Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **512,** 1127–1140 (2022) https://doi.org/10.1093/mnras/stac366
Advance Access publication 2022 February 11

A pixelated approach to galaxy catalogue incompleteness: improving the dark siren measurement of the Hubble constant

R. Gray [1,2]★ C. Messenger [2] and J. Veitch [2]

[1] *Department of Physics and Astronomy, Queen Mary University of London, Mile End Road, London E1 4NS, UK*
[2] *SUPA, School of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, UK*

Accepted 2022 February 8. Received 2022 February 4; in original form 2021 December 16

ABSTRACT

The use of gravitational wave standard sirens for cosmological analyses is becoming well known, with particular interest in measuring the Hubble constant, H_0, and in shedding light on the current tension between early- and late-time measurements. The current tension is over 4σ and standard sirens will be able to provide a completely independent measurement. Dark sirens (binary black hole or neutron star mergers with no electromagnetic counterparts) can be informative if the missing redshift information is provided through the use of galaxy catalogues to identify potential host galaxies of the merger. However, galaxy catalogue incompleteness affects this analysis, and accurate modelling of it is essential for obtaining an unbiased measurement of H_0. Previously most methods have assumed uniform completeness within the sky area of a gravitational wave event. This paper presents an updated methodology in which the completeness of the galaxy catalogue is estimated in a directionally dependent matter, by pixelating the sky and computing the completeness of the galaxy catalogue along each line of sight. The H_0 inference for a single event is carried out on a pixel-by-pixel basis, and the pixels are combined for the final result. A reanalysis of the events in the first gravitational wave transient catalogue leads to an improvement on the measured value of H_0 of approximately 5 per cent compared to the 68.3 per cent highest density interval of the equivalent LIGO and Virgo result, with $H_0 = 68.8^{+15.9}_{-7.8}$ km s^{-1} Mpc^{-1}.

Key words: gravitational waves – methods: data analysis – catalogues – cosmological parameters.

1 INTRODUCTION

The use of gravitational wave (GW) signals from the mergers of compact binaries (black holes or neutron stars) to constrain cosmological parameters has gained traction in recent years. The tension between early- and late-time measurements of the Hubble constant (H_0) remains above 4σ at the time of writing (Riess et al. 2019; Planck Collaboration VI 2020), making novel and independent measurements of H_0 of particular interest at present. GWs have the luminosity distance of the source directly encoded within them, without requiring external calibration, making them independent distance measures (Schutz 1986). Preferring the early-time measurements would indicate that the source of the tension lies in systematic errors between the different measurement techniques, while preferring the local, late-time measurement would indicate that the source of the tension lies in fundamental physics, and that the current cosmological model, lambda cold dark matter (ΛCDM), does not adequately fit our Universe to the level of accuracy at which it is now being measured.

The first three observing runs of Advanced LIGO and Virgo have provided over 50 GW detections (Abbott et al. 2019, 2021a,c). Of these, only one [the well-known binary neutron star (BNS) GW170817] has been observed with a confirmed electromagnetic counterpart (Abbott et al. 2017a,c). In order to use the remaining detections for a measurement of H_0, a different method of obtaining redshift information about the sources is required. Of particular interest is the *statistical* or *galaxy catalogue* dark siren method, in which galaxy catalogues are used to identify the galaxies within the localization volume of a GW event, and they are all treated as potential hosts (Schutz 1986; Del Pozzo 2012). The contribution from an individual event is less informative than in the counterpart case, but combining information from multiple events reduces the uncertainty and allows an additional constraint on H_0 to be made. [See, e.g. MacLeod & Hogan (2008), Chen, Fishbach & Holz (2018), Fishbach et al. (2019), Soares-Santos et al. (2019), Gray et al. (2020), Abbott et al. (2020), Palmese et al. (2020), Vasylyev & Filippenko (2020), and Finke et al. (2021). An alternative, cross-correlating GW events with galaxies of known redshift, is presented in Mukherjee et al. (2021)]. One important aspect of this method is acknowledging that galaxy catalogues are incomplete, and therefore may not contain the real host galaxy of the GW event. In order to account for this, an incompleteness correction must be applied. The galaxy catalogue methodology, as implemented in the GWCOSMO codebase (Gray et al. 2020) was applied to the GW detections from the first GW transient catalogue (GWTC-1) (Abbott et al. 2019) in Abbott et al. (2021b), leading to a measurement of $H_0 = 68.7^{+17.0}_{-7.8}$ km s^{-1} Mpc^{-1}.[1]

The result in Abbott et al. (2021b) showed proof of principle and highlighted many of the upcoming challenges within the field. The

★ E-mail: r.gray@qmul.ac.uk

[1] GWCOSMO is available at https://git.ligo.org/lscsoft/gwcosmo.

analysis also had several limitations. One of note was the assumption that galaxy catalogue completeness [due to the limited sensitivity of the telescope(s) carrying out the survey] is uniform with the sky area of each GW event. However, the GLADE catalogue (Dálya et al. 2018) [used to inform the redshift prior for the majority of the GWTC-1 events in Abbott et al. (2021b)] is a composite catalogue, made up of multiple different surveys, and as a result has highly variable completeness across the sky. Even for a galaxy catalogue that is made from a single survey, limited by a single telescope's sensitivity, the Milky Way band means that completeness cannot ever be truly uniform across the whole sky. Another approximation of note was that the sky and distance information from a GW event were treated as uncorrelated, rather than being treated as fully 3D. This was done in order to aid computational efficiency, at the expense of losing some constraining power from the GW events.

The issue of non-uniform catalogue incompleteness was addressed in Finke et al. (2021), in which areas of the sky with similar levels of completeness were grouped together, and a completeness correction was applied. The methods of completeness correction used in Finke et al. (2021) differ to the one used here (and in Gray et al. 2020; Abbott et al. 2021b), where the limiting magnitude of the telescope is used to assess the probability of any galaxy being 'seen'. Under the same set of assumptions these two methods should give consistent results; however, the methodology outlined here allows for a more complete treatment of the GW mass priors (see Section 3.1 for details), including consistency between the prior applied to individual events and the prior used to compute GW selection effects, which is necessary for an unbiased estimate of H_0.

The assumption of uniform catalogue completeness is problematic in two ways. When a uniform completeness correction is applied across a non-uniform patch of sky, there will be areas for which the completeness is overestimated, and areas for which it is underestimated. The contribution from galaxies in areas where it is overestimated will be artificially inflated, which risks biasing the result if the host galaxy is not inside the galaxy catalogue. Where it is underestimated, useful information about the redshift distribution of galaxies at higher redshifts will be diluted unnecessarily by the completeness correction.

Fortunately, there is an extension to the method which addresses both of these limitations at once. This extension entails pixelating the sky into equally sized pieces, and analysing each independently, using a line-of-sight (LOS) distance distribution for the GW event within each pixel, and an estimate of the completeness within that section of the sky. Both of the approximations described above are removed, making the analysis more robust and, theoretically, more informative.

This paper builds upon the work in Gray et al. (2020) to allow variations in galaxy catalogue completeness to be robustly estimated. Section 2 outlines the Bayesian implementation of the pixelated method. Section 3 discusses the practicalities of implementing it in the GWCOSMOcodebase and applying it to real data. Section 4 reanalyses the GWTC-1 detections using the pixelated method, while keeping all other assumptions the same as in Abbott et al. (2021b), and quantifies the improvement to the results. Section 5 discusses potential extensions for the pixelated method in the future, and concludes the paper.

2 METHODOLOGY

A measurement of H_0 can be made using a set of $N_{\rm det}$ detected GW events. Expressing this in a Bayesian form, the posterior probability on H_0 can be written as follows:

$$p(H_0|\{x_{\rm GW}\},\{D_{\rm GW}\},I) \propto p(H_0|I)p(N_{\rm det}|H_0,I) \times \prod_i^{N_{\rm det}} p(x_{{\rm GW}i}|D_{{\rm GW}i},H_0,I). \quad (1)$$

Here, $\{x_{\rm GW}\}$ is a set of GW data, and each $x_{\rm GW}$ corresponds to a GW detection (denoted by $D_{\rm GW}$). In general, an event is deemed detected if some detection statistic associated with $x_{\rm GW}$ passes a threshold, such as the signal-to-noise ratio (SNR) of the event passing some SNR threshold for the GW detector network. The term $p(H_0|I)$ is the prior on H_0, and $p(N_{\rm det}|H_0,I)$ is the probability of detecting $N_{\rm det}$ events for a given value of H_0 [intrinsically linked to the astrophysical rate of events, R, but independent of H_0 when a prior of $1/R$ is used (Fishbach, Holz & Farr 2018)]. The final term is a product over $N_{\rm det}$ individual event likelihoods. The parameter I is a placeholder which contains any additional information which is not explicitly stated (such as the underlying cosmological and population models).

In Gray et al. (2020), when considering the likelihood for a single GW event, the first step was to marginalize over the probability that the host galaxy of the event is, or is not, inside the catalogue. Here, it is first necessary to consider the likelihood's dependence on sky direction, Ω. This can be written as follows:

$$p(x_{\rm GW}|D_{\rm GW},H_0,I) = \int p(x_{\rm GW},\Omega|D_{\rm GW},H_0,I)\,{\rm d}\Omega. \quad (2)$$

Instead of considering the continuous variable Ω, it is necessary to switch to a discrete approximation: that of splitting the sky into $N_{\rm pix}$ equally sized pieces, which are later summed. Doing so, equation (2) becomes

$$\begin{aligned} &p(x_{\rm GW}|D_{\rm GW},H_0,I) \\ &= \sum_i^{N_{\rm pix}} p(x_{\rm GW}|\Omega_i,D_{\rm GW},H_0,I)p(\Omega_i|D_{\rm GW},H_0,I), \\ &= \sum_i^{N_{\rm pix}} p(x_{\rm GW}|\Omega_i,D_{\rm GW},H_0,I)\frac{p(D_{\rm GW}|\Omega_i,H_0,I)p(\Omega_i|H_0,I)}{p(D_{\rm GW}|H_0,I)}. \end{aligned} \quad (3)$$

It is possible to further simplify this expression if the probability of detection is assumed to be uniform over the sky – a reasonable assumption, as it is averaged over the full length of an observing run, and so the rotation of the Earth blurs out much of the sky dependence (Chen et al. 2017). The term $p(D_{\rm GW}|\Omega_i,H_0,I)$ loses its dependence on Ω_i and cancels with the denominator. Making this approximation ignores the mild declination dependence that the probability of detection retains, which is expected to have only a minor impact on the result (a full investigation of which is left for the future). Acknowledging that $p(\Omega_i|H_0,I)$ is independent of H_0 in an isotropic universe gives the following:

$$\begin{aligned} p(x_{\rm GW}|D_{\rm GW},H_0,I) &= \sum_i^{N_{\rm pix}} p(x_{\rm GW}|\Omega_i,D_{\rm GW},H_0,I)p(\Omega_i|I), \\ &= \frac{1}{N_{\rm pix}}\sum_i^{N_{\rm pix}} p(x_{\rm GW}|\Omega_i,D_{\rm GW},H_0,I), \end{aligned} \quad (4)$$

where the equally sized pixels mean that $p(\Omega_i|I)$ becomes a constant corresponding to the fraction of the surface area of a sphere which one pixel covers. It is worth noting that any pixels with zero GW support will evaluate to zero, and so the sum over $N_{\rm pix}$ can be reduced to a sum over $N_{\rm GWpix}$ with no impact to the result (for clarity, the pre-factor would remain $1/N_{\rm pix}$ in this case).

Now that the likelihood has been pixelated, it can be marginalized over the possibility of the host being inside (G) or outside ($\bar{G}$) the galaxy catalogue. This leads to the likelihood within a single pixel taking the following form:

$$\begin{aligned} &p(x_{\rm GW}|\Omega_i, D_{\rm GW}, H_0, I) \\ &\quad = p(x_{\rm GW}|\Omega_i, G, D_{\rm GW}, H_0, I)p(G|\Omega_i, D_{\rm GW}, H_0, I) \\ &\qquad + p(x_{\rm GW}|\Omega_i, \bar{G}, D_{\rm GW}, H_0, I)p(\bar{G}|\Omega_i, D_{\rm GW}, H_0, I), \end{aligned} \tag{5}$$

Here, $p(x_{\rm GW}|\Omega_i, G, D_{\rm GW}, H_0, I)$ is the likelihood when the host galaxy is inside the catalogue, for which the redshift prior will consist of the galaxies in the galaxy catalogue which lie within pixel i. The redshift uncertainty of each galaxy is assumed to be Gaussian, the standard deviation of which is provided by the galaxy catalogue, and is marginalized over. The term $p(x_{\rm GW}|\Omega_i, \bar{G}, D_{\rm GW}, H_0, I)$ is the likelihood when the host galaxy is *not* inside the catalogue, for which the redshift prior will be, in the simplest case, an uninformative uniform in co-moving volume distribution. Alternatively, a more complex GW rate evolution model, which allows the merger rate of binaries to be redshift dependent, can be assumed. These likelihoods are weighted by the probability that the host galaxy is inside the catalogue or outside it (assuming it lies in pixel i), which is determined using the apparent magnitude threshold ($m_{\rm th}$) of that pixel, as well as an assumption about the luminosity distribution of galaxies in the Universe – that they follow, e.g. a Schechter function (Schechter 1976). See the appendix of Gray et al. (2020) for more details.

The value of $m_{\rm th}$ can now be chosen on a by-pixel basis, allowing varying catalogue incompleteness to be accurately taken into account, modulo the choice of pixel size. There is also the possibility that some pixels will be empty (contain no galaxies due to e.g. obscuration by the Milky Way band), in which case equation (5) simplifies to the 'empty catalogue' case in which $m_{\rm th} \to -\infty$, leading to $p(G|\Omega_i, D_{\rm GW}, H_0, I) = 0$ and $p(\bar{G}|\Omega_i, D_{\rm GW}, H_0, I) = 1$. The contribution of that pixel then comes fully from the $p(x_{\rm GW}|\Omega_i, \bar{G}, D_{\rm GW}, H_0, I)$ term.

Additionally, for each pixel $p(x_{\rm GW}|\Omega_i, H_0, I)$ can be approximated as the GW information corresponding to the patch of sky covered by pixel i, meaning there is no longer a requirement to separate GW sky and distance information. The distance posterior along the line of sight of each pixel can be estimated, making this method inherently '3D'.

3 PIXELATING GWCOSMO: PRACTICALITIES

In order to implement the pixel-based method described in Section 2, HEALPY, a PYTHON implementation of HEALPIX which handles pixelated data on a sphere, is used (Górski et al. 2005; Zonca et al. 2019).[2] It allows the user to split the sky into equally sized pixels, for a choice of resolutions, where the resolution is set by a parameter called nside. The total number of pixels the sky is divided into is set by $12 \times \text{nside}^2$, where nside is must be a power of 2. At its lowest resolution, the sky is divided into 12 pixels of equal area. A one-step increase in resolution (which corresponds to doubling nside) divides each pixel into four further pixels.

In GWCOSMO's case, HEALPY allows data points with known right ascension (RA) and declination (dec.) (e.g. GW posterior samples, or galaxies) to be uniquely associated with a specific pixel. As increasing or decreasing the resolution by one step involves either dividing one existing pixel into 4, or combining four pixels into one, this opens up the possibility of combining different resolution HEALPY maps within the same analysis, without the danger of double-counting data. This is useful because variations in the LOS distance estimate for a GW across its sky area, and variations in the $m_{\rm th}$ of a galaxy catalogue across the same area, are not (necessarily) on the same scale.

Take, for example, the GLADE 2.4 catalogue (Dálya et al. 2018). It is a composite catalogue, made up of many different surveys. Those surveys cover different patches of the sky, and have overlapped in some areas. Taking a HEALPY nside of 32 divides the catalogue into 12 288 pixels, each covering approximately 3.36 deg^2. The top panel of Fig. 1 visualizes this, and shows how the number density of galaxies varies on that scale. Many sharp features can be seen, and are adequately represented by this resolution, though they could be better represented with a higher one (e.g. nside = 64, the lower left panel), while a lower resolution starts to blur out useful information (e.g. nside = 16, the lower right panel). However, looking at Table 1 that summarizes the GWTC-1 binary black hole (BBH) detections, it can be seen that GW151226, a relatively nearby event with reasonable catalogue support, has a 90 per cent sky area of 1033 deg^2. So even a conservative analysis that only considered the 90 per cent most probable sky area would still require 308 pixels to represent it with an nside of 32. To cover the 99.9 per cent sky area would push this to $\mathcal{O}(1000)$. The computational time for this analysis increases approximately linearly with the number of pixels, meaning that it could take 1000 times longer to analyse the same event as it did when using the original Gray et al. (2020) method, or would require 1000 computer cores to analyse it on a similar time-scale, if the analysis was parallelized. However, dividing the GW sky area into 1000 pixels is an unnecessarily high resolution for adequately representing the changes in the event's distance distribution.

The method of choosing which resolution to treat the GW data with, and computing the LOS distance distributions for it, is discussed in Section 3.1. The method of combining GW data with a galaxy catalogue represented by a higher resolution is presented in Section 3.2.

3.1 Line-of-sight luminosity distance estimates

The GW data used for this cosmological analysis come in the form of posterior samples, which are drawn from the posterior distributions of GW parameters for each event, including the luminosity distance, sky location, and detector frame masses. As was discussed in Abbott et al. (2021b), it is important to re-weight the GW posterior samples in order to remove the detector-frame mass prior used for parameter estimation, and apply the desired source-frame mass prior in order to match the priors used to compute the normalizing probability of detection term. This importance arises due to the cosmological information that is inherently encoded with a GW detection, in the relation between source-frame and detector-frame mass. The intrinsic mass of a merger cannot be measured directly, but requires converting detector-frame (redshifted) masses to source-frame, which is H_0-dependent. This cosmological information can be particularly informative (with enough BBH detections) when the source frame mass distribution of BBHs contains sharp features such as peaks or cut-offs (Farr et al. 2019; Mastrogiovanni et al. 2021). The re-weighting of an event's posterior samples not only affects the shape of the event's distance distribution, but also its normalization as a function of H_0, as samples may become inconsistent with the source-frame mass prior for certain values of H_0. As such, when implementing the pixelated method into GWCOSMO, the GW mass

[2] http://healpix.sf.net

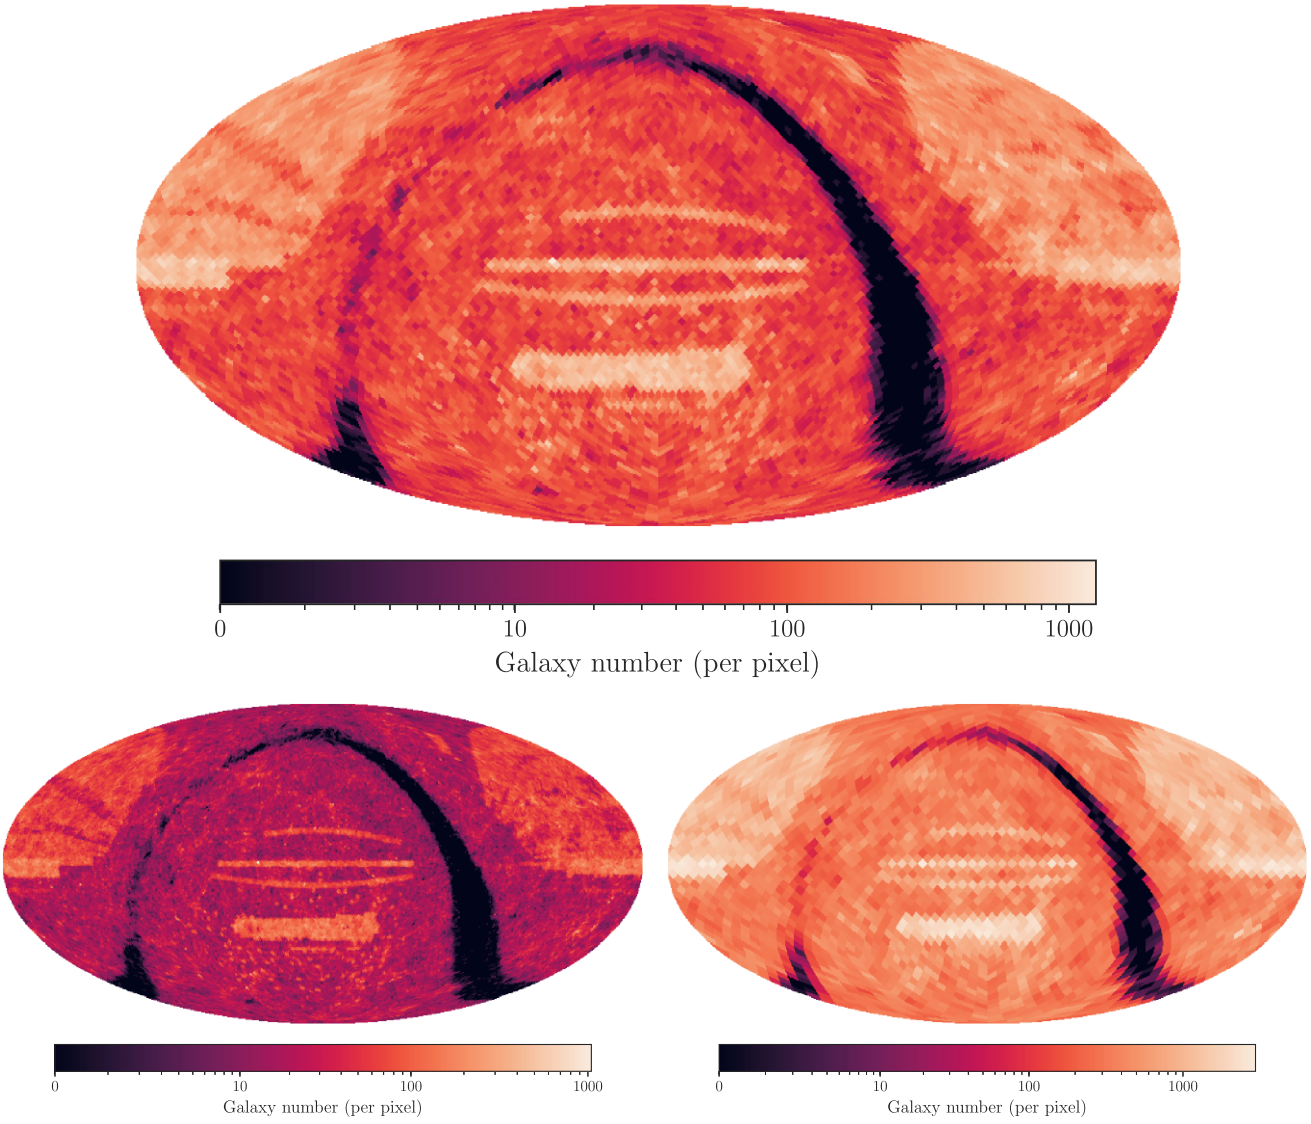

Figure 1. Different resolutions of the galaxy number density map for GLADE 2.4 (Dálya et al. 2018). Brighter colours correspond to a higher number density, darker colours to a lower one. Black pixels are empty, covered by the Milky Way band. *Top panel:* nside = 32. The sky is divided into 12 288 pixels, each covering an area of 3.36 deg^2. *Bottom left panel:* nside = 64. The sky is divided into 49 152 pixels, each covering an area of 0.84 deg^2. *Bottom right panel:* nside = 16. The sky is divided into 3072 pixels, each covering an area of 13.4 deg^2.

and distance information needs to be preserved for each pixel. This means that, even though 3D skymaps exist which contain the LOS GW distance information for each pixel, these cannot be used here because they include the marginalized-over priors applied during parameter estimation.

In order to retain the necessary information, the posterior samples are used directly to compute the LOS distance estimate for each pixel. A crude way to do this would be to choose some resolution for which to grid up the sky, then divide up the posterior samples between pixels, and create the LOS distance estimate for each pixel from the samples within it. However, this attempt fails at the first hurdle, where the finite number of samples comes into play. Most events have $\mathcal{O}(100,000)$ samples, of which $\mathcal{O}(100)$ are required to get a reasonable distance estimate (and more samples is preferable, as the estimate will be more reliable). As the samples cover the full parameter space, including RA and Dec., the edges of the event's sky area are particularly poorly sampled. Even for a relatively low-resolution pixelation of the sky, this leaves pixels in the event's 99.9 per cent sky area with as few as $\mathcal{O}(1)$ samples.

In order to avoid this, while still remaining compatible with a HEALPY implementation, the LOS distance estimate for each pixel is found by selecting all the samples within a certain angular radius of the centre of a pixel (determined by the resolution of the pixels). If the number of samples exceeds some threshold to be deemed 'enough samples' (taken to be 100 in this case), then a kernel density estimate (KDE) on these samples is used to create the LOS distance estimate. If there are not enough samples within the selected area, the angular radius is incrementally increased until the number of samples passes the threshold. KDEs are normalized by default, so each pixel's LOS distance distribution needs to be additionally weighted by the sky probability within that pixel of the GW skymap. By selecting samples in this way, the necessary information required to re-weight the samples by their source-frame mass priors is retained. Pixels in the most probable sky areas will have $\mathcal{O}(10,000)$ samples,

Table 1. Relevant parameters of the BBHs from GWTC-1: 90 per cent sky localization region $\Delta\Omega$ (deg^2), luminosity distance d_L (Mpc, median with 90 percent credible intervals), and estimated redshift z_{event} (median with 90 per cent range assuming Planck 2015 cosmology) from Abbott et al. (2019). The final column gives each event's 90 per cent 3D localization comoving volume.

Event	$\Delta\Omega$ (deg^2)	d_L (Mpc)	z_{event}	V (Mpc^3)
GW150914	182	440^{+150}_{-170}	$0.09^{+0.03}_{-0.03}$	3.5×10^6
GW151012	1523	1080^{+550}_{-490}	$0.21^{+0.09}_{-0.09}$	5.8×10^8
GW151226	1033	450^{+180}_{-190}	$0.09^{+0.04}_{-0.04}$	2.4×10^7
GW170104	921	990^{+440}_{-430}	$0.20^{+0.08}_{-0.08}$	2.4×10^8
GW170608	392	320^{+120}_{-110}	$0.07^{+0.02}_{-0.02}$	3.4×10^6
GW170729	1041	2840^{+1400}_{-1360}	$0.49^{+0.19}_{-0.21}$	8.7×10^9
GW170809	308	1030^{+320}_{-390}	$0.20^{+0.05}_{-0.07}$	9.1×10^7
GW170814	87	600^{+150}_{-220}	$0.12^{+0.03}_{-0.04}$	4.0×10^6
GW170818	39	1060^{+420}_{-380}	$0.21^{+0.07}_{-0.07}$	1.5×10^7
GW170823	1666	1940^{+970}_{-900}	$0.35^{+0.15}_{-0.15}$	3.5×10^9

while those in the lowest probability regions will only have 100, taken from a larger sky patch. The decrease in reliability of these low-probability distance estimates will be compensated for by the down-weighting of their contribution to the final result. Practically, GWCOSMO uses SCIPY's `Gaussian_kde` (Virtanen et al. 2020) for the LOS distance estimates. This KDE is effectively a weighted sum of Gaussians centred at the location of each posterior sample in the parameter space. The width of these Gaussians, and hence the overall smoothing of the KDE, which determines how accurately underlying structure is represented, is computed using the Scott method (Scott 1992).

The choice of resolution used to represent the GW data should depend on the area of the sky covered by the GW within some probability threshold. Events with small sky areas will need smaller pixels to adequately represent the changing of their distance distribution across the sky, and vice versa. As such, for GWCOSMO the resolution is chosen by defining a threshold on the sky area of the event and determining the nside resolution that would be necessary to split that patch of the sky into (at least) a minimum number of pixels. Choosing a minimum pixel number of 30 to cover the 99.9 per cent sky area of each event produces the LOS distance (redshift) estimates shown in Fig. 2. The left-hand panels show the breakdown of the redshift distribution along the line of sight of each pixel for the six BBH events from GWTC-1 that were used in the main analysis result of Abbott et al. (2021b) (GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814).[3] The redshift at which the distribution peaks varies from pixel to pixel, indicating that galaxies at different redshifts will be favoured for different lines of sight. The right-hand panel of Fig. 2 demonstrates that the overall distribution on redshift when summing the contribution from individual pixels is very close to the original distance distribution, estimated from all of the samples. As both are marginalized over Ω they should theoretically be identical, but in practice the pixelation process introduces minor variation. This variation is not large enough to cause noticeable impact to the inference of H_0 (see Section 3.1.1). Theoretically, a larger number of pixels would more accurately capture the event's 3D localization (see Section 4.1), up to the limit where the smallness of the pixels would require samples to be reused for the LOS estimates in the most probable pixels.

Table 2 summarizes the HEALPY resolution ($\text{nside}_{\text{low}}$) and number of pixels ($N_{\text{pix}}$) required for each event in order to meet the criteria of at least 30 pixels to cover the event's 99.9 per cent sky area. Additionally, the number of posterior samples in the most probable pixel is recorded (N_{samples}) which shows that for every event the most probable pixels, which will be contributing most strongly to the result, have at least 1000 samples, and in some cases as high as tens of thousands.[4] The impact of changing the sky area under consideration, or the minimum number of pixels required to cover it (and hence the resolution with which the GW data are treated) is examined in greater detail in Section 4.1.

3.1.1 Sanity check: the empty catalogue case

To further prove that this method of estimating the LOS distance is robust, and to ensure that the H_0 dependence of the distribution has been correctly propagated, it is worth coming back to the empty catalogue analysis, in which no galaxy catalogue is used and the redshift prior is taken to be uniform in co-moving volume. In this case, the results are relatively (but not completely) uninformative. Small amounts of information come from the mass and distance distributions of the GW sources, leading to individual event likelihoods that are not completely flat in H_0. The pixelated empty catalogue analysis should return results equivalent an empty catalogue analysis that is done over the whole sky – because no catalogue information enters this analysis, the order in which the marginalization over Ω takes place is irrelevant.

The analysis assumes a power-law source frame mass distribution for the BBHs, with the primary mass $p(m_1) \propto m_1^{-\alpha}$, where $\alpha = 1.6$ and the secondary mass uniform between M_{min} and m_1. Additional constraints are defined such that $M_{\text{min}} = 5\,\text{M}_\odot$ and $M_{\text{max}} = 100\,\text{M}_\odot$ and the network SNR threshold assumed for computing GW selection effects is 12. Results for GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814 can be seen in Fig. 3. The left-hand panels show the fractional difference between the likelihoods computed using the pixelated method, and the likelihoods computed over the whole sky at once. For every event, the difference is less than 1 per cent across all values of H_0. The right-hand panels show the contribution to the pixelated empty catalogue likelihood from each pixel covering the 99.9 per cent sky area of the event. A close examination of the right-hand panels shows that individual pixel contributions do not have the same H_0 dependence, independent of the scale (GW150914 in particular shows a good example of this). This difference in the slope is due to the change in the GW LOS distance distribution for different pixels. The fact that the different methods produce only fractional changes to the final likelihoods demonstrates that the GW data is well represented by this pixelated approach.

3.2 Varying m_{th} within an event's sky area

The second major improvement the pixelated method offers is the ability to treat the apparent magnitude threshold, m_{th}, of a galaxy catalogue as a function of sky direction.

[3] Fig. 2 is shown as a function of redshift, as opposed to luminosity distance, simply because in GWCOSMO the KDE is done on redshift samples.

[4] While this was true for all of the GWTC-1 events it should not be assumed true in general, and number of pixels or the threshold for the sky area may need to be adjusted to meet this criteria.

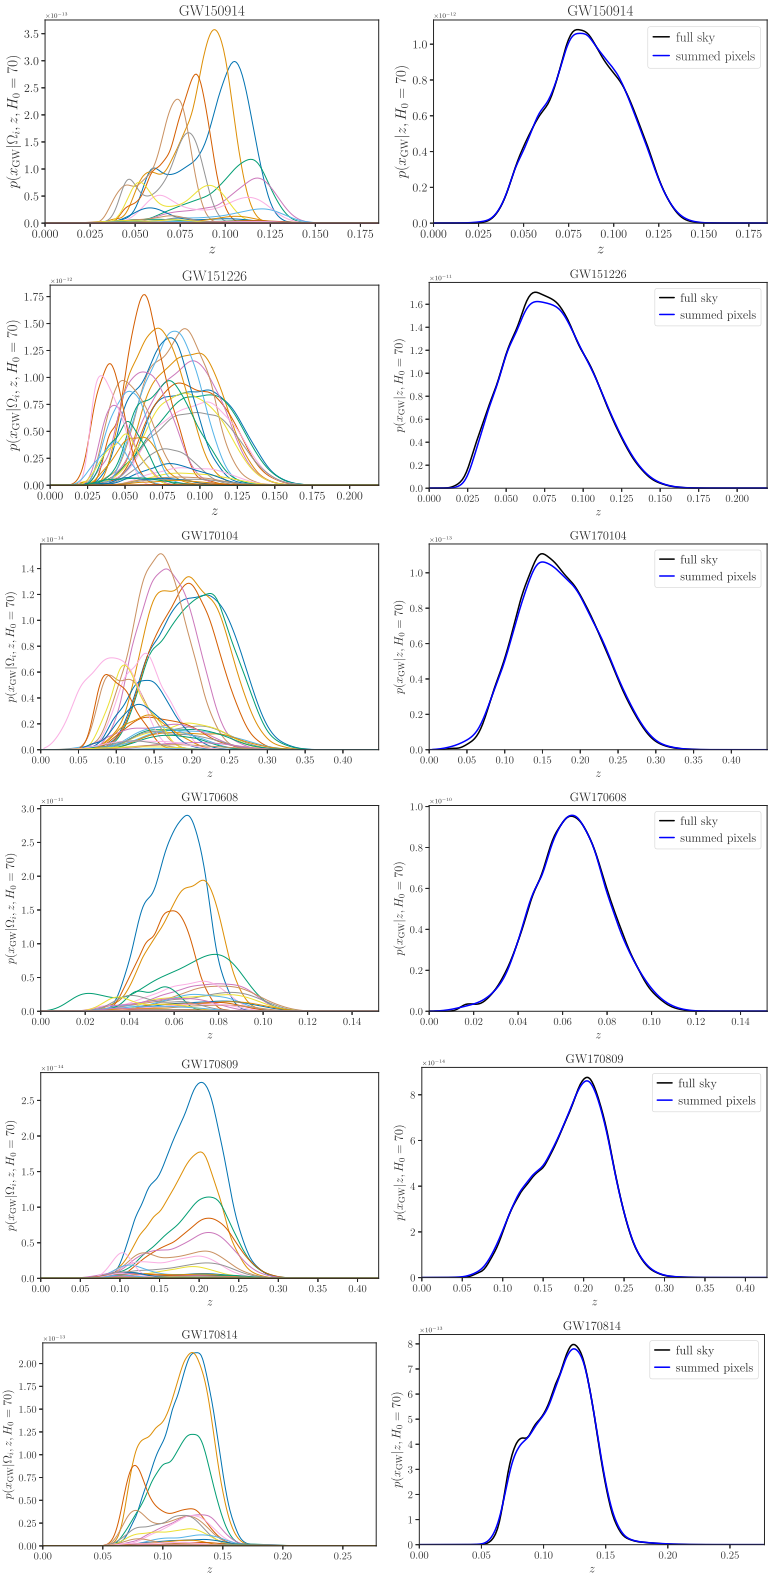

Figure 2. LOS redshift estimates for GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814. *Left-hand panels:* The LOS redshift distribution of the event within each pixel, where the number of pixels and pixel nside for each event are specified in the $N_{\rm pix}$ (low-res) and $n_{\rm low}$ columns of Table 2. *Right-hand panels:* The full-sky redshift distribution for the event. The black curve shows the estimate from doing a KDE on all the samples, while the blue curve shows the summed curves from the left-hand panel.

Table 2. A summary of the resolutions used for the pixelated analysis for the GWTC-1 BBHs. The second and third columns show the lowest nside that satisfies the criteria that at least 30 pixels must cover the 99.9 per cent sky area of the event, $\text{nside}_{\text{low}}$, and the number of pixels, N_{pix}, required to do so. The most probable pixel for each event contains N_{samples} samples. Each of the pixels is broken into $N_{\text{sub-pix}}$ sub-pixels in order to reach a galaxy catalogue resolution of nside = 32, and the total number of sub-pixels which cover the event's 99.9 per cent sky area is $N_{\text{sub-pix}}$ (total).

Event	$\text{nside}_{\text{low}}$	N_{pix}	N_{samples}	$N_{\text{sub-pix}}$ (pixel)	$N_{\text{sub-pix}}$ (total)
GW150914	8	35	7759	16	560
GW151012	4	40	1941	64	2560
GW151226	4	34	3737	64	2176
GW170104	4	30	4804	64	1920
GW170608	8	38	4010	16	608
GW170729	4	32	22 582	64	2048
GW170809	8	34	30 882	16	544
GW170814	16	54	12 750	4	216
GW170818	16	72	6417	4	288
GW170823	4	43	6130	64	2752

Fig. 4 shows the variation in *B*-band apparent magnitude threshold across the 99.9 per cent sky area of GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814. The resolution of the larger pixels, which determine which patch of the sky is considered for the analysis (in colour, where the grey is what is excluded), is determined by the resolution required to cover the 99.9 per cent sky area with at least 30 pixels. For GW150914, GW170608, and GW170809 this was an nside of 8. For GW151226 and GW170104, which were particularly poorly localized, this reduces to an nside of 4, while GW170814, which is much better-localized, adopts an nside of 16. An nside of 32 is chosen to represent the galaxy catalogue information, meaning that the pixels for each event needs to be split into a number of sub-pixels in order to reach the required resolution. If $\text{nside}_{\text{low}}$ is the low-resolution nside used to represent the GW data, and $\text{nside}_{\text{high}}$ is the nside used to represent the galaxy catalogue, the number of sub-pixels that a single pixel must be divided into, $N_{\text{sub-pix}}$, is given by 4^k, where $k = \log_2(\text{nside}_{\text{high}}/\text{nside}_{\text{low}})$. This information is summarized for each event in the last columns of Table 2.

The method of determining m_{th} within each sub-pixel is the same as Abbott et al. (2021b) – by taking the median *B*-band apparent magnitude of galaxies within the pixel. Again, looking at Fig. 4, it is clear that this method captures a large variation in the apparent magnitude threshold within each event's sky area. Picking the median apparent magnitude is a conservative choice, and results which use this method will give less support to the in-catalogue part of the analysis than a method which calculates m_{th} more robustly. However, for consistency of comparisons with Abbott et al. (2021b), the same choice is made here.

The variation of m_{th} over the sky translates directly into a variation in the probability that the host galaxy of the event is inside the catalogue, with lower thresholds corresponding to lower in-catalogue probabilities. Fig. 5 shows how the probability that the host is in the catalogue varies within each event's sky area. The orange curves correspond to the pixels that contain 50 per cent of the GW event's sky probability, with darker curves corresponding to the pixels with higher probabilities. The yellow shaded area shows the full range covered by the pixels that make up the 99.9 per cent sky area. This extends to a probability of 0 at $z = 0$ if there are pixels within the event's 99.9 per cent sky area which are empty. In general, these plots show how the probability that the host is in the catalogue compares between the pixelated case and the Abbott et al. (2021b) case that is shown by the dashed blue line. This curve always lies within the extremes of the pixelated case (within the yellow area). However, the orange curves indicate where the bulk of the GW probability is lying. For GW150914, for example, the bulk of the GW probability corresponds to lower in-catalogue probability than previously, no doubt driven by the fact that part of the event is obscured by the Milky Way band, which will cause a low m_{th} in the adjacent pixels (pixels which are entirely empty do not show up on this plot). Conversely, GW170608's probability is clustered in an area of higher in-catalogue completeness than the average, boosting the in-catalogue support. For GW170814, the most probable pixels overlap with the Abbott et al. (2021b) curve, as the m_{th} estimation for the DES-Y1 catalogue varies very little within the sky area of the event.

4 RESULTS

Taking the method outlined in Section 2 and applying it to the six GWTC-1 BBHs that pass a network SNR threshold of 12 (GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814) produces the individual event likelihoods on H_0 shown in Fig. 6. The same mass distribution and detection threshold as outlined in Section 3.1.1 is used. The *B*-band luminosity function of galaxies (used to match those in the GLADE catalogue) is assumed to follow a Schechter function with slope $\alpha = -1.07$, a characteristic absolute magnitude of $M^*(H_0) = -19.7 + 5\log h$ and a maximum limit on the faintest galaxies of $-12.2 + 5\log h$, where $h \equiv H_0/100$, which follows Gehrels et al. (2016). The *g*-band luminosity function (used to match the DES-Y1 galaxy catalogue; Abbott et al. 2018; Drlica-Wagner et al. 2018) is modelled using a Schechter function with $\alpha = -0.89$, $M^*(H_0) = -19.39 + 5\log_{10}h$, and a limit on faint galaxies of $-16.1 + 5\log_{10}h$ based on Blanton et al. (2003). Both of these Schechter functions are chosen to match the assumptions made in Abbott et al. (2021b). GW170814 is analysed with the DES-Y1 catalogue,[5] while the remaining events are analysed with the GLADE 2.4 catalogue.

The most interesting event is GW170608, due to its high in-catalogue probability. Looking at its right-hand panel in Fig. 6, which shows the contribution to the final likelihood on H_0 from each individual pixel, it is clear that for low values of H_0 the in-catalogue contribution is dominating. The final likelihood shows more structure than the Abbott et al. (2021b) result, including increased posterior support around $H_0 \sim 70$ km s^{-1} Mpc^{-1}.

While the likelihoods on H_0 from GW150914 and GW151226 do not show much structure, both have more support around central values of H_0, and have reduced support at high H_0, compared to the Abbott et al. (2021b) case. For these events, the contributions from the individual pixels are more interesting than the combined likelihood as they show, especially at the low-H_0 end (which corresponds to low redshifts and therefore increased catalogue support) structure which corresponds to real redshift and luminosity information from the GLADE catalogue galaxies. It is interesting that GW150914's most probable pixel peaks around $H_0 \sim 70$ km s^{-1} Mpc^{-1} which could indicate an overdensity of galaxies at the relevant redshift, although it remains much more likely that the host galaxy for this event is not contained within the catalogue (based on the top panel of Fig. 5).

The other event of interest is GW170814, which was analysed with the DES-Y1 catalogue, and for which the likelihood remains

[5]The DES-Y1 catalogue is available at https://des.ncsa.illinois.edu/releases/y1a1.

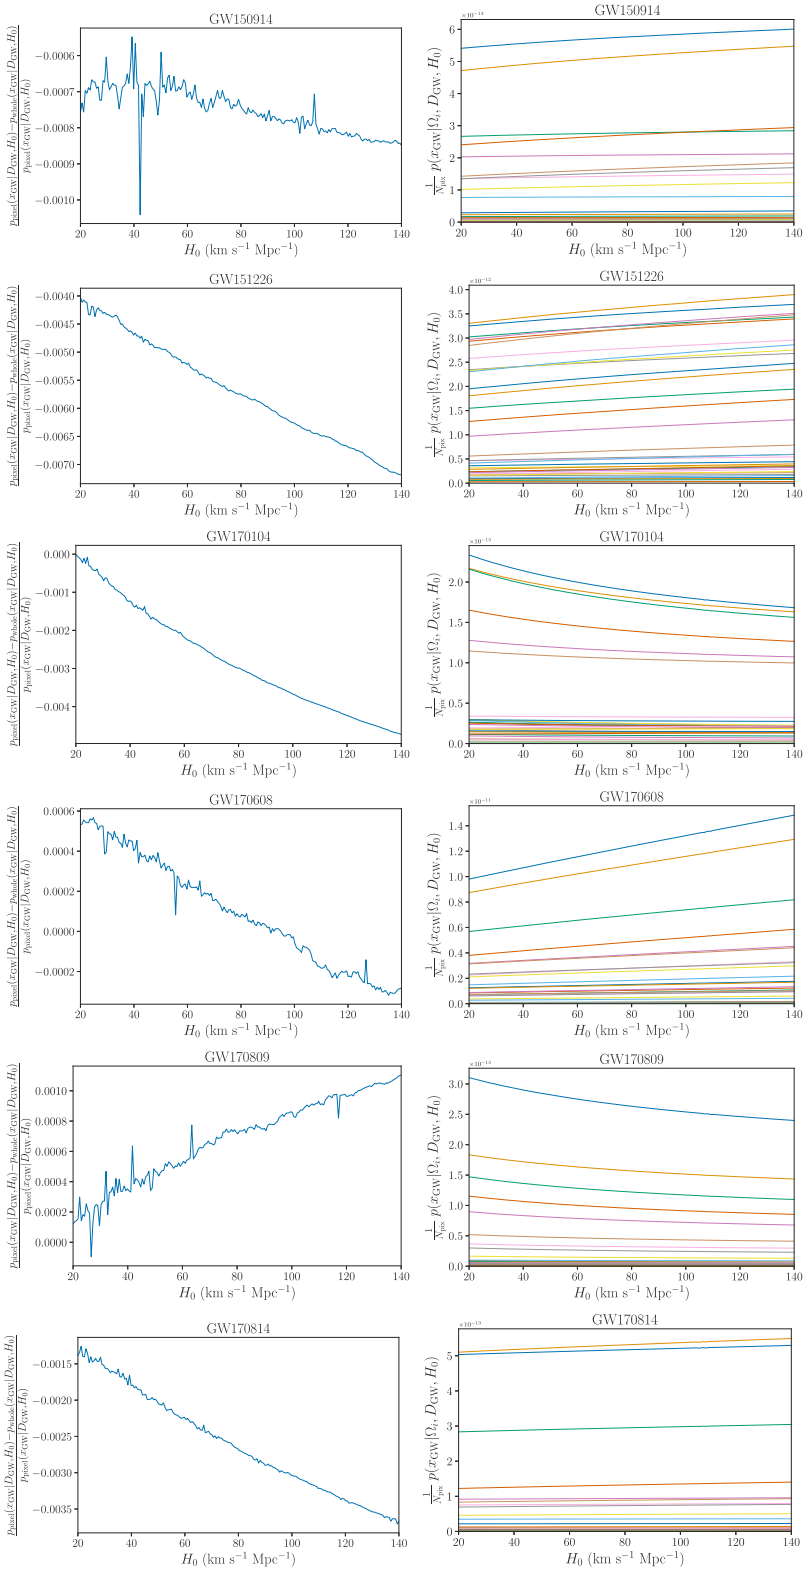

Figure 3. Likelihoods on H_0 for the empty catalogue analysis with GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814. *Left-hand panels:* Fractional difference between the pixelated empty catalogue likelihood on H_0 and the non-pixelated (whole sky) likelihood, for each event. *Right-hand panels:* A breakdown of the pixelated empty catalogue likelihood by pixel.

relatively unchanged with the pixelated method, though with a fraction more support around central values of H_0. Looking at the LOS estimates for this event (Fig. 2), it is clear that the bulk of the information is coming from a small number of pixels, which peak at approximately the same redshifts. The apparent magnitude threshold within the event's sky area remains relatively uniform (see Fig. 4). From the likelihood breakdown of the event in Fig. 6, it seems likely that the most probable pixels overlap with the same overdensity

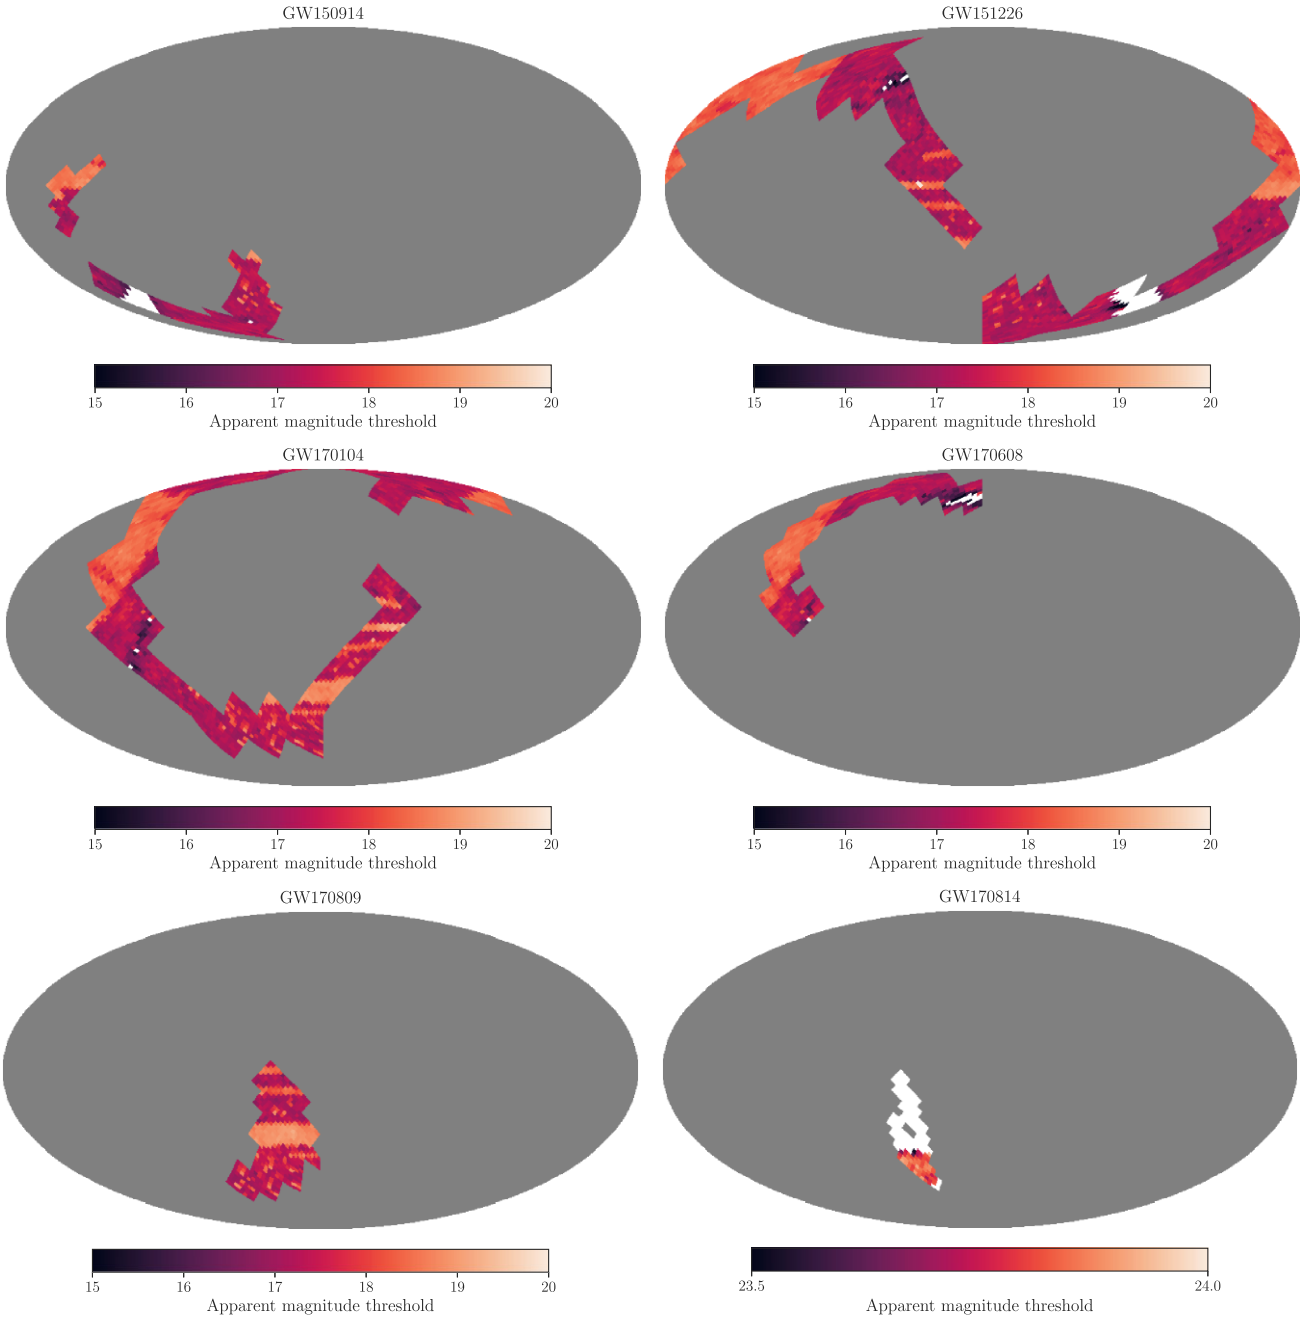

Figure 4. Variation of the apparent magnitude threshold within the 99.9 per cent sky area of GW150914, GW151226, GW170104, GW170608, and GW170809 [using the Galaxy List for the Advanced Detector Era (GLADE) catalogue], and GW170814 [using the Dark Energy Survey (DES)-Y1 catalogue]. The galaxy catalogue resolution (small pixels) is set with an nside of 32. The grey indicates a part of the sky that is not in the 99.9 per cent sky area of the event. The white pixels contain fewer than 10 galaxies either due to obscuration by the Milky Way band, or a lack of survey data, and are taken to be empty.

of galaxies in the catalogue, leading to a relatively unchanged result.

Taking these six BBHs and combining them with the result from GW170817 and its counterpart (Abbott et al. 2017b, c) gives the posterior on H_0 shown in Fig. 7. The pixelated method gives a result of $H_0 = 68.8^{+15.9}_{-7.8}$ km s^{-1} Mpc^{-1}, (maximum a posteriori and 68.3 per cent highest density interval, with a flat-in-log prior on H_0). This is approximately a 5 per cent improvement over the Abbott et al. (2021b) result. The improvement is driven by the additional information from GW170608, which has greater support around an H_0 of 70 km s^{-1} Mpc^{-1} due to a more informative catalogue contribution, as well as by minor improvements from GW150914, GW151226, and GW170814, all of which have marginally increased support at middling values of H_0 under this new method.

4.1 The impact of resolution choices on the result

The pixelated method introduces several choices to the analysis: the threshold for the sky area that will be analysed, the number of pixels that will be used to cover that sky area (and hence the size of those pixels) and the resolution of the galaxy catalogue. This section investigates the impact of varying those choices.

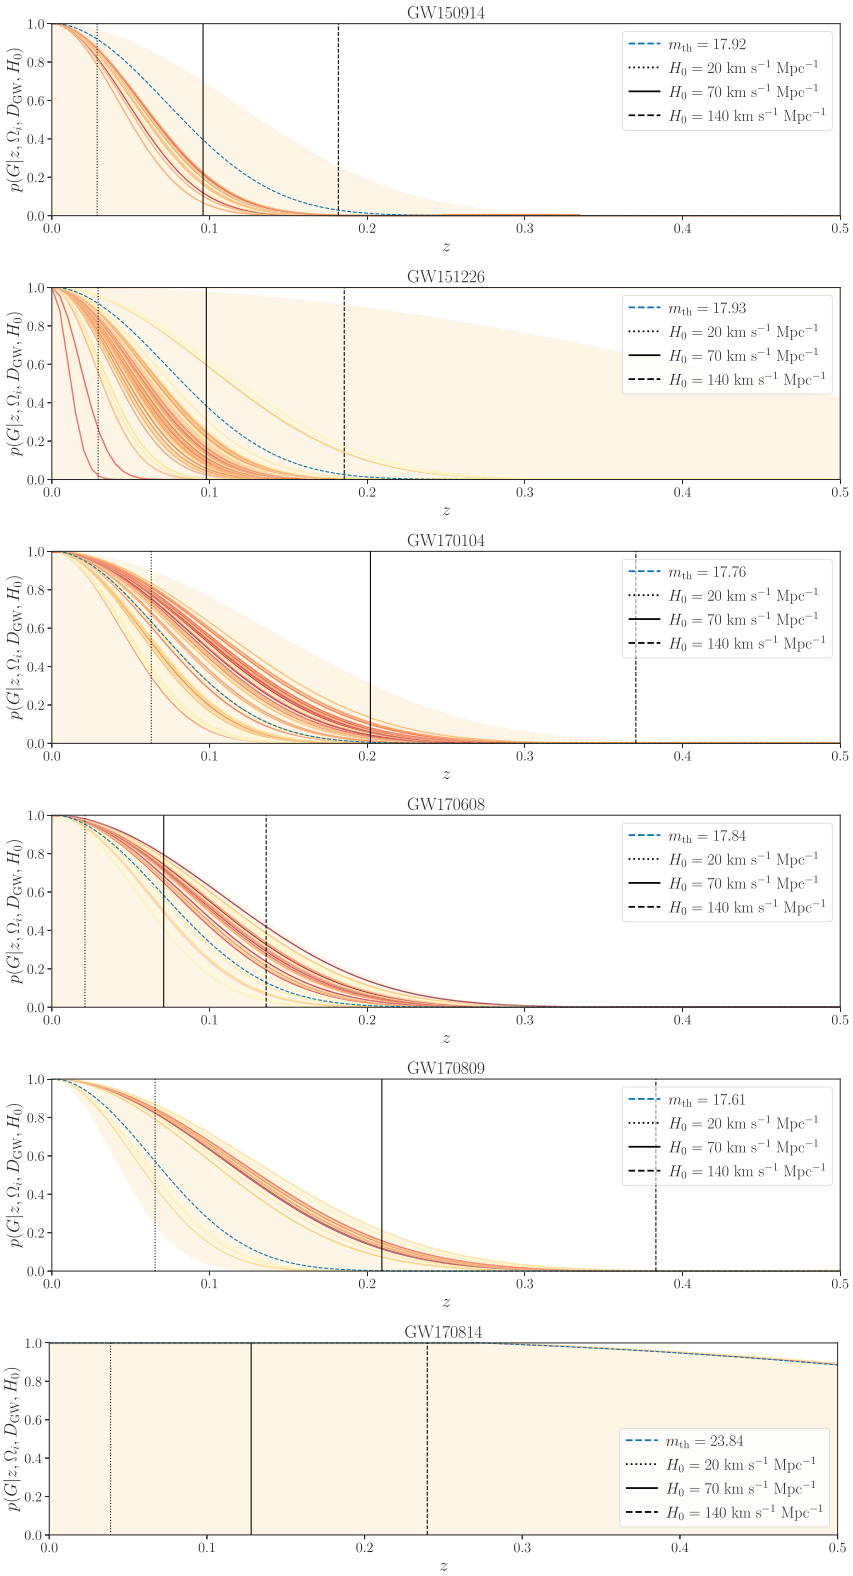

Figure 5. Probability that the host galaxy is in the galaxy catalogue as a function of redshift, for the pixelated catalogue case. Solid orange curves show the probability that the host is inside the galaxy catalogue along the line of sight of the pixels that cover the 50 per cent sky area of the event (darker orange corresponds to pixels which contain higher values of the GW sky probability). The shaded yellow area covers the range between the minimum and maximum apparent magnitude threshold within the 99.9 per cent sky area of the event. Vertical black lines show the median redshift of the event for different values of H_0. The blue dashed line gives the probability that the host is in the galaxy catalogue assuming the $m_{\rm th}$ value used in Abbott et al. (2021b).

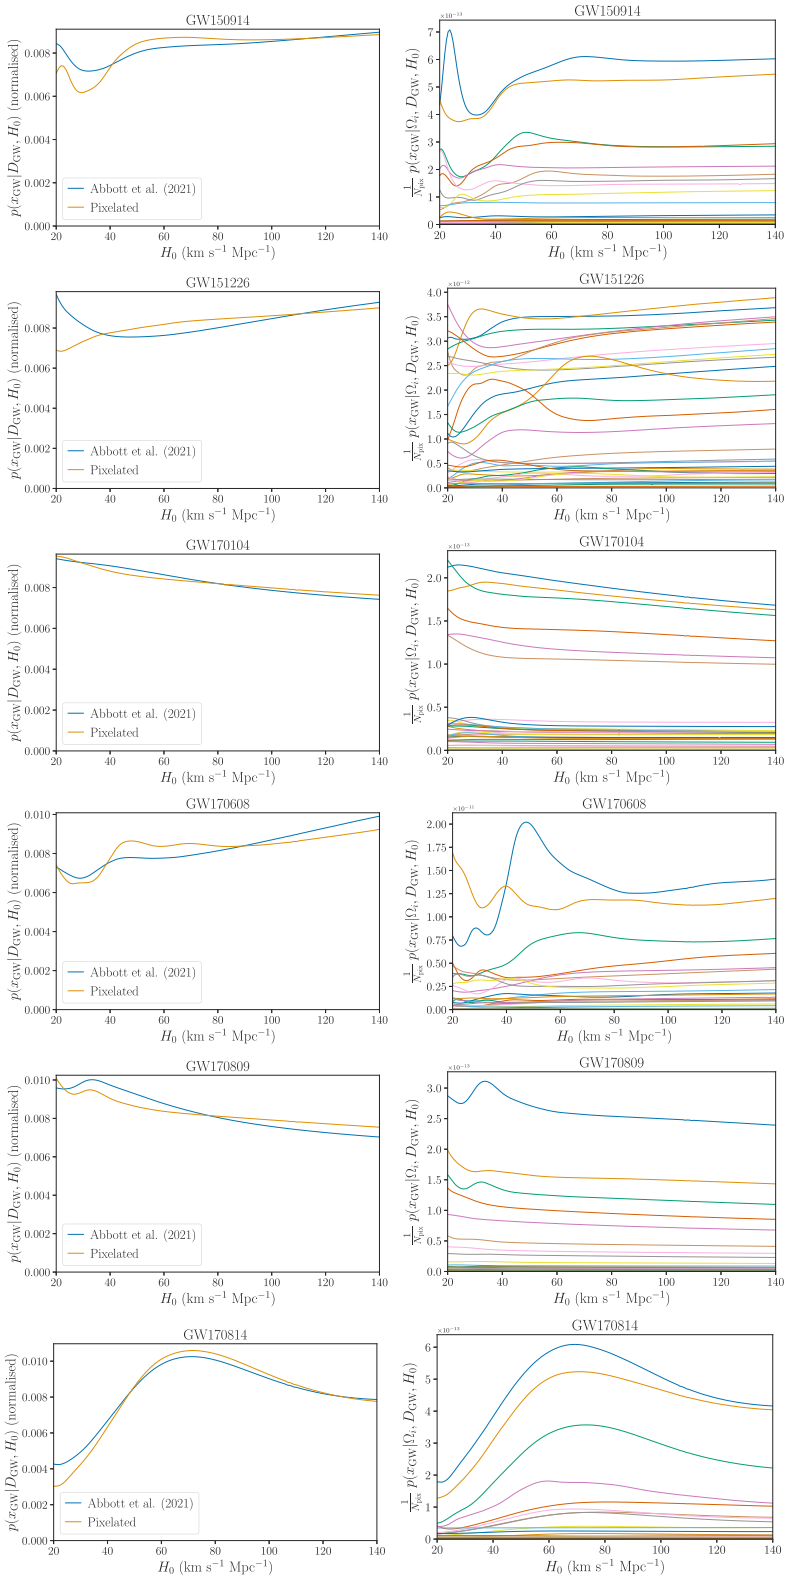

Figure 6. Likelihoods on H_0 for the pixelated analysis with GW150914, GW151226, GW170104, GW170608, GW170809, and GW170814. *Left-hand panels:* Comparison of the (normalized) likelihoods between the Abbott et al. (2021b) result (blue) and the pixelated result (orange). *Right-hand panels:* A breakdown of the pixelated likelihood by low resolution pixel. The pixelated likelihood in the left-hand panel is the sum of the curves in the right-hand panel (then normalized).

In the first instance, the threshold on the GW sky area is varied. The results in Section 4 made use of the 99.9 per cent sky area. Fig. 8 demonstrates the change to the final posterior on H_0 when the area is reduced to 99 per cent and 90 per cent for the BBHs under consideration. The change to the final posterior is marginal, with a slight increase in the height of the peak corresponding to the smaller sky areas. This is expected as the less informative edges of the GW sky distribution are being discarded. In this case, more informative is

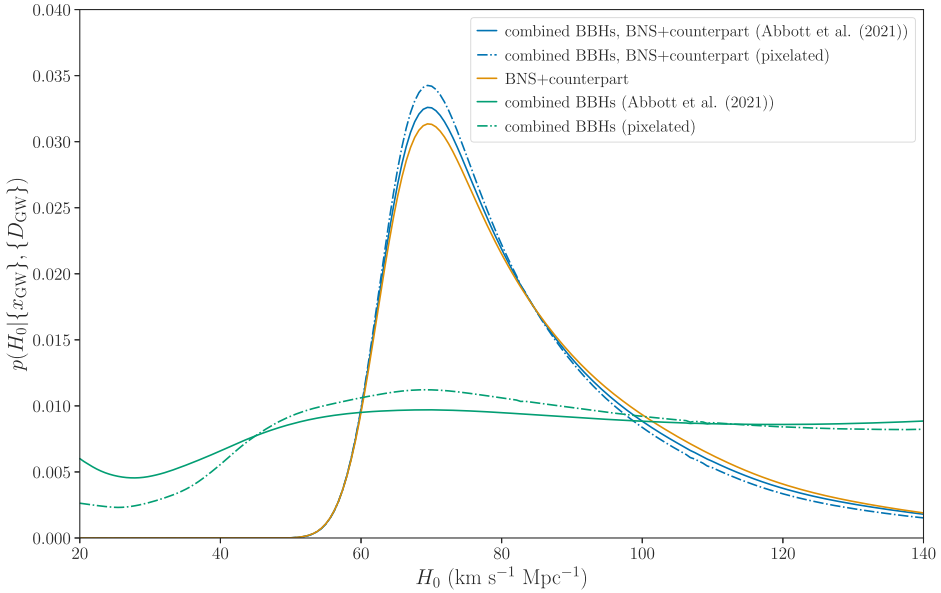

Figure 7. Comparison of the posterior on H_0 between the pixelated and Abbott et al. (2021b) analyses. The green lines show the combined posterior from six BBHs (GW150914, GW151226, GW170104, GW170608, and GW170809 analysed with the GLADE 2.4 catalogue, and GW170814 analysed with the DES-Y1 catalogue). The orange line shows the contribution from GW170817 and its counterpart. Blue lines show the combined posterior on H_0 from the six BBHs and the BNS with counterpart. Solid lines correspond to the Abbott et al. (2021b) results and dot–dashed lines show the pixelated results.

not necessarily a good thing, as there is the possibility that discarding this additional information could eventually introduce some level of bias to the result. While not important here, that impact should be reassessed in the future, when results from large numbers of GW events are being combined.

Next, the threshold is reverted to the 99.9 per cent sky area, but the number of pixels which covers it is increased. The nside that determines the resolution of the GW data (column 2 of Table 2) was doubled for each event, which leads to a factor of 4 increase in the number of pixels covering the 99.9 per cent sky area.[6] The impact on the H_0 posterior, shown in Fig. 9, is a marginal increase in the height of the peak when the higher resolution is applied to the GW data. The fact that the difference is so small should be taken as additional confirmation that a relatively low number of pixels can adequately represent the variation in the GW LOS distance distribution.[7]

Finally, the resolution of the galaxy catalogue is investigated, by increasing the nside of the sub-pixels from 32 to 64. This allows for a better representation of the hard edges of the GLADE catalogue where it is intersected by the Milky Way band, as well as better representation of the variation of $m_{\rm th}$ across different parts of the survey. The results are shown in Fig. 10, which shows a very small increase in the height of the peak for the higher resolution results. As the difference is again very small, it is safe to assume that the resolution from nside = 32 is adequate to represent the variation in galaxy catalogue completeness over the sky.

[6]This is approximate, not exact, as variations in how the GW sky probability is distributed between the higher resolution pixels means that not all will necessarily be required in order to reach the 99.9 per cent threshold.

[7]It is also worth noting that the results in Figs 8 and 9 are correlated, as reducing the sky area of the event under consideration results, for some events, in an increase in pixel resolution in order to keep the number of pixels covering the sky area roughly the same.

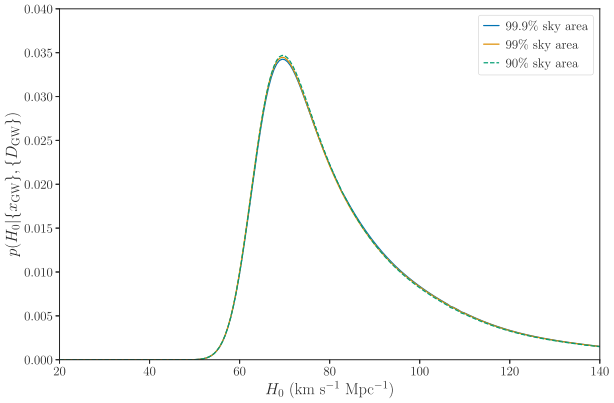

Figure 8. Comparison of the posterior on H_0 when the threshold on the GW sky area used in the pixelated analysis is reduced. The blue line shows the main result for six BBHs and the BNS, where the 99.9 per cent sky area for each BBH is analysed. The orange line corresponds to the 99 per cent sky area, and the green line to the 90 per cent sky area.

5 CONCLUSION

There are two major benefits of using the pixelated method presented in this paper for the measurement of H_0 using standard sirens and galaxy catalogues. The first is that full use is now made of the GW data by estimating a separate distance distribution for each pixel which makes up the event's sky area. The fact that these distributions will peak at different distances for different lines of sight – and will therefore pick out galaxies at different redshifts depending on how they align – increases the information available for the dark siren analysis. The second benefit is that the pixelated method allows for an accurate estimation of how the apparent magnitude threshold – and hence the completeness – of a galaxy catalogue changes across the sky. This is particularly important

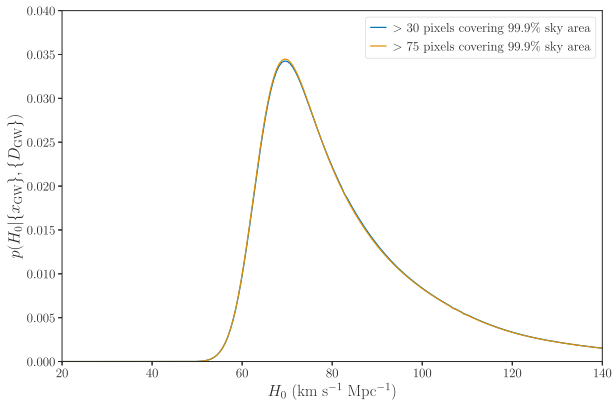

Figure 9. Comparison of the posterior on H_0 when the minimum number of pixels covering the 99.9 per cent GW sky area is increased. The blue line shows the main result for six BBHs and the BNS, where at least 30 pixels cover the 99.9 per cent sky area for each BBH is analysed. The orange line corresponds to an analysis where the GW data is analysed at one resolution step higher, where at least 75 pixels are used to cover the 99.9 per cent sky area.

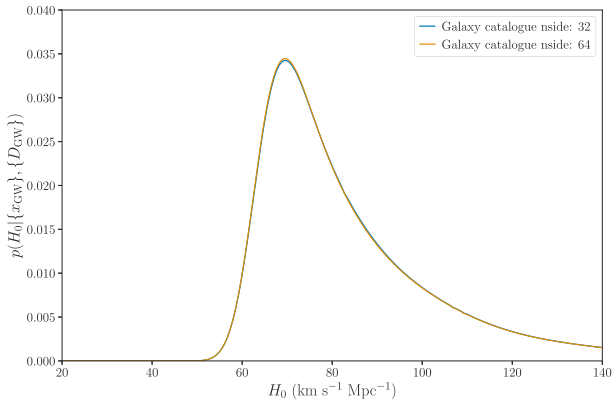

Figure 10. Comparison of the posterior on H_0 when the galaxy catalogue resolution is increased. The blue line shows the main result for 6 BBHs and the BNS, where an nside of 32 is assumed for both the GLADE and DES catalogue. The orange line shows the same, but where the nside of the both galaxy catalogues has been increased to 64.

around the Milky Way band, where telescope visibility is limited, and at any boundaries between different observing surveys, between which telescope sensitivities may have changed. Given how dominant the out-of-catalogue contribution is to the dark siren result for the majority of GWTC-1 events, and will continue to be for the third observing run and beyond, this is an important milestone.

The results presented in Section 4 show a clear improvement over the results in Abbott et al. (2021b) – around that of 5 per cent. Not only does the pixelated method make use of the data in a way that is more technically correct, but this leads to a direct increase in the informativeness of the results it produces. These results are both more robust and have more to say. That said, it is worth remembering that both the results in this paper and in Abbott et al. (2021b) are sensitive to the choice of GW population model that, for a large number of dark sirens, can play a major role in the measurement if not suitably marginalized over. See e.g. The LIGO Scientific Collaboration (2021), in which the impact of population assumptions on the GW measurement of H_0 using GWTC-3 events is explored in detail.

An investigation into how various choices for the resolution of the pixelated analysis impacts the result revealed that, for the most part, the results are insensitive to reasonable changes. Variations from increasing the resolution of pixels used to represent the GW data were small, as were those from increasing the resolution of the galaxy catalogue. These changes may become more important as more events are analysed and the goal to reach 1 per cent measurement uncertainty on H_0 comes into reach. In the future it may be worth improving the analysis outlined in this paper to use multiresolution maps to represent the GW data: well-localized highly probable areas could be treated with high-resolution pixels in order to capture the fine detail, and the tail ends of the distribution – the large areas with very little probability – could be treated on a lower resolution. This would allow for the most information to be gained from the GW data, without having to discard the low-probability areas, which it may be necessary to include in order to avoid introducing bias to the result.

Of course, increasing the size of the GW pixels means that, in order to reach the resolution set by the galaxy catalogue, each pixel needs to be split into more sub-pixels, which increases the computation time. Currently, the computational cost of each pixel is approximately proportional to the number of sub-pixels that it is divided into (which means that the pixels from events with large sky areas have a longer run-time than those which are well localized, even though these events are unlikely to contribute much information to the final result). The solution to this, which will radically reduce computation time for large pixels, is to merge sub-pixels which have similar apparent magnitude thresholds. Looking back at the top left panel of Fig. 4: within the sky area of GW150914, $m_{\rm th}$ varies between approximately 16 and 19 mag – a huge variation in terms of completeness. However, it is also clear that the majority of pixels have a threshold around 17.3 (reddish-pink pixels). Pixels in the top left have higher thresholds, around 18.8. Pixels on the edge of the Milky Way band dip below 17. Overall, the variation of $m_{\rm th}$ within GW150914's sky area could be adequately represented by a handful of different thresholds. The costly out-of-catalogue contribution within each pixel, which has to be calculated for every value of $m_{\rm th}$, would then be reduced from 16, in this case, down to only several. And if the resolution of the catalogue was increased at some point in the future, the number of out-of-catalogue calculations required would not increase, allowing incredibly high resolution of features such as empty patches and boundaries between surveys, with no extra computational cost. Following a similar approach to Finke et al. (2021), and computing a completeness map (but in terms of $m_{\rm th}$) that assigns pixels to groups of similar completeness would be one way to do this.

In summary, the pixelated analysis demonstrated in this paper shows a clear improvement on the analysis in Abbott et al. (2021b). It improves the final constraint on H_0 through better, more effective use of the GW and galaxy catalogue data available. The implementation demonstrated here has clear avenues for future development, both in terms of further improving the accuracy of the analysis, and improving its efficiency. This pixelated approach has potential to become the main method for any future H_0 analysis with dark standard sirens and galaxy catalogues.

ACKNOWLEDGEMENTS

The authors are grateful to members of the LIGO, Virgo, and KAGRA collaborations and specifically those in the cosmology working group for their valuable input, and to Maciej Bilicki and Archisman Ghosh in particular for their helpful comments. The authors are grateful for computational resources provided by the LIGO Laboratory. RG was

supported by the Science and Technology Facilities Council (award number 1947165) and by the European Research Council (starting grant SHADE 949572). JV and CM were supported by the Science and Technology Research Council (grant no. ST/V005634/1).

DATA AVAILABILITY

GWTC-1 data are available at https://www.gw-openscience.org/GWTC-1. The GLADE 2.4 catalogue is available at http://glade.elte.hu, and the DES-Y1 catalogue is available at https://des.ncsa.illinois.edu/releases/y1a1. GWCOSMO is available at https://git.ligo.org/lscsoft/gwcosmo.

REFERENCES

Abbott B. P. et al., 2017a, Phys. Rev. Lett., 119, 161101
Abbott B. P. et al., 2017b, Nature, 551, 85
Abbott B. P. et al., 2017c, ApJ, 848, L12
Abbott T. M. C. et al., 2018, ApJS, 239, 18
Abbott B. P. et al., 2019, Phys. Rev. X, 9, 031040
Abbott R. et al., 2020, ApJ, 896, L44
Abbott R. et al., 2021a, Phys. Rev. X, 11, 021053
Abbott B. P. et al., 2021b, ApJ, 909, 218
Abbott R. et al., 2021c, ApJ, 915, L5
Blanton M. R. et al., 2003, ApJ, 592, 819
Chen H.-Y., Essick R., Vitale S., Holz D. E., Katsavounidis E., 2017, ApJ, 835, 31
Chen H.-Y., Fishbach M., Holz D. E., 2018, Nature, 562, 545
Dálya G. et al., 2018, MNRAS, 479, 2374
Del Pozzo W., 2012, Phys. Rev. D, 86, 043011
Drlica-Wagner A. et al., 2018, ApJS, 235, 33
Farr W. M., Fishbach M., Ye J., Holz D., 2019, ApJ, 883, L42
Finke A., Foffa S., Iacovelli F., Maggiore M., Mancarella M., 2021, J. Cosmol. Astropart. Phys., 08, 026
Fishbach M., Holz D. E., Farr W. M., 2018, ApJ, 863, L41
Fishbach M. et al., 2019, ApJ, 871, L13
Gehrels N., Cannizzo J. K., Kanner J., Kasliwal M. M., Nissanke S., Singer L. P., 2016, ApJ, 820, 136
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Gray R., et al., 2020, Phys. Rev. D, 101, 122001
MacLeod C. L., Hogan C. J., 2008, Phys. Rev., D77, 043512
Mastrogiovanni S. et al., 2021, Phys. Rev. D, 104, 062009
Mukherjee S., Wandelt B. D., Nissanke S. M., Silvestri A., 2021, prd, 103, 043520
Palmese A. et al., 2020, ApJ, 900, L33
Planck Collaboration VI, 2020, A&A, 641, A6
Riess A. G., Casertano S., Yuan W., Macri L. M., Scolnic D., 2019, ApJ, 876, 85
Schechter P., 1976, ApJ, 203, 297
Schutz B. F., 1986, Nature, 323, 310
Scott D. W., 1992, Multivariate Density Estimation: Theory, Practice and Visualization. Wiley, New York
Soares-Santos M., et al., 2019, ApJ, 876, L7
The LIGO Scientific Collaboration, 2021, preprint (arXiv:2111.03604)
Vasylyev S. S., Filippenko A. V., 2020, ApJ, 902, 149
Virtanen P. et al., 2020, Nat. Methods, 17, 261
Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, J. Open Source Softw., 4, 1298

This paper has been typeset from a TEX/LATEX file prepared by the author.