Draft version January 15, 2015
Preprint typeset using LaTeX style emulateapj v. 11/10/09

ENABLING HIGH CONFIDENCE DETECTIONS OF GRAVITATIONAL-WAVE BURSTS

Tyson
Sweet Home

Jonah
Sunny

and

Neil
Big Sky

Draft version January 15, 2015

ABSTRACT

We show that a simple model demonstrates some key properties of the signal to glitch bayes factor. Moreover, we show that background rates using BayesWave are predictable.

Subject headings: gravitational waves

1. THE SIGNAL TO GLITCH BAYES FACTOR

Past searches for Burst signals are based on frameworks that employ detection statistics to measure the likelihood that Gaussian noise could produce the data [cite burst search papers]. After the detection statistic is calculated, various cuts are applied to separate glitches from astrophysical signals. In these searches, the false alarm rate (FAR), as measured by time-slide studies, is difficult to predict in advance. Typically, the searches succeed in identifying times that are inconsistent with Gaussian noise, but the FAR is dominated by glitches that are not removed by the the applied cuts. In many cases, the distribution of these glitches as a function of the detection statistic has a long "tail", suggesting that false alarms are possible at any value of the detection statistic. This leads to a dissapointing conclusion: essentially no gravitational wave signal, no matter how loud, could be declared a high-confidence detection.

In contrast, the `BayesWave` algorithm includes a glitch model, and computes the evidence ratio that a given trigger is a gltich or an astrophysical signal. In this section, we propose using this evidence, known as the Bayes Factor, as a detection statistic, and explore its propoerties. Given minimal assumptions about the underlying glitch population, we find that this statistic has some desirable properties:

- The FAR always goes to lower values with increasing detection statistic.
- The expectation value of the loudest event is fairly insesitive to the underlying glitch population, so it is possible to predict it in advance.
- A simple, robust model limits the FAR, so that it may be extrapalated to very high confidence levels.

1.1. *The Occam Factor*

For each trigger, `BayesLine` estimates the evidence for each of three models: signal, glitch, or noise. We can then construct the Bayes Factor comparing any two models as the ratio of supporting evidence. For example, the Bayes Factor for data $\mathbf{s}$ comparing the signal (S) and glitch (G) models is

$$B_{SG} = \frac{p(\mathbf{s}|S)}{p(\mathbf{s}|G)} \tag{1}$$

The evidence for either model may be written as a marginalization over the parameters of the model, $\mathbf{h}$. For example, for the signal model:

$$p(\mathbf{s}|S) = \int p(\mathbf{s}|\mathbf{h}, S)p(\mathbf{h}|S)d\mathbf{h}. \tag{2}$$

For signals with enough SNR to be strong candidates for detections as Burst Signals (around SNR ≥ 10) we expect the integrand of Equation (2) to be strongly peaked at the most likely parameter values of the waveform $\mathbf{h}_{\mathrm{MP}}$. If we take the peak to have an area $\sigma_{\mathbf{h}|S}$ in parameter space, then, the integral may be approximated as

$$p(\mathbf{s}|S) = p(\mathbf{s}|\mathbf{h_{MP}}, S) \times p(\mathbf{h_{MP}}|S)\sigma_{\mathbf{h}|S} \equiv \Lambda_S \times \mathrm{O}_S, \tag{3}$$

where the first term represents the best fit likelihood for the signal model (Λ_S), and the following term is known as the Occam Factor (O_S). The Occam Factor represents the fraction of parameter space where the likelihood is significantly high. In the parlance of the literature, this is described as the posterior accessible volume as a fraction of the prior accessible volume.

We may know write the signal-to-glitch Bayes Factor in this approximation:

$$B_{SG} = \frac{\Lambda_S \times \mathrm{O}_S}{\Lambda_G \times \mathrm{O}_G}, \tag{4}$$

Again limiting ourselves to the case of potential Burst detection candidates (SNR ≥ 10), we expect strong peaks in the likelihood at the best fit waveform parameters. In `BayesWave`, the signal model is described as a geocenter collection of wavelets projected on to each detector in the network, where the glitch model allows wavelets at each detector that are independent of wavelets in other detectors. This means that, as seen in the detectors, *the*

tyson@gmail.com

allowed waveforms in the signal model are a subset of the waveforms allowed in the glitch model. Since the glitch model has more freedom to fit the data than the signal model, it must be true that $(\Lambda_G \geq \Lambda_S)$, which implies:

$$B_{SG} \leq \frac{\mathrm{O}_S}{\mathrm{O}_G}. \tag{5}$$

For glitches that are not consistent with astrophysical signals (i.e. look very different in different detectors), $\Lambda_G \gg \Lambda_S$, and so the Bayes Factor will be much smaller than the limit in equation 5. On the other hand, data roughly consistent with some astrophysical waveform will give $B_{SG} \simeq \mathrm{O}_S/\mathrm{O}_G$, regardless of whether it was actually caused by an astrophysical or local source. This means that, for strong signal candidates, `Bayeswave` uses primarily the Occam Factor to distinguish signals that are very likely astrophysical from those that are consistent with detector artifacts. For this reason, it is possible to develop an expectation for the Bayes Factor associated with certain types of signals by estimating the Occam Factor alone.

1.2. *Estimating the Occam Factor*

We will make an estimate for the Occam Factors associated with signals observed by a two detector network. For the moment, we will limit ourselves to the case of a single wavelet in each detector. Then, the glitch model has five parameters in each detector, for a total of ten parameters. Using subscripts to denote detectors 1 and 2, we can write these as $(A_1, f_{01}, Q_1, t_{01}, \phi_{01}, A_2, f_{02}, Q_2, t_{02}, \phi_{02})$. We choose Detector 1 to be whichever detector has a higher SNR signal.

For the signal model, five parameters are used to construct the wavelet at the geocenter: $(A_1, f_{01}, Q_1, t_{01}, \phi_{01})$. Since we chose Dectector 1 to be where we see the higher SNR waveform, we can expect the posterior volume for these parameters in the signal model to be similar to the posterior volume of the corresponding parameters in the glitch model. Then, neglecting correlations between parameters, we can expect terms related to these parameters to appear in both the numerator and denominator when calculating the ratio of Occam Factors in Equation 5, and so should approximately cancel.

The signal model also uses four extrinsic parameters, $(\theta, \phi, \epsilon, \psi)$, which allow some freedom to choose how the geocenter waveform will be projected onto Detector 2. In particular, for the signal model, A_2, t_{02}, and ϕ_{02} are functions of the extrinsic parameters. Given that both the signal and glitch models have some freedom choose these parameters, it seems a reasonable approximation that the associated accessible volume fractions cancel from Equation 5. This may not be a good approximation in all cases, but should be reasonable for an order-of-magnitude estimate.

On the other hand, the signal model gives no freedom in choosing how the central frequency and quality factor of the geocenter waveform will be projected onto Detector 2. So, the glitch model has two parameters with no counterpart in the signal model: f_{02} and Q_2. In cases of strong Burst candidate detections, we expect this extra freedom in the glitch model waveform to dominate the calculation of the Bayes Factor. For this special case of strong candidates in a two detector network, we apply this approximation to Equation 5

$$\frac{\mathrm{O}_S}{\mathrm{O}_G} \simeq \frac{1}{p(f_{02}|G)\sigma_{f_{02}} p(Q_2|G)\sigma_{Q_2}} \tag{6}$$

where $\sigma_{f_{02}}$ and σ_{Q_2} represent the posterior uncertainty on the central frequency and quality factor of the glitch in the Detector 2. Since `BayesWave` uses flat priors on these parameters, $p(f_{02}|G) = 1/V_{f_{02}}$, where $V_{f_{02}}$ is the allowed range of f_{02}. Applying the same convention to Q_2, we write:

$$\frac{\mathrm{O}_S}{\mathrm{O}_G} \simeq \frac{V_{f_{02}} V_{Q_2}}{\sigma_{f_{02}} \sigma_{Q_2}} \tag{7}$$

We see that the ratio of Occam Factors, then, is the fraction of the f-Q plane where the signal in Detector 2 is consistent with the signal in Detector 1. This gives a simple interpertation of this quantity. Given some glitch in Detector 1, and a coincident glitch in Detector 2, the Occam Factor quantifies the chances that the two glitches will have matching parameters consistent with an astrophysical signal. The calculation of this Occam Factor is relatively robust against different glitch populations - the main assumption is that glitches are created with random central frequency and quality factor.

The qauntities $\sigma_{f_{02}}$ and σ_{Q_2} represent the posterior uncertainty on the corresponding parameters. In the high SNR limit, we may approximate these values using the diagonal terms of the Fisher Matrix for a Morlet-Gabor waveform. These are given without proof in Shourov Chatterji's Ph.D. dissertation[1] (See Eqn 3.21):

$$\sigma_{f0} = \frac{2f_0}{A\sqrt{2+Q^2}} \tag{8}$$

$$\sigma_Q = \frac{\sqrt{2}Q}{A} \tag{9}$$

$$\frac{\mathrm{O}_S}{\mathrm{O}_G} \simeq \frac{A^2 V_Q V_{f_0}}{2\sqrt{2} f_0} \tag{10}$$

where A, f_0, and Q represent the parameters of the waveform in the second detector.

As an example, we used `BayesWave` to recover injections of a single wavelet with $f_0 = 153Hz$ and $Q = 9$. We added these signals to real data from the LIGO H1 and L1 instruments collected during 2010. We used a frequency range up to 512 Hz and Q = [0, 40]. For such signals, we can calculate the expected Occam Factor using Equations (7) - (10):

$$\ln \frac{\mathrm{O}_S}{\mathrm{O}_G} \simeq \ln 47 + 2 \ln A_2 \tag{11}$$

This model is compared with the results using `BayesWave` in Figure 1.

2. WHERE'S THE MAGIC?

Notice in Figure 1 and Equation (7) that the scaling of $\ln B_{SG}$ with SNR is relatively weak at high values - the detection statistic has similar values at SNR 10 - 40. This is a robust feature of the model, and will be

[1] http://hdl.handle.net/1721.1/34388

true both for glitches and for astrophysical signals. Since even very high SNR signals give only a modest increase in the expected value of the Bayes Factor, we might ask what could be done to construct a signal with a very high Bayes Factor. The answer is to create a signal with a more complex structure that requires more than one wavelet to model. Equation 7 may be extended to multiple wavelets by taking a product of the Occum Factor for each wavelet. In order to see how this scales, we imagine each of N wavelets has a typical frequency, quality factor, and amplitude of f, Q, and A, respectively:

$$\frac{\mathrm{O}_S}{\mathrm{O}_G} \propto \left[\frac{A^2 V_Q V_{f_0}}{2\sqrt{2} f_0}\right]^N \tag{12}$$

Now it is clear that the detection statistic, $\ln B_{SG}$, scales linearly with the number of wavelets, but has only a logarithmic dependence on other parameters, including SNR. In fact, Equation (12) is conservative, since parameters like t_{02} will also increase the Occam Factor ratio. So, *in order to achieve a gold plated detection*, `BayesWave` *requires a signal with complex time-frequency structure, rather than a very loud signal.* This is a significant difference from other Burst pipelines. In the next section, we will show that even two wavelets is enough for a high confidence detection in many cases, while a single wavelet will almost always be plausibly explained by coincident glitches. The fact that other pipelines fail to quantify the importance of structure in forming their detection statistic is the main reason that `BayesWave` can identify high confidence detections in the presence of loud glitches, while other pipelines cannot.

In order to demonstrate this effect, we injected band-limited white noise burst signals into archived LIGO data. These signals have a time-frequency structure that requires many wavelets to reconstruct. However, at low SNR, only a fraction of the signal power can be recovered. The result is that higher SNR signals require a larger number of wavelets, and so we expect that the Occum Factor will increase with increasing SNR. This behavior is clearly seen in Figure 2.

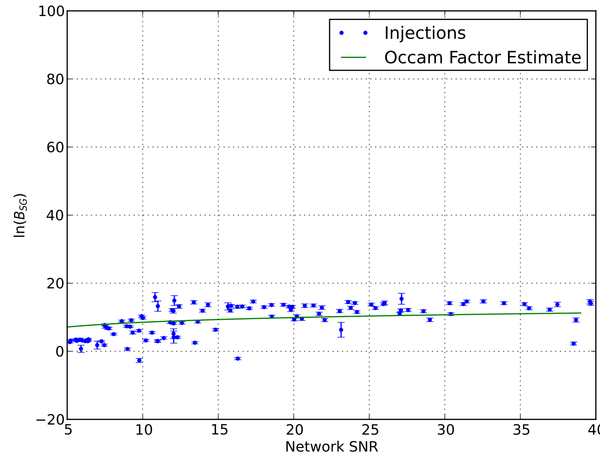

Figure 1. The signal-to-glitch Bayes Factor, as computed by BayesWave, for simulated single wavelet signals. Also shown is the estimate for the Occam Factor presented in this work. At high SNR, the behavior of the Bayes Factor is broadly similar to expectations for the simple model.

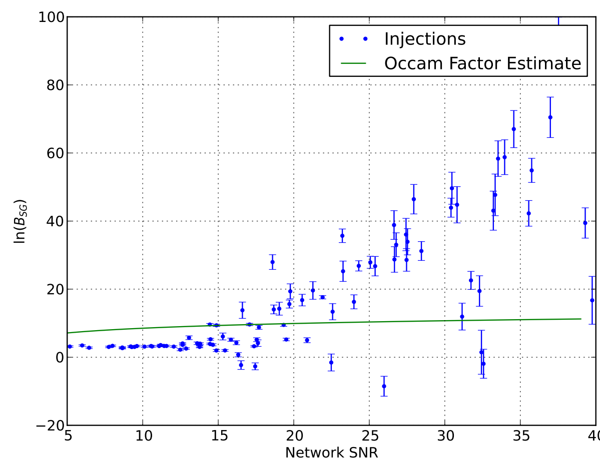

Figure 2. Simulated white noise bursts, which require many wavelets to reconstruct at high SNR. For comparison, the single wavelet model of the Occum Factor is also shown. Notice that the signals with complex time-frequency structure typically give Bayes Factors which are much higher than the single wavelet estimate.

3. BACKGROUND ESTIMATION

Equation 5 shows that, for strong Burst candidates, the Occam factor provides a estimate for the Bayes Factor. The fact that this estimate captures some features of `BayesWave` is demonstrated in Figures 1 and 2. Now, we see what implications this has for a population of glitches.

For background estimation, the key feature of Equation 7 is that the Occam Factor represents the fraction of the glitch model parameter space which is consistent with the signal model. As a heuristic picture, imagine that the waveform in Detector 1 is well represented by a single wavelet with parameters f_{01} and Q_1. Now imagine that there is a coincident glitch in Detector 2. If the signal is astrophysical in nature, the waveform in Detector 2 must have parameters Q_2 and f_{02} that are consistent with f_{01} and Q_{01}, within the measurement uncertainties $\sigma_{f_{02}}$ and σ_{Q_2}. On the other hand, if the data represent coincident glitches, then *a priori* there is no reason for the glitch in Detector 2 to match the parameters in Detector 1. Instead, the wavelet in Detector 2 is chosen at random, and the Occum Factor quantifies the odds of selecting parameters that match Detector 1. It is as if we are blindly throwing darts at the f-Q plane, and hoping to hit a target of size $\sigma_{f_{02}}$ by σ_{Q_2}. "Missing the target" means the glitch in Detector 2 is not consistent with the signal model, and so the signal model likelihood will be lower than the glitch model likelihood, thus reducing the value of the Bayes Factor well below the limit in Equation 5, and so rejecting the event as a glitch.

This interpertation has an important implication for the background rate. Let's assume, in a given two detector data set, there are N_{gl} coincident glitches. We can use the interpertation of the Occam Factor to place a limit on the expectation value of the detection statistic for the loudest false alarm in the data set. Glitches with low values of the Occam Factor will also have low values of the detection statistic, as seen in Equation 5. However, glitches with a high value of the Occam Factor have a high probability of "missing the target", and so will more often acquire low likelihood and whence low Bayes Factor. Then, in order to maximize the expectation value of the Bayes Factor of the loudest background event, the "worst case scenario" would be if *all* of the glitches had an Occam Factor at the maximum value, such that one glitch on average would be consistent with a true gravitational wave. In this limiting case, the expectation value for the Bayes Factor of the loudest event would be equal to the number of coincident glitches.

$$\langle B_{SG}^{\rm MAX} \rangle \leq N_{gl} \tag{13}$$

Notice this limit is not a statement about the population of glitches, beyond the assumption that the parameters f_0 and Q are chosen at random for glitches in each detector. If glitches exist in the population which have an Occam Factor greater than the limit in 13, the expectation is that none of them will have parameters consistent with the signal model. It is also important to point out this is a conservative estimate. Most glitches are at low SNR in any realistic glitch population, and so low values of the Occam Factor will likely be much more common than high values. Notice, also, that this limit could be applied equally well to a population of truly coincident glitches or a to a population of glitches found to be coincident in a number of timeslides. In this way, it could be used to set a "five-sigma" detection threshold for any data set.

To see how this could work in practice, we must set a single parameter based on our knowledge of the detectors. This is the rate of coincident glitches, R_{gl}. Fortunately, this rate is carefully studied within LIGO. The single detector glitch rate is known to typically have values between 1 and 0.1 Hz (See S6 detchar paper, LIGO-P1000142). The light travel time between LIGO detectors is 10 ms, leading to a coincident glitch rate of $R_{gl} \sim 1$ Hz$\times$1 Hz$\times$0.01 s = 0.01Hz.

As an example, imagine an early run of the Advanced LIGO detectors. Such a run might last for around three months. A so-called "three sigma" detection requires an event louder than the expected background in ~ 300 time slides. This means the background data would have 75 years of livetime, and $N_{gl} \sim 2 \times 10^7$. Equation 13 then suggests that events with $\ln B_{SG}^{\rm MAX} \simeq 16$ would be marginally detectable. This is similar to the Bayes Factors for the single wavelet injections in Figure 1. The conclusion, then, is that with a two detector network, single wavelet signals are marginally detectable with `BayesWave` at any reasonable SNR value. This is very similar to the performance seen in past Burst searches with a two detector network.

On the other hand, what is required for a "five-sigma" detection? For this case, we seek a p-value of less than 3×10^{-7}, and so demand our event be louder than the loudest event in 3×10^6 time slides. For our hypothetical three month observing window, this leads to $N_{gl} \sim 2 \times 10^{11}$, for an expected loudest event $\ln B_{SG}^{MAX} \sim 26$. We have already seen that single wavelet events can not reach this level at any reasonable SNR. However, applying the scaling law in Equation (12), we find that such a "gold-plated" detection could be achieved at reasonable SNR with as few as two or three wavelets. This is an important feature of the `BayesWave` pipeline that is distinct from other Burst detection schemes: *gold-plated detections are possible even in the presence of a significant glitch population.*

An equivalent interpertation of Equation (13) is to use the glitch rate to set a Bayesian *prior* on the glitch model, and then use the *posterior* to express our confidence in a signal candidate. A simliar approach was used by [James Clark thesis, I think]. The glitch rate expresses our prior belief in the chances that a stretch of gravitational wave data will contain a glitch. While we could use knowledge of astrophysics to set the prior on the signal model, a reasonable guess at this stage is that any observation period contains ~ 1 true gravitational wave signal. After all, if the rate was higher, we would have already seen one. If we really believed the rate was much lower, then there is no reason to analyze the data. We can the calculate the posterior as

$$\frac{p(S|\mathbf{s})}{p(G|\mathbf{s})} = \frac{1}{R_{gl}L} B_{SG} \tag{14}$$

where L is the livetime of the network. In the limit of strong candidates, equation 13 will make this description equivalent to the frequentist scheme described above.

In order to validate Equation 13, it would be natural to study the performance of `BayesWave` on a number of background time-slides. In general, it would be prudent to do this on future data sets. In fact, in most cases we expect B_{SG}^{MAX} to be less than the limit, so performing time-slide analysis may actually allow a lower detection threshold. Here, as a demonstration, we show the performance of `BayesWave` on a collection of the loudest one hundered `coherentWaveBurst` (cWB) triggers from the two detector S6D analysis. This can be seen in Figure 3. The S6D cWB background study used 300 timeslides on 27 days of livetime, for a total of 22 years of livetime in the background set. The loudest event in this background set has a `BayesWave` detection statistic of $\ln B_{SG} = 12$. The conservative limit in Equation (13) predicts this value will be less than 15.7.

For comparison, we also run a simplistic simulation, where we draw simulated glitches from ad hoc populations, compute the Occum Factor using Equation (12), and then reject glitches with with probablity O_G/O_S. We use a population of single wavelet glitches with SNRs chosen randomly between 1 and 50, SNRs randomly chosen between 1 and 1000, and a population where half the glitches have one wavelet and half have two wavelets, with SNRs randomly chosen between 1 and 50. The results are that all of these populations have a loudest event with $11 \leq \ln B_{SG} \leq 13$. This highlights that the limit for the loudest event in Equation (13) is relatively robust to different glitch populations.

TEST (Sutton 2009) (Mackay 2003)

REFERENCES

P. J. Sutton, (2009), arXiv:arXiv:0905.4089 [physics.data-an].

D. J. C. Mackay, *Information Theory, Inference and Learning Algorithms, by David J. C. MacKay, pp. 640. ISBN 0521642981. Cambridge, UK: Cambridge University Press, October 2003.* (2003)

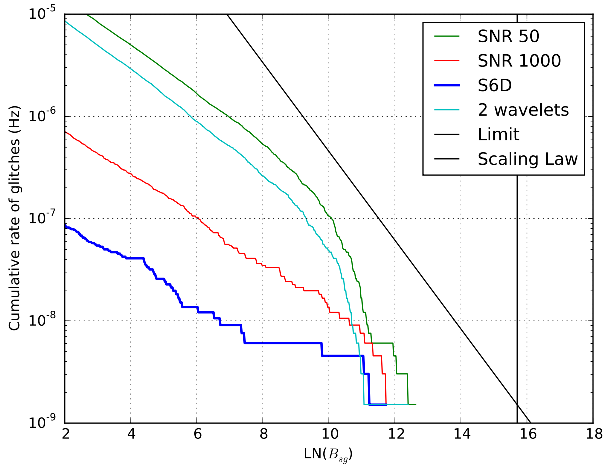

Figure 3. The resutls of time-slides studie using BayesWave and the S6D data set are shown as the blue trace. For comparison, a simplistic simulation using several hypothetical glitch populations is shown in other colors. The black, vertical line shows the expected value for the loudest event using the limit in Equation 13. In all cases, both real and simulated, the loudest event is seen to be between 10 and 13 in the detection statistic, all below the expected limit of 15.7. Since this represents 300 timeslides of the data set, we see that essentially all of the sine-gaussian injections above network SNR 10 were marginally detected, as were most WNB above network SNR 20. Moreover, most WNBs above network SNR 25 were "gold-plated" detections. These resutls are consistent with expectations described in this work.