



Gravitational-wave Geodesy: A New Tool for Validating Detection of the Stochastic Gravitational-wave Background

T. A. Callister¹, M. W. Coughlin¹, and J. B. Kanner¹

LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA; tcallist@caltech.edu

Received 2018 August 13; revised 2018 November 20; accepted 2018 November 23; published 2018 December 14

Abstract

A valuable target for advanced gravitational-wave detectors is the stochastic gravitational-wave background. The stochastic background imparts a weak correlated signal into networks of gravitational-wave detectors, and so standard searches for the gravitational-wave background rely on measuring cross-correlations between pairs of widely separated detectors. Stochastic searches, however, can be affected by any other correlated effects that may also be present, including correlated frequency combs and magnetic Schumann resonances. As stochastic searches become sensitive to ever-weaker signals, it is increasingly important to develop methods to separate a true astrophysical signal from other spurious and/or terrestrial signals. Here, we describe a novel method to achieve this goal—gravitational-wave geodesy. Just as radio geodesy allows for the localization of radio telescopes, so too can observations of the gravitational-wave background be used to infer the positions and orientations of gravitational-wave detectors. By demanding that a true observation of the gravitational-wave background yield constraints that are consistent with the baseline’s known geometry, we demonstrate that we can successfully validate true observations of the gravitational-wave background while rejecting spurious signals due to correlated terrestrial effects.

Key words: gravitational waves – methods: data analysis – methods: statistical

1. Introduction

The recent Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO–Virgo observations of binary black hole (Abbott et al. 2016a, 2017a, 2017b, 2017c) and binary neutron star (Abbott et al. 2017d) mergers suggest that the astrophysical stochastic gravitational-wave background may soon be within reach (Abbott et al. 2016b, 2017e, 2017f, 2018a). As the superposition of all gravitational-wave signals too weak to individually detect, the stochastic gravitational-wave background is expected to be dominated by compact binary mergers at cosmological distances (Regimbau & Mandic 2008; Rosado 2011; Zhu et al. 2011, 2013; Wu et al. 2012; Callister et al. 2016). Although the stochastic background is orders of magnitude weaker than instrumental detector noise, it will impart a weak *correlated* signal to pairs of gravitational-wave detectors. The stochastic background may therefore be detected in the form of excess correlations between widely separated gravitational-wave detectors (Christensen 1992; Allen & Romano 1999; Romano & Cornish 2017).

Cross-correlation searches for the stochastic background rely on the assumption that, in the absence of a gravitational-wave signal, the outputs of different gravitational-wave detectors are fundamentally uncorrelated. The LIGO-Hanford and LIGO-Livingston detectors, for instance, are separated by 3000 km, with a light travel time of ≈ 0.01 s between sites. One might therefore reasonably expect them to be safely uncorrelated at $\sim \mathcal{O}(100$ Hz), in the frequency band of interest for ground-based detectors.

In reality, however, terrestrial gravitational-wave detectors are *not* truly uncorrelated. Hanford-Livingston coherence spectra consistently show correlated features that, if not properly identified and removed, can severely contaminate searches for the stochastic gravitational-wave background (Covas et al. 2018). Schumann resonances are one expected source of terrestrial correlation (Schumann 1952a, 1952b). Global electromagnetic excitations in

the cavity formed by the Earth and its ionosphere, Schumann resonances may magnetically couple to Advanced LIGO and Advanced Virgo’s test mass suspensions and induce a correlated signal between detectors (Christensen 1992; Thrane et al. 2013, 2014; Coughlin et al. 2016, 2018). Another expected source of correlation is the joint synchronization of electronics at each detector to Global Positioning System (GPS) time. In Advanced LIGO’s O1 observing run, for instance, a strongly correlated 1 Hz comb was traced to blinking LED indicators on timing systems independently synchronized to GPS (Covas et al. 2018).

Any undiagnosed terrestrial correlations may yield a false-positive detection of the stochastic gravitational-wave background. While Schumann resonances and frequency combs represent two known classes of correlation, others may also exist. The validation of any apparent observation of the stochastic background will therefore require us to answer the following question. *How likely is an observed correlated signal to be of astrophysical origin, rather than a yet-unidentified source of terrestrial correlation?*

We currently lack the tools to quantitatively answer this question. Searches for gravitational-wave transients can address this issue through the use of time-slides: the artificial time-shifting of data from one detector relative to another’s. This process eliminates any coherent gravitational-wave signals while preserving all other properties of the data, allowing for accurate estimation of the false-positive detection rate. In cross-correlation searches for the stochastic background, however, time-slides would not only remove a gravitational-wave signal but also any correlated terrestrial contamination. Time-slides are therefore of limited use in searches for the gravitational-wave background.

Using techniques borrowed from the field of radio geodesy, here we develop a novel method to evaluate the astrophysical significance of an apparent correlated stochastic signal. Just as interferometric observations of the radio sky can serve to

precisely localize radio telescopes on the Earth, we demonstrate that measurements of the gravitational-wave background can be similarly reverse-engineered to infer the separations and relative orientations of gravitational-wave detectors. By demanding that a true gravitational-wave background yield results consistent with the *known* geometry of our detectors, we can separate true gravitational-wave signals from spurious terrestrial correlations.

First, in Section 2, we review search methods for the stochastic gravitational-wave background and introduce gravitational-wave geodesy. In Section 3, we use geodesy as the basis of a Bayesian test with which to reject non-astrophysical signals, and in Section 4, we demonstrate this procedure using simulated measurements of both a gravitational-wave background and terrestrial sources of correlation. Finally, in Appendix C, we discuss potential complications and outline directions for future work.

2. Gravitational-wave Geodesy

The stochastic background is typically described via its energy-density spectrum $\Omega(f)$, defined as the energy density $d\rho_{\text{GW}}$ of gravitational waves per logarithmic frequency interval $d \ln f$ (Allen & Romano 1999; Romano & Cornish 2017):

$$\Omega(f) = \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}. \quad (1)$$

The energy-density spectrum is made dimensionless by dividing by the universe’s closure energy density $\rho_c = 3H_0^2 c^2 / (8\pi G)$, where H_0 is the Hubble constant, c is the speed of light, and G is Newton’s constant.

Searches for the stochastic background seek to measure $\Omega(f)$ by computing the cross-correlation spectrum $\hat{C}(f)$ between pairs of gravitational-wave detectors:

$$\hat{C}(f) = \frac{1}{\Delta T} \frac{20\pi^2}{3H_0^2} f^3 \text{Re} [\tilde{s}_1^*(f) \tilde{s}_2(f)], \quad (2)$$

where ΔT is the time duration of data analyzed, and $\tilde{s}_i(f)$ is the (Fourier domain) strain measured by detector i . Equation (2) is normalized such that, for Advanced LIGO, the expectation value of $\hat{C}(f)$ is (Allen & Romano 1999)

$$\langle \hat{C}(f) \rangle = \gamma(f) \Omega(f). \quad (3)$$

In the weak signal limit, the variance of $\hat{C}(f)$ is given by $\langle \hat{C}(f) \hat{C}(f') \rangle = \delta(f - f') \sigma^2(f)$, with

$$\sigma^2(f) = \frac{1}{\Delta T} \left(\frac{10\pi^2}{3H_0^2} \right)^2 f^6 P_1(f) P_2(f), \quad (4)$$

where $P_i(f)$ is the one-sided noise power spectral density of detector i . Given a model $C_{\mathcal{H}}(f)$ for the energy-density spectrum of the background, the signal-to-noise ratio (S/N) of a stochastic measurement $\hat{C}(f)$ is given by the inner product $S/N^2 = \langle \hat{C} | C_{\mathcal{H}} \rangle$, where

$$(A|B) = 2 \int_0^\infty \frac{A^*(f) B(f)}{\sigma^2(f)} df. \quad (5)$$

The factor $\gamma(f)$ appearing in Equation (3), called the normalized overlap reduction function, encodes the dependence of the measured correlations on the detector baseline geometry—the




Figure 1. Overlap reduction function $\gamma(f)$ (blue) for the Advanced LIGO’s Hanford-Livingston detector baseline. Alternative baseline geometries have different overlap reduction functions as illustrated by the collection of gray curves, which show overlap reduction functions between hypothetical detectors randomly positioned on Earth’s surface.

detectors’ locations and relative orientations (Christensen 1992). Advanced LIGO’s normalized overlap reduction function is given by

$$\gamma(f) = \frac{5}{8\pi} \sum_A \int_{\text{Sky}} F_1^A(\hat{n}) F_2^A(\hat{n}) e^{2\pi i f \Delta \mathbf{x} \cdot \hat{n} / c} d\hat{n}. \quad (6)$$

Here, $F_i^A(\hat{n})$ is the antenna response of detector i to gravitational waves of polarization A , and $\Delta \mathbf{x}$ is the separation vector between detectors. The integral is performed over all sky directions \hat{n} and a sum is taken over both the “plus” and “cross” gravitational-wave polarizations. The leading factor of $5/8\pi$ normalizes the overlap reduction function such that identical, coincident, and co-aligned detectors would have $\gamma(f) = 1$.

Overlap reduction functions are strongly dependent upon baseline geometry—different pairs of gravitational-wave detectors generically have very different overlap reduction functions. To illustrate this, the overlap reduction function for the LIGO-Hanford and LIGO-Livingston baseline is shown in blue in Figure 1. The collection of gray curves, meanwhile, illustrates alternative overlap reduction functions for hypothetical pairs of detectors placed randomly on the surface of the Earth.

The strong dependence of $\gamma(f)$ on baseline geometry raises an interesting possibility. Given cross-correlation measurements $\hat{C}(f)$ between two detectors, we can use the measurements themselves to infer the baseline’s geometry. In the electromagnetic domain, a very similar technique has long been used in the field of geodesy: the experimental study of Earth’s geometry. While most commonly used to study the radio sky, very-long baseline interferometry can instead be used to precisely localize radio telescopes on the Earth, allowing for measurements of tectonic motion to better than $\sim 0.1 \text{ mm yr}^{-1}$ (Sovers et al. 1998; Schuh & Behrend 2012). Similarly, here we will use the *gravitational-wave sky* to determine our detectors’ relative positions and orientations.

As an initial demonstration, the left panel of Figure 2 illustrates a simulated observation of the stochastic gravitational-wave background with design-sensitivity Advanced LIGO. We assume a stochastic energy-density spectrum $\Omega(f) = 3.3 \times 10^{-9} (f/25 \text{ Hz})^{2/3}$, chosen to yield $S/N = 10$ after three years




Figure 2. Left panel: simulated Advanced LIGO cross-correlation measurements (blue) following a three-year observation of an isotropic stochastic gravitational-wave background. The injected background has energy-density $\Omega(f) = 3.33 \times 10^{-9} (f/25 \text{ Hz})^{2/3}$, corresponding to an expected S/N of 10 after three years of observation. The dashed curve shows the expected cross-correlation in the absence of measurement noise, and the gray band indicates $\pm 1\sigma$ uncertainties. Right panel: posterior on the distance between the LIGO Hanford and Livingston detectors, obtained using the simulated cross-correlation measurements shown on the left. The dashed line indicates the distance prior used and the vertical black line marks the true Hanford-Livingston separation. Using the gravitational-wave sky, we self-consistently recover a posterior compatible with the true distance between detectors. Details regarding parameter estimation are explained in Section 3 below.

of observation. The dashed curve indicates the mean correlation spectrum $\langle \hat{C}(f) \rangle$ corresponding to this injection, while the solid trace shows a simulated cross-correlation spectrum $\hat{C}(f)$ after three years of observation. By fitting to $\hat{C}(f)$ (as will be described below in Section 3), we can attempt to estimate the geometry of the LIGO Hanford-Livingston baseline. The resulting posterior on the separation between the LIGO Hanford and Livingston detectors is shown in the right panel of Figure 2. This posterior is consistent with the true separation between detectors (≈ 3000 km).

3. Model Selection

Of course, the physical separations between current gravitational-wave detectors are already known to far better accuracy than can be obtained through gravitational-wave geodesy. Nevertheless, the ability to measure baseline geometry with the gravitational-wave sky suggests a powerful consistency test for any possible detection of the gravitational-wave background.

In the presence of an isotropic, astrophysical stochastic background, the measured cross-correlation spectrum $\hat{C}(f)$ *must* exhibit amplitude modulations and zero-crossings consistent with the baseline’s overlap reduction function. Thus, when using the data $\hat{C}(f)$ to infer the baseline’s geometry, we must obtain results that are consistent with the known separations and orientations of the detectors. In contrast, spurious sources of terrestrial correlation are *not* bound to trace the overlap reduction function. Hence, there is no a priori reason that a correlated terrestrial signal should prefer the true baseline geometry over any other possible detector configuration.

We can more rigorously define this consistency check within the framework of Bayesian hypothesis testing. Given a measured cross-correlation spectrum $\hat{C}(f)$, we will ask which of the following hypotheses better describes the data.

1. Hypothesis \mathcal{H}_γ : the measured cross-correlation is consistent with the true baseline geometry (and hence the baseline’s true overlap reduction function).
2. Hypothesis $\mathcal{H}_{\text{Free}}$: the cross-correlation spectrum is consistent with a model in which we *do not* impose the

baseline’s known geometry, instead (unphysically) treating the detectors’ positions and orientations as free variables to be determined by the data.

An isotropic, astrophysical stochastic signal will be consistent with both \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ (assuming that the true baseline geometry is among the possible configurations supported in $\mathcal{H}_{\text{Free}}$). As the simpler hypothesis, however, \mathcal{H}_γ will be favored by the Bayesian “Occam’s factor” that penalizes the more complex model. So a true isotropic astrophysical stochastic background will favor \mathcal{H}_γ . A generic terrestrial signal, on the other hand, is unlikely to follow the baseline’s true overlap reduction function, and so will be better fit by the additional degrees of freedom allowed in $\mathcal{H}_{\text{Free}}$. Terrestrial sources of correlation are therefore likely to favor $\mathcal{H}_{\text{Free}}$.

This procedure is similar to the “sky scramble” technique used in pulsar timing searches for very low-frequency gravitational waves (Cornish & Sampson 2016; Taylor et al. 2017). In pulsar timing experiments, the analog to the overlap reduction function is the Hellings and Downs curve, which quantifies the expected correlations between pulsars as a function of their angular separation on the sky (Hellings & Downs 1983). By artificially shifting pulsar positions on the sky, one can seek to disrupt this spatial correlation and produce null data devoid of gravitational-wave signal but that retains other (possibly correlated) noise features.

Given a tentative detection of the stochastic background, we can compute a Bayes factor \mathcal{B} between hypotheses \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ to determine which is favored by the data. Due to the large number of time segments analyzed in stochastic searches, cross-correlation measurements are well described by Gaussian statistics. We therefore assume Gaussian likelihoods, such that the probability of measuring $\hat{C}(f)$ given a model spectrum $C_{\mathcal{H}}(\Theta; f)$ with parameters Θ is

$$p(\{\hat{C}\}|\Theta, \mathcal{H}) \propto \exp\left[-\frac{1}{2}(\hat{C} - C_{\mathcal{H}}(\Theta))|\hat{C} - C_{\mathcal{H}}(\Theta)\right], \quad (7)$$

using the inner product defined in Equation (5).

For both hypotheses, we adopt a power-law form for the background’s energy-density spectrum, defined by a reference




Figure 3. Parametrized geometry of an arbitrary detector baseline on the Earth’s surface. We initially choose coordinates such that the detectors lie in the x - z plane, with one detector at the pole. The remaining degrees of freedom are the polar angle θ between detectors, and the rotation angles ϕ_1 and ϕ_2 specifying the orientation of each detector.

amplitude Ω_0 and a spectral index α :

$$\Omega(f) = \Omega_0 \left(\frac{f}{25 \text{ Hz}} \right)^\alpha. \quad (8)$$

Our model for the cross-correlation spectrum under \mathcal{H}_γ is therefore

$$C_\gamma(\Omega_0, \alpha; f) = \gamma_{\text{True}}(f) \Omega_0 (f/25 \text{ Hz})^\alpha, \quad (9)$$

where $\gamma_{\text{True}}(f)$ is the true overlap reduction function for the given baseline.

For $\mathcal{H}_{\text{Free}}$, we additionally need a parametrized model for possible baseline geometries. We use the scheme illustrated in Figure 3. Given two detectors on the surface of the Earth (which we approximate as a sphere of radius $R_\oplus = 6.4 \times 10^6$ m), one can choose coordinates such that the first detector lies at the pole and the second along the meridian (in the x - z plane). We then have three remaining degrees of freedom: the polar angle θ between detectors, and the angles ϕ_1 and ϕ_2 specifying the rotation of each detector about its local zenith. Specifically, $\phi_{1/2}$ are the angles between the detectors’ \hat{v} arms (see Figure 3) and the y -axis. For convenience, below we will work in terms of the distance $\Delta x = 2R_\oplus \sin \theta/2$ between detectors, rather than the polar angle. All together, the model cross-correlation spectrum under hypothesis $\mathcal{H}_{\text{Free}}$ is

$$C_{\text{Free}}(\Omega_0, \alpha, \Delta x, \phi_1, \phi_2; f) = \gamma(\Delta x, \phi_1, \phi_2; f) \Omega_0 (f/25 \text{ Hz})^\alpha. \quad (10)$$

We choose a log-uniform prior on Ω_0 between $(10^{-12}, 10^{-6})$, extending well above and well below Advanced LIGO’s sensitivity, and uniform priors on ϕ_1 and ϕ_2 on $(0, 2\pi)$. Similarly, we use a uniform prior on $\cos \theta$ between $(-1, 1)$, corresponding to a prior $p(\Delta x) \propto \Delta x$ on the distance between detectors. We adopt a triangular prior on the background’s spectral index: $p(\alpha) \propto 1 - |\alpha|/\alpha_{\text{Max}}$, with $\alpha_{\text{Max}} = 6$. This prior penalizes very steeply sloped backgrounds, while still accommodating backgrounds that are much steeper than those predicted from known sources.

4. Demonstration

To explore our ability to differentiate terrestrial correlation from an astrophysical background, we will simulate Advanced LIGO measurements of three different sources of correlation: an isotropic stochastic background, a correlated frequency comb, and magnetic Schumann resonances. These latter two sources are terrestrial, and hence should disfavor \mathcal{H}_γ over $\mathcal{H}_{\text{Free}}$. We note that there exist dedicated strategies for identifying and mitigating combs and Schumann resonances (Thrane et al. 2014; Covas et al. 2018). Here, we use combs and Schumann resonances simply as proxies for any as-of-yet *unknown* sources of terrestrial correlation that could contaminate stochastic search efforts.

Below, we describe the model cross-correlation spectra adopted for each test case:

1. *Isotropic stochastic gravitational-wave background:* We assume that the stochastic gravitational-wave background is well described by a power law with spectral index $\alpha = 2/3$, as predicted for compact binary mergers. The corresponding expected cross-correlation spectrum is

$$\langle C(f) \rangle_{\text{Stoch}} = \gamma_{\text{LIGO}}(f) \Omega_0 \left(\frac{f}{25 \text{ Hz}} \right)^{2/3}, \quad (11)$$

where $\gamma_{\text{LIGO}}(f)$ is the overlap reduction function for the Hanford-Livingston baseline (shown in Figure 1).

2. *Frequency comb:* We consider a correlated comb of uniformly spaced lines, separated in frequency by Δf and with heights set by C_0 :

$$\langle C(f) \rangle_{\text{Comb}} = C_0 \Delta f \sum_{n=0}^{\infty} \delta(f - n\Delta f). \quad (12)$$

Note that the leading factor of Δf in Equation (12) ensures that C_0 is dimensionless. In the examples below, we use a comb spacing of $\Delta f = 2$ Hz.

3. *Magnetic Schumann resonances:* Given an environmental magnetic field $\tilde{m}(f)$, the strain induced in a gravitational-wave detector is $\tilde{s}(f) = T(f)\tilde{m}(f)$, where $T(f)$ is a transfer function with units (strain/Tesla). If there exists a correlated magnetic power spectrum $M(f) = \langle \tilde{m}_1^*(f)\tilde{m}_2(f) \rangle$ between the sites of two gravitational-wave detectors, then from Equation (2) the resulting strain correlation will be of the form $\hat{C}(f) \propto f^3 |T(f)|^2 \text{Re } M(f)$. We take $M(f)$ to be the median Schumann auto-power spectrum measured at the Hylaty station in Poland, as reported by Coughlin et al. (2018). This may not exactly match the magnetic cross-power spectrum between LIGO-Hanford and LIGO-Livingston. Most notably, we take $\text{Re } M(f)$ to be everywhere positive, as the (potentially frequency-dependent) sign of the Schumann cross-power between the LIGO detectors is not well known. Nevertheless, this model captures the qualitative features expected of a Schumann signal. The magnetic transfer functions for the LIGO detectors are expected to be power laws, but their spectral indices are also not well known; we somewhat arbitrarily choose $T(f) \propto f^{-2}$. Our Schumann signal model is therefore

$$\langle C(f) \rangle_{\text{Schumann}} = S_0 \left(\frac{f}{25 \text{ Hz}} \right)^{-1} \frac{\text{Re } M(f)}{\text{Re } M(25 \text{ Hz})}, \quad (13)$$

normalized so that S_0 is the cross-correlation measured at the reference frequency 25 Hz.

The mean cross-correlation spectra for the astrophysical, Schumann, and comb models are shown in Figure 4. For each




Figure 4. Mean cross-correlation spectra used to simulate stochastic search measurements with the Advanced LIGO Hanford and Livingston detectors. We consider an isotropic astrophysical stochastic background, with energy density $\Omega(f) \propto f^{2/3}$ (blue; Equation (11)). We additionally consider two sources of terrestrial, non-astrophysical correlation: a signal due to magnetic Schumann resonances (red; Equation (13)) and a correlated frequency comb with $\Delta f = 2$ Hz spacing (green; Equation (12)). The amplitudes of the spectra have been scaled such that each is expected to be detected with $S/N = 10$ after three years of observation with design-sensitivity Advanced LIGO. For comparison, the gray band illustrates the $\pm 1\sigma$ uncertainties of a cross-correlation search after three years of integration.

source of correlation, we simulate Advanced LIGO measurements of 300 injected signals, with expected S/N s ranging from 0.1 to 100. To produce each realization, we scale the amplitude parameters (Ω_0 , C_0 , and S_0) to obtain the desired S/N and add random Gaussian measurement noise $\delta C(f)$ with variance given by Equation (4). For each simulated measurement, we compute a Bayes factor \mathcal{B} between \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ to determine whether the data physically favors the correct detector geometry, or unphysically favors some alternate geometry. We compute Bayesian evidences using MultiNest (Feroz & Hobson 2008; Feroz et al. 2009), an implementation of the nested sampling algorithm (Skilling 2004, 2006). We make use of PyMultiNest, which provides a Python interface to MultiNest (Buchner et al. 2014).

The resulting Bayes factors are plotted in Figure 5 as a function of injected signal amplitude. As physically distinct parameters, the power-law, Schumann, and comb amplitudes should not be directly compared to one another. Instead, we show the injections' expected S/N s (which *can* be directly compared) on the upper horizontal axes. To compute these S/N s, we assume recovery with a power-law model of slope $\alpha = 2/3$. Thus the S/N s of the power-law injections are optimal. While S/N s for the comb and Schumann injections are *not* optimal (as the recovery model and injections are not identical), they do represent the S/N s at which such signals would contaminate searches for the stochastic background.

At $S/N \lesssim 1$ the log-Bayes factors for all three sources of correlation cluster near $\ln \mathcal{B} \sim 0$. For an astrophysical signal above $S/N \sim 1$, $\ln \mathcal{B}$ becomes positive, growing approximately linearly with $\log \Omega_0$. In contrast, $\ln \mathcal{B}$ falls exponentially to large negative values as we increase the amplitude of Schumann and comb injections. In Appendix A, we illustrate






Figure 5. Log-Bayes factors between the physical and unphysical hypotheses \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ as a function of injection strength for isotropic astrophysical backgrounds, Schumann resonances, and correlated combs (Equations (11)–(13)). To enable a direct comparison between injection types, the upper horizontal axes show the S/N s of these injections. $\ln \mathcal{B}$ increases linearly with the strength of an astrophysical injection, indicating consistency with the correct (known) detector geometry. Meanwhile, $\ln \mathcal{B}$ decreases exponentially for the terrestrial sources of correlation, *disfavoring* the correct geometry. In these cases, at least, $\ln \mathcal{B}$ therefore successfully discriminates between astrophysical and terrestrial sources of measured cross-correlation.

how the Laplace approximation can be used to derive these approximate scaling relations.

It is instructive to look at parameter estimation results for specific astrophysical, comb, and Schumann injections. In Appendix B, we show posteriors on the parameters of $\mathcal{H}_{\text{Free}}$ obtained using simulated observations of the stochastic background, a frequency comb, and a Schumann signal. As suggested by Figure 5, an observation of an isotropic stochastic background yields posteriors consistent with Advanced LIGO’s correct geometry. The comb and Schumann observations, on the other hand, produce unphysical posteriors on the positions and orientations of the Advanced LIGO detectors.

Figure 5 demonstrates that gravitational-wave geodesy can be used to successfully reject cross-correlation spectra that are inconsistent with Advanced LIGO’s overlap reduction function. However, there still remains the possibility of false positives: non-astrophysical correlation spectra that, purely by chance, yield posteriors consistent with Advanced LIGO’s geometry. To carefully calculate the probability of a false positive at a particular \mathcal{B} , one could analyze a set of random cross-correlation spectra (e.g., drawn from the space of spectra supported by $\mathcal{H}_{\text{Free}}$) and construct a null distribution of the resulting Bayes factors. Alternatively, we can quickly estimate the probability of false positives at a given $\ln \mathcal{B}$ using Figure 5(a). Given equal prior odds for \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$, the Bayes factors in Figure 5(a) may be directly interpreted as odds ratios. A Bayes factor of $\ln \mathcal{B} = 4$ (corresponding to $S/N \approx 10$), for example, indicates $e^4 : 1$ odds that the given data is drawn from \mathcal{H}_γ versus $\mathcal{H}_{\text{Free}}$. If taken at face value, this implies that we would need to simulate $e^4 + 1 \approx 56$ random correlation spectra with $S/N = 10$ before finding one that yields $\ln \mathcal{B} \gtrsim 4$ by chance. In this way, our formalism not only offers a means of rejecting non-astrophysical correlations, but can bolster the statistical significance of a real stochastic signal.

5. Discussion and Conclusion

As searches for the stochastic gravitational-wave background grow increasingly sensitive, we may be nearing the first detection of the background. This prospect, though, comes with significant risk, namely the high cost of a false positive detection. To minimize this risk, it will be important to develop methods to validate tentative detections of the gravitational-wave background. Specifically, when assessing any apparent detection, it will be necessary to argue not just that an observed correlation is statistically significant, but that it is *astrophysical*—that it is due to gravitational waves and not other terrestrial process. While well-developed methods exist to quantify the statistical significance of measured correlations, until now no generic method has existed to gauge whether or not a statistically significant cross correlation is indeed astrophysical.

In this Letter, we explored how gravitational-wave geodesy—the use of the stochastic gravitational-wave background itself to determine the positions and orientations of gravitational-wave detectors—can form the basis for a novel consistency check on an apparent detection of the background. If the measured correlation between detectors truly represents a gravitational-wave signal, then the reconstructed detector orientations and positions must be compatible with their true known values. Correlations due to any terrestrial source, on the other hand, have no reason to prefer the baseline’s true geometry over any other possible arrangement. By demanding that gravitational-wave geodesy yield results consistent with the true baseline geometry, we can discriminate between astrophysical and terrestrial sources of correlation. Used in this

fashion, gravitational-wave geodesy provides a second independent measure of detection significance besides a standard S/N .

Our analysis has relied on several important assumptions. First, we have assumed that our model energy-density spectrum (a power law) is a good descriptor of the true stochastic background. If our model were actually a poor fit to the true signal, the geodesy technique might conceivably reject a true astrophysical background. In Appendix C.1, we explore how our results are affected if we mistakenly assume an *incorrect* form for the background’s energy-density spectrum. We find that our analysis remains robust, correctly classifying astrophysical signals despite significant mismatches between our model spectrum and the true stochastic signal.

Second, our definition of the overlap reduction function (Equation (6)) makes several further implicit assumptions about the nature of the stochastic background—that it is isotropic, unpolarized, and composed purely of the tensor “plus” and “cross” modes predicted by general relativity. These assumptions may not be correct. Structure in the local universe may yield anisotropies in the stochastic background (Cusin et al. 2018b; Jenkins et al. 2018), and parity violations in the early universe may give rise to polarization asymmetries (Crowder et al. 2013). Modified theories of gravity, meanwhile, generally predict the presence of additional non-standard gravitational-wave polarizations (Callister et al. 2017; Abbott et al. 2018b).

The failure of any of these assumptions would modify the true overlap reduction function, which might be naively misinterpreted as disagreement with the detectors’ true geometry. The current geodesy test, then, should only be applied to searches for signals consistent with our assumptions (Abbott et al. 2017e). Nevertheless, as discussed in Appendix C.2, deviations from these assumptions are most likely small and therefore should minimally impact our analysis. Better yet, the geodesy test can be straightforwardly extended to handle the case of more complex stochastic backgrounds. As an example, in Appendix C.2 we outline how our analysis can be modified to accommodate an anisotropic stochastic signal. It remains a valuable future exercise to quantitatively explore exactly how much anisotropy, parity violation, and/or non-standard polarizations can be accommodated by our isotropic geodesy test.

We would like to thank Sharan Banagiri, Jan Harms, Andrew Matas, Joe Romano, Colm Talbot, Steve Taylor, Eric Thrane, Alan Weinstein, and members of the LIGO/Virgo Stochastic Data Analysis Group for useful comments and conversation. We also thank our anonymous referees, whose feedback has greatly enhanced the quality and content of the text. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. M.C. was supported by the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. This paper carries LIGO Document Number LIGO-P1800226.

Appendix A Bayes Factor Scaling

The behavior of the Bayes factors in Figure 5 can be understood using the Laplace approximation. The Laplace approximation involves the following two assumptions.

1. Our prior $p(\Theta|\mathcal{H})$ on the parameters of hypothesis \mathcal{H} is flat over a range $\Delta\Theta$, so that $p(\Theta|\mathcal{H}) = 1/\Delta\Theta$.
2. The likelihood $p(\hat{C}|\Theta, \mathcal{H})$ is strongly peaked about maximum-likelihood parameter values $\bar{\Theta}$ and a peak value $\bar{\mathcal{L}}$. The width of the peak is $\delta\Theta$.

Under these assumptions, a Bayesian evidence may be approximated as

$$\begin{aligned} p(\hat{C}|\mathcal{H}) &= \int p(\hat{C}|\Theta, \mathcal{H})p(\Theta|\mathcal{H})d\Theta \\ &\approx \frac{\delta\Theta}{\Delta\Theta}\bar{\mathcal{L}}. \end{aligned} \quad (14)$$

The leading term $\delta\Theta/\Delta\Theta$ can be interpreted as the volume of the available parameter space that is compatible with the measured data. Given two hypotheses \mathcal{H}_A and \mathcal{H}_B , the Bayes factor between them becomes

$$\begin{aligned} \mathcal{B}_B^A &= \frac{p(\hat{C}|\mathcal{H}_A)}{p(\hat{C}|\mathcal{H}_B)} \\ &\approx \frac{\delta\Theta_A/\Delta\Theta_A \bar{\mathcal{L}}_A}{\delta\Theta_B/\Delta\Theta_B \bar{\mathcal{L}}_B}. \end{aligned} \quad (15)$$

The ratio $\bar{\mathcal{L}}_A/\bar{\mathcal{L}}_B$ is the standard maximum likelihood ratio between \mathcal{H}_A and \mathcal{H}_B . The preceding term, known as the ‘‘Occam’s factor,’’ penalizes the more complex hypothesis with the larger available parameter space. Using the Laplace approximation, our Bayes factor between hypotheses \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ may be written

$$\begin{aligned} \mathcal{B} &= \frac{p(\hat{C}|\mathcal{H}_\gamma)}{p(\hat{C}|\mathcal{H}_{\text{Free}})} \\ &\approx \left[\frac{\delta\Omega_0}{\Delta\Omega_0} \frac{\delta\alpha}{\Delta\alpha} \right]_\gamma \left[\frac{\delta\Omega_0}{\Delta\Omega_0} \frac{\delta\alpha}{\Delta\alpha} \frac{\delta\theta}{\Delta\theta} \frac{\delta\phi_1}{\Delta\phi_1} \frac{\delta\phi_2}{\Delta\phi_2} \right]_{\text{Free}}^{-1} \\ &\quad \times \frac{\exp\left[-\frac{1}{2}(\hat{C} - \bar{\mathcal{C}}_\gamma|\hat{C} - \bar{\mathcal{C}}_\gamma)\right]}{\exp\left[-\frac{1}{2}(\hat{C} - \bar{\mathcal{C}}_{\text{Free}}|\hat{C} - \bar{\mathcal{C}}_{\text{Free}})\right]}, \end{aligned} \quad (16)$$

where $\bar{\mathcal{C}}_\gamma$, for instance, is the maximum-likelihood fit to the data under the \mathcal{H}_γ hypothesis.

First, consider the case of an isotropic astrophysical background of amplitude Ω_0 . In this case, both hypotheses \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ can successfully fit the resulting cross-correlation spectrum. Then $\hat{C} - \bar{\mathcal{C}}_\gamma \approx \hat{C} - \bar{\mathcal{C}}_{\text{Free}} \approx 0$ and the likelihood ratio in Equation (16) is approximately one. Because both models can fit the data, posteriors on each parameter (of each hypothesis) are peaked, with fractional widths (e.g., $\delta\theta/\Delta\theta$) that scale as $S/N^{-1} \propto \Omega_0^{-1}$. Then, in the presence of an astrophysical background, we expect Equation (16) to scale as $\mathcal{B} \propto \Omega_0^3$, or

$$\ln \mathcal{B} \sim 3 \log \Omega_0, \quad (17)$$

up to additive constants.

Next, consider a correlated signal of terrestrial origin, characterized by some amplitude C_0 . We will assume that \mathcal{H}_γ is unable to accommodate the measured correlations, but that

$\mathcal{H}_{\text{Free}}$, with a greater number of free parameters, *can* successfully fit the data to some extent. Then $\hat{C} - \bar{\mathcal{C}}_{\text{Free}} \approx 0$ but $\hat{C} - \bar{\mathcal{C}}_\gamma \neq 0$. So the likelihood term in Equation (16) is not constant, but will depend exponentially on C_0 . Ignoring the leading Occam’s factors (which can scale at most as a power law in C_0), our Bayes factor becomes

$$\begin{aligned} \mathcal{B} &\propto \exp\left[-\frac{1}{2}(\hat{C} - \bar{\mathcal{C}}_\gamma|\hat{C} - \bar{\mathcal{C}}_\gamma)\right] \\ &\propto \exp\left[-\frac{1}{2}(\hat{C}|\hat{C}) + (\hat{C}|\bar{\mathcal{C}}_\gamma) - \frac{1}{2}(\bar{\mathcal{C}}_\gamma|\bar{\mathcal{C}}_\gamma)\right], \end{aligned} \quad (18)$$

giving

$$\ln \mathcal{B} \sim -\frac{1}{2}(\hat{C}|\hat{C}) + (\hat{C}|\bar{\mathcal{C}}_\gamma) - \frac{1}{2}(\bar{\mathcal{C}}_\gamma|\bar{\mathcal{C}}_\gamma). \quad (19)$$

The maximum likelihood value of Ω_0 (the amplitude of our model spectrum $C_\gamma(f)$) is given by (Callister et al. 2016)

$$\bar{\Omega}_0 = \frac{(f^{2/3}|\hat{C})}{(f^{2/3}|f^{2/3})}. \quad (20)$$

Although this does scale proportionally with C_0 , in this scenario our measured correlation $\hat{C}(f)$ is assumed to have a very different shape from an astrophysical power law. The inner product $(f^{2/3}|\hat{C})$ may therefore be small, in which case the cross term $(\hat{C}|\bar{\mathcal{C}}_\gamma)$ in Equation (19) may be neglected. As a result, $\ln \mathcal{B} \propto -C_0^2$, or

$$\ln \mathcal{B} \propto -10^{2 \log C_0}. \quad (21)$$

Appendix B Parameter Estimation Results

In this section we show example parameter estimation results obtained when analyzing simulated observations of a stochastic gravitational-wave background, a correlated frequency comb, and Schumann resonances, each with $S/N = 10$. For each injection, we perform parameter estimation under the $\mathcal{H}_{\text{Free}}$ hypothesis, allowing the detector positions and orientations to (unphysically) vary to best match the observed cross-correlation spectrum. We implement parameter estimation using MultiNest and PyMultiNest.

Figure 6 shows the three injections as well as the posteriors obtained on each cross-correlation spectrum. With the five free parameters afforded by $\mathcal{H}_{\text{Free}}$, we succeed in reasonably fitting each of the three spectra. Note that, although we appear to poorly recover the correlated comb injection, the posterior on $C(f)$ closely matches the *frequency-averaged* correlation.

Although the gravitational-wave background, comb, and Schumann injections are all reasonably well fit under $\mathcal{H}_{\text{Free}}$, they yield very different posteriors on Advanced LIGO’s baseline length Δx and detector orientation angles ϕ_1 and ϕ_2 . Figure 7 shows the parameter posteriors given by the simulated gravitational-wave background. The diagonal subplots show marginalized one-dimensional posteriors on each parameter, while the central subplots show joint posteriors between each pair of parameters.




Figure 6. Reconstructed cross-correlation spectra using simulated Advanced LIGO observations of an isotropic gravitational-wave background (top panel), a correlated frequency comb (middle panel), and Schumann resonances (bottom panel). The blue, green, and red curves show the injected gravitational-wave, comb, and Schumann spectra, respectively, while the shaded bands indicate the $\pm 1\sigma$ uncertainty region on the simulated measurements. We perform parameter estimation on each injection using the $\mathcal{H}_{\text{Free}}$ hypothesis, fitting simultaneously for the spectral shape of a presumed stochastic background as well as the detectors’ separation and orientations. The collections of gray curves show the resulting posteriors on the injected cross-correlation spectrum. Posteriors on the model parameters themselves are shown in Figures 7–9.

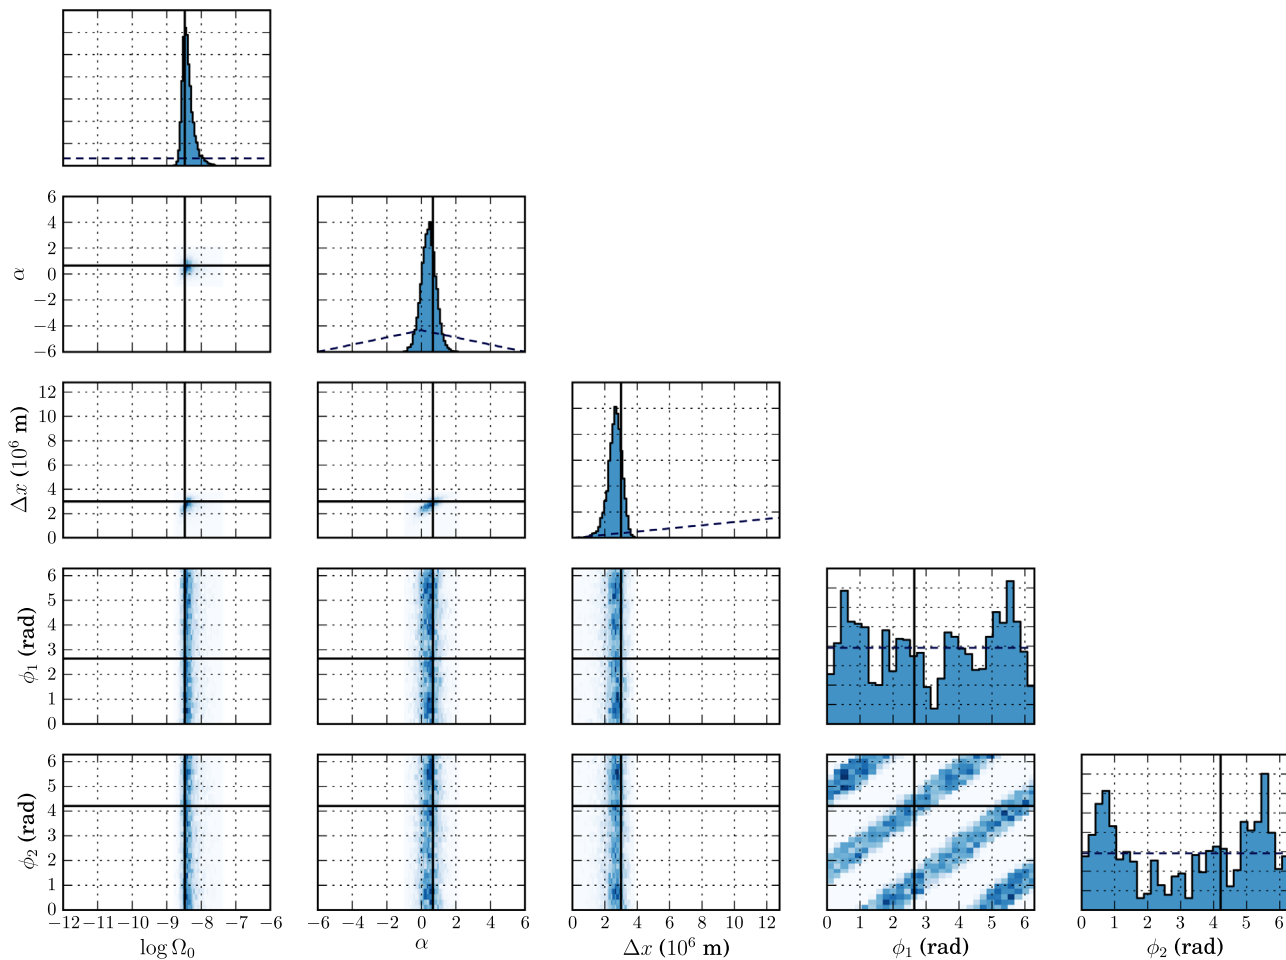


Figure 7. Example posterior on the stochastic background amplitude Ω_0 and spectral index α , as well as the separation Δx and rotation angles ϕ_1 and ϕ_2 of the Advanced LIGO detectors, given a simulated three-year observation of an isotropic astrophysical background. The injected signal has spectral index $\alpha = 2/3$ and amplitude $\Omega_0 = 3.33 \times 10^{-9}$, with an expected S/N of 10. Dashed lines in the one-dimensional marginalized posteriors show the prior adopted for each parameter, while solid black lines mark the injected background parameters and the true Advanced LIGO geometry. In addition to recovering the amplitude and spectral index of the injected stochastic signal, we obtain posteriors consistent with the true separation and rotation angles of the Advanced LIGO detectors.

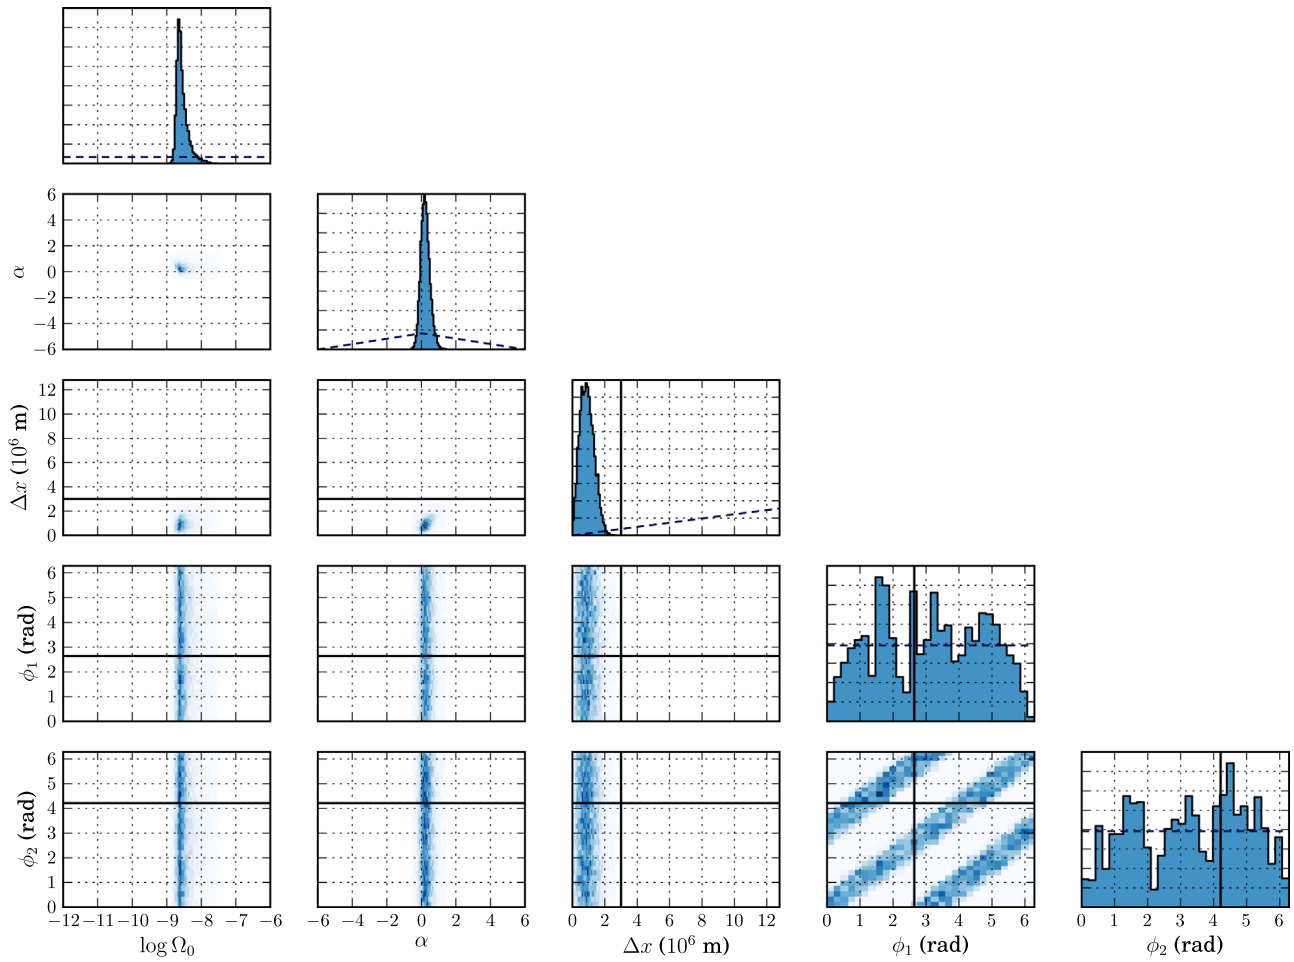


Figure 8. Same as in Figure 7 above, but for a simulated measurement of a correlated frequency comb with spacing $\Delta f = 2$ Hz and height $C_0 = 7.83 \times 10^{-9}$. The comb's amplitude is chosen so that it has $S/N = 10$ after three years of observation. The correlated comb is not well fit by the Advanced LIGO overlap reduction function, and so our recovered posteriors on Hanford and Livingston's separation and rotation angles are inconsistent with their known values (solid black lines).

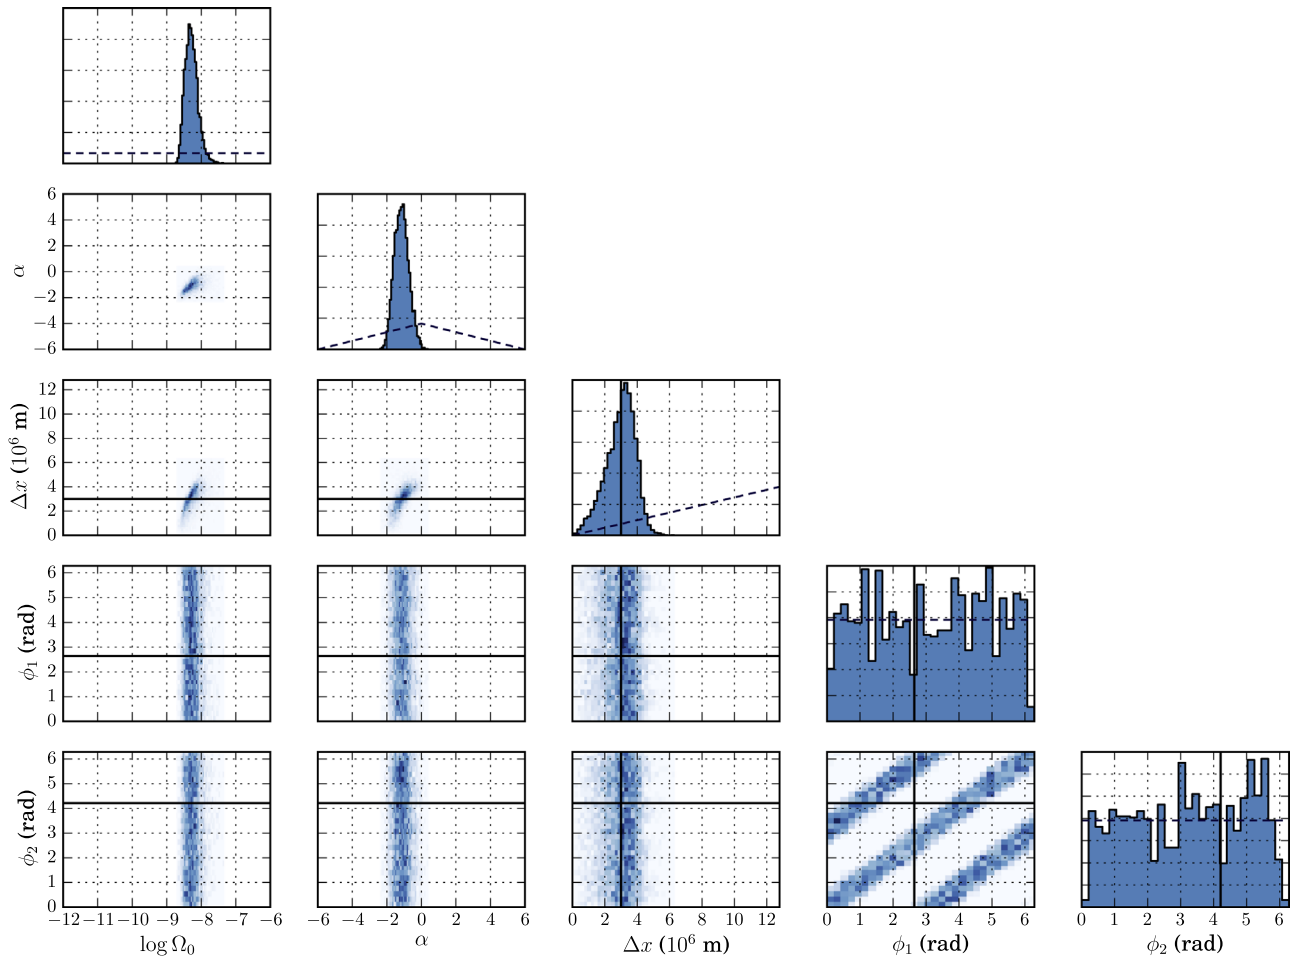


Figure 9. Same as in Figures 7 and 8 above, for a simulated observation of a correlated Schumann signal of height $S_0 = 2.33 \times 10^{-9}$, chosen to yield $S/N = 10$ after three years of observation. While the posterior does encompass the correct Hanford-Livingston separation, it is incompatible with the detectors’ true rotation angles.

The solid black lines indicate true parameter values and dashed curves show the priors placed on each parameter. We recover posteriors consistent with the amplitude and spectral index of the injected stochastic signal. More importantly, we also obtain well-behaved posteriors on Advanced LIGO’s geometry, with a distance posterior (the same as shown in Figure 2) consistent with the true separation between detectors. Interestingly, although neither ϕ_1 nor ϕ_2 are well constrained, their *difference* is well measured. This can be seen in the joint posterior between both angles, which strongly supports diagonal bands of constant $\phi_1 - \phi_2$, including the true rotation angles of Hanford and Livingston. We therefore have strong support for the correct detector geometry, yielding a log-Bayes factor $\ln \mathcal{B} = 3.6$ ($\mathcal{B} = 36.6$) in favor of \mathcal{H}_γ .

Figure 8, meanwhile, shows parameter estimation results obtained for the comb injection. As seen in Figure 6 above, we have enough freedom to fit the (average) cross-correlation spectrum, yielding reasonably peaked posteriors in Figure 8. However, the posteriors on detector separation and orientation are unphysical, excluding the known Hanford-Livingston geometry. We therefore obtain $\ln \mathcal{B} = -58.5$ ($\mathcal{B} = 3.9 \times 10^{-26}$). Similarly, Figure 9 gives parameter estimation results for the Schumann injection. Interestingly, the distance posterior for this injection is consistent with the true Hanford-Livingston separation. The rotation angle posteriors, though, again exclude the true detector orientations, yielding $\ln \mathcal{B} = -62.7$ ($\mathcal{B} = 5.9 \times 10^{-28}$).

Appendix C Complications

We demonstrated in Section 4 that gravitational-wave geodesy can be successfully used to discriminate between a true stochastic gravitational-wave background and non-astrophysical, terrestrial sources of correlation. Here, we highlight important assumptions that have been made in our analysis and discuss what to do should these assumptions not hold.

C.1. Non-power-law Energy-density Spectra

In the main text, we have assumed that our model energy-density spectrum (a power law) is a good description of the true stochastic background. This assumption was guaranteed by design, as our injected stochastic energy-density spectrum *was* a power law. While most gravitational-wave sources are predicted to yield power-law energy-density spectra in the Advanced LIGO and Virgo band, there do exist speculative sources like superradiant axion clouds (Brito et al. 2017a, 2017b) that may instead yield more complex spectra.

It is worthwhile to investigate how our method fares given more complex energy-density spectra. Specifically, we will consider observations of a broken power-law background with




Figure 10. Log-Bayes factors between hypotheses \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ obtained when analyzing simulated astrophysical broken power-law signals (a) and magnetic Schumann correlations (b), assuming a broken power law for our model energy-density spectrum. For each set of simulations, we assume three years of observation with design-sensitivity Advanced LIGO. When adopting this more complex model, our Bayes factors still scale as expected, with the astrophysical signal preferring \mathcal{H}_γ and the Schumann correlations preferring $\mathcal{H}_{\text{Free}}$.




Figure 11. Log-Bayes factors between \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ when deliberately analyzing astrophysical broken power-law signals with an *incorrect* power-law model. Each injected signal has a knee frequency of $f_0 = 30$ Hz and an amplitude Ω_0 scaled such that the signal has $S/N = 5$ after three years of observation with design-sensitivity Advanced LIGO. Despite the signal-model mismatch, we correctly classify the majority of the simulated signals, with no evidence of increased false-dismissals due to the mismatch.

energy density

$$\Omega(f) = \begin{cases} \Omega_0(f/f_0)^{\alpha_1} & f \leq f_0 \\ \Omega_0(f/f_0)^{\alpha_2} & f > f_0. \end{cases} \quad (22)$$

Correspondingly, we will adopt a broken power-law model for $\Omega(f)$ (with free parameters Ω_0 , f_0 , α_1 and α_2) in both hypotheses \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$. We simulate broken power-law signals with slopes $\alpha_1 = 1$, $\alpha_2 = -1$, a knee frequency $f_0 = 30$ Hz, and amplitudes Ω_0 ranging from 10^{-11} to 10^{-7} . The recovered Bayes factors between \mathcal{H}_γ and $\mathcal{H}_{\text{Free}}$ are shown in Figure 10(a). We see that, even given a more complex signal and model, our method remains effective.

With a more complex model, is it also true that we can still correctly reject terrestrial sources of correlation? To verify that the additional free parameters afforded by the broken power-law model do not lead to false acceptances, we again apply the

broken power-law model to Schumann correlations of various strengths. As shown in Figure 10(b), the resulting Bayes factors behave as expected, indicating inconsistency with Advanced LIGO’s correct geometry.

We have shown that geodesy is a successful discriminator between astrophysical and terrestrial correlations, even when using a model that is more complicated than a simple power-law spectrum. Crucially, though, we have still assumed the *correct* energy-density spectrum, using the same model (a broken power law) to both inject and recover simulated signals. The most troubling case is the possibility of an *incorrect* model—one that is a poor descriptor of the true stochastic background. In this case, would we risk rejecting a real stochastic background as a terrestrial signal?

To test this, we again simulate observations of a broken power-law background, but recover them using an ordinary power-law model, deliberately choosing an incorrect description of the simulated signal. Figure 11 illustrates the resulting Bayes factors for simulated signals with α_1 and α_2 ranging between -4 and 4 . For each injection we chose $f_0 = 30$ Hz, placing the broken power-law’s “knee” in the center of the stochastic sensitivity band, and scaled the amplitudes Ω_0 such that each observation has $S/N = 5$ when naively recovered with an ordinary power law. The vast majority of these simulations yield positive log-Bayes factors, correctly classifying these signals despite our poor choice of model. Note that the injections falling along the line $\alpha_1 = \alpha_2$ are power laws. If the signal-model mismatch significantly degraded our ability to classify stochastic signals, then Figure 11 would exhibit a color gradient as we move perpendicularly off the $\alpha_1 = \alpha_2$ line, away from power laws and toward increasingly sharp signal spectra. Instead, Figure 11 shows no such gradient, and our method remains robust even in the case of poorly fitting models.

We attribute this robustness to the fact that the isotropic energy-density spectrum and baseline geometry have very different effects on the expected cross-correlation spectrum $\langle C(f) \rangle = \gamma(f)\Omega(f)$. The energy density spectrum $\Omega(f)$ is everywhere positive, and so different energy-density spectra can change only the *amplitude* of $C(f)$, not its sign. The sign of $C(f)$ is set by the overlap reduction function, which alternates

between positive and negative values with zero-crossings fixed by the baseline geometry. Even if our model for $C(f)$ assumes an incorrect energy-density spectrum (as above), our \mathcal{H}_γ hypothesis nevertheless predicts the correct zero-crossings of the observed cross-correlation spectrum. This offers some robustness against false-dismissal of a true stochastic signal, even if our model energy-density spectrum is imperfect. At the same time, it prevents us from *over-fitting* spurious terrestrial correlations (whose sign is unrelated to the sign of $\gamma(f)$), mitigating the risk of false positives.

C.2. Anisotropy, Polarization, and Modified Theories of Gravity

We have additionally assumed that the stochastic gravitational-wave background is isotropic, unpolarized, and free of the non-standard “vector” and “scalar” gravitational-wave polarizations predicted by modified theories of gravity. These assumptions are unlikely to all be strictly true. The stochastic background may be polarized by a variety of early universe effects (Crowder et al. 2013), as well as the scattering of gravitational waves by massive objects during propagation (Cusin et al. 2018a). Meanwhile, the solar system’s motion with respect to the cosmic microwave background will likely impart a small apparent dipole moment to the stochastic gravitational-wave background. Additional anisotropies might arise from structure in the local universe (Cusin et al. 2018b; Jenkins et al. 2018), together with the fact that, over a finite integration time, we observe only a discrete set of gravitational-wave events (Meacher et al. 2014).

A stochastic background containing anisotropies, polarization asymmetries, or non-standard polarizations would yield correlations that are *not* consistent with the standard overlap reduction function, but that instead obey some different effective overlap reduction function. If we naively analyzed such a signal with the method presented in the main text, we would likely find a preference for the (unphysical) hypothesis $\mathcal{H}_{\text{Free}}$ over \mathcal{H}_γ and risk rejecting the signal as terrestrial.

In practice, deviations from our ideal stochastic background model are expected to be small, and so these complicating factors are unlikely to significantly affect our analysis. For example, the solar system moves with speed $v_\oplus \approx 370 \text{ km s}^{-1}$ with respect to the cosmic microwave background, and so the stochastic background’s apparent dipole moment is a factor of $v_\oplus/c \sim 10^{-3}$ weaker than the isotropic monopole moment. The intrinsic anisotropy and polarization of the astrophysical background are also predicted to be small. Considering multipole moments up to $l = 20$ (the approximate angular resolution limit of the LIGO Hanford-Livingston baseline; Thrane et al. 2009), the observed energy density is expected to vary by no more than $\sim 10\%$ with direction (Cusin et al. 2018b; Jenkins et al. 2018). Any net polarization arising from scattering is predicted to be further suppressed by many orders of magnitude in the frequency band of Advanced LIGO and Virgo (Cusin et al. 2018a).

If any of these complications were a significant concern, however, the formalism of Section 3 can be straightforwardly extended to accommodate these effects. As an example, here we demonstrate how to extend our formalism to the case of an anisotropic stochastic background.

When allowing for anisotropy, the observed energy-density of the stochastic background will generically have directional dependence on our viewing angle \hat{n} . It is generally assumed

that an anisotropic energy-density spectrum can be factored via $\Omega(\hat{n}, f) = H(f)\mathcal{P}(\hat{n})$, where $H(f)$ and $\mathcal{P}(\hat{n})$ encode the frequency and directional dependence of $\Omega(\hat{n}, f)$, respectively. We can further decompose $\mathcal{P}(\hat{n})$ into a sum of spherical harmonics $Y_{lm}(\hat{n})$, giving (Allen & Ottewill 1997; Thrane et al. 2009; Abbott et al. 2017f)

$$\Omega(\hat{n}, f) = H(f) \sum_{l,m} \mathcal{P}^{lm} Y_{lm}(\hat{n}) \quad (23)$$

for some set of coefficients \mathcal{P}^{lm} . We use the normalization convention $\int |Y_{lm}(\hat{n})|^2 d\hat{n} = 1$.

Over the course of a sidereal day, gravitational-wave detectors have varying sensitivities to different sky directions \hat{n} . In the presence of an anisotropic background, the expected cross-correlation between detectors is therefore time-dependent:

$$\langle C(f, t) \rangle = H(f) \sum_{l,m} \mathcal{P}^{lm} \gamma_{lm}(t, f), \quad (24)$$

where t is periodic over a sidereal day. This expression is similar in form to Equation (3), but with a sum over spherical harmonics and distinct (time-dependent) overlap reduction functions for each spherical harmonic (Allen & Ottewill 1997; Thrane et al. 2009)

$$\begin{aligned} \gamma_{lm}(t, f) &= \frac{5}{2\sqrt{4\pi}} \sum_A \int_{\text{Sky}} Y_{lm}(\hat{n}) F_1^A(\hat{n}, t) F_2^A(\hat{n}, t) \\ &\quad \times e^{2\pi i f \Delta \mathbf{x}(t) \cdot \hat{n} / c} d\hat{n}. \end{aligned} \quad (25)$$

In Equation (25), the detectors’ antenna patterns $F_i^A(\hat{n}, t)$ and separation vector $\Delta \mathbf{x}(t)$ are time-dependent, rotating with the Earth over the course of a sidereal day. The normalization of Equation (25) is chosen such that monopole overlap reduction function $\gamma_{00}(t, f)$ reduces to Equation (6) above. The time-dependence of Equation (25) can be conveniently factored out via (Allen & Ottewill 1997; Thrane et al. 2009)

$$\gamma_{lm}(t, f) = \gamma_{lm}(0, f) e^{2\pi i m(t/T)}, \quad (26)$$

where T is the length of one sidereal day.

If we incorrectly assumed an isotropic background and averaged our cross-correlation measurements over a sidereal day, we would measure cross correlation

$$\begin{aligned} \langle C(f) \rangle &= \frac{1}{T} \int_0^T \langle C(f, t) \rangle dt \\ &= H(f) \sum_{l,m} \mathcal{P}^{lm} \gamma_{lm}(0, f) \frac{1}{T} \int_0^T e^{2\pi i m(t/T)} dt \\ &= H(f) \sum_l \mathcal{P}^{l0} \gamma_{l0}(0, f), \end{aligned} \quad (27)$$

where the integral vanishes for all $m \neq 0$. Equation (27) does not trace the isotropic overlap reduction function, but instead follows a linear combination of the anisotropic $\gamma_{l0}(f)$ ’s. Thus, if the background were significantly anisotropic (with some \mathcal{P}^{l0} comparable in magnitude to the monopole amplitude \mathcal{P}^{00}), we would incorrectly conclude that the resulting correlated signal is incompatible with our detector geometry and dismiss it as terrestrial.

In analogy to Equation (9), one could define hypothesis \mathcal{H}_γ via the model

$$C_\gamma(\Theta, \mathcal{P}^{lm}; f) = H(\Theta; f) \sum_{l,m} \mathcal{P}^{lm} \gamma_{lm}^{\text{True}}(t, f), \quad (28)$$

where $\gamma_{lm}^{\text{True}}(t, f)$ is the baseline's known overlap reduction function for spherical harmonic (l, m) and Θ represents the variables parametrizing $H(f)$. Similarly, the unphysical hypothesis $\mathcal{H}_{\text{Free}}$ would become

$$\begin{aligned} C_{\text{Free}}(\Theta, \mathcal{P}^{lm}, \theta, \phi_1, \phi_2; f) \\ = H(\Theta; f) \sum_{l,m} \mathcal{P}^{lm} \gamma_{lm}(\theta, \phi_1, \phi_2; t, f). \end{aligned} \quad (29)$$

ORCID iDs

T. A. Callister  <https://orcid.org/0000-0001-9892-177X>
M. W. Coughlin  <https://orcid.org/0000-0002-8262-2924>
J. B. Kanner  <https://orcid.org/0000-0001-8115-0577>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *PhRvX*, **6**, 041015
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvL*, **116**, 131102
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *PhRvL*, **118**, 221101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, *ApJL*, **851**, L35
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017c, *PhRvL*, **119**, 141101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017d, *PhRvL*, **119**, 161101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017e, *PhRvL*, **118**, 121101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017f, *PhRvL*, **118**, 121102
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018a, *PhRvL*, **120**, 091101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018b, *PhRvL*, **120**, 201102
Allen, B., & Ottewill, A. C. 1997, *PhRvD*, **56**, 545
Allen, B., & Romano, J. D. 1999, *PhRvD*, **59**, 102001
Brito, R., Ghosh, S., Barausse, E., et al. 2017a, *PhRvD*, **96**, 064050
Brito, R., Ghosh, S., Barausse, E., et al. 2017b, *PhRvL*, **119**, 131101
Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
Callister, T., Biscoveanu, A. S., Christensen, N., et al. 2017, *PhRvX*, **7**, 041058
Callister, T., Sammut, L., Qiu, S., Mandel, I., & Thrane, E. 2016, *PhRvX*, **6**, 031018
Christensen, N. 1992, *PhRvD*, **46**, 5250
Cornish, N. J., & Sampson, L. 2016, *PhRvD*, **93**, 104047
Coughlin, M. W., Christensen, N. L., De Rosa, R., et al. 2016, *CQGra*, **33**, 224003
Coughlin, M. W., Cirone, A., Meyers, P., et al. 2018, *PhRvD*, **97**, 102007
Covas, P. B., Effler, A., Goetz, E., et al. 2018, *PhRvD*, **97**, 082002
Crowder, S., Namba, R., Mandic, V., Mukohyama, S., & Peloso, M. 2013, *PhLB*, **726**, 66
Cusin, G., Durrer, R., & Ferreira, P. G. 2018a, **4**, 1, arXiv:1807.10620
Cusin, G., Dvorkin, I., Pitrou, C., & Uzan, J.-P. 2018b, *PhRvL*, **120**, 231101
Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39
Jenkins, A. C., Sakellariadou, M., Regimbau, T., & Slezak, E. 2018, *PhRvD*, **98**, 063501
Meacher, D., Thrane, E., & Regimbau, T. 2014, *PhRvD*, **89**, 084063
Regimbau, T., & Mandic, V. 2008, *CQGra*, **25**, 184018
Romano, J. D., & Cornish, N. J. 2017, *LRR*, **20**, 2
Rosado, P. A. 2011, *PhRvD*, **84**, 084004
Schuh, H., & Behrend, D. 2012, *JGeo*, **61**, 68
Schumann, W. O. 1952a, *ZNatA*, **7**, 149
Schumann, W. O. 1952b, *ZNatA*, **7**, 250
Skilling, J. 2004, in AIP Conf. Proc. 735, 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer, R. Preuss, & U. von Toussaint (Melville, NY: AIP), 395
Skilling, J. 2006, *BayAn*, **1**, 833
Sovers, O. J., Fanselow, J. L., & Jacobs, C. S. 1998, *RvMP*, **70**, 1393
Taylor, S. R., Lentati, L., Babak, S., et al. 2017, *PhRvD*, **95**, 042002
Thrane, E., Ballmer, S., Romano, J. D., et al. 2009, *PhRvD*, **80**, 122002
Thrane, E., Christensen, N., & Schofield, R. M. S. 2013, *PhRvD*, **87**, 123009
Thrane, E., Christensen, N., Schofield, R. M. S., & Effler, A. 2014, *PhRvD*, **90**, 023013
Wu, C., Mandic, V., & Regimbau, T. 2012, *PhRvD*, **85**, 104024
Zhu, X.-J., Howell, E., Regimbau, T., Blair, D., & Zhu, Z.-H. 2011, *ApJ*, **739**, 86
Zhu, X.-J., Howell, E. J., Blair, D. G., & Zhu, Z.-H. 2013, *MNRAS*, **431**, 882