

Optimal Follow-Up of Gravitational-Wave Events with the UltraViolet EXplorer (UVEX)

LEO P. SINGER,^{1,2} ALEXANDER W. CRISWELL,^{3,4,5,6} SYDNEY C. LEGGIO,^{3,4} R. WEIZMANN KIENDREBEOGO,^{7,8,4}
MICHAEL W. COUGHLIN,^{3,4} HANNAH P. EARNSHAW,⁹ SUVI GEZARI,^{10,11} BRIAN W. GREFENSTETTE,⁹ FIONA A. HARRISON,⁹
MANSI M. KASLIWAL,⁹ BRETT M. MORRIS,¹¹ ERIK TOLLERUD,¹¹ AND S. BRADLEY CENKO^{1,2}

¹*Astroparticle Physics Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA*

²*Joint Space-Science Institute, University of Maryland, College Park, MD 20742, USA*

³*Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, MN 55455, USA*

⁴*School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA*

⁵*Department of Physics and Astronomy, Vanderbilt University, Nashville, TN 37240, USA*

⁶*Department of Life and Physical Sciences, Fisk University, Nashville, TN 37208, USA*

⁷*Laboratoire de Physique et de Chimie de l'Environnement, Université Joseph KI-ZERBO, Ouagadougou, Burkina Faso*

⁸*Artemis, Observatoire de la Côte d'Azur, Université Côte d'Azur, Boulevard de l'Observatoire, F-06304 Nice, France*

⁹*Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*

¹⁰*Department of Physics and Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, USA*

¹¹*Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA*

Submitted to PASP

ABSTRACT

The UltraViolet EXplorer (UVEX) is a wide-field ultraviolet space telescope selected as a NASA Medium-Class Explorer (MIDEX) mission for launch in 2030. UVEX will undertake deep, cadenced surveys of the entire sky to probe low mass galaxies and explore the ultraviolet (UV) time-domain sky, and it will carry the first rapidly deployable UV spectroscopic capability for a broad range of science applications. One of UVEX's prime objectives is to follow up gravitational wave (GW) binary neutron star mergers as targets of opportunity (ToOs), rapidly scanning across their localization regions to search for their kilonova (KN) counterparts. Early-time multiband ultraviolet light curves of KNe are key to explaining the interplay between jet and ejecta in binary neutron star mergers. Owing to high Galactic extinction in the ultraviolet and the variation of GW distance estimates over the sky, the sensitivity to kilonovae can vary significantly across the GW localization and even across the footprint of a single image given UVEX's large field of view. Good ToO observing strategies to trade off between area and depth are neither simple nor obvious. We present an optimal strategy for GW follow-up with UVEX in which exposure time is adjusted dynamically for each field individually to maximize the overall probability of detection. We model the scheduling problem using the expressive and powerful mathematical framework of mixed integer linear programming (MILP), and employ a state-of-the-art MILP solver to automatically generate observing plan timelines that achieve high probabilities of kilonova detection. We have implemented this strategy in an open-source astronomical scheduling software package called the Multi-Mission Multi-Messenger Observation Planning Toolkit (M⁴OPT), on GitHub at <https://github.com/m4opt/m4opt>.

Keywords: Computational methods (1965) — Gravitational wave astronomy (675) — Open source software (1866) — Ultraviolet observatories (1739) — Ultraviolet transient sources (1854) — Wide-field telescopes (1800)

1. INTRODUCTION

In 2017, the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo (Abbott et al. 2017a) detected

a long-duration gravitational wave (GW) inspiral signal, GW170817, at the same time that Fermi (Goldstein et al. 2017) and INTEGRAL (Savchenko et al. 2017) recorded a short gamma-ray burst (GRB), GRB170817A. The alert sprang traps that had been set by hundreds of telescopes worldwide (Aasi et al. 2014; Abbott et al. 2016) which

quickly found the optical counterpart (Coulter et al. 2017), AT2017gfo. As a measure of the degree to which the event focused the efforts of astronomers everywhere, the author list of Abbott et al. (2017b) runs to 24 pages!

The scientific harvest from this one event was remarkable. It fulfilled a three-decade-old dream of using GWs as “standard sirens” to measure the Hubble constant (Schutz 1986; Abbott et al. 2017c). Moreover, it proved once and for all the hypotheses that neutron star (NS) mergers are the central engines of short GRBs (Fong & Berger 2013) and the main cosmic factories of heavy r -process elements (Freiburghaus et al. 1999).

It had long been understood that such mergers would tidally disrupt their NSs, and that radioactive decay of the heavy elements synthesized in their hot neutron-rich ejecta would fuel transients (Lattimer & Schramm 1974; Eichler et al. 1989; Li & Paczyński 1998) that came to be called kilonovae (KNe).

In those early days, people assumed that the ejecta would have opacities similar to those in supernovae (SNe) and predicted fairly bright KN light curves that peaked in the optical or ultraviolet (UV) and that would be fairly easy to detect. However, further study of the atomic structure of lanthanides led to the realization that their dense absorption spectra would lead to line-blanketing in the optical, containing the radiation and only letting it leak out much more slowly and at longer wavelengths, in the infrared (Kasen et al. 2013). Observers grimly realized that although KNe were still among the most promising counterparts of NS mergers (Metzger & Berger 2012), they would be much dimmer, redder, and harder to detect than previously expected.

Although these later and more sober predictions agreed remarkably well with the observed spectral sequence of AT2017gfo at times later than a few days (Pian et al. 2017; Kasen et al. 2017; Villar et al. 2017), contrary to those expectations it was quite blue and featureless at the earliest observed times, less than a day after the merger (Shappee et al. 2017).

In the fallout from GW170817, the cause of this early optical and UV emission remains one of the most enduring mysteries. The blue emission could be radioactively powered but result from a geometrically distinct outflow component with higher velocity and/or lower lanthanide fraction (Nicholl et al. 2017) or could result from the shock caused by the “cocoon” interaction between the ejecta and the nascent jet (Kasliwal et al. 2017; Gottlieb et al. 2018). Lacking for GW170817, early-time UV observations, less than 12 hours after merger, could handily settle the debate (Arcavi 2018).

1.1. The Coming UV Time Domain Revolution

More generally, there is a recognized gap in transient discovery capability in the UV, and an acknowledged need for a




Figure 1. A rendering of UVEX. Reproduced from the CSR.

space-based UV wide-field time-domain survey (Sagiv et al. 2014). To meet this need, some of the authors proposed Dorado (née GUCI, Cenko 2019; Dorsman et al. 2023) to NASA as a Mission of Opportunity (MoO).

Although NASA made no MoO selection in that cycle, shortly thereafter the National Academies of Sciences, Engineering, and Medicine (2021) recommended in the 2020 decadal survey that “NASA should establish a time-domain program to realize and sustain the necessary suite of space-based electromagnetic capabilities required to study transient and time-variable phenomena, and to follow up multi-messenger events.” NASA soon selected a much larger and more capable mission called the UltraViolet EXplorer (UVEX; Kulkarni et al. 2021; see rendering in Fig. 1)¹ as the next Medium-Class Explorer (MIDEX).

UVEX will have a wide field of view (FOV) camera that will take images simultaneously in both a near ultraviolet (NUV) and a far ultraviolet (FUV) band as well as a long-slit spectrograph. The imaging point spread function (PSF) diameter of about $\sim 2''$ will be well-matched to ground-based follow-up. UVEX will perform a mix of surveys with different sky coverage and cadence, and will observe the entire sky to at least >25.8 mag over the duration of the prime mission. UVEX will not just perform a transformative all-sky time-domain UV survey. It will also be able to perform targets of opportunity (ToOs) to follow up GW mergers and other multi-messenger phenomena (see Fig. 2). With its wide FOV and

¹ <https://www.uvex.caltech.edu>




Figure 2. An example of a ToO observation sequence with UVEX to follow up a GW event. Adapted from Fig. D-8 from the UVEX CSR submitted to NASA HQ, which is also Fig. 14 in Kulkarni et al. (2021).

two UV bandpasses, it should be able to finally discern the physical mechanism of early-time emission in KNe.

NASA has also partnered with Israel to provide a launch for ULTRASAT (Shvartzvald et al. 2024). ULTRASAT will have an even larger FOV than UVEX and will be ready to launch a few years before, but with some tradeoffs: it is much less sensitive, has a much larger PSF, only a single imaging band, and no spectrograph. Still, UVEX and ULTRASAT together will make a potent discovery engine for UV transients in general and GW counterparts in particular.

1.2. A Multi-Mission Multi-Messenger Observation Planning Toolkit

NASA selects Explorer-class missions on the basis of a concept study report (CSR) which undergoes several reviews, culminating in an in-person site visit (SV). In Criswell et al. (2025), we elaborated upon the case for GW follow-up with UVEX more or less as we presented it to NASA in the CSR. That UVEX study leveraged the same GW ToO analysis and strategy that we had developed for the Dorado CSR. We envisioned our observation planning software,

*dorado-scheduling*², as an early draft of what would eventually evolve into real ground software and a part of the guest observer toolkit. Therefore, both studies benefited from an unusual level of fidelity and realism for so early in their respective mission lifecycles.

The *dorado-scheduling* package considers, for any given ToO, the time-varying field of regard (FOR) constraints of the spacecraft, the spacecraft’s slew time between any two target fields, the footprint of the instrument’s FOV, and the probability distribution of the true but unknown position of the source — the GW localization map. Its purpose is to find the sequence of observations that maximizes the probability of detecting the counterpart. This can be surprisingly challenging: as with similar “hard” optimization problems (e.g. the “traveling salesman” problem), the combinatorial scale makes a brute-force search of all possible observation sequences out of the question. On the other hand, simple strategies can yield unacceptable results. For example, observing the fields in descending order of probability may be a poor strategy if it causes one to miss fields that pass out of the FOR early on, or if it results in long slew times. Moreover, it may be labor-intensive to adjust or rewrite any handmade strategy to re-tune it for changes to the mission design. Instead of such sub-optimal heuristics, *dorado-scheduling* directly models the constraints and objective using the versatile and expressive mathematical framework of mixed integer linear programming (MILP), and employs IBM’s state-of-the-art MILP solver, CPLEX,³ to search for the global optimum.

Responding to queries during the Dorado SV about varying extinction and foregrounds, we added to *dorado-scheduling* a powerful capability to dynamically optimize the exposure time of each field to adjust for spatial variations in sensitivity. This refinement brought in several new possible tradeoffs to increase detection probability. For example, the GW localization is a distribution over both sky location and distance, and there is a Malmquist-like tendency for regions of high probability on the sky to favor large distances (Singer et al. 2016a), so at times it is advantageous to observe lower probability fields that require less exposure time. Furthermore, UV sensitivity varies significantly across the sky. For instance, it may be beneficial to spend less time observing high-probability fields that also have high Milky Way dust extinction. The MILP framework allowed us to take advantage of these additional degrees of freedom without needing to meticulously craft the observing strategy itself.

In our simulations, adaptive exposure time greatly improved the detection efficiency of the mission. In earlier re-

² <https://github.com/nasa/dorado-scheduling>

³ <https://www.ibm.com/products/ilog-cplex-optimization-studio>

lated works, [Coughlin & Stubbs \(2016\)](#) derived power-law expressions for optimal exposure time under simplified conditions. [Chan et al. \(2017\)](#) varied exposure times but without considering field-to-field variation in distance, variation in sensitivity, or FOR constraints. [Herner et al. \(2020\)](#) employed exposure time maps to select feasible fields for ToO observations. [Liu et al. \(2021\)](#) varied exposure time in accordance with the expected time-dependence of the light curve. However, we believe that our *adaptive optimization of exposure time for each field, accounting for spatial variation in distance estimate and sensitivity, to globally maximize the probability of detection, is novel*. Unfortunately, this code was not used in the UVEX study or in [Criswell et al. \(2025\)](#).

The `dorado-scheduling` code also had the major limitation that it was released under the obscure and anachronistic NASA Open Source Agreement (NOSA), which placed severe obstacles to our own collaborators using or contributing to it. The lawyers at Goddard Space Flight Center (GSFC) continue to require all GSFC scientists to employ NOSA even though it has been rejected by the [Free Software Foundation \(2006\)](#), the [National Academies of Sciences, Engineering, and Medicine \(2018\)](#), and even the [NASA Science Mission Directorate \(2022\)](#). Fortunately, NASA HQ intervened in this case and we were permitted to establish a new, truly open-source, permissively licensed software project, which anyone can contribute to and use. However, this required a rewrite which took several years.

In this paper, we finally present the Multi-Mission Multi-Messenger Observation Planning Toolkit (M⁴OPT), which has many other advances over earlier work:

- It is released under a permissive, mainstream license (the Berkeley Software Distribution, or BSD, license) to promote adoption and contribution by the community.
- It is designed from the start to support multiple missions, including UVEX and ULTRASAT. We envision supporting both space- and ground-based observatories in the future.
- It can adjust exposure times given the anticipated absolute magnitude and the three-dimensional GW localization distribution in sky position and distance.
- It can dynamically vary the exposure time of each field to pierce through spatial variations in foregrounds (zodiacal light, Galactic diffuse emission) and dust extinction.
- It can be given the anticipated absolute magnitude of the source as either a point estimate or a distribution with Gaussian uncertainty.
- The dynamic exposure times are made possible by a `numpy` ([Harris et al. 2020](#)) vectorized exposure time calculator (ETC) which enables large parameter sweeps of synthetic photometry calculations which are otherwise prohibitively slow with `synphot` ([STScI Development Team 2018](#)) alone.
- It models additional spacecraft dynamics effects, including the roll angle of the telescope which is determined by solar power requirements.
- It already complies fully with NASA Procedural Requirements (NPR) 7150 ([NASA 2022](#)) software engineering practices for “Class C” software, and has 95% test coverage.
- It is deeply integrated with the Astropy ([Astropy Collaboration et al. 2013, 2018](#)) ecosystem, and has an interface that is based on Astropy coordinates, units, and model classes.

In this paper, we describe the mathematical approach of M⁴OPT, and then use it to produce realistic ToO observing sequences with UVEX for simulated GW events. An example M⁴OPT observing plan is shown in [Fig. 3](#) and [Table 1](#). We show that the new dynamic exposure time capabilities dramatically increases the probability of KN detection.

2. BASICS OF MILP

Linear programming (LP) is a mathematical optimization formalism in which one represents an objective as a linear combination of decision variables, subject to constraints that take the form of a system of linear inequalities. The canonical form a LP is

$$\begin{aligned}
 &\text{Find} && \mathbf{x} \in \mathbb{R}^N \\
 &\text{that maximizes} && \mathbf{c}^T \mathbf{x} \\
 &\text{subject to} && \mathbf{A} \mathbf{x} \leq \mathbf{0} \\
 &\text{and} && \mathbf{x} \geq \mathbf{0}.
 \end{aligned}$$

LP is useful for problems of resource allocation. If instead of being reals, certain of the decision variables are required to be integers, then the problem is called mixed integer linear programming (MILP). MILP allows one to model situations where some decision variables represent alternative courses of action or where some constraints are Boolean in nature. The textbook chapters of [Chen et al. \(2009\)](#) and [Williams \(2013\)](#) serve as useful translation dictionaries of sorts from logical constraints to integer linear inequalities.

There are a variety of free and open source MILP solvers; the most readily available for Python users is HiGHS ([Huangfu & Hall 2018](#)) which now can be called directly




Figure 3. An example M⁴OPT observing plan. The top panel shows the footprints of the observations, the 90% credible region of the GW localization, the region that is always outside of the FOR as a deep blue filled region, and the region that is outside of the instantaneous FOR. The bottom panel shows the timeline of the accumulation of detection probability and area. The region that has been covered by one visit is shown in pale red and green, whereas the region that has been covered by two visits is shown in deep red and green. This figure is available as an animation.

from SciPy (Virtanen et al. 2020). However, the best commercial MILP solvers are much faster and can handle much larger problems than any current open-source software. Although M⁴OPT itself is open source, it calls the commercial MILP solver CPLEX by IBM. CPLEX is free for academic and research use by students, faculty, and staff at academic institutions. For non-academic users, it is priced affordably, comparable in cost to mid-range computer-aided design software.

Common applications of MILP include planning and scheduling. There are already many noteworthy and successful uses of MILP to astronomical observation planning. Las Cumbres Observatory (LCO) uses it to allocate observations on a global queue-scheduled network of telescopes (Saunders et al. 2014). Zwicky Transient Facility (ZTF) uses it to multiplex several sky surveys with different coverage and cadence

requirements on a single telescope (Bellm et al. 2019), and our colleagues have proposed using it to schedule ToO observations with ZTF (Parazin et al. 2022). It has been applied to scheduling observing programs on Atacama Large Millimeter/Submillimeter Array (ALMA; Solar et al. 2016) and exoplanet searches on Keck (Handley et al. 2024).

The classic maximum weighted coverage problem (MWC) has a MILP representation that is at the core of our scheduler. In MWC, one has a finite sequence of real-valued weights, $(w_j)_j$, and a finite set of sets, $S = \{S_i\}_i$, over the integers, $\forall i : S_i \subset \mathbb{Z}$. The objective is to find a subset $S' \subseteq S$, with a maximum cardinality $|S'| \leq k$, that maximizes the sum over all of the weights $\sum_{j \in \cup S'} w_j$. This is illustrated in Fig. 4.


Table 1. An example M⁴OPT observing plan (same as Fig. 3).

Start time (UTC)	Dur. (s)	Action	Spacecraft location (ITRS, km)			Target (ICRS, deg)		
			x	y	z	RA	Dec	Roll
2012-06-01 10:23:42	325	observe	-183412	141613	-49650	31.6571	-23.5560	-138
2012-06-01 10:29:06	120	slew	—	—	—	—	—	—
2012-06-01 10:31:06	767	observe	-178552	148032	-49680	32.4155	-28.4429	-141
2012-06-01 10:43:53	127	slew	—	—	—	—	—	—
2012-06-01 10:46:00	1348	observe	-168105	160424	-49740	36.9122	-33.0462	-146
2012-06-01 11:08:28	107	slew	—	—	—	—	—	—
2012-06-01 11:10:15	1529	observe	-149278	178994	-49835	39.3087	-35.3375	-149
2012-06-01 11:35:44	127	slew	—	—	—	—	—	—
2012-06-01 11:37:51	3006	observe	-125392	197415	-49939	34.6192	-30.7459	-143
2012-06-01 12:27:57	120	slew	—	—	—	—	—	—
2012-06-01 12:29:57	2166	observe	-74501	223271	-50125	33.7574	-25.7998	-141
2012-06-01 13:06:04	108	slew	—	—	—	—	—	—
2012-06-01 13:07:51	924	observe	-34127	233984	-50251	35.9256	-28.0474	-144
2012-06-01 13:23:15	107	slew	—	—	—	—	—	—
2012-06-01 13:25:02	767	observe	-15274	236459	-50306	32.4155	-28.4429	-141
2012-06-01 13:37:49	120	slew	—	—	—	—	—	—
2012-06-01 13:39:49	325	observe	1080	237371	-50352	31.6571	-23.5560	-138
2012-06-01 13:45:14	108	slew	—	—	—	—	—	—
2012-06-01 13:47:01	2166	observe	9080	237405	-50374	33.7574	-25.7998	-141
2012-06-01 14:23:08	108	slew	—	—	—	—	—	—
2012-06-01 14:24:55	924	observe	50975	233149	-50485	35.9256	-28.0474	-144
2012-06-01 14:40:19	104	slew	—	—	—	—	—	—
2012-06-01 14:42:04	3006	observe	69587	228794	-50533	34.6192	-30.7459	-143
2012-06-01 15:32:10	107	slew	—	—	—	—	—	—
2012-06-01 15:33:57	1348	observe	123168	206692	-50669	36.9122	-33.0462	-146
2012-06-01 15:56:25	107	slew	—	—	—	—	—	—
2012-06-01 15:58:12	1529	observe	146130	192006	-50728	39.3087	-35.3375	-149

The MWC has a straightforward MILP representation:

$$\begin{aligned}
& \text{Maximize} && \sum_j y_j w_j \\
& \text{subject to the constraints} && \sum_i x_i \leq k \\
& && \text{and} \quad \sum_{i|j \in S_i} x_i \geq y_j.
\end{aligned}$$

The most basic version of our ToO problem is that we want to maximize the total integrated probability that the true, but unknown, position of the source is within one or more of the footprints that we select to observe. Each GW alert comes with a localization sky map that gives the posterior probability distribution of the sky position as an image that is sampled on a Hierarchical Equal Area isoLatitude Pixelization (HEALPix; Górski et al. 2005) grid. The weights $(w_j)_j$ are the probability values and the subsets $(S_i)_i$ are the HEALPix

**Figure 4.** Illustration of MWC. Four sets $S = \{S_1, S_2, S_3, S_4\}$ are represented as regions with colored borders. The elements of those sets are represented by black numbered circles.




Figure 5. MWC applied to coverage of an astronomical ROI by multiple partially overlapping fields. The centers of HEALPix pixels are marked with gray dots.

pixels contained within the footprints of the FOV on a reference grid of allowed pointings of the telescope (see Fig. 5).

3. MILP PROBLEM FORMULATION

Now we develop an MILP formulation of the problem of planning and scheduling GW ToO observations on a large FOV telescope. In Section 3.1, we start with a simple but realistic problem in which we require multiple visits of each selected field that must all obey the FOR and slew speed limitations of the telescope. In Sections 3.2 and 3.3, we add progressively more detail, finally arriving at the full dynamic exposure time problem. All three formulations from Sections 3.1, 3.2, and 3.3 are implemented in M⁴OPT.

3.1. Problem 1: Fixed Exposure Time

We receive a HEALPix probability sky map that describes the probability distribution of the true but unknown position of a target of interest as a function of position on the sky. There is a delay between the time that the event occurred and when we can start observations due to the time it takes to uplink commands to the spacecraft, and there is a deadline by which we must complete our observations.

Our telescope can observe any of a set of n_J fields at predetermined sky locations in order to tile the sky map. For each field that we select, our telescope must visit the field at least n_K times. We have a cadence requirement that each visit of a given field must occur at least a time γ after the previous visit. Multiple visits with a minimum cadence are essential

in short timescale transient searches to rule out moving solar system objects, which otherwise are a major contaminant.

Every visit takes a certain amount of exposure time, and it takes a known amount of time to slew between different fields. We may only visit a field when it is within the FOR, the region that constrains where the telescope may point at any given instant of time.

3.1.1. Data Preparation

1. Construct a discrete 1D grid of times that stretch from the delayed start of observations up to the deadline.
2. Propagate the orbit of the spacecraft to calculate the position of the spacecraft at each time step.
3. For each field and each time step, test whether the field is within the instantaneous FOR, creating an observability bit map.
4. Transform the observability bit map into a list of time segments during which each field is observable.
5. Discard segments shorter than the exposure time.
6. Discard fields that have no observable segments.
7. For each field, find the HEALPix pixel indices that are within the field's footprint.
8. Select the 50 fields that contain the greatest probability, summed over the respective HEALPix pixels.
9. Discard pixels that are not contained in any field.
10. Calculate the slew times between all pairs of distinct fields.

3.1.2. Problem Setup

Index sets —

pixels	$I = \{0, 1, \dots, n_I - 1\}$
fields	$J = \{0, 1, \dots, n_J - 1\}$
visits	$K = \{0, 1, \dots, n_K - 1\}$
observable segments	$(M_j = \{0, 1, \dots, n_{M_j}\})_{j \in J}$
fields containing pixel i	$(J_i = \{0, 1, \dots, n_{J_i}\})_{i \in I}$

Parameters —

probability of pixel i	$(\rho_i)_{i \in I}$
slew time from field j to j'	$(\sigma_{jj'})_{j \in J, j' \in J}$
start times of segments	$(\alpha_{jm})_{j \in J, m \in M}$
end times of segments	$(\omega_{jm})_{j \in J, m \in M}$
exposure time	ϵ
cadence, time between visits	γ
delay	β
deadline	δ

Binary decision variables —

pixel i is in any selected field	$(p_i)_{i \in I}$
field j is selected	$(r_j)_{j \in J}$
field j visit k is in segment m	$(s_{jkm})_{j \in J, k \in K, m \in M n_{Mj} > 1}$

Continuous decision variables —

mid time of field j visit k	$(t_{jk})_{j \in J, k \in K}$
---------------------------------	-------------------------------

3.1.3. Constraints

Containment—Only count pixels that are in one or more selected fields.

$$\forall i : \quad p_i \leq \sum_{j \in J_i} r_j \quad (1)$$

Cadence—If a field is selected for observation, then enforce a minimum time between visits.

$$\forall k > 1, j : \quad t_{jk} - t_{j,k-1} \geq (\epsilon + \gamma)r_j \quad (2)$$

No overlap—Observations cannot overlap in time; they must be separated by at least the exposure time plus the slew time.

$$\forall j' > j, k, k' : \quad |t_{jk} - t_{j'k'}| \geq (\sigma_{jj'} + \epsilon)(r_j + r_{j'} - 1) \quad (3)$$

FOR—An observation of a field can only occur while the coordinates of the field are within the FOR. For fields that have one observable segment ($n_{Mj} = 1$), this constraint is simply an inequality:

$$\forall j, k, m \mid n_{Mj} = 1 : \quad \alpha_{jm} + \epsilon/2 \leq t_{jk} \leq \omega_{jm} - \epsilon/2 \quad (4)$$

For fields that have more than one observable segment ($n_{Mj} > 1$), we use the decision variable s_{jkm} to determine which inequality is satisfied:

$$\forall j, k, m \mid n_{Mj} > 1 : \quad s_{jkm} = 1 \Rightarrow \alpha_{jm} + \epsilon/2 \leq t_{jk} \leq \omega_{jm} - \epsilon/2 \quad (5)$$

$$\sum_m s_{jkm} \geq 1 \quad (6)$$

3.1.4. Cuts

Total exposure time—Although it is implied by other constraints, the constraint that the total exposure time cannot exceed the total available time is found to speed up the search. We add it as a cut: an inequality that the MILP may use to help guide its search but that is not checked when evaluating integer feasibility.

$$\sum_{j \in J} r_j \leq \frac{\delta - \beta}{\epsilon n_K} \quad (7)$$

3.1.5. Objective

Maximize the sum of the probability of all of the pixels that are contained within selected fields:

$$\sum_{i \in I} \rho_i p_i \quad (8)$$

3.2. Problem 2: Variable Exposure Time

In this variation, we have a sky map of the exposure time required to detect the source as a function of its position on the sky. We permit the exposure time to vary for each field. A given pixel counts toward the objective value only if the exposure time of a field that contains that pixel exceeds the pixel's exposure time.

*3.2.1. Problem Setup**Additional parameters —*

min exposure time to detect a source in pixel i	$(\epsilon_i)_{i \in I}$
min allowed exposure time	$\epsilon_{\min}$
max allowed exposure time	$\epsilon_{\max}$

Additional, semicontinuous decision variables —

exposure time of field j

$$(e_j)_{j \in J}, \forall j \in J : e_j = 0 \text{ or } \epsilon_{\min} \leq e_j \leq \epsilon_{\max}$$

3.2.2. Constraints

The constraints are slightly different.

Depth—Only count pixels that are observed to sufficient exposure time.

$$\forall i \in I : \quad p_i = 1 \Rightarrow \max_{j \in J_i} e_j \geq \epsilon_i \quad (9)$$

Exposure time—If a field's exposure time is nonzero, then it is selected for observation.

$$\forall j \in J : \quad \epsilon_{\max} r_j \geq e_j \quad (10)$$

Cadence—This is similar to Eq. (2), except that we replace the right-hand side of the inequality.

$$\forall k > 1, j : t_{jk} - t_{j,k-1} \geq \gamma r_j + e_j \quad (11)$$

No overlap—This is also similar to Eq. (3), except with a slightly different right-hand side.

$$\forall j' > j, k, k' :$$

$$|t_{jk} - t_{j'k'}| \geq \sigma_{jj'} (r_j + r_{j'} - 1) + (e_j + e_{j'})/2 \quad (12)$$

FOR—This is similar to Eqs. (4, 5), except that we replace ϵ with e_j . For fields that have one observable segment:

$$\forall j, k, m \mid n_{Mj} = 1 :$$

$$\alpha_{jm} + e_j/2 \leq t_{jk} \leq \omega_{jm} - e_j/2 \quad (13)$$

For fields that have more than one observable segment:

$$\forall j, k, m \mid n_{Mj} > 1 :$$

$$s_{jkm} = 1 \Rightarrow \alpha_{jm} + e_j/2 \leq t_{jk} \leq \omega_{jm} - e_j/2 \quad (14)$$

$$\sum_m s_{jkm} \geq 1 \quad (15)$$

3.2.3. Cuts

Total exposure time—Replace Eq. (7) with:

$$\sum_{j \in J} r_j \leq \frac{\delta - \beta}{\epsilon_{\min} n_K} \quad (16)$$

$$\sum_{j \in J} e_j \leq \frac{\delta - \beta}{n_K} \quad (17)$$

3.2.4. Objective

Same as in Section 3.1.5 above.

3.3. Problem 3: Variable Exposure Time with Prior Distribution of Absolute Magnitude

In this variation, we don't know the precise absolute magnitude X of the source. In the case of KNe, our prior knowledge about the absolute magnitude is scant; for the sake of mathematical convenience, we assume that the absolute magnitude has a normal distribution, $X \sim \mathcal{N}[\mu_X, \sigma_X]$. We need to compute the distribution of *apparent* magnitudes x in order to determine the probability of detection as a function of exposure time for each pixel.

Gravitational-wave sky maps provide the posterior distribution of distance, as a parametric ansatz distribution (Singer et al. 2016a,b),

$$p(r) = \frac{N}{\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2} \left(\frac{r - \mu}{\sigma} \right)^2 \right] r^2,$$

with the location parameter μ , scale parameter σ , and normalization N tabulated for each pixel. This is an inconvenient distribution for integration, so instead we construct




Figure 6. Piecewise linear approximation of the detection efficiency for a given pixel. In this example, we have assumed that the exposure time is inversely proportional to the square root of the flux, valid for sky background dominated imaging.

a log-normal distance distribution with the same mean and standard deviation as the ansatz distribution.

We calculate the mean m and standard deviation s from μ and σ using the function `parameters_to_moments` from `ligo.skymap`⁴. Then, the location and scale parameters of the log-normal distribution are given by

$$\mu_{\ln r} = \ln m - \frac{1}{2} \ln \left(1 + \frac{s^2}{m^2} \right) \quad (18)$$

$$\sigma_{\ln r}^2 = \ln \left(1 + \frac{s^2}{m^2} \right). \quad (19)$$

The logarithm of the distance then has the distribution $\ln r \sim \mathcal{N}[\mu_{\ln r}, \sigma_{\ln r}]$. The apparent magnitude is related to the absolute magnitude through $x = X + 5 \log_{10} r + 25$, assuming that r is in the units of Mpc. Therefore the apparent magnitude has the distribution $x \sim \mathcal{N}[\mu_x, \sigma_x]$, with

$$\mu_x = \mu_X + \left(\frac{5}{\ln 10} \right) \mu_{\ln r} + 25 \quad (20)$$

$$\sigma_x^2 = \sigma_X^2 + \left(\frac{5}{\ln 10} \right)^2 \sigma_{\ln r}^2. \quad (21)$$

With this Gaussian distribution of apparent magnitudes, we can now calculate the detection efficiency for each pixel: the probability that we detect the source assuming that the source is in that pixel, as a function of exposure time. For the purpose of implementation of this function in a MILP, we approximate it with a piecewise linear function, as illustrated in Fig 6.

⁴ <https://lscsoft.docs.ligo.org/ligo.skymap/distance/>

3.3.1. Additional Data Preparation

1. Use the function `parameters_to_moments` and Eqs. (18, 20) to calculate the mean and standard deviation of the apparent magnitude in each pixel.
2. Select the desired quantiles for the piecewise linear approximation of the detection efficiency curve: for example, $(0, 0.05, 0.275, 0.5, 0.725, 0.95)$. For each pixel, calculate the exposure time required to achieve the specified detection efficiency.

3.3.2. Problem Setup

Additional index sets —

indices of quantiles $N = \{0, 1, \dots, n_N\}$

Additional parameters —

quantiles of detection efficiency $(\xi_n)_{n \in N}$

exposure time of quantiles $(\epsilon_{in})_{i \in I, n \in N}$

piecewise linear functions $(f_i : \mathbb{R}_{\geq 0} \rightarrow [0, 1])_{i \in I}$

Additional, continuous decision variables —

change from binary to continuous $(p_i)_{i \in I}$

3.3.3. Additional Constraints

Depth—Replace Eq. 9 with:

$$\forall i \in I : \max_{j \in J_i} e_j \geq f_i(p_i)$$

3.3.4. Objective

Same as in Section 3.1.5 above.

4. EXPOSURE TIME ESTIMATION

The reference UVEX ETC is still closed source because the instrument is still evolving in small details. We use M⁴OPT’s open-source toy model of the instrument performance that roughly reproduces the public sensitivity curve plot⁵ (see Fig. 7). The parameters of the toy model are listed in Table 2.

We need to perform about 1 million ETC calculations at different sky positions and apparent magnitudes for each event. The `synphot` Python package (STScI Development Team 2018) is the standard tool for synthetic photometry. Unfortunately, it would be prohibitively time-consuming to set up and evaluate 1 million `synphot` scenarios due to the

Table 2. ETC toy model parameters.

Quantity	Value
Aperture diameter	75 cm
Pixel scale	1'' pix ⁻¹
PSF sharpness ^a	1/(4π)
Gain	0.85
Read noise	2 ct
Dark noise	10 ⁻³ ct s ⁻¹
NUV response ^b	0.2√2πϕ[(λ − 2300 Å)/(180 Å)]
FUV response ^b	0.15√2πϕ[(λ − 1600 Å)/(100 Å)]

^aWe assume optimal PSF photometry with Nyquist pixel sampling, that is, a sharpness of 1/4π (Mighell 2005).

^bThe function $\phi(x)$ is the standard normal distribution.




Figure 7. Median limiting magnitude, averaged over target coordinates and observation time.

overhead of creating all of the Python objects involved. Our `m4opt.synphot` module accelerates synthetic photometry calculations by using the symbolic algebra package `sympy` (Meurer et al. 2017) to rearrange the `synphot` model tree into additive components with separable spatial dependence that can be integrated separately and then added back together. For nonlinear components (i.e., extinction), we do a parameter sweep and employ interpolation. The result is many orders of magnitude faster than `synphot` and contributes only about a second of run time to the scheduler.

There are three spatially-dependent spectral components in the ETC:

Zodiacal light—We modeled zodiacal background light by taking the “high” sky background spectrum from Hubble Space Telescope Imaging Spectrograph (STIS) Instrument

⁵ <https://www.uvex.caltech.edu/page/for-astronomers>; accessed 2025-01-26

Handbook (Rickman & Brown 2024) and scaling it by a spatially-dependent factor resulting from bilinear interpolation of from the tables of Leinert et al. (1998). In the future, we plan to replace the spatial dependence with `zodipy` (San et al. 2022; San 2024), which has a higher spatial resolution and is valid for locations throughout the Solar System beyond Earth orbit.

Galactic diffuse background—Murthy (2014) provides piecewise cosecant fits to the surface brightness of Galactic diffuse emission in the two GALEX bands. We employ these fits to get the spatial dependence, and obtain the wavelength dependence by interpolating and extrapolating linearly through the two bands.

Dust extinction—To model extinction due to dust in the Milky Way, we use the Gordon et al. (2023) model as implemented in the `dust_extinction` Python package (Gordon 2024). We employ a fixed total-to-selective extinction ratio $R(V) = 3.1$ and obtain $E(B - V)$ reddening values as a function of sky position using the Planck Collaboration et al. (2016) model as implemented by the `dustmaps` Python package (Green 2018).

5. CASE STUDY: GW OBSERVATIONS WITH UVEX

Here we explain the setup of M⁴OPT for UVEX.

GW localizations—We started with the same simulated GW localizations as Criswell et al. (2025), which covers LIGO, Virgo, and KAGRA’s fifth (O5) and sixth (O6) observing runs. The data are publicly archived in Kiendrebeogo (2025). These simulated events were generated using the same methodology as Petrov et al. (2022) and Kiendrebeogo et al. (2023), except that the signal to noise ratio (S/N) threshold for GW detection is set to 10. The localizations were generated with the rapid localization engine BAYESTAR (Singer & Price 2016) and consist of 3D posterior probability distributions of sky location and luminosity distance (Singer et al. 2016a,b). Like Criswell et al. (2025), we selected only those events for which the rest frame secondary (lighter) compact object mass m_2 was $\leq 3M_\odot$ to pick only events that could plausibly be binary neutron star (BNS) or neutron star–black hole (NS–BH) binaries.

KN absolute magnitude—Appendix E.2 of Kulkarni et al. (2021) specifies fiducial parameter ranges for nucleosynthesis-powered or shock-powered KN models and 90% credible intervals for the absolute magnitude in each band. These absolute magnitude ranges are reproduced in the Table 3. UVEX observes in both the NUV and FUV filters simultaneously, and our objective is to detect the source in both bands. Therefore we should plan observations using the fainter of the two models and the fainter of the two bands: the nucleosynthesis-powered model in FUV, with an

Table 3. Ranges of peak absolute magnitudes of KNe. Adapted from Appendix E.2 of Kulkarni et al. (2021).

Model	Absolute magnitude range	
	NUV	FUV
Nucleosynthesis powered	[-15.6, -12.4]	[-14.5, -10.2]
Shock powered	[-17.8, -15.3]	[-17.9, -15.0]

absolute magnitude range of $[-14.5, -10.2]$. Assuming that this is the 90% credible interval of a Gaussian distribution, the absolute magnitude has the approximate distribution

$$M_{\text{NUV}} \sim \mathcal{N}(-12.4, 1.3). \quad (22)$$

Follow-up time window—Criswell et al. (2025) required a single epoch of UVEX observations to take 3 hours or less. To match this choice, we configure M⁴OPT to plan two visits of each field with a minimum cadence of 30 min between repeated visits, with a total elapsed time limit of 6 hours.

Exposure time limits—The exposure time is allowed to vary adaptively for each field, with a minimum exposure time of $\epsilon_{\text{min}} = 300$ s. The minimum exposure time corresponds to a single standard UVEX imaging exposure. (A standard survey dwell consists of 3 consecutive stacked 300 s exposures.)

FOV—Like most space telescopes, UVEX has solar panels that rotate on a solar array drive assembly (SADA) perpendicular to the telescope boresight (see Fig. 1). The position angle of observations is fixed to the nominal roll angle that allows the spacecraft to orient the solar panels perpendicular to the Sun while keeping the cold side of the spacecraft facing away from the Sun. For any given target position, the roll angle goes through one revolution per year. For targets at the ecliptic poles, the nominal roll angle varies linearly with time. For targets in the ecliptic plane, the nominal roll angle flips by 180° at the solstices. At all intermediate ecliptic latitudes, the roll angle oscillates smoothly in a manner that interpolates between these extremes (see Fig. 8). Because the roll angle changes slowly on the timescale of a day, we calculate the nominal roll angle for each field at the time of the event, and leave it fixed at that value for the duration of the observing plan.

Slew time—The optimal “bang-bang” slew trajectory between any two points consists of an acceleration phase at the maximum acceleration, possibly a coasting phase at the maximum angular velocity, a deceleration phase at the maximum acceleration, and settling (see Fig. 9). We employ a simplified slew model in which the spacecraft has a fixed maximum




Figure 8. Nominal roll angle as a function of time for selected ecliptic latitudes.




Figure 9. The optimal slew consists of an acceleration phase at the maximum acceleration, possibly a coasting phase at the maximum angular velocity, and a deceleration phase at the maximum acceleration.

angular acceleration of $\alpha = 0.006 \text{ s}^{-2}$, maximum angular rate of 0.6 s^{-1} , and settling time of 60 s, irrespective of axis. With these parameters, it takes 108 s to slew by the width of the FOV and 460 s to slew by 180° .




Figure 10. The footprint of the UVEX FOV discretized on a HEALPix grid with $n_{\text{side}} = 128$.

HEALPix resolution—We discretize the footprint of the FOV on a HEALPix grid with $n_{\text{side}} = 128$ (see Fig. 10).

Field grid—The centers of the fields are the vertices of a $\{3, 5+\}_{21,4}$ icosahedral geodesic polyhedron (see Fig. 11) generated using the `antiprism` software by Adrian Rossiter⁶. This grid ensures that all of the fields cover the sky without gaps, regardless of roll angle.

Run duration—As in Criswell et al. (2025), we assumed 1.5 years of overlap between the UVEX prime mission and the GW observing run.

Follow-up selection criterion—We ran `M4OPT` on all simulated events. We considered an event selected for follow-up if the scheduler’s objective value P was less than $P^* = 0.1$. Recall from Section 3.3 that the scheduler’s objective value is (a numerical approximation of) the probability of detection—integrated over the absolute magnitude, sky position, and distance, none of which are known to the scheduler. We stress that our strategy for selecting which events to trigger is to simply run the scheduler for every single event and proceed if the predicted detection probability is at least 10%.

This is conceptually very different from the selection criteria in Criswell et al. (2025) which is based on the 90% credible area, $A_{90\%}$, and luminosity distance d_L of the event. However, we can predict analytical thresholds on both of these quantities that will be roughly equivalent to our strat-

⁶ <https://github.com/antiprism/antiprism>




Figure 11. Footprints of the UVEX FOV on the geodesic reference grid. This figure is available as an animation that shows the variation in position angle of each field with time.

egy:

$$d_L < d_L^* = 10^{\frac{1}{5}(x^* - \mu_X + \sigma_X \Phi^{-1}(1-P^*) - 25)} \text{ Mpc} \quad (23)$$

$$A_Q < A_Q^* = \left(\frac{\Psi^{-1}(Q)}{\Psi^{-1}(P^*)} \right)^2 \left(\frac{\delta - \beta}{\epsilon_{\min} n_K} \right) A_{\text{FOV}} \quad (24)$$

$$\frac{A_Q}{A_{\text{FOV}}} < \left(\frac{d_L}{d_L^*} \right)^{-4} \quad (25)$$

A_Q is the Q th percentile credible region, x^* is the faintest limiting magnitude at any point on the sky, A_{FOV} is the area of the FOV, $\Phi^{-1}(x)$ is the inverse of the cumulative distribution function (CDF) of the standard normal distribution, and $\Psi^{-1}(x) = -2 \ln(1 - x)$ is the inverse of the CDF of a χ^2 distribution with 2 degrees of freedom.

5.1. Results

All simulated events are listed in the online version of Table 4. Because these are simulated events, we can determine the probability of detection for each event given its true sky location and distance. To calculate the detection probability, we calculate the limiting absolute magnitude at that true position and distance for the longest exposure that contains that position (or an exposure time of zero if the true sky position is not contained in any planned observation). The detection probability is simply the CDF of the absolute magnitude distribution, Eq. (22), evaluated at that limiting absolute magnitude.

In Figs. 12 and 13, we show both the scheduler objective value and the detection probability on a scatter plot of

the 90% credible area and the luminosity distance of events. As we would expect, both the objective value and detection probability increase as the area and distance of the events decrease.

In Table 5, we list the expected numbers of events selected for follow-up and the number of KNe detected over 1.5 years of observation. The number of events selected is simply the number of observing plans with objective value ≥ 0.1 . The expected number of KNe detected is the sum of the detection probabilities of all of the events.

As in Petrov et al. (2022) and in Kiendrebeogo et al. (2023), this table gives central 90% credible intervals about the median, incorporating a lognormal uncertainty in the BNS merger rate of $210_{-120}^{240} \text{ Gpc}^{-3} \text{ yr}^{-1}$ and Poissonian variation in the number of events over a finite time duration.

If instead of optimizing for a detection in both filters, we optimize for a detection in *at least one filter*, then both the number of events selected and the number of events detected approximately double.

6. CONCLUSION

Unfortunately, direct comparisons of the rates in Table 5 to Criswell et al. (2025), which used dorado-scheduling and a fixed exposure time for each event, are not meaningful because we made many refinements to the starting assumptions. The following change would lead to a higher detection rate:

- We used a slightly brighter distribution of absolute magnitudes that is consistent with the stated KN model ranges in Kulkarni et al. (2021).

The following change would lead to a lower detection rate:

- We used a much denser and more heavily overlapping grid of fields in order to ensure complete coverage of the sky at all roll angles.

Our trigger and detection rates in Table 5 are comparable with the most optimistic estimates in Criswell et al. (2025). Although by itself this comparison is not informative on the merits of these two strategies, it is a positive result for the prospects of KN detection with UVEX.

More quantitatively, we can show that the dynamic exposure time strategy results in a higher detection probability than any given fixed exposure time for a given event. Fig. 14 shows a typical example. This is a plot of detection probability (without knowledge of the position of the source) versus exposure time for observing plans for a single event generated using the formulation from Section 3.1. The detection probability for the dynamic exposure time strategy from Section 3.3 is shown as a horizontal dashed line.

Table 4. Simulated events.

Run	Event ID	Source frame masses		True position			$A_{90\%} / \text{deg}^2$	Objective value	Detection prob.
		$m_1/M_\odot$	$m_2/M_\odot$	α/deg	δ/deg	d_L/Mpc			
O5	14	1.976	1.597	217.5627	+50.1543	204	3	0.95	0.92
O5	27	1.605	1.477	9.8234	+23.2808	685	571	0.00	0.00
O5	30	1.318	1.280	337.1588	-36.0109	253	192	0.16	0.00
O5	41	1.979	1.772	194.5206	-19.3125	488	129	0.16	0.32
O5	46	6.405	2.379	139.4866	-25.4656	1277	766	0.00	0.00
O5	52	6.754	2.218	335.1421	-3.1377	1457	90	0.00	0.00
O5	54	1.961	1.583	131.7071	-38.2637	679	1143	0.00	0.00
O5	60	25.417	2.561	91.4538	+78.0966	800	4579	0.00	0.00
...	...	...	...	...	...	...	...	...	...

NOTE—Table 4 is published in its entirety in a machine-readable format. A portion is shown here for guidance regarding its form and content.




Figure 12. The distribution events selected for follow-up and detected for O5. **Bottom left panel:** scatter plot of 90% credible area versus distance. Events that were selected for follow-up with UVEX are represented by colored dots. The color of the dot represents the detection probability and the area of the dot represents the scheduler objective value. Events that were not selected for follow-up are marked with gray dots. The blue boundary represents the analytical predictor of the detection threshold given by Eqs. (23, 24, 25). **Upper panel:** distribution of luminosity distance, with 90th percentiles marked. **Right panel:** distribution of 90% credible area.




Figure 13. Same as Fig. 12, but for O6.

Table 5. Expected number of events.

	O5	O6
Number of events selected	29^{+39}_{-18}	43^{+56}_{-26}
Number of events detected	12^{+18}_{-9}	17^{+24}_{-11}




Figure 14. Detection probability versus exposure time for a single example event, resulting from the fixed exposure time strategy (see Section 3.1). The horizontal dashed line shows the detection probability resulting from the variable exposure time strategy incorporating a distribution of absolute magnitudes (see Section 3.3). The variable exposure time strategy is more likely to detect the counterpart than the fixed exposure time strategy, for any exposure time.

This novel dynamic exposure time observing strategy is promising for GW follow-up with any large FOV imaging instrument. Although we have focused on UVEX, we encourage ULTRASAT to try it out using the preliminary support for that mission that we have added to M⁴OPT.

The M⁴OPT model itself is not specific to space-based observations. With additional FOR and foreground models it could be extended to ground-based telescopes or even heterogeneous combinations of ground and space telescopes. In the future, we plan to extend it to ZTF, Vera Rubin Observatory, and other ground-based facilities. Ultimately, our vision is for M⁴OPT to evolve into a scheduling toolkit featuring (1) a composable library of observing constraints inspired by the interface of the `astrowplan` package (Morris et al. 2018), (2) uniform support for observers on the ground and in space, and (3) a globally optimizing scheduler.

Our probabilistic approach to scheduling and triggering follow-up observations can in principle allow for a more detailed treatment of the selection effects incurred during GW follow-up campaigns and astrophysical inference from the population of KNe observed by missions like UVEX. De-

veloping a framework for KN inference with UVEX observations triggered by this method is an exciting direction of future work.

Some practical implementation details remain as future work to incorporate M⁴OPT into the operations of a real observatory. For example:

- Our slow time constraints must conservatively model the dynamics of the actual spacecraft, including differing angular rates about different axes and keep-out constraints that must hold through the entire slew trajectory.
- We must be able to retrieve the spacecraft ephemeris from a live data source rather than evaluating it from a predefined orbital model.
- We must incorporate pre-defined segments of observing dead time for housekeeping events such as pre-planned ground contacts and momentum dumping maneuvers.

Like the ZTF scheduler (Bellm et al. 2019) and the Vera Rubin Observatory scheduler (Naghib et al. 2019), we see M⁴OPT as both a core component of the ground software system that enables fully autonomous observation planning and execution, and also an approachable and well-documented part of the guest observer science tools for a mission. Its permissive open-source licensing and open development model, free from the encumbrances of NOSA, are essential to those goals.

We hope to build from it a community tool that serves many missions from the concept phase all the way through operations. Also, integration of M⁴OPT with community follow-up platforms such as the Gravitational Wave Treasure Map (Wyatt et al. 2020) could open possibilities for more effective and more meaningfully coordinated follow-up observations by multiple independently-operated facilities.

1 This work was performed in part at the Aspen Center for
2 Physics, which is supported by the U.S. National Science
3 Foundation (NSF) grant PHY-2210452.

4 This work used Expanse at the San Diego Supercom-
5 puting Center (SDSC) and Delta at the National Center
6 for Supercomputing Applications (NCSA) through alloca-
7 tion AST200029, “Towards a complete catalog of vari-
8 able sources to support efficient searches for compact bi-
9 nary mergers and their products,” from the Advanced Cy-
10 berinfrastructure Coordination Ecosystem: Services & Sup-
11 port (ACCESS) program, which is supported by NSF grants
12 #2138259, #2138286, #2138307, #2137603, and #2138296.

13 SCL and MWC acknowledge support from NSF with grant
14 numbers PHY-2308862 and PHY-2117997.

15 We thank Steve Crawford at NASA Headquarters for help
16 and support with the NASA software release process.

17 All of the code and data to reproduce the figures and tables
18 in this paper are available from Zenodo (Singer 2025).

19 This is LIGO document P2500008-v1.

Software: `astropy` (Astropy Collaboration et al. 2013, 2018), `astroquery` (Ginsburg et al. 2019), `dust_extinction` (Gordon 2024), `dustmaps` (Green 2018), `healpix` (Górski et al. 2005), `healpy` (Zonca et al. 2019), `ligo.skymap` (Singer & Price 2016; Singer et al. 2016a,b), `matplotlib` (Hunter 2007), `numpy` (Harris et al. 2020), `regions` (Bradley et al. 2022), `scipy` (Virtanen et al. 2020), `spiceypy` (Annex et al. 2020), `sympy` (Meurer et al. 2017), `synphot` (STScI Development Team 2018)

REFERENCES

- Aasi, J., Abadie, J., Abbott, B. P., et al. 2014, *ApJS*, 211, 7, doi: [10.1088/0067-0049/211/1/7](https://doi.org/10.1088/0067-0049/211/1/7)
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *ApJL*, 826, L13, doi: [10.3847/2041-8205/826/1/L13](https://doi.org/10.3847/2041-8205/826/1/L13)
- . 2017a, *PhRvL*, 119, 161101, doi: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101)
- . 2017b, *ApJL*, 848, L12, doi: [10.3847/2041-8213/aa91c9](https://doi.org/10.3847/2041-8213/aa91c9)
- . 2017c, *Nature*, 551, 85, doi: [10.1038/nature24471](https://doi.org/10.1038/nature24471)
- Annex, A., Pearson, B., Seignovet, B., et al. 2020, *JOSS*, 5, 2050, doi: [10.21105/joss.02050](https://doi.org/10.21105/joss.02050)
- Arcavi, I. 2018, *ApJL*, 855, L23, doi: [10.3847/2041-8213/aab267](https://doi.org/10.3847/2041-8213/aab267)
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019, *PASP*, 131, 068003, doi: [10.1088/1538-3873/ab0c2a](https://doi.org/10.1088/1538-3873/ab0c2a)
- Bradley, L., Deil, C., Patra, S., et al. 2022, *astropy/regions*: v0.6, Zenodo, doi: [10.5281/zenodo.6374572](https://doi.org/10.5281/zenodo.6374572)
- Cenko, S. B. 2019, in *American Astronomical Society Meeting Abstracts*, Vol. 234, American Astronomical Society Meeting Abstracts #234, 212.03
- Chan, M. L., Hu, Y.-M., Messenger, C., Hendry, M., & Heng, I. S. 2017, *ApJ*, 834, 84, doi: [10.3847/1538-4357/834/1/84](https://doi.org/10.3847/1538-4357/834/1/84)
- Chen, D.-S., Batson, R. G., & Dang, Y. 2009, *Transformation Using 0–1 Variables* (John Wiley & Sons, Ltd), 54–78, doi: <https://doi.org/10.1002/9781118166000.ch3>
- Coughlin, M., & Stubbs, C. 2016, *Experimental Astronomy*, 42, 165, doi: [10.1007/s10686-016-9503-4](https://doi.org/10.1007/s10686-016-9503-4)
- Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, *Science*, 358, 1556, doi: [10.1126/science.aap9811](https://doi.org/10.1126/science.aap9811)
- Criswell, A. W., Leggio, S. C., Coughlin, M. W., et al. 2025, *arXiv e-prints*, arXiv:2501.14109. <https://arxiv.org/abs/2501.14109>
- Dorsman, B., Raaijmakers, G., Cenko, S. B., et al. 2023, *ApJ*, 944, 126, doi: [10.3847/1538-4357/aca9e](https://doi.org/10.3847/1538-4357/aca9e)
- Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, *Nature*, 340, 126, doi: [10.1038/340126a0](https://doi.org/10.1038/340126a0)
- Fong, W., & Berger, E. 2013, *ApJ*, 776, 18, doi: [10.1088/0004-637X/776/1/18](https://doi.org/10.1088/0004-637X/776/1/18)
- Free Software Foundation. 2006, Various Licenses and Comments about Them. <https://www.gnu.org/licenses/license-list.html#NASA>
- Freiburghaus, C., Rosswog, S., & Thielemann, F. K. 1999, *ApJL*, 525, L121, doi: [10.1086/312343](https://doi.org/10.1086/312343)
- Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, 157, 98, doi: [10.3847/1538-3881/aafc33](https://doi.org/10.3847/1538-3881/aafc33)
- Goldstein, A., Veres, P., Burns, E., et al. 2017, *ApJL*, 848, L14, doi: [10.3847/2041-8213/aa8f41](https://doi.org/10.3847/2041-8213/aa8f41)
- Gordon, K. 2024, *JOSS*, 9, 7023, doi: [10.21105/joss.07023](https://doi.org/10.21105/joss.07023)
- Gordon, K. D., Clayton, G. C., Decleir, M., et al. 2023, *ApJ*, 950, 86, doi: [10.3847/1538-4357/accb59](https://doi.org/10.3847/1538-4357/accb59)
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759, doi: [10.1086/427976](https://doi.org/10.1086/427976)
- Gottlieb, O., Nakar, E., Piran, T., & Hotokezaka, K. 2018, *MNRAS*, 479, 588, doi: [10.1093/mnras/sty1462](https://doi.org/10.1093/mnras/sty1462)

- Green, G. M. 2018, *JOSS*, 3, 695, doi: [10.21105/joss.00695](https://doi.org/10.21105/joss.00695)
- Handley, L. B., Petigura, E. A., & Mišić, V. V. 2024, *AJ*, 167, 33, doi: [10.3847/1538-3881/ad0dfb](https://doi.org/10.3847/1538-3881/ad0dfb)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Herner, K., Annis, J., Brout, D., et al. 2020, *A&C*, 33, 100425, doi: [10.1016/j.ascom.2020.100425](https://doi.org/10.1016/j.ascom.2020.100425)
- Huangfu, Q., & Hall, J. J. 2018, *Mathematical Programming Computation*, 10, 119
- Hunter, J. D. 2007, *CSE*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Kasen, D., Badnell, N. R., & Barnes, J. 2013, *ApJ*, 774, 25, doi: [10.1088/0004-637X/774/1/25](https://doi.org/10.1088/0004-637X/774/1/25)
- Kasen, D., Metzger, B., Barnes, J., Quataert, E., & Ramirez-Ruiz, E. 2017, *Nature*, 551, 80, doi: [10.1038/nature24453](https://doi.org/10.1038/nature24453)
- Kasliwal, M. M., Nakar, E., Singer, L. P., et al. 2017, *Science*, 358, 1559, doi: [10.1126/science.aap9455](https://doi.org/10.1126/science.aap9455)
- Kiendrebeogo, R. W. 2025, Simulation of HLVK Configuration for O5 and O6 Observation Runs: Gravitational Wave Detection Rates, Skymaps, and Multimessenger Parameters (November 2024 edition), Zenodo, doi: [10.5281/zenodo.14585837](https://doi.org/10.5281/zenodo.14585837)
- Kiendrebeogo, R. W., Farah, A. M., Foley, E. M., et al. 2023, *ApJ*, 958, 158, doi: [10.3847/1538-4357/acfcb1](https://doi.org/10.3847/1538-4357/acfcb1)
- Kulkarni, S. R., Harrison, F. A., Grefenstette, B. W., et al. 2021, arXiv e-prints, arXiv:2111.15608, doi: [10.48550/arXiv.2111.15608](https://doi.org/10.48550/arXiv.2111.15608)
- Lattimer, J. M., & Schramm, D. N. 1974, *ApJL*, 192, L145, doi: [10.1086/181612](https://doi.org/10.1086/181612)
- Leinert, C., Bowyer, S., Haikala, L. K., et al. 1998, *A&AS*, 127, 1, doi: [10.1051/aas:1998105](https://doi.org/10.1051/aas:1998105)
- Li, L.-X., & Paczyński, B. 1998, *ApJL*, 507, L59, doi: [10.1086/311680](https://doi.org/10.1086/311680)
- Liu, M.-X., Tong, H., Hu, Y.-M., et al. 2021, *Research in Astronomy and Astrophysics*, 21, 308, doi: [10.1088/1674-4527/21/12/308](https://doi.org/10.1088/1674-4527/21/12/308)
- Metzger, B. D., & Berger, E. 2012, *ApJ*, 746, 48, doi: [10.1088/0004-637X/746/1/48](https://doi.org/10.1088/0004-637X/746/1/48)
- Meurer, A., Smith, C. P., Paprocki, M., et al. 2017, *PeerJ Computer Science*, 3, e103, doi: [10.7717/peerj-cs.103](https://doi.org/10.7717/peerj-cs.103)
- Mighell, K. J. 2005, *MNRAS*, 361, 861, doi: [10.1111/j.1365-2966.2005.09208.x](https://doi.org/10.1111/j.1365-2966.2005.09208.x)
- Morris, B. M., Tollerud, E., Sipőcz, B., et al. 2018, *AJ*, 155, 128, doi: [10.3847/1538-3881/aaa47e](https://doi.org/10.3847/1538-3881/aaa47e)
- Murthy, J. 2014, *ApJS*, 213, 32, doi: [10.1088/0067-0049/213/2/32](https://doi.org/10.1088/0067-0049/213/2/32)
- Naghib, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, *AJ*, 157, 151, doi: [10.3847/1538-3881/aafce](https://doi.org/10.3847/1538-3881/aafce)
- NASA. 2022, NASA Software Engineering Requirements, Tech. Rep. NPR 7150.2D, NASA Office of the Chief Engineer. <https://nodis3.gsfc.nasa.gov/displayDir.cfm?t=NPR&c=7150&s=2D>
- NASA Science Mission Directorate. 2022, SMD Policy Document SPD-41a. <https://science.nasa.gov/spd-41>
- National Academies of Sciences, Engineering, and Medicine. 2018, *Open Source Software Policy Options for NASA Earth and Space Sciences* (Washington, DC: The National Academies Press), doi: [10.17226/25217](https://doi.org/10.17226/25217)
- . 2021, *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*, doi: [10.17226/26141](https://doi.org/10.17226/26141)
- Nicholl, M., Berger, E., Kasen, D., et al. 2017, *ApJL*, 848, L18, doi: [10.3847/2041-8213/aa9029](https://doi.org/10.3847/2041-8213/aa9029)
- Parazin, B., Coughlin, M. W., Singer, L. P., Gupta, V., & Anand, S. 2022, *ApJ*, 935, 87, doi: [10.3847/1538-4357/ac7fa2](https://doi.org/10.3847/1538-4357/ac7fa2)
- Petrov, P., Singer, L. P., Coughlin, M. W., et al. 2022, *ApJ*, 924, 54, doi: [10.3847/1538-4357/ac366d](https://doi.org/10.3847/1538-4357/ac366d)
- Pian, E., D'Avanzo, P., Benetti, S., et al. 2017, *Nature*, 551, 67, doi: [10.1038/nature24298](https://doi.org/10.1038/nature24298)
- Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2016, *A&A*, 596, A109, doi: [10.1051/0004-6361/201629022](https://doi.org/10.1051/0004-6361/201629022)
- Rickman, E., & Brown, J. 2024, *Space Telescope Imaging Spectrograph Instrument Handbook for Cycle 33, Hubble Space Telescope User Documentation, STIS Instrument Handbook, Version 24.0*, (Baltimore: STScI)
- Sagiv, I., Gal-Yam, A., Ofek, E. O., et al. 2014, *AJ*, 147, 79, doi: [10.1088/0004-6256/147/4/79](https://doi.org/10.1088/0004-6256/147/4/79)
- San, M. 2024, *JOSS*, 9, 6648, doi: [10.21105/joss.06648](https://doi.org/10.21105/joss.06648)
- San, M., Herman, D., Erikstad, G. B., Galloway, M., & Watts, D. 2022, *A&A*, 666, A107, doi: [10.1051/0004-6361/202244133](https://doi.org/10.1051/0004-6361/202244133)
- Saunders, E. S., Lampoudi, S., Lister, T. A., Norbury, M., & Walker, Z. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9149, *Observatory Operations: Strategies, Processes, and Systems V*, ed. A. B. Peck, C. R. Benn, & R. L. Seaman, 91490E, doi: [10.1117/12.2056642](https://doi.org/10.1117/12.2056642)
- Savchenko, V., Ferrigno, C., Kuulkers, E., et al. 2017, *ApJL*, 848, L15, doi: [10.3847/2041-8213/aa8f94](https://doi.org/10.3847/2041-8213/aa8f94)
- Schutz, B. F. 1986, *Nature*, 323, 310, doi: [10.1038/323310a0](https://doi.org/10.1038/323310a0)
- Shappee, B. J., Simon, J. D., Drout, M. R., et al. 2017, *Science*, 358, 1574, doi: [10.1126/science.aaq0186](https://doi.org/10.1126/science.aaq0186)
- Shvartzvald, Y., Waxman, E., Gal-Yam, A., et al. 2024, *ApJ*, 964, 74, doi: [10.3847/1538-4357/ad2704](https://doi.org/10.3847/1538-4357/ad2704)
- Singer, L. 2025, Data Release: Optimal Follow-Up of Gravitational- Wave Events with the UltraViolet EXplorer (UVEX), Zenodo, doi: [10.5281/zenodo.14825979](https://doi.org/10.5281/zenodo.14825979)
- Singer, L. P., & Price, L. R. 2016, *PhRvD*, 93, 024013, doi: [10.1103/PhysRevD.93.024013](https://doi.org/10.1103/PhysRevD.93.024013)
- Singer, L. P., Chen, H.-Y., Holz, D. E., et al. 2016a, *ApJL*, 829, L15, doi: [10.3847/2041-8205/829/1/L15](https://doi.org/10.3847/2041-8205/829/1/L15)
- . 2016b, *ApJS*, 226, 10, doi: [10.3847/0067-0049/226/1/10](https://doi.org/10.3847/0067-0049/226/1/10)
- Solar, M., Michelon, P., Avarias, J., & Garcés, M. 2016, *A&C*, 15, 90, doi: [10.1016/j.ascom.2016.02.005](https://doi.org/10.1016/j.ascom.2016.02.005)

- STScI Development Team. 2018, synphot: Synthetic photometry using Astropy, Astrophysics Source Code Library, record ascl:1811.001
- Villar, V. A., Guillochon, J., Berger, E., et al. 2017, ApJL, 851, L21, doi: [10.3847/2041-8213/aa9c84](https://doi.org/10.3847/2041-8213/aa9c84)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Williams, H. 2013, Building integer programming models I (Wiley), 165–201. <https://learning.oreilly.com/library/view/model-building-in/9781118506172/>
- Wyatt, S. D., Tohuvavohu, A., Arcavi, I., et al. 2020, ApJ, 894, 127, doi: [10.3847/1538-4357/ab855e](https://doi.org/10.3847/1538-4357/ab855e)
- Zonca, A., Singer, L., Lenz, D., et al. 2019, JOSS, 4, 1298, doi: [10.21105/joss.01298](https://doi.org/10.21105/joss.01298)