

# Panning for gold with the *Neil Gehrels Swift Observatory*: an optimal strategy for finding the counterparts to gravitational wave events

R. A. J. Eyles-Ferris<sup>1</sup>,<sup>1\*</sup> P. A. Evans<sup>1</sup>, A. A. Breeveld,<sup>2</sup> S. B. Cenko,<sup>3,4</sup> S. Dichiara<sup>5</sup>, J. A. Kennea,<sup>5</sup> N. J. Klingler,<sup>3,6,7</sup> N. P. M. Kuin,<sup>2</sup> F. E. Marshall,<sup>3</sup> S. R. Oates<sup>8</sup>, M. J. Page,<sup>2</sup> G. Raman,<sup>5</sup> S. Ronchini,<sup>5,9</sup> M. H. Siegel,<sup>5</sup> A. Tohuvavohu,<sup>10,11,12</sup> S. Campana,<sup>13</sup> V. D’Elia,<sup>14,15</sup> D. H. Hartmann,<sup>16</sup> J. P. Osborne<sup>1</sup>, K. L. Page<sup>1</sup>, M. De Pasquale<sup>17</sup> and E. Troja<sup>18,19</sup>

<sup>1</sup>*School of Physics and Astronomy, University of Leicester, University Road, Leicester LE1 7RH, UK*

<sup>2</sup>*Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK*

<sup>3</sup>*Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA*

<sup>4</sup>*Joint Space-Science Institute, Computer and Space Sciences Building, University of Maryland, College Park, MD 20742, USA*

<sup>5</sup>*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA*

<sup>6</sup>*Center for Space Sciences and Technology, University of Maryland, Baltimore County, Baltimore, MD 21250, USA*

<sup>7</sup>*Center for Research and Exploration in Space Science and Technology, NASA Goddard Space Flight Center, Greenbelt, Maryland 20771, USA*

<sup>8</sup>*Physics Department, Lancaster University, Bailrigg, Lancaster LA1 4YB, UK*

<sup>9</sup>*Institute for Gravitation & the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA*

<sup>10</sup>*David A. Dunlap Department of Astronomy and Astrophysics, University of Toronto, Toronto M5S 3H7, Canada*

<sup>11</sup>*Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto M5S 3H7, Canada*

<sup>12</sup>*Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA*

<sup>13</sup>*INAF - Osservatorio Astronomico di Brera, Via Bianchi 46, I-23807 Merate (LC), Italy*

<sup>14</sup>*Space Science Data Center (SSDC) - Agenzia Spaziale Italiana (ASI), I-00133 Roma, Italy*

<sup>15</sup>*INAF - Osservatorio Astronomico di Roma, via Frascati 33, I-00040 Monte Porzio Catone, Italy*

<sup>16</sup>*Department of Physics and Astronomy, Clemson University, Clemson, SC 29634-0978, USA*

<sup>17</sup>*MIFT Department, University of Messina, Vi F. S. D’Alcontres 31, I-98166 Messina, Italy*

<sup>18</sup>*Dipartimento di Fisica, Università Degli Studi di Roma – Tor Vergata, via della Ricerca Scientifica 1, I-00100 Rome, Italy*

<sup>19</sup>*INAF – National Institute of Astrophysics, 00136 Rome, Italy*

Accepted 2024 December 18. Received 2024 December 17; in original form 2024 November 7

## ABSTRACT

The LIGO, Virgo, and KAGRA (LVK) gravitational wave observatories are currently undertaking their O4 observing run offering the opportunity to discover new electromagnetic counterparts to gravitational wave events. We examine the capability of the *Neil Gehrels Swift Observatory* (*Swift*) to respond to these triggers, primarily binary neutron star mergers, with both the UV/Optical Telescope (UVOT) and the X-ray Telescope (XRT). We simulate *Swift*’s response to a trigger under different strategies using model skymaps, convolving these with the 2MPZ catalogue to produce an ordered list of observing fields, deriving the time taken for *Swift* to reach the correct field and simulating the instrumental responses to modelled kilonovae and short gamma-ray burst afterglows. We find that UVOT, using the *u* filter with an exposure time of order 120 s, is optimal for most follow-up observations and that we are likely to detect counterparts in  $\sim 6$  per cent of all binary neutron star triggers detectable by LVK in O4. We find that the gravitational wave 90 per cent error area and measured distance to the trigger allow us to select optimal triggers to follow-up. Focussing on sources less than 300 Mpc away, or 500 Mpc if the error area is less than a few hundred square degrees, distances greater than previously assumed, offer the best opportunity for discovery by *Swift* with  $\sim 5$ –30 per cent of triggers having detection probabilities  $\geq 0.5$ . At even greater distances, we can further optimise our follow-up by adopting a longer 250 s or 500 s exposure time.

**Key words:** gravitational waves – methods: observational – gamma-ray bursts – neutron star mergers.

## 1 INTRODUCTION

GW 170817 was a true watershed moment in astronomy, the first time a gravitational wave (GW) detection was found to have an

electromagnetic (EM) counterpart (e.g. Abbott et al. 2017; Covino et al. 2017; Cowperthwaite et al. 2017; Evans et al. 2017; Goldstein et al. 2017; Kilpatrick et al. 2017; Savchenko et al. 2017; Smartt et al. 2017; Tanvir et al. 2017; Troja et al. 2017). However, at time of writing, it remains a unique event. The Laser Interferometer Gravitational-Wave Observatory (LIGO), the Virgo interferometer and Kamioka Gravitational Wave Detector (KAGRA), together

\* E-mail: [raje1@leicester.ac.uk](mailto:raje1@leicester.ac.uk)


referred to in this work as LVK, commenced their O4 run on 24 May 2023 offering the opportunity to finally make GW 170817 commonplace rather than unique. It is therefore crucial to devise an optimal strategy to follow-up GW detections and quickly discover any EM counterparts. Here, we build on the work of Evans et al. (2016a, hereafter [Paper I](#)) and Evans et al. (2016b, hereafter [Paper II](#), see also Evans et al. 2019), which discuss the strategy for previous LVK runs. However, both works assumed that the afterglow from a short GRB would be the signal most likely to be detected by *Swift* and therefore concentrated on *Swift*'s X-Ray Telescope (XRT) as the discovery instrument. However, the afterglow of GW 170817 was not detected by the XRT<sup>1</sup> but did have a UV bright kilonova (Evans et al. 2017) detected by the UV/Optical Telescope (UVOT).

A short GRB and a kilonova are the two EM transients that are typically expected to accompany the merger of a binary neutron star (BNS) or neutron star-black hole (NSBH) system, in addition to the gravitational waves detectable by LVK. Short GRBs result in two signals, the prompt gamma-ray emission that emerges from within a relativistic jet (the burst itself) and a synchrotron powered broad-band afterglow resulting from jet interactions with the circumburst medium. However due to Doppler beaming, both these components are, at least initially, strongly directional, particularly at the wavelengths accessible to *Swift*. While structured jets or cocoon emission (Nakar & Piran 2017; Piro & Kollmeier 2018) do result in more isotropic emission, the relatively narrow opening angle of the jet in a short GRB ( $\sim 15^\circ$ , Fong et al. 2015) therefore means that only a very small fraction of the merger population detectable by LVK through the far more isotropic gravitational wave emission will have bright and detectable afterglows. Kilonovae, on the other hand, are powered by *r*-process nucleosynthesis in the merger ejecta and emit much more isotropically (e.g. Metzger 2019). While there is still significant viewing angle dependence due to the nature of the ejecta, a search focused on kilonova emission might therefore be most effective. In this work, therefore, we focus primarily on the UVOT but also apply our analysis to the XRT.

In this paper, we examine the UVOT's capability in detecting a counterpart from a GW trigger during O4. We concentrate primarily on BNS mergers as their resultant kilonovae are significantly brighter and easier to detect than those expected from NSBH mergers (e.g. Gompertz et al. 2023), although we also briefly discuss these sources. Although we generally find the probability of a significant detection relatively low, we can improve our chances by using specific criteria such as the 90 per cent error area and the luminosity distance. This was already used to inform *Swift*'s strategy and this work has aided in updating these criteria and the strategy accordingly. Selecting events using these criteria significantly enhances the chance of UVOT promptly discovering an EM counterpart and ensures *Swift* remains effective across all aspects of its mission.

In Section 2, we detail the setup of our trigger simulations and their underlying data. We also describe *Swift*'s response to a trigger. Section 3 presents our modelling methodology for the afterglow and kilonova components of the light curve. We evaluate *Swift*'s expected performance and how to optimise follow-up in Section 4. We discuss additional factors that could impact follow-up in Section 5 and present our conclusions in Section 6. Unless otherwise mentioned, we adopt a cosmology with  $H_0 = 71 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ,  $\Omega_m = 0.27$ , and  $\Omega_\Lambda = 0.73$  throughout this paper.

<sup>1</sup>Although we note that *Swift* did not observe its position until  $\sim 14$  h after the trigger.



**Figure 1.** The distribution of luminosity distance for the simulation sample, normalised with the luminosity distance cubed.

## 2 SIMULATED TRIGGERS


To evaluate and optimise *Swift*'s observing plan, we simulated 3688 triggers. To summarise, from simulations of observed GW events and their associated skymaps, we seed the event in a host galaxy or random point across the sky. The skymap is then shifted and convolved with known galaxies to produce an ordered list of fields to be observed. Our full procedure is detailed in the rest of this section.

### 2.1 Gravitational wave simulations

To examine BNS mergers, we used the data set of Coughlin et al. (2022), which builds on the work and simulations of Singer et al. (2014, 2016). This consists of 8258 simulations of compact binary mergers, including both neutron stars (NSs) and black holes (BHs). Each of these simulations has an associated skymap derived using a range of instrument combinations based on the predicted duty cycle. The assumed sensitivity of each instrument is based on the predictions for O4. To limit our sample to mergers most capable of producing kilonovae, i.e. BNSs, we selected the simulations where the mass of each component was less than  $3M_\odot$ , a total of 922 simulations. Each simulation was used to simulate four triggers hence our 3688 total triggers.

We checked the properties of our sample of simulations and confirmed that, like the overall population, the mergers were randomly placed across the sky with an isotropic distribution. We found the mergers to be distributed uniformly in distance cubed out to  $\sim 350$  Mpc with a rapid decline at higher luminosity distances as shown in Fig. 1. The sensitivities of the observatories were expected to increase to 160–190 Mpc for LIGO-Hanford and Livingston and 80–115 Mpc for Virgo for the merger of two  $1.4 M_\odot$  NSs between O3 and O4. The simulation sample therefore covers a suitable range with higher distances attributable to the larger NS masses allowed. The inclinations of the systems have a bimodal distribution with peaks at  $\sim 30^\circ$  and  $\sim 150^\circ$ . This is due to the dependency of the GW signal on inclination and is consistent with the distribution expected due to the Malmquist bias (e.g. Schutz 2011).

The distributions of the masses of each component and their ratios are shown in Fig. 2. Component 1 is always more massive by definition and this leads to its relatively flat distribution and the low peak of component 2's distribution. Observationally, BNS systems within the Milky Way have been found to have mass ratios up to  $\sim 1.3$



**Figure 2.** The distributions of component masses (left) and mass ratio (right) for the simulation sample. Note that component 1 is, by definition, the more massive of the two.

(e.g. Tauris et al. 2017; Ferdman et al. 2020) and theory predicts ratios up to  $\sim 1.7$  (e.g. Wagg et al. 2022). The mass ratio distribution of the simulated triggers is somewhat broader than this. However, greater mass ratios may still be physically viable and they represent only a small minority within our sample. The masses and their ratios therefore cover a wide range of BNSs with diverse properties.

Overall, our sample is consistent with expectations for the BNS population and is ideal for exploring the possibilities of such mergers across O4.

## 2.2 Host selection

A real BNS merger will, rather than be positioned randomly in space, be associated with a host galaxy that may be known or unknown. We therefore used the Two Micron All Sky Survey Photometric Redshift catalogue (2MPZ; Bilicki et al. 2014) to determine if the source is placed in a catalogued galaxy and if it is, to select a suitable host for each simulated trigger to combine with its skymap. The completeness of the galaxy catalogue must therefore be taken into account when seeding a trigger into a galaxy. For 2MPZ, the completeness was derived in Paper II following Gehrels et al. (2016), by comparing the luminosity observed in the catalogue to that predicted and we summarise their method here.

The predicted integrated luminosity per unit volume of the local universe is described well by the integral of a Schechter function (Schechter 1976)

$$\frac{d\mathcal{L}}{dV} = \int_0^\infty \phi^* L^* x^{\alpha+1} e^{-x} dx, \quad (1)$$

where  $x = L/L^*$ ,  $\mathcal{L}$  is the total integrated luminosity across the universe,  $L^*$  is the characteristic galaxy luminosity at which the luminosity function steepens significantly,  $\phi^*$  is a number density and  $\alpha$  describes the slope for  $L < L^*$ . These last three parameters are all measurable from observations. For the IR selected 2MPZ, the absolute  $K$  band magnitudes ( $M^*$  for a galaxy with luminosity  $L^*$ ) are used to derive these parameters:  $M_K^* = -23.39 + 5 \log h$ ,  $\alpha_K = -1.09$ , and  $\phi_K^* = 0.0116h^3$  (Kochanek et al. 2001), where we have assumed  $h = H_0/100 = 0.7$ . Equation (1) can be further manipulated to derive the total luminosity within volume  $V$  as

$$\frac{\mathcal{L}}{L^*} = \int_0^\infty \phi^* x^{\alpha+1} e^{-x} dx V = \phi^* \Gamma(\alpha + 2, 0) V, \quad (2)$$

where  $\Gamma$  is the incomplete gamma function.<sup>2</sup>

Rather than the total IR magnitudes, 2MPZ contains the magnitude measured to the 20 mag arcsec<sup>-2</sup> isophote. We therefore corrected the catalogued values by  $\Delta M_K = -0.2$  in order to match the total magnitude (Kochanek et al. 2001; Bilicki et al. 2011). The Galactic extinction observed for each source must also be accounted for. At higher extinctions, there is a natural bias towards redder sources as bluer galaxies are more likely to be obscured. The colours and therefore photometric redshifts are accordingly less reliable. For 2MPZ, this results in the photometric redshift,  $z_{\text{phot}}$ , being underestimated by 0.004 compared to the spectroscopic redshift,<sup>3</sup>  $z_{\text{spec}}$  for sources with  $0.5 \leq E_{B-V} < 1.0$ . At even higher extinctions, this offset becomes greater and the redshift calibration becomes very unreliable. To alleviate these issues, the 9638 (1.0 per cent of the catalogue) sources with  $E_{B-V} > 1.0$  were cut from the sample. This also had the effect of eliminating sources affected by high stellar densities, which introduces an additional colour bias and affect the UVOT’s UV filters more significantly than optical wavelengths. This final sample was used to derive the completeness of 2MPZ with distance,  $C$ , shown in fig. 5 of Paper II. We note that Paper II did identify that a small discrepancy between  $z_{\text{phot}}$  and  $z_{\text{spec}}$ ,  $\sigma_{z_{\text{photo}}} = z_{\text{spec}} - z_{\text{photo}}$ , remained in this final sample, characterised as a broken power law

$$\sigma_{z_{\text{photo}}} = \begin{cases} 0.043 z_{\text{photo}}^{0.402} & z_{\text{photo}} < 0.10 \\ 0.023 z_{\text{photo}}^{0.140} & z_{\text{photo}} \geq 0.10. \end{cases} \quad (3)$$

However, because it is more likely that the photometric redshifts are underestimates rather than overestimates, the completeness is actually greater than calculated.

For a given trigger, therefore, the probability of it being in a known galaxy is simply the completeness at that distance i.e.  $P_{\text{gal}} = C$ . If probabilistically determined to be located in a known galaxy,<sup>4</sup> a

<sup>2</sup>We note that this is only valid for  $\alpha > -2$  and diverges for smaller values (Gehrels et al. 2016).

<sup>3</sup>In our chosen cosmology, this increases the luminosity distances by  $\sim 18$  Mpc.

<sup>4</sup>Due to high natal kick velocities, a trigger could occur a significant distance ‘outside’ its host galaxy (e.g. Mandhai et al. 2022). However, at the luminosity

catalogued galaxy is randomly selected to be the host weighted by the  $K$  band luminosity as a proxy for the galaxies' stellar masses. We also require the host to be at an appropriate distance and a similar declination to the original position of the trigger. The original skymap is then shifted across the sky to match the trigger's new location (as the real position of the source is known relative to the skymap). By maintaining a similar declination, we minimise the effect of the shift on the GW localisation due to the locations of the GW detectors. This means we effectively only change the (fictional) time that a given trigger was detected which has no impact on the follow-up strategy. For sources that were not selected to be in a host galaxy, their position and associated skymap were left unchanged from the original simulation.

### 2.3 Simulated follow-up

Having derived a trigger, we also model *Swift*'s planned response. This is again detailed in [Paper II](#) and its corresponding erratum (Evans et al. 2019), and we give an overview of it here. We note that this response is modelled as if it were a real trigger e.g. we do not know if the trigger is in a known host galaxy.

The GW skymaps are given in HEALPIX format and each individual pixel,  $p$ , in the skymap has a probability of the trigger occurring within it,  $P_{\text{GW},p}$ . Because a given trigger will be associated with a host galaxy this skymap can be convolved with a galaxy catalogue to more accurately capture the relative probability of each pixel. However, this means the completeness of the chosen galaxy catalogue also has a significant effect on the follow-up strategy as the trigger could be in an uncatalogued or a catalogued galaxy.<sup>5</sup> This means the effective probability for a given pixel is

$$P_{\text{GW},p} = P_{\text{nogal},p} + P_{\text{gal},p}, \quad (4)$$

where  $P_{\text{nogal},p}$  and  $P_{\text{gal},p}$  are the probabilities of the event occurring in an uncatalogued and a catalogued galaxy within  $p$ , respectively.

Each pixel in the skymap has its own probability distribution of the distance to the trigger,  $D$ . Consequently, the completeness of each pixel is

$$C_p = \frac{\int P_p(D)C(D)dD}{\int P_p(D)dD}, \quad (5)$$

where  $P_p(D)$  is the probability distribution of distance in  $p$ . The probability of the GW event occurring in an uncatalogued galaxy within a pixel is therefore simply


$$P_{\text{nogal},p} = P_{\text{GW},p}(1 - C_p). \quad (6)$$

The probability of the trigger occurring in a catalogued galaxy in pixel  $p$  is slightly more complex. Essentially,  $P_{\text{gal},p}$  is the probability of the pixel in the skymap multiplied by the fraction of total galaxy luminosity within the volume defined by the pixel and its distance distribution. It is given by

$$P_{\text{gal},p} = P_{\text{GW},p}C_pN \sum_g \left( \mathcal{P}(g|P_p(D)) \frac{L_g}{L_{\text{tot}}} \right), \quad (7)$$

distances involved, this offset is relatively insignificant and can therefore be neglected in our analysis.

<sup>5</sup>There is also a possibility of chance alignment between a galaxy and the trigger. However, the chance is extremely low, particularly due to the fact that the distance is also constrained, and thus this effect is likely to be negligible and therefore ignored in our analysis.



**Figure 3.** The cumulative distribution of simulated triggers against the field in which the source is contained.

where the sum is over all galaxies in the pixel,  $L_g$  is the luminosity of the galaxy divided by the number of pixels it covers,  $N$  is a normalization factor such that  $\sum_p \mathcal{P}_{\text{gal},p} = C$ , and  $L_{\text{tot}}$  is the total catalogued galaxy luminosity within the volume defined by the skymap and distance distribution. This means that  $\sum_g \frac{L_g}{L_{\text{tot}}}$  is effectively the relative probability of the galaxies in a specific pixel compared to any other pixel's galaxies.  $N$  and  $L_{\text{tot}}$  are given by

$$N = \frac{\sum_p P_{\text{GW},p}C_p}{\sum_p \left( P_{\text{GW},p}C_p \sum_g \left( \mathcal{P}(g|P_p(D)) \frac{L_g}{L_{\text{tot}}} \right) \right)} \quad (8)$$

and

$$L_{\text{tot}} = \sum_p \sum_g \mathcal{P}(g|P_p(D))L_g, \quad (9)$$

where  $\mathcal{P}(g|P_p(D)) = \int P_p(D)P_g(D)dD$  is the probability that a given galaxy  $g$  resides at an appropriate distance and  $P_g(D)$  is the probability distribution of  $g$ 's distance. Because 2MPZ does not include errors on its redshifts, we assumed these distributions to be Gaussians. For sources with photometric redshifts, the error was taken from equation (3) while for spectroscopic redshifts, the peculiar velocity was assumed to be the dominant source of error. Assuming a characteristic velocity of  $500 \text{ km s}^{-1}$  gives a  $\sigma$  of  $500/H_0 = 7.4 \text{ Mpc}$  (where  $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$  has been assumed) to assign to the Gaussian.

The convolution of the skymap and 2MPZ results in a new probability map. Following [Paper I](#), this map can be divided into fields, containing a large number of pixels from the convolved map, that *Swift* would observe in probability order. We can then model *Swift*'s progress over these fields to determine the time at which it will reach the field containing the trigger's position,  $t_{\text{reach}}$ . We simulate only to the 6000th field – if *Swift* has not yet reached the transient, it can safely be assumed it would not be discovered as this will be several days after the trigger and any transient fades very rapidly. In [Fig. 3](#), we plot the distribution of the field which contains the source finding that 33.2 per cent of triggers lie within the first 1000 fields while 50.5 per cent of triggers are reached within the full 6000 fields simulated. This is somewhat lower than expected as in theory we use the same parameters to order galaxies when both seeding and searching. However, to ensure the skymaps are realistic relative to the positions of the LVK detectors, the seed galaxy is on roughly the same declination as the trigger in the original sky map. This means

the seed galaxy is selected from a subset of galaxies rather than the much larger full population that is searched over and could introduce a small systematic bias that accounts for the proportion of triggers reached. It is also possible that flattening in the mass function of 2MPZ galaxies means that this is actually a purely stochastic effect; however, it is difficult to quantify which effect is more significant and whether one dominates over the other. We also ignore triggering criteria at this point – including these will significantly improve these fractions as discussed below.

To make our simulations as realistic as possible, we selected trigger times for each event and calculated the visibility windows of each field. Using realistic models of *Swift*'s slewing capability we then created an optimised observing plan, thus obtaining the times at which each field would be observed in a real-world scenario,  $t_{\text{reach}}$ .

There are three components to  $t_{\text{reach}}$ : the time for *Swift* to start observing, the cumulative time for it to slew between each field, and the cumulative exposure time to observe each field. Previously, in Paper I, it was assumed that the time for *Swift* to start observing,  $t_{\text{start}}$ , was best drawn from a Lorentz distribution based on the follow-up of IceCube neutrino triggers (Evans et al. 2015). This distribution was indeed similar to the real values during O2 (Klingler et al. 2019) and O3 (Page et al. 2020; Oates et al. 2021). This delay primarily came from the time for the observing command upload to arrive at *Swift*. However, the *Swift* Mission Operations Centre have developed a system of ‘continuous commanding’ over the Tracking and Data Relay Satellite System (TDRSS), removing the need for a ground-station pass and dramatically reducing the time between the trigger and the observing command upload (Tohuvavohu et al. 2024). We therefore drew  $t_{\text{start}}$  from a normal distribution

$$t_{\text{start}} \sim \mathcal{N}(3600, 1800) \text{ s} \quad (10)$$

with a lower limit of 1200 s to account for the processing of the skymap and construction of the list of fields.

The time for *Swift* to slew has also changed since the previous LVK runs. Following a reaction wheel failure in early 2022 (Cenko 2022), the response speed of *Swift* is reduced and we therefore approximate the time to slew between two fields as

$$t_{\text{slew}} = 25 + 2\delta \text{ s}, \quad (11)$$

where  $\delta$  is the angular separation in degrees between the two fields.

Finally, the exposure time,  $t_{\text{exp}}$ , for each field is controlled as part of the follow-up strategy. The exposure times and how they should vary in time to optimise *Swift*'s follow-up are examined in Section 4.2.


For a given trigger and the field it is contained within, the time to reach it is therefore

$$t_{\text{reach}} = t_{\text{start}} + \sum_s t_{\text{slew}} + \sum_e t_{\text{exp}} \quad (12)$$

summing over all slews,  $s$ , and exposures,  $e$ . We do note that our follow-up is likely to be interrupted to investigate candidates reported by other teams and  $t_{\text{reach}}$  increased accordingly. However, this is impossible to quantify and we therefore ignore this factor in the following analysis.

### 3 LIGHT-CURVE MODELLING

Having determined when *Swift* will reach and observe the trigger, we must next evaluate what the electromagnetic transient will look like at that time. We assume there to be two components to this transient, a kilonova and an afterglow from an accompanying short GRB. We model these components independently and combine them to produce a final combined count rate for the transient. For simplicity,



**Figure 4.** A sketch of the kilonova model assumed in N21. The colours indicate the viewing angles of the various components and the black point represents the merger remnant.

we assume the centre time of the observation to be representative of the transient’s behaviour through the observation.

#### 3.1 Kilonova component

Their tiny observed population mean that kilonovae remain poorly understood and models of their emission are diverse in behaviour (e.g. Rossi et al. 2022; Troja 2023). For instance, toy models with assumed ejecta velocity distributions and opacities can approximate much of the behaviour of a kilonova (Metzger 2019). More advanced radiation transfer models, on the other hand, can fully predict the angular dependence of luminosity and spectra (Kasen et al. 2017; Bulla 2019; Korobkin et al. 2021; Wollaeger et al. 2021; Bulla 2023). Arguably the most successful models, however, combine inferences and predictions from both the GW and EM data (Coughlin et al. 2019; Dietrich et al. 2020; Breschi et al. 2021). Here, we use the analytical BNS models of Nicholl et al. (2021, hereafter N21) which offer a distinct advantage over these options – they are designed to generate a light curve from only parameters inferred from a GW signal. These models are therefore ideal for use in our planning and during O4 itself.

The models are implemented in the MOSFiT v.1.1.8 package (Guillochon et al. 2018), specifically as the `bns_generative` model, building on and utilizing previous MOSFiT models of kilonovae (Cowperthwaite et al. 2017; Villar et al. 2017). The merger of the neutron stars produces ejecta components primarily consisting of dynamical ejecta and post-merger ejecta. From its primary inputs of the redshift, luminosity distance,  $D_L$ , chirp mass,  $M_c$ , mass ratio,  $q$ , and observation angle,  $\theta_{\text{obs}}$ , the model analytically predicts the masses, velocities, and opacities of these components. The models of Cowperthwaite et al. (2017) and Villar et al. (2017) are then used to predict the luminosities of each component. We show a sketch of the modelled kilonova in Fig. 4 (see N21 for a significantly more detailed version). For this work, the parameters above were taken

from the GW simulations, deriving  $\theta_{\text{obs}}$  from the binary inclination angle,  $\theta_{\text{inc}}$ . As the kilonova model assumes vertical symmetry, for  $\theta_{\text{inc}} < 90^\circ$ ,  $\theta_{\text{obs}} = \theta_{\text{inc}}$ , and for  $\theta_{\text{inc}} > 90^\circ$ ,  $\theta_{\text{obs}} = 180^\circ - \theta_{\text{inc}}$ .

The dynamical ejecta result from two interactions: the extreme tidal forces overcoming the NSs' self gravity as they orbit each other and material ejected by shocks at the contact interface. N21 use the prescription of Dietrich & Ujevic (2017) for the total mass of this material. This is derived from 172 numerical relativity simulations covering 21 equations of state and mass and compactness ranges of  $1 < M_i/M_\odot < 2$  and  $0.1 < C_i < 0.23$  where  $M_i$  is the mass of the  $i$ th NS,  $R_i$  is its radius and  $C_i \equiv GM_i/R_i c^2$  is its compactness (Bauswein, Goriely & Janka 2013; Hotokezaka et al. 2013; Dietrich et al. 2015, 2017; Lehner et al. 2016; Sekiguchi et al. 2016). The compactness typically depends on the equation of state but in the generative model it is fixed to the value for an NS with  $M = 1.4M_\odot$  and radius 10.7 km. While the mass range used in the `bns_generative` model does not include some of our higher mass simulated mergers, the results should still be accurate enough for our purposes. Dietrich & Ujevic found the total dynamical mass to follow

$$\frac{M_{\text{dyn}}}{10^{-3}M_\odot} = \left[ a \left( \frac{M_2}{M_1} \right)^{1/3} \left( \frac{1-2C_1}{C_1} \right) + b \left( \frac{M_2}{M_1} \right)^n + c \left( 1 - \frac{M_1}{M_1^*} \right) \right] M_1^* + \left[ a \left( \frac{M_1}{M_2} \right)^{1/3} \left( \frac{1-2C_2}{C_2} \right) + b \left( \frac{M_1}{M_2} \right)^n c \left( 1 - \frac{M_2}{M_2^*} \right) \right] M_2^* + d \quad (13)$$

where  $M_i^* = M_i + 0.08M_i^2$  is the baryonic mass of the  $i$ th NS and  $a$ ,  $b$ ,  $c$ ,  $d$ , and  $n$  are fitted parameters. The velocity of the ejecta is given by

$$v = \left[ a_v \left( \frac{M_1}{M_2} \right) (1 + c_v C_1) \right] + \left[ a_v \left( \frac{M_2}{M_1} \right) (1 + c_v C_2) \right] + b_v, \quad (14)$$

where  $a_v$ ,  $b_v$ , and  $c_v$  are fitted parameters with different values in the polar and orbital plane directions.

The dynamical ejecta varies in electron fraction according to its geometry. The equatorial ejecta are dominated by tidal ejecta and have a low electron fraction  $Y_e < 0.25$ . The material is therefore more capable of producing heavy  $r$ -process nuclei with high opacity (Barnes & Kasen 2013) and hence is termed the red component. In the polar direction, shocks and neutrino heating mean the material has a significantly higher electron fraction  $Y_e > 0.25$  and can therefore only produce lower mass  $r$ -process nuclei. The resultant lower opacity leads to bluer emission from this component. In N21's model, the UV/optical opacities of these components are fixed to  $\kappa_{\text{red}} = 10$  and  $\kappa_{\text{blue}} = 0.5 \text{ cm}^2 \text{ g}^{-1}$  following Radice et al. (2018). The ratio between the masses of these red and blue components has been found to depend strongly on the mass ratio between the progenitor NSs (Radice et al. 2018) and N21 use the simulations of Sekiguchi et al. (2016) to derive the relative masses of each component in each simulated kilonova (e.g. in their fig. 2).

In addition to dynamical ejecta, kilonovae can also be driven by material ejected after the merger driven by winds from the merger remnant and its accretion disc. The merger remnant can take the form of a stable NS, a supermassive NS that eventually collapses to a black hole, or could immediately collapse. NS remnants are predicted to produce larger disc masses with a higher electron fraction and therefore bluer emission (Metzger & Fernández 2014; Lippuner

**Table 1.** The parameter space drawn from to generate the kilonova light curve where  $\mathcal{N}(\mu, \sigma^2)$  refers to a normal distribution and  $\mathcal{U}(\text{min}, \text{max})$  refers to a uniform distribution.

Parameter	Value
Redshift	From GW simulation
$D_L$	From GW simulation
$\theta_{\text{obs}}$	From GW simulation
$M_c$	From GW simulation
$q$	From GW simulation
$\epsilon_{\text{disc}}$	$\mathcal{N}(0.20, 0.03^2)$
$\alpha$	$\mathcal{U}(0.1, 1.0)$
$\theta_{\text{cocoon}}$	$\mathcal{U}(0^\circ, 45^\circ)$

et al. 2017). The mass of the disc is found to follow

$$\log_{10}(M_{\text{disc}}) = \max \left[ -1, a \left( 1 + b \tanh \left( \frac{c - M_{\text{tot}}/M_{\text{th}}}{d} \right) \right) \right], \quad (15)$$

where  $a$ ,  $b$ ,  $c$ , and  $d$  are fitted parameters,  $M_{\text{tot}} = M_1 + M_2$ . The threshold mass for a prompt collapse to a black hole is then given by

$$M_{\text{th}} = \left( 2.38 - 3.606 \frac{M_{\text{TOV}}}{R_{1.4}} \right) M_{\text{TOV}}, \quad (16)$$

where  $M_{\text{TOV}}$  is the Tolman–Oppenheimer–Volkoff mass (the upper mass limit for a non-rotating NS) and  $R_{1.4}$  is the radius of a  $1.4 M_\odot$  NS. Based on N21, we use  $M_{\text{TOV}} = 2.17M_\odot$  and  $R_{1.4} = 10.7 \text{ km}$ , which yields  $M_{\text{th}} = 3.58M_\odot$ .

A fraction of the disc mass, given by  $\epsilon_{\text{disc}}$  and varying from 0.16 to 0.23, is ejected by viscously driven winds. For our simulations, we therefore draw it from a suitable normal distribution (see Table 1). N21 derive the final mass of this component from  $M_{\text{tot}}$ ,  $M_{\text{TOV}}$ , and  $M_{\text{th}}$  following Metzger (2019). The material ejected from the disc is defined as a ‘purple’ component (marked as  $M_{\text{purple}}$  in Fig. 4) and has an intermediate UV/optical opacity of  $\kappa_{\text{purple}}$ , in turn derived from the simulations of Lippuner et al. (2017) and the  $Y_e - \kappa$  relations of Tanaka et al. (2020) and found to range from 1.5 to  $5.5 \text{ cm}^2 \text{ g}^{-1}$  depending on the remnant's lifetime.

Winds, specifically those magnetically or neutrino driven, can also induce an additional enhancement in the blue ejecta (e.g. Metzger, Thompson & Quataert 2018; Radice et al. 2018; Ciolfi & Kalinani 2020) if the remnant does not promptly collapse, i.e.  $M_{\text{tot}} < M_{\text{th}}$ . This is implemented in the model using the parameter  $\alpha$ , defined as the ratio between the dynamical blue ejecta and total blue ejecta, i.e.  $M_{\text{blue,dyn}}/M_{\text{blue,tot}}$ . Setting  $\alpha < 1$  therefore provides the enhancement. Of the 922 simulations, 357 (37.8 per cent) have a total merger mass less than this threshold and we draw a value for  $\alpha$  from a uniform distribution between 0.1 and 1.0.

The geometry of all these different components leads to a strong angular dependence. N21 build on the prescription of Villar et al. (2017) using that of Darbha & Kasen (2020). The kilonova is modelled as a sphere with conical polar caps with half-opening angles of  $\theta_{\text{open}} = 45^\circ$ . For a given  $\theta_{\text{obs}}$ , the luminosity of each component is simply scaled using the projected area of the caps for the blue and purple ejecta or remaining sphere for the red ejecta.

The luminosities are derived using the MOSFIT module of Cowperthwaite et al. (2017) and Villar et al. (2017). The ejecta expand homologously with a grey opacity and a velocity  $v_{\text{ej}}$ . The temperature of the photosphere is therefore initially calculated from the Stefan–Boltzmann law until the ejecta expand sufficiently to cool and begin to recombine with recombination temperature  $T_{\text{rec}} = 2500 \text{ K}$ . The luminosity of each component is scaled proportionally to its mass.

A final possible emission component derives from the short GRB itself. As the jet expands, it can shock-heat the surrounding ejecta and form a cocoon. Following Piro & Kollmeier (2018) and Nakar & Piran (2017), N21 assume the shock deposits constant energy per decade of velocity in the ejecta and therefore the luminosity is proportional to the mass of the shocked ejecta. This is a fraction  $\theta_{\text{cocoon}}^2/2$  of the total ejecta and for our simulations we draw  $\theta_{\text{cocoon}}$  from a uniform distribution between  $0^\circ$  and  $45^\circ$ .

Between the GW simulations and drawing from our selected distributions, we therefore explore a diverse sample of possible kilonovae and effectively evaluate strategies for follow-up. The parameters we use for the simulated light curves are summarised in Table 1.

### 3.2 Afterglow component

There is also the possibility of significant afterglow contribution to the light curves at UV wavelengths. This, however, has a much stronger dependency on  $\theta_{\text{obs}}$  than the kilonova component.

We evaluate the afterglow contribution using the semi-analytical code AFTERGLOWPY v.0.7.3 (Ryan et al. 2020) which uses the single shell approximation of van Eerten, Zhang & MacFadyen (2010) and van Eerten (2018). In this model, the ejecta, contact discontinuity and forward shock are integrated over and treated as a single radially uniform fluid element. To interpolate between the ultra-relativistic and non-relativistic regimes, a trans-relativistic equation of state is employed (van Eerten 2013; Nava et al. 2013). The final model is calibrated to the BOXFIT code (van Eerten, van der Horst & MacFadyen 2012). This uses full high resolution 2D numerical hydrodynamics simulations of jets to derive afterglows.

Traditionally, GRB jets are assumed to be ‘top hat’ jets, with roughly constant energy across the jet and a precipitous decline at its edge. However, it is likely that jets have significantly more structure and AFTERGLOWPY includes several different possibilities. For instance, both Gaussian and power law structures are modelled. In our case, we assume the jets to have a Gaussian structure, as was likely the case for GW 170817 (e.g. Lamb et al. 2019; Salafia et al. 2019; Troja et al. 2019; Ryan et al. 2020), and therefore the energy of the jet varies with angle  $\theta$  as

$$E(\theta) = E_0 \exp\left(-\frac{\theta^2}{2\theta_c^2}\right), \quad (17)$$

where we have taken  $\theta_c$  as the half opening angle of the jet<sup>6</sup> i.e.  $\theta_c = \theta_j/2$  and the truncation angle as  $\theta_w = 4\theta_c$  to a maximum of  $45^\circ$ . To model a structured jet, AFTERGLOWPY breaks it down into many top hat components. The properties of the blast wave for each of these components, primarily the radial position of the shock, dimensionless four-velocity and time-dependent jet opening angle, are described by a set of ordinary differential equations. These are numerically integrated and solved and the top hat components are summed over to derive the resultant afterglow of the GRB.

For each simulation, some input parameters to the afterglow model can be taken from the GW simulation, primarily the observing conditions. The remaining parameters of the short GRB powering the afterglow are drawn from the distributions observed in the *Swift* SGRB population (Fong et al. 2015). This parameter space is

<sup>6</sup>We note that this is a slightly different assumption than that made in Ryan et al. (2020) and may result in the afterglow contribution being slightly overestimated. However, the effect of this is minimal as the vast majority of sources are still dominated by kilonova emission in the  $u$  band.

**Table 2.** The parameter space drawn from to generate afterglow contributions to the light curve where  $\mathcal{N}(\mu, \sigma^2)$  refers to a normal distribution. Note that  $\theta_j$  is the jet opening angle.

Parameter	Value
Redshift	From GW simulation
$D_L$	From GW simulation
$\theta_{\text{obs}}$	From GW simulation
$\log(E_{\gamma, \text{iso}})$	$\mathcal{N}(51.3, 0.7^2)$ erg
$\theta_j$	$\mathcal{N}(16, 10^2)^\circ$
$P$	$\mathcal{N}(2.37, 0.25^2)$
$\log(n_0)^a$	$\begin{cases} P = 0.57, \mathcal{N}(-3.81, 0.43^2) \\ P = 0.43, \mathcal{N}(-0.02, 0.92^2) \end{cases}$ cm <sup>-3</sup>
$\epsilon_e$	0.1
$\epsilon_B$	0.01

<sup>a</sup> $\log(n_0)$  has a bimodal distribution, which we model as two normal distributions with relative probabilities of being drawn from of 0.57 and 0.43.

summarised in Table 2. Again, drawing from a range of parameters allows a properly diverse sample of afterglow contributions to be evaluated and our observing strategy to be fully optimised.

### 3.3 *Swift*'s instrumental response


The next step is to combine the individual components from our kilonova and afterglow modelling in Section 3 and calculate the responses from *Swift*'s XRT and UVOT. For the XRT, we assume the contribution from any kilonova would be negligible and therefore only the afterglow emission would be relevant. The range of spectra we model are similar enough that a simple conversion from the 1 keV flux density calculated by AFTERGLOWPY yields sufficiently representative count rates (CR) in the full XRT band. This conversion, which assumes  $N_{\text{H}} = 3 \times 10^{20}$  cm<sup>-2</sup> and  $\Gamma = 2$ , gives for the flux density in Jy,

$$\text{CR}_{0.3-10 \text{ keV}} = S_{1 \text{ keV}} \times 2.07 \times 10^5 \text{ ct s}^{-1}. \quad (18)$$

For the UVOT, we calculate the kilonova and afterglow components for the  $u$ ,  $b$ ,  $v$ ,  $uvw1$ ,  $uvm2$ , and  $uvw2$  bands. We convert the AB magnitudes computed by AFTERGLOWPY and MOSFIT to count rates using each filter's AB zero-point (Breeveld et al. 2011) and apply an additional calibration factor  $\lesssim 1$  to account for the decrease in sensitivity of UVOT over *Swift*'s life.<sup>7</sup> To derive the total number of counts for each simulation, we multiply the total count rate for both the kilonova and the afterglow components by the exposure time.

It is relatively straightforward to determine if a source is likely to be detected by the XRT, with a threshold of six counts. However, it is more complicated for UVOT as the background is significantly greater and can vary by more than a factor of 2, depending on Galactic latitude, ecliptic latitude, and Earth or Sun limb angle. To fully model the UVOT's response, we use example images of exposure times approximately equal ( $\pm 15$  s) to those used in our strategy. These exposures are selected from those far outside the Galactic plane to avoid diffuse Galactic background or Zodiacal contamination, and are also chosen so that we had a ‘median’ background and ‘high’ background to test. Artificial sources are injected into the UVOT images with their total number of counts taken from the results of our simulations. This is repeated 25 times for each simulation then the UVOTDETECT utility was used to extract the sources. When running UVOTDETECT, zerobkg is set to  $-1$ , which calculates the local

<sup>7</sup>[https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/swift/docs/uvot/uvotcaldb\\_throughput\\_07.pdf](https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/swift/docs/uvot/uvotcaldb_throughput_07.pdf)



**Figure 5.** The *uvm2* (yellow), *u* band (blue), *b* band (purple), *v* band (red) light curves of an example kilonova at 100 Mpc, driven by the merger of two  $1.4 M_{\odot}$  NSs. It is viewed at three observation angles,  $0^{\circ}$  (solid),  $45^{\circ}$  (dashed), and  $90^{\circ}$  (dotted).

background for all points in the image, and the detection threshold is set to  $2\sigma$  as these are the typical parameters used by the UVOT team. The fraction of the 25 injected artificial sources recovered is then taken to be the detection probability.

#### 4 OPTIMISING *Swift*'S OBSERVING STRATEGY

There are a number of factors we can control in terms of *Swift*'s observing strategy. In terms of the observations themselves, it is important to select the filter and exposure time correctly. It is also crucial to use *Swift* efficiently for all science, not just looking for EM counterparts to LVK detections, in particular using *Swift*'s unique UV capability. We must therefore ensure we focus on those triggers where *Swift* can be effective. There are three pieces of information that LVK release that can help us make this decision – the object-type probabilities, the distance to the event, and the area of 90 per cent error region. In this section, we explore each factor that affects our strategy in turn.

##### 4.1 Filter selection

Kilonovae are primarily thermally driven and peak at relatively low temperatures. The vast majority of events will be seen at high observation angles and therefore be very ‘red’, with fast colour evolution away from blue, as seen in Fig. 5, although there will also be a significant fraction of events that will be viewed relatively ‘on-axis’ and therefore appear ‘blue’ for an appreciable period of time. The UVOT is optimised for bluer sources, likely to be the minority of triggers particularly at later times, and not wide areas. This means that ensuring *Swift* reaches the triggers on a rapid time-scale is more important than, for instance, a wide field optical instrument.

From our preliminary results, we find the *u*, *b*, *uvw1*, *uvm2*, and *uvw2* counts are broadly comparable at early times for various strategies, as shown in Fig. 6, and consistent with the example kilonova light curve in Fig. 5. In the *v* band, the counts are significantly lower. We therefore discount this band for our follow-up. We found this trend was consistent across all our simulations, regardless of the properties of the individual GW event.

We also needed to consider the background and therefore likely signal to noise in each filter. The background of the *b* and *v* bands

is significantly greater than that of the *u*, *uvw1*, *uvm2*, and *uvw2* bands (Breeveld et al. 2010). As the *u* band has a small apparent advantage in Fig. 6 and the *Swift* Gravitational Wave Galaxy Survey (Tohuvavohu et al. 2019) will provide image subtraction templates to aid in identification of new sources, we therefore concentrate on it for our follow-up strategy. Finally, the NUV wavelengths covered by the *u* band ( $\sim 3500 \text{ \AA}$ ; Poole et al. 2008) cannot be covered by ground-based telescopes which are, in turn, better optimised than *Swift* to search in the visible filters. For the *u* band, the median and high backgrounds described in Section 3.3 are  $0.01 \text{ ct s}^{-1} \text{ pixel}^{-1}$  and  $0.015 \text{ ct s}^{-1} \text{ pixel}^{-1}$ . We only use the procedure detailed there for sources where the expected number of counts is at least 10, however, as fewer counts are essentially undetectable.

##### 4.2 Exposure time

The biggest factor in optimizing *Swift*'s strategy is the time taken observing each field – the exposure time. It impacts two things – the time it takes to get to the correct field (see equation 12) and how likely we are to detect the source when it is reached.

To simplify the process of responding to triggers and the required controlling, we concentrate on strategies with a fixed exposure time per field. These exposure times are 120 s, 250 s, 500 s, and 1000 s with 120 s being approximately the minimum feasible time for a *Swift* observation. We did briefly examine strategies where the exposure time varied based on the time since the trigger but found no significant differences compared with the simpler strategy.

###### 4.2.1 $t_{\text{reach}}$

From the list of fields for each simulation, we calculated the  $t_{\text{reach}}$  from equation (12) for each possible  $t_{\text{exp}}$ . The cumulative distribution for the simulations is shown in Fig. 7, where we set an upper limit of 72 h as by this time, any EM signals are too faint for *Swift* to detect in the exposure times being probed here.

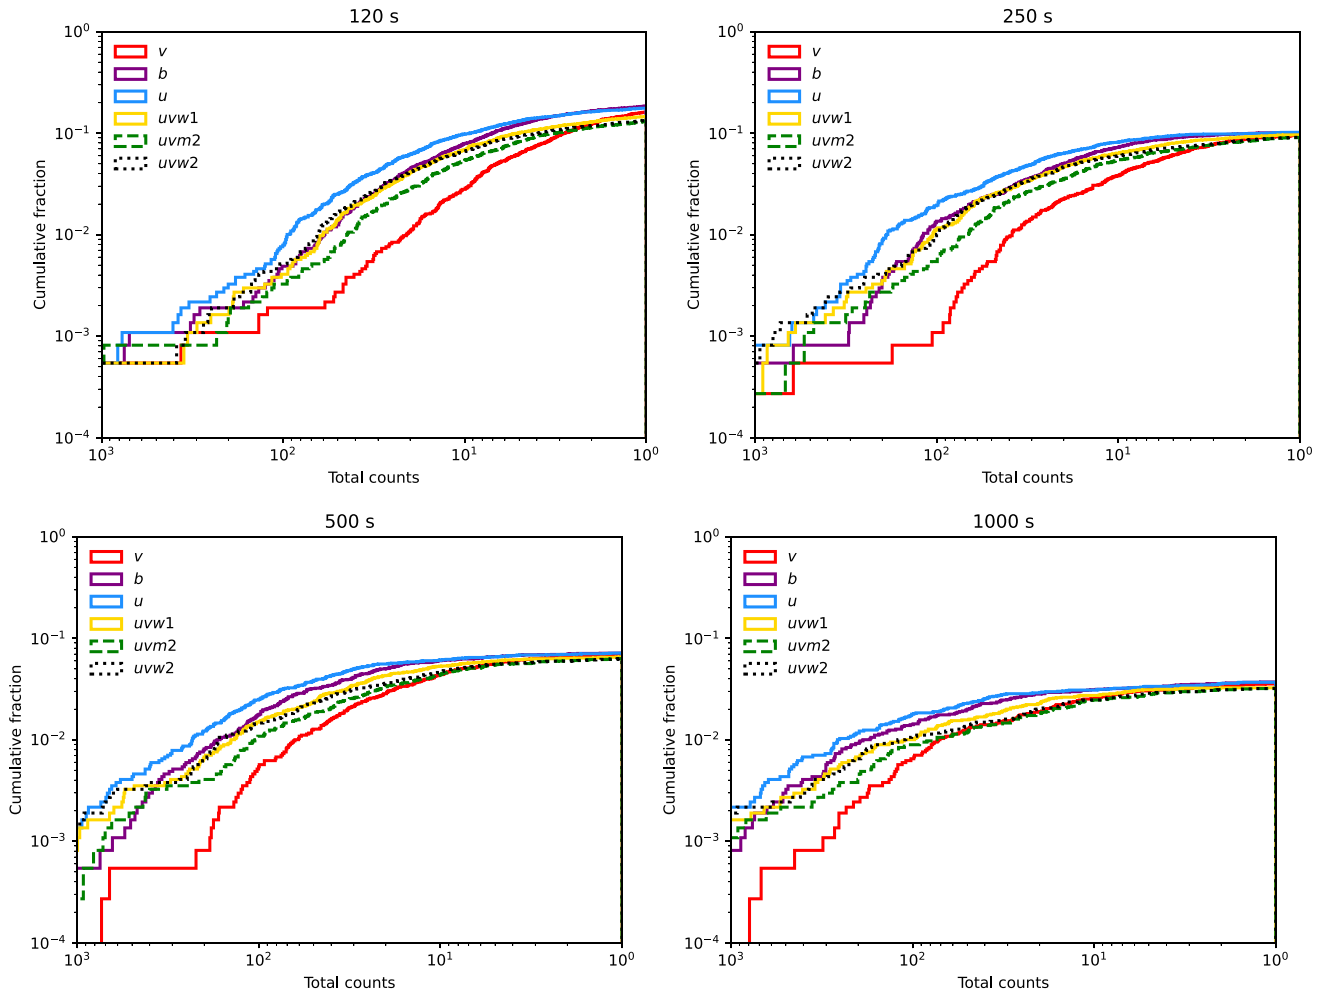
Unsurprisingly, the 120 s exposure strategy reaches the correct fields most rapidly achieving this within 24 h for 26.5 per cent of our simulations. However, the equivalent fractions for the 250 s, 500 s, and 1000 s are 21.3, 17.3, and 13.6 per cent, higher than might be expected if the exposure time was the dominant factor in  $t_{\text{reach}}$ . Instead, the total slew time dominates. We plot the distribution of individual slew times in Fig. 8. The majority of slews are reasonably short ( $\sim 50$  per cent are less than 60 s and  $\sim 80$  per cent are less than 150 s) but there is a significant fraction of slews that take hundreds of seconds. This is likely a result of skymaps with disparate regions of high probability. While the field-selection algorithm does de-weight long slews, maximizing the probability covered while accounting for observing constraints means they are sometimes necessary. However, because often thousands of fields must be observed before *Swift* reaches the correct location, the number of slews dominates over their individual lengths.

###### 4.2.2 Counterpart detection


We must also consider how likely each instrument is able to detect a given source. Their<sup>8</sup> sensitivities vary as approximately  $t_{\text{exp}}^{1/2}$ , hence, for slowly evolving light-curve components, the exposure time will have a greater impact than the time to reach the field containing

<sup>8</sup>We note that this is not necessarily the case for the XRT which tends to be photon dominated at these relatively short exposures and hence the sensitivity varies roughly as  $t$ .







**Figure 6.** Results from our simulations, the cumulative fraction of simulated triggers against total counts. The title of each panel refers to the exposure time for each observation i.e.  $t_{\text{exp}}$ . Note the scale on the  $x$ -axis goes from large to small.



**Figure 7.** The cumulative fraction of the 3688 simulated triggers where *Swift* reaches the correct field against  $t_{\text{reach}}$  for the first 72 h following the trigger for each exposure time.



**Figure 8.** The cumulative fraction of individual slews against their length.



**Figure 9.** The cumulative fraction of simulated triggers against detection probability in the  $u$  band for the median background (top) and higher background (bottom). Note the  $x$  axis goes from large to small and that the  $y$  axis is in log scale.

the source. In Fig. 9, we show the cumulative distribution of the 3688 simulated triggers against their detection probability. For both backgrounds, we find that 120, 250, and 500 s are roughly comparable with 1000 s significantly worse. This is probably because for the 1000 s strategy, the sources are inherently fainter having been reached at a later time, and therefore the signal to noise when observed is significantly lower. From our results, it appears that 120 s is the optimal time for the majority of triggers. We note that this is subject to small number statistics with only a relatively small sample of suitable simulated events.

Our results also show that for the 120 s strategy,  $\sim 5.7$  per cent of triggers have a detection probability  $\geq 0.5$ . Because it is impractical for *Swift* to follow-up every trigger due to its other observing duties, it is therefore crucial to focus our follow-up on those sources whose properties maximise our chance of detection.<sup>9</sup>

### 4.3 Trigger selection

The notices issued by LVK when a trigger is detected provide valuable information on whether or not to follow-up. For instance,

<sup>9</sup>We note that *Swift* does already have such criteria in place; however, we are refining and adding to these criteria here.

we can derive the probability that the event contains a disrupted NS:

$$P_{\text{DNS}} = \text{PROB\_REMNANT} \times (1 - \text{PROB\_TERRES}), \quad (19)$$

where `PROB_REMNANT` and `PROB_TERRES` are the probabilities of a merger remnant or that the trigger was terrestrial (noise).<sup>10</sup> These are taken directly from the notices. If  $P_{\text{DNS}}$  is low, the chance of an EM counterpart, particularly a bright one, is accordingly small and therefore only the best candidates are likely to be worth following up.

In addition, they include a skymap encoding both probability and distance information. The skymaps are processed by our pipeline and the 90 per cent probability area is extracted. From Fig. 10, it is clear that both probability area and distance have significant impacts on how likely *Swift* is to detect a counterpart. Here, we therefore examine in greater detail how the area and distance affect how the EM sources are recovered and which triggers we should follow-up.

#### 4.3.1 90 per cent error area

We binned the simulated skymaps based on the 90 per cent error area with bins of  $<150 \text{ deg}^2$ ,  $150\text{--}300 \text{ deg}^2$ ,  $300\text{--}500 \text{ deg}^2$ ,  $500\text{--}1000 \text{ deg}^2$ , and  $1000\text{--}5000 \text{ deg}^2$ .

In Fig. 11, we plot the cumulative fraction of simulated triggers against detection probability for each bin. We are again subject to small number statistics, however, for very well localised sources ( $\lesssim 150 \text{ deg}^2$ ), the fraction of sources that have detection probabilities greater than 0.5 is of the order of 20–25 per cent for both 120 and 250 s strategies. For larger areas ( $\lesssim 300 \text{ deg}^2$ ), there is a decrease but we might expect around 10 per cent of triggers to have detectable counterparts. The probability of *Swift* detecting a counterpart does continue to decline, however, there is a non-negligible chance of recovering a counterpart<sup>11</sup> even with a sky area of order  $1000 \text{ deg}^2$ . Our results do show that 120 s is an optimal strategy for sources with smaller 90 per cent error areas; however at  $\sim 300 \text{ deg}^2$ , the greater sensitivity achieved in a 250 s exposure means this strategy is more suitable.


During O4a, only the two LIGO detectors were online and as at least three detectors are required to achieve a localization better than  $\sim 1000 \text{ deg}^2$ , *Swift* did not follow up any triggers. This is a significant issue that affects all narrow-field telescopes including the UVOT. However, the addition of Virgo during O4b has greatly improved the localization. For instance, the candidate NSBH trigger S240422ed (Ligo Scientific Collaboration, VIRGO Collaboration & Kagra Collaboration 2024) had a 90 per cent error area of only  $250 \text{ deg}^2$  and was therefore followed up by *Swift*. We do note, however, that S240422ed has been subsequently downgraded and is most likely terrestrial. Nevertheless, it implies that more suitable triggers will be detected when more GW detectors are available.

#### 4.3.2 Distance

The distance to an event will also have a significant impact. While the ‘real’ distances to each simulated trigger are known, for actual events this will not be the case; we therefore investigate detection as a function of the calculated distance reported via the `DISTMEAN` keyword in the skymap.

<sup>10</sup>See <https://emfollow.docs.ligo.org/userguide/index.html> for details of these parameters.

<sup>11</sup>Although we note that other instruments are much more optimised to tackle a search of this nature.



**Figure 10.** The cumulative fraction of simulated triggers with detection probability  $\geq 0.5$  against 90 per cent error area (left panel) and measured  $D_L$  (right panel). The median and higher backgrounds are indicated with solid and dashed lines, respectively. Note the axes are in log scale.

In Fig. 12, we plot the cumulative fraction of triggers against detection probability for four measured distance bins:  $<150$  Mpc,  $150\text{--}300$ ,  $300\text{--}500$ , and  $500\text{--}1000$  Mpc. At lower distances ( $\lesssim 300$  Mpc) our results are consistent with an optimal  $t_{\text{exp}}$  of 120 s and we might expect 20–30 per cent of triggers resulting in detections. At greater distances up to  $\sim 500$  Mpc, a 250 s strategy appears more optimal but there is a significant decrease in the probability of a successful detection to a few per cent. Beyond this point the probability of a detection is very low and attempting to follow up these sources is inadvisable. We note that it appears that distance has a greater impact on *Swift*'s discovery ability than the sky localization. There is also likely to be a contributory factor from the completeness of the chosen galaxy catalogue. However, as shown by Paper II, 2MPZ is relatively complete over the distances relevant here and while newer catalogues have been released (e.g. GLADE+; Dálya et al. 2022), the additional completeness is not overly significant.

#### 4.3.3 Combined error area and distance

While independently, error area and distance can provide some information on whether to follow-up an event, it is necessary to examine both to fully assess the trigger. In particular, the error area and distance are related – a more nearby trigger is more likely to result in better localization.

In Fig. 13, we plot a grid of distance against the sky localization region size. We also compare the fraction of triggers with detection probability  $\geq 0.5$  for 120 and 250 s exposure times in Table 3. As our results above showed, the recovery fraction for triggers in the largest area and distance bins is negligible and we therefore do not examine it here. While each bin is especially susceptible to small number statistics here, we can draw some conclusions. In particular, there appears to be a  $\geq 25$  per cent chance of detecting a trigger's EM counterpart if it is within 150 Mpc irrespective of the error area. For luminosity distances between 150 and 300 Mpc, the chances are somewhat lower and decrease further for areas greater than  $300 \text{ deg}^2$ . At distances greater than 300 Mpc, our general strategy, particularly with larger error regions, appears less effective. However, increasing the exposure time to 250 or 500 s may allow us to follow-up these triggers effectively. The increased exposure time will counteract the much fainter nature of these counterparts. This effect was not seen

in Fig. 12, likely due to that covering all error areas. In addition, the increased time to reach the correct field for areas greater than  $1000 \text{ deg}^2$  means the source is undetectable when it is reached, regardless of the exposure time.

Overall, therefore, we can significantly boost our chances of discovering an EM counterpart by being selective about which events we follow-up and, depending on the parameters of the trigger, by careful choice of exposure time. Selecting to prioritise sources with a high  $P_{\text{DNS}}$  will assist with this. This will also ensure *Swift* remains available for other science.

## 5 DISCUSSION

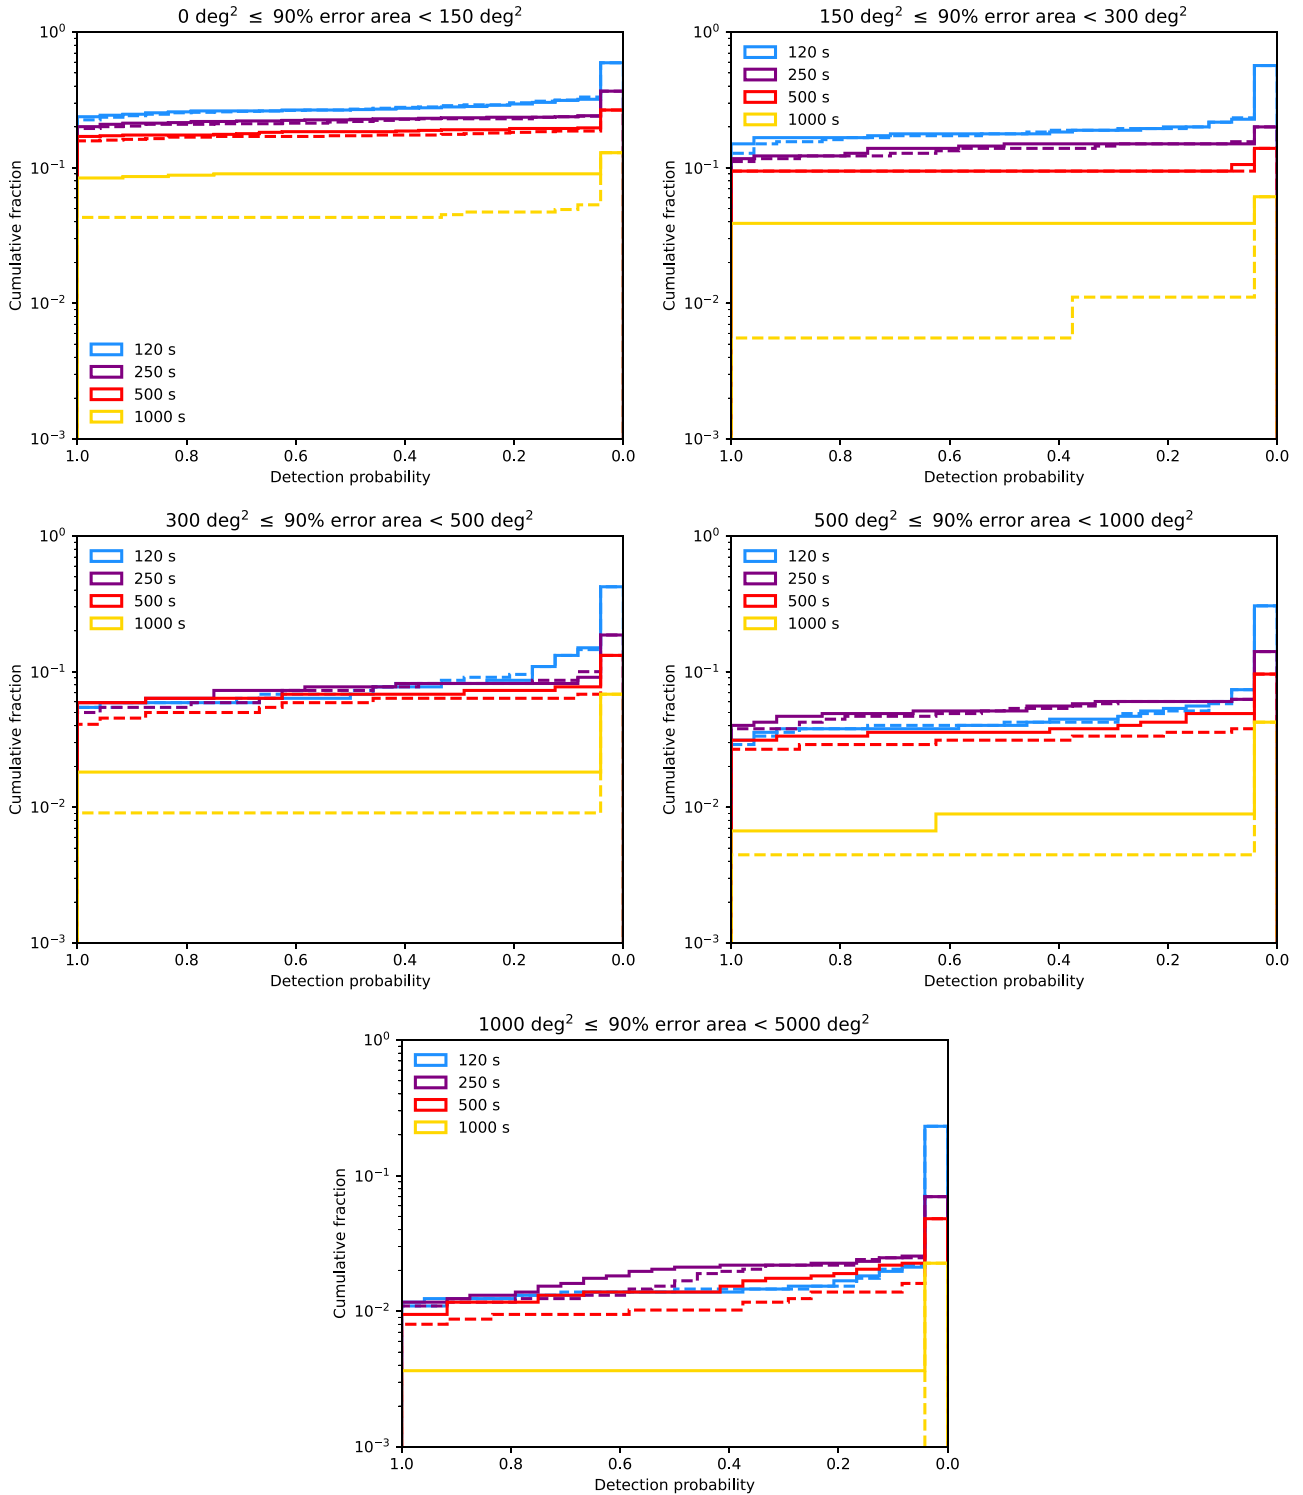
While our method is thorough and replicates the real response to notices, there are a few points that should be addressed. These include whether and how *Swift* should respond to NSBH GW triggers and what effect extinction might have on our results.

### 5.1 NSBH mergers

Significant EM emission is also expected from some NSBH merger in a similar way to a BNS merger. While in practice *Swift*'s strategy is unlikely to differ when following-up an NSBH trigger versus a BNS trigger, we briefly examine possible differences here.

MOSFIT has recently been updated to include a semi-analytic forward NSBH merger model, `nsbh_generative` (Gompertz et al. 2023). Similarly to the `bns_generative` model, light curves are predicted from the initial state of the binary and assumed equation of state of the NS. While we do not repeat our full analysis for this model, we can still draw significant conclusions on how *Swift* will be able to respond to these triggers.

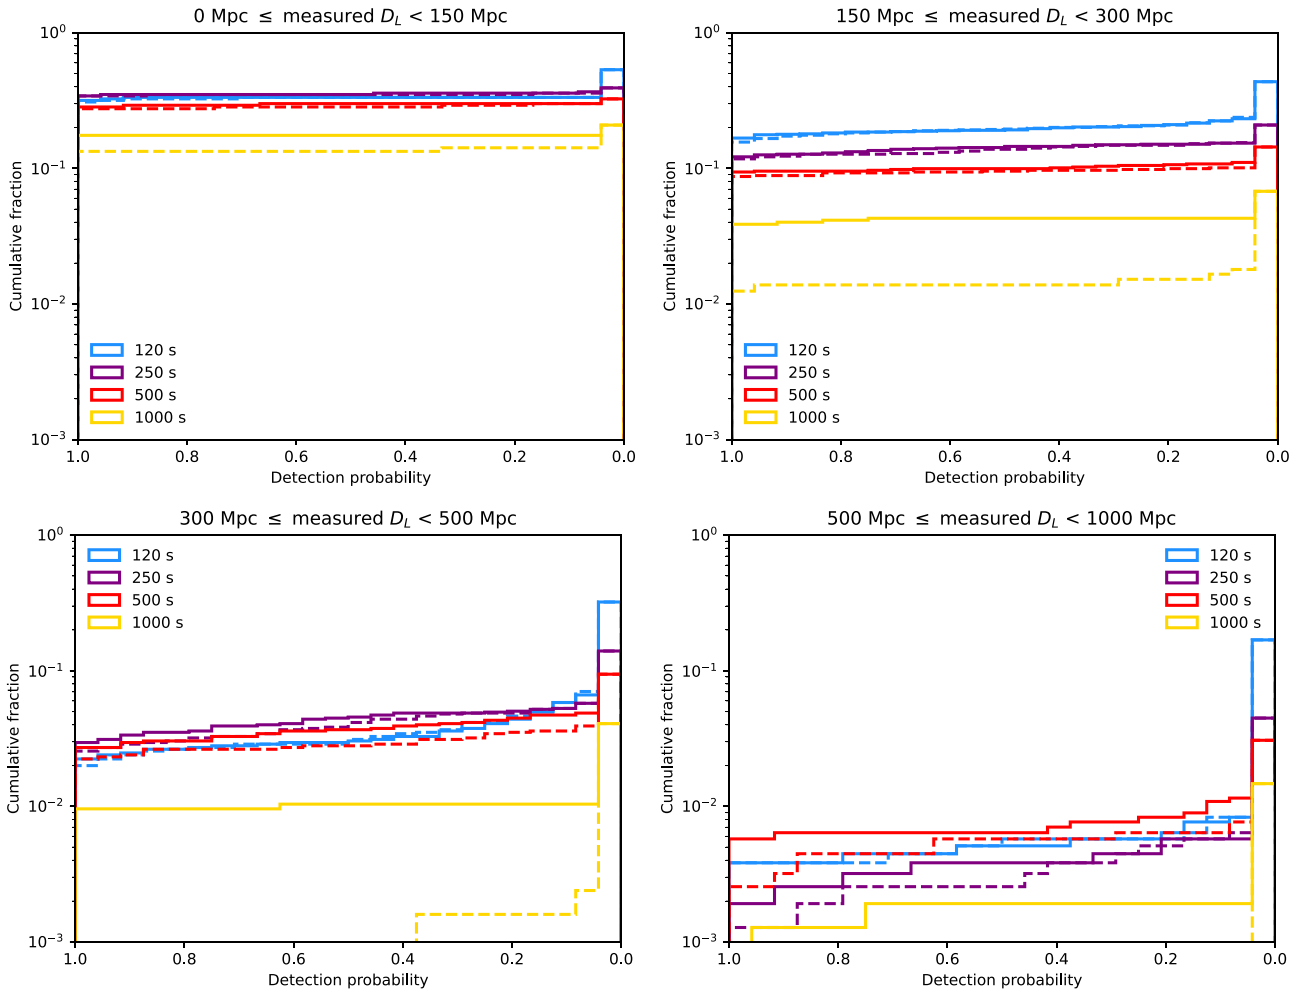
Like in a BNS merger, the ejecta in the NSBH merger can be divided into several components. These include tidal debris, similar to the BNS case, and two disc winds, one thermally driven and one magnetically driven. Both wind components have relatively low electron fractions with the thermal wind having a grey opacity and the magnetically driven wind being redder. A further significant difference between BNS and NSBH mergers is the impossibility of a post-merger NS remnant in the latter case. As shown in Section 3.1, a significant proportion of BNS mergers will leave such a remnant



**Figure 11.** The cumulative fraction of simulated triggers against detection probability in the  $u$  band for triggers with different 90 per cent error areas as indicated by the panel titles. The median and higher backgrounds are indicated with solid and dashed lines, respectively. Note the  $x$ -axis goes from large to small and that the  $y$ -axis is in log scale.

which drives a neutrino wind to increase the ejecta’s electron fraction and make the kilonova bluer. However, the effects of the magnetically driven wind and lack of possible remnant ensure the kilonova from an NSBH remains red with steady, little changing colours (see fig. 6 of Gompertz et al. 2023).

Such a source is even more poorly suited to the UVOT filters than a BNS kilonova. This can be seen in fig. 5 of Gompertz et al. (2023) which compares the peak magnitude of kilonovae from BNS and NSBH mergers. Even in the redder optical band presented there, NSBH kilonovae are much fainter by two to three magnitudes. This



**Figure 12.** The cumulative fraction of simulated triggers against detection probability in the  $u$  band for triggers with different measured distances as indicated by the panel titles. The median and higher backgrounds are indicated with solid and dashed lines, respectively. Note the  $x$ -axis goes from large to small and that the  $y$ -axis is in log scale.

effect is likely to be greater for *Swift*'s  $u$  filter, for instance, and will mean many NSBH mergers will be too faint to detect. Thus, while we will respond to NSBH triggers during O4, we expect our chances of discovery to be extremely low. The strategy devised for BNS mergers is likely to be most effective as difficulty in merger classification means that our best chance is hoping the NSBH is a BNS in disguise.

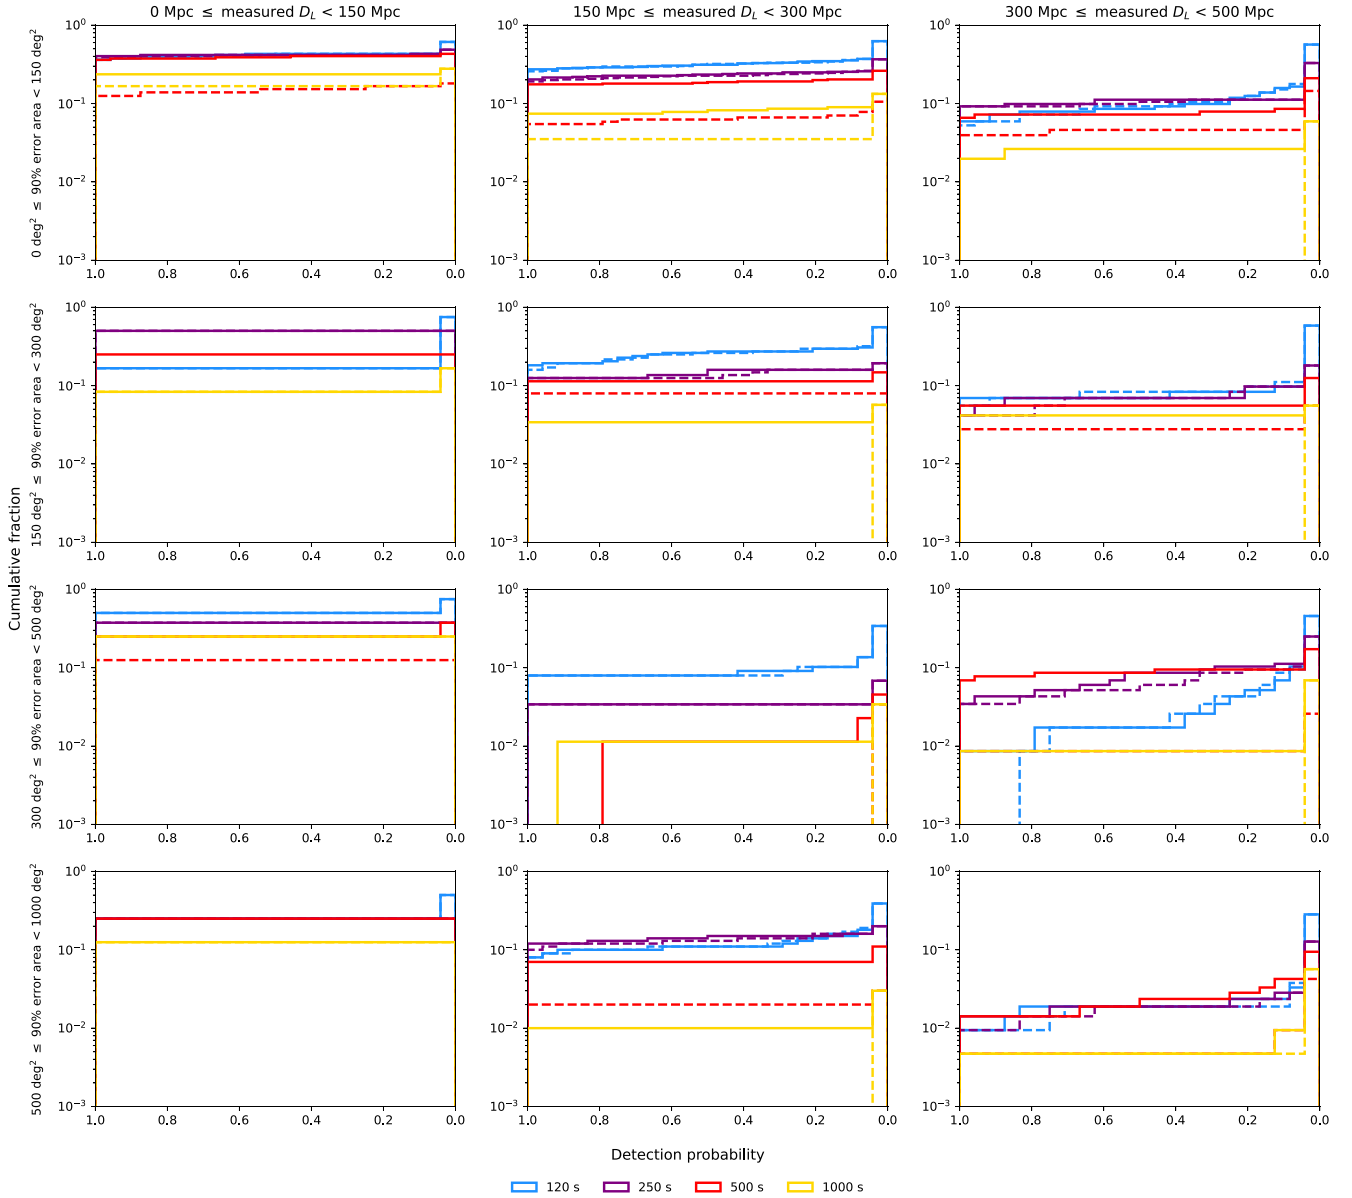
## 5.2 Binary black hole mergers

To this point we have only considered the gravitational wave sources detectable by LVK that are expected to produce significant EM emission: BNS and NSBH mergers. However, the vast majority of sources detected by LVK are binary black hole (BBH) mergers. Rather than the kilonova and short GRB afterglow emission components examined here, the possible EM emission from a BBH merger is far more unconstrained. Due to the inherent nature of a BBH system, EM emission can only be produced if the merger occurs in a dense environment, for instance, the disc of an active galactic nucleus (e.g. Kelly et al. 2017; McKernan et al. 2019). No confirmed EM counterpart to a BBH merger has thus far been discovered with only

a few examples of possible candidates (e.g. Connaughton et al. 2016; Graham et al. 2020, 2023; Cabrera et al. 2024, see also Greiner et al. 2016 and Connaughton et al. 2018). Due to the lack of constraints on these triggers' emission, *Swift* follows a different strategy to that explored in this work. Instead *Swift* follows-up only the best localised events with 90 per cent error areas  $\leq 30 \text{ deg}^2$  and using minimal length 80 s exposures. This ensures the area is covered rapidly and minimises disruption to *Swift*'s other activities. During O4, *Swift* has followed up several BBH mergers meeting the criterion above, with no likely EM counterpart candidates so far.

## 5.3 Extinction

Another important aspect to consider is the reddening from dust and other interstellar and intergalactic media. The wavelengths of the UVOT's filters mean that they will be affected to a large degree by such reddening and thus there will be a major impact on *Swift*'s ability to detect an EM counterpart. We can minimise Galactic extinction by avoiding the Galactic plane where possible. This is already accounted for in the tiling algorithm, which avoids the Galactic plane as much as possible. However, there is still the problem of extinction from the



**Figure 13.** The cumulative fraction of simulated triggers against detection probability in the  $u$  band for triggers with different 90 per cent error areas and measured luminosity distances as indicated by the row and column labels, respectively. The median and higher backgrounds are indicated with solid and dashed lines, respectively. Note the  $x$ -axis goes from large to small.

**Table 3.** The fraction of simulated triggers with detection probability  $\geq 0.5$  in each measured  $D_L$  and 90 per cent error area bin for 120 s/250 s/500 s exposure times and a median background. Note that these bins are highly vulnerable to small number statistics, particular at low measured  $D_L$  e.g. the 0–150 Mpc/300–500  $\text{deg}^2$  bin.

90% error area \ Measured $D_L$	Measured $D_L$		
	0–150 Mpc	150–300 Mpc	300–500 Mpc
0–150 $\text{deg}^2$	0.43/0.42/0.39	0.31/0.23/0.18	0.09/0.11/0.07
150–300 $\text{deg}^2$	0.17/0.50/0.25	0.26/0.14/0.11	0.07/0.07/0.06
300–500 $\text{deg}^2$	0.50/0.38/0.25	0.08/0.03/0.01	0.02/0.09/0.09
500–1000 $\text{deg}^2$	0.25/0.25/0.25	0.11/0.14/0.07	0.02/0.02/0.02

environment the merger took place in. This is impossible to predict, as a wide variety of environments and extinctions have been observed in the short GRB population (e.g. Fong et al. 2015). Because the effect of this extinction cannot be eliminated, it is best to maintain our

strategy without alterations. If another observatory were to identify an EM counterpart in a reddened environment, however, this would have to be accounted for in the targeted UVOT (or XRT, see Asquini et al. 2019) follow-up that would occur.

## 6 CONCLUSIONS

The discovery of EM counterparts to GW events is vital to increase our understanding of both fundamental physics and the impact these sources have on the universe, for instance, in terms of elemental abundances.

*Swift* has distinct advantages offered by its nature as a satellite. It does not need to wait for night and is, at present, the only way to catch the early blue and UV emission. Its rapid response offers a window unprobed by other observatories. However, the UVOT's relatively narrow field of view means often thousands of fields will need to be observed and the red nature of the predicted source may make it hard to discern. *Swift* is vitally important for many other fields of X-ray and UV astronomy and it is important to make use of it effectively.

We have therefore shown that we can maximise *Swift*'s efficiency in discovering counterparts to GW triggers by careful target selection. In particular, targeting only triggers with inferred distances  $\leq 300$  Mpc and a sky localization  $\leq 500$  deg<sup>2</sup> should result in  $\sim 25$  per cent of counterparts being discovered. Extending this to greater distances requires cutting the area to  $< 150$  deg<sup>2</sup>, which still reduces the proportion that may be detected significantly. For triggers with a low  $P_{\text{DNS}}$ , limiting our follow-up to only the closest and therefore brightest sources is also necessary. We further find the optimal exposure time to be 120 s, although this is dependent on the properties of the trigger, and that the  $u$  filter offers the best balance between source brightness and background. We note that strictly obeying these criteria does restrict follow-up to a minority of triggers. However, as detector sensitivity improves and particularly as Virgo and KAGRA begin observations in earnest, sky localization will improve dramatically and the fraction of triggers to which our criteria apply will significantly increase. This will continue as LVK approach their final sensitivity in future observing runs and while many of our conclusions will still apply, it will be important to revisit how they are affected by these changes.

In addition to improvements to the gravitational wave detector network, the launch of other missions, particularly UV focused, will have a significant impact on how GW events are followed up. Our results show that even a narrow field instrument like the UVOT can effectively follow-up GW triggers. This bodes extremely well for the discovery potential of sensitive, wide-field missions such as the *Ultraviolet Explorer* (UVEX; Kulkarni et al. 2021) or the *Ultraviolet Transient Astronomy Satellite* (ULTRASAT; Shvartzvald et al. 2024). Even with these missions operational; however, *Swift*'s UVOT will still be crucial for characterizing the spectral and temporal evolution of kilonovae, particularly at early times. Following up the most likely candidates from such missions may therefore be the best strategy in future LVK observing runs.

While to date O4 has been somewhat disappointing in terms of the triggers the universe has supplied, there is still a significant amount of time remaining and new EM bright events could occur at any time. In this work, we have ensured *Swift* is in an optimal position to respond to future events.

## ACKNOWLEDGEMENTS

We thank the referee for their helpful comments.

This work made use of data supplied by the UK *Swift* Science Data Centre at the University of Leicester; and the Wide Field Astronomy Unit (WFAU).

RAJEF acknowledges support from the UK Space Agency and the European Union's Horizon 2020 Programme under the AHEAD2020

project (grant agreement number 871158). ET acknowledges support by the European Research Council through the Consolidator grant BHianca (Grant agreement ID: 101002761).

## DATA AVAILABILITY

The simulated compact binary mergers are available at <https://zenodo.org/record/7026209#.Y90gIKfP1hE>.

The Two Micron All Sky Survey Photometric Redshift catalogue (2MPZ) is available at <http://ssa.roe.ac.uk/TWOMPZ.html>.

## REFERENCES

- Abbott B. P. et al., 2017, *ApJ*, 848, L12  
 Asquini L. et al., 2019, *A&A*, 625, A6  
 Barnes J., Kasen D., 2013, *ApJ*, 775, 18  
 Bauswein A., Goriely S., Janka H. T., 2013, *ApJ*, 773, 78  
 Bilicki M., Chodorowski M., Jarrett T., Mamon G. A., 2011, *ApJ*, 741, 31  
 Bilicki M., Jarrett T. H., Peacock J. A., Cluver M. E., Steward L., 2014, *ApJS*, 210, 9  
 Breeveld A. A. et al., 2010, *MNRAS*, 406, 1687  
 Breeveld A. A., Landsman W., Holland S. T., Roming P., Kuin N. P. M., Page M. J., 2011, in McEnery J. E., Racusin J. L., Gehrels N., eds, AIP Conf. Ser. Vol. 1358, Gamma Ray Bursts 2010. Am. Inst. Phys., New York, p. 373  
 Breschi M., Perego A., Bernuzzi S., Del Pozzo W., Nedora V., Radice D., Vescovi D., 2021, *MNRAS*, 505, 1661  
 Bulla M., 2019, *MNRAS*, 489, 5037  
 Bulla M., 2023, *MNRAS*, 520, 2558  
 Cabrera T. et al., 2024, preprint (arXiv:2407.10698)  
 Cenko B., 2022, GRB Coordinates Network, 31500, 1  
 Ciolfi R., Kalinani J. V., 2020, *ApJ*, 900, L35  
 Connaughton V. et al., 2016, *ApJ*, 826, L6  
 Connaughton V. et al., 2018, *ApJ*, 853, L9  
 Coughlin M. W., Dietrich T., Margalit B., Metzger B. D., 2019, *MNRAS*, 489, L91  
 Coughlin M. W., Farah A., Foley E., Gray A., Singer L. P., Weizmann R., 2022, LIGO/Virgo/KAGRA Observing Capabilities: Simulated Detections and Localization for O4 and O5 (October 2022 edition). Zenodo, <https://doi.org/10.5281/zenodo.7026209>  
 Covino S. et al., 2017, *Nat. Astron.*, 1, 791  
 Cowperthwaite P. S. et al., 2017, *ApJ*, 848, L17  
 Dálya G. et al., 2022, *MNRAS*, 514, 1403  
 Darbha S., Kasen D., 2020, *ApJ*, 897, 150  
 Dietrich T., Ujevic M., 2017, *Class. Quantum Gravity*, 34, 105014  
 Dietrich T., Bernuzzi S., Ujevic M., Brüggemann B., 2015, *Phys. Rev. D*, 91, 124041  
 Dietrich T., Ujevic M., Tichy W., Bernuzzi S., Brüggemann B., 2017, *Phys. Rev. D*, 95, 024029  
 Dietrich T., Coughlin M. W., Pang P. T. H., Bulla M., Heinzel J., Issa L., Tews I., Antier S., 2020, *Science*, 370, 1450  
 Evans P. A. et al., 2015, *MNRAS*, 448, 2210  
 Evans P. A. et al., 2016a, *MNRAS*, 455, 1522 (Paper I)  
 Evans P. A. et al., 2016b, *MNRAS*, 462, 1591 (Paper II)  
 Evans P. A. et al., 2017, *Science*, 358, 1565  
 Evans P. A. et al., 2019, *MNRAS*, 484, 2362  
 Ferdman R. D. et al., 2020, *Nature*, 583, 211  
 Fong W., Berger E., Margutti R., Zauderer B. A., 2015, *ApJ*, 815, 102  
 Gehrels N., Cannizzo J. K., Kanner J., Kasliwal M. M., Nissanke S., Singer L. P., 2016, *ApJ*, 820, 136  
 Goldstein A. et al., 2017, *ApJ*, 848, L14  
 Gompertz B. P., Nicholl M., Smith J. C., Harisankar S., Pratten G., Schmidt P., Smith G. P., 2023, *MNRAS*, 526, 4585  
 Graham M. J. et al., 2020, *Phys. Rev. Lett.*, 124, 251102  
 Graham M. J. et al., 2023, *ApJ*, 942, 99  
 Greiner J., Burgess J. M., Savchenko V., Yu H. F., 2016, *ApJ*, 827, L38

- Guillochon J., Nicholl M., Villar V. A., Mockler B., Narayan G., Mandel K. S., Berger E., Williams P. K. G., 2018, *ApJS*, 236, 6
- Hotokezaka K., Kiuchi K., Kyutoku K., Okawa H., Sekiguchi Y.-i., Shibata M., Taniguchi K., 2013, *Phys. Rev. D*, 87, 024001
- Kasen D., Metzger B., Barnes J., Quataert E., Ramirez-Ruiz E., 2017, *Nature*, 551, 80
- Kelly B. J., Baker J. G., Etienne Z. B., Giacomazzo B., Schnittman J., 2017, *Phys. Rev. D*, 96, 123003
- Kilpatrick C. D. et al., 2017, *Science*, 358, 1583
- Klingler N. J. et al., 2019, *ApJS*, 245, 15
- Kochanek C. S. et al., 2001, *ApJ*, 560, 566
- Korobkin O. et al., 2021, *ApJ*, 910, 116
- Kulkarni S. R. et al., 2021, preprint (arXiv:2111.15608)
- Lamb G. P. et al., 2019, *ApJ*, 870, L15
- Lehner L., Liebling S. L., Palenzuela C., Caballero O. L., O'Connor E., Anderson M., Neilsen D., 2016, *Class. Quantum Gravity*, 33, 184002
- LIGO Scientific Collaboration, VIRGO Collaboration, Kagra Collaboration, 2024, GRB Coordinates Network, 36236, 1
- Lippuner J., Fernández R., Roberts L. F., Foucart F., Kasen D., Metzger B. D., Ott C. D., 2017, *MNRAS*, 472, 904
- Mandhai S., Lamb G. P., Tanvir N. R., Bray J., Nixon C. J., Eyles-Ferris R. A. J., Levan A. J., Gompertz B. P., 2022, *MNRAS*, 514, 2716
- McKernan B. et al., 2019, *ApJ*, 884, L50
- Metzger B. D., 2019, *Living Rev. Relat.*, 23, 1
- Metzger B. D., Fernández R., 2014, *MNRAS*, 441, 3444
- Metzger B. D., Thompson T. A., Quataert E., 2018, *ApJ*, 856, 101
- van Eerten H., Zhang W., MacFadyen A., 2010, *ApJ*, 722, 235
- Nakar E., Piran T., 2017, *ApJ*, 834, 28
- Nava L., Sironi L., Ghisellini G., Celotti A., Ghirlanda G., 2013, *MNRAS*, 433, 2107
- Nicholl M., Margalit B., Schmidt P., Smith G. P., Ridley E. J., Nuttall J., 2021, *MNRAS*, 505, 3016
- Oates S. R. et al., 2021, *MNRAS*, 507, 1296
- Page K. L. et al., 2020, *MNRAS*, 499, 3459
- Piro A. L., Kollmeier J. A., 2018, *ApJ*, 855, 103
- Poole T. S. et al., 2008, *MNRAS*, 383, 627
- Radice D., Perego A., Hotokezaka K., Bernuzzi S., Fromm S. A., Roberts L. F., 2018, *ApJ*, 869, L35
- Rossi A. et al., 2022, *ApJ*, 932, 1
- Ryan G., van Eerten H., Piro L., Troja E., 2020, *ApJ*, 896, 166
- Salafia O. S., Ghirlanda G., Ascenzi S., Ghisellini G., 2019, *A&A*, 628, A18
- Savchenko V. et al., 2017, *ApJ*, 848, L15
- Schechter P., 1976, *ApJ*, 203, 297
- Schutz B. F., 2011, *Class. Quantum Gravity*, 28, 125023
- Sekiguchi Y., Kiuchi K., Kyutoku K., Shibata M., Taniguchi K., 2016, *Phys. Rev. D*, 93, 124046
- Shvartzvald Y. et al., 2024, *ApJ*, 964, 74
- Singer L. P. et al., 2014, *ApJ*, 795, 105
- Singer L. P. et al., 2016, *ApJ*, 829, L15
- Smartt S. J. et al., 2017, *Nature*, 551, 75
- Tanaka M., Kato D., Gaigalas G., Kawaguchi K., 2020, *MNRAS*, 496, 1369
- Tanvir N. R. et al., 2017, *ApJ*, 848, L27
- Tauris T. M. et al., 2017, *ApJ*, 846, 170
- Tohuvavohu A., Evans P., Kennea J., Cenko S. B., 2019, in American Astronomical Society Meeting Abstracts #233. p. 210.08
- Tohuvavohu A. et al., 2024, *ApJ*, 975, L19
- Troja E., 2023, *Universe*, 9, 245
- Troja E. et al., 2017, *Nature*, 551, 71
- Troja E. et al., 2019, *MNRAS*, 489, 1919
- van Eerten H., 2018, *Int. J. Mod. Phys. D*, 27, 1842002
- van Eerten H., van der Horst A., MacFadyen A., 2012, *ApJ*, 749, 44
- van Eerten H., 2013, preprint (arXiv:1309.3869)
- Villar V. A. et al., 2017, *ApJ*, 851, L21
- Wagg T., Broekgaarden F. S., de Mink S. E., Frankel N., van Son L. A. C., Justham S., 2022, *ApJ*, 937, 118
- Wollaeger R. T. et al., 2021, *ApJ*, 918, 10

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