

A coincidence null test for Poisson-distributed eventsReed Essick^{*}*Kavli Institute for Cosmological Physics, The University of Chicago,
5640 South Ellis Avenue, Chicago, Illinois 60637, USA*Geoffrey Mo[⊗] and Erik Katsavounidis*LIGO Laboratory, Massachusetts Institute of Technology,
77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA
and Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts
Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA*

(Received 27 November 2020; accepted 7 January 2021; published 5 February 2021)

When transient events are observed with multiple sensors, it is often necessary to establish the significance of coincident events. We derive a universal null test for an arbitrary number of sensors motivated by the archetypal detection problem for independent Poisson-distributed events in gravitational-wave detectors such as LIGO and Virgo. In these detectors, transient events may be witnessed by myriad channels that record interferometric signals and the surrounding physical environment. We apply our null test to a broad set of simulated gravitational-wave events as well as to a real gravitational-wave detection to determine which auxiliary channels do and do not witness real gravitational waves, and therefore which are safe to use when constructing vetoes. We also describe how our approach can be used to study detector artifacts and their origin, as well as to quantify the statistical independence of candidate gravitational-wave signals from noise artifacts observed in auxiliary channels.

DOI: [10.1103/PhysRevD.103.042003](https://doi.org/10.1103/PhysRevD.103.042003)**I. INTRODUCTION**

Establishing correlations between separate sets of possibly related transient events, or demonstrating the vanishingly small probability of observing a set of coincidences in uncorrelated events, is of general interest. This situation arises naturally in a variety of contexts in astrophysics, especially in the correlation of events across multiple detectors. Examples include the multiple elements of cosmic-ray arrays (e.g., [1,2]); neutrino detections from IceCube [3], Antares [4], and the Supernova Early-Warning System [5]; correlated magnetic noise from Schumann resonances in distant detectors [6–8]; as well as coincidences between gravitational waves (GWs), electromagnetic waves, cosmic rays, and neutrinos in multimessenger astrophysics (see, e.g., [9]). Additionally, such questions are often asked while characterizing noise artifacts in detectors and searching for their sources (e.g., [10–12]). This problem also appears outside the physical sciences, for example, in understanding internet server loads and distributed denial-of-service attacks (e.g., [13]), aspects of quantitative finance (such as identifying causation in price impact [14] or factors associated with large price movements in risk assessment [15]), and many others.

We provide a general null hypothesis test to determine whether coincidences are random in nature.

Throughout this work, we focus on the identification of noise transients (either terrestrial or instrumental) in ground-based GW interferometers such as the advanced LIGO [16] and Virgo [17] detectors. GW interferometry provides a rich assortment of example analyses. We refer to any short-duration ($\lesssim \mathcal{O}(1 \text{ sec})$) transient within the instruments' sensitive frequency band (20 Hz to a few kHz) as an *event*. When individual events can be accurately modeled and the transfer function between the event's source and the observed data is well understood, more complicated consistency checks can provide powerful tests of causal connections (e.g., [18–20]). This has been used to fantastic effect by combining accurate waveform predictions for GW signals expected from the coalescence of compact binary systems (e.g., [21]) with precise knowledge of the detectors' interferometric response [22,23] in order to detect and infer the properties of some of the most extreme events in the Universe [24,25].

However, this is often not the case. Whether direct measurements of the couplings between different sensors are too difficult to obtain, models of the transients are difficult to construct, or whether there are simply too many possible couplings between large numbers of sensors to accurately catalog, we are often faced with the general

*reed.essick@gmail.com

problem of determining the significance of coincidences between sets of discretized events without knowledge of their expected shapes.

Within GW astronomy, a variety of event trigger generators (ETGs; typically distinguished by the wavelet transform employed; see [26,27]) process discretely sampled time series, or *channels*, into tabular summaries of excess signal power that are well localized in the time-frequency plane. Each event is generally associated with a measure of the event’s time, duration, and amplitude or signal-to-noise ratio (ρ , often representative of how rare the event is). Within GW interferometers, there are typically $\mathcal{O}(10^4)$ separate channels that are sampled at frequencies fast enough to record noise within the detectors’ sensitive band.¹ Astrophysical GW signals are recorded with the highest sensitivity by the designated “GW channel” (a measure of the antisymmetric motion of the L-shaped interferometer arms [22,28]), but their passing may also be recorded by other auxiliary channels. Conversely, terrestrial and instrumental noise recorded by auxiliary channels may be present in the GW channel. Using auxiliary data to infer the nonastrophysical origin of noise artifacts in the GW channel can increase our confidence in the astrophysical nature of candidate events [12,29], but we generally do not have accurate models for many sources of noise, nor do we know the precise transfer functions between auxiliary channels and the GW channel (e.g., [30,31], although some monitors of the physical environment are given particular care [32]). Again, without models that describe the relationship between events’ waveforms (as they appear in multiple channels) and the GW channel, we are faced with the prospect of determining the significance of coincidences with only information about the events’ time, ρ , and approximate rate at which they occur.

This problem has been studied within the context of GW interferometers’ data quality in some detail. References [11] and [33] estimate the probability of observing a number of coincidences within a prespecified time window given an estimate of the rate of Poisson-distributed events. That is, they construct a counting experiment to determine the number of GW noise artifacts that are within small time windows surrounding a series of events in an auxiliary channel, estimating the significance with the Poisson probability for the observed number of events given a point-estimate of the rate of events. Reference [34] explores a similar approach, although their metric for the significance of coincidences is not derived from Poisson statistics. These algorithms implement costly, direct searches over both the size of the time window and the rarity of the auxiliary artifacts about which the windows are placed. They also intentionally learn to ignore many of the auxiliary channels through a supervised training regimen, extracting only a

subset of auxiliary channels shown to correlate with noise in the GW channel. If the selected auxiliary channel suddenly stops witnessing the source of noise (e.g., the sensor is unplugged), as occasionally happens, these algorithms lose all predictive power and must be retrained from scratch. Additionally, they often do not clearly quantify null results, or the false dismissal probability. That is, they do not provide a measure of how likely the observed data is if there is no observed coincidence within the time windows chosen *a priori*.

More complicated approaches that utilize general machine learning algorithms have also been explored (e.g., [10,12]; see [35] for a review), but have been met with mixed success. These may be improved with additional features based on our physical understanding of the situation. For example, Ref. [10] found that the most important features were the time separating auxiliary events from noise artifacts in the GW channel and how rare the auxiliary events were, implying that the algorithm spent most of its time learning how to determine the significance of coincidences.

This motivates the particular statistic investigated in this study. In addition to providing additional features for existing algorithms, we find it generally useful to determine the false dismissal probability directly based on the proximity of the coincidence and the rarity of the events involved. This can be determined from the receiver operating characteristic curves for some existing algorithms, but the interpretation thereof may depend on the properties of the signals being sought. That is to say, the false dismissal probability may depend on the selection criteria employed to identify the targeted events (loud noise artifacts may be fewer in number and easier to identify than more common quiet artifacts). Essentially, we want to be able to quantify a statement like “there is no important correlation or coincidence” rather than stating that “we did not find anything, but we do not know whether we would have found everything if anything was there.”

As such, we develop a specific null test in an attempt to directly quantify the significance of any putative coincidence. In particular, we seek the following:

- (i) A quantifiable null test for the hypothesis that the time of interest is randomly drawn (i.e., uncorrelated with a stationary Poisson process).
- (ii) A more computationally efficient algorithm than a direct search over time windows within counting experiments.
- (iii) A more natural interpretation for our results when we have a single event of interest, rather than a large set of possible coincidences.

We present a general null test under the assumption that the events are distributed according to stationary Poisson processes. This improves upon previous approaches to estimate the significance of coincidences between gamma-ray bursts and GW events (Appendix B of [36]), starting

¹Current interferometers record in excess of 2×10^5 channels, but the majority are recorded at sample rates ≤ 16 Hz.

from first-principles distributions for noise artifacts and explicitly allowing for the possibility of multiple coincident artifacts for any particular time of interest. This motivates our *pointy statistic*, so named because it is a very pointed test of a specific null hypothesis with well-specified prior assumptions. We then show its usefulness in a variety of situations, even when our motivating assumptions do not hold.

This paper is structured as follows. Section II motivates and derives the specific form of the pointy statistic. Section III then presents several examples, each of which demonstrates a different way to utilize the pointy statistic. We focus on examples from ground-based GW interferometry to demonstrate general techniques, but, as mentioned above, our techniques are immediately applicable within broader contexts. We conclude in Sec. IV.

II. FORMALISM

To formulate a general null test, we begin with a set of assumptions about individual channels. Specifically, we assume that events in each channel are independent of events in other channels and that each set of events is distributed according to a stationary Poisson process (i.e., with a constant rate). Although both of these assumptions are violated in many practical applications, they are generally good local approximations. Additionally, they provide a starting point for more general analyses.

We are interested in the probability that a coincidence between an event and a (separate event at a) random time would be as close or closer than the observed separation. That is to say, we want to know the probability that no events would occur within a window that is as wide as the observed coincidence. To do this, we need to know how the events are distributed in time. Specifically, we assume that the time between consecutive events Δt in a single channel with a Poisson rate of λ is distributed according to

$$p(\Delta t|\lambda) = \lambda e^{-\lambda \Delta t}. \quad (1)$$

Given this assumption, we naturally define the probability of observing an event as close or closer than τ to an uncorrelated time-of-interest to be

$$P(\Delta t \leq \tau|\lambda) = 1 - e^{-2\lambda\tau}, \quad (2)$$

where the factor of 2 is introduced by the fact that we search acausally (both backward and forward in time) to find the nearest event.

However, we do not observe the rate λ directly and instead estimate it from a counting experiment over a wider window. We marginalize over the uncertainty in λ such that

$$P(\Delta t \leq \tau|N, T) = \int d\lambda p(\lambda|N, T) P(\Delta t \leq \tau|\lambda), \quad (3)$$

where

$$p(\lambda|N, T) = \frac{T}{N!} (\lambda T)^N e^{-\lambda T}, \quad (4)$$

which assumes a uniform prior on λ and that the observed number of events (N) over the wider window (T) is Poisson distributed. We then obtain

$$P(\Delta t \leq \tau|N, T) = 1 - \left(1 + \frac{2\tau}{T}\right)^{-(N+1)}. \quad (5)$$

Alternatively, one could also assume a Jeffreys prior $p(\lambda) \propto \lambda^{-1/2}$, putting more weight on smaller rates *a priori*. However, we typically use large enough windows that $N \gg 1$, rendering the precise prior assumptions unimportant. For that matter, simply using the point estimate $\hat{\lambda} = N/T$ within $P(\Delta t \leq \tau|\lambda)$ instead of marginalizing gives similar results, although there are differences in the tails. Specifically, in the interesting limit of small $\tau (\ll T)$, we have

$$\lim_{\tau \ll T} P(\Delta t \leq \tau|N, T) \sim 2\tau \frac{N+1}{T}, \quad (6)$$

whereas

$$\lim_{\tau \ll T} P(\Delta t \leq \tau|\lambda = N/T) \sim 2\tau \frac{N}{T}. \quad (7)$$

If $N \gg 1$, then $N+1 \approx N$ and the rate is measured well enough that marginalization has little effect. However, the marginalized statistic [Eq. (6)] will not vanish, even if $N=0$, as it accounts for the nonvanishing likelihood that the true rate is nonzero.

As a final practical consideration, we often limit the maximum τ considered to be within some relatively large coincidence window ($\tau \leq \mathcal{T} \ll T$) so that

$$P(\Delta t \leq \tau|\tau \leq \mathcal{T}, N, T) = \frac{1 - (1 + 2\tau/T)^{-(N+1)}}{1 - (1 + 2\mathcal{T}/T)^{-(N+1)}}. \quad (8)$$

Equation (8), then, defines the pointy statistic for a particular interval τ defined between a time-of-interest and single event. It measures the probability of obtaining a coincidence as close or closer than the one observed assuming the events in the channel are distributed according to a stationary Poisson process and are uncorrelated with the time-of-interest. If this probability (p -value) is small, then such coincidences are rare. In this case, we may reject the null hypothesis that the event is uncorrelated with the time-of-interest, thereby inferring a causal connection.

Generally, the rate at which events occur within a channel also depends on their significance, or how loud they are, often characterized by a signal-to-noise ratio ρ . More extreme excursions (larger ρ) are less likely than




FIG. 1. Left: an example P_{\min} time series for a seismic isolation control channel in the LLO x-arm’s input test mass’s suspension system (L1:ISI-ITMX_CDMON_ST2_V1_V_IN1_DQ) between 8 and 128 Hz, centered at 22:35:30 GMT on Tue September 3, 2019. The sequence of peaks in the time series corresponds to a sequence of events in this channel, and the width of each peak roughly corresponds to the estimated Poisson rate. Right: a cumulative histogram of P_{\min} , from which we can directly measure the probability of observing smaller P_{\min} by random chance (i.e., the false alarm probability).

smaller ones. That is to say, the Poisson rate depends on ρ , with $d\lambda/d\rho \leq 0$. This presents a dilemma. If the rate of quiet events is high enough, including the sea of quiet events within Eq. (8) could drown out the significance of a loud coincident event. We resolve this by considering multiple subsets of events corresponding to multiple thresholds on ρ , computing the pointy statistic for each subset. To wit, we select subsets such that $\rho \geq \rho_{\text{thr}}$ and then minimize the p -value over ρ_{thr} such that

$$P_{\min}(\Delta t \leq \tau) = \min_{\rho_{\text{thr}}} \{P(\Delta t \leq \tau | \tau \leq T, N(\rho \geq \rho_{\text{thr}}), T)\}, \quad (9)$$

where $N(\rho \geq \rho_{\text{thr}})$ is the number of events that satisfy $\rho \geq \rho_{\text{thr}}$. Typically, a handful of thresholds are chosen to reduce the computational cost of the direct minimization over ρ_{thr} .² In this way, P_{\min} is sensitive to the fact that rare, loud events (small λ) may constitute a more significant coincidence than quieter events (large λ) even if the louder events are further away (larger τ). On the other hand, quiet events (large λ) that are exceptionally coincident (small τ) may be more significant than exceptionally loud events (small λ) that are very far away (large τ).

It is also worth noting that estimates of the rate at different thresholds extracted from the same underlying set of events will be correlated [37]. As such, a more complete analysis should include this correlated uncertainty during the minimization over ρ_{thr} or otherwise account for the

²Note that other algorithms [11,33,34] must marginalize directly over both ρ_{thr} and a time window.

correlated knowledge of coincidences’ significance at different thresholds. However, direct minimization works well in practice, and we leave further theoretical exploration to future work.

Nonetheless, minimization implies that our statistic may not be distributed as one would naïvely expect for a p -value when selecting random times. That is to say, they are not necessarily distributed uniformly between 0 and 1. While this slightly complicates the statistical interpretation of P_{\min} , it does not pose a significant problem; we directly measure the distribution of P_{\min} from randomly selected times to establish rigorous false alarm probabilities (FAPs).

Equipped with Eq. (9), we evaluate P_{\min} at every point in a regularly sampled time series (see below), obtaining results like those in Fig. 1. We additionally impose a minimum τ based on a fraction (F) of the most significant event’s durations (Δ). This approximates our uncertainty in the individual events’ central times, as these are not known precisely but should have an uncertainty that scales with their durations. Specifically, we assign a value,

$$\tau = \min_{i \in N(\geq \rho_{\text{thr}})} \{\max\{|t - t_i|, F\Delta_i\}\}, \quad (10)$$

to each time t in the time series, associating it with the nearest event above ρ_{thr} , and then compute P_{\min} via Eq. (9) while marginalizing over the unobserved Poisson rate. This produces small P_{\min} near the locations of events that decay approximately exponentially. Minimization over different ρ_{thr} , and therefore different estimates of the Poisson rate, produces peaks of different widths in the time series. Common, quiet events have relatively narrow peaks as

they occur at a high rate and therefore their significance drops off quickly. Loud, rare events form much broader (possibly deeper) peaks because they remain significant even with broader coincidence windows.

Our analyses focus on different ways to combine P_{\min} from a variety of auxiliary channels and times. Typically, we assume statistical independence between the different auxiliary channels within naïve Bayes classification schemes, finding that the assumption of independence renders analyses tractable while still performing well. Specifically, we generate P_{\min} separately for a collection (\mathcal{S}) of separate channels and/or times, combining them under the assumption of statistical independence so that

$$P_{\text{joint}} = \prod_{i \in \mathcal{S}} P_{\min}^{(i)}. \quad (11)$$

We refer to this practice of combining P_{\min} into P_{joint} as *stacking*. Section III describes a few ways in which we select sets \mathcal{S} that are of physical relevance within GW data analysis.

III. APPLICATIONS

We demonstrate the pointy statistic with two examples. We will first establish channel *safety* in interferometric GW detectors in Sec. III A, meaning whether veto conditions based on auxiliary channels could systematically reject real GW signals. This is a key aspect of any veto condition in GW transient searches. This application explores situations where we can identify repeated experiments, implying that we can stack P_{\min} for each channel from different times to learn more about channels separately. Section III B presents a second example of our pointy statistic where it constructs a veto condition based on the simultaneous use of information from multiple safe auxiliary channels. This allows us to investigate particular times in GW detectors and quantify the possible correlations between noise in the GW channel and auxiliary sensors.

A. Veto safety in gravitational-wave searches

We first focus on establishing channel safety in interferometric GW detectors, and, more specifically, the application of veto conditions in GW transient searches. GW interferometers, like advanced LIGO and Virgo, record not only the astrophysical GW signal but also a plethora of auxiliary channels that monitor the instruments' states and their physical environments [32,38–41]. Information from these auxiliary channels can be used as vetoes for putative astrophysical events. That is to say, noise artifacts that pollute the GW channel may also appear in auxiliary channels. However, although the main GW channel records the astrophysical signal with the highest sensitivity, some auxiliary channels may also witness GW signals either inadvertently or because they are part of the detector's

feedback control scheme. This can confuse the veto inference; if veto conditions are based on channels that have some sensitivity to GW signals, then these conditions may systematically reject real astrophysical events,³ the most extreme example of which would be constructing a veto based on the GW channel itself (e.g., [42]). In many analyses then, it is of the utmost importance to determine which auxiliary channels witness GW signals and are therefore *unsafe* for veto conditions. Most approaches to data quality and vetoes limit themselves *a priori* to the subset of auxiliary channels that can be demonstrated to be safe [12].

Due to the complexity of kilometer scale interferometers, it is not possible to directly measure every coupling between the over 2×10^5 auxiliary channels at each interferometer [or even the $\mathcal{O}(10^4)$ sampled faster than 16 Hz]. Therefore, we determine correlations probabilistically. This has historically been done with a set of transient injections performed in hardware by directly manipulating the instruments. Typically, an external force of known strength is applied to the interferometer's test masses, inducing differential arm motion that is read out analogously to an astrophysical signal. We additionally demonstrate how actual astrophysical events provide another method to independently “inject” signals into the interferometers, allowing one to establish auxiliary channel safety.

1. Safety studies with hardware injections

Hardware injections in GW interferometers directly manipulate the length of one of the detector's arms, producing a change in the differential arm length that mimics the effects of a GW signal. If the injected signal appears in an auxiliary channel, then it is also possible for a real GW to appear in that channel, and the channel is declared unsafe. Furthermore, couplings may depend on both the GW signal's frequency (e.g., different seismic isolation subsystems intentionally have very different responses at low and high frequencies) and its amplitude (a quiet GW signal may not show up significantly louder than the Gaussian noise in an auxiliary channel, and therefore may not produce a trigger with a particular ETG). As we can control the number and the properties of the hardware injections, but do not know the transfer functions between the injection site and all auxiliary channels, we adopt a statistical approach to safety. In this situation, we are interested in the behavior of each auxiliary channel separately, which provides the perfect example of how to combine independent, identically distributed trials with the pointy statistic. As such, we produce a time series

³It is worth noting that the concept of safety may be defined in terms of the raw auxiliary channel time series themselves or in terms of the events produced with a specific ETG. The former is often what is sought, but we are forced to use a single ETG due to computational limitations and must settle for some notion of the latter.

representing the minimized pointy statistic [Eq. (9)], computing a joint p -value [Eq. (11)] by stacking the individual P_{\min} associated with the time of each injection in a sequence (\mathcal{S}).

We summarize results obtained with hardware injections performed in the LIGO Livingston Observatory (LLO) during LIGO’s third observing run (O3) [43]. We injected sine-Gaussian signals as a general proxy for GW transient events that are compact in the time-frequency plane (spanning a short duration and limited to a narrow frequency bandwidth), allowing us to isolate possible couplings at different frequencies. The injections varied in amplitude, corresponding to expected ρ from 15 to 500 and with central frequencies (f) between 20 and 700 Hz. Each f - ρ combination was repeated 3 times, with each injection separated by 5 sec. Over 5500 auxiliary channels at LLO were analyzed for all injections using an ETG based on a dyadic Haar wavelet decomposition (KleineWelle: KW [26,44]).

Following the formalism of Sec. II, we used a time window ($T \sim 5000$ sec) much larger than the injection durations (each injection lasted for ~ 50 msec and the sequence of 84 injections spanned 435 sec) and the number of observed events above a set of thresholds [$N(\rho \geq \rho_{\text{thr}})$] in each auxiliary channel separately to marginalize over the Poisson rate (λ) for each threshold [Eq. (3)] before minimizing the P_{\min} over ρ_{thr} [Eq. (9)] to identify the most significant coincident event. This procedure was conducted repeatedly to build up a P_{\min} time series sampled at 128 Hz spanning the full 5000 sec window. The P_{\min} corresponding to the times of each triplet of hardware injections with the same f and ρ was then extracted via linear interpolation, constituting a set (\mathcal{S}) that we stack to compute P_{joint} [Eq. (11)].

Figure 2 demonstrates the observed distributions of P_{joint} for triplets of hardware injections and randomly selected sets of 3 times not associated with hardware injections that empirically determine the pointy statistic’s background distribution. Specifically, we select 3 random times repeatedly for each channel, computing P_{joint} for each random triplet. The total ensemble then empirically determines the expected FAP at each P_{joint} . The pointy statistic also normalizes all auxiliary channels so that we can fairly compare the significance of coincidences in each. *De facto*, this is done by not only considering Δt but also by incorporating knowledge of the Poisson rate of accidental coincidences. Indeed, for this reason, estimates of the background distributions for individual channels closely resemble our background estimate that averages over all channels. We note that we expect $\text{FAP} \sim P_{\min}$ by construction, as Fig. 1 demonstrates.

If the interferometric data were truly Poisson distributed and the repeated injections were exactly identical and independent, then P_{joint} should be a sufficient statistic. That is to say, no additional information would be available

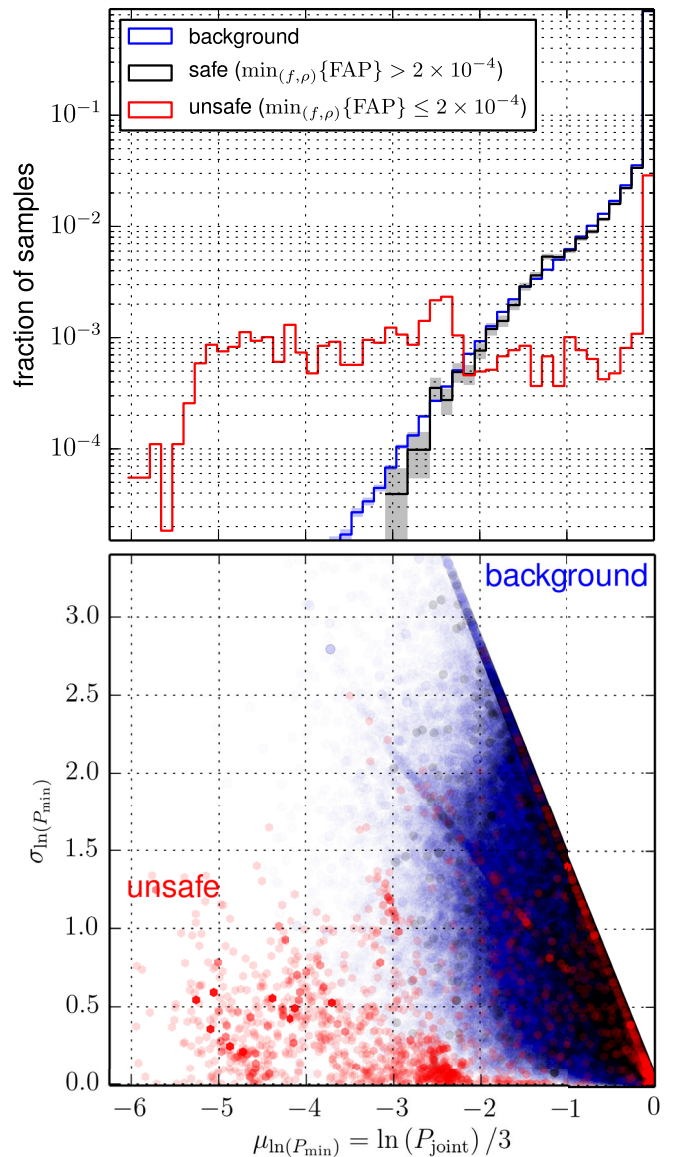


FIG. 2. Summary of results from hardware injections at LLO. Top: distributions of P_{joint} for channels declared unsafe (red), channels declared safe (black), and an estimate of the expected background (blue). Shaded bands represent $1\text{-}\sigma$ uncertainty estimates from counting statistics, showing good agreement between the background and safe distributions. We present distributions over all (f , ρ)-injection triplets, and safety was determined by the minimum FAP observed for each channel over all triplets. The unsafe distribution extends to $P_{\text{joint}} \sim 1$ because some channels that are unsafe at low frequencies and large ρ are safe at high frequencies and/or small ρ . Bottom: joint distributions of P_{joint} and the standard deviation of P_{\min} between the three repeated injections in each triplet ($\sigma_{\ln(P_{\min})}$). Although there is clearly additional structure in the two-dimensional distributions for both the unsafe (red hexagons) and background (blue circles), we find that safety criteria based solely on P_{joint} accurately capture all the relevant information.

to an analysis that considered the set of three P_{\min} for each injection triplet compared to one that only considered the corresponding P_{joint} . Figure 2 shows projected histograms

of P_{joint} for both the hardware injection triplets and for the background distribution. Indeed, we see that most of the information available is captured in this statistic. However, because the motivating assumptions behind the pointy statistic may be violated in practice (see the Appendix), we additionally consider a higher-dimensional inference. Specifically, we consider the joint distribution between P_{joint} and $\sigma_{\ln(P_{\text{min}})}$ (the standard deviation obtained from the triplet of the natural logarithm of P_{min}). For deterministic couplings, we would expect nearly identical P_{min} for each injection and therefore small $\sigma_{\ln(P_{\text{min}})}$. However, for rare but nonetheless accidental coincidences, we expect large $\sigma_{\ln(P_{\text{min}})}$ as rare coincidences are not likely to be repeated. Indeed, this is exactly the behavior we see in the background distribution: $\ln(P_{\text{joint}}) \propto \sigma_{\ln(P_{\text{min}})}$.

It is noteworthy that, even though real noise artifacts are not perfectly Poisson distributed, the information available in the joint two-dimensional inference is nonetheless captured by the one-dimensional inference over only P_{joint} with the additional caveat that a channel that is found to be unsafe for any f - ρ pair is declared unsafe for all f - ρ pairs.⁴ That is to say, unsafe channels identified in the two-dimensional analysis for a single f - ρ pair but were safe in the one-dimensional analysis of the same f - ρ pair were either associated with different frequency bands in the ETG (KW divides “raw channels” into several overlapping bandpasses, producing a set of events for each) that witnessed an injection less efficiently or with frequency-dependent transfer functions between the injection site and the auxiliary channel, which typically caused injections to appear with lower significance at higher frequencies. Both effects meant that the significance of coincidences was compared against an elevated rate of lower- ρ background artifacts in each channel, thereby increasing P_{joint} above the one-dimensional analysis’s threshold. However, because the injection set spanned a large range of frequencies, these channels were always identified as unsafe by either a separate injection triplet or a separate KW bandpass produced from the same raw channel.

Another important feature of our analysis is the direct assignment of a FAP to each channel. While final lists of safe and unsafe channels were still collated based on a hard threshold, our analysis allows that threshold to be chosen based on the expected number of false positives rather than an arbitrary detection statistic without an immediate physical interpretation. Within this study, we divided the channels into three sets: those found to be confidently unsafe ($\text{FAP} \lesssim 2 \times 10^{-4}$), those found to be confidently safe ($\text{FAP} \gtrsim 2 \times 10^{-3}$), and suspicious channels in between ($2 \times 10^{-4} \lesssim \text{FAP} \lesssim 2 \times 10^{-3}$). These

thresholds were chosen so that the expected number of false positives was ~ 1 for the confidently unsafe class and $\lesssim 10$ for the suspicious class based on the fact that we analyzed ≈ 5500 channels. Indeed, inspecting the identified channels’ behavior by hand, we found 1.9 ± 1.1 false positives (mean and standard deviation between different injection triplets) within the channels identified with $\text{FAP} \lesssim 2 \times 10^{-4}$ and 10.3 ± 3.4 channels with $2 \times 10^{-4} \lesssim \text{FAP} \lesssim 2 \times 10^{-3}$, in agreement with expectations for a basic counting experiment. Figure 2 shows the suspicious and confidently safe sets together, and we see that it closely follows the background distribution.

This analysis identified all 69 channels considered unsafe throughout O3 based on statistical evidence from hardware injections, including all channels known to be unsafe *a priori* as well as several that were not. Additional channels were also declared unsafe out of an abundance of caution because they either witnessed similar physical signals to channels that were identified as statistically unsafe (64 channels), in which case small changes to physical couplings could render the uncorrelated channels unsafe, or had historically been declared unsafe even if they no longer could be shown to correlate with hardware injections (45 channels). Furthermore, the consistency of the observed rate of false positives with expectations based on our FAP thresholds suggests that our analysis was able to cleanly separate the entire unsafe channel population, implying that all statistically unsafe channels were robustly identified. If this was not the case, then we would expect an excess of suspicious channels. Indeed, this emphasizes the pointy statistic’s utility as it provides a way to efficiently extract all relevant information about coincidences, even when its motivating assumptions are not perfectly satisfied, and our ability to extract more information about individual channels by examining repeated experiments simultaneously.

2. Safety studies with astrophysical events

While hardware injections allow us to perform repeated experiments, they inevitably require the interferometers to be manipulated in ways they would not be by a bona fide GW signal. For example, the channels that record the excitations introduced into the interferometer will obviously correlate with the injected signals, and therefore hardware injections cannot be used to determine the safety of the excitation channels. If one can inject signals into the interferometer in multiple ways (e.g., by driving the length of each arm separately), then one could perform multiple sets of injections, using each to cross-check the channels used during the other set.⁵ However, we present an

⁴While our analysis identifies both the frequency and the amplitude at which auxiliary channels become unsafe, current applications of safety information do not account for these dependencies. We therefore adopt the logic that channels are safe if and only if they are shown to always be safe.

⁵We expect excitation channels to be safe *a priori* in that the presence of coincident transients in the excitation channels would call into doubt the astrophysical nature of any GW candidate. Nonetheless, one may want to confirm this out of an abundance of caution.

alternative that uses the “perfect” injections available through confident GW detections.

As of the time-of-writing, the advanced LIGO and Virgo detectors have published several dozen confident detections of binary black hole (BBH) coalescences [24,25]. The astrophysical nature of these events is not in doubt, and therefore we can exploit them as natural experiments to determine which auxiliary channels witness real GW signals. Previous work has shown how known astrophysical sources can be used to calibrate networks of GW interferometers [23], and herein we extend this approach to auxiliary channel safety.

However, using astrophysical events to determine veto safety presents some limitations. High-mass BBH mergers detectable in the LIGO interferometers are characterized by rapid frequency evolution and only extend up to relatively low frequencies [$\mathcal{O}(100$ Hz)]. Although these astrophysical signals do not require direct manipulation of the interferometers, they also cannot be repeated. Each BBH merger comes from a separate astrophysical system with different component masses and spins, meaning that the maximum frequency reached during the coalescence and the signal amplitude will differ for each system. As channel safety is known to depend on the signal frequency and amplitude, we cannot stack multiple astrophysical events without extreme care.

Nonetheless, we examine an individual confident BBH detection to establish channel safety as a proof-of-principle. We analyze GW170814 [45,46] in LLO, where it was detected with a $\rho = 13.7$, much less than the $\rho \sim 500$ achieved at similar frequencies with hardware injections. Again, we processed approximately 5500 auxiliary channels with KW, each sampled faster than 16 Hz, generating pointy time series from the resulting sequences of events. We again estimated the background distribution of P_{\min} from accidental coincidences by randomly selecting times from the ~ 5000 sec surrounding GW170814 for each channel. We additionally require that these random times be at least 1 sec away from loud noise artifacts in the GW channel to avoid contaminating the background distribution with auxiliary channels that may witness such noise. Figure 3 shows the result.

We see a clear excess at small P_{\min} within the on-source distribution. However, the background distribution for a single random time has significant support down to even lower P_{\min} . This implies that the FAP associated with the unsafe channels ($P_{\min} \lesssim e^{-4}$) is as high as 5×10^{-3} , implying we would be forced to accommodate ~ 25 false positives if we were to identify all channels known to be unsafe in this way. Therefore, we conclude that it is unlikely that a single BBH will produce on-source coincidences significant enough to overcome the large trials factor from the $\mathcal{O}(10^4)$ auxiliary channels considered. The confident identification of channel safety with low contamination therefore requires repeated hardware injections. These should be conducted regularly throughout observing

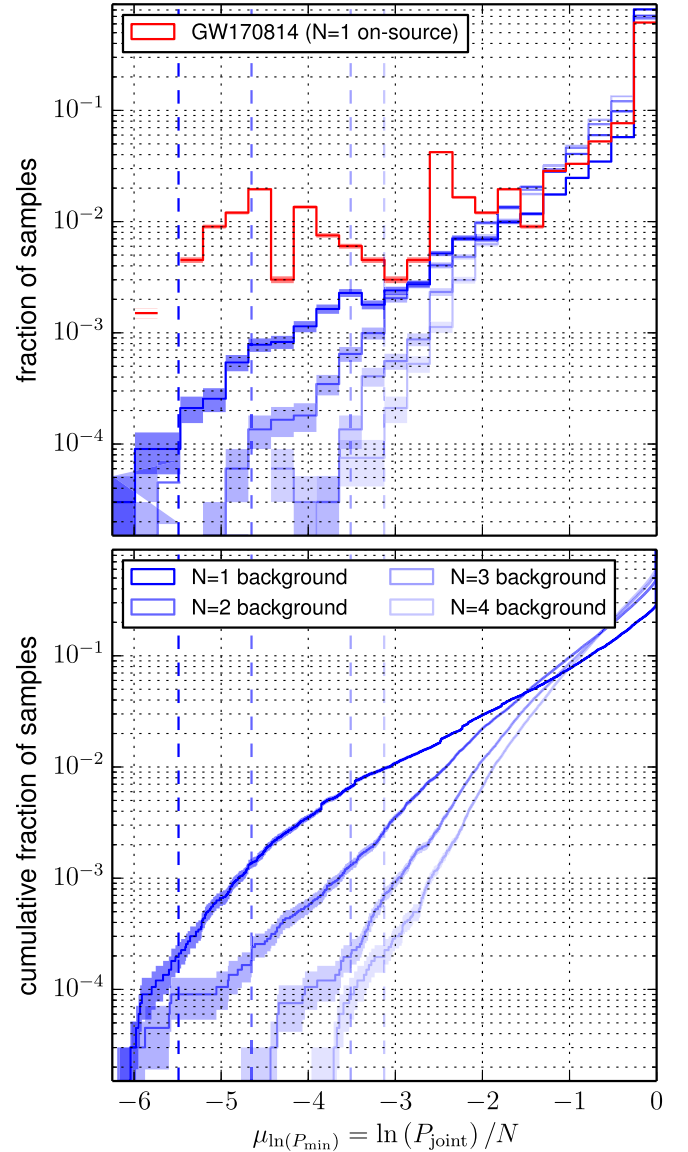


FIG. 3. Top: P_{joint} distributions of background estimates obtained by stacking different numbers of random times (N), with the observed distribution from GW170814 superimposed. Vertical lines denote the approximate value of P_{joint} required to confidently detect an unsafe channel (FAP $\lesssim 2 \times 10^{-4}$) given the large number of auxiliary channels considered. Bottom: cumulative P_{joint} distributions for the same background estimates. While GW170814 clearly shows an excess of channels with small p -values, we find that we need to combine at least three repeated observations before the corresponding P_{joint} would be clearly separated from the background.

runs, which has not been the case to-date during the advanced detector era.

Figure 3 goes further, estimating the background that would be obtained by stacking more and more events. We see that, if we require FAP $\lesssim 10^{-4}$ at $P_{\text{joint}} \sim e^{-4}$ to limit the expected number of false positives to $\lesssim 1$, we need to stack at least three repeated experiments. The hardware

injection campaign described in Sec. III A 1, then, performed the minimum number of injections needed to cleanly separate safe and unsafe channels.

B. Glitch identification via naïve Bayes classification

Section III A described how to determine which auxiliary channels are safe and therefore can be used within veto conditions. We now demonstrate the construction of such a veto condition tailored to a specific class of non-Gaussian noise artifact, colloquially referred to as *whistles*. While pointy analyses have also targeted other classes of noise, most notably a study of *blip* glitches during LIGO's first observing run (O1), these have since been investigated in some detail in other work [30]. We focus on whistles to demonstrate unique advantages of the naïve Bayes classification with the pointy statistic and to avoid repeating conclusions already presented elsewhere.

Our specific use case was motivated by a small statistical excess of foreground events in a search for unmodeled GW transients during the second observing run (O2) [47,48]. While no single event in that search was significant enough to claim a detection, a collection of events in the tail of the distribution constituted a small excess above what was expected, which prompted further follow-up. Inspecting a collection of the most significant foreground events, we determined that the tail was mostly composed of either blip or whistle glitches. While it is known that the majority of blip glitches do not have clear auxiliary witnesses (see, e.g., [30]), whistles typically do have clear auxiliary witnesses and should be straightforward to veto.

As such, we analyzed the full set of auxiliary channels surrounding four whistles from the tail of the background distribution. For each whistle, we identified the set of auxiliary channels found to be in coincidence with the background event with $P_{\min} \leq 3 \times 10^{-2}$ (the precise value has little impact), corresponding to a relatively high FAP ($\gtrsim 2\%$) and a large number of expected accidental coincidences given the large number of channels analyzed. This procedure identified ~ 130 channels that are of interest around each whistle (for some as few as 87, for others as many as 208), most of which were likely accidental coincidences. While some individual auxiliary witnesses [typically angular sensing and control (ASC) and length sensing and control (LSC) channels] produced P_{\min} small enough to possibly veto these background events separately,⁶ we note that there was also an excess of channels with slightly larger P_{\min} beyond what one would expect from a basic counting experiment (cf., sanity checks in Sec. III A 1). As such, we were additionally interested in using information from all auxiliary witnesses simultaneously rather than only the individual channels that were

most significant. In this situation, the set \mathcal{S} in Eq. (11) represents a collection of separate auxiliary channels evaluated for a single time-of-interest, rather than P_{\min} from a single channel evaluated at multiple times.

Using all channels identified as possibly significant around each whistle, we construct a naïve Bayes statistic by assuming the channels were independent (though they likely are not) and multiplying their individual P_{\min} to




FIG. 4. An example whistle glitch identified in the background of a search for unmodeled transients during O2. Top: a time-frequency representation [26] of 20 sec of the strain data from LLO, with insets highlighting the noise identified by P_{joint} , which have characteristic durations of $\lesssim 50$ ms. (*in order below*). We additionally show P_{joint} from the 208 auxiliary channels found in coincidence with the background event as well as P_{\min} from two individual channels, one found to have a low FAP (L1:ASC-AS_B_RF26_Q_PIT_OUT_DQ, a control channel associated with angular alignment of the interferometer) and one randomly selected (L1:SUS-SR2_M3_MASTER_OUT_UL_DQ, a sensor in the signal-recycling mirror suspension system). Projected histograms show the cumulative distribution of the pointy statistics for each time series, and gray lines denote the time of the background event, the statistic at that time, and the corresponding FAP. We note that, although it was not identified in the search background, there is another, louder whistle approximately 2 sec before the background event. P_{joint} correctly identifies both with FAP $< 10^{-5}$.

⁶Some individual channels produced $P_{\min} \leq 10^{-3}$, implying FAP $< 10^{-4}$ (see the $N = 1$ background estimate Fig. 3) and an expected number of accidental coincidences $\lesssim 1$.

obtain P_{joint} . While many of the channels were accidental coincidences, if a subset of auxiliary channels did indeed witness the noise source, then their individual significances should stack together to provide a more confident veto than any individual channel would produce. Indeed, this is exactly the behavior we observe.

Figure 4 shows P_{joint} for a background whistle identified at 03:03:15.05 GMT on February 21, 2017 in the LIGO Livingston detector as well as the individual P_{min} time series from significant and insignificant individual auxiliary channels with events found in coincidence. From the time series, we see that there is a clear peak in both P_{joint} and the significant channel's P_{min} at the time of the whistle, but the FAP associated with P_{joint} ($\lesssim 10^{-5}$) is lower than the FAP associated with individual P_{min} ($\gtrsim 10^{-4}$). This behavior suggests that each individual witness is polluted by other sources of noise that do not couple to the GW channel in the same way, thereby confounding the inference if only a single channel is used. However, the additional noise sources are uncorrelated in each auxiliary channel, and the probability of these uncorrelated noise sources occurring

simultaneously is low. Therefore, we can be even more confident that there is a whistle in the GW channel when all the auxiliary channels contain a significant transient. Additionally, if a single auxiliary channel suddenly stops witnessing whistles, this procedure retains its predictive power, unlike other algorithms that rely on correlations between the GW channel and single auxiliary channels [11,33,34].

This behavior repeats itself for several whistles identified in the search background, typically with similar auxiliary witnesses. This begs the question of whether a subset of auxiliary witnesses were always coincident with this type of noise, or whether detector nonstationarity led to modified couplings throughout the observing run. Taking the intersection of each set of ~ 130 auxiliary witnesses identified with $P_{\text{min}} \leq 3 \times 10^{-2}$ around each whistle produces a much smaller list of only 37 channels. We then produced a naïve Bayes P_{joint} time series using only this subset of channels throughout O2 to evaluate its performance.

While the resulting statistic correctly identified the whistles already known from the search background, it




FIG. 5. The targeted whistle veto's performance around two of the events identified in the search's background. (Top): P_{joint} time series throughout the ~ 2.5 hr stretch. Gray shading corresponds to the specific example times called out below. (Middle): expanded views of P_{joint} around four example local minima. (Bottom): time-frequency representations [26] of the GW channel at each example. We see that a variety of behavior is identified by the veto, including nonstationary coupling to the GW channel: Left to right: a pair of whistles with different ρ , no apparent non-Gaussian noise, three whistles, and a whistle followed by low-frequency scattering arches (during which there seems to be some activity in a subset of the auxiliary channels).

showed mixed performance at other times throughout O2. Indeed, two of the original four events identified in the search background came from the same stretch of data. Figure 5 shows the behavior of P_{joint} surrounding those events. During this ~ 2.5 hr period between 22:21:53 GMT on January 5, 2017 and 00:52:57 GMT on January 6, 2017, which contains a plethora of non-Gaussian noise, we typically find that the P_{joint} derived from this subset of channels clearly identifies whistles at LLO. However, there appears to be significant variation in the signal-to-noise ratio (ρ) in the GW channel associated with those whistles, somewhat independent of the significance of the P_{joint} minima. This, perhaps, suggests that the auxiliary channels consistently witness a possible source of noise in the GW channel, but the coupling between that noise and the GW channel varies in time, potentially even multiple times per hour (see Fig. 5). Anecdotally, we also found some evidence that the frequency content of noise identified in the GW channel by this set of channels also varied, at times resembling the archetypal whistles shown in Fig. 4 but at other times identifying low-frequency features, not dissimilar to scattering arches [49]. This type of nonstationarity implies that, even though a particular time may correspond to an exceptionally rare P_{joint} , it does not necessarily imply the presence of noise in the GW channel. As such, the generalization error prevents us from applying this subset of channels to a broader stretch of O2, although the associated FAP could be limited to $\lesssim 10^{-3}$ while identifying a nontrivial fraction of whistles, and we may be forced to identify different sets of witnesses through periodic retraining (cf. [12]).

Nonetheless, we find that the naïve Bayes statistic constructed from the intersection of channels identified around background events correctly identifies the two whistles assigned the largest significance within the search’s foreground. Specifically, P_{joint} corresponds to $\text{FAP} \lesssim 3 \times 10^{-4}$ in the neighborhood of each foreground event, even though those events are at least 3.5 weeks away from the nearest of the four background events considered.

Additionally, it is difficult to quantify the fraction of whistles identified by the veto condition due to the difficulty in identifying all whistles present in the detector. While automatic glitch-classification schemes based on time-frequency representations of the GW channel have been explored (e.g., [31]), these tools do not provide a simple estimate of the completeness of their catalogs. That is to say, they do not assign a detection efficiency for identifying individual glitch classes. Nonetheless, anecdotal evidence suggests that the identified subset of auxiliary channels jointly provide a useful veto condition that identified whistles throughout O2. Indeed, it does not require deterministic couplings or perfect witnesses, instead relying on the wisdom of a collection of witnesses rather than a single expert.

As a final note, we remark that likelihood ratio tests [50] may obviate the need to down-select the list of channels in

any way. However, we leave a full exploration of likelihood ratio tests based on the pointy statistic to future work.

IV. DISCUSSION AND CONCLUSIONS

This work was motivated within the context of GW experiments, in particular the need for a null test to determine the significance of coincidences between GW candidates and auxiliary channels at particular times. Our statistic can accommodate an arbitrary set of channels with disparate event rates and amplitudes, as it reduces all measurements to an intuitive probabilistic statement that is directly comparable across all channels. Even though the formal derivation of our null test assumes events are independent and Poisson distributed in each channel separately, and this assumption is known to be violated in real-world situations, we nonetheless show our test’s utility in several settings. The power of our null test in establishing the statistical significance of a measurement is retained, and we precisely quantify that significance by directly measuring the background distribution using the same data.

We have shown how our null test can establish which channels in GW detectors can safely be used to construct vetoes by both combining multiple, repeated measurements in a controlled experiment (like hardware injections) and by analyzing bona fide astrophysical events. We find that individual astrophysical events will not be able to overcome the large trials factor from $\mathcal{O}(10^4)$ channels, and typically at least three repeated injections will be needed to limit the expected number of false positives to $\lesssim 1$. It can also enable targeted studies of transients at specific times in GW detectors, thereby illuminating their couplings and ultimate origin. We have shown one such case, focusing on non-Gaussian noise artifacts called whistles. We identify a subset of auxiliary channels that individually correlate with whistles, but together provide a particularly efficient identification scheme during parts of O2.

The utility of such null tests has also proven useful when confirming the astrophysical nature of individual GW candidate events. Combining the list of auxiliary channels identified as safe with hardware injections with the type of analysis we demonstrated for whistles, one can directly compute the FAP associated with all auxiliary events coincident with a GW candidate. If all FAPs are large, this constitutes evidence that the behavior of each auxiliary channel is uncorrelated with the GW channel at that time. Indeed, such analyses were conducted by-hand for the first few GW detections, but further parallelization is needed to scale the pointy analysis to higher detection rates expected from planned detector upgrades [51].

In conclusion, the pointy statistic and methods developed herein are applicable to a wide range of problems beyond GW astrophysics. Indeed, our null test captures the relevant features of coincidences and naturally quantifies their significance as intuitive probabilistic statements, which are of general use.

ACKNOWLEDGMENTS

The authors gratefully thank Lindy Blackburn for helpful discussions in which the idea of a safety measurements with real astrophysical signals was first suggested. The work of R. E. was supported at the University of Chicago by the Kavli Institute for Cosmological Physics through an endowment from the Kavli Foundation and its founder Fred Kavli. The work of G. M. and E. K. was supported by the National Science Foundation (NSF) through Grant No. PHY-1764464 to the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the NSF and operates under cooperative agreement PHY-1764464. The authors also gratefully acknowledge the computational resources provided by the LIGO Laboratory and supported by NSF Grants No. PHY-0757058 and No. PHY-0823459.

APPENDIX: ON THE ASSUMPTIONS OF STATIONARITY, POISSONIANITY, AND INDEPENDENCE

The pointy statistic is based on the assumptions of stationarity and Poissonianity. In practice, both of these may break down. Although the pointy statistic remains useful in such cases, we briefly discuss a few diagnostics that may help address which assumption is breaking down.

Specifically, the assumptions of stationarity and Poissonianity are encoded in the distribution of time between consecutive events [Eq. (1)]. By directly examining the observed distribution of the time between events, and modeling how this distribution changes over time, one could empirically remove the need for these assumptions. However, data sets may be sparse enough that some level of modeling will be needed, whether in the form of the distribution or how rapidly it can vary over time. A stationary Poisson process is a reasonable approximation for the data we consider.

Correlograms, or histograms of the time between events, can be useful diagnostic tools. Correlograms either compare all events within a single set (autocorrelogram) or events from two separate sets (crosscorrelogram). Figure 6 demonstrates examples of each. Comparing these with our expectations from Poissonianity and independence provides a valuable sanity check. For a Poisson process, all the pairwise time differences between events, i.e., the autocorrelogram, are expected to form a uniform distribution. This remains the case for the crosscorrelogram when constructed from two *independent and uncorrelated* Poisson processes. For example, channels may be Poisson distributed but correlated, in which case the autocorrelograms would agree with expectations but the crosscorrelogram would not. Similarly, non-Poissonianity in a single channel would appear in that channel's autocorrelogram, but the crosscorrelogram could still appear as expected if the channels are independent.

Furthermore, we sometimes assume statistical independence between different channels. This allows us to analyze

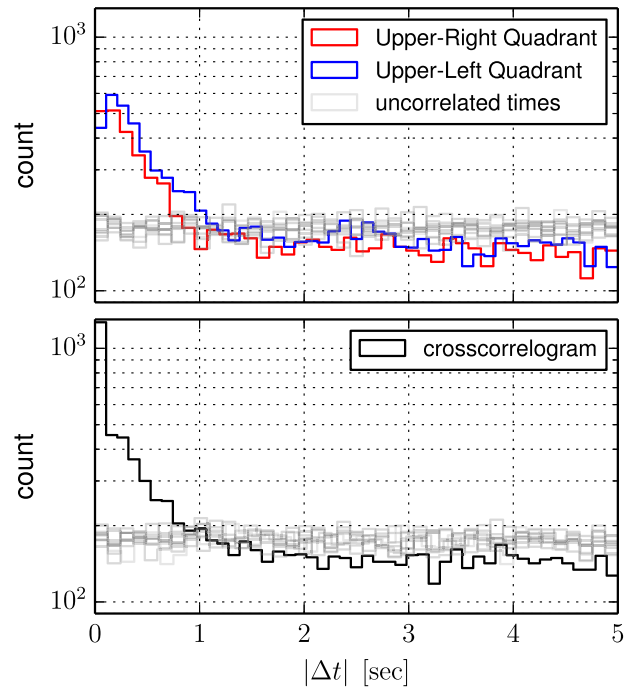


FIG. 6. Top: separate autocorrelograms for events from two channels measuring the current through two nearby electrostatic actuators in the power-recycling mirror suspension system of LIGO Livingston Observatory between 8–128 Hz: red: upper-right and blue: upper-left quadrants of the second stage: L1:SUS-PR2_M2_FASTMON_{U,L}R_OUT_DQ along with (gray) distributions from the same number of randomly distributed events. Bottom: the crosscorrelogram between the two channels. Although there is an excess of nearby events compared to the randomly distributed times (the channels produce correlated clusters of events), the channels are well modeled as independent Poisson processes on time scales $\gtrsim \mathcal{O}(1 \text{ sec})$.

each channel independently and then stack p -values by multiplying results obtained from different channels to create a naïve Bayes detection statistic: P_{joint} . This assumption may hold quite well in many cases, such as different detector subsystems that are isolated and do not interact (e.g., sensors in different end stations within kilometer scale GW interferometers). Nonetheless, this assumption can often break down in practice. For example, many auxiliary channels within GW detectors record data from nearly identical sensors (e.g., photodiodes are often divided in half or in quadrants). Such channels will almost certainly witness correlated signals and therefore produce correlated sequences of events. Similarly, multiple sensors often observe the same suspension system within ground-based GW interferometers, meaning that a single jolt of excess ground motion could appear in all of them. As an extreme example, actuation signals within control loops are often explicitly constructed as linear combinations of the sensor data. Reference [12] discusses the fully connected probabilistic graphical model that describes the complicated situation within real detectors.




FIG. 7. Left: observed distribution of the time between consecutive events in the upper-right quadrant of the power-recycling mirror suspension system shown in Fig. 6. Shaded regions approximate $1\text{-}\sigma$ counting uncertainty. We note that the data are roughly exponentially distributed, as expected, but with an excess at $\Delta t \lesssim 1$ sec, in agreement with Fig. 6. Right: the distribution of the first digit (*red*) and the analytic predictions from Benford's law (*black circles*).

Benford's law [52,53], or the relative frequency of the first digit (d) of each number in a large set,⁷ provides an additional test of whether a channel's behavior is expected. In a base-10 number system, Benford's law predicts that the relative frequency of the first digit in a number drawn is given by

$$P(d) = \log_{10} \left(1 + \frac{1}{d} \right). \quad (\text{A1})$$

⁷For example, $d = 1$ for 13 and $d = 8$ for 831.

Indeed, the time-between-events in a Poisson distribution is exponential and should closely follow Benford's law [54]. Figure 7 demonstrates this for one of the channels shown in Fig. 6. As with the correlogram, we see that there are deviations from the expected distribution, but the data are reasonably approximated by the analytic prediction nonetheless.

While the pointy statistic's assumptions may break down, Sec. III shows how to directly measure how P_{\min} and/or P_{joint} is distributed in the presence of correlated channels within real data sets. This means we can still make rigorous, statistically precise statements regardless of whether or not the assumptions that motivate the functional form of the pointy statistic and our naïve Bayes techniques hold exactly.

We would be remiss if we did not discuss at least one additional subtlety. The pointy statistic's assumption of Poissonianity implies events within a channel are completely uncorrelated. If this is violated by, for example, a preference for events to occur regularly on integer second boundaries, then selecting time that are also regularly spaced with the same periodicity may introduce spuriously small P_{\min} . That is to say, the pointy statistic detects that events occur too close to the selected times to be due to random chance, but this is because we have accidentally aligned our selected times with preexisting correlations within that channel. Indeed, correlation does not imply causation. Because we often wish to infer a common cause based on coincidence null tests, care should be taken to randomly distribute the selected times as much as possible to avoid accidentally coinciding with existing periodic behavior within series of events that are not Poisson distributed. This is particularly relevant for hardware injections to determine auxiliary channel safety within GW interferometers, where analysts have complete control over the parameters of each injection.

-
- [1] The Pierre Auger Collaboration, *Nucl. Instrum. Methods Phys. Res., Sect. A* **798**, 172 (2015).
 - [2] T. DeYoung, *Nucl. Instrum. Methods Phys. Res., Sect. A* **692**, 72 (2012); 3rd Roma International Conference on Astroparticle Physics.
 - [3] F. Halzen and S. R. Klein, *Rev. Sci. Instrum.* **81**, 081101 (2010).
 - [4] M. Ageron *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **656**, 11 (2011).
 - [5] P. Antonioli *et al.*, *New J. Phys.* **6**, 114 (2004).
 - [6] W. O. Schumann, *Z. Naturforsch. A* **7**, 149 (1952).
 - [7] W. O. Schumann, *Z. Naturforsch. A* **7**, 250 (1952).
 - [8] E. Thrane, N. Christensen, and R. M. S. Schofield, *Phys. Rev. D* **87**, 123009 (2013).
 - [9] B. P. Abbott *et al.*, *Astrophys. J.* **848**, L12 (2017).
 - [10] R. Biswas *et al.*, *Phys. Rev. D* **88**, 062003 (2013).
 - [11] R. Essick, L. Blackburn, and E. Katsavounidis, *Classical Quant. Grav.* **30**, 155010 (2013).
 - [12] R. Essick, P. Godwin, C. Hanna, L. Blackburn, and E. Katsavounidis, *Mach. Learn. Sci. Technol.* **2**, 015004 (2020).
 - [13] A. Bhandari, A. L. Sangal, and K. Kumar, *Secur. Commun. Networks* **9**, 2222 (2016).
 - [14] J. P. Bouchaud, *arXiv:0903.2428*.
 - [15] S. Poon, M. Rockinger, and J. Tawn, *Rev. Financ. Stud.* **17**, 581 (2004).
 - [16] J. Aasi *et al.*, *Classical Quant. Grav.* **32**, 074001 (2015).

- [17] F. Acernese *et al.* (Virgo Collaboration), *Classical Quant. Grav.* **32**, 024001 (2015).
- [18] M. Isi, R. Smith, S. Vitale, T. J. Massinger, J. Kanner, and A. Vajpeyi, *Phys. Rev. D* **98**, 042007 (2018).
- [19] N. J. Cornish and T. B. Littenberg, *Classical Quant. Grav.* **32**, 135012 (2015).
- [20] J. B. Kanner, T. B. Littenberg, N. Cornish, M. Millhouse, E. Xhakaj, F. Salemi, M. Drago, G. Vedovato, and S. Klimenko, *Phys. Rev. D* **93**, 022002 (2016).
- [21] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
- [22] C. Cahillane *et al.*, *Phys. Rev. D* **96**, 102001 (2017).
- [23] R. Essick and D. E. Holz, *Classical Quant. Grav.* **36**, 125002 (2019).
- [24] B. P. Abbott *et al.*, *Phys. Rev. X* **9**, 031040 (2019).
- [25] R. Abbott *et al.*, [arXiv:2010.14527](https://arxiv.org/abs/2010.14527).
- [26] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *Classical Quant. Grav.* **21**, S1809 (2004).
- [27] P. Godwin, Low-latency statistical data quality in the era of multi-messenger astronomy (2020), <https://inspirehep.net/literature/1817585>.
- [28] L. Sun *et al.*, *Classical Quant. Grav.* **37**, 225008 (2020).
- [29] B. P. Abbott *et al.*, *Classical Quant. Grav.* **35**, 065010 (2018).
- [30] M. Cabero *et al.*, *Classical Quant. Grav.* **36**, 155010 (2019).
- [31] M. Zevin *et al.*, *Classical Quant. Grav.* **34**, 064003 (2017).
- [32] A. Effler, R. M. S. Schofield, V. V. Frolov, G. González, K. Kawabe, J. R. Smith, J. Birch, and R. McCarthy, *Classical Quant. Grav.* **32**, 035017 (2015).
- [33] J. R. Smith, T. Abbott, E. Hirose, N. Leroy, D. MacLeod, J. McIver, P. Saulson, and P. Shawhan, *Classical Quant. Grav.* **28**, 235005 (2011).
- [34] T. Isogai *et al.*, *J. Phys.* **243**, 012005 (2010).
- [35] E. Cuoco *et al.*, *Mach. Learn. Sci. Technol.* **2**, 011002 (2020).
- [36] V. Connaughton *et al.*, *Astrophys. J.* **826**, L6 (2016).
- [37] R. Lynch, S. Vitale, and E. Katsavounidis, [arXiv:1811.01297](https://arxiv.org/abs/1811.01297).
- [38] C. L. Mueller *et al.*, *Rev. Sci. Instrum.* **87**, 014502 (2016).
- [39] F. Matchard *et al.*, *Classical Quant. Grav.* **32**, 185003 (2015).
- [40] A. Staley, Ph.D. Thesis, Columbia University, 2015.
- [41] J. Graef Rollins, [arXiv:1604.01456](https://arxiv.org/abs/1604.01456).
- [42] B. Zackay *et al.*, [arXiv:1908.05644](https://arxiv.org/abs/1908.05644).
- [43] R. Essick, Detchar safety hardware injections starting at 21:00:00 utc 3 sep 2019, <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=48277> (2019).
- [44] L. Blackburn, Kleine-welle algorithm, <https://dcc.ligo.org/public/0027/T060221/000/T060221-00.pdf> (2007).
- [45] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **119**, 141101 (2017).
- [46] The LIGO Scientific Collaboration, Auxiliary channel three hour release, <https://www.gw-openscience.org/auxiliary/GW170814/> (2020).
- [47] R. Lynch, S. Vitale, R. Essick, E. Katsavounidis, and F. Robinet, *Phys. Rev. D* **95**, 104046 (2017).
- [48] B. P. Abbott *et al.*, *Phys. Rev. D* **100**, 024017 (2019).
- [49] S. Soni *et al.*, *Classical Quant. Grav.* **38**, 025016 (2021).
- [50] J. Neyman and E. S. Pearson, *Phil. Trans. R. Soc. A* **231**, 289 (1933).
- [51] B. P. Abbott *et al.*, *Living Rev. Relativity* **21**, 3 (2018).
- [52] S. Newcomb, *Am. J. Math.* **4**, 39 (1881).
- [53] F. Benford, *Proc. Am. Phil. Soc.* **78**, 551 (1938), <https://www.jstor.org/stable/984802?seq=1>.
- [54] M. Cong, C. Li, and B.-Q. Ma, *Phys. Lett. A* **383**, 1836 (2019).