A Test of Universal Isotropy With Gravitational Wave Sources

Mor Evron, Simone Mastrogiovanni

July 30, 2024

Abstract

We develop angular models to constrain anisotropies in the sky distribution of compact binary merger events. Beginning with a simple dipole model, we simulate upper bounds on the dipole magnitude with an increasing number of observations. We find that as the number of simulated samples increases, the upper bound decreases according to the inverse square root of the number of samples. Using current GW events from LVK, we derive an upper bound from real data that agrees with the simulation and calculate a Bayes factor of 1.12 in preference of isotropy over the dipole model. We also develop a quadrupole model, for which we devise a value for the upper limit and perform a similar analysis, arriving at the same $\frac{1}{\sqrt{N}}$ relationship.

1 Bayesian Statistics

For many years, the standard statistical approach to scientific methods has been that of frequentist statistics. Frequentist philosophy rationalizes statistics and probabilities as the results of repeated experiments. In other words, should an experiment be repeated a large number of times, the data should obey a distribution with some precisely defined parameters (e.g. mean, variance, etc.). Bayesian statistics presents an alternate mode of thinking regarding these distribution parameters. In Bayesian philosophy, the parameters themselves follow a distribution. When scientific experiments are performed, the resulting population parameters are drawn from this distribution. More rigorously, Bayesian statistics is derived from Bayes's Theorem. Begin with:

$$P(A, B) = P(A|B)P(B)$$

Similarly,

$$P(A, B) = P(B|A)P(A)$$

Setting equal and rearranging, we arrive at Bayes's Theorem

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

From here, let's say that A is the distribution parameters θ and B the data collected x. We rewrite as:

$$P(\theta|x) = \frac{P(x|\theta)P(\theta)}{P(x)} \tag{1}$$

The left hand side of the equation becomes what is known as a **posterior distribution**. It is the distribution of the parameters given the collected data. From here, $P(x|\theta)$ is typically rewritten as $\mathcal{L}(x|\theta)$ and is known as the **likelihood function**, which represents the probability of the data being collected given a set value for the parameters. $P(\theta)$ is rewritten as $\pi(\theta)$ and referred to as the **prior distribution**. It represents prior knowledge of the parameters inserted by researchers. Last, $P(x)$ is the **evidence**. Because it behaves as a normalization constant, the governing equation is often written as:

$$p(\theta|x) \propto \mathcal{L}(x|\theta)\pi(\theta)$$

The implementation of Bayesian statistical methods can be beneficial for many scientific fields. It allows for a rather straightforward method of model comparison, which will be discussed, and it is more consistent philosophically with how researchers conduct experiments. When planning and performing experiments, researchers always have some educated preconceptions regarding the possible values for a parameter, hence the existence of priors. Additionally, studies are often expanded upon after their completion. When this occurs, it is common practice in Bayesian methodology to utilize the posterior distributions acquired in previous experiments as priors for subsequent ones. Historically, the primary roadblock preventing Bayesian statistics from entering the mainstream was computational power, which is not as major of an obstacle now as it once was. With modern computing capabilities, the last few decades have seen Bayesian statistics reach a level of practicality for mainstream science. Due to the advantages is presents in addition to its growing accessibility and practicality, Bayesian statistics is quickly gaining popularity. As the scientific community continues advancing, it is possible that the Bayesian approach may one day surpass its frequentist counterpart.

In gravitational wave astronomy, the Bayesian approach provides an additional benefit. Gravitational wave detections, though greatly increasing in number, are still very few and far between. Therefore, it is illogical to consider the frequentist idea of 'repeated experiment' when experiments cannot be repeated or reasonably replicated. In gravitational wave experiment, Bayesian statistics allows researchers to maximize the value of the little data available for analysis.

Bayes Factor

To compare models M_1 and M_2, it is sufficient to evaluate

$$\frac{p(M_1|x)}{p(M_2|x)} = \frac{\mathcal{L}(x|M_1)}{\mathcal{L}(x|M_2)} \frac{\pi(M_1)}{\pi(M_2)} \tag{2}$$

observing that $p(x)$ is independent of model and therefore cancels. Notice that the $\mathcal{L}(x|M)$ terms in Eq. 2 are actually identical to the $P(x)$ evidence term from Eq. 1, as there is a model implicit throughout all of Eq. 1.

Practically speaking, $\mathcal{L}(x|M)$ is numerically straightforward to calculate; it is just the normalization factor for $\mathcal{L}(x|\theta, M)\pi(\theta, M)$, meaning $\mathcal{L}(x|M) = \int \mathcal{L}(x|\theta, M)\pi(\theta, M)d\theta$.

The prior ratio represents prior belief in the models, which is simply 1 if there is no preference. The likelihood ratio is known as the Bayes factor $\mathcal{B}^{M_1}_{M_2}$ and is used in model comparison.

2 ICAROGW and Hierarchical Bayesian Inference

When studying population properties, gravitational wave data sets are exceptionally complicated to work with and present a number of challenges to researchers. The two most significant of these challenges are the following:

> **Selection Bias:** The probability of detecting a gravitational wave event is not constant for every event. Detector sensitivities introduce biases in which events may be detected. These biases must be accounted for, or statistical methods will return inaccurate results.
>
> **Noise:** Gravitational wave data is *heterogeneous*, meaning there is noise involved that requires advanced statistical methods to work with.

Currently, many population inference methods employ the use of **hierarchical Bayesian inference**, which allows us to analyze data that is both biased and heterogeneous. **I**nferring **C**osmology and **A**st**R**ophysics with **O**bservations of **G**ravitational **W**aves (ICAROGW) [6] is a Python package developed to perform hierarchical Bayesian inference using gravitational wave data produced by compact binary coalescences (CBCs). More specifically, ICAROGW allows us to produce a value of the hierarchical likelihood for a CBC dataset. To do this, ICAROGW requires three user inputs:

1. A set of posterior samples from true gravitational wave observations.

2. A set of injections (simulated data) to evaluate selection biases.

3. A model for parameter distribution, also known as a rate.

From Ref. [7], the hierarchical likelihood can be expressed as:

$$\mathcal{L}(\{x\}|\Lambda) \propto e^{-N_{\rm exp}(\Lambda)} \prod_{i=1}^{N_{\rm obs}} T_{\rm obs} \int \mathcal{L}(x_i|\theta,\Lambda) \frac{dN}{dtd\theta} d\theta$$

where $\{x\}$ is the set of data, x_i are individual data points, θ is the source parameters, Λ is the hyperparameters, $T_{\rm obs}$ is the observation time, $N_{\rm obs}$ is the number of observed GW events, $N_{\rm exp}$ is the expected number of GW events, and $\frac{dN}{dtd\theta}$ is the event rate.

The two parameters of interest for this study are α and δ, which define angular sky position similarly to the spherical coordinates ϕ and θ. They will be discussed in section 3.

2.1 Single Event Likelihood

The right half of this expression is the rate integral of the single event likelihood. It is atypical to know the actual closed form of this likelihood, and instead we are provided with $N_{s,i}$ parameter estimation (PE) samples. By sampling from the set of posterior data provided, it is possible to approximate the likelihood integral. The posterior data are sampled from $p(\theta|x_i,\Lambda) \propto \mathcal{L}(x_i|\theta,\Lambda)\pi_{\rm PE}(\theta|\Lambda)$, meaning the integral can be rearranged to be

$$\int \mathcal{L}(x_i|\theta,\Lambda) \frac{dN}{dtd\theta}(\Lambda) d\theta \propto \int \frac{p(\theta|x_i,\Lambda)}{\pi_{\rm PE}(\theta|\Lambda)} \frac{dN}{dtd\theta}(\Lambda) d\theta$$

This is an expectation value of $\frac{1}{\pi_{\rm PE}(\theta|\Lambda)} \frac{dN}{dtd\theta}(\Lambda)$ over θ space, which can be estimated using

$$\frac{1}{N_{s,i}} \sum_{j=1}^{N_{s,i}} \frac{1}{\pi_{\rm PE}(\theta_{i,j}|\Lambda)} \left. \frac{dN}{dtd\theta}(\Lambda) \right|_{i,j}$$

where i represents the event in question and j is the index for the posterior samples that are summed over. It is typical to define a weight $w_{i,j}$ that is equivalent to the expression within the summation, and so the whole integral is typically written as

$$\int \mathcal{L}(x_i|\theta,\Lambda) \frac{dN}{dtd\theta}(\Lambda) d\theta \approx \frac{1}{N_{s,i}} \sum_{j=1}^{N_{s,i}} w_{i,j}$$

2.2 Expected Number of Events

We begin with the assumption that merger rate does not depend on when the data is taken, only how long the data was taken for. Then it is possible to write

$$N_{\rm exp}(\Lambda) = T_{obs} \int p_{\rm det}(\theta,\Lambda) \frac{dN}{dtd\theta}(\Lambda) d\theta$$

where $p_{\rm det}$ is the probability of detecting a GW event given its parameters and the underlying hyperparameters. Unfortunately, detection probability is not an easily accessible quantity, and so must be estimated. This is done through the use of *injections*. Injections are artificial data produced from some prior $\pi_{\rm inj}(\theta)$. ICAROGW takes $N_{\rm det}$ injections out of an overall generated $N_{\rm gen}$ injections and, using the prior, estimates

$$\begin{aligned} N_{\rm exp}(\Lambda) &\approx \frac{T_{\rm obs}}{N_{\rm gen}} \sum_{j=1}^{N_{\rm det}} \frac{1}{\pi_{\rm inj}(\theta_j} \left. \frac{dN}{dtd\theta}(\Lambda) \right|_j \\ &\equiv \frac{T_{\rm obs}}{N_{\rm gen}} \sum_{j=1}^{N_{\rm det}} s_j \end{aligned}$$

once again defining a weight for the value within the summation. Here, the index j indicates the detected injections.

2.3 Putting it Together

We now have everything we need to write an expression for the hierarchical likelihood. It is as follows:

$$\mathcal{L}(\{x\}|\Lambda) \approx e^{-\frac{T_{\rm obs}}{N_{\rm gen}} \sum_{j=1}^{N_{\rm det}} s_j} \prod_{i=1}^{N_{\rm obs}} \frac{T_{\rm obs}}{N_{s,i}} \sum_{j=1}^{N_{s,i}} w_{i,j}$$

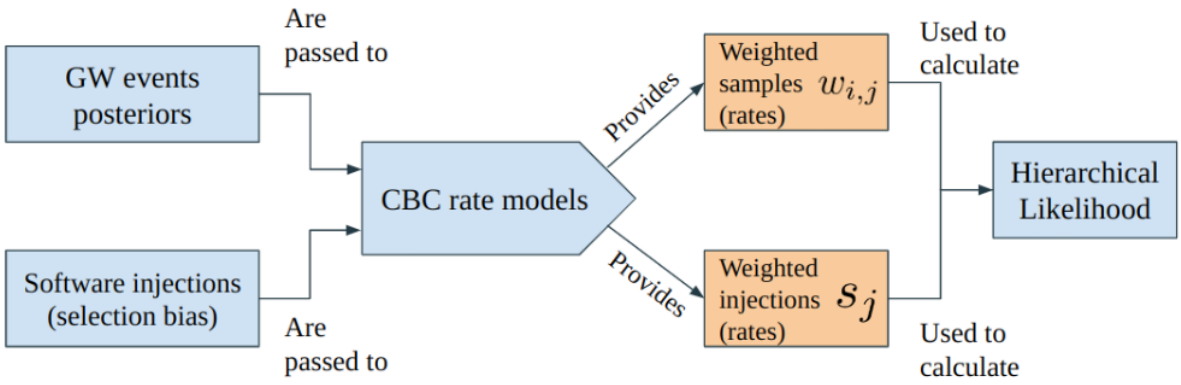

Figure 1: The logical structure of ICAROGW, lifted from Ref. [6]

However, it is much more common practice to work with the logarithm of the likelihood, which also simplifies this equation visually.

$$\ln \mathcal{L}(\{x\}|\Lambda) \approx -\frac{T_{\rm obs}}{N_{\rm gen}} \sum_{j=1}^{N_{\rm det}} s_j + \sum_{i=1}^{N_{\rm obs}} \ln \left(\frac{T_{\rm obs}}{N_{s,i}} \sum_{j=1}^{N_{s,i}} w_{i,j} \right) \tag{3}$$

Fig. 1 summarizes the overall structure of ICAROGW. The PE posteriors and injections are passed to the CBC rate model provided by the user, which ICAROGW uses to calculate the two weights defined in this section. The weights are then used in calculation of the hierarchical likelihood.

3 Angular Isotropy

As the number of detected gravitational wave events grows, it becomes increasingly feasible to study the parameters that characterize these events. The parameters of interest in this study simply define the sky location of the merger events. They are defined similarly (but not identically) to the typical ϕ and θ of spherical coordinates.

Right Ascension, represented with the letter α, is defined as the typical ϕ angle from spherical coordinates, ranging from 0 to 2π radians.

Declination, written as δ, is similar to the typical θ angle. However, the 0 angle is at the equator, $\frac{\pi}{2}$ radians is at the north pole, and $-\frac{\pi}{2}$ radians is at the south pole. Together, these two coordinates define a single point in the sky.

It is typical to refer to both angles together as Ω.

The simplest model of angular dependence is isotropy, pictured in the top left of Fig. 2, under which all sky positions are equally likely. In an isotropic model,

$$p_{\rm iso}(\Omega) = \frac{1}{4\pi}$$

It seems quite logical that gravitational wave events be distributed this way. After all, there is nothing special about the Earth's position in the universe, so merger locations should be random. Yet, there may be a case for some slight anisotropies, and studying them can be beneficial for multiple reasons. Investigation into isotropy can assist in:

1. Studying clustering around galaxies and observing which types of galaxies are most likely to contain events.

2. Detection of a kinematic dipole caused by the Earth's motion through the universe [5].

3. Probing new physics through detection of unexplained anisotropy.

3.1 The Dipole Model

The most basic of the anisotropic models is the dipole, shown in the top right of Fig. 2. In a dipole model, there is one direction of highest event production, opposite from which is the direction of lowest production. Mathematically, this can be expressed as

$$\frac{dN}{d\Omega}(b,\Omega) = \frac{N}{4\pi}(1 + b[\hat{r}(\Omega)\cdot\hat{\Omega}_{\text{dip}}])$$

where Ω represents the two angular parameters, b is the magnitude of the dipole, which is restricted to between 0 and 1, $\hat{r}$ is the unit vector in the direction of the given angles, and $\hat{\Omega}_{\text{dip}}$ is the unit vector in the dipole direction. Notice that the domain restriction on b forces the production rate to remain nonnegative. In practice, this is used as a probability distribution function in the form

$$p_{\text{dip}}(\Omega|\Lambda) = \frac{1}{4\pi}(1 + b\big[\hat{r}(\Omega)\cdot\hat{\Omega}_{\text{dip}}\big]) \tag{4}$$

and fed to ICAROGW. This model takes three parameters: the dipole magnitude b and the right ascension and declination angles defining the dipole direction.

3.2 The Quadrupole Model

Unfortunately, models more complex than the dipole are far more tedious to express algebraically. Because we are representing a probability distribution on the sphere, the most logically basis for these functions is that of the **spherical harmonics**.

The typical spherical harmonics, represented as $Y_l^m(\theta,\phi)$, are commonly found in quantum mechanical fields. The l typically corresponds to total orbital angular momentum and m to z angular orbital momentum. There is contained within the formulation for these functions the requirement that $|m| \leq l$. Because the spherical harmonics form a basis for functions on the sphere, they are ideal for representing arbitrary probability distributions on that same domain.

When $l = 0$, there is only one term, a constant, which for our purposes of probability normalization is scaled to $\frac{1}{4\pi}$. The $l = 1$ terms have already been covered; together with the constant term they represent the dipole model. Therefore, the next step is to approach the $l = 2$ terms, which make up the quadrupole, shown in the bottom left of Fig. 2.

Because probability densities cannot have imaginary parts, it is useful to consider the real spherical harmonics, written as

$$Y_{l,m} = \begin{cases} \sqrt{2}(-1)^m \Im[Y_l^{|m|}] & \text{if } m < 0 \\ Y_l^0 & \text{if } m = 0 \\ \sqrt{2}(-1)^m \Re[Y_l^m] & \text{if } m > 0 \end{cases}$$

These functions behave similarly to the classic spherical harmonics and still represent a basis for functions on the sphere. However, they do not require complex number calculations to handle.

We can now define an eight parameter model for the quadrupole as

$$p_{\text{quad}}(\Omega|\Lambda) = \frac{1}{\sqrt{4\pi}}\left(\sum_{l=0}^{2}\sum_{m=-1}^{l} \mathbf{b_l^m} Y_{l,m}(\Omega)\right) \tag{5}$$

where the $\mathbf{b_l^m}$ are the distribution parameters and $\mathbf{b_0^0} = 1$. Recall that $p(\Omega|\Lambda)$ must remain nonnegative, which creates a complicated space of valid parameter combinations.

Unlike the simple dipole model, there is no immediately obvious metric for the magnitude of the quadrupole, so we must devise some logical expression to occupy that role. The value used was

$$\begin{aligned} b_{\text{ani}}^2 = 3\Big[(b_1^{-1})^2 + (b_1^0)^2 + (b_1^1)^2\Big] \\ + \frac{15}{4}\Big[(b_2^{-2})^2 + (b_2^{-1})^2 + \frac{1}{3}(b_2^0)^2 + (b_2^1)^2 + (b_2^2)^2\Big] \end{aligned} \tag{6}$$

Although this expression appears arbitrary, it has some very useful properties:

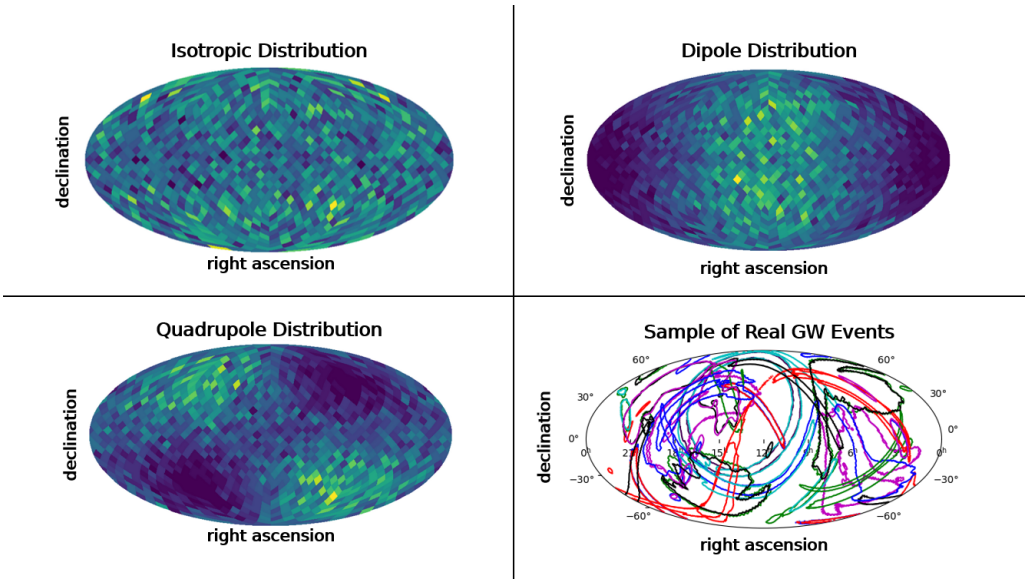

Figure 2: Some skymaps of different distributions, including simulated isotropic (top left), dipole (top right), and quadrupole (bottom left) models. Bright pixels represent regions of high production, dark pixels represent low production. In the bottom right is a sample of some real GW events.

- If all but one of the parameters are set to 0, the value of b_{ani} scales with the magnitude of the remaining parameter and is restricted to fall between 0 and 1.

- If all the b_2^m values are set to 0, b_{ani} reduces down to the expression for the dipole magnitude expressed in spherical harmonics.

It is worth mentioning that the parameters are multiplied by different values due to the manner in which the spherical harmonics are normalized. For our applications, the functions must be normalized as probability distributions, while the harmonics are typically normalized for their inner products.

Priors and Parameter Space

Unlike the dipole model, the quadrupole model does not possess a simple method to integrate a measure of magnitude within the model itself. Because of this, a large proportion of the population parameter combinations can produce negative probability values. While ICAROGW can handle these parameter choices by simply associating with them a probability of zero, it is computationally inefficient.

To examine the efficiency, we produce sets of parameters from uniform priors and test which combinations return negative likelihoods. Ultimately, we find that less than 1% of the prior combinations result in likelihoods above 0. The same procedure was undertaken using Gaussian priors, and a much larger proportion of of the prior draws (approximately 50%) had valid likelihoods. Results are shown in Fig. 3 and in Fig. 4. Notice that once cuts are applied, the remaining distributions are not identical between the Gaussian and uniform priors.

However, when samplers were run using both the uniform and gaussian priors, it was found that the posterior distributions were almost identical as long as the sampler had enough steps to converge, indicating that while the Gaussian priors may be more computationally efficient, they do not produce different posterior results from the uniform priors. In other

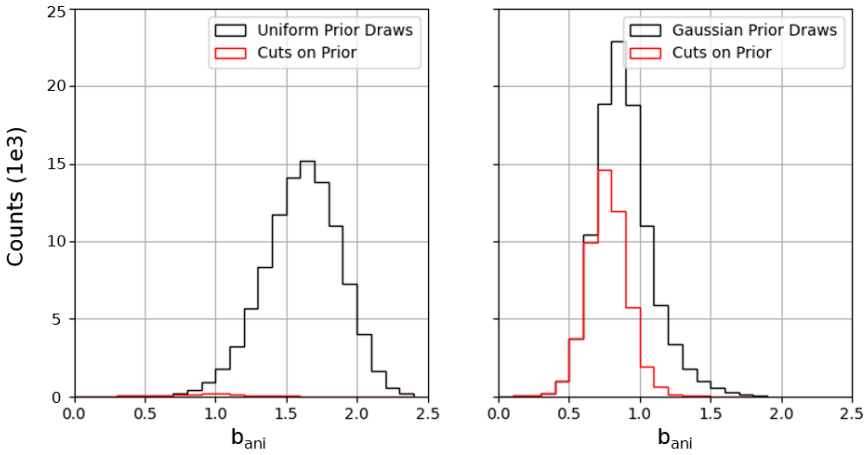

Figure 3: The simulated prior distributions of b_{ani} for both uniform and Gaussian priors with 100,000 prior draws. Each prior draw was analyzed by ICAROGW, and those with negative likelihood were cut. Approximately 1% of the uniform draws remained compared to 50% of the Gaussian draws.

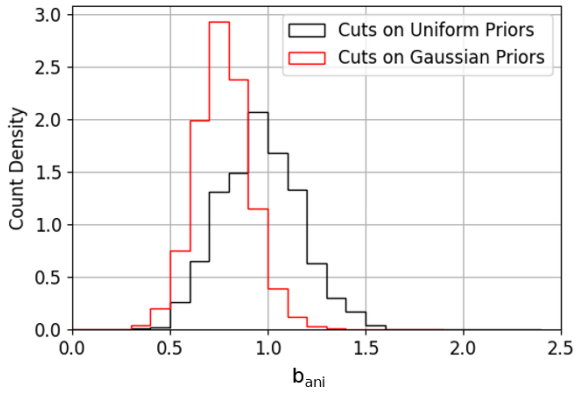

Figure 4: Density plots of the two post-cut prior distributions.

words, the likelihood term dominates, and any constraints we set on the quadrupole coefficients come from the data, not from the priors.

The only major difference observed between the uniform and gaussian priors in regards to real CBC merger events was in the Bayes factors they produced, which differed significantly. This, however, is an indication in the failings of the Bayes factor in this situation, as it is overly reliant on the priors and can therefore be artificially inflated. For this reason, it is not wise to use the Bayes factor in analysis of the quadrupole model. A similar phenomenon was observed in Ref. [4].

4 Bounds On Anisotropies

4.1 Dipole

The first portion of this project was dedicated to examining the relationship between quantity of data and effectiveness of Bayesian inference in placing upper bounds on the dipole magnitude b. By performing the Bayesian analysis using an increasing number of simulated data points N_{sim}, a relationship could be established. By quantifying this relationship, we

can make predictions regarding the future of isotropy testing in gravitational wave astronomy.

Recall that ICAROGW requires three inputs to calculate a likelihood: posterior parameter estimation samples, injections, and a rate model. The process for working with these is outlined in Fig. 5. Both the PE samples and software injections were simulated isotropically, and the dipole rate model was used. Once ICAROGW calculated the likelihood, it was inserted into a 10,000 step MCMC sampler along with a set of priors, generating a posterior distribution for each of the three parameters. The upper bound for the dipole coefficient b was chosen such that 95% of the posterior distribution was below it. Importantly, the b prior was a logarithmic uniform prior between 10^{-6} and 1. The choice of bounds for this prior affects where the 95% point is, so it is important to standardize.

It is worth mentioning that this procedure makes one major simplification of sky position data. Typical CBC merger event data contains a range of possible locations rather than one localized point. The simulations we perform produce GW events with a precise location, meaning the analysis has more information than would normally be provided by the same number of true CBC merger data sets. Regardless, we believe that this distinction will have only a minor impact, and that the procedure outlined here serves as a satisfactory proxy through which to simulate upper bounds.

Plotting $\log b$ on a logarithmic scale for $N_{\rm sim}$ as in Fig. 6 displays a clear negative trend, indicating that the upper bounds that can be set do indeed decrease with additional samples. Fitting linearly, we arrive at a slope of $-0.480(77)$, which agrees nicely with the expected slope of -0.5 according to a proportionality of $\frac{1}{\sqrt{N_{\rm sim}}}$ outside of the logarithm.

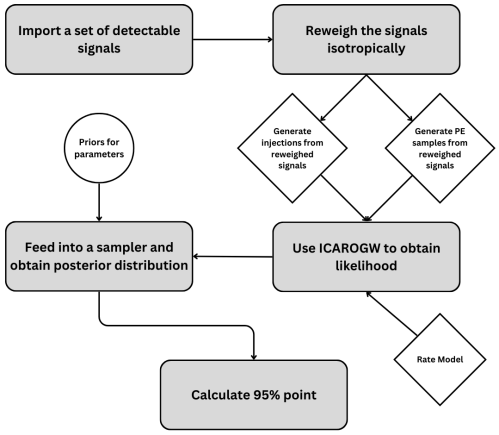

Figure 5: Procedure for calculated upper bounds using simulated data.

Real Data

For this analysis, CBC merger events with false alarm rates (FAR) of $< 1\text{yr}^{-1}$ were used as recommended by Ref. [3] as posterior PE samples. There are 74 such events. Injections were sourced from O3 search sensitivity estimates, released alongside Ref. [3], which inserted fake GW data into real background noise and sent the signals through the usual data analysis pipeline. Using these posteriors and injections we can set an upper bound on b in the same fashion as the simulated data. The resulting upper bound is $\ln b_u = -0.82$, which agrees with the linear expectation of -0.61 for 74 samples we derived from the simulated data to within 1σ of a typical data point.

We can also calculate the Bayes factor between the dipole and isotropic models. This value is $\mathcal{B}^{\rm iso}_{\rm dip} = 1.12$. A value of 1 would indicate that the models are equally supported by the data, while a value above one favors isotropy. Although this result displays a slight preference for isotropy, $\mathcal{B}^{\rm iso}_{\rm dip}$ is so close to 1 that the models are still essentially equal.

At first, this result may be discouraging. However, it is exactly what we expected before any analysis was performed. The current catalog of gravitational wave events is simply not extensive enough to effectively limit anisotropies. As the size of the catalog increases, the same framework discussed in this paper will be able to determine which model better describes the distribution of mergers. Fig. 6 marks the number of merger events we expect to measure through different run periods. In the LVK O4 and O5 periods [1], we expect to detect up to 5,000 GW

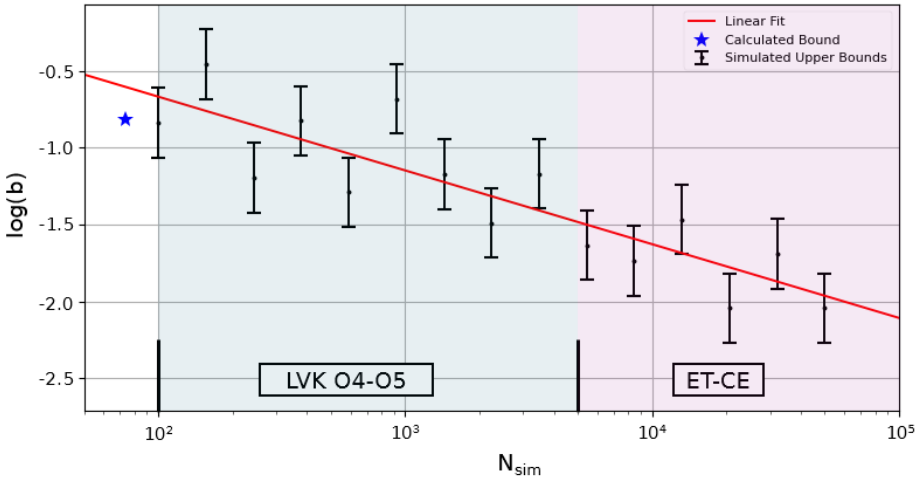

Figure 6: The simulated relationship between number of posterior samples available in Bayesian inference and the 95% upper bound on the value of b. Uncertainty is standard (.23) and is the standard deviation from 10 repeated simulations on a single point. The linear fit is $\log(b) = 0.29 - 0.480 \log(N_{\rm sim})$. The blue star is the upper bound derived from the currently available gravitational wave events.

events. Past that, we expect Cosmic Explorer (CE) and Einstein Telescope (ET) [2] to increase further the size of the merger catalog.

4.2 Quadrupole

As with the dipole model, we simulated the effect that an increasing number of data samples would have on the ability to place upper bounds on the magnitude $b_{\rm ani}$. Unlike the dipole model, however, we could not place a logarithmically uniform prior on the measure of magnitude and calculate the 95% point when it drops off. Because of this, it was necessary to find an alternate metric for the upper bound, defined in Eq. 6.

It was observed that the posterior distribution for $\log b_{\rm ani}$ (produced by using the posterior distributions of the parameters) was largely symmetrical, with a peak in the middle. Additionally, as the number of samples increased, this peak occurred at lower values of $\log b_{\rm ani}$. Note that this peak does not truly indicate a 'correct' value of $\log b_{\rm ani}$, it merely measures the simulated data's ability to overcome the prior distribution for the parameters, which produce a prior for the magnitude that is not logarithmically uniform as the dipole magnitude's was. Therefore, we observe that as the amount of data increases, the distribution's peak approaches 0, as the simulated data follows an isotropic distribution.

It is then this peak that we use as a measure of the upper bound. Because the posteriors of $\log b_{\rm ani}$ are symmetrical, it is convenient to calculate the peak as the mean of the distribution.

Ultimately, we used uniform priors on each of the distribution parameters, and due to computation time concerns, only 1,000 steps of the Markov chain Monte Carlo sampler were performed. Still, a very compelling linear relationship, as seen in Fig. 7, could be established with a slope of $-0.436(23)$, which is within 3σ of the expected -0.5. Additionally, the upper bound derived from the currently available GW data sets agrees with expectation derived from the linear relation to within 1σ, so we may be convinced that this metric and methodology produce adequate results.

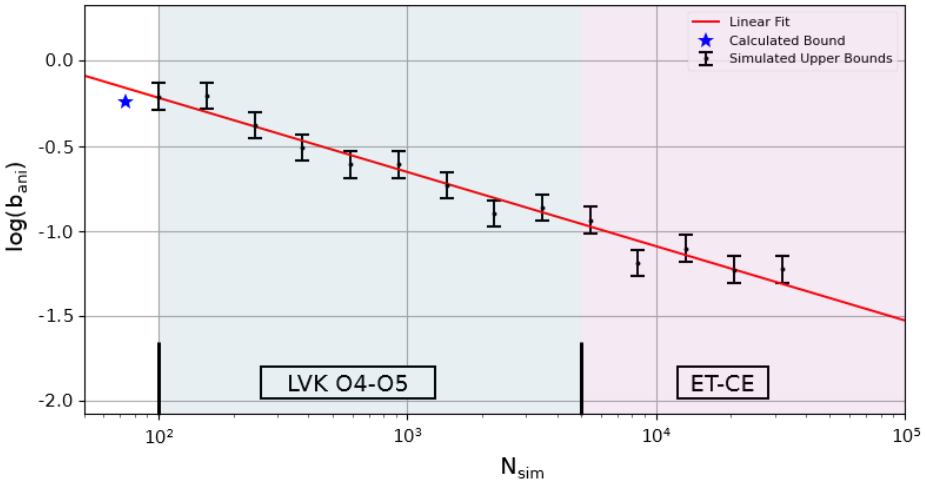

Figure 7: The simulated relationship between number of posterior samples available in Bayesian inference and the upper bound metric for the value of b_{ani}. Uncertainty is standard (.08) and is the standard deviation from 10 repeated simulations on a single point. The linear fit is $\log(b_{\text{ani}}) = 0.654 - 0.436\log(N_{\text{sim}})$. The blue star is the upper bound derived from the currently available gravitational wave events.

5 Discussion

We have leveraged the Bayesian inference package ICAROGW to analyze the statistical behavior of angular distributions of CBC merger events. Using isotropically simulated events, we established a relationship between the number of simulated events and the upper bounds placed on the magnitude metric for both the dipole and quadrupole models. This relationship was found to be $b_{\text{u}} \propto \frac{1}{\sqrt{N_{sim}}}$ for both models, which aligned with expectations.

Additionally, real GW data from the LVK collaborations was used to calculate an upper bound for both magnitude metrics. These values agreed within error with the relationship established using the simulated samples.

The primary achievement of this study has been to develop reliable metrics through which to quantify upper bounds on dipole and quadrupole models as well as processes to calculate these upper bounds. As the volume of CBC merger data increases, the upper bounds that can be set using available data should follow the relationships established in this paper. Significant deviation from these trends could be an indication of anisotropies.

6 Acknowledgements

I would like to thank the organizers of The University of Florida International REU in Gravitational Physics, which is funded by NSF grant 2348913. In particular, I'd like to thank Paul Fulda, Peter Wass, and Sujata Krishna for their work in ensuring my placement and research went smoothly. I'd also like to thank Simone Mastrogiovanni, my supervisor, for his constant support throughout this project. Last, I'd like to extend my gratitude to my fellow summer researchers Brody Nies and Christian Lainez-Artiga for their camaraderie and support throughout the two months of this program.

References

[1] B. P. Abbott et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Reviews in Relativity*, 21:3, Apr. 2018. ADS Bibcode: 2018LRR....21....3A.

[2] M. Branchesi et al. Science with the Einstein Telescope: a comparison of different designs. *Journal of Cosmology and Astroparticle Physics*, 2023:068, July 2023. Publisher: IOP ADS Bibcode: 2023JCAP...07..068B.

[3] T. L. S. Collaboration et al. The population of merging compact binaries inferred using gravitational waves through GWTC-3, Feb. 2022. arXiv:2111.03634 [astro-ph, physics:gr-qc].

[4] R. Essick, W. M. Farr, M. Fishbach, D. E. Holz, and E. Katsavounidis. An Isotropy Measurement with Gravitational Wave Observations. *Physical Review D*, 107(4):043016, Feb. 2023. arXiv:2207.05792 [astro-ph].

[5] S. Mastrogiovanni, C. Bonvin, G. Cusin, and S. Foffa. Detection and estimation of the cosmic dipole with the Einstein Telescope and Cosmic Explorer. *Monthly Notices of the Royal Astronomical Society*, 521(1):984–994, Mar. 2023. arXiv:2209.11658 [astro-ph].

[6] S. Mastrogiovanni, G. Pierra, S. Perriès, D. Laghi, G. C. Santoro, A. Ghosh, R. Gray, C. Karathanasis, and K. Leyde. ICAROGW: A python package for inference of astrophysical population properties of noisy, heterogeneous and incomplete observations, Dec. 2023. arXiv:2305.17973 [astro-ph, physics:gr-qc].

[7] S. Vitale, D. Gerosa, W. M. Farr, and S. R. Taylor. Inferring the properties of a population of compact binaries in presence of selection effects. pages 1–60. 2021. arXiv:2007.05579 [astro-ph, physics:gr-qc].