

SAGUARO: Time-domain Infrastructure for the Fourth Gravitational-wave Observing Run and Beyond

Griffin Hosseinzadeh¹ , Kerry Paterson² , Jillian C. Rastinejad³ , Manisha Shrestha¹ , Philip N. Daly¹, Michael J. Lundquist⁴ , David J. Sand¹ , Wen-fai Fong³ , K. Azalee Bostroem^{1,9} , Saarah Hall³ , Samuel D. Wyatt⁵ , Alex R. Gibbs⁶ , Eric Christensen⁷ , William Lindstrom⁸ , Jonathan Nation⁸ , Joseph Chatelain⁸ , and Curtis McCully⁸

¹ Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721-0065, USA; griffin0@arizona.edu

² Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, D-85748 Garching, Germany

³ Center for Interdisciplinary Exploration and Research in Astrophysics and Department of Physics and Astronomy, Northwestern University, 1800 Sherman Avenue, Floor 8, Evanston, IL 60201-3497, USA

⁴ W.M. Keck Observatory, 65-1120 Māmalahoa Highway, Kamuela, HI 96743-8431, USA

⁵ Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195-0002, USA

⁶ Lunar and Planetary Laboratory, University of Arizona, 1629 East University Boulevard, Tucson, AZ 85721-0092, USA

⁷ Rubin Observatory Project Office, 950 North Cherry Avenue, Tucson, AZ 85719-4933, USA

⁸ Las Cumbres Observatory, 6740 Cortona Drive, Suite 102, Goleta, CA 93117-5575, USA

Received 2023 October 12; revised 2024 January 16; accepted 2024 January 18; published 2024 March 13

Abstract

We present upgraded infrastructure for Searches After Gravitational waves Using ARizona Observatories (SAGUARO) during LIGO, Virgo, and KAGRA’s fourth gravitational-wave (GW) observing run (O4). These upgrades implement many of the lessons we learned after a comprehensive analysis of potential electromagnetic counterparts to the GWs discovered during the previous observing run. We have developed a new web-based target and observation manager (TOM) that allows us to coordinate sky surveys, vet potential counterparts, and trigger follow-up observations from one centralized portal. The TOM includes software that aggregates all publicly available information on the light curves and possible host galaxies of targets, allowing us to rule out potential contaminants like active galactic nuclei, variable stars, solar system objects, and preexisting supernovae, as well as to assess the viability of any plausible counterparts. We have also upgraded our image-subtraction pipeline by assembling deeper reference images and training a new neural-network-based real–bogus classifier. These infrastructure upgrades will aid coordination by enabling the prompt reporting of observations, discoveries, and analysis to the GW follow-up community, and put SAGUARO in an advantageous position to discover kilonovae in the remainder of O4 and beyond. Many elements of our open-source software stack have broad utility beyond multimessenger astronomy, and will be particularly relevant in the “big data” era of transient discoveries by the Vera C. Rubin Observatory.

Unified Astronomy Thesaurus concepts: [Astronomy software \(1855\)](#); [Transient detection \(1957\)](#); [Sky surveys \(1464\)](#); [Neutron stars \(1108\)](#); [Gravitational waves \(678\)](#); [Compact binary stars \(283\)](#)

Supporting material: interactive figures

1. Introduction

The discovery of a γ -ray burst (GRB) and a kilonova associated with gravitational waves (GWs) from the binary neutron-star (BNS) merger GW170817 ushered in a new era of multimessenger astronomy (Arcavi et al. 2017; Coulter et al. 2017; LIGO Scientific Collaboration et al. 2017b; Lipunov et al. 2017; Soares-Santos et al. 2017; Tanvir et al. 2017; Valenti et al. 2017). Among many other accomplishments, these observations confirmed that at least some GRBs come from neutron-star mergers (Goldstein et al. 2017; LIGO Scientific Collaboration et al. 2017c; Savchenko et al. 2017), revealed evidence of r -process nucleosynthesis in compact-object mergers (Chornock et al. 2017; McCully et al. 2017; Nicholl et al. 2017; Pian et al. 2017; Shappee et al. 2017; Smartt et al. 2017; Tanvir et al. 2017; Watson et al. 2019), and

yielded the first measurement of the Hubble constant via the standard-siren method (LIGO Scientific Collaboration et al. 2017a). Despite this breakthrough, little is known observationally about what diversity exists among kilonovae, caused either by differences in the progenitor system or by viewing-angle effects. The handful of kilonovae discovered in association with GRBs, and therefore viewed nearly pole-on, hints at an intrinsic luminosity span of up to two orders of magnitude (Gompertz et al. 2018; Rastinejad et al. 2021, 2022a; Troja et al. 2022; Yang et al. 2022, 2023; Gillanders et al. 2023; Levan et al. 2023). A larger sample of off-axis kilonovae, discovered in association with GW events, is needed to fill in this parameter space.

The Laser Interferometer GW Observatory (LIGO) and the Virgo GW detectors undertook their third observing run (O3) from 2019 April 1 to 2020 March 27, yielding detections of dozens of compact-object mergers, most of which were mergers of binary black holes (BBH; LIGO Scientific Collaboration & Virgo Collaboration 2021a, 2021b; LIGO Scientific Collaboration et al. 2023). Figure 1 summarizes the real-time alerts sent in O3 and a final assessment of their sources. Five events were promising targets for follow-up with electromagnetic (EM)

⁹ LSSTC Catalyst Fellow.

 Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

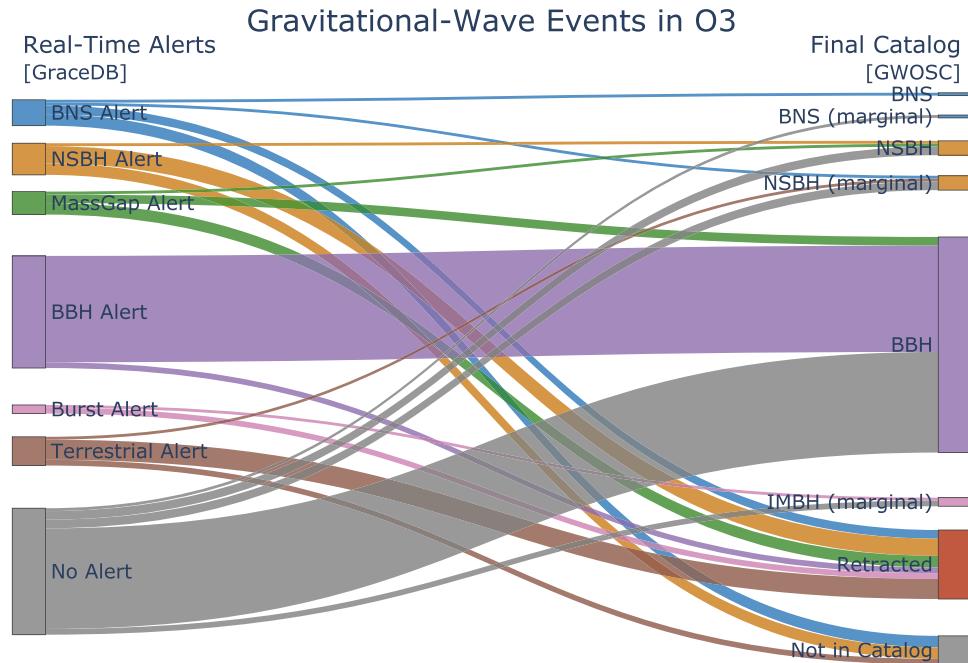


Figure 1. Sankey diagram showing the outcome of real-time alerts and the provenance of astrophysical GW events during LIGO and Virgo’s third observing run, using data from the Gravitational-Wave Candidate Event Database (<https://gracedb.ligo.org/superevents/public/O3/>) and the Gravitational Wave Open Science Center (<https://gwosc.org/eventapi/html/query/show?min-gps-time=1230336018&lastver=true>). Compact binary coalescence alerts are labeled by their highest probability source classification. Events with data-quality issues identified in real time are labeled as “retracted.” Events with a probability of astrophysical origin $p_{\text{astro}} < 50\%$ after further analysis are labeled as “marginal.” Events with a false-alarm rate of $\text{FAR} > 2 \text{ yr}^{-1}$ after further analysis are not included in the final catalog (LIGO Scientific Collaboration & Virgo Collaboration 2021a; LIGO Scientific Collaboration et al. 2023). Only three confirmed and two marginal events containing an NS had alerts allowing follow-up in real time. SAGUARO covered a portion of the localization region for four of these events and discovered a single possible counterpart, which was unlikely to be a kilonova based on its luminosity (Paterson et al. 2021). An interactive version of this figure that includes counts and event names for each flow is available.

telescopes, in that the lower-mass component of the binary had a mass $< 3 M_{\odot}$, and a detection alert was sent in nearly real time. Of those five, two “confident” and two “marginal” (referring to astrophysical significance) events were confirmed to be neutron-star–black-hole (NSBH) mergers: GW190426_152155, GW190814, GW200105_162426, and GW200115_042309 (LIGO Scientific Collaboration & Virgo Collaboration 2020, 2021a, 2021b; LIGO Scientific Collaboration et al. 2021, 2023). The fifth (GW190425) was confirmed as a confident BNS merger (Abbott et al. 2020). Despite copious follow-up, no EM transients were confirmed as GW counterparts (Andreoni et al. 2019, 2020; Coughlin et al. 2019; Dobie et al. 2019, 2022; Goldstein et al. 2019; Gomez et al. 2019; Hosseinzadeh et al. 2019; Lundquist et al. 2019; Antier et al. 2020a, 2020b; Ackley et al. 2020; Garcia et al. 2020; Gompertz et al. 2020; Kasliwal et al. 2020; Morgan et al. 2020; Pozanenko et al. 2020; Thakur et al. 2020; Vieira et al. 2020; Watson et al. 2020; Alexander et al. 2021; Anand et al. 2021; Becerra et al. 2021; Bhakta et al. 2021; Chang et al. 2021; de Wet et al. 2021; Dichiara et al. 2021; Kilpatrick et al. 2021; Oates et al. 2021; Ohgami et al. 2021; Paterson et al. 2021; de Jaeger et al. 2022; Tucker et al. 2022; see Rastinejad et al. 2022b for a comprehensive discussion).

Our EM follow-up program, Searches After Gravitational waves Using ARizona Observatories (SAGUARO), surveys the localization regions of GW events by partnering with the Catalina Sky Survey (CSS), which routinely searches for near-Earth asteroids (Christensen et al. 2019). CSS’s 1.5 m telescope on Mt. Lemmon, Arizona, surveys fixed 5 deg^2 fields with a typical unfiltered limiting magnitude of ~ 21.5 mag. In the event of a GW alert requiring urgent follow-up, we have the ability to send a ranked list of fields to the CSS’s scheduling

software, which will interrupt normal survey operations on the 1.5 m to observe the GW localization as soon as possible. Serendipitous CSS observations may also cover parts of GW localization regions without our request.

If we or others discover a plausible EM counterpart, we will trigger targeted follow-up observations under our approved program on any or all of the following observatories:

1. the 2×8.4 m Large Binocular Telescope on Mt. Graham, Arizona, USA,
2. the 6.5 m Magellan Telescopes at Las Campanas Observatory, Chile,
3. the 6.5 m MMT on Mt. Hopkins, Arizona, USA,
4. the 2.3 m Bok Telescope on Kitt Peak, Arizona, USA,
5. the 1.8 m Vatican Advanced Technology Telescope on Mt. Graham, Arizona, USA, and
6. the 1.5 m Kuiper Telescope on Mt. Bigelow, Arizona, USA.

We also participate in active GW follow-up programs on several facilities worldwide, which protects our effort against bad weather at a single site and provides us the ability to obtain follow-up photometry and spectroscopy of potential kilonovae within minutes to hours of discovery.

Lundquist et al. (2019) present an overview of the SAGUARO’s infrastructure as it stood at the beginning of O3, and Paterson et al. (2021) present an analysis of the O3 follow-up effort, including 17 GW events with CSS data either triggered by our program or observed serendipitously as part of the normal near-Earth asteroid survey. We continue this series of papers by describing new and improved hardware and software infrastructure to be used by SAGUARO during LIGO,

Virgo, and KAGRA’s fourth observing run (O4), which started on 2023 May 24. In Section 2, we present our new target and observation manager (TOM), an interactive web front end and database back end for carrying out the project. In Section 3, we describe updates to our image-subtraction pipeline and machine-learning algorithm for real-bogus classification. In Section 4, we outline our future plans, including progress toward adding a new instrument to our program. In Section 5, we conclude by discussing our prospects for discovering the second kilonova during O4. All of our software, including installation instructions, is publicly available on GitHub¹⁰ and archived on Zenodo (Daly et al. 2023; Hosseinzadeh et al. 2023b; Hosseinzadeh & Shrestha 2023; Paterson et al. 2023; Rastinejad & Hosseinzadeh 2023).

2. Target and Observation Manager

The complexity of our observing program and the necessity of responding in real time to alerts and incoming data require that we centralize all of our tools for planning observations, monitoring data, vetting targets, and reporting results in a single location that is accessible to all team members at all times. Current transient surveys and large follow-up programs use similar systems, including the Asteroid Terrestrial-impact Last Alert System (ATLAS) Transient Science Server (Smith et al. 2020), the Global Rapid Advanced Network Devoted to the Multi-messenger Addicts’ Interface and Communication for Addicts of the Rapid follow-up in multi-messenger Era (Antier et al. 2020a), the Global Relay of Observatories Watching Transients Happen Marshal (Kasliwal et al. 2019), the Global Supernova Project’s Supernova Exchange (Howell & Global Supernova Project 2017), the Public ESO Spectroscopic Survey of Transient Objects Marshall (Smartt et al. 2015), the Young Supernova Experiment’s YSE-PZ (Coulter et al. 2023), and the Zwicky Transient Facility’s Fritz (based on SkyPortal; Coughlin et al. 2023).

We call our portal a target and observation manager (TOM), after Street et al. (2018). The SAGUARO TOM is a browser-based front end on top of a PostgreSQL 14 database,¹¹ built using the Django¹²-based TOM Toolkit (Lindstrom et al. 2022) and its affiliated packages. The TOM is hosted on a dedicated server located on the University of Arizona campus and is available over the web to all members of our team with login credentials. In the following sections, we describe the various aspects of our observing campaign that have now been integrated into the TOM.

2.1. Receiving GW Alerts

Our workflow begins with the receipt of an alert from the LIGO/Virgo/KAGRA Collaboration notifying us of a GW event. Starting in O4, the International Gravitational Wave Observatory Network¹³ has been sending alerts over Hopskotch,¹⁴ an Apache Kafka¹⁵-based alert platform developed by the Scalable Cyberinfrastructure to support Multi-Messenger Astrophysics (SCIMMA) project.¹⁶ Two packages affiliated

with the TOM Toolkit allow us to receive and handle these alerts. The `tom-alertstreams` app¹⁷ provides a management command `readstreams` that listens for Hopskotch alerts and passes the alert contents to a user-defined handler function. The `tom_nonlocalizedevents` app¹⁸ provides basic handler functions that store alert information in the TOM database and web templates to display information for a given GW event, including the contents of related circulars from the General Coordinates Network (GCN)¹⁹ retrieved via SCIMMA’s Hermes message exchange service²⁰ (Figure 2; see also Section 2.7.3).

We customized the alert handler functions in our TOM to provide additional functionality. We have the ability to send human-readable summaries of each event over short message service (SMS) through Twilio²¹ and Slack through a custom app²² (Figure 3). We divide incoming alerts into five streams, with the conditions in parentheses evaluated in the following order:

1. test alerts (if the event name starts with “MS”),
2. subthreshold alerts (if `significant=False`),
3. burst alerts (if `group='Burst'`),
4. BBH alerts (if `HasNS < 1%` and `NSBH < 1%` and `BNS < 1%`), and
5. NS alerts (otherwise).

Each user can choose which SMS streams to subscribe to via their profile page in the TOM, and streams 2–5 each have dedicated channels in our Slack workspace. NS alerts push notifications to all members of the channel with the `@channel` command.

The new Kafka alerts contain multiorder coverage (MOC; Fernique et al. 2014) maps of the localization area in the Hierarchical Equal Area isoLatitude Pixelization (HEALPix) format (Górski et al. 2005). The `tom_nonlocalizedevents` package stores these maps in the TOM database with the help of HEALPix Alchemy (Singer et al. 2022). For alerts containing a new sky map, including maps generated by external coincidence with a GRB, we calculate and store the innermost credible region (among 25%, 50%, 75%, 90%, and 95%) that contains the coordinates of each linked target (see Section 2.4) and the center of each survey field. For computational efficiency, we use the center of the field as a proxy for all future sources we discover in that field. We also calculate and store the 90% credible region contour itself, which we use for display (see Section 2.2). Last, we calculate and store the probability contained within the footprints of each of our survey fields, which we use to prioritize our observation requests. This last calculation requires rasterization of the MOC map using `healpy` (Zonca et al. 2019) and `ligo.skymap` (Singer & Price 2016).

2.2. Planning and Triggering a Search

Once we decide to follow up a GW event, we use the TOM to plan our survey and send instructions to the CSS 1.5 m telescope. CSS observes in groups of 12 adjacent fields as part of their near-

¹⁰ <https://github.com/SAGUARO-MMA/>

¹¹ <https://www.postgresql.org/>

¹² <https://www.djangoproject.com/>

¹³ <https://emfollow.docs.ligo.org/userguide/>

¹⁴ <https://scimma.org/hopskotch>

¹⁵ <https://kafka.apache.org/>

¹⁶ <https://scimma.org/>

¹⁷ <https://github.com/TOMToolkit/tom-alertstreams>

¹⁸ https://github.com/TOMToolkit/tom_nonlocalizedevents

¹⁹ <https://gcn.nasa.gov/>

²⁰ <https://hermes.lco.global/?topic=gcn.circular>

²¹ <https://www.twilio.com/>

²² <https://api.slack.com/start/apps>

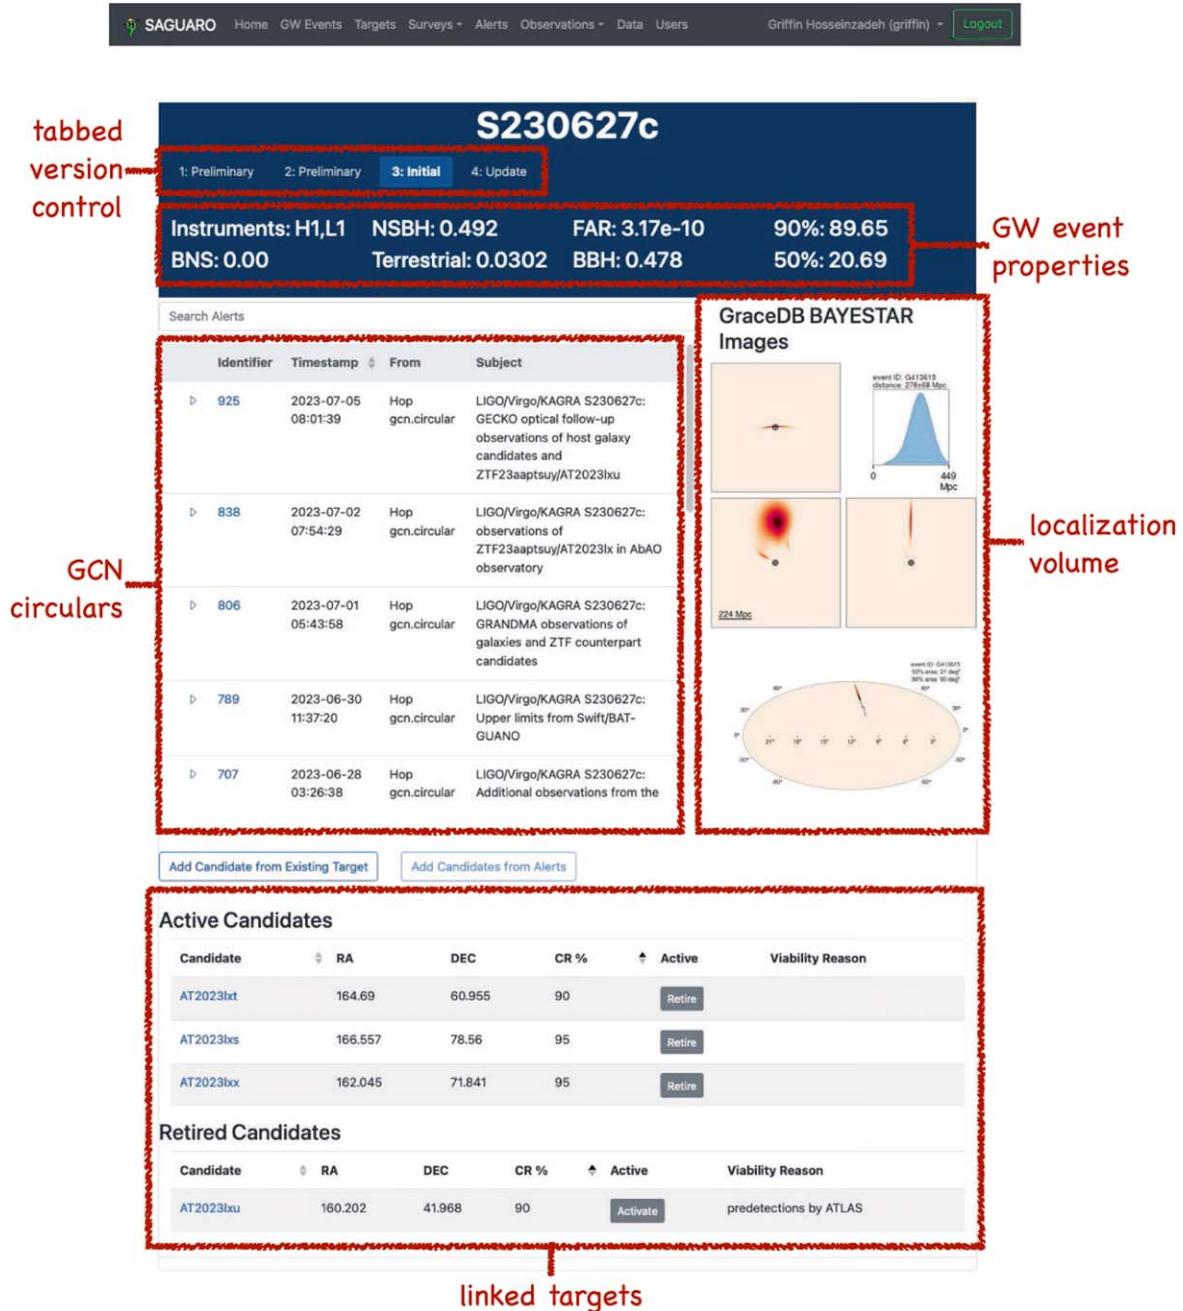


Figure 2. GW event detail page from the SAGUARO TOM, showing the source classification probabilities, false-alarm rate, localization areas, and GCN circulars for event S230627c. The tabbed interface switches between different versions of the alert information. The active and retired candidates lists at the bottom of the page link to targets that have been suggested as low-probability EM counterparts to the GW event by Anumarlapudi et al. (2023).

Earth asteroid survey, so we generate survey plans using the following algorithm: starting with the highest probability field in the localization region, we select the next highest probability adjacent field that is observable between now and the next sunrise (using 12° twilight), taking into account an airmass limit of 1.75 (to allow time for the full 12-field pattern to complete) and a phase-dependent moon exclusion region of $(3 + 42\phi)^\circ$ while the moon is above the horizon, where $\phi \in [0, 1]$ is the moon phase. We use the Astropy coordinates package (Astropy Collaboration 2022) for visibility calculations and Astroplan (Morris et al. 2018) to look up the moon phase. We iterate, excluding fields for which we do not have reference images (see Section 3.1), until we reach 12 fields (60 deg^2). Depending on the size and position of the

localization region, we may generate and submit multiple sets of 12 fields. To visualize our survey plans, we embed version 3 of the Aladin Lite sky viewer (Baumann et al. 2022), which natively reads HEALPix maps, into our TOM’s survey planning page. An example is shown in Figure 4.

Survey plans can be manually edited by selecting or deselecting fields using checkboxes in the table. Users can then click the yellow “Submit” button on the survey planning page to submit the list of fields to CSS. This button opens one or more files on CSS’s server using Paramiko²³ and writes lists of 12 field names. These fields are then ingested by CSS’s

²³ <https://www.paramiko.org/>

GW Alerts APP 15:39

BBH PRELIMINARY v1
S230927be (significant)
2023-09-27T15:38:32.919Z
1/FAR = 69435 Qyr
Distance = 1.1 ± 0.3 Gpc
50% Area = 119 deg^2
90% Area = 539 deg^2
Has NS = 0%
Has Mass Gap = 0%
Has Remnant = 0%
BNS = 0%
NSBH = 0%
BBH = 100%
Terrestrial = 0%
[SAGUARO Hermes GraceDB Treasure Map](#)

BBH PRELIMINARY v2
S230927be (significant)
2023-09-27T15:38:32.919Z
1/FAR = 100 yr
Distance = 1.1 ± 0.3 Gpc
50% Area = 136 deg^2
90% Area = 527 deg^2
Has NS = 0%
Has Mass Gap = 0%
Has Remnant = 0%
BNS = 0%
NSBH = 0%

+ Message #alerts-bbh

Home DMs Activity

Figure 3. Real-time human-readable alert summary sent by the SAGUARO TOM, as viewed on the Slack mobile app. The message time stamp shows that the first alert arrived within ~ 1 minute of the detection. The summary contains helpful information for determining whether follow-up is appropriate, including the inverse of the false-alarm rate (1/FAR), the distance to the source, the areas of the 50% and 90% localization regions, and, for compact binary coalescences, the probabilities of the progenitor system containing at least one object with mass $< 3 M_\odot$ (Has NS) or $3-5 M_\odot$ (Has Mass Gap), the event having an EM counterpart (Has Remnant), and the signal coming from a BNS merger, NSBH merger, BBH merger, or terrestrial noise source. More information is easily accessed by following the hyperlinks at the bottom of the summary.

scheduling software and observed as soon as they are visible. Users can also click the teal “Export” button to download a copy of the field lists to their local computer.

As we have plans to expand to other telescopes (see Section 4), we have taken care to make our survey planning infrastructure as general as possible. Additional survey fields can be added for different facilities, including fields centered on galaxies for imagers with small fields of view. We plan to

release this module as a `tom_surveys` app for the TOM Toolkit. The core of this app consists of the `SurveyField` and `SurveyObservationRecord` models, which are analogous to the `Target` and `ObservationRecord` models included with the base TOM Toolkit but without the connection to a single pair of coordinates implied by a “target.”

2.3. Scanning Survey Candidates

We copy all CSS images to our data-reduction server at 5 minute intervals, regardless of whether we requested them. In some cases, their ongoing near-Earth asteroid survey will overlap with GW localization regions by coincidence; we refer to these as “serendipitous” observations. Paterson et al. (2021) describe our data-reduction pipeline in detail; it performs image subtraction and detects new sources in the difference images, which we refer to as “survey candidates” because, at this stage, they may either be real transients or artifacts from the subtraction. We use the optimal image-subtraction algorithm of Zackay et al. (2016), which takes in a new image and a deep reference image of the same field and produces a raw difference image and a significance image corrected for astrometric noise (called S_{corr}). We refer to this set of four images as “image quadruplets” below. Updates to the reference images and real-bogus machine-learning classifier are discussed in Section 3.

Once detections are in the database, they can be viewed, filtered, and sorted in the TOM. The survey candidates page (Figure 5) shows the observed fields and candidate point sources overlaid on the localization region in the Aladin sky map for a chosen GW event, followed by a table of sources with their real–bogus scores (see Section 3.2) and image quadruplets (see Section 3.3). By default, we display sources within the 95% localization region of the event observed between 0 and 3 days after the merger time, sorted by decreasing “real” score, although further filtering and sorting is available via the form above the table. It is also possible to remove the filtering by localization region and view all detections from the previous night, as would be done for a typical all-sky time-domain survey. If a detection looks real, the user can link it to the GW event by clicking the blue “Link” button below its temporary target name, effectively bookmarking it for further vetting.

2.4. Vetting Potential Kilonovae

Motivated by the large number of contaminating sources found during searches following O3 merger events, we incorporate a number of publicly available tools into our infrastructure that automatically vet potential kilonovae, from both our own survey and others. In Rastinejad et al. (2022b), we evaluated the performance of each tool on potential O3 kilonovae and made recommendations for including automatic vetting in O4 follow-up. Here, we describe the components of our Kilonova Candidate Vetting package,²⁴ which implements the procedures described in our previous work (Rastinejad et al. 2022b) for SAGUARO’s O4 follow-up effort.

We have downloaded and ingested each of the catalogs we use for vetting into the TOM database, so that we can quickly run queries for large numbers of targets, and indexed each catalog table using the Quad Tree Cube (Q3C; Koposov & Bartunov 2006) PostgreSQL extension for efficient cone

²⁴ <https://github.com/SAGUARO-MMA/kne-cand-vetting>

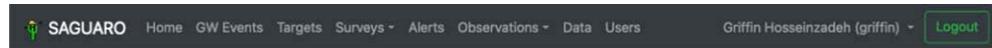


Figure 4. Survey planning page on the SAGUARO TOM, including an Aladin Lite viewer showing three groups of 12 planned fields (red, green, and blue), the 90% credible region contour for S230904n (purple), and the current positions of the Sun (gold) and moon (silver). The radii of the Sun and moon markers roughly represent their respective exclusion regions, although these are distorted with respect to the Mollweide projection. The purple dot near the top left of the map shows AT 2023rkw, suggested as a possible EM counterpart by Necker et al. (2023) but then classified as a Type Ia supernova by Anand et al. (2023). The planned fields are listed in a table at the bottom of the page, with their probability, rank order in our plan, time requested (if any), and time first observable in UT.

searches. The entire procedure below, with the exception of ATLAS forced photometry, has been wrapped into a single function, which is run automatically in two situations: when a button is pressed on the target detail page, useful for potential kilonovae from our own survey that already exist in the TOM, or when a new target is added through the TOM web interface, useful for potential kilonovae reported by other surveys. Because ATLAS forced photometry can take several minutes, we run it asynchronously on the press of a separate button.

2.4.1. Point-source Matching

We begin by crossmatching each target with several point-source catalogs to identify known variable sources (namely, quasars, active nuclei, and variable stars) that may masquerade as kilonovae. We perform a cone search on each catalog using a $2''$ radius, chosen to match the typical seeing of ground-based surveys, and then apply additional cuts based on the specification of each catalog. Where a match is determined, we display the angular offset, catalog probability score (where available), and

source class (e.g., variable star type, where available) on the target detail page (see Figure 6). This page also includes a view of the field around the target via an embedded Aladin Lite viewer (version 2; Boch & Fernique 2014), which we can use to identify nearby point sources or false matches by eye.

To identify quasar matches, we utilize the Million Quasars Catalog (Milliquas; Flesch 2021) compilation, which includes quasars from SDSS DR16 (Ahumada et al. 2020). Following Flesch (2015) we rule out targets as potential kilonovae if they have a $>97\%$ probability of being a quasar. To search for variable star contaminants, we employ three unique catalogs. First, we look for matches within the Gaia Data Release 3 variable star catalog (DR3; Gaia Collaboration 2023; Rimoldini et al. 2023; updated from using eDR3 in Rastinejad et al. 2022b). In addition, we crossmatch with the All-Sky Automated Survey for Supernovae (ASAS-SN; Jayasinghe et al. 2019) variable star catalog. We rule out all targets as kilonovae that have matches in either catalog. Finally, we search the Pan-STARRS1 (PS1) 3π Survey (Chambers et al. 2016) point-source catalog (Tachibana & Miller 2018). The

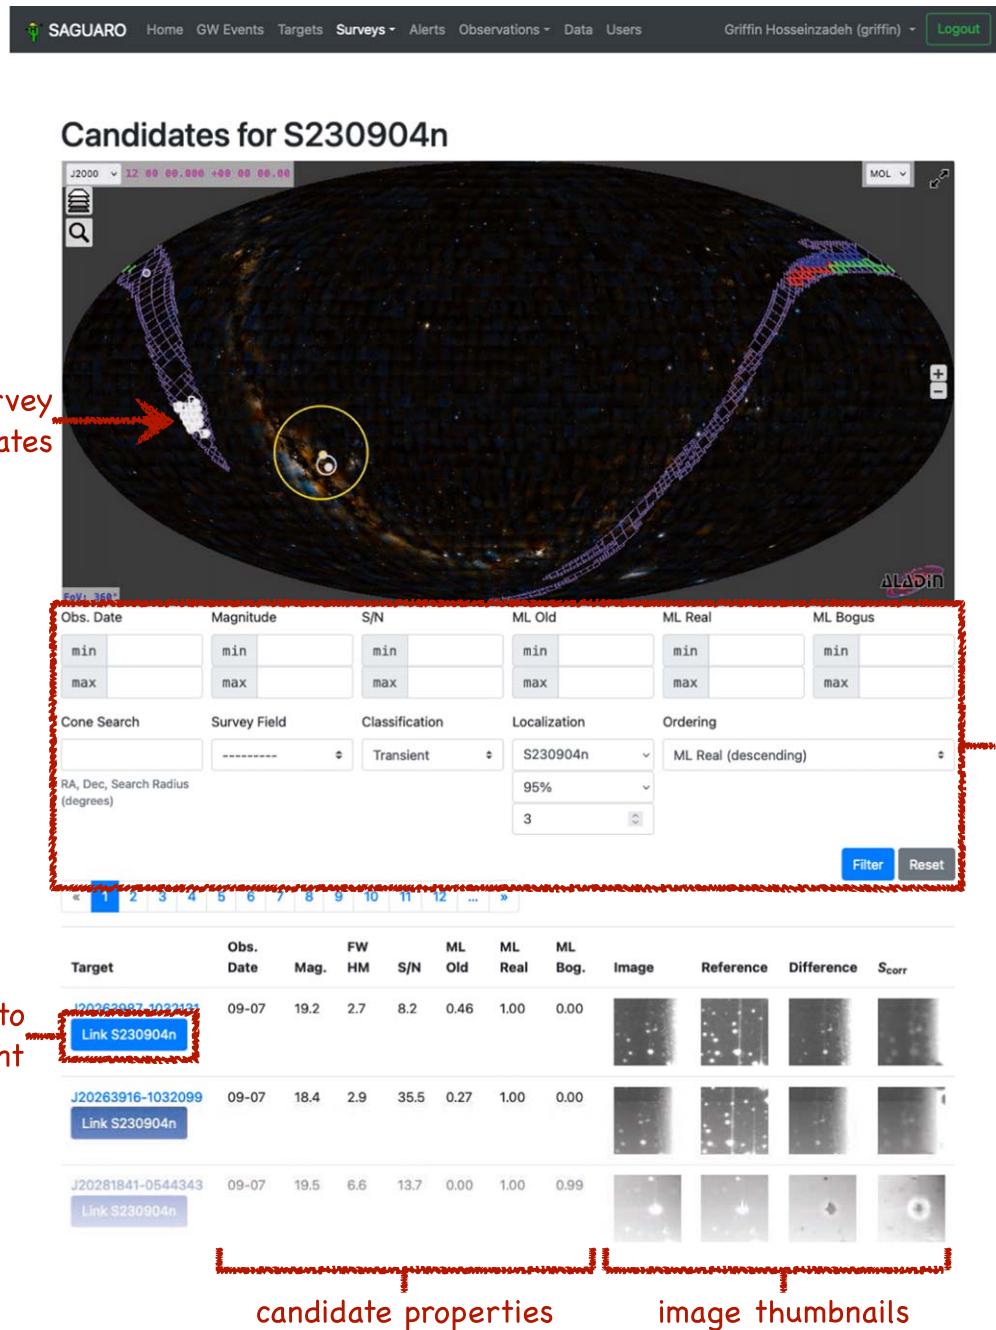


Figure 5. Survey candidates page for S230904n in the SAGUARO TOM. As on the survey planning page in Figure 4, the Aladin sky map shows the 90% localization region (purple) and current Sun (gold) and moon (silver) position. Observed fields and detections are shown in white. Below, the survey candidates table shows a temporary coordinate-based name, the observation month and day, magnitude, full width at half maximum (FWHM), signal-to-noise ratio (S/N), three real-bogus machine-learning (ML) scores, and thumbnails of the image quadruplet. The blue “Link” buttons allow the user to bookmark promising sources for further vetting in association with the chosen GW event.

catalog provides a point-source probability score for nearly all sources in PS1 using a random-forest machine-learning (ML) algorithm (Breiman 2001). Following Tachibana & Miller (2018), we utilize an ML score cutoff of 0.83 to determine if matches are point sources and thus unrelated to the GW merger event.

2.4.2. Host-galaxy Matching

We incorporate five individual galaxy catalogs to further contextualize each target and determine if its distance and luminosity are consistent with expectations of a kilonova counterpart. In Rastinejad et al. (2022b) we demonstrated that host-galaxy

information was powerful in determining if targets were associated with GW events, as it ruled out $>30\%$ of potential O3 kilonovae without the need for any follow-up observations. We begin by querying each catalog for a match within 3°44 of the target. This radius translates to a 100 kpc offset at a distance of 100 Mpc, consistent with the upper end of short GRB offsets (though the median offset lies ~ 8 kpc; e.g., Fong et al. 2022). To determine the likelihood the target is associated with each source, we utilize the probability of chance coincidence method (P_{cc} ; Bloom et al. 2002). P_{cc} combines the galaxy’s offset and magnitude, and is used to associate GRBs (including NS merger-origin GRBs) to host

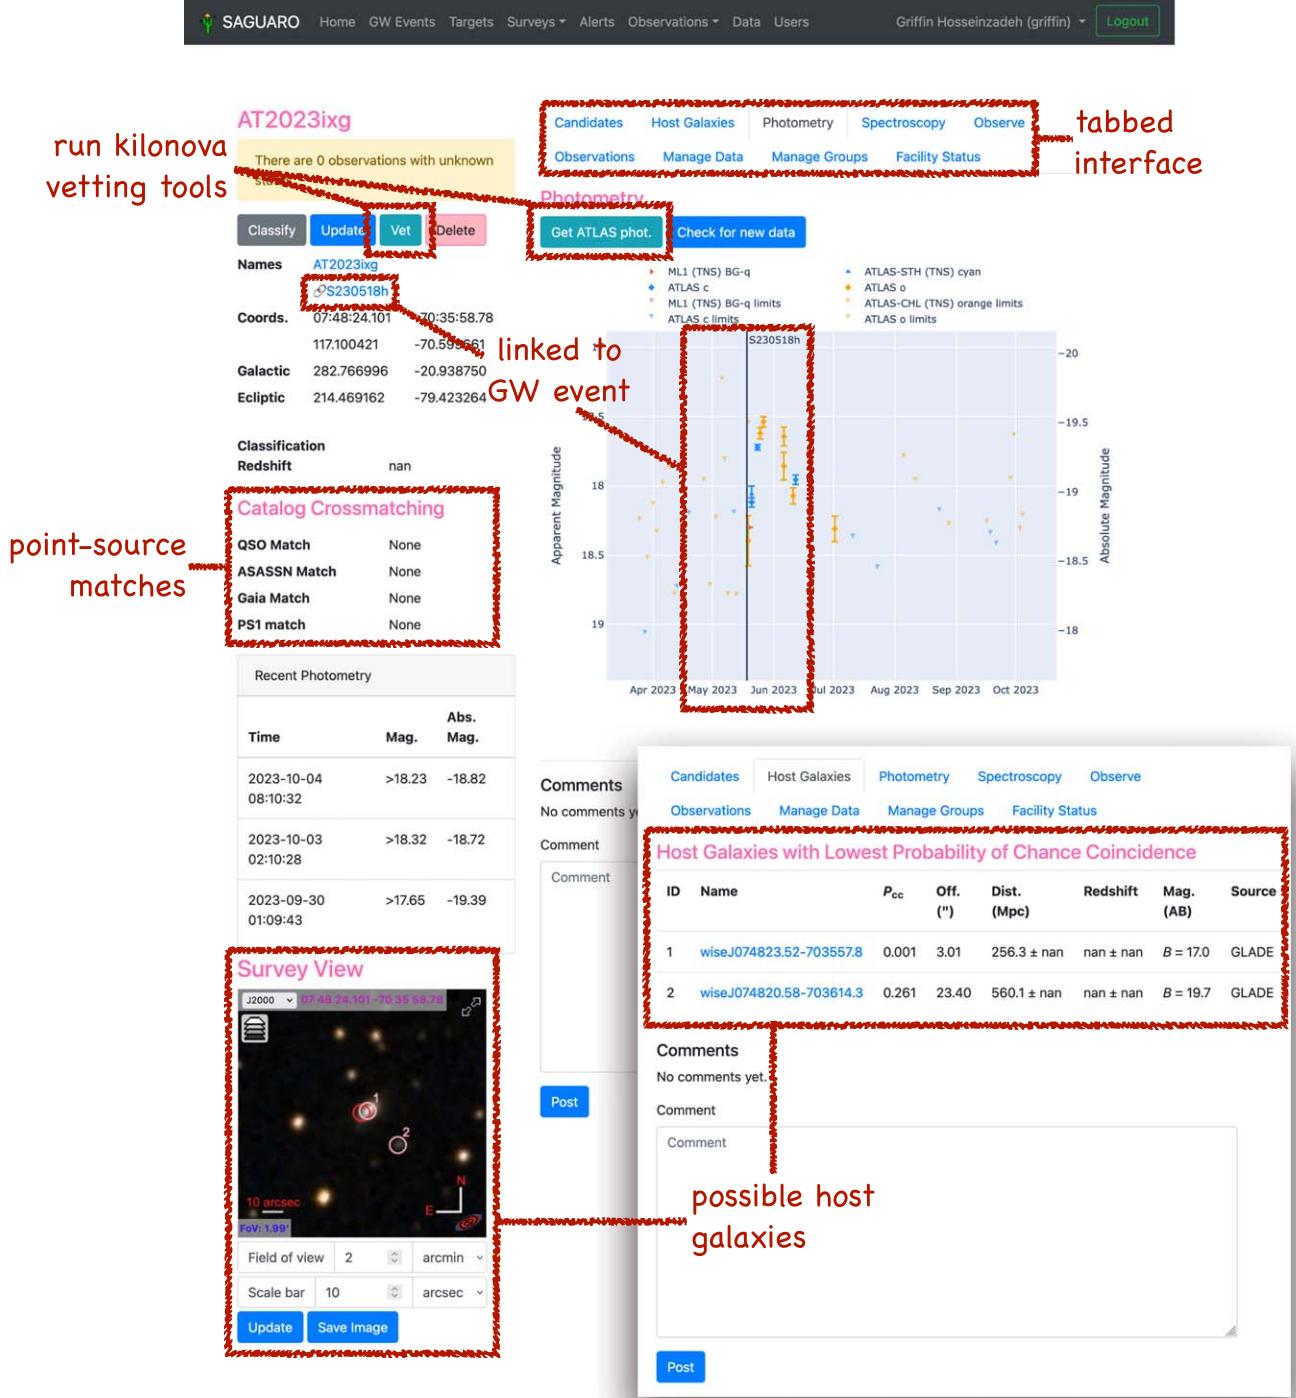


Figure 6. Two tabs of the target detail page for AT 2023ixg, which was suggested as a possible EM counterpart to GW event S230518h by Johnston et al. (2023). The teal “Vet” button near the top left of the page runs our kilonova vetting code. The results are displayed in the “Catalog Crossmatching” section at the center left of the page, which indicates no point-source matches but several galaxy matches in this case. The “Photometry” tab shows a plot of all known photometric detections and upper limits for this target, including public data from the ATLAS forced photometry server (obtained by pressing the teal “Get ATLAS phot.” button), the Transient Name Server, and the Zwicky Transient Facility. Because this target has been linked with GW candidate S230518h (note the chain icon under the target names), the time of merger of that event is indicated by a vertical line on the photometry plot. The “Host Galaxies” tab (inset) shows a table of potential host galaxies, ordered by increasing probability of chance coincidence (P_{cc}). These are also marked on the Aladin viewer at the bottom left of the page.

galaxies (e.g., Bloom et al. 2002). Where available, we calculate P_{cc} using the galaxy’s r -band magnitude, as it is generally the most widely available filter, though in cases where this filter is not provided we employ a band with similar wavelength coverage. We display the 10 galaxies with the lowest P_{cc} values and $P_{cc} < 0.8$ on

each target’s page and circle them in the Aladin viewer. For each potential host galaxy, we display its name, P_{cc} , offset from the target, distance or redshift (depending on what the catalog provides), magnitude, and catalog source (see Figure 6, bottom right).

As we expect counterparts to an O4 GW event to originate in a low-redshift ($z \lesssim 0.05$) galaxy, we employ three galaxy catalogs specifically designed for the local Universe. Using the method described above, we query the Gravitational Wave Galaxy Catalog (GWGC; White et al. 2011), the Galaxy List for the Advanced Detector Era (GLADE+; Dálya et al. 2018, 2022), and the Heraklion Extragalactic Catalog (HECATE; Kovlakas et al. 2021) for galaxies with known distances. Using B -band luminosity density, the GLADE+ catalog is complete out to $d_L = 47^{+4}_{-2}$ Mpc (Dálya et al. 2022) while HECATE is $>50\%$ complete up to $d_L = 170$ Mpc (Kovlakas et al. 2021). In addition, to identify contaminants that may be associated with higher-redshift galaxies, we use photometric redshifts from the Sloan Digital Sky Survey Data Release 12 (SDSS DR12; Alam et al. 2015) and the Pan-STARRS Source Types and Redshifts with Machine Learning (PS1-STRM; Beck et al. 2021) catalog. SDSS DR12 also provides a number of spectroscopic redshifts that we incorporate when available.

2.4.3. Survey Photometry Matching

Finally, we search for and display any additional survey photometry for each target. This survey data is useful for determining if the transient was detected prior to the merger (ruling it out as a counterpart) and provides information on the color and temporal evolution. We again employ a search radius of $2''$ to match the typical seeing at ZTF’s Palomar Observatory. We automatically ingest photometry from the public ZTF alert stream (Bellm et al. 2019; Masci et al. 2019; Patterson et al. 2019) and from the Transient Name Server (TNS; Gal-Yam 2021). We also gather ATLAS data using their forced photometry service (Tonry et al. 2018a; Smith et al. 2020; Shingles et al. 2021), and stack and sigma-clip observations from the same night using the script they provide.²⁵ Each of these streams provides photometry from difference images.

2.5. Coordinating Follow-up

Once the target has been vetted, rapid photometric and spectroscopic follow-up is crucial to confirm the association of the EM counterpart and gather early data for later analysis of the kilonova. The TOM Toolkit already includes the infrastructure to submit observations to all facilities in the Astrophysical Events Observatories Network (AEON; Street et al. 2020). Currently, this includes optical and near-infrared imagers and spectrographs on Las Cumbres Observatory (Brown et al. 2013), SOAR (Clemens et al. 2004; Schlawin et al. 2014), and Gemini (Hook et al. 2004; Elias et al. 2006a, 2006b; Eikenberry et al. 2006), all of which have active GW follow-up programs in which we participate (see Section 1).

In our TOM, we have added the ability to submit observations to two imaging spectrographs on the MMT: Binospec in the optical (Fabricant et al. 2019) and the MMT and Magellan Infrared Spectrograph in the near-infrared (MMIRS; McLeod et al. 2012). The `tom_mmt` app²⁶ is an “observatory module” for the TOM Toolkit based on PyMMT (Wyatt et al. 2023; Shrestha et al. 2024), a recently released

Python package that communicates with the application programming interface (API) of the MMT Observatory Manager (Gibson & Porter 2018). The `tom_mmt` module provides observation request forms (Figure 7, bottom), which are prefilled with default settings based on observations of AT 2017gfo (Coulter et al. 2017) and theoretical models of kilonovae (e.g., Kasen et al. 2017), along with the most recent magnitude of the target in question. The module also adds the MMT to the TOM’s airmass charts (Figure 7, top) and “Facility Status” page, allowing the user to quickly determine which instrument is available for the upcoming night. Last, it allows the user to check the status of requested observations without leaving the TOM.

2.6. Visualizing Data

Regardless of whether they were requested through the TOM or if they were obtained by us or others, we have the ability to upload reduced data products for a given target. They will then be displayed on interactive light-curve and spectral plots on the target detail page, produced using Plotly’s Python graphing library.²⁷ These allow us to keep track of the evolution of targets in real time and plan or cancel follow-up accordingly. We have tested this functionality by using it to manually aggregate GRB data from GCN circulars (e.g., Figure 8). However, we strongly encourage other observers to switch to machine-readable communications as soon as possible, so that they can be ingested into TOMs automatically (see Section 2.7.3). In the future, we hope to include preliminary analysis tools, such as spectral-line plotting and comparison to light-curve or spectral templates, directly in the TOM.

2.7. Reporting Results

2.7.1. Treasure Map

The Gravitational Wave Treasure Map is a tool to coordinate, visualize, and assess the EM follow-up of GW events (Wyatt et al. 2020). We report CSS pointings within the 95% localization region of any “active” GW events to the Treasure Map. We consider a GW event to be active during a given pointing if

1. the event is astrophysically significant,
2. the event has not been retracted, and
3. the pointing occurred between 0 and 3 days after the most recent estimate of the merger time.

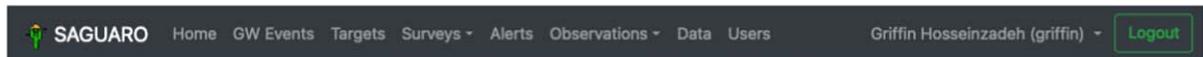
Since the beginning of 2023 September, when we restarted our reporting, we have sent 287 reports of 283 serendipitous pointings for 12 BBH mergers to the Treasure Map (some observations apply to multiple events), corresponding to 1435 deg^2 of localization-region coverage.

To accomplish this, we have developed the `tom_treasuremap` app, which, when used alongside the `tom_non-localizedevents` (Section 2.1) and `tom_surveys` (Section 2.2) apps, provides a `TreasureMapPointing` model connecting a `NonLocalizedEvent` to a `SurveyObservationRecord` and labeling them with a Treasure Map pointing ID number and a status (planned or completed). When we submit observation requests to CSS, we automatically send these to the Treasure Map as planned pointings and store the pointing IDs in our database. The

²⁵ <https://gist.github.com/thespacedoctor/86777fa5a9567b7939e8d84fd8cf6a76>

²⁶ https://github.com/griffin-h/tom_mmt

²⁷ <https://plotly.com/python/>



Submit an Observation to MMT

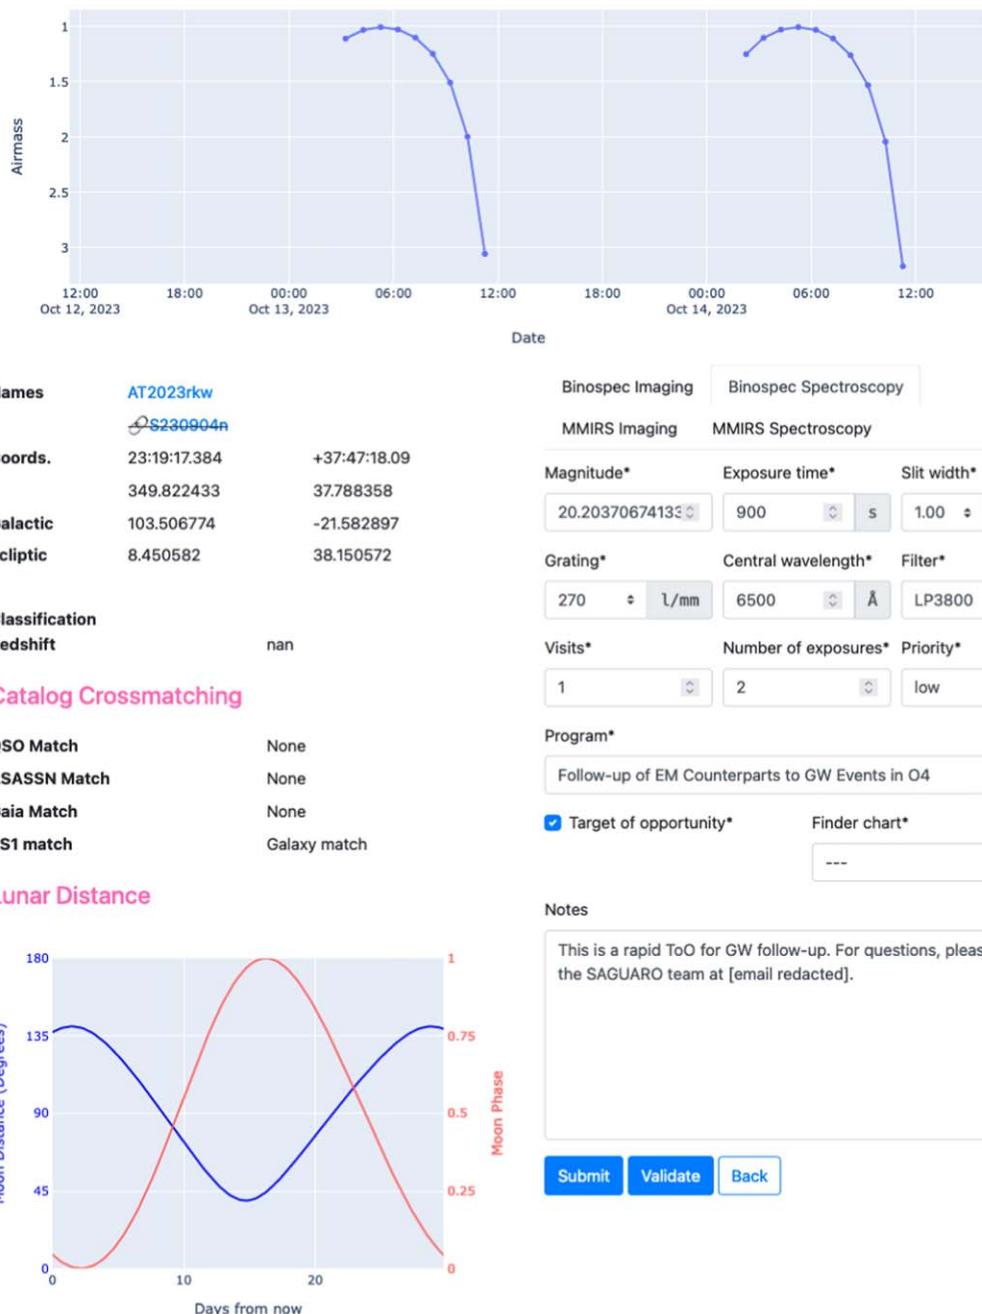


Figure 7. MMT observation request page for AT 2023rk on the SAGUARO TOM, provided by our `tom_mmt` module. AT 2023rk was suggested as a possible EM counterpart to S230904n by Necker et al. (2023) but then classified as a Type Ia supernova by Anand et al. (2023). Therefore, the link to the GW event at the left of the page is struck through. An airmass chart (top) and lunar distance chart (bottom left) show the visibility of the target. The form at the bottom right of the page allows the user to specify the instrument settings, choose a program, upload a finder chart, and provide free-text notes. The tabbed interface provides four forms, for imaging and spectroscopy on Binospec and MMIRS.

app also provides the management command `report_pointings`, which queries the database for completed observations (requested or serendipitous) in the localization

regions of any active GW events and reports them to the Treasure Map. We run this command hourly via a cronjob on the TOM server.

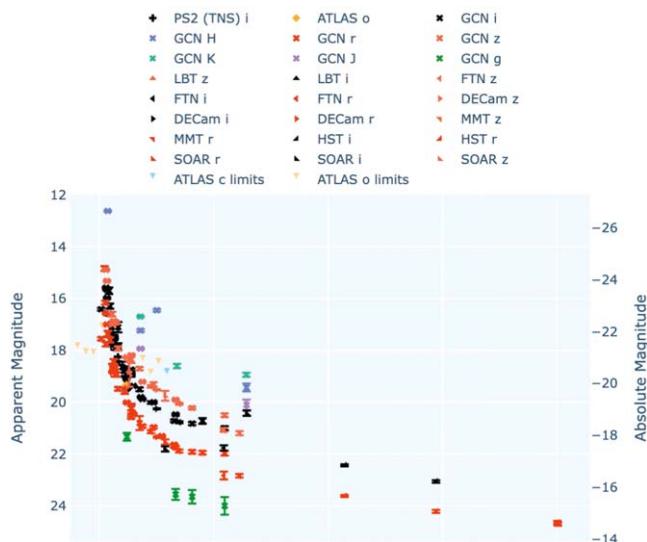


Figure 8. Light curve of GRB221009A as analyzed by Shrestha et al. (2023). These data were aggregated from GCN circulars and uploaded to the TOM in real time. An interactive version of this figure, extracted from the SAGUARO TOM, is available. The plot area can be zoomed and panned, and specific data series can be turned on or off by clicking in the legend.

2.7.2. Transient Name Server

The IAU Transient Name Server (TNS) is the database of record for the discovery and classification of astronomical transients (Gal-Yam 2021). When we discover real transients in the course of our GW follow-up survey, whether or not they are real EM counterparts to GW events, we wish to report their coordinates and magnitude to the TNS and receive a permanent astronomical transient (AT) name for them. To facilitate this, we have built into our TOM a web form (Figure 9, top), accessed by clicking the gray “Report” button on the target detail page, that sends a discovery report to the TNS API without leaving the TOM. If the transient was discovered by our pipeline, the form will automatically be filled out with all the required information. Users can review the pre-filled form and add authors or comments before submitting. After submission, the target will have been renamed with a permanent AT name in the TOM.

Likewise, if we obtain a classification spectrum of a transient and upload it to the TOM, we will want to report the classification to the TNS. If a target name starts with “AT,” a gray “Classify” button will replace the “Report” button on the target detail page (e.g., Figure 6). Clicking this button leads to a classification report form, which is also pre-filled to the extent possible (Figure 9, bottom). Submitting the form will send the report and upload its associated files to the TNS. After submission, the target’s classification and redshift, if provided, will appear on the target detail page. If the target’s prefix is changed (e.g., from “AT” to “SN”) as a result of the classification, its name will also be updated.

2.7.3. Hermes

In addition to programmatic reporting of observations to these two databases, the GW follow-up community has traditionally relied on free-text GCN circulars to communicate information about potential EM counterparts, including their discovery, their likelihood of association with particular GW events, and the results of any vetting or follow-up observations.

SCIMMA’s Hermes service²⁸ is designed to supplement GCN circulars by providing several semistructured message schemata for communicating targets and data in a machine-readable format, alongside free text. Messages sent through Hermes can also be reduced to plain text and sent as GCNs at the user’s request. We plan to integrate Hermes messaging into the SAGUARO TOM in a similar method to our TNS reporting—providing a form prefilled with details of targets and data from the TOM that can be quickly reviewed and sent to the Hermes API. Eventually, this will be our preferred method of composing GCN circulars, but as Hermes has been under active development during the same time period as our TOM, we have not made this a priority for early O4. Until this is integrated into the TOM, we plan to send any GCN circulars manually through the Hermes web interface.

3. Image-subtraction Pipeline

3.1. Deeper Reference Images

With the accumulation of more than 4 yr of data since the creation of templates for SAGUARO in 2019, we look to create new templates for image subtraction using the additional data for use during O4. Previously, we had access to a total of 5385 fields, arrayed in the sky accessible in southern Arizona, excluding areas near the Galactic plane (see Figure 1 of Lundquist et al. 2019) with reference templates. We have new data for 4916 of these fields, allowing us to build deeper templates for 91% of our previous fields, plus additional data for 25 new fields, expanding the number of templates. We build these new templates using the same criteria as our prior work, requiring the number of stars detected in an image to be greater than 2000, and the sky brightness lower than 20 mag (except in the cases where the total number of images is less than four, for which we implement no cuts). Compared to the previous sample of image templates that were created with a median of 60 individual images (see Lundquist et al. 2019 for details), the new sample of image templates (which includes the deeper template for fields that had additional data, the old templates for the fields that were not observed again since 2019, and the new additional fields) were created with a median of 149 images. On the low end, those with fewer than five images make up less than 1% (compared to 9% previously) and 74% of templates were made with >90 images (compared to the previous 5%). A histogram of the number of individual images that make up both the old and new templates is shown in Figure 10. We expect these deeper reference images to result in fewer image-subtraction artifacts and allow the recovery of fainter transients. We leave a full analysis of their effect for future work, after they have been in use for a longer period.

3.2. Machine Learning

Previously, SAGUARO made use of a random-forest ML algorithm (Breiman 2001), as implemented in `scikit-learn` (Pedregosa et al. 2011), to classify candidates from the image subtraction as real (astrophysical transients, variable stars, or slow-moving objects) or bogus (artifacts caused by bad subtraction, bright stars, or detector effects). The random forest was trained on 10×10 pixels taken from the difference images centered on the source (equal to a $15.^{\circ}4 \times 15.^{\circ}4$ box; see Paterson et al. (2021) for details) and was able to recover 94%

²⁸ <https://hermes.lco.global/>


SAGUARO
[Home](#)
[GW Events](#)
[Targets](#)
[Surveys](#)
[Alerts](#)
[Observations](#)
[Data](#)
[Users](#)

Griffin Hosseinzadeh (griffin) [Logout](#)

Report J19044308+4447388 to the TNS

Reporter*

Reporting group*

Discovery data source*

R.A.*

Dec.*

Discovery date*

AT type*

Discovery remarks

Discovery Photometry

Observation date*	Instrument*	Limiting flux	Observer
<input type="text" value="2023-09-21 05:43:05"/>	<input type="text" value="Other"/>	<input type="text"/>	<input type="text"/>
Flux*	Flux error*	Flux units*	Filter*
<input type="text" value="18.62343406677246"/>	<input type="text" value="0.2121419757604599"/>	<input type="text" value="ABMag"/>	<input type="text" value="Clear"/>

Photometry remarks

Last Nondetection

Archive*	Archival remarks*	Nondetection remarks
<input type="text" value="Other"/>	<input type="text" value="CSS"/>	<input type="text"/>


SAGUARO
[Home](#)
[GW Events](#)
[Targets](#)
[Surveys](#)
[Alerts](#)
[Observations](#)
[Data](#)
[Users](#)

Griffin Hosseinzadeh (griffin) [Logout](#)

Classify AT2022oxk on the TNS

Classifier*

Group*

Name*

Classification*

Redshift

Classification remarks

Classification Spectrum

Observation date*	Observer*	Reducer
<input type="text" value="2019-09-05 00:00:00"/>	<input type="text"/>	<input type="text"/>
Instrument*	Exposure time	Spectrum type*
<input type="text" value="Other"/>	<input type="text"/>	<input type="text" value="Object"/>

ASCII file*

FITS file

Spectrum remarks

Figure 9. Forms for sending transient discovery (top) and classification (bottom) reports to the TNS from within the SAGUARO TOM. These forms are prefilled with all relevant information that is available within the TOM, including photometric and spectroscopic data. This allows the user to quickly review and submit the report with little to no typing.

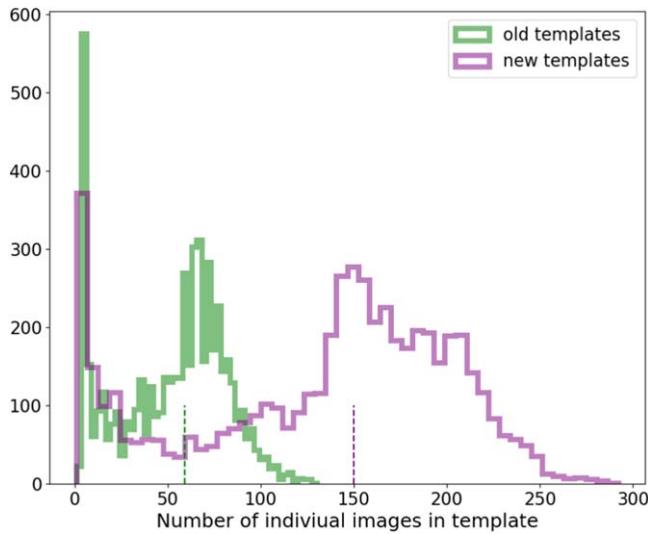


Figure 10. Histograms of the number of individual images that were used to create the reference templates for the old templates vs. the newly rebuilt templates. Dashed lines indicate the median values.

of the real transients used for validation (where the recovery threshold was set to 0.7). In an attempt to improve the real–bogus score for O4, we looked at enhancing the method used for the ML. First we made use of the Zooniverse portal²⁹ to classify good images for the transients and moving objects (through a private project). We found a sample of 383 and 2480 images for each, respectively. To increase the number of these images in our training set, we included each image rotated by 90°, 180°, and 270°, as well as the relative inversions through the x - and y -axes, which also ensures the algorithm is not sensitive to the orientation of the candidate. We then extracted a subset of images flagged as variable stars from our previous observations to add to the transient and moving-object sample of real detections, as well as a subset of images flagged as bad for the bogus sample in our training set.

Convolutional neural networks (CNNs; Fukushima 1980) are often used for image classification problems and have been used extensively in astronomy for object classification, including the real–bogus classification for transients (e.g., Cabrera-Vives et al. 2017; Gieseke et al. 2017; Turpin et al. 2020; Goode et al. 2022; Acero-Cuellar et al. 2023; Chen et al. 2023; Liu et al. 2023). Using TensorFlow (Abadi et al. 2016a, 2016b), we set up a basic CNN consisting of multiple 2D convolution layers, maxpooling, dropouts, and Dense layers (all with the relu activation), and finally ending with a Dense layer with a sigmoid activation. For the training data, we tried a combination of arrays extracted from the new, reference, difference, and corrected significance (S_{corr}) images derived from the image subtraction. We tried training on the full 64×64 pixel ($99'' \times 99''$) thumbnails created from the image-subtraction pipeline, as well as on the inner 16×16 pixels ($25'' \times 25''$) centered on the source. Using the CNN, we found using the arrays just from the S_{corr} images provided the highest accuracy, while including the other arrays often resulted in a drop of accuracy due to overfitting. While the initial results were similar for the 64×64 and 16×16 pixel cutouts, further testing found that using the full 64×64 pixel array increased the false positive rate for images where bright stars were

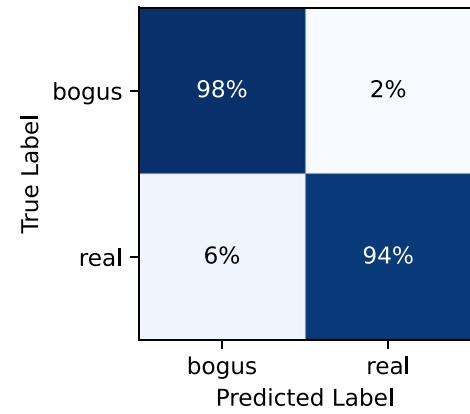


Figure 11. Confusion matrix for the new CNN ML algorithm, based on a highest probability classification; 98% of bogus and 94% of real detections are classified correctly.

present on the edge of the cutout, while the 16×16 pixel array avoided this. We thus decided to use 16×16 pixels from the S_{corr} image as the final training data. We then looked to optimize the hyperparameters of the CNN model using this data. The final CNN model comprises two blocks of two 2D convolution layers of increasing filter size (16 for the first block and 32 for the second), followed by a 2×2 maxpooling, a 0.25 dropout, two 2D convolution layers with a filter size of 64, a reshaping filter to flatten the array, and two Dense layers with values of 64 and 32. All of these layers are applied using the relu activation. A final Dense layer with a value of 2 and a sigmoid activation is then applied, with a final dropout of 0.5.

The final CNN returns an array of two values containing the probability (from 0 to 1) of the candidate being real and the probability of the candidate being bogus. While the classification of real–bogus is set by whichever class has the highest probability, looking at both values can provide additional insight into candidates that the ML had trouble classifying. Using the highest probability value to classify the candidates, we correctly recover 94% of real objects and label 98% of the bad data as bogus (see Figure 11 for the confusion matrix).

3.3. Candidate Ingestion

Although the postprocessing procedure has not changed significantly compared to the steps described by Lundquist et al. (2019) and Paterson et al. (2021), we have upgraded this portion of the pipeline to increase its speed and decrease its memory usage. We typically run the pipeline in a multithreaded mode so that we can process up to 16 images at once. In refactoring the pipeline, we have moved the reading of large files (i.e., the trained random-forest classifier and the moving-object catalog) outside of the threads, so that they are only accessed once at the beginning of the night. The trained CNN classifier is reloaded for each image, due to limitations of TensorFlow in the Python multiprocessing environment. We have also vectorized the postprocessing operations to the greatest extent possible to avoid overheads associated with Python loops. With the exception of thumbnail creation, which requires writing to the hard disk, the duration of these steps is now nearly independent of the length of the source catalog. We describe each step below.

After extracting sources from the difference images using Source Extractor (Bertin & Arnouts 1996), the pipeline crossmatches the source catalog with 1.3 million solar system

²⁹ <https://www.zooniverse.org/>

objects known to the Minor Planet Center. We download an up-to-date catalog of orbital elements at the beginning of each night and use PyEphem (Rhodes 2011) to calculate the equatorial coordinates of each body at the time of each image. Then we use the Astropy coordinates module (Astropy Collaboration 2022) to crossmatch the source catalog with the solar system catalog, using a matching radius of $25''$ to account for uncertainties in the orbits. In total, this step takes about 50 s per image, almost all of which is spent calculating the ephemerides. Any sources with matches in the solar system catalog are marked as moving objects in our database and not displayed by default on the survey candidates page.

We then run both our old random-forest and our new CNN real–bogus classifiers, each of which takes $\lesssim 1$ s per source catalog. Finally, for each source in the catalog, we produce thumbnails of the image quadruplet (see Section 2.3), which are saved on the data-reduction server and hosted with a lightweight Flask³⁰ web server so that they can be displayed in the TOM. Saving the thumbnails takes a few milliseconds per source, for a total of 1–50 s per image, depending on the length of the catalog.

The pipeline records our observations in four places in the TOM database:

1. the observation record table, which stores the field name, time of observation, and other metadata for each image,
2. the targets table, which keep tracks of distinct pairs of coordinates using a $2''$ matching radius,
3. a custom candidates table, which logs each detection of a point source along with its moving-object match and real–bogus scores, and
4. the reduced data table, which records a photometry point for each detection.

The new source catalog is crossmatched with the existing targets table in PostgreSQL using the Q3C extension (Koposov & Bartunov 2006) and insertions are done in bulk for each source catalog, allowing us to complete this step in a few seconds per catalog.

In total, the ingestion of candidates typically takes 1–2 minutes per image. We have also begun logging the time at which candidates are ingested into the TOM database. We find that, absent technical or network issues, the typical delay from shutter open of the first exposure to the appearance of candidates on the web page is $\lesssim 30$ minutes (Figure 12).

4. Future Plans

4.1. Bok 90Prime Upgrades

While the workhorse CSS 1.5 m telescope is a very efficient survey machine, its typical limiting magnitude (~ 21.5 mag) is not much deeper than the expected brightness of kilonovae (~ 20.5 mag) at LIGO’s distance horizon for BNS mergers in O4 (~ 200 Mpc). Therefore, we are in the process of adding to our program a second wide-field imager on a larger telescope: 90Prime, the prime-focus imager on the 2.3 m (90 in) Bok Telescope on Kitt Peak, Arizona (Williams et al. 2004). 90Prime has a ~ 1 deg 2 field of view and has been used for deep survey programs like the Beijing–Arizona Sky Survey (BASS; Zou et al. 2017). Equipped with standard SDSS *ugriz* filters, it also has the advantage of providing color information of

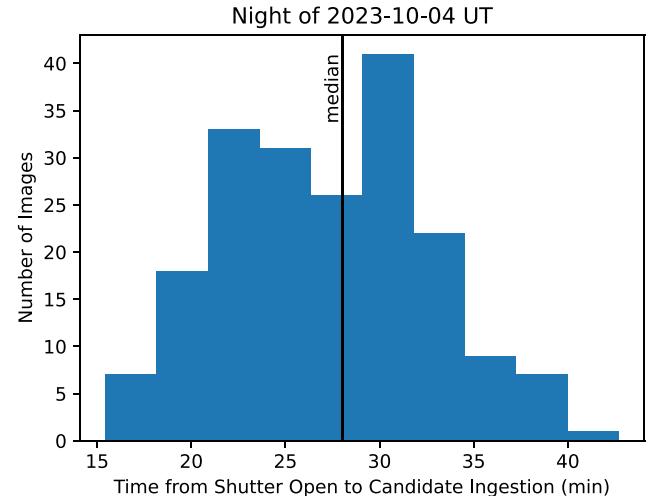


Figure 12. Histogram of the delay time between shutter open of the first of four exposures and the appearance of survey candidates in the TOM in our updated pipeline. The median delay is under 30 minutes.

transients. Over the past several years, 90Prime has undergone an upgrade, replacing its CCDs with higher-quality, blue-sensitive chips, along with its liquid nitrogen dewar, shutter, and guider camera. Along with these hardware upgrades come upgrades to the instrument control software, as well as a new pipeline for image preprocessing and data-quality monitoring, which we discuss in Section 4.2. The upgraded 90Prime achieved engineering first light in Spring of 2023 but is now undergoing further electronics work before being released to observers. D. Sand et al. (2024, in preparation) will present a complete description of the 90Prime upgrade in a future work.

4.2. Image Processing

CSS provides SAGUARO with images that have already been bias-subtracted, flat-fielded, astrometrically aligned, and photometrically calibrated. We need the equivalent preprocessing for 90Prime and any other instruments we wish to add to our survey. For this, we have adopted a customized version of Beautiful Algorithms for Normalizing Zillions of Astronomical Images (BANZAI; McCully et al. 2018b), the preprocessing pipeline developed and used by Las Cumbres Observatory. BANZAI is designed to be flexible so that it can process data from Las Cumbres’ several different imagers, and our usage is a proof of concept that it can also be applied to extramural instruments with minimal modification (described below).

In addition to installing the pipeline itself, deploying BANZAI on Kitt Peak involves setting up three Docker containers (Merkel 2014) containing a PostgreSQL database to log observations and processed data, an astrometry service based on the Astrometry.net framework (Lang et al. 2010) with index files generated from Gaia Data Release 2 (Gaia Collaboration 2018; Lindegren et al. 2018), and a photometric calibration service based on the ATLAS All-sky Stellar Reference Catalog (Tonry et al. 2018b). The main customization to the pipeline for 90Prime relates to the fact that the 90Prime focal plane consists of four chips, each with four amplifiers, producing a 16-extension FITS image, whereas all Las Cumbres imagers contain a single four-amplifier chip. Our solution involves splitting each image into four four-extension FITS files and passing them to BANZAI to be processed in parallel. Relatedly, 90Prime’s depth and large field of view can

³⁰ <https://flask.palletsprojects.com/>

produce very crowded fields, which were not always a good match for the Astrometry.net configuration included with BANZAI by default. We will refine these settings further as we test the pipeline on a larger quantity of images.

4.3. Survey Coordination

Once the upgraded 90Prime is regularly available, we will begin to integrate it into our survey strategy. This will involve reoptimizing our field selection algorithm when 90Prime is available (usually for 1–3 weeks per lunation) to simultaneously plan observations on two telescopes. As discussed in Section 2.2, we have kept the `tom_surveys` app as general as possible with this future in mind.

Likewise, we will have to adapt our image-subtraction pipeline to run on images from a new instrument, possibly using reference images from BASS. Our initial plan is to run a second instance of the pipeline on the BANZAI server on Kitt Peak to avoid the bottleneck of transferring data off the mountain. In that case, the pipeline will remotely ingest survey candidates into our TOM database and transfer only the thumbnail images off the mountain. Because of the database schema in which “targets” (coordinates) are tracked separately from “candidates” (detections; see Section 3.3), multiple instances of the pipeline can detect the same transient, and these will automatically be associated in the TOM if they are within 2[”] of each other.

4.4. Forced Photometry

Rastinejad et al. (2022b) and Section 2.4 describe the utility of public photometry sources (e.g., ATLAS, TNS, ZTF) for vetting potential kilonovae. After several years of processing CSS data as part of SAGUARO, we ourselves have a large archive of images that we can use to rule out associations with GW events through premerger detections. We plan to set up a forced photometry service for SAGUARO–CSS images and integrate this into the TOM as part of our standard kilonova vetting procedure.

4.5. Additional Follow-up Facilities

In the longer term, we hope to add functionality to the TOM to trigger additional facilities accessible via API. First, we plan to add the remainder of the queue-scheduled instruments on the MMT (Hectospec and Hectochelle; Fabricant et al. 2005; Szentgyorgyi et al. 2011) to the `tom_mmt` app. There is also ongoing work to build a TOM Toolkit facility module around the target-of-opportunity API³¹ for the Neil Gehrels Swift Observatory (Gehrels et al. 2004), which would be particularly useful for constraining kilonova models via very early UV photometry (Arcavi 2018). The 4 m Víctor M. Blanco Telescope at CTIO, Chile, may also become part of AEON, allowing it to be scheduled through Las Cumbres Observatory’s observing portal, providing a critical publicly accessible survey instrument in the southern hemisphere.

5. Summary and Outlook

We have presented our completed and ongoing upgrades to SAGUARO infrastructure for optical follow-up of GW events. These include the development of a TOM that lets users receive

alerts, plan and trigger a survey, scan for new survey discoveries, vet potential EM counterparts, coordinate follow-up observations, visualize reduced data, and report results. We have also described improvements to our image-subtraction pipeline, including deeper reference images, a new CNN real-bogus classifier, and streamlined candidate ingestion. Last, we described upgrades to the 90Prime wide-field imager on the 2.3 m Bok Telescope, including improvements to camera hardware, computing hardware, and software, and outlined our plans to integrate it into our search when upgrades are completed. Together, these upgrades place SAGUARO in a good position to discover and characterize EM counterparts to any BNS or NSBH mergers in the remainder of O4 and beyond, if they are visible from Arizona.

The infrastructure we have developed here has applications far beyond GW follow-up. During our design and testing, we have used the SAGUARO TOM to observe young supernovae (Shrestha et al. 2024) and GRB afterglows (Shrestha et al. 2023), and we plan to add astrophysical neutrino localizations (e.g., IceCube Collaboration et al. 2017) soon. As with the other follow-up systems we mention in Section 2, the SAGUARO TOM will grow into a centralized portal for all transient science in our research collaboration. Likewise, we plan to extend our implementation of BANZAI to run on smaller Arizona telescopes without dedicated image processing pipelines.

We also view the SAGUARO TOM as a case study in the type of software infrastructure that will be necessary for time-domain science in the era of Vera C. Rubin Observatory (Ivezić et al. 2019), when millions of transients will be discovered each year. While brokers like Alert Management, Photometry, and Evaluation of Light Curves (Nordin et al. 2019), Arizona-NOIRLab Temporal Analysis and Response to Events System (Matheson et al. 2021), Automatic Learning for the Rapid Classification of Events (Förster et al. 2021), Fink (Möller et al. 2021), and Lasair (Smith et al. 2019) will handle alert distribution, classification, and filtering, TOMs will play a critical role in managing the follow-up observations and data gathered in response. Thoughtful architecture of these systems, in coordination with other time-domain researchers, will help us maximize the scientific return from both large and small observatories.

Acknowledgments

We thank the organizers and attendees of the GW-EM Workshop at the Schwartz/Reismann Institute for Theoretical Physics, the Gravitational Wave Physics and Astronomy Workshop 2022, and the Windows on the Universe Workshop (The 2023 Windows on the Universe Workshop White Paper Working Group 2024) for helpful discussions about GW alerts and coordination of EM follow-up.

Time-domain research by the University of Arizona team, M.S., and D.J.S. is supported by NSF grants AST-1821987, 1813466, 1908972, 2108032, and 2308181, and by the Heising-Simons Foundation under grant #2020-1864. The Fong Group at Northwestern acknowledges support by the National Science Foundation under grant Nos. AST-1909358, AST-2308182, and CAREER grant No. AST-2047919. K.A.B. is supported by an LSSTC Catalyst Fellowship; this publication was thus made possible through the support of grant 62192 from the John Templeton Foundation to LSSTC. The opinions expressed in this publication are those of the authors and do not

³¹ https://www.swift.psu.edu/too_api/

necessarily reflect the views of LSSTC or the John Templeton Foundation. W.F. gratefully acknowledges support by the David and Lucile Packard Foundation, the Alfred P. Sloan Foundation, and the Research Corporation for Science Advancement through Cottrell Scholar Award #28284. The TOM Toolkit is managed by Las Cumbres Observatory, with generous financial support from the Heising-Simons Foundation, the Zegar Family Foundation, and the National Science Foundation under grant No. 2209852.

Facility: ADS, Bok, GWTM, SO:1.5m, SO:Kuiper, TNS.

Software: Astrometry.net (Lang et al. 2010), astroplan (Morris et al. 2018), Astropy (Astropy Collaboration 2013, 2018, 2022), astropy-healpix (Astropy Collaboration 2022), Astroquery (Ginsburg et al. 2019), Astro-SCRAPPY (McCullly et al. 2018a), BANZAI (McCullly et al. 2018b), Beautiful Soup (Richardson 2023), crispy-bootstrap4 (Smith 2022), dateutil (Niemeyer et al. 2021), Django (Django Software Foundation 2023), django-bootstrap4 (Verheul 2021), Django “excontrib” Comments (Django Software Foundation 2022), django-crispy-forms (Araujo 2023), Django Extensions (Trier & van Oostveen 2023), Django Filter (Gibson 2021), django-gravatar (Waddington 2020), django-guardian (Balcerzak 2021), django-phonenumbers-field (Foulis & Freitag 2023), Django REST Framework (Christie 2022), django-webpack-loader (Lone 2022), Docker (Merkel 2014), fastavro (Tebeka 2022), fits2image (Nation et al. 2022), Flask (Pallets Projects 2022), Flask-SQLAlchemy (Pallets Projects 2021), fundamentals (Young 2023), gracedb-sdk (Singer 2022), HEALPix Alchemy (Singer et al. 2022), healpy (Zonca et al. 2019), Hop Client (Godwin 2022), Hopskotch (SCiMMA Project 2023), Light Curve Fitting (Hosseinzadeh et al. 2023a), ligo.skymap (Singer & Price 2016), lmfit (Newville et al. 2023), Apache Kafka (Apache Software Foundation 2023), Kilonova Candidate Vetting (Rastinejad & Hosseinzadeh 2023), Matplotlib (Hunter 2007), MMT Observatory Manager (Gibson & Porter 2018), MOCPy (Baumann et al. 2023), NumPy (Harris et al. 2020), Paramiko (Forcier et al. 2023), Photutils (Bradley et al. 2022), Pillow (Murray & van Kemenade 2023), plotly.py (Plotly 2023), PostgreSQL (PostgreSQL Global Development Group 2022), Psycopg (Di Gregorio & Varrazzo 2023), PyEphem (Rhodes 2011), pyFFTW (Gomersall 2016), PyMMT (Wyatt et al. 2023), Python-Markdown (Stienstra et al. 2023), Q3C (Koposov & Bartunov 2006), Requests (Reitz 2023b), SAGUARO Pipeline (Paterson et al. 2023), SAGUARO TOM (Hosseinzadeh et al. 2023b), SASSy Q3C Models (Daly et al. 2023), SCAMP (Bertin 2006, 2010a), scikit-image (van der Walt et al. 2014), scikit-learn (Pedregosa et al. 2011), SciPy (Virtanen et al. 2020), setuptools (Python Packaging Authority 2023), sip_tpv (Shupe et al. 2012), Slack (Slack Technologies 2023), slackclient (Slack Technologies 2022), Source Extractor (Bertin & Arnouts 1996, 2010), specutils (Earl et al. 2023), SQLAlchemy (Bayer 2012), SWarp (Bertin 2010b), TensorFlow (Abadi et al. 2016a, 2016b; TensorFlow Developers 2023), TOM Toolkit (Collom et al. 2020; Lindstrom et al. 2022), tom-alert-streams (TOM Toolkit Projectet al. 2023), tom_mmt (Hosseinzadeh & Shrestha 2023), tom_nonlocalizedevents (TOM Toolkit Projectet al. 2023), twilio-python (Twilio 2023), urllib3 (Petrov et al. 2023), voevent-

parse (Staley 2014), Watchdog (Mangalapilly 2023), White-Noise (Evans 2023), ZOGY (Pallets Projects 2023).

ORCID iDs

Griffin Hosseinzadeh <https://orcid.org/0000-0002-0832-2974>

Kerry Paterson <https://orcid.org/0000-0001-8340-3486>

Jillian C. Rastinejad <https://orcid.org/0000-0002-9267-6213>

Manisha Shrestha <https://orcid.org/0000-0002-4022-1874>

Michael J. Lundquist <https://orcid.org/0000-0001-9589-3793>

David J. Sand <https://orcid.org/0000-0003-4102-380X>

Wen-fai Fong <https://orcid.org/0000-0002-7374-935X>

K. Azalee Bostroem <https://orcid.org/0000-0002-4924-444X>

Saarah Hall <https://orcid.org/0000-0002-3841-380X>

Samuel D. Wyatt <https://orcid.org/0000-0003-2732-4956>

Alex R. Gibbs <https://orcid.org/0000-0002-2575-2618>

Joseph Chatelain <https://orcid.org/0000-0002-1278-5998>

Curtis McCullly <https://orcid.org/0000-0001-5807-7893>

References

Abadi, M., Agarwal, A., Barham, P., et al. 2016a, arXiv:[1603.04467](https://arxiv.org/abs/1603.04467)

Abadi, M., Barham, P., Chen, J., et al. 2016b, arXiv:[1605.08695](https://arxiv.org/abs/1605.08695)

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020, *ApJL*, **892**, L3

Acero-Cuellar, T., Bianco, F., Dobler, G., et al. 2023, *AJ*, **166**, 115

Ackley, K., Amati, L., Barbieri, C., et al. 2020, *A&A*, **643**, A113

Ahumada, R., Prieto, C. A., Almeida, A., et al. 2020, *ApJS*, **249**, 3

Alam, S., Albareti, F. D., Prieto, C. A., et al. 2015, *ApJS*, **219**, 12

Alexander, K. D., Schroeder, G., Paterson, K., et al. 2021, *ApJ*, **923**, 66

Anand, S., Coughlin, M. W., Kasliwal, M. M., et al. 2021, *NatAs*, **5**, 46

Anand, S., Necker, J., Earley, N., et al. 2023, GCN, 34751, <https://gcn.nasa.gov/circulars/34751.txt>

Andreoni, I., Goldstein, D. A., Anand, S., et al. 2019, *ApJL*, **881**, L16

Andreoni, I., Goldstein, D. A., Kasliwal, M. M., et al. 2020, *ApJ*, **890**, 131

Antier, S., Agayeva, S., Aivazyan, V., et al. 2020a, *MNRAS*, **492**, 3904

Antier, S., Agayeva, S., Almualla, M., et al. 2020b, *MNRAS*, **497**, 5518

Anumarlapudi, A., Ahumada, T., Kasliwal, M., et al. 2023, GCN, 34089, <https://gcn.nasa.gov/circulars/34089.txt>

Apache Software Foundation 2023, Apache Kafka v3.5.1, <https://kafka.apache.org/>

Araujo, M. 2023, django-crispy-forms v2.0, Github, <https://github.com/django-crispy-forms/django-crispy-forms>

Arcavi, I. 2018, *ApJL*, **855**, L23

Arcavi, I., Hosseinzadeh, G., Howell, D. A., et al. 2017, *Natur*, **551**, 64

Astropy Collaboration 2013, *A&A*, **558**, A33

Astropy Collaboration 2018, *AJ*, **156**, 123

Astropy Collaboration 2022, *ApJ*, **935**, 167

Balcerzak, L. 2021, django-guardian v2.4.0, GitHub, <https://github.com/django-guardian/django-guardian>

Baumann, M., Boch, T., Pineau, F.-X., et al. 2022, *ASPC*, **532**, 7

Baumann, M., Manon, P. F.-X., et al. 2023, cds-astro/mocpy: Release v0.13.0, v0.13.0, Zenodo, doi:[10.5281/zenodo.8297730](https://doi.org/10.5281/zenodo.8297730)

Bayer, M. 2012, in The Architecture of Open Source Applications Volume II: Structure, Scale, and a Few More Fearless Hacks, ed. A. Brown & G. Wilson (Mountain View, CA: AOSA), 291

Becerra, R. L., Dichiara, S., Watson, A. M., et al. 2021, *MNRAS*, **507**, 1401

Beck, R., Szapudi, I., Flewelling, H., et al. 2021, *MNRAS*, **500**, 1633

Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002

Bertin, E. 2006, *ASPC*, **351**, 112

Bertin, E. 2010a, SCAMP: Automatic Astrometric and Photometric Calibration, Astrophysics Source Code Library, ascl:[1010.063](https://ascl.net/1010.063)

Bertin, E. 2010b, SWarp: Resampling and Co-adding FITS Images Together, Astrophysics Source Code Library, ascl:[1010.068](https://ascl.net/1010.068)

Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393

Bertin, E., & Arnouts, S. 2010, SExtractor: Source Extractor, Astrophysics Source Code Library, ascl:[1010.064](https://ascl.net/1010.064)

Bhakta, D., Mooley, K. P., Corsi, A., et al. 2021, *ApJ*, **911**, 77

Bloom, J. S., Kulkarni, S. R., & Djorgovski, S. G. 2002, *AJ*, **123**, 1111

Boch, T., & Fernique, P. 2014, *ASPC*, **485**, 277

Bradley, L., Sipőcz, B., Robitaille, T., et al. 2022, *astropy/photutils*: 1.5.0, v1.5.0, Zenodo, doi:10.5281/zenodo.6825092

Breiman, L. 2001, *Mach. Learn.*, **45**, 5

Brown, T. M., Baliber, N., Bianco, F. B., et al. 2013, *PASP*, **125**, 1031

Cabrera-Vives, G., Reyes, I., Förster, F., Estévez, P. A., & Maureira, J.-C. 2017, *ApJ*, **836**, 97

Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560

Chang, S.-W., Onken, C. A., Wolf, C., et al. 2021, *PASA*, **38**, e024

Chen, Z., Zhou, W., Sun, G., et al. 2023, arXiv:2309.09937

Chornock, R., Berger, E., Kasen, D., et al. 2017, *ApJL*, **848**, L19

Christensen, E., Africano, B., Farneth, G., et al. 2019, *EPSC*, **13**, 1912

Christie, T. 2022, Django REST Framework v3.14.0, GitHub, <https://github.com/encode/django-rest-framework>

Clemens, J. C., Crain, J. A., & Anderson, R. 2004, *Proc. SPIE*, **5492**, 331

Collom, D., Lindstrom, L., Riba, A., et al. 2020, The TOM Toolkit, v2.0.0, Zenodo, doi:10.5281/zenodo.4437764

Coughlin, M. W., Ahumada, T., Anand, S., et al. 2019, *ApJL*, **885**, L19

Coughlin, M. W., Bloom, J. S., Nir, G., et al. 2023, *ApJS*, **267**, 31

Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, *Sci*, **358**, 1556

Coulter, D. A., Jones, D. O., McGill, P., et al. 2023, *PASP*, **135**, 064501

Daly, P. N., Bostroem, K. A., & Hosseinzadeh, G. 2023, SASSy Q3C Models v1.4.0, v1.4.0, Zenodo, doi:10.5281/zenodo.8436176

Dálya, G., Díaz, R., Bouchet, F. R., et al. 2022, *MNRAS*, **514**, 1403

Dálya, G., Galgócz, G., Dobos, L., et al. 2018, *MNRAS*, **479**, 2374

de Jaeger, T., Shappee, B. J., Kochanek, C. S., et al. 2022, *MNRAS*, **509**, 3427

de Wet, S., Groot, P. J., Bloemen, S., et al. 2021, *A&A*, **649**, A72

Dichiara, S., Becerra, R. L., Chase, E. A., et al. 2021, *ApJL*, **923**, L32

Di Gregorio, F., & Varrazzo, D. 2023, Psycopg v2.9.6, GitHub, <https://github.com/psycopg/psycopg2>

Django Software Foundation 2022, Django “excontribute” Comments v2.2.0, GitHub, <https://github.com/django/django-contrib-comments>

Django Software Foundation 2023, Django v4.2, GitHub, <https://github.com/django/django/>

Dobie, D., Stewart, A., Hotokezaka, K., et al. 2022, *MNRAS*, **510**, 3794

Dobie, D., Stewart, A., Murphy, T., et al. 2019, *ApJL*, **887**, L13

Earl, N., Tollerud, E., O’Steen, R., et al. 2023, *astropy/specutils*: v1.11.0, v1.11.0, Zenodo, doi:10.5281/zenodo.8049033

Eikenberry, S., Elston, R., Raines, S. N., et al. 2006, *Proc. SPIE*, **6269**, 409

Elias, J. H., Joyce, R. R., Liang, M., et al. 2006a, *Proc. SPIE*, **6269**, 1389

Elias, J. H., Rodgers, B., Joyce, R. R., et al. 2006b, *Proc. SPIE*, **6269**, 374

Evans, D. 2023, WhiteNoise v6.5.0, GitHub, <https://github.com/evansd/whitenoise>

Fabricant, D., Fata, R., Epps, H., et al. 2019, *PASP*, **131**, 075004

Fabricant, D., Fata, R., Roll, J., et al. 2005, *PASP*, **117**, 1411

Fernique, P., Boch, T., Donaldson, T., et al. 2014, IVOA Recommendation 02 June 2014

Flesch, E. W. 2015, *PASA*, **32**, e010

Flesch, E. W. 2021, arXiv:2105.12985

Fong, W., Nugent, A. E., Dong, Y., et al. 2022, *ApJ*, **940**, 56

Forcier, J. 2023, Paramiko v3.1.0, GitHub, <https://github.com/paramiko/paramiko>

Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, *AJ*, **161**, 242

Foulis, S., & Freitag, F. 2023, django-phonenumbers v7.1.0, GitHub, <https://github.com/stefanfoulis/django-phonenumbers>

Fukushima, K. 1980, *Biol. Cybern.*, **36**, 193

Gaia Collaboration 2018, *A&A*, **616**, A1

Gaia Collaboration 2023, *A&A*, **674**, A1

Gal-Yam, A. 2021, AAS Meeting, **237**, 423.05

Garcia, A., Morgan, R., Herner, K., et al. 2020, *ApJ*, **903**, 75

Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *ApJ*, **611**, 1005

Gibson, J. D., & Porter, D. 2018, *Proc. SPIE*, **10707**, 1070710

Gibson, N. 2021, Django Filter v21.1, GitHub, <https://github.com/carltonborg/django-filter>

Gieseke, F., Bloemen, S., van den Bogaard, C., et al. 2017, *MNRAS*, **472**, 3101

Gillanders, J. H., Troja, E., Fryer, C. L., et al. 2023, arXiv:2308.00633

Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, **157**, 98

Godwin, P. 2022, Hop Client v0.8.0, GitHub, <https://github.com/scimma/hop-client>

Goldstein, A., Veres, P., Burns, E., et al. 2017, *ApJL*, **848**, L14

Goldstein, D. A., Andreoni, I., Nugent, P. E., et al. 2019, *ApJL*, **881**, L7

Gomersall, H. 2016, pyFFTW, v1, Zenodo, doi:10.5281/zenodo.59508

Gomez, S., Hosseinzadeh, G., Cowperthwaite, P. S., et al. 2019, *ApJL*, **884**, L55

Gompertz, B. P., Cutler, R., Steeghs, D., et al. 2020, *MNRAS*, **497**, 726

Gompertz, B. P., Levan, A. J., Tanvir, N. R., et al. 2018, *ApJ*, **860**, 62

Goode, S., Cooke, J., Zhang, J., et al. 2022, *MNRAS*, **513**, 1742

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357

Hook, I. M., Jørgensen, I., Allington-Smith, J. R., et al. 2004, *PASP*, **116**, 425

Hosseinzadeh, G., Bostroem, K. A., & Gomez, S. 2023a, Light Curve Fitting, v0.9.0, Zenodo, doi:10.5281/zenodo.8049154

Hosseinzadeh, G., Cowperthwaite, P. S., Gomez, S., et al. 2019, *ApJL*, **880**, L4

Hosseinzadeh, G., Rastinnejad, J., & Shrestha, M. 2023b, SAGUARO Target and Observation Manager v1.0.0, v1.0.0, Zenodo, doi:10.5281/zenodo.8436090

Hosseinzadeh, G., & Shrestha, M. 2023, TOM MMT Observatory Module v0.1.2, v0.1.2, Zenodo, doi:10.5281/zenodo.8436192

Howell, D. A. & Global Supernova Project 2017, AAS meeting 230, **318**.03

Hunter, J. D. 2007, *CSE*, **9**, 90

IceCube Collaboration, ASAS-SN, Astrophysical Multimessenger Observatory Network, et al. 2017, *A&A*, **607**, A115

Ivezić, Ž., Kahn, S. M., & Tyson, J. A. 2019, *ApJ*, **873**, 111

Jayasinghe, T., Stanek, K. Z., Kochanek, C. S., et al. 2019, *MNRAS*, **485**, 961

Johnston, C., Groot, P., Wet, S. D., et al. 2023, *TNSAN*, **121**, 1

Kasen, D., Metzger, B., Barnes, J., Quataert, E., & Ramirez-Ruiz, E. 2017, *Natur*, **551**, 80

Kasliwal, M. M., Anand, S., Ahumada, T., et al. 2020, *ApJ*, **905**, 145

Kasliwal, M. M., Cannella, C., Bagdasaryan, A., et al. 2019, *PASP*, **131**, 038003

Kilpatrick, C. D., Coulter, D. A., Arcavi, I., et al. 2021, *ApJ*, **923**, 258

Koposov, S., & Bartunov, O. 2006, in ASP Conf. Ser. 351, Astronomical Data Analysis Software and Systems XV, ed. C. Gabriel (San Francisco, CA: ASP), **735**

Kovlakas, K., Zezas, A., Andrews, J. J., et al. 2021, *MNRAS*, **506**, 1896

Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, *AJ*, **139**, 1782

Levan, A., Gompertz, B. P., Salafia, O. S., et al. 2023, arXiv:2307.02098

LIGO Scientific Collaboration, Virgo Collaboration 2020, *ApJL*, **896**, L44

LIGO Scientific Collaboration, Virgo Collaboration 2021a, *PhRvX*, **11**, 021053

LIGO Scientific Collaboration, Virgo Collaboration 2021b, arXiv:2108.01045

LIGO Scientific Collaboration, Virgo Collaboration, 1M2H Collaboration, et al. 2017a, *Natur*, **551**, 85

LIGO Scientific Collaboration, Virgo Collaboration, Fermi Gamma-Ray Burst Monitor, et al. 2017c, *ApJL*, **848**, L13

LIGO Scientific Collaboration, Virgo Collaboration, Fermi GBM, et al. 2017b, *ApJL*, **848**, L12

LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration 2021, *ApJL*, **915**, L5

LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration 2023, *PhRvX*, **13**, 041039

Lindegren, L., Hernández, J., Bombrun, A., et al. 2018, *A&A*, **616**, A2

Lindstrom, W., Chatelain, J., Collom, D., et al. 2022, TOM Toolkit: Target and Observation Manager Toolkit, Astrophysics Source Code Library, ascl:2208.004

Lipunov, V. M., Gorbovskoy, E., Kornilov, V. G., et al. 2017, *ApJL*, **850**, L1

Liu, J. J., Boongoen, T., & Iam-On, N. 2023, *Data Inf. Manag.*, in press

Lone, O. 2022, django-webpack-loader v1.6.0, GitHub, <https://github.com/django-webpack/django-webpack-loader>

Lundquist, M. J., Paterson, K., Fong, W., et al. 2019, *ApJL*, **881**, L26

Mangalapilly, M. J. 2023, Watchdog v3.0.0, GitHub, <https://github.com/gorakhargosh/watchdog>

Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003

Matheson, T., Stubens, C., Wolf, N., et al. 2021, *AJ*, **161**, 107

McCully, C., Crawford, S., Kovacs, G., et al. 2018a, *astropy/astroscrappy*: v1.0.5 Zenodo Release, v1.0.5, Zenodo, doi:10.5281/zenodo.1482019

McCully, C., Hiramatsu, D., Howell, D. A., et al. 2017, *ApJL*, **848**, L32

McCully, C., Turner, M., Volgenau, N., et al. 2018b, LCOGT/banzai: Initial Release, v0.9.4, Zenodo, doi:10.5281/zenodo.1257560

McLeod, B., Fabricant, D., Nystrom, G., et al. 2012, *PASP*, **124**, 1318

Merkel, D. 2014, *Linux j.*, 2014, 2

Möller, A., Peloton, J., Ishida, E. E. O., et al. 2021, *MNRAS*, **501**, 3272

Morgan, R., Soares-Santos, M., Annis, J., et al. 2020, *ApJ*, **901**, 83

Morris, B. M., Tollerud, E., Sipőcz, B., et al. 2018, *AJ*, **155**, 128

Murray, A., van Kemenade, H., & wiredfool 2023, python-pillow/Pillow: 10.0.1, v10.0.1, Zenodo, doi:10.5281/zenodo.8349181

Nation, J. 2022, LCOGT/fits2image v0.4.6, GitHub, <https://github.com/LCOGT/fits2image>

Necker, J., Anand, S., Reusch, S., et al. 2023, GCN, 34688, <https://gcn.nasa.gov/circulars/34688.txt>

Newville, M., Otten, R., Nelson, A., et al. 2023, lmfit/lmfit-py: 1.2.2, v1.2.2, Zenodo, doi:[10.5281/zenodo.8145703](https://doi.org/10.5281/zenodo.8145703)

Nicholl, M., Berger, E., Kasen, D., et al. 2017, *ApJL*, **848**, L18

Niemeyer, M., Pieviläinen, T., Leeuw, Y., Ganssle, P., et al. 2021, dateutil / dateutil v2.8.2, GitHub, <https://github.com/dateutil/dateutil>

Nordin, J., Brinnel, V., van Santen, J., et al. 2019, *A&A*, **631**, A147

Oates, S. R., Marshall, F. E., Breeveld, A. A., et al. 2021, *MNRAS*, **507**, 1296

Ohgami, T., Tominaga, N., Utsumi, Y., et al. 2021, *PASJ*, **73**, 350

Pallets Projects 2021, Flask-SQLAlchemy v2.5.1, GitHub, <https://github.com/pallets-eco/flask-sqlalchemy>

Pallets Projects 2022, Flask v2.2.2, GitHub, <https://github.com/pallets/flask>

Paterson, K. 2023, ZOGY v0.64, GitHub, <https://github.com/KerryPaterson/ZOGY>

Paterson, K., Lundquist, M., & Hosseinzadeh, G. 2023, SAGUARO Pipeline v2.1.0, v2.1.0, Zenodo, doi:[10.5281/zenodo.8436113](https://doi.org/10.5281/zenodo.8436113)

Paterson, K., Lundquist, M. J., Rastinejad, J. C., et al. 2021, *ApJ*, **912**, 128

Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, *PASP*, **131**, 018001

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J.Mach. Learn. Res.*, **12**, 2825

Petrov, A., Larson, S. M., Pradet, Q., et al. 2023, urllib3/urllib3 v1.26.15, GitHub, <https://github.com/urllib3/urllib3>

Pian, E., D'Avanzo, P., Benetti, S., et al. 2017, *Natur*, **551**, 67

Plotly 2023, plotly.py v5.14.1, GitHub, <https://github.com/plotly/plotly.py>

PostgreSQL Global Development Group 2022, PostgreSQL v14.6, <https://www.postgresql.org/>

Pozanenko, A. S., Minaev, P. Y., Grebenev, S. A., & Chelovekov, I. V. 2020, *AstL*, **45**, 710

Python Packaging Authority 2023, setuptools v67.6.1, GitHub, <https://github.com/pypa/setuptools>

Rastinejad, J., & Hosseinzadeh, G. 2023, Kilonova Candidate Vetting v1.0.0, v1.0.0, Zenodo, doi:[10.5281/zenodo.8436164](https://doi.org/10.5281/zenodo.8436164)

Rastinejad, J. C., Fong, W., Kilpatrick, C. D., et al. 2021, *ApJ*, **916**, 89

Rastinejad, J. C., Gompertz, B. P., Levan, A. J., et al. 2022a, *Natur*, **612**, 223

Rastinejad, J. C., Paterson, K., Fong, W., et al. 2022b, *ApJ*, **927**, 50

Reitz, K. 2023b, Requests v2.28.2, GitHub, <https://github.com/psf/requests>

Rhodes, B. C. 2011, PyEphem: Astronomical Ephemeris for Python, Astrophysics Source Code Library, ascl:[2012.014](https://doi.org/10.3843/ascl.2012.014)

Richardson, L. 2023, Beautiful Soup v4.12.2, Launchpad, <https://launchpad.net/beautifulsoup/>

Rimoldini, L., Holl, B., Gavras, P., et al. 2023, *A&A*, **674**, A14

Savchenko, V., Ferrigno, C., Kuulkers, E., et al. 2017, *ApJL*, **848**, L15

Schlawin, E., Herter, T. L., Henderson, C., et al. 2014, *Proc. SPIE*, **9147**, 2H

SCiMMA Project 2023, Hopskotch, <https://scimma.org/hopskotch.html>

Shappee, B. J., Simon, J. D., Drout, M. R., et al. 2017, *Sci*, **358**, 1574

Shingles, L., Smith, K. W., Young, D. R., et al. 2021, Transient Name Server AstroNote TNS [2021-7](https://tns.astro.washington.edu/)

Shrestha, M., Pearson, J., Wyatt, S., et al. 2024, *ApJ*, **961**, 247

Shrestha, M., Sand, D. J., Alexander, K. D., et al. 2023, *ApJL*, **946**, L25

Shupe, D. L., Laher, R. R., Storrie-Lombardi, L., et al. 2012, *Proc. SPIE*, **8451**, E1M

Singer, L. 2022, gracedb-sdk v0.1.7, <https://git.ligo.org/emfollow/gracedb-sdk>

Singer, L. P., Parazin, B., Coughlin, M. W., et al. 2022, *AJ*, **163**, 209

Singer, L. P., & Price, L. R. 2016, *PhRvD*, **93**, 024013

Slack Technologies 2022, slackclient v2.9.4, GitHub, <https://github.com/slackapi/python-slack-sdk>

Slack Technologies 2023, Slack v4.34.121, <https://slack.com>

Smartt, S. J., Chen, T.-W., Jerkstrand, A., et al. 2017, *Natur*, **551**, 75

Smartt, S. J., Valenti, S., Fraser, M., et al. 2015, *A&A*, **579**, A40

Smith, D. 2022, crispy-bootstrap4 v2022.1, GitHub, <https://github.com/django-crispy-forms/crispy-bootstrap4>

Smith, K. W., Smartt, S. J., Young, D. R., et al. 2020, *PASP*, **132**, 085002

Smith, K. W., Williams, R. D., Young, D. R., et al. 2019, *RNAAS*, **3**, 26

Soares-Santos, M., Holz, D. E., Annis, J., et al. 2017, *ApJL*, **848**, L16

Staley, T. D. 2014, voevent-parse: Parse, manipulate, and generate VOEvent XML packets, Astrophysics Source Code Library, ascl:[1411.003](https://doi.org/10.3843/ascl.1411.003)

Stienstra, M., Takhteyev, Y., & Limberg, W. 2023, Python-Markdown v3.4.3, GitHub, <https://github.com/Python-Markdown/markdown>

Street, R. A., Adamson, A., Blakeslee, J. P., et al. 2020, *Proc. SPIE*, **11449**, 1144925

Street, R. A., Bowman, M., Saunders, E. S., & Boroson, T. 2018, *Proc. SPIE*, **10707**, 1070711

Szentygyorgyi, A., Furesz, G., Cheimets, P., et al. 2011, *PASP*, **123**, 1188

Tachibana, Y., & Miller, A. A. 2018, *PASP*, **130**, 128001

Tanvir, N. R., Levan, A. J., González-Fernández, C., et al. 2017, *ApJL*, **848**, L27

Tebeka, M. 2022, fastavro v1.6.1, GitHub, <https://github.com/fastavro/fastavro>

TensorFlow Developers 2023, TensorFlow, v2.14.0, Zenodo, doi:[10.5281/zenodo.8381573](https://doi.org/10.5281/zenodo.8381573)

Thakur, A. L., Dichiara, S., Troja, E., et al. 2020, *MNRAS*, **499**, 3868

The 2023 Windows on the Universe Workshop White Paper Working Group 2024, arXiv:[2401.02063](https://arxiv.org/abs/2401.02063)

TOM Toolkit Project, Collom, D., Lindstrom, L., & Nation, J. 2023, tom_nonlocalizedevents v0.7.7, GitHub, https://github.com/TOMToolkit/tom_nonlocalizedevents

TOM Toolkit Project, Lindstrom, L. 2023, tom_nonlocalizedevents v0.7.7, GitHub, <https://github.com/TOMToolkit/tom-alertstreams>

Tonry, J., Stalder, B., Denneau, L., et al. 2018a, Transient Name Server Discovery Report TNS [2018-707](https://tns.astro.washington.edu/)

Tonry, J. L., Denneau, L., Flewelling, H., et al. 2018b, *ApJ*, **867**, 105

Trier, M., & van Oostveen, B. 2023, Django Extensions v3.2.3, GitHub, <https://github.com/django-extensions/django-extensions>

Troja, E., Fryer, C. L., O'Connor, B., et al. 2022, *Natur*, **612**, 228

Tucker, D. L., Wiesner, M. P., Allam, S. S., et al. 2022, *ApJ*, **929**, 115

Turpin, D., Ganet, M., Antier, S., et al. 2020, *MNRAS*, **497**, 2641

Twilio 2023, twilio-python v8.1.0, GitHub, <https://github.com/twilio/twilio-python>

Valenti, S., David, J. S., Yang, S., et al. 2017, *ApJL*, **848**, L24

van der Walt, S. J., Schönberger, J. L., Nunez-Iglesias, J., et al. 2014, *PeerJ*, **2**, e453

Verheul, D. 2021, django-bootstrap4 v3.0.1, GitHub, <https://github.com/zostera/django-bootstrap4>

Vieira, N., Ruan, J. J., Haggard, D., et al. 2020, *ApJ*, **895**, 96

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261

Waddington, T. 2020, django-gravatar v1.4.4, GitHub, <https://github.com/twaddington/django-gravatar>

Watson, A. M., Butler, N. R., Lee, W. H., et al. 2020, *MNRAS*, **492**, 5916

Watson, D., Hansen, C. J., Selsing, J., et al. 2019, *Natur*, **574**, 497

White, D. J., Daw, E. J., & Dhillon, V. S. 2011, *CQGra*, **28**, 085016

Williams, G. G., Olszewski, E., Lesser, M. P., & Burge, J. H. 2004, *Proc. SPIE*, **5492**, 787

Wyatt, S., Shrestha, M., & Hosseinzadeh, G. 2023, PyMMT v1.0.0, v1.0.0, Zenodo, doi:[10.5281/zenodo.8322354](https://doi.org/10.5281/zenodo.8322354)

Wyatt, S. D., Tohuvavohu, A., Arcavi, I., et al. 2020, *ApJ*, **894**, 127

Yang, J., Ai, S., Zhang, B.-B., et al. 2022, *Natur*, **612**, 232

Yang, Y.-H., Troja, E., O'Connor, B., et al. 2023, arXiv:[2308.00638](https://arxiv.org/abs/2308.00638)

Young, D. 2023, fundamentals, v2.4.1, Zenodo, doi:[10.5281/zenodo.8037510](https://doi.org/10.5281/zenodo.8037510)

Zackay, B., Ofek, E. O., & Gal-Yam, A. 2016, *ApJ*, **830**, 27

Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298

Zou, H., Zhou, X., Fan, X., et al. 2017, *PASP*, **129**, 064101