

Optimizing searches for electromagnetic counterparts of gravitational wave triggers

Michael W. Coughlin,^{1★} Duo Tao,² Man Leong Chan,³ Deep Chatterjee,⁴
Nelson Christensen,^{2,5} Shaon Ghosh,⁴ Giuseppe Greco,^{6,7} Yiming Hu,⁸
Shasvath Kapadia,⁴ Javed Rana,⁹ Om Sharan Salafia¹⁰ and Christopher W. Stubbs^{11,12}

¹*Division of Physics, Math, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*

²*Carleton College, Northfield, MN 55057, USA*

³*University of Glasgow, Glasgow G12 8QQ, UK*

⁴*University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA*

⁵*Artemis, Université Côte d'Azur, Observatoire Côte d'Azur, CNRS, CS 34229, F-06304 Nice Cedex 4, France*

⁶*INFN, Sezione di Firenze, I-50019 Sesto Fiorentino, Firenze, Italy*

⁷*Università degli Studi di Urbino 'Carlo Bo', I-61029 Urbino, Italy*

⁸*TianQin Research Center for Gravitational Physics, Sun Yat-sen University, Tangjiawan, Zhuhai 519082, Guangdong, P. R. China*

⁹*Inter-University Centre for Astronomy and Astrophysics, Pune 411007, India*

¹⁰*INAF – Osservatorio Astronomico di Brera Merate, via E. Bianchi 46, I-23807 Merate, Italy*

¹¹*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

¹²*Department of Astronomy, Harvard University, Cambridge, MA 02138, USA*

Accepted 2018 April 24. Received 2018 April 23; in original form 2018 March 6

ABSTRACT

With the detection of a binary neutron star system and its corresponding electromagnetic counterparts, a new window of transient astronomy has opened. Due to the size of the sky localization regions, which can span hundreds to thousands of square degrees, there are significant benefits to optimizing tilings for these large sky areas. The rich science promised by gravitational wave astronomy has led to the proposal for a variety of proposed tiling and time allocation schemes, and for the first time, we make a systematic comparison of some of these methods. We find that differences of a factor of 2 or more in efficiency are possible, depending on the algorithm employed. For this reason, with future surveys searching for electromagnetic counterparts, care should be taken when selecting tiling, time allocation, and scheduling algorithms to optimize counterpart detection.

Key words: gravitational waves.

1 INTRODUCTION

The era of multimessenger gravitational wave astronomy has arrived with the detection of GW170817 (Abbott et al. 2017b) by Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acernese et al. 2015) coincident with the detection of both a short gamma-ray burst (SGRB) (Abbott et al. 2017d; Goldstein et al. 2017; Savchenko et al. 2017) and a kilonova detected in coincidence (Abbott et al. 2017e; Coulter et al. 2017; Smartt et al. 2017). This work is the culmination of significant effort expended in the search for the electromagnetic counterpart of the gravitational waves found by compact binary black hole systems (Abbott et al. 2016b,c, 2017a) (see Abbott et al. 2016d for an overview of the search for an electromagnetic counterpart to GW150914).

It has been known for some time that there are potential electromagnetic counterparts to binary neutron star and black hole – neutron star systems across durations and wavelengths (Nakar 2007; Metzger & Berger 2012). For example, a kilonova, arising from subrelativistic ejecta, in particular has predicted bolometric luminosities of $\approx 10^{40} - 10^{42} \text{ ergs s}^{-1}$ (Metzger et al. 2015; Barnes & Kasen 2013) (GW170817 peaked at $\approx 10^{42} \text{ ergs s}^{-1}$; Smartt et al. 2017) and optical and near-infrared colours and durations that depend on the physical conditions of the merger (Metzger et al. 2010; Barnes & Kasen 2013; Kasen, Badnell & Barnes 2013; Tanaka & Hotokezaka 2013; Kasen, Fernandez & Metzger 2015; Barnes et al. 2016; Metzger 2017).

The scientific output from a joint gravitational wave and electromagnetic observation is significant, as the detection of a kilonova coincident with a gravitational wave observation allows for the exploration of the neutron star equation of state (Bauswein, Baumgarte & Janka 2013) and r-process nucleosynthesis in the unbound ejecta

* E-mail: mcoughli@caltech.edu

from a merger involving a neutron star (Just et al. 2015; Metzger et al. 2015; Wu et al. 2016; Roberts et al. 2017). The gravitational wave posteriors alone, including mass and tidal parameter information, allow for estimates of the contribution of dynamical ejecta to the optical counterpart (Abbott et al. 2017f). Among others, Smartt et al. (2017) use photometry of GW170817 to place constraints on the ejecta mass, velocity, and effective opacity. Radice et al. (2018) use photometry of GW170817 in conjunction with kilonova models and numerical relativity results to place a lower bound on the tidal deformability parameter. The identification of the host galaxy allows for a distance-ladder independent measurement of the Hubble Constant (Abbott et al. 2017c). In addition, the joint observation with an SGRB not only confirms that these phenomena are driven by compact binary mergers, but also allows for the study of their beaming, energetics, and galactic environment (Metzger & Berger 2012).

To facilitate the detection of gravitational wave counterparts, probability skymaps as a function of sky direction and distance are released for gravitational wave triggers produced by the LIGO and Virgo detectors (Singer et al. 2014; Berry et al. 2015; Singer et al. 2016). Due to the significant sky coverage required to observe the gravitational wave sky localization regions, usually spanning $\approx 100 \text{ deg}^2$, techniques to optimize the follow-up efforts are of significant utility (Fairhurst 2009, 2011; Wen & Chen 2010; Grover et al. 2014; Sidery et al. 2014; Singer et al. 2014; Berry et al. 2015; Cornish & Littenberg 2015; Essick et al. 2015; Klimentko et al. 2016). Given the large sky localization regions involved, wide-field survey telescopes have the best opportunities to make a detection. The Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) (Morgan et al. 2012), Asteroid Terrestrial-impact Last Alert System (ATLAS) (Tonry 2011), the intermediate Palomar Transient Factory (PTF) (Rau et al. 2009) and (what will become) the Zwicky Transient Facility (ZTF), BlackGEM (Bloemen et al. 2015), and the Large Synoptic Survey Telescope (LSST) (Ivezic et al. 2008) are all examples of such systems. For example, Pan-STARRS has a 7 deg^2 field of view (FOV), achieving a 5σ limit of 21.5 (AB mag) in the i band in a 45 s exposure. ATLAS has a 29.2 deg^2 FOV, achieving a 5σ limit of 18.7 in the cyan band in a 30 s exposure. For comparison, LSST will have a 9.6 deg^2 FOV and will require a 21 s r -band exposure length to reach 22 mag.

Due to the significant difference in telescope configurations, including FOV, filter, typical exposure times, and limiting magnitudes, in addition to placement on the earth and therefore different seeing and sky conditions, optimizing gravitational wave follow-ups for an arbitrary telescope is difficult. In the following, we will take the telescopes mentioned above as examples. For this reason, we have created a codebase named *gwemopt* (Gravitational Wave – ElectroMagnetic OPTimization) that utilizes methods from a variety of recent papers geared towards optimizing efforts of follow-up. We employ methods to read gravitational wave skymaps and the associated information made available from GraceDB,¹ which is the gravitational wave data base from which one can access information about gravitational wave trigger candidates (Abbott et al. 2016d). We also use information about the telescopes to tile the sky, allocate available telescope time to the chosen tiles, and schedule the telescope time. This is done in a way that is optimized based on the telescope configurations and the light curves, i.e. the time evolution of luminosity, of the transients they are expected to detect. We

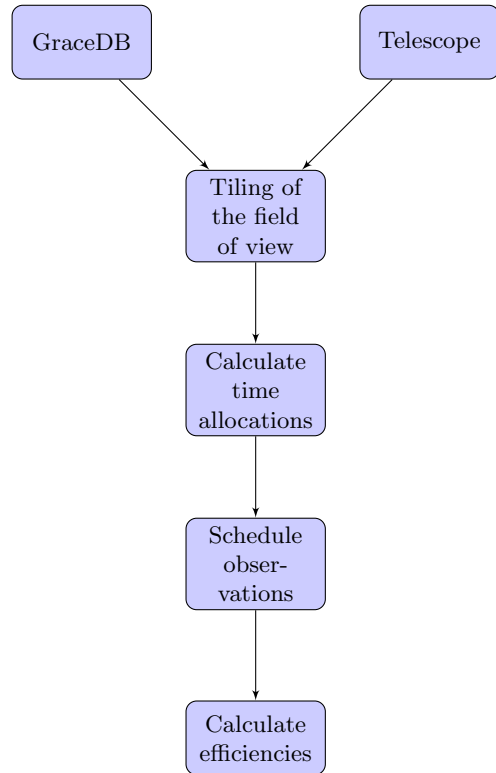


Figure 1. A flow chart of the *gwemopt* pipeline.

will describe algorithms that use the gravitational wave probability skymaps, some of them with right ascension and declination information only and some of them with distance information as well, to perform optimizations. In Section 2, we describe the algorithm. In Section 3, we describe the performance of the algorithms. In Section 4, we offer concluding remarks and suggest directions for future research.

2 ALGORITHM

Fig. 1 shows the flow chart for the *gwemopt* pipeline, developed to optimize the efforts of electromagnetic follow-up of gravitational wave events. *gwemopt* is developed in `python`, which has the benefit of an interface to LIGO and Virgo’s gravitational wave candidate event data base (GraceDB). It internally uses HEALPIX (Hierarchical Equal Area isoLatitude Pixelization) (Górski et al. 2005), the format in which LIGO and Virgo report skymaps, when performing optimization calculations.

The general sequence of the pipeline is as follows. *gwemopt* uses events provided by GraceDB in addition to information about the telescopes for creating tiles and optimizing time allocations in the fields. It uses information about potential light curves from electromagnetic counterparts to schedule the available telescope time. In the following, we will describe the calculations that go into creating tiles, time allocations, and observing sequences from the skymaps. We will account for both diurnal and observational constraints and have the possibility of imaging over many nights.

We will show the command line syntaxes required to reproduce the results at the beginning of each section. By way of an outline of the algorithm to be discussed, the subsections in this section are as follows:

¹<https://gracedb.ligo.org>




Figure 2. The gravitational wave posterior probability distribution $L_{\text{GW}}(\alpha, \delta)$ (marginalized over distance) for GW170104.

(i) GraceDB (Section 2.1): Loading HEALPIX maps from the gravitational wave server.

(ii) Telescope configurations (Section 2.2): Configuration files giving required information about the telescopes for optimization purposes.

(iii) Tiling (Section 2.3): Algorithms for determining what the best fields to observe are.

(iv) Time allocation (Section 2.4): Algorithms for determining how long the exposure times are and the number of exposures for each of the fields.

(v) Scheduling (Section 2.5): Algorithms for determining the order in which each of the fields is observed.

(vi) Efficiency (Section 2.6): Probability of detecting transients given the specific characteristics of the skymap, the telescope, and the tiling, time allocation, and scheduling choices.

2.1 GraceDB

```
python gwemopt_run -doEvent -do3D -
event G268556
```

GraceDB is a service that provides information on candidate gravitational wave events and the multimessenger follow-ups performed on them. An application programming interface (API) is made available that allows for access to this information. *gwemopt* uses this API to access information pertinent for gravitational wave follow-ups. First of all, it downloads the gravitational wave skymap for a given event; an example is shown in Fig. 2. In addition, information such as the time of the event and the time delay between the time of arrivals at the detectors is noted.

2.2 Telescope configuration

```
python gwemopt_run -doEvent -do3D -
telescope LSST
```

gwemopt relies on standardized configuration files for the telescopes to be analysed (please see² for examples for the telescopes in Table 1). The information in these files includes the filter being used, the limiting magnitude of the instrument, the exposure time required to achieve that magnitude, site location information, and information about the shape and size of the FOV. Different telescopes have different FOV shapes. For this, two options, square and

circle, are available, with the FOV being specified by the length of the square side and the radius of the circle. In addition, a tessellation file, which encodes the placement of the fields on the sky, is requested. This is especially useful for telescopes such as ZTF, which use fixed telescope pointings. Fixed telescope pointings can be useful for ensuring the availability of reference images, which are useful when performing image subtraction to look for transients. In case a tessellation file is not available, one is automatically generated, and this output is described in the next section. Configuration files for ATLAS, BlackGEM, LSST, PS1, and ZTF are available. Table 1 provides the information assumed for these telescopes.

2.3 Skymap tiling

```
python gwemopt_run -doEvent -do3D -doTiles -
doPlots -tileType ranked
```

Once a telescope configuration has been determined, the next step of the analysis is to generate the skymap tiling. There are a variety of algorithms in the literature for skymap tiling, and the ones implemented in *gwemopt* will be detailed below. The idea is to cover the sky with tiles the size of the telescope’s FOV with minimal overlap between the tiles. In some cases, these tiles are pre-determined by survey constraints in order to simplify difference imaging, where a reference image is subtracted from a science image to facilitate transient identification. In other cases, it is possible to optimize the tile locations based on the gravitational wave skymaps, such that the tiles maximize the probability contained. Due to the FOV for these telescopes being in general much smaller than the probability region, the effect is expected to be relatively minimal. For example, the number of tiles chosen to allocate time to will be dominated by choices such as whether there is utility in filling in chip gaps (gaps in the CCD arrays that compose the entire camera) and thereby having overlapping tiles, as opposed to simply maximizing the probability contained. Gravitational wave skymaps in general contain metrics that report the spatial probability of a gravitational wave source lying within a certain location. They are composed of HEALPIX arrays that encode either the 2D probability, in right ascension and declination, or 3D probability, which includes probability distributions for the luminosity distance of the transient. These are reported in maps with a particular number of pixels, usually $N_{\text{side}} = 512$. This can introduce quantization errors, which arise from using a map of limited resolution, especially for small FOV telescopes. The `nside` option in the code allows for the upsampling and downsampling of the skymaps in the analysis.

There are four options related to skymap tiling currently available and defined below: ranked, MOC (multiorder coverage), hierarchical, and greedy. In the following, we will summarize the key features of each implementation and refer the reader to the literature for further details. The goal is to create a common mathematical formalism for straightforward comparisons between algorithms. An example tiling for each scheme is shown in Fig 3.

Ranked. The ranked tiling scheme, described in Ghosh et al. (2016), uses pre-defined sky cells (see fig. 2 of Ghosh et al. 2016 for a visualization of how this tiling is performed). This tiling scheme is based on a grid system with grids of equal sizes. The sizes of the grids are the same as the size of the telescope’s FOV. For each tile in the grid at (α_i, δ_i) , where α_i is the right ascension and δ_i is the declination, we calculate a double integral that accumulates the

²<https://github.com/mcoughlin/gwemopt/tree/master/config>

Table 1. Configuration of telescopes.

Telescope	Latitude (deg)	Longitude (deg)	Elevation (m)	FOV (deg)	FOV shape	Filter	Exp. time (s)	Lim. mag.
ATLAS	20.7204	-156.1552	3055.0	5.46	Square	<i>c</i>	30.0	18.7
Pan-STARRS	20.7204	-156.1552	3055.0	1.4	Circle	<i>i</i>	45.0	21.5
BlackGEM	-29.2612	-70.7313	2400.0	2.85	Square	<i>g</i>	300.0	23.0
LSST	-30.1716	-70.8009	2207.0	1.75	Circle	<i>r</i>	30.0	24.4
ZTF	33.3563	-116.8648	1742.0	6.86	Square	<i>r</i>	30.0	20.4

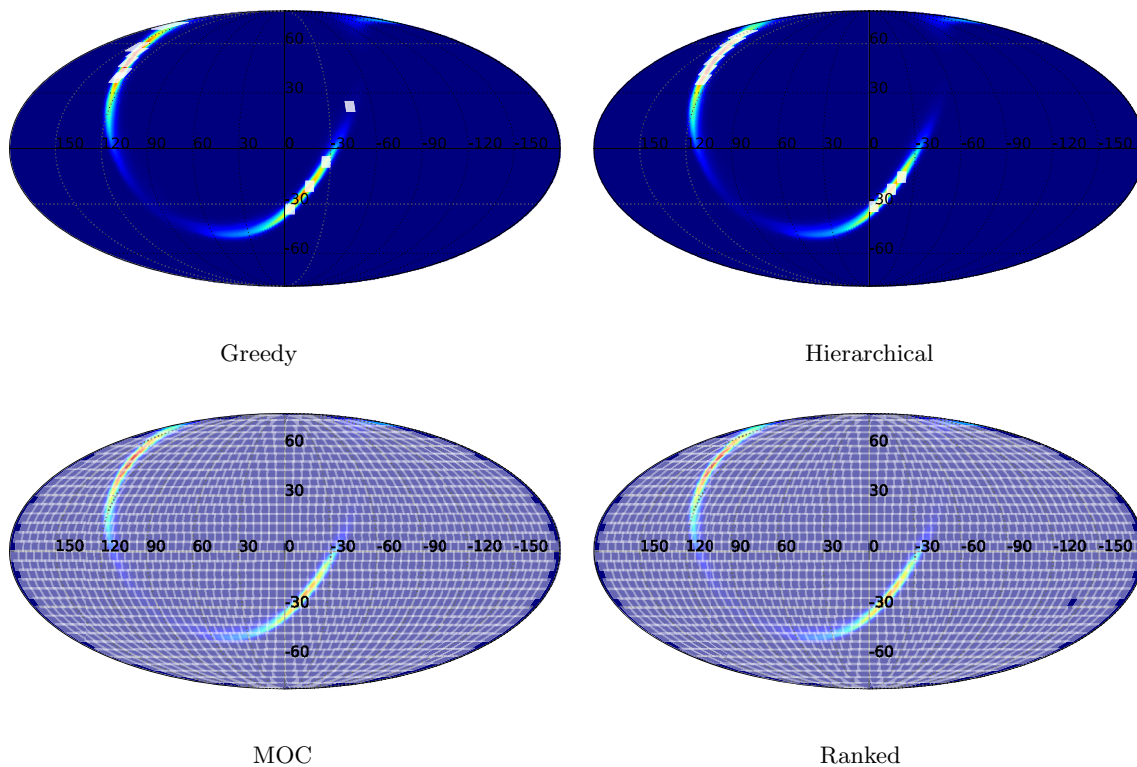


Figure 3. Example outputs of different tiling algorithms. On the top left is the greedy version with $N_{\text{tiles}} = 10$, where N_{tiles} is the number of tiles employed, on the top right is the hierarchical version with the same, on the bottom left is the multiorder coverage (MOC) skymap, and on the bottom right is the ranked skymap tiling. The top rows of skymaps are similar, given that 10 tiles are used for both algorithms. The differences between the two schemes, one of which places the tiles all at once and the other places them sequentially, when integrating the probabilities, lead to minor differences in tile placement. The bottom rows, with two schemes that cover the entire grid, are identical up to small numerical differences. This is expected as MOC is an on-the-fly scheme, and ranked pre-computes the tile locations.

probability distribution in this tile, shown in equation (1)

$$T_{ij} = \int_{\alpha_i}^{\alpha_i + \Delta\alpha} \int_{\delta_i}^{\delta_i + \Delta\delta} L_{\text{GW}}(\alpha, \delta) d\Omega, \quad (1)$$

where $L_{\text{GW}}(\alpha, \delta)$ is the sky location probability of the event as a function of right ascension and declination, as derived from the analysis of gravitational wave data. Then, we rank all the tiles with their T_{ij} and select from the top of the rankings until we reach the target-integrated probability desired. In the following, we will take 95 per cent as our target-integrated probability, which is a reasonable trade-off between capturing as many potential counterparts as possible while also limiting the sky coverage required. Other integrated probabilities could also be considered, and the ideal choice will depend on a trade-off between a few different priorities. These could include the desire to take at least one image of a potential counterpart by covering nearly the entire probability region; another possibility is to take enough images of the same

fields to determine the candidates that are appropriately fading and reddening by mapping out their luminosity and colour and time evolution.

MOC. The ‘MOC’ tiling scheme, based on the MOC of HEALPIX maps, hierarchically pre-defines cells in order to specify arbitrary sky regions (Fernique et al. 2014). MOC is proposed in order to provide ‘fast set operations’ between regions on the sky, which are designed to minimize computational time for standard set operations (unions, intersections, etc.). In MOC, the spherical sky is recursively divided into four regions and each region is a diamond. The division stops according to the resolution necessary for a particular usage. MOC is particularly useful in the case that most information is contained in just a handful of pixels, and therefore it is desired to save that information at high resolution for those pixels, and the remainder of the map at low resolution. This is important to be both memory efficient and retain all the relevant data.

Here are three relevant implementation details about MOC.

- (i) MOC uses an equatorial coordinate system.
- (ii) MOC divides the sphere recursively into four diamonds.
- (iii) MOC indexes each tile as follows: the initial tile is numbered 0 on level 0. Then, when divided, we get tile indices of 0, 1, 2, 3 on level 1. More generally, if we start from a tile numbered M , its children will be numbered $M \times 4$, $M \times 4 + 1$, $M \times 4 + 2$, $M \times 4 + 3$ on the next level.

The scheme for integrating probability in tiles is the same as in ‘ranked’ above.

Hierarchical. The hierarchical tiling scheme, which is a MultiNest-based (Feroz, Hobson & Bridges 2009b; Feroz et al. 2009a; Buchner et al. 2014) optimization, chooses tiles for a given skymap by placing them sequentially on the skymap and maximizing the probability at each step. This is distinctive from the ranked scheme in that it does not use pre-defined sky cells. MultiNest was chosen because with many live points available, multiple live points could simultaneously explore multiple portions of the skymap at once. This helps to overcome any potential issues that arise from natural bimodality of gravitational wave skymaps, where a single chain could be caught on a single island. We note that any algorithm designed for high-dimensional sampling is possible here. This method starts by selecting the tile that covers the most probability. Then, it sets the probability in that tile to be zero before going to the next iteration, when it again selects the tile that covers the most probability. The algorithm stops when a user-specified number of tiles, N_{tiles} , are selected. The tiles selected might overlap on the corners when there are higher probability distributions around that corner.

Greedy. The greedy tiling scheme, an EMCEE-based algorithm (Foreman-Mackey et al. 2013), optimizes tiles for a given skymap by placing all tiles simultaneously on the map. This allows all tile locations to vary at the same time. Similar to the ‘hierarchical’ case, any high-dimensional sampler would be reasonable to use here. The algorithm works by taking the user-specified number of tiles, N_{tiles} , and allowing the tile locations to vary. It uses a likelihood that maximizes the integrated probability. It prevents double counting by setting the probability of a given tile to zero when integrating the probability. This method selects tiles that cover the highest integrated probability from the skymap. It ranks all possible tiles and selects from the top. Thus, the tiles selected by greedy methods might overlap significantly when the probability distribution is concentrated.

2.4 Time allocations

```
python gwemopt_run -doEvent -doPlots -
doTiles -doSchedule -timeallocationType cov-
erage
```

Having decided on the preferred set of tiles, exposure time allocation forms the next important constraint to address. Because telescopes have FOVs that are in general smaller than the probability region and typical exposure times for these telescopes are of order minutes (see Table 1), it is therefore not possible to image the entire probability region to interesting limiting magnitudes in a reasonable amount of time. There are further constraints that arise from the diurnal cycle, observing time available for follow-up, limitations on the pointing that a particular telescope is capable of (such as horizon limits), and the rising and setting of tiles. The follow-

ing algorithms use a variety of methods to optimize the probability of imaging a counterpart with the constraint of limited telescope resources.

The amount of time allocated is defined with a few constraints. First of all, time segments are generated based on the observing time allocated after the gravitational wave event. In the following analyses, we will assume that 72 h following the event are available for follow-up (the code can account for limited target of opportunity time by breaking the time into segments). The segments are then intersected with nighttime at the site of the particular telescope, which defines the segments that can be used for observations. This assumes implicitly that the electromagnetic counterpart has not faded beyond detection limits in the time available. Some of the time allocation algorithms below use models to determine the detectability of the light curve. There are four options related to time allocations as a function of sky location, coverage, power law, WAW (where and when), and PEM (probability of electromagnetic counterpart). Fig. 4 shows examples of the power law, WAW, and PEM types.

Coverage. The ‘coverage’ option is one whereby coverage from existing surveys, including the right ascension and declination of the pointing and the limiting magnitude, is used. The benefit of this mode is to establish efficiencies of detection of kilonovae in existing surveys. For example, this code will be used to determine the efficiencies of the Pan-STARRS and ATLAS surveys to both gravitational wave and gamma-ray burst triggers. It can be used as a way to prioritize examination of survey data for serendipitous observations of these transients, where both reference images and the time of the transients are known.

Power law. The ‘power law’ option is one where scaling relations are applied to the probability map to determine the time allocation. Many authors (Coughlin & Stubbs 2016; Ghosh et al. 2016; Chan et al. 2017) have proposed a variation on simply scaling the time allocation proportional to the probability skymaps, a technique employed in the power-law method below. Pan-STARRS and ATLAS based searches have employed this technique (Smartt et al. 2016a,b; Stalder et al. 2017). Coughlin & Stubbs (2016) derived scaling relations for the time allocated to any given field, t_i , given the gravitational wave posterior probability distribution. While the power-law-based analysis is straightforward, it does not account for the fact that the telescope must be sensitive enough to detect the counterpart. In this sense, this algorithm is the least model dependent. The detectability of a given transient is model dependent, depending both on the distances predicted by the gravitational wave data analysis and the absolute magnitude of the sources; the following algorithms account for this dependence in multiple ways.

The *power-law* algorithm optimizes the probability of detecting the transient with N observations, which is simply the sum of the probability of each observation. The expression is shown in equation (2):

$$p_{\text{tot}} = \sum_{i=1}^N \frac{M_i}{M_{\text{tot}}} \frac{L_{\text{GW}}(\alpha_i, \delta_i, R_i)}{L_{\text{GW,tot}}} \frac{F_i(t_i)}{a(\alpha_i, \delta_i)}, \quad (2)$$

where R_i is the distance to the galaxy, M_i is the mass for galaxy i , M_{tot} is the total mass of galaxies in the field, $L_{\text{GW}}(\alpha_i, \delta_i, R_i)$ is the posterior probability distribution of the gravitational wave source in this galaxy, $F(t)$ is the luminosity as a function of allocated time, and $a(\alpha_i, \delta_i)$ is the Galactic extinction. In the following, we will simply scale the gravitational wave posterior probability distribution as $L_{\text{GW}}(\alpha_i, \delta_i)^{2/3}$ (as shown by the analysis of Coughlin & Stubbs 2016 to be optimal given their assumptions). Within this formalism, independent scaling of the distance is also possible (such as with R^4 to maintain a constant electromagnetic signal-to-noise ratio for a

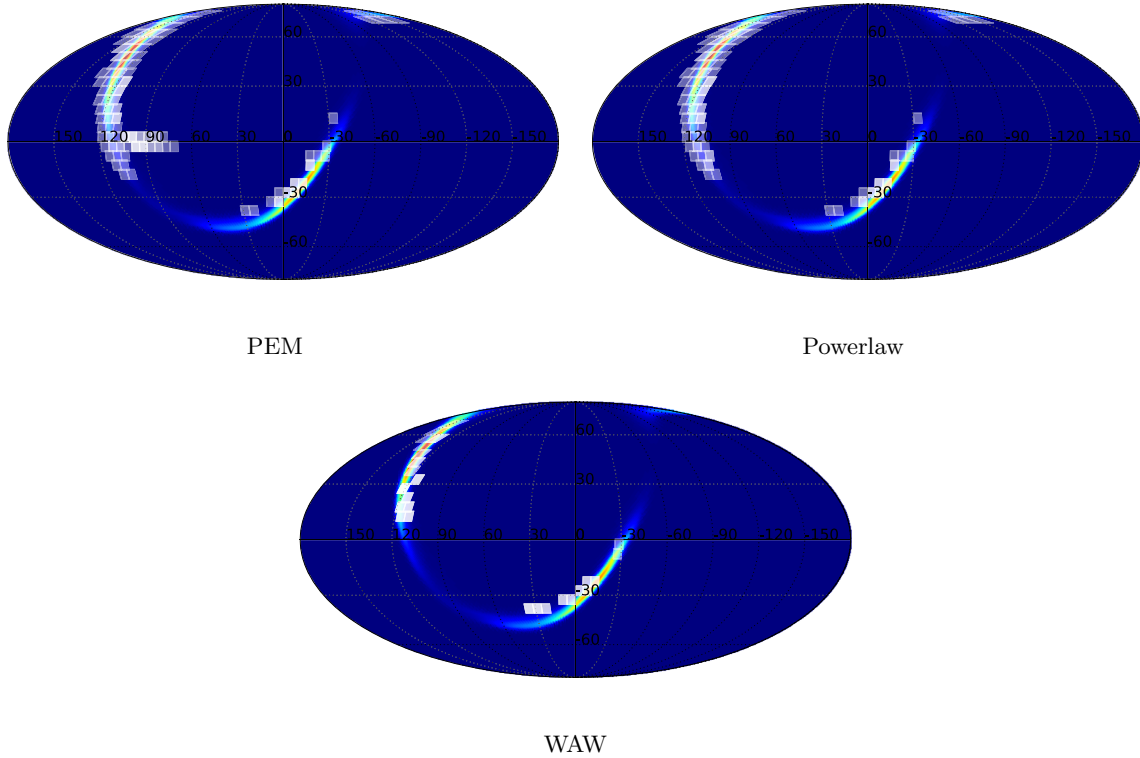


Figure 4. Example outputs of different time allocation algorithms. On the top left is the tile coverage with the PEM algorithm. On the top right is the tile coverage with the power-law algorithm. On the bottom is the tile coverage with the WAW algorithm. In generating all of the plots, the MOC algorithm is used. Given the differences in the algorithms discussed in the text, it is unsurprising that some differences in the time allocations are seen.

particular target absolute magnitude), but this dependence has been accounted for in calculating the skymap probability’s associated power law. It is possible that this approximation could be improved using galaxy catalogues, although this introduces concerns about galaxy catalogue completeness. Equation (2) is optimized with the constraint that the total observation time is limited, shown as

$$\sum_{i=1}^N t_i = T, \quad (3)$$

where T is the total observation time.

WAW. The ‘When and Where’ algorithm, defined in Salafia et al. (2017), uses counterpart light-curve models in the optical, infrared, and radio constructed from information from the gravitational wave signals to create a time- and sky location-dependent probability for detecting electromagnetic transients. The WAW approach introduces time into the model by defining a concept of detectability. Detectability is the probability of detecting a light flux greater than the flux limit at position α and time t . These flux limits are computed using a combination of light-curve models and knowledge of distance to the counterparts for compact binary coalescences. Thus, by having detectability introduced, the algorithm can optimize ‘where’ and ‘when’ to schedule the observation based on α and t with a greedy approach. The procedures of the algorithm are shown below.

- (i) The tiles are generated covering the confidence region based on the probability distribution, which comes from the gravitational wave signal.
- (ii) The algorithm takes in the information encoded in the gravitational waves and computes the light curve $F_i(t)$ for each tile.
- (iii) Then, the algorithm computes the detectability as

$$P(F(t) > F_{\text{lim}}|\alpha, S) \approx \sum_{i=1}^N \omega_i H(F_i(t) - F_{\text{lim}}), \quad (4)$$

where H is the Heaviside function where if $F_i(t)$ is greater than F_{lim} , it is 1; otherwise it is 0. $F_i(t)$ is the light flux for position sample i at time t . F_{lim} is the limiting flux, which is the minimum detectable flux by the instrument. ω_i is the ‘inverse distance weight’ that gives the contribution of the sample i to position α . The further away sample i is from α , the less it contributes. ω_i is normalized so that $\sum_{i=1}^N \omega_i = 1$.

(iv) For each tile, we find a time interval $[t_{E,\lambda}, t_{L,\lambda}]$ when the detectability is greater than a threshold λ .

(v) We start from the tiles that cover the most probability and arrange their observation times $[t_{E,\lambda}, t_{L,\lambda}]$ if the time is available.

This method optimizes the search by introducing detectability, defined as equation (4) over the 3D observation volume of direction and time, with the constraint that only one location can be observed at the same time.

PEM. The ‘Probability of electromagnetic counterpart’ algorithm, defined in Chan et al. (2017), optimizes the number of fields to observe and their time allocations by adopting priors on the intrinsic luminosities of the sources and using knowledge of distance to the counterparts for compact binary coalescences. More concretely, its input is the sky localization map and information about the telescope. It selects the tiles to observe with a greedy algorithm and allocates the observation time for each tile to maximize the probability of detecting the EM counterpart of the GW event.

The procedures of the algorithm are shown below, and further details can be found in Chan et al. (2017).

- (i) Based on the sky localization map, we locate the tiles that cover the region enclosed by the contour of the target confidence level.

(ii) These N tiles are ranked based on the total probability covered.

(iii) We optimize the number of tiles selected and then the time allocation for each selected tile. For all k from 1 to n , we do the following:

(a) The top k tiles from the rankings are selected.

(b) Equation (5) is the total detection probability of all the tiles and optimized with Lagrange multipliers with the constraint of equation (6), which encodes the limited observational time on the telescope:

$$P(D_{\text{EM}}|k) = \sum_{i=1}^{k \leq n} P(D_{\text{EM}}|\omega_i^{(k)}, \tau_i^{(i)}, I) \quad (5)$$

$$kT_0 + \sum_{i=1}^k \tau_i^{(i)} = T, \quad (6)$$

where ω_i is the probability density, τ_i is the time allocated and I are the parameters of the telescope, T_0 is the time to adjust the telescope before each observation, τ_i is the time allocated to each tile, and T is the total observation time.

(c) Equation (5) is maximized with the constraint of equation (6) to determine the best k , the tiles and the time allocation.

(iv) The optimal tiles $\{\omega_i\}$ and their allocated times $\{\tau_i\}$ for $i = 1 \dots k$ are the output.

unnumlist

2.5 Scheduling

```
python gwemopt_run -doEvent -doPlots -
doTiles -doSchedule -scheduleType weighted
```

Once the time allocated to each tile has been set, the next task is to schedule the observations that both best represent the time requested and optimize the times that are chosen in some way. For example, the tiles could be re-imaged at an approximately fixed cadence so as to measure possible light-curve evolution. Another option is to simply go as deep as possible in one field to ensure detection of a counterpart there. Other optimizations might employ ordering based on airmass (or the amount of atmosphere a telescope observes through), as sources imaged through higher airmass will have lower signal-to-noise ratios.

For each of the algorithms that perform the scheduling, the time that each tile is available for observation above the altitude limit, which corresponds to the lowest altitude a telescope can observe (assumed to be 30° in this analysis), is computed. These tile-specific times are intersected with the set of times available to the telescopes to form a set of visibility segments for each tile. This has the benefit of avoiding issues related to simply tracking the rise and set times of each tile. To account for lunar sky brightness, we use a model from Coughlin, Stubbs & Claver (2016). Any tile whose sky brightness is increased by at least 1 mag is excluded.

There are three options related to scheduling observations: greedy, SEAR, and weighted. *Greedy*. The ‘greedy’ algorithm, which is the simplest version in the code package, employs a schedule on the basis of probability contained. The idea is that higher ranked tiles are observed before lower ranked tiles based on this ranking scheme. Rana et al. (2017) implement a greedy algorithm whereby the field with the highest probability region in a given time window is observed. As the Rana et al. (2017) implementation does not include the possibility of multiple exposures for each pointing,

it is modified in this analysis to include multiple exposures. The algorithm is as follows:

(i) Construct a list of the tiles and number of exposures for each tile based on the time allocation algorithm utilized.

(ii) For each window, find the sky tiles that are in the current window: $T_0 + (j - 1)T_{\text{exp}}$ and $T_0 + jT_{\text{exp}}$, where T_0 is the start time for the observation, T_{exp} is the exposure time, and j is the index of the window.

(iii) Allocate the window to the sky tile with the greatest probability, and increment the number of exposures for that tile down by 1.

SEAR (Setting Array). The ‘Setting Array’ algorithm overcomes the shortcoming of the greedy algorithm to include site visibility. This motivates re-ordering the sequence such that as many tiles can be imaged as possible. Rana et al. (2017) also implement a version whereby the rising and setting of tiles were accounted for. SEAR prioritizes observing high-probability tiles first, subject to the condition that each tile from the observing sequence must be observed before it sets. The concept of imaging ‘windows’ is used in this algorithm, where the available observational time is broken up into segments referred to as windows. We denote the i th window as W_i . We denote the tile that will be imaged in window W_i as the tile S_i . The algorithm uses the recursive relation between the optimal observation arrangement between the first k windows $S_1 \dots S_k$ and the $k + 1$ window S_{k+1} . The details are shown below.

(i) Consider the first window W_1 and initialize S_1 to be the tile that has the highest probability for W_1 .

(ii) Move on to W_2 and find the two tiles c_1 and c_2 that have the greatest probability density.

(iii) Compare c_1 and c_2 with S_1 and act depending on the following conditions:

(a) If both c_1 and c_2 contain greater probability than S_1 , set $S_1 = c_1$ and $S_2 = c_2$.

(b) Otherwise, put the tile with higher probability coverage between c_1 and c_2 into S_2 ; c_2 is then assigned to S_3 .

We can see that either way, S_1 and S_2 will have the two highest probability tiles observable.

(iv) Move on to the next observation window and repeat until the last one is reached.

(v) Return the last set S_w , where w is the total number of observation windows.

Weighted. A different scheme, known as ‘weighted’, is an algorithm where each tile is weighted based on both gravitational wave posterior probability distribution enclosed, the number of exposures required for that tile, and the number of available slots for it to be imaged. This is an alternative to greedy and SEAR, as it is impossible to observe all of the tiles as they rise and set given the requirement of using multiple exposures per tile.

The idea is that all of the tiles are given a weight that depends on the requested number of exposures. Therefore, we define the weights w_i as

$$w_i = L_{\text{GW}}(\alpha_i, \delta_i) \times \frac{N_{\text{R}}}{N_{\text{A}}}, \quad (7)$$

where N_{R} is the number of remaining images to be taken for a given tile and N_{A} is the number of allocated exposures for that tile. In this way, $\frac{N_{\text{R}}}{N_{\text{A}}}$ prioritizes the tiles with the most remaining exposures. Therefore, for each exposure segment, we calculate the weight for

each possible tile and select the tile with the highest weight to fill that slot.

2.6 Efficiency

```
python gwemopt_run -doEvent -doPlots -
doTiles -doSchedule -doEfficiency
```

We are able to test and compare the performance of these algorithms by performing simulated observations. We adopt observational constraints as follows. We use an observing limit of an altitude of 30° , corresponding to an airmass of 2.0. We assume observations are available to begin at twilight and dawn, corresponding to when the sun is 12° below the western and eastern horizons. We do not point away from the moon or account for sky brightness.

The simplest metric of success is the ‘efficiency’, which is defined as the number of transients detected over the number of transients injected. To estimate the efficiency for the ‘detection’ of the electromagnetic counterparts to gravitational wave transients, we perform simulated injections of supplied light curves, which corresponds to the absolute magnitude in the colour requested in the telescope configuration file. The number of transients injected is proportional to the values in the sky localization probability map to account for the selection effects of the gravitational wave detector network. We provide example light curves for a variety of light-curve models, including

- (i) Tanaka et al. (2014): Simulations of binary systems showing ejecta morphology and resulting light curves. These simulations led to analytical models for black hole neutron star systems from Kawaguchi et al. (2016) and Dietrich & Ujevic (2017).
- (ii) Kawaguchi et al. (2016): Analytical models for black hole neutron star systems based on Tanaka et al. (2014).
- (iii) Dietrich & Ujevic (2017): Analogue to Kawaguchi et al. (2016) for binary neutron star systems.
- (iv) Barnes et al. (2016): Simulations of binary systems studying the emission profiles of radioactive decay products from the merger.
- (v) Metzger et al. (2015): Blue ‘precursor’ to the kilonova driven by β -decay of the ejecta mass.
- (vi) Metzger (2017): Toy model with grey opacity for lanthanide-free matter expanding with a range of velocities with a mass density profile $M(<v) = v^{-1}$.

The requirements for ‘detection’ of the electromagnetic counterparts to gravitational wave transients are as follows. We require that the transient appear in two images over two nights. In each image, the transient must exceed the limiting magnitude in that image. The colour of the transient is estimated from the filter given in the configuration file. We simulate the transients at a variety of locations and distances consistent with the gravitational wave probability skymap.

3 PERFORMANCE

In this section, we compare the efficiency of the algorithms based on simulated information about what percentage of the events the algorithm can detect. According to the workflow given in Fig. 1 and the algorithms given in the sections above, we will have four options for tiling algorithms, three options for time allocation algorithms, and another three options for scheduling algorithms. This combines to 36 total options for the whole workflow. We want to know which combination has the best efficiency and then analyse

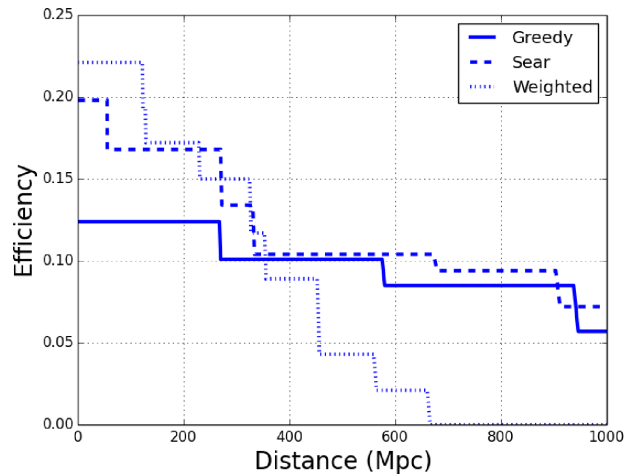


Figure 5. Example plot of efficiency for Metzger (2017) injections, comparing the scheduling algorithms. The constraint for each algorithm was on the total observing time available. The greedy algorithm is used for the tiling and PEM algorithm is used for time allocation. The greedy and the SEAR algorithms have similar performance for long distances and both perform better than weighted. This difference is also reflected in Fig. 6. As the algorithm accounts for observability from a site, including both whether tiles are visible from the site of interest and diurnal effects, efficiencies are expected to peak at around 25 per cent for an event that fades quickly and has a probability region with peaks in both the north and south.

and compare the algorithms individually. We take as an example the ATLAS instrument in the following. As all of the instruments are large aperture, wide-field instruments, the results will not strongly depend on the instrument chosen and the way the efficiencies scale will be the same for the level of detail considered here. Including information such as slew and readout time of the individual telescopes will change the results in case of lengthy readout times or characteristically slow slew times. For telescopes with either readout or slew times on the order of the typical length of exposures, the efficiencies could be impacted by up to a factor of 2. For the survey telescopes we consider in this analysis, they are designed to have short readout and slew times so as to minimize this overhead, and so the loss in efficiency will be much smaller.

3.1 Method

We will focus on the model by Metzger (2017) to compare the efficiency. All the efficiency values in the (logarithmically spaced) distance range between 10^{-1} and 10^3 Mpc are calculated and plotted. Thus, we will have a plot of efficiency versus distance for each of the 36 algorithm combinations. An example efficiency plot is shown in Fig. 5, where different time allocation algorithms are compared. The greedy algorithm is used for tiling and the PEM algorithm is used for time allocation. It can be seen that greedy and SEAR scheduling do better than weighted for long distances.

In Fig. 5, we show an example where we hold the tiling and time allocation algorithms fixed (to disentangle their effects from the scheduling algorithm), and show the efficiency as a function of distance for the scheduling algorithms discussed in this paper. We find that both Greedy and SEAR perform better than weighted at larger distances. Weighted performs best at the lowest distances considered because it targets the highest probability tiles, ensuring that they will be scheduled for multiple exposures so as to meet the detection criteria. Greedy and SEAR, both of which schedule tiles




Figure 6. Plot of the weighted efficiency metric for each of the 36 options. On the horizontal axis are tiling algorithms and scheduling algorithms and on the vertical axis are the time allocation algorithms. Abbreviations are used for the algorithms. The first capital letter stands for the tiling algorithm and the second letter stands for the scheduling algorithm. The abbreviations are the first letters of the algorithms: G, greedy; H, hierarchical; M, MOC; R, ranked; S, SEAR; W, weighted. The grids are coloured such that highest efficiency combinations are darker and lower efficiency ones are lighter, with the highest being completely blue and the lowest one being completely white.

to maximize the number of fields imaged, with SEAR accounting for the rising and setting of the tiles, perform significantly better at greater distances. This is because they sacrifice some efficiency for nearby sources by effectively exploring more fields, but have a more constant detection efficiency with distance.

In order to compare the 36 efficiencies as plotted in Fig. 5, we use a single statistic to reflect the overall performance of the algorithms based on the efficiency for each distance in the range of 10^{-1} – 10^3 Mpc. Thus, we come to a metric that reflects the percentage of events that can be detected in a spherical volume of radius 10^3 Mpc. The events are evenly distributed in the volume. Suppose the event density per volume is ρ and the distance is r . Sampling a distance at r corresponds to a shell with volume $4\pi r^2 dr$. Assuming that the density is ρ , then the total events on that shell are $4\pi\rho r^2 dr$. Thus, if the efficiency is e , the expected number of detected events on the shell will be $4\pi\rho r^2 e dr$. From this we can see that the efficiency at r is weighted by r^2 . If we treat the efficiency at each distance as an individual sample, a weighted average on the squared radius will then be a good metric of the overall efficiency; it reflects how well the algorithm detects events uniformly distributed in a spherical volume of radius 10^3 Mpc. A consequence of this metric is that the weight of r^2 makes the long-range efficiency more important than the short-range efficiency. Under this metric, an algorithm that performs well at further distances would be better than an algorithm that does better at short distances but whose performance deteriorates quickly as the distance increases.

3.2 Performance and algorithms

For each of the 36 algorithm options, we compute the efficiency metric as described above, which results in 36 numerical efficiency values. The results are plotted in Fig. 6. On the horizontal axis are the combined options for the tiling algorithm and scheduling algorithm. There are four tiling algorithms and three scheduling algorithms, so they combine to 12 columns on the horizontal axis. Abbreviations are used for the algorithms. The first capital letter stands for the tiling algorithm and the second letter stands for the scheduling algorithm. On the vertical axis are the time allocation algorithms. The colour in the 36 boxes shows the efficiency as measured above. The colours are scaled to the efficiency such that higher efficiencies are more darkly coloured. The highest efficiency of 0.19 is achieved by a combination of ranked tiling, power-law time allocation, and greedy scheduling. Compared to the lowest efficiency of 0.01, it can detect approximately 19 times more events

within a range of 10^{-1} – 10^3 Mpc. That corresponds to the darkest box in the 10th column and the second row in Fig. 6.

Also, from Fig. 6, we can compare the efficiencies of the individual algorithms. We concentrate on the SEAR scheduling algorithm in the following. For example, in the case of the use of the MOC tiling algorithm, where all (non-overlapping) tiles are available for allocation, we can look at the ‘MS’ column. The performance of PEM and power law are very similar, and this fact holds true for all of the configurations. This similarity may be surprising, given that the power-law algorithm does not use the distance posteriors like PEM does. Instead, the power-law algorithm encodes the way the probability should be allocated with distance in its power-law term, and when considering distances where the ‘shot noise’ of galaxies disappears, they become equivalent. WAW suffers somewhat from the lack of the asset that makes it most valuable: inclusion of the dependence of light curves on the inclination angle of the original binary. Therefore, it is not surprising that PEM outperforms WAW in this instance, as PEM is optimal in the presence of distance information, while WAW requires inclination information in order to be optimal. However, in case inclination and distance were available in low latency, WAW may be best. The lower efficiencies in the case of overlapping tiles of the greedy and hierarchical tiling algorithms are due to the upper limit of the number of tiles posed (to be returned to in the next section), while performing similarly to one another. This is unsurprising, as both algorithms differ only in the order in which tiles are determined (greedy optimizes the location all at once, while hierarchical places one tile after the other).

3.3 Performance and the number of tiles

We also study how the number of tiles affects the performance. This study in particular will depend on the FOV of the instrument in question. It is only relevant in the hierarchical and greedy tiling algorithm. The hierarchical algorithm is used for this study, which selects tiles that cover the highest probability and masks the tiles once they are selected. It stops when a user-defined number of tiles is selected. The efficiency is computed based on the simulation of 5000 injections.

From the plot on the left of Fig. 7, it can be seen that efficiency increases at first when we increase the number of tiles, but after a peak is seen around 35 tiles, the efficiency starts decreasing. The right plot shows the efficiency curves for different number of tiles. Blue indicates high efficiency; white indicates low efficiency. It can be seen from the plot that efficiency at the smallest distances continues to improve as the number of tiles is increased. However, as the number of tiles increases, the efficiency for distant sources first rises and then falls (starting around 35 tiles). This occurs because as more tiles are chosen, less overall time can be allocated to the sources, which would otherwise require longer exposure times. The most distant sources are found when using approximately 35 tiles and this corresponds to the peak we see in the left figure. It is important to note that just because a tile is selected by the tiling algorithm, it is not assured that time will be allocated to that tile or it will be scheduled. It is also important to note that the algorithms presented here do not account for readout or slew times, which will have some effect on these results. For an alternative study that accounts for these effects, please see Chan et al. (2017).

4 CONCLUSION

The detection of GW170817 (Abbott et al. 2017b) has invigorated the search for improved strategies for associating gravita-




Figure 7. On the left is the plot of efficiency weighted average as a function of the number of tiles considered. On the right is efficiency as a function of distance for different number of tiles. In the simulation, hierarchical tiling algorithm, power-law time allocation algorithm, and greedy scheduling algorithm are used to detect 5000 random injections for number of tiles ranging from 1 to 100. The standard deviation is calculated and the 99 per cent confidence interval is plotted as the grey shaded region in the top figure. White indicates lower efficiency and blue indicates higher efficiency. We can see that as the number of tiles increases, low distance gets better, but high distance gets worse. We see the least blue at high distance around 35 tiles, which corresponds to the peak in the left plot.

tional waves with electromagnetic counterparts. Due to the large uncertainty footprint, which can range from 10 to 1000 deg² (Abbott et al. 2016a), efficiently scanning sky areas of this size in search of an electromagnetic counterpart is challenging. However, in this paper, we have described and compared a number of algorithms in the literature available for significantly improving upon the most naive approaches. We have shown comparisons between the algorithms, describing the limits in which they are the most effective.

One potential improvement to the analysis considered here is using the locations of known galaxies in the gravitational wave sensitivity distance, which was ≈ 100 Mpc for GW170817 (Abbott et al. 2017b) and will extend to ≈ 300 Mpc at design sensitivity (Abbott et al. 2016a). Recent improvements in galaxy catalogue completeness have made this effort possible. For example, the Galaxy List for the Advanced Detector Era (GLADE) galaxy catalogue is complete (with respect to a Schechter function) out to ≈ 300 Mpc for galaxies brighter than the median Schechter function galaxy luminosity.³ The Census of the Local Universe (CLU) catalogue (Cook et al. 2017) is complete to 85 per cent in star formation and 70 per cent in stellar mass at 200 Mpc. Within these local volumes, the sky area coverage of galaxies is ≈ 1 per cent (Cook et al. 2017), bringing the sky areas searched down by a factor of 100, which makes the possibility of targeted galaxy pointing tractable, especially for small FOV telescopes (see Arcavi et al. 2017 for an example).

Unfortunately, it is not immediately clear how the algorithms presented here will perform relative to the schemes that use galaxy catalogues. One expects that there will be two regimes where one technique will be optimized over the other. For galaxy-targeted searches, nearby sources with small localization volumes are most optimal, as these dedicated searches will be more sensitive to intrinsically fainter sources than the wide-field surveys. On the other hand, wide-field surveys will be more successful across large localization regions where the distance is not necessarily well constrained. As future work, we intend to explore the benefit of combining the power of both techniques. This may be done by modifying

the probability map to include the effects of a discrete mass distribution, similar to that proposed by Arcavi et al. (2017) for assigning priorities to individual galaxies. Instead, the probability map will become a sum of the contributions of the galaxies in any given pixel. For nearby events, this map will likely be very pixelated, given the limited number of galaxies that will contribute significantly. For further away events, this map will closely resemble the original probability map.

Also not addressed in this work is how to determine which sources are the best for allocating telescope time to. This is less of an issue for the wide-field, all-sky survey instruments, where target of opportunity observations change the cadence of the survey but are not strictly time lost to the overall endeavour. On the other hand, this is an essential question for target of opportunity observations on narrow FOV telescopes where the number of sources that can be followed up is often significantly limited. There have been a number of studies along these lines recently. For example, Del Pozzo et al. (2018) showed that the localization volume depends strongly on the signal-to-noise ratio of the gravitational wave event. Therefore, it is likely easier to make a successful observation of the counterpart associated with events with larger signal-to-noise ratio, and therefore it may be best to wait for the loudest events (Chen & Holz 2016). This was also addressed in Lynch et al. (2018), where it was pointed out that the rate of false positives also significantly increases as a function of signal-to-noise ratio. Going forward, optimizing the choice of events to follow up will be important.

A code to produce the results in this paper is available at <https://github.com/mcoughlin/gwemopt> for public download.

ACKNOWLEDGEMENTS

MC is supported by the David and Ellen Lee Prize Postdoctoral Fellowship at the California Institute of Technology. DT and NC are supported by the National Science Foundation, under NSF grant number PHY 1505373. YH is supported by the National Natural Science Foundation of China, under NSFC 11703098. CWS is grateful to the DOE Office of Science for their support under award DESC 0007881.

³<http://aquarius.elte.hu/glade/>

REFERENCES

- Aasi J. et al. 2015, *Class. Quantum Gravity*, 32, 074001
- Abbott B. P. et al. 2016a, *Living Rev. Relativ.*, 19, 1
- Abbott B. P. et al. 2016b, *Phys. Rev. Lett.*, 116, 061102
- Abbott B. P. et al. 2016c, *Phys. Rev. Lett.*, 116, 241103
- Abbott B. P. et al. 2016d, *ApJ*, 826, L13
- Abbott B. P. et al., 2017a, *Phys. Rev. Lett.*, 118, 221101
- Abbott B. P. et al. 2017b, *Phys. Rev. Lett.*, 119, 161101
- Abbott B. P. et al., 2017c, *Nature*, 551, 85
- Abbott B. P. et al. 2017d, *ApJ*, 848, L13
- Abbott B. P. et al., 2017e, *ApJ*, 848, L12
- Abbott B. P. et al., 2017f, *ApJ*, 850, L39
- Acernese F. et al. 2015, *Class. Quantum Gravity*, 32, 024001
- Arcavi I. et al., 2017, *ApJ*, 848, L33
- Barnes J., Kasen D., 2013, *ApJ*, 775, 18
- Barnes J., Kasen D., Wu M.-R., Martínez-Pinedo G., 2016, *ApJ*, 829, 110
- Bauswein A., Baumgarte T. W., Janka H.-T., 2013, *Phys. Rev. Lett.*, 111, 131101
- Berry C. P. L. et al., 2015, *ApJ*, 804, 114
- Bloemen S., Groot P., Nelemans G., Klein-Wolt M., 2015, in Rucinski S. M., Torres G., Zejda M., eds, ASP Conf. Ser. Vol. 496, Living Together: Planets, Host Stars and Binaries. Astron. Soc. Pac., San Francisco, p. 254
- Buchner J. et al., 2014, *A&A*, 564, A125
- Chan M. L., Hu Y.-M., Messenger C., Hendry M., Heng I. S., 2017, *ApJ*, 834, 84
- Chen H.-Y., Holz D. E., 2016, preprint ([arXiv:1612.01471](https://arxiv.org/abs/1612.01471))
- Cook D. O. et al., 2017, preprint ([arXiv:1710.05016](https://arxiv.org/abs/1710.05016))
- Cornish N. J., Littenberg T. B., 2015, *Class. Quantum Gravity*, 32, 135012
- Coughlin M., Stubbs C., 2016, *Exp. Astron.*, 42, 1
- Coughlin M., Stubbs C., Claver C., 2016, *Exp. Astron.*, 41, 393
- Coulter D. A. et al. 2017, *Science*, 358, 1556
- Del Pozzo W., Berry C., Ghosh A., Haines T., Vecchio A., 2018, preprint ([arXiv:1801.08009](https://arxiv.org/abs/1801.08009))
- Dietrich T., Ujevic M., 2017, *Class. Quantum Gravity*, 34, 105014
- Essick R., Vitale S., Katsavounidis E., Vedovato G., Klimentenko S., 2015, *ApJ*, 800, 81
- Fairhurst S., 2009, *New J. Phys.*, 11, 123006
- Fairhurst S., 2011, *Class. Quantum Gravity*, 28, 105021
- Fernique P., Boch T., Donaldson T., Durand D., O'Mullane W., Reinecke M., Taylor M., 2014, MOC – HEALPix Multi-Order Coverage map Version 1.0, IVOA Recommendation 02 June 2014, International Virtual Observatory Alliance
- Feroz F., Gair J. R., Hobson M. P., Porter E. K., 2009a, *Class. Quantum Gravity*, 26, 215003
- Feroz F., Hobson M. P., Bridges M., 2009b, *MNRAS*, 398, 1601
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Ghosh S., Bloemen S., Nelemans G., Groot P. J., Price L. R., 2016, *A&A*, 592, A82
- Goldstein et al. 2017, *ApJ*, 848, L14
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Grover K., Fairhurst S., Farr B. F., Mandel I., Rodriguez C., Sidery T., Vecchio A., 2014, *Phys. Rev. D*, 89, 042004
- Ivezic Z., Tyson J. A., Allsman R., Andrew J., Angel R., 2008, preprint ([arXiv:0805.2366](https://arxiv.org/abs/0805.2366))
- Just O., Bauswein A., Pulpillo R. A., Goriely S., Janka H.-T., 2015, *MNRAS*, 448, 541
- Kasen D., Badnell N. R., Barnes J., 2013, *ApJ*, 774, 25
- Kasen D., Fernandez R., Metzger B., 2015, *MNRAS*, 450, 1777
- Kawaguchi K., Kyutoku K., Shibata M., Tanaka M., 2016, *ApJ*, 825, 52
- Klimentenko S. et al., 2016, *Phys. Rev. D*, 93, 042004
- Lynch R., Coughlin M., Vitale S., Stubbs C. W., Katsavounidis E., 2018
- Metzger B. D., 2017, *Living Rev. Relativ.*, 20, 3
- Metzger B. D., Berger E., 2012, *ApJ*, 746, 48
- Metzger B. D. et al., 2010, *MNRAS*, 406, 2650
- Metzger B. D., Bauswein A., Goriely S., Kasen D., 2015, *MNRAS*, 446, 1115
- Morgan J. S., Kaiser N., Moreau V., Anderson D., Burgett W., 2012, in Proc. SPIE Conf. Ser. Vol. 8444, Ground-based and Airborne Telescopes IV, SPIE Bellingham, p. 8440H
- Nakar E., 2007, *Phys. Rep.*, 442, 166
- Radice D., Perego A., Zappa F., Bernuzzi S., 2018, *ApJ*, 852, L29
- Rana J., Singhal A., Gadre B., Bhalerao V., Bose S., 2017, *ApJ*, 838, 108
- Rau A. et al., 2009, *PASP*, 121, 1334
- Roberts L. F. et al., 2017, *MNRAS*, 464, 3907
- Salafia O. S., Colpi M., Branchesi M., Chassande-Mottin E., Ghirlanda G., Ghisellini G., Vergani S. D., 2017, *ApJ*, 846, 62
- Savchenko V. et al., 2017, *ApJ*, 848, L15
- Sidery T. et al., 2014, *Phys. Rev. D*, 89, 084060
- Singer L. P. et al. 2016, , *ApJL*, 829, 1, submitted
- Singer L. P. et al., 2014, *ApJ*, 795, 105
- Smartt S. J. et al. 2016a, *MNRAS*, 462, 4094
- Smartt S. J. et al. 2016b, *ApJ*, 827, L40
- Smartt S. J. et al. 2017, *Nature*, 551, 75
- Stalder B. et al. 2017, *ApJ*, 850, 149
- Tanaka M., Hotokezaka K., 2013, *ApJ*, 775, 113
- Tanaka M., Hotokezaka K., Kyutoku K., Wanajo S., Kiuchi K., Sekiguchi Y., Shibata M., 2014, *ApJ*, 780, 31
- Tonry J. L., 2011, *PASP*, 123, 58
- Wen L., Chen Y., 2010, *Phys. Rev. D*, 81, 082001
- Wu M.-R., Fernández R., Martínez-Pinedo G., Metzger B. D., 2016, *MNRAS*, 463, 2323

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.