

August 2022

Searching for Gravitational Wave Associations with High-Energy Astrophysical Transients

Brandon Joseph Piotrkowski
University of Wisconsin-Milwaukee

Follow this and additional works at: <https://dc.uwm.edu/etd>



Part of the [Astrophysics and Astronomy Commons](#), and the [Physics Commons](#)

Recommended Citation

Piotrkowski, Brandon Joseph, "Searching for Gravitational Wave Associations with High-Energy Astrophysical Transients" (2022). *Theses and Dissertations*. 3060.
<https://dc.uwm.edu/etd/3060>

This Dissertation is brought to you for free and open access by UWM Digital Commons. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of UWM Digital Commons. For more information, please contact scholarlycommunicationteam-group@uwm.edu.

SEARCHING FOR GRAVITATIONAL WAVE
ASSOCIATIONS WITH HIGH-ENERGY
ASTROPHYSICAL TRANSIENTS

by

Brandon Piotrkowski

A Dissertation Submitted in
Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
in Physics

at

The University of Wisconsin-Milwaukee
August 2022

ABSTRACT

SEARCHING FOR GRAVITATIONAL WAVE ASSOCIATIONS WITH HIGH-ENERGY ASTROPHYSICAL TRANSIENTS

by

Brandon Piotrkowski

The University of Wisconsin-Milwaukee, 2022
Under the Supervision of Professor Patrick Brady

Gravitational waves (GW) have become an invaluable tool in modern astronomy, especially in conjunction with other astronomical observations. GWs are created in highly dynamical systems such as compact binary coalescences (CBC) which are comprised of black holes and/or neutron stars. The Laser Interferometer Gravitational Wave Observatory (LIGO), Virgo Observatory, and KAGRA have now collectively identified almost a hundred of these events. GWs have also been predicted to come from core collapse supernovae. Both of these types of systems have been shown to produce other detectable transients, such as gamma-ray bursts (GRB) and neutrino bursts. Observations of the same astrophysical system with multiple messengers are proven to be invaluable to understanding these systems and constraining crucial physical constants. Therefore, efforts to find more multi-messenger events are increasingly important as detection sensitivities and expected rates increase. It is especially important to identify multi-messenger events as soon as possible in order to alert other astronomers and to facilitate additional follow-up. In this dissertation, I describe my work in the LIGO-Virgo-KAGRA (LVK) Collaboration developing both infrastructure and techniques for multi-messenger searches involving GWs.

© Copyright by Brandon Piotrkowski, 2022
All Rights Reserved

To my friends, family,
and my partner.

TABLE OF CONTENTS

1	Introduction	1
1.1	The Aftermath of GW170817	2
1.2	Format of Dissertation	3
2	Messengers in Modern Astronomy	6
2.1	Gravitational Waves	6
2.1.1	Compact Binary Coalescences	8
2.1.2	Other Transients	9
2.2	Gamma-Ray Bursts	10
2.3	Neutrino Bursts	11
2.4	Fast Radio Bursts	13
3	RAVEN: The Low-Latency Gravitational Wave Focused Multi-messenger Pipeline	15
3.1	Motivation	15
3.2	Design	16
3.3	Joint Significance	19
3.3.1	Untargeted GW-GRB Analysis	20
3.3.2	Targeted GW-GRB Analysis	22
3.3.3	GW-GRB Publishing Conditions	23
3.3.4	GW Burst-SNEWS	25
3.4	FAR Simulations	25
3.5	Discussion	29
4	A Joint Ranking Statistic for Multi-messenger Astronomical Searches with Gravitational Waves	32
4.1	Motivation	32
4.2	Derivation	34
4.2.1	General Multi-messenger Candidate	35
4.2.2	GW-GRB Candidate	38
4.2.3	Coincidence Bayes Factors	40
4.3	Examples	41
4.3.1	GW170817-GRB 170817A: Using 3D Sky Maps	42
4.3.2	GW170817-GRB 170817A: Using Posterior Samples	44
4.3.3	GBM-GW150914	48
4.3.4	Simulation	49

4.4	Conclusion	51
5	Testing Ranking Statistics for Joint Gravitational Wave and Gamma-Ray Burst Searches	53
5.1	Motivation	53
5.2	Ranking Statistics	54
5.3	Simulation	55
5.4	Results	58
5.5	Discussion	60
6	Contributions to the Third LIGO-Virgo-KAGRA Operating Run	62
6.1	Joint FAR Derivation	62
6.2	Correcting Skewness in the Joint FAR for Offline Searches	66
6.3	GW-FRB Search w/ CHIME	67
6.3.1	Cross-validation with PyGRB and X-pipeline	67
6.3.2	Using Extended Search Windows	68
6.4	GWCosmo	70
6.4.1	Partial Loading of Galaxy Catalogues	70
7	Contributions to the Fourth LIGO-Virgo-KAGRA Operating Run	78
7.1	Additional Improvements to RAVEN	78
7.1.1	Multi-ordered (MOC) Sky Maps	78
7.1.2	Constant Testing and Validation	80
7.1.3	Other Improvements	81
7.2	Sub-threshold Search w/ <i>Swift</i>	82
7.3	Offline PyCBC Sky Map Generation	85
7.3.1	Creating a Workflow Using DAGs	85
7.4	Speed of Gravity in Low-latency	86
7.4.1	Probability of Association	87
8	Conclusion	89
8.1	Future Work	89
8.2	Looking Forward	90
	Bibliography	91
	Appendix	99
	A.1 Coincidence Info in LVK GCN VOEvents	99
	A.2 Coincidence Info in LVK GCN Circulars	99
	Curriculum Vitae	103

LIST OF FIGURES

1.1	Sky Localizations of GW170817 and GRB 170817A	2
1.2	Lightcurves of GRB 170817A and Spectrogram of GW170817	5
2.1	Simple Michelson Interferometer	7
3.1	RAVEN Pipeline Logic Flowchart	18
3.2	Fermi/GBM Joint False Alarm Rates	26
3.3	Fermi/GBM Sky Map Overlap Values	27
3.4	<i>Swift</i> /BAT Joint False Alarm Rates	28
3.5	<i>Swift</i> /BAT Sky Map Overlap Values	29
4.1	ROC Curves from Preliminary Joint Simulation	50
5.1	ROC Curves from Threshold Joint Simulation	59
5.2	ROC Curves from Sub-threshold Joint Simulation	60
6.1	Simulation of Joint FAR using Various Terms	64
6.2	Simulation of Untargeted RAVEN Joint FAR Method	72
6.3	Comparison of Use of Sky Map Overlap in Joint FAR	73
6.4	Corrected Joint FAR with Fermi/GBM	74
6.5	Corrected Joint FAR with <i>Swift</i> /BAT	75
6.6	Results of Joint FAR from FRB Search Using Standard Time Windows	76
6.7	Results of Joint FAR from FRB Search Using Extended Time Windows	77
7.1	Example of Adaptive Mesh Refinement for GW170817	80
7.2	Workflow of the Targeted Search with <i>Swift</i> /BAT	83
7.3	O2 Joint Targeted Search with <i>Swift</i> /BAT-GUANO	84
7.4	Workflow of Parallelized PyCBC Sky Map Creation	86
A.1	RAVEN VOEvent Additional Fields	100

LIST OF TABLES

2.1	Properties of GRB Experiments	12
3.1	RAVEN Coincident Time Windows	21
4.1	Overlap Integral Values for GW170817-GRB 170817A Using HEALPix Sky Maps	46
4.2	Overlap Integral Values for GW170817-GRB 170817A Using Posterior Samples	46
4.3	Latency in Computing Overlap Integral Values for GW170817-GRB 170817A	47
4.4	Overlap Integral Values for GBM-GW150914 Using Posterior Samples . . .	49
6.1	Results from Extended Window CBC-FRB Search	69
7.1	Latency in Computing Overlap Integrals Using Both MOC and Flattened Sky Maps for GW170817-GRB 170817A	81

ACKNOWLEDGEMENTS

Firstly, I'd like to thank my advisor Patrick Brady for all the direction, advice, and opportunities. This especially includes the period of time switching from numerical relativity to multi-messenger astronomy, and even now moving to the next stage of my career. I thank Min-a Cho for her patience and availability to train me in the various operations of RAVEN. I greatly appreciate Shaon Ghosh's help to get me involved in LIGO, as well all the instruction and attention Deep Chatterjee has devoted to make sure I was able to begin working. Deep, I thank you for your insight and assistance in designing the RAVEN alert pipeline, even up to the end of my grad school career. I thank Ignacio Magaña-Hernandez for bouncing off ideas of projects we'll never find the time to finish. I am grateful for my friendship and all the support from Caitlin Rose, as well as her now husband Christian Parsons. And I still fondly recall late nights working on homework assignments with classmates. Thank you Siddharth Mohite, Logan Prust, and many many other graduates students for your friendships during this time. I look forward to seeing how Naresh will improve upon my work with RAVEN and all of the junior graduate students as they walk down a similar path as we have before them.

I thank my family for their support: my mom Carol for believing in me despite the difficulties, my brother Taylor for his friendship, and my brother Connor for his companionship and many intellectual discussions. I still recall the terror of missing any homework assignments instilled by my late sixth grade teach Sister Mary Jude, who sent me down the path of academic excellence. I am grateful for my grandparents, Bill and Diane, who have demonstrated incredible generosity in all facets of their time and money. I thank my friend Chris Simaz for encouraging me throughout the years, and especially when suggesting physics when I wasn't sure what I wanted to do in undergrad. And not in the least, I thank my partner and my ashegham Elaheh Kheirandish, who has also been my light during difficult times.

This dissertation has been designated the document ID P2200173 by the LIGO DCC.

This work was supported by NSF awards PHY-1607585 and PHY-1912649. I am grateful for computational resources provided by the LIGO Laboratory and the support from the National Science Foundation Grants PHY-0757058 and PHY-0823459, and those provided by the Leonard E Parker Center for Gravitation, Cosmology and Astrophysics at the University of Wisconsin-Milwaukee.

The only thing worth writing about is the human heart in
conflict with itself.

William Faulkner

Ask, and it will be given to you; seek, and you will find;
knock, and it will be opened to you.

Matthew 7:7 NASB

Chapter 1

Introduction

Multi-messenger events are an important aspect of modern astronomy, with the ability to give more insight into a progenitor event than an individual detection could. This is because different astronomical messengers typically carry information about distinct physical processes, and therefore the combination gives a more complete picture of the underlying system. This is similar to using different senses to perceive the world, such as sight and hearing; each messenger of information has its advantages, and combining them leads to a more comprehensive view of the world. The first example of multi-messenger astronomy was the joint optical and neutrino burst detection of supernova SN 1987A, which helped confirm supernova models and provided upper limits on both neutrino flavors and mass ([Hirata et al., 1987](#); [Arnett & Rosner, 1987](#); [Lunardini & Smirnov, 2004](#)).

Astronomers have also long expected gravitational waves (GW) from binary neutron star (BNS) systems to be accompanied by prompt short gamma-ray bursts (GRB) ([Blinnikov et al., 1984](#); [Paczynski, 1986](#); [Eichler et al., 1989](#)). Multi-messenger searches were created in anticipation of such an event, looking at real-time data to inform other astronomers of discoveries (online) or by analyzing data using more sensitive (and generally more computationally expensive; offline) methods ([Aasi et al., 2014](#); [Singer et al., 2014](#); [Abbott et al., 2017d](#); [Urban, 2016](#); [Cho, 2019](#)). The association between GWs and GRBs

was confirmed with the confident joint detection of GW170817 (Abbott et al., 2017b,a) and GRB 170817A (Goldstein et al., 2017). An optical kilonova transient AT 2017gfo was also found approximately 11 hours afterward, localizing the progenitor to the galaxy NGC 4993 (Valenti et al., 2017).¹ Visualization on the sky localizations and data can be found in figures 1.1 and 1.2 respectively.

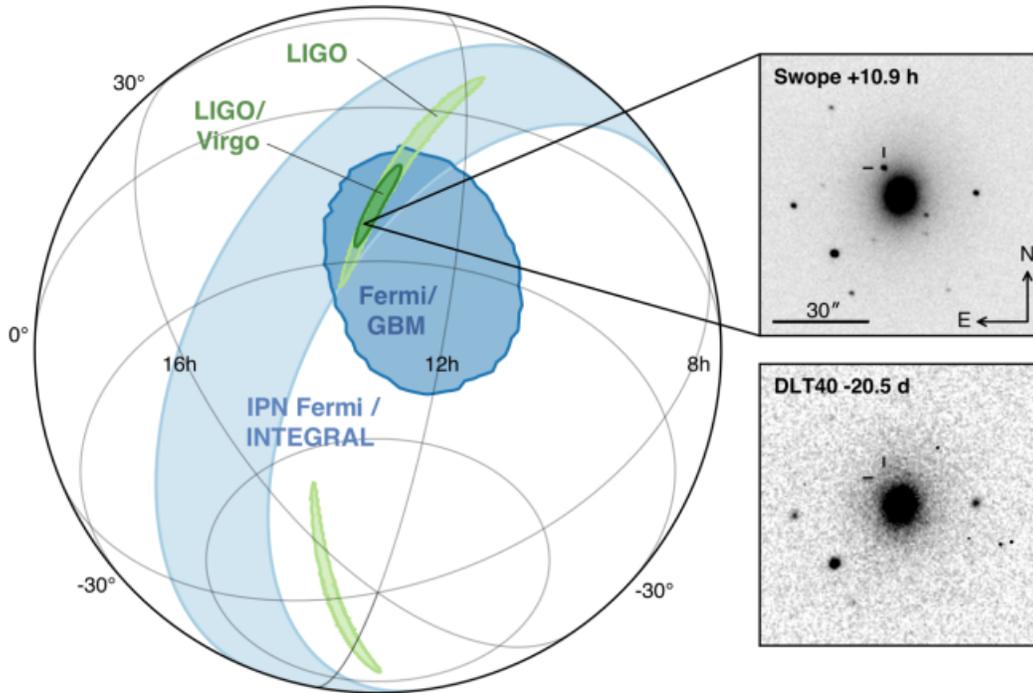


Figure 1.1: Sky localizations of GW170817 detected by the LIGO-Virgo Collaboration, both with and without Virgo, and GRB 170817A detected by Fermi. This figure was reproduced from Abbott et al. (2017c). Also included is the localization due to the detection timing difference between Fermi and INTEGRAL (IPN). On the right we can see the optical transient appearing within 11 hours after the merger while not being present 20.5 days before.

1.1 The Aftermath of GW170817

The number of direct astrophysical measurements and constraints as a result of the multi-messenger detections of GW170817 has been staggering. A strong constraint was made on

¹<https://gcn.gsfc.nasa.gov/other/G298048.gcn3>

the speed of gravity to be nearly identical to the speed of light (Abbott et al., 2017a) due to a delay of 1.7 s between the GW and GRB detections (seen in figure 1.2), even with both traveling a distance of roughly 40 Mpc. The nuclear equation of state, which should determine the mass-radius relation in neutron stars, was measured using just GW170817 (Abbott et al., 2018) but was further improved by including the other detections (Coughlin et al., 2019). The simultaneous measurement of distance by the GW and redshift of its host galaxy led to a constraint of the Hubble constant (Collaboration et al., 2017). Also, the subluminal detection of GRB170817A along with the measurement of inclination by the GW has given support for wider short GRB jet models than expected (Howell et al., 2019; Farah et al., 2020). Finding more coincidences similar to GW170817 or others we haven't seen yet (described in Chapter 2) will lead to insights into the universe that would not be possible otherwise.

1.2 Format of Dissertation

This dissertation details my work in searching for and studying coincidences between gravitational waves and other high energy astrophysical transients. The text is all my own and the work presented, namely the RAVEN alert pipeline, derivation of the Odds ratio ranking statistic, development of a multi-messenger pipeline simulation, and the numerous algorithms, are my own.

This thesis is organized in the following manner. In Chapter 2, background regarding the various messengers of interest are discussed, including gravitational waves, neutrino bursts, gamma-ray bursts, and fast radio bursts. In Chapter 3, my work in expanding and improving the capabilities of the low-latency multi-messenger search RAVEN is detailed. In Chapter 4, an association statistic to determine whether two candidates are correlated is derived and then tested with some known coincidence candidates. In Chapter 5, a simulation of a joint search is described with the purpose of testing the performance of

various ranking statistics. In Chapters 6 and 7, I detail my contributions to LVK operations and analysis during the O3 and O4 runs respectively. Finally in Chapter 8, I conclude and describe future work.

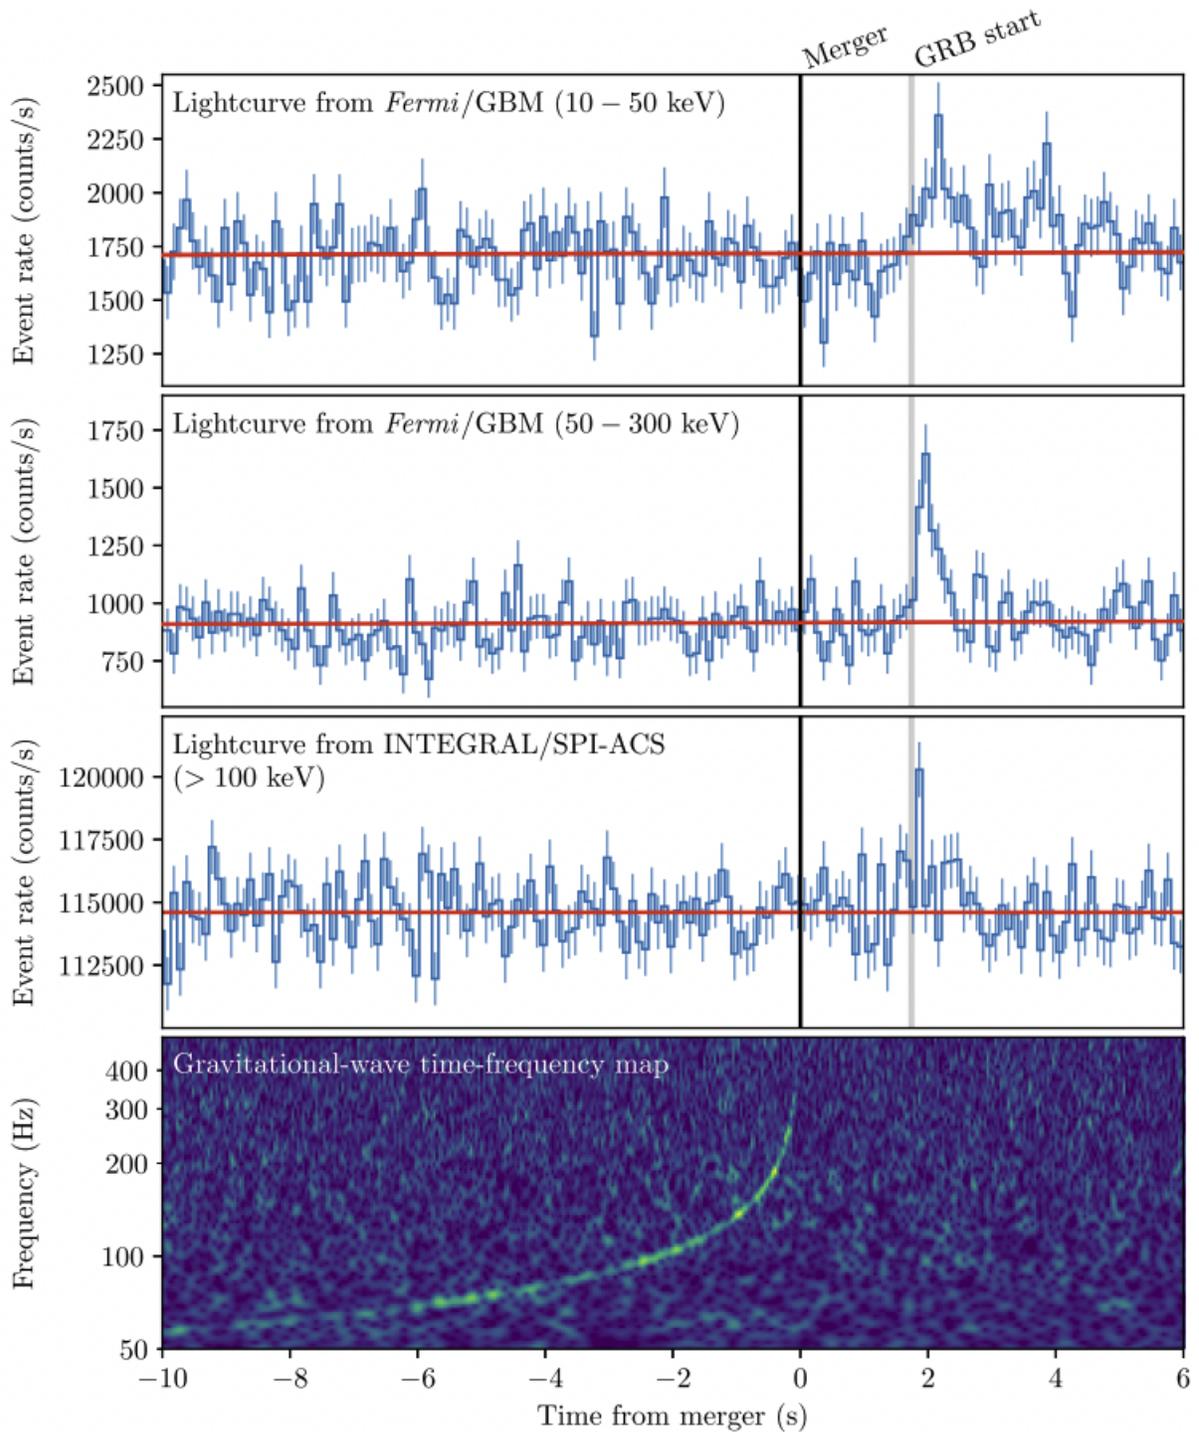


Figure 1.2: Lightcurves of GRB 170817A from both *Fermi*/GBM and INTEGRAL, as well as a spectrogram of GW170817 where the color represents relative amplitude in the strain. This figure was reproduced from [Abbott et al. \(2017c\)](#). We can see the “chirp” signal in GW170817 terminating just 1.7 seconds before the start of the GRB signal.

Chapter 2

Messengers in Modern Astronomy

In this chapter we will go over all the astronomical messengers involved in the various coincidence searches detailed in this thesis. Note there are many other transients, such as active galactic nuclei (AGN) flares and cosmic rays, that are relevant to other progenitor systems.

2.1 Gravitational Waves

Gravitational waves (GW) are perturbations in space-time that can travel and transfer energy (Creighton & Anderson, 2011). They cannot be removed or explained away by coordinate transformations since they have non-zero Riemann tensor components, representing real curvature in space-time. GWs are transverse and have two distinct polarizations, usually indicated by plus (+) and cross (\times). At lowest order, GWs are produced whenever a system has a non-zero second time derivative of its quadrupole mass-energy tensor. A classical example of this is a rotating dumbbell, a rough model of the compact binary systems. We will discuss compact binary systems later in section 2.1.1.

Gravitational waves provide a unique insight into the universe compared to other astrophysical messengers. Some progenitor systems are thought not to produce any other types of messengers, such as binary black hole (BBH) mergers, potentially making GWs

the sole observational tool for these systems (Zwart & McMillan, 1999). Due to their low interactability with matter, GWs travel through the universe largely unimpeded before detection, with the exception of gravitational lensing (Takahashi & Nakamura, 2003). Since gravitational wave sources are generally well modeled by general relativity, the underlying parameters of these systems can be well and independently measured. An example would be the measurement of distance to a system, which is independent from the cosmological distance ladder. This is an advantageous trait for multi-messenger studies and makes gravitational wave detections a good compliment to electromagnetic observations.

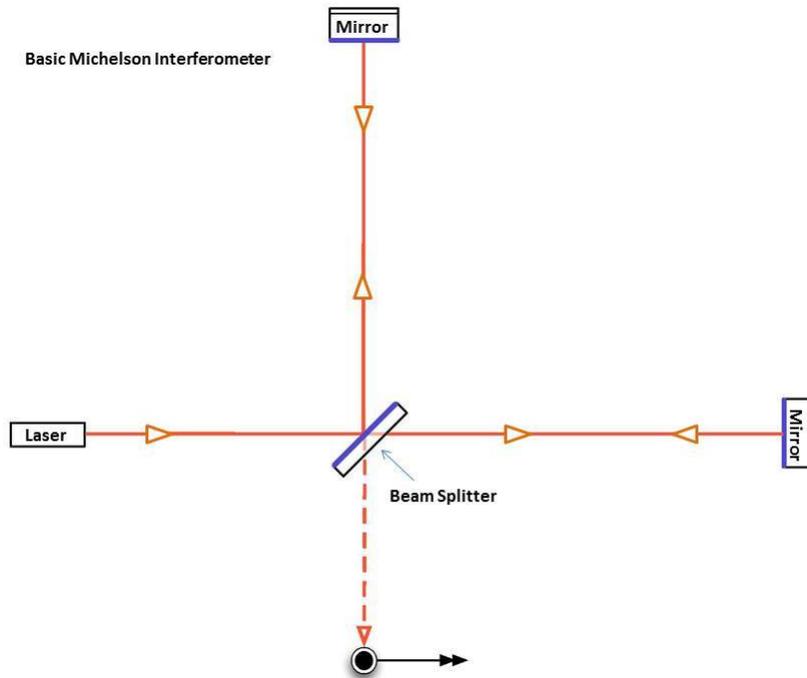


Figure 2.1: A schematic of a simple Michelson interferometer, the type used by LIGO Abramovici et al. (1992). Monochromatic coherent light is sent simultaneously down each of its arms so that changes in detector armlength create a detectable phase differences when recombined.

Since gravitational waves interact little with matter and the resulting strain caused by astrophysical events is so small, detection is a difficult task. The strain caused by gravitational waves is very small and impractical to measure using everyday objects such as rulers. Therefore, all modern gravitational wave experiments, such as the LIGO-Virgo-KAGRA

(LVK) collaboration ([Abramovici et al., 1992](#); [Abbott et al., 2020a](#)), Laser Interferometer Space Antenna (LISA) ([Amaro-Seoane et al., 2017](#)), and North American Nanohertz Observatory for Gravitational Waves (NANOGrav) ([McLaughlin, 2013](#)), use light to measure the differences in distance or timing since the speed of light is invariant. We can see this with example of an interferometer used by LIGO in [figure 2.1](#), where different distances in the length of the arms due to GWs result in timing differences (and hence phase differences for monochromatic coherent light) when the light is recombined. The detectors in the LVK are sensitive to relatively high frequency (> 10 Hz) GW transients and are of primary interest for this thesis ([Abramovici et al., 1992](#); [Acernese et al., 2014](#); [Akutsu et al., 2018](#)).

2.1.1 Compact Binary Coalescences

Compact binary coalescences (CBC) have been the archetypal detectable gravitational wave signals over the past few years ([Abbott et al., 2021a](#)). These signals occur when two orbiting compact objects come close enough that their orbital mechanics become dominated by the radiation of gravitational waves and then finally merge ([Mandel & O’Shaughnessy, 2010](#)). Although the methods of bringing compact objects close enough is still being studied, this could occur if two stars in a binary system exchange mass, perhaps via a common envelope, where drag forces could remove much of the angular momentum and reduce the separation distance ([Tauris et al., 2017](#)).

Once the orbital evolution is dominated by gravitational wave radiation, the remaining phase before merger is thought to be largely deterministic and be well explained by general relativity. For example, by modeling two compact objects as points we can derive the frequency (f) evolution as ([Creighton & Anderson, 2011](#))

$$\frac{df}{dt} \propto \mathcal{M}_c^{5/3} f^{11/3} \tag{2.1}$$

where \mathcal{M}_c is the chirp mass, defined as $\mathcal{M}_c = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$, where m_1 and m_2

are the masses of each compact object. This produces an obvious “chirp” signature as seen in figure 1.2, which is rather unique compared to most obfuscating noise sources. Since CBC signals can be modeled prior to detection, this allows matched filtering via template banks to be used in searches, affording us greater sensitivity to weak signals than otherwise.

Of primary interest to multi-messenger astronomy are systems involving neutron stars, since interaction of matter can produce electromagnetic and neutrino emissions. Neutron stars in CBC systems are thought to have masses roughly between $1M_{\odot}$ and $2.5M_{\odot}$, where this upper limit depends on the equation of state at nuclear densities (Abbott et al., 2021b). Neutron stars can be in either binary neutron star (BNS) systems or neutron star-black hole systems (NSBH), where the former has a strong multi-messenger case with GW170817 (Abbott et al., 2017b,a) and the latter has had two detections with no counterparts yet (Abbott et al., 2021a), creating some doubt in NSBH multi-messenger prospects (Fragione, 2021). Regardless, most multi-messenger searches tend to be agnostic and include both in their search.

2.1.2 Other Transients

There are also other potential GW signatures that have been proposed and that the LVK could detect. Although we used the example of rotating dumbbells with CBCs earlier in this section, we again note here that any system with non-zero second derivatives of its quadrupole tensor will create gravitational waves. For instance, core collapse supernova through either an asymmetric shock and/or neutrino emission have been proposed as GW progenitors (Abdikamalov et al., 2021). Core collapse supernova are prime candidates for multi-messenger detections, potentially creating a detectable long gamma-ray burst or neutrino burst as discussed later in sections 2.2 and 2.3 respectively.

Due to the numerous number of proposed models and to account for unknown models, other GW searches look instead for excess signal or energy after removing glitches and coherently combining the detectors (Abbott et al., 2009). Coherent Wave Burst (cWB)

uses this unmodeled approach, as well as other specific models in its search (Klimenko et al., 2008). Another search, X-pipeline, triggers off of GRBs for its unmodeled search and uses more computationally expensive methods in determining significance (Sutton et al., 2010). Although unmodeled searches have been improving over time, the detection distance going in O4 will still be mostly confined to our local galaxy, or at the most optimistic, relatively nearby galaxies (Abbott et al., 2020a).

2.2 Gamma-Ray Bursts

Gamma-ray bursts (GRB) are relatively short duration ($< 10^3$ s) pulses of high frequency light, with prompt emission in the gamma-ray/X-ray regime (> 1 keV) and afterglows in x-rays down to radio waves (Piran, 2005; Meszaros, 2006). These can be incredibly energetic, with isotropic luminosities up to 10^{53} erg/s, the same regime as GWs. GRBs are classically defined based on their duration, where a duration of less than 2 s is a short GRB (sGRB) while any longer is considered a long GRB (lGRB). The spectrum of a GRB is important as well, where hard GRBs skew to higher energy while soft GRBs skew to lower energy (note that sGRBs typically have harder spectrum than lGRBs) (Meszaros, 2006).

While the various mechanisms to create GRBs are still being understood, we will present a general scenario. A system with very high magnetic fields and outward moving gravitationally unbound matter can direct charged matter along these magnetic field lines. Shells of matter moving at different speeds would collide, and perhaps due to a combination of synchrotron radiation and inverse Compton scattering, matter is accelerated to relativistic speeds and emit photons at high energies (Piran, 2005). The presence of magnetic fields and relativistic speeds leads to the light being fairly collimated. We can see this since the relativistic velocity-addition formula states the opening angle θ_v should be reduced by a factor of $1/\Gamma$, where $\Gamma = \frac{1}{\sqrt{1-v^2/c^2}}$, v is the speed of the jet, and c is the speed of light.

GRBs are able to be detected from incredibly far distances, with the farthest detected having a redshift of 9.4 (Cucchiara et al., 2011). This jet of matter could then hit the surrounding medium, creating additional shocks and emitting from X-rays to radio waves as kinetic energy is lost due to this radiation.

Long and short GRBs have been predicted to come from different progenitors, where sGRBs are thought to originate from CBCs and magnetar flares, while lGRBs are thought to come from collapsars (MacFadyen & Woosley, 1999). The duration that a central engine can be sustained via accretion is thought to determine the length of a GRB, hence why a large collapsar should result in longer emission than a more compact CBC. We note that recent observations have brought this classical view into question, including a proposed association with a lGRB and kilonova (Rastinejad et al., 2022), as well as a sGRB thought to come from unsuccessful jet forming out of a collapsar (Rossi et al., 2021). This has emphasized the importance of using the spectrum and lightcurve of GRBs to categorize rather than only rely on the duration.

Since the Earth’s atmosphere is fairly opaque to gamma-rays, the detection of astrophysical GRBs have been focused on space-based missions. We have detailed the relevant instruments to this work in table 2.1. These collectively cover a wide field-of-view and energy range, giving a fairly complete view of the high-energy electromagnetic sky. We will discuss these experiments in conjunction with a joint search later in chapter 3. Overall, we expect to see roughly $\mathcal{O}(10^0)$ joint BNS-GRB detection per year during O4, with errors of an order of magnitude on the upper and lower bounds (Howell et al., 2019). Other potential coincidences are highly uncertain but worth preparing for.

2.3 Neutrino Bursts

Neutrinos are incredibly light particles ($< 10^{-36}$ kg) that are created by interactions of the weak force during different radioactive decays (Spurio et al., 2018; Aker et al., 2022).

Experiment	Instrument	Energy range	FOV (sr)	Citation
Fermi	LAT	20 MeV – 300 GeV	2.4	(Atwood et al., 2009)
	GBM	8 keV - 40 MeV	12.6	(Meegan et al., 2009)
<i>Swift</i>	BAT	15 keV - 150 keV	1.4	(Barthelmy et al., 2005)
	XRT	.2 keV - 10 keV	$1.3 \cdot 10^{-8}$	(Gehrels et al., 2004)
	UVOT	2 eV - 7.3 eV	$6.8 \cdot 10^{-9}$	(Gehrels et al., 2004)
INTEGRAL	IBIS	15 keV - 10 MeV	.11	(Ubertini et al., 2003)
AGILE	MCAL	300 keV - 100 MeV	12.6	(Labanti et al., 2009)

Table 2.1: List of relevant GRB experiments and their properties. The field-of-view (FOV) listed are half-coded when applicable. Some of these instruments can detect X-rays or UV rays since the former can be part of the initial burst or both in the afterglow. Note these FOVs don't include any occultation that the Earth or other bodies may cause.

Similarly to gravitational waves, they interact little with matter and must be indirectly observed. For example, the IceCube observatory looks for Cherenkov radiation characteristic of a charged particle moving faster than the phase velocity of ice, likely accelerated via inverse beta decay in a high-energy neutrino collision (Aartsen et al., 2017). Neutrinos have proven very useful in multi-messenger astronomy with both co-detections of supernova SN1987A (Hirata et al., 1987) and blazar TXS 0506+056 (Collaboration et al., 2018), giving insight to the most energetic moments of these events.

One of the primary progenitors we are interested in are core collapse supernovae, which alongside neutrinos, can potentially produce gravitational wave bursts as described in section 2.1.2 and long GRBs as described in section 2.2. Several scenarios can result in core collapse, including an iron core nearing the Chandrasekhar mass limit no longer able to support the inward pressure via electron degeneracy pressure. At a high enough density, electron (e^-) capture via protons (p) will produce neutrons (n) and electron neutrinos (ν_e)

$$p + e^- \rightarrow n + \nu_e \quad . \quad (2.2)$$

At high enough temperatures neutrino-antineutrino pairs will be created, believed to make up the majority of neutrinos that escape the system. The observable neutrinos from these

interactions are expected to have energies of 10-20 MeV (Athar & Singh, 2020).

Although there exists numerous active neutrino experiments with different energy ranges and observing goals, of primary interest to us is the SuperNova Early Warning System (SNEWS) (Antonioli et al., 2004; Al Kharusi et al., 2021). SNEWS is a collaboration between neutrino experiments to provide highly significant joint neutrino detections, which can potentially provide early warning of a supernova. The detection range is roughly confined to our own galaxy, similar to GW bursts, and the rate of such an event is roughly on the order of once per century (Antonioli et al., 2004; Al Kharusi et al., 2021). Although such a joint detection will be incredibly rare, the potential insight into the mechanics of supernovae would be invaluable.

2.4 Fast Radio Bursts

Fast radio bursts (FRB) are very short ($< .15$ s) bursts of low-energy light. Although the exact mechanism (or mechanisms) that produce FRBs are unknown, the fact they tend to be polarized and coherent gives us some insight (Zhang, 2020). FRBs experience significant dispersion, potentially due to interstellar or intergalactic medium, that appears to be consistent with redshift measurements of localized galaxies (Zhang, 2020). FRBs can either be repeating or non-repeating, where differences in bandwidths and intrinsic burst widths suggests these are two distinct populations (Amiri et al., 2021). Contributing greatly to the effort of detecting FRBs is the Canadian Hydrogen Intensity Mapping Experiment (CHIME) (Amiri et al., 2018), which has now reported over 500 detections (Amiri et al., 2021).

The association of an FRB with a soft gamma repeater (SGR) has firmly established magnetars as the progenitors at least some FRBs (Bochenek et al., 2020). This is consistent with the short time variability of FRBs indicating that they must come from compact objects. Although the origin of FRBs is still an active area of research, young magne-

tars formed from core collapse supernova and CBCs are still possible progenitors for non-repeaters ([Bhandari et al., 2020](#); [Amiri et al., 2021](#)). This gives motivation to search for coincidences between gravitational wave transients and FRBs. Such a search has been developed in conjunction with CHIME and has so far yielded no significant candidates during O3a ([Abbott et al., 2022](#)). We elaborate on our contributions to this effort in section [6.3](#). Still, there are incredible potential significant science gains if such an association is established in the future.

Chapter 3

RAVEN: The Low-Latency Gravitational Wave Focused Multi-messenger Pipeline

This is a preliminary version of a paper in preparation.

3.1 Motivation

With the coincident detection of the binary neutron star (BNS) gravitational wave (GW) signal GW170817 and gamma-ray burst GRB 170817A ([Abbott et al., 2017b,a](#); [Goldstein et al., 2017](#)), multi-messenger astronomy with gravitational waves has squarely entered the realm of observational science. This detection was found rather early by the multi-messenger pipeline referred to as the Rapid, on-source VOEvent Coincident Monitor (RAVEN), internally alerting the LIGO-Virgo collaboration (LVC) of this coincidence within a latency of 6 seconds from being called and 7 minutes of the GW merger.¹ Initial localizations were later provided by Fermi/GBM (delay of 25 min), and then by the LVC first using both the Hanford and Livingston detectors (delay of 50 min), and then later

¹<https://gracedb.ligo.org/events/view/G298048>

including the Virgo detector (delay of 4.5 hours).² This prompt detection and the subsequent localization led to the detection of the optical transient AT 2017gfo, announced around 11 hours after the merger (Valenti et al., 2017).² Unfortunately, this did not include the rise of the optical light curve which could have given stronger constraints on kilonova models (Arcavi, 2018). Since then, there has been a tremendous effort to reduce the latency of the LVK alert system. At the moment of writing, the total delay from the merger time to sending an alert is on the order of a *minute* or less, getting close to making early warning detection possible in O4 (Magee et al., 2021). RAVEN has also received separate major improvements, now being able to autonomously send alerts of coincident events and analyze *sub-threshold* GRBs.

RAVEN now has two primary functions: i.) find and determine the significance of GW events coincident with GRBs or neutrino bursts, including *sub-threshold* GW events, which could be of interest to the astronomical community and ii.) alert the external astronomical community and deliver data products in a latency low enough to facilitate electromagnetic follow-up. RAVEN can identify additional GW candidates via their proximity to external astrophysical transients and overlap in sky localizations.

In this paper, we will go over the design of RAVEN as a part of the LVK low-latency infrastructure in section 3.2, describe how RAVEN assigns significance to joint coincident candidates in section 3.3, present the results of simulations for this analysis in section 3.4, and conclude with a discussion in 3.5.

3.2 Design

The RAVEN pipeline is built as a module in `gwcelery`, the hub of the LVK low-latency analysis, and uses a few simple functions from the standalone python package `ligo-raven`. This package consists primarily of three functions: i.) a query of the Gravitational-Wave Candidate Event Database (GraceDB) to look for coincident GW events and external

²<https://gcn.gsfc.nasa.gov/other/G298048.gcn3>

events within a given time window, ii.) a calculation of joint significance using (3.1) or (3.8) depending on the type of search, and iii.) a calculation of the overlap between sky maps if they are available, described later in section 3.3.

The basic design of the RAVEN pipeline is to perform queries of the GraceDB database whenever either a GW event or external event is uploaded to GraceDB. The collection of data concerning a single GW candidate is organized into a superevent in GraceDB, which includes triggers from the various GW pipelines within a time window, while an external event is created via a Gamma-ray Coordination Network (GCN) notice (Moe et al., 2016). The query in GraceDB checks whether there is a complementary event within a predetermined coincidence time window, chosen based on both the type of event queried on and being searched for as seen in table 3.1. If a coincident event is found, a set of publishing conditions are then checked and if met, an automatic alert is sent out via a GCN notice. See figure 3.1 for a detailed flowchart of this pipeline and sections 3.3.3 and 3.3.4 for the full set of publishing conditions. Additionally, the *ligo-raven* package can be configured to run to perform rapid offline searches.

During O3, RAVEN listened for GRB candidates via GCN from Fermi/GBM (Meegan et al., 2009), *Swift*/BAT (Barthelmy et al., 2005), INTEGRAL (Winkler et al., 2003), and AGILE MCAL (Tavani et al., 2009), as well as neutrino burst candidates from SNEWS (Antonioli et al., 2004). Participating compact binary coalescence (CBC) focused GW pipelines include gstLAL (Sachdev et al., 2019), PyCBC (Usman et al., 2016), MBTAOnline (Adams et al., 2016), and SPIIR (Chu, 2017). In addition, the cWB pipeline also participates, primarily searching for unmodeled GW transients (Klimenko et al., 2008).

RAVEN ran using its query functionality during the entirety of O3. The ability to send alerts was added in O3b on January 14th with the deployment of gwcelery 0.9.2, the use of sky maps on February 11th with gwcelery 0.10.0, and the targeted analysis described in section 3.3.2 on March 10th with gwcelery 0.12.0. All of these changes have been thoroughly tested using built-in pytests, multiple rounds of manual reviews, and injections put into

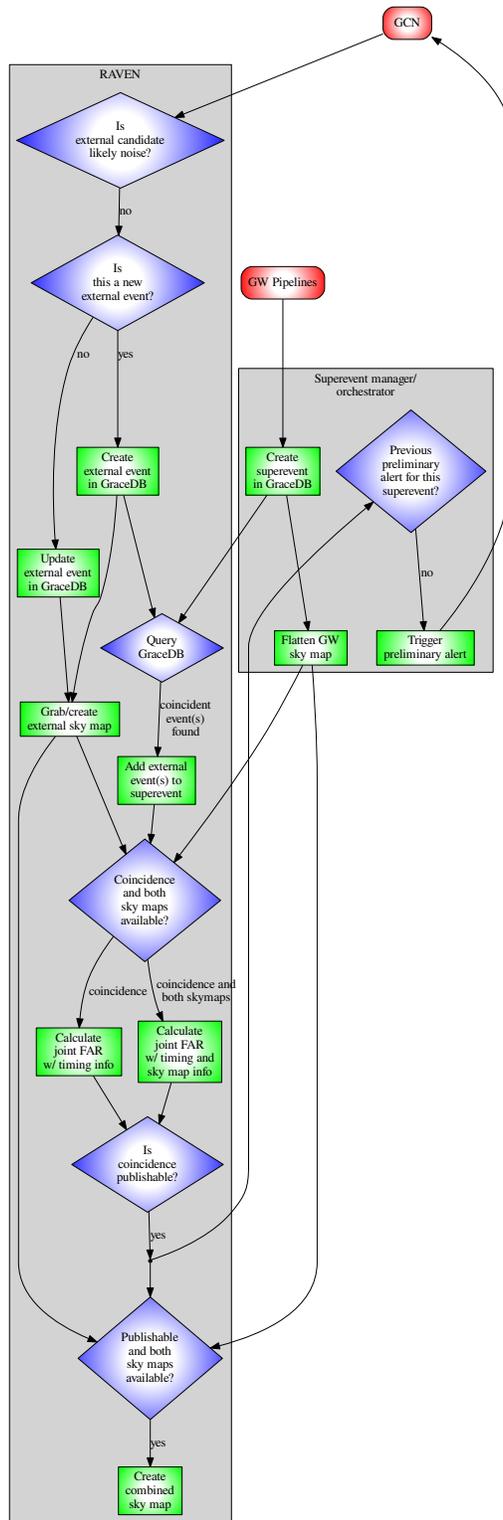


Figure 3.1: A flowchart of the logic used for GW-GRB coincidences within the RAVEN pipeline. Note that the logic with SNEWS coincidences is identical without the use of sky maps.

identical test GraceDB servers.

3.3 Joint Significance

Determining the joint significance of multi-messenger events has been an area of active research following GW170817. There have been many joint ranking statistics used in searches such as odds ratios and Bayes factors, but many of these approaches suffer from the fact that they require complete datasets to draw comparison against the background of signals and are therefore likely better suited to offline analysis (Ashton et al., 2018a; Stachie et al., 2020a; Hamburg et al., 2020). Instead for the purposes of a low-latency search like RAVEN, the false alarm rate (FAR) is likely a more useful statistic since this comparison is to the *immediate* surrounding background which is readily available. The LVK already uses FARs to assess the significance of GW candidates in data that is dominated by noise (Abbott et al., 2016b). In multi-messenger searches, the FAR is not just of either candidate but of the joint coincident event. A joint candidate with a significant joint FAR does not guarantee either candidate is significant or even has particularly useful parameters to be derived, say from an MCMC parameter estimation analysis. A high joint significance also does not guarantee the two candidates are truly associated, which is the goal of the coincidence odds ratio (Ashton et al., 2018a; Piotrkowski et al., 2022). A low FAR simply means that a similar false joint candidate does not occur often and likely should be of interest for additional follow-up.

RAVEN has two methods of assigning joint significance, separated into an untargeted method described in section 3.3.1 and a targeted method in section 3.3.2. The untargeted method involves taking in both streams of GW events and external events independently, performing queries when either are ingested. The targeted method instead requires sending moderately significant GW events (with FARs of less than one per day) to external partners, such as Fermi and *Swift*, where they then search their sub-threshold data for possible

coincidence. If a coincidence is found, the joint event is sent back and an alert is prepared. These methods are also differentiated by the nature of the GRB candidates considered. The untargeted method considers highly significant GRBs while the targeted considers lower significance GRBs well in the sub-threshold regime. This is reflected in the term dominating the rates of GRB candidates used, with the expected detectable rate of astrophysical GRBs being used in (3.1) and the FAR reported by the GRB pipeline in (3.8).

3.3.1 Untargeted GW-GRB Analysis

The test statistic used in the untargeted analysis that we assert is the joint (FAR_c) is

$$FAR_c = FAR_{gw} \frac{R_{grb} \Delta t}{\mathcal{I}_\Omega} \quad (3.1)$$

where FAR_{gw} is the FAR given by one of the GW pipelines, R_{grb} is the rate of unique GRBs detected by any combination of Fermi/GBM, *Swift*/BAT, INTEGRAL, and AGILE, Δt is the total coincident time window, and \mathcal{I}_Ω is the sky map overlap integral (Ashton et al., 2018a). This FAR method was first derived in Urban (2016) and then modified in Cho (2019). The GRB rate here has been measured as 310/*yr* after considering joint detections between GRB experiments (Urban, 2016; Bhat et al., 2016; Sakamoto et al., 2011; Vianello et al., 2009; Ursi et al., 2019). The coincident time window is chosen based on the given event and desired search, described in table 3.1. The GRB candidates considered here are highly significant, with the exception of sub-threshold Fermi GBM candidates with an additional rate of 65/*yr*, and are all published via GCN prior to being processed by RAVEN (Kocevski et al., 2018). We note these sub-threshold GRBs considered are still significant compared to the background, where filtering based on classification helps rule most of the noise transients out, so we still expect the real detection rate to dominate.

The sky map overlap integral in (3.1), first defined in Ashton et al. (2018a), can be

Search	Pipeline(s)	Untargeted	Targeted
CBC-GRB	Fermi/GBM	[-1, +5]	[-1, +10]
	<i>Swift</i> /BAT	[-1, +5]	[-10, +20]
	INTEGRAL	[-1, +5]	N/A
	AGILE	[-1, +5]	N/A
GW Burst-GRB	All GRB	[-60, +600]	N/A
GW Burst-Neutrino	SNEWS	[-10, +10]	N/A

Table 3.1: Coincident time windows in seconds chosen based on search and pipeline, centered on the GW merger/burst time. Wider windows were chosen for GW bursts due to the lack of confident tight delay models for GRBs from supernova and other progenitors of GW bursts (Zhang, 2019). Slightly wider windows were chosen for the targeted searches in order to include models that predict additional delay (Zhang, 2019), with *Swift* being confident to rule out noise events due to their highly precise localizations. The tighter window for SNEWS compensates for the time delay due to the mass of neutrinos emitted just outside our galaxy.

written as

$$\mathcal{I}_\Omega(x_{gw}, x_{grb}) = \int \frac{p(\Omega|x_{gw}, \mathcal{H}_{gw}^s)p(\Omega|x_{grb}, \mathcal{H}_{grb}^s)}{p(\Omega|\mathcal{H}^s)} d\Omega \quad (3.2)$$

where \mathcal{H}_a^s is the hypothesis that the data x_a is from a real astrophysical signal, $p(\Omega|x_a, \mathcal{H}_a^s)$ is the normalized probability density of the event a at coordinates Ω , and $p(\Omega|\mathcal{H}^s)$ is the prior on sky position, which we take as uniform so that $p(\Omega|\mathcal{H}^s) = 1/4\pi$. This uniform prior could likely be improved by considering the sensitivity of the experiments involved at the time of detection. To compute (3.2) with Hierarchical Equal Area isoLatitude Pixelization (HEALPix) sky maps (Górski et al., 2005b; Zonca et al., 2019), we can subdivide the sky into pixels with equal area $d\Omega \approx \Delta A$ as in Cho (2019). We denote each pixel by the index i so that $p(\Omega|x_a, \mathcal{H}_a^s)\Delta A \approx P(i|x_a, \mathcal{H}_a^s)$ and $4\pi \cdot p(\Omega|x_a, \mathcal{H}_a^s) \approx N_{pix} \cdot P(i|x_a, \mathcal{H}_a^s)$. This all together yields

$$\mathcal{I}_\Omega(x_{gw}, x_{grb}) = N_{pix} \sum_{i=1}^{N_{pix}} P(i|x_{gw}, \mathcal{H}_{gw}^s) \cdot P(i|x_{grb}, \mathcal{H}_{grb}^s) \quad (3.3)$$

where we have assumed each sky map has the same resolution with N_{pix} being the number

of pixels in a single sky map.

We expect the statistic in (3.1) to generally have the properties of a FAR, assuming FAR_{gw} has these properties. If we ignore sky maps for the moment, we can see that $R_{em}\Delta t$ is the expected number of random coincidences per GW candidates, which compensates for the reduction of joint candidates we expect to see by increasing their significance. We discuss the importance of this property in section 6.1. In addition, we performed simulations that show that using uncorrelated sky maps gives $\langle \mathcal{I}_\Omega \rangle \approx 1$ in tables 3.3 and 3.5. Overall this means we are convolving a FAR distribution with random variable of mean 1, giving us a distribution that overall represents a FAR but could have skews in certain regions (see section 3.4).

3.3.2 Targeted GW-GRB Analysis

RAVEN is also used as an alert system for both the joint LVK-Fermi/GBM and LVK-*Swift* targeted searches. This involves the LVK sending moderately significant GW events to these experiments ($FAR < 1/day$), these external pipelines looking through their sub-threshold data for joint events, and uploading a new external event to GraceDB if a joint candidate is found. RAVEN will then determine whether the joint event is publishable and then alert both LVK members and the external astronomical community of these events. The analysis used by RAVEN and these external experiments are meant to be identical, using the same search windows and joint FAR calculation. Since this type of analysis is dominated by noise-noise coincidences, we can write the joint ranking statistic Z as

$$Z = FAR_{gw}FAR_{grb}\Delta t \tag{3.4}$$

where FAR_{gw} and FAR_{grb} are the FARs of their respective experiments. This statistic can be mapped to a joint FAR since each individual FAR is drawn from a uniform distribution

$$FAR_t(Z) = \int_0^{Z_{max}} P(FAR_{grb}\Delta t \leq Z/FAR_{gw})dFAR_{gw} \quad (3.5)$$

$$= \int_0^Z dFAR_{gw} + \int_Z^{Z_{max}} Z/FAR_{gw}dFAR_{gw} \quad (3.6)$$

$$= Z(1 + \ln(Z_{max}/Z)) \quad (3.7)$$

where $Z_{max} = FAR_{gw,max}FAR_{grb,max}\Delta t$, given from the max FAR thresholds of each experiment. A final coincidence FAR_c with sky map information can be made by dividing by (3.3) similarly to (3.1) to get

$$FAR_c = FAR_t/\mathcal{I}_\Omega. \quad (3.8)$$

We note that sky maps from these external experiments may be either delayed or entirely absent, so RAVEN will try to create an approximate Gaussian sky map for a GRB based on the sky localization information from GCN and known systematics, or simply calculate the joint FAR without this information if unavailable. Next, for a coincidence to be published the joint FAR (either (3.1) or (3.8) depending on the search) must be under predetermined thresholds.

3.3.3 GW-GRB Publishing Conditions

RAVEN is designed to be agnostic when searching for coincidences, but stringent when deciding to publish them. This is to alert LVK members internally of any potentially interesting candidates, while avoiding sending false alarms to the external astronomical community. Regardless of whether a joint candidate passes the following publishing conditions for automatic alerts, LVK EM advocates have the ability to publish or retract candidates manually if the need arises.

RAVEN will automatically publish joint GW-GRB candidates that pass the following conditions:

1. The joint FAR, including trials factors, is lower than the public alert threshold (CBC: 1 per 2 months, GW Burst: 1 per year). The sky map overlap in (3.3) will be included if both sky maps are available. GRB candidates from *Swift* require both sky maps to be available.
2. The GRB candidate is not sub-threshold, i.e. considered significant apart from RAVEN. Sub-threshold GRBs require further analysis and approval of experts from the corresponding experiment before manually publishing.
3. The GRB candidate is not likely from an unrelated noise source, i.e. if classifications are available then there is at least a 50% probability the candidate is a GRB.
4. An alert for the GW candidate has not been sent. If a GW candidate is publishable based on both its joint significance and its own significance, one alert will be sent containing additional information on the coincidence unless this coincidence is discovered after the initial GW-only alert, which then another alert will be sent about this coincidence manually.

The trials factors used here are worth expanding further upon. Since RAVEN both listens to multiple GW pipelines and is also calibrated to submit triggers at the rate of a GW pipeline, the trials factor used is the product of these two effects. In other words if N is the number of independent GW searches aside from RAVEN (4 CBC searches; 3 GW burst searches), the total trials factor is then $N(N + 1)$. This product both accounts for the number of GW pipelines used in the RAVEN analysis (N) as well as calibrate to the total number of pipelines publishing alerts which now includes RAVEN ($N + 1$). This is a rather conservative estimate, approximately true at low significance, but at high significance multiple GW searches tend to find the same candidates and thus are not

truly independent. The various GRB experiments have been combined by using the joint detected GRB rate as described in section 3.3.1.

3.3.4 GW Burst-SNEWS

Neutrino burst candidates from SNEWS have a very low FAR ($< 1/100yr$) and also lack a promptly available sky localization (Antonioli et al., 2004), so the only determiner of the joint significance is the FAR of the GW candidate. We have chosen to publish joint GW Burst-SNEWS events if the GW candidate has a FAR of less than 1 per day after including trials factors. We also acknowledge the related low-latency GW-neutrino search, the Low-latency algorithm for multi-messenger astrophysics (LLAMA), designed specifically for this type of search (Countryman et al., 2019; Keivani et al., 2019).

3.4 FAR Simulations

In order to test whether both (3.1) and (3.8) have the the properties of a FAR, we created a background dataset by performing simulations of random coincidences. We simulated GW and Fermi/GBM-like GRB candidates with random event times during a period of 100 years, looking for events that fall within the given time window according to table 3.1. We repeated this analysis with *Swift*/BAT-like GRB candidates as well. Each simulation was run 50 times to measure the spread of results per inverse FAR bin. The FAR for each individual candidate was drawn from a uniform distribution, with max FAR values of $FAR_{GW,max} = 1/hr$ for the untargeted search (Magee et al., 2021). The Fermi/GBM targeted search used $FAR_{GW,max} = 1/day$ and $FAR_{GRB,max} = 1/10000 s$, while the *Swift* targeted search used $FAR_{GW,max} = 2/day$ and $FAR_{GRB,max} = 1/1000 s$.

We also wanted to understand the effect of sky maps in these simulations. We used sky maps from both sub-threshold GW candidates during O3 and real Fermi/GBM candidates (Bhat et al., 2016; Connaughton et al., 2015), indicative of both the targeted and untargeted

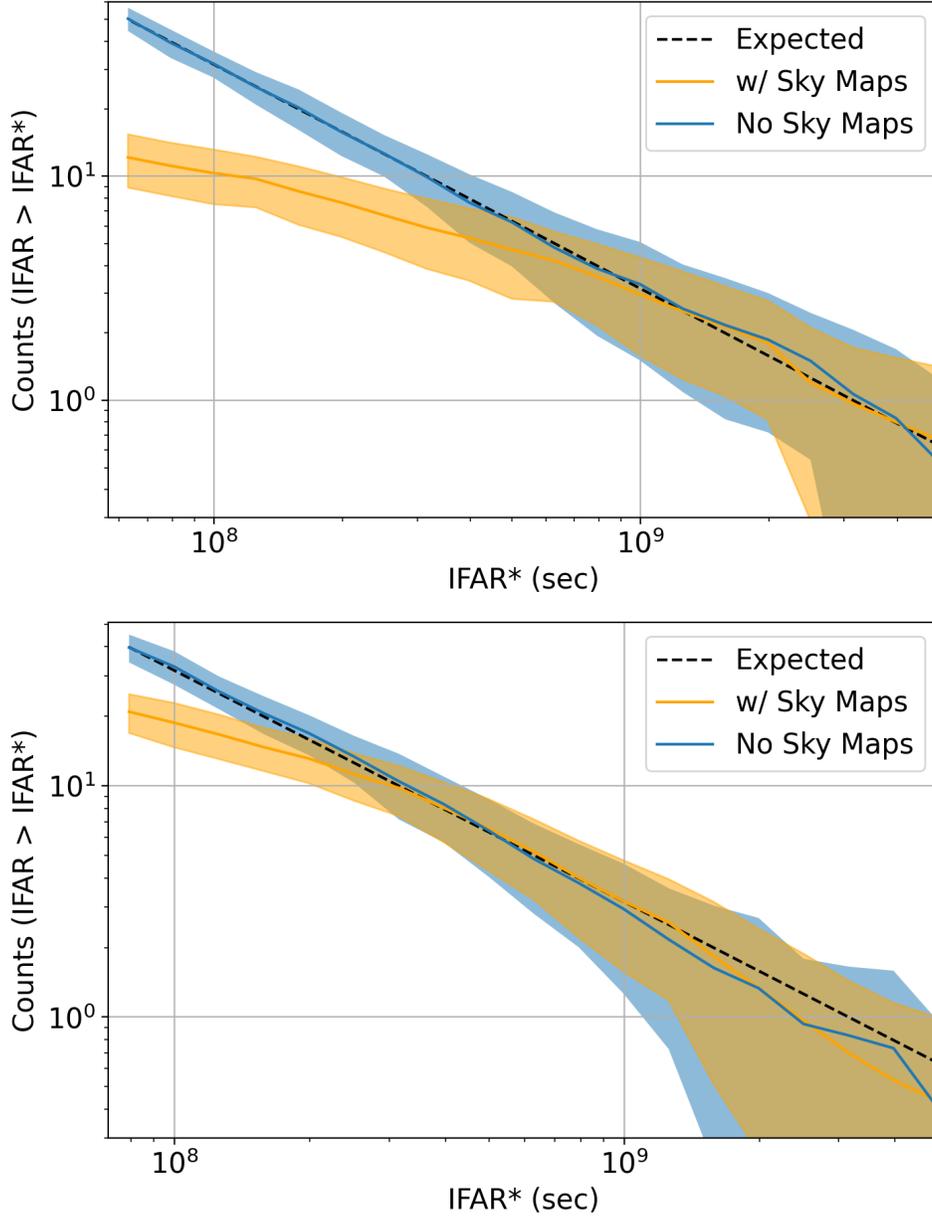


Figure 3.2: Cumulative counts of the inverse joint false alarm rate (IFAR) calculated using sky maps (*orange*) and no sky maps (*blue*), with the top figure being simulated from the untargeted Fermi/GBM search using (3.1) and the bottom figure simulated from the targeted Fermi/GBM search using (3.8). The colored regions show the 1-sigma uncertainty around the mean, represented by a solid line. There is a skew due a number of instances with small overlap as seen in figure 3.3. This skew is less pronounced for the targeted search due to the low-significance GRB sky maps being less informative and more poorly localized.

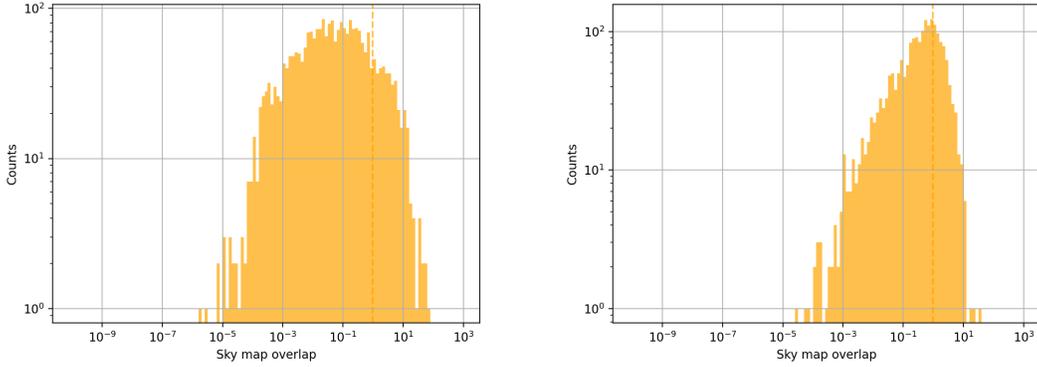


Figure 3.3: Sky map overlap integral calculated with (3.3) using Fermi/GBM-like sky maps consistent with each search. There are a number of very small overlap values due to astrophysically motivated localizations that are highly peaked and inconsistent with each other. Note that the mean of each method is consistent with 1.

search respectively, and then compared this approach to using no sky map information (i.e. $\mathcal{I}_\Omega = 1$). For the *Swift*/BAT searches, due to the highly localized nature of these GRBs, we modeled these as single pixel sky maps. For every coincidence we drew a random sky map from each respective set, effectively meaning the GRB and LVC sky maps were uncorrelated.

These results can be found in figures 3.2 and 3.4 where we see that both the untargeted and targeted methods generally result in sensible FARs without sky maps, while including this information causes the FAR to be underestimated due to large number of low overlap joint events (see figures 3.3 and 3.5). This effect is even more pronounced for LVC-*Swift* events due to the general lower values of sky map overlap compared to LVC-Fermi, as well as for the untargeted searches. This skew is inherent with the current method, as we have a uniform distribution (FAR) being convolved with another distribution (sky map overlap), which is not guaranteed to create another uniform distribution (joint FAR). In the case where we know what the two distributions will be beforehand, we can correct for skewness with an analytical transformation as in equation (3.7) where we initially had two uniform distributions. However the sky map overlap is dependent on the detectors and specifics of the individual candidate searches at the time, and therefore has no obvious predetermined

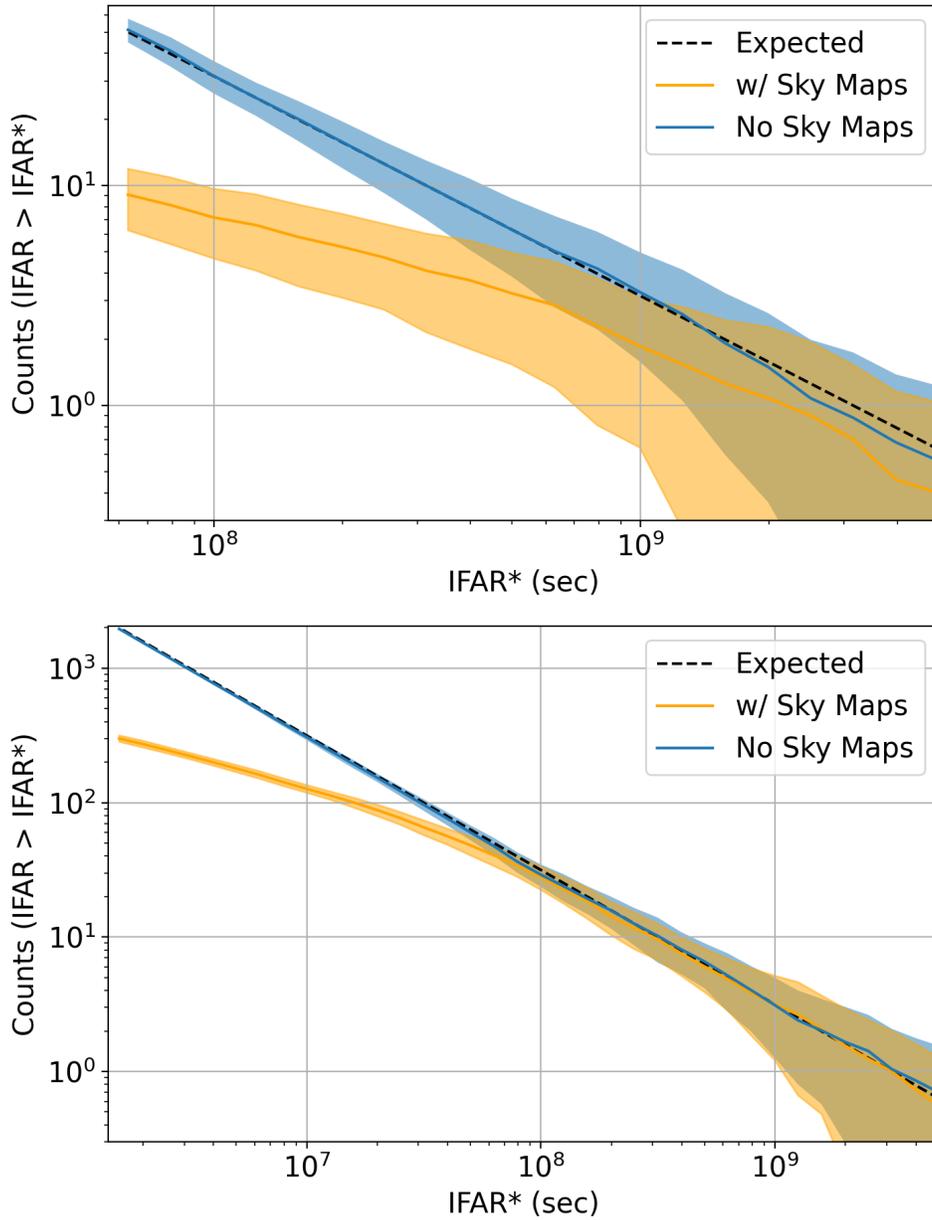


Figure 3.4: Similar to figure 3.2, cumulative counts of the inverse joint false alarm rate (IFAR) calculated using sky maps (*orange*) and no sky maps (*blue*), with the top figure being simulated from the untargeted *Swift*/BAT search using (3.1) and the bottom figure simulated from the targeted *Swift*/BAT search using (3.8). There is even greater skew than present in 3.2 due to the large population of incredibly low sky map overlap values seen in figure 3.5.

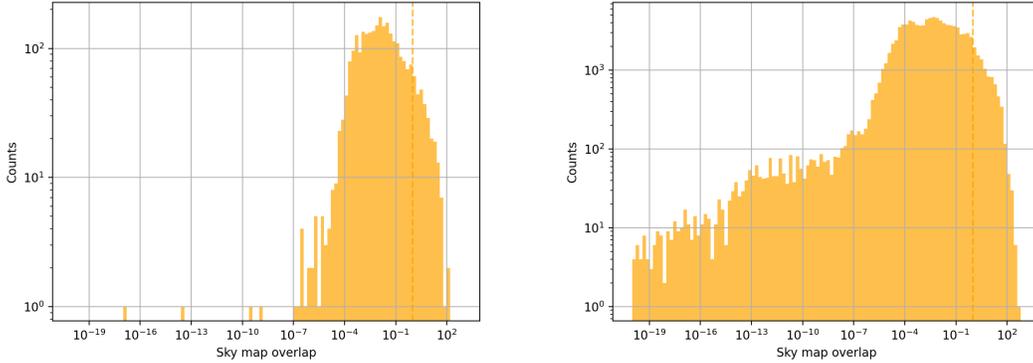


Figure 3.5: Sky map overlap integral calculated with (3.3) using *Swift*/BAT-like sky maps consistent with each search. Compared to figure 3.3, there are even lower overlap values due to how highly localized *Swift* events are. Note as in figure 3.3, the mean of each method is consistent with 1.

shape (see figures 3.3 and 3.5). The only trivial second distribution that results in a uniform distribution is that of unity, which we have done by setting $\mathcal{I}_\Omega = 1$ and labeled this as no sky map information in figures 3.2 and 3.4. We could also numerically correct for this skew after an operating run when all the candidates have been collected, similar to Stachie et al. (2020a) by doing a simple remapping to the expected joint FAR. However at the time of writing, there is no clear solution without knowing the sky map overlap distribution *a priori*. Even in the best case, using pre-computed distributions presents a difficult problem to avoid additional biases if changes to any of the searches are made midway through an operating run or, worse yet, the wrong assumptions are made. We explore this type of remapping further in section 6.2.

3.5 Discussion

We described the low-latency coincidence pipeline RAVEN, whose primary functions are to find interesting coincident astrophysical events and then to report these to the external astronomical community in a latency low enough to facilitate additional follow-up. We discussed the design of this pipeline, the methods to assign significance for both the targeted

and untargeted searches, and ran simulations to measure their validity.

The choice to make RAVEN as agnostic as possible does come at the cost of sensitivity. For instance, we could choose to filter out likely binary black hole GW candidates by considering mass, effectively reducing our trials and increasing significance of the remaining candidates. We could also reduce our time windows in the wake of the 1.7 second delay between GW170817 and GRB 170817A, achieving a similar effect. However since there has been one confident GW-GRB detection, which has only begun to constrain coincidence models, we have decided to retain sensitivity to events that may be different than predicted by these models.

We also recognize the shortcomings of this method, especially with regards to the skewness in the joint FAR as discussed in section 3.4, but don't believe this invalidates the analysis. We must first recognize that this makes our method more conservative with regards to lower significance events, which is much more palatable compared to overestimating high significance events. This could lead to a loss of potentially interesting joint candidates, but will also definitely reduce the number of published false alarms as well. We emphasize that coincidences of any significance will be reviewed manually by pipeline experts as they were in O3. Secondly, we must recall that the purpose of RAVEN is to rapidly alert astronomers, which already requires approximations for the sake of brevity. There exists more sophisticated methods of assigning significance with PyGRB and X-pipeline, but these operate with the goal being the more robust rather than computable in low latency (Abbott et al., 2017d). Just like how GW pipelines use quick online template searches and then more thorough offline searches exist, RAVEN has use its niche as a rapid joint pipeline. We will keep exploring improvements to this method, including potentially using distance and inclination information, as well as using the sensitivity of the experiments to sky localization.

We also note that the joint FAR is likely not the best determiner of whether two events are truly coincident but instead whether the GW candidate should receive electromagnetic

follow-up. For instance, in the limit where the GW candidate is very significant any association with an external event will likely pass the alert threshold while the sky map overlap integral may strongly argue against this being a real coincidence. We can see this issue is prevalent in a search described in section [6.3.2](#). Even in this case, the joint FAR is still useful for identifying events that warrant further follow-up (whether individually significant or already accompanied by another detection), the primary purpose of RAVEN.

RAVEN will be an especially important tool in O4 and onward, potentially significantly contributing to the number of BNS detections and maybe even helping to discover the first unmodeled GW transient.

Chapter 4

A Joint Ranking Statistic for Multi-messenger Astronomical Searches with Gravitational Waves

This chapter is reproduced in part from [Piotrkowski et al. \(2022\)](#), which has been accepted for publication in *Classical and Quantum Gravity*. An online record can be found at <https://doi.org/10.1088/1361-6382/ac5c00>.

4.1 Motivation

One area of active research for joint searches has been the ranking statistic used, whose function is to rank candidates based on how likely they are to have a real joint origin rather than be randomly associated, an especially difficult task in the sub-threshold regime. The types of statistics that have been used include modified signal-to-noise ratios ([Nitz et al., 2019](#)), false alarm rates ([Urban, 2016](#); [Cho, 2019](#)), and other generic ranking statistics ([Hamburg et al., 2020](#)). Among this last group, [Ashton et al. \(2018b\)](#) put forward a general

methodology that considered the overlap of time and sky localizations as the primary determiners of significance. This seems intuitive since real coincidences must share the same underlying parameters. Next, [Stachie et al. \(2020b\)](#) improved this by considering additional noise hypotheses and including the significance of the individual candidates, albeit with some approximations. This allowed them to start to extend this method to the sub-threshold regime, where many coincidence events may still be distinguishable from noise ([Burns, 2019](#)).

In this work we expand on these methods further by considering the overlap of additional parameters (distance and inclination for BNS-GRBs) and show that filtering using coincidence models emerges naturally, shown by avoiding approximations until warranted by specific use cases. We also demonstrate the use of a galaxy catalogue as a prior for sky localization and distance in the case of GW-GRBs. Although this statistic will be inevitably more computationally expensive than previous, we derive a version with minimal additional computational cost and intended to be used in low-latency targeted searches, as well as a best-possible version using posterior samples given from parameter estimation.

This work is meant to be a Bayesian derivation of the odds ratio for whether two events are correlated (real coincidence) versus non-correlated (random coincidence), primarily based on the results of [Ashton et al. \(2018b\)](#) and [Stachie et al. \(2020b\)](#) while similarly applying this method specifically to GW-GRB coincidences. We note that other Bayesian approaches have been developed for multi-messenger searches in the context of GW-neutrino bursts ([Bartos et al., 2019](#); [Veske et al., 2020](#)). We will first write out the odds ratio using Bayes rule, expand on the Bayes factors and prior odds, write out these terms in the context of GW-GRB candidates, and then conclude with a number of example demonstrations of this statistic.

4.2 Derivation

The joint ranking statistic of interest here will be the odds ratio between whether two candidates are truly correlated (\mathcal{H}^c), meaning they have the same underlying shared parameters θ (Ashton et al., 2018b), or whether they are uncorrelated ($\mathcal{H}^{\text{uncor}}$), meaning the association is from random chance. To understand the uncorrelated hypothesis as described in Stachie et al. (2020b), let us first consider a single isolated astrophysical event candidate. This candidate has only a couple of possibilities for its origin: either this event is from the intended astrophysical source (\mathcal{H}^s) or is an artifact of noise (\mathcal{H}^n), instrumental or otherwise. Next, if we consider two unrelated candidates detected by different observatories, then the joint hypothesis between them is then comprised of the two individual hypotheses. In other words, since either of the two events could be an unrelated astrophysical signal or a noise transient, this means the joint candidate could be any of the four permutations \mathcal{H}^{XY} where $X, Y \in \{s, n\}$. Therefore the total uncorrelated hypothesis is the combination of these four possibilities, although they may be not all have the same probability and thus require different relative weights.

We can write out the odds ratio between the correlated and uncorrelated hypothesis generally as

$$\mathcal{O}_{c/\text{uncor}} = \frac{P(\mathcal{H}^c|x_a, x_b)}{P(\mathcal{H}^{\text{uncor}}|x_a, x_b)} \quad (4.1)$$

where x_a and x_b are independent data sets (e.g. data from a GW and a GRB candidate respectively). Using Bayes' theorem this can be refactored to

$$\begin{aligned} \mathcal{O}_{c/\text{uncor}} &= \frac{P(x_a, x_b|\mathcal{H}^c)}{P(x_a, x_b|\mathcal{H}^{\text{uncor}})} \frac{P(\mathcal{H}^c)}{P(\mathcal{H}^{\text{uncor}})} \\ &= B_{c/\text{uncor}}(x_a, x_b) \frac{P(\mathcal{H}^c)}{P(\mathcal{H}^{\text{uncor}})} \end{aligned} \quad (4.2)$$

where this first term $B_{c/\text{uncor}}(x_a, x_b)$ is the joint Bayes factor, and the second is the prior odds. Next we will expand these two factors in terms of statistics that can be provided by

astronomical experiments.

4.2.1 General Multi-messenger Candidate

Let us first work with the joint Bayes factor. Since the four components of $\mathcal{H}_{\text{uncor}}$ are mutually exclusive of each other, we can write the total probability as the sum of each part (Stachie et al., 2020b). We do this by first using Bayes' theorem and writing

$$P(x_a, x_b | \mathcal{H}^{\text{uncor}}) = \frac{P(\mathcal{H}^{\text{uncor}} | x_a, x_b) P(x_a, x_b)}{P(\mathcal{H}^{\text{uncor}})} \quad (4.3)$$

$$= \frac{\sum_{X, Y \in \{s, n\}} P(\mathcal{H}^{XY} | x_a, x_b) P(x_a, x_b)}{P(\mathcal{H}^{\text{uncor}})} . \quad (4.4)$$

Now using Bayes' theorem again, we get

$$P(x_a, x_b | \mathcal{H}^{\text{uncor}}) = \frac{\sum_{X, Y \in \{s, n\}} P(x_a, x_b | \mathcal{H}^{XY}) P(\mathcal{H}^{XY})}{P(\mathcal{H}^{\text{uncor}})} \quad (4.5)$$

Thus the total uncorrelated likelihood can be written as the prior weighted sum of each component likelihood. Next let us examine the correlated hypothesis in (4.2), where we can marginalize over the set of all shared parameters θ between x_a and x_b to get

$$P(x_a, x_b | \mathcal{H}^c) = \int_{\Theta_c} P(x_a, x_b | \theta, \mathcal{H}^c) P(\theta | \mathcal{H}^c) d\theta . \quad (4.6)$$

We note that we have restricted our domain of θ to the subset Θ_c where $P(\theta | \mathcal{H}^c) > 0$.

Let us expand the first term in the integrand as

$$P(x_a, x_b | \theta, \mathcal{H}^c) = P(x_a | \theta, \mathcal{H}^c) P(x_b | \theta, \mathcal{H}^c) \quad (4.7)$$

since x_a and x_b are independent data sets. Now using Bayes' Theorem we get

$$P(x_a, x_b | \theta, \mathcal{H}^c) = \frac{P(x_a | \mathcal{H}^c) P(\theta | x_a, \mathcal{H}^c)}{P(\theta | \mathcal{H}^c)} \frac{P(x_b | \mathcal{H}^c) P(\theta | x_b, \mathcal{H}^c)}{P(\theta | \mathcal{H}^c)}. \quad (4.8)$$

It's worth noting that we have used \mathcal{H}^c in conjunction with data sets from a single experiment here, which still has a clear meaning since the parameter spaces for correlated and generic detections are not identical. We will expand further on such terms in section 4.2.3.

Putting the above equation back into (4.6), we are left with

$$P(x_a, x_b | \mathcal{H}^c) = P(x_a | \mathcal{H}^c) P(x_b | \mathcal{H}^c) \int_{\Theta_c} \frac{P(\theta | x_a, \mathcal{H}^c) P(\theta | x_b, \mathcal{H}^c)}{P(\theta | \mathcal{H}^c)} d\theta \quad (4.9)$$

where we will refer to this integral as the overlap integral $\mathcal{I}_\theta(x_a, x_b)$, as similarly derived in Ashton et al. (2018b).

Since the likelihoods in (4.5) are composed of the independent individual likelihoods of each experiment, we can write these as the product

$$P(x_a, x_b | \mathcal{H}^{XY}) = P(x_a | \mathcal{H}_a^X) P(x_b | \mathcal{H}_b^Y). \quad (4.10)$$

We can then write out the overall Bayes factor in (4.2) using (4.5) and (4.9) to get

$$B_{c/\text{uncor}}(x_a, x_b) = \frac{P(\mathcal{H}^{\text{uncor}})}{\sum_{X,Y \in \{s,n\}} B_{XY/c}(x_a, x_b) P(\mathcal{H}^{XY})} \quad (4.11)$$

where each individual Bayes factor, defined as the ratio of (4.10) and (4.9), is given by

$$\begin{aligned} B_{XY/c}(x_a, x_b) &= \frac{P(x_a | \mathcal{H}_a^X) P(x_b | \mathcal{H}_b^Y)}{P(x_a | \mathcal{H}^c) P(x_b | \mathcal{H}^c)} / \mathcal{I}_\theta(x_a, x_b) \\ &= B_{X/c}(x_a) B_{Y/c}(x_b) / \mathcal{I}_\theta(x_a, x_b) \end{aligned} \quad (4.12)$$

and where we have reduced this to the Bayes factor of each individual event along with

the overlap integral between their shared parameters. Note that the noise vs coincidence Bayes factor can be separated out as

$$B_{n/c}(x_a) = B_{n/s}(x_a)B_{s/c}(x_a). \quad (4.13)$$

with $B_{n/s}(x_a)$ being the noise vs signal Bayes factor, which can be determined using the event and properties of the detector, while $B_{s/c}(x_a)$ is the uncorrelated signal vs correlated signal Bayes factor, which can be determined by whether the measured parameters are consistent with a given coincidence model. If we plug (4.11) into (4.2), using (4.12) and (4.13), we get

$$\mathcal{O}_{c/\text{uncor}}(x_a, x_b) = \frac{P(\mathcal{H}^e)\mathcal{I}_\theta(x_a, x_b)B_{c/s}(x_a)B_{c/s}(x_b)}{P(\mathcal{H}^{ss}) + B_{n/s}(x_a)P(\mathcal{H}^{ns}) + B_{n/s}(x_b)P(\mathcal{H}^{sn}) + B_{n/s}(x_a)B_{n/s}(x_b)P(\mathcal{H}^{nn})}. \quad (4.14)$$

The terms in (4.14) are worth discussing in more detail. This statistic can be interpreted as having contributions from three distinct parts: 1.) the significance of the individual candidates weighted by their priors, 2.) the overlap of the shared parameters between the candidates, and 3.) the evidence for whether the data for each candidate favors a coincident model rather than a generic detection model. We see in the limit where each candidate is infinitely significant, there is still the possibility that they are randomly associated (i.e. \mathcal{H}^{ss}), leaving the overlap integral $\mathcal{I}_\theta(x_a, x_b)$ and coincidence Bayes factors as the sole determiners of joint significance. Previous methods have not considered these coincidence Bayes factors and have set $B_{c/s}(x_a) = 1$ (Ashton et al., 2018b; Stachie et al., 2020b). However this is not strictly correct, as there are regions of parameter space with different expectations of whether a coincidence is possible. We will discuss these specific terms using a specific example in section 4.2.3.

4.2.2 GW-GRB Candidate

We will next focus on writing down the various terms in (4.14) for the specific case of GW-GRB coincidences, starting with the prior terms. The probability of detecting a single event should follow a Poisson distribution (Urban, 2016; Ashton et al., 2018b) so that

$$P(\mathcal{H}^X) = R_X T e^{-R_X T}. \quad (4.15)$$

where T is the co-observing time and R_X is the rate. If we consider the co-observing time T for an individual candidate as the co-observing time window Δt , typically at least two orders of magnitude higher than the corresponding rate, we can take the approximation that

$$P(\mathcal{H}^X) \approx R_X T. \quad (4.16)$$

For the case of two independent events, the probability that they both occur is the product

$$P(\mathcal{H}^{XY}) \approx R_X R_Y T^2. \quad (4.17)$$

Meanwhile the noise vs signal Bayes factor $B_{n/s}(x_a)$ can be supplied by the given experiment or calculated by using a proxy detection statistic once a number of events have been collected (Stachie et al., 2020b). Another example of this is calculated by `ligo.skymap`, although this assumes Gaussian noise rather than the non-stationary noise artifacts typically found in LVK data.

The overlap integral $\mathcal{I}_\theta(x_a, x_b)$ derived in (4.9) measures the overlap between the common set of parameters θ of two experiments. For joint GW-GRB detections (under the \mathcal{H}^c

hypothesis) this is

$$\mathcal{I}_\theta(x_{gw}, x_{grb}) = \int_{\Theta_c} \frac{P(\theta|x_{gw}, \mathcal{H}^c)P(\theta|x_{grb}, \mathcal{H}^c)}{P(\theta|\mathcal{H}^c)} d\theta \quad (4.18)$$

$$= \int_{\Theta_c} dD_L d\iota d\Omega dt_c \frac{P(D_L, \iota, \Omega, t_c|x_{gw}, \mathcal{H}^c)}{P(D_L, \iota, \Omega, t_c|\mathcal{H}^c)} P(D_L, \iota, \Omega, t_c + t_d|x_{grb}, \mathcal{H}^c) \quad (4.19)$$

where the shared parameters are $\theta = \{D_L, \iota, \Omega, t_c\}$. More explicitly, D_L is the luminosity distance, ι the inclination angle, Ω the sky location of the source and t_c the coalescence time. The two parameters more to relevant GRB observations, redshift z and viewing angle θ_v , can be simply converted to D_L and ι by assuming a cosmology and reflecting about 90° respectively. We note that the overlap in distance has been used before in the case of GW190521-ZTF19abahr (Ashton et al., 2020). The parameter t_d is the time delay between the the GW and GRB arrival times, which we allow to span the possible range of t_l before to t_h after the coalescence time, meaning the total co-observing time window is $\Delta t = |t_h - t_l|$. If the observations are disjoint, with inconsistent measurements of the parameters, the overlap integral above will be negligible and thus will heavily weight against the correlated hypothesis. From the generic overlap integral \mathcal{I}_θ we can separate out the overlap in time

$$\mathcal{I}_{t_c} = \begin{cases} T/\Delta t & (t_l < t_d < t_h) \\ 0 & (otherwise) \end{cases}$$

as in Ashton et al. (2018b) and Stachie et al. (2020b) to write (4.14) as

$$\begin{aligned} \mathcal{O}_{c/uncor}(x_{gw}, x_{grb}) = & \\ & \frac{R_{gw,grb}^c \mathcal{I}_{\Omega, D_L, \iota}(x_{gw}, x_{grb}) B_{c/s}(x_{gw}) B_{c/s}(x_{grb})}{\Delta t [R_{gw}^s R_{grb}^s + B_{n/s}(x_{gw}) R_{gw}^n R_{grb}^s + B_{n/s}(x_{grb}) R_{gw}^s R_{grb}^n + B_{n/s}(x_{gw}) B_{n/s}(x_{grb}) R_{gw}^n R_{grb}^n]} \end{aligned} \quad (4.20)$$

where the R_a^n terms are the rates of noise triggers that are considered in a given joint

search, the R_a^s terms are expected rates of detected real astrophysical triggers, and $R_{a,b}^c$ is the expected rate of detected real coincidences. We note that (4.20) in principle should perform well in sub-threshold searches because for a joint candidate to have a higher significance, both individual candidates need a higher individual significance. This effect will be more pronounced than in Stachie et al. (2020b) due to the inclusion of rates, helping to penalize obvious noise events and potentially bring real coincidences to the foreground.

We consider (4.20) the primary result of this paper and later detail how to calculate this for specific examples in section 4.3.

4.2.3 Coincidence Bayes Factors

The presence of coincident vs signal Bayes factors $B_{c/s}(x_a)$ in (4.20) is one way that differentiates this statistic from previous efforts. This was claimed in (Cho, 2019) to always be unity, but ignored the differences in priors between coincidences and generic detections. We can see this by expanding over the parameters θ_a

$$B_{c/s}(x_a) = \frac{\int P(x_a|\theta_a, \mathcal{H}^c)P(\theta_a|\mathcal{H}^c)d\theta_a}{\int P(x_a|\theta_a, \mathcal{H}_a^s)P(\theta_a|\mathcal{H}_a^s)d\theta_a} \quad (4.21)$$

using the likelihood $P(x_a|\theta_a, \mathcal{H}_a^s)$ and prior $P(\theta_a|\mathcal{H}_a^s)$. The coincidence Bayes factor should be interpreted as how much the measured parameters of an *individual* experiment support a coincidence by comparing between coincidence and generic detection models. For example, if the masses of a GW candidate are both above $5M_\odot$ we would think the progenitor is a binary black hole system and not capable of producing a GRB, hence we could set $B_{c/s}(x_{gw}) = 0$. For GRBs, an example would be placing a cutoff on the duration and require $T_{90} \leq 2.0s$ as to only consider short GRBs. This is already possible to include in multi-messenger searches by simply screening out candidates that don't meet certain parameter requirements. In other words, one could define $B_{c/s}(x_a)$ as a convolution of step functions by giving 1 if the parameters fall within expected bounds and 0 if not,

an especially useful method for a low-latency search. However, a more explicit calculation could be used to gain more sensitivity and we will show such an example by using posterior distributions in the next section.

4.3 Examples

Let us test our method by applying (4.20) to a number of situations, including a couple of known GW-GRB coincident candidates, GW170817-GRB 170817A and GBM-GW150914, as well as a simulation of many coincidences. If the individual events are significant enough so that the noise-to-signal Bayes factors $B_{n/s}(x_a) \approx 0$, then our odds ratio from (4.20) simplifies to

$$\mathcal{O}_{c/\text{uncor}}(x_{gw}, x_{grb}) \approx \frac{R_{gw,grb}^c}{R_{gw}^s R_{grb}^s \Delta t} \mathcal{I}_{\Omega, D_L, \iota} B_{c/s}(x_{gw}) B_{c/s}(x_{grb}), \quad (4.22)$$

where this result is similar to that in Stachie et al. (2020b) with the addition of the coincidence Bayes factors $B_{c/s}(x_{grb})$ and overlap $\mathcal{I}_{\Omega, D_L, \iota}$ with distance and inclination. Let us briefly detail the various fixed values in (4.22) in the case of GW170817-GRB 170817A. The measured expected rate of BNS triggers in O2 was $R_{gw}^s \approx 0.8_{-0.6}^{+1.7}$ /year while the expected rate of *significant* GW-GRB coincidences was $R_{gw,grb}^c \approx 0.14_{-0.11}^{+0.30}$ /year using 90% confidence bounds (Howell et al., 2019). We also take the detected rate of Fermi-GBM short GRBs to be $R_{grb}^s \approx 40$ /year (Howell et al., 2019), although including additional GRB experiments would increase this. The coincidence window we consider is the standard GW-GRB $[-1, +5]$ s centered on the merger time, which means $\Delta t = 6$ s (Abbott et al., 2017d). We also make the practical approximation $P(x_a|\theta_a, \mathcal{H}^c) \approx P(x_a|\theta_a, \mathcal{H}_a^s)$, valid in the coincident parameter regime. This also gives terms like (4.21) a clear interpretation: the coincident hypothesis has more support if the likelihood favors the more constrained coincident prior.

4.3.1 GW170817-GRB 170817A: Using 3D Sky Maps

In order for this statistic to be useful in a low-latency targeted search, it must only use information immediately available to these searches and itself be calculable in low-latency. To this end we restricted ourselves to only using HEALPix sky maps, with the GW sky map including distance information (Górski et al., 2005a; Singer & Price, 2016). We also used the GLADE v2.4 galaxy catalogue as a prior over sky-localization and distance, both with the motivation of being more astrophysically motivated and also to speed up the calculation, although the latter depends on the sky area for a particular joint event (Dályá et al., 2018). While loading even partial catalogue may take some time, to the point of possibly requiring to load this prior, we find evaluating on these galaxies is much faster. We have ignored inclination in this case since this information isn't currently available to these searches using BAYESTAR, but requires posterior samples produced from much higher latency.

We can write the joint sky-localization/distance overlap integral (4.19) as

$$\begin{aligned} \mathcal{I}_{\Omega, D_L} &= \int d\Omega dD_L \frac{p(\Omega, D_L | x_{gw}, \mathcal{H}^c)}{p(\Omega, D_L | \mathcal{H}^c)} \frac{p(\Omega, D_L | x_{grb}, \mathcal{H}^c)}{p(\Omega, D_L | \mathcal{H}^c)} p(\Omega, D_L | \mathcal{H}^c) \\ &\approx \int d\Omega dD_L \frac{p(\Omega, D_L | x_{gw}, \mathcal{H}_{gw}^s)}{p(\Omega, D_L | \mathcal{H}_{gw}^s)} \frac{p(\Omega | x_{grb}, \mathcal{H}_{grb}^s)}{p(\Omega | \mathcal{H}_{grb}^s)} p(\Omega, D_L | \mathcal{H}^c) \end{aligned} \quad (4.23)$$

where we assume the GRB has no distance information and have taken an approximation to use posterior information from individual experiments. We let $p(\Omega, D_L | \mathcal{H}^c)$ be the galaxy catalogue prior, written as

$$p(\Omega, D_L | \mathcal{H}^c) = \sum_{j=1}^{N_{gal}} w_j \delta(\Omega - \Omega_j) \delta(D_L - D_{L,j}) . \quad (4.24)$$

Here w_j are the weights to the j -th galaxy and N_{gal} is the total number of galaxies in the catalogue within the priors. Note if we used uniform weights then $w_j = 1/N_{gal}$, or we could weight each galaxy by its luminosity so that $w_j = L_j / \sum_k L_k$.

We can evaluate (4.23) using (4.24) on a set of discrete number of pixels according to the HEALPix standard with probabilities $P(\Omega_j|x_a, \mathcal{H}_a^s)$. We assume uniform priors on sky localization, then $P(\Omega_j|\mathcal{H}_{gw}^s) = P(\Omega_j|\mathcal{H}_{grb}^s) = 1/N_{pix}$, where N_{pix} is the number of pixels in an individual sky map. This finally gives us

$$\mathcal{I}_{\Omega, D_L} = N_{pix}^2 \sum_{j=1}^{N_{gal}} w_j P(\Omega_j|x_{gw}, \mathcal{H}_{gw}^s) P(\Omega_j|x_{grb}, \mathcal{H}_{grb}^s) \frac{p(D_{L,j}|\Omega_j, x_{gw}, \mathcal{H}_{gw}^s)}{p(D_{L,j}|\mathcal{H}_{gw}^s)} . \quad (4.25)$$

Since the prior $p(D_{L,j}|\mathcal{H}_{gw}^s)$ assumes the event is already detected, the probability should match the distribution of relatively local galaxies. Therefore we used a D_L^2 prior from 10 Mpc to 60 Mpc and then calculated the line of sight distance posterior using the conditional posteriors in `ligo.skymap` (Singer et al., 2016). We stress that the choice of the distance prior needs to follow the probability distribution fairly tightly, as setting the maximum distance too large will artificially increase the overlap in distance. One way to test that this was done properly is ensuring that $\mathcal{I}_{\Omega} \approx \mathcal{I}_{\Omega, D_L}$ using a uniform catalogue, where each pixel contains one galaxy and the distance is sampled from the D_L^2 prior. We also explored the utility of making distance and luminosity cuts, as well as weighting each galaxy by its luminosity (Fishbach et al., 2019), which we found to have improved our measurement.

While the inclusion of the overlap in distance may seem strange since we assume the GRB will not have a distance measurement, we offer another explanation. Considering now just the terms including distance in (4.23), we see that these are measuring whether the posterior over distance supports the galaxy catalogue vs the general D_L^2 prior. This difference might be quite significant when the localization is small, as is the case with GW170817, and we see support of the catalogue in tables 4.1 and 4.2. Also in the situations where there is a distance measurement from afterglow measurements or else, as may be in the case of a detection by *Swift* (Gehrels et al., 2004), this could be included and likely

would improve this calculation.

The results of this method are summarized in table 4.1. If we include galaxies from the whole GLADE catalogue after distance and luminosity cuts, we get $\mathcal{I}_{\Omega, D_L} \approx 226$. This is a notable improvement over using a non-informative (uniform) galaxy catalogue ($\mathcal{I}_{\Omega, D_L} \approx 37.7$), where each pixel contains one galaxy and the distance is sampled from the prior. We should note that the non-informative case is numerically consistent with and analytically simplifies to the spatial overlap integral used in the RAVEN pipeline (Urban, 2016; Cho, 2019) and separately calculated in Ashton et al. (2018b). If we only include NGC 4993 in our catalogue, the galaxy that GW170817 was localized to (Valenti et al., 2017; Abbott et al., 2017c), we find that $\mathcal{I}_{\Omega, D_L} \approx 380,000$. An interesting double check and alternate application of this method is to rank galaxies as hosts using (4.25). If we only consider sky-localization, we find that the host galaxy NGC 4993 is ranked 5th in the catalogue, while including distance information changes this rank to 3rd.

Using (4.22) we can directly compare with the results of Ashton et al. (2018b), albeit using updated rates and sky maps. We will ignore the coincident Bayes factors $B_{c/s}(x_a)$ momentarily as is in Ashton et al. (2018b). We find we get $\mathcal{O}_{c/\text{uncor}}(x_{gw}, x_{grb}) \approx 5,200,000^{+30,000}_{-740,000}$ compared to the previous method $\mathcal{O}_{c/\text{uncor}}(x_{gw}, x_{grb}) \approx 870,000^{+4,900}_{-120,000}$, an improvement of a factor of about 6 due to the inclusion of distance and the galaxy catalogue prior.

4.3.2 GW170817-GRB 170817A: Using Posterior Samples

We can also calculate (4.22) using posterior samples to see how well our statistic performs using the best available information. We used the low-spin posterior samples for GW170817 from Abbott et al. (2019a) since these samples assumed a uniform sky prior. However, posterior samples aren't typically produced for GRBs, so we were limited to using the sky localization and instead used models for the remaining shared parameters. Since inclination information is available from GW samples, we can go through a similar derivation to (4.25)

again with this additional parameter

$$\mathcal{I}_{\Omega, D_L, \iota} = \int d\Omega dD_L d\iota \frac{p(\Omega, D_L, \iota | x_{gw}, \mathcal{H}_{gw}^s)}{p(\Omega, D_L, \iota | \mathcal{H}_{gw}^s)} \frac{p(\Omega, D_L, \iota | x_{grb}, \mathcal{H}_{grb}^s)}{p(\Omega, D_L, \iota | \mathcal{H}_{grb}^s)} p(\Omega, D_L, \iota | \mathcal{H}^c) \quad (4.26)$$

and then using (4.24) we find

$$\begin{aligned} \mathcal{I}_{\Omega, D_L, \iota} &= N_{pix}^2 \sum_{j=1}^{N_{gal}} w_j P(\Omega_j | x_{gw}, \mathcal{H}_{gw}^s) P(\Omega_j | x_{grb}, \mathcal{H}_{grb}^s) \\ &\times \int d\iota \frac{p(D_{L,j}, \iota | \Omega_j, x_{gw}, \mathcal{H}_{gw}^s)}{p(D_{L,j} | \mathcal{H}_{gw}^s) p(\iota | \mathcal{H}_{gw}^s)} \frac{p(D_{L,j}, \iota | x_{grb}, \mathcal{H}_{grb}^s)}{p(D_L | \mathcal{H}_{grb}^s) p(\iota | \mathcal{H}_{grb}^s)} p(\iota | \mathcal{H}^c). \end{aligned} \quad (4.27)$$

Since we again don't expect GRBs to contain distance information and note that the coincidence inclination prior should be identical to that of GRBs, we finally get

$$\begin{aligned} \mathcal{I}_{\Omega, D_L, \iota} &= N_{pix}^2 \sum_{j=1}^{N_{gal}} w_j P(\Omega_j | x_{gw}, \mathcal{H}_{gw}^s) P(\Omega_j | x_{grb}, \mathcal{H}_{grb}^s) \\ &\times \int d\iota \frac{p(D_{L,j}, \iota | \Omega_j, x_{gw}, \mathcal{H}_{gw}^s)}{p(D_{L,j} | \mathcal{H}_{gw}^s)} \frac{p(\iota | D_{L,j}, x_{grb}, \mathcal{H}_{grb}^s)}{p(\iota | \mathcal{H}_{gw}^s)}. \end{aligned} \quad (4.28)$$

We took $p(\iota | \mathcal{H}_{gw}^s)$ as a uniform prior over $\sin \iota$, with ι ranging from 0 to π , while we modeled $p(\iota | D_{L,j}, x_{grb}, \mathcal{H}_{grb}^s)$ similarly except now constrained by the maximum viewing angle given by the distance, as done in (Howell et al., 2019). This relationship between maximum viewing angle and distance is both dependent on the GRB jet model and on the properties of the the specific GRB experiment, and we note that this is somewhat circular as the parameters from the GRB jet model used in Howell et al. (2019) are constrained themselves by GW170817-GRB 170817A. In the future, jet models will be improved and made more robust with additional GW-GRB joint detections (Farah et al., 2020). To check robustness against this jet model, we also calculated the overlap using a constant maximum viewing angle of 25° and found generally similar results, varying roughly on the order of 10%. The conditional GW distance/inclination posterior was calculated using a

Galaxy Catalogue	\mathcal{I}_Ω	$\mathcal{I}_{\Omega,DL}$
Uniform	37.8	37.7
GLADE (un-weighted)	152	217
GLADE ($L_B > .05L_{B*}$)	183	226
Only NGC 4993	210,000	380,000

Table 4.1: Overlap integrals for GW170817-GRB 170817A using the low-latency method from (4.25). We see that including distance, with the exception of uninformative uniform catalogue, improves significance due to galaxy clustering around areas of higher probability. Also including luminosity cuts of GLADE where $L_B > .05L_{B*}$, which gives us near 100% completeness at 60 Mpc (Fishbach et al., 2019), in addition to weighting based on B -band luminosity improves this further.

Galaxy Catalogue	\mathcal{I}_Ω	$\mathcal{I}_{\Omega,DL}$	$\mathcal{I}_{\Omega,DL,\iota}$
Uniform	37.8	37.7	116
GLADE (un-weighted)	152	266	929
GLADE ($L_B > .05L_{B*}$)	183	350	1510
Only NGC 4993	210,000	473,000	729,000

Table 4.2: Overlap integrals for GW170817-GRB 170817A using the posterior sample method from (4.28). As in table 4.1, including distance and using the GLADE catalogue both increases significance. Also, the inclusion of inclination improves the overlap further due to the samples favoring a higher inclination.

4-dimensional kernel density estimation (KDE), checked to match the original posterior samples using resampling. The GW HEALPix sky map was calculated from the samples using `ligo.skymap`.

Since ι is not directly measured by GRB experiments we could instead calculate its contribution using the GW coincident Bayes factor $B_{c/s}(x_{gw})$ by expanding over ι . This should be an equivalent approach because both the overlap integral and coincidence Bayes factor enter (4.22) similarly. We obtained the likelihood in (4.21) by re-weighting the posterior samples using the inverse of the inclination prior as weights. Similar to what is done with the overlap integral, we used a $\sin(\iota)$ prior bounded to less than the maximum viewing angle 25° . This gave us the result of $B_{c/s}(x_{gw}) = 3.5$ for the GW coincident Bayes factor. From table 4.2 we can see that this is roughly the factor between $\mathcal{I}_{\Omega,DL}$ and $\mathcal{I}_{\Omega,DL,\iota}$, so we should be fairly indifferent where we include this contribution. In principle a similar calculation could also be done for other parameters such as mass using a combined BNS,

Galaxy Catalogue	\mathcal{I}_Ω	$\mathcal{I}_{\Omega, D_L}$ (LL)	$\mathcal{I}_{\Omega, D_L}$ (post)	$\mathcal{I}_{\Omega, D_L, \iota}$
Uniform	.03	.03	.81	1160
GLADE (un-weighted)	.03	.03	.04	27.1
GLADE ($L_B > .05L_{B^*}$)	.02	.02	.03	21.0
Only NGC 4993	.02	.04	.03	1.99

Table 4.3: Timing in seconds for the various overlap integrals for GW170817-GRB 170817A, using both the low-latency (LL) method in table 4.1 and posterior method (post) in table 4.2. We see that including distance information doesn’t generally lead to a significant increase in computation time, but including inclination does. This is likely because we are computing numerical integral in a linear fashion rather than parallelize this task, which could significantly improve this. Also, note that the computational times for other coincidences may be much longer due to a greater amount of loaded galaxies, occurring for more distant or poorly localized events.

binary black holes (BBH), and neutron star-black hole (NSBH) mass population model.

We also calculated a contribution from the GRB coincident Bayes factor $B_{c/s}(x_{grb})$ by expanding over the duration T_{90} , which for GRB 170817A was measured at $T_{90} = 2.0 \pm 0.5s$ (Goldstein et al., 2017). We estimated the likelihood as a Gaussian with these values. We used a KDE over T_{90} values from Fermi on HEASARC for the general prior (von Kienlin et al., 2020) and restricted these to $T_{90} < 2.5s$ events for the coincident prior, giving us $B_{c/s}(x_{grb}) = 12.9$. As in the case of GWs, doing similar calculations for other parameters could improve this statistic further.

Using this method with posterior samples we find that $\mathcal{O}_{c/uncor}(x_{gw}, x_{grb}) \approx 35,000,000_{-5,000,000}^{+200,000}$ if we neglect the coincident Bayes factors, while including these gives $\mathcal{O}_{c/uncor}(x_{gw}, x_{grb}) \approx 450,000,000_{-64,000,000}^{+2,600,000}$, leading to an improvement of a factor of 40 and 516 respectively. This implies candidates might be better stratified, as the highly significant joint event GW170817-GRB 170817A is pushed much further out into the foreground compared to the more abundant uncorrelated joint candidates that will typically fail to have overlap in their parameters. We see in table 4.3 that using the posterior method doesn’t lead to greater latency with the exception of including inclination. Although we have included uncertainty due to rates, we note that using different waveform models could give varying results that include distances, similarly to Ashton et al. (2020). We neglect this since our

goal is to show the difference between the different methods presented here rather than a comprehensive measurement of this coincidence.

4.3.3 GBM-GW150914

In addition to the very confident joint detection of GW170817-GRB 170817A, we also wanted to test with a more tenuous association in order to find the limitations of this method. We therefore computed the significance of the first binary black hole (BBH) GW detection GW150914 with the corresponding temporally coincident Fermi-GBM GRB candidate (Connaughton et al., 2016). The astrophysical nature of this GRB has been called into question (Greiner et al., 2016), so we will be wary as we make an evaluation of this joint event using similar methods as in section 4.3.2. We used the GW150914 posterior samples given in the O1 data release (Abbott et al., 2016a) and the GRB candidate sky map produced by the Fermi-GBM team (Connaughton et al., 2016).

Unlike in section 4.3.2, we should be cautious about the use of a galaxy catalogue as a prior due to the amount of incompleteness at these distances. Nonetheless we used a number of luminosity cuts and weights to see how this method fairs, detailed in table 4.4. We can see that in general this coincidence is less significant than GW170817-GRB 170817A, which is expected given the general lack of confidence in comparison. We also see that including distance and inclination information generally improves confidence by roughly the same factors (3 – 5) as in table 4.2. This implies that this information may be shared generally among significant gravitational wave candidates due to a detection bias towards more face-on candidates. Inclination information from the GW alone may possibly not be helpful in distinguishing between joint candidates, although more investigation is needed to confirm this. We also note that due to the much larger number of galaxies contained inside the parameter regime compared to GW170817-GRB 170817A, the computational time was of the order of hours to days compared to just seconds in table 4.3

In conclusion for this event, we note that despite some non-negligible overlap in the

Galaxy Catalogue	\mathcal{I}_Ω	$\mathcal{I}_{\Omega, D_L}$	$\mathcal{I}_{\Omega, D_L, \iota}$
Uniform	7.07	7.03	36.5
GLADE (un-weighted)	5.13	8.02	25.7
GLADE (weighted)	3.58	5.52	20.9
GLADE ($L_B > 2L_{B*}$)	2.98	4.49	18.5

Table 4.4: Overlap integrals for GBM-GW150914 using the posterior sample method similar to table 4.2. In general we see that increasing catalogue information decreases the significance of this coincidence in contrast to GW170817-GRB 170817A, a potential additional argument this may not be a real coincidence. We also used a $L_B > .05L_{B*}$ cut, as was done in table 4.2, but this returned identical values to the weighted full GLADE catalogue due to only to very close galaxies being cut, which had no support from the posterior samples. We note that none of the cuts here give 100% completeness, but in the case of $L_B > 2L_{B*}$ we wanted to show the effects of an extreme cut.

parameters, our statistic should return $\mathcal{O}_{c/\text{uncor}}(x_{gw}, x_{grb}) \approx 0$ here. This is because there is little expectation of a significant joint BBH-GRB detection rate and we could set $B_{c/s}(x_{gw}) = 0$ just due to the measured masses of GW150914, even before including the lower significance of the GRB candidate (Greiner et al., 2016).

4.3.4 Simulation

Lastly, we performed a simplistic simulation of a joint GW-GRB search to quantify the potential improvements compared to previous methods. We modeled the Bayes factor $B_{s/n}(x_{gw})$ as Gaussian distributions in log space with means of 1 and 10 for the random and real sets respectively, and standard deviations of 10^1 and $10^{1.5}$ respectively. For the other Bayes factor $B_{s/n}(x_{grb})$, we instead used power laws of -2 and -1 for the random and real sets respectively, with a lower bound at zero. Finally, we modeled the sky map overlap integral \mathcal{I}_Ω as Gaussians in log space with means 10^{-1} and 10 for the random and real sets respectively, and both with standard deviations of 10. These were chosen to roughly match the outputs expected from pipelines based on our preliminary investigations of doing a more robust simulation. We also used the rates detailed earlier in section 4.3. We didn't use many of the additional terms introduced in this work, such as the additional overlap integrals, the galaxy catalogue prior, or the coincidence Bayes factors; the only

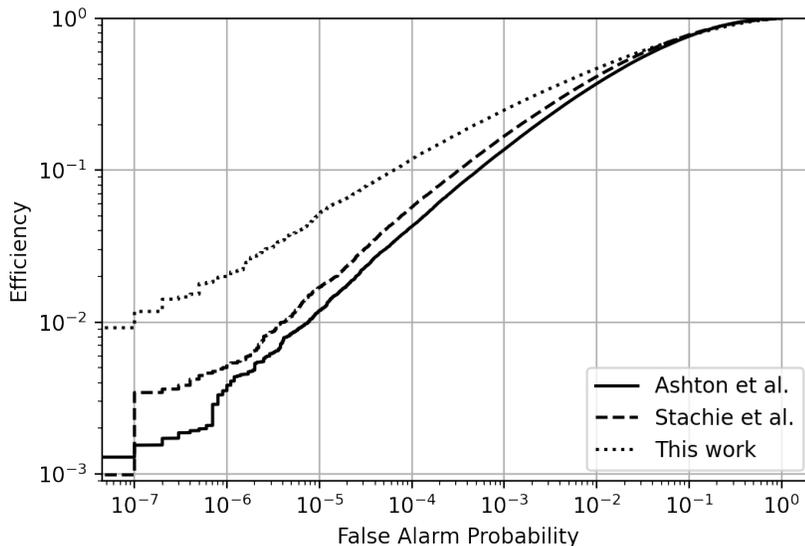


Figure 4.1: Results for the simulations described in section 4.3.4, using the statistics from Ashton et al. (2018b), Stachie et al. (2020b), and equation (4.20) in this work. The horizontal axis is the fraction of random coincidences above a given statistic value while the vertical axis is the fraction of real coincidences above this value. The statistics here vary only whether or not they include the significance from individual candidates, weighted based on rates in the case of equation (4.20). We see generally that given a false alarm probability, our statistic returns more real coincidences.

difference between the statistics tested was the rate-weighted denominator in equation (4.20), compared to the weightless detection Bayes factors Stachie et al. (2020b) and unity Ashton et al. (2018b). This was done in order to show how this simple change could greatly increase sensitivity.

We can see the results of this simulation in figure 4.1. At most given false alarm probabilities, equation (4.20) recovers more real injections compared to previous methods. We also tried a number of reasonable parameters for our injected distributions and this trend tended to hold. This likely occurs because (4.20) tends to downweight random coincidences compared to previous because of this stricter condition of requiring individual candidates to be more significant. In general, we found that (4.20) outperformed the others in cases where the individual detection Bayes factors distributions were more distinct and performed worse when these Bayes factors were not informative, although we argue the

former case is more realistic since random coincidences are dominated by lower significance noise events (\mathcal{H}^{nn}). We should be a bit wary of these results since this simulation is rather simplistic, and we emphasize that more robust simulations are required to definitively establish that this statistic should be preferred to previous iterations.

4.4 Conclusion

We derived a joint ranking statistic in a similar manner to [Ashton et al. \(2018b\)](#) and [Stachie et al. \(2020b\)](#) while adding additional parameters in the overlap integral, using a galaxy catalogue as a prior, and including the previously neglected coincidence-vs-signal Bayes factors, while presenting some evidence that this statistic may perform better than these previous iterations. This statistic placed GW170817-GRB 170817A farther into the foreground by a factor of 6 – 516, based on the method used. We showed that this statistic minimized the significance of the more dubious association of GBM-GW15091, while also present some possible short-comings in this case. We performed a basic simulation that argues in favor of this statistic, although an additional study simulating real and background joint events is needed to confirm this and to determine which additional terms will be most useful for searches.

We acknowledge that some of the new terms introduced with equation (4.20) may not be helpful in distinguishing between real and random coincidences in real searches. We showed in the case of GRB-GW150914 that including distance and inclination information improved the joint significance similarly to GW170817-GRB 170817A, meaning that this improvement may be shared by significant GW detections. We also note that results from the coincident Bayes factors would be shared with other reasonable joint candidates, $B_{c/s}(x_{grb})$ with any other short GRB, and $B_{c/s}(x_{gw})$ with other close face-on candidates, so the addition of these could be ignored and replaced by filter in an actual search such as in [Stachie et al. \(2020b\)](#). If an EM redshift is available, as may be in the case of Swift/XRT

follow-up (Gehrels et al., 2004), there will likely be an improvement in the distance overlap. However in the case of other experiments such as Fermi-GBM (Meegan et al., 2009), the approach detailed in this work is likely the best we can do at the moment. We also note that more care could be given to distinguishing between the coincidence and generic detection terms, as well as towards more informative priors (e.g. using antenna factors in the sky localization). Finally, although there are estimates for the coincident rates in our priors, there are significant uncertainties due the poorly constrained BNS rate (Abbott et al., 2017a; Howell et al., 2019) and GRB jet models (Farah et al., 2020). We argue that while including detection rates improves significance as in section 4.3.4, a spread of values should be used until more observations are done. An alternative would be to use the expected ratio of rates, which would circumvent the uncertainty of detected BNS events.

Chapter 5

Testing Ranking Statistics for Joint Gravitational Wave and Gamma-Ray Burst Searches

This chapter is similar to that in ([Piotrkowski, 2021](#)), where an online record can be found at [10.17307/wsc.v1i1.350](#). Follow-up work will include various other improvements, including a gstLAL injection campaign and improved methods.

5.1 Motivation

As mentioned in chapter 4, joint ranking statistics have been an active area of research and function to categorize candidates based on how likely they are to have a real joint origin rather than be randomly associated. This is an especially difficult task in the sub-threshold regime where more noise events are present. The types of statistics that have been used include modified signal-to-noise ratios ([Nitz et al., 2019](#)), false alarm rates ([Urban, 2016](#);

Cho, 2019), and other generic ranking statistics (Ashton et al., 2018b; Stachie et al., 2020b; Piotrkowski et al., 2022). However, the advantages of each approach has not been well explored, which may mean that searches could be operating at lower sensitivities than otherwise possible.

In this chapter we test these various statistical methods by using a single simulation between LIGO-Virgo GWs and Fermi/GBM GRBs, making both coincidences that resemble real multi-messenger events and those created by random chance. We detail the statistics used in section 5.2 along with the nature of the simulation in section 5.3. We then discuss how efficiently these methods separated out real from random coincidences in section 5.4 and then conclude with a discussion in section 5.5.

5.2 Ranking Statistics

Ideally, a ranking statistic should quantify the significance of a joint event, separating real joint events from random coincidences as best as possible. While there exists many approaches, we limited ourselves to those currently used or could be used in low-latency searches with the LIGO-Virgo-KAGRA collaboration that operate on the minutes or less time scales, as apposed to higher-latency searches that inject gravitational waves signals in the surrounding data to determine significance (Abbott et al., 2020b). We considered the odds ratio from Ashton et al. (2018b)

$$\mathcal{O}_A = \frac{R_{gw,grb}^c \mathcal{I}_\Omega}{\Delta t R_{gw}^s R_{grb}^s}, \quad (5.1)$$

where R_a^s is the rate of real detections for experiment a , $R_{gw,grb}^c$ is the joint BNS-GRB detection rate, Δt is the total time coincidence window, and \mathcal{I}_Ω is the sky map overlap integral. We also considered the odds ratio from Stachie et al. (2020b)

$$\mathcal{O}_S = \frac{\mathcal{I}_\Omega}{\Delta t [1 + B_{n/s}(x_{gw}) + B_{n/s}(x_{grb}) + B_{n/s}(x_{gw})B_{n/s}(x_{grb})]}. \quad (5.2)$$

Here $B_{n/s}(x_a)$ is the detection Bayes factor for experiment a , weighing whether the data x_a is likely from a noise transient (\mathcal{H}_a^n) or real signal (\mathcal{H}_a^s) as

$$B_{n/s}(x_a) = \frac{P(x_a|\mathcal{H}_a^n)}{P(x_a|\mathcal{H}_a^s)} \quad (5.3)$$

where $P(x_a|\mathcal{H}_a^n)$ is probability of the data x_a given that it is noise and $P(x_a|\mathcal{H}_a^s)$ is probability of the data x_a given that it is a real signal. We also used the odds ratio from (4.20)

$$\mathcal{O}_P = \frac{R_{gw,grb}^c \mathcal{I}_\Omega}{\Delta t [R_{gw}^s R_{grb}^s + B_{n/s}(x_{gw}) R_{gw}^n R_{grb}^s + B_{n/s}(x_{grb}) R_{gw}^s R_{grb}^n + B_{n/s}(x_{gw}) B_{n/s}(x_{grb}) R_{gw}^n R_{grb}^n]} \quad (5.4)$$

where R_a^n is the rate of noise detections for a , and we have omitted some overlap terms and coincident Bayes factors. Note that all these statistics so far have similar derivations but varying levels of approximation, meaning testing these against each other will determine whether additional terms regarding these detection Bayes factors $B_{n/s}(x_a)$ improve sensitivity. Another statistic of interest is the joint false alarm rate (FAR) used by RAVEN from (3.1)

$$FAR_c = FAR_{gw} R_{grb}^s \Delta t / \mathcal{I}_\Omega \quad (5.5)$$

where FAR_{gw} is the GW false alarm rate.

5.3 Simulation

Simulating multi-messenger searches is difficult due to the unique task of trying to combine two disparate experiments in a coherent manner, especially when real coincidences should have the same underlying parameters (Ashton et al., 2018b). This is still largely an unsolved problem with BNS systems due to the lack of understanding of BNS systems, especially with how underlying parameters affect GRB emission. For instance, there is a lack of strong

constraints on how variations in viewing angle affects the luminosity of sGRBs (Farah et al., 2020). Also, other seemingly straightforward relations such as how the neutron star masses affect GRB luminosity remain elusive. For simplicity in the case of real coincidences, we assumed many of the underlying parameters were independent from one another, such as inclination and masses, while we set others, such as timing and sky location, to be the same between each experiment. In addition, all the parameters for random coincidence were always set to be independent. This approach is not strictly correct, but is likely the best we can do until we better understand these systems.

In order to test these statistics against each other in the most fair manner possible, we simultaneously calculated these for the same set of events. This meant computing several data products to be shared between the various statistics, such as sky maps and detection Bayes factors. We also needed to establish what characterizes a real multi-messenger detection versus a coincidence from random association. As in Ashton et al. (2018b), we define real joint events as having consistent shared parameters (e.g. timing, sky position) *and* both events being detectable. This latter condition is required in order to relieve the ambiguity with how to handle sub-threshold events, i.e. events that could have real artificial origins but aren't significant enough to be confident detections. Many of these sub-threshold events will inevitably be indistinguishable from noise, since one or more events may be undetectable for a real joint event due to distance/inclination/etc. We explore this by considering two datasets: one that only includes individually significant events (threshold), and another with events that may not be significant individually but their joint significance could be above a joint detection threshold (sub-threshold). We note that as in chapter 4, the random coincidence datasets will be made up of three types of false coincidences: two real candidates that are in reality unrelated, one real and one noise candidate, and two noise candidates. These were proportioned according to the rates we soon define.

Based on these definitions we simulated our various event datasets to be similar to the

third operating run (O3) of the LVK. For our BNS events, we created a threshold dataset by setting a network signal-to-noise ratio (SNR) (between the Livingston, Hanford, and Virgo detectors) threshold of 12 while our sub-threshold dataset had a threshold of 8. We used tools of `ligo.skymap` to create a set of injection parameters, using the BNS astrophysical set. We then used `ligo.skymap` to create GW localizations using its method of matched filtering, simulating a search done by the LVK. For the GRBs, we drew from given log likelihood ratio (LLR) distributions, with minimum thresholds of 25 for the threshold set and 5 for the sub-threshold set (Stachie et al., 2020b; Goldstein et al., 2019). We modeled the GRB sky probability for each event as a Gaussian on the sky as recommended by Fermi/GBM (Connaughton et al., 2015) using the tools developed for RAVEN in `gwcelery` (Collaboration, 2018a). The standard deviation for these Gaussians was given by the LLR (Goldstein et al., 2019) and we then offset from the true source by using the probability sky map generated at this true source as weights. Again, we acknowledge that there should be some correlation between the LLR of the GRB and the injection parameters, but at the time of writing this exact relationship isn't well understood. In addition, we chose to restrict our coincidence dataset to only include GW candidates with viewing angles of less than 25° (Howell et al., 2019).

We used the rates of a multi-messenger pipeline during O3, although we note that detection rates are quite uncertain at this time. We set the coincident rate $R_{gw,grb}^c$ to $.58/yr$, the GW BNS rate R_{gw}^s to $5.3/yr$, the short GRB rate R_{gw}^s to $33/yr$, the rate of noise GW candidates to $1/hr$, the rate of noise GRB candidates to $1/10000s$, and the total time coincident window Δt to $6s$ (Howell et al., 2019; Abbott et al., 2017b; Collaboration, 2018a). We ran these simulations for a period of roughly 600 years in order to accumulate enough events to start to represent the underlying distributions.

We calculated the sky map overlap integral using (3.3) and computed our detection Bayes factors similarly to Stachie et al. (2020b) by using kernel density estimation (KDE) over intermediary detection statistics, LLR for the GRBs and the detection Bayes factor

given in the GW sky maps. To get FAR_{gw} for equation (5.5), we monotonically mapped from these intermediary detection Bayes factor in the GW sky maps. This approach is not strictly correct since a FAR calculated by a pipeline would emerge from a comparison of the real surrounding background while the detection Bayes factors we used here assumes a Gaussian noise background. Having a more realistic calculation of this FAR_{gw} would give more confidence to the results involving equation (5.5).

5.4 Results

Having computed our various statistics in our simulated multi-messenger search, we now can analyze their relative performance. We do this by a simple receiver operating characteristic (ROC) analysis, comparing the number of real coincidences (true positives) above a given threshold to the allowed number of random coincidences (false positives). This allows us to test every statistic regardless of their mathematical form and scale, as well as letting us see how they perform at different significance regimes.

We can see our results of our threshold dataset in figure 5.1 and sub-threshold dataset in figure 5.2. We see that in the case of highly significant individual events with the threshold dataset, statistics with the additional complexity of individual event significance have much more sensitivity, especially in the case of equation (5.4). This is likely because the detection Bayes factors $B_{n/s}(x_a)$ are more informative here due to largely separated individual real and noise populations, consistent with the observations in chapter 4. Furthermore, we also see this same stratification in the sub-threshold dataset to a lesser effect since the detection Bayes factors are less informative. We see that the joint FAR from equation (5.5) performs the best in the sub-threshold case, although it's not clear whether this is due to the peculiarities with how we calculated FAR_{gw} . Worth noting is that like equation (5.4), equation (5.5) uses the significance of the GW candidate but weights this much more highly, meaning that even in this sub-threshold data there may be more of a difference

between our real and noise GW datasets than expected. Regardless, we should be wary of the results involving equation (5.5) until this is better understood.

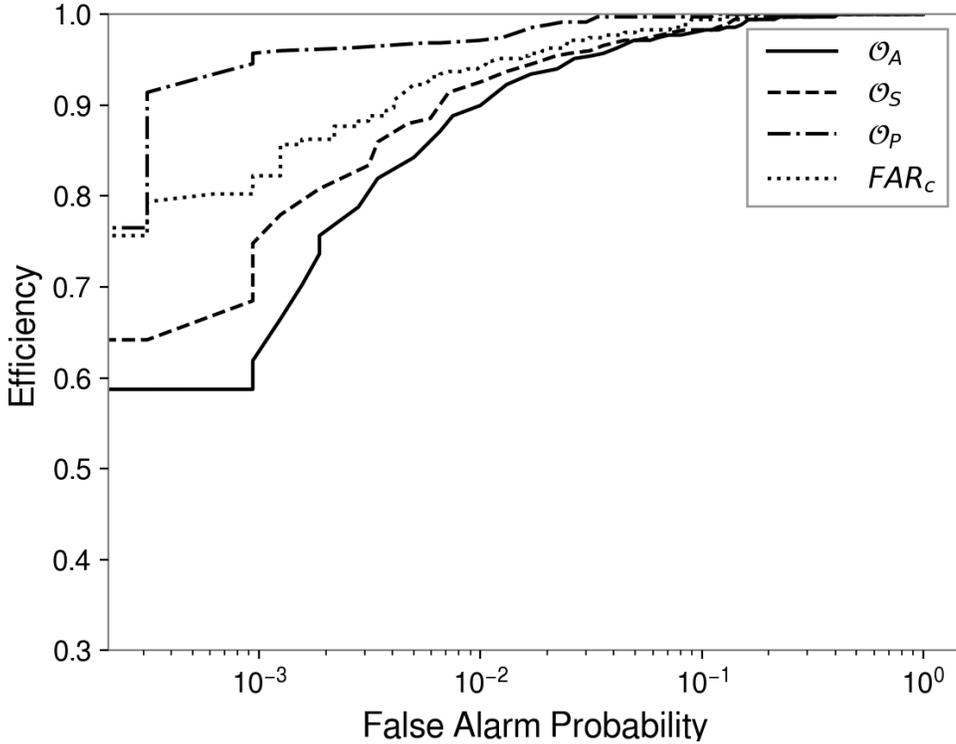


Figure 5.1: Simulation results for the threshold dataset described in section 5.3, testing the statistics from section 5.2. The horizontal axis is the fraction of random coincidences above a given statistic value while the vertical axis is the fraction of real coincidences above this value. We see that the additional terms in Piotrkowski et al. (2022) (\mathcal{O}_P) and Stachie et al. (2020b) (\mathcal{O}_S) using the individual significance of events resulted in higher sensitivity, especially for \mathcal{O}_P , compared to Ashton et al. (2018b) (\mathcal{O}_A) which does not include these. We note that higher false alarm probabilities correspond to higher efficiencies in finding real events due to the fact that the threshold to consider events of both type is lower. This is why GW searches usually have higher thresholds: although the number of real detected events is low, the number of false positives is also low in order to have high confidence in the events that do pass.

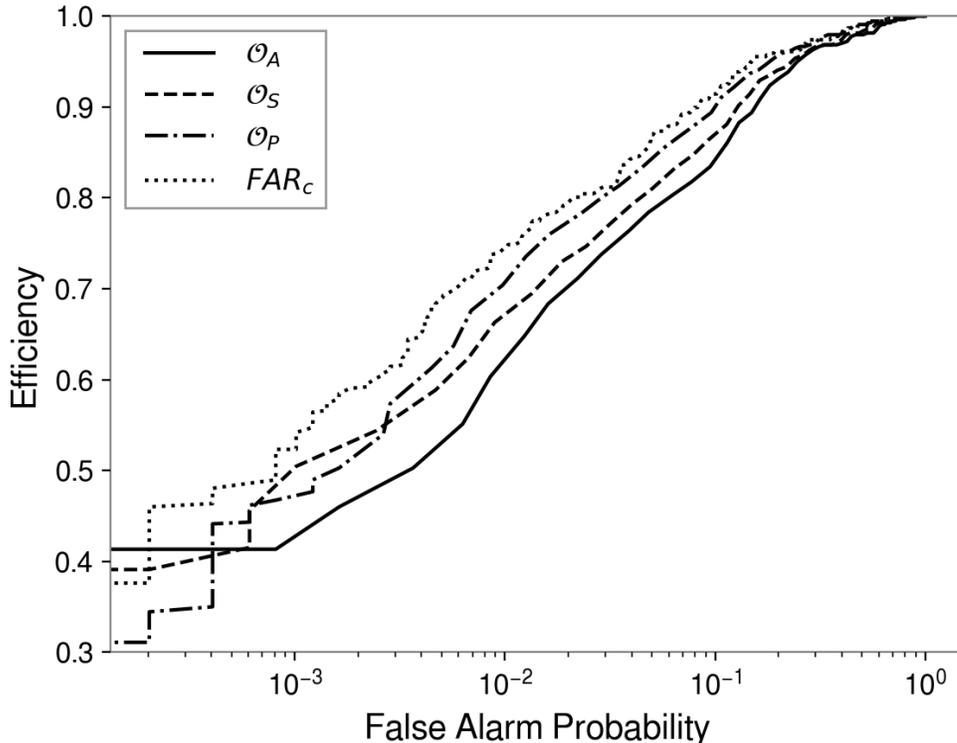


Figure 5.2: Simulation results for the sub-threshold dataset described in section 5.3, testing the statistics from section 5.2, similar to figure 5.1. We see that the additional terms in Piotrkowski et al. (2022) (\mathcal{O}_P) compared to Ashton et al. (2018b) (\mathcal{O}_A) and Stachie et al. (2020b) (\mathcal{O}_S) still resulted in higher sensitivity at lower significance but this reversed at the highest significance. We also see that FAR_c from equation (5.5) performed the best at nearly all regimes but it’s not clear whether this is due to the peculiarities of the calculation of FAR_{gw} in section 5.3.

5.5 Discussion

In this work we created a simulated multi-messenger search between GWs from the LVK collaboration and GRBs from Fermi/GBM, the purpose being to test various joint ranking statistics in section 5.2. We described the nature of this simulation and the various methods to produce the necessary data products. Then we discussed the results for both a threshold and sub-threshold dataset, where we showed some evidence that statistics including the significance of the individual candidates can give more sensitivity to searches.

A major criticism of this work is the fundamental data used and the methods used

to handle this data. As opposed to creating events while using LVK flavored Gaussian noise, a much more realistic GW simulation would be to do an injection campaign with real LVK noise, returning both known real injected events as well as noise transients to create the background. This will ultimately mean using robust search pipeline such as GstLAL (Sachdev et al., 2019; Cannon et al., 2021). This also would give us a more realistic FAR_{gw} since this is a natural by-product for such a method. Using a GW search with real noise would better simulate realistic GW search conditions and give more weight to any conclusion made with that data.

The results in this study are also limited by the fact that they would not apply to a different GRB experiment such as *Swift*, which has a much smaller localization area and therefore a different sky map overlap distribution with the LVK (Gehrels et al., 2004). Since we omitted the additional overlap terms of distance and inclination introduced in prior statistics (Piotrkowski et al., 2022; Ashton et al., 2020), we cannot give any indication of whether these increase sensitivity. Worth noting is that while equation (5.4) and (5.2) seem to give more sensitivity than equation (5.1), the need of a detection Bayes factor that requires KDEs and therefore all the events after an observing run means these aren't as helpful in low-latency searches unless a different method to approximate a detection Bayes factor is used.

As mentioned in section 5.3, the inability to completely produce a set of coherent GW and GRBs just from a set of injection parameters is currently a flaw in simulating joint BNS-GRB events. Unfortunately this will require better modeling and future constraints, both helped for the most part by more multi-messenger detections. For instance, just the relationship between viewing angle and a short GRB's luminosity still requires numerous more detections to strongly constrain this and rule out models (Farah et al., 2020). Regardless, more observations will continue to grow the field of multi-messenger astronomy and further constrain fundamental physics.

Chapter 6

Contributions to the Third

LIGO-Virgo-KAGRA Operating Run

6.1 Joint FAR Derivation

We will be discussing a derivation of the joint FARs introduced in section 3.3. We note that both Urban (2016) and Cho (2019) worked through similar but independent derivations to ours, but each came to different conclusions or made assumptions we disagree with. Since our goal is to analyze whether GW candidates should receive additional follow-up due to proximity to external triggers, we can say that the joint FAR (FAR_c) should have the form

$$FAR_c = FAR_{gw}\Lambda \tag{6.1}$$

where FAR_{gw} is the GW FAR and $\Lambda(x_{gw}, x_{em})$ is some association statistic that may depend on the data of the GW candidate and/or the external trigger.

Since a FAR is a uniform distribution, the CDF of a FAR should be linear and have the property

$$N(FAR < FAR^*) = FAR^* \times T \tag{6.2}$$

where N is the expected number of events of higher significance than a chosen FAR^* , and T is the total co-observing time. This property gives us the expected background shown in our various simulations, such as in figure 6.2. Note that the joint FAR using just timing information will be a subset of the GW FAR, reducing the number by a factor of the expected number of false coincidences per GW candidate $\langle N_{uncor} \rangle$. This can be written as

$$N_c = N_{gw} \langle N_{uncor} \rangle \quad (6.3)$$

where N_{gw} is the number of expected GW candidates and N_c is the expected number of background joint candidates. Therefore, if FAR_{gw} has the property (6.2), then for FAR_c to also satisfy (6.2) means that

$$\Lambda(t_{gw}, t_{em}) = \langle N_{uncor} \rangle . \quad (6.4)$$

If this wasn't true, Λ would introduce a noticeable skew on the entire distribution and would no longer fall on the expected line defined in (6.2). Another more intuitive explanation can be found by analyzing any plot containing both the GW FAR and joint FAR, such as figure 6.2: the proportion of the maximum GW candidates to the maximum joint candidates is given by the expected value of false coincidences per GW candidate, meaning to stay on the expected line we must decrease the FAR by the same amount. This implies the relationship between the maximum FAR thresholds

$$FAR_{c,thresh} = FAR_{gw,thresh} \langle N_{uncor} \rangle \quad (6.5)$$

which is given directly from (6.2) and (6.3), and also nicely mirrors the definition of the joint FAR we are trying to establish.

The expected number of false coincidences per GW candidate can be calculated by the

expected value of the Poisson distribution for k associated candidates

$$\langle N_{uncor} \rangle = \sum_{k=1}^{\infty} k \cdot P(X = k) \quad (6.6)$$

$$= \sum_{k=1}^{\infty} k \frac{(R_{em}\Delta t)^k}{k!} e^{-R_{em}\Delta t} \quad (6.7)$$

$$= R_{em}\Delta t \sum_{k=1}^{\infty} \frac{(R_{em}\Delta t)^{k-1}}{(k-1)!} e^{-R_{em}\Delta t} \quad (6.8)$$

$$= R_{em}\Delta t \quad (6.9)$$

due to the normalization of the Poisson distribution, using the Taylor series expansion of an exponential. This means the association statistic with timing information should be written as

$$\Lambda(t_{gw}, t_{em}) = R_{em}\Delta t . \quad (6.10)$$

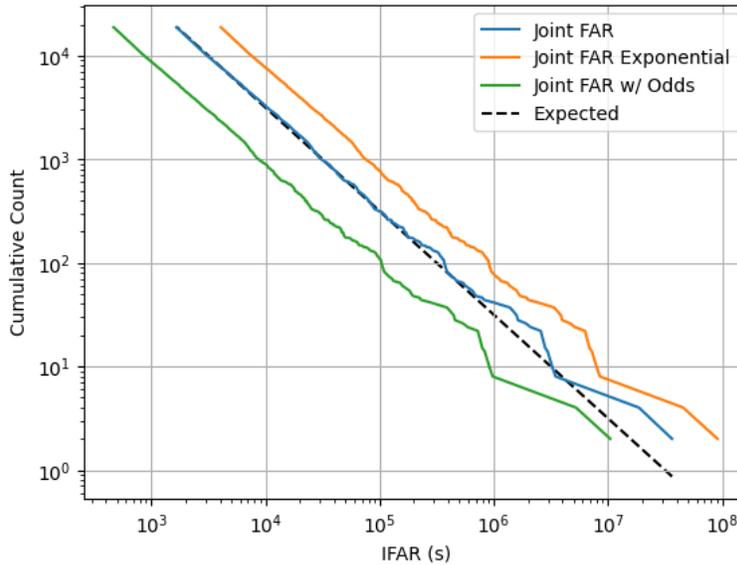


Figure 6.1: Simulation of using the various methods for the temporal joint FAR described in section 6.1 (blue), including (6.11), the exponential probability (orange) suggested in Urban (2016), and a version using the odds ratio from Cho (2019) (green). The expected value approach in (6.10) was the only to succeed when the number of false coincidences per GW candidate was high ($R_{em}\Delta t = 2$).

Let us next discuss some previous efforts to establish $\Lambda(t_{gw}, t_{em})$ and the issues with them. As noted in [Urban \(2016\)](#), the probability that a candidate has any type of coincidence is

$$P(\mathcal{H}_{uncor}) = 1 - P(\mathcal{H}_0) = 1 - e^{-R_{em}\Delta t} \quad (6.11)$$

where $P(\mathcal{H}_0)$ is the probability there is no coincidence, R_{em} is the rate of external candidates, and Δt is the coincident time window. This was believed to be the correct term for $\Lambda(t_{gw}, t_{em})$ but an approximation for high significance events was made, reducing this down to [\(3.1\)](#) since

$$1 - e^{-R_{em}\Delta t} = 1 - [1 - R_{em}\Delta t + \mathcal{O}((R_{em}\Delta t)^2)] \quad (6.12)$$

$$\approx R_{em}\Delta t \quad (R_{em}\Delta t \ll 1) \quad . \quad (6.13)$$

Also, [Cho \(2019\)](#) asserted that $\Lambda = P(\mathcal{H}_{uncor}|x_{gw}, x_{em})/P(\mathcal{H}_c|x_{gw}, x_{em})$, the odds ratio we have derived in chapter [4](#). Again we note this assertion is true in the high significance limit since in this regime

$$P(\mathcal{H}_c|x_{gw}, x_{em}) \approx 1 - P(\mathcal{H}_{ss}|x_{gw}, x_{em}) \quad (6.14)$$

because \mathcal{H}_{sn} , \mathcal{H}_{ns} , and \mathcal{H}_{nn} are negligible. Therefore

$$\mathcal{O}_{ss/c} \approx \frac{P(\mathcal{H}_{ss}|x_{gw}, x_{em})}{P(\mathcal{H}_c|x_{gw}, x_{em})} \approx \frac{P(\mathcal{H}_{ss}|x_{gw}, x_{em})}{1 - P(\mathcal{H}_{ss}|x_{gw}, x_{em})} \approx P(\mathcal{H}_{ss}|x_{gw}, x_{em}) \approx R_{em}\Delta t/\mathcal{I}_\Omega \quad . \quad (6.15)$$

We performed simulations to see which of these methods works in different significance regimes and we can see these results in [figure 6.1](#). In the high significance regime all these approaches are the same, as expected, but when the expected number of false coincidences

is fairly large only (6.4) correctly describes the joint FAR. We also find that this continues to be true in section 6.3.2, where (6.4) again appears to work well when $R_{em}\Delta t \approx 2$.

Although there are issues with using the odds ratio, it does give insight into likely one of the best ways to incorporate sky localization information: the sky map overlap integral \mathcal{I}_Ω . In all simulations we have performed, regardless of the experiments or searches involved, we find that $\langle \mathcal{I}_\Omega \rangle \approx 1$. This means the inverse FAR is convolved by a random variable with expected value of 1, largely just rearranging high significant candidates. We see that this property is relevant to the inverse FAR rather than the FAR in figure 6.3. Whether overlap integral is the correct or even the best term to include sky localization information isn't clear and still an open problem, and we discuss issues including sky localizations in sections 3.4 and 6.2. We note the sky localization term in Urban (2016) is likely incorrect due to not satisfying (6.4). Putting this all together, we can say that the general joint FAR can be written as

$$FAR_c = FAR_{gw}R_{em}\Delta t/\mathcal{I}_\Omega \tag{6.16}$$

where the overlap integral \mathcal{I}_Ω is calculated using (3.3) or (7.2).

6.2 Correcting Skewness in the Joint FAR for Offline Searches

We have noted a limitation with the current methods of including sky maps into the FAR calculation as described in section 3.4 and shown in numerous figures (figs 3.2, 3.4, and 7.3). To summarize, the issue is trying to convolve a random variable drawn from a uniform distribution (FAR_{gw}) with a variable drawn from an unknown distribution (I_Ω), trying to get a result that also as if drawn from a uniform distribution (FAR_c). Since I_Ω is highly sensitive to each instrument's configurations and the nature of the background in each, the distribution must be carefully computed for the specific search being carried out.

Although it's unlikely that this skew can be addressed for online searches, we can

correct this for offline searches once we have the full set of FARs. This can be done similarly to [Stachie et al. \(2020b\)](#) by mapping the FARs of false coincidences that represent the background to the expected background. This background of false coincidences can be created in real searches by time sliding outside the the expected time windows. The expected background is easy to compute given the total time the search took place. In our simulations we binned up each distribution, and given a joint FAR value, chose the corresponding remapped value with the same array index. To create these arrays for remapping we ran a set of simulations for each search and then another set of simulations (with the same parameters) to test this method.

Results of this remapping method can be found in figures [6.4](#) and [6.5](#) where we can see that each type of search was corrected within acceptable bounds. However, this remapping method is currently not well suited for online searches. Firstly, if there are meaningful differences between the simulations and actual search pipelines, this method will introduce a bias and additional skewness, defeating the point of including this in the first place. Secondly, fixing issues may be more complicated due to the need to rerun simulations. For instance, if Fermi or the LVK introduce changes that improve sky localization, our simulations would need to be modified to include these and a new version of RAVEN would need to be released.

6.3 GW-FRB Search w/ CHIME

Alongside searching for coincident GRBs and neutrino bursts, work was done in O3 to find fast radio bursts (FRB) coincident with GWs.

6.3.1 Cross-validation with PyGRB and X-pipeline

Rather than perform the full GW-FRB search alone, RAVEN was instead used in conjunction with both PyGRB ([Abbott et al., 2019b](#)) and X-pipeline ([Sutton et al., 2010](#)). The goal

of each of these pipelines is to look for plausible, sub-threshold GW candidates around given trigger times. These two pipelines primarily analyze in a much higher latency than RAVEN, although by using injections and measures of coherence, they are typically more sensitive. Hence, RAVEN was used to check whether any *previously known* GW candidates were coincident with the given CHIME trigger list over the entirety of O3a.

A similar calculation was performed to (3.1) except the rate of FRBs from CHIME R_{frb} was approximated as the number of triggers over the time of the search ($\approx 2/day$). An extended time window was used (CBC: $[-2, +10]$ s; GW Burst: $[120, +600]$ s, centered on the merger time) compared to a GW-GRB search due to our ignorance of the underlying physics of the source.

This search, detailed in (Abbott et al., 2022), did not find any significant joint candidates. RAVEN only found one CBC joint candidate, but a poor sky map overlap integral of $\approx 10^{-3}$ meant it was not significant enough to warrant further followup. Eight GW Burst candidates were found as well, but all had insufficient joint FARs and most had FRB candidates at distances beyond GW detection.

6.3.2 Using Extended Search Windows

There has been speculation that the delay between GW and FRB candidates could be much longer to account for various emission models (Platts et al., 2019). Using an extended search window of $[-2, +24]$ hr, a joint candidate involving GW190425 has been claimed (Moroianu et al., 2021). We investigated by running RAVEN over all of O3a using this extended window.

The overall results of this search can be found in figure 6.7. We see that joint FAR and GW FAR populations are largely consistent, meaning there is insufficient evidence of there being a significant population of GW-FRBs associations. We also see that including sky maps generally reduced the significance of these joint events further, consistent with random associations. Another seemingly strange result is that there were about double (≈ 2.05) the

Superevent	FRB ID	FAR_c (Hz)	FAR_c w/ sky (Hz)	\mathcal{I}_Ω	Time Delay (hr)
S190408an	FRB36025223	5.026e-18	1.179e-11	4.265e-07	+9.13
	FRB36090381	5.026e-18	8.065e-20	62.32	+23.01
GW190425	FRB37888771	8.114e-13	7.835e-13	1.036	+2.5
	FRB37948191	8.114e-13	4.402e-12	0.184	+9.79
	FRB37984665	8.114e-13	4.619e-12	0.176	+17.16
S190519bj	FRB39275710	1.019e-08	9.992e-10	10.20	+2.24
	FRB39278375	1.019e-08	1.640e-09	6.217	+2.55
	FRB39278417	1.019e-08	1.640e-09	6.217	+2.58
S190630ag	FRB43098433	2.907e-13	7.711e-12	0.0377	-1.04
	FRB43123299	2.907e-13	1.070e-12	0.272	+4.13
	FRB43168950	2.907e-13	4.825e-12	0.060	+14.97
S190727h	FRB45920865	2.791e-10	1.135e-10	2.460	+5.19
S190814bv	FRB49083177	4.117e-33	1.030e-22	3.995e-11	+11.38

Table 6.1: Results of significant individual joint candidates of extended window CBC-FRB search detailed in section 6.3.2. This list only includes events that pass the normal publishing threshold of FAR_c w/ sky $< 3.71\text{e-}8$ Hz (1/year). Note that the only GW candidate that is likely from a BNS is GW190425 while all the other candidates have masses consistent with BBH systems. In addition, every skymap overlap with any positive evidence (greater than 1) has been boldened.

joint GW-FRB candidates than GW candidates, meaning most FRB candidate had more than one associated GW. We can see that this was expected by computing the number of expected random associations

$$\langle N_{uncor} \rangle = R_{frb} \Delta t = 1.87/day \cdot 1.08 \text{ day} = 2.03 . \quad (6.17)$$

We also can see the candidates with joint FARs (when including skymaps) of less than $1/yr$ in table 6.1. We notice that every joint candidate includes a GW candidate that was significant enough to be an issued a public alert on its own, meaning there were no sub-threshold GW candidates elevated due to an adjacent FRB (note that the search windows were sufficiently wide so that the temporal joint FAR *decreased* significance compared to the GW-only FAR). Additionally, there was only one GW candidate that had masses consistent with a neutron star, GW190425. However, the most plausible associated FRB with this (FRB37888771) results in $I_\Omega \approx 1$ (computed from (3.3)), meaning there is only neutral at best evidence to associate these two. Although some mechanisms have been

proposed for a BBH system to produce an FRB (Platts et al., 2019), these require charged black holes which may not exist in nature. Overall, we can say there were no confident GW-FRB associations during O3a and insufficient evidence to support a GW-FRB subthreshold population.

6.4 GWCosmo

One important result of the second and third LIGO-Virgo-KAGRA operating runs has been an independent measurement of the Hubble constant H_0 (Collaboration et al., 2017; Fishbach et al., 2019). This can be very roughly understood using Hubble’s Law, which draws a connection between the redshift z due to a recessional velocity and proper distance D as

$$H_0 = \frac{zc}{D} \tag{6.18}$$

where c is the speed of light. Since the speed of light is known and the distance can be measured from a GW, this leaves the redshift to be determined. The redshift can be accounted for by two methods: 1.) a counterpart method that assumes a galaxy has been identified and hence the redshift is known or 2.) a statistical method that uses numerous galaxies with known redshifts. This second method relies on galaxy catalogues, such as GLADE (Dalya et al., 2018), but suffers from excessive computational time if the same computation is made for every galaxy.

6.4.1 Partial Loading of Galaxy Catalogues

One method of reducing the computational cost of the statistical method is to only use galaxies within the GW localization. We designed a method to do this, which works in the following manner:

1. load the GW sky map, sort by probability, and find the indices of every pixel up to a predetermined percentile limit (up to 99.9% of total probability by default)

2. load just the coordinates of each galaxy and only use if it falls within an accepted GW pixel

This algorithm can be very effective at reducing computational cost when the GW localization is small. For example, if we only use the 99.9% localization area of GW170817, this reduces the required pixels to below 1% of the sky. This could be improved by additionally adding distance as was done in chapter 4. In the example of GW170817, adding a distance cut of < 60 Mpc meant only 83 galaxies remained to give a sky map overlap integral numerically equivalent to using the entire catalogue.

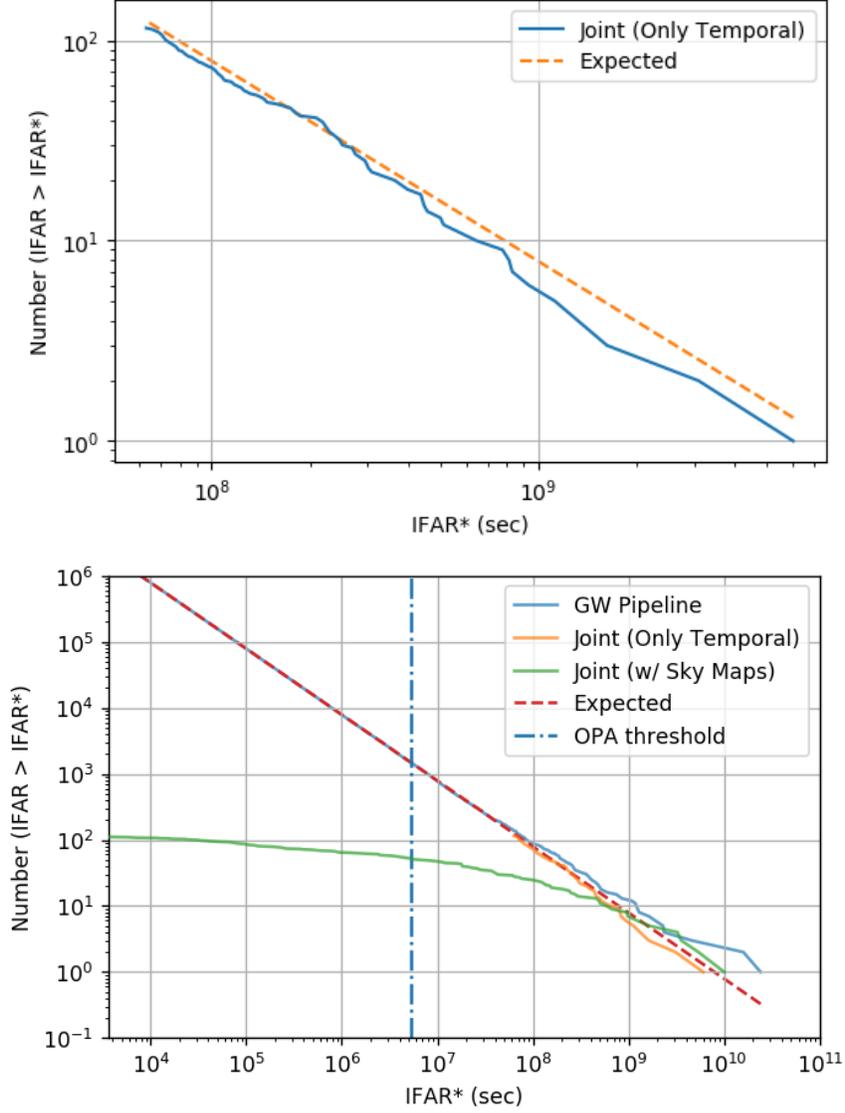


Figure 6.2: The cumulative counts vs inverse FAR (IFAR) for the joint FAR in (3.1). We simulated with 250 years worth of triggers, each given a random time, and looked for coincidences using the time window $[-1, +5]$ s centered on the GW time. For every coincidence we calculated the joint far both using the sky map overlap integral (w/ Sky Maps) and excluding it (Only Temporal). We used $FAR_{GW,thresh} = 1/hr$ and $R_{grb} = 300/yr$ to get the total number of triggers for each type (Collaboration, 2018a). The open public alert threshold (OPA) displayed here is $1/(2\ months)$ (Collaboration, 2018b). We see poor convergence of the joint FAR with sky map information at low IFAR due to low sky map overlap outliers ($I_{\Omega} \approx 10^{-7} - 10^{-5}$). These come about because the GW sky maps used in this study are astrophysically motivated and typically well localized, leading to very little overlap if the two sky maps are inconsistent. However, sky maps made from noise triggers are often poorly localized and therefore we expect much less inconsistent outliers of this severity in real analysis, which should lead to better convergence of this joint statistic.

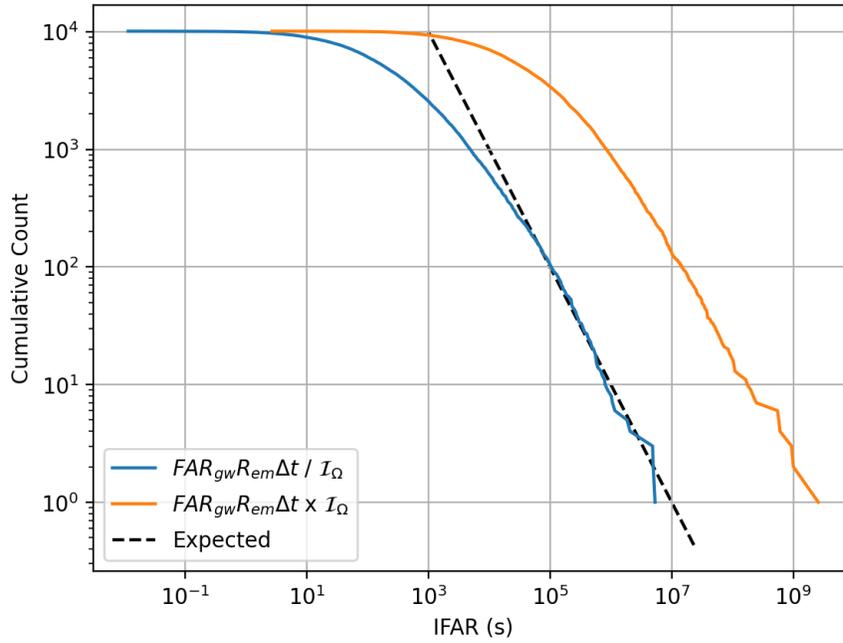


Figure 6.3: The cumulative counts vs inverse FAR (IFAR) using both an inverse (blue) and a proportional sky map overlap (orange). The overlap was sampled from a Gaussian in log space with an expected value of 1 while the inverse overlap had an expected value of ≈ 200 . We see that the inverse overlap generally gives the correct distribution, implying this expected value is crucial for the inverse FAR rather than the FAR.

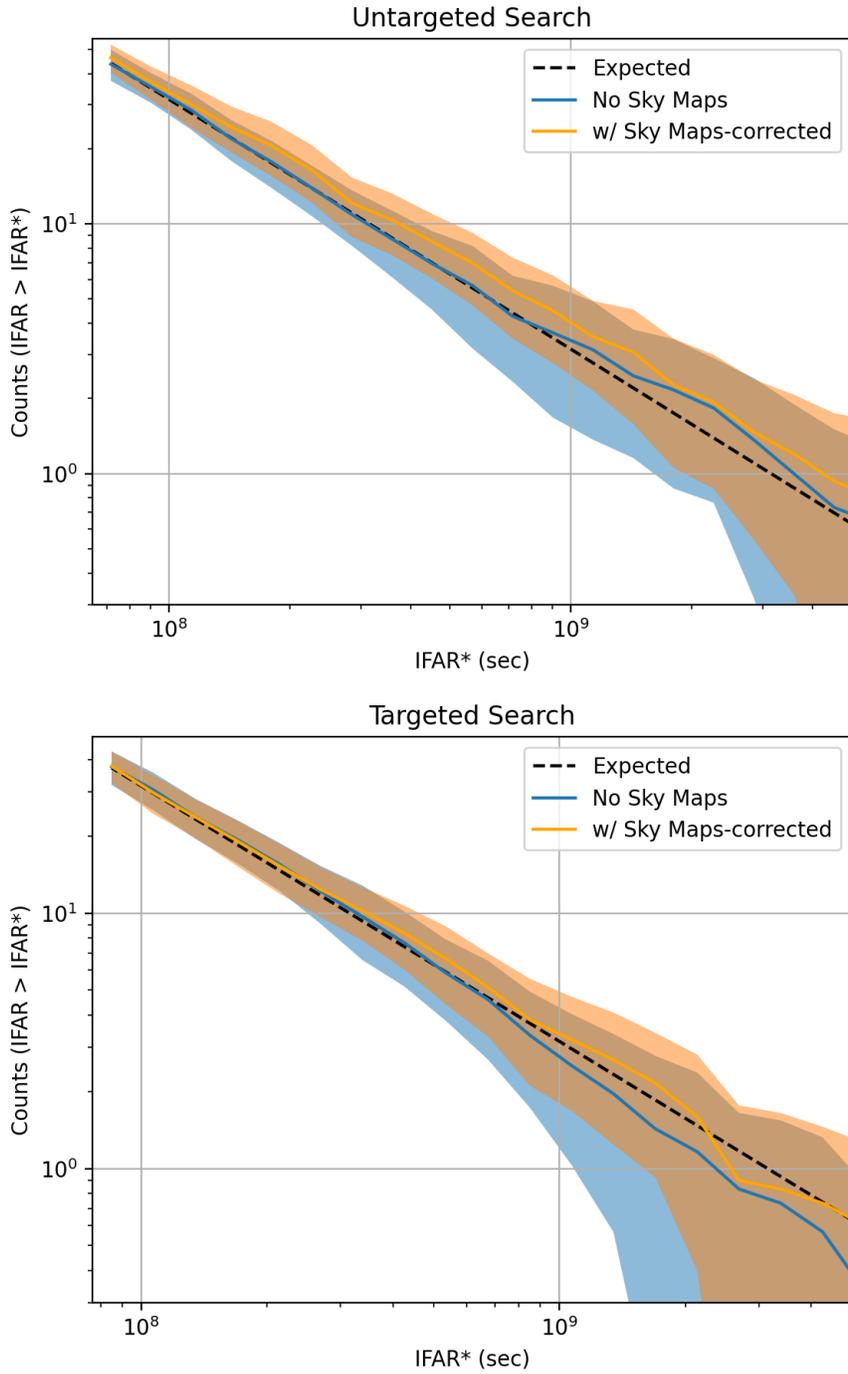


Figure 6.4: Corrected joint FAR with Fermi/GBM-like GRB candidates (same simulations as in figure 3.2) using the simple remapping as described in section 6.2.

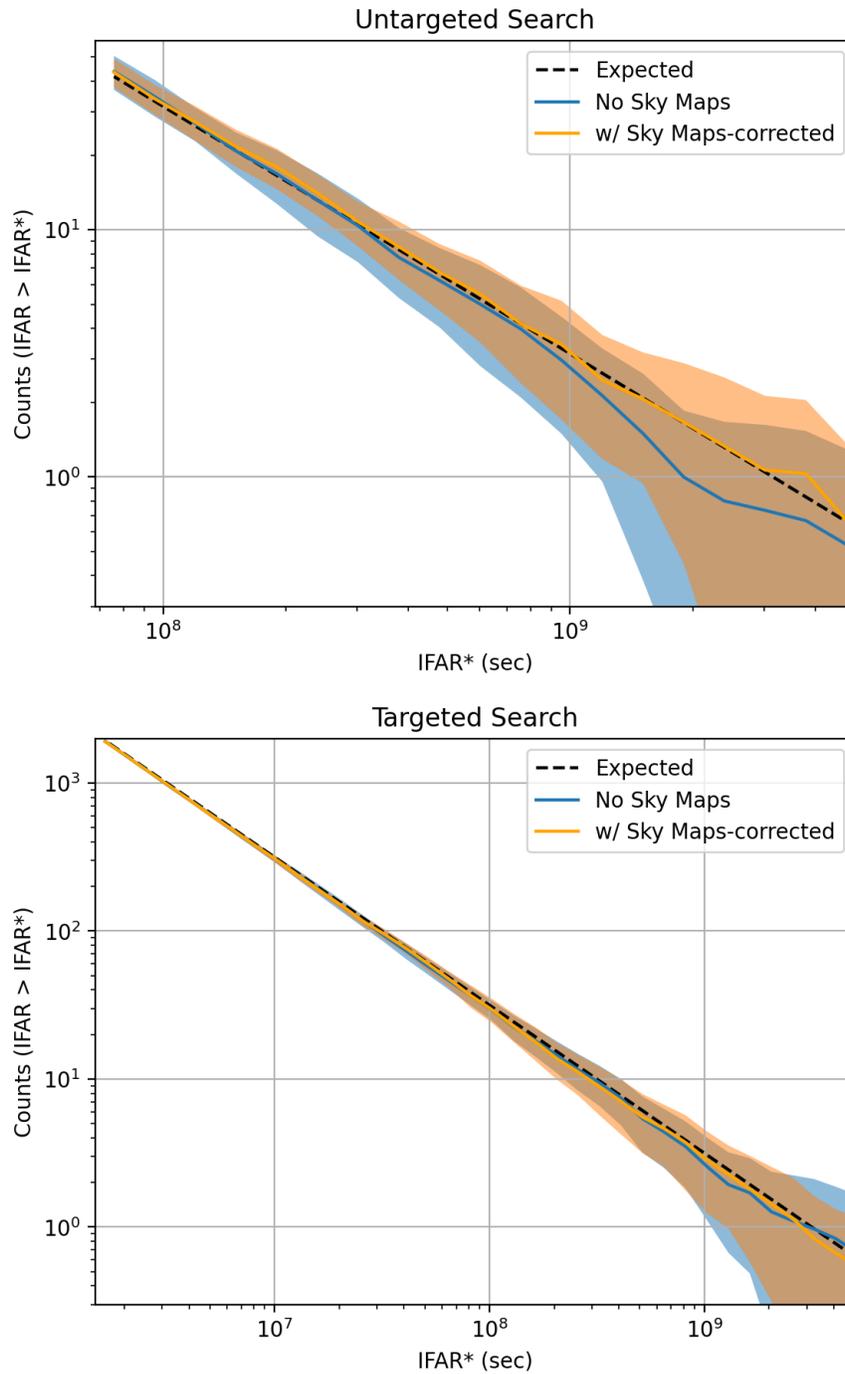


Figure 6.5: Corrected joint FAR with *Swift*/BAT-like GRB candidates (same simulations as in figure 3.4) using the simple remapping as described in section 6.2.

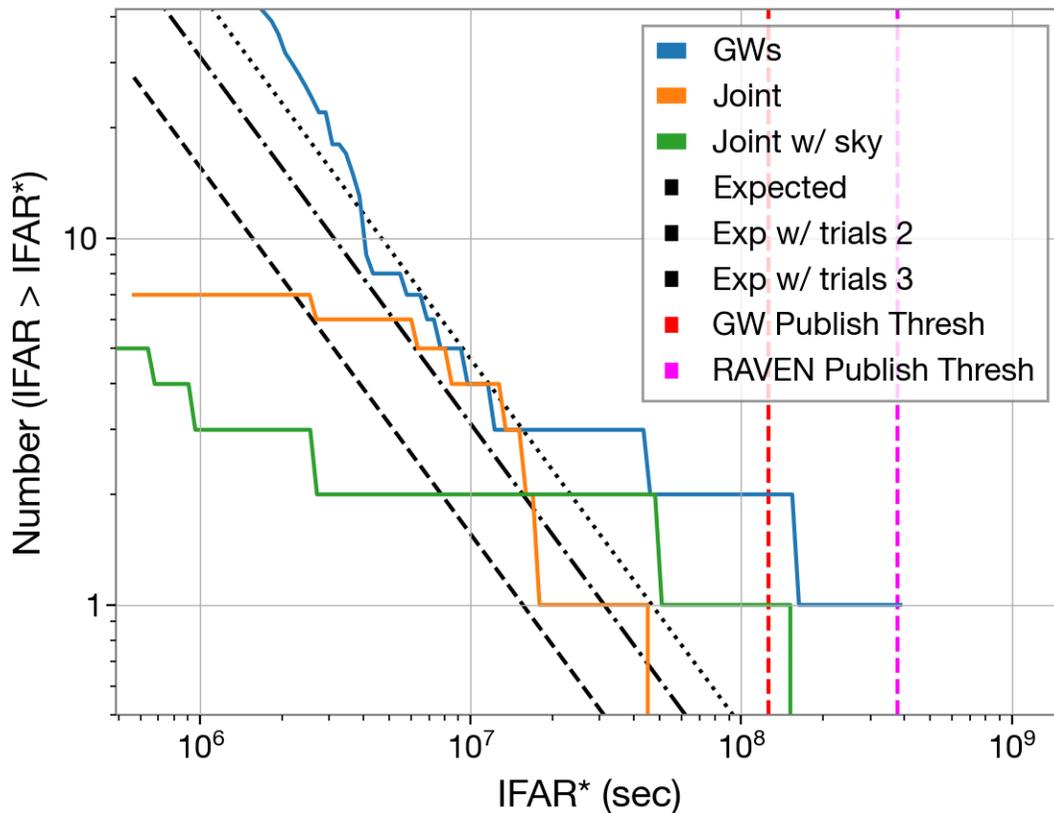


Figure 6.6: Results from the standard joint GW-FRB search detailed in section 6.3. The black lines here are the expected background assuming trials factors of ascending order through the y-axis. We see that the joint FAR not using sky maps seems to be expected from background while including sky maps only marginally improves this. None of these candidates passed the alert threshold, which agrees with results of PyGRB and X-pipeline [Abbott et al. \(2022\)](#).

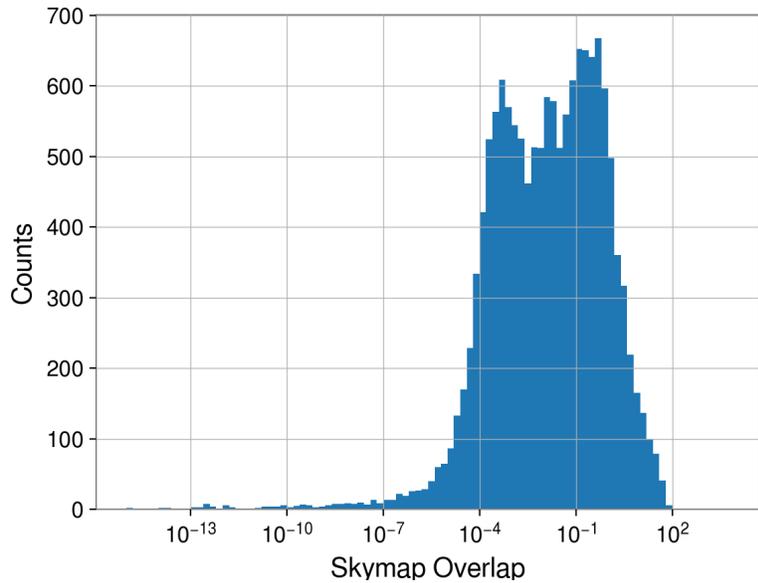
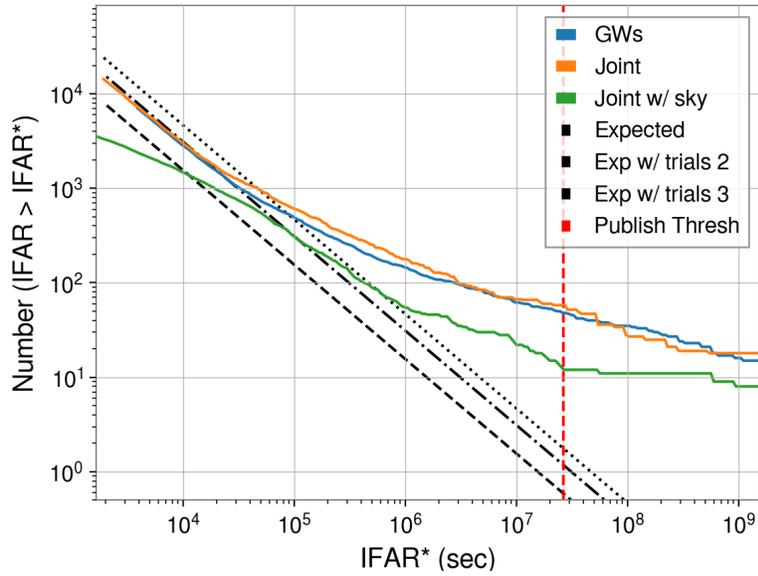


Figure 6.7: Results from the extended window joint GW-FRB search detailed in section 6.3.2. The black lines here are the expected background assuming trials factors of ascending order through the y-axis. We see that GW and temporal joint FARs are nearly identical, with the latter essentially being a superset of the former. Note that there are more joint candidates than GW candidates, due to the large time window leading to an average of more than one GW candidate associated with each FRB. We also see that including sky maps generally reduces the significance of events, consistent with random associations.

Chapter 7

Contributions to the Fourth

LIGO-Virgo-KAGRA Operating Run

7.1 Additional Improvements to RAVEN

Along with the numerous changes detailed in section 3, additional changes have been made in preparation for the fourth LVK observing run (O4). As demonstrated in [Magee et al. \(2021\)](#), the total latency of the LVK alert system is approximately 10 – 20s with plans to reduce it further. This means that the 30s maximum latency of the RAVEN pipeline when using sky maps will be inadequate in comparison. In addition, some critical bugs present RAVEN during O3 need to be fixed. Therefore, the primary goals for RAVEN in preparation for O4 have been:

1. Reduce latency to $\approx 5s$ when using sky maps.
2. Introduce strict code validation to minimize the chance of serious bugs.

7.1.1 Multi-ordered (MOC) Sky Maps

Arguably the most important feature for the next observing run is the ability to use multi-ordered (also known as MOC, i.e. Multi-Ordered-Coverage) sky maps. These sky maps use

the UNIQ ordering system with pixels that can vary in resolution, as opposed to the equal-area pixels normally used with the HEALPix standard (Collaboration, 2018b). This allows for greater resolution around places of higher probability while resulting in a reduced file size, and therefore faster data transfer. Instead of identifying pixels solely by their standard index i_{pix} , the *uniq* index also incorporates *nside*, a measure of resolution

$$uniq = i_{pix} + 4 nside^2. \quad (7.1)$$

This presents a challenge of computing the sky map overlap integral from (3.3), since the indices of two sky maps would no longer trivially match. Fortunately, the coordinates (i.e. right ascension and declination) allow us to match pixels of sky maps with any resolution or configuration. Therefore we developed the following algorithm for handling MOC sky maps

1. For the first sky map, convert the *uniq* index into *nside* ($nside_1 = 2^{\lfloor \log_2(uniq_1/4) \rfloor}$) and then to pixel area ($\Delta A_1 = 4\pi/N_{pix,1}$ where the total number of pixels $N_{pix,1} = 12 nside_1^2$) (Zonca & et al., 2022).
2. Find the coordinates (RA, dec) of a given pixel Ω_1 using the index ($i_1 = uniq - 4nside^2$) (Collaboration, 2018b).
3. Compute the coordinates of the second sky map Ω_2 and then find the closest corresponding pixel where $\Omega_2 \approx \Omega_1$.

With each pixel in the first sky map matched one-to-one to another in the second, the sky map overlap can be computed as

$$I_{\Omega}(x_1, x_2) \approx 4\pi \sum_{i_1=1}^{N_{pix,1}} p(\Omega_1(i_1)|x_1, \mathcal{H}^s) p(\Omega_1(i_1)|x_2, \mathcal{H}^s) \Delta A(i_1) \quad (7.2)$$

where $p(\Omega_1(i)|x_a, \mathcal{H}^s)$ is a probability density normalized so that $\sum_i p(\Omega_1(i)|x_a, \mathcal{H}^s) \Delta A =$

1. This algorithm can be extended to use a standard HEALPix secondary sky map by

setting $p(\Omega_1(i_1)|x_2, \mathcal{H}^s) = P(\Omega_1(i_1)|x_2, \mathcal{H}^s)/\Delta A_2$ where $P(\Omega_1(i_1)|x_2, \mathcal{H}^s)$ is the probability in pixel i_1 and ΔA_2 is the pixel area of the second sky map.

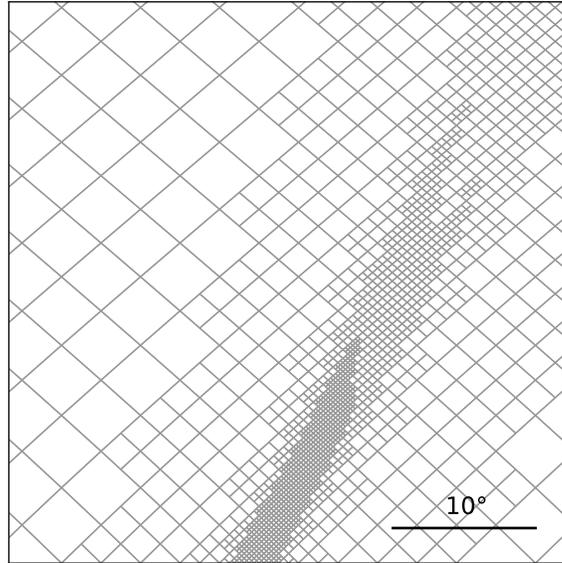


Figure 7.1: An example of adaptive mesh refinement using the MOC standard with GW170817 (Collaboration, 2018b). Resolution increases around areas of higher probability density to capture the relevant structure while minimizing file size, a standard numerical technique.

This method has been confirmed to give the same results as Eq. (3.3) at a significantly lower latency, as seen in table 7.1, although this does not include the time involved with producing or loading sky maps ($\sim 1s$). This will also eliminate the need to wait for the GW sky map to be flattened (i.e. changed from MOC to a standard HEALPix format), a process that can take on the order of seconds.

7.1.2 Constant Testing and Validation

One problem with the operation of RAVEN during O3 was the presence of serious bugs causing false alarms and other issues. These include:

- Query not removing events with incorrect group or search field
- No events returned if querying with pipeline field

Method	Latency (s)
MOC-MOC	$1.07^{+.45}_{-.39}$
MOC-flat	$.16^{+.23}_{-.07}$
flat-flat	$9.76^{+2.34}_{-2.10}$

Table 7.1: Timing in seconds for the sky map overlap integrals for GW170817-GRB 170817A, using both multi-ordered (MOC) and standard HEALPix (flat) sky maps. The uncertainties shown here are the 10 and 90 percentile respectively over ten computations. We see using MOC sky maps generally reduces latency by about an order of magnitude, with the MOC-flat being nearly two orders of magnitude faster. However, we expect file creation/transfer for this flat sky map to eliminate any advantage MOC-flat would have over MOC-MOC.

- Double uploads of coincidence FAR results

as well as numerous other minor issues. To prevent similar bugs from being in the next observing run, we have implemented a number of stages of testing:

- Outputs of all functions must agree with expected results before a change can be moved to production (i.e. unit tests).
- A simulation is run to ensure (3.1) and (3.8) give expected results before a code release.
- Formal review of code by a committee is required prior to implementation.
- Constant end-to-end pipeline validation using mock events.
- Constant ingestion of real and test alerts from GCN.
- Error monitoring during deployment and live testing.

By performing this thorough testing prior to the start of the next observing run, the chance of major bugs running in production will be significantly reduced.

7.1.3 Other Improvements

In addition to decreasing latency via MOC sky maps as described in section 7.1.1, effort has been made to decrease latency elsewhere. This has centered on reducing calls to GraceDB

whenever possible by passing the results of functions and state of the related pipelines (indicated by labels) onto the next function when possible rather than call GraceDB to get this information.

Another change for O4 is to improve and finalize the alert system. While the system in O3 could send an alert for the first coincidence, there was no way to handle subsequent coincidences. This could be an issue if multiple GRB experiments detect the same candidate as was the case with GRB170817A (Goldstein et al., 2017). We have developed a system to choose which coincidence to report based on the lowest joint FAR, where joint FARs including sky maps have precedence. Work has also been done to include combined sky maps in alerts (which did not work correctly in O3), although this is still pending.

7.2 Sub-threshold Search w/ *Swift*

As mentioned in chapter 3, we have been developing an additional search that involves sending GW candidates to other experiments to search for coincident candidates within their sub-threshold data. This allows us to take advantage of data that would otherwise be dumped due to lack of local storage space or computing resources to process it. Prioritizing certain regions of data lets us circumvent these issues and increase sensitivity, as demonstrated by the Gamma-Ray Urgent Archiver for Novel Opportunities (GUANO) system with the Non-Imaging Transient Reconstruction And TEmporal Search (NITRATES) used by *Swift*/BAT (Tohuvavohu et al., 2020; DeLaunay & Tohuvavohu, 2021). We have developed the targeted method in section 3.3.2 explicitly for this use case.

This search has been demonstrated by GUANO to increase their sensitivity by up to 400% (Tohuvavohu et al., 2020), mainly to off-axis events similar to GRB 170817A (Goldstein et al., 2017). Although a universal sGRB jet profile has not been well measured yet (Farah et al., 2020), off-axis detection of CBC events is much more viable compared to GRBs, giving motivation to use GW candidates to trigger sub-threshold sGRB searches.

A schematic of the system to communicate with *Swift*/BAT can be see in figure 7.2. The majority of this system has been finished, including the integration into RAVEN and the GUANO system itself, with the exception of communicating alerts back and forth with GUANO via SCiMMA as of writing. This final part is left as future work for other developers.

We have tested the method in (3.8) by performing a targeted search in conjunction with GstLAL and *Swift*/BAT-GUANO over a period of roughly 65 days during O2. This led to 139 joint candidates with joint FARs consistent with background, as seen in figure 7.3. No events passed the public alert threshold of 1/year (see section 3.3.3). This should give us confidence that the targeted search method from section 3.3.2 is well calibrated, properly handling a set of background triggers. If there is truly a population of sub-threshold off-axis joint candidates, we would expect to find a surplus of significant candidates above the background. Further results are still forthcoming.

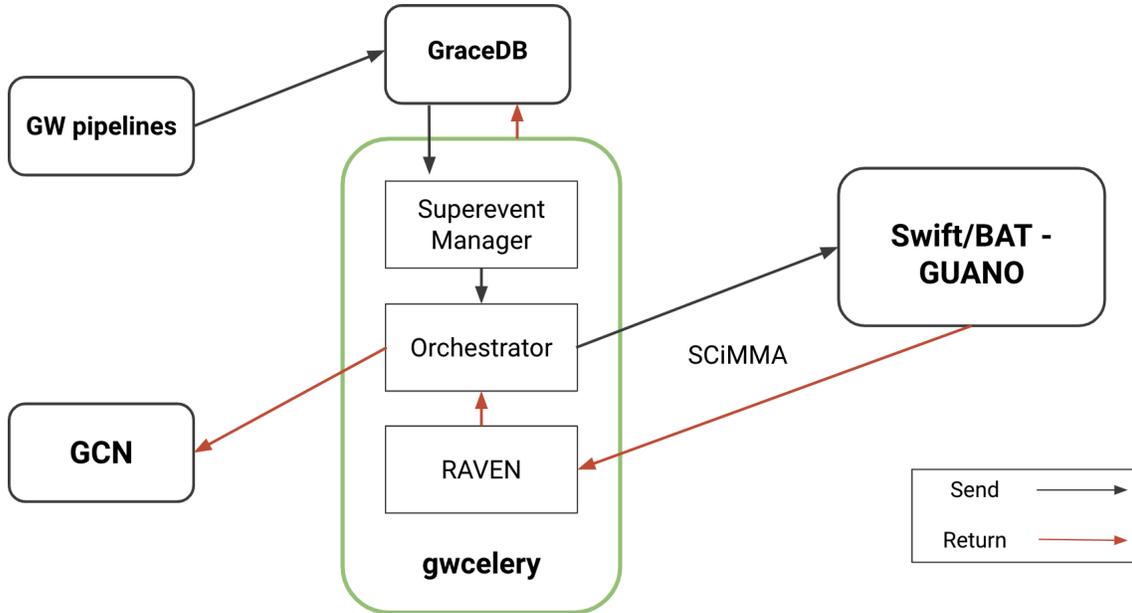


Figure 7.2: Workflow of the targeted search with *Swift*/BAT. GW candidates are found independently and passed along to *Swift*/BAT if they have a FAR less than 2/day. *Swift*/BAT will then send back the GRB candidate if the joint event passes the joint FAR threshold detailed in section 3.3.3. The GRB candidate is then checked for all publishing conditions by RAVEN and if met, an alert is sent to GCN

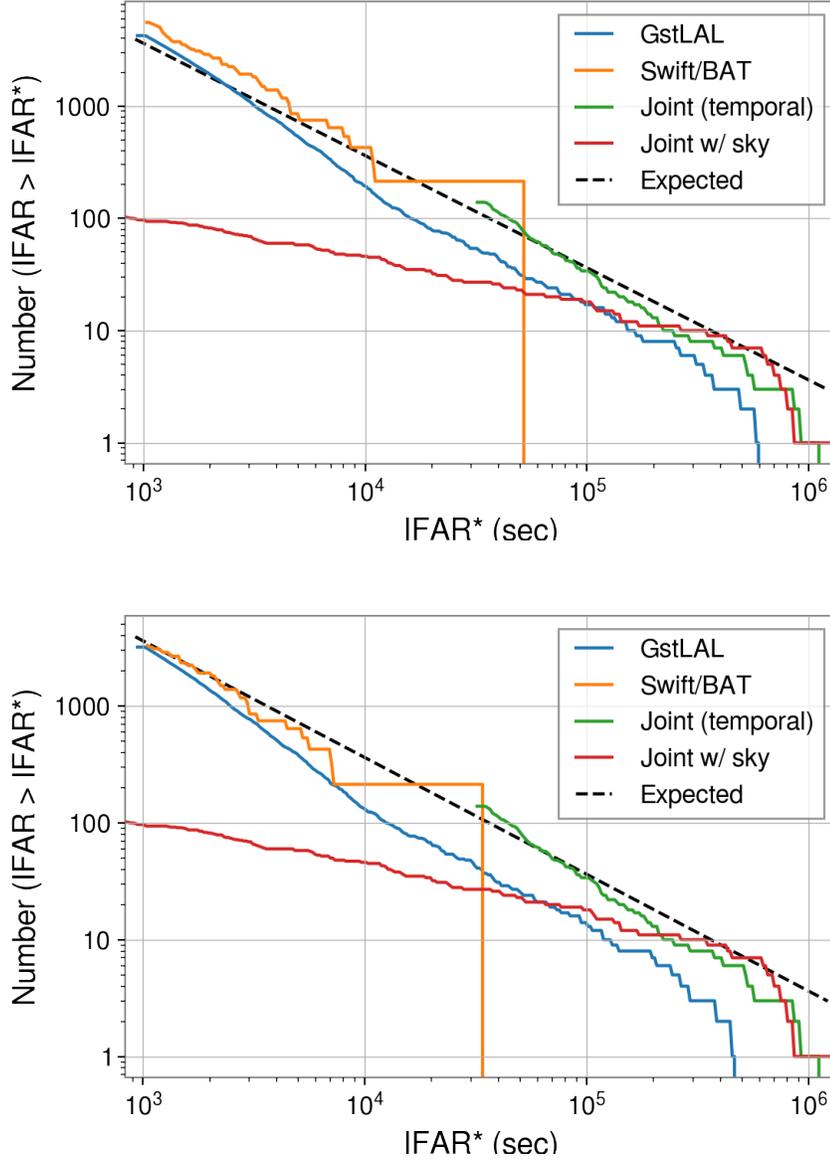


Figure 7.3: Results from an O2 joint targeted search with *Swift*/BAT-GUANO. The top plot shows IFARS directly from each pipeline vs the expected number from (6.2) (with an uptime of 3628065.6 s). The bottom plot shifts the FARs of the individual pipelines by the fraction of their uptime to the joint search uptime, which shows these pipelines are reasonably well calibrated. We see that we get joint FARs that are consistent with an expected background, including the skew when including sky maps as discussed in section 6.2.

7.3 Offline PyCBC Sky Map Generation

Although sky maps are generated in the online workflow of the PyCBC pipeline ([Usman et al., 2016](#)), there has been still work to do with regards to constructing an offline workflow. This is especially important for offline joint searches such as [Stachie et al. \(2020b\)](#) which use PyCBC in conjunction with Fermi-GBM to examine an observing run’s worth of candidates at a time (roughly a year; $\mathcal{O}(10^3)$ events). To this end, we were commissioned to create a work flow that produces a large amount of sky maps as quickly as possible.

7.3.1 Creating a Workflow Using DAGs

Fortunately, sky maps are completely independent from one another, which means this task can be widely parallelized. This means we can leverage the many cores and threads available in a computing cluster via HTCondor and DAGman ([Bockelman et al., 2015](#)). The workflow is shown in figure [7.4](#). Parameters of each candidate recovered by PyCBC (masses, spins, and gpstimes of the max SNR at each detector) are loaded and a job is created for each candidate. We intentionally throttle the number of concurrent sky maps being generated to 50 due to excessive querying of the LIGO segment database. Once every sky map has been generated, excess files (either logs or dag files) are either moved or deleted.

Using this workflow, we created sky maps for the entirety of O2 by working on one chunk (roughly a twentieth of O2) at a time, taking roughly 2-4 hrs per chunk. This means that the entirety of O2 could be finished within 2-4 days provided sufficient computing resources are available. This workflow is intended to be used in conjunction with PyCBC offline runs generally, as well as with a future study using O3 data similar to [Stachie et al. \(2020b\)](#).

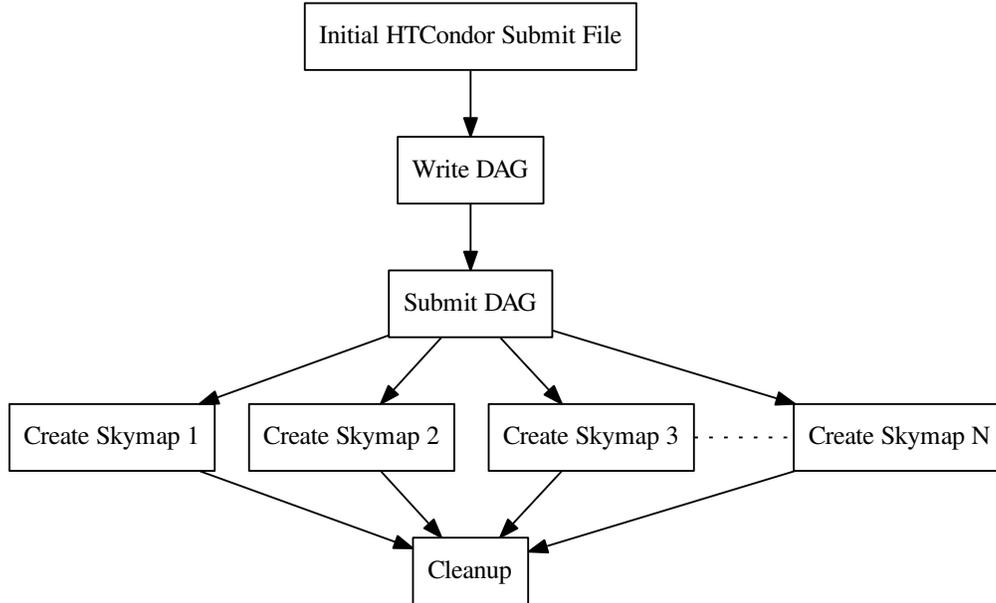


Figure 7.4: Workflow of parallelized PyCBC sky map creation. Writing the DAG within the workflow allows an indefinite number of sky maps to be produced in a parallel manner.

7.4 Speed of Gravity in Low-latency

Since the co-detection of GW170817 and GRB 170817A led to constraints on the speed of gravity (Abbott et al., 2017a), there has been work to perform a similar calculation for the next BNS-GRB detection. With multiple events, coming for larger distances in future runs, we will improve on current measurements. Since the calculations for the speed of gravity limits are relatively straightforward, they can be performed rapidly following a detection by RAVEN and be submitted to publication on the order of a day, provided a template paper is produced beforehand. Similarly to Abbott et al. (2017a), we will describe a method to calculate the probability of association using the data products of and available to RAVEN.

7.4.1 Probability of Association

We start from the results in section 6.1. Recall that a coincidence is either due to a real astrophysical event (\mathcal{H}_c) or random chance (\mathcal{H}_{uncor}), which means we can write the probability of the former by

$$P(\mathcal{H}_c|x_{gw}, x_{grb}) = 1 - P(\mathcal{H}_{uncor}|x_{gw}, x_{grb}) \quad (7.3)$$

where x_{gw} and x_{grb} are the data for the GW and GRB experiments respectively. The probability of random association due to timing can be given from the Poisson distribution

$$P(\mathcal{H}_{uncor}|t_{gw}, t_{grb}) = \begin{cases} R_{grb}\Delta t e^{-R_{grb}\Delta t} \approx R_{grb}\Delta t & (-1 < t_{gw} - t_{grb} < 5) \\ 1 & (otherwise) \end{cases}$$

where R_{grb} is the rate of GRBs, t_{gw} is the GW time, t_{grb} is the GRB time, and $\Delta t = |t_{gw} - t_{grb}|$. We can take this approximation due to the relatively low rate of GRBs and tight coincidence window as discussed in chapter 3. Note that if a coincidence is not found then this equation is trivially 1.

Regarding sky localizations, [Abbott et al. \(2017a\)](#) compared the sky map overlap integral (\mathcal{I}_Ω) of the joint candidate of interest to those made from a distribution of over a hundred sGRB GBM localizations, rotating and shifting each ten times. The probability of being uncorrelated is given by counting the proportion of random associations more significant. However, we have discussed the potential issues with using data products of injections or other detections when computing significance in section 6.2, so we have decided to opt for a self-consistent method instead. We instead rotate one of our sky maps from our coincident candidate 1000 times to create the background of \mathcal{I}_Ω values to compare to, denoting this probability of random association by P_Ω . We find this method gives the same result of $P(\mathcal{H}_{uncor}|\Omega_{gw}, \Omega_{grb}) \approx .01$ for GW170817-GRB 170817A as in ([Abbott et al., 2017a](#)).

Putting this all together we find that

$$\begin{aligned}
P(\mathcal{H}_{uncor}|x_{gw}, x_{grb}) &= P(\mathcal{H}_{uncor}|t_{gw}, t_{grb})P(\mathcal{H}_{uncor}|\Omega_{gw}, \Omega_{grb}) \\
P(\mathcal{H}_{uncor}|x_{gw}, x_{grb}) &\approx R_{grb}\Delta t \cdot P_{\Omega}(\mathcal{I}_{\Omega})
\end{aligned}
\tag{7.4}$$

so that

$$P(\mathcal{H}_c|x_{gw}, x_{grb}) \approx 1 - R_{grb}\Delta t \cdot P_{\Omega}(\mathcal{I}_{\Omega}) .
\tag{7.5}$$

We may think this result contradicts the result of section 6.1, but we assert that the goal here was to recreate the analysis used in [Abbott et al. \(2017a\)](#). Computing P_{Ω} is much more computationally costly than just I_{Ω} since it requires creating a background of these overlap values, meaning it is not well suited for low-latency analysis. It is worth noting that there should be a regime where $1/\mathcal{I}_{\Omega} \approx P_{\Omega}(\mathcal{I}_{\Omega})$ for $\mathcal{I}_{\Omega} \gg 1$ similarly to the corollary in section 6.1. We can see this with GW170817-GRB 170817A where

$$1/\mathcal{I}_{\Omega} = 1/37.7 \approx .027 \approx 0.01 (P_{\Omega}) .
\tag{7.6}$$

However, very high \mathcal{I}_{Ω} values will not be typical, as we see in figures 3.3 and 3.5. Values above just 100 will be very rare for most GW-GRB analysis, with the possible exception of real coincidences involving *Swift*.

Chapter 8

Conclusion

In this dissertation we have developed numerous joint searches and techniques concerning multi-messenger astronomy with gravitational waves. We described advancements we made to the low-latency joint pipeline RAVEN for O3 in chapter 3 and for O4 in section 7.1. We developed new statistics to create new searches and better separate real from false coincidences, including a new joint FAR in 3.3.2 and odds ratio in chapter 4. We tested these methods against others and found some evidence of improvement in chapter 5. We have also been participating in exciting new searches, detailed in sections 6.3 and 7.2.

8.1 Future Work

Although efforts to significantly reduce latency in RAVEN was discussed in section 7.1, this could always be improved as the code is further refined and new technologies develop. Due to CHIME FRB candidates now being available via a private GCN channel, work could be done in creating a low-latency GW-FRB search using RAVEN. Some functions and scripts have been developed for offline searches, but these would be need to be standardized and generalized to be officially included in RAVEN.

We introduced a new joint ranking statistic in chapter 4 with some additional terms that may or may not increases sensitivity in searches. As of writing, there is insufficient ev-

idence that including distance or inclination helps joint searches (see [Ashton et al. \(2020\)](#)). However distance may be helpful for either GW-FRB searches ([Abbott et al., 2022](#)) or GRB candidates with Swift-XRT followup where distance information is available for the EM candidate ([Gehrels et al., 2004](#)). The utility and useful form of the coincidence Bayes factors introduced also requires additional investigation.

Finally, the coincident simulation presented in chapter 5 is far from ideal and requires a more mature iteration to be considered authoritative. Effort could be made to make both the simulated GW and GRB pipelines more realistic, perhaps by using injections or templates from real search pipelines such as `gstLAL` ([Sachdev et al., 2019](#)). In general, coincident simulations are in their infancy since we don't fully understand the underlying models and thus lack the ability to coherently create injections from a shared parameter list. More joint detections are needed to constrain these underlying models.

8.2 Looking Forward

Efforts to improve joint search techniques now at a time when coincidences are incredibly rare will be the most impactful once sensitivity increases. Compared to O3, the rate of BNS detections for O4 and O5 should increase by roughly a factor of 3 and 20 respectively ([Abbott et al., 2020a](#)), although there is still roughly two orders of magnitude uncertainty in the overall rate of BNS detections ([Howell et al., 2019](#); [Abbott et al., 2021](#)). Speculative estimates show that sub-threshold joint searches could increase joint events by a factor of 3 ([Burns, 2019](#)) or 4 ([Tohuvavohu et al., 2020](#)).

The inverse relationship between sensitivity and accommodating for uncertainty greatly incentivizes joint searches to find more candidates, even if very sub-threshold. The most obvious example are the time windows in table 3.1, where if we discover that sGRBs are all delayed around 1.7s similar to GW170817. We could then confidently reduce these time windows, rejecting more false alarms and increasing the significance of the candidates that

remain. In general, the better we understand the underlying models that define multi-messenger candidates, the more specific and sensitive we can design our searches.

As we move towards the next generation of ground-based gravitational detectors and the number of credible GW candidates increases, the likelihood of finding another coincidence like GW170817-GRB 170817A or another type of coincidences will increase. The next generation of joint searches will be faster, more sensitive, and have deeper integration between their underlying experiments. Some new joint searches will be looking for coincidences that are more speculative but could have enormous insights into these systems and fundamental physics as a whole.

Bibliography

- Aartsen, M. G., Ackermann, M., Adams, J., et al. 2017, *Journal of Instrumentation*, 12, P03012
- Aasi, J., Abbott, B., Abbott, R., et al. 2014, *Physical Review D*, 89, 122004
- Abbott, B., Abbott, R., Abbott, T., et al. 2016a, *Physical Review X*, 6, doi:10.1103/physrevx.6.041015
- . 2018, *Physical review letters*, 121, 161101
- . 2019a, *Physical Review X*, 9, 011001
- Abbott, B. P., Abbott, R., Adhikari, R., et al. 2009, *Physical Review D*, 80, 102001
- Abbott, B. P., Abbott, R., Abbott, T., et al. 2016b, *Physical Review D*, 93, 122003
- . 2017a, *The Astrophysical Journal Letters*, 848, L13
- . 2017b, *Physical Review Letters*, 119, 161101
- Abbott, B. P., Bloemen, S., Canizares, P., et al. 2017c, *Astrophysical Journal Letters*
- Abbott, B. P., Abbott, R., Abbott, T., et al. 2017d, *The Astrophysical Journal*, 841, 89
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, *The Astrophysical Journal*, 886, 75
- Abbott, B. P., Abbott, R., Abbott, T., et al. 2020a, *Living reviews in relativity*, 23, 1
- Abbott, R., Abbott, T., Abraham, S., et al. 2020b, arXiv preprint arXiv:2010.14550
- Abbott, R., Abbott, T., Acernese, F., et al. 2021a, arXiv preprint arXiv:2111.03606
- . 2021b, arXiv preprint arXiv:2111.03634
- Abbott, R., Abbott, T. D., Acernese, F., et al. 2021, arXiv preprint arXiv:2111.03634
- Abbott, R., Abbott, T., Acernese, F., et al. 2022, arXiv preprint arXiv:2203.12038

- Abdikamalov, E., Pagliaroli, G., & Radice, D. 2021, in Handbook of Gravitational Wave Astronomy (Springer Singapore), 1–37
- Abramovici, A., Althouse, W. E., Drever, R. W., et al. 1992, *science*, 256, 325
- Acernese, F. a., Agathos, M., Agatsuma, K., et al. 2014, *Classical and Quantum Gravity*, 32, 024001
- Adams, T., Buskulic, D., Germain, V., et al. 2016, *Classical and Quantum Gravity*, 33, 175012
- Aker, M., Beglarian, A., Behrens, J., et al. 2022, *Nature Physics*, 18
- Akutsu, T., Ando, M., Arai, K., et al. 2018, arXiv preprint arXiv:1811.08079
- Al Kharusi, S., BenZvi, S., Bobowski, J., et al. 2021, *New Journal of Physics*, 23, 031201
- Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv preprint arXiv:1702.00786
- Amiri, M., Bandura, K., Berger, P., et al. 2018, *The Astrophysical Journal*, 863, 48
- Amiri, M., Andersen, B. C., Bandura, K., et al. 2021, *The Astrophysical Journal Supplement Series*, 257, 59
- Antonioli, P., Fienberg, R. T., Fleurot, F., et al. 2004, *New Journal of Physics*, 6, 114–114
- Arcavi, I. 2018, *The Astrophysical Journal Letters*, 855, L23
- Arnett, W. D., & Rosner, J. L. 1987, *Physical review letters*, 58, 1906
- Ashton, G., Ackley, K., Hernandez, I. M., & Piotrkowski, B. 2020, arXiv preprint arXiv:2009.12346
- Ashton, G., Burns, E., Dal Canton, T., et al. 2018a, *The Astrophysical Journal*, 860, 6
- . 2018b, *The Astrophysical Journal*, 860, 6
- Athar, M. S., & Singh, S. 2020, *The physics of neutrino interactions* (Cambridge University Press)
- Atwood, W., Abdo, A. A., Ackermann, M., et al. 2009, *The Astrophysical Journal*, 697, 1071
- Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *Space Science Reviews*, 120, 143

Bartos, I., Veske, D., Keivani, A., et al. 2019, *Physical Review D*, 100, 083017

Bhandari, S., Sadler, E. M., Prochaska, J. X., et al. 2020, *The Astrophysical Journal*, 895, L37

Bhat, P. N., Meegan, C. A., von Kienlin, A., et al. 2016, *The Astrophysical Journal Supplement Series*, 223, 28

Blinnikov, S., Novikov, I., Perevodchikova, T., & Polnarev, A. 1984, *Soviet Astronomy Letters*, 10, 177

Bochenek, C. D., Ravi, V., Belov, K. V., et al. 2020, *Nature*, 587, 59

Bockelman, B., Cartwright, T., Frey, J., et al. 2015, *Journal of Physics: Conference Series*, 664, 062003

Burns, E. 2019, arXiv preprint arXiv:1909.06085

Cannon, K., Caudill, S., Chan, C., et al. 2021, *SoftwareX*, 14, 100680

Cho, M.-A. 2019, PhD thesis, University of Maryland

Chu, Q. 2017, PhDT

Collaboration, I., Fermi-LAT, MAGIC, et al. 2018, *Science*, 361, eaat1378

Collaboration, L. S., Collaboration, V., Collaboration, M., et al. 2017, *Nature*, 551, 85

Collaboration, L.-V.-K. 2018a, GWcelery documentation, <https://igwn.readthedocs.io/projects/gwcelery/en/latest/gwcelery.conf.html>, accessed 01-June-2021

—. 2018b, LIGO/Virgo Public Alerts User Guide, https://emfollow.docs.ligo.org/userguide/tutorial/multiorder_skymaps.html, accessed 22-March-2022

Connaughton, V., Briggs, M., Goldstein, A., et al. 2015, *The Astrophysical Journal Supplement Series*, 216, 32

Connaughton, V., Burns, E., Goldstein, A., et al. 2016, *The Astrophysical Journal Letters*, 826, L6

Coughlin, M. W., Dietrich, T., Margalit, B., & Metzger, B. D. 2019, *Monthly Notices of the Royal Astronomical Society: Letters*, 489, L91

Countryman, S., Keivani, A., Bartos, I., et al. 2019, arXiv preprint arXiv:1901.05486

- Creighton, J. D., & Anderson, W. G. 2011, Wiley-VCH, 2, 1
- Cucchiara, A., Levan, A. J., Fox, D. B., et al. 2011, *Astrophysical Journal*, 736, 7
- Dály, G., Galgóczi, G., Dobos, L., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 2374
- DeLaunay, J., & Tohuvavohu, A. 2021, arXiv preprint arXiv:2111.01769
- Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, *Nature*, 340, 126
- Farah, A., Essick, R., Doctor, Z., Fishbach, M., & Holz, D. E. 2020, *The Astrophysical Journal*, 895, 108
- Fishbach, M., Gray, R., Hernandez, I. M., et al. 2019, *The Astrophysical Journal Letters*, 871, L13
- Fragione, G. 2021, *The Astrophysical Journal Letters*, 923, L2
- Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *The Astrophysical Journal*, 611, 1005
- Goldstein, A., Veres, P., Burns, E., et al. 2017, *The Astrophysical Journal Letters*, 848, L14
- Goldstein, A., Hamburg, R., Wood, J., et al. 2019, arXiv preprint arXiv:1903.12597
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005a, *Astrophysical Journal*, 622, 759
- . 2005b, *Astrophysical Journal*, 622, 759
- Greiner, J., Burgess, J. M., Savchenko, V., & Yu, H.-F. 2016, *The Astrophysical Journal Letters*, 827, L38
- Hamburg, R., Fletcher, C., Burns, E., et al. 2020, *The Astrophysical Journal*, 893, 100
- Hirata, K., Kajita, T., Koshiba, M., et al. 1987, *Physical Review Letters*, 58, 1490
- Howell, E., Ackley, K., Rowlinson, A., & Coward, D. 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 1435
- Keivani, A., Veske, D., Countryman, S., et al. 2019, arXiv preprint arXiv:1908.04996
- Klimenko, S., Yakushin, I., Mercer, A., & Mitselmakher, G. 2008, *Classical and Quantum Gravity*, 25, 114029
- Kocevski, D., Burns, E., Goldstein, A., et al. 2018, *The Astrophysical Journal*, 862, 152

Labanti, C., Marisaldi, M., Fuschino, F., et al. 2009, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 598, 470

Lunardini, C., & Smirnov, A. Y. 2004, Astroparticle Physics, 21, 703

MacFadyen, A. I., & Woosley, S. E. 1999, The Astrophysical Journal, 524, 262

Magee, R., Chatterjee, D., Singer, L. P., et al. 2021, First demonstration of early warning gravitational wave alerts, arXiv:2102.04555

Mandel, I., & O’Shaughnessy, R. 2010, Classical and Quantum Gravity, 27, 114007

McLaughlin, M. A. 2013, Classical and Quantum Gravity, 30, 224008

Meegan, C., Lichti, G., Bhat, P., et al. 2009, The Astrophysical Journal, 702, 791

Meszaros, P. 2006, Reports on Progress in Physics, 69, 2259

Moe, B., Stephens, B., & Brady, P. 2016, GraceDB—Gravitational Wave Candidate Event Database

Moroianu, A., Wen, L., Panther, F., & Kovalam, M. 2021, Candidate Fast Radio Burst Associated with a Binary Neutron Star Merger, <https://dcc.ligo.org/LIGO-G2101883>

Nitz, A. H., Nielsen, A. B., & Capano, C. D. 2019, The Astrophysical Journal Letters, 876, L4

Paczynski, B. 1986, The Astrophysical Journal, 308, L43

Piotrkowski, B. 2021, Proceedings of the Wisconsin Space Conference, 1

Piotrkowski, B., Baylor, A., & Hernandez, I. M. 2022, Classical and Quantum Gravity, 39, 085010

Piran, T. 2005, Reviews of Modern Physics, 76, 1143

Platts, E., Weltman, A., Walters, A., et al. 2019, Physics Reports, 821, 1

Rastinejad, J. C., Gompertz, B. P., Levan, A. J., et al. 2022, A Kilonova Following a Long-Duration Gamma-Ray Burst at 350 Mpc, doi:10.48550/ARXIV.2204.10864

Rossi, A., Rothberg, B., Palazzi, E., et al. 2021, arXiv preprint arXiv:2105.03829

Sachdev, S., Caudill, S., Fong, H., et al. 2019, arXiv preprint arXiv:1901.08580

Sakamoto, T., Barthelmy, S., Baumgartner, W., et al. 2011, *The Astrophysical Journal Supplement Series*, 195, 2

Singer, L. P., & Price, L. R. 2016, *Physical Review D*, 93, 024013

Singer, L. P., Price, L. R., Farr, B., et al. 2014, *The Astrophysical Journal*, 795, 105

Singer, L. P., Chen, H.-Y., Holz, D. E., et al. 2016, *The Astrophysical Journal Supplement Series*, 226, 10

Spurio, M., Spurio, & Bellantone. 2018, *Probes of Multimessenger Astrophysics* (Springer)

Stachie, C., Canton, T. D., Burns, E., et al. 2020a, arXiv preprint arXiv:2001.01462

Stachie, C., Dal Canton, T., Burns, E., et al. 2020b, *Classical and Quantum Gravity*, 37, 175001

Sutton, P. J., Jones, G., Chatterji, S., et al. 2010, *New Journal of Physics*, 12, 053034

Takahashi, R., & Nakamura, T. 2003, *The Astrophysical Journal*, 595, 1039

Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *The Astrophysical Journal*, 846, 170

Tavani, M., Barbiellini, G., Argan, A., et al. 2009, *Astronomy & Astrophysics*, 502, 995

Tohuvavohu, A., Kennea, J. A., DeLaunay, J., et al. 2020, *The Astrophysical Journal*, 900, 35

Ubertini, P., Lebrun, F., Di Cocco, G., et al. 2003, *Astronomy & Astrophysics*, 411, L131

Urban, A. L. 2016, PhD thesis, University of Wisconsin-Milwaukee

Ursi, A., Tavani, M., Verrecchia, F., et al. 2019, *The Astrophysical Journal*, 871, 27

Usman, S. A., Nitz, A. H., Harry, I. W., et al. 2016, *Classical and Quantum Gravity*, 33, 215004

Valenti, S., David, J., Yang, S., et al. 2017, *The Astrophysical Journal Letters*, 848, L24

Veske, D., Márka, Z., Bartos, I., & Márka, S. 2020, arXiv preprint arXiv:2010.04162

Vianello, G., Götz, D., & Mereghetti, S. 2009, *Astronomy & Astrophysics*, 495, 1005

von Kienlin, A., Meegan, C., Paciesas, W., et al. 2020, *The Astrophysical Journal*, 893, 46

Winkler, C., Di Cocco, G., Gehrels, N., et al. 2003, *Astronomy & Astrophysics*, 411, L1

Zhang, B. 2019, *Frontiers of Physics*, 14, 64402

—. 2020, *Nature*, 587, 45

Zonca, & et al. 2022, Healpy documentation, <https://healpy.readthedocs.io/en/latest/index.html>, accessed 22-March-2022

Zonca, A., Singer, L., Lenz, D., et al. 2019, *Journal of Open Source Software*, 4, 1298

Zwart, S. F. P., & McMillan, S. L. 1999, *The Astrophysical Journal*, 528, L17

Appendix

A.1 Coincidence Info in LVK GCN VOEvents

To communicate coincidences with other astronomers, RAVEN adds additional information to the standard GW alerts issued by the LVK. The chief purpose of these messages is to inform astronomers of the latest information for a candidate in order to make informed decisions when or even deciding to perform additional follow-up. These alerts are currently sent via machine-readable xml files called VOEvents.

This means we need to just to add additional fields to the VOEvent, which we have summarized in figure [A.1](#). These fields include information on the external event, the significance of the joint event, and a combined sky map to help facilitate follow-up searches. Note that the generation of these fields was made flexible, so that sky map information can be omitted if not available. This could occur for SNEWS events since they don't produce sky localizations or perhaps a sky localization isn't available yet from the respective external experiment.

A.2 Coincidence Info in LVK GCN Circulars

We have also worked on the human-readable emails sent through GCN, also known as circulars. Although this work included many patches and bug fixes, the primary additions we made were templates for RAVEN detections or co-detections alongside GW pipelines. For example, the following is an example for a coincidence involving a *sub-threshold* GRB:

External GCN Notice ID	{583417860, 583327924}
External IVORN	External IVORN identification field
External Observatory	{Fermi, Swift}
External Search	{GRB, SubGRB}
Time Coincidence FAR	Estimated coincidence false alarm rate in Hz using timing
Time and Sky Position Coincidence FAR	Estimated coincidence false alarm rate in Hz using timing and sky position
Combined Sky Map	URL of combined GW-External HEALPix FITS sky localization file
Time Difference	Time between source and external event in seconds

Figure A.1: Additional VOEvent fields for a coincidence found by RAVEN. These include descriptions or examples of the data products discussed in chapter 3.

SUBJECT: Swift BAT trigger with ID 876016: LIGO/Virgo identification of a possible sub-threshold GW compact binary merger counterpart, S5678

****NOTE TO ADVOCATES:** If our circular is the first one to mention the sub-threshold GRB, please use the subject line:

SUBJECT: LIGO/Virgo S5678: Identification of a GW compact binary merger candidate possibly associated with sub-threshold Swift BAT trigger with ID 876016**

The LIGO Scientific Collaboration and the Virgo Collaboration report: We identified the compact binary merger candidate S5678 during real-time processing of data from LIGO Hanford Observatory (H1) and LIGO Livingston Observatory (L1) at 2018-06-28 03:08:04.741 UTC (GPS time:

1214190502.741). The candidate was found by the MBTAOnline [1] analysis pipeline.

Based on the analysis of gravitational-wave data alone, this candidate does not meet our criteria for a public alert. However, a search performed by the RAVEN pipeline found a significant coincidence between this candidate and 876016. The event's properties can be found at this URL: <https://gracedb.invalid/superevents/S5678>

Two GW-only sky maps are available at this time and can be retrieved from the GraceDB event page:

- * bayestar.fits.gz, an initial localization generated by BAYESTAR [2], distributed via GCN notice about 10 hours after the candidate
- * LALInference.fits.gz,0, an updated localization generated by LALInference [3], distributed via GCN notice about 11 hours after the candidate

The preferred sky map at this time is LALInference.fits.gz,0. For the LALInference.fits.gz,0 sky map, the 90% credible region is well fit by an ellipse with an area of 82 deg² described by the following DS9 region (right ascension, declination, semi-major axis, semi-minor axis, position angle of the semi-minor axis):

```
icrs; ellipse(03h08m23s, -45d07m51s, 9d, 3d, 112d)
```

Marginalized over the whole sky, the a posteriori luminosity distance estimate is 522 +/- 102 Mpc (a posteriori mean +/- standard deviation).

A search performed by the RAVEN pipeline found a temporal coincidence between S5678 and a sub-threshold Swift BAT trigger with ID 876016

****CITE ORIGINAL GCN FOR THE EXTERNAL TRIGGER FROM**

https://gcn.gsfc.nasa.gov/gcn3_archive.html, e.g., (Bhalerao et al.,

GCN Circular XXXXX)**. The GRB trigger time is 0.0 seconds before the GW candidate event. The estimated joint false alarm rate for the time coincidence is $3e-11$ Hz, or about one in $1e3$ years.

The classification of the GW signal, in order of descending probability, is BNS (78%), NSBH (22%), BBH (<1%), or MassGap (<1%). Assuming the candidate is astrophysical in origin, there is strong evidence for the lighter compact object having a mass < 3 solar masses (HasNS: >99%). Using the masses and spins inferred from the signal, there is strong evidence for matter outside the final compact object (HasRemnant: >99%).

A combined sky map is also available:

* bayestar-gbm.fits.gz, the normalized product of the GW and GRB localizations.

For the bayestar-gbm.fits.gz sky map, the 90% credible region is 52 deg². The spatial and temporal coincidence is not yet available.

For further information about analysis methodology and the contents of this alert, refer to the LIGO/Virgo Public Alerts User Guide <<https://emfollow.docs.ligo.org/userguide/>>.

- [1] Adams et al. CQG 33, 175012 (2016)
- [2] Singer & Price PRD 93, 024013 (2016)
- [3] Veitch et al. PRD 91, 042003 (2015)

BRANDON PIOTRZKOWSKI

EDUCATION

PhD in Physics, University of Wisconsin-Milwaukee July 2022
Advisor: Patrick Brady
GPA: 3.560

B.S. in Physics and Mathematics, Grand Valley State University May 2016
Advisors: Brett Bolen, Shane Larson
GPA: 3.649

SKILLS

Languages	Python
Scheduling tools	HTCondor, Celery
Software Development	Git, Unix shell, Continuous integration/deployment

ACHIEVEMENTS

- Reviewer for LIGO collaboration publications (O3b catalogue)
- WSGC Graduate & Professional Research Fellowship (2021)
- Excellence in Physics Research and UWM Chancellor's Awards (2016-2021)
- NSF STEM scholarship (2014-2016)

PROJECTS

ligo-raven python [GitLab](#)

- Developed algorithm for multi-resolution sky localizations.
- Expanded query functionality.
- Developed association statistics.

gwcelery python, celery, redis [GitLab](#)

- Reduced latency of coincidence search by up to 10x.
- Created automated alert pipeline.
- Developed new targeted search with *Swift*/BAT.

ligo-followup-advocate python, jinja [GitLab](#)

- Automated updates to other astronomers.
- Automated notice of coincidences.

EXPERIENCE

Graduate Research Assistant UW-Milwaukee Dept. of Physics Fall 2019, 2020, & 2021, Spring 2022

Graduate Teaching Assistant UW-Milwaukee Dept. of Physics Fall 2016-Spring 2019, Spring 2020 & 2021

Undergraduate Research Assistant University of Mississippi Summer 2015

OUTREACH

- Run Physics/Mathematics Educational Youtube Channel (Biscuit Physics)
- Leader of UWM Physics graduate student advocacy group (2018-2022)

- Media Coordinator and presenter of Coffeeshop Astrophysics (2019-2020)
 - The Physics of Everyday Objects (5/12/2019)
 - Alternative Telescopes (11/16/2019)
- Active participant and presenter for UWM CGCA journal club (2018-2022)

TALKS

- Searching for Sub-threshold Gravitational Wave Candidates with RAVEN. APS April Meeting 2022. New York City, NY.
- Testing the Performance of Various Multi-messenger Statistics Used in Gravitational Wave Searches. Midwest Relativity Meeting 2021. Urbana-Champaign, IL.
- Testing Ranking Statistics for Joint Gravitational Wave and Gamma-Ray Burst Searches. Wisconsin Space Conference 2021. Milwaukee, WI.
- Recent developments in gravitational wave multi-messenger searches. Midwest Relativity Meeting 2020. Notre Dame, MI.
- Testing strong cosmic censorship for near extremal Reissner-Nordstrom-de Sitter spacetime.. Midwest Relativity Meeting 2018. Milwaukee, WI.
- Spurious Acceleration Noise on the LISA Configuration Due to Solar Irradiance. APS April Meeting 2016. Salt Lake City, UT.
- Spurious Acceleration Noise on the LISA Configuration Due to Solar Irradiance. Midwest Relativity Meeting 2015. Evanston, IL.

RESEARCH INTERESTS

I am primarily interested in coincidences between gravitational waves (GW) and other high-energy astronomical transients such as gamma-ray bursts (GRB), neutrino bursts, and fast radio bursts (FRB). These are thought to come from mergers of neutron stars or perhaps core collapse supernova. Whenever we have detected these types of events with multiple messengers, it has led to incredible constraints of fundamental physics and insights into these systems.

I have developed searches, participated in searches, and developed association statistics in order to find these sorts of coincidences. I have worked as a member of the Laser Interferometer Gravitational wave Observatory (LIGO), Virgo, and Kamioka Gravitational Wave Detector (KAGRA) to develop the low-latency system to alert astronomers of GW detections. I have also helped development in experiments that use these coincidences, such as measuring the speed of gravity and Hubble constant.

PUBLICATIONS

Short Author

- **Piotrzkowski, B.**, Baylor, A., & Hernandez, I. M. (2022). *A joint ranking statistic for multi-messenger astronomical searches with gravitational waves*. *Classical and Quantum Gravity*, 39(8), 085010.
- **Piotrzkowski, B.** (2021). *Testing Ranking Statistics for Joint Gravitational Wave and Gamma-Ray Burst Searches*. In *Proceedings of the Wisconsin Space Conference* (Vol. 1, No. 1).
- Magee, R., Chatterjee, D., Singer, L. P., Sachdev, S., Kovalam, M., Mo, G., ... **Piotrzkowski, B.** ... & Zweizig, J. (2021). *First demonstration of early warning gravitational-wave alerts*. *The Astrophysical Journal Letters*, 910(2), L21.
- Ashton, G., Ackley, K., Hernandez, I. M., & **Piotrzkowski, B.** (2020). *Current observations are insufficient to confidently associate the binary black hole merger GW190521 with AGN J124942. 3+ 344929*. arXiv preprint arXiv:2009.12346.
- Frank, B. M., **Piotrzkowski, B.**, Bolen, B., Cavaglia, M., & Larson, S. L. (2020). *Modeling spurious forces on the LISA spacecraft across a full solar cycle*. *Classical and Quantum Gravity*, 37(17), 175007.
- **Piotrzkowski, B.**, et al. (2022). *RAVEN: The Low-Latency Gravitational Wave Multi-messenger Pipeline* (In Preparation).

Contributing

- Abbott, R., Abbott, T. D., Acernese, F., Ackley, K., Adams, C., Adhikari, N., ... & Bouffanais, Y. (2022). *Search for Gravitational Waves Associated with Fast Radio Bursts Detected by CHIME/FRB During the LIGO–Virgo Observing Run O3a*. arXiv preprint arXiv:2203.12038.
- Abbott, R., Abbott, T. D., Acernese, F., Ackley, K., Adams, C., Adhikari, N., ... & Bossilkov, V. (2021). *GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run*. arXiv preprint arXiv:2111.03606.