Enabling discovery of gravitationally lensed explosive transients: a new method to build an all-sky watch list of groups and clusters of galaxies

Dan Ryczanowski [ID],[1]★ Graham P. Smith [ID],[1] Matteo Bianconi [ID],[1] Sean McGee,[1] Andrew Robertson [ID],[2] Richard Massey [ID][3] and Mathilde Jauzac[3,4,5,6]

[1]*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK*
[2]*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA*
[3]*Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*
[4]*Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK*
[5]*Astrophysics Research Centre, University of KwaZulu-Natal, Westville Campus, Durban 4041, South Africa*
[6]*School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Durban 4041, South Africa*

ABSTRACT

Cross-referencing a watch list of galaxy groups and clusters with transient detections from real-time streams of wide-field survey data is a promising method for discovering gravitationally lensed explosive transients, including supernovae, kilonovae, gravitational waves, and gamma-ray bursts in the next 10 yr. However, currently there exists no catalogue of objects with both sufficient angular extent and depth to adequately perform such a search. In this study, we develop a cluster-finding method capable of creating an all-sky list of galaxy group- and cluster-scale objects out to $z \simeq 1$ based on their lens-plane properties and using only existing data from wide-field infrared surveys such as VISTA Hemisphere Survey and UKIRT Hemisphere Survey, and all-sky *WISE* (*Wide-field Infrared Survey Explorer*) data. In testing this method, we recover 91 per cent of a sample containing known and candidate lensing objects with Einstein radii of $\theta_E \geq 5$ arcsec. We also search the surrounding regions of this test sample for other groups and clusters using our method and verify the existence of any significant findings by visual inspection, deriving estimates of the false positive rate that are as low as 6 per cent. The method is also tested on simulated Rubin data from their Data Preview 0 programme, which yields complementary results of a good recovery rate of $\gtrsim 80$ per cent for $M_{200} \geq 7 \times 10^{13}\,M_\odot$ clusters and with no false positives produced in our test region. Importantly, our method is positioned to create a watch list in advance of Rubin's LSST, as it utilizes only existing data, therefore enabling the discovery of lensed transients early within the survey's lifetime.

Key words: gravitational lensing: strong – galaxies: clusters: general – transients: supernovae – transients: gamma-ray bursts – gravitational waves.

1 INTRODUCTION

The study of strongly lensed explosive transients is an emerging field with many applications for modern astronomy and cosmology (see Oguri 2019, for a review). The magnification associated with gravitationally lensed sources acts as a free telescope upgrade, allowing us to observe objects that are fainter and farther from us than would normally be possible. Strong lensing also produces multiple images of the same source, which is valuable if the multiply-imaged source is an explosive transient (a short-time-scale and non-repeating transient – hereafter just referred to as a 'transient'), because it can be used to further constrain the mass distribution of the lens (e.g. Sharon & Johnson 2015) and precisely measure the Hubble constant (or other cosmological parameters) independently of the cosmic distance ladder (Refsdal 1964; Treu & Marshall 2016). In addition, there are currently very few confirmed detections of multiply-

imaged transients – of which all are supernovae (SNe). Within the next decade, we anticipate expansion to other flavours of transient, including gravitational waves and their kilonova (KN) counterparts (Smith et al. 2019a, 2022), as well as gamma-ray bursts (Ahlgren & Larsson 2020). Not only this, but lensed transient detections are expected to increase vastly in number thanks to deep wide-field surveys such as Rubin's LSST (LSST Science Collaboration 2009; Oguri & Marshall 2010; Goldstein, Nugent & Goobar 2019; Wojtak, Hjorth & Gall 2019), which is expected to discover millions of SNe and thousands of KNe over its lifetime (see table 8.2 in LSST Science Collaboration 2009) – a fraction of which will be lensed.

The first resolved and spectroscopically confirmed strongly lensed transient (dubbed SN Refsdal) was discovered by Kelly et al. (2015) during *Hubble* observations of galaxy cluster MACS J1149.6+2223, in which one of the cluster members had lensed a Type II SN (hosted in a background galaxy at $z = 1.49$; Smith et al. 2009) into a four-image Einstein cross-configuration. The modelling of this cluster as a lens provided strong evidence for the existence of two additional images of SN Refsdal – one that should have appeared about 10 yr

★ E-mail: danr@star.sr.bham.ac.uk

prior to this discovery, and one predicted to appear up to around 2 yr in the future (Jauzac et al. 2016; Treu et al. 2016). This future image was indeed detected in concordance with these forecasts about a year later (Kelly et al. 2016), enabling measurements of the Hubble constant by several groups (Vega-Ferrero et al. 2018; Grillo et al. 2020).

The second resolved multiply-imaged SN, iPTF16geu, was of particular interest first because it was discovered within the data stream sourced from a wide-field survey, and secondly for its classification as a Type Ia (Goobar et al. 2017). SNe Ia have standardizable luminosities, meaning it was possible to more precisely estimate the associated lensing magnification. However, this SN could not be used for estimating cosmological parameters as this requires a measurement of the time delays between the appearance of successive images of the lensed SN (see Treu & Marshall 2016, for a review). In general, time delays caused by isolated galaxy lenses (such as iPTF16geu) are shorter than for objects lensed by clusters, and so require a much greater survey cadence in order to precisely capture when new images arrive. Because of its short time delay, all four images of iPTF16geu were discovered simultaneously, no time delay measurement could be made, and therefore no cosmological parameter inference could be completed (More et al. 2017). Another similarly highly magnified SN Ia has recently been discovered with a very short time delay (Goobar et al. 2022).

An intriguing candidate multiply-imaged SN, AT2016jka (or SN Requiem; Rodney et al. 2021), has a much longer predicted time delay of ∼20 yr between its first and final images. This final image of the SN is predicted to appear in an image of the lensed host galaxy close to the core of the cluster lens. Any possible future estimates of H_0 from this system would benefit from the smaller fractional uncertainty achievable from such a long arrival time difference. More broadly, lensed transients with more precise and accurate measurements will help to suppress the uncertainties in H_0 obtained from time delay cosmography (e.g. Birrer & Treu 2021).

In addition to these, many singly-imaged – but gravitationally magnified – lensed SNe have also been reported in the literature (Goobar et al. 2009; Amanullah et al. 2011; Patel et al. 2014; Rodney et al. 2015; Rubin et al. 2018). Detections of these objects still benefit from lens magnification and can still be used to help constrain models of their lens; however, single images do not enable measurements of cosmological parameters.

One observing strategy for finding lensed transients follows a watch list-based approach, whereby wide-field survey telescopes routinely scan the sky for transient events and compare their locations with a list of lens coordinates. A complete lens watch list would encompass as many potential lenses as possible across the entire mass range from massive galaxies to clusters, and should span the entire sky in order to maximize prospects of finding lensed sources. However, previous contents of potential watch lists are limited by either the footprint of particular surveys or lenses selected according to source-plane selection methods – i.e. objects identified from magnitude-limited searches for arc features. This motivates the need for a more general method to identify potential lenses, in order to populate a lensed transient watch list.

When transient detections have been made sufficiently nearby a lens within the watch list, they are flagged for further investigation and follow-up observations. This technique works to filter out candidate lensed transients from the increasingly large number of events detected nightly by wide-field surveys like those conducted currently by the Zwicky Transient Facility (ZTF; Bellm 2014), and that will soon commence with Rubin's LSST. In addition, the watch list method allows lensed transients to be identified even if they have multiple images arriving on time-scales shorter than the discovery survey's cadence, or if multiply-imaged transients have only one detectable image due to variances in magnification. This is advantageous, as lensed transients with these properties would evade detection by a search that instead looks for multiple images by way of detecting spatially coincident transient events within a wide-field survey that are separated by some time delay.

An important requirement of a lens watch list is that it contains objects across the full sky. This is because many transients are now being discovered by facilities that are able to monitor the entire celestial sphere, including gravitational wave interferometers and gamma-ray burst satellites. This is in addition to the optical surveys that have been, and will continue to monitor a significant fraction of the sky like ZTF and Rubin's LSST. Therefore, current catalogues of lenses that only reside in the footprints of particular surveys are insufficient for maximizing lensed transient discoveries. In other words, a significant fraction of the mass function covering the range $10^{12}\,\mathrm{M_\odot} \lesssim M_{200} \lesssim 10^{15}\,\mathrm{M_\odot}$ (Robertson et al. 2020) is missing from the tools used to find lensed transients due to the lack of an all-sky lens catalogue. This is especially important in the Southern hemisphere, which Rubin will begin to survey in the next few years to unprecedented depths. Ensuring such a list of lenses is available as soon as Rubin's operations begin is ideal so that searches for lensed transients can be optimized immediately – maximizing early science prospects and the baseline over which these discoveries can be made.

It is also crucial to emphasize that the dark matter haloes that are efficient lenses for the population of transients in question do not have to have been previously identified as lenses (for example, by identification of arcs or multiple images) to be considered a valid entry in a watch list. This is because, in general, the detectability of lensed transients is independent of whether or not the lensed galaxies that host them can be detected in magnitude-limited surveys – even optical transients such as SNe can easily be brighter than their host galaxy, so host detection does not always come in tandem. Therefore, the selection of objects for a complete strongly lensed transient watch list should not just consist of known lenses, but in fact any object capable of lensing. In other words, lenses should be selected in the lens-plane (i.e. based on their lensing ability), rather than in the source-plane (i.e. based on a chance alignment with source-plane objects detected in magnitude-limited searches for arcs). As shown in Ryczanowski et al. (2020), for the case of group- and cluster-scale objects, ∼95 per cent of $10^{14}\,\mathrm{M_\odot}$ clusters and ∼40 per cent of $10^{15}\,\mathrm{M_\odot}$ clusters would not be identified as lenses based on source-plane selection at the sensitivity of Rubin's early data releases, despite being capable lenses of transients. Therefore, even with next-generation surveys, source-plane lens selection methods will miss a significant fraction of the objects capable of lensing transients. Alternative lens-plane selection methods have been explored previously (e.g. Wong et al. 2013; Stapelberg, Carrasco & Maturi 2019); however, a catalogue of these objects does not yet exist that fulfils both criteria to be all-sky and sufficiently deep to include the majority of the lenses associated with the high-redshift population of lensed transients.

Before discussing lens-plane selection methods further, it is first important to consider what objects should populate the watch list. Ray tracing through hydrodynamical simulations of large-scale structure in the Universe indicates that the optical depth to high magnification ($|\mu| \geq 10$) spans a broad range of halo masses, with ∼50 per cent of it contributed by objects of group and cluster mass

scales ($M_{200} \gtrsim 10^{13}$ M$_{\odot}$; Robertson et al. 2020). This is a key region of parameter space because at fixed lens magnification clusters produce longer and more precisely measured arrival time differences than galaxy-scale lenses due to the former's larger Einstein radii and denser and flatter density profiles (Smith et al. 2022). It is also relatively underexplored in relation to very wide field time domain surveys, with search strategies for lensed SNe concentrating on watch lists comprising individual galaxy lenses (e.g. Goldstein & Nugent 2017).

Previous studies aiming to detect galaxy clusters have historically used a variety of methods. Direct detection is possible through detection of X-rays emitted by the hot intra-cluster gas (e.g. Pacaud et al. 2016), or distortion of the cosmic microwave background by the same medium – known as the thermal Sunyaev–Zeldovich effect (Sunyaev & Zeldovich 1972; Planck Collaboration XXVII 2016b). Clusters can also be discovered through observations of galaxies by the detection of a cluster red sequence (e.g. Rykoff et al. 2014), or by utilizing galaxy positions and photometric redshifts to test for clustering in three-dimensional space (e.g. Eisenhardt et al. 2008). More recently, cluster detection in future wide-field surveys was thoroughly tested using a variety of methods in preparation for the wide survey component that is to be conducted by *Euclid* (Euclid Collaboration 2019). These methods each presented their own merits and challenges, and overall sported great success. Inspired by these methods, we set out to develop our own that satisfies our requirements for an all-sky lens-plane selected cluster/group catalogue. The method of Gonzalez et al. (2019) was of particular interest due to its small number of fundamental requirements and capability to be used alongside existing all-sky survey data, aiming to utilize all-sky *J*-band data from the UKIRT Hemisphere Survey (UHS; Dye et al. 2018) in the north and the VISTA (Visible and Infrared Telescope for Astronomy) Hemisphere Survey (VHS; McMahon et al. 2013) in the south, as well as *W1*-band data from the *WISE* (*Wide-field Infrared Survey Explorer*) mission (Wright et al. 2010).

Therefore, in this paper we describe and test a method with the capability to produce an all-sky watch list of galaxy group- and cluster-scale objects ($M \gtrsim 10^{13}$ M$_{\odot}$) out to $z \simeq 1$, with the intention of using the list to aid in the discovery of gravitationally lensed transients. Our method is based on the principles of Gonzalez et al. (2019) and can detect clusters to sufficient depth using only existing all-sky near-infrared data. The main difference between the method presented here and that of Gonzalez et al. is that our method is tuned to locate groups and clusters out to $z \lesssim 1$ for the purpose of populating an all-sky lens watch list. In Gonzalez et al., the cluster catalogue assembled is limited to the footprints of the surveys used, and focuses on the most massive systems at $z \gtrsim 1$, which are less efficient lenses for the predicted populations of transients. In addition, we formulate a different method of extracting cluster detections and estimating the significance of these detections. Furthermore, our testing concentrates on the southern sky that is due to be surveyed by Rubin in the coming years. Prioritizing this region allows a curated watch list to be available to find lensed transients as soon as Rubin begins operations in ~2024.

The structure of this paper is as follows: Section 2 gives an overview of the surveys and data used in this study. Section 3 describes the cluster-finding method, and Section 4 explains how the method was tested, using both real data and state-of-the-art simulated data, before summarizing in Section 5. Where relevant, and unless otherwise stated, we have assumed a flat cosmology with $H_0 = 67.74$ km s^{-1} Mpc^{-1}, and $\Omega_{\Lambda} = 0.693$ (Planck Collaboration XIII 2016a) and give magnitudes in the Vega system.

2 SURVEYS AND DATA

2.1 Overview of surveys

The main data used in this study are from wide-field surveys conducted with the VISTA and *WISE* instruments. These surveys are used due to their complete coverage in the Southern hemisphere – the region due to be surveyed by the Rubin Observatory once in operation – and also for their access to the wavebands sensitive to galaxies in cluster environments out to $z \simeq 1$. Specifically, we make use of *J*-band photometry in the Southern hemisphere from VISTA and all-sky *W1*-band photometry from *WISE* in order to create maps of the number density of galaxies on the sky. It should be noted that suitable *J*-band data are also available for the Northern hemisphere from the UHS, but we prioritize testing in the Southern hemisphere due to slightly better magnitude depths and to cover the field of Rubin's LSST. Sources are first matched between the *WISE* and VISTA catalogues, giving multiband coverage of each galaxy – this ensures that all detections are robust, and allows estimates of further properties from $J - W1$ colours such as the redshift. The following sections describe the data sets in more detail.

2.2 VISTA data

The VISTA is a 4.1 m survey telescope located at Paranal Observatory, Chile. It has conducted a variety of wide-field surveys in the sky above the Southern hemisphere and equator, and between these the entire southern sky has collectively been surveyed. In this work, we directly use data from the largest of these surveys: the VHS ($\sim$20 000 deg^2; McMahon et al. 2013) and the VISTA Kilo-Degree Infrared Galaxy Survey ($\sim$1500 deg^2; Edge et al. 2013). These surveys provide complete coverage in two near-infrared bands, *J* and *Ks*, with additional coverage in other bands in specific regions. In creating our maps, we utilize the *J*-band (1.25 μm) data that have a 5σ detection limit of at least $J(5\sigma) = 20.1$ across the surveys. All VISTA data used are publicly available through the online VISTA science archive.

2.3 *WISE* data

The NASA *Wide-Field Survey Telescope Explorer* (*WISE*) is a space-based 0.4 m infrared telescope that surveys the entire sky in four infrared bands, named *W1* to *W4* with wavelengths of 3.4, 4.6, 12, and 22 μm, respectively. We use sources detected specifically in the *W1* band – the most sensitive of the *WISE* passbands – from the CatWISE2020 catalogue (Marocco et al. 2021), a compilation of almost 2 billion sources collected by the instrument from its first operation in January 2010 up to December 2018.

The addition of the *WISE* data is useful for assuring robust galaxy detections by requiring each *WISE* source to match to a nearby *J*-band detection from VISTA, and to ensure that the galaxies have infrared colours consistent with old stellar populations out to $z \simeq$ 1. Fig. 1 shows the predicted evolution of an $L^{\star}$ galaxy's[1] $J - W1$ colour using the EzGal modelling tool (Mancone & Gonzalez 2012).[2] This shows that with the magnitude limits of each survey, detecting cluster galaxies down to at least $L^{\star}$ is attainable out to $z \simeq 1$. When modelling the galaxy evolution, we assume a Bruzual and Charlot

---

[1]L∗ is a characteristic luminosity scale present in the Schechter (1976) luminosity function, typically representing the value below which the number density grows exponentially.
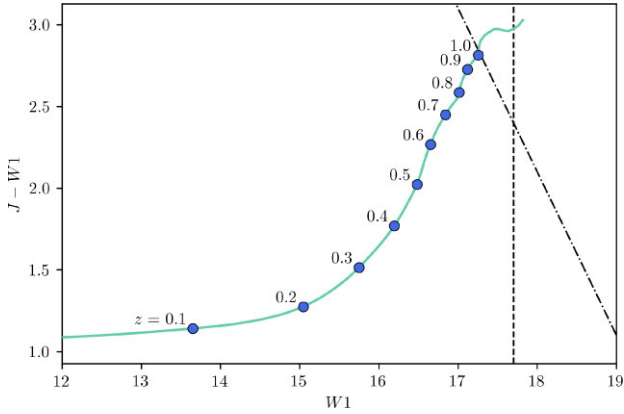[2]http://www.baryons.org/ezgal/

Figure 1. Predicted $J - W1$ colour evolution of an $L^\star$ cluster galaxy as a function of apparent $W1$ magnitude (blue solid line). Orange points highlight the values at specific redshifts during this evolution. The vertical black dashed line represents the 5σ $W1$-band magnitude limit of CatWISE ($W1 = 17.7$), and the dot–dashed diagonal line represents the 5σ J-band magnitude limit of VHS ($J = 20.1$), the largest of the VISTA surveys. Magnitudes are obtained using the EzGal tool assuming a Bruzual and Charlot evolution, with an SSP star formation history, Chabrier initial mass function, and solar metallicity. Given that $z = 1$ $L^\star$ cluster galaxies reside within these detection limits, the detection of clusters out to this redshift is tractable.

model (Bruzual & Charlot 2003), with a single delta burst of star formation at $z = 3$ that follows a Chabrier initial mass function (Chabrier 2003) and contains stars of solar metallicity. We also normalize the luminosity of an $L^\star$ galaxy to the Lin, Mohr & Stanford (2004) cluster sample, which has $K_S^\star = 15.5$ at $z = 0.25$. $W1$ was chosen for this purpose over the K band, as $W1$ has full coverage in both hemispheres, and the K band in current surveys is not sensitive enough to reach $z = 1$.

One important consideration when utilizing the *WISE* data is the large point spread function (PSF) [$W1$-band full width at half-maximum (FWHM) $\simeq 6$ arcsec] compared to the VISTA data (J-band FWHM $\simeq 1$ arcsec). This means that in dense cluster cores, there is a possibility that multiple galaxies detected by VISTA are blended into a single *WISE* detection. To take account of this, when matching sources between the two catalogues, the number of VISTA sources within one PSF half-width of a *WISE* detection is counted and used to weight the contribution of that detection to the galaxy map. The existence of blended sources in dense cluster cores also makes it difficult to determine accurate colours for these central galaxies and their neighbours. For example, within our sample of known strong lensing clusters described in Section 4.1, approximately 10 per cent of *WISE* galaxies within 1 arcmin of their respective cluster centre have multiple matches in the J band. However, as we do not make any explicit cuts based on the colours, this does not directly affect the density maps, but does have a visible effect on colour–magnitude diagrams we produce during testing in Section 4.3.

Our selection criteria for galaxies are relatively straightforward, and consist of a flat cut of the 5σ detection limit in the relevant bands in both surveys ($J = 20.1$ for VHS, the largest of the VISTA surveys and $W1 = 17.7$ for *WISE*) and a cut on the VISTA catalogue's *pGalaxy* attribute that estimates the probability that a source is a galaxy. We require *pGalaxy* > 0.9 to eliminate any interloper stars, although the majority of the sample has *pGalaxy* > 0.99, indicating that these are secure detections of extended sources.

3 METHOD

3.1 Background

Galaxy cluster members are typically early-type galaxies that emit strongly in the near-infrared, due to the spectrum of metal-poor population II stars ($T \simeq 3000$K) that dominate early-type galaxies' peaks at a rest-frame wavelength of $\lambda \simeq 1\,\mu$m. Utilizing the $W1$ band of *WISE* and the J band of VISTA, we produce density maps of galaxies detected in both catalogues. This is done following a similar methodology to Gonzalez et al. (2019), whereby a raw map of galaxy positions is convolved with a difference-of-Gaussians smoothing kernel, but we use slightly different data sets to produce the maps. This convolved map then provides an estimate of the local density of galaxies within that region of sky, and peaks within the map can hence be used to unveil candidate galaxy groups and clusters within the wide-field data. We then develop a new method to estimate the significance of peaks referring to cluster detections within the maps.

Given that no specific colour or redshift information is taken into account when selecting galaxies, it is entirely possible that multiple chance alignments of smaller groups can produce signals in the maps comparable to those from richer clusters. While it might initially appear that such detections are problematic, such cases are still considered to be high-density lines of sight and are therefore still valuable from a strong lensing perspective, and hence belong in a lensed transient watch list. Therefore, we can motivate using this method to assemble a list of the densest regions of the sky as traced by collections of galaxies detected in infrared data. Furthermore, we are more concerned with the fast creation of a watch list in advance of Rubin, which can then be refined based on specific halo properties once higher quality data are obtained following commencement of the LSST. These lines of sight may not contribute the same lens potential as a singular massive halo, but filtering them would risk reducing the number of interesting lines of sight too significantly – of the 270 serendipitous detections used for estimating false positive rates in Section 4.3, 40 ($\sim$15 per cent) show 2 or more distinct photometric redshift peaks, indicative of multiple haloes aligned along the line of sight. Identification of such cases also requires additional photometric redshift data that are not available all-sky, so removing these detections on a large scale is not feasible with current data.

3.2 Creating density maps with kernel convolution

To create a raw density map, galaxies matched between the two catalogues that pass the selection criteria are placed into a 0.5×0.5 deg pixel grid with 7.5 arcsec/pix scale. This is almost identical to the approach used by Gonzalez et al. (2019), except using a slightly different pixel scale and we assign our maps to cover smaller regions, allowing for multiple disconnected regions to be mapped. This is done specifically so we can test the method on smaller regions containing known group- and cluster-scale objects. Each pixel contains a value representing how many galaxies exist within that region of space, and we take into account the large PSF of *WISE* by weighting the contribution of a single *WISE* detection by the number of J-band detections from VISTA within the *WISE* PSF radius. This ensures we are not underestimating the number density due to blended sources in dense cluster cores where there can be up to to of order 10 blended galaxies.

Once the raw map has been constructed, it is convolved with a kernel designed to smooth out contributions from components that do not match those of typical cluster core scales. The kernel is described
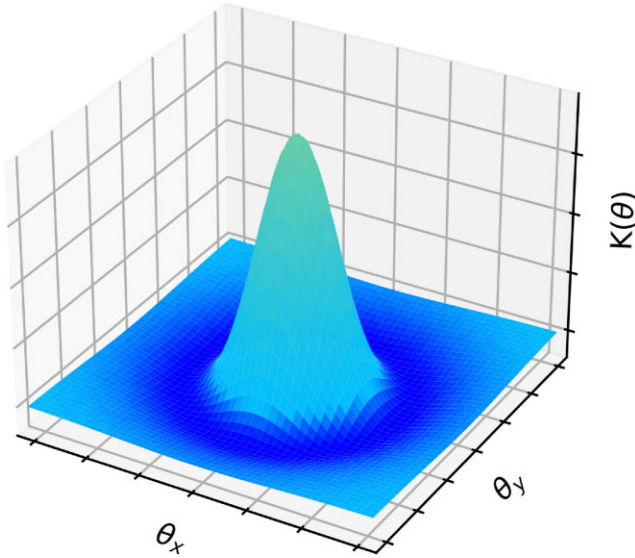
Figure 2. Example surface plot of a normalized difference-of-Gaussians kernel, showing the positive central peak and the negative region surrounding it. This is used as a smoothing function to identify cluster-like structures within pixellated galaxy maps. The angular extent of the interior width is tuned to that of cluster cores, and acts to smooth out contributions from scales larger and smaller than this.

by a difference of two normalized Gaussian functions with different widths given by

$$K(\theta) = \frac{1}{2\pi\,\sigma_{in}^2\sigma_{out}^2}\left[\sigma_{out}^2\exp\left(\frac{-\theta^2}{2\sigma_{in}^2}\right) - \sigma_{in}^2\exp\left(\frac{-\theta^2}{2\sigma_{out}^2}\right)\right], \qquad (1)$$

where θ is the angular separation and σ_{in} and σ_{out} are the widths of the inner and outer Gaussians, respectively. By definition, $\sigma_{in} < \sigma_{out}$. Each Gaussian is normalized to have unit volume, such that the kernel integrates to zero. Similarly to Gonzalez et al. (2019), we select $\sigma_{in} = 45$ arcsec, corresponding to the scale of the dense cores of galaxy clusters, and choose the ratio of σ_{in} to σ_{out} to be 1:6. Multiple values were experimented with for the filter widths, anticipating an effect based on the changing scales of clusters at different redshifts – however, the final results were insensitive to any change of these parameters. Fig. 2 shows an example of a difference-of-Gaussians function as given by equation (1).

Applying the convolution produces a map whose pixels correspond to the local galaxy number density; Fig. 3 shows an example map centred on known galaxy cluster and strong lens Abell 1689.

3.3 Map size and edge effects

Within each density map, a peak-finding algorithm is utilized to identify positions of the densest regions of the map. To reduce the impact of edge effects caused by the convolution on outer regions of a map, we ignore any peaks that are within 24 pixels of the map border. This effectively reduces the size of each map to 0.3×0.3 deg, but ensures that no bias is introduced to outer regions. The reduced size of map is a fairly arbitrary choice that has been seen to work well on test regions containing known objects – the average pixel values around the edge of the maps are not significantly different to typical values closer to the centre, signifying that the edge effects are not prominent. The (reduced) map size can be optimized for efficiency in a blind search by determining the maximal usable area for a given map that is not noticeably affected by any edge effects.
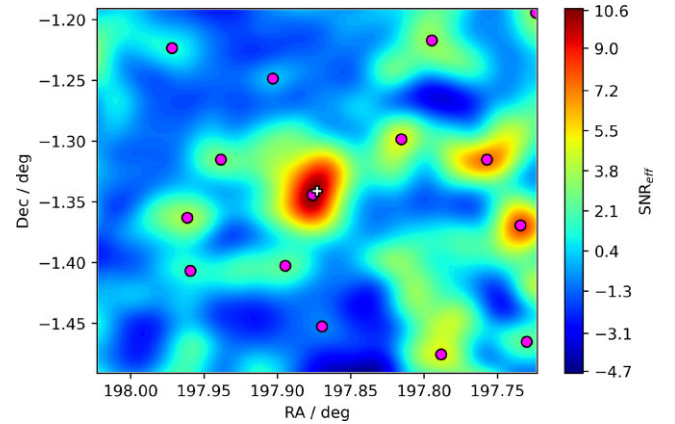

Figure 3. A 0.3×0.3 deg map produced by convolving a difference-of-Gaussians kernel with a pixel grid of galaxy positions surrounding the coordinates of known cluster and strong lens Abell 1689. Purple dots highlight peaks in the galaxy number density distribution, and the white plus marks the coordinates of the cluster centre, which in this case is located very close to the largest peak. The colour bar scale shows the SNR_{eff} of each pixel, a quantity introduced in Section 3.4 that indicates the significance of a detection, and is normalized between the maximum and minimum values within the map. Data to produce this map are taken from an overlapping 0.5×0.5 deg region, extending an additional 0.1 deg from each edge of the map. This is to ensure there are sufficient data beyond the map region to prevent any edge effects caused by the convolution. The cluster in the centre is recovered with high significance, and a few other regions are highlighted as having a high density along the line of sight and hence represent candidate groups/clusters.

This can be done by looking at how the average pixel value changes in successively larger annuli further from the centre of the post-convolution map, and choosing a suitable tolerance for how much edge effects can vary pixel values, but we leave the investigation of this for future work.

The limits on the size of a map are based only on the memory limits of the machine processing the data. We chose to use the size we did to efficiently create maps for many disconnected regions containing known group- and cluster-scale objects, which is efficient for testing. However, a blind search over a larger continuous footprint could in theory be done as a smaller number of much larger maps on a more powerful machine, which would significantly reduce the fraction of regions affected by edge effects.

3.4 Estimating overdensity significance

Once the peak-finding algorithm has been run on the reduced map, we quantify the significance of each overdensity in a map. This is done by adopting a quantity similar in form to a signal-to-noise ratio (SNR), which we call the 'effective SNR', SNR_{eff}, and calculating this for each pixel in the convolved map. We then take the largest SNR_{eff} pixel value of each local maximum to represent the significance of a detection at that position. This is a new way to formulate the significance of each candidate cluster. To calculate this quantity, we first assume that the background of non-clustered galaxies is randomly and uniformly distributed. Under this assumption, we can produce a large set of maps of randomly distributed galaxies, where each map contains the same total number of galaxies as the real map now randomly positioned across the same size patch of sky. These random maps are then convolved by the same difference-of-Gaussians kernel, and can be used as a metric to estimate the background noise. By calculating the mean and standard deviation

of pixels within these random maps, we define SNR_{eff} for each pixel within real maps as

$$SNR_{eff} = \frac{s - \mu}{\sigma} = \frac{s}{\sigma}, \qquad (2)$$

where s is a pixel value, and μ and σ are the pixel mean and standard deviations, respectively, determined from the sample of random maps. By construction, any pixel map convolved with a kernel of the form given by equation (1) will have $\mu = 0$, due to the kernel's property of integrating to zero, hence the additional simplification in the equation and thus leaving only σ to be determined. In a random map containing N_{gal} galaxies, the pixel values follow a Poisson distribution and so $\sigma \propto \sqrt{N_{gal}}$. Given that the smoothing kernel acts only to spread the pixel values from a raw map over a larger region, one can expect the standard deviation of pixel values within the entire convolved map to follow the same distribution. Therefore, we can precisely measure the standard deviation σ' for a single arbitrary number of galaxies N'_{gal} and calculate SNR_{eff} in proportion for the relevant N_{gal}:

$$SNR_{eff} = \frac{s}{\sigma'} \sqrt{\frac{N'_{gal}}{N_{gal}}}. \qquad (3)$$

Consequently, random maps only need to be created once, with new random maps being required only if a fundamental property of the method, such as the kernel size, is changed. By creating 10^4 random maps each containing $N'_{gal} = 2000$ galaxies, we determine $\sigma' = (8.7 \pm 0.4) \times 10^{-3}$ galaxies per pixel. These values are then used to determine SNR_{eff} within the pixels of the real maps, which is in turn used to assess the confidence that a given local maximum within a map corresponds to a real group- or cluster-scale object, and comparatively rank any peak determined by this method by way of proxy for the mass or richness of the cluster core. We discuss in Section 4 the values of SNR_{eff} that categorize the robustness of detections. It should be stressed that SNR_{eff} is not a true SNR, but rather an estimator based on correlated nearby pixels that aims to evaluate the overdensity of a region of galaxies under simple assumptions.

Given that the non-clustered background galaxy distribution will not be truly random but rather correlated on small scales to, on average, be denser than random, SNR_{eff} will be a slight underestimate due to underestimating the magnitude of the noise. However, we believe this effect to be small due to the amplitude of the galaxy two-point correlation function at scales similar to those of our maps. Fig. 15 in Wang, Brunner & Dolence (2013) shows a determination of this quantity – their faint population of galaxies is approximately representative of the typical L∗ cluster galaxies we expect to be detectable with *WISE*, based on colours inferred from the same EzGal models introduced in Section 2.3. Based on their results, galaxies at scales similar to our maps (∼0.1 deg) are ∼5 per cent denser than a purely random distribution. Therefore, we expect the difference made due to our simplifying assumptions to be small. In addition, the rank ordering of objects based on SNR_{eff} is arguably more useful than the specific values of SNR_{eff} alone, and this would largely remain unchanged as a result of this assumption.

4 TESTING THE METHOD

4.1 Known cluster/group-scale lens test sample

We use two samples of lenses to test our method's ability to recover known groups and clusters – 130 spectroscopically confirmed cluster-scale lenses assembled by Smith et al. (2018b), and 98 galaxy-

and group-scale lenses assembled by Carrasco et al. (2018). The cluster-scale lenses have been utilized in previous searches for lensed gravitational waves and SNe with a prototype watch list (Smith et al. 2018a, 2019b; Ryczanowski et al. 2021; Bianconi et al. 2022). These 130 objects are some of the best studied strong lenses in the literature, many of which have been observed by the *Hubble Space Telescope*. This sample ensures that our test sample extends to the most extreme clusters in terms of mass, as the Einstein radii range from 3 arcsec ≲ θ_E ≲ 60 arcsec, and hence includes individual objects with some of the largest strong lensing cross-sections known. The distribution of these clusters peaks at $z \sim 0.35$, but the upper tail stretches to $z \sim 1$, which suitably matches the distributions of objects we aim to detect, as outlined in Section 1.

Objects from the Carrasco et al. (2018) sample are lenses and lens candidates that were found within the Canada–France–Hawaii Telescope (CFHT) Lensing Survey (CFHTLenS) data. CFHTLenS is a deep optical ($u*/g'/r'/i'/z'$ band) survey spanning 154 deg^2 observed by the CFHT as part of the Canada–France–Hawaii Telescope Legacy Survey. The survey's primary objectives involve studies of weak lensing effects; however, the high-quality data along with accurate photometric redshifts make it useful for some strong-lensing science as well. These group- and galaxy-scale objects typically have smaller Einstein radii than those in the cluster sample, ranging from 3 arcsec ≤ θ_E ≤ 18 arcsec, which is valuable for testing the limits of what the method can reliably detect. The redshift range of $0.2 < z < 0.9$ is also well matched to our aims and the cluster sample.

We were able to retrieve the VISTA data for 63 out of the 98 objects within the Carrasco et al. (2018) sample. The majority of the objects that do not have the VISTA data reside in the third CFHTLenS region around $\delta = 55$ deg – far outside of the coverage of the surveys. We were also only able to obtain data for 62 out of the 130 samples of known lenses, giving a total of 125 objects in our final test sample. Ultimately, the only information we utilize in our method beyond selection is the coordinates of objects from the samples and use them to obtain VISTA and *WISE* data from the surrounding regions. This allows us to test our method fully, extract the SNR_{eff} distribution for regions containing known objects, and hence estimate the recovery rate of the method.

4.2 Results of the known lens test

For each of our sample objects, we produce a convolved density map using the difference-of-Gaussians kernel and the procedure as described in Section 3.2. Within each map, we locate peaks within $5\sigma_{in}$ (225 arcsec) of the object's coordinates given in the sample and calculate SNR_{eff} for each of these peaks using the same method discussed in Section 3.4. If the SNR_{eff} around the peak closest to the object is above some designated threshold (which we initially define to be 5), then the object is considered to be recovered with high significance.

Fig. 4 shows a plot of the calculated SNR_{eff} against estimated Einstein radius, θ_E, for our 125 sample objects. It should be noted that the θ_E values for objects in the Carrasco et al. (2018) sample are taken directly from their catalogue and are only estimates found, as they describe in their paper, by finding the average distance between the bright arc (or candidate arc in the unconfirmed cases) and the object's centre. We use θ_E as a proxy for the lensing cross-section (and equivalently, mass) of each object as is commonplace in many lens models. Objects with small θ_E tend to produce smaller SNR_{eff}, making up the majority of non-recovered objects with SNR_{eff} < 5. In addition, above $\theta_E \sim 5$ arcsec, there is a slight correlation between Einstein radius and SNR_{eff}, with the data in Fig. 4 showing
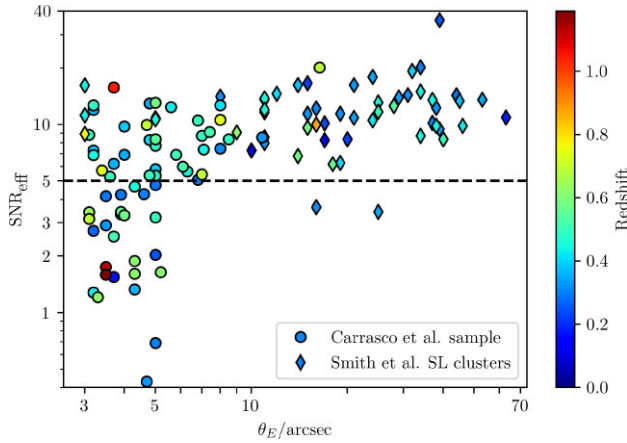
Figure 4. Distribution of SNR_{eff} against Einstein radius for the 125 test objects, made up of those from the Carrasco et al. (2018) sample (circles) and the Smith et al. (2018b) SL sample (diamonds). Horizontal dashed line at $SNR_{eff} = 5$ marks the threshold above which an object is classed as recovered with high significance. The colour of the points represents the redshift of the object, signified by the colour bar. Recovery appears to be related to Einstein radius, as the majority of non-recovered objects have small θ_E ($\theta_E \lesssim 5$ arcsec). Recovery also does not appear to depend strongly on redshift, as shown by the absence of any obvious trends with redshift.
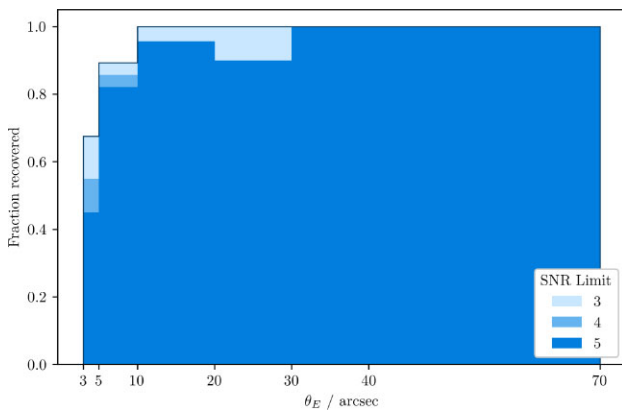
Figure 5. Histogram showing the recovery fraction of objects in bins of Einstein radius. Different colour bars represent different SNR_{eff} thresholds to define recovery. Even with the stringent requirement of $SNR_{eff} > 5$, the recovery fraction of our sample is >80 per cent for $\theta_E > 5$ arcsec, and 100 per cent of $\theta_E > 30$ arcsec objects are recovered, with an average of 91 per cent recovered for $SNR_{eff} = 5$. The reason for the drop-off below 5 arcsec is partially due to the sample containing isolated galaxy lenses as well as smaller galaxy groups, and this method is naturally less sensitive to objects with fewer members. Using a lower SNR_{eff} threshold of 3 allows for complete recovery of all $\theta_E > 10$ arcsec objects, but this will come at the expense of a greater number of false positive detections.

a Spearman correlation coefficient of $r_S = 0.504$, indicating that clusters with larger Einstein radii generally produce slightly higher SNR_{eff}. As there is inherent scatter in the Einstein radii for clusters of a given mass or richness in both the Smith et al. (2018b) and Carrasco et al. samples, we would not expect a perfect correlation between the two variables. These results are further highlighted by Fig. 5, which shows the fraction of objects recovered in bins of θ_E. The recovery fraction drops significantly for $\theta_E < 5$ arcsec, due to the inclusion of single galaxy and small group lenses within the Carrasco et al. sample, but is much higher (>80 per cent recovery) above this.
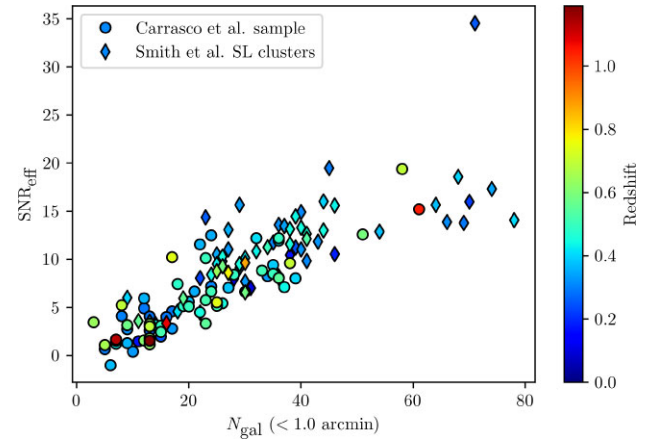

Figure 6. SNR_{eff} of objects in the two test samples against the weighted number of galaxies detected within 1 arcmin of the corresponding map peak. The colour of each point indicates the redshift of the object. The SNR_{eff} appears to correlate with the number of galaxies, although there is some clear scatter present. Therefore, SNR_{eff} is a reasonable estimator for the richness of an object.

Given that $SNR_{eff} = 5$ is fairly arbitrary, since SNR_{eff} is only an estimator for an object's overdensity, we can vary the cut to see the effect on recovery. A cut at $SNR_{eff} = 3$ increases the recovery rate at low Einstein radii fairly significantly, and allows the remainder of the larger radius objects to be recovered, giving 100 per cent recovery for $\theta_E > 10$ arcsec. However, decreasing the SNR_{eff} threshold will increase the proportion of false positives in a blind search, as we explore in the next section, so this must be done with caution.

The drop in recovery fraction for the $\theta_E < 5$ arcsec objects is unsurprising, as some of these objects are single galaxy lenses or small galaxy groups, which will naturally be either impossible or very difficult to detect by any method sensitive to number density. It is also interesting to note that Fig. 4 suggests the redshift of the object does not appear to have a major impact on whether the object is recovered or not. There are objects with a wide range of redshifts both above and below $SNR_{eff} = 5$ – although the impact on SNR_{eff} due to redshift at the upper end of the range ($z \gtrsim 0.8$) is not well tested due to the scarcity of these objects.

Fig. 6 shows the SNR_{eff} against the weighted number of galaxies detected in the matched VISTA–*WISE* catalogues within 1 arcmin of the map peak closest to the coordinates of each lens in our sample. As explained in Section 3.2, the weighting is done based on the number of blended *J*-band detected galaxies within the PSF of each *WISE* detection. This affirms that the method is more sensitive to objects with a larger member number density, given that there appears to be a scattered but approximately linear relationship between the number of selected galaxies within 1 arcmin of the peak and the corresponding SNR_{eff} (with Spearman correlation coefficient $r_S = 0.870$). We suspect that the larger r_S for this relation compared to SNR_{eff} against Einstein radius in Fig. 4 is because the cluster richness will directly influence the pixel values within the map that drive the SNR_{eff} calculation, and additionally because there is inherent scatter between Einstein radius values for a given cluster mass or richness. The calculations of N_{gal} here ignore the fact that each 1 arcmin region will be contaminated by galaxies that are not part of the group or cluster; however, by assuming contaminant galaxies are randomly distributed (or close to it), this should only contribute to the scatter that would be overshadowed by the cluster overdensity.
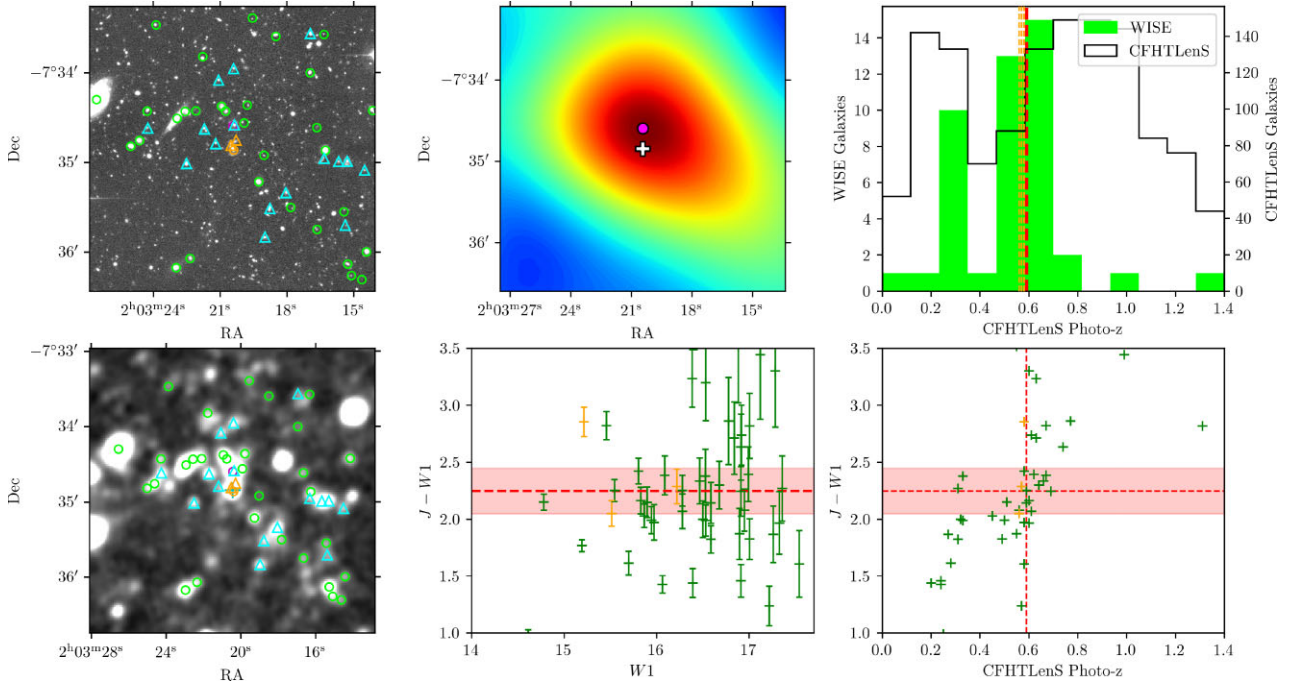
Figure 7. Example of a series of plots used to confirm the existence of objects within our test sample, at the coordinates provided by the original source. This series represents object SA6 from the Carrasco et al. (2018) sample. Top left: An image of the region from optical CFHTLenS data showing the galaxies detected within the *J* and *W*1 data. Top centre: A zoomed-in version of the density map, centred on the largest peak, highlighting the closeness of the peak to the recorded coordinates of the object. Top right: Histograms of the redshifts of galaxies within the CFHTLenS catalogue (black), and a subset of those that also appear in the *WISE* data (green). This shows two distinct redshift peaks at $z \sim 0.3$ and 0.6, indicating the presence of real collapsed objects at these redshifts. Bottom left: Same as top right, except using *WISE* imagery. Bottom centre: A *J*–*W*1 colour–magnitude diagram for the galaxies within this region, showing a cluster red sequence in accordance with the predicted colour for the object. Bottom right: Colour against photometric redshift for galaxies in the colour–magnitude diagram. This shows a cluster of points around the predicted colour and recorded photometric redshift of this object, furthering the evidence for a real object at this location.

Within Figs 4 and 6, there are two clear anomalous data points that stand out from the rest. First is the high redshift ($z \sim 1$) and high SNR_{eff} ($\simeq 15$) point. This particular object is a high-redshift group that is closely aligned with a larger and closer cluster ($z \sim 0.3$) along the line of sight. Therefore, the detection appears with an inflated N_{gal} and hence SNR_{eff} compared to objects at a similar redshift. The other is Abell 1835 – the object at $SNR_{eff} \simeq 35$, which is significantly higher than any other object in our sample, especially considering there are other clusters with similar N_{gal}. This object in particular has a large number of highly weighted galaxies near to the cluster core; hence, our method assigns this cluster a high SNR_{eff} even among the large N_{gal} objects.

4.3 False positive rate estimates

Our method is based on counting galaxies without considering their colours. This is well matched to the available all-sky data given that the colour of galaxies in crowded cluster and group cores is challenging to measure accurately given the mismatch in angular resolution between ground-based *J*-band ($\lambda = 1.25\,\mu m$) and space-based *W*1 ($\lambda = 3.4\,\mu m$) photometry. In contrast, modern galaxy-based cluster- and group-finding algorithms typically incorporate colour information (e.g. Rykoff et al. 2014; Gonzalez et al. 2019). Objects detected by our method may therefore include projections of more than one group or cluster along the line of sight, and chance alignments of galaxies at very different redshifts that are indistinguishable from real groups or clusters when colour information

is not taken into account. Given that our focus is on identifying group- and cluster-scale dark matter haloes capable of strong lensing, we regard the former (multiple groups and clusters) as a secure detection, and the latter (chance projection of field galaxies) as false positives.

We estimate the false positive rate from the population of detections in the sky regions surrounding the known lenses from the Carrasco et al. (2018) sample. Specifically, we check whether there are obvious peaks in the CFHTLenS photometric redshift catalogue associated with detections other than the known lens in our 0.3×0.3 deg *J*/*W*1-based maps centred on these lenses. In practice, this involved inspecting figures following the same format as Fig. 7, which contains a series of images and plots produced using data from CFHTLenS, as well as *J* and *W*1 photometry, in order to confirm whether the peaks correspond to real objects.

The left-hand column of Fig. 7 shows two images of the region in optical (top, from CFHTLenS) and in *W*1 (bottom, *WISE*). Cyan triangles mark the most likely cluster members, as these galaxies have colours consistent with the EzGal model, while green circles mark those that have colours inconsistent with the model, and are therefore less likely to be members. Orange symbols are galaxies within 10 arcsec of the provided centre coordinates for this object, and are candidates for the brightest central galaxies of the cluster. The top centre panel shows a zoomed-in version of the density map (3.75 arcmin × 3.75 arcmin), centred on the largest peak. The white plus signifies the coordinates of the object centre and the purple circle is the location of the peak. The top

right panel shows histograms of the redshifts of galaxies within the CFHTLenS catalogue (black), and a subset of those that also appear in the *WISE* data (green). The object's recorded redshift is marked with a red dashed line, and the photometric redshifts of the most central galaxies are marked with orange dashed lines. The bottom centre panel shows a $J–W1$ colour–magnitude diagram, with the red dashed line indicating the colour predicted by the EzGal model for an $L^\star$ cluster galaxy at the object's recorded redshift. The red transparent band marks a fiducial ±0.2 mag error on this value. The bottom right panel shows colour against the CFHTLenS photometric redshift for galaxies in the colour–magnitude diagram, with identical red lines to the bottom centre and top right plots showing the expected colour for the object and the recorded photometric redshift.

The primary resource for identifying whether the object is indeed a true detection is using the top right plot and counting the number of peaks in the redshift distribution. In this case, using object SA6 that appears in Fig. 7 as an example known lens, we would count two peaks, one at $z \simeq 0.3$ and one at $z \simeq 0.6$. We also stress that qualitatively the serendipitous detections near to SA6 are typical of those found in the outskirts of the maps. Note the broad distribution of $(J - W1)$ colour at the latter redshift, which arises due to our method in which we cope with the multiple possible matches for galaxies at the centre of dense cluster cores, occurring due to the PSF of *WISE* being distinctly larger than that of VISTA. In the absence of relying on galaxy colours, galaxies are matched based on the closest entry in the other catalogue, up to a maximum separation of 1.4 arcsec. Multiple matches are therefore common in high-density cluster cores where multiple J-band detections are separated by less than the *WISE* PSF. This leads to inconsistent colour measurements either because a $W1$-band flux measurement contains contributions from multiple individual J-band sources (leading to an overestimate of $J - W1$) or because galaxies in different wavebands are simply matched incorrectly.

In total, we examine 269 serendipitous detections with $\mathrm{SNR_{eff}} \geq 3$, of which 70 have $\mathrm{SNR_{eff}} \geq 5$. At $\mathrm{SNR_{eff}} \geq 5$, 66 detections are associated with one or more redshift peaks and therefore are classified as secure detections in that there is clear evidence of one or more collapsed group/cluster-scale haloes along the line of sight. This translates to a false positive rate of $\simeq6$ per cent at $\mathrm{SNR_{eff}} \geq 5$. At $3 \leq \mathrm{SNR_{eff}} < 5$, we find that 158/199 serendipitous detections are associated with one or more redshift peaks, giving 45 total detections for $\mathrm{SNR_{eff}} \geq 3$ that are not associated with redshift peaks. This implies a false detection rate of $\simeq17$ per cent at $\mathrm{SNR_{eff}} \geq 3$. We also find no evidence of a strong trend in false positive rate at $3 \leq \mathrm{SNR_{eff}} < 5$, and expect an increase in false positive rate below $\mathrm{SNR_{eff}} = 3$ due to a large increase in the number of detections. To summarize, including $\mathrm{SNR_{eff}} > 3$ detections only increases the false positive rate of detections to 17 per cent while producing almost four times as many real detections than if we include only the $\mathrm{SNR_{eff}} > 5$ detections. In addition, when comparing to our tests on objects from our test sample (Section 4.2), extending the detection threshold down to $\mathrm{SNR_{eff}} \geq 3$ increases the recovery of lenses with small Einstein radii, 3 arcsec $< \theta_E \leq 5$ arcsec from $\simeq40$ to $\simeq70$ per cent (Fig. 5). This is encouraging for building a watch list based on J- and $W1$-band photometry in advance of LSST, especially bearing in mind that photometric redshifts from LSST photometry and spectroscopic redshifts from surveys taking advantage of instruments such as 4MOST will be efficient tools to subsequently suppress the false positives. Therefore, we believe $\mathrm{SNR_{eff}} > 3$ to be a sufficient threshold for selecting detections of real objects to include in a lensed transient watch list.
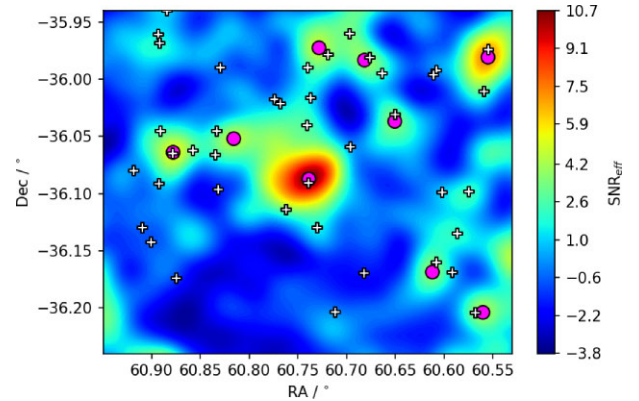

Figure 8. The map produced for the test region of simulated Rubin data. Magenta dots mark the locations of peaks with $\mathrm{SNR_{eff}} > 3$, and white pluses mark centres of dark matter haloes with mass $M \geq 10^{13}\,\mathrm{M_\odot}$. Qualitatively, the distribution of dark matter haloes follows the high-$\mathrm{SNR_{eff}}$ regions of the map, which is indicative that the method is functioning as expected. The nine peaks are all located within 1 arcmin of a dark matter halo, indicating that the method has a low false positive rate when cutting at $\mathrm{SNR_{eff}} = 3$.

4.4 Testing on Rubin DP0 data

The Rubin Observatory Data Preview 0 (DP0) is a programme providing access to simulated LSST-like data and images. These data were produced following the Dark Energy Science Collaboration's (DESC) second data challenge (Korytov et al. 2019; LSST Dark Energy Science Collaboration (LSST DESC) et al. 2021), whereby the results of *N*-body simulations are passed through a model of the Rubin hardware before being processed by the official LSST pipelines, producing a catalogue of objects that is representative of what is to be expected from the Rubin telescope. It includes data for 300 square deg of sky in all six Rubin bands and is based on stacked images containing 5 yr of observation data. The objects within the simulation reach out to $z = 3$; however, not every object will necessarily be detected following the processing steps. The base simulated data set contains information on individual dark matter haloes and the galaxies bound to them, presenting an opportunity to test our method on data where the true underlying distribution of dark matter haloes is known.

We applied our method to a single 0.5×0.5 region of the DP0 catalogue that was known to contain a significant number of massive and less massive haloes. We select galaxies by choosing extended objects that are detected in the y band brighter than 21.1 mag, which corresponds to the J-band detection limit of VISTA corrected to the y band using the expected $y - J$ colour for an $L^\star$ galaxy as calculated using the same EzGal model as described in Section 2.3. Our selection is therefore tuned to be as similar as possible to the current detection limits of VISTA J band, and makes for a reasonable comparison despite the DP0 data set being significantly deeper. All other parameters used in creating and convolving maps are identical to those used with the VISTA and *WISE* data as described in Section 3.

Within the simulation test region, there are 41 dark matter haloes with $M \geq 10^{13}\,\mathrm{M_\odot}$ at $z < 1$. We restrict to this redshift as the selection is tuned to the detection limits of $L^\star$ cluster galaxies in the J band as shown in Fig. 1, which we estimate to be detectable up to $z \simeq 1$. Fig. 8 shows the map for the test region, with the locations of the 41 dark matter haloes overlaid. First, it is reassuring to note that the qualitative distribution of these dark matter haloes mostly follows the higher $\mathrm{SNR_{eff}}$ regions of the map, which matches expectations. The map contains nine peaks with $\mathrm{SNR_{eff}} > 3$, all of which are
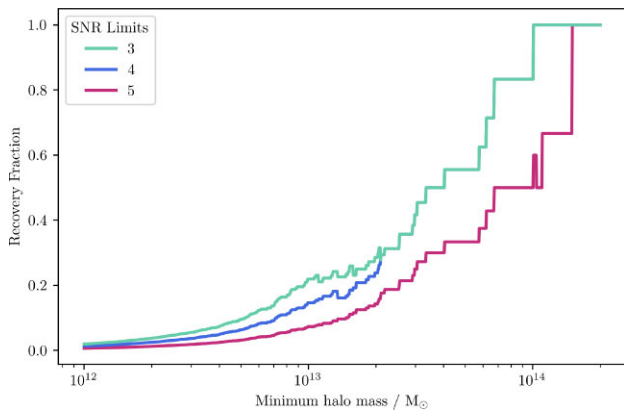
Figure 9. The recovery fraction of dark matter haloes within the test region of simulated Rubin data as a function of halo mass. We calculate the fraction of haloes greater than a given mass that are recovered by the method, where an object is considered to be recovered if it is the closest object within 1 arcmin of an $SNR_{eff} > 3$, 4, or 5 peak (note that the $SNR_{eff} = 3$ and 4 curves are identical above $M \sim 2 \times 10^{13} M_{\odot}$). The recovery rate using the simulated data matches well to the real data, with a good rate at high masses, which falls off at smaller values due to the difficulty in detecting those with fewer members. The ability to recover objects associated with massive dark matter haloes does not explicitly match the goals of using the method, but it is sufficiently related, easy to test, and gives a reasonable metric for the method's capability.

within 1 arcmin of at least one of these haloes. Therefore none of these peaks can be considered a false positive detection, based on the definition established in the previous section, which further affirms and is consistent with our previous evaluation that the false positive rate of the method is low when cutting at $SNR_{eff} \geq 3$.

In order to quantify the recovery fraction of objects within the test region, we assume that the halo closest to each peak is the one that provides the greatest contribution to that peak's signal, and hence is the object considered to be recovered by the presence of the peak. In reality, any given peak arises from the contribution of many objects that may be related to multiple different haloes; however, this simplification allows us to relate the recovery to a single parameter – the halo mass. We vary the minimum mass of the haloes we consider, and plot the recovery fraction (as defined above) as a function of the halo mass in Fig. 9, and for this we extend the lower mass cut for haloes down to $10^{12} M_{\odot}$. This shows that the recovery fraction is >80 per cent for $M_{200} > 7 \times 10^{13} M_{\odot}$ clusters and drops as we extend to include the less rich lower mass haloes that do not produce as strong of a signal in the map.

5 SUMMARY

Watch list-based searches for gravitationally lensed transients provide a promising route towards the discovery of many candidates, including SNe, KNe, gravitational waves, and gamma-ray bursts from various wide-field survey data streams. The watch list-based approach allows transients detected near watch list objects to be quickly highlighted as candidates, although the lensing nature of these events would need to be confirmed by utilizing further follow-up observations and/or analysis of the data, depending on the transient in question.

A wide-field survey of particular interest is the upcoming Rubin Observatory LSST, which is expected to find approximately millions of SNe and approximately thousands of KNe during its operation – some of which will be lensed. However, in order to maximize

the opportunities for discovery of the rare lensed cases, and to enable science early in Rubin's operations, the watch list must be assembled in advance. Sufficient data already exist in the *J* (from VISTA/UKIRT) and *W*1 (from *WISE*) bands to produce such a watch list that covers the entire sky, and is deep enough to include the majority of the lenses associated with the high-redshift population of lensed transients. We have described a proof of concept for a method that utilizes these data to discover galaxy groups and clusters across full sky out to $z \sim 1$. In testing, our method successfully recovers between 80 and 100 per cent of a sample of large-Einstein radius ($\theta_E \geq 5$ arcsec) galaxy clusters, and 40–70 per cent of smaller Einstein radius objects (dependent on individual radii and the SNR_{eff} detection threshold used). We also estimated the false positive rate of the method by investigating the surrounding regions of test objects that the method identifies as likely to contain a group or cluster. By inspecting telescope images and by producing and inspecting both colour–magnitude diagrams and plots of the photometric redshift distribution of galaxies local to these candidates, we searched for evidence of real group or cluster objects at these locations. Using these, we were able to constrain the false positive rate for significant detections to be between approximately 6 and 17 per cent, depending on the applied SNR_{eff} detection threshold.

Our investigations clearly show that the applied value of the SNR_{eff} threshold changes both the level of completeness and the false positive rate of the method. Using a higher $SNR_{eff} = 5$ threshold results in secure detections of most large-Einstein radius clusters with a very low false positive rate, but conversely produces a less complete catalogue due to many groups and clusters with smaller Einstein radii being detected with lower SNR_{eff}. We concluded from our testing that a lower $SNR_{eff} = 3$ threshold is optimal, as the increased number of recovered objects outweighs the increase in false positives. We also anticipate being able to suppress the number of false positives in the future once data from Rubin become available.

We also tested the method on simulated data from the Rubin DP0 programme, which was selected to be comparable to currently available data sets. Testing on these data has the advantage of knowing the true underlying distribution of the dark matter haloes that drive the formation of lenses in our universe. This test provides similar results, indicating a good recovery rate (>80 per cent for haloes above $10^{14} M_{\odot}$) and with no false positives produced in our test region.

Future work will focus on producing a catalogue of cluster detections across the full sky using the method described here, and will re-evaluate both recovery and false positive rates by comparison of other large cluster catalogues such as those assembled by Rykoff et al. (2014), Wen, Han & Yang (2018), and Finoguenov et al. (2020), ensuring to collectively encompass multiple detection methods including optical/red sequence, galaxy overdensity, and X-ray detection.

DATA AVAILABILITY

The VISTA and *WISE* data underlying this article are available from the VISTA science archive (http://horus.roe.ac.uk/vsa/index.html) and the NASA/IPAC infrared science archive (https://irsa.ipac.caltech.edu/Missions/wise.html), respectively. The sources of data corresponding to the samples of galaxy clusters are available through the relevant publications cited: Smith et al. (2018b) and Carrasco et al. (2018). Rubin DP0 data are not publicly available and are only accessible to those with Rubin data rights and that are DP0 delegates. More information can be found on the Rubin Observatory community forum (https://community.lsst.org).

REFERENCES

Ahlgren B., Larsson J., 2020, ApJ, 897, 178

Amanullah R. et al., 2011, ApJ, 742, L7

Bellm E., 2014, in Wozniak P. R., Graham M. J., Mahabal A. A., Seaman R., eds, The Third Hot-Wiring the Transient Universe Workshop. Stanford University, p. 27

Bianconi M. et al., 2022, preprint (arXiv:2204.12978)

Birrer S., Treu T., 2021, A&A, 649, A61

Bruzual G., Charlot S., 2003, MNRAS, 344, 1000

Carrasco M., Stapelberg S., Maturi M., Bartelmann M., Seidel G., Erben T., 2018, preprint (arXiv:1807.03793)

Chabrier G., 2003, PASP, 115, 763

Dye S. et al., 2018, MNRAS, 473, 5113

Edge A., Sutherland W., Kuijken K., Driver S., McMahon R., Eales S., Emerson J. P., 2013, The Messenger, 154, 32

Eisenhardt P. R. M. et al., 2008, ApJ, 684, 905

Euclid Collaboration, 2019, A&A, 627, A23

Finoguenov A. et al., 2020, A&A, 638, A114

Goldstein D. A., Nugent P. E., 2017, ApJ, 834, L5

Goldstein D. A., Nugent P. E., Goobar A., 2019, ApJS, 243, 6

Gonzalez A. H. et al., 2019, ApJS, 240, 33

Goobar A. et al., 2009, A&A, 507, 71

Goobar A. et al., 2017, Science, 356, 291

Goobar A. et al., 2022, preprint (arXiv:2211.00656)

Grillo C., Rosati P., Suyu S. H., Caminha G. B., Mercurio A., Halkola A., 2020, ApJ, 898, 87

Jauzac M. et al., 2016, MNRAS, 457, 2029

Kelly P. L. et al., 2015, Science, 347, 1123

Kelly P. L. et al., 2016, ApJ, 819, L8

Korytov D. et al., 2019, ApJS, 245, 26

Lin Y.-T., Mohr J. J., Stanford S. A., 2004, ApJ, 610, 745

LSST Dark Energy Science Collaboration (LSST DESC), 2021, ApJS, 253, 31

LSST Science Collaboration, 2009, preprint (arXiv:0912.0201)

Mancone C. L., Gonzalez A. H., 2012, PASP, 124, 606

Marocco F. et al., 2021, ApJS, 253, 8

McMahon R. G., Banerji M., Gonzalez E., Koposov S. E., Bejar V. J., Lodieu N., Rebolo R., VHS Collaboration, 2013, The Messenger, 154, 35

More A., Suyu S. H., Oguri M., More S., Lee C.-H., 2017, ApJ, 835, L25

Oguri M., 2019, Rep. Prog. Phys., 82, 126901

Oguri M., Marshall P. J., 2010, MNRAS, 405, 2579

Pacaud F. et al., 2016, A&A, 592, A2

Patel B. et al., 2014, ApJ, 786, 9

Planck Collaboration XIII, 2016a, A&A, 594, A13

Planck Collaboration XXVII, 2016b, A&A, 594, A27

Refsdal S., 1964, MNRAS, 128, 307

Robertson A., Smith G. P., Massey R., Eke V., Jauzac M., Bianconi M., Ryczanowski D., 2020, MNRAS, 495, 3727

Rodney S. A. et al., 2015, ApJ, 811, 70

Rodney S. A., Brammer G. B., Pierel J. D. R., Richard J., Toft S., O'Connor K. F., Akhshik M., Whitaker K. E., 2021, Nat. Astron., 5, 1118

Rubin D. et al., 2018, ApJ, 866, 65

Ryczanowski D., Smith G. P., Bianconi M., Massey R., Robertson A., Jauzac M., 2020, MNRAS, 495, 1666

Ryczanowski D. et al., 2021, Transient Name Serv. AstroNote, 231, 1

Rykoff E. S. et al., 2014, ApJ, 785, 104

Schechter P., 1976, ApJ, 203, 297

Sharon K., Johnson T. L., 2015, ApJ, 800, L26

Smith G. P. et al., 2009, ApJ, 707, L163

Smith G. P. et al., 2018a, in Rodriguez Espinosa J. M., ed., Proc. IAU Symp. 338, Gravitational Wave Astrophysics: Early Results from Gravitational Wave Searches and Electromagnetic Counterparts. Kluwer, Dordrecht , p. 98

Smith G. P., Jauzac M., Veitch J., Farr W. M., Massey R., Richard J., 2018b, MNRAS, 475, 3823

Smith G. P., Robertson A., Bianconi M., Jauzac M., 2019a, preprint (arXiv:1902.05140)

Smith G. P. et al., 2019b, MNRAS, 485, 5180

Smith G. P. et al. 2022, preprint (arXiv:2204.12977)

Stapelberg S., Carrasco M., Maturi M., 2019, MNRAS, 482, 1824

Sunyaev R. A., Zeldovich Y. B., 1972, Comments Astrophys. Space Phys., 4, 173

Treu T., Marshall P. J., 2016, A&AR, 24, 11

Treu T. et al., 2016, ApJ, 817, 60

Vega-Ferrero J., Diego J. M., Miranda V., Bernstein G. M., 2018, ApJ, 853, L31

Wang Y., Brunner R. J., Dolence J. C., 2013, MNRAS, 432, 1961

Wen Z. L., Han J. L., Yang F., 2018, MNRAS, 475, 343

Wojtak R., Hjorth J., Gall C., 2019, MNRAS, 487, 3342

Wong K. C., Zabludoff A. I., Ammons S. M., Keeton C. R., Hogg D. W., Gonzalez A. H., 2013, ApJ, 769, 52

Wright E. L. et al., 2010, AJ, 140, 1868

This paper has been typeset from a TeX/LaTeX file prepared by the author.