# Ensemble Monte Carlo Markov Chain Methods for Testing General Relativity in Gravitational Wave Signals

Noah E. Wolfe, Jacob Golomb, Colm Talbot

(Dated: 15 May 2021)

The population of observed gravitational wave transients continues to grow, and with it, our ability to further constrain deviations from our current understanding of gravity. However, our current procedures for computing these constraints will not successfully scale with future transient catalogs. Thus, we will leverage modern statistical methods, like ensemble Monte Carlo Markov Chain sampling, to provide more efficient and more complete investigations of the parameter space of deviations from general relativity given gravitational wave observations of binary black hole mergers.

## I. INTRODUCTION

General relativity is currently our most successful theory of gravity. However, modern developments in both theoretical and observational physics may provide hints that general relativity may not be a complete theory of gravity. For example, mathematical developments in string theory and quantum gravity have the potential to provide a unified theory of gravitation across length scales. Some of the most popular alternatives to general relativity that address theoretical and observational concerns include scalar-tensor theories [10] and dynamical Cherns-Simons gravity [1].

General relativity has been rigorously tested in the non-dynamical, weak-field regime, through experiments like the Gravity Probe B experiment and time-delay measurements with the Cassini space probe [18]. It has passed these tests with flying colors. Thus, the next gravitational frontier lies in the highly dynamical, strong-field regime of compact object mergers. LIGO has already begun tests in this regime, analyzing both single-events [12, 13] and the burgeoning population of gravitational wave transients [11, 14, 15], and to date has not yet identified deviations from general relativity.

The number of observed mergers will only continue to grow, further enhancing our resolution on key parameters describing the strong-field regime. However, this also necessitates that our statistical and computational instrumentation can support larger and more complex analyses. Current population analyses like [15], based on the procedure of single-event analyses like [12], are approaching limits of reasonable computational efficiency. Thus, the goal of this work will be to improve existing analysis procedures, specifically targeted at the analysis of deviations from general relativity, with modern statistical and computational methods. This will allow us to scale our analysis as the population of observed mergers grows, and further constrain minor deviations in gravitational wave signals predicted by general relativity.

In this paper, we will outline a proposal for summer research in the LIGO SURF 2021 program. In Section II, we will provide relevant background on the physical context of astrophysical gravitational wave generation, computational and statistical methods that we will use to explore gravity in these contexts, and relevant previous work. In Section III, we will detail our particular objectives for summer 2021, and in Section IV we will detail the necessary tools to achieve these objectives. Finally, in Section VI, we propose a timeline and procedure for achieving our objectives.

## II. BACKGROUND

Gravitational waves are periodic ripples in spacetime due to asymmetric acceleration of mass. Astrophysically, they are generated by violent events involving compact objects, most notably in the context of compact object mergers. As the vast majority of current gravitational wave observations are the result of binary black hole (BBH) mergers, we will focus on BBH mergers in particular.

### A. Anatomy of a Binary Black Hole Merger

A BBH merger can be divided into three time-domains, in order: the *inspiral*, *merger*, and *ringdown* [15]. The inspiral phase begins when the black holes have formed a binary system. This phase is characterized by quasi-circular orbits, and lasts until weak-field approximations like the post-Newtonian expansion breakdown. Observationally, this frequency is usually determined to end at a gravitational wave frequency between 100 and 200 Hz [15]. Following the inspiral is the merger phase, which begins with a "plunge". In the plunge, the quasi-circular orbits suddenly become unstable, and the horizons merge. This phase is only possible to describe through numerical methods. Finally, the ringdown phase is the asymptotic relaxation of the combined black hole to a stable, isolated black hole state. This phase is well-described by analytical quasi-normal modes (QNMs).

Each regime is traditionally modeled through different mathematical and computational methods, owing to their differing physical time and length scales. The inspiral phase can be modeled analytically, through a parameterized post-Newtonian (PN) expansion of the gravitational potential. This formalism was introduced in its modern form by [17], and specified for gravitational wave emission from compact object mergers by [5] and [4]. As initially identified by [17], each term in the PN expansion can vary depending on the underlying theory of gravity and so is one of the key quantitative models for differentiating between alternative theories of gravity. The merger phase can only be fully understood through numerical relativity simulations, as it is described with nonlinear field equations in the chosen theory of gravity. See, for example, the work of [9] to model the gravitational wave signal from a BBH merger in dynamical Cherns-Simons gravity, which identifies a dipole component in the gravitational radiation produced. Finally, quasi-normal modes can be understood analytically, and are specified by mode frequencies and damping times that depend on the assumed model of gravity [15? ].

### B. Bayesian Inference

For our description of Bayesian inference, we will begin by following the description and notation of [16]. In Bayesian inference, we consider a set of parameters $\theta$ in the context of the data $d$. For a physical situation of interest, we will have a model parameterized in terms of $\theta$. For example, in this work, we will have a set of parameters which includes black hole binary properties (like mass and spin), with additional parameters to denote deviations from the predictions of general relativity. Our goal is to constrain these parameters given the data; statistically, we want to construct the posterior distribution $p(\theta|d)$, read as "the probability of getting a particular set of parameter values given observed data."

Bayesian inference allows us to construct the posterior distribution via Bayes' theorem, which relates $p(\theta|d)$ to our observations:

$$p(\theta|d) = \frac{\mathcal{L}(d|\theta)\pi(\theta)}{\mathcal{Z}} \tag{1}$$

where $\mathcal{L}(d|\theta)$ is known as the likelihood, $\pi(\theta)$ is the prior distribution. The normalization factor $\mathcal{Z}$ is known as the evidence, i.e. the probability of observing the data given the parameteric model we choose:

$$\mathcal{Z} \equiv \int d\theta \mathcal{L}(d|\theta)\pi(\theta) \tag{2}$$

The likelihood is a model for our observations under different parameters $\theta$, and includes a model for noise; for gravitational wave observations, we most commonly use a likelihood that assumes Gaussian detector noise. The prior distribution is, effectively, an initial guess for the distribution of parameter values; when making this guess, we need to ensure that we do not bias the posterior distribution, and so may choose, for example, a uniform prior distribution.

We observe that $p(\theta|d)$ provides a distribution on the entire (multi-dimensional) set of parameters $\theta$. To extract information on specific parameters of interest $\theta_i$, we must "marginalize", i.e. integrate, over the rest of the parameters:

$$p(\theta_i|d) = \int \left( \prod_{k\neq i} d\theta_k \right) p(\theta|d) \tag{3}$$

We observe that this integration may be difficult to compute through standard numerical methods, especially if we have a high-dimensional parameter space. One common method is to use a Monte Carlo Markov Chain (MCMC) approach to approximate the posterior distribution. In this approach, a series of "walkers" explore the space of parameters $\theta$ in such a way that, given enough iterations, their paths will produce a representative sample of the posterior distribution.

Finally, in order to combine observations from multiple events that we suspect are part of a contiguous population (e.g. a population of binary black hole mergers) we will use a hierarchical analysis. In a hierarchical analysis, we assume that the parameters we infer given each observation are drawn from the same underlying distribution. For example, we will see that we assume parameterization of deviations from general relativity are common among all binary black hole mergers, and thus these deviations are drawn from the same posterior *hyper*parameter distribution. The mechanics of hierarchical inference rely on the same mechanics of standard Bayesian inference, wherein we seek to infer hyper-parameters given a population of observations through application of Bayes theorem and numerical techniques like MCMC.

### C. Previous Parametric Tests of General Relativity

A key challenge in comparing the results of general relativity with alternative theories of gravity in the context of gravitational wave emission is that alternative theories of gravity are often ill-posed in the strong field regime [15]. Therefore, we will ultimately follow the approach of previous methods in using a mathematical framework that quantifies deviations *away* from the predictions of general relativity, as opposed to measuring agreement with any particular alternative theory of gravitation. This additionally affords flexibility in testing multiple theories simultaneously.

Previous work with this approach includes tests with single events (a single BBH [12] and a single neutron star-neutron star merger [13]), as well as tests with observing run O1 [11], the first gravitational wave transient catalog (GWTC-1) [14], and most recently with GWTC-2 [15]. We will briefly focus on the last study, as we intend to improve upon this analysis.

In Section V, subsection A, of [15], the authors constrain a set of parameters $\delta\hat{\varphi}_i$ that denote deviations from coefficients in the PN expansion during the inspiral phase (see equation 4 of [15]). Using, alternatively, the `SEOBNRv4_ROM` and `IMRPhenomPv2` models, and

`LALInference`, [15] varies each parameter $\hat{\varphi}_i$. Their approach only varies one parameter at a time, as they found that multiple parameter variation reduced the information content of their posteriors. From the modifications to the predicted waveforms with these variations, they construct posterior distributions on the parameterized violations of general relativity (see Figure 6 of [15] in particular). Ultimately, they find no evidence for violations of general relativity.

## III. OBJECTIVES

There are two key limitations in the method employed by [15] that we seek to remedy.

1. The method in [15] is computationally expensive. They begin with a set of 15 parameters from general relativity, with an additional variation parameter $\delta\hat{\phi}_i$. So, in total, they must infer a set of 16 parameters for each variation $\delta\hat{\phi}_i$, requiring many evaluations of the signal model in order to explore the parameter space of physically relevant $\delta\hat{\phi}$.

2. The method in [15] only varies one parameter $\hat{\varphi}_i$ at a time, due to the computational expensive involved in one parameter variation and an observed loss of information in posterior distributions with multiple parameter variation. Although there is reason to believe that variations in one parameter may be sufficient [15], there remains the possibility that violations of general relativity will only appear in multi-coefficient deviations from the PN expansion.

We intend to use advanced statistical techniques to simultaneously improve the efficiency of inspiral-phase analyses and effectively increase the dimensionality of inspiral-phase analyses, in the context of identifying violations of general relativity from populations of BBH mergers.

## IV. METHODS AND PROCEDURES

To achieve our first objective, we will use ensemble Monte Carlo-Markov Chain (MCMC) methods. These build upon existing MCMC methods by replacing a single walker, as used in traditional approaches like the Metropolis-Hastings algorithm, with an ensemble of walkers that explore the parameter space in parallel [3]. A key feature of this approach is that we can reduce the number of samples we need to generate for our Bayesian inference methods, as at any single step the ensemble of walkers provides us a representation of the target posterior distribution. Additionally, ensemble MCMC methods allow us to use the posterior distribution assuming GR is correct as the initial condition for our walkers. We will conduct an ensemble MCMC analysis over the set of general relativity and PN expansion correction parameters with the `emcee` python code, implemented through the `BILBY` API. This will allow us to efficiently compute posterior distributions on the correction parameters.

We will achieve our second objective after demonstrating our capacity to replicate the single-parameter variation analysis of [15] with an ensemble MCMC method. With a more efficient method to construct posterior distributions on each parameter $\delta\hat{\phi}_i$, we will not be constrained by computation time, and thus be able to sample over multiple parameter variations as well.

## V. PROGRESS UPDATE

### A. Completed Objectives

To date, I have completed key startup activities and tutorials relevant to contextual astrophysics and software tools necessary for completing our project objectives. I have attended four LIGO Gravitational Wave Science Seminars, covering the basics principles of LIGO detector science and gravitational wave astrophysics. Simultaneously, I have completed the Gravitational Wave Open Data Workshop 4 tutorials [6]. These tutorials described the basic usage of relevant software tools for gravitational wave analysis, from time- and frequency-domain signal processing to Bayesian parameter estimation. In these tutorials, I learned how to use the `PyCBC` software library to generate model waveforms from general relativity approximants. I also learned the basics of parameter estimation for gravitational wave signals with the `BILBY` software library. These tutorials were completed on a personal laptop running Ubuntu 18.04, with the `igwn-py38` Python environment. Work for this project can be found in the associated GitLab repository [7].

#### *1. TaylorF2 Waveform Generation*

I have successfully implemented the `TaylorF2` waveform approximant in Python. This is a PN expansion of the gravitational wave strain in the frequency domain, broken into expansions of the amplitude and phase of the signal.

$$\tilde{h}(f) = A_{\rm PN} e^{-i\psi(f)} \tag{4}$$

where $A$ is the amplitude and $\psi$ is the phase; both may depend on intrinsic or extrinsic parameters of the BBH. The form of the amplitude expansion, in geometric units ($c = G = 1$) is

$$A_{\rm PN} = A_0 \sum_i \mathcal{A}_i (\pi f)^{i/3} \tag{5}$$

Typically, for `TaylorF2`, only one term in the expansion is used. This yields, with physical units,

$$A_{\rm PN} = \sqrt{\frac{5}{24}\frac{\pi}{\eta}}\frac{m_1 m_2}{d_L}\left(\pi M f \frac{G}{c^3}\right)^{-7/6} \tag{6}$$

where $m_1, m_2$ are the component masses of the BBH (by convention, $m_1 \geq m_2$), $M = m_1 + m_2$, $\eta = m_1 m_2/M^2$ is the symmetric mass ratio, and $d_L$ is the luminosity distance to the BBH.

The phase is calculated with an expansion of the form to 3.5PN order, adapted from [2],

$$\psi(f) = -\frac{\pi}{4} + \frac{3}{128\eta v^5}\sum_i^7 \varphi_i v^i \tag{7}$$

where $v = \left(\pi M f G c^{-3}\right)^{1/3}$ (we nominally set the time, $t_c$, and phase, $\phi_c$, of coalescence to zero). The coefficients $\varphi_i$ can be found in Equation 3.18 of [2]; for example,

$$\varphi_1 = v^{-1} \tag{8}$$

$$\varphi_2 = \frac{20}{9}\left(\frac{743}{336} + \frac{11}{4}\eta\right) \tag{9}$$

Although this approximant is already implemented in `PyCBC`, I am working to manually generate these waveforms so that I can directly add non-GR corrections to the phase terms, effectively yielding

$$\psi(f) = -\frac{\pi}{4} + \frac{3}{128\eta v^5}\sum_i^7 \left(\varphi_i + \delta\varphi_i\right) v^i \tag{10}$$

with non-GR corrections $\delta\varphi_i$. In future, I could apply similar corrections to the amplitude terms $\mathcal{A}_i$, although LIGO is less sensitive to these deviations.

I have successfully replicated the amplitude evolution, as shown in Figure 1, as well as the phase evolution, as shown in Figures 2 and 3. This progress is notable, as of the last interim report I had not yet successfully replicated the phase evolution.

### 2. *Overlap Calculation*

With correct `TaylorF2` waveforms in hand, I then calculated the overlap between GR (i.e. `PyCBC` implementation) waveforms and waveforms generated with non-GR corrections $\delta\varphi_i$. Ultimately, our goal is to enforce a minimum overlap condition when sampling for non-GR parameters, to make our final analysis more efficient.

The complex overlap between two frequency-domain waveforms, $\tilde{h}_1(f), \tilde{h}_2(f)$, is calculated in the following manner

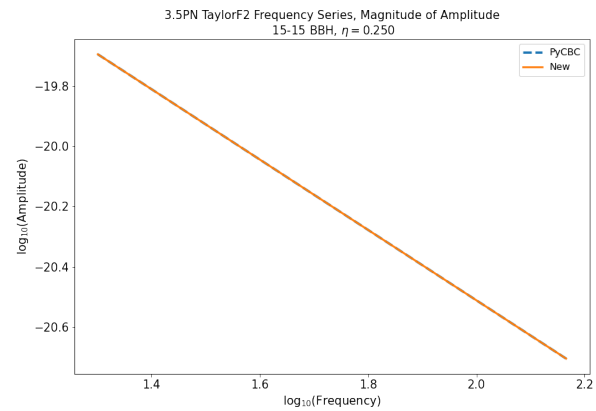


FIG. 1: Log-log plot of waveform amplitude verus frequency, demonstrating the agreement between our manual calculation of the `TaylorF2` amplitude and that implemented in `PyCBC`.

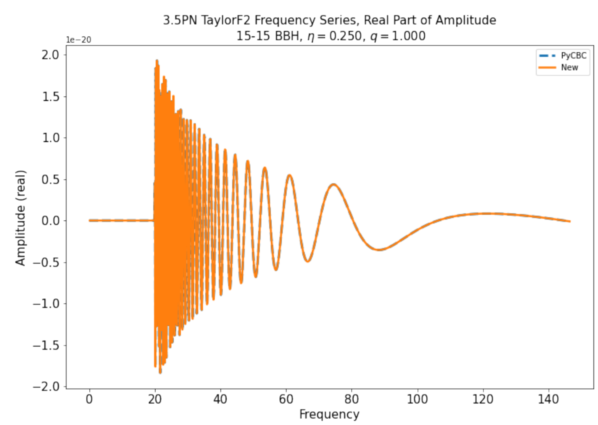


FIG. 2: Plot of the real part of the waveform amplitude verus frequency, demonstrating agreement between our manual calculation of the `TaylorF2` amplitude and that implemented in `PyCBC`.

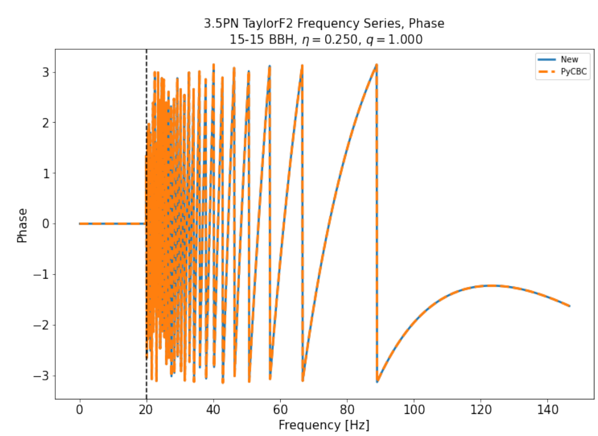


FIG. 3: Plot of the phase evolution of our manual calculation of the `TaylorF2` waveform and that implemented in `PyCBC`.

$$\mathcal{O} = \frac{\left\langle \tilde{h}_1(f), \tilde{h}_2(f) \right\rangle}{\sqrt{\left\langle \tilde{h}_1(f), \tilde{h}_1(f) \right\rangle \left\langle \tilde{h}_2(f), \tilde{h}_2(f) \right\rangle}} \tag{11}$$

where $\langle \cdot, \cdot \rangle$ denotes an inner product of the form

$$\left\langle \tilde{h}_1(f), \tilde{h}_2(f) \right\rangle = 4\Delta f \sum_i^N \frac{\tilde{h}_{1,i}\tilde{h}^*_{2,i}}{s_i} \tag{12}$$

where we sample each waveform, as well as a detector power spectral density $s$, at $N$ discrete sampling frequencies, evenly spaced by $\Delta f$. I performed an initial investigation of how the overlap varies with mass ratio $q$ at fixed values of the total mass $M$, qualitatively identifying how a minimum overlap of $\mathcal{O} \geq 0.9$ would limit the parameter space. As shown in Figure 4, I observed that as the total mass $M$ increases, so does the spread in $\mathcal{O}$ versus $\delta\varphi_2$, meaning that more values of $\delta\varphi_2$ would be allowed with a cutoff of $\mathcal{O} \geq 0.9$.

#### 3. *Hybrid Sampling with a Generalized Normal Distribution*

Finally, I have also implemented a basic version of the hybrid sampling method that we hope to ultimately apply to parameter estimation for variations $\delta\varphi_2$ in BBH waveforms. Broadly, our hybrid method first uses a nested sampling approach to estimate GR parametetrs like mass and spin, and uses the nested sampling results as the input to an ensemble method for estimating the remaining non-GR parameters.

Our basic test cases uses the generalized normal distribution to illustrate this approach. The generalized normal distribution has the following PDF:

$$P(x) = \frac{\beta}{2\alpha\Gamma(1/\beta)} e^{-(|x-\mu|/\alpha)^\beta} \tag{13}$$

where $\alpha$ takes a role akin to the standard deviation $\sigma$ in a normal distribution. We note that, in the limit of $\beta = 2$, this becomes a normal distribution (and we have $\alpha = \sqrt{2}\sigma$).

Starting with the nested sampler `dynesty`, I inferred the mean $\mu$ and standard deviation $\sigma$ of a normal distribution generated with $\mu = 3$ and $\sigma = 4$. The results of this first step are shown in Figure 5.

One key output of `dynesty` are the "posterior weights" $p_i$ that determine how each sample of $\mu$, $\sigma$ is weighted in constructing the final posterior. I constructed tempered versions of these weights by raising $p_i$ to a series of powers $\beta_T$. The tempered posterior weights are shown in Figure 6; we note that $\beta_T = 1$ corresponds to the original posterior weights and in the limit $\beta_T = 0$ we recover the prior.

For each $\beta_T$, I then performed rejection sampling on the normalized weights $p_i^{\beta_T} / \sum_i p_i^{\beta_T}$ to obtain representative initial positions for parallel tempered walkers in `ptemcee`. The results of this analysis are shown in 7. From this figure, we achieve $\beta = 2$ for a normal distribution, as expected, and approximately recover $\mu = 3$ and $\sigma = \alpha/\sqrt{2} = 4$ ($5.67/\sqrt{2} \approx 4$).

### B. Current Objectives

Having implemented the basic machinery for generating non-GR waveforms and using a hybrid sampling approach, my current objective is to combine these tools. I am currently working to combine these tools with the `BILBY` API; in particular, I am developing a modified source method that will return non-GR waveforms given variations $\delta\varphi_i$, in the style of existing `BILBY` source methods like `lal_binary_black_hole` [8]. Then, I will use this modified source method to test parameter estimation with our hybrid sampling method given an injected GR signal, and an injected non-GR signal.

### C. Future Plans and Expected Challenges

Immediately, I will validate my source function using the `dynesty` sampler on a purely GR waveform, inferring the luminosity distance and masses of the system. I will inspect the time required to evaluate the likelihood function with this source function, and perform any speedup improvements that I identify in concert with my mentors.

Following the completion and validation of the source method, I will perform parameter estimation on injected BBH signals for the luminosity distance and masses of the system using the hybrid sampling scheme for both pure GR and non-GR waveforms. I will validate these results by using the `dynesty` sampler alone to estimate the same set of parameters. Finally, given success in these initial tests, I will perform the same parameter estimation on the full suite of 15 GR parameters plus 8 non-GR parameters using `bilby_pipe` on the LIGO computing cluster.

## VI. WORK PLAN

There will be three main phases to the project. In the first phase, we will attempt to replicate the results of [15] using the ensemble MCMC approach on a single event, implemented via `BILBY` and `emcee`. Then, we will repeat this analysis with a hierarchical approach to combine data from multiple observations. In the second phase, we will extend this analysis to include simultaneous variations of multiple parameters, and explore methods to ensure that our posterior distributions remain informative. In the third and final phase, given time, we will work to implement this with more recent parameterizations of

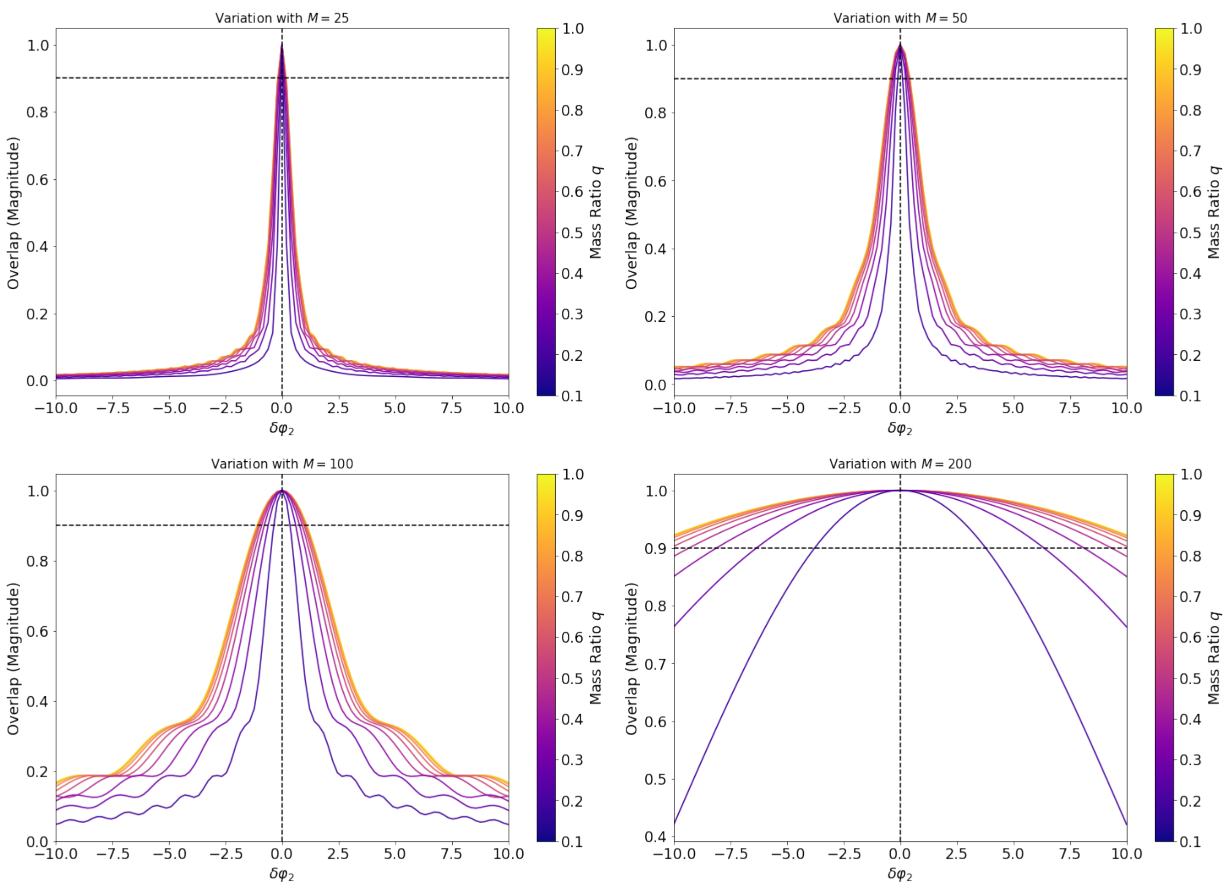


FIG. 4: Overlap with respect to variations in the mass ratio $q$, at four fixed total masses $M$. The horizontal dotted line shows the cutoff $\mathcal{O} = 0.9$, and the vertical dotted line shows $\delta\varphi_2 = 0$.

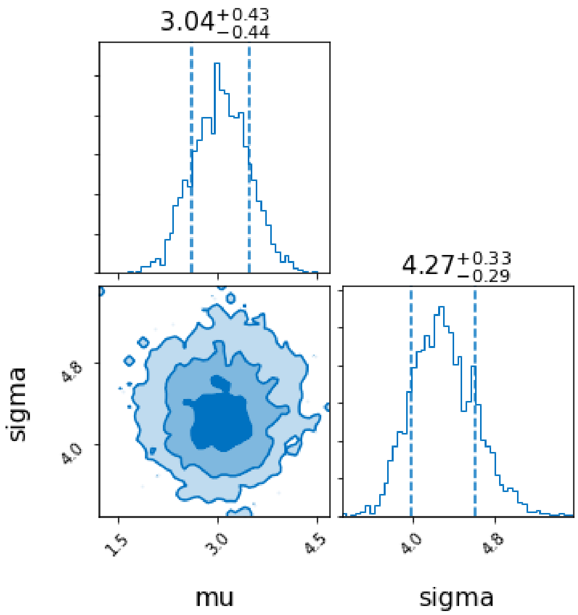


FIG. 5: Marginal and joint distributions on $\mu$ and $\sigma$ inferred with `dynesty`.

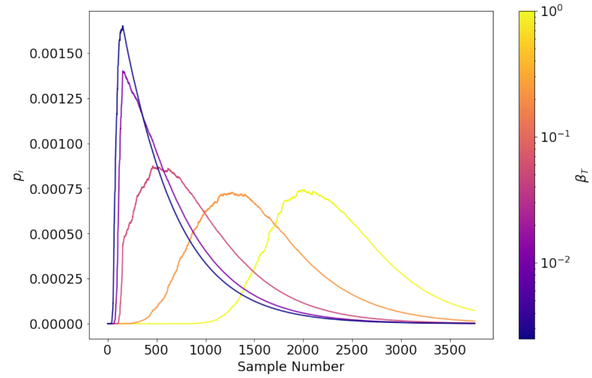


FIG. 6: Tempered posterior weights at five temperatures, chosen by the `ptemcee` `make_ladder` method.

deviations from general relativity, like the parameterized post-Einsteinian framework [19]. The timeline for this work plan can be found in Table I.

| **Date** | **Event** |
|---|---|
| June 15, 2021 | Start of LIGO SURF 2021 |
| June 21 - 25, 2021 (Week 2) | Basic tutorials in `emcee`, BILBY. Implement emcee with a test data set. |
| July 5 - 11, 2021 (Week 4) | Interim Report 1. Successful `emcee` & `BILBY` single parameter inference for a single event. |
| July 12 - 16, 2021 (Week 5) | Successful `emcee` & `BILBY` hierarchical analysis on multiple events. |
| July 26 - August 1, 2021 (Week 7) | Interim Report 2. |
| August 9 - August 13, 2021 (Week 9) | Implement multiple simultaneous parameter variations, including relevant physical posterior constraints. |
| August 18 - 20, 2021 (Week 10) | Final Presentation. |
| September 24, 2021 | Deadline for submitting final report. |

TABLE I: Timeline for the work plan.

[1] Stephon Alexander and Nicolás Yunes. Chern-Simons modified general relativity. , 480(1-2):1–55, August 2009.
[2] Alessandra Buonanno, Bala R. Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev. D*, 80(8):084043, October 2009.
[3] Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. emcee: The MCMC Hammer. , 125(925):306, March 2013.
[4] T. G. F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, J. Veitch, K. Grover, T. Sidery, R. Sturani, and A. Vecchio. Towards a generic test of the strong field dynamics of general relativity using compact binary coalescence. *Phys. Rev. D*, 85(8):082003, April 2012.
[5] Chandra Kant Mishra, K. G. Arun, Bala R. Iyer, and B. S. Sathyaprakash. Parametrized tests of post-Newtonian theory using Advanced LIGO and Einstein Telescope. *Phys. Rev. D*, 82(6):064010, September 2010.
[6] `https://www.gw-openscience.org/static/workshop4/`.
[7] `https://git.ligo.org/noah.wolfe/beyond-gr-pe`.
[8] https://git.ligo.org/lscsoft/bilby/-/blob/master/bilby/gw/source.py.
[9] Maria Okounkova, Leo C. Stein, Mark A. Scheel, and Daniel A. Hemberger. Numerical binary black hole mergers in dynamical Chern-Simons gravity: Scalar field. *Phys. Rev. D*, 96(4):044020, 2017.
[10] Israel Quiros. Selected topics in scalar-tensor theories and beyond. *International Journal of Modern Physics D*, 28(7):1930012–156, January 2019.
[11] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, et al. Binary black hole mergers in the first advanced ligo observing run. *Phys. Rev. X*, 6:041015, Oct 2016.
[12] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, et al. Tests of general relativity with gw150914. *Phys. Rev. Lett.*, 116:221101, May 2016.
[13] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, et al. Tests of general relativity with gw170817. *Phys. Rev. Lett.*, 123:011102, Jul 2019.
[14] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, et al. Tests of general relativity with the binary black hole signals from the ligo-virgo catalog

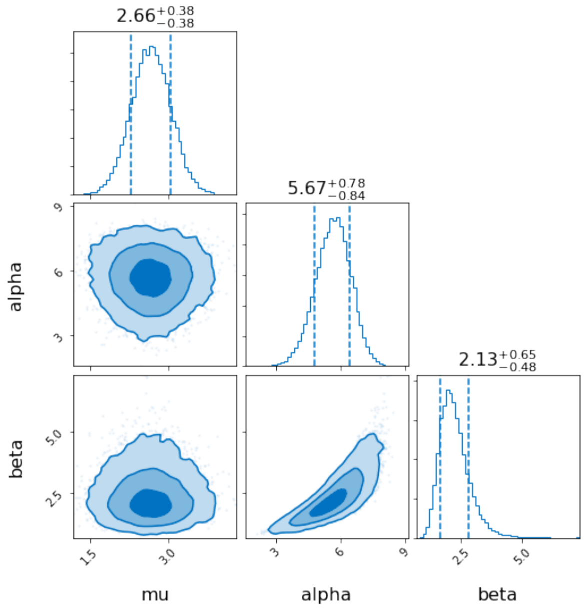


FIG. 7: Marginal and joint distributions on $\mu$, $\alpha$, and $\beta$ inferred with `ptemcee`, using tempered `dynesty` input weights.


gwtc-1. *Phys. Rev. D*, 100:104036, Nov 2019.

[15] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, et al. Tests of General Relativity with Binary Black Holes from the second LIGO-Virgo Gravitational-Wave Transient Catalog. *arXiv e-prints*, page arXiv:2010.14529, October 2020.

[16] Eric Thrane and Colm Talbot. An introduction to Bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models. , 36:e010, March 2019.

[17] Clifford M. Will. Theoretical Frameworks for Testing Relativistic Gravity. II. Parametrized Post-Newtonian Hydrodynamics, and the Nordtvedt Effect. *Astrophys. J.*, 163:611, February 1971.

[18] Clifford M. Will. The Confrontation between General Relativity and Experiment. *Living Reviews in Relativity*, 17(1):4, December 2014.

[19] Nicolás Yunes and Frans Pretorius. Fundamental theoretical bias in gravitational wave astrophysics and the parametrized post-Einsteinian framework. *Phys. Rev. D*, 80(12):122003, December 2009.