(\theta_{k \leq 3}|d)q(\theta'_{k \leq 3}, \theta'_{k > 3}) = p(\theta'_{k \leq 3}, \theta'_{k > 3}|d)q(\theta_{k \leq 3})\alpha, \quad (\text{B7})$$

where $p(\theta|d)$ is the target distribution, i.e., posteriors in Bayesian inference, $q(\theta)$ is the proposal for samples in Markov chain Monte Carlo sampling and α is the acceptance probability. This makes use of the trade-off between higher dimension proposals in the left-hand side and higher dimension posteriors in the right-hand side.

With the implementation of ghost parameters, we artificially add extra dimensions for posteriors and proposals in both side with the detailed balance equation written as

$$p(\theta_{k \leq 3}, \theta_{k > 3}|d)q(\theta'_{k \leq 3}, \theta'_{k > 3}) = p(\theta'_{k \leq 3}, \theta'_{k > 3}|d)q(\theta_{k \leq 3}, \theta_{k > 3})\alpha. \quad (\text{B8})$$

As we show above, in the case where $\theta_{k > 3}$ are not used in the evaluation of the likelihood, the posteriors could be written as two independent parts

$$p(\theta_{k \leq 3}, \theta_{k > 3}|d) = p(\theta_{k \leq 3}|d) \times \pi(\theta'_{k > 3}|\theta_{k \leq 3}). \quad (\text{B9})$$

Thus, if we choose a proper reversible proposal distribution to make

$$\begin{aligned} q(\theta'_{k \leq 3}, \theta'_{k > 3}) &= q(\theta'_{k \leq 3})q(\theta'_{k > 3}|\theta'_{k \leq 3}) \\ &= q(\theta'_{k \leq 3})\pi(\theta_{k > 3}|\theta_{k \leq 3}), \end{aligned}$$

then the detailed balance can be written as

$$\begin{aligned} p(\theta_{k \leq 3}|d)\pi(\theta_{k > 3}|\theta_{k \leq 3})q(\theta'_{k \leq 3}, \theta'_{k > 3}) &= \\ p(\theta'_{k \leq 3}, \theta'_{k > 3}|d)q(\theta'_{k \leq 3})\pi(\theta_{k > 3}|\theta_{k \leq 3})\alpha. \end{aligned} \quad (\text{B10})$$

This reduces to Eq.B7 where we do not implement ghost parameters. In fact, it is not necessary to set up special proposals for ghost parameters. With arbitrary proposal distributions, the sampling result with the implementation of ghost parameters will always be consistent with the situation without ghost parameters as the statistical average of acceptance rate α over the entire parameters space.

C. THE ORDER STATISTICS OF A UNIFORM DISTRIBUTION

In this Appendix, we describe the order statistics of a uniform distribution. This formalism imposes an order of different components while maintaining the uniform distribution of the unordered statistics. From [David & Nagaraja \(2004\)](#), if one draws N samples from a uniform distribution on the interval $[0, 1]$ and arranges them in ascending order, for $i < j$, the joint probability density function of the i^{th} and j^{th} samples (u, v) can be shown to be

$$f(u, v|N) = N! \frac{u^{i-1}}{(i-1)!} \frac{(v-u)^{j-i-1}}{(j-i-1)!} \frac{(1-v)^{n-j}}{(n-j)!}. \quad (\text{C11})$$

Since the u variables are ordered by construction, the conditional prior for u_i depends on u_{i-1} . But it does not depend on u_j explicitly for $j < i-1$. It is already taken care of that $u_i < u_{i-2}$ because $u_{i-1} < u_{i-2}$. We can use the fact

$$f(u_i|u_{i-1}, u_{i-2}, \dots, u_0, N) = f(u_i|u_{i-1}, N), \quad (\text{C12})$$

to sample u_i in order given Eq. C11. The joint density of all N variables is constant

$$f(u_1, u_2, \dots, u_N|N) = N!. \quad (\text{C13})$$

The normalization factor accounts for the decreased parameter space by the factor of number of permutations $N!$.

The marginalized likelihoods, i.e., evidence values of those two kinds of priors, are the same, which indicates no additional biases are introduced to our transdimensional sampling by the order statistics uniform prior. The overall evidence using uniform priors $F(u_1, u_2, \dots, u_N)$ without ordering can be calculated by

$$\mathcal{Z} = \int \mathcal{L}(d|u_1, u_2, \dots, u_N) F(u_1, u_2, \dots, u_N) du. \quad (\text{C14})$$

This integral can be refactored to be the summation of integrals in the parameter spaces by different permutations of u_i ; i.e.,

$$\begin{aligned} \mathcal{Z} &= \int \mathcal{L}(d|u_1, u_2, \dots, u_N) F(u_1, u_2, \dots, u_N) du \\ &= \int_{u_1 > u_2 > u_3 > \dots > u_N} \mathcal{L}(d|u_1, u_2, \dots, u_N) F(u_1, u_2, \dots, u_N) du + \\ &\quad \int_{u_2 > u_1 > u_3 > \dots > u_N} \mathcal{L}(d|u_1, u_2, \dots, u_N) F(u_1, u_2, \dots, u_N) du + \dots \end{aligned} \quad (\text{C15})$$

Each integral of parameter space by one permutation will lead to the same value, i.e., $\mathcal{Z}/N!$.

For the evidence obtained with the priors, $f(u_1, u_2, \dots, u_N|N)$ is distributed according to the order statistics of a uniform distribution, the integral happens in one parameter subspace, which follows the permutation. However, the density of the priors increase by a factor of $N!$, as we show in Eq. C13, so that the evidence $\mathcal{Z}'$ remains the same

$$\begin{aligned} \mathcal{Z}' &= \int_{u_1 > u_2 > u_3 > \dots > u_N} \mathcal{L}(d|u_1, u_2, \dots, u_N) f(u_1, u_2, \dots, u_N|N) du \\ &= N! \times \int_{u_1 > u_2 > u_3 > \dots > u_N} \mathcal{L}(d|u_1, u_2, \dots, u_N) F(u_1, u_2, \dots, u_N) du \\ &= N! \frac{\mathcal{Z}}{N!} \\ &= \mathcal{Z} \end{aligned} \quad (\text{C16})$$

Finally, we note that one might be tempted to design a prior where u_1 is uniform on the interval $[0, 1]$, u_2 is uniform on the interval $[0, u_1]$, and so on. Naively, one might expect this to reproduce the distribution described above, but it does not. The prior for μ_k rails against 1 for large values of k and the distribution is not equivalent to what one obtains by drawing N uniform numbers and then ordering them after the fact.

	ln Bayes factor	
N	SNR ordering	mean ordering
1	8.12	8.16
2	13.03	13.08
3	14.33	14.01
4	13.94	13.91
5	13.06	13.13
6	11.82	12.13

Table 2. Natural log Bayes factors relative to noise obtained with assumptions of different number of Gaussian pulses and different ordering schemes in priors.

D. ALTERNATIVE ORDERING SCHEMES

In this Appendix, we explore alternative ordering schemes. While our examples above employ ordering in SNR, one can, in principle, order in any component function parameter. Mathematically, the prior is the same no matter how one chooses to order the parameters, and so the ordering scheme is important only to ensure a converged result and for interpretability.

For example, let us reanalyse the superposition of Gaussian pulses but ordering by center time μ . Let us denote SNR-ordered parameters with subscript numbers $(\mu_1, \mu_2, \dots, \rho_1, \dots)$. We contrast this with μ -ordered parameters with subscript letters: $(\mu_A, \mu_B, \dots, \rho_A, \dots)$. The likelihood of the data given different ordering schemes is the same since different ordering schemes just relabel the parameters:

$$\mathcal{L}(d|\mu_1, \mu_2, \dots, \rho_1, \dots) = \mathcal{L}(d|\mu_A, \mu_B, \dots, \rho_A, \dots). \quad (\text{D17})$$

Likewise, ordering schemes do not change the shape of the prior distribution; though, they increase the overall probability density by $N!$ compared to the unordered prior. This compensates the decrease in parameter space that occurs when completely degenerate modes are eliminated with ordering. Thus, different ordering schemes lead to the same prior values. For example, if we draw samples from the SNR-ordered prior

$$\pi(\mu_1, \mu_2, \dots, \rho_1, \dots), \quad (\text{D18})$$

and convert the samples to μ -ordering, they will match samples drawn from the μ -ordered prior:

$$\pi(\mu_A, \mu_B, \dots, \rho_A, \dots). \quad (\text{D19})$$

We demonstrate this in Fig. 10.

Since the choice of ordering does not affect the likelihood or the prior, the posteriors for different ordering schemes are mathematically equivalent. In Fig. 11, we relabel the posterior samples from SNR ordering priors by ascending μ . The resultant distribution (green) is consistent with the posterior by μ -ordered priors (orange). The Bayes factors are also consistent; see Table 2.

As noted above, the fact that different ordering schemes are mathematically equivalent does not imply they are all equally useful in practice. To illustrate this, we revisit the reconstruction of GW150914, but this time employing a frequency-ordered prior. The sampling process gets stuck at $N=2$ with a natural log maximum likelihood value ~ 100 less than we obtain with SNR-ordered runs, which favor samples with $N = 4$. The sampler successfully captures the first two most significant features, i.e., the components with highest SNRs, which have frequencies around 85 Hz and 140 Hz. This result is consistent with what we find when we analyse GW150914 data with $N = 2$. However, the central frequency of the third component is $\lesssim 140$ Hz, which is unlikely to be identified in the frequency-ordered scheme because the second wavelet to be identified already has $f > 140$ Hz. Therefore it is hard for the sampler to escape from the $N = 2$ local likelihood maximum.

E. SINE-GAUSSIAN INJECTION STUDY

We perform an injection study to validate our transdimensional result of wavelets fit. The distribution of our simulated parameters is equivalent to the prior used for our GW150914 analysis with some exceptions. The distribution




Figure 10. A corner plot shows the distribution of different order schemes. Blue is ordering by SNRs and uniform means priors while orange is ordering by means and uniform SNRs. We reconstruct the distribution in green by resorting the samples from the distribution in orange by descending SNRs. Note the blue almost completely overlaps the green as we expect.

of non-ghost SNR $k \leq N$ are distributed according to Eqs. 10-12 for ordered statistics uniformly distributed on the interval $[10, 20]$ where the bulk of GW150914 posterior support lies. Ghost parameter SNRs $k > N$ are uniformly distributed on the interval $[10, 20]$. The quality factor Q is taken from a uniform distribution on the interval $[2, 40]$. We exclude the small $Q < 2$ region deliberately because it is hard to fit the delta function like signal for our sampler currently. The prior for the number of wavelets is uniform on $[3, 6]$. We use the standard LIGO O4 noise power spectral density. Also, to limit computational costs, we only consider a single detector. We expect that multiple detectors would produce similar results.

We find that 190 out of 200 injections the true value of N is strongly preferred ($\gtrsim 99\%$ credibility). In the cases where the sampler fails to correctly identify the injected number of wavelets, there is a strong preference for $N_{\text{true}} - 1$. Upon further investigation, we determine that all 10 of these injections fit the same pattern. There are two wavelets,




Figure 11. Posteriors of Gaussian pulses using different order schemes when N fixed at 3. Blue is using SNR ordering priors while orange is using mean ordering priors. We reconstruct the distribution in green by resorting the samples from distribution in blue by ascending means.

overlapping in time, which can be approximately described by one wavelet (see Fig. 12 for example). However, the reconstructed waveform produces a reasonable match to the injected signal; See Fig. 13.

Various solutions may be possible to eliminate this corner case. One option is to employ proximity priors, which prevent two wavelets from falling on top of each other so as to produce a signal that looks like one wavelet. Since wavelet models are phenomenological, it is worth designing them to be minimise the challenges for the sampler. Another possibility is to employ custom jumps so that the sampler more efficiently explores the possibility of replacing one wavelet with a superposition of two.

REFERENCES




Figure 12. Example of the failure when the sampler can't identify the number of wavelets correctly. Blue and orange represents the two separate sine-Gaussians in the simulated data. Green is the combination of these two.




Figure 13. Example of the reconstructed waveform when the sampler can't identify the number of wavelets correctly. Blue is the injected signal and red is the fitting result.

Abbott, B. P., et al. 2016a, *Phys. Rev. Lett.*, 116, 061102
— 2016b, *Phys. Rev. Lett.*, 116, 241102
— 2016c, *Phys. Rev. Lett.*, 116, 221101
— 2017, *Nature*, 551, 85
— 2018, *Phys. Rev. Lett.*, 121, 161101
Abbott, R., et al. 2021, *Astrophys. J. Lett.*, 913, L7
— 2023, *Phys. Rev. X*, 13, 011048
Acernese, F., et al. 2015, *Class. Quant. Grav.*, 32, 024001
Agazie, G., et al. 2023, *Astrophys. J. Lett.*, 956, L3
Antoniadis, J., et al. 2023, *Astron. Astrophys.*, 678, A50
Ashton, G., & Dietrich, T. 2022, *Nature Astronomy*, 6, 961
Ashton, G., et al. 2019, *Astrophys. J. Supp.*, 241, 27
Aso, Y., Michimura, Y., Somiya, K., et al. 2013, *Phys. Rev. D*, 88, 043007
Brewer, B. J., Huijser, D., & Lewis, G. F. 2015, *Mon. Not. R. Ast. Soc.*, 455, 1819

Buscicchio, R., Roebber, E., Goldstein, J. M., & Moore, C. J. 2019, *Physical Review D*, 100
Chatziioannou, K., Cornish, N., Wijngaarden, M., & Littenberg, T. B. 2021a, *Physical Review D*, 103, 044013
Chatziioannou, K., Isi, M., Haster, C.-J., & Littenberg, T. B. 2021b, *Phys. Rev. D*, 104, 044005
Cornish, N. J., & Littenberg, T. B. 2015, *Class. Quantum Grav.*, 32, 135012
Cornish, N. J., Littenberg, T. B., Bècsy, B., et al. 2021, *Phys. Rev. D*, 103, 044006
Dàlya, G., Raffai, P., & Bècsy, B. 2021, Bayesian reconstruction of gravitational-wave signals from binary black holes with nonzero eccentricities
David, H. A., & Nagaraja, H. N. 2004, *Order Statistics, Wiley Series in Probability and Statistics (Wiley)*

- Davis, D., Littenberg, T. B., Romero-Shaw, I. M., et al. 2022, *Class. Quantum Grav.*, 39, 245013
- Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022, *Astrophys. J.*, 924, 101
- Edelman, B., Farr, B., & Doctor, Z. 2023, *Astrophys. J.*, 946, 16
- Ellis, J., & Cornish, N. 2016, *Phys. Rev. D*, 93, 084048
- Green, P. J. 1995, *Biometrika*, 82, 711
- Gupta, T., & Cornish, N. 2023, *Phys. Rev. D*, 109, 064040
- Hotokezaka, K., Nakar, E., Gottlieb, O., et al. 2019, *Nature*, 3, 940
- Hourihane, S., Chatzioannou, K., Wijngaarden, M., et al. 2022, *Physical Review D*, 106, 042006
- Johnson-McDaniel, N. K., Ghosh, A., Ghonge, S., et al. 2022, *Phys. Rev. D*, 105, 044020
- Karnesis, N., Katz, M. L., Korsakova, N., Gair, J. R., & Stergioulas, N. 2023, *Eryn: A multi-purpose sampler for Bayesian inference*
- Klimenko, S., Yakushin, I., Mercer, A., & Mitselmakher, G. 2008, *Class. Quantum Grav.*, 25, 114029
- Kronland-Martinet, R., Morlet, J., & Grossmann, A. 1987, *International Journal of Pattern Recognition and Artificial Intelligence*, 1, 273
- LIGO Scientific Collaboration, Virgo Collaboration and KAGRA Collaboration. 2018, *GWTC-1 Data Release*, <https://gwosc.org/GWTC-1/>
- Littenberg, T., Cornish, N., Lackeos, K., & Robson, T. 2020, *Phys. Rev. D*, 101, 123021
- Littenberg, T. B., & Cornish, N. J. 2010, *Phys. Rev. D*, 82, 103007
- . 2015, *Phys. Rev. D*, 91, 084034
- Littenberg, T. B., Kanner, J. B., Cornish, N. J., & Millhouse, M. 2016, *Phys. Rev. D*, 94, 044050
- Liu, A., Wong, I. C. F., Leong, S. H. W., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 4149
- Lodewyckx, T., Kim, W., Lee, M. D., et al. 2011, *Journal of Mathematical Psychology*, 55, 331
- Millhouse, M., Cornish, N. J., & Littenberg, T. 2018, *Phys. Rev. D*, 97, 104057
- Miravet-Tenés, M., Florencia L. Castillo, R. D. P., Cerdá-Durán, P., & Font, J. A. 2023, *Phys. Rev. D*, 107, 103053
- Nathan, R. S., et al. 2023, *Mon. Not. R. Ast. Soc.*, 523, 4405
- Pankow, C., et al. 2018, *Phys. Rev. D*, 98, 084016
- Pannarale, F., Macas1, R., & Sutton, P. J. 2021, *Class. Quantum Grav.*, 36, 035011
- Payne, E., Hourihane, S., Golomb, J., et al. 2022, *Phys. Rev. D*, 106, 104017
- Pratten, G., et al. 2021, *Phys. Rev. D*, 103, 104056
- Raza, N., McIver, J., Dàlya, G., & Raffai, P. 2022, *Phys. Rev. D*, 106, 063014
- Reardon, D. J., et al. 2023, *Astrophys. J. Lett.*, 951, L6
- Romero-Shaw, I. M., et al. 2020, *Mon. Not. R. Ast. Soc.*, 499, 3295
- Speagle, J. S. 2020, *Mon. Not. R. Ast. Soc.*, 493, 3132
- Tiwari, V., & Fairhurst, S. 2021, *Astrophys. J. Lett.*, 913, L19
- Toubiana, A., Katz, M. L., & Gair, J. R. 2023, *Mon. Not. R. Ast. Soc.*, 524, 5844
- Vousden, W. D., Farr, W. M., & Mandel, I. 2015, *mnras*, 455, 1919
- Xu, H., et al. 2023, *Res. Astrono. Astrophys.*, 23, 075024
- Yi Shuen C. Lee, M. M., & Melatos, A. 2021, *Phys. Rev. D*, 103, 062002