

(An)isotropy measurement with gravitational wave observationsReed Essick^{*}

*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario N2L 2Y5, Canada;
 Canadian Institute for Theoretical Astrophysics, University of Toronto,
 60 St. George Street, Toronto, Ontario M5S 3H8, Canada;
 Department of Physics, University of Toronto, 60 St. George Street, Toronto, Ontario M5S 1A7, Canada;
 and David A. Dunlap Department of Astronomy, University of Toronto,
 50 St. George Street, Toronto, Ontario M5S 3H4, Canada*

Will M. Farr

*Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York 11794, USA
 and Center for Computational Astrophysics, Flatiron Institute, New York, New York 10010, USA*

Maya Fishbach

*NASA Hubble Fellowship Program Einstein Fellow, Center for Interdisciplinary Exploration and
 Research in Astrophysics (CIERA) and Department of Physics and Astronomy, Northwestern University,
 1800 Sherman Ave, Evanston, Illinois 60201, USA;
 Canadian Institute for Theoretical Astrophysics, University of Toronto,
 60 St. George Street, Toronto, Ontario M5S 3H8, Canada;
 Department of Physics, University of Toronto, 60 St. George Street, Toronto, Ontario M5S 1A7, Canada;
 and David A. Dunlap Department of Astronomy, University of Toronto, 50 St. George Street,
 Toronto, Ontario M5S 3H4, Canada*

Daniel E. Holz

*Department of Physics, Department of Astronomy and Astrophysics, Enrico Fermi Institute,
 and Kavli Institute for Cosmological Physics, The University of Chicago,
 5640 South Ellis Avenue, Chicago, Illinois 60637, USA*

Erik Katsavounidis

*LIGO Laboratory and Kavli Institute for Astrophysics and Space Research,
 Massachusetts Institute of Technology, 185 Albany St. Cambridge, Massachusetts 02139, USA*



(Received 14 July 2022; accepted 8 December 2022; published 14 February 2023)

We constrain the distribution of merging compact binaries across the celestial sphere using the GWTC-3 catalog from the LIGO-Virgo-KAGRA Collaborations' (LVK) third observing run. With 63 confident detections from O3, we constrain the relative variability (standard deviation) of the rate density across the sky to be $\lesssim 16\%$ at 90% confidence assuming the logarithm of the rate density is described by a Gaussian random field with correlation length $\geq 10^\circ$. This tightens to $\lesssim 3.5\%$ when the correlation length is $\geq 20^\circ$. While the new data provides the tightest constraints on anisotropies available to date, we do not find overwhelming evidence in favor of isotropy, either. A simple counting experiment favors an isotropic distribution by a factor of $\mathcal{B}_{\text{ani}}^{\text{iso}} = 3.7$, which is nonetheless an improvement of more than a factor of two compared to analogous analyses based on only the LVK's first and second observing runs.

DOI: [10.1103/PhysRevD.107.043016](https://doi.org/10.1103/PhysRevD.107.043016)**I. INTRODUCTION**

The observation of gravitational waves (GWs) from the coalescence of compact binaries provides a new way to study how these systems form, evolve, and are distributed throughout the Universe (see Ref. [1] and references

therein). In particular, the spatial distribution of GW sources can test the cosmological principle—is the universe statistically homogeneous and isotropic? Deviations from perfect homogeneity have already been proposed as a way to infer cosmological parameters through cross-correlations of clustering within GW and electromagnetic observations (see, e.g., [2,3]). However, these studies assume *a priori* that GW events follow anisotropies

^{*}essick@cita.utoronto.ca

measured from electromagnetic surveys. That is, they do not directly measure anisotropies from the GW data. Our goal in this paper is to constrain anisotropies in the population of merging binaries using only GW data.

Although large deviations from isotropy are not expected, it behooves us nevertheless to check this, similar to the motivation within Ref. [4]. Directly constraining anisotropies with GW catalogs may be of interest in several astrophysical situations. For example, resolving clustering scales from GW data alone may be used to test the assumption that GW sources are always associated with galaxies. Along these lines, GW sources could be used to directly trace clustering scales; see, e.g., Refs. [5,6] for discussion of this in the context of third-generation detectors. Similarly, the identification of individual host galaxies for specific events and/or the statistical association between the full GW catalog and different types of galaxies may suggest, perhaps through the mass-dependent galaxy-clustering scale, which types of galaxies most often host compact binary coalescences [7,8]. This could be combined with knowledge of the star formation history to in turn constrain the delay time distribution between binary formation and coalescence [9]. See, e.g., Ref. [10] for a similar application to short gamma-ray bursts (GRBs).

It is also worth remembering that the LIGO-Virgo-KAGRA (LVK) Collaborations [11,12] search for unmodeled “burst” events in addition to compact binaries [13]. Given that the source of such events will not be known *a priori*, their spatial distribution will likely provide crucial clues as to their origins. Indeed, determining whether burst events correlate with local structure will inform the distance to the sources and therefore their intrinsic energy scales, analogous to GRBs [14] and other high-energy astrophysical phenomena.

The detection of anisotropies within the distribution of merging binaries could be the signature of more exotic physics, such as wormholes that may effectively tunnel to larger volumes and therefore a higher number of merging binaries [15] or lensed events, which appear as repeated signals from the same part of the sky [16,17]. In particular, strong lensing may distort the shape of the waveform, particularly the relative phasing between different harmonics [18]. These effects may be difficult to distinguish from more general alternative theories of gravity [19], and the identification of anisotropies may be a cleaner signature of lensing than the waveform’s phasing alone. Indeed, many searches for lensed events begin with overlaps on the sky.

Several authors have already studied the distribution of merging binaries with the LIGO-Virgo Collaborations’ [11,12] first catalog of 11 detections (GWTC-1 [20]). Specifically, Ref. [21] modeled anisotropies with 12 pixels of equal area and a set of Euler angles that rotated the pixelization across the sky. Using an approximation of the catalog’s sensitivity that assumed constant and equal power spectral densities for both LIGO detectors throughout the run, neglecting the presence of Virgo, but accounting for the

diurnal cycle and correlations in when the LIGO interferometers recorded science-quality data [22], they found weak evidence in favor of isotropy. Similarly, Ref. [23] used the same 11 events but a different estimate of survey sensitivity to constrain anisotropies with a model constructed from a low-order spherical-harmonic expansion. They considered several models with different numbers of harmonics up to $l_{\max} = 5$, finding equivalently weak evidence in favor of isotropy regardless of $l_{\max}$. Finally, Ref. [24] attempted to measure the two-point correlation function of GW events with a spherical-harmonic decomposition of the sum of individual event localizations while assuming the sensitivity of the detector network was uniform over the entire sky. They also found no evidence for an excess of correlation at any angular scale.

Maps of upper limits on anisotropies in the stochastic GW background are routinely produced under various assumptions in either the pixel or spherical-harmonic domains. Although no statistically significant detection has been made, these analyses typically make assumptions about the power spectrum of the stochastic GW background and produce maximum likelihood estimates of the angular distribution of the intensity. See Ref. [25] for a review. While there has been no unambiguous detection of the stochastic GW background to date, let alone the detection of anisotropies, there may still be information about the distribution of merging binaries at high redshift encoded in the nondetection (see, e.g., Ref. [26]).

Additionally, anisotropies are of general interest in other high-energy astrophysical phenomena. Analyses of GRBs show that they are consistent with isotropic distributions, regardless of how the catalog is subdivided [14], the distribution of fast-radio bursts (FRBs) is an active area of research [27], and multiple groups have claimed detections of anisotropies in the arrival directions of cosmic rays [28–31].

Therefore, it is of general interest to develop methods to constrain the rate of mergers as a function of their position on the celestial sphere. We use hierarchical Bayesian inference to construct posterior processes for the distribution of merging compact binaries over the sky using 63 confidently detected binaries, including binary black hole (BBH), neutron star-black hole (NSBH), and binary neutron star (BNS) sources, from the LVK’s third observing run (O3, 1 April 2019–27 March 2022 [32–34]). In addition to the nearly six-fold increase in sample size from GWTC-1, our analysis benefits from estimates of survey sensitivity derived from simulated signals injected into real detector noise and processed directly with the searches used to construct the catalog [35]. These injections implicitly account for variability in each detector’s sensitivity and correlations between the times when detectors record data.¹

¹Appendix B shows that the O3 catalog’s sensitivity is nearly uniform over the entire sky, although measurable deviations exist.

This improves upon previous estimates of survey sensitivity, which depended on approximations with poorly quantified systematic uncertainties [21,23]. We also self-consistently incorporate realistic models of the masses, spins, and redshift distributions of merging binaries derived from GW observation [1,36].

We find mild evidence in favor of isotropy. This agrees with Refs. [21,23], but we place tighter constraints on anisotropies because of the larger sample size now available. In fact, we find Bayes factors in favor of isotropy ($\mathcal{B}_{\text{ani}}^{\text{iso}}$) similar to Refs. [21,23] when we use only events from GWTC-1, and these increase by a factor of two when we only use the 63 events from O3. Although there are a few persistent “hot pixels” from O3 in all our models on average, we cannot confidently bound the rate density in these directions to be inconsistent with isotropy. Indeed, we bound the relative variability (standard deviation) in the rate density to $\lesssim 16\%$ of the isotropic rate at 90% credibility if the correlation length scale in the rate density is $\geq 10^\circ$, and this is improved to $\lesssim 3.5\%$ if the length scale is $\geq 20^\circ$.

The rest of this paper is structured as follows. In Sec. II, we perform a simple counting experiment by dividing the sky into hemispheres, showing that most of the information about (an)isotropy comes from the best-localized events (less than half our catalog). Section III presents additional models of varying complexity, including pixelized representations like Ref. [21] (Sec. III B 1) and representations based on low-order spherical-harmonic expansions like Ref. [23] (Sec. III B 2), culminating in a nonparametric description of the rate density as a Gaussian random field (Sec. III C). We discuss implications of current constraints and conclude in Sec. IV.

II. COUNTING EXPERIMENTS

We begin with a simple counting experiment: divide the sky into two hemispheres and “count” the number of events that fall within each.² In the context of GW catalogs, this simple model is useful because of the symmetry inherent in the sensitivity for current interferometers. Each interferometer’s sensitivity has even parity when reflected across the plane defined by its arms. This means that the sensitivity to each hemisphere will be equal regardless of how many interferometers participate in the survey and exactly where hemispheres are drawn as long as they divide the sky in half equally.

To wit, we construct a model that divides the sky in half, assigning a different rate density to each hemisphere; the expected fraction of events coming from the “northern” hemisphere is f , and the corresponding fraction from the “southern” hemisphere is $1 - f$. We also consider all possible hemispheres by sampling over Euler angles that rotate the simple north-south hemisphere model into




FIG. 1. Marginal prior (blue) and posterior (orange) for the mixing fraction in the rotated-hemisphere model (Table II). The data favor isotropy (gray, $f = 0.5$) by a factor of $\mathcal{B}_{\text{ani}}^{\text{iso}} = 3.7$ and rule out anisotropies bigger than 3:1.

arbitrarily oriented hemispheres. This rotated model is similar to the approach in Ref. [21]. Inference proceeds by effectively counting the number of events consistent with each hemisphere and inferring the expected fraction of events and the hemispheres’ orientation most consistent with the observations. Uncertainty in the fraction of events in a particular hemisphere, then, roughly corresponds to counting uncertainty from a binomial distribution. We list the prior ranges for the rotated-hemisphere model and compare it to others in Table II.

Much of the information about isotropy comes from the best-localized events, and we find that the data prefer equal fractions of events from each hemisphere ($f = 0.5$) by a factor of $\mathcal{B}_{\text{ani}}^{\text{iso}} = 3.7$, assuming uniform priors for f and the Euler angles and calculating $\mathcal{B}_{\text{ani}}^{\text{iso}}$ via the Savage-Dickey density ratio [37] (Fig. 1). The model also finds no preference for specific rotations, which is expected if $f \sim 0.5$. Furthermore, the number of events in one hemisphere is binomially distributed, and the uncertainty in the fraction of events will be $\sigma_f^2 = f(1 - f)/N$ with N events. With our selection of 63 events and assuming isotropy, we expect $\sigma_f = 6.3\%$. However, this is significantly smaller than the actual standard deviation observed in Fig. 1, which corresponds to only 23.5 effective events ($\sigma_f = 10.3\%$). While this could be due in part to the trials factor associated with sampling over possible rotations, it is likely because many of the events in our catalog have uninformative broad localization uncertainties. Indeed, if we only use the 25 best-localized events from our catalog,³ we find $\sigma_f = 12.6\%$ and $\mathcal{B}_{\text{ani}}^{\text{iso}} = 2.7$, only slightly less constraining than the uncertainty obtained with the full catalog. Similarly, if we only include the ten best-localized events, we obtain $\sigma_f = 17.5\%$ and $\mathcal{B}_{\text{ani}}^{\text{iso}} = 1.9$, only slightly worse than expected from the binomial distribution ($\sigma_f = 15.8\%$ with ten events).

²Because GW events often have very broad localizations, we always employ hierarchical Bayesian inference to account for measurement uncertainty. See Sec. III A for more details.

³In general, selecting events in this way may significantly complicate our estimate of the catalog’s sensitivity. However, our Rotated Hemisphere model is immune to such considerations because of the symmetry of the interferometer antenna patterns.

This should be contrasted with the constraints obtained using only GWTC-1: $\mathcal{B}_{\text{ani}}^{\text{iso}} = 1.3$ and $\sigma_f = 18.9\%$.⁴ Reference [21] found $\mathcal{B}_{\text{ani}}^{\text{iso}} = 1.3$ and Ref. [23] quote $\mathcal{B}_{\text{ani}}^{\text{iso}} \sim 1.1\text{--}1.6$ depending on how many spherical harmonics they include. We see, then, that our larger sample size provides the tightest constraints to date.

Note that one could combine all events from O1, O2, and O3 in order to obtain an even tighter constraint on anisotropy under the rotated-hemisphere model without accurate estimates of search sensitivity because of the symmetry in the model. We avoid this because of the expectation that systematic uncertainty from the relatively strong assumptions behind the shape of anisotropies allowed by this model will be more important than any improved statistical uncertainty, and therefore focus on more flexible models of isotropy, for which accurate estimates of search sensitivity are more important, in what follows.

III. HIERARCHICAL MODELS

We now consider several additional representations of the distribution of merging binaries and construct maps of the merger rate across the sky with each. Section III A briefly reviews hierarchical Bayesian inference before Secs. III B, III C, and III D provide our results. Comparing different modeling choices allows us to examine, to some extent, which features are constrained by the data and which are dominated by our modeling choices.

In what follows, we assume fixed distributions for the source-frame component masses, redshift, spins, inclination, orientation, and arrival time of GWs from binary systems. These are described in Table I. In order to focus on isotropy, we only infer the parameters of the distribution over right ascension (α) and declination (δ). While we do not expect the assumption of, e.g., fixed masses, redshift, and spin distributions to affect our conclusions about (an)isotropy, it would be worthwhile to check this. However, we leave studies of possible correlations between the direction to the source and other source properties to future work.

We only consider events from O3 (GWTC-2 [32], GWTC-2.1 [33], and GWTC-3 [34]), as the publicly available set of simulated signals processed with real searches used to estimate the catalog’s sensitivity only covers O3 [35]. However, we consider *all* events from O3, including BNS, NSBH, and BBH systems. We approximate the catalog selection by requiring the false alarm rate (FAR) from at least one pipeline within GWTC-3 to be $\leq 1/\text{year}$.⁵

⁴Although our sensitivity estimates only cover O3, we can analyze GWTC-1 without accounting for selection effects because of the symmetry in this model.

⁵We include all searches present in GWTC-3: both modeled (GstLAL, MBTA, PyCBC broad, and PyCBC BBH) and unmodeled (cWB) searches. See Ref. [34] for more details about individual searches.

With this selection threshold, we retain 63 events from O3. See Appendix B for more details.

A. Formalism

We employ hierarchical Bayesian inference to infer the rate density of merging compact binaries,

$$\frac{dN}{d\theta} = \mathcal{R}p(\theta|\Lambda), \quad (1)$$

where each event is described by parameters θ (masses, redshift, right ascension, declination, etc.), the population distribution $p(\theta|\Lambda)$ is described by some set of parameters Λ (minimum and maximum masses, anisotropy parameters, etc.), and $\mathcal{R}$ acts as an overall normalization constant.

Specifically, we sample from the rate-marginalized inhomogeneous Poisson likelihood for the observed data $\{D_i\}$ from N events

$$p(\{D_i\}|\Lambda) = \prod_i^N \frac{\int d\theta p(D_i|\theta)p(\theta|\Lambda)}{\int d\theta P(\text{det}|\theta)p(\theta|\Lambda)} \quad (2)$$

with a corresponding prior for Λ . Here, $P(\text{det}|\theta)$ is the (time-averaged) probability of detecting a signal with parameters θ . Equation (2) implicitly assumes $p(\mathcal{R}) \sim 1/\mathcal{R}$ within the marginalization over $\mathcal{R}$. We estimate the numerators in Eq. (2) via Monte Carlo importance sampling of single-event posterior samples for each event, and the denominator with a set of detected simulated signals (Appendix B). See, e.g., Refs. [39–42] and references therein for more details.

B. Cartography

Within the hierarchical framework, we consider several different representations of the distributions over the sky. Broadly, these can be classified as those based on pixelizations (like the rotated-hemisphere model in Sec. II) and those based on spherical-harmonic decompositions. Table II summarizes our models, their parameters, and the priors chosen for those parameters. While there is no fundamental difference between the two approaches, each introduces different priors on the types of variation over the sky. Nonetheless, as we will see, we obtain comparable results regardless of the precise model choices.

1. Pixelized representations

The simple rotated-hemisphere model (Sec. II) found weak evidence for isotropy, but this could be due to the assumptions implicit in the model. We now focus on pixelization schemes that allow for more complex anisotropies. Figure 2 shows Mollweide projections of $p(\Omega)$ derived from different modeling assumptions. Specifically, we employ the Healpix pixelization scheme [43] and consider models with 12, 48, and 192 pixels, respectively.

TABLE I. Fixed population models for the source-frame primary mass (m_1), secondary mass ($m_2 \leq m_1$), Cartesian-spin vectors for each component ($\vec{s}_1, \vec{s}_2$), and redshift (z). We employ the *maximum a posteriori* values for the Broken Power-Law + dip model from Ref. [36] as well as a flat Λ CDM cosmology with $H_0 = 67.32$ km/s/Mpc, $\Omega_M = 0.3158$, and $\Omega_\Lambda = 1 - \Omega_M$ (first column of Table 1 in Ref. [38]). While there is evidence that the (BBH) population evolves with redshift and the spins are not isotropically distributed, these effects are not expected to strongly influence our inference for the right ascension and declination. We also assume events' orbital inclinations are isotropically distributed, events' phases at coalescence and polarization angles are uniformly distributed throughout their physical ranges, and that events' arrival times are uniformly distributed throughout the duration of the experiment.

Variates	Name/Description	Functional form	Visualization
m_1, m_2	Broken power-law + dip (BPL + Dip)	$p(m_1, m_2) \propto p(m_1)p(m_2) \times p_{\text{pair}}(m_2, m_2/m_1)$ see Refs. [1,36]	
z	Uniform in comoving volume and source-frame time	$p(z) \propto (dV_c/dz)/(1+z)$	
$\vec{s}_1, \vec{s}_2$	Uniform in magnitude and isotropic orientation	$p(\vec{s}_i) = 1/4\pi \vec{s}_i ^2$	

In each model, the rate density in each pixel is independently, exponentially distributed *a priori*. The exponential distribution is convenient because it only has support for non-negative rate densities. Furthermore, the independent priors for each pixel give models with more pixels more freedom. That is, as the number of pixels increases, the prior not only allows for, but actually prefers increased variation across the sky.






We additionally consider a radically different pixelization scheme; the 88 constellations⁶ defined by the

⁶Although there are only 88 constellations, we fit the rate in 89 pixels, dividing Serpens (Ser) into its two disjoint regions.

International Astronomical Union (IAU) [44]. Although we expect the IAU constellations to be completely unrelated to GW events, which come from much greater distances than the stars that make up the constellations, they provide a convenient and memorable alternative pixelization. As with the Healpix models, the rate in each constellation is independently exponentially distributed *a priori*. Because we parametrize the model in terms of the rate (count per steradian) within each constellation, this implies that the expected number of events from a constellation scales with the constellation's area.

Unlike our rotated-hemisphere model and Ref. [21], we do not consider rotations of these pixelizations. Reference [21] only used 12 pixels and introduced three

TABLE II. Population models for the distribution over right ascension (α) and declination (δ). See text for more detailed definitions of each model's parameters. We denote the uniform distribution between X and Y as $U(X, Y)$, the exponential distribution with scale parameter Z as $\text{Exp}[Z]$ ($p(x) = Z^{-1}e^{-x/Z}$), and the multivariate normal distribution with mean vector μ and covariance matrix Ξ as $\mathcal{N}(\mu, \Xi)$. Where relevant, we denote the area of pixel i with A_i .

Variates	Name	Parameters	Functional form	Example
$\Omega \equiv \alpha, \delta$	Rotated Hemisphere (RH)	$f \sim U(0, 1)$ $\phi \sim U(0, 2\pi)$ $\theta \sim U(0, 2\pi)$ $\psi = 0$	Rotate by Euler angles (ϕ, θ, ψ) . $\alpha, \delta \rightarrow \tilde{\alpha}, \tilde{\delta}$ In the rotated frame $p = (f\Theta(\tilde{\delta} > 0) + (1-f)\Theta(\tilde{\delta} < 0))/2\pi$	
	Simple Dipole (SD)	$ \vec{b} \sim U(0, 1)$ $\arctan(b_y/b_x) = \phi \sim U(0, 2\pi)$ $b_z/ b = \cos \theta \sim U(-1, +1)$	$p = (1 + \vec{b} \cdot \hat{\Omega})/4\pi$ $ \vec{b} \leq 1$	
	Healpix pixelization (HP: $N_{\text{pix}} = 12, 48, 192$)			
	IAU constellations (IAU: $N_{\text{pix}} = 89$)	$f_i \sim \text{Exp}(A_i^{-1})$ $\forall i \in [1, \dots, N_{\text{pix}}]$	$p = \sum_i^{N_{\text{pix}}} f_i A_i^{-1} \Theta((\alpha, \delta) \in A_i)$ $\sum_i^{N_{\text{pix}}} f_i = 1$	
	Gaussian random field (GRF)	$\log(f_i/A_i) \sim \mathcal{N}(0, \Xi_{ij})$ $\Xi_{ij} = \sigma_{\text{wn}}^2 \delta_{ij} + \sigma^2 e^{-(\Delta\theta_{ij})^2/\theta^2}$ $\vartheta \sim U(\vartheta_{\text{min}}, \pi/3)$ $\sigma \sim U(0, 3)$ $\sigma_{\text{wn}} = \sigma/10$		
	Exponentiated spherical harmonics (ESH: $l_{\text{max}} = 1, 2, 3, 4$)	$\mathbb{R}\{b_{l \leq l_{\text{max}}}^{m=0}\} \sim U(-10, +10)$ $\mathbb{I}\{b_{l \leq l_{\text{max}}}^{m=0}\} = 0$ $\mathbb{R}\{b_{l \leq l_{\text{max}}}^{m>0}\} \sim U(-10, +10)$ $\mathbb{I}\{b_{l \leq l_{\text{max}}}^{m>0}\} \sim U(-10, +10)$	$p \propto \exp\left(\sum_{lm} b_l^m Y_l^m(\alpha, \delta)\right)$ $b_l^{-m} = (b_l^{+m})^*$	

Euler angles to attempt to control for model systematics associated with the placement of the 12 pixels. We instead use models with more pixels and different methods of partitioning the sky to test for model systematics.

Figure 2 shows one-dimensional summary statistics defined for each direction on the sky. We show the average *a posteriori* rate for each direction (Ω)

$$\langle p(\Omega) \rangle = \int d\Lambda p(\Lambda | \{D_i\}) p(\Omega | \Lambda), \quad (3)$$

where $p(\Lambda | \{D_i\})$ is the hyperposterior distribution inferred via Eq. (2). This average is normalized by the equivalent rate for an isotropic distribution: $p(\Omega | \text{iso}) = 1/4\pi$. We also show the difference $\langle p(\Omega) \rangle - p(\Omega | \text{iso})$ divided by the standard deviation of the one-dimensional marginal posterior distribution: $\sigma_{p(\Omega)}$.

We find similar features with all models. Although there are “hot pixels” throughout the sky for each, on average (left column of Fig. 2) there is a consistently hot pixel near $(\alpha, \delta) = (-45^\circ, +10^\circ)$. This hot spot lies within Equuleus

(Equ: the “little horse”), which has the second smallest area of any of the IAU constellations and is associated with a handful of relatively well localized events (see Appendix B). Even though the horse is little, at face value it may play a big role in GW anisotropy measurements. However, while hot pixels can at times correspond to rates several times larger than the rate for an isotropic distribution, there is still significant uncertainty in the posterior. In fact, the expected value of the rate in any pixel is always less than ± 1.5 standard deviations away from isotropy *a posteriori* (right column of Fig. 2).

We note that the size of the deviations from isotropy are less significant within models with more pixels. This is because there are fewer expected events per pixel and therefore greater relative uncertainty in each pixel's rate density. Indeed, while there are always some pixels for which the rate is not confidently bounded away from zero, there are (many) more poorly constrained pixels for models with larger N_{pix} . See Appendix A for more discussion.

While comparisons based on only one-dimensional marginal posteriors do not actually represent a full test of isotropy (the rate density must be consistent with




FIG. 2. Mollweide projections of the posterior for the rate density across the sky with pixelized representations. (top to bottom) Healpix pixelizations with 12, 48, and 192 pixels as well as a pixelization based on the 88 IAU Constellations (89 pixels; Serpens is divided into two disjoint regions). (left) The average rate density *a posteriori* scaled by the rate for an isotropic distribution. (right) A measure of statistical significance; the difference between the average rate density and the rate for an isotropic distribution scaled by the standard deviation of the rate density in each pixel *a posteriori*.

isotropy in all pixels simultaneously, not just separately for each pixel), this is nevertheless suggestive. We also eschew $\mathcal{B}_{\text{ani}}^{\text{iso}}$ for these models due to the possibly strong dependence on our prior choices (see discussions in, e.g., Refs. [45,46]). We quantify constraints on anisotropies in more detail in Sec. III D.

2. Spherical-harmonic representations

We now turn our attention to representations of $p(\Omega)$ based on spherical-harmonic decompositions. There are many ways to construct a representation of a positive semidefinite function defined on $\mathbb{S}_2$ in terms of spherical harmonics. We explored several models of the form

$$p(\Omega) = F\left(\sum_{l=0}^{l_{\max}} \sum_{m=-l}^{m=+l} b_l^m Y_l^m(\Omega)\right) \quad (4)$$

with the additional constraint $b_l^{-m} = (b_l^m)^*$ to ensure the sum is real, where $(\cdot)^*$ denotes complex conjugation.

To begin, we consider a Simple Dipole model described by a single vector $\vec{b}$ so that

$$p(\Omega) = \frac{1}{4\pi} (1 + \vec{b} \cdot \hat{\Omega}), \quad (5)$$

with $|\vec{b}| \leq 1$. This model is similar to the rotated-hemisphere model from Sec. II, but avoids sharp features in the rate density. It corresponds to $F(x) = x$ and $l_{\max} = 1$ in Eq. (4). With a uniform prior over $|\vec{b}|$ and isotropic priors on its orientation, we constrain $|\vec{b}| \leq 0.5$ at 90% credibility, in agreement with the rotated-hemisphere model. We also find $\mathcal{B}_{\text{ani}}^{\text{iso}} = 2.5$, slightly smaller than the rotated-hemisphere model. This is because the Simple Dipole lacks sharp features in $p(\Omega)$ and therefore larger anisotropies are harder to constrain.

We now additionally consider $l_{\max} \geq 1$. Although we find consistent results with several choices of $F(x)$,⁷ we focus on $F(x) = e^x$, which we refer to as the Exponentiated Spherical-Harmonic (ESH) model. That is, we model the logarithm of the probability density with a spherical-harmonic decomposition. This preserves the parity of all Y_l^m , thereby removing many of the degeneracies introduced by other choices and simplifying the interpretation of posterior constraints for b_l^m .

We consider independent, uniform priors for the real and imaginary parts of b_l^m (subject to the reality constraint) up to several maximum harmonic numbers ($l_{\max}$). Just as larger N_{pix} allow for more model freedom, larger $l_{\max}$ allow the spherical-harmonic decomposition to represent more complex distributions over the sky. Figure 3 shows maps constructed with this Exponentiated Spherical-Harmonic model for $l_{\max} = 1, 2$, and 3.

As a rule of thumb, constraints on low- l coefficients weaken as $l_{\max}$ increases. However, we consistently find that the $l = 1$ (dipole) coefficients are constrained to be rather small, consistent with the rotated-hemisphere and Simple Dipole models. Constraints on higher harmonics are weaker, but they are also all constrained to be relatively

small. Figure 4 shows the prior and posterior for individual b_l^m when $l_{\max} = 2$. We again eschew $\mathcal{B}_{\text{ani}}^{\text{iso}}$ for this model because of ambiguity in the choices for the prior bounds on the $\{b_l^m\}$. Indeed, because the posterior is consistent with isotropy, we can make $\mathcal{B}_{\text{ani}}^{\text{iso}}$ as large as we like by simply increasing the extent of the prior.

When $l_{\max} \geq 2$, we begin to see structure appear across the sky on average *a posteriori* (Fig. 3). This resembles the structure observed with pixelized representation, and, like the pixelized representations, there are large fluctuations in the posterior that render the difference between the marginal means and isotropy statistically insignificant. Another way to view this is to examine the power in each harmonic. Figure 4 shows these distributions as well. Indeed, the power allowed in each harmonic *a posteriori* is larger for higher harmonics, but it is always much smaller than the prior, showing that the data favor isotropy.

C. Gaussian random fields

We complete our survey of the impact of model choices by modeling the (logarithm of the) rate density as a Gaussian random field (GRF), also known as a Gaussian process [47].

Specifically, we assume a Healpix pixelization scheme with many pixels but, importantly, do not assign independent priors to each pixel. Instead, we assume the rate density in each pixel is correlated with neighboring pixels according to a covariance kernel

$$\begin{aligned} \text{Cov}[\ln(p(\Omega_i)), \ln(p(\Omega_j))] &\equiv \Xi_{ij} \\ &= \sigma_{\text{wn}}^2 \delta_{ij} + \sigma^2 \exp\left(-\frac{(\Delta\theta_{ij})^2}{\vartheta^2}\right), \end{aligned} \quad (6)$$

composed of a white-noise component (uncorrelated variance within each pixel scaled by σ_{wn}^2) and a squared exponential component (described by a marginal variance σ^2 and correlation length scale ϑ) that correlates neighboring pixels based on the angular separation between their centers ($\Delta\theta_{ij}$). We fix $\sigma_{\text{wn}} = \sigma/10$, as we wish pixels to be significantly correlated and only include the white-noise variance for numerical stability. While this choice was made primarily to guarantee numerical stability within Cholesky decompositions of (at times) ill-conditioned covariance matrices with large ϑ , it also introduces a natural resolution scale at which σ_{wn} from many pixels tends to dominate the variance in $p(\Omega)$ over the sky. For large ϑ (strong squared-exponential correlations), we expect σ_{wn} to contribute a significant fraction of the overall variability when $N_{\text{pix}} \gtrsim 100$ if $\sigma_{\text{wn}} = \sigma/10$. However, for $\vartheta \sim 60^\circ$, this increases to $N_{\text{pix}} \sim 1100$, and for $\vartheta \sim 10^\circ$ it increases to $N_{\text{pix}} \sim 41,000$. We therefore expect our results to not depend strongly on the choice $\sigma_{\text{wn}} = \sigma/10$ given the range of ϑ included and the number of pixels used.

⁷Reference [23] chose $F(x) = x^2$, and we also explored $F(x) = |x|$. However, both of these approaches complicate the interpretation of the model as they introduce strong degeneracies. That is, multiple distinct sets of b_l^m can produce similar $p(\Omega)$. For example, a distribution with only $b_1^0 \neq 0$ produces similar $p(\Omega)$ to a distribution with only $b_2^0 \neq 0$. It is this mixing between different l can be difficult to interpret. These degeneracies render the posterior for $\{b_l^m\}$ multimodal, which in part motivated Ref. [23]'s choice to limit the magnitude of b_l^m to small values *a priori*.




FIG. 3. Mollweide projections of the posterior rate density across the sky when it is represented by a spherical-harmonic decomposition. Analogous to Fig. 2. (top to bottom) Spherical harmonics are included up to $l_{\max} = 1, 2$, and 3. As more harmonics are included, we observe structure across the sky similar to what is found with the pixelized representations. However, as in Fig. 2, the fluctuations in the posterior are larger than the difference between the mean *a posteriori* and an isotropic distribution.

We confirmed this by also investigating $\sigma_{\text{wn}} = \sigma/100$ and $\sigma_{\text{wn}} = \sigma/1000$, finding consistent results.

Just as our GRF model is related to our pixelized models with a different prior, it can also be expressed in terms of a spherical-harmonic representation. Specifically, the b_l^m are independently, normally distributed within a GRF prior, and their individual variances depend on the form of the covariance kernel (see, e.g., Refs. [48,49]). The GRF prior controls how the prior uncertainty in b_l^m decreases as l increases; the contribution of high- l modes are limited and the resulting rate density is smooth. Similarly, the same prior controls how quickly the rate density is allowed to vary from pixel to pixel.

The key advantages of the GRF approach are that it is straightforward to learn the correlation parameters at the same time we fit the data and that it does not depend strongly on how many pixels or harmonics are included. That is, we need not make strong (and poorly understood)

prior choices about how many pixels or b_l^m to include. The data itself will determine which correlations are preferred. Figure 5 shows the resulting posteriors for the GRF parameters. The data prefer small σ when $\vartheta \gtrsim 15^\circ$, and are consistent with the isotropic limit ($\sigma \rightarrow 0$) for all ϑ .

The data do not strongly constrain the correlation length, although the constraints on σ are more stringent for longer ϑ . In other words, if neighboring pixels are significantly correlated, then the data are less consistent with large fluctuations in the rate density across the sky. This is similar to the fact that we are able to more tightly constrain the low- l coefficients in the spherical-harmonic model compared to high- l coefficients.

Finally, Fig. 6 shows Mollweide projections analogous to Figs. 2 and 3 when we impose several lower limits on the correlation length ($\vartheta \geq \vartheta_{\min}$). The key differences between Figs. 6 and 2 are that the most extreme excursions of the posterior's mean are smaller for the GRF models due to




FIG. 4. Prior (blue) and posterior (orange) for the spherical-harmonic coefficients (b_l^m) with $l_{\max} = 2$. Other $l_{\max}$ produce similar behavior. (lower left) Distributions over b_l^m . Although the marginal posteriors for some $b_{l>0}^m$ peak at nonzero values, they are all consistent with isotropy. Our priors are, perhaps, unrealistically broad, but were intentionally chosen to be much broader than the posterior. (upper right) Distributions over the power in each angular harmonic. Even though isotropy is strongly disfavored *a priori*, the data constrain the power in higher harmonics to be small compared to the prior.

the correlations between neighboring pixels from the prior. We also note that, correspondingly, the fluctuations that do occur in the GRF model appear even less significant.

Nevertheless, we see features reminiscent of individual events within the posterior process's mean when $\vartheta_{\min}$ is small. This is not unexpected, as the posterior mean is related to the sum of individual events' posteriors when the anisotropies are small. Appendix A describes this in more detail.

D. Quantifying constraints on anisotropy

As we have discussed, it can be difficult to interpret Bayes factors within our models because of ambiguity in the choices of prior ranges. For example, $\mathcal{B}_{\text{ani}}^{\text{iso}}$ can be made as large or as small as one would like within the Exponentiated Spherical-Harmonic model by changing the prior ranges allowed for each b_l^m . Therefore, we propose a more direct measure of the extent of anisotropies: the variance of the rate density across the sky,




FIG. 5. Posterior probability for GRF parameters (σ : marginal uncertainty; ϑ : correlation length) assuming independent, uniform priors for each (see Table II) when we use Healpix decompositions with (light blue) 192, (dark blue) 768, and (green) 3072 pixels. The distributions do not depend strongly on the number of pixels used.

$$\sigma_{p(\Omega)}^2 \equiv \frac{1}{4\pi} \int d\Omega \left(p(\Omega) - \frac{1}{4\pi} \right)^2. \quad (7)$$

This is closely related to the GRF model's σ parameter. Isotropy corresponds to the limit $\sigma_{p(\Omega)} = 0$.

In both the Healpix and Exponentiated Spherical-Harmonic models, we find that the posterior supports larger amounts of variability as we increase N_{pix} or l_{max} . That is, the data does not strongly constrain rapid oscillations within the distribution over the sky, and the variability in the inferred distribution is dominated by the prior induced over these high- l modes. In particular, both the Healpix and exponentiated spherical harmonic *a priori* have vanishingly small support for small variance over the sky. This carries over to the posterior, and any upper limit on the variance will strongly depend on the prior.

The GRF model, on the other hand, naturally avoids this issue by simultaneously sampling over both the correlation parameters and the distribution over the sky. Because the GRF model contains support for all correlation lengths (ϑ), as opposed to a fixed choice of N_{pix} or l_{max} , we do not observe vanishing support for small variances. Indeed, we obtain a consistent upper limit within the GRF model when $\vartheta_{\text{min}} \gtrsim 10^\circ$ regardless of the number of pixels used. The distribution of merging binaries produces $\sigma_{p(\Omega)} \lesssim 16\%$ of the isotropic rate at 90% credibility when $\vartheta \geq 10^\circ$. That is, the rate density fluctuates by $\lesssim 16\%$ across the sky. When $\vartheta \geq 20^\circ$, this is reduced to $\lesssim 3.5\%$.

Additionally, we compare multiple representations of the rate-density in order to assess possible model systematics associated with each. These systematics can be thought of as correlations within the prior process for the distribution over the sky that determine the allowed shapes of the distribution. See Ref. [50] for more discussion in the context of the neutron star equation of state. These correlations can be very high dimensional and therefore difficult to visualize. What is more, statistics based on one-dimensional marginal prior distributions analogous to those shown in Figs. 2, 3, and 6 are uninformative; our priors were intentionally designed to have the same marginal distribution for the rate in all directions (although this is only approximately true for the ESH models).

Because the high-dimensional correlations induced by different model assumptions are difficult to visualize directly, we only compare a few summaries of the posterior process over the sky. Nevertheless, these show that prior assumptions can strongly affect the types of anisotropies inferred *a posteriori*. While we leave a full investigation to future work, Appendix A presents techniques to help diagnose which features are constrained by the data and which are constrained by the prior.

IV. DISCUSSION

Using 63 confidently detected GW sources from the LVK's third observing run, including BNS, NSBH, and BBH systems, we constrained the distribution of merging binaries across the celestial sphere. Our constraints improve upon previous work that used the 11 events from GWTC-1, finding constraints on anisotropies ($\mathcal{B}_{\text{ani}}^{\text{iso}}$) that are a factor of a few stronger. However, because of ambiguity in the interpretation of $\mathcal{B}_{\text{ani}}^{\text{iso}}$ due to arbitrary prior choices, we instead quantify constraints on anisotropies with a direct measure of how much $p(\Omega)$ varies over the sky. Modeling anisotropies as a Gaussian random field, we constrain the fluctuations to be $\lesssim 16\%$ if the field is correlated with a length scale $\geq 10^\circ$. That is, the distribution of merging binaries varies by $\lesssim 16\%$ of the isotropic rate at 90% credibility.

We also observe consistently hot pixels within all of our models of $p(\Omega)$. While none of these are statistically significant, the brightest pixel is in the direction of the constellation Equuleus. Our hot pixels tend to fall near the equator on average, and they do not match the hot spots found in previous work with GWTC-1 [21,23]. This is consistent with the expectation that the distribution is isotropic, and we are in effect “fitting noise” when we construct maps of the mean $p(\Omega)$.

Nonetheless, it may be interesting to extend this work in the future. In particular, we have only studied isotropy, and the cosmological principle also predicts homogeneity. It may be of interest to not only consider clustering in three




FIG. 6. Mollweide projections of the mean and significance of the rate density marginalized over correlation parameters within the GRF model with $N_{\text{pix}} = 3072$. We observe generally consistent results with the rest of our models, with a decrease in the difference between the posterior mean and isotropy as the minimum allowed correlation length (ϑ_{min}) increases.

spatial dimensions,⁸ but also correlations between intrinsic source properties (masses, spins, etc.) and extrinsic properties (location, orientation, etc.). Furthermore, correlating anisotropies and/or inhomogeneities in GW catalogs with other catalogs will be of increasing importance. Under the assumption that GW events only come from galaxies, current galaxy catalogs have been used to provide a weak constraint on the Hubble parameter [2]. With larger GW catalogs, it may be possible to directly test the assumption that GWs only come from galaxies, or to determine which types of galaxies are more likely to host GW sources [7,8]. More generally, this may constrain cosmic structure, and clustering scales in GW catalogs could connect to the mass scales of typical host galaxies [9]. Of course, there may also be synergies from connecting the distribution of nearby,

⁸The rate of GW sources almost certainly evolves over cosmic time [1,51]. This means we will need to consider the effect of lookback time when considering inhomogeneities in the spatial distribution.

well-resolved systems with the stochastic GW background from the myriad more distant sources [26].

The LVK also searches for unmodeled “burst” events. If such events are detected and no obvious source presents itself, determining whether the sources are isotropically distributed or correlated with local structure can inform the distance to the events and therefore their energy scale. Demonstrating the ability to perform such a measurement and determining the size of the catalog needed to rule out isotropy⁹ may be worth establishing before such events are detected.

For all these reasons and more, it is worth studying in greater detail which properties of individual events make them informative and over what angular scales. Indeed, as the size of the catalog grows, we may wish to know whether the isotropy constraints will always be dominated by the best-localized events or whether the legion of poorly

⁹Note that even a single event may rule out the correlation with local structure, and therefore ruling out isotropy when the events do correlate with local structure is likely to be of greater interest.

localized events will eventually dominate through sheer force of numbers. We have also shown that most of the information about (an)isotropy is carried by the best-localized events. As searches become more sensitive to quieter events and/or events detected in a single interferometer, we may expect the rate at which isotropy constraints improve to slow as a larger fraction of GW catalogs will have large, uninformative localizations.¹⁰

Here we used 63 confident BNS, NSBH, and BBH detections from O3 to place limits on the anisotropy of gravitational wave events on the sky. We do not find any evidence for anisotropy. On the contrary, using flexible and data-driven models we bound the variability of the gravitational wave merger rate over the sky to $\lesssim 16\%$ on scales larger than 10° , or $\lesssim 3.5\%$ on scales larger than 20° . As the GW catalog continues to grow, our methodology will lead to more definitive measurements. Understanding the homogeneity and isotropy of GW sources is an important astrophysical and cosmological probe of this newly discovered population. There are still many unknowns about the distribution of merging binaries, and future catalogs will continue to provide surprises if we continue to look for them.

ACKNOWLEDGMENTS

R. E. thanks the Canadian Institute for Advanced Research (CIFAR) for support. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. M. F. was supported by NASA through NASA Hubble Fellowship Grant No. HST-HF2-51455.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under

NASA Contract No. NAS5-26555. W.M.F. is partially supported by the Simons Foundation. D. E. H. is supported by NSF Grants No. PHY-2006645, No. PHY-2110507, and No. AST-2006645, as well as the Kavli Institute for Cosmological Physics through an endowment from the Kavli Foundation and its founder Fred Kavli. D. E. H. also gratefully acknowledges support from the Marion and Stuart Rice Award. E. K. is supported by the National Science Foundation (NSF) through Award No. PHY-1764464 to the LIGO Laboratory. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. This work would not have been possible without the following software: NumPy [52], SciPy [53], MATPLOTLIB [54], JAX [55], and NumPyro [56,57].

APPENDIX A: PERTURBATIVE ANALYSIS FOR SMALL ANISOTROPIES

We consider in detail how the data constrain different degrees of freedom in the distribution of merging binaries. Reference [58] introduced expressions for the Fisher information matrix that describes how constraining the data is expected to be on average. We instead consider the constraints from a particular realization of data by perturbing the likelihood directly.

In particular, we construct a model that perturbs a “base distribution” over the single-event parameters θ by a small amount. That is, we consider a rate density

$$\frac{dN}{d\theta} = \mathcal{R}p(\theta|\Lambda)(1 + \eta(\theta)) \quad (\text{A1})$$

with $|\eta| \ll 1 \ \forall \ \theta$. Inserting this into the inhomogeneous Poisson likelihood [39–41,58], we obtain

$$\begin{aligned} \ln p(\{D_i\}|\mathcal{R}, \Lambda, \eta) &= N \ln[\mathcal{R}] - \mathcal{R} \int d\theta P(\det|\theta) p(\theta|\Lambda)(1 + \eta(\theta)) + \sum_i^N \ln \left[\int d\theta p(D_i|\theta) p(\theta|\Lambda)(1 + \eta(\theta)) \right] \\ &= N \ln[\mathcal{R}] - \mathcal{R} \int d\theta P(\det|\theta) p(\theta|\Lambda) + \sum_i^N \ln \left[\int d\theta p(D_i|\theta) p(\theta|\Lambda) \right] \\ &\quad - \mathcal{R} \int d\theta P(\det|\theta) p(\theta|\Lambda) \eta(\theta) + \sum_i^N \left[\frac{\int d\theta p(D_i|\theta) p(\theta|\Lambda) \eta(\theta)}{\int d\theta p(D_i|\theta) p(\theta|\Lambda)} - \frac{1}{2} \left(\frac{\int d\theta p(D_i|\theta) p(\theta|\Lambda) \eta(\theta)}{\int d\theta p(D_i|\theta) p(\theta|\Lambda)} \right)^2 + \dots \right]. \end{aligned} \quad (\text{A2})$$

We recognize

$$\frac{p(D_i|\theta) p(\theta|\Lambda)}{\int d\theta p(D_i|\theta) p(\theta|\Lambda)} = p(\theta|D_i, \Lambda) \quad (\text{A3})$$

¹⁰Appendix A introduces an eigenvalue analysis of which anisotropies can be best constrained with current data. The magnitude of the eigenvalues rapidly decreases, suggesting that it may take many more events to precisely constrain high- l modes compared to low- l modes.




FIG. 7. Nonvanishing eigenvalues and example eigenvectors from the perturbative analysis of 63 events from O3. (top, left to right) The best-constrained eigenvector and two less constrained eigenvectors. (bottom) The distribution of eigenvalues, which decays roughly exponentially up to the 63rd eigenvalue. After that, the eigenvalues for the remaining eigenvectors are many orders of magnitude smaller.

and, retaining only terms up to second order in η , obtain

$$\begin{aligned} \ln p(\{D_i\}|\mathcal{R}, \Lambda, \eta) - p(\{D_i\}|\mathcal{R}, \Lambda, \eta = 0) = & - \int d\theta \eta(\theta) \left[\mathcal{R}P(\det|\Lambda) p(\theta|\det, \Lambda) - \sum_i^N p(\theta|D_i, \Lambda) \right] \\ & - \frac{1}{2} \int d\theta d\theta' \eta(\theta) \left[\sum_i^N p(\theta|D_i, \Lambda) p(\theta'|D_i, \Lambda) \right] \eta(\theta') \end{aligned} \quad (\text{A4})$$

We see, then, that the inhomogeneous Poisson likelihood naturally induces a Gaussian process over small perturbations away from a base distribution. In particular, the Gaussian process has a positive semidefinite inverse covariance matrix

$$\text{Cov}^{-1}[\eta(\theta), \eta(\theta')] = \sum_i^N p(\theta|D_i, \Lambda) p(\theta'|D_i, \Lambda) \quad (\text{A5})$$

Examining the mean vector in more detail, we see that it is proportional to the difference of two terms. Taking the maximum likelihood estimate for $\mathcal{R}$ conditioned on Λ and N , we expect $\mathcal{R}P(\det|\Lambda) = N$. Therefore, the mean vector is proportional to the difference between the distribution over θ for detectable sources and the average of the single-event posteriors. This is intuitively appealing and explains why stacking (adding) posteriors can often produce useful diagnostics even if it is not the correct way to perform a hierarchical inference [39]. This is also why the mean of the GRF model in Sec. III C at times displays features reminiscent of individual events. Some events are well localized

relative to $p(\theta|\det, \Lambda)$ and therefore the mean looks as if we simply summed the posteriors of each event (compare Figs. 6 and 9).

We can also consider which types of features are constrained by the data by examining the eigenvectors and eigenvalues of the inverse covariance matrix. While this analysis is completely general,¹¹ we specialize to the case at hand: an isotropic base distribution with masses, spins, and redshifts distributed as in Table I. We also only perturb the distribution over the sky. Immediately, we see that the only pixels that are constrained by the data are

¹¹Similar “semiparametric” analyses have been conducted for the mass distribution, although they implemented a spline model for the deviations from the base model [59]. However, considering the full Gaussian process induced by the likelihood and adopting a conjugate prior may allow for a clearer determination of exactly which features are driven by the data and which are driven by the prior assumptions, particularly when the correlations in the prior span high-dimensional spaces. Furthermore, this type of perturbative analysis can be conducted completely *post hoc* given any base distribution, even semiparametric or non-parametric representations of $p(\theta|\Lambda)$.




FIG. 8. Distribution of detected injections from O3 [35] assuming an isotropic population with mass, spin, and redshift distributions listed in Table I. (left) Scatter plot of detected events in a Mollweide projection. (right) Marginal distributions of the detected population (red) and the isotropic distribution (black) for reference. Shaded regions correspond to 1, 2, and 3 σ uncertainty on the detected distribution's marginals from the finite number of injections.

those that have nonzero probability of containing at least one event under the base model [$p(\theta|D_i, \Lambda) \neq 0$ for at least one D_i]. It is natural to control these poorly constrained eigenvectors with a Gaussian process prior, like the one introduced in Sec. III C. Indeed, if we fix the base distribution (including the rate), then we can construct a posterior for η analytically.

What is more, the magnitude of the inverse-covariance matrix's eigenvalues rapidly decays. Figure 7 demonstrates this with our selection of 63 events from O3. As such, we can always expect there to be many eigenvectors that are dominated by the prior for any finite catalog. Figure 7 also shows a few eigenvectors. Typically, the best-constrained eigenvectors resemble well-localized individual events, or just a few pixels on the sky, while less constrained eigenvectors resemble the overlap of multiple events.

APPENDIX B: SELECTED EVENTS AND CATALOG SENSITIVITY

For completeness, we present the estimates of our survey's sensitivity (false alarm rate for any search $\leq 1/\text{year}$) across the sky assuming the mass, spin, and redshift populations in Table I. As the injected distribution within Ref. [35] is somewhat complicated and not particularly astrophysically motivated, we have reweighed the injections so the injected distribution matches the distributions listed in Table I. Figure 8 shows the distribution of detected events from an isotropic source distribution.

While the search sensitivity is nearly uniform, we do observe slight excesses of detected injections from the midlatitudes and a dearth of detections near the equator, in agreement with Fig. 1 of Ref. [22]. We also note that the diurnal cycle identified by Ref. [22] during the first observing run (O1) is not apparent in O3. This is likely due to a combination of factors: the detector duty cycles

were higher in O3 than in O1 [60], and O3 lasted for nearly a full calendar year, thereby washing out the impact of a diurnal cycle (determined by the Earth's orientation to the Sun) when projected on the celestial sphere.

Figure 9 shows the superposition of all the individual events' localizations assuming an isotropic distribution and reweighing individual events' posterior samples to match the mass, spin, and redshift populations listed in Table I. Overdensities of points correspond to hotspots in Fig. 6, as expected based on the analysis in Appendix A. Table III shows the medians and 90% symmetric credible intervals for the component masses, spins, and redshifts of each of the 63 selected events after reweighing the posteriors samples so the prior matches the distributions in Table I and an isotropic distribution over the sky.

Single-event posterior samples for all events detected during the first half of O3 (O3a; GWTC-2 [32]) were taken


















FIG. 9. Superposition of localization estimates from all 63 events considered in this study. Each point is a fair draw from one event's posterior assuming an isotropic distribution over the sky and the mass, spin, and redshift distributions in Table I. Darker shading corresponds to areas with many overlapping events or extremely well localized events, and roughly correspond to the hot-spots seen in the posterior means in Fig. 6. See Table III for individual events' localizations.

TABLE III. Medians and 90% symmetric credible regions for each of our 63 events assuming an isotropic source distribution and the mass, spin, and redshift distributions from Table I. We also show the smallest area on the sky that contains 90% of the posterior probability and a scatter plot of posterior samples. Brighter colors in scatter plots correspond to higher likelihoods, and each point is a fair draw from the posterior.

Name	$m_1 [M_\odot]$	$m_2 [M_\odot]$	z	D_L [Mpc]	$\Delta\Omega_{90\%} [\text{deg}^2]$	Skymap
GW190408_181802	$22.89^{+4.14}_{-2.11}$	$19.88^{+2.60}_{-3.20}$	$0.29^{+0.07}_{-0.10}$	$1540.36^{+427.70}_{-601.78}$	179.2	
GW190412	$26.10^{+5.91}_{-6.26}$	$9.45^{+3.06}_{-1.53}$	$0.14^{+0.03}_{-0.04}$	$694.54^{+174.19}_{-217.11}$	83.3	
GW190413_052954	$29.96^{+8.25}_{-5.29}$	$25.65^{+6.03}_{-5.34}$	$0.60^{+0.26}_{-0.24}$	$3633.56^{+2041.16}_{-1627.01}$	1425.1	
GW190413_134308	$41.19^{+10.54}_{-7.08}$	$35.24^{+8.75}_{-7.96}$	$0.74^{+0.30}_{-0.31}$	$4734.55^{+2403.83}_{-2238.56}$	562.5	
GW190421_213856	$38.18^{+7.54}_{-5.33}$	$33.24^{+6.73}_{-5.93}$	$0.50^{+0.19}_{-0.20}$	$2936.10^{+1363.16}_{-1339.61}$	1033.4	
GW190424_180648	$37.30^{+7.63}_{-5.57}$	$32.78^{+6.24}_{-5.79}$	$0.41^{+0.21}_{-0.19}$	$2297.44^{+1483.64}_{-1174.10}$	25972.7	
GW190425	$1.90^{+0.45}_{-0.23}$	$1.44^{+0.18}_{-0.25}$	$0.04^{+0.01}_{-0.02}$	$161.23^{+65.85}_{-72.41}$	8517.3	
GW190503_185404	$38.26^{+7.94}_{-5.17}$	$31.87^{+5.56}_{-7.00}$	$0.29^{+0.11}_{-0.12}$	$1519.47^{+692.77}_{-695.95}$	108.3	
GW190512_180714	$18.86^{+5.65}_{-2.52}$	$15.12^{+2.33}_{-3.04}$	$0.28^{+0.09}_{-0.10}$	$1470.52^{+577.61}_{-595.76}$	245.9	
GW190513_205428	$28.33^{+7.97}_{-3.93}$	$23.24^{+4.30}_{-5.91}$	$0.38^{+0.13}_{-0.15}$	$2096.74^{+900.07}_{-941.55}$	462.5	
GW190517_055101	$33.08^{+7.38}_{-5.58}$	$27.98^{+5.03}_{-5.61}$	$0.36^{+0.23}_{-0.15}$	$1998.60^{+1622.41}_{-940.58}$	429.2	
GW190519_153544	$59.04^{+11.70}_{-12.54}$	$44.10^{+9.68}_{-10.02}$	$0.49^{+0.30}_{-0.16}$	$2835.40^{+2212.13}_{-1092.74}$	820.9	
GW190521	$82.90^{+19.55}_{-12.45}$	$70.36^{+17.39}_{-13.88}$	$0.72^{+0.29}_{-0.28}$	$4514.92^{+2346.49}_{-2003.55}$	887.6	
GW190521_074359	$40.25^{+5.25}_{-3.87}$	$34.61^{+4.23}_{-5.26}$	$0.24^{+0.07}_{-0.10}$	$1252.81^{+405.74}_{-552.87}$	491.7	
GW190527_092055	$30.70^{+8.05}_{-5.44}$	$25.67^{+6.36}_{-6.12}$	$0.44^{+0.27}_{-0.20}$	$2506.95^{+1986.41}_{-1268.74}$	3329.4	
GW190602_175927	$61.27^{+12.96}_{-9.71}$	$51.57^{+10.20}_{-11.95}$	$0.51^{+0.25}_{-0.18}$	$2982.27^{+1881.03}_{-1207.87}$	725.1	
















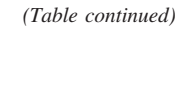
(Table continued)

TABLE III. (*Continued*)

Name	$m_1 [M_\odot]$	$m_2 [M_\odot]$	z	D_L [Mpc]	$\Delta\Omega_{90\%} [\text{deg}^2]$	Skymap
GW190620_030421	$48.09^{+12.52}_{-7.76}$	$39.74^{+8.43}_{-9.34}$	$0.54^{+0.21}_{-0.21}$	$3185.49^{+1588.11}_{-1422.76}$	6158.8	
GW190630_185205	$31.73^{+5.66}_{-3.57}$	$26.65^{+3.24}_{-4.89}$	$0.17^{+0.11}_{-0.06}$	$837.61^{+619.59}_{-336.30}$	1558.4	
GW190701_203306	$49.52^{+8.74}_{-5.85}$	$43.15^{+6.51}_{-8.13}$	$0.38^{+0.11}_{-0.11}$	$2096.04^{+738.49}_{-712.16}$	66.7	
GW190706_222641	$55.49^{+15.88}_{-10.15}$	$43.59^{+10.47}_{-10.37}$	$0.82^{+0.29}_{-0.31}$	$5335.31^{+2430.24}_{-2363.03}$	620.9	
GW190707_093326	$10.58^{+1.95}_{-0.97}$	$9.06^{+0.90}_{-1.32}$	$0.17^{+0.06}_{-0.08}$	$857.19^{+333.14}_{-427.25}$	1416.8	
GW190708_232457	$16.31^{+2.76}_{-1.32}$	$14.20^{+1.37}_{-2.16}$	$0.17^{+0.06}_{-0.07}$	$876.65^{+334.85}_{-383.05}$	10021.6	
GW190719_215514	$28.97^{+9.29}_{-5.41}$	$24.50^{+6.37}_{-5.60}$	$0.64^{+0.31}_{-0.28}$	$3942.61^{+2475.88}_{-1938.22}$	2579.4	
GW190720_000836	$10.94^{+2.69}_{-1.29}$	$9.21^{+1.07}_{-1.71}$	$0.18^{+0.11}_{-0.07}$	$882.93^{+668.84}_{-372.22}$	616.7	
GW190725_174728	$9.32^{+2.69}_{-1.01}$	$7.88^{+0.99}_{-1.64}$	$0.20^{+0.10}_{-0.09}$	$1034.40^{+598.20}_{-476.09}$	2162.7	
GW190727_060333	$35.37^{+6.95}_{-4.66}$	$30.87^{+5.82}_{-5.50}$	$0.56^{+0.20}_{-0.22}$	$3367.62^{+1488.39}_{-1502.40}$	741.7	
GW190728_064510	$10.74^{+2.26}_{-0.85}$	$9.25^{+0.87}_{-1.49}$	$0.17^{+0.05}_{-0.07}$	$856.36^{+271.90}_{-355.00}$	325.0	
GW190731_140936	$36.95^{+9.29}_{-6.92}$	$31.37^{+7.43}_{-7.46}$	$0.58^{+0.31}_{-0.26}$	$3535.48^{+2373.90}_{-1769.26}$	3091.9	
GW190803_022701	$33.81^{+7.51}_{-5.18}$	$29.19^{+6.34}_{-5.74}$	$0.57^{+0.24}_{-0.24}$	$3412.79^{+1822.16}_{-1610.66}$	1458.4	
GW190805_211137	$41.34^{+11.95}_{-8.35}$	$34.74^{+9.88}_{-8.26}$	$0.89^{+0.44}_{-0.39}$	$5914.41^{+3760.37}_{-2975.96}$	3533.6	
GW190814	$23.41^{+1.12}_{-0.92}$	$2.57^{+0.08}_{-0.09}$	$0.05^{+0.01}_{-0.01}$	$245.77^{+40.13}_{-43.26}$	29.2	
GW190828_063405	$30.67^{+4.91}_{-3.15}$	$27.12^{+4.26}_{-3.70}$	$0.38^{+0.10}_{-0.15}$	$2128.25^{+668.61}_{-928.63}$	475.0	















(*Table continued*)

TABLE III. (*Continued*)

Name	$m_1 [M_\odot]$	$m_2 [M_\odot]$	z	D_L [Mpc]	$\Delta\Omega_{90\%} [\text{deg}^2]$	Skymap
GW190828_065509	$17.39^{+6.81}_{-2.57}$	$13.44^{+2.19}_{-3.22}$	$0.32^{+0.10}_{-0.11}$	$1714.51^{+658.82}_{-682.49}$	745.9	
GW190910_112807	$41.62^{+6.16}_{-5.53}$	$36.51^{+5.31}_{-5.82}$	$0.29^{+0.16}_{-0.10}$	$1535.74^{+1021.35}_{-604.01}$	9838.2	
GW190915_235702	$30.84^{+5.22}_{-3.61}$	$26.76^{+3.91}_{-4.28}$	$0.31^{+0.10}_{-0.11}$	$1699.15^{+678.11}_{-649.29}$	362.5	
GW190924_021846	$7.20^{+1.90}_{-0.59}$	$6.06^{+0.56}_{-1.11}$	$0.12^{+0.04}_{-0.05}$	$572.77^{+215.26}_{-229.08}$	358.4	
GW190929_012149	$50.07^{+17.25}_{-11.85}$	$35.38^{+11.44}_{-11.85}$	$0.71^{+0.35}_{-0.30}$	$4440.24^{+2803.54}_{-2138.08}$	1954.3	
GW190930_133541	$10.58^{+2.24}_{-0.95}$	$9.02^{+0.91}_{-1.48}$	$0.15^{+0.06}_{-0.06}$	$758.74^{+351.63}_{-328.50}$	1683.5	
GW191103_012549	$10.33^{+2.24}_{-0.97}$	$8.98^{+1.01}_{-1.61}$	$0.19^{+0.08}_{-0.09}$	$968.08^{+488.21}_{-468.99}$	2558.5	
GW191105_143521	$9.69^{+1.96}_{-0.86}$	$8.41^{+0.89}_{-1.42}$	$0.22^{+0.07}_{-0.09}$	$1147.37^{+406.73}_{-479.83}$	820.9	
GW191109_010717	$60.18^{+10.44}_{-10.59}$	$48.18^{+12.55}_{-11.35}$	$0.27^{+0.22}_{-0.13}$	$1409.53^{+1453.76}_{-760.26}$	1579.3	
GW191127_050227	$32.60^{+13.87}_{-6.59}$	$28.11^{+10.78}_{-7.04}$	$0.52^{+0.42}_{-0.33}$	$3041.05^{+3244.28}_{-2094.05}$	1029.2	
GW191129_134029	$9.18^{+2.47}_{-0.79}$	$7.77^{+0.85}_{-1.58}$	$0.15^{+0.05}_{-0.06}$	$765.83^{+269.14}_{-328.52}$	1333.4	
GW191204_171526	$10.82^{+2.41}_{-0.95}$	$9.03^{+0.85}_{-1.54}$	$0.13^{+0.04}_{-0.05}$	$632.63^{+192.64}_{-234.48}$	329.2	
GW191215_223052	$22.85^{+4.53}_{-2.76}$	$19.76^{+2.98}_{-3.37}$	$0.34^{+0.14}_{-0.14}$	$1879.79^{+940.85}_{-851.88}$	562.5	
GW191216_213338	$10.40^{+2.76}_{-0.83}$	$8.82^{+0.73}_{-1.70}$	$0.07^{+0.02}_{-0.03}$	$338.99^{+113.86}_{-133.57}$	241.7	
GW191222_033537	$41.82^{+8.59}_{-6.58}$	$36.40^{+7.70}_{-7.21}$	$0.50^{+0.24}_{-0.24}$	$2903.37^{+1792.76}_{-1574.71}$	1991.8	
GW191230_180458	$44.21^{+9.88}_{-7.17}$	$38.46^{+8.55}_{-7.67}$	$0.72^{+0.25}_{-0.27}$	$4529.72^{+2043.75}_{-1969.25}$	1100.1	

(Table continued)

TABLE III. (*Continued*)

Name	$m_1 [M_\odot]$	$m_2 [M_\odot]$	z	D_L [Mpc]	$\Delta\Omega_{90\%} [\text{deg}^2]$	Skymap
GW200105_162426	$9.13^{+2.40}_{-1.57}$	$1.89^{+0.29}_{-0.31}$	$0.06^{+0.02}_{-0.02}$	$271.86^{+113.79}_{-112.24}$	7496.4	
GW200112_155838	$33.67^{+5.33}_{-3.29}$	$29.29^{+3.40}_{-4.90}$	$0.24^{+0.07}_{-0.09}$	$1276.29^{+431.75}_{-490.32}$	3204.4	
GW200115_042309	$6.32^{+1.19}_{-1.08}$	$1.37^{+0.23}_{-0.16}$	$0.06^{+0.03}_{-0.02}$	$288.40^{+128.24}_{-93.89}$	366.7	
GW200128_022011	$38.41^{+9.03}_{-6.73}$	$33.34^{+8.24}_{-6.95}$	$0.57^{+0.28}_{-0.27}$	$3399.78^{+2147.79}_{-1801.03}$	2466.8	
GW200129_065458	$33.66^{+5.33}_{-2.63}$	$29.55^{+3.27}_{-5.16}$	$0.19^{+0.04}_{-0.07}$	$942.17^{+250.98}_{-387.45}$	45.8	
GW200202_154313	$9.17^{+1.81}_{-0.60}$	$8.09^{+0.59}_{-1.28}$	$0.09^{+0.03}_{-0.04}$	$411.93^{+140.20}_{-176.53}$	158.3	
GW200208_130117	$34.53^{+6.52}_{-4.44}$	$29.72^{+4.72}_{-5.62}$	$0.40^{+0.14}_{-0.14}$	$2244.01^{+992.89}_{-901.31}$	33.3	
GW200209_085452	$32.47^{+7.19}_{-5.42}$	$27.99^{+6.79}_{-5.49}$	$0.56^{+0.24}_{-0.29}$	$3352.15^{+1852.25}_{-1931.87}$	1025.1	
GW200216_220804	$41.65^{+11.55}_{-7.80}$	$35.68^{+8.75}_{-9.53}$	$0.68^{+0.35}_{-0.34}$	$4272.53^{+2811.53}_{-2392.08}$	3104.4	
GW200219_094415	$34.11^{+7.64}_{-5.13}$	$29.40^{+6.34}_{-5.71}$	$0.58^{+0.22}_{-0.24}$	$3522.91^{+1680.98}_{-1655.16}$	745.9	
GW200224_222234	$38.35^{+5.64}_{-3.64}$	$33.67^{+4.29}_{-5.22}$	$0.32^{+0.08}_{-0.12}$	$1725.50^{+488.84}_{-691.54}$	50.0	
GW200225_060421	$17.91^{+3.22}_{-2.11}$	$15.44^{+1.96}_{-2.93}$	$0.22^{+0.09}_{-0.09}$	$1104.17^{+526.97}_{-493.85}$	616.7	
GW200302_015811	$31.47^{+8.84}_{-5.68}$	$24.26^{+5.11}_{-6.64}$	$0.30^{+0.16}_{-0.13}$	$1608.82^{+1041.31}_{-780.03}$	8684.0	
GW200311_115853	$32.60^{+4.69}_{-3.04}$	$28.83^{+3.49}_{-4.53}$	$0.23^{+0.05}_{-0.08}$	$1169.68^{+284.52}_{-433.83}$	45.8	
GW200316_215756	$11.00^{+3.10}_{-1.17}$	$9.20^{+0.97}_{-1.77}$	$0.22^{+0.08}_{-0.08}$	$1120.21^{+448.67}_{-437.91}$	370.9	

from Ref. [61], with the exception of two events first published in GWTC-2.1 [33]: GW190725_174728 and GW190805_211137. Posterior samples for these events

are available in Ref. [62]. Samples for events from the second half of O3 (O3b; GWTC-3 [34]) are available within Ref. [63].

Although only used to benchmark our constraints from O3, posterior samples for events from GWTC-1 are available in Ref. [64]. Because these samples do not include all

the Cartesian spin components and because spin inference largely decouples from localization, we do not include the prior for the spin within analyses of GWTC-1.

-
- [1] R. Abbott *et al.*, The population of merging compact binaries inferred using gravitational waves through GWTC-3, [arXiv:2111.03634](#).
 - [2] R. Abbott *et al.*, Constraints on the cosmic expansion history from GWTC-3, [arXiv:2111.03604](#).
 - [3] S. Mukherjee, B.D. Wandelt, S.M. Nissanke, and A. Silvestri, Accurate precision cosmology with redshift unknown gravitational wave sources, *Phys. Rev. D* **103**, 043520 (2021).
 - [4] S. Vitale, S. Biscoveanu, and C. Talbot, The orientations of the binary black holes in GWTC-3, [arXiv:2204.00968](#).
 - [5] S. Libanore, M.C. Artale, D. Karagiannis, M. Liguori, N. Bartolo, Y. Bouffanais, N. Giacobbo, M. Mapelli, and S. Matarrese, Gravitational wave mergers as tracers of large scale structures, *J. Cosmol. Astropart. Phys.* **02** (2021) 035.
 - [6] X. Shao, Z. Cao, X. Fan, and S. Wu, Probing the large-scale structure of the universe through gravitational wave observations, *Res. Astron. Astrophys.* **22**, 015006 (2022).
 - [7] H.-Y. Chen and D. E. Holz, Finding the one: Identifying the host galaxies of gravitational-wave sources, [arXiv:1612.01471](#).
 - [8] L. P. Singer *et al.*, Going the distance: Mapping host galaxies of LIGO and Virgo sources in three dimensions using local cosmography and targeted follow-up, *Astrophys. J.* **829**, L15 (2016).
 - [9] M. Fishbach and V. Kalogera, The time delay distribution and formation metallicity of LIGO-Virgos binary black holes, *Astrophys. J. Lett.* **914**, L30 (2021).
 - [10] M. Zevin, A. E. Nugent, S. Adhikari, W.-f. Fong, D. E. Holz, and L. Z. Kelley, Observational inference on the delay time distribution of short gamma-ray bursts, [arXiv:2206.02814](#).
 - [11] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [12] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [13] R. Abbott *et al.* (The LIGO Scientific, the Virgo, and the KAGRA Collaborations), All-sky search for short gravitational-wave bursts in the third Advanced LIGO and Advanced Virgo run, *Phys. Rev. D* **104**, 122004 (2021).
 - [14] U. Andrade, C. A. P. Bengaly, J. S. Alcaniz, and S. Capozziello, Revisiting the statistical isotropy of GRB sky distribution, *Mon. Not. R. Astron. Soc.* **490**, 4481 (2019).
 - [15] J. B. Dent, W. E. Gabella, K. Holley-Bockelmann, and T. W. Kephart, Gravitational waves from a black hole orbiting in a wormhole geometry, *Phys. Rev. D* **104**, 044030 (2021).
 - [16] F. Xu, J. M. Ezquiaga, and D. E. Holz, Please repeat: Strong lensing of gravitational waves as a probe of compact binary and galaxy populations, *Astrophys. J.* **929**, 9 (2022).
 - [17] M. Çalışkan, J. María Ezquiaga, O. A. Hannuksela, and D. E. Holz, Lensing or luck? False alarm probabilities for gravitational lensing of gravitational waves, [arXiv:2201.04619](#).
 - [18] J. M. Ezquiaga, W. Hu, M. Lagos, and M.-X. Lin, Gravitational wave propagation beyond general relativity: Waveform distortions and echoes, *J. Cosmol. Astropart. Phys.* **11** (2021) 048.
 - [19] J. M. Ezquiaga, W. Hu, M. Lagos, M.-X. Lin, and F. Xu, Modified gravitational wave propagation with higher modes and its degeneracies with lensing, [arXiv:2203.13252](#).
 - [20] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo During the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
 - [21] R. Stiskalek, J. Veitch, and C. Messenger, Are stellar-mass binary black hole mergers isotropically distributed?, *Mon. Not. R. Astron. Soc.* **501**, 970 (2020).
 - [22] H.-Y. Chen, R. Essick, S. Vitale, D. E. Holz, and E. Katsavounidis, Observational selection effects with ground-based gravitational wave detectors, *Astrophys. J.* **835**, 34 (2017).
 - [23] E. Payne, S. Banagiri, P. D. Lasky, and E. Thrane, Searching for anisotropy in the distribution of binary black hole mergers, *Phys. Rev. D* **102**, 102004 (2020).
 - [24] M. Cavaglia and A. Modi, Two-dimensional correlation function of binary black hole coalescences, *Universe* **6**, 93 (2020).
 - [25] A. I. Renzini, B. Goncharov, A. C. Jenkins, and P. M. Meyers, Stochastic gravitational-wave backgrounds: Current detection efforts and future prospects, *Galaxies* **10**, 34 (2022).
 - [26] T. Callister, M. Fishbach, D. E. Holz, and W. M. Farr, Shouts and murmurs: Combining individual gravitational-wave sources with the stochastic background to measure the history of binary black hole mergers, *Astrophys. J.* **896**, L32 (2020).
 - [27] J. M. Cordes and S. Chatterjee, Fast radio bursts: An extragalactic enigma, *Annu. Rev. Astron. Astrophys.* **57**, 417 (2019).
 - [28] A. U. Abeysekara *et al.*, Observation of small-scale anisotropy in the arrival direction distribution of TeV cosmic rays with HAWC, *Astrophys. J.* **796**, 108 (2014).

- [29] G. Illuminati, Study of the high energy cosmic rays large scale anisotropies with the ANTARES neutrino telescope, *J. Phys. Conf. Ser.* **689**, 012012 (2016).
- [30] A. Aab *et al.*, Large-scale cosmic-ray anisotropies above 4 EeV measured by the Pierre Auger Observatory, *Astrophys. J.* **868**, 4 (2018).
- [31] F. McNally, R. Abbasi, P. Desiati, J. C. Díaz Vélez, T. Aguado, K. Gruchot, A. Moy, A. Simmons, A. Thorpe, and H. Woodward, Observation of cosmic ray anisotropy with nine years of IceCube data, [arXiv:2107.11454](https://arxiv.org/abs/2107.11454).
- [32] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, *Phys. Rev. X* **11**, 021053 (2021).
- [33] R. Abbott *et al.*, GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
- [34] R. Abbott *et al.*, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
- [35] R. Abbott *et al.*, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run—O3 search sensitivity estimates, [10.5281/zenodo.5546676](https://doi.org/10.5281/zenodo.5546676) (2021).
- [36] A. M. Farah, M. Fishbach, R. Essick, D. E. Holz, and S. Galadage, Bridging the gap: Categorizing gravitational-wave events at the transition between neutron stars and black holes, *Astrophys. J.* **931**, 108 (2022).
- [37] E.-J. Wagenmakers, T. Lodewyckx, H. Kuriyal, and R. Grasman, Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method, *Cogn. Psychol.* **60**, 158 (2010).
- [38] N. Aghanim *et al.*, Planck 2018 results—VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [39] T. J. Loredo, Accounting for source uncertainties in analyses of astronomical survey data, in *Proceedings of the Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, American Institute of Physics Conference Series Vol. 735, edited by R. Fischer, R. Preuss, and U. V. Toussaint (American Institute of Physics Conference Series, 2004), pp. 195–206.
- [40] I. Mandel, Parameter estimation on gravitational waves from multiple coalescing binaries, *Phys. Rev. D* **81**, 084029 (2010).
- [41] I. Mandel, W. M. Farr, and J. R. Gair, Extracting distribution parameters from multiple uncertain observations with selection biases, *Mon. Not. R. Astron. Soc.* **486**, 1086 (2019).
- [42] R. Essick and W. Farr, Precision requirements for Monte Carlo sums within hierarchical Bayesian inference, [arXiv:2204.00461](https://arxiv.org/abs/2204.00461).
- [43] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [44] E. Delporte and I. A. Union, *Délimitation Scientifique des Constellations: (tables et cartes)*, Report/Commission 3 of the International Astronomical Union (At the University Press, Cambridge, UK, 1930).
- [45] R. Essick and P. Landry, Discriminating between neutron stars and black holes with imperfect knowledge of the maximum neutron star mass, *Astrophys. J.* **904**, 80 (2020).
- [46] M. Isi, W. M. Farr, and K. Chatziioannou, Comparing Bayes factors and hierarchical inference for testing general relativity with gravitational waves, [arXiv:2204.10742](https://arxiv.org/abs/2204.10742).
- [47] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning series (MIT Press, Cambridge, MA, 2005).
- [48] A. Lang and C. Schwab, Isotropic Gaussian random fields on the sphere: Regularity, fast simulation and stochastic partial differential equations, *Ann. Appl. Probab.* **25**, 3047 (2015).
- [49] B. Farr, W. M. Farr, N. B. Cowan, H. M. Haggard, and T. Robinson, exocartographer: A Bayesian framework for mapping exoplanets in reflected light, *Astron. J.* **156**, 146 (2018).
- [50] I. Legred, K. Chatziioannou, R. Essick, and P. Landry, Implicit correlations within phenomenological parametric models of the neutron star equation of state, *Phys. Rev. D* **105**, 043016 (2022).
- [51] M. Fishbach, Z. Doctor, T. Callister, B. Edelman, J. Ye, R. Essick, W. M. Farr, B. Farr, and D. E. Holz, When are LIGO/Virgo’s big black hole mergers?, *Astrophys. J.* **912**, 98 (2021).
- [52] C. R. Harris *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [53] P. Virtanen *et al.* (SciPy 1.0 Contributors), SciPy 1.0: Fundamental algorithms for scientific computing in PYTHON, *Nat. Methods* **17**, 261 (2020).
- [54] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [55] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necoara, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python + NumPy programs (2018), <https://github.com/google/jax# citing-jax>.
- [56] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. A. Szerlip, P. Horsfall, and N. D. Goodman, Pyro: Deep universal probabilistic programming, *J. Mach. Learn. Res.* **20**, 28:1 (2019).
- [57] D. Phan, N. Pradhan, and M. Jankowiak, Composable effects for flexible and accelerated probabilistic programming in NumPyro, [arXiv:1912.11554](https://arxiv.org/abs/1912.11554).
- [58] J. R. Gair, A. Antonelli, and R. Barbieri, A Fisher matrix for gravitational-wave population inference, [arXiv:2205.07893](https://arxiv.org/abs/2205.07893).
- [59] B. Edelman, Z. Doctor, J. Godfrey, and B. Farr, Ain’t no mountain high enough: Semiparametric modeling of LIGO–Virgo’s binary black hole mass distribution, *Astrophys. J.* **924**, 101 (2022).
- [60] A. Buikema *et al.*, Sensitivity and performance of the Advanced LIGO detectors in the third observing run, *Phys. Rev. D* **102**, 062003 (2020).

- [61] R. Abbott *et al.*, GWTC-2 data release: Parameter estimation samples and skymaps, <https://dcc.ligo.org/LIGO-P2000223-v7/public> (2021).
- [62] R. Abbott *et al.*, GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run—parameter estimation data release, [10.5281/zenodo.5117703](https://zenodo.org/record/5117703) (2021).
- [63] LIGO Scientific, Virgo, and KAGRA Collaborations, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run—parameter estimation data release, [10.5281/zenodo.5546663](https://zenodo.org/record/5546663) (2021).
- [64] R. Abbott *et al.*, Parameter estimation sample release for gWTC-1, <https://dcc.ligo.org/LIGO-P1800370/public> (2020).