

The COSMOS-Web Lens Survey (COWLS) I: discovery of >100 high redshift strong lenses in contiguous *JWST* imaging

James W. Nightingale^{1,★}, Guillaume Mahler^{2,3,4}, Jacqueline McCleary⁵, Qiuhan He^{6,4}, Natalie B. Hogg⁶, Aristeidis Amvrosiadis⁴, Ghassem Gozaliasl^{7,8}, Wilfried Mercier⁹, Diana Scognamiglio¹⁰, Edward Berman⁵, Gavin Leroy^{3,4}, Daizhong Liu¹¹, Richard J. Massey^{3,4}, Marko Shuntov^{12,13}, Maximilian von Wietersheim-Kramsta^{3,4}, Maximilien Franco¹⁴, Louise Paquereau¹⁵, Olivier Ilbert⁹, Natalie Allen^{12,13}, Sune Toft^{12,13}, Hollis B. Akins¹⁴, Caitlin M. Casey^{12,14,16}, Jeyhan S. Kartaltepe¹⁷, Anton M. Koekemoer¹⁸, Henry Joy McCracken¹⁵, Jason D. Rhodes¹⁰, Brant E. Robertson¹⁹, Nicole E. Drakos²⁰, Andreas L. Faisst²¹ and Shuowen Jin^{12,22}

Affiliations are listed at the end of the paper

Accepted 2025 July 22. Received 2025 July 22; in original form 2025 March 17

ABSTRACT

We present the COSMOS-Web Lens Survey (COWLS), a sample of over 100 strong lens candidates from the 0.54 deg² COSMOS-Web survey, discovered using exquisite *James Webb* Space Telescope (*JWST*) imaging across four wavebands. Following two rounds of visual inspection, over 100 candidates were ranked as ‘high confidence’ or ‘likely’ by at least 50 per cent of inspectors. The COWLS sample has several notable properties: (i) magnified source galaxies spanning redshifts $z \sim 0.1$ to $z \sim 9$, which therefore extend into the epoch of reionization; (ii) the highest-redshift lens galaxies known, pushing galaxy density profile evolution studies beyond $z \sim 2$; (iii) all lenses are distributed within a contiguous 0.54 deg² region, allowing for joint strong and weak lensing analyses; and (iv) a subset exhibits lensed source emission ray-traced near the lens galaxy centres, enabling studies of supermassive black holes and dust absorption. A key innovation of our approach is the use of lens modelling to aid in identifying lenses that may otherwise be missed. This paper is accompanied by the first COWLS public release, providing *JWST* NIRCcam imaging in four bands, lens models, pixelized source reconstructions and lens redshift estimates.

Key words: astroparticle physics – gravitational lensing: strong – dark matter.

1 INTRODUCTION

Galaxy-scale strong gravitational lensing, where a background source is multiply imaged by a foreground deflector galaxy, has facilitated numerous advancements in astrophysics and cosmology. This phenomenon has provided magnified views of distant galaxies that would otherwise be inaccessible (Vieira et al. 2013; Swinbank et al. 2015; Rizzo et al. 2020, 2021; Liu et al. 2024; Amvrosiadis et al. 2025), offered insights into the mass distribution of galaxies up to a redshift of $z \sim 2$ (Koopmans et al. 2009; Auger et al. 2010; Shajib et al. 2021; Etherington et al. 2023; Tan et al. 2024), and enabled the measurement of cosmological parameters (Collett & Auger 2014; Vegetti et al. 2014; Birrer et al. 2020; Hogg 2024). Our ability to leverage strong lensing as a means to understand the Universe is ultimately limited by the number of known lenses, how high in redshift their lens and source galaxies are and the quality of the imaging data.

Finding strong lenses begins by building a sample of candidate lenses through one of three selection techniques: (i) searching for arcs

or rings in large wide-field ground-based imaging surveys (Gavazzi et al. 2007; Sonnenfeld et al. 2013; Jacobs et al. 2019; Tran et al. 2022); (ii) detecting multiple emission lines at different redshifts in the spectra of galaxies (Bolton et al. 2006, 2008, 2012; Shu et al. 2016), or (iii) applying flux-density cuts on point-source catalogues from wide-area submillimetre surveys to isolate galaxies that are so bright they must have been magnified by lensing (Vieira et al. 2010; Negrello et al. 2014; Harrington et al. 2016). In all cases, the initial sample of galaxies inspected exceeds millions, to overcome the rarity of strong lensing occurring. For all three methods, confirming that candidates are strong lenses requires expensive high-resolution imaging follow-up, historically from the *Hubble* Space Telescope (*HST*). These collective efforts have led to the discovery of hundreds of galaxy-scale strong lenses.

The *James Webb* Space Telescope (*JWST*) offers a powerful approach to finding strong lenses through high-resolution imaging at near-infrared (NIR) and infrared (IR) wavelengths. With its 6.5-metre mirror, *JWST*’s increased depth observes a significant fraction of distant and faint galaxies in a single exposure, even with modest exposure times. Consequently, within the relatively small 2.2×2.2 arcmin field of view, there is still a high probability that an image will contain multiple detectable strong lenses (Holloway

* E-mail: James.Nightingale@newcastle.ac.uk

et al. 2023; Pearson et al. 2023; Casey et al. 2024). Any strong lens discovered this way comes with exquisite multiwavelength *JWST* imaging ‘for free’, which would be challenging and expensive to acquire through individual follow-up observations.

We therefore search for strong lenses in the *JWST* Cosmic Origins Survey (COSMOS-Web) survey (Casey et al. 2024), a 0.54 deg^2 field observed in 4 NIR and IR bands using *JWST*, which completed observations in 2024 May. Forecasts predict that the survey contains between 30 and 100 strong lenses (Holloway et al. 2023; Casey et al. 2024). COWLS Paper III, Hogg et al. (2025a, hereafter H25), uses the properties of the candidates found in this paper to infer the total number of lenses the COSMOS-Web field contains – inferring a value of 107, indicating our sample has high completeness. We use 4 wavebands of Near-Infrared Camera (NIRCam) imaging (F115W, F150W, F277W, and F444W), prioritizing this shorter wavelength data as it provides the highest spatial resolution and is able to resolve the lower Einstein radius strong lenses. Our lens search therefore omits the higher wavelength COSMOS-Web Mid-Infrared Imager (MIRI) data, which likely contains additional strong lenses (Pearson et al. 2023). However, MIRI data are available for many of the lens candidates we present.

We name this the COSMOS-Web Lens Survey (COWLS), which stands out among lens surveys due to several unique characteristics that shape its primary scientific goals. The exceptional survey depth allows the lensed source population to extend beyond reionization to $z > 6$, surpassing previous lens surveys, as shown in H25. These magnified sources provide an unprecedented view of distant galaxy structure and morphology. Lens modeling enables detailed pixel-level reconstructions of every unlensed source galaxy (Nightingale, Dye & Massey 2018), with reconstructions across all four NIR-Cam bands included in the public data release accompanying this paper.

The COWLS lens population is also distinct, featuring the highest-redshift lenses known (above $z \sim 2$), lower-mass systems, and rare morphologies such as disc galaxies. This enables studies of galaxy density profile evolution at higher redshifts (Shajib et al. 2021; Etherington et al. 2023; Tan et al. 2024) and makes strong lens dark matter detection studies more sensitive to line-of-sight haloes (Vegetti et al. 2014; Despali et al. 2018, 2022; Ritondale et al. 2019; He et al. 2022b; Nightingale et al. 2024). Unlike other lens surveys that span large sky areas, all COWLS candidates reside within a 0.54 deg^2 region with contiguous *JWST* imaging, providing a unique opportunity to study strong lenses within their broader cosmological environment (Peng et al. 2010). COSMOS-Web’s deep *JWST* imaging enables high-precision weak lensing analysis, and the dense concentration of strong lenses makes this dataset ideal for joint strong and weak lensing studies to measure cosmic shear (Birrer et al. 2017; Birrer, Refregier & Amara 2018; Fleury, Larena & Uzan 2021; Hogg et al. 2023; Duboscq et al. 2024; Hogg et al. 2025b), provided systematic uncertainties in strong lens shear measurements are addressed (Etherington et al. 2024). Additionally, we identify lenses where source emission passes within 0.2 arcsec of the lens centre, potentially probing the influence of supermassive black holes, as demonstrated in Abell 1201 where the source-lens separation was 0.38 arcsec (Nightingale et al. 2023b).

An important aspect of the lens-finding approach used in this study is the use of traditional lens modelling, via the open-source software PyAutoLens.¹ (Nightingale & Dye 2015; Nightingale et al. 2018, 2021b), to help determine whether candidates are genuinely strong

lenses. Only a small fraction of lens-finding studies have applied lens modelling in this manner (e.g. Sonnenfeld et al. 2018, 2020; Rojas et al. 2022), using low-resolution ground-based imaging which limits the information a lens model can extract. Acevedo Barroso et al. (2025) used lens modelling to help find lenses with high resolution *Euclid* early release observations, and we compare our approach to theirs. The lens modelling pipeline is fully automated after initial data preprocessing and has been successfully applied to simulated lenses (Cao et al. 2022; He et al. 2023) and *HST* data sets (Etherington et al. 2022, 2023; He et al. 2024; Nightingale et al. 2024). In this work, we fit over 400 candidates across all four imaging wavebands and illustrate multiple ways in which it allows us to find high quality lens candidates we would have otherwise missed.

This work was performed in tandem with COWLS Paper II, Mahler et al. (2025, hereafter M25), who report on the 17 most ‘spectacular’ strong lenses in COWLS, showing how *JWST*’s depth, resolution and NIR coverage makes a subset of strong lenses easily identifiable through basic visual inspection. COWLS Paper III, H25, compares the lens sample we found to a forecasting analysis, to address whether the lens sample is consistent with that expected based on known galaxy populations and cosmology.

The 1.64 deg^2 *HST* COSMOS field, which contains the 0.54 deg^2 COSMOS-Web field, has previously been searched for strong lenses, both via human visual inspection (Faure et al. 2008; Jackson 2008) and using machine learning (Pourrahmani, Nayyeri & Cooray 2018). There are 64 candidates presented in Faure et al. (2008), 112 in Jackson (2008), and 92 by Pourrahmani et al. (2018). Where COSMOS-Web overlaps with the *HST* COSMOS field, we compare candidates. Boosts in submillimetre counts and far IR emission of COSMOS data have also produced strong lens candidates (Aretxaga et al. 2011; Jin et al. 2018). The COSMOS field has also had individual strong lenses discovered and discussed, all of which are in our high ranking sample (Guzzo et al. 2007; More et al. 2012; van der Wel et al. 2013; Pearson et al. 2023; Jin et al. 2024). COWLS includes the COSMOS-Web ring (van Dokkum et al. 2024; Mercier et al. 2024), recently spectroscopically confirmed to be the highest redshift known galaxy-scale lens galaxy ($z = 2.02$) and have a $z = 5.10$ source galaxy (Shuntov et al. 2025b).

With the recent launch of *Euclid*, strong lenses are now being found with high resolution space based imaging Acevedo Barroso et al. (2025); Nagam et al. (2025); O’Riordan et al. (2025); Pearce-Casey et al. (2025); *Euclid* Collaboration et al. (2025a, b, c) spanning large patches of the Universe. The COWLS samples complements perfectly the tens of thousands of lenses *Euclid* is poised to discover, with the two samples essentially covering different regions of lens and source redshift parameter space.

Data for all lens candidates is publicly available², including fits data files, point spread functions, PYAUTOLENS modelling results, unlensed source reconstructions, redshift estimates and other metadata.

This paper is structured as follows. In Section 2, we describe the COSMOS-Web survey and *JWST* data reduction. In Section 3, we describe the PYAUTOLENS method and model fits performed in this work. In Section 4, we describe the visual inspection process used to find lens candidates. In Section 5, we present the lens candidates found after visual inspection and lens modelling. In Section 6, we discuss the implications of our measurements, and we give a summary in Section 7.

¹<https://github.com/Jammy2211/PyAutoLens>

²https://github.com/Jammy2211/COWLS_COSMOS_Web_Lens_Survey

2 DATA

2.1 NIRCам imaging and catalogue

We use *JWST* data from the cycle 1 GO program COSMOS-Web (PID 1727; PIs: Kartaltepe & Casey; Casey et al. 2024), a 255 h *JWST* treasury program. The full survey was completed in 2024 May and maps a 0.54 deg^2 area using NIRCам (we do not use the 0.19 deg^2 COSMOS-Web MIRI data which partially overlap the NIRCам footprint). We use the complete COSMOS-Web dataset including the NIRCам F115W, F150W, F277W, and F444W filters over the full area. It is the largest *JWST* programme both in terms of contiguous area covered at this depth and GO time allocated. The depths of the NIRCам data are measured to be 26.6–27.3 AB (F115W), 26.9–27.7 (F150W), 27.5–28.2 (F277W), and 27.5–28.2 (F444W) for 5σ point sources calculated within 0.15 arcsec radius apertures. Data for each waveband were reduced to two pixel scales, 0.03 and 0.06 arcsec . To minimize correlated noise, lens modelling uses data closest to NIRCам’s native pixel-scale, therefore 0.03 arcsec for F115W and F150W, 0.06 arcsec for F277W and F444W.

The COSMOS-Web team has constructed a model-based photometric catalogue combining ground and space-based data for broad general use; this catalogue is described in the forthcoming paper by Shuntov et al. (in preparation; see also Shuntov et al. 2025a). To briefly describe the catalogue construction, the four NIRCам imaging filters were PSF-homogenized to F444W’s PSF then combined in a $\sqrt{\chi^2}$ image (Szalay, Connolly & Szokoly 1999; Drlica-Wagner et al. 2018) truncated towards positive values only. The detection was performed on this $\sqrt{\chi^2}$ image following a hot/cold methodology (Koekemoer et al. 2011) using SEP (Barbary 2016). The hot/cold catalogue is then used as input to construct 2D Sérsic models for each galaxy, fit simultaneously to the four NIRCам filters using SOURCEEXTRACTOR++ (SE++ ; Bertin et al. 2020; Kümmel et al. 2020). These Sérsic subtracted images are used in the visual inspection campaign. Once the model shapes are derived, SE++ is used once more to extract photometry from 30+ bands across the field (ground and space) to produce the final photometric catalogue. Photometric redshifts and physical parameters are derived with LEPHARE (Arnouts et al. 1999; Ilbert et al. 2006). From these stellar masses and photometric redshifts, we then select candidate lenses for visual vetting. COSMOS-Web data were collected as ~ 304 fields of view (FOVs) across the whole COSMOS-Web field, taken together in 152 visits. Data reduction and cataloguing reduced and processed these as a single FOV in 2023 January, 10 FOVs in 2023 April, and 10 FOVs in 2024 January, producing 21 FOVs in total.

2.2 Point spread function

Point spread functions (PSFs) describe the impulse response of an optical system like a telescope to light. Because the distortion produced by the PSF mimics and dilutes the gravitational lensing shear of interest, the PSF must be modelled and mitigated before further analysis. The NIRCам PSF varies with time, bandpass, and across the field of view and moreover is non-trivially affected by the mosaicing procedure. Accordingly, we create PSF models empirically from the observations themselves, using stars as reference points with which to build the model.

We begin by using SourceExtractor (Bertin et al. 2002) to generate astronomical source catalogues on the *i2d*-format resampled mosaics produced by the *JWST* Science Calibration Pipeline, using the

WHT extension as a weight image and a 3×3 Gaussian convolution kernel with a two-pixel standard deviation.

A preliminary star sample is created from this catalogue using size (specifically, FLUX_RADIUS) and magnitude cuts ($19 < \text{MAG_AUTO} < 25$) designed to isolate the stellar locus in a size–magnitude diagram. This initial sample is scrubbed of stars with saturated or otherwise bad pixels indicated by sentinel value of 0 in the ERR extension of the *i2d* mosaics. We then pass this star catalogue sample to PSFEX to obtain PSF fits (Bertin 2011). We use the PIXEL basis, fix the PSFEX sampling step to 0.5, and fit first-order polynomials in (X, Y) , which we find sufficient to capture the spatial variation across the FOV of a single resampled mosaic. Once models are obtained, the PSFEX³PYTHON package is used to create a rendering of the PSF at the location of lens galaxies, which can then be passed to PYAUTOLENS for further analysis. This procedure is repeated in all four bandpasses for each of the 21 NIRCам tiles.

2.3 Redshifts

The redshift of the lens is a crucial piece of information for strong lensing studies. We collected both spectroscopic and photometric redshift measurements for the foreground lensing galaxies in our sample. For spectroscopic redshifts, we cross-matched our candidates with the DR10 catalogue of the DESI Legacy Imaging Surveys (Dey et al. 2019). We found that approximately one-third of our candidates have spectroscopic redshifts from previous surveys. For photometric redshifts, we used the COSMOS-Web survey catalogue, where redshifts were derived from SED fitting to multi-band images using LEPHARE (Arnouts et al. 1999; Ilbert et al. 2006). We show and compare these redshifts in the results section after presenting our lens candidate sample.

3 LENS MODELLING

3.1 Overview

We use version 2024.11.13.2 of the lens modelling software PYAUTOLENS⁴ (Nightingale et al. 2021c). PYAUTOLENS fits the lens galaxy’s light and mass and the source galaxy simultaneously. The method assumes a model for the lens’s foreground light (a multi-Gaussian expansion He et al. 2024), which accounts for blurring by the instrumental PSF and is subtracted from the observed image. A mass model (an isothermal mass distribution) ray-traces image-pixels from the image-plane to the source-plane and a pixelized source reconstruction, using an adaptive Voronoi mesh, is performed.

Fig. 1 shows an overview of a PYAUTOLENS lens model using the COSMOS-Web ring (Dokkum et al. 2024; Mercier et al. 2024), also one of M25’s ‘spectacular’ lenses, where models for the image-plane lens galaxy emission, lensed source, source-plane source reconstruction and mass model convergence are shown. This image is used to help inspectors determine if a candidate is a strong lens, as discussed in Section 4.3.

We now describe each step in more detail. The following link contains Jupyter notebooks providing a visual step-by-step guide of the PYAUTOLENS likelihood function used in this work⁵:

³<https://github.com/esheldon/psfex?tab=readme-ov-file>

⁴<https://github.com/Jammy2211/PyAutoLens>

⁵https://github.com/Jammy2211/autolens_likelihood_function




Figure 1. An overview of the lens modelling performed to aid the inspection of strong lens candidates, for the confirmed strong lens COSJ100024+015334, also named the COSMOS-Web ring (Dokkum et al. 2024; Mercier et al. 2024; Shuntov et al. 2025b). Each row shows the imaging data for the F115W, F150W, F277W, and F444W wavebands. From left to right the columns show: (i) the imaging data with a $\log_{10}$ colour scale; (ii) the data where a Multi-Gaussian Expansion (MGE) model subtracts the foreground lens emission; (iii) the image-plane lensed source model; (iv) the source-plane source reconstruction using a Voronoi mesh; (v) the MGE lens light model on a $\log_{10}$ colour scale and; (vi) the inferred lens model convergence. In Section 5.4, we will discuss in detail how this image enables high confidence strong lens candidates to be found.

3.1.1 Lens light subtraction

To subtract the foreground emission of the lens galaxy we use a Multi-Gaussian Expansion (Cappellari 2002, hereafter MGE). This decomposes a galaxy’s emission into a set of 2D elliptical Gaussian light profiles, unlike many lensing studies which use an elliptical Sérsic profile. The MGE is implemented within the semi-linear inversion framework for lens modelling (Warren & Dye 2003), with the implementation described fully in He et al. (2024, hereafter H24). Its enhanced flexibility for fitting galaxy emission is key to subtracting the complex morphological features seen in lens galaxies when they are observed with deep space based observations (e.g. *HST* or *JWST*).

The MGE decomposes the light of a galaxy into sets of Gaussians whose overall intensity is given as

$$I_{\text{set}}(x, y) = \sum_i^N G_i(x, y), \quad (1)$$

with G_i the i th Gaussian profile,

$$G_i(x, y) = I_i \cdot \exp\left(-\frac{R_i^2(x, y)}{2\sigma_i^2}\right), \quad (2)$$

containing I_i as the intensity normalization factor and σ_i as the full-width at half-maximum of the Gaussian profile. $R_i(x, y)$ is the Gaussian elliptical radius and given as

$$R_i(x, y) = \sqrt{x'^2 + \left(\frac{y'}{q_i}\right)^2} \quad (3)$$

$$x' = \cos \phi_i \cdot (x - x_i^c) + \sin \phi_i \cdot (y - y_i^c)$$

$$y' = \cos \phi_i \cdot (y - y_i^c) - \sin \phi_i \cdot (x - x_i^c),$$

where q_i is the axis ratio, ϕ_i the position angle, and (x_i^c, y_i^c) is the centre of the Gaussian. The intensity of each Gaussian is solved for linearly, using the semi-linear inversion formalism for lens analysis (Warren & Dye 2003). We enforce positivity on the solution using a modified version of the fast non-negative least-square (fnnls) algorithm (Bro & De Jong 1997),⁶ because H24 show that allowing for negative intensities produces unphysical solutions whereby the Gaussians alternate between large positive and negative values.

Gaussians are grouped in ‘sets’, where a set of Gaussians share the same centres, position angles, axis ratios, and their σ values

⁶The fnnls code we are using is modified from <https://github.com/jvendrow/fnnls>.




Figure 2. The foreground lens light subtracted images performed using a Multi-Gaussian Expansion (MGE) He et al. 2024 and no lens mass modelling or source reconstruction, provided to inspectors in the second round of visual inspection. This figure shows an example using the lens COSJ095921 + 020638. Each row shows the imaging data for the F115W, F150W, F277W and F444W wavebands. The first and second columns show the data before the MGE subtraction with a $\log_{10}$ and linear colourmap, to highlight the lens galaxy in different ways. The remaining four columns show the lensed source with different colour scalings applied, to highlight the lensed source emission in different ways and aid the visual inspection. By performing the MGE subtraction independently of the lens model, an inspector can assert if features in the lensed source are genuinely present in the data as opposed to an artefact of lens modelling. Not including the source in the model can produce residuals in the centre of the lens subtraction, most clearly seen above for the F150W and F277W filters, which the full lens modelling pipeline fits successfully. The right hand column also shows white patches where the $\log_{10}$ colour scale cuts out negative values. This is an intentional choice to provide different information on the candidate source emission.

are fixed to preset values which evenly increase in $\log_{10}$ -spaced intervals between one-fifth of the image pixel scale and the circular radius of the mask applied to the lens image. We use two sets of 30 Gaussians for every lens galaxy, where all 60 Gaussians share the same centre. We include an additional set of 10 Gaussians to model central point-like emission, where the σ values of these Gaussians are evenly $\log_{10}$ -spaced between 0.01 arcsec and twice the pixel scale of the data. The centre of these 10 Gaussians is shared with the other 60. The grouping of Gaussians into sets is done to reduce the number of non-linear free parameters in the MGE, which for this study consists of just $N = 6$ free parameters for the lens light (see H24).

Each candidate strong lens is fitted twice across all four wavebands using an MGE for the lens light, under two different assumptions. First, an MGE fit is performed without a lens mass model or lensed source model. This produces a lens-light subtracted image used for preprocessing before lens modelling and visual inspection. Fig. 2 shows an example of such an image, with different colour scales highlighting the lens and source emission. Secondly, an MGE fit is performed simultaneously with the lens mass model and source model, as shown in Fig. 1.

3.2 Mass profiles

For the lens mass model (superscript ‘mass’), we first assume a singular isothermal ellipsoid (SIE) density profile representing the total mass of the lens (e.g. stars and dark matter) of the form

$$\kappa(\xi) = \frac{1}{1 + q_{\text{mass}}} \left(\frac{\theta_{\text{E}}^{\text{mass}}}{\xi} \right)^2, \quad (4)$$

where $\theta_{\text{E}}^{\text{mass}}$ is the model Einstein radius in arc-seconds. Deflection angles are computed via an implementation of the method of Tessore & Metcalf (2015) in PYAUTOLENS.

An external shear (superscript ‘ext’) field is included and parametrized as two elliptical components ($\gamma_1^{\text{ext}}, \gamma_2^{\text{ext}}$). The shear magnitude, γ^{ext} , and orientation measured counter-clockwise from north, ϕ^{ext} , are given by

$$\gamma^{\text{ext}} = \sqrt{\gamma_1^{\text{ext}2} + \gamma_2^{\text{ext}2}}, \quad \tan 2\phi^{\text{ext}} = \frac{\gamma_2^{\text{ext}}}{\gamma_1^{\text{ext}}}. \quad (5)$$

The deflection angles due to the external shear are computed analytically. A recent study by Etherington et al. (2024) suggests

that this external shear component represents missing complexity in the lens mass distribution.

As an alternative, we also perform fits using an MGE mass model, where the convergence of each Gaussian is related to its intensity as

$$\kappa_i(\xi) = \Psi G_i(x, y), \quad (6)$$

where Ψ is the mass-to-light ratio of the Gaussian and is fixed to the same value across all Gaussians. Every other parameter of each Gaussian in the MGE mass model (I, q, ϕ, σ) comes from the lens light fit. Deflection angles are computed following Shajib (2019). We do not include a dark matter component. The MGE mass model has been used to perform lens mass modelling in a small number of previous studies (e.g. Collett et al. 2018; Melo-Carneiro, Furlanetto & Chies-Santos 2023), and in all cases, it was combined with stellar dynamics, which are not used in this study.

The SIE mass model is independent of the MGE lens light model, such that it may infer solutions where the mass and light are geometrically offset and misaligned. In contrast, the MGE mass model is tied to the lens light MGE model, meaning that the stellar mass must be geometrically aligned with the lens light. The two mass models therefore provide different and complementary information, to help the visual inspectors determine if a candidate is a genuine strong lens.

We fit an MGE mass model without an explicit dark matter component (e.g. a Navarro–Frenk–White (NFW) profile; Navarro, Frenk & White 1996). Our intended use of the MGE mass model is to illustrate whether a mass model which ties lens light to lens mass can still produce a reliable lens model fit. When a dark matter component is included, solutions are often inferred where the dark matter becomes the dominant mass component, meaning that this model can, analogous to the SIE, infer solutions where the lens’s mass does not trace its light. We verified that the MGE-only mass model provides physically plausible lens model solutions to the spectacular lenses found in M25, and therefore verified whilst this assumption may not be fully realistic it is sufficient to add information that aids inspectors in judging strong lens candidates.

3.3 Source model

After subtracting the foreground lens emission and ray-tracing coordinates to the source-plane via the mass model, the source is reconstructed in the source-plane. The first stage of our lens modelling pipeline uses an MGE to model the light of the lensed source. As discussed in H24, the MGE source model is an effective tool for automated and efficient lens modelling because it can fit the data in a highly flexible way whilst retaining a relatively low number of non-linear free parameters. For the source, we use only one set of 30 Gaussians, whose σ values span $\log_{10}$ increments from 0.001'' to 1.0''.

However, an MGE source galaxy model is still symmetric around the centre of the Gaussians and thus cannot fit the complex and irregular source morphologies typical of high redshift galaxies. The later stages of our lens modelling pipeline therefore use an adaptive Voronoi mesh which can reconstruct irregular and asymmetric source morphologies. A complete description of the process, including the linear algebra formalism, interpolation scheme and regularisation, is given in Nightingale et al. (2024) and appendix A of H24. An example source reconstruction using this Voronoi mesh is shown in the fourth column of Fig. 1. The Voronoi mesh uses natural neighbour interpolation (Sibson 1981) and the cross-like regularisation scheme introduced in H24.

Table 1. Pipeline composition used in the analysis built using PyAutoLens.

Pipeline	Component	Model	Prior info
Source	Lens light	MGE	–
Parametric (SP)	Lens mass	SIE + Shear	–
	Source light	MGE	–
Source	Lens light	MGE	SP (fixed)
Pixelization 1 (SPix1)	Lens mass	SIE + Shear	SP
	Source light	MPR	–
Source	Lens light	MGE	SP (fixed)
Pixelization 2 (SPix2)	Lens mass	SIE + Shear	–
	Source light	Voronoi	–
Mass1 (MP1)	Lens light	MGE	SP1
	Lens mass	SIE	SP1
	Source light	Voronoi	SP2
Mass2 (MP2)	Lens light	MGE	SP1
	Lens mass	MGE	–
	Source light	Voronoi	SP2

3.4 Lens modelling pipeline

Table 1 outlines our automated lens modelling pipeline, which iteratively fits various combinations of light, mass, and source models. The pipeline chains together five lens model fits. The pipeline initially fits a simpler model using an MGE source for efficient and robust convergence towards accurate results. Subsequent stages employ the more complex Voronoi source reconstruction. The pipeline performs five chained fits, the first two and final two use the nested sampler NAUTILUS⁷ (Lange 2023) and the third stage uses the nested sampler dynesty (Speagle 2020a). The pipeline is based on the PYAUTOLENS SLAM (Source, Light, and Mass) pipelines used by various other studies (e.g. Cao et al. 2022; Etherington et al. 2022; He et al. 2022a; Nightingale et al. 2023b, 2024). A break down of the different fits are as follows:

Source parametric (SP): The SP pipeline computes an accurate initial estimate of the lens model parameters. The lens and source light are both modelled using a MGE (see Section 3.1.1) and the lens mass uses an SIE mass profile with shear (see Section 3.2). The non-linear parameter space has $N = 17$ free parameters.

Source pixelized 1 (SPix1): The SPix1 pipeline uses a Voronoi source reconstruction, where the regularization coefficients of the source are free. The lens light uses an MGE, where the centres, axis-ratios and position angles of all Gaussians are fixed to the results of the SP pipeline, but their intensities are linearly solved simultaneously with the source. The mass model is an SIE plus shear where all parameters are free (with priors based on the inferred model from SP).

Source pixelized 2 (SPix2): The lens light MGE and mass model parameters are fixed to the maximum likelihood values of SPix1. The only free parameters are associated with the Voronoi source, which adapt to the source’s unlensed morphology (see Nightingale et al. 2018) to enable a sharper and more detailed reconstruction of a compact and complex source.

Mass pipeline 1 (MP1): The lens light MGE parameters are fixed to the maximum likelihood values of SPix1 and Voronoi source parameters fixed to SPix2. The only free parameters are the SIE mass model and external shear parameters.

⁷<https://github.com/johannesulf/nautilus>

Mass pipeline 2 (MP2): Identical to MP1, but instead of an SIE plus shear the only free parameters are the MGE mass-to-light ratio Ψ and external shear parameters.

3.5 Multi-wavelength Fits

For results used during visual inspection, the lens modelling pipeline was run independently on imaging from all four wavebands. As a result, the mass model fitted to each waveband is independent, explaining the slight differences in convergence maps in the right-hand column of Fig. 1. Since lensing is achromatic, mass variations across wavelength are unphysical. This was intentional, as it allows visual inspectors to assess the consistency of mass models across wavebands, providing more information to judge whether a candidate is a lens. The role of this strategy in identifying candidates will be discussed in Section 5.

A drawback of fitting each waveband independently is that the source reconstruction also varies across wavelengths, complicating studies of the magnified source. To address this in the COWLS public data release, all candidates were re-fitted after visual inspection using a pipeline that constrained the mass model to a single waveband – the ‘primary waveband’ – selected by JWN based on the clarity of the candidate’s source emission in the initial fit. This mass model was then applied to other wavebands, with the source regularisation and MGE lens light refitted as free parameters, alongside two additional parameters accounting for arc-second filter offsets, ensuring alignment across reconstructions. These aligned models are included in the first COWLS public data release.

When quoting or plotting lens mass model values (e.g. Einstein radius), we refer to those inferred from the primary waveband fit. Magnitudes are computed using models that fit each waveband independently. Lens light magnitudes are derived from the total flux of each MGE lens light model, lensed source magnitudes from the total flux of the source emission in the image plane, and source magnitudes from the sum of all source-plane Voronoi mesh pixels. Magnifications are calculated as the ratio of the total lensed source flux in the image plane to the total source flux in the source plane.

3.6 User inputs

Fig. 3 illustrates the manual user inputs required before processing by the lens modelling pipeline. All steps are performed manually by JWN, unlike visual inspection tasks which are conducted by multiple inspectors. JWN’s subjective determination of what constitutes a candidate’s lensed source emission may therefore impact the lens model visuals used later by inspectors to rank lenses. These pre-processing steps are performed using images across all four wavebands in the following order:

- (i) **Mask:** A circular mask is applied to the imaging dataset, such that only image pixels within this mask are fitted by the lens model. Fig. 3 shows an example 3.5 arcsec circular mask with a black circle. The mask’s size is chosen to ensure it contains the candidate lens galaxy’s emission and all lensed source multiple images.
- (ii) **Contaminant removal:** Emission from nearby line-of-sight galaxies expected to not be directly associated with the candidate strong lens may fall within the circular mask. If this emission is not removed, the lens model may incorrectly fit it as part of the lensed source. A GUI is therefore used to ‘spray-paint’ any regions of the image with this contaminant emission. The GUI replaces the emission with zeros and increases the RMS noise map to such large




Figure 3. Illustration of preprocessing steps applied to each lens before modelling, using the lens COSJ100013 + 023424. The left panel shows an RGB image, where colours help distinguish the candidate lensed source (e.g. red arcs) from the foreground lens. The right panel displays the processed F444W image, with the following steps applied: (i) A circular mask (black) is set to encompass the lens and lensed source, in this example of radius 3.5 arcsec. (ii) A GUI ‘spray-paint’ tool removes light from nearby galaxies by setting their values to zero and inflating RMS noise to prevent them from affecting the fit; (iii) The foreground lens light inside this mask is subtracted using an MGE-only model to enhance source visibility independently for all four wavebands and; (iv) a GUI marks the (x, y) coordinates of multiple lensed images (black crosses), which during lens modelling must map within $0.1''$ in the source plane; otherwise, a penalty term is applied to the likelihood, removing unphysical solutions (Maresca, Dye & Li 2021).

values that the lens model fit ignores it. This technique is commonly applied in lens modelling (e.g. Etherington et al. 2022; Nightingale et al. 2024).

(iii) **MGE lens subtraction:** The MGE-only lens light fit described in Section 3.1.1 is performed on all four wavebands of data using this mask and with the contaminants removed. This foreground subtracted image is used to make candidate lensed source multiple images more clearly visible.

(iv) **Multiple image positions:** Using a graphical user interface (GUI), the (x, y) coordinates of each multiple image of the candidate lensed source are input, shown by black crosses in Fig. 3. Each fit of the lens modelling pipeline requires that the mass model traces all multiple images within a 0.1 arcsec threshold of one in the source plane, or else a penalty term is applied to the likelihood, down-weighting that solution. This technique is commonly used in lens modelling to remove unphysical solutions, as described in Maresca et al. (2021).

4 VISUAL INSPECTION

4.1 First round

The COSMOS-Web survey contains millions of galaxies, of which only the most massive are anticipated to act as strong gravitational lenses. The first round of visual inspection therefore isolated the 42 660 most massive galaxies in the survey, using luminosity cuts. This used a magnitude cut at $m_{AB} = 23$ in the F277W filter and required a stellarity below 1, the latter removing stars which made up approximately 16 per cent of objects brighter than 23 mag. Based on SED fitting, a magnitude cut of 23 in the F277W band represents a median and 68th percentile stellar mass of $(3.8 \pm 3.3) \times 10^8 M_{\odot}$, $(2.9 \pm 2.6) \times 10^9 M_{\odot}$, and $(6 \pm 6) \times 10^9 M_{\odot}$ at $z = 0.3 \pm 0.05$, $z = 1.0 \pm 0.1$, and $z = 2.0 \pm 0.1$, respectively.

Visual inspection of the sample was conducted by multiple inspectors using images similar to those in Fig. 4. The first column




Figure 4. Example image shown to inspectors during the first round of visual inspection, featuring the confirmed strong lens COSJ100024 + 015334, also named the COSMOS-Web ring (Dokkum et al. 2024; Mercier et al. 2024; Shuntov et al. 2025b). The left column shows an RGB cutout centred on the candidate lens. The central column shows the same cutout after subtracting one or more Sérsic profiles using the COSMOS-Web photometry pipeline. The right column displays the model Sérsic light profiles used for this subtraction. The top and bottom rows use different RGB weightings to emphasise red and blue emission. Inspectors reviewed similar images for all candidates and classified them as either ‘Y’ (Yes, this is a lens), ‘M’ (Maybe, this is a lens), or ‘N’ (No, this is not a lens).

shows a postage stamp cutout of each galaxy, where RGB colour scaling helps distinguish red and blue emission. The second column presents the same image after subtracting Sérsic light profiles fitted to the central galaxy and nearby sources as part of the COSMOS-Web catalogue (Shuntov et al. 2025a). The third column displays the Sérsic model images used for this subtraction. The top and bottom rows apply different RGB scaling using different sets of *JWST* and NIRC*am* filters, to highlight red and blue emission. In this example, the Sérsic photometry successfully removed the central lens galaxy and two of the lensed source’s multiple images. Inspectors were advised to account for such subtractions when assessing whether a candidate was a strong lens.

Inspectors evaluated all 42 660 galaxies using this image format, selecting one of three options via a GUI: ‘Y–Yes, this is a lens’, ‘M–Maybe this is a lens’, or ‘N–No, this is not a lens’. An optional text box allowed for comments, and inspectors could revisit candidates later. Given the large number of galaxies, the GUI and ranking system enabled responses at a cadence of approximately one second per object. For brevity, we refer to these responses as ‘Yes’, ‘Maybe’, and ‘No’ throughout the paper. Visual inspection was paired with the 21 FOVs described in Section 2.1, with the 2023 January data was inspected from May 2023 to June 2023, the 2023 April data from August 2023 to September 2023, and the January 2024 data from April 2024 to May 2024.

4.2 Spectacular lenses

After the first round of visual inspection, 32 objects were marked as ‘Yes’ by over 50 per cent of inspectors. These objects are analysed in M25, where inspectors re-evaluated the cut-outs in Fig. 4 to identify those considered ‘spectacular’ lenses. These objects are included in our lens modelling analysis but are not forwarded to the second round of inspection and are labelled ‘M25’ throughout this work. The remaining 15 objects, not marked as spectacular in M25, are still potential lenses and proceed to the second round of inspection.

4.3 Second round

Objects advancing to the second round of visual inspection were selected based on two criteria: (i) those where at least 50 per cent of inspectors responded ‘Yes’ or ‘Maybe’, and (ii) ‘edge case’ objects that fell below this threshold but received at least one ‘Yes’ or ‘Maybe’, where JWN reviewed all qualifying objects and advanced any deemed strong candidates for further inspection.

All candidates meeting either selection criterion were fitted with an MGE-only foreground subtraction across all wavebands (Section 3.1.1) and had independent lens models fitted in each waveband (Section 3.4). The resulting images were provided to inspectors in the second round of visual inspection, viewed in the following sequence:

- (i) The initial cut-out image from the first round of visual inspection (see Fig. 4).
- (ii) Plots of the candidate lens and source in all four wavebands, displayed with linear and base-10 logarithmic colour scales, as well as RGB images.
- (iii) The MGE-subtracted image (see Fig. 2), where different colour scales highlight the lensed source emission in each waveband.
- (iv) The PYAUTOLENS model-fit to all four wavebands (see Fig. 1), including model images of the lens galaxy, lensed source, source-plane reconstruction, and mass model convergence. Separate figures show fits using both SIE and MGE mass models.

Additional images were available in a separate folder, allowing for detailed inspections of specific candidates. These included images where the image-source plane mappings were visualised using coloured polygons, to assess model consistency. Inspectors did not often use these extra images during the review process, with them only required for detailed inspection of select candidates.

Inspectors ranked each lens into one of the following five categories:

- (i) **A:** High confidence this is a strong lens.
- (ii) **B:** Likely a strong lens, but there is ambiguity.
- (iii) **U:** Unlikely to be a strong lens, but not impossible.
- (iv) **S:** A singly imaged strong lens feature/arc (e.g. without an observed counter image).
- (v) **X:** Not a lens.
- (vi) **I (Optional and independent):** The candidate is interesting and warrants follow up investigation.

Each candidate is assigned a final score based on their ranking, with 2 points for an A, 1 point for a B, and 1 point for an S. For example, a candidate ranked ‘AAABSX’ would score 8 points. All candidates in the second round of visual inspection were assessed by 6 inspectors, giving them a total score out of 12. This score is displayed alongside all candidate visualisations and listings in tables. Plots are colour-coded based on the scores: black for the M25 sample, dark green for scores above 7, yellow for scores of 5 or 6, orange for scores of 3 or 4, and maroon for scores below 3. Some plots are ordered by descending score. The second round of visual inspection took place from August 2nd to August 30th, 2024.

5 RESULTS

5.1 Visual inspection

5.1.1 First round

The results of the first round of visual inspection are provided in Table 2. In total, 42 660 images were inspected by 4 or 5 inspectors. 32 objects had at least 50 per cent of inspectors rank them in the ‘Yes’

Table 2. The results of the first round of visual inspection. The first column lists the time of acquisition of observations the results correspond to, with the bottom providing total values for each group. The second column shows the number of inspectors for that data set group. The third column shows how many objects are in this dataset group because they met the luminosity selection criteria described in Section 4.1. The fourth column lists the number of objects which are in the M25 spectacular lens sample. Column 5 shows how many objects at least 50 per cent of visual inspections indicated ‘Y–Yes, this is a lens’ which were not included in M25. The sixth column lists the number of objects for which at least 50 per cent of visual inspections indicated ‘Y–Yes, this is a lens’ or ‘M–Maybe this is a lens’ (including the 15 objects in column 5), meaning they were forwarded to the second round of visual inspection. The seventh column lists how many objects at least one inspector identified as a potential lens but less than 50 per cent of inspectors overall did, meaning that JWN reinspected it to determine if it is an edge case that should be forwarded to the second round of visual inspection. The final column lists how many of these objects JWN forwarded as edge cases. Bottom row values in bold are the candidate numbers which made it through to the second round of visual inspection. The percentages in brackets for each row are relative to the total number of objects in that dataset group, (column 3).

Data set group	Total inspectors	Total objects	M25 Spectacular lenses	At least 50 % Yes (Excluding M25)	At least 50 % Maybe (Excluding M25)	Over 1 Yes or Maybe and less than 50 % Yes or Maybe	Edge cases
2023 Jan	4	847	0 (0.0 %)	0 (0.0 %)	8 (0.94 %)	101 (11.92 %)	10 (1.18 %)
2023 April	4	14 688	8 (0.05 %)	6 (0.05 %)	152 (1.03 %)	1120 (7.62 %)	48 (0.33 %)
Jan 2024	5	27 125	9 (0.03 %)	9 (0.03 %)	105 (0.39 %)	2476 (9.13 %)	96 (0.35 %)
Total	4 or 5	42 660	17 (0.04 %)	15 (0.04 %)	265 (0.62 %)	3697 (11.53 %)	154 (0.40 %)

category, which were inspected further in M25. Of these 32 objects, M25 presented 17 objects as spectacular lenses, 0.04 per cent of the total number of galaxies inspected. A total of 419 objects made it through to the second round of visual inspection (excluding the 17 M25 lenses). 265 candidates met the criteria that at least 50 per cent of inspectors ranked them as a ‘Yes’ or ‘Maybe’. There are 3697 objects that at least one inspector said ‘Yes’ or ‘Maybe’ but not 50 per cent of inspectors. JWN inspected all 3697 objects and decided to put forward 154 as edge cases. In Section A, we discuss the first round of visual inspection results in more detail.

5.1.2 Second round

Table 3 presents the scores for the 419 candidates inspected in the second round of visual inspection. Two candidates received the highest possible score of 12, indicating unanimous agreement from all inspectors that they are strong lenses. The third column shows the cumulative total of candidates with scores above certain thresholds, for example, 17 candidates scored 10 or higher, 62 scored 7 or higher, and 143 scored 5 or higher. The final two columns of Table 3 indicate whether candidates came from the first round of visual inspection as a candidate or an edge case. Edge cases are common; for example, of the 143 candidates scoring 5 or higher, 55 (38.4 per cent) were edge cases, meaning only one inspector initially flagged them as a lens or a possible lens in the first round of visual inspection. A significant portion of high-scoring candidates would have been missed if edge cases had not been carried forward. The high scores of these edge cases suggest that lens modelling provided additional information, allowing inspectors to rank them higher, as discussed in Section 5.4.

5.2 Candidate images

Fig. 5 shows RGB colour composite images and foreground galaxy subtracted images inferred via lens modelling for the M25 spectacular lenses and next 21 highest ranked candidates. Single wavelength images which most clearly show the candidate lensed source emission are shown. For conciseness, only two postage stamp cut-out images are shown for each candidate; however, inspectors had access to more information when grading these lenses, including data from all four wavelengths and lens models. As a result, some objects with high scores might not appear to be convincing strong lenses based

Table 3. Candidate scores for the second round of visual inspection. The first column lists the score out of 12, as described in Section 4.3. The second, third, and fourth columns show the total number of objects for each score, the cumulative sum of candidates with scores equal to or above that score, and the percentage of the total 419 second-round candidates corresponding to each score. The fifth column indicates how many candidates reached the second round as a result of the first round of visual inspection, meaning at least 50 per cent of inspectors flagged them as ‘Y–Yes, this is a lens’ or ‘M–Maybe this is a lens’. The sixth column shows the number of edge cases, where less than 50 per cent of inspectors flagged the object but at least one identified it as a potential lens. These edge cases were reinspected by JWN and forwarded for further evaluation.

Score	Total Candidates	Cumulative	Percent	Round 1 Candidate	Round 1 Edge
S12	2	2	0.48 %	1	1
S11	3	5	0.72 %	3	0
S10	12	17	2.86 %	9	3
S09	11	28	2.63 %	9	2
S08	14	42	3.34 %	9	5
S07	20	62	4.77 %	11	9
S06	37	99	8.83 %	19	18
S05	44	143	10.5 %	27	17
S04	58	201	13.84 %	34	24
S03	62	263	14.8 %	36	26
S02	66	329	15.75 %	45	21
S01	46	375	10.98 %	32	14
S00	44	419	10.5 %	30	14

solely on these single cut-outs but are more persuasive when viewed with all the available data, and vice versa. All images used for the second round of inspection can be accessed at the following URL⁸.

Fig. 5 shows the M25 lenses followed by the candidates in descending order of score, from the top left, with score decreasing right and then down. Readers can therefore scan across and down this figure, and assess how the quality of candidates changes with decreasing score. The majority of candidates show features distinctive of strong lensing, for example emission from a central lens galaxy, multiple images in locations consistent with a strong lens model and arcs tangentially stretched relative to the lens galaxy centre. However, especially

⁸https://github.com/Jammy2211/COWLS_COSMOS_Web_Lens_Survey




Figure 5. RGB composite images showing the lens and source emission, and a foreground-lens-subtracted image highlighting the lensed source, of the highest ranking COWLS candidates. The first 15 lenses belong to the COWLS Paper II spectacular lens sample of *M25*. Two *M25* lenses, COSJ100024+015334 and COSJ095921+020638, are excluded as they appear in Figs 1 and 2, respectively. The 143 highest-ranked candidates from the second round of visual inspection (all scoring 5 or above) are then shown. This indicates most inspectors selected ‘A: High confidence this is a strong lens’ or ‘B: Likely a strong lens, but there is ambiguity’. Images are arranged in descending order of rank, starting from the top left, with a maximum score of 12 if all six inspectors gave an ‘A’ grade. Inspectors had access to all four wavelengths and lens models during grading, making some candidates more convincing with full data.




Figure 6. Spectroscopic and photometric redshift measurements for candidate lens galaxy emission in the COWLS sample are shown here. The full source of redshift information is described in Section 2.3, which includes DESI DR1 spectroscopic data and photometric redshift estimates from COSMOS archive ground-based and *HST* samples. The left panel compares spectroscopic and photometric redshift estimates for 243 candidates where both measurements are available. The near one-to-one correspondence reaffirms that the majority of photometric redshift measurements are reliable. The central panel shows photometric redshift estimates for 409 candidates where available, plotted as a function of their second-round visual inspection score. The right panel shows spectroscopic redshifts for 245 candidates with available measurements. Scores are colour-coded: high-ranked systems in green, mid-ranked in orange, and low-ranked in red (see Section 4.3). Spectroscopic data is unavailable above redshifts of $z \sim 2$, while photometric estimates extend to redshifts around $z \sim 4$. There is no correlation between candidate score and redshift.

for the lower scoring candidates, many of these features are hard to discern from complex and irregular morphological structures often seen in deep imaging of high redshift galaxies (e.g. spiral arms, rings).

5.3 The COWLS sample properties

5.3.1 Lens redshifts

Photometric and spectroscopic redshift estimates for candidate lens galaxies are available in the COSMOS archive (see Section 2.3), and these are shown for all candidates from the second round of visual inspection in Fig. 6. The left panel compares photometric and spectroscopic estimates for 243 candidates where both measurements are available. The near one-to-one correspondence between the two confirms the majority of photometric redshift estimates are accurate, particularly up to $z \sim 2$, where spectroscopic data is accessible. Beyond this redshift, spectroscopic catalogues (e.g. DESI, see Section 2.3) lack data. Overall, the sample has reliable redshift information for most of its candidate lenses.

The central and right panels show the relationship between photometric and spectroscopic redshifts and second-round scores. Based on photometric redshifts, about half of the sample's lens galaxies are above $z = 1$, with some candidates exceeding $z = 2$. There is no observable trend between candidate score and redshift estimate, indicating that high-ranking candidates are not more likely to be more distant lens galaxies.

5.3.2 Weak lensing mass map

Fig. 7 shows the distribution of strong lens candidates scoring five and above, plotted over the *HST* weak lensing convergence map inferred by Massey et al. (2007). Higher convergences indicate denser regions of dark matter, with two high-density peaks visible to the north of the COSMOS-Web field, along with filamentary structures connecting them. The red rectangle outlines the 0.54 deg² COSMOS-Web footprint, where *JWST* imaging is available, while weak lensing analysis outside this rectangle uses *HST* imaging. Our strong lens




Figure 7. The weak lensing convergence estimated via the *Hubble Space Telescope* weak lensing mass map analysis of Massey et al. (2007). High-ranked strong lens candidates are marked as green (M25 spectacular lenses and scores of 7 and above) and yellow (scores of 5 and 6) crosses. The red rectangle represents the coverage of the COSMOS-Web *JWST* imaging, and strong lenses appear only within the red rectangle since they are exclusively identified through COSMOS-Web *JWST* data. The distribution of candidates shows correlation with regions of higher weak lensing convergence, but there are also many candidates in lower density regions.

search was conducted using only the *JWST* data, so all strong lens candidates (marked with crosses) fall within this red rectangle.

There is a noticeable clustering of strong lens candidates over the higher-density peaks of dark matter inferred from weak lensing, especially at the top-left and top-right corners of the red rectangle. However, many candidates are also found in lower-density regions. A detailed statistical comparison to determine if lens candidates are more likely to be found in denser regions of dark matter large-scale




Figure 8. The Einstein radius (arcsec), lens galaxy AB magnitudes (F444W waveband), lensed source galaxy image-plane AB magnitudes (highest S/N waveband), source galaxy source-plane AB magnitudes (highest S/N waveband), magnification ratios and stellar masses as $\log_{10}(M_*/M_\odot)$ for all candidates in the second round of visual inspection. The x -axis of each plot shows each candidate’s score in the second round of visual inspection. Scores are coloured such that high ranked systems are green, mid ranked orange and low ranked red (see Section 4.3). There is no correlation between candidate score and any quantity, except stellar mass, where higher scoring candidates have higher stellar masses.

structure would require a comparison to the full sample of inspected candidates, which is beyond the scope of this work. Nevertheless, this figure highlights that all lenses have high-quality weak lensing data available, which can help inform their local environment, constrain their mass distribution on larger scales, and enable unique studies of the strong lens candidates in the COWLS sample.

5.3.3 Lens model properties

Fig. 8 presents the Einstein radii (arcsec), AB magnitudes of lens galaxies (F444W waveband), lensed source galaxy magnitudes in both the image and source planes (highest S/N waveband), magnification ratios and stellar masses (collected from the COSMOS-Web catalogue) for all candidates from the second round of visual inspection. Most candidates, regardless of their score, have Einstein radii below $1.0''$, lens galaxy magnitudes ranging from 16 to 22, and source magnitudes between 20 and 25. These values differ from existing strong lens samples, such as SLACS (Bolton et al. 2008; Shu et al. 2017), where Einstein radii are typically above $0.8''$ and lens magnitudes range from 15 to 17. The smaller Einstein radii and fainter lens galaxies are consistent with higher-redshift lens galaxies, as discussed in the redshift analysis. SLACS source magnitudes peak around 24 for the F814W *HST* filter, so the candidate sources in our sample have similar magnitudes but observed in *JWST*’s longer waveband filters (Newton et al. 2011). SLACS magnifications are

higher, with a median of $\mu = 8.8$ and peak values above 40 (Newton et al. 2011); their higher magnifications could partly explain why SLACS sources have similar magnitudes despite being selected based on SDSS spectroscopy. The stellar masses of SLACS lenses are nearly all above $\log_{10}(M_*/M_\odot) = 11$ for a Chabrier initial mass function, indicating that the majority of the COWLS sample consists of lower-mass lenses.

No correlation is observed between candidate score and the quantities shown in Fig. 8, except for stellar mass, where the highest ranked candidates tend to correspond to the highest stellar mass lens galaxies. The lack of correlation between Einstein Radius and score further suggests that lens models fitted to false positives can yield physically plausible values, given that the H25 forecast suggests around ~ 100 of the 419 are genuine lenses. While lens modelling provides crucial information for inspectors to better assess whether a system is a lens, quantitative cuts in modelling results alone therefore cannot be broadly applied across the sample to eliminate false positives.

5.3.4 Source-lens image separation

The locations of the lensed source images in a number of candidates are close to the centre of their respective lens galaxies. To quantify this, we define R_{ls} as the arcsecond radial distance between the lens galaxy centre (based on the MGE lens light model fit) and the closest




Figure 9. Lens subtracted images of two candidates which scored 7 or above where R_{ls} , the arcsecond radial distance between the lens galaxy centre (based on the MGE lens light model fit) and closest lensed source image, is below $0.15''$. The lens galaxy centres are shown with red crosses and the lensed source emission centres as cyan crosses.




Figure 10. R_{ls} , the arcsecond radial distance between the lens galaxy centre (based on the MGE lens light model fit) and closest lensed source image, for 37 candidates in the second round of visual inspection which scored 5 and above and had an R_{ls} below 0.4. The x-axis of each plot shows each candidate's score in the second round of visual inspection. The black dashed line shows $R_{ls} = 0.38$ arcsec for the strong lens Abell 1201 (Nightingale et al. 2023b), where the separation was sufficiently small that the lens galaxy's super massive black hole was detected and measured to have a mass of $M_{BH} = 3.27 \pm 2.12 \times 10^{10} M_{\odot}$. The red dashed line shows this separation for the SLACS strong lens SDSSJ1032+5322, which of the sample presented in (Etherington et al. 2022) has the lowest $R_{ls} = 0.38$ arcsec value. The COWLS lenses have some of the lowest R_{ls} values of any known strong lens.

lensed source image. Fig. 9 shows two example candidates which scored 7 or above, where $R_{ls} < 0.15$ arcsec and where in all cases the emission near the lens centre corresponds to the counter image of a larger arc located farther from the candidate lens on the opposite side. These candidate sources are therefore near the radial caustic in the source plane, which leads to these configurations.

A total of 37 candidates scored 5 or above and have $R_{ls} \leq 0.4$ arcsec. Fig. 10 shows R_{ls} versus score for these candidates, confirming that 15 have $R_{ls} < 0.2$ arcsec. The black dashed line shows the value $R_{ls} = 0.38$ arcsec, for the strong lens Abell 1201 (Nightingale et al. 2023b), where the separation was small enough




Figure 11. Two example candidates from the second round of visual inspection show evidence for singly imaged, almost-strong lensing. These objects had at least one inspector input ‘S: A singly imaged strong lens feature/arc (e.g. without an observed counter image)’, scored above 5, and after inspection of the source reconstruction, it was confirmed that the majority of the emission lies outside both caustics. Each candidate is displayed three times: the left image shows the RGB cut-out, the next image the foreground lens-subtracted image, and the right image displays the source-plane reconstruction, using the single wavelength image that most clearly shows the candidate lensed source emission. Inspectors had access to more information when grading these lenses, including data from all four wavelengths, which for these systems further supports their singly imaged nature. There are at least 8 singly imaged candidates in total.

for the lens galaxy's supermassive black hole to be detected via lens modelling, with a mass measurement of $M_{BH} = 3.27 \pm 2.12 \times 10^{10} M_{\odot}$. The red dashed line indicates this separation for the SLACS strong lens SDSSJ1032+5322, which, of the sample presented in Etherington et al. (2022), has the smallest value of $R_{ls} = 0.25$ arcsec. The COWLS sample therefore contains candidates where the source emission passes closer to the lens centre than in other surveys, which in Section 6 we argue is a selection effect.

5.3.5 Singly imaged arcs

Fig. 11 shows two example candidates from the second round of visual inspection, which display evidence for singly imaged almost-strong lensing. These objects had at least one inspector input ‘S: a singly imaged strong lens feature/arc (e.g. without an observed counter image)’, and inspection of the source reconstruction by JWN confirmed that the majority of the emission is located outside both caustics. The candidate lensed source arcs exhibit tangential shearing consistent with strong gravitational lensing and colours distinct from the lens galaxy. There are at least 8 candidates showing this behaviour. Singly imaged systems of this nature have been discovered and studied in only a few previous studies (Shu et al. 2015; Smith, Lucey & Collier 2018). Some systems may be genuine multiply imaged strong lenses, where the counter image is not observed due to it being too faint, blending with the foreground lens, or obscured by dust absorption.

5.4 Lens modelling

We now illustrate how lens modelling helped visual inspection.




Figure 12. How lens modelling enhances inspector rankings during the second round of visual inspection. The first column shows the postage stamp cut-out images used in the initial round of inspection, including one with foreground emission and another with a single Sérsic fit foreground subtraction. The following six columns present additional information provided to inspectors during the second round, which for each of the four wavebands fitted (F115W, F150W, F277W, F444W) consists of: (i) a cleaner foreground-subtracted image using a Multi-Gaussian expansion (MGE) rather than a single Sérsic fit; (ii) the model lensed source in the image plane, highlighting features that the lens model can accurately reproduce; and (iii) source-plane reconstructions, demonstrating whether the candidate has a plausible source configuration. The black and white curves represent tangential and radial critical curves and caustics, providing further insights to assess the likelihood of the candidate being a strong lens. The example lens shown was an edge case in the first round of visual inspection, meaning only one inspector thought it might be a lens, but in the second round of visual inspection scored the maximum score of all A's. It is therefore the clearest example of how lens modelling provided information which changed the inspectors' opinion on if a candidate is a lens.

5.4.1 Extra information

For many images inspected in the first round of visual inspection, the emission from the candidate source galaxy is visible but is close to and blended with the much brighter emission from the candidate lens galaxy. This made it challenging and time-consuming to get a clear view of the source emission, even after single Sérsic subtractions and careful adjustments to visualisation RGB colour maps and scalings. Fig. 12 illustrates this for the highest ranked candidate in the second round of visual inspection, which was an edge case after the first round. The cut-outs shown in the left column were used in the first round of visual inspection. These images do not clearly show the candidate's lensed source emission and have artefacts from the single Sérsic subtraction, further complicating detailed inspection for lensing features, explaining why this high-confidence candidate was ranked poorly in the first round of visual inspection.

The remaining columns in Fig. 12 show some of the additional images provided by lens modelling in the second round of visual inspection. In each waveband, lens modelling effectively deblends the lens and source, producing clean images of the lensed source galaxy, which makes it easier to judge whether the system is a lens. The lensed source model image, source-plane reconstruction, and critical curves and caustics, all shown in Fig. 12, further assist the inspector by confirming that the system fits a lens model. While we cannot quantify how many high-ranked candidates relied on this additional information, the large number of edge cases that were ranked highly in the second round of visual inspection (38.4 per cent of candidates scoring 5 and above) suggests that it is crucial for many candidates.

5.4.2 Revealing counter images

The left column of Fig. 13 shows cut-outs from the first round of visual inspection of a lens where there is emission near the candidate lens galaxy that could be a lensed source. However, its counter image is too faint and outshone by the bright lens emission, making them impossible to visually separate, even after the Sérsic subtraction. The remaining columns of Fig. 13 display the results of lens modelling, which effectively deblends the candidate lens and source emission, revealing additional emission indicative of the counter image of the lensed source. The emission appears in multiple wavebands,

confirming that it is not an artefact of the modelling process. For this system, carefully tuned RGB multicolour images did not make the counter image visible. The ability of lens modelling to reveal counter images in this case required detailed deblending, made possible only by lens modelling and its comprehensive treatment of each waveband's point spread function.

Lens modelling can therefore reveal a lensed source's faint counter image, turning candidates that initially show no signs of multiple imaging – therefore likely to be ranked low – into candidates showing multiple images and thus becoming highly ranked. In total, approximately 30 candidates scoring 5 and above had counter images that were not visible in RGB images and only became visible after lens modelling, highlighting this as a crucial way lens modelling enables lens discovery. A large number of lenses with low R_{ls} values discussed in Section 5.3.4 were only identified because of this capability. Many previous lens searches do not use lens modelling to deblend candidate lens and source emission, which is why in Section 6 we argue that our approach overcomes a selection effect that leads other surveys to miss these types of systems.

5.4.3 Compact multiple image pairing

Fig. 14 illustrates a candidate where multiply imaged compact and clumpy emissions appears in the candidate lensed source, *in addition to more extended arc/ring like features*. Furthermore, the SIE plus shear mass model can produce an image-plane to source-plane mapping which simultaneously fits the extended ring-like emission and the compact features, with the latter shown in Fig. 14 using yellow boxes. These image-plane pairings (or even groupings of four images) are strong indicators of genuine lens systems and allow inspectors to effectively distinguish true lenses from false positives, like spiral arms, ring galaxies or nearby galaxies that happen to be aligned, which are unlikely to exhibit image pairings consistent with a strong lens model. This method of identifying lenses is particularly effective in the F115W filter, which has the highest spatial resolution, allowing these clumps to be resolved, and operates at the shortest wavelength, where galaxies generally appear more clumpy. In the second round of visual inspection, inspectors were provided with figures with polygons overlaid showing mappings like this. In total,




Figure 13. Illustration of how lens modelling reveals counter images of candidates in the second round of visual inspection. The first column shows the postage stamp cut-out images used in the first round of inspection, including cut-outs with foreground emission and where a single Sérsic fit subtracts it. An arc to the left, indicative of a candidate lensed source, is present, but the corresponding counter image cannot be seen in either image, including after the single Sérsic subtraction. The remaining six columns present for each of the four wavebands fitted (F115W, F150W, F277W, F444W): (i) a cleaner foreground-subtracted image using a Multi-Gaussian expansion; (ii) the model lensed source in the image plane; and (iii) source-plane reconstructions. The black and white curves represent tangential and radial critical curves and caustics. The lens model reveals faint emission near the candidate lens’s centre indicative of the source’s counter image, with this feature visible in all four wavebands.




Figure 14. Illustration of how lens modelling enables inspectors to pair compact multiple images between the image and source planes. The first column displays the postage stamp cut-out images from the initial inspection round. The subsequent six columns, for the F115W and F277W wavebands, include: (i) a cleaner foreground-subtracted image using a Multi-Gaussian expansion (MGE); (ii) the model lensed source in the image plane; and (iii) source-plane reconstructions. Black and white curves represent tangential and radial critical curves and caustics. In the F115W images, two yellow boxes overlay the model source, indicating two multiple images in the image-plane that map to a single emission clump in the source plane. Due to lower spatial resolution, this clump is not resolved in the F277W images. These examples show how detailed inspection in high-resolution filters like F115W can reveal image-plane pairings (or even groupings of four images), which are strong indicators of genuine lens systems. This method effectively distinguishes true lenses from false positives, such as spiral arms, which are unlikely to exhibit image pairings consistent with a strong lens model.

four candidates ranked 5 and above showed evidence of this compact multiple image pairing.

5.4.4 Other remarks

After the second round of visual inspection, the inspectors met and discussed other aspects of how lens modelling information impacted their candidate ranking, with a summary as follows:

(i) The projected lens light image and SIE mass model convergence should show a degree of agreement for genuine lenses. Inspectors noted this information, but often found it hard to guide their judgement as even in genuine strong lenses the two can appear somewhat different (e.g. due to dark matter).

(ii) The MGE mass model often gave poor fits, even for high confidence lenses, because the lens galaxy was often too faint to constrain the stellar mass distribution and due to the omission of dark matter in the mass model. Reviewers tended to therefore not use this information and relied mostly on the SIE plus shear mass models.

(iii) ‘Successful’ lens models were fitted to the majority of candidates, in the sense that they focused the candidate source emission into a single region of the source-plane and produced plausible critical curves and caustics. A lens model being successful was therefore not used by itself to rank candidates highly, it was the detailed inspection of the model thereafter that was key.

5.5 Previous searches

The 1.64 deg^2 *HST* COSMOS field, which encompasses the 0.54 deg^2 COSMOS-Web field, has been extensively searched for strong lenses, yielding 64 candidates in Faure et al. (2008), 112 in Jackson (2008), and 92 in Pourrahmani et al. (2018), alongside individual discoveries (Guzzo et al. 2007; More et al. 2012; Wel et al. 2013; Jin et al. 2018; Pearson et al. 2023; Dokkum et al. 2024; Mercier et al. 2024). Of the 17 ‘spectacular’ lenses in M25, nine were previously identified, while only 8 of the 419 second-round visual inspection candidates have been reported in prior studies. These include high-scoring lenses such as COSJ095930+021351 and COSJ100050+020357 (score of 10) from Faure et al. (2008), COSJ100103+020159 (score of 9) from Jackson (2008), COSJ100028+021600 (Pourrahmani et al. 2018), COSJ095940+023253 (Faure et al. 2008), and COSJ100023+021652 (Jackson 2008), all scoring 6, as well as COSJ095943+022829 and COSJ095939+023043 from Pourrahmani et al. (2018), both scoring 4.

The majority of COWLS candidates are absent from prior *HST*-based searches, suggesting most sources were too faint for detection with *HST* and instead benefit from *JWST*’s unprecedented depth. Additionally, some source galaxies are extremely red, appearing only in F444W and/or F277W bands, making them ‘*HST*-dark’. *HST* F814W imaging is available for most candidates, and visual inspection confirms that lensed source emission is rarely visible. A quantitative analysis, including lens modelling, is beyond this

work’s scope but would further characterize these newly uncovered systems.

Previous searches also include 14 candidates from Faure et al. (2008), 18 from Jackson (2008), and 18 from Pourrahmani et al. (2018) that did not advance to our second inspection round. We generated *JWST* images and performed lens modelling for all of these, which are now available in the COWLS public data release. JWN individually inspected these candidates, finding most to be false positives due to: (i) spiral/barred galaxies that, at *HST* depth, mimic strong lenses but reveal their structure in *JWST* imaging, and (ii) candidate systems with a nearby line-of-sight galaxy but lacking counter-images in lens modelling. However, two candidates from Pourrahmani et al. (2018) would have been graded A by JWN, and therefore have a reasonable probability of being genuine strong lenses the COWLS search missed.

6 DISCUSSION

6.1 The COSMOS-Web lens survey

This work presents the COSMOS-Web Lens Survey (COWLS). It includes over 100 highly ranked strong lens candidates, including the spectacular lenses discussed in M25. Selected via *JWST*, the sample exhibits unique characteristics compared to other lens surveys, which drive a variety of science cases which will only be possible with this sample for years to come:

(i) **Highest redshift source galaxies:** While source redshifts remain unmeasured, the high lens redshifts and results of COWLS paper III imply the COWLS sample contains some of the most distant galaxy-scale sources which extend beyond $z > 6$ and into the epoch of reionization. Sources are imaged in remarkable detail with *JWST*’s deep multiband data, for example at $z = 6$ the F115W filter will achieve effective ~ 100 pc resolution. Being galaxy-scale lenses, this allows precise reconstructions of their unlensed structure, enabling unprecedented studies of high-redshift galaxy morphology, as done for sources up to $z \sim 3$ (Swinbank et al. 2015).

(ii) **Highest redshift lens galaxies:** Spectroscopic and photometric data indicate that half the lenses lie at $z > 1$, with some pushing beyond $z > 2$. This regime is crucial for studying galaxy density profile evolution (Shajib et al. 2021; Etherington et al. 2024; Tan et al. 2024), dark matter substructure (Vegetti et al. 2014; Nightingale et al. 2024) and leveraging large lens samples for cosmology (Li et al. 2024; Geng et al. 2025).

(iii) **Strong lenses in a single cosmic volume:** All lenses reside within a contiguously imaged 0.54 deg^2 region, enabling the combination of strong and weak lensing to measure cosmic shear with unparalleled precision (Birrer et al. 2017; Fleury et al. 2021; Hogg et al. 2023, 2025b), and the study of lens environments within their large-scale structure (Peng et al. 2010).

(iv) **Lensed sources near lens centres:** In a subset of COWLS candidates, lensed emission passes within 0.25 arcsec of the lens galaxy centre, closer than most previously known lenses. This may allow detection of the influence of the lens’s supermassive black hole on the lensing signal (Nightingale et al. 2023b) and provide insights into dust absorption effects (Krekel et al. 2013; Barone et al. 2024).

The survey’s top priority is confirming which candidates are genuine strong lenses, best achieved by measuring spectroscopic redshifts of the source galaxies. Many lensed sources are faint, visible only due to *JWST*’s deep imaging, making NIRSpect follow-up or observations with eight-metre-class ground-based telescopes necessary – though feasible for only a subset of the lenses. Candidates

benefit from extensive multiwavelength COSMOS data, including MIRI F770W, *HST* F814W, and numerous ground-based exposures, which enabled the precise photometric lens redshifts presented here. Whilst faint lensed sources may not have wide wavelength coverage, photometric redshift analysis could still help determine whether the candidate lensed emission originates at a different redshift than the candidate lens galaxy, given that the lens redshift is known.

The impact of source redshift availability varies by science case. For instance, studying the density slope evolution of lens galaxies requires only lens redshifts (Etherington et al. 2024), while cosmology combining strong and weak lensing is not critically dependent on a complete set of source redshifts (Hogg et al. 2023). Dark matter subhalos and SMBHs near the lens galaxy can be detected without a source redshift, but converting their mass to physical units requires one. Similarly, source galaxy studies can identify key targets using multiwavelength reconstructions and photometric data, but require redshifts to know the physical scales and time in the Universe’s formation of the galaxy.

Another priority is studying the source galaxy population and identifying the most intriguing objects for detailed individual study, leveraging lensing magnification to extract more information than would otherwise be possible for galaxies in the distant Universe. To support this, the public data release includes unlensed source reconstructions for all candidates across four NIRC2 bands and additional detected wavebands. Preliminary analysis (see M25) confirms a fraction of *HST*-dark red galaxies (Barrufet et al. 2023; Pérez-González et al. 2023), indicative of intense star formation beyond $z > 4$. The sample also contains compact, red lensed sources, potentially belonging to the ‘red nugget’ population – progenitors of massive ellipticals (Barro et al. 2013; Dekel & Burkert 2014; Oldham et al. 2017) – or ‘little red dots’ at even higher redshifts (Matthee et al. 2024).

6.2 Undiscovered lenses observed with The James Webb space telescope

We identified numerous high-quality strong lens candidates within just 0.54 deg^2 of *JWST* data. Given that *JWST* has conducted science observations for over three years, its archive now contains extensive high-quality imaging across much larger areas. Yet, only one galaxy-scale *JWST* strong lens discovery has been reported to date, which is included in our survey (Dokkum et al. 2024; Mercier et al. 2024).

This suggests that *JWST* has likely already observed hundreds more undiscovered strong lenses. Some may have been spotted by scientists during data analysis but not formally reported. As demonstrated by COWLS, these lenses form a unique sample, distinct from existing catalogs and future wide-field surveys like *Euclid* (Collett & Auger 2014), with many likely probing the highest lens and source redshifts. We encourage further searches of the *JWST* archive, using machine learning or citizen science approaches, as successfully done with *HST* (Garvin et al. 2022).

M25 focuses on the 17 most ‘Spectacular’ lenses identified by COWLS, which are immediately recognizable by eye as strong lenses in *JWST* data. M25 also estimates how many similar high-quality lens candidates are likely already detected but not yet reported.

6.3 Lens modelling for lens finding

A significant fraction of high-ranked lens candidates would not have been identified without lens modelling. For example, 38.4 per cent of candidates that scored highly in the second round of visual inspection (with lens modelling information) were edge cases, meaning only one

inspector flagged them as potential lenses in the first round (without this information). The improved visualization and detailed insights provided by lens models allowed inspectors to rank these candidates more accurately, even those that were previously overlooked. We also documented that 30 candidates exhibited faint counter-images near the centre of the lens galaxy, which were not visible by eye (even in manually tuned RGB images) but became apparent after lens modelling. This is because it cleanly deblends the lens and source emission, fully utilising the PSF in each waveband.

Our results suggest that ongoing lens-finding efforts in wide-field surveys, such as *Euclid* – expected to discover over 100 000 strong lenses (Collett 2015; Holloway et al. 2023; Ferrami & Wythe 2024) – will significantly benefit from incorporating lens modelling into their processes. Lens-finding studies using ground-based data and modelling (e.g. Sonnenfeld et al. 2018, 2020; Rojas et al. 2022) have highlighted several advantages, such as improved arc visibility through foreground lens subtraction and the ability to reject false positives by testing whether a candidate fits a lens model. Our study identifies two additional benefits: (i) deblending faint counter-images near the lens galaxy center, and (ii) matching compact source features to assess the plausibility of a lens candidate. These improvements are likely enabled by a combination of more advanced modeling techniques and the higher resolution of space-based imaging. The modeling approach used here was recently applied to *Euclid* Q1 data (Euclid Collaboration 2025a, b, c), where both methods of lens identification were again found to be effective.

However, lens modelling does not definitively determine whether a candidate is a lens. We found that most candidates could be fitted with a physically plausible lens model, even when other evidence (such as colours) suggested they were unlikely to be strong lenses. Measured lens model parameters (e.g. Einstein radius, magnification) showed no correlation with candidate rankings, indicating that lens models fitted to both high- and low-ranked candidates were statistically indistinguishable. Additionally, fit quality metrics – such as reduced chi-squared values – did not effectively differentiate between high- and low-ranked candidates.

Thus, our study highlights that lens modelling should serve as a tool to provide visual inspectors with enhanced information for ranking candidates, rather than relying solely on the successful fitting of a lens model as a diagnostic. This approach helps to identify and properly rank the most promising lens candidates.

6.3.1 Comparison to euclid lens search experiment

Acevedo Barroso et al. (2025) used lens modeling of 16 candidates from high-resolution *Euclid* early release observations to identify lenses based on three criteria: (1) whether an SIE critical curve can enclose or exclude the correct number of bright components in the image plane; (2) whether the critical curve, as an isodensity contour, aligns with the lens galaxy light profile; and (3) whether the reconstructed source surface brightness distribution resembles a compact, focused object within a caustic.

In our modelling of 419 candidates, we found that defining strict criteria or metrics for identifying lenses is unreliable. Many objects that meet such criteria score low and are likely not lenses based on other factors, such as colors. While lens modeling significantly aids candidate assessment, we believe its usefulness is more subjective and does not adhere to rigid rules. Notably, lens modeling excels at revealing counter-images and multiple images, providing crucial information beyond straightforward criteria.

6.4 Selection effects

The COWLS sample highlights a selection effect that previous lens samples built by inspecting wide-field ground-based surveys are subject to, for example, SL2S, and DES (Gavazzi et al. 2012, 2014; Sonnenfeld et al. 2013; Jacobs et al. 2019; Tran et al. 2022). The separation between their lens centres and the closest lensed source image is constrained by the resolution of the instrument, which for ground-based surveys is around 0.27 arcsec. The higher resolution of the NIRCам allowed us to find lenses where this separation was as low as 0.1 arcsec. This required that lens modelling provided a clean deblending of the lens and source light. We therefore recommend that future lens surveys based on wide-field imaging, most notably *Euclid* with a resolution of ~ 0.15 arcsec, incorporate lens modelling to ensure they recover these low-separation strong lenses.

This also raises the question of what selection effects COWLS itself is subject to, given that the selection function of strong lenses is now well characterized (Sonnenfeld et al. 2023; Sonnenfeld 2024). COWLS Paper III investigates this by forecasting lens properties such as Einstein radii, redshifts, and magnitudes, comparing these to the values measured in this paper. This analysis confirms that typical lensing selection biases are at play, for example, the tendency for lens galaxies to be more massive than galaxies in photometric catalogues.

Future scientific studies will benefit from a more complete characterization of our selection function, noting that achieving this for a complete sample enables an unbiased comparison to the general galaxy population (Sonnenfeld 2022). This would require a probabilistic assessment of how inspectors identified lenses and which ones were overlooked. However, owing to its relatively small size and the high value of its *JWST* data, COWLS could serve as an ideal test case for performing such an analysis on a smaller scale before applying it to much larger lens samples.

7 SUMMARY

The *James Webb* Space Telescope (*JWST*) provides a powerful new avenue for discovering strong gravitational lenses, thanks to its high-resolution imaging capabilities and the unprecedented depth offered by its 6.5 metre mirror. To capitalize on this, we performed a systematic search for strong lenses in COSMOS-Web (Casey et al. 2024), *JWST*'s largest contiguous survey, covering 0.54 deg^2 in four NIRCам bands (F115W, F150W, F277W, and F444W). We present the COSMOS-Web Lens Survey (COWLS), a sample of over 100 strong lens candidates. This includes 17 ‘spectacular’ lenses reported in COWLS Paper II by Mahler et al. (2025) and is consistent with the lens forecasts presented in COWLS Paper III by Hogg et al. (2025a). A key aspect of our approach was the incorporation of traditional lens modelling, using the open-source software PYAUTOLENS (Nightingale & Dye 2015; Nightingale et al. 2018, 2021b), which enabled us to identify many lenses that would otherwise have been missed.

The COWLS sample stands out due to several unique features that enhance its scientific potential. First, the sample contains some of the most distant source galaxies of galaxy-scale strong lenses ever found, with redshifts anticipated to extend beyond $z > 6$ and into the epoch of reionization. Being galaxy-scale lenses, this allows precise reconstructions of their unlensed structure, providing an unprecedented opportunity to study high-redshift galaxy morphology. Secondly, the lens galaxies are at higher redshifts than previous samples, with approximately half at $z > 1$, and some reaching beyond $z > 2$. This redshift range is crucial for understanding the evolution of galaxy density profiles and strong lensing studies of

cosmology. All COWLS lenses reside within a contiguously imaged 0.54 deg^2 region, enabling the combination of strong and weak lensing to measure cosmic shear with unparalleled precision and to investigate lens environments within their large-scale structure. Finally, a subset of COWLS candidates features lensed emission passing within 0.25 arcsec of the lens galaxy centre, offering the potential to detect the influence of the supermassive black hole on the lensing signal (Nightingale et al. 2023b), while also providing new insights into dust absorption effects.

JWST has been operating for over three years, and the COWLS sample represents just the tip of the iceberg in terms of the strong lenses it has observed, many of which remain undiscovered. Future COWLS studies will demonstrate the potential of these lenses for probing the high-redshift Universe, highlighting the need for dedicated efforts to uncover and exploit the vast reservoir of strong lenses hidden within the *JWST* archive.

SOFTWARE CITATIONS

This work uses the following software packages:

- (i) *PyAutoFit* (Nightingale, Hayes & Griffiths 2021a)
- (ii) *PyAutoGalaxy* (Nightingale et al. 2023a)
- (iii) *PyAutoLens* (Nightingale & Dye 2015; Nightingale et al. 2018, 2021c)
- (iv) *Astropy* (Astropy Collaboration 2013; Price-Whelan et al. 2018)
- (v) *Corner.py* (Foreman-Mackey 2016)
- (vi) *Dynesty* (Speagle 2020b)
- (vii) *Matplotlib* (Hunter 2007)
- (viii) *Numba* (Lam, Pitrou & Seibert 2015)
- (ix) *NumPy* (van der Walt, Colbert & Varoquaux 2011)
- (x) *Python* (Van Rossum & Drake 2009)
- (xi) *Scikit-image* (Van der Walt et al. 2014)
- (xii) *Scikit-learn* (Pedregosa et al. 2011)
- (xiii) *Scipy* (Virtanen et al. 2020)
- (xiv) *SQLite* (Hipp 2020)

ACKNOWLEDGEMENTS

JWN is supported by an STFC/UKRI Ernest Rutherford Fellowship, Project Reference: ST/X003086/1. AA and QH acknowledge support from the European Research Council (ERC) through Advanced Investigator grant DMIDAS (GA 786910). GM, MvW-K, and RM were supported in Durham by STFC via grant ST/X001075/1, and the UK Space Agency via grant ST/X001997/1.

This work used both the Cambridge Service for Data Driven Discovery (CSD3) and the DiRAC Data-Centric system, project code dp195, which are operated by the University of Cambridge and Durham University on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). These were funded by BIS capital grant ST/K00042X/1, STFC capital grants ST/P002307/1, ST/R002452/1, ST/H008519/1, ST/K00087X/1, STFC Operations grants ST/K003267/1, ST/K003267/1, and Durham University. DiRAC is part of the UK National E-Infrastructure.

This work was made possible by utilizing the CANDIDE cluster at the Institut d'Astrophysique de Paris. The cluster was funded through grants from the PNCG, CNES, DIM-ACAV, the Euclid Consortium, and the Danish National Research Foundation Cosmic Dawn Center (DNRF140). It is maintained by Stephane Rouberol. The French contingent of the COSMOS team is partly supported by the Centre National d'Etudes Spatiales (CNES). OI acknowledges

the funding of the French Agence Nationale de la Recherche for the project iMAGE (grant ANR-22-CE31-0007). DS acknowledge the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

DATA AVAILABILITY

Analysis results are publicly available at <https://github.com/Jammy2211/COWLS-COSMOS-Web-Lens-Survey>.

REFERENCES

- Acevedo Barroso J. A. et al., 2025, *A&A*, 697, A14
 Amvrosiadis A. et al., 2025, *MNRAS*, 537, 1163
 Aretxaga I. et al., 2011, *MNRAS*, 415, 3831
 Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, *MNRAS*, 310, 540
 Astropy Collaboration, 2013, *A&A*, 558, A33
 Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V., Marshall P. J., Moustakas L. A., Burles S., 2010, *ApJ*, 724, 511
 Barbary K., 2016, *J. Open Source Softw.*, 1, 58
 Barone T. M. et al., 2024, *Nat. Commun. Phys.*, 7, 1
 Barro G. et al., 2013, *ApJ*, 765, 11
 Barrufet L. et al., 2023, *MNRAS*, 522, 449
 Bertin E., 2011, in Evans I. N., Accomazzi A., Mink D. J., Rots A. H., eds, ASP Conf. Ser. Vol. 442, Astronomical data analysis software and systems XX, Astron. Soc. Pac., San Francisco, p. 435
 Bertin E., Mellier Y., Radovich M., Missonnier G., Didelon P., Morin B., 2002, in Bohlender D. A., Durand D., Handley T. H., eds, ASP Conf. Ser. Vol. 281, Astronomical Data Analysis Software and Systems XI, Astron. Soc. Pac., San Francisco, p. 228
 Bertin E., Schefer M., Apostolakis N., Álvarez-Ayllón A., Dubath P., Kümmel M., 2020, in Pizzo R., Deul E. R., Mol J. D., de Plaa J., Verkouter H., eds, ASP Conf. Ser. Vol. 527, Astronomical Data Analysis Software and Systems XXIX, Astron. Soc. Pac., San Francisco, p. 461
 Birrer S., Welschen C., Amara A., Refregier A., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 049
 Birrer S., Refregier A., Amara A., 2018, *ApJ*, 852, L14
 Birrer S. et al., 2020, *A&A*, 643, A165
 Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Moustakas L. A., 2006, *ApJ*, 638, 703
 Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Gavazzi R., Moustakas L. A., Wayth R., Schlegel D. J., 2008, *ApJ*, 682, 964
 Bolton A. S. et al., 2012, *ApJ*, 757, 82
 Bro R., De Jong S., 1997, *J. Chemometr.*, 11, 393
 Cao X. et al., 2022, *Res. Astron. Astrophys.*, 22, 30
 Cappellari M., 2002, *MNRAS*, 333, 400
 Casey C. M. et al., 2024, *ApJ*, 965, 98
 Collett T. E., 2015, *ApJ*, 811, 20
 Collett T. E., Auger M. W., 2014, *MNRAS*, 443, 969
 Collett T. E., Oldham L. J., Smith R. J., Auger M. W., Westfall K. B., Bacon D., Nichol R. C., Masters K. L., 2018, *Science*, 360, 1342
 Dekel A., Burkert A., 2014, *MNRAS*, 438, 1870
 Despali G., Vegetti S., White S. D. M., Giocoli C., van den Bosch F. C., 2018, *MNRAS*, 475, 5424
 Despali G., Vegetti S., White S. D., Powell D. M., Stacey H. R., Fassnacht C. D., Rizzo F., Enzi W., 2022, *MNRAS*, 510, 2480
 Dey A. et al., 2019, *AJ*, 157, 168
 Drlica-Wagner A. et al., 2018, *ApJS*, 235, 33
 Duboscq T., Hogg N. B., Fleury P., Larena J., 2024, *J. Cosmol. Astropart. Phys.*, 2024, 021
 Etherington A. et al., 2022, *MNRAS*, 517, 3275
 Etherington A. et al., 2023, *MNRAS*, 521, 6005
 Etherington A. et al., 2024, *MNRAS*, 531, 3684
 Euclid Collaboration, 2025a, preprint (arXiv:2503.15324)

- Euclid Collaboration, 2025b, preprint ([arXiv:2503.15326](https://arxiv.org/abs/2503.15326))
- Euclid Collaboration, 2025c, preprint ([arXiv:2503.15325](https://arxiv.org/abs/2503.15325))
- Faure C. et al., 2008, *ApJS*, 176, 19
- Ferrami G., Wyithe J. S. B., 2024, *MNRAS*, 532, 1832
- Fleury P., Larena J., Uzan J.-P., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 024
- Foreman-Mackey D., 2016, *J. Open Source Softw.*, 1, 24
- Garvin E. O., Kruk S., Cornen C., Bhatawdekar R., Cañameras R., Merín B., 2022, *A&A*, 667, 141
- Gavazzi R., Treu T., Rhodes J. D., Koopmans L. V. E., Bolton A. S., Burles S., Massey R. J., Moustakas L. A., 2007, *ApJ*, 667, 176
- Gavazzi R., Treu T., Marshall P. J., Brault F., Ruff A., 2012, *ApJ*, 761, 170
- Gavazzi R., Marshall P. J., Treu T., Sonnenfeld A., 2014, *ApJ*, 785, 144
- Geng S., Grespan M., Thuruthipilly H., Harikumar S., Pollo A., Biesiada M., 2025, *A&A*, 694, 15
- Guzzo L. et al., 2007, *ApJSS*, 172, 254
- Harrington K. C. et al., 2016, *MNRAS*, 458, 4383
- He Q. et al., 2022a, *MNRAS*, 511, 3046
- He Q. et al., 2022b, *MNRAS*, 512, 5862
- He Q. et al., 2023, *MNRAS*, 518, 220
- He Q. et al., 2024, *MNRAS*, 532, 2441 (H24)
- Hipp R. D., 2020, SQLite. Available at: <https://www.sqlite.org/index.html>
- Hogg N. B., 2024, *MNRAS: Lett.*, 528, L95
- Hogg N. B., Fleury P., Larena J., Martinelli M., 2023, *MNRAS*, 520, 5982
- Hogg N. B. et al., 2025a, preprint ([arXiv:2503.08785](https://arxiv.org/abs/2503.08785)) (H25)
- Hogg N. B., Shajib A. J., Johnson D., Larena J., 2025b, preprint ([arXiv:2501.16292](https://arxiv.org/abs/2501.16292))
- Holloway P., Verma A., Marshall P. J., More A., Tecza M., 2023, *MNRAS*, 525, 2341
- Hunter J. D., 2007, *Comput Sci Eng.*, 9, 90
- Ilbert O. et al., 2006, *A&A*, 457, 841
- Jackson N., 2008, *MNRAS*, 389, 1311
- Jacobs C. et al., 2019, *MNRAS*, 484, 5330
- Jin S. et al., 2018, *ApJ*, 864, 56
- Jin S. et al., 2024, *A&A*, 690, L16
- Koekemoer A. M. et al., 2011, *ApJSS*, 197, 3535
- Koopmans L. V. et al., 2009, *ApJ*, 703, L51
- Krekel K. et al., 2013, *ApJ*, 771, 62
- Kümmel M., Bertin E., Schefer M., Apostolakis N., Álvarez-Ayllón A., Dubath P., 2020, in Pizzo R., Deul E. R., Mol J. D., de Plaa J., Verkouter H., eds, ASP Conf. Ser. Vol. 527, Astronomical Data Analysis Software and Systems XXIX. Astron. Soc. Pac., San Francisco, p. 29
- Lam S. K., Pitrou A., Seibert S., 2015, *Proc. 2nd Workshop on the LLVM Compiler Infrastructure in HPC – LLVM’15*. ACM, p. 1
- Lange J. U., 2023, *MNRAS*, 525, 3181
- Li T., Collett T. E., Krawczyk C. M., Enzi W., 2024, *MNRAS*, 527, 5311
- Liu D. et al., 2024, *Nat. Astron.*, 8, 1181
- Mahler G. et al., 2025, preprint ([arXiv:2503.08782](https://arxiv.org/abs/2503.08782)) (M25)
- Maresca J., Dye S., Li N., 2021, *MNRAS*, 503, 2229
- Massey R. et al., 2007, *ApJS*, 172, 239
- Matthee J. et al., 2024, *ApJ*, 963, 129
- Melo-Carneiro C. R., Furlanetto C., Chies-Santos A. L., 2023, *MNRAS*, 520, 1613
- Mercier W. et al., 2024, *A&A*, 687, A61
- More A., Cabanac R., More S., Alard C., Limousin M., Kneib J. P., Gavazzi R., Motta V., 2012, *ApJ*, 749, 38
- Nagam B. C. et al., 2025, preprint ([arXiv:2502.09802](https://arxiv.org/abs/2502.09802))
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Negrello M. et al., 2014, *MNRAS*, 440, 1999
- Newton E. R., Marshall P. J., Treu T., Auger M. W., Gavazzi R., Bolton A. S., Koopmans L. V., Moustakas L. A., 2011, *ApJ*, 734, 104
- Nightingale J. W., Dye S., 2015, *MNRAS*, 452, 2940
- Nightingale J., Dye S., Massey R., 2018, *MNRAS*, 478, 4738
- Nightingale J. W., Hayes R. G., Griffiths M., 2021a, *J. Open Source Softw.*, 6, 2550
- Nightingale J. W. et al., 2021b, *J. Open Source Softw.*, 6, 2825
- Nightingale J. et al., 2021c, *J. Open Source Softw.*, 6, 2825
- Nightingale J. W. et al., 2023a, *J. Open Source Softw.*, 8, 4475
- Nightingale J. W. et al., 2023b, *MNRAS*, 521, 3298
- Nightingale J. W. et al., 2024, *MNRAS*, 527, 10480
- O’Riordan C. M. et al., 2025, *A&A*, 694, A145
- Oldham L. et al., 2017, *MNRAS*, 465, 3185
- Pearce-Casey R. et al., 2025, *A&A*, 696, A214
- Pearson J. et al., 2023, *MNRAS*, 527, 12044
- Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825
- Peng C. Y., Ho L. C., Impey C. D., Rix H. W., 2010, *AJ*, 139, 2097
- Pérez-González P. G. et al., 2023, *ApJL*, 946, L16
- Pourrahmani M., Nayyeri H., Cooray A., 2018, *ApJ*, 856, 68
- Price-Whelan A. M. et al., 2018, *AJ*, 156, 123
- Ritondale E., Vegetti S., Despali G., Auger M. W., Koopmans L. V., McKean J. P., 2019, *MNRAS*, 485, 2179
- Rizzo F., Vegetti S., Powell D., Fraternali F., McKean J. P., Stacey H. R., White S. D., 2020, *Nature*, 584, 201
- Rizzo F., Vegetti S., Fraternali F., Stacey H. R., Powell D., 2021, *MNRAS*, 507, 3952
- Rojas K. et al., 2022, *A&A*, 668, 1
- Shajib A. J., 2019, *MNRAS*, 488, 1387
- Shajib A. J., Treu T., Birrer S., Sonnenfeld A., 2021, *MNRAS*, 503, 2380
- Shu Y. et al., 2015, *ApJ*, 803, 1
- Shu Y. et al., 2016, *ApJ*, 824, 86
- Shu Y. et al., 2017, *ApJ*, 851, 48
- Shuntov M. et al., 2025a, *A&A*, 695, A20
- Shuntov M. et al., 2025b, *A&A*, 696, L14
- Sibson R., 1981, *Interpreting Multivariate Data*. John Wiley and Sons, New York
- Smith R. J., Lucey J. R., Collier W. P., 2018, *MNRAS*, 481, 2115
- Sonnenfeld A., 2022, *A&A*, 659, A132
- Sonnenfeld A., 2024, *A&A*, 690, A325
- Sonnenfeld A., Gavazzi R., Suyu S. H., Treu T., Marshall P. J., 2013, *ApJ*, 777, 97
- Sonnenfeld A. et al., 2018, *PASJ*, 70, 1
- Sonnenfeld A. et al., 2020, *A&A*, 642, A148
- Sonnenfeld A., Li S. S., Despali G., Gavazzi R., Shajib A. J., Taylor E. N., 2023, *A&A*, 678, 1
- Speagle J. S., 2020a, *MNRAS*, 493, 3132
- Speagle J. S., 2020b, *MNRAS*, 493, 3132
- Swinbank A. M. et al., 2015, *ApJ: Lett.*, 806, L17
- Szalay A. S., Connolly A. J., Szokoly G. P., 1999, *AJ*, 117, 68
- Tan C. Y. et al., 2024, *MNRAS*, 530, 1474
- Tessore N., Metcalf R. B., 2015, *A&A*, 580, A79
- Tran K.-V. H. et al., 2022, *AJ*, 164, 148
- van Dokkum P., Brammer G., Wang B., Leja J., Conroy C., 2024, *Nat. Astron.*, 8, 119
- 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 der Wel A. et al., 2013, *ApJ*, 777, L17
- Van Rossum G., Drake F. L., 2009, *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA
- Vegetti S., Koopmans L. V. E., Auger M. W., Treu T., Bolton A. S., 2014, *MNRAS*, 442, 2017
- Vieira J. D. et al., 2010, *ApJ*, 719, 763
- Vieira J. D. et al., 2013, *Nature*, 495, 344
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Warren S., Dye S., 2003, *ApJ*, 590, 673

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

COSMOSWeb Image Supplementary (2).pdf

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: VISUAL INSPECTION DETAILS

Table A1 shows the distribution of ‘Yes’ and ‘Maybe’ inputs from the 5 inspectors. Inspector 4 was notably optimistic, classifying over 8 percent of objects as ‘Yes’ or ‘Maybe’. There is also significant variation among the remaining inspectors. The final column highlights the number of unique objects each inspector flagged as ‘Yes’ or ‘Maybe’ that no other inspector did. Inspectors 1 to 5 had 345, 518, 99, 2950, and 92 unique candidates each, illustrating that it is common for inspectors to disagree.

Table 2 displays the distribution of categories for the three data set groups: 2023 January, 2023 April, and 2024 January. Notably, the At Least 50 percent ‘Maybe’ category, which dictates if a candidate makes it to the second round of visual inspection, contains over double the percentage of candidates in 2024 April compared to 2024 January. This is because, for 2023 April (and 2023 January), there were only 4 inspectors, including the optimistic inspector, making it more likely a candidate reached this category as only one of the other 3 inspectors had to agree with the optimistic inspector. With 5 inspectors, 2 inspectors had to agree, making it less likely for candidates to make it into this category.

The optimistic inspector influenced the number of candidates advancing to the second round of visual inspection. This effect primarily comes from lenses making it into the At Least 50 percent ‘Maybe’ category. However, they also contributed a number of the edge cases put forward after reinspection by JWN. Out of the 29 edge cases that became highly ranked, 14 qualified as edge cases due to the optimistic inspector.

Table A1. The table presents the number and percentage of objects each inspector assigned the inputs ‘Y—Yes, this is a lens’ and ‘M—Maybe this is a lens’ during the first round of visual inspection. Inspectors 1–4 inspected all 42 660 objects, whereas inspector 5 only inspected the 2024 January data comprising 27 125 objects, which is reflected in the percentages. The final column provides the number of candidates an inspector flagged as ‘Yes’ or ‘Maybe’ that are unique because no other inspector flagged them.

Inspector	Yes	Maybe	Unique
1	19 (0.05 %)	615 (1.44 %)	345 (0.81 %)
2	106 (0.25 %)	868 (2.03 %)	518 (1.21 %)
3	23 (0.05 %)	247 (0.57 %)	99 (0.22 %)
4	1082 (2.50 %)	2500 (5.86 %)	2950 (6.91 %)
5	45 (0.17 %)	188 (0.69 %)	92 (0.34 %)

APPENDIX B: CATALOGUE

The full catalogue of candidate lenses, visual inspection scores and measured properties is given in the supplementary materials.

- ¹*School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, Newcastle-upon-Tyne, NE1 7RU, UK*
- ²*STAR Institute, Quartier Agora - Allée du six Août, 19c B-4000 Liège, Belgium*
- ³*Department of Physics, Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK*
- ⁴*Department of Physics, Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*
- ⁵*Department of Physics, Northeastern University, 360 Huntington Ave., Boston, MA 02115, USA*
- ⁶*Laboratoire Univers et Particules de Montpellier, CNRS & Université de , Parvis Alexander Grothendieck, Montpellier F-34090, France*
- ⁷*Department of Computer Science, Aalto University, PO Box 15400, Espoo FI-00 076, Finland*
- ⁸*Department of Physics, Faculty of Science, University of Helsinki, Helsinki FI-00014, Finland*
- ⁹*Laboratoire d’astrophysique de Marseille, 38 Rue Frédéric Joliot Curie, Aix Marseille University, CNRS, CNES, 13013, Marseille, France*
- ¹⁰*Jet Propulsion Laboratory, California Institute of Technology, 4800, Oak Grove Drive, Pasadena, CA, 91011, USA*
- ¹¹*Purple Mountain Observatory, Chinese Academy of Sciences, 10 Yuanhua Road, Nanjing 210023, China*
- ¹²*Cosmic Dawn Centre (DAWN), Ørstedts Plads 342, 2800 Kongens Lyngby, Denmark*
- ¹³*Niels Bohr Institute, University of Copenhagen, Jagtvej 128, DK-2200 Copenhagen, Denmark*
- ¹⁴*Department of Astronomy, The University of Texas at Austin, 2515 Speedway, Austin, TX 78712,, USA*
- ¹⁵*Institut d’Astrophysique de Paris, UMR 7095, CNRS, Sorbonne Université, 98 bis boulevard Arago, F-75014 Paris, France*
- ¹⁶*Department of Physics, University of California, Santa Barbara, CA 93106, USA*
- ¹⁷*Laboratory for Multiwavelength Astrophysics, School of Physics and Astronomy, Rochester Institute of Technology, 84 Lomb Memorial Drive, Rochester, NY 14623, USA*
- ¹⁸*Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA*
- ¹⁹*Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA*
- ²⁰*Department of Physics and Astronomy, University of Hawaii, Hilo, 200 W Kawili St, Hilo, HI 96720, USA*
- ²¹*Caltech/IPAC, 1200 East California Boulevard, Pasadena, CA 91125, USA*
- ²²*DTU-Space, Technical University of Denmark, Elektrovej 327, DK-2800 Kgs. Lyngby, Denmark*

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