

# Low-Latency Algorithm for Multi-messenger Astrophysics (LLAMA) with Gravitational-Wave and High-Energy Neutrino Candidates

Stefan Countryman,<sup>1,\*</sup> Azadeh Keivani,<sup>1</sup> Imre Bartos,<sup>1,2</sup> Zsuzsa Marka,<sup>3</sup>  
Thomas Kintscher,<sup>4</sup> K. Rainer Corley,<sup>1</sup> Erik Blaufuss,<sup>5</sup> Chad Finley,<sup>6</sup> and Szabolcs Marka<sup>1</sup>

<sup>1</sup>*Department of Physics, Columbia University, New York, NY 10027*

<sup>2</sup>*Department of Physics, University of Florida, Gainesville, FL 32611*

<sup>3</sup>*Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027*

<sup>4</sup>*DESY, D-15738 Zeuthen, Germany*

<sup>5</sup>*Department of Physics, University of Maryland, College Park, MD 20742, USA*

<sup>6</sup>*Oskar Klein Centre and Department of Physics, Stockholm University, SE-10691 Stockholm, Sweden*

(Dated: January 18, 2019)

We describe in detail the online data analysis pipeline that was used in the multi-messenger search for common sources of gravitational waves (GWs) and high-energy neutrinos (HENs) during the second observing period (O2) of Advanced LIGO and Advanced Virgo. Beyond providing added scientific insight into source events, low-latency coincident HENs can offer better localization than GWs alone, allowing for faster electromagnetic follow-up. Transitioning GW+HEN analyses to low-latency, automated pipelines is therefore mission-critical for future multi-messenger efforts. The O2 Low-Latency Algorithm for Multi-messenger Astrophysics (LLAMA) also served as a proof-of-concept for future online GW+HEN searches and led to a codebase that can handle other messengers as well. During O2, the pipeline was used to take LIGO/Virgo GW candidates as triggers and search in realtime for temporally coincident HEN candidates provided by the IceCube Collaboration that fell within the 90% confidence region of the reconstructed GW skymaps. The algorithm used NASA's Gamma-ray Coordinates Network to report coincident alerts to LIGO/Virgo's electromagnetic follow-up partners.

## I. INTRODUCTION

Recent multi-messenger discoveries of GWs in coincidence with electromagnetic (EM) observations have opened up new windows to the Universe. The multi-messenger science reach of the GW detectors had been enabled by decades of effort preceding the discovery [1–33]. The detection of a short gamma-ray burst (sGRB) 1.7 seconds after the GW detection from a binary neutron star (BNS) merger on August 17, 2017 (GW170817) was the first multi-messenger/multi-wavelength observation of GWs [29, 34]. The LIGO/Virgo detectors recorded the GW170817 signal, which was followed by the detection of GRB 170817A by *Fermi*-GBM [35, 36] and *INTEGRAL* [37, 38], that were both spatially and temporally coincident with GW170817. This event was subsequently followed up by several different observatories in a broad range of wavelengths and cosmic messengers [29, 30].

Another recent multi-messenger discovery is related to the detection of a high-energy neutrino (HEN; IceCube-170922A) by the IceCube Neutrino Observatory, which is the first  $3\sigma$  correlation with EM emissions from a flaring blazar, TXS 0506+056 [39, 40]. The flaring blazar was in an active phase in high-energy and very high-energy  $\gamma$ -rays and showed variabilities in X-rays and radio bands. This confirms that HENs are produced by cosmic accelerators such as blazars.

These two recent multi-messenger discoveries enable us to better understand and explore the origin of cosmic particles, the astrophysical mechanisms that produce them, and the physical implications of their sources.

The detection of GWs [34, 41–45] and HENs [46–48]

have been separately reported and confirmed over a few years of detector operations, though no astrophysical source has yet been observed simultaneously with both messengers [16, 28, 30, 31, 49]. Observing even a single joint source of GWs and HENs in both messengers could transform our understanding of the underlying mechanisms that create them [21, 24].

A previous paper by Baret et al. [18] presented a joint GW+HEN analysis algorithm that has been used for previous searches and which was adapted for the Low-Latency Algorithm for Multi-Messenger Astrophysics (LLAMA). It was used in a joint search using Initial LIGO/Initial Virgo and the partially completed IceCube detector data, which placed upper limits on the source rate [26]. The algorithm described below builds on this work. There had also been earlier efforts on designing joint searches for GWs and HENs, such as [5] and [50].

The rapid EM follow-up of GW events from compact binary mergers has been made significantly easier by LIGO/Virgo's low latency alert distribution to partner observatories [38, 51, 52]. However, the typically large source localization uncertainty from GW data creates challenges for EM follow-ups due to EM telescopes' relatively small fields of view. Another problem with limited localization is the large number of foreground transients, mostly supernovae, that are cumbersome to differentiate from GW counterparts [53, 54]. By contrast, HENs from IceCube provide localizations of a median of  $\sim 0.5^\circ$  for an  $E^{-2}$  signal neutrino spectrum [55]. This, in addition to high duty cycle of neutrino observatories, makes HENs also a well-suited messenger to enhance GW studies. The rapid identification of GW+HEN coincidences

will provide significantly more precise localization than GWs alone, enabling faster and more efficient EM follow-up observations with a high scientific payoff.

Several sources capable of generating HENs and GWs have been proposed, including core-collapse supernovae (CCSN) [19, 56], GRBs (see e.g. [57, 58]), BNS mergers [59], neutron star-black hole (NS-BH) mergers [60], soft gamma repeaters [61, 62], and microquasars [50]. Core-collapse supernovae have long been considered sources of GW emission (e.g. [19, 63, 64]). These sources create relativistic outflows capable of producing HENs that can travel unimpeded for billions of light years before reaching Earth, providing useful information when they interact with the detectors. TeV neutrinos are also expected to be detected in the HEN detectors from a galactic SN  $\sim 0.1 - 10$  days after detections of GWs and MeV neutrinos [56].

Among the most promising multi-messenger sources are the progenitors of GRBs. sGRBs associated with BNS (or NS-BH) mergers are known to be GW emitters [29] and are considered to be HEN sources [58]. Binary Black Hole (BBH) mergers emitting GWs (e.g. [41, 42]) are typically not expected to be strong emitters of EM and neutrino radiation, although there could be exceptions [65–69].

Besides the known sources, searching for astrophysical GW+HEN signal might reveal unknown sources or production mechanisms.

This paper describes LLAMA in the configuration that searched for joint GW+HEN sources from LIGO/Virgo during LIGO’s second observing run (O2). In Sec. II we describe the state of the LIGO and Virgo detectors as well as the IceCube Neutrino Observatory during O2. We then summarize past offline searches in Sec. III. In sec. IV we explain the details of the GW+HEN online search, including the data analysis method and software implementation. We conclude in Sec. V with lessons learned from the GW+HEN pipeline in O2.

## II. DETECTORS AND DATA

During O2, LLAMA combined data from LIGO, Virgo, and IceCube, using NASA’s Gamma-Ray Coordinates Network (GCN; [70]) as a mediator to receive triggers and disseminate results. Results were stored on LIGO/Virgo’s candidate database (GraceDB<sup>1</sup>) and on [gw-astronomy.org](http://gw-astronomy.org). Sec. II C describes this flow of data between partners (summarized in Fig. 1).

### A. LIGO and Virgo

LIGO consists of two interferometers, each with 4-km-long arms. One site is in Livingston, Louisiana and the

other is in Hanford, Washington [71]. Virgo is also a similar interferometer located near Pisa, Italy [72]. The arm length of Virgo is 3 kilometers. Both LIGO and Virgo are sensitive to GWs in a frequency band of 10–10,000 Hz [73]. The Initial LIGO upgrades started in 2010 resulted in operational commencement of LIGO [74] in 2015. LIGO underwent a series of upgrades between its first and second observing run (O1 and O2, respectively; [75]).

The second observing run (O2) of the LIGO detectors started in November 30, 2016, with Virgo [76] joining on August 1, 2017. O2 finished on August 25, 2017. Four GW discoveries were reported from this period: GW170104 [43], GW170608 [44], GW170814 [45], and GW170817 [34]. The first three were emitted from BBH mergers and the last one as mentioned in Sec. I resulted from a BNS merger. Triggers from all LIGO/Virgo search pipelines (including cWB [77, 78], GSTLAL [51], PyCBC [79], oLIB [80], and MBTAOnline [81]) were included in the search as long as they had been manually approved for EM follow-up by LIGO/Virgo experts (indicated by GraceDB’s EM\_READY flag) and had had a GCN Notice sent out (acting as a redundant check of the event’s validity). LIGO/Virgo have recently announced a new catalog of GWs including four additional GW candidates detected during O2 [82].

### B. IceCube Neutrino Observatory

IceCube is a gigaton neutrino detector located under the geographical South Pole in Antarctica [83]. IceCube has been continuously detecting neutrinos with its complete 86-string configuration since 2010 [84].

IceCube HENs interact with the ice through neutral-current and charged-current interactions. In a charged-current interaction, most of the neutrino energy is transferred to its associated leptons. Muons are capable of traveling several kilometers (unlike electrons, which lose energy rapidly, and taus, which decay quickly), emitting Cherenkov light along their path that is detected by the IceCube digital optical modules (DOMs) placed in the polar ice [83]. The direction of the muon and hence its parent neutrino is reconstructed using the timing information of photons captured by different DOMs along the muon track. The reconstruction methods cannot distinguish between neutrinos and anti-neutrinos and are sensitive to their combined flux. IceCube can identify the direction of these through-going muons with a precision of  $\sim 0.5^\circ$  [85]. The IceCube Collaboration provided us with a stream of realtime muon data [85] for testing this pipeline.

The primary background source in IceCube comes from the muons that are produced by cosmic ray interactions in the atmosphere [83]. The up-going events are more likely to be from neutrino interactions, as the Earth attenuates the cosmic-ray induced atmospheric muons. The mis-reconstructed muons are the main backgrounds

<sup>1</sup> <https://gracedb.ligo.org>

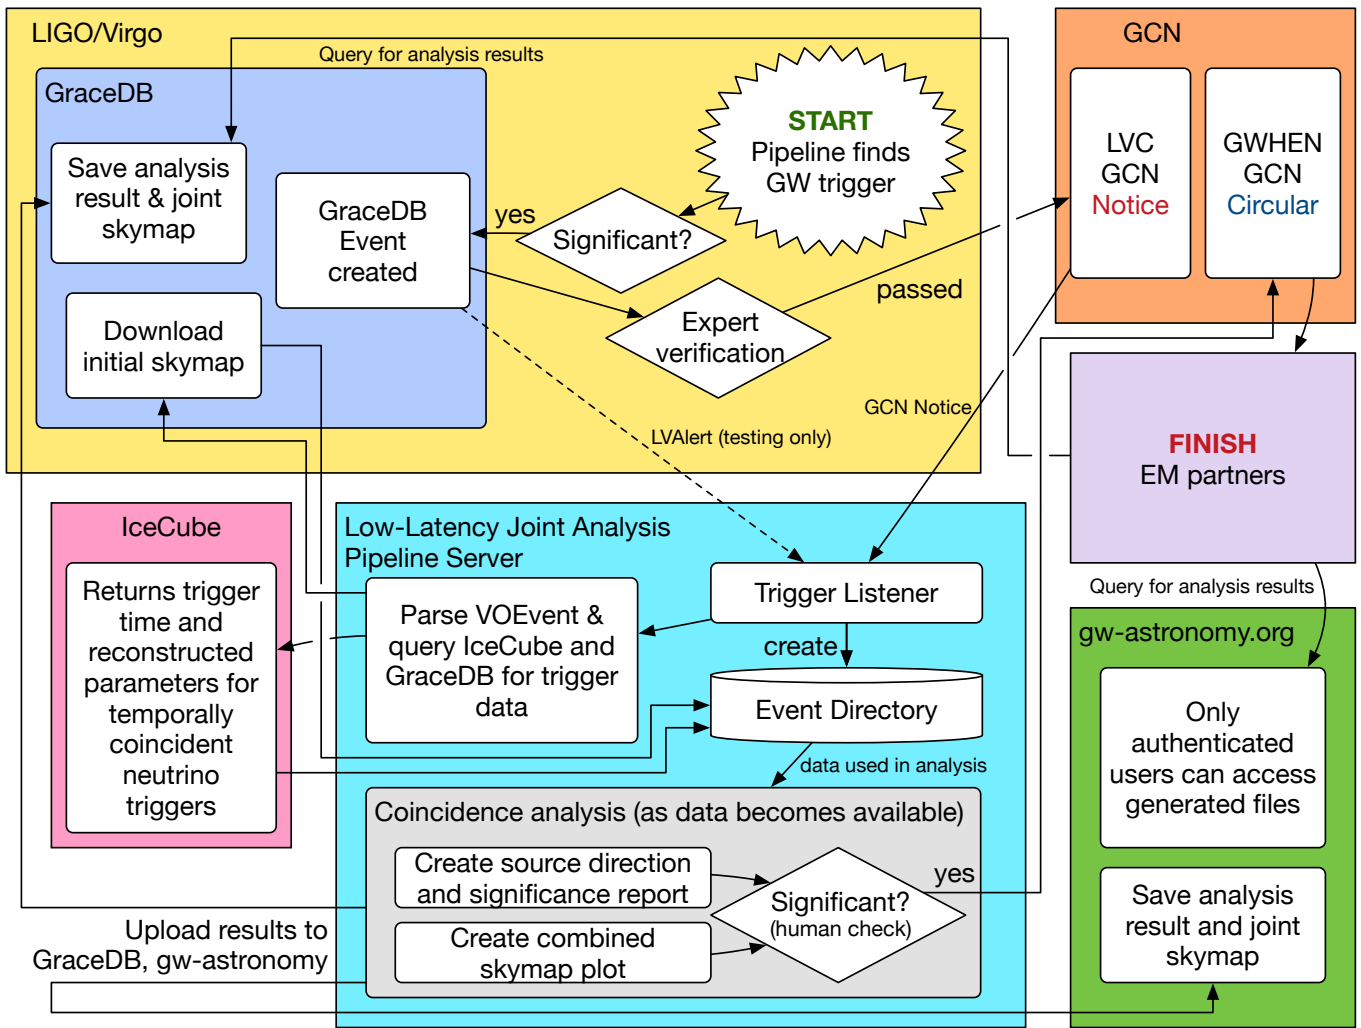


FIG. 1. Information flow in the pipeline. This diagram shows trigger and data sources and destinations used during O2, though the pipeline itself can readily accommodate new sources and destinations for both. See Sec. II C for more information.

left in the northern hemisphere. IceCube uses multivariate selection techniques (Boosted Decision Tree, BDT; [86]) to remove the poorly reconstructed events to distinguish the signal (muon products from neutrino interactions in ice) from the background [55, 84, 87]. The other source of all-sky background consists of atmospheric neutrinos which follow a softer energy spectrum compared to the typically expected astrophysical neutrino spectrum.

In the southern sky, the background is dominated by the atmospheric cosmic-ray muons which are well-reconstructed. Therefore more restrictive cuts are applied to reduce the background in the southern hemisphere. In this study, in addition to the upgoing neutrinos, we accept the down-going events with a BDT score greater than 0.1 [87].

### C. Data Flow to and from the Pipeline

LLAMA relied on LIGO/Virgo and IceCube data for input; GCN (as an intermediary) for GW triggers; LIGO/Virgo and [gw-astronomy.org](https://www.gw-astronomy.org) for data result storage; and GCN for dissemination of results via GCN Circular to EM followup partners (see Fig. 1 for a diagram of data).

LLAMA used GW triggers deemed significant by LIGO/Virgo as input. These triggers were generated by LIGO/Virgo detection pipelines (Sec. II A), stored on GraceDB along with reconstructed skymaps, and human-vetted (Sec. II A). If deemed significant, they were sent in the form of a VOEvent [88] to GCN and distributed as a GCN Notice. The pipeline received and parsed these GCN Notices, used the metadata they contained to fetch GW and neutrino localizations from LIGO/Virgo and IceCube, and used those localizations to run the joint analysis (see Sec. IV C for a description of the analysis

method and Sec. IV D for details on the software implementation of the internal part of the analysis).

Data products from the joint analysis (in the form of a joint skymap and neutrino candidate data) were uploaded to GraceDB and [gw-astronomy.org](http://gw-astronomy.org) to facilitate data archiving and access by EM followup partners. A draft GW+HEN GCN Circular was automatically generated by the pipeline and distributed for verification and (if necessary) modification. Once it was deemed ready for distribution, it was distributed as a GW+HEN GCN Circular following the GW trigger’s LIGO/Virgo GCN Circular; this final step was a necessary step to avoid confusion among EM followup partners unfamiliar with LLAMA.

### III. PAST OFFLINE SEARCHES

An offline GW+HEN search was conducted for Initial LIGO and Initial Virgo using HENs from the partially completed IceCube [26]. The joint search ran during Initial LIGO’s 5th and 6th observing runs (S5 and S6, respectively) and Initial Virgo’s first three observing runs (SR1, SR2, and SR3); IceCube ran in its 22, 59, and 79-string configurations. A joint search was also conducted with Initial LIGO S5/Initial Virgo SR1 and ANTARES in its 5-string configuration [22]. The complete configuration of ANTARES with data in 2009-2010 coincided with Initial Virgo SR2-3 and Initial LIGO S6 runs and significantly improved the GW+HEN search sensitivity [89]. These searches used sub-threshold GW event candidates (i.e. triggers whose significance was too low for LIGO/Virgo to claim a detection) from the entire run and calculated a joint test statistic, taking advantage of the improved sensitivity of a joint search to GW+HEN events. No significant coincident events were found in either search, leading to upper limits on joint GW+HEN event rates.

HEN search results and GW+HEN science foundations have also been published for several GW detections from LIGO’s 1st and 2nd observing runs (O1 and O2, respectively) [2–33] and Virgo’s first observing run. HEN candidates for these searches were provided by IceCube, and ANTARES. Data from the Pierre Auger Observatory was also included in the case of the first detected BNS merger, GW170817 [30]. Other analyzed events were BBH mergers GW150914 [28], GW151226 and LVT151012 [31]. These searches ran in response to high-significance GW candidates (all but LVT151012 were claimed as detections); full searches for O1 that include sub-threshold GW events are reported in [49]. Full searches for O2 are ongoing.

### IV. GW+HEN ONLINE SEARCH

The low-latency joint GW+HEN event search was enabled in response to new event candidates. The rapid

response and low-latency analysis expanded capabilities of GW+HEN searches compared to previous archival searches. Beyond its obvious discovery potential, a low-latency search also offers numerous advantages for EM follow-up, including improved localization (Sec. IV A 1), low-latency sub-threshold source detection (Sec. IV A 2), and the ability to handle increased event rates through automation (Sec. IV A 3). Achieving these goals within O2 constraints required nominal analysis times of less than 30 minutes (Sec. IV B) to run the data analysis procedure described in Sec. IV C. The implementation of LLAMA, described in Sec. IV D, achieved this goal, providing joint GW+HEN localization and testing key methods for future low-latency sub-threshold searches and the planned full automation of GW+HEN joint searches.

#### A. LLAMA Advantages for EM Follow-ups

##### 1. Improved localization with neutrinos

The GW search area size is a limiting and costly factor in the speed of EM follow-up efforts for all but the highest energy photons [90, 91]. Typical EM observatories have viewing areas smaller than  $10 \text{ deg}^2$ , whereas LIGO/Virgo GW skymaps typically have 90% confidence regions ranging from tens to thousands of  $\text{deg}^2$  (depending on the number of detectors included in a trigger and the loudness of the GW signal). IceCube neutrinos, by contrast, typically have error regions of  $0.5 \text{ deg}^2$ . A joint GW+neutrino signal’s localization will be determined by better localized messenger, i.e. neutrinos. The typical  $0.5 \text{ deg}^2$  localization makes these joint events easier to follow up, particularly when the original GW localization was poor (Fig. 2).

A low-latency GW+HEN search can identify neutrinos that come from the GW source and provide their localizations to EM partners, allowing for rapid EM source identification without having to tile large areas of the sky. This is particularly important in cases when only one or two GW detectors are operational, as the typical 90% confidence region will typically be hundreds to thousands of square degrees in these cases. 1 or 2 detector triggers are a likely scenario given typical LIGO/Virgo duty cycles of 50-70% ([51] and references therein).

Even when all detectors are operational, issues with data transmission and glitch removal (among other things) can prevent all operating GW detectors from providing data for a low-latency skymap. This specific scenario occurred during the detection of GW170817, the first direct detection of a BNS merger; a glitch in the LIGO Livingston observatory and an issue with Virgo data transmission caused the earliest available skymap to have 1-detector localization, with improved localizations (with 2 and 3 detectors) following later on as data became available. Had a neutrino been detected from the merger by IceCube, LLAMA would have provided a faster improved localization than that of the GW/GRB



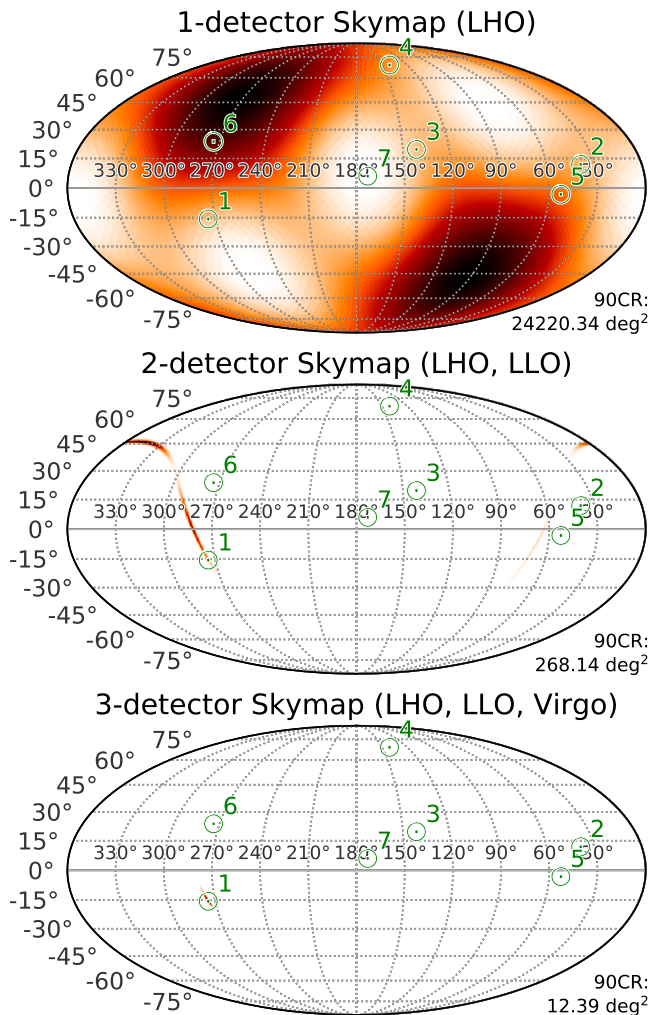


FIG. 2. **Simulated** joint skymaps showing what a joint detection would look like for 1, 2, and 3 detector skymaps from a BNS merger with LIGO/Virgo at design sensitivity. The brightest regions of the skymaps are the likeliest GW source directions; sizes of the 90% credible regions are noted at bottom right. Both progenitors have masses of  $1.4M_{\odot}$ . The neutrinos are located at the green dots (surrounding circles added for emphasis). A fake coincident neutrino has been inserted into the figures as Neutrino 1 to show what the skymap would have looked like for a coincident detection. Though the neutrino localization is better in all three cases, the improvement provided by a joint localization is much greater in the 2-detector skymap. 1-detector localization is too poor to pick out the correct neutrino from direction alone, though other factors (like neutrino time/energy or other non-GW+HEN observations, e.g. GRBs) can boost the joint significance enough to make a given neutrino trigger worth investigating.

skymaps alone. Providing rapid joint localization was a primary goal of the GW+HEN pipeline during O2.

Fig. 2 shows how the number of contributing GW detectors affects source localization and how a joint GW+HEN detection can provide much faster source-direction recovery than a GW skymap alone.

## 2. Low-latency sub-threshold search

LLAMA calculates the joint significance of an event from the significances of the individual triggers. The coincident events are expected to achieve higher significances than the individual events. Sub-threshold triggers, which, alone have lower significances (in comparison with the rare high significance events) are by themselves unrecoverable as astrophysical GW signals but can sometimes be identified in the company of another sub-threshold signal in a different messenger channel. This means that sub-threshold GW and neutrino triggers can be run through a joint analysis with the intention of finding events whose significances are high enough that the resulting multi-messenger event candidates exceed the detection threshold.

Past sub-threshold search efforts have been offline, taking place months after the triggers were identified and thus precluding any chance at prompt EM follow-up. An online sub-threshold search could identify joint events that would otherwise not be broadcast to the EM follow-up community because of low significance of the GW trigger alone. The O2 GW+HEN pipeline is capable of searching for sub-threshold GW+HEN events in real-time, although this feature was not utilized during O2 due to data use restrictions. (It was successfully tested internally, however.)

## 3. Automation needed for higher event rate

Event rates for both GW and neutrino searches are expected to climb in the coming years [75, 92]. In addition, sub-threshold searches must include significantly higher numbers of triggers than standard joint searches. As analysis methods and observational data APIs converge and stabilize, it becomes feasible to automate and speed up these searches and thus avoid analysis backlogs.

## B. Target Timeline for an Online Search

During O1 and O2, there was typically a latency of about half an hour (sometimes significantly more) between LIGO/Virgo GW trigger identification and alert dissemination via a GCN Notice [70] describing the GW event in VOEvent [88] format (see fig. 3). Most of this delay was due to the LIGO policy of human-in-the-loop verification of trigger quality. After each VOEvent GCN Alert, a human-readable GCN Circular would follow after a time period of half an hour or greater.

The O2 multi-messenger pipeline disseminated results in the form of human-readable GCN Circulars sent out after the LIGO/Virgo GCN Circular. A GW+HEN GCN Circular was released after a LIGO/Virgo GCN Circular had been distributed in order to avoid confusing EM partners and to add a third confirmation that the event was valid (in addition to the EM\_READY GraceDB tag and

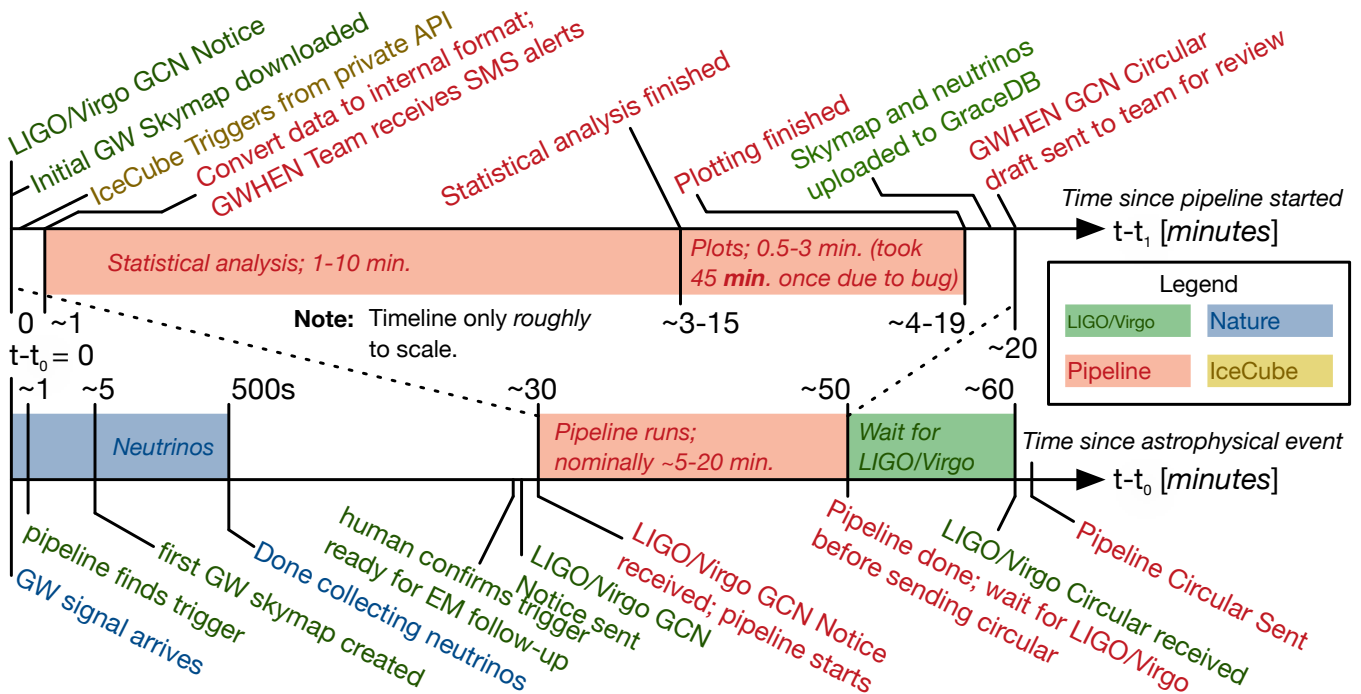


FIG. 3. A typical timeline for a LIGO/Virgo trigger. The bottom timeline shows the full chain of events, particularly the role of LIGO/Virgo. The party responsible for each event or data product is color-coded according to the legend. The time of signal arrival at Earth is denoted  $t_0$ . Since the O2 multi-messenger pipeline only ran on triggers marked by LIGO/Virgo as ready for EM followup (which required a human in the loop during O2 due to LIGO policy), there was a significant delay (usually 30 minutes) before processing began. The top timeline shows the detailed steps performed by the O2 multi-messenger pipeline during a typical event and their associated delays. The time at which the pipeline starts processing a trigger is denoted  $t_1$ . The calculation of the joint likelihood  $\mathcal{L}(\vec{x}_s)$  and the plotting of the joint skymap typically took up the longest (and most variable) amount of time. Retrieving and disseminating data took negligible amounts of time in comparison. There were instances during O2 where edge cases, changing data formats, and bugs caused exceptionally long delays (often requiring human intervention), but these are considered atypical and not included in this timeline since their causes were easily fixed. Minute details of the pipeline’s execution are also omitted from the timeline.

GCN Notice). This constraint allowed a comfortable target timeline of half an hour to run the combined analysis and human-in-the-loop checks while still being ready to distribute our Circular as soon as the LIGO/Virgo Circular was sent. By adhering to this timeline, the search would not introduce any extra latency compared to the localization provided by GW skymaps alone.

### C. Data Analysis Method

The analysis method used in this paper was an upgraded version of Baret et al. [18]. A HEN event, provided by the IceCube Collaboration, was automatically internally marked as coincident with a GW if it was detected within the  $t_{\text{GW}} \pm 500\text{s}$  time window (where  $t_{\text{GW}}$  is the time of the GW event) and if any region of the GW’s 90% confidence region had a neutrino signal likelihood density greater than  $10^{-4} \text{ deg}^{-2}$ . A joint skymap plot was also generated for each event and uploaded to the internal LIGO database (GraceDB). Interpretation of observations was then performed by a human and de-

scribed in an automatically-generated GCN circular for each event. Most of the terms required for a joint significance calculation (described in [18]) were calculated by LLAMA as described below. P-value calculations were not reported during O2.

#### 1. Coincidence Time Window

The  $t_{\text{GW}} \pm 500\text{s}$  time window was chosen to account for upper limits on observed GRB durations combined with central engine breakout time and precursor delay (as described in [12]). This time window makes no further model assumptions and allows for high energy neutrino and gravitational wave emissions in both the precursor and GRB.

#### 2. Neutrino Data Processing

In this study, we used a sample of through-going muons that were originally designed for the IceCube gamma-

ray follow-up (GFU) program [55, 85] as described in Sec. II B. We received these events in realtime with a latency of  $\sim 1$ min and accepted any event reconstructed as an up-going neutrino as well as events that were reconstructed as down-going neutrinos by applying a lower cut of 0.1 on their BDT score [87]. The neutrino arrival time, direction, angular uncertainty, BDT score, point spread function (PSF) and false alarm rate are used in our correlation analysis. The PSF ( $\mathcal{F}_\nu(\vec{x}_s)$ ) is the probability distribution of the neutrino source direction and can be explained with a Gaussian distribution around the true source location ( $\vec{x}_s$ ) as described in Eq. 6 of [18]. The HEN PSF term is incorporated in the likelihood function of the joint GW+HEN analysis. We define the signal likelihood as being equal to the PSF,

$$\mathcal{S}_\nu(\vec{x}_s) = \mathcal{F}_\nu(\vec{x}_s), \quad (1)$$

and the background likelihood is considered to be a flat distribution,

$$\mathcal{B}_\nu = \frac{1}{2\pi}. \quad (2)$$

### 3. Gravitational Wave Data Processing

The GW skymap  $\mathcal{F}_{\text{GW}}(\vec{x}_s)$  gives the probability per  $\text{deg}^2$  that the GW came from sky direction  $\vec{x}_s$ . These skymaps are calculated automatically by LIGO/Virgo in response to gravitational wave event triggers (discussed here). For performance reasons, probability was set to 0 in regions outside the 90% confidence region of the GW skymap (defined as the smallest region from which the GW had a 90% probability of originating), yielding a reduced skymap  $\mathcal{F}_{\text{GW}}^{90}(\vec{x}_s)$ . In the LLAMA code, pixels outside this region were removed before the main analysis was run (effectively setting them to zero).

We defined the signal likelihood as being equal to the reduced skymap,

$$\mathcal{S}_{\text{GW}}(\vec{x}_s) = \mathcal{F}_{\text{GW}}^{90}(\vec{x}_s). \quad (3)$$

Assuming that gravitational wave triggers caused by background noise have an isotropic distribution, the gravitational wave background likelihood is simply

$$\mathcal{B}_{\text{GW}} = \frac{1}{4\pi}. \quad (4)$$

### 4. Joint Likelihood and Coincidence

LLAMA calculated a joint likelihood ratio,

$$\mathcal{L}(\vec{x}_s) = \frac{\mathcal{S}_{\text{GW}}(\vec{x}_s)\mathcal{S}_\nu(\vec{x}_s)}{\mathcal{B}_{\text{GW}}\mathcal{B}_\nu}, \quad (5)$$

for each temporally coincident neutrino. This is a modified version of the formula described in [12] without a joint p-value calculation.

There were two other features present in the codebase and method paper that were not used during O2, namely, the neutrino clustering and galaxy catalog features. The clustering code, which accounted for the possibility of multiple neutrinos coming from the same source, was not used due to the low predicted multi-neutrino detection rate for O2. The galaxy catalog code, which used information on galaxy locations to further constrain source direction, was not used due to the lack of a galaxy catalog whose range matched LIGO's maximum BNS detection range in O2.

## D. Software Implementation

### 1. Overview

LLAMA (Fig. 4) triggered on LIGO/Virgo GCN Notices. Upon receiving a GCN Notice for a new LIGO/Virgo event, the pipeline responded by creating an event directory on the analysis server's filesystem for the new trigger. It then pulled data from LIGO/Virgo (via GraceDB) and IceCube (via their GFU API) as data became available; this included event metadata (in VO-Event form) as well as actual skymaps and reconstructed parameters. The pipeline would alert GW+HEN team members about the event; put input data into an internal representation; run the analysis described in IV C; upload results (neutrino triggers and a joint skymap plot) to GraceDB and [gw-astronomy.org](http://gw-astronomy.org); and, finally, send a joint skymap and GCN Circular draft to team members, who would send out the GCN Circular once the LIGO/Virgo GCN Circular had been distributed.

### 2. Architecture

The O2 pipeline was mostly implemented as a Python library (with the exception of the analysis described in IV C, which was a modified version of a reviewed MATLAB analysis codebase used in previous offline searches, see e.g., [1, 4, 5, 8, 17, 18, 26, 28, 30, 31, 49].) The steps in the pipeline (and their dependencies on one another) are described in software by a Directed Acyclic Graph (DAG), a simplified version of which is illustrated in Fig. 4. Two Python scripts running as daemon processes (`gcmd` and `gwhend`, described below) ran throughout O2 and automatically processed incoming events using the tools provided by the Python library.

Trigger acquisition was the first step of the pipeline and was accomplished by `gcmd`. This daemon would parse incoming GCN notices in order to identify new LIGO/Virgo gravitational wave triggers. Once a suitable trigger was received, `gcmd` would make a directory on the file system to hold data associated with that trigger, save the VOEvent received from GCN in that directory, and extract GW trigger metadata into a new file

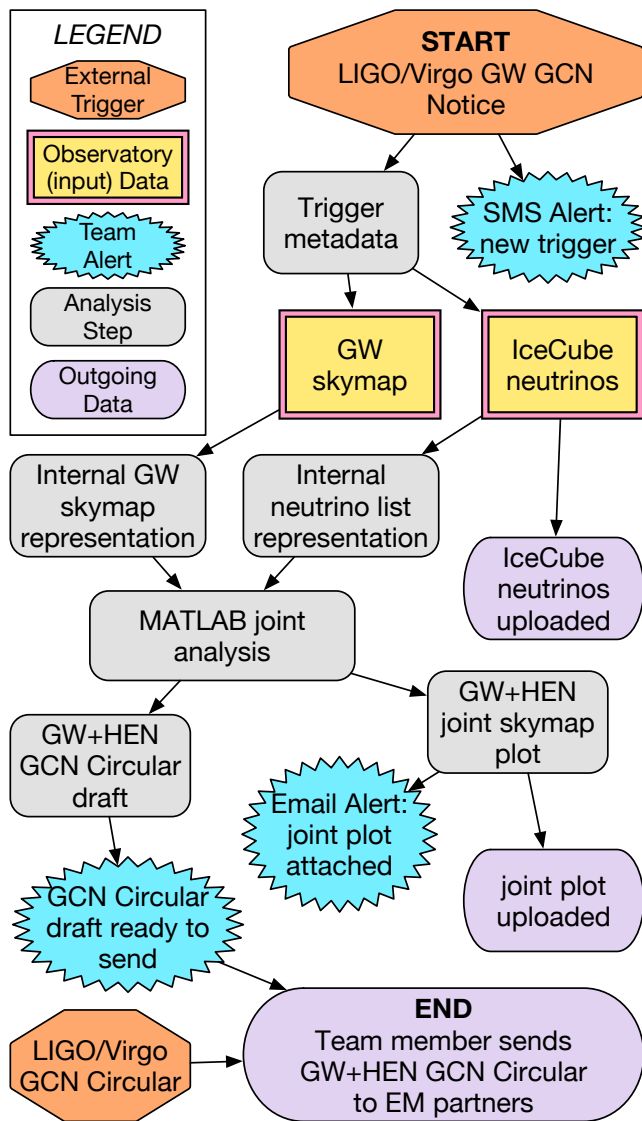


FIG. 4. The pipeline’s analysis steps as a DAG. Each node corresponds to an analysis output file (auxiliary steps have been omitted for clarity). *External triggers* are created outside the pipeline but ingested and stored internally. The pipeline will repeatedly try to download *observatory data* if not immediately available. *Team alerts* are used internally. *Outgoing data* are uploaded to GraceDB and [gw-astronomy.org](http://gw-astronomy.org) (except for the pipeline’s GCN Circular, which is distributed to GCN after internal review).

(also in the trigger directory) following LLAMA’s internal format. The pipeline is agnostic to trigger type and source and can accept triggers from alternative sources using a suitable trigger listener alongside (or instead of) `gcnd`. The combination with GWs was tested using a stream of realtime muon data provided by the IceCube Collaboration.

All subsequent steps of the pipeline were run automatically by `gwhend`, which would monitor the directory where all triggers were stored looking for partially-

analyzed triggers and seeing if any further steps of their analysis could be run. It accomplished this by periodically checking whether any non-existing output files from the DAG (Fig. 4) had input files available and then generating any such files. By iterating on this process, `gwhend` would push the analysis for any recent event triggers as far as possible while waiting for new data. Every other step described in [IV D 1](#) and Fig. 4 (data analysis, team alerts, and analysis output uploads) was run in this way with the exception of the final GW+HEN GCN Circular submission.

This final step was performed manually by a team member both to guard against spurious submission and to allow for tweaks to the GCN Circular manuscript to satisfy LIGO requirements. The pipeline would generate a draft GW+HEN GCN Circular and send it to team members, who could then modify it as needed and email it to GCN for distribution once the LIGO/Virgo GCN Circular (Sec. [IV B](#)) for that trigger had been distributed.

### 3. Features

The pipeline was required to be highly extensible, reliable, and reproducible. The input dependencies and generation procedure for each node/file in the DAG (Fig. 4) were defined within a Python class in the main LLAMA library.

Because each step in the procedure was self-contained with explicitly-defined inputs, new analysis steps and procedures could be added without affecting the overall stability of the pipeline. Furthermore, having explicit dependencies and procedures made it possible for `gwhend` to automatically execute steps of the pipeline once data became available (rather than executing them in a specific order), recover gracefully from errors, and precisely define pipeline state (in terms of existing vs. non-existing files) without introducing a complex development framework.

The pipeline would also send out detailed stack traces and error logs to the LLAMA team when exceptions occurred. This, combined with the independent nature of the analysis steps, proved useful at several points during O2, e.g., when external sources provided wrongly-formatted skymap data or when external API changes caused exceptions while fetching data. In each case, the pipeline was able to continue processing unaffected parts of each trigger’s analysis while fixes were implemented, minimizing delays.

Using a DAG also facilitated offline and manual analyses. Analyses could be rerun with updated data without regenerating all inputs or having to run the entire pipeline, and subcomponents of the analysis could be manually run offline even if unneeded inputs were missing.

The pipeline’s ability to handle arbitrary trigger sources (once trigger data is put into an internal format) will allow for both new types of analysis as well as



redundancy in trigger acquisition; a more flexible gravitational wave trigger-acquisition script using LIGO/Virgo's LVAlert system was tested alongside `gcnd` during O2 to confirm this.

In general, the flexibility and robustness of LLAMA will enable other planned upgrades to the pipeline, including parallel file generation (for increased performance), atomic step execution (for maximally robust fault recovery and easier development), and pipeline state snapshots (for rerunning analyses with different input data or analysis procedures).

## V. CONCLUSIONS

The O2 multi-messenger pipeline established a robust platform for discoveries and enabled a low-latency GW/HEN search during O2, ensuring that improved GW source localization could be provided in the event of a joint neutrino detection and that joint-detection candidates could be quickly found and analyzed.

No joint events were found during O2. Nonetheless, O2 provided a valuable opportunity to implement, test, and refine LLAMA, laying the groundwork for fast, reliable joint GW+HEN+GRB+MMA searches. during O2 have inspired important new features and performance improvements. Furthermore, the proof-of-concept provided by the O2 multi-messenger pipeline makes the framework a credibly reliable and performant component of more complex multi-messenger search strategies in the future.

The performance of the pipeline was within the speci-

fications required to ensure low-latency followup and was sufficiently fast to make the O2 pipeline the fastest online multi-messenger search in the O2 era, with GW+HEN GCN Circulars usually coming out immediately following the corresponding LIGO/Virgo GCN Circulars. Nonetheless, there are numerous straightforward optimizations that can further improve LLAMA performance by an order of magnitude. In particular, less conservative pipeline triggering (by subscribing to trigger updates directly from LIGO/Virgo) and optimized analysis code (using a faster implementation of the analysis method) will enable this goal.

## VI. ACKNOWLEDGMENTS

The authors thank the many researchers working in LIGO/Virgo, IceCube, and other astrophysics projects that enabled this pipeline to operate successfully in O2. In particular, they thank Scott Barthelmy and Leo Singer of NASA for providing helpful code and advice for working with GCN.

The authors are grateful to the IceCube Collaboration for providing neutrino dataset and support for using it for testing this algorithm. The Columbia Experimental Gravity group is grateful for the generous support of Columbia University in the City of New York and the National Science Foundation under grants PHY-1404462 and PHY-1708028. The authors are thankful for the generous support of the University of Florida and Columbia University in the City of New York.

---

\* Email: [stefan.countryman@ligo.org](mailto:stefan.countryman@ligo.org); <http://markalab.org>

- [1] Z. Marka *et al.*, (2006), LIGO Document G060660.
- [2] R. Seaman *et al.*, “Sky Event Reporting Metadata (VO-Event) Version 1.11,” IVOA Recommendation 1 November 2006 (2006).
- [3] S. Marka *et al.*, in *American Astronomical Society Meeting Abstracts*, Bulletin of the American Astronomical Society, Vol. 39 (2007) p. 868.
- [4] Y. Aso *et al.*, in *APS April Meeting Abstracts* (2008) p. E8.006.
- [5] Y. o. Aso, *Classical and Quantum Gravity* **25**, 114039 (2008), [arXiv:0711.0107](https://arxiv.org/abs/0711.0107).
- [6] B. Abbott *et al.*, *Classical and Quantum Gravity* **25**, 114051 (2008), [arXiv:0802.4320](https://arxiv.org/abs/0802.4320) [gr-qc].
- [7] S. Marka *et al.*, in *AAS/High Energy Astrophysics Division #10*, AAS/High Energy Astrophysics Division, Vol. 10 (2008) p. 15.01.
- [8] V. van Elewyck *et al.*, *International Journal of Modern Physics D* **18**, 1655 (2009), [arXiv:0906.4957](https://arxiv.org/abs/0906.4957) [astro-ph.IM].
- [9] Z. Marka *et al.*, in *American Astronomical Society Meeting Abstracts #215*, Bulletin of the American Astronomical Society, Vol. 42 (2010) p. 230.
- [10] Z. Marka, in *APS Meeting Abstracts* (2010) p. K13.004.
- [11] S. Márka *et al.*, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 243 (2010) p. 012001.
- [12] B. Baret *et al.*, *Astroparticle Physics* **35**, 1 (2011), [arXiv:1101.4669](https://arxiv.org/abs/1101.4669) [astro-ph.HE].
- [13] S. Márka *et al.*, *Classical and Quantum Gravity* **28**, 114013 (2011).
- [14] E. Chassande-Mottin *et al.*, *General Relativity and Gravitation* **43**, 437 (2011), [arXiv:1004.1964](https://arxiv.org/abs/1004.1964) [gr-qc].
- [15] R. Seaman *et al.*, “Sky Event Reporting Metadata Version 2.0,” IVOA Recommendation 11 July 2011 (2011), [arXiv:1110.0523](https://arxiv.org/abs/1110.0523) [astro-ph.IM].
- [16] I. Bartos *et al.*, *Physical Review Letters* **107**, 251101 (2011), [arXiv:1108.3001](https://arxiv.org/abs/1108.3001) [astro-ph.HE].
- [17] B. Baret *et al.*, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 363 (2012) p. 012022.
- [18] B. Baret *et al.*, *Phys. Rev. D* **85**, 103004 (2012).
- [19] I. Bartos, B. Dasgupta, and S. Márka, *Phys. Rev. D* **86**, 083007 (2012), [arXiv:1206.0764](https://arxiv.org/abs/1206.0764) [astro-ph.HE].
- [20] M. W. E. Smith *et al.*, *Astroparticle Physics* **45**, 56 (2013), [arXiv:1211.5602](https://arxiv.org/abs/1211.5602) [astro-ph.HE].
- [21] I. Bartos, P. Brady, and S. Márka, *Classical and Quantum Gravity* **30**, 123001 (2013), [arXiv:1212.2289](https://arxiv.org/abs/1212.2289) [astro-ph.CO].

- [22] S. Adrián-Martínez *et al.*, *J. Cosmology Astropart. Phys.* **6**, 008 (2013), arXiv:1205.3018 [astro-ph.HE].
- [23] I. Bartos *et al.*, *Physical Review Letters* **110**, 241101 (2013), arXiv:1301.4232 [astro-ph.HE].
- [24] S. Ando *et al.*, *Reviews of Modern Physics* **85**, 1401 (2013), arXiv:1203.5192 [astro-ph.HE].
- [25] I. Bartos and S. Márka, *Phys. Rev. D* **90**, 101301 (2014), arXiv:1409.1217 [astro-ph.HE].
- [26] M. G. Aartsen *et al.*, *Phys. Rev. D* **90**, 102002 (2014), arXiv:1407.1042 [astro-ph.HE].
- [27] I. Bartos and S. Márka, *Physical Review Letters* **115**, 231101 (2015), arXiv:1508.07810 [astro-ph.HE].
- [28] S. Adrián-Martínez *et al.*, *Phys. Rev. D* **93**, 122010 (2016), arXiv:1602.05411 [astro-ph.HE].
- [29] B. P. Abbott *et al.*, *ApJ* **848**, L12 (2017), arXiv:1710.05833 [astro-ph.HE].
- [30] A. Albert *et al.*, *ApJ* **850**, L35 (2017), arXiv:1710.05839 [astro-ph.HE].
- [31] A. Albert *et al.*, *Phys. Rev. D* **96**, 022005 (2017), arXiv:1703.06298 [astro-ph.HE].
- [32] I. Bartos *et al.*, *Phys. Rev. D* **96**, 023003 (2017), arXiv:1611.03861 [astro-ph.HE].
- [33] S. Márka, in *42nd COSPAR Scientific Assembly*, COSPAR Meeting, Vol. 42 (2018) pp. E1.14–12–18.
- [34] B. P. Abbott *et al.*, *Physical Review Letters* **119**, 161101 (2017), arXiv:1710.05832 [gr-qc].
- [35] A. von Kienlin, C. Meegan, and A. Goldstein, *GCN* **21520** (2017).
- [36] C. Meegan *et al.*, *ApJ* **702**, 791 (2009), arXiv:0908.0450 [astro-ph.IM].
- [37] V. Savchenko *et al.*, *GCN* **21507** (2017).
- [38] V. Savchenko *et al.*, *ApJ* **848**, L15 (2017), arXiv:1710.05449 [astro-ph.HE].
- [39] M. G. Aartsen *et al.*, *Science* **361**, eaat1378 (2018), arXiv:1807.08816 [astro-ph.HE].
- [40] A. Keivani *et al.*, *ApJ* **864**, 84 (2018), arXiv:1807.04537 [astro-ph.HE].
- [41] B. P. Abbott *et al.*, *Physical Review Letters* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [42] B. P. Abbott *et al.*, *Physical Review Letters* **116**, 241103 (2016), arXiv:1606.04855 [gr-qc].
- [43] B. P. Abbott *et al.*, *Physical Review Letters* **118**, 221101 (2017), arXiv:1706.01812 [gr-qc].
- [44] B. P. Abbott *et al.*, *ApJ* **851**, L35 (2017), arXiv:1711.05578 [astro-ph.HE].
- [45] B. P. Abbott *et al.*, *Physical Review Letters* **119**, 141101 (2017), arXiv:1709.09660 [gr-qc].
- [46] M. G. Aartsen *et al.*, *Physical Review Letters* **113**, 101101 (2014), arXiv:1405.5303 [astro-ph.HE].
- [47] IceCube Collaboration, *Science* **342**, 1242856 (2013), arXiv:1311.5238 [astro-ph.HE].
- [48] IceCube Collaboration, in *35th International Cosmic Ray Conference (ICRC2015)* (2017) p. 981.
- [49] A. Albert *et al.*, ArXiv e-prints (2018), arXiv:1810.10693 [astro-ph.HE].
- [50] T. Pradier, *Nuclear Instruments and Methods in Physics Research A* **602**, 268 (2009), arXiv:0807.2562.
- [51] C. Messick *et al.*, *Phys. Rev. D* **95**, 042001 (2017), arXiv:1604.04324 [astro-ph.IM].
- [52] A. Goldstein *et al.*, *ApJ* **848**, L14 (2017), arXiv:1710.05446 [astro-ph.HE].
- [53] I. Bartos, A. Crotts, and S. Márka, *ApJ* **801**, L1 (2015), arXiv:1410.0677 [astro-ph.HE].
- [54] P. Cowperthwaite *et al.*, *ApJ* **848**, L17 (2017), arXiv:1710.05840 [astro-ph.HE].
- [55] M. G. Aartsen *et al.*, *Journal of Instrumentation* **11**, P11009 (2016), arXiv:1610.01814 [hep-ex].
- [56] K. Murase, *Phys. Rev. D* **97**, 081301 (2018), arXiv:1705.04750 [astro-ph.HE].
- [57] K. Murase *et al.*, *ApJ* **651**, L5 (2006), astro-ph/0607104.
- [58] P. Mészáros, *Astroparticle Physics* **43**, 134 (2013), arXiv:1204.1897 [astro-ph.HE].
- [59] S. S. Kimura *et al.*, ArXiv e-prints (2018), arXiv:1805.11613 [astro-ph.HE].
- [60] S. S. Kimura *et al.*, *ApJ* **848**, L4 (2017), arXiv:1708.07075 [astro-ph.HE].
- [61] K. Ioka *et al.*, *ApJ* **633**, 1013 (2005), astro-ph/0503279.
- [62] S. Mereghetti, *A&A Rev.* **15**, 225 (2008), arXiv:0804.0250.
- [63] S. Razzaque, P. Mészáros, and E. Waxman, *Physical Review Letters* **93**, 181101 (2004), astro-ph/0407064.
- [64] S. Ando and J. F. Beacom, *Physical Review Letters* **95**, 061103 (2005), astro-ph/0502521.
- [65] I. Bartos *et al.*, *ApJ* **835**, 165 (2017), arXiv:1602.03831 [astro-ph.HE].
- [66] K. Murase *et al.*, *ApJ* **822**, L9 (2016), arXiv:1602.06938 [astro-ph.HE].
- [67] N. C. Stone, B. D. Metzger, and Z. Haiman, *MNRAS* **464**, 946 (2017), arXiv:1602.04226.
- [68] R. Perna, D. Lazzati, and B. Giacomazzo, *ApJ* **821**, L18 (2016), arXiv:1602.05140 [astro-ph.HE].
- [69] I. Bartos *et al.*, *Nature Communications* **8**, 831 (2017), arXiv:1701.02328 [astro-ph.HE].
- [70] S. D. Barthelmy *et al.*, *Ap&SS* **231**, 235 (1995).
- [71] B. P. Abbott *et al.*, *Reports on Progress in Physics* **72**, 076901 (2009), arXiv:0711.3041 [gr-qc].
- [72] T. Accadia *et al.*, *Journal of Instrumentation* **7**, P03012 (2012).
- [73] D. V. Martynov *et al.*, *Phys. Rev. D* **93**, 112004 (2016), arXiv:1604.00439 [astro-ph.IM].
- [74] J. Aasi *et al.*, *Classical and Quantum Gravity* **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
- [75] B. P. Abbott *et al.*, *Living Reviews in Relativity* **21**, 3 (2018), arXiv:1304.0670 [gr-qc].
- [76] F. Acernese *et al.*, *Classical and Quantum Gravity* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [77] S. Klimenko *et al.*, *Classical and Quantum Gravity* **25**, 114029 (2008), arXiv:0802.3232 [gr-qc].
- [78] S. Klimenko *et al.*, *Phys. Rev. D* **93**, 042004 (2016), arXiv:1511.05999 [gr-qc].
- [79] S. A. Usman *et al.*, *Classical and Quantum Gravity* **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [80] R. Lynch *et al.*, ArXiv e-prints (2015), arXiv:1511.05955 [gr-qc].
- [81] T. Adams *et al.*, *Classical and Quantum Gravity* **33**, 175012 (2016), arXiv:1512.02864 [gr-qc].
- [82] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, and *et al.*, ArXiv e-prints (2018), arXiv:1811.12907 [astro-ph.HE].
- [83] M. G. Aartsen *et al.*, *Journal of Instrumentation* **12**, P03012 (2017), arXiv:1612.05093 [astro-ph.IM].
- [84] M. G. Aartsen *et al.*, *ApJ* **835**, 151 (2017), arXiv:1609.04981 [astro-ph.HE].
- [85] M. G. Aartsen *et al.*, *Astroparticle Physics* **92**, 30 (2017), arXiv:1612.06028 [astro-ph.HE].
- [86] Y. Freund and E. Schapire, *Journal of Computer and*

- [System Sciences](#) **55**, 119 (1997).
- [87] M. G. Aartsen *et al.*, *ApJ* **796**, 109 (2014), [arXiv:1406.6757 \[astro-ph.HE\]](#).
- [88] R. D. Williams and R. Seaman, in *Astronomical Data Analysis Software and Systems XV*, Astronomical Society of the Pacific Conference Series, Vol. 351, edited by C. Gabriel, C. Arviset, D. Ponz, and S. Enrique (2006) p. 637.
- [89] V. van Elewyck *et al.*, in *ICRC2015*, International Cosmic Ray Conference, Vol. 34 (2015) p. 1112.
- [90] I. Bartos, P. Veres, D. Nieto, V. Connaughton, B. Humensky, K. Hurley, S. Márka, P. Mészáros, R. Mukherjee, P. O'Brien, and J. P. Osborne, *MNRAS* **443**, 738 (2014), [arXiv:1403.6119 \[astro-ph.HE\]](#).
- [91] I. Bartos, T. Di Girolamo, J. R. Gair, M. Hendry, I. S. Heng, T. B. Humensky, S. Márka, Z. Márka, C. Messenger, R. Mukherjee, D. Nieto, P. O'Brien, and M. Santander, *MNRAS* **477**, 639 (2018), [arXiv:1802.00446 \[astro-ph.HE\]](#).
- [92] M. G. Aartsen *et al.*, *ArXiv e-prints* (2015), [arXiv:1510.05228 \[astro-ph.IM\]](#).