Testing strong lensing subhalo detection with a cosmological simulation

Qiuhan He [ORCID],[1]★ James Nightingale [ORCID],[1,2] Andrew Robertson [ORCID],[3] Aristeidis Amvrosiadis,[1] Shaun Cole,[1] Carlos S. Frenk,[1] Richard Massey [ORCID],[1,2] Ran Li,[4,5] Nicola C. Amorisco,[1] R. Benton Metcalf,[6] Xiaoyue Cao[4,5] and Amy Etherington[1,2]

[1]*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*
[2]*Centre for Extragalactic Astronomy, Department of Physics, Durham University, South Road, Durham, DH1 3LE, UK*
[3]*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA*
[4]*National Astronomical Observatories, Chinese Academy of Sciences, 20A Datun Road, Chaoyang District, Beijing 100012, China*
[5]*School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing 100049, China*
[6]*Dipartimento di Fisica e Astronomia 'Augusto Righi' - Alma Mater Studiorum Università di Bologna, via Piero Gobetti 93/2, I-40129 Bologna, Italy*

ABSTRACT

Strong gravitational lensing offers a compelling test of the cold dark matter paradigm, as it allows for subhaloes with masses of $\sim 10^9$ M$_\odot$ and below to be detected. We test commonly used techniques for detecting subhaloes superposed in images of strongly lensed galaxies. For the lens we take a simulated galaxy in a $\sim 10^{13}$ M$_\odot$ halo grown in a high-resolution cosmological hydrodynamical simulation, which we view from two different directions. Though the resolution is high, we note the simulated galaxy still has an artificial core which adds additional complexity to the baryon dominated region. To remove particle noise, we represent the projected galaxy mass distribution by a series of Gaussian profiles which precisely capture the features of the projected galaxy. We first model the lens mass as a (broken) power-law density profile and then search for small haloes. Of the two projections, one has a regular elliptical shape, while the other has distinct deviations from an elliptical shape. For the former, the broken power-law model gives no false positives and correctly recovers the mass of the superposed small halo; however, for the latter we find false positives and the inferred halo mass is overestimated by ~ 4–5 times. We then use a more complex model in which the lens mass is decomposed into stellar and dark matter components. In this case, we show that we can capture the simulated galaxy's complex projected structures and correctly infer the input small halo.

Key words: *(cosmology:)* dark matter – gravitational lensing: strong.

1 INTRODUCTION

The cold dark matter (CDM) model predicts the existence of a vast population of dark matter haloes, from the scale of galaxy clusters down to Earth masses and below. Their mass function is characterized by a simple power law with an exponential cutoff at the very high mass end (Frenk & White 2012; Wang et al. 2020). For large masses these predictions have been verified by large sky surveys (Frenk et al. 1990; Rozo et al. 2010). At lower masses, where dark matter haloes are too small to host a luminous galaxy (Efstathiou 1992; Benson et al. 2002; Benitez-Llambay & Frenk 2020), it remains unclear whether the prediction still holds true. Alternative dark matter models predict a cutoff of the halo mass function. For example, warm dark matter (WDM) with a dark matter particle mass of around a few keV predicts a cutoff in the range $10^6 - 10^9$ M$_\odot$. Pushing constraints on the halo mass function towards this smaller mass range can distinguish different dark matter models.

Strong gravitational lensing serves as a promising tool to probe the existence of small invisible dark matter haloes. These 'dark' haloes perturb the images of lensed galaxies when they fall along the path of light from the source to the observer (Koopmans 2005; Vegetti & Koopmans 2009a, b). One can statistically study the lensing perturbation of ensembles of small dark haloes, and directly put constraints on the halo mass function (Gilman et al. 2019; He et al. 2022). Gilman et al. (2020a) constrained the 'half mode mass'[1] to be below $10^{7.8}$ M$_\odot$ by analysing flux ratio anomalies in eight strongly lensed quasar systems.

Individual subhaloes can be detected by analysing luminous strong lensing arcs (Vegetti & Koopmans 2009b). Li et al. (2016b) has shown that with ~ 50 high-quality strong lensing images, one can put stringent constraints on the cutoff mass and rule out CDM if no subhalo is detected. More recent work by Li et al. (2017) shows that the existence of line-of-sight haloes can boost the number of detection by a factor of several which improves the constraining power on the identity of dark matter (see also Despali & Vegetti 2017; Amorisco et al. 2022; He et al. 2022).

---

[1]This is a characteristic mass related to the k-mode where the dark matter power spectrum has an amplitude half the size expected with CDM. It can be considered as a 'cutoff' mass, below which the halo mass function is strongly suppressed with respect to CDM.

★ E-mail: qiuhan.he@durham.ac.uk

Three such detections of dark haloes with pseudo-Jaffe[2] masses below 10^{10} M$_\odot$ have been made so far. The first was made by Vegetti et al. (2010), who detected a $3.51 \pm 0.15 \times 10^9$ M$_\odot$ subhalo using *Hubble Space Telescope* (*HST*) imaging. The second was made by Vegetti et al. (2012) via Keck adaptive optics, with an inferred mass of $1.9 \pm 0.1 \times 10^8$ M$_\odot$. Finally, Hezaveh et al. (2016) found a $9.1 \pm 2.5 \times 10^8$ M$_\odot$ subhalo using Atacama Large Millimeter/submillimeter Array (ALMA) interferometer observations.

Attempts to detect dark haloes through gravitational imaging have been made in approximately 30 lenses (Vegetti & Vogelsberger 2014; Ritondale et al. 2019); however, no other clean detections have been made. Due to the small number of detections, constraints on the halo mass function are somewhat loose, constraining the mass function cutoff to be below $\sim 10^{10.9}$ M$_\odot$ (Enzi et al. 2021). This constraint is not yet competitive with other probes, owing to the limited data quality and sample size. Over the next decade, high-quality strong lensing observations from space telescopes such as the *James Webb Space Telescope*, *Euclid Space Telescope*, and *China Space Station Telescope* will allow these constraints to push to much lower cutoff masses on much larger lens samples (Collett 2015).

Substructure detection places stringent requirements on the model of the lens galaxy's mass distribution. The subhalo mass is usually less than 0.1 per cent that of the lens galaxy, necessitating per cent level accuracy of the main lens's mass. Inaccuracies in the lens mass model may create 'false-positive' detections, where the subhalo 'fills-in' for the mass model's missing complexity. Previous studies have discussed false-positive subhalo detections (Vegetti et al. 2010; Ritondale et al. 2019), where they apply strict criteria to ensure all subhalo detections are genuine. This includes requiring a high enough increase of Bayesian evidence that tests on mock data demonstrate that the signal cannot be due to an inaccurate mass model (Vegetti et al. 2012) and verifying that a consistent subhalo detection is made when pixelized corrections to the gravitational potential are applied (Koopmans 2005; Vegetti & Koopmans 2009a; Vegetti et al. 2010, 2012). In certain lenses these potential corrections clearly account for missing complexity in the lens galaxy's mass, thereby correctly flagging a candidate subhalo detection as a false positive. Mass model complexity is not the only contributor to false positives (Vegetti & Vogelsberger 2014).

This motivates the investigation of more complex lens models, which could improve the subhalo inference by accounting for this missing complexity in the lens galaxy's mass. The subhalo detections listed previously assume a simple parametric model for the lens's mass, the elliptical power law (Tessore & Metcalf 2015) with an external shear [Hezaveh et al. (2016) also included a fourth-order multipole term]. However, recent studies have highlighted deficiencies with this model. In a companion paper of this work, Cao et al. (2021) fitted this model to strong lenses simulated using mass models derived from dynamical models of nearby SDSS-IV MaNGA (Bundy et al. 2015) early-type galaxies and showed this can bias the measurement of the local density slopes around the Einstein ring by 13 per cent. Gomer & Williams (2021) and Van de Vyvere et al. (2022) discuss how departures from ellipticity symmetry may affect H_0 inference in lensed quasars. Nightingale et al. (2019) have also showed that departures from elliptical symmetry are observed in the luminous emission of three strong lenses.

In this work, we therefore use a hydrodynamic simulation to test the robustness of different parametric lens mass models, focusing on their efficacy for the task of detecting individual subhaloes. By using a simulation, we can compare the lens galaxy's true complex mass distribution to the lens model we fit and if it fails understand why. We perform two tests: (i) We do not add a subhalo to the lens galaxy when generating the mock data and investigate whether a lens model with a subhalo produces a false-positive signal and (ii) we include a subhalo when creating the mock data and test how accurately its mass and position are recovered. We first apply an extension of the commonly used power-law profile to fit the main lens (O'Riordan, Warren & Mortlock 2019, 2020, 2021), followed by a 'decomposed' model which models the lens mass as a combination of stars and dark matter (Dye & Warren 2005; Nightingale et al. 2019). Our goal is to understand whether modelling the lens mass as a power-law profile is sufficient for detecting subhaloes, and if not, whether there is a better model that can provide a correct inference.

Hydrodynamic simulations have previously been used to simulate galaxy-galaxy strong lensing images (Metcalf & Petkova 2014; Xu et al. 2017; Mukherjee et al. 2018, 2021; Despali et al. 2020; Enzi et al. 2020; Ding et al. 2021). Converting particle data into a corresponding deflection angle field (necessary for lensing) is non-trivial. A common problem is that the mass profiles of galaxies found in hydrodynamic simulations have a sub-kpc core in their centre. The strong lens imaging then produces a bright central image feature, which is not observed in real strong lenses (Bolton et al. 2012; Shu et al. 2016). These cores are believed to be due to the limited resolution of the simulations, with previous works assuming a particle resolution of $\sim 10^5$ M$_\odot$. Our simulation, which has a particle resolution ~ 10 times that of Illustris-1, still forms a core and central image. We incorporate this feature into our lens modelling such that we can still investigate dark matter subhalo detection. We also mitigate systematic effects related to particle noise in the simulation (Xu et al. 2009) and truncation effects which introduce an artificial shear (Van de Vyvere et al. 2020; Ding et al. 2021).

Testing with galaxies from the Illustris simulation (Vogelsberger et al. 2014), Xu et al. (2017) demonstrated that deviations of simulated galaxies from a simple elliptical power-law profile affect inference on the Hubble constant. More recently, Enzi et al. (2020) used 10 galaxies from the Illustris-1 simulation to test the power-law lens assumption for substructure lensing, and showed no degeneracy between the complexity of the true mass distribution of their mock lenses and the inferred substructure abundance. However, their work focused on the statistical properties of subhaloes' signals and did not test individual subhalo detection. To fully understand how the use of simple parametric lens models affects the detection of individual substructure, testing with mock lenses extracted from simulations is necessary.

This paper is structured as follows: In Section 2, we introduce our simulation data and the way we simulate strong lensing images from particle data. In Section 3, we introduce how we model the lensing images and search for subhaloes. In Section 4, we show the power-law fitting results. In Section 5, we introduce a more complex lens model where we model the lens' stellar and dark components separately and then we show how it behaves for our tests. In Section 6, we discuss our results. Finally, in Section 7, we summarize our results. All the computations, if not specified, are done by the state-of-art open-source strong lensing software PYAUTOLENS[3] (Nightingale,

---

[2]The masses reported are masses of pseudo-Jaffe profiles (Muñoz, Kochanek & Keeton 2001). If modelled by an NFW, the virial mass obtained is usually 0.5~1.0 dex higher than the pseudo-Jaffe mass (Despali & Vegetti 2017).

[3]The PyAutoLens software is open source and available from https://github.com/Jammy2211/PyAutoLens.
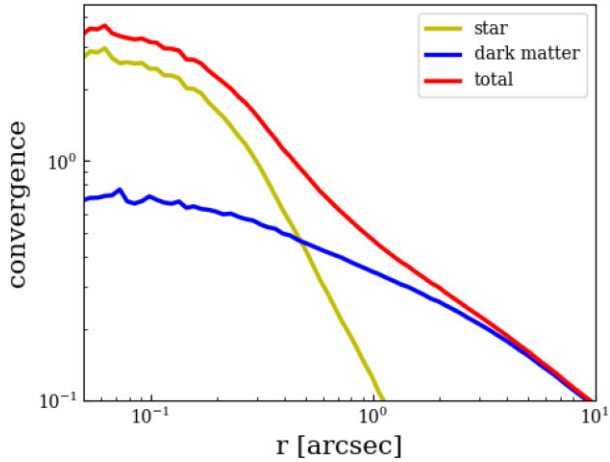
Figure 1. Convergence of different components of the simulated galaxy. Yellow, blue, and red represent the stars, dark matter, and total matter respectively, assuming the lens at $z = 0.2$ and source at $z = 2.5$. At the redshift of the lens $z = 0.2$, 1 arcsec corresponds to 3.3 kpc in angular size.

Dye & Massey 2018; Nightingale et al. 2021b). Throughout the paper, we adopt the Planck cosmology (Planck Collaboration XIII 2016), of which $H_0 = 67.7$ km s^{-1} Mpc^{-1}, $\Omega_m = 0.307$, and $\Omega_\Lambda = 0.693$.

2 MOCK LENSING IMAGES

2.1 Particle data

We create our lens galaxy by using data from a cosmological hydrodynamical zoom-in simulation of a $\sim 10^{13}$ M$_\odot$ galaxy group (Richings et al. 2021). The simulated galaxy is selected from the EAGLE 100 Mpc box-size simulation (Schaye et al. 2015) and was first identified by Despali & Vegetti (2017) as having similar properties to lenses from the Sloan Lens ACS (SLACS) survey (Bolton et al. 2006). The friends-of-friends ID of the halo is 129. To resolve dark matter haloes with masses down to $\sim 10^6$ M$_\odot$, this zoom-in simulation applies a novel technique whereby there are many more dark matter particles than gas particles. Unlike the common construction of initial conditions in hydrodynamic simulations, where each dark matter particle in a dark-matter-only simulation is split into a pair of dark matter and gas particles, the simulation we use initializes seven dark matter particles per gas particle, resulting in dark matter and gas particle mass of 8.3×10^4 and 10.7×10^4 M$_\odot$, respectively. At $z = 0$, the dark matter halo of the zoom-in simulation's galaxy group has a mass of $m_{200} = 10^{13.14}$ M$_\odot$ and size of $r_{200} = 506$ kpc.[4] The Plummer-equivalent gravitational softening length is 0.05 kpc.

In Fig. 1, we show the convergence profile (projected density divided by a constant lensing critical density) of different components of the simulated galaxy assuming the lens and source galaxy to be at $z = 0.2$ and 2.5, respectively. Inside the central ~ 0.7 arcsec, the baryonic mass is larger than the dark matter mass, and the central density of stellar mass is around 4 times higher than that of the dark matter. A constant-density core with a size of ~ 0.1 arcsec exists in the central region, which is a result of the finite resolution of the

[4]m_{200} is the mass enclosed within a radius of r_{200}, where r_{200} is determined as the radius at which the mean enclosed density is 200 times the critical density of the universe.

simulations. This phenomenon has been seen in several other studies that simulate strong lens images from simulation data (Mukherjee et al. 2018; Enzi et al. 2020; Ding et al. 2021), and it can produce a dim central image in simulated strong lensing images that is rarely seen in real observations (Winn, Rusin & Kochanek 2004; Quinn et al. 2016). The core feature introduces additional complexity beyond realistic massive ellipticals and thus might lead to an overestimation on the baryonic effects in our tests. Fortunately, for the lens configurations considered in this work, the central image is sufficiently dim and small that one can mask it out without it impacting the lens modelling and subhalo inference, an approach also followed by Enzi et al. (2020). To do this, we artificially increase the assumed error on the flux in the region containing the central image to such high values that they are effectively removed from the goodness-of-fit measurement.[5]

2.2 Simulating strong lensing images

2.2.1 Mock lenses

To simulate images that are strongly lensed by the particle distribution from a hydrodynamical simulation one needs a method which can determine the corresponding deflection angle map. There are two common ways of approaching this: (i) Derive the projected density distribution of the particle data and solve for its potential via a fast Fourier transform (FFT) or (ii) assume analytic profiles representing each particle enabling deflection angles to be easily computed, such that the overall deflection field is the sum over all particles. For the latter method, the computational cost can be greatly reduced by using a k-d tree algorithm (Bentley 1975), making it faster than the FFT method at comparable resolution (Metcalf & Petkova 2014; Petkova, Metcalf & Giocoli 2014). However, none of these methods offers a well-posed way of quantifying particle noise in the deflection angles, which can closely resemble the deflection angles of a dark matter subhalo in a strong lens (Xu et al. 2009). Besides the particle noise, these above methods also face the boundary truncation effect, which is that when truncating the particle data in an improper way (e.g. a square boundary applied to an elliptically shaped galaxy), an artificial shear component is introduced (Van de Vyvere et al. 2020). The shear magnitude depends on the galaxy's profile, the truncation area size, and the truncation scheme used, and an improper truncation on the particle mass data can induce several per cent bias to H_0 inference (Van de Vyvere et al. 2020; Ding et al. 2021).

To avoid particle noise and the boundary truncation effect, we therefore instead fit analytic profiles to the simulated galaxy's particle data and use these profiles to compute our lens galaxy's deflection angle field. We approximate the projected mass distribution of the simulated galaxy using the multiple Gaussian expansion (MGE) method, which is widely used for modelling galaxy surface brightness profiles in studies of galaxy stellar dynamics (e.g. Cappellari 2008; Li et al. 2016a, 2019; He et al. 2020). Li et al. (2016a) applied the expansion method to both galaxies and dark matter haloes in Illustris simulation, showing that it has flexibility to capture irregular and asymmetric features in a galaxy's light or mass distribution. The deflection angles (and other lensing quantities) of an elliptical

[5]We previously attempted to remove the central image by applying a mask that removed the data in the central region altogether. However, we found that this introduced systematics due to edge effects associated with the source-plane pixelization. Pixels at the edge of the mask (which have non-negligible flux due to the central image) were not appropriately regularized because their neighbours were not traced to the source plane (see Nightingale et al. 2018).
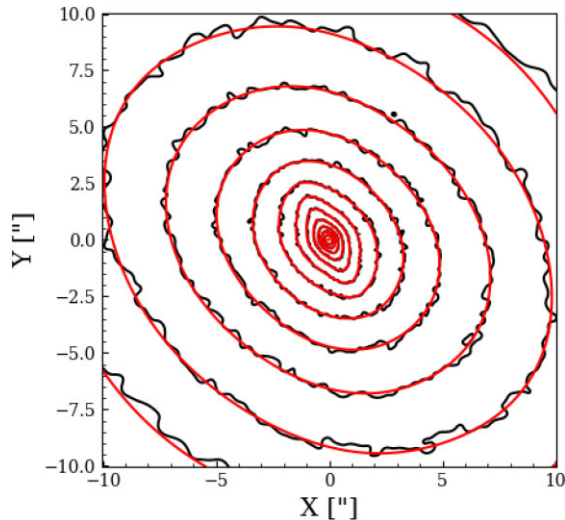
In the top-left panel of Fig. 3, we show the convergence (i.e. the projected density divided by the critical surface density for lensing) of the MGE representation of the simulated galaxy, where a pointy 'American football-like' shape can be seen. To investigate how lens model fits change depending on the shape of the convergence, we rotate the same galaxy to view it along a different line of sight, intentionally choosing a viewing angle that produces a rounder convergence map, which is shown on the bottom-left panel of Fig. 3. Following equation (43) of Shajib (2019), we compute MGE deflection angles for both projections, which are then used to simulate strong lensing images. For an accurate computation of each image pixel's flux, we treat every pixel with a 4×4 subgrid so that for each image pixel 16 light rays are traced to the source plane, with the pixel flux set to their mean value.

2.2.2 Mock sources

We simulate source galaxies using the cored Sérsic profile (Trujillo et al. 2004):

$$I(r) = I' \exp\left[-b_n \left(\frac{r^2 + r_c^2}{r_e^2}\right)^{1/(2n)}\right], \tag{1}$$

where I' is the scale density, r_c is the core size, r_e is the effective radius, n is the Sérsic index, and b_n is a dimensionless parameter fully determined by n (Graham & Driver 2005). Our input sources are a single elliptical cored Sérsic profile (the ellipticity is introduced using $r = \sqrt{(x/q)^2 + y^2}$), which is simple compared to observed lenses where source galaxies are more complex and show features such as multiple star-forming clumps, spiral structures, and extremely compact centres. Our choice to assume a simple source profile is to make it straight forward to test the effects of using different lens mass models. Our lens modelling procedure uses pixelized source reconstructions which are able to fit the more complex sources seen in real data (Nightingale et al. 2019).

2.2.3 Subhaloes

Some of our mock lensed images include a dark matter subhalo in the lens galaxy near one of the arcs. We represent subhaloes using the spherical Navarro–Frenk–White (NFW) profile (Navarro, Frenk & White 1996):

$$\rho(r) = \frac{M_0}{4\pi r (r_s + r)^2}, \tag{2}$$

where M_0 is the scale mass and r_s is the scale radius. Following equation (A.18) of Baltz, Marshall & Oguri (2009), we analytically compute its deflection angles and add it to that inferred via the MGE fit to the stellar and dark components. We assume the NFW halo follows the mass–concentration relation given by Ludlow et al. (2016), allowing us to parametrize it with only its mass, m_{200} (and its position). Mock lenses are generated with an input subhalo of two different masses, $m_{200} = 5 \times 10^8 \, \mathrm{M_\odot}$ or $m_{200} = 5 \times 10^9 \, \mathrm{M_\odot}$.

2.2.4 Data quality

We simulate mock images similar to observations of the *HST*/Advanced Camera for Surveys Wide Field Camera, with a pixel size of 0.05 arcsec and a Gaussian point spread function (PSF) with
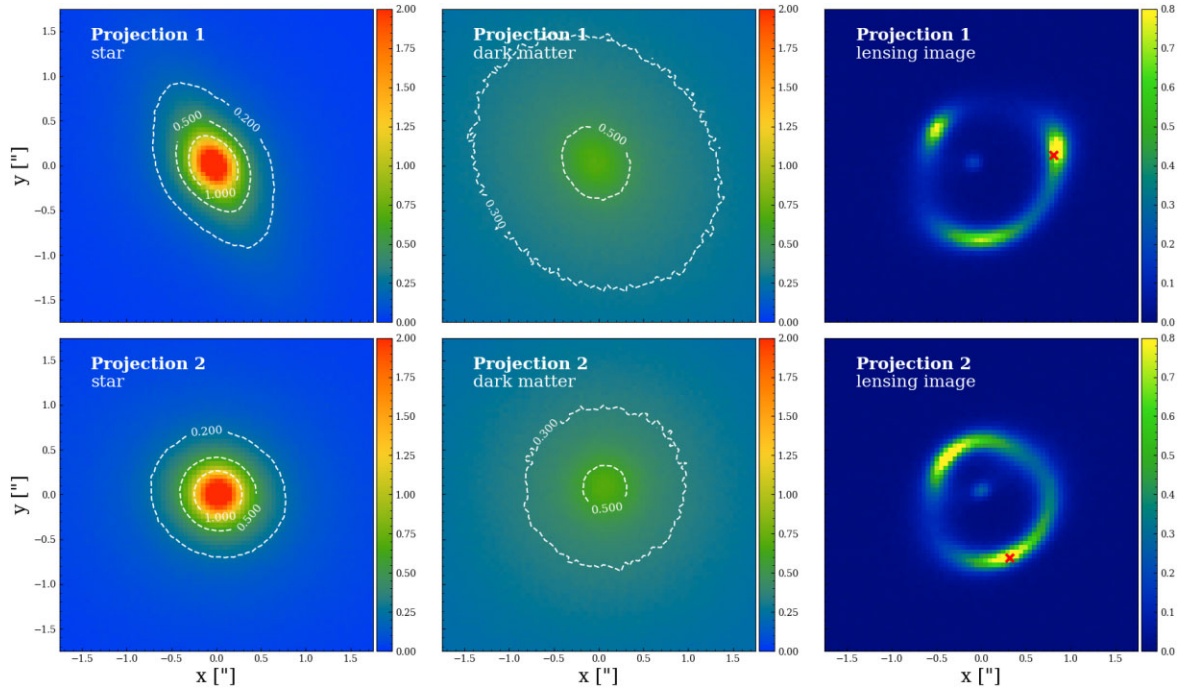

Figure 2. Isodensity contours of the projected stellar mass distribution of our simulated galaxy. The contours measured directly from the particle data using GLAMER (Metcalf & Petkova 2014; Petkova et al. 2014) are shown in black, while those for the best-fitting MGE are in red. The contours are evenly log-spaced in projected density. From inside to outside, each contour decreases by 0.4 order. Note that only very inner part (within 5 arcsec) of this image is observable and the reason we plot it on a much larger region is to show that the MGEs represent the lens' stellar mass well to a very large range. For clarity, later plots are all zoomed in to the region around the Einstein radius.

Gaussian profile can be easily computed (Shajib 2019), making lens simulations convenient and fast.

We compute the deflection angles of the simulated galaxy separately for its stellar and dark matter components, and then add them together to get the total deflection angles. The gas component is omitted because its contribution to the total mass in the galaxy's central region is negligible. We will add subhaloes to the deflection angle map via an analytic mass profile and therefore must ensure no subhaloes in the particle data are included in our simulation process. We therefore only use particles belonging to the main halo identified by the SUBFIND algorithm (Springel et al. 2001).

We set the lens galaxy to be at redshift $z = 0.2$, and the source galaxy at $z = 2.5$. We first use GLAMER (Metcalf & Petkova 2014; Petkova et al. 2014) to generate convergence maps, where each particle is represented by a smoothed b-spline in three-dimension (3D). For each star particle the smoothing length is the distance to its eighth nearest (stellar) neighbour and for each dark matter particles it is the distance to its 64th nearest (dark matter) neighbour. We then use the MGE code of Cappellari (2002) to decompose the convergence maps into multiple Gaussian profiles, where the Gaussian components share the same centre but are free to have different amplitudes, sizes, position angles, and axis ratios. As an example, Fig. 2 shows contours tracing the particle data input (dark lines) and best-fitting MGE (red lines) of the simulation's projected stellar mass distribution. The MGE-fitting code decomposes the stellar component into 13 individual Gaussian profiles and Fig. 2 shows asymmetric features such as the twist in ellipticity are well captured by the MGEs. The relative errors between the input profile and best-fitting MGEs are smaller than ~ 5 per cent. We apply the same routine to the simulation's dark matter particles and then add the best-fitting Gaussian profiles together to represent the simulated galaxy's total projected mass distribution.

Figure 3. Input stellar (left-hand column) and dark matter (middle column) convergence maps of our simulated galaxy along two different lines of sight (the two rows). The corresponding strong lensing images are shown in the right-hand column. The top row is for the line of sight that produces pointy shaped iso-convergence contours, and the bottom row shows the projection with a rounder convergence field. The red crosses on the mock images mark the positions where we will later place subhaloes.

a standard deviation of 0.05 arcsec.[6] We set the normalization of the source's surface brightness to give a signal-to-noise ratio (S/N) of ∼80 in the brightest pixel of the lensed source's image, whilst adjusting the background noise level to closely match that expected from a few *HST* orbits. For our mocks, the background sky noise is 0.1 e$^-$ pixel^{-1} s^{-1}. This S/N represents observations that are around double the highest S/N sources observed currently with Hubble, e.g. the SLACS sample. Using such high S/N data is a choice we made to ensure our tests of deficiencies in the lens mass model are easier to distinguish from noise in the mock data. The right-hand column of Fig. 3 shows the two mock images, where a source galaxy is lensed by the two different line-of-sight projections (the corresponding projected densities are shown on the left hand). For mock data sets which include a subhalo, the positions marked by red crosses show the locations of the subhalo that we add. In Table 1, we summarize the relevant parameters used to simulate these images.

3 METHOD

3.1 Mass models

3.1.1 Broken power law

The simulated lens galaxy has an artificial ∼0.1 kpc constant density core, which forms a spurious central, demagnified image. We mask this central image by manually decreasing the contribution of central

---

[6]The actual PSF σ for *HST*/Advanced Camera for Surveys is ∼0.034 arcsec and the pixel size is 0.04 arcsec, so our tests are slightly worse than real *HST* observations in terms of resolution. However, in terms of the PSF modelling, we assume we have perfect knowledge of the PSF, which goes in the other direction of being optimistic.

Table 1. Parameters used to simulate the mock lensing images.

	Projection 1	Projection 2
Input lenses		
Stellar MGE number	13	11
Dark MGE number	5	5
Redshift	0.2	
Input sources	**Cored Sérsic**	
Centre (x, y) [(arcsec, arcsec)]	(0.08, −0.03)	
Axial ratio	0.55	
Position angle (°)	30	
$I^{'}$ (e$^-$ pix^{-1} s^{-1})	2.0	
r_e (arcsec)	0.11	
n	2.0	
r_c (arcsec)	0.01	
Redshift	2.5	
Input subhaloes (if added)	**Spherical NFW profile**	
Centre (x, y) [(arcsec, arcsec)]	(0.81, 0.12)	(0.32, −0.71)
m_{200} (M$_\odot$)	5×10^8 or 5×10^9	
Mass-concentration relation	Ludlow et al. (2016)	
Redshift	0.2	
Image settings		
Pixel size (arcsec)	0.05	
PSF σ (arcsec)	0.05	
Background noise level (e$^-$ pix^{-1} s^{-1})	0.1	
Exposure time (s)	8000	
Maximum pixel S/N	∼80	

image pixels to the likelihood calculation; however, we must also ensure that our mass model parametrization is able to represent the cored density, to avoid biasing our reconstruction of the lensed source's arcs (Enzi et al. 2020). We therefore assume the elliptical broken power-law (eBPL) profile (O'Riordan et al. 2019, 2020, 2021) with convergence

$$\kappa(r) = \begin{cases} \kappa_b \, (r_b/r)^{t_1} \,, \ r \le r_b \\ \kappa_b \, (r_b/r)^{t_2} \,, \ r > r_b \end{cases}, \quad (3)$$

where r_b is the break radius, κ_b is the convergence at the break radius, t_1 is the inner slope, and t_2 is the outside slope. When $r_b = 0$, the eBPL reduces to the standard power-law profile with 3D density $\rho(r) \propto r^{-\gamma}$, as used in many lens studies. We introduce ellipticity by setting $r = \sqrt{qx^2 + y^2/q}$, where q is the axial ratio. In practice, we parametrize a model's axial ratio and position angle, θ, in terms of two components of ellipticity:

$$e_1 = \frac{1-q}{1+q} \sin 2\theta, \quad e_2 = \frac{1-q}{1+q} \cos 2\theta. \quad (4)$$

With the additional two parameters describing the profile's centre, the eBPL model has eight free parameters. Degeneracies between certain parameters in the eBPL profile, e.g. the two different slopes, make it challenging to fit efficiently and avoid inferring local maxima. We therefore assume priors that lessen these degeneracies and simplify parameter space, where we constrain $r_b \le 0.4$ arcsec, $t_1 \le 0.5$, and $t_2 > 0.5$. For some cases, we further limit the Einstein radius to be larger than 0.5 arcsec. All eBPL models are fitted with an additional external shear in the lens model, which provides further flexibility in stretching and squeezing of the mass profile that can capture asymmetric features in the lens's convergence (Cao et al. 2021). Similar to the ellipticity parametrization, the external shear is also parametrized with two components γ_{1ext} and γ_{2ext}, where the shear's magnitude, γ_{ext}, and position angle, θ_{ext}, can be recovered as

$$\gamma_{ext} = \sqrt{\gamma_{1ext}^2 + \gamma_{2ext}^2}, \ \tan 2\theta_{ext} = \frac{\gamma_{2ext}}{\gamma_{1ext}}. \quad (5)$$

For modelling of a subhalo, we take the same NFW form we use to simulate the image.

3.1.2 Source reconstruction

The final lens model of our analysis – from which all results in the main content of this paper are taken – reconstructs the source galaxy using a pixelization that adapts to the source's surface brightness distribution (see Nightingale et al. 2018, for a discussion of systematics this approach removes compared to other pixelizations). However, before using this pixelized source, a number of initial fits is performed which estimates the parameters of the lens mass model efficiently (for details see the next subsection). These fits assume either a parametric source which is modelled using the Sérsic profile ($r_c = 0$ in equation 1) or a pixelized source where the density of pixels adapts to the magnification, leading to smaller pixels in more magnified areas of the source plane. In Appendix A, we also show that we reproduce our main conclusions assuming a parametric cored Sérsic source model.

3.2 Fitting procedure

We use PyAutoLens (Nightingale et al. 2021b) to model the simulated lens data sets, which is described in Nightingale et al. (2018, hereafter N18) and builds on the works of Warren & Dye (2003), Suyu et al. (2006), and Nightingale & Dye (2015).

PyAutoLens uses a technique called 'non-linear search chaining' to compose pipelines which break the lens modelling procedure into a series of simpler model fits. This allows us to begin modelling our data with a simple lens model (e.g. an isothermal mass profile and a Sérsic source) and via a sequence of non-linear searches gradually increase the model complexity, so as to eventually fit the desired more complex lens model (in this work, mass models which include a dark matter subhalo and with a source reconstructed on the brightness-based pixelization). Non-linear search chaining is implemented in PyAutoLens via the probabilistic programming language PyAutoFit[7] (Nightingale, Hayes & Griffiths 2021a). We use the nested sampling algorithm dynesty to perform all model fits.

We employ the Source, Light and Mass (SLaM) pipelines that are distributed with PyAutoLens.[8] The SLaM pipelines were used in the work of Cao et al. (2021) and Etherington et al. (2022) and our analysis closely follows theirs, albeit we end with an additional pipeline that determines whether including a subhalo in the lens model increases the Bayesian evidence relative to the model without a subhalo. Like in Cao et al. (2021), we do not need to model the lens light; and therefore, we employ a model fitting procedure consisting of four distinct pipelines each one of which focuses on fitting a specific aspect of the model. These pipelines are, in order: (i) the parametric source pipeline; (ii) the pixelized source pipeline; (iii) the lens mass pipeline; and (iv) the subhalo pipeline. Each pipeline consists of one or more non-linear searches that fit a unique lens model parametrization, which Fig. 4 shows a flow chart of, which we will now explain in detail.

3.2.1 Parametric source pipeline

The parametric source pipeline aims to initialize a robust model for the lens galaxy's mass by fitting a source galaxy that has a smooth analytic form. The primary reason for this pipeline is that a robust model for the lens galaxy's mass is necessary to avoid pixelized source reconstructions inferring the unphysical solutions described by Maresca, Dye & Li (2021), where the reconstruction inferred is a demagnified version of the lens data. This pipeline assumes a singular isothermal ellipsoid mass model [where in equation (3), r_b is set to be 0 and t_2 is fixed to be 1.0] with an external shear and a Sérsic profile for the source surface brightness.

3.2.2 Pixelized source pipeline

The pixelized source pipeline is composed of four search phases. The first search fits for parameters describing the resolution of the magnification based pixelization and the regularization coefficient of the constant regularization scheme, with the lens mass model fixed to the result of the parametric source pipeline. The second search re-fits the lens mass model using the pixelization and regularization inferred previously. The third search fits for parameters that derive the surface brightness based pixelization and the luminosity weighted regularization scheme, where the lens mass model is fixed to the best-fitting values inferred in the previous search. The fourth search again re-optimizes the lens mass model now using the brightness based pixelization and regularization and we finally re-fit the pixelization and regularization parameters again one last time, ensuring that the
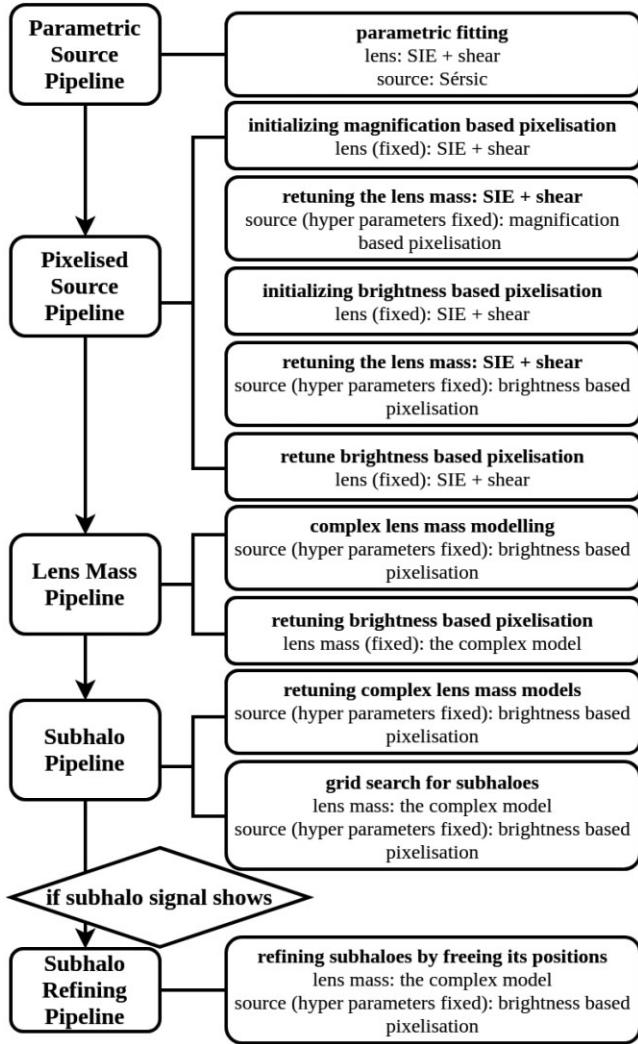

Figure 4. The fitting procedure we use to search for subhaloes.

source reconstruction is tailored to the properties of the source it is fitting.

3.2.3 Mass pipeline

This pipeline fits a more complex lens mass, either the eBPL model plus an external shear or the decomposed model that separately models the stellar and dark components plus an external shear. This pipeline consists of two searches. It first fits the new lens mass model with fixed source pixelization parameters. The priors of the (broken) power-law model's centres, elliptical components, and Einstein radius are updated using information of the previous best-fitting models. We set those priors to be Gaussian priors centring on corresponding best-fitting values of previous models and their widths are set manually using values which balance reducing the size of parameter space to ensure an efficient fit whilst being broad enough not to remove physically plausible solutions. For other parameters (like the break radius, inner (outer) slopes, and external shear) we assume broad uniform priors that are not informed by the previous mass model fits. Having now fitted this more complex mass model, we again update the source pixelization and regularization parameters using the best-fitting lens mass model of the first step. This is the final

fit which updates the pixelization and regularization parameters, with all remaining fits focusing on the lens (and subhalo) mass models.

3.2.4 Subhalo pipeline

This pipeline performs Bayesian model comparison to determine if a lens model with a subhalo is preferred over a lens model without a subhalo. The pipeline begins by fitting the same lens mass model (with fixed source pixelization and regularization parameters) inferred at the end of the mass pipeline, with all priors inherited from this fit. This provides us with an estimate of the Bayesian evidence of the lens model without a subhalo. We then fit lens models which include an NFW subhalo. For the subhalo's mass, we assume a uniform prior on $\log_{10}(m_{200}/M_\odot)$ between 6 and 11.

Due to the complexity of our parameter space (which consists of the mass models of both the main lens galaxy and a subhalo) we found it was common for the inferred posterior to correspond to a local likelihood maximum (as opposed to the global maximum). To mitigate this, we scan for subhaloes using a grid of non-linear searches, where each search confines the (x, y) image-plane coordinates of the subhalo to a small 2D square segment of the image-plane. We perform 25 independent model fits, corresponding to a 5×5 grid, which divides the image region between -1.0 and 1.0 arcsec into subregions with sizes of 0.4 arcsec $\times$ 0.4 arcsec. The parameters of the main lens are fit for simultaneously along with the subhalo parameters in each of these 25 fits.

To determine whether the lens model with a subhalo is favoured by the data over the model without a subhalo, we must choose a statistical quantity with which to compare them. Obvious choices are the Bayesian evidence or differences in maximum log likelihood values. We use the maximum log likelihood to compare models which do and do not include a subhalo. However, the Bayesian evidence is as an output of `dynesty` and we have verified that our results are unchanged using this quantity. We denote the difference between the two maximum log likelihoods as ΔL, such that if ΔL in certain cells of the subhalo search is large, it suggests the existence of a subhalo within one of those certain grids. Instead, if all log likelihood differences are very small, then it indicates no subhaloes of a sufficiently high mass to be detected are present in the image. For this paper, we take the threshold as follows: If $\Delta L \leq 5$, we call it a non-detection; if $5 < \Delta L \leq 10$, we call it a plausible detection; and if $\Delta L > 10$, we call it a detection.

If the subhalo grid search has a plausible detection ($\Delta L > 5$) the subhalo pipeline performs one more fit, which fits for both the main lens and subhalo parameters. The subhalo's (x, y) position is no longer confined to a square segment of the grid search and we instead place a Gaussian prior on the x and y positions. The 2D Gaussian prior is centred at the maximum-likelihood subhalo position inferred previously using the grid search, with a relatively large standard deviation of 0.5 arcsec. For the subhalo's mass, we retain a prior uniform in $\log_{10}M$ between 10^6 and 10^{11} $M_\odot$.

4 POWER-LAW TESTS

We first use our simulated lenses to test the broken power-law profile, which is commonly assumed in strong lensing studies to model the mass distribution of the lens galaxy (Collett & Auger 2014; Vegetti & Vogelsberger 2014; Dye et al. 2015; Ene et al. 2018). Our tests are divided into two parts: (i) How do power-law fits behave for the case where no subhalo is present in the mock data and (ii) can the power-law correctly recover the subhalo's properties when there is
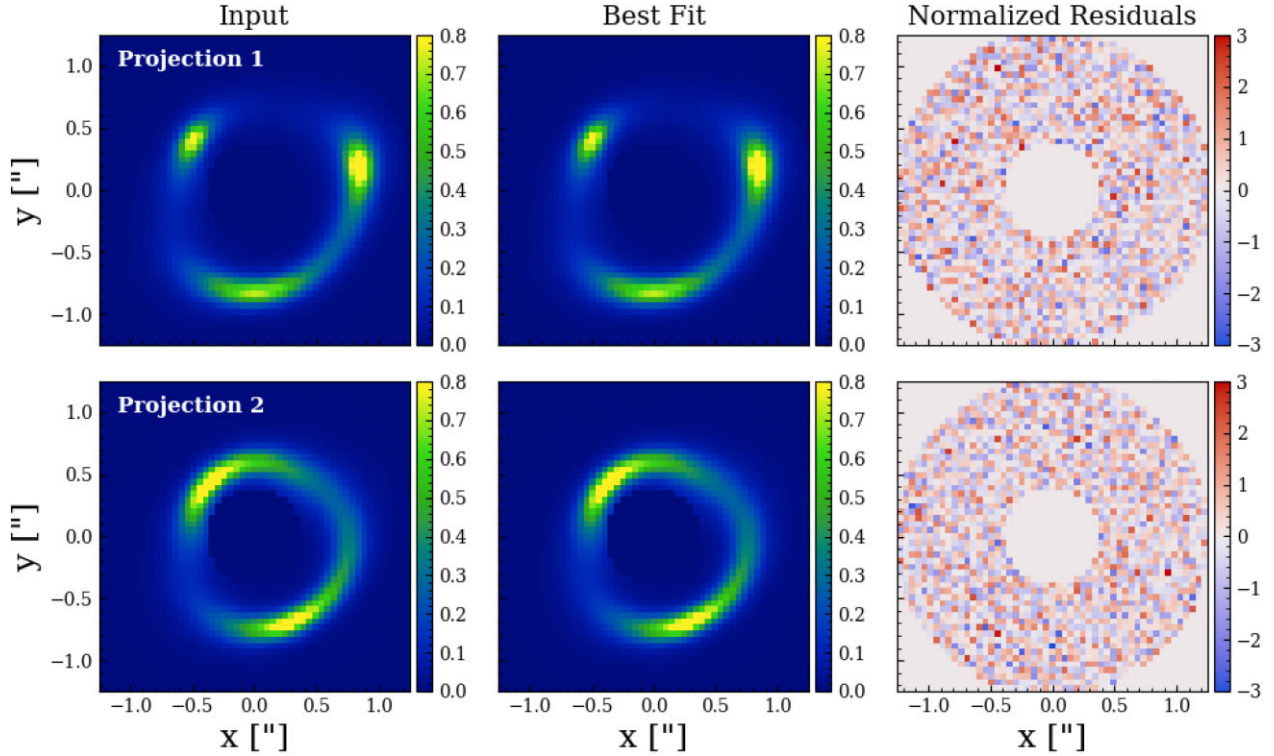
Figure 5. Left-hand column: The mock lensing images. Middle column: The best recovered images using an eBPL to model the lens mass. Right-hand column: The corresponding normalized residuals (residuals divided by the noises). The top row shows the case of Projection 1 and the bottom row shows the case of Projection 2. The colour bar unit for the left two column images is e^- pix^{-1} s^{-1}. The units of the y and x axes are arcsec.

one present in the mock data. For convenience, we call the tests where no subhalo is added 'smooth tests', and tests where there is a subhalo added 'subhalo tests'.

4.1 Smooth test results

In Fig. 5, we show the input and recovered images from the best-fitting smooth model when fitted to simulated images that do not include a subhalo (this corresponds to the first model fit in the subhalo pipeline and will act as the model we compare to models including a subhalo in a moment). For visual clarity, we have removed the central image caused by the core of the simulated galaxy; however, note that this region is included in the model-fit with high error values. For both projections, the reconstructed images in the middle panel are similar to the input images shown on the left hand. The normalized residuals (residuals divided by the noises) shown in the right-hand panel confirm the good fit, showing no clear or obvious correlated residuals. It is noted that the best-fitting eBPL model's break radius for Projection 1 and 2 are ∼0.2 and 0.1 arcsec, respectively, which confirms that the core is able to affect the lensing even though the central image has been masked out. Using an eBPL model is therefore necessary to account for the core.

We now consider the results of the subhalo search. The left-hand column of Fig. 6 shows the results of the subhalo phase, using the quantity ΔL (defined in Section 3.2.4) inferred in every cell of the subhalo-position grid. The upper and lower panels show the results of Projection 1 and 2, respectively. For Projection 1, where the input galaxy has a pointy shaped convergence, grids around the top-left luminous arc have ΔL over 10, and the highest ΔL is ∼21.4 for the left-most grid cell of the third row from bottom. For the grid cell with

$\Delta L \sim 21.4$, a subhalo with m_{200} of $10^{9.8^{+0.4}_{-0.5}}$ M$_\odot$ is inferred around that region. Given that the simulated lens galaxy we fitted here does not contain a subhalo, this signal is a false-positive. However, for Projection 2 which has a rounder convergence, no grid has a $\Delta L > 5$. Assuming our criteria of requiring $\Delta L > 5$ the inclusion of an additional subhalo model using the eBPL is therefore correctly not favoured by the data and the eBPL gives the correct answer for this projection. However, it should be noted ΔL values of ∼3–4 are still visible, indicating that at a very low level the subhalo is still improving the fit to the data.

4.2 Subhalo test results

Having shown the performance of using an eBPL to fit images without a subhalo present, we now test whether the same pipeline can correctly recover a subhalo's properties when a subhalo is included when generating the mock data. For both projections, we add an NFW-like subhalo of $m_{200} = 5 \times 10^8$ M$_\odot$ or 5×10^9 M$_\odot$ at the positions marked by the red crosses in Fig. 3.

Similar to our earlier analysis, we first check the ΔL maps. The middle column of Fig. 6 shows ΔL maps for the cases where a 5×10^8 M$_\odot$ subhalo is added and the right-hand panels show the results for a 5×10^9 M$_\odot$ subhalo. The upper and lower panels show the results of Projection 1 and 2, respectively. For Projection 1, grid cells near the subhalo's true input location (marked as white triangles in the image) show clear increases in ΔL. For the 5×10^8 M$_\odot$ subhalo case, the maximum ΔL is ∼24.0 (for the rightmost cell on the fourth row from bottom) and for the 5×10^9 M$_\odot$ case, the maximum ΔL is ∼81.4 (also for the rightmost cell on the fourth row from bottom).
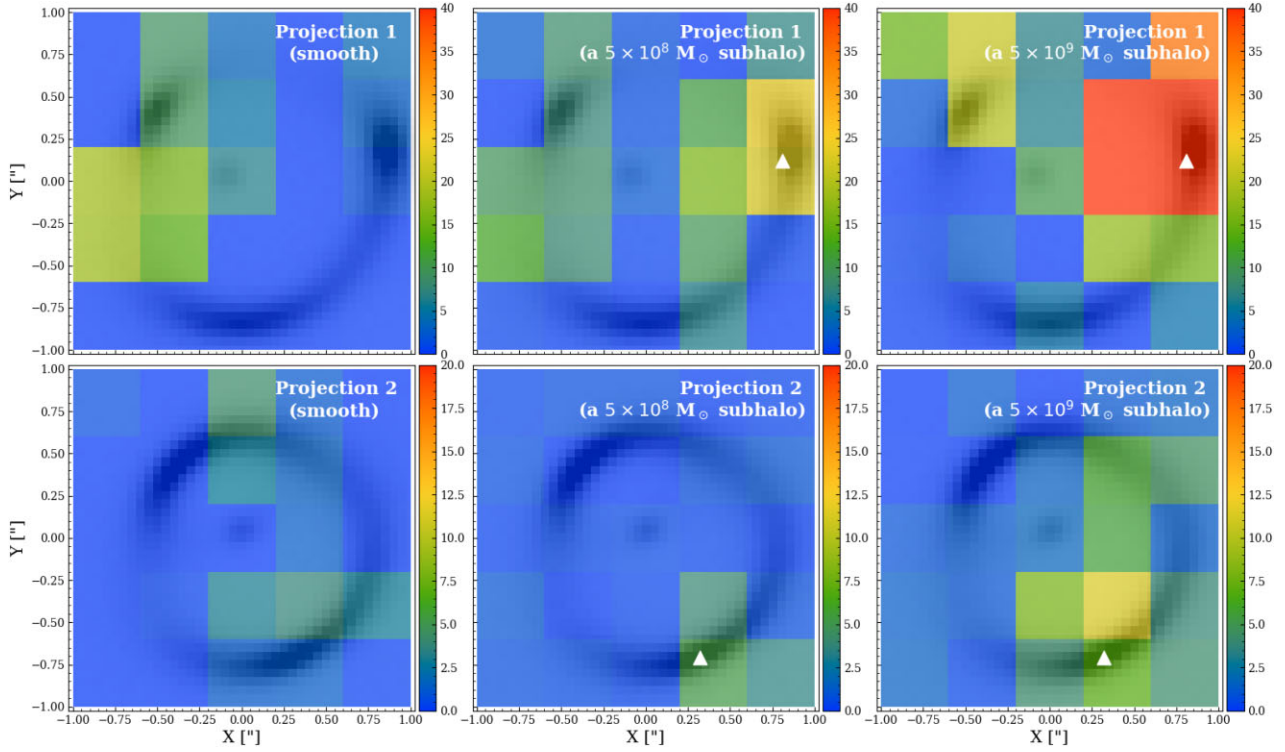
Figure 6. When modelling the lens mass distribution with an eBPL, substructures are (too) easily detected. Colours indicate the increase in maximum log likelihood, ΔL, when a subhalo is included inside 0.4 arcsec × 0.4 arcsec squares during a fit to a lens that has: no subhaloes (left-hand panel), a subhalo of mass 5×10^8 M$_\odot$ (middle panel), or a subhalo of mass 5×10^9 M$_\odot$ (right-hand panel). Top and bottom rows show the results for Projection 1 and 2 (with different colour scales). White triangles mark the true locations of the subhaloes. Note the false-positive detections in the left-hand panels; the best-fitting subhalo masses in the other panels are also overestimated by a factor 4–5.

Note that the colour bar saturates (for cells with $\Delta L > 40$) in the top-right panel.

Based on the detections shown in the ΔL maps, we continue the subhalo pipeline and fit a model where we no longer confine the subhalo within a particular square cell, and instead use a 2D Gaussian prior on the subhalo position, centred on the best-fitting position from the 'grid-search' phase, with a standard deviation of 0.5 arcsec. In Fig. 7, we show the posterior of the subhalo parameters for both the case with a 5×10^8 M$_\odot$ subhalo (red) and a 5×10^9 M$_\odot$ subhalo (blue). The true input values are marked by the dashed lines. As shown, for both cases, the subhalo's mass is significantly overestimated and the true input subhalo masses are excluded by 99 per cent confidence regions. When the input subhalo has a mass of 5×10^8 M$_\odot$, the recovered subhalo mass is overestimated by $\sim$5 times with a value of $10^{9.4^{+0.4}_{-0.4}}$ M$_\odot$ inferred, whereas for the input subhalo with 5×10^9 M$_\odot$, the recovered mass is overestimated by around 4 times and has a value of $10^{10.3^{+0.2}_{-0.2}}$ M$_\odot$.

For Projection 2, with an input subhalo of 5×10^8 M$_\odot$, the maximum ΔL is only 3.4, therefore no subhalo is detected and we do not analyse the posterior on the subhalo properties. When the input subhalo mass is 5×10^9 M$_\odot$, the maximum ΔL is 9.7 providing us with a plausible detection. We take this plausible detection and refine the fit, with the resulting posterior for the subhalo parameters plotted in Fig. 8. For Projection 2, the subhalo's mass is recovered to be $10^{10.1^{+0.5}_{-1.2}}$ M$_\odot$ where the errors mark the 99 per cent confidence regions. It is noted that in Fig. 8, although the subhalo's true parameters are recovered within 99 per cent confidence regions (the light blue regions), the best-fitting m_{200} and y coordinate are
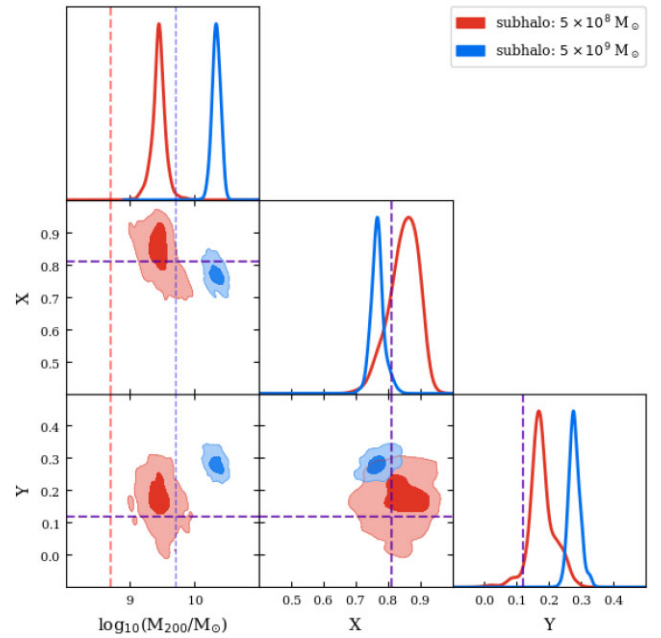

Figure 7. Posterior probability distribution of the subhalo model parameters, after the subhalo refining phase, when modelling the Projection 1 lens mass with an eBPL. Red and blue correspond to the cases of a 5×10^8 and 5×10^9 M$_\odot$ input subhalo, respectively. The 2D contours cover the 68 and 99 per cent credible regions. For 1D posteriors, the vertical dashed lines mark the true input values.
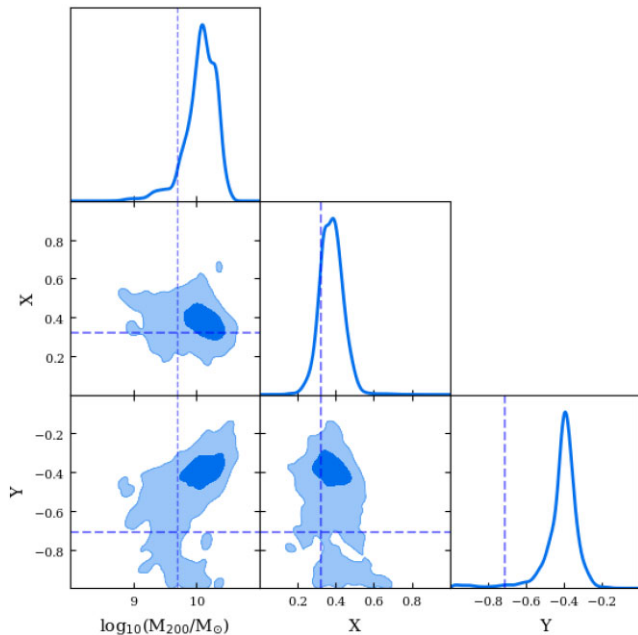
Figure 8. Subhalo parameters' posteriors of the subhalo refining phase, when modelling the Projection 2 lens mass with an eBPL. Only the case of a 5×10^9 $M_\odot$ input subhalo is shown (the case of a 5×10^8 $M_\odot$ subhalo does not lead to a clear detection). The 2D contours cover the 68 and 99 per cent confidence regions. For 1D posteriors, the vertical dashed lines mark the true input values.

clearly offset from the true input and in a case of smaller errors (higher S/N images), the true inputs might be ruled out.

4.3 Parametric source

To verify that our conclusions are not a result of a systematic associated with our pixelized source model, in Appendix A we re-perform all of the above fits assuming a cored Sérsic profile for the source. For Projection 1, we see nearly identical behaviour in terms of false positives and the subhalo inference; however, the ΔL values are much larger (of order $\sim$250 compared to the values of $\sim$25 seen for the pixelized source). This is expected, as the greater flexibility of the pixelized source reduces our sensitivity to a subhalo and therefore also false positives (Gilman, Birrer & Treu 2020b). For Projection 2, fits to the smooth data now infer a false positive with $\Delta L = 18.1$. This does not contradict the results using a pixelized source above, instead the values of $\Delta L = \sim$3–4 shown in Fig. 6 have simply been boosted above our threshold value of $\Delta L = 5$ because fitting a cored Sérsic increases our sensitivity to subhaloes (and false positives). Thus, the eBPL does still produce false-positive detections when it fits Projection 2; however, to see these using a pixelized source one would likely require much higher S/N data.

4.4 Summary

For Projection 1, we saw false-positive detections and an inability to recover an input subhalo's mass correctly. In contrast, for Projection 2, we did not infer a false-positive detection (when assuming a pixelized source) and the subhalo's true mass is covered by the posterior. Taking into account the different convergence shapes of the two projections (see Fig. 3), we speculate that the inaccurate inferences on subhaloes for Projection 1 are caused by the clear

mismatch in the shape of the eBPL and the more elliptical input profile. We also speculate that the better performance seen for Projection 2 is because its rounder convergence is easier for the eBPL to model. However, due to the limited number of projections available, we cannot generalize these conclusions any further. We only saw the eBPL produced false positives in Projection 1 and there is a possibility that it is a different property of the lens driving this result. When analysing real lens systems we will look to see whether departures from ellipticity in the lens galaxy's light (Nightingale et al. 2022) are correlated with subhalo detections, possibility indicating a false-positive signal.

5 A DECOMPOSED MODEL

Motivated by the inability of the eBPL to provide a robust subhalo inference, we now consider the decomposed model, which models a galaxy's stellar and dark matter mass separately. This includes sufficient freedom to capture complex features such as a pointy convergence profile, or other departures from elliptical symmetry.

5.1 Model introduction

In most strong lens images, we observe not only the lensed source's light, but also the light emitted from the lens galaxy, which should approximately trace its stellar mass distribution. For example, through inspection of the lens galaxy's light profile, we can estimate the position angle and axial ratio of the lens's stellar mass profile. More detailed light profile fits can provide us with a more detailed model of the stellar mass distribution. We now explore the potential of utilizing this information and if it can allow us to correctly recover the subhalo information hidden in the source's lensed images. We fit the lens's mass using a decomposed model which treats the lens galaxy's stellar mass and its dark matter mass separately. This type of model has been fitted in many previous studies (Dye & Warren 2005; Suyu et al. 2014; Wong et al. 2017) and Nightingale et al. (2019) showed using *HST* imaging of three SLACS lenses that such models capture variations in ellipticity and position angle within a galaxy that are indicative of pointy mass distributions.

For the stellar mass, we assume it exactly traces the stellar light, which allows us to directly transform between the two by multiplying by a constant mass-to-light ratio (M/L) parameter, which can be described as

$$\kappa(r) = \Psi \cdot I(r) , \qquad (6)$$

where $I(r)$ corresponds to the light profile and Ψ is its 'M/L'. For simplicity, we directly take the input stellar mass of the simulation's particle data as our lens light and therefore do not consider a more realistic galaxy light simulation process. In that sense, the '$I(r)$' is equivalent to the convergence profile and thus Ψ becomes a dimensionless quantity and is set to be 1.0.

To utilize the 'lens light' information, we model the 'lens light' with three cored Sérsic profiles as described by equation (1). We opt for the cored Sérsic because of the simulated galaxy's core; for real lenses we anticipate that the regular non-cored Sérsic profile will suffice. We impose that the 3 cored Sérsic profiles share the same centre; however, allow for them to have different position angles and axis ratios. We use three profiles because fits using two profiles do not fully capture the features of the 'lens light' (e.g. clear spatially correlated normalized residuals are seen when the best-fitting 2 cored Sérsic model is subtracted from the true stellar mass distribution). In Fig. 9, we show the input lens light (left-hand column), best-fitting 3 cored Sérsic profiles (middle column), and
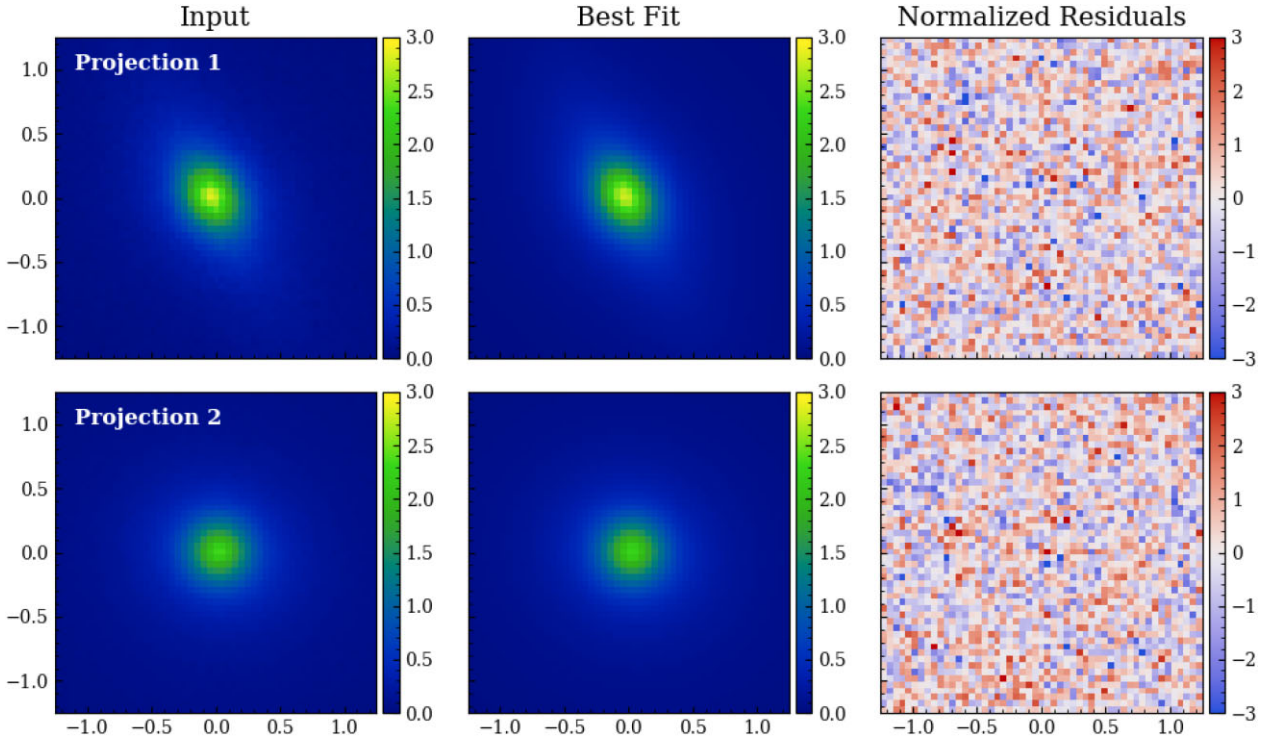
Figure 9. Left-hand column: The mock lens light images. Middle column: The best recovered images using 3 cored Sérsics to model the lens light. Right column: The corresponding normalized residuals. The top row shows the case of Projection 1 and the bottom row shows the case of Projection 2. The units of the *y* and *x* axes are arcsec.

corresponding normalized residuals (right-hand column). For both projections (upper row corresponds to the Projection 1 and the lower one is for the Projection 2) the light is well fit by 3 cored Sérsics. Later, in our lens mass modelling, we fix the stellar mass distribution to be exactly the same as the best-fitting 3 cored Sérsic profiles obtained from fitting the lens light, except for a free Ψ which changes the overall normalization of the projected stellar mass distribution.

In addition to the stellar mass, we include an elliptical NFW profile into the lens model (to account for the dark matter). This has six free parameters: a scale radius, r_s, and scale convergence, κ_{NFW}; two ellipticity components; and the 2D coordinates of the halo centre. As in the eBPL case, we include an external shear in the decomposed lens model.

Neither cored Sérsic nor elliptical NFW profiles have analytical formulae for their deflection angles. For fast computation, we follow Shajib (2019) and use a sum of 2D Gaussian profiles to approximate the cored Sérsic and elliptical NFW profiles. The resulting deflection angles are simply a sum of the deflection angles of the individual Gaussian profiles, which can be efficiently computed using analytical formulae. To be specific, in our work, in most cases we approximate a cored Sérsic profile by 30 Gaussian profiles with their standard deviations uniformly distributed in the $\log_{10}$ space between 0.01 and $50r_e$, where r_e is the effective radius of the cored Sérsic. Similarly, for an elliptical NFW profile, we also approximate it with 30 Gaussians and the standard deviations of those Gaussians are uniformly distributed in the $\log_{10}$ space between 0.0005 and $30r_s$. We noticed that one of the best-fitting cored Sérsic components to the 'lens light' of Projection 2 has a Sérsic index of 0.51 and for that profile the decomposition formula (equation 5 of Shajib 2019) becomes numerically unstable. For that one particular case, we instead decompose the Sérsic profile into a sum of Gaussians using Cappellari (2002)'s method, which optimizes the standard deviations

and amplitudes of those Gaussians at the same time. We have tested our choices of the parameters of the Gaussian decomposition method across a large variety of cored Sérsics and elliptical NFW profiles to ensure that errors of approximating the deflection angles are much smaller than the perturbation of a subhalo of interest. In Table 2, we summarize our lens model parameters.

The approach we follow cannot be straightforwardly translated to real data. For example, we have modelled the lens's light in the absence of the source light and ignored potential complications such as a radial gradient in the M/L. The goal of this work is not to present a method that can be directly transferred to the fitting of real data, but simply to show that when sufficient complexity is added to the lens mass model one's inference on subhalo properties improves. Nevertheless, Nightingale et al. (2019) have already shown how `PyAutoLens` can fit this type of model to real data and we expand on this further in Section 6.3.

5.2 Results

We now present results using the decomposed model, following the same structure we used for the eBPL results, whereby we begin with the smooth test results (where no subhalo is present in the simulated data) followed by results where the simulated data include a subhalo.

In Fig. 10, we compare the input and best-fitting model images for smooth cases. As shown by the normalized residuals in the third column, no clear correlated residuals exist, which indicates an overall good fit with the decomposed model. Comparing the results with the equivalent BPL results in Fig. 5, we see that the BPL results are indistinguishable from the decomposed model results in terms of the residuals, which confirms again that 'subhalo-like' perturbations

Table 2. Parameters and priors for the decomposed model. Parameters with values shown in '()' or '{}' are fixed during the modelling. Parameters with values shown as '[a, b]' are fit for, with a uniform prior between a and b.

	Projection 1	Projection 2
Stellar mass	**3 Core Sérsics**	
Centre (x, y) [(arcsec, arcsec)]	$(0.008, -0.036)$	$(0.003, 0.022)$
I'	$\{0.44, 0.60, 0.31\}$	$\{0.32, 0.50, 1.06\}$
r_e (arcsec)	$\{0.65, 0.02, 4.42\}$	$\{0.11, 0.18, 2.39\}$
r_c (arcsec)	$\{0.14, 0.27, 0.25\}$	$\{0.35, 0.13, 0.02\}$
n	$\{1.44, 4.36, 4.91\}$	$\{2.64, 0.51, 2.31\}$
Position angle (°)	$\{-62, -59, -45\}$	$\{-27, 69, -73\}$
Axial ratio	$\{0.33, 0.82, 0.82\}$	$\{0.90, 0.89, 0.90\}$
Ψ	$[0.8, 1.2]$	
Redshift	0.2	
MGE $\{n, r_{min}, r_{max}\}$	$\{30, 0.01r_e, 50r_e\}$	
Dark matter mass	**NFW**	
Centre (x, y) (arcsec, arcsec)	$[-0.1, 0.1]$	
$\log_{10} \kappa_{NFW}$	$[-2, 0.3]$	
r_s (arcsec)	$[10, 50]$	
e_1	$[-1.0, 1.0]$	
e_2	$[-1.0, 1.0]$	
Redshift	0.2	
MGE $\{n, r_{min}, r_{max}\}$	$\{30, 0.0005r_s, 30r_s\}$	
External shear		
γ_{1ext}	$[-0.2, 0.2]$	
γ_{2ext}	$[-0.2, 0.2]$	
Subhalo	**Spherical NFW**	
Centre (x, y) [(arcsec, arcsec)]	$([-1.0, 1.0], [-1.0, 1.0])$	
$\log_{10} m_{200}$ (M$_\odot$)	$[6, 11]$	
Mass-concentration relation	Ludlow et al. (2016)	

cannot be detected visually from the residual maps and we have to rely on careful statistical comparisons to make inferences about subhaloes.

In the left-hand column of Fig. 11, we first show the maximum log likelihood difference maps when modelling the smooth image with the decomposed model described above. For both projections the decomposed model fits the image accurately with a maximum ΔL value below 5, correctly indicating that no subhalo exists in the lens galaxy. Unlike the eBPL, the decomposed model does not give false-positive signals in our 'smooth tests'.

In the middle and right-hand columns of Fig. 11, we show the ΔL maps when a 5×10^8 or 5×10^9 M$_\odot$ subhalo is added to the lens galaxy at the positions marked by the white triangles. For Projection 1 (upper panels), the region where we detect the maximum ΔL is consistent with the position of each input subhalo. For an input subhalo of 5×10^8 M$_\odot$, the result shows a plausible detection where the maximum ΔL is 9.3, whereas for an input subhalo of mass 5×10^9 M$_\odot$, the detection is even clearer with a maximum ΔL of 36.0. Having successfully detected the subhalo in each case, we continue on to the subhalo refining fit, with Fig. 12 showing the inferred posteriors of the subhalo parameters. For both cases, the subhalo parameters are correctly recovered within 99 per cent credible regions. For a 5×10^8 M$_\odot$ subhalo, the recovered value is $10^{8.9^{+0.8}_{-2.6}}$ M$_\odot$, and for a 5×10^9 M$\odot$ subhalo, the recovered value is $10^{9.5^{+0.5}_{-0.4}}$ M$_\odot$.

For Projection 2, we only get a detection when the true subhalo mass is 5×10^9 M$_\odot$, with a maximum ΔL of 15.8 (in the middle cell of the bottom row). With a 5×10^8 M$_\odot$ subhalo, all subhalo-

position cells have $\Delta L < 5$, corresponding to no detection. In Fig. 13, we show the subhalo posteriors obtained from the subhalo refining phase for the 5×10^9 M$_\odot$ case. We recover the input subhalo mass, with a 99 per cent credible region on m_{200} of $10^{9.5^{+0.4}_{-0.7}}$ M$_\odot$. For the non-detection of the 5×10^8 M$_\odot$ subhalo, we do not believe this is a failure of the decomposed model, but instead a limitation of the data quality. In fact, if we check the inferred subhalo parameters for the sub grid cell which contains the input subhalo, the inferred subhalo's mass is $10^{8.7^{+0.7}_{-2.5}}$ M$_\odot$, which is still consistent with the true input mass. Thus, our inference on the subhalo's parameters is consistent with the truth; however, we have insufficient S/N for the model to be favoured in terms of ΔL.

We note that in Figs 12 and 13, the posterior distributions are not smooth, and having a 'patchy' appearance in the 2D marginalized posteriors and 'wiggles' in the 1D posteriors. These arise due to the pixelized source plane. The source-plane pixelization is created from a Voronoi tessellation of generating points, where the generating points are first placed in the image plane and then mapped into the source plane. Changes to the mass model change the mapping from image plane back to the source plane, such that the positions and shapes of the source-plane pixels vary as the lens mass model is changed. Certain locations for pixel boundaries may be more or less able to reproduce the observed data, leading to small changes to the mass model parameters capable of macroscopic changes to the likelihood. This phenomenon is more significant for more complex mass models which have more parameters and freedom to allocate those source pixel grids on the source plane. As a result, we get unsmooth posteriors for our decomposed model. The work of Etherington et al. (2022) discusses this further and presents a solution using a cap on the log likelihood of the model-fit.

5.3 Parametric source

In Appendix A, we again verify that our conclusions hold when we assume a cored Sérsic profile for the source. For both projections, the decomposed model does not give a false positive; with the highest value of $\Delta L = 2.5$. Note that, for a cored Sérsic source, false positives were detected for both projections. Given that fits assuming a cored Sérsic for the source give a much higher sensitivity to subhaloes and false positives, this further strengthens our conclusion that by adding the right type of complexity to the decomposed mass model removes the presence of false positives. As a result of this increased sensitivity, the 5×10^8 M$_\odot$ subhalo is also detected successfully in Projection 2, which is not the case for the pixelized source.

5.4 Offset true positive detections

In the right-hand panels of Fig. 11, we note increases of $\Delta L = \sim 10$–15 away from the true location of the $m_{200} = 5 \times 10^9$ M$_\odot$ subhalo for both projections. These are solutions where an offset dark matter subhalo closely mimics the perturbing effects of the actual subhalo in the data. However, it is not a perfect representation of the actual subhalo, which is why fits at the true location infer higher overall log likelihood values. We do not consider these as false-positive detections because they are caused by the true presence of a detectable subhalo in the data. Should this behaviour be seen in real data we therefore should not discount the signal as a false positive. In fact, a candidate subhalo detection should be made and followed up with a second subhalo search which includes the first subhalo in the model, so as to validate the detection.
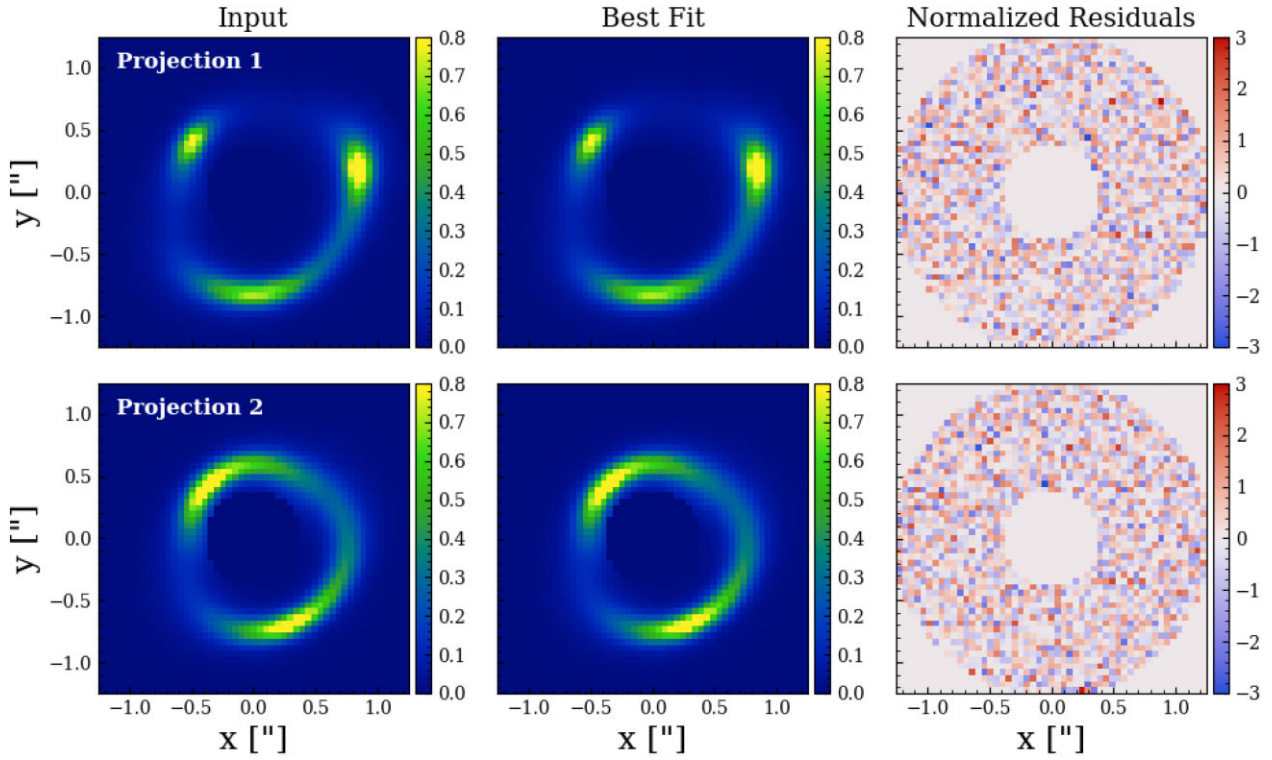
Figure 10. Left-hand column: The mock lensing images. Middle column: The best recovered images using the decomposed model to fit the lens mass. Right-hand column: The corresponding normalized residuals (residuals divided by the noise). The top row shows the case of Projection 1 and the bottom row shows the case of Projection 2. The colour bars for the left two columns are in units of $e^- \, pix^{-1} \, s^{-1}$.
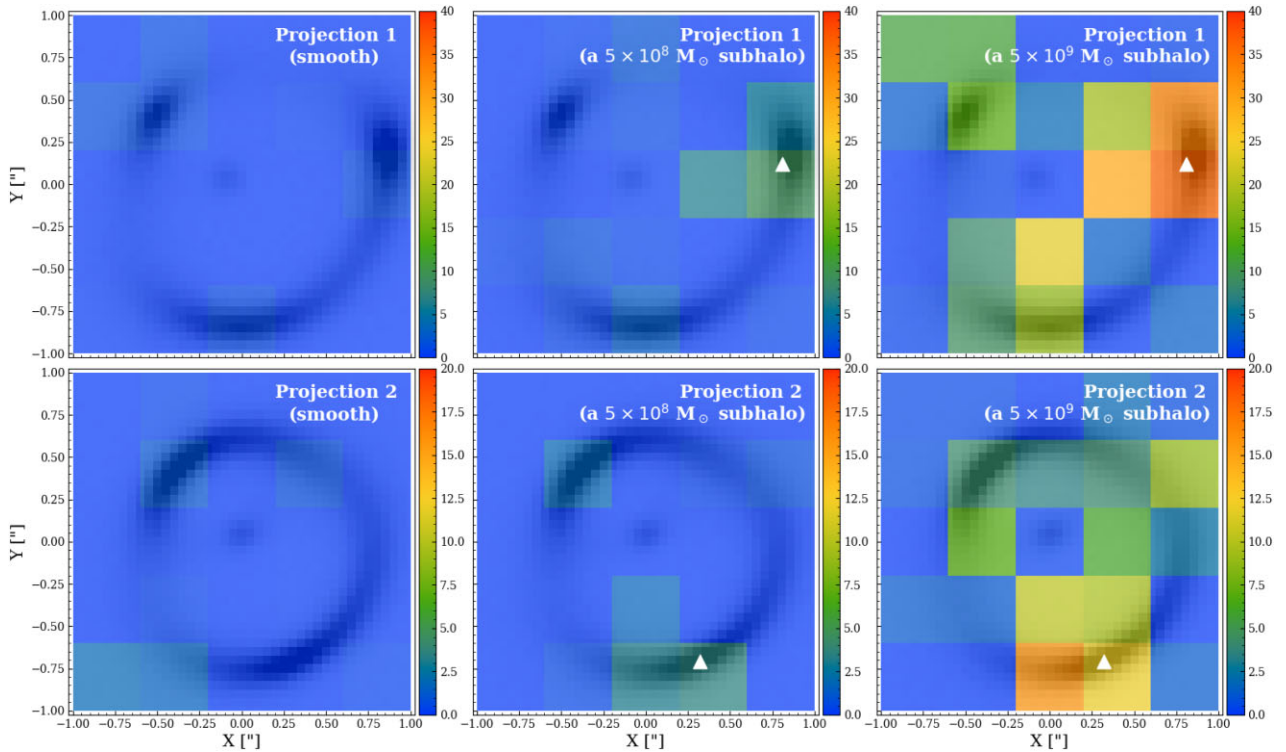

Figure 11. Modelling the lens with a decomposed stellar + dark matter model removes false-positive detections, and yields correct subhalo masses. Colours indicate the increase in maximum log likelihood, ΔL, when a subhalo is included in the fit to a lens that: has no subhaloes (left-hand panel), has a subhalo of mass $5 \times 10^8 \, M_\odot$ (middle panel), or has a subhalo of mass $5 \times 10^9 \, M_\odot$ (right-hand panel). The position of the subhalo in the fit is free to vary within squares of side 0.4 arcsec. The top and bottom rows show the results for Projection 1 and 2 (with different colour scales). White triangles mark the true locations of the subhaloes.
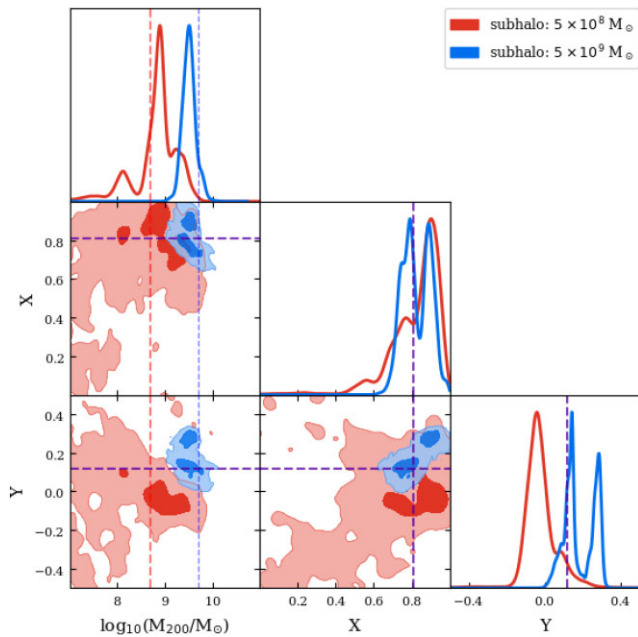
Figure 12. The posteriors on the subhalo parameters from the subhalo refining phase, fitting the decomposed model to mock data generated using Projection 1. Red and blue colours show the cases with an input subhalo mass of 5×10^8 and 5×10^9 $M_\odot$, respectively. The 2D contours cover the 68 and 99 per cent confidence regions. For 1D posteriors, the vertical dashed lines mark the true input values.
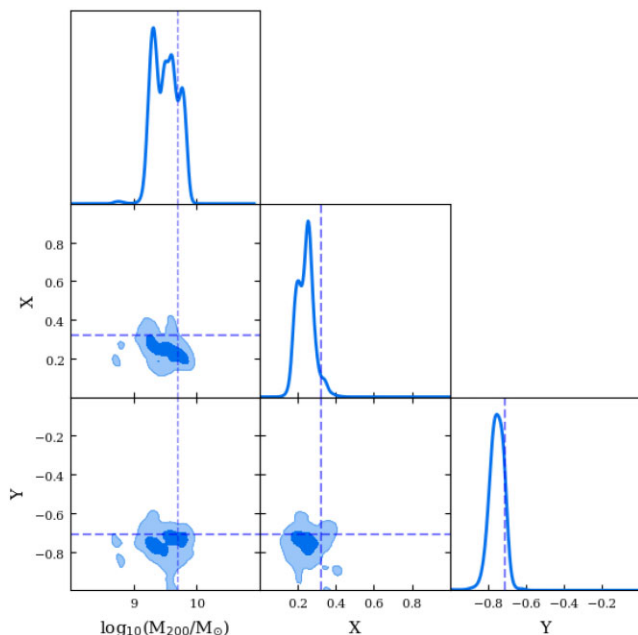

Figure 13. The same as Fig. 12, but for Projection 2. Only the case with a 5×10^9 $M_\odot$ subhalo is shown (a 5×10^8 $M_\odot$ subhalo is not clearly detected).

5.5 Summary

For both projections the decomposed model is a success. When we do not include a subhalo in the input lens galaxy, it returns no detections. When a subhalo is included in the mock data, it is able to correctly infer the existence of the subhalo through an increase in ΔL (at least for three of the four cases we tried). Furthermore, it

recovers the masses and positions of the subhaloes within 99 per cent credible regions. By utilizing (idealized) 'lens light' information, the decomposed model therefore successfully captures complexity in the mass profile (e.g. the non-elliptical shape) that the eBPL could not. The success of the decomposed model confirms that for subhalo detection, it is vital to model the lens galaxy's mass accurately.

6 DISCUSSION

6.1 Implications for strong lensing subhalo detection

Our results confirm that if a dark matter subhalo is located near the emission of a strongly lensed source galaxy, its perturbing effects mean that its presence can be inferred. For gravitational imaging,[9] our work demonstrates this for first time by simulating the lens galaxy using a mass distribution derived from the particle data of a cosmological simulation, which therefore does not make idealized assumptions like a single axis of ellipticity. However, we also showed that assuming an overly simplistic mass model for the lens galaxy which lacks certain complexity compared to the true underlying mass distribution has two negative effects on the subhalo inference: (i) It may lead to false-positive detections of a dark matter subhalo even though a subhalo is not present in the data and (ii) when a subhalo is truly present in the data it may lead to systematic biases on the inferred subhalo mass by a factor of 4–5.

The notion that a mismatch in mass profile shape could lead to false-positive subhalo detections supports the analysis of Ritondale et al. (2019), who noted several false-positive signals found in real lensing systems in the BELLS-GALLERY sample. For example, they noted an increase in log Bayesian evidence of 72 in the lens SDSSJ0755+3445, but demonstrated – using a potential correction technique (Koopmans 2005; Vegetti & Koopmans 2009b) – that the mass model could be improved by small corrections over a large angular scale, as opposed to a localised correction reminiscent of a subhalo. This indicates that the subhalo-like signal is probably due to the mismatch in the macro models, as we saw in our tests. False positives are also partly the reason why Vegetti & Vogelsberger (2014) and Vegetti et al. (2018) require Bayesian evidence increases of 50 and 100 to claim a dark matter detection; values below this threshold may be false positives (the authors also require validation via potential corrections). Whilst the false positives in this work did not create Bayesian evidence increases above 30, the overall size of the increase depends on the properties of the strong lens and sources simulated, the S/N of the data and model used to fit the data. In Appendix A, false positives with evidence increases above 200 are inferred. Therefore, 'our results do not indicate that previous detections of dark matter subhaloes in strong lenses are false positives'. Instead, they show the importance of techniques like the potential corrections and we provide insight on why these methods are able to distinguish a subhalo detection from missing complexity in the mass model.

6.2 Is our simulated lens galaxy realistic?

It is important to consider how realistic the simulated galaxy used in this work is. As discussed previously, the galaxy was selected to be similar to lens galaxies from the SLACS survey. It has a typical halo mass for a SLACS lens and its stellar mass and size (i.e.

[9]Similar tests on flux ratios have previously been explored by Hsueh et al. (2018).

the half-light radius) follow recent observations of massive galaxies (Huang et al. 2018). The central galaxy has a complex shape, with isophotes that change shape when viewed from different directions, and where the shapes of the isophotes can vary with radius for a fixed viewing direction. We speculated that this departure from elliptical symmetry drove the false positives, because they are only seen for the projection where the mass distribution is highly elliptical. A varying isophotal shape with radius is commonly seen in observations of massive elliptical galaxies with comparable mass to SLACS strong lenses. For example, over one-third of galaxies with stellar masses above $10^{11.5}$ $M_{\odot}$ taken from the MASSIVE survey show isophotal position-angle rotations (Goullaud et al. 2018), known as 'isophotal twists' (see Oh, Greene & Lackner 2017 for similar results in lower mass early-type galaxies). Similar features are also reported in three strong lenses by Nightingale et al. (2019). We therefore believe this aspect of our simulation is representative of real strong lenses and is a plausible cause of some of the false positives in the SLACS and BELLS-GALLERY lenses discussed previously.

The simulated galaxy also has a sub-kpc core, which generates a central image in our mock lens images. This phenomenon is seen in other works which simulate strong lenses from cosmological simulations (Mukherjee et al. 2018; Despali et al. 2020; Ding et al. 2021), with the core due to insufficient simulation resolution. Central images of this brightness are not seen in real observations of strong lenses, therefore such a large core is unrealistic. To ensure it does not impact our tests, the mass model parametrizations fitted in this work all included cores. We masked the central image so as to ensure the mass models did not utilize additional information that is not present in real images of strong lenses. Whilst this aspect of the simulated lens is therefore not realistic, the mass modelling performed in this work ensures that we can generalize our conclusions to the analysis of real data.

6.3 Application to real data

Our next step is applying the decomposed model to real data. We expect that we will be able to fit mass models which omit parameters that account for a core, given that the core feature is a consequence of the inadequate simulation resolution. For the decomposed model, we will likely fit regular Sérsic functions instead of the cored Sérsics fitted in this work.

The decomposed model verified that if a mass model can accurately capture the lens galaxy's complexity, it will improve the subhalo inference. This work used information from the simulation that is not available when analysing real data, e.g. we utilized our true knowledge of the lens's stellar mass distribution. Nevertheless, we believe these models can be translated to real data, where the light emitted from the lens galaxy acts as a tracer for the stellar mass, information which is often omitted when modelling a strong lens (e.g. by assuming a power-law mass model). This approach to lens modelling was explored in Nightingale et al. (2019), who fitted a decomposed stellar plus dark matter to three strong lenses. The authors showed that all three lenses showed isophotal twists in their stellar emission and that when this was modelled using two stellar components with different ellipticities and position angles it improved the mass model compared to a model assuming a single elliptical geometry. We are now investigating whether these lens systems produce subhalo detections, which would be indicative of a false positive.

The decomposed model must also make assumptions in converting light to mass. For example, whether the Sérsic profiles representing each stellar component share the same M/L or whether each ratio is

a free parameter in the model. For each component, one must also choose whether the lens model accounts for a radially varying M/L (Napolitano et al. 2005; Tortora et al. 2011; Ge et al. 2021). The assumption of an elliptical NFW profile to describe the dark matter poses another possible mismatch. The main concern on small scales is whether the central slope, which in simulations is affected by the presence of baryons, is equal to the NFW one. To take this into account when modelling real data, we could model the dark matter as a profile with a free central slope, e.g. a generalized NFW profile (Zhao 1996), or explicitly model the way baryons are expected to alter the dark matter distribution (Callingham et al. 2020; Cautun et al. 2020). We do not expect this to be a significant issue since for galaxy-galaxy strong lensing, the dark matter mass is typically subdominant in the region of interest (e.g. Li et al. 2016a, 2019). In future work, we will seek to understand the importance of all these different assumptions with a view to improving the dark matter subhalo inference.

6.4 Subhalo sensitivity

If the decomposed model can be successfully fitted to real data, it also has implications for how sensitive strong lensing is to low mass dark matter subhaloes. Firstly, if the method is able to reduce or remove the Bayesian evidence thresholds applied by works like Vegetti & Vogelsberger (2014) to remove false positives, this will make us sensitive to lower mass subhaloes (which produce smaller evidence increases). Furthermore, because the decomposed model uses the stellar light as additional information which constrains the mass model, this may further boost one's sensitivity to subhaloes by reducing the degeneracy between the lens galaxy's mass model and subhalo. This will require that sensitivity mapping of a strong lens, which quantifies what mass subhaloes one will detect if truly present in the data (Despali et al. 2020; Amorisco et al. 2022; He et al. 2022), is performed using the decomposed model, as opposed to the power-law model assumed in previous studies. The same level of care will be necessary in understanding how robust assumptions associated with the M/L and dark matter are.

6.5 Other lensing studies

A mass model mismatch has also been discussed in the analysis of strongly lensed quasars. Hsueh et al. (2017) showed that the flux ratio anomalies observed in lens system CLASS B0712+472 can be largely resolved by additionally adding a disk profile to the lensing model. The works of Gomer & Williams (2020, 2021), Cao et al. (2021), and Van de Vyvere et al. (2022) show that such mismatches can impact on the inference of the Hubble constant via time-delay cosmology.

7 CONCLUSIONS

With a large increase in the number of observed galaxy-galaxy strong lenses expected within this decade, strong lensing could soon push the constraints on the halo mass function to low enough masses that it provides evidence in favour of or against WDM models. However, detecting dark matter subhaloes through strong lensing is a challenging problem due to the complexity of the lens galaxy's mass distribution. In this work, we use a massive elliptical galaxy extracted from a state-of-the-art hydrodynamic simulation to create mock strong lens images. We represent the simulated galaxy's projected mass distribution as a sum of elliptical Gaussian profiles, which shows departures from the idealized elliptically symmetric

mass models typically employed to analyse strong lenses [e.g. the power-law profile (Tessore & Metcalf 2015)]. We project the same simulated galaxy along two different line-of-sight directions, with one projection producing a pointy 'American football like'-shape and the other one appearing rounder.

For each projection, we simulate three strong lens imaging data sets. The first data set does not include a dark matter subhalo, whereas the other two include a $m_{200} = 5 \times 10^8$ M$_\odot$ and $m_{200} = 5 \times 10^9$ M$_\odot$ dark matter subhalo near the lensed source's light. To every data set, we fit two lens mass models: (i) an eBPL mass model (O'Riordan et al. 2019) which represents the overall mass distribution of the lens galaxy (e.g. stars and dark matter) and (ii) a decomposed model that models the stellar and dark matter mass separately (using the stellar particle data from the simulation to constrain part of the stellar mass model). For both models, we investigate fits which include a dark matter subhalo in the lens mass model, and therefore quantify whether we can accurately recover a dark matter subhalo when it is included in the simulation as well as whether we incorrectly infer the presence of a subhalo when it is not truly there; a false positive.

Our main results can be summarized as follows:

(i) When using an eBPL model to fit the lens mass to the pointy projection without a dark matter subhalo, a false-positive detection is inferred at over 3σ confidence. For the same projection, when a 5×10^8 or 5×10^9 M$_\odot$ subhalo is added to the mock lens, the fit correctly recovers the subhalo but overestimates its mass by a factor of 4–5, with the true input mass outside the inferred 99 per cent credible regions. However, when modelling data from the projection with a rounder convergence, the eBPL model does not give a false positive and recovers the input 5×10^9 M$_\odot$ subhalo's mass accurately (the 5×10^8 M$_\odot$ subhalo is not detected due to insufficient data quality).

(ii) When using the decomposed model to fit the lens mass, for both projections, we get no false positives and correctly recover the properties of an input subhalo when there is sufficient data quality to detect it.

The eBPL total mass model therefore shows undesirable results, including false positives and an inaccurate estimate of the subhalo mass, which the decomposed mass model does not. We speculate that this is because the eBPL parametrization does not capture aspects of the simulated lens's mass distribution. In particular, the eBPL does not capture the varying ellipticity and orientation seen in the pointy projection's mass distribution. The decomposed mass model does not assume a single elliptical mass distribution and can therefore account for this variation in ellipticity and orientation. Its improved model of the lens galaxy's mass therefore offers an improved subhalo inference which does not suffer false-positive detections.

Our results do not imply that previous detections of dark matter subhaloes in strong lenses are false positives (e.g. Vegetti & Vogelsberger 2014). These studies are fully aware of the false-positive phenomena and they require a subhalo detection to pass stringent criteria to be considered a genuine dark matter subhalo. This includes a pixel-based correction to the gravitational potential (Koopmans 2005) which accounts for the types of deficiencies in the mass model discussed in this work. In fact, our work demonstrates that dark matter substructures can be successfully detected in images of strong lenses, even when the lens galaxy's mass distribution is more complex than the mass model assumed to fit it.

Our work highlights the benefits of using cosmological simulations to test strong lens modelling methodology. When the eBPL showed inaccurate results, we were able to compare directly to the simulation's particle data in order to understand what complexity the model is missing. This is not possible when analysing real images of strong lenses. We are now looking to apply what we have learned in this study to real data, and fit strong lenses from existing lens samples with decomposed mass models which, crucially, relax the assumption of a single axis of ellipticity. Applying the models to real data has challenges, e.g. instead of relying on the simulation's stellar particle data we will need to use the lens's light to constrain the stellar mass (Nightingale et al. 2019). However, the payoff could be huge, allowing us to more reliably detect lower mass dark matter substructures, that could potentially push our sensitivity down to pivotal masses of $m_{200} = 5 \times 10^8$ M$_\odot$ where many viable alternatives to the CDM model begin to make different, testable predictions.

SOFTWARE

This work used the following software packages: Astropy (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018), Colossus (Diemer 2018), Corner.py (Foreman-Mackey 2016), Dynesty (Speagle 2020), GLAMER (Metcalf & Petkova 2014; Petkova et al. 2014; Metcalf et al. 2020), Matplotlib (Hunter 2007), Numba (Lam, Pitrou & Seibert 2015), NumPy (van der Walt, Colbert & Varoquaux 2011), PyAutoFit (Nightingale et al. 2021a), PyAutoLens (Nightingale & Dye 2015; Nightingale et al. 2018, 2021b), Pyquad (Kelly 2020), Python (Van Rossum & Drake 2009), Scikit-image (Van der Walt et al. 2014), Scikit-learn (Pedregosa et al. 2011), Scipy (Virtanen et al. 2020), and SQLite (Hipp 2020).

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Amorisco N. C. et al., 2022, MNRAS, 510, 2464
Astropy Collaboration et al., 2013, A&A, 558, A33

Baltz E. A., Marshall P., Oguri M., 2009, J. Cosmol. Astropart. Phys., 2009, 015

Benitez-Llambay A., Frenk C., 2020, MNRAS, 498, 4887

Benson A. J., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2002, MNRAS, 333, 177

Bentley J. L., 1975, Commun. ACM, 18, 509

Bolton A. S. et al., 2012, ApJ, 757, 82

Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Moustakas L. A., 2006, ApJ, 638, 703

Bundy K. et al., 2015, ApJ, 798, 7

Callingham T. M., Cautun M., Deason A. J., Frenk C. S., Grand R. J. J., Marinacci F., Pakmor R., 2020, MNRAS, 495, 12

Cao X. et al., 2021, Res. Astron. Astrophys., 22, 025014

Cappellari M., 2002, MNRAS, 333, 400

Cappellari M., 2008, MNRAS, 390, 71

Cautun M. et al., 2020, MNRAS, 494, 4291

Collett T. E., 2015, ApJ, 811, 20

Collett T. E., Auger M. W., 2014, MNRAS, 443, 969

Despali G., Lovell M., Vegetti S., Crain R. A., Oppenheimer B. D., 2020, MNRAS, 491, 1295

Despali G., Vegetti S., 2017, MNRAS, 469, 1997

Diemer B., 2018, Astrophys. J. Suppl. Ser., 239, 35

Ding X. et al., 2021, MNRAS, 503, 1096

Dye S. et al., 2015, MNRAS, 452, 2258

Dye S., Warren S. J., 2005, ApJ, 623, 31

Efstathiou G., 1992, MNRAS, 256, 43P

Ene I. et al., 2018, MNRAS, 479, 2810

Enzi W. et al., 2021, MNRAS, 506, 5848

Enzi W., Vegetti S., Despali G., Hsueh J.-W., Metcalf R. B., 2020, MNRAS, 496, 1718

Etherington A. et al., 2022, MNRAS, 517, 3275

Foreman-Mackey D., 2016, J. Open Source Softw., 1, 24

Frenk C. S., White S. D. M., 2012, Ann. Phys., Lpz., 524, 507

Frenk C. S., White S. D. M., Efstathiou G., Davis M., 1990, ApJ, 351, 10

Ge J., Mao S., Lu Y., Cappellari M., Long R. J., Yan R., 2021, MNRAS, 507, 2488

Gilman D., Birrer S., Nierenberg A., Treu T., Du X., Benson A., 2020a, MNRAS, 491, 6077

Gilman D., Birrer S., Treu T., 2020b, A&A, 642, A194

Gilman D., Birrer S., Treu T., Nierenberg A., Benson A., 2019, MNRAS, 487, 5721

Gomer M. R., Williams L. L. R., 2021, MNRAS, 504, 1340

Gomer M., Williams L. L. R., 2020, J. Cosmol. Astropart. Phys., 2020, 045

Goullaud C. F., Jensen J. B., Blakeslee J. P., Ma C.-P., Greene J. E., Thomas J., 2018, ApJ, 856, 11

Graham A. W., Driver S. P., 2005, Publ. Astron. Soc. Aust., 22, 118

He Q. et al., 2020, MNRAS, 496, 4717

He Q. et al., 2022a, MNRAS, 511, 3046

He Q. et al., 2022b, MNRAS, 512, 5862

Hezaveh Y. D. et al., 2016, ApJ, 823, 37

Hipp R. D., 2020, SQLite, Available at: https://www.sqlite.org/index.html

Hsueh J. W. et al., 2017, MNRAS, 469, 3713

Hsueh J.-W., Despali G., Vegetti S., Xu D., Fassnacht C. D., Metcalf R. B., 2018, MNRAS, 475, 2438

Huang S. et al., 2018, MNRAS, 480, 521

Hunter J. D., 2007, Comput. Sci. Eng., 9, 90

Kelly A. J., 2020, pyquad, Available at: https://github.com/AshKelly/pyquad

Koopmans L. V. E., 2005, MNRAS, 363, 1136

Lam S. K., Pitrou A., Seibert S., 2015, Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, Numba: a LLVM-based Python JIT compiler, 1–6

Li H., Li R., Mao S., Xu D., Long R. J., Emsellem E., 2016a, MNRAS, 455, 3680

Li R. et al., 2019, MNRAS, 490, 2124

Li R., Frenk C. S., Cole S., Gao L., Bose S., Hellwing W. A., 2016b, MNRAS, 460, 363

Li R., Frenk C. S., Cole S., Wang Q., Gao L., 2017, MNRAS, 468, 1426

Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, MNRAS, 460, 1214

Maresca J., Dye S., Li N., 2021, MNRAS, 503, 2229

Metcalf B., Petkova M., Tessore N., Leier D., Giocoli C., Nugier F., 2020, glenco/glamer,

Metcalf R. B., Petkova M., 2014, MNRAS, 445, 1942

Mukherjee S. et al., 2018, MNRAS, 479, 4108

Mukherjee S., Koopmans L. V. E., Metcalf R. B., Tortora C., Schaller M., Schaye J., Vernardos G., Bellagamba F., 2021, MNRAS, 504, 3455

Muñoz J. A., Kochanek C. S., Keeton C. R., 2001, ApJ, 558, 657

Napolitano N. R. et al., 2005, MNRAS, 357, 691

Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563

Nightingale J. W. et al., 2021b, J. Open Source Softw., 6, 2825

Nightingale J. W. et al. 2022, preprint (arXiv:2209.10566)

Nightingale J. W., Dye S., 2015, MNRAS, 452, 2940

Nightingale J. W., Dye S., Massey R. J., 2018, MNRAS, 478, 4738(N18)

Nightingale J. W., Hayes R. G., Griffiths M., 2021a, J. Open Source Softw., 6, 2550

Nightingale J. W., Massey R. J., Harvey D. R., Cooper A. P., Etherington A., Tam S.-I., Hayes R. G., 2019, MNRAS, 489, 2049

O'Riordan C. M., Warren S. J., Mortlock D. J., 2019, MNRAS, 487, 5143

O'Riordan C. M., Warren S. J., Mortlock D. J., 2020, MNRAS, 496, 3424

O'Riordan C. M., Warren S. J., Mortlock D. J., 2021, MNRAS, 501, 3687

Oh S., Greene J. E., Lackner C. N., 2017, ApJ, 836, 115

Pedregosa F. et al., 2011, J. Mach. Learn. Res., 12, 2825

Petkova M., Metcalf R. B., Giocoli C., 2014, MNRAS, 445, 1954

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

Price-Whelan A. M. et al., 2018, AJ, 156, 123

Quinn J. et al., 2016, MNRAS, 459, 2394

Richings J., Frenk C., Jenkins A., Robertson A., Schaller M., 2021, MNRAS, 501, 4657

Ritondale E., Vegetti S., Despali G., Auger M. W., Koopmans L. V. E., McKean J. P., 2019, MNRAS, 485, 2179

Rozo E. et al., 2010, ApJ, 708, 645

Schaye J. et al., 2015, MNRAS, 446, 521

Shajib A. J., 2019, MNRAS, 488, 1387

Shu Y. et al., 2016, ApJ, 833, 264

Speagle J. S., 2020, MNRAS, 493, 3132

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726

Suyu S. H. et al., 2014, ApJ, 788, L35

Suyu S. H., Marshall P. J., Hobson M. P., Blandford R. D., 2006, MNRAS, 371, 983

Tessore N., Metcalf R. B., 2015, A&A, 580, A79

Tortora C., Napolitano N. R., Romanowsky A. J., Jetzer P., Cardone V. F., Capaccioli M., 2011, MNRAS, 418, 1557

Trujillo I., Erwin P., Asensio Ramos A., Graham A. W., 2004, AJ, 127, 1917

Van de Vyvere L., Gomer M. R., Sluse D., Xu D., Birrer S., Galan A., Vernardos G., 2022, A&A, 659, A127

Van de Vyvere L., Sluse D., Mukherjee S., Xu D., Birrer S., 2020, A&A, 644, A108

van der Walt S., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 22

Van der Walt S., Schönberger J. L., Nunez-Iglesias J., Boulogne F., Warner J. D., Yager N., Gouillart E., Yu T., 2014, PeerJ, 2, e453

Van Rossum G., Drake F. L., 2009, Python 3 Reference Manual. CreateSpace, Scotts Valley, CA

Vegetti S., Despali G., Lovell M. R., Enzi W., 2018, MNRAS, 481, 3661

Vegetti S., Koopmans L. V. E., 2009a, MNRAS, 392, 945

Vegetti S., Koopmans L. V. E., 2009b, MNRAS, 400, 1583

Vegetti S., Koopmans L. V. E., Bolton A., Treu T., Gavazzi R., 2010, MNRAS, 408, 1969

Vegetti S., Lagattuta D. J., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., 2012, Nature, 481, 341

Vegetti S., Vogelsberger M., 2014, MNRAS, 442, 3598

Virtanen P. et al., 2020, Nature Methods, 17, 261

Vogelsberger M. et al., 2014, MNRAS, 444, 1518

Wang J., Bose S., Frenk C. S., Gao L., Jenkins A., Springel V., White S. D. M., 2020, Nature, 585, 39

Warren S. J., Dye S., 2003, ApJ, 590, 673

Winn J. N., Rusin D., Kochanek C. S., 2004, Nature, 427, 613

Wong K. C. et al., 2017, MNRAS, 465, 4895

Xu D. D. et al., 2009, MNRAS, 398, 1235

Xu D., Springel V., Sluse D., Schneider P., Sonnenfeld A., Nelson D., Vogelsberger M., Hernquist L., 2017, MNRAS, 469, 1824

Zhao H., 1996, MNRAS, 278, 488

APPENDIX A: PARAMETRIC SOURCE RESULTS

This work primarily focuses on the lens mass distribution. The effects of source modelling on the subhalo inference are seldom discussed. In this section, to give a brief idea on how our results would be affected by source modelling, we fit the same mock data with the same mass models discussed above but with a parametric source model. To be specific, the source model we apply here has the same form used to simulate the data, which is an elliptical cored Sérsic profile. When simulating mock data, we have fixed its break radius to be 0.01 arcsec; however, when using it as a source model, we set its break radius to be a free parameter.

The increase in log likelihood for many model-fits including a subhalo, ΔL, are higher when we assume that the source is an elliptical cored Sérsic profile as opposed to a pixelized source. This is because the pixelized source models have a much higher level of freedom in how they fit the data. If a mass model provides a good – but not perfect – fit, the pixelization can make small adjustments to the source pixel values to fit the data equally well (Gilman et al. 2020b). This is appropriately penalized using a Bayesian framework
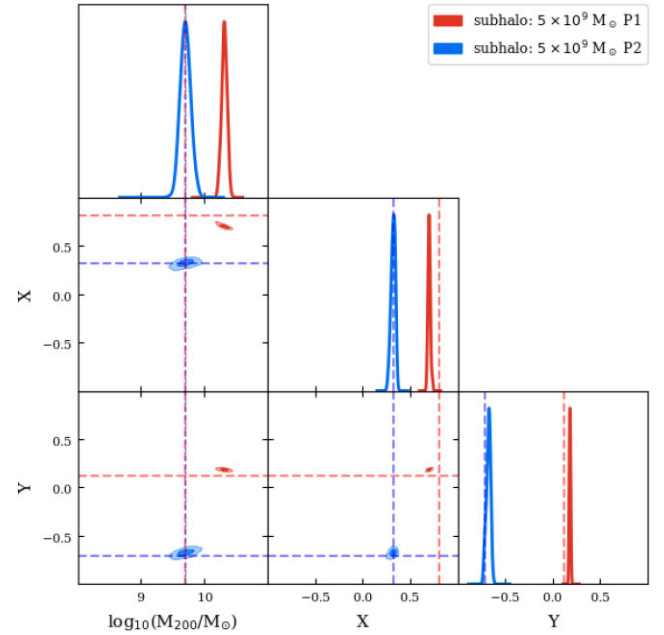

Figure A2. Posteriors of detected subhalo parameters of eBPL + cored Sérsic model. The red posteriors show the results for an input of a $5 \times 10^9 \, M_\odot$ subhalo in Projection 1. The blue posteriors show the results for an input of a $5 \times 10^9 \, M_\odot$ subhalo in Projection 2. The 2D contours cover the 68 and 99 per cent confidence regions. The dashed lines in corresponding colours marked the true input values.
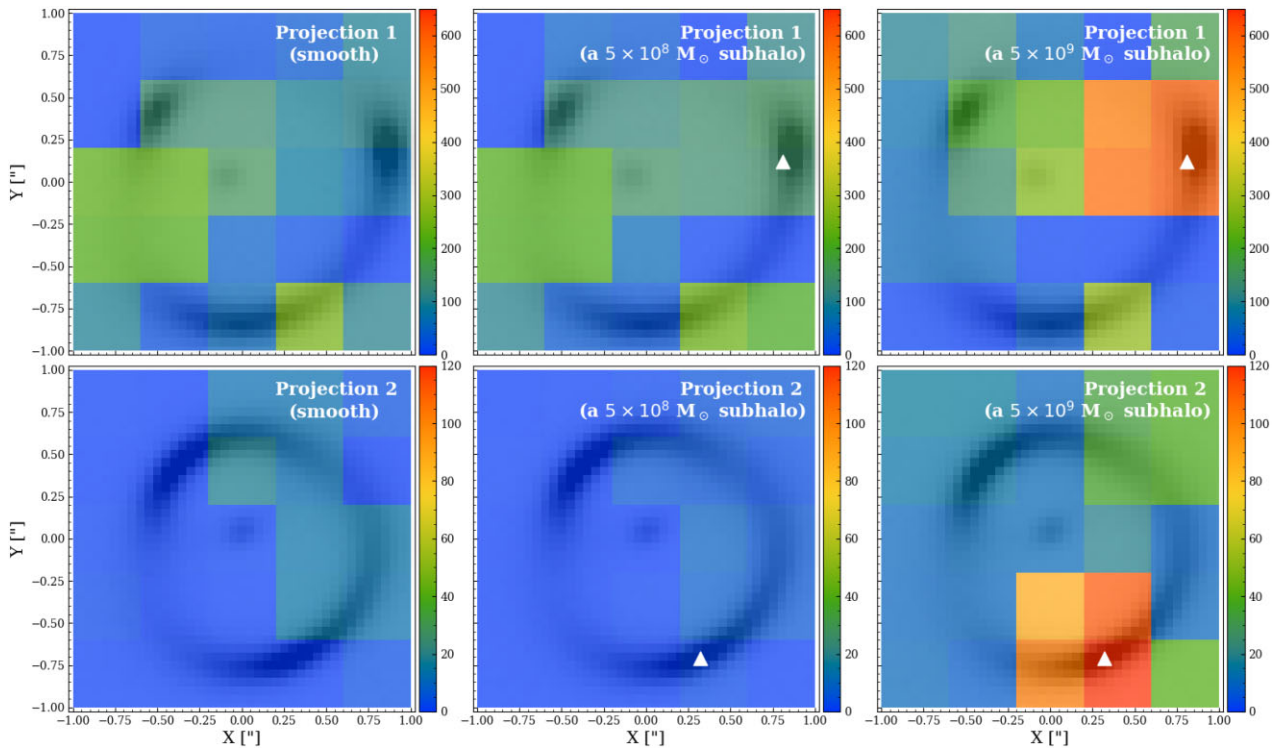

Figure A1. ΔL maps of using eBPL + cored Sérsic source model. Colours indicate the increase in maximum log likelihood, ΔL, when a subhalo is included inside 0.4 arcsec × 0.4 arcsec squares during a fit to a lens that has: no subhaloes (left-hand panel), a subhalo of mass $5 \times 10^8 \, M_\odot$ (middle panel), or a subhalo of mass $5 \times 10^9 \, M_\odot$ (right-hand panel). Top and bottom rows show the results for Projection 1 and 2 (with different colour scales). White triangles mark the true locations of the subhaloes.
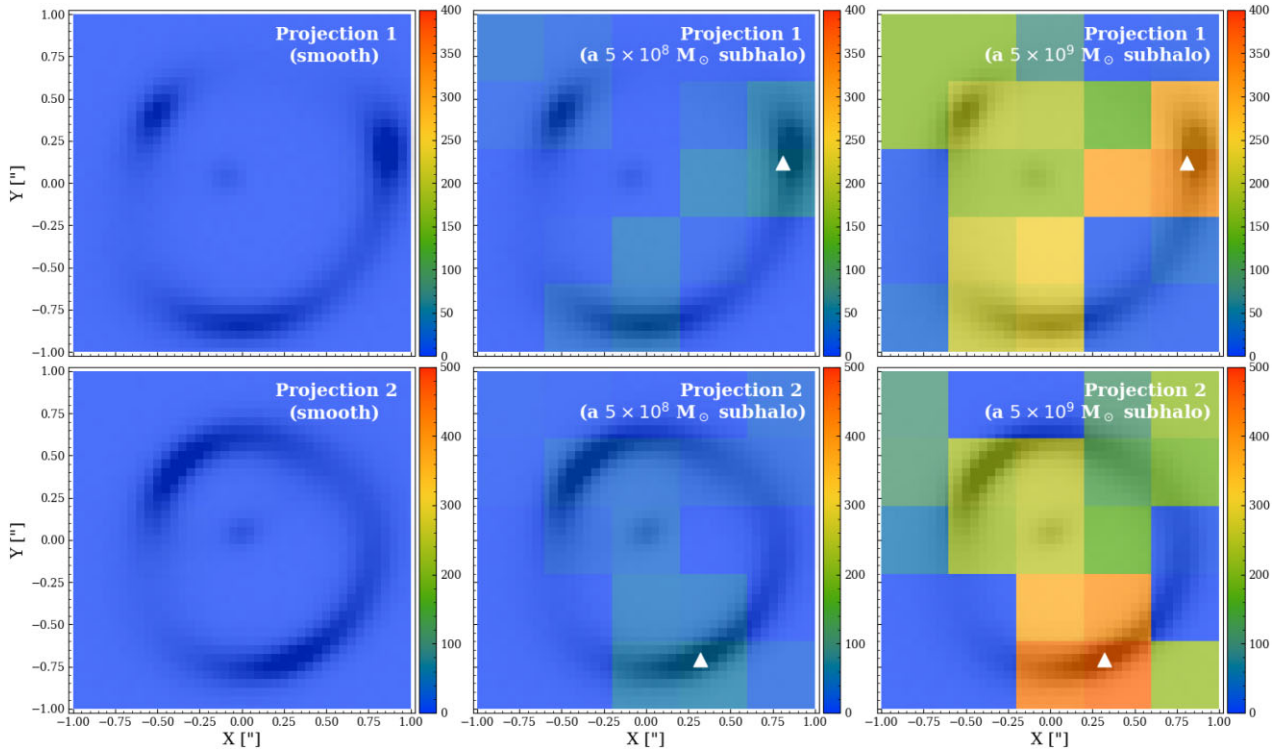
Figure A3. ΔL maps of using stellar + dark matter + cored Sérsic source model. Colours indicate the increase in maximum log likelihood, ΔL, when a subhalo is included inside 0.4 arcsec × 0.4 arcsec squares during a fit to a lens that has: no subhaloes (left-hand panel), a subhalo of mass 5×10^8 M$_\odot$ (middle panel), or a subhalo of mass 5×10^9 M$_\odot$ (right-hand panel). Top and bottom rows show the results for Projection 1 and 2 (with different colour scales). White triangles mark the true locations of the subhaloes.
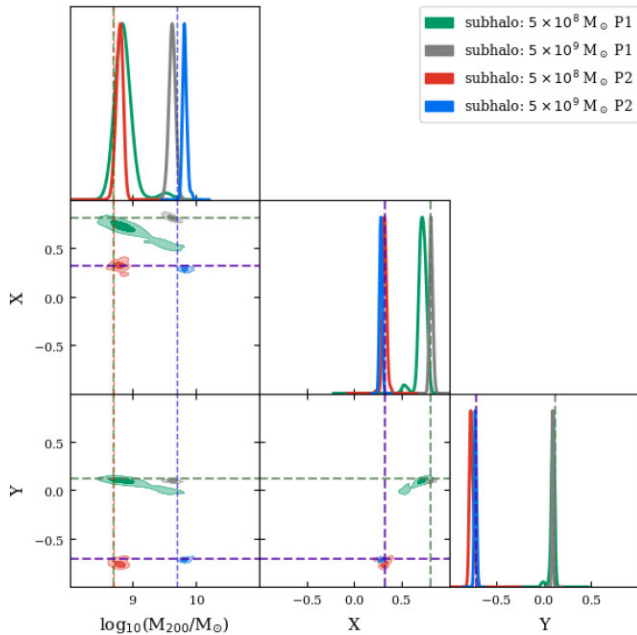

Figure A4. Posteriors of detected subhalo parameters of the stellar + dark matter + cored Sérsic source model. The green, grey, red, and blue posteriors, respectively, show the results for: an input of a 5×10^8 M$_\odot$ subhalo in Projection 1; an input of a 5×10^9 M$_\odot$ subhalo in Projection 1; an input of a 5×10^8 M$_\odot$ subhalo in Projection 2; an input of a 5×10^9 M$_\odot$ subhalo in Projection 2. The 2D contours cover the 68 and 99 per cent confidence regions. The dashed lines in corresponding colours marked the true input values.

(see Suyu et al. 2006 and N18), but nevertheless produces smaller likelihood contrasts than fitting a parametric source model like the cored Sérsic profile, which has a lot less freedom in adjusting its parameters in order to account for an inaccurate mass model. This is also dependent on the fact that the elliptical cored Sérsic profile was used to both simulate and fit the mock strong lenses; had there been a mismatch here parametric fits would likely not give such large ΔL values.

In Fig. A1, we show ΔL when fitting the data with an eBPL profile. For Projection 1, for both the smooth case and a 5×10^8 M$_\odot$ subhalo input case, the eBPL plus cored Sérsic source model returns similar results, with a highest ΔL giving ∼250 at the middle-left region indicating the existence of a $10^{10.1^{+0.2}_{-0.1}}$ M$_\odot$ subhalo, which is not consistent with our input (e.g. it is a false positive). For the case of a 5×10^9 M$_\odot$ input subhalo, the highest ΔL is ∼600 around the middle-right region, which is consistent with our input. For this case, we further model the subhalo by freeing its position and the posterior we get is shown in colour red in Fig. A2. We see that although the position is estimated around the true input, the subhalo's mass is overestimated by around 4 times, which is similar to our previous findings for a pixelized source.

For projection 2, we see that for the smooth test case, false-positive signals show up in upper right regions with the highest ΔL to be ∼18. Read from the grid of highest ΔL, the best-fitting subhalo's mass is $10^{9.3^{+0.4}_{-0.5}}$ M$_\odot$. False positives were not detected for this projection using a pixelized source. For the case of an input subhalo of 5×10^8 M$_\odot$, there are some plausible signals around the middle-right regions with the highest ΔL to be ∼8. The mass of the plausible subhalo obtained in this case is $10^{9.6^{+0.2}_{-0.5}}$ M$_\odot$. For the third case where a 5×10^9 M$_\odot$ subhalo added, the ΔL map returns the

correct answer with the highest ΔL to be 120 at the place where we input the subhalo. For this one, similarly, we further model the subhalo by freeing its position. The posterior is shown in colour blue in Fig. A2. We see that the input subhalo can be well recovered in this case.

In Fig. A3, we show ΔL maps of modelling the data with the decomposed model plus a cored Sérsic source. We see that the results are similar to the pixelization results: For smooth tests, no clear false-positive signals show up. For subhalo tests, the highest ΔL is consistent with the region of an input subhalo. In Fig. A4, we further plot the posteriors obtained for the detected subhaloes. Overall, input subhaloes can be recovered to a good level although for the 5×10^9 M$_\odot$ subhalo cases, the recovered masses are slightly offset to the true value, albeit this is close enough that it could simply be due to noise in the mock observation.

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