

Exploiting the diversity of modeling methods to probe systematic biases in strong lensing analyses

A. Galan^{1,2,3,*}, G. Vernardos^{4,5}, Q. Minor^{6,5}, D. Sluse⁷, L. Van de Vyvere⁷, and M. Gomer⁷

¹ Technical University of Munich, TUM School of Natural Sciences, Department of Physics, James-Franck-Str 1, 85748 Garching, Germany

² Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany

³ Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, 1290 Versoix, Switzerland

⁴ Department of Physics and Astronomy, Lehman College of the CUNY, Bronx, NY 10468, USA

⁵ American Museum of Natural History, Department of Astrophysics, New York, NY 10024, USA

⁶ Borough of Manhattan Community College, City University of New York, Department of Science, New York, NY 10007, USA

⁷ STAR Institute, University of Liège, Quartier Agora, Allée du six Août 19c, 4000 Liège, Belgium

Received 13 June 2024 / Accepted 23 September 2024

ABSTRACT

Challenges inherent to high-resolution and high signal-to-noise data as well as model degeneracies can cause systematic biases in analyses of strong lens systems. In the past decade, the number of lens modeling methods has significantly increased, from purely analytical methods, to pixelated and non-parametric ones, or ones based on deep learning. We embraced this diversity by selecting different software packages and use them to blindly model independently simulated *Hubble* Space Telescope (HST) imaging data. To overcome the difficulties arising from using different codes and conventions, we used the COde-independent Organized LENS STandard (COOLEST) to store, compare, and release all models in a self-consistent and human-readable manner. From an ensemble of six modeling methods, we studied the recovery of the lens potential parameters and properties of the reconstructed source. In particular, we simulated and inferred parameters of an elliptical power-law mass distribution embedded in a shear field for the lens, while each modeling method reconstructs the source differently. We find that, overall, both lens and source properties are recovered reasonably well, but systematic biases arise in all methods. Interestingly, we do not observe that a single method is significantly more accurate than others, and the amount of bias largely depends on the specific lens or source property of interest. By combining posterior distributions from individual methods using equal weights, the maximal systematic biases on lens model parameters inferred from individual models are reduced by a factor of 5.4 on average. We investigated a selection of modeling effects that partly explain the observed biases, such as the cuspy nature of the background source and the accuracy of the point spread function. This work introduces, for the first time, a generic framework to compare and ease the combination of models obtained from different codes and methods, which will be key to retain accuracy in future strong lensing analyses.

Key words. methods: data analysis – methods: statistical – galaxies: elliptical and lenticular, cD – galaxies: structure – cosmological parameters – cosmology: observations

1. Introduction

Understanding the evolutionary path of galaxies over cosmic times continues to be a major challenge in astrophysics. In this context, strong gravitational lensing enables the observation of galaxies lying at different redshifts in a single observation, making it an inescapable tool to constrain galaxy evolution models. Strong gravitational lensing arises when a foreground distant galaxy – the lens, or deflector – is coincidentally aligned with a more distant background galaxy – the source – causing the appearance of multiple and magnified images of the latter. The typical redshift range for lens galaxies lies between $z_d \sim 0.2$ and 1.5, while source galaxies are often found between redshifts $z_s \sim 1$ and 4 (Oguri & Marshall 2010; Collett 2015), such that strong lensing systems can display a wide variety of galaxy morphologies and evolutionary stages.

Besides galaxy evolution studies, strong lensing has several important applications in cosmology. As it is dictated by the total mass of galaxies, one can use this effect to put con-

straints on their dark matter halo. In particular, strong lensing data enables the separation of the baryonic and dark components of galaxies (e.g., Suyu et al. 2012; Shajib et al. 2021), and the detection of dark matter subhalos and other invisible masses along the line of sight (e.g., Vegetti et al. 2010; Sengül et al. 2022; Nightingale et al. 2024). When combined with time-varying sources, strong lenses can also be used to measure cosmological parameters, including the *Hubble* constant (H_0 e.g., Wong et al. 2020; Birrer et al. 2020; Kelly et al. 2023) and density parameters (e.g., with multiplane lensing systems, Collett & Auger 2014). All these applications rely heavily on a precise characterization of both the azimuthal and radial mass profiles of lens galaxies.

The central step when analyzing strong lensing data is lens modeling. The goal of this step is to model both the mass and light distribution of the lens galaxy, while simultaneously reconstructing an unlensed version of the source galaxy. Lens modeling is a challenging task because inverting the lensing effect is an ill-posed problem, in particular due to known degeneracies between the lens mass distribution and the source

* Corresponding author; aymeric.galan@gmail.com

morphology. For example, the infamous mass-sheet degeneracy (MSD, Falco et al. 1985; Schneider & Sluse 2013), a mathematically exact degeneracy between the lens mass density and the source scaling, has been studied both theoretically and practically (e.g., Birrer et al. 2016; Unruh et al. 2017; Wagner 2018; Gomer & Williams 2020; Cao et al. 2022) and can be mitigated using complementary data sets (e.g., Birrer et al. 2020; Yıldırım et al. 2023; Khadka et al. 2024). In the past twenty years, many different lens modeling techniques, ranging from analytical to pixelated techniques and neural networks, have been developed and successfully applied to real images (e.g., Warren & Dye 2003; Suyu et al. 2006; Vegetti & Koopmans 2009; Birrer et al. 2016; Nightingale et al. 2018; Galan et al. 2024). In general, these techniques have been developed with specific lensing systems, data sets, and science goals in mind, and then have been extended to cover more use cases. Consequently, it is crucial to assess how these different methods compare to each other, and if their combination is warranted to improve the robustness of lensing analyses. Such a comparison enables the quantification of possible systematic biases. Additionally, if several methods lead to consistent results, those can be combined together, improving the overall accuracy.

So far, lens modeling comparison analyses have been rare. For cluster-scale systems, a prominent work has been initiated by Meneghetti et al. (2017) by performing an extensive comparison of several modeling approaches on both simulated and real *Hubble* Frontier Fields clusters. However, there have not been comparable efforts for galaxy-scale strong lens systems. While some works have reanalyzed archival data with alternative modeling software packages (e.g., Birrer et al. 2016; Shajib et al. 2021), it is only recently that more quantitative comparisons between different methods have been reported (Shajib et al. 2022; Etherington et al. 2022). The Time Delay Lens Modeling Challenge (TDLMC) compared the output of different modeling and inference strategies, but focusing only on the recovery of H_0 (Ding et al. 2021). Different lens modeling codes have been compared in Lefor & Futamase (2015) and Pascale et al. (2024), although using point-like multiple images rather than extended gravitational arcs as constraints. Finally, recent works from Schuldt et al. (2023a) and Gawade et al. (2024) compared neural network predictions with more classical approaches, although on with ground-based imaging data.

Our goal here is to analyze imaging data similar to those obtained with the *Hubble* Space Telescope (HST) with different modeling and inference techniques, and study the recovery of a given set of lens parameters while reconstructing the lensed source in different ways. To our knowledge, this is the first time a systematic and self-consistent comparison between a large number (six) of state-of-the-art galaxy-scale lens modeling methods has been conducted. In order to maintain this novel kind of analysis tractable, we have restricted the assumptions regarding the description of the lens mass distribution and properties of the data, although remaining reasonably realistic. In particular, we limit our scope to the commonly used power-law elliptical mass distribution embedded in a shear field. This description of the lens deflection field has proven to be a minimal but efficient prescription for modeling the observed strong lensing effect caused by large elliptical galaxies (e.g., Koopmans et al. 2006; Suyu et al. 2013; Millon et al. 2020; Shajib et al. 2021; Etherington et al. 2022; Tan et al. 2024, to cite only a few), although the simplicity of this model has known limitations (Sonnenfeld 2018; Gomer & Williams 2021; Cao et al. 2022; Etherington et al. 2024; Ruan & Keeton 2023). We also note that recent analyses of strong lenses found evidence for multipolar

deviations to the elliptical power-law profile, but a higher resolution than HST is warranted for robust detection (Powell et al. 2022; Stacey et al. 2024). While most lensing analyses focus on the properties of the galaxies acting as lenses, the morphology of the lensed galaxies also hold important information about galaxy formation and evolution. Current high-resolution images of strong lenses such as those from HST showcase highly structured lensed sources (e.g., Bolton et al. 2006; Garvin et al. 2022; Wang et al. 2022). Consequently, it is also crucial to assess the ability of lens modeling codes to recover the morphology of extended lensed sources.

We first selected different lens modeling software packages and modeling methods that are well suited to model high-resolution and high signal-to-noise (S/N) data. Since each software package typically follows different parameter definitions and model conventions, it is not possible to directly compare the modeling results. We overcame this challenge by using the COde-independent Organized LEns STandard (COOLEST, Galan et al. 2023), an open-source standard that enables storage, sharing, and analysis of all lens modeling products in a uniform manner, regardless of the modeling code originally used to perform the lens modeling tasks. Included in this standard is an analysis interface allowing us to compute important quantities (e.g., effective radii and profile slopes) and visualize lens modeling results. We have extensively used COOLEST in this work, both for releasing the models and data, as well as performing the analysis of the results and generating the figures.

The paper is organized as follows. In Sect. 2 we briefly recall the strong lensing formalism we follow. In Sect. 3, we present the different lens modeling methods, in particular their commonalities and differences. We explain how the data was simulated using an independent software in Sect. 4, and the standardized comparison and analysis framework is introduced in Sect. 5. The modeling results after unblinding are visualized and described in Sect. 6, followed by an exploration of possible sources of systematics in Sect. 7. In Sect. 8 we discuss our results and place them in a broader context, and Sect. 9 concludes our work.

2. Formalism of strong gravitational lensing

We give for completeness a brief overview of the mathematical formalism to describe strong lensing data and models. More background details can be found in recent reviews such as Vegetti et al. (2023), Shajib (2024), Saha et al. (2024).

The main strong lensing observables are the positions and intensities of multiply lensed images of features in a background sources. These features can either be unresolved (i.e., point sources) or spatially extended. In the latter case, the lensed source appears as several arcs or as an Einstein ring surrounding the lens object and typically covering many pixels in high-resolution imaging data. We call the (observable) plane of the sky where lensed images appear the “image plane”, that we place at the redshift z_d of the foreground lens, also called the main deflector. Observed features in the image plane are localized with a two-dimensional angular position vector, θ . For conciseness, we interchangeably use the standard Cartesian coordinates (x, y) to describe a position, θ , in the image plane. Each feature in the lensed images has a corresponding (unobservable) angular position, β , in the “source plane” placed at the redshift of the background object, z_s .

The central equation in gravitational lensing is the lens equation, which gives the relationship between β and θ :

$$\beta = \theta - \nabla\psi(\theta), \quad (1)$$

where $\nabla\psi \equiv \alpha$ is the deflection field originating from the lens potential, ψ , the latter being a rescaled and projected version of the underlying three-dimensional gravitational potential of the lens galaxy. Usually, a more physically relevant quantity is the projected mass density of the lens, characterized by the so-called lens convergence, κ (dimensionless), obtained with a combination of second derivatives of the lens potential:

$$\kappa = \frac{1}{2} \nabla^2 \psi. \quad (2)$$

As it will be useful for the discussion (Sect. 8.4), we also recall the formula of the Fermat potential, mostly relevant for time-varying sources. The Fermat potential, Φ_i , and the Fermat potential difference, $\Delta\Phi_{ij}$, between a pair of lensed images i and j , are defined as

$$\Phi(\theta_i) = \frac{(\theta_i - \beta)^2}{2} + \psi(\theta_i), \quad (3)$$

$$\Delta\Phi_{ij} \equiv \Phi(\theta_i) - \Phi(\theta_j), \quad (4)$$

where θ_i and θ_j are the positions of images i and j , respectively.

In this work, we consider parametrized forms for the lens mass distribution, while the surface brightness of the lensed galaxy is described following a variety of techniques. We give more details about the modeling of these different components in Sect. 3.

3. Lens modeling methods and assumptions

We have considered an ensemble of “modeling methods” that each differ on two aspects: modeling assumptions and inference techniques. For instance, modeling assumptions are typically specific choices of model components (mass and light profiles, fixed or not), regularization strategies for pixelated models and necessary hyper-parameters. Inference techniques are typically minimization and sampling algorithms to obtain best-fit parameters and estimate their posterior distributions, or sequences of distinct steps (e.g., preliminary coarse and fast model fits) to converge to the best-fit solution.

In practice, a given lens modeling software package can be considered as a modeling method, as specific choices regarding the code structure, model types and optimization techniques have been made throughout its development. For this work, we selected a subset of software packages that are sufficiently different to be considered as distinct modeling methods: LENSTRONOMY (Birrer & Amara 2018; Birrer et al. 2021), the Very Knotty Lenser (VKL, Vernardos & Koopmans 2022), HERCULENS (Galan et al. 2022, 2024) and QLENS (Minor et al., in prep.). Other software packages used in several published analyses so far are GLEE (Suyu & Halkola 2010; Suyu et al. 2012), PYAUTOLENS (Nightingale & Dye 2015; Nightingale et al. 2018, 2021), GLAFIC (Oguri 2010) and methods from Vegetti & Koopmans (2009, and subsequent works). However, for practical reasons we only use the first set of methods, which already form a representative sample of the various modeling methods that are currently available, from fully analytical to pixelated models, with or without adaptive grids. Such methods are referred to as classical methods, in contrast to deep learning methods that we do not consider in this work (e.g., Schuldt et al. 2023b; Adam et al. 2023; Gentile et al. 2023, for some recent works), as these would require additional assumptions regarding training sets and network architectures beyond our scope. Nevertheless, we encourage future works to conduct

self-consistent comparison analyses similar to ours, that involve both classical and deep learning methods (see e.g., Schuldt et al. 2023a).

This remaining of this section presents the general modeling strategy we adopt throughout this work. We first describe modeling assumptions that are common to all methods, then give more details regarding each of these modeling methods, and finally mention extra choices that are left free to the modelers.

3.1. Common modeling aspects

Throughout this work, we reasonably assume that the noise in the imaging data, \mathbf{d} , follows a Gaussian distribution with covariance matrix, C_d . In this setting, we can write the negative log-probability of the data likelihood as

$$-\log \mathcal{L}(\eta) = \frac{1}{2} \left[\mathbf{m}(\eta) - \mathbf{d} \right]^T C_d^{-1} \left[\mathbf{m}(\eta) - \mathbf{d} \right] + \log \left(2\pi \sqrt{\det C_d} \right), \quad (5)$$

where \mathbf{m} is the predicted image (i.e., the model) and η represents a generic vector of model parameters. We note that C_d is assumed to be diagonal with contributions from both background noise and photon noise. In other words, we follow the widely used assumption that the noise is uncorrelated and normally distributed. With simulated data, we have access to the true matrix C_d . As in this work we do not explore the effects of inaccurate assumptions regarding noise characteristics, we give to the modelers the true matrix C_d and use it in all lens models.

While the likelihood term in Eq. (5) is common to all models considered here, specific modeling assumptions such as morphological properties of the source galaxy are encoded as additional priors. Such priors can either be explicitly incorporated in the inference via a regularization term written as the negative of the log-prior $-\log \mathcal{P}$, or they can be implicitly defined through a choice of parametrization such as an analytical functions. Summing the log-likelihood and log-prior terms gives the full penalty or loss function, L , which is directly proportional to the log-posterior and minimized during the inference of model parameters:

$$L(\eta) = -\log \mathcal{L}(\eta) - \log \mathcal{P}(\eta). \quad (6)$$

Modeling methods generally describe the lensing of photons from the source by casting the lens equation into a lensing operator, \mathbf{L} , which depends on the lens potential parameters that we denote by η_ψ . This operator acts on a model of the source, s , described by parameters, η_s , which can be either analytical, pixelated or a representation in function basis set, as per

$$\mathbf{m}(\eta) \equiv \mathbf{m}(\eta_\psi, \eta_s) = \mathbf{R} \mathbf{B} \mathbf{L}(\eta_\psi) s(\eta_s). \quad (7)$$

so that we get a model image, \mathbf{m} , that has the same pixel size as the data, after possible downsampling by the operator \mathbf{R} and blurring by the operator \mathbf{B} . The latter incorporates the effect of the point spread function (PSF) of the instrument and seeing conditions. This PSF is assumed to be known with the same spatial sampling as the data and available to all modelers. As stated in Sect. 3.6, no constraints are imposed to modelers regarding optional supersampling of ray-tracing and convolution operations. The light distribution of the lens galaxy is not modeled because we assume that the lens light has been perfectly subtracted from the data beforehand.

For modeling the lens mass distribution of the lensing galaxy – alternatively, its lens potential – we consider the commonly used power-law elliptical mass distribution (PEMD) with an additional shear field component. This shear component has been commonly referred to as “external” shear as it captures the net effect masses external to the lens along the line of sight; we use this wording within the scope of our work but it does not necessarily holds true when modeling real systems (see e.g., Etherington et al. 2024).

The assumption of a PEMD with external shear is common to all modeling methods, namely all modelers use the same lens potential parameter vector η_ψ . The convergence of the PEMD is described by (Barkana 1998; Tessore & Metcalf 2015):

$$\kappa_{\text{PEMD}}(x, y) = \frac{3 - \gamma}{2} \left(\frac{\theta_E}{\sqrt{q_m x^2 + y^2 / q_m}} \right)^{\gamma-1} \quad (8)$$

where γ is the logarithmic power-law slope ($\gamma = 2$ corresponding to an isothermal profile), q_m is the axis ratio and the coordinate system (x, y) has been rotated by a position angle ϕ_m around the lens center (x_0, y_0) . The lens potential generated by an external shear can be easily expressed in polar coordinates with the following formula (e.g., Etherington et al. 2024):

$$\psi_{\text{ext}}(x, y) \equiv \psi_{\text{ext}}(r, \phi) = \frac{r^2}{2} \gamma_{\text{ext}} \cos \left[2(\phi - \phi_{\text{ext}}) \right], \quad (9)$$

where γ_{ext} is the strength of the external shear, and ϕ_{ext} its position angle. We note that Eqs. (8) and (9) follow the parameters conventions used in the COOLEST (see Sect. 5 and the online documentation for other conventions¹).

We note that the loss function of Eq. (6) has a non-linear response to lens mass parameters, η_ψ . However, the set of source parameters, η_s , can formally be split into linear (light profile amplitudes) and non-linear parameters. This property is explicitly exploited by some of the modeling methods we use in this work.

The modeling methods considered in this work thus mainly differ in the assumptions regarding the light distribution of the source, $s(\eta_s)$, which we summarize in Table 1 and describe in more detail in the next subsections.

3.2. Smooth modeling with Sérsic and shapelets

We used a modeling method implemented in the multipurpose LENSSTRONOMY package. We followed the baseline model presented in Sect. 3.1, and which consists of a PEMD plus external shear for the lens. The source was modeled with a Sérsic profile, to which were added shapelets basis functions, which are capture additional complexity of the source light distribution. We implemented this modeling strategy using version 1.11.3 of the multi-purpose open-source software package LENSSTRONOMY (Birrer & Amara 2018; Birrer et al. 2021). This tool, regularly enhanced with new user-contributed capabilities, provides a large family of lens mass distribution and light profiles. We refer the reader to Birrer et al. (2015) for a formal description of the shapelets model in the context of lens modeling.

The procedure used in LENSSTRONOMY to derive the posterior distribution on the parameters is sequential. The optimal linear parameters are found through matrix inversion, given

¹ <https://coolest.readthedocs.io/en/latest/conventions.html>

values of non-linear parameters. First, a suitable region in non-linear parameter space that minimizes the loss function defined in Eq. (6) is found via a Particle Swarm Optimization algorithm (PSO, Kennedy & Eberhart 2002). Second, the parameters space is sampled using a Monte-Carlo Markov Chain (MCMC). The parameters of the optimal model found previously are randomly perturbed, and used to start the chain. We use the MCMC sampler emCEE, which is the most used so far among the LENSSTRONOMY user community (Foreman-Mackey et al. 2013). In this work, as it is also a common practice, the model investigated during the optimization step is a simplified version of the final model, retaining only the main model components that enable to reproduce the largest fraction of the data pixel values. Components that yield small changes of the loss function, such as the source shapelets, are added only during the MCMC sampling. This hierarchy in the significance of model parameters, while not explicitly formalized in the code, is similar to the methodology developed by several automatized lens modeling efforts (e.g., Etherington et al. 2022; Ertl et al. 2023; Tan et al. 2024). We give more technical details regarding this method in Appendix B.1.

3.3. Adaptive grid source modeling

If one assumes that the free parameters of the source are its brightness values cast on a grid of pixels s (instead of being defined from a continuous analytical profile), then the likelihood of Eq. (5) becomes a quadratic function of s . The benefit of such quadratic functions is that their derivative can be calculated analytically and have a unique minimum (Warren & Dye 2003). However, with just the likelihood term this leads to an ill-posed problem and the addition of a (quadratic) regularization term is required, which has the following generic form:

$$-\log \mathcal{P}(\lambda, \mathbf{g}, s) = \frac{1}{2} \lambda s^T C_s^{-1}(\mathbf{g}) s, \quad (10)$$

where C_s is some covariance kernel of the source as a function of parameters, \mathbf{g} . In this form, the source parameters can be obtained analytically from $\nabla_s L = 0$ once the lens parameters, η_ψ , regularization strength, λ , and covariance parameters, \mathbf{g} , are given. This approach is referred to as semi-linear inversion. In comparison with forward methods, which may treat more source parameters as non-linear parameters, there are only a few additional parameters that require sampling, λ and \mathbf{g} (usually corresponding to just one or two parameters). The linear source parameters, in other words the pixel brightness values, are obtained using matrix inversion.

A key assumption in such inverse problems is the choice of regularization, which can be interpreted in a Bayesian way as a prior imposed on the source. Traditionally, one may choose to impose smoothness to the solution through its derivatives, where the matrix C_s^{-1} is constructed from the numerical derivative coefficients computed on the pixelated grid (Warren & Dye 2003; Suyu et al. 2006; Vegetti & Koopmans 2009). Alternatively, more physically motivated covariance kernels obtained from real galaxy brightness distributions have been shown to perform better and lead to less biased results (Vernardos & Koopmans 2022). A quite generic such example is the Matérn kernel that has the following analytic form:

$$C(r_{ij}|l, \nu) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{r_{ij} \sqrt{2\nu}}{l} \right)^\nu K_\nu \left(\frac{r_{ij} \sqrt{2\nu}}{l} \right), \quad (11)$$

where r_{ij} is the distance between any two source pixels and ν, l correspond to the non-linear parameters \mathbf{g} (the latter can be

Table 1. Labels used to identify the six modeling methods of this work.

Modeling method label	Description	Corresponding source modeling technique
Sérsic+Shapelets	Sect. 3.2	Concentric elliptical Sérsic profile and shapelets basis functions
Adaptive+Matérn	Sect. 3.3	Adaptive Voronoi grid, regularized with a Matérn kernel
Cluster+Exp	Sect. 3.3	Clustered adaptive Voronoi grid, regularized with an exponential kernel
Cluster+Exp+Lumweight	Sect. 3.3	Similar to Cluster+Exp with regularization weighted by source luminosity
Sparsity+Wavelets	Sect. 3.4	Cartesian grid, regularized with sparsity in wavelets space and non-negativity constraints
Correlated Field	Sect. 3.5	Cartesian grid with Gaussian processes, regularized with a parametric power-spectrum

Notes. These modeling methods mainly differ in their source modeling techniques, but share other differences (see Sect. 3 for details).

interpreted as a correlation length). The type of regularization can be objectively chosen based on the Bayesian evidence.

Another feature employed in semi-linear inversion implementations is the use of adaptive, non-regular grids for the source. This is because smaller regions in the source plane can contribute a higher fraction of the total flux, especially in the regions of high magnification near caustics, hence higher resolution is required. Using a high resolution fixed regular grid will lead to more computationally demanding matrix inversions, but adapting the resolution to the lensing magnification (which is a function of the lens model parameters η) will increase the resolution in those regions of the source plane where it is needed without adding more degrees of freedom. Such adaptive grids can be constructed simply by tracing a subset of the data pixels back to the source (Vegetti & Koopmans 2009; Vernardos & Koopmans 2022) and using them as grid vertices. A more sophisticated way of constructing the adaptive grid, pioneered by Nightingale & Dye (2015), is to split each pixel into $N \times N$ subpixels, ray-trace the subpixels to the source plane, and then use a k -means clustering algorithm to determine the location of the source grid vertices. The latter approach has the advantage of minimizing aliasing effects, at the cost of additional overhead due to the clustering algorithm.

The regularization of an adaptive grid may be extended to allow for greater fluctuations in surface brightness in the inner regions of the source with high S/N, while keeping the outer regions of the source relatively smooth. Luminosity-weighted regularization has been explored by Nightingale et al. (2018) in the context of gradient regularization. In this work, we explored luminosity-weighted covariance kernels to achieve a similar effect. To implement this, we define the luminosity-weighted kernel as:

$$C_{ij,\text{lum}} = C_{ij} W_i W_j, \quad (12)$$

with the weighting function W_i given by:

$$W_i = \exp[-\rho(1 - s_{i,0}/s_{\text{max}})], \quad (13)$$

where ρ is a free parameter to be varied, $s_{i,0}$ is an approximate surface brightness of the i -th pixel before luminosity weighting, and s_{max} is the maximum surface brightness of all the source pixels. Pixel values $s_{i,0}$ can be taken from the best-fit model of a previous fit altogether, or it can be estimated during each likelihood evaluation by doing an inversion without luminosity weighting first; here, we adopted the latter approach. The advantage of the form given in Eq. (12) is that it ensures that the kernel remains positive-definite, which is critical given the quadratic form of the regularization term. Although more sophisticated forms of the weighting function W_i are possible, we only explored the single parameter function given by Eq. (13) in this work.

Here, we used two implementations of the semi-linear inversion technique, VKL (Vernardos & Koopmans 2022) and QLENS (Minor et al., in prep.), which can handle different regularization schemes and recipes for constructing the adaptive source-pixel grid. In the VKL modeling run, which we label Adaptive+Matérn, a Matérn kernel will be used for regularization without supersampling of the image plane; whereas in the QLENS modeling runs, an exponential kernel will be used (equivalent to Matérn with $\nu = 0.5$), image plane supersampling with k -means clustering will be used to construct the adaptive source grid, and models without and with luminosity weighted regularization will be used (Cluster+Exp and Cluster+Exp+Lumweight, respectively). Both implementations use Nested Sampling (as implemented by the MultiNest algorithm, Feroz et al. 2009) to sample the parameter space and converge to the maximum a posteriori solution, in addition to calculating the Bayesian evidence. We outline the specific modeling choices that are considered as three different source modeling techniques in Sects B.2 and B.3.

3.4. Multi-scale regularization with wavelets

Regularization does not necessarily have to be quadratic so that it can lead to a linear solution for the source. A non-linear example is the method by Joseph et al. (2019), Galan et al. (2021) that is based on sparsity constraints in the wavelet domain. We refer to this regularization as a multi-scale regularization (and use “ms” as subscripts in the corresponding equations). We briefly describe the method here, but refer the reader to the original papers for the full mathematical treatment.

Given the vector representation of the source expressed on a regular (non-adaptive) pixelated grid, the regularization term to be minimized (i.e., the second term in Eq. (6)) is:

$$-\log \mathcal{P}(\eta, s, \lambda_{\text{ms}}) = \lambda_{\text{ms}} \left\| \mathbf{W}_{\text{ms}}(\eta_\psi, \eta_s) \odot \Phi^\top s \right\|_1 + i_{\geq 0}(s), \quad (14)$$

where λ_{ms} is a global (scalar) regularization parameter, Φ^\top is the wavelet transform operator that transforms s into its wavelets coefficients, and \mathbf{W}_{ms} is a matrix that scales the regularization strength for each these coefficients. Operations $\|\cdot\|_1$ and \odot refer to the ℓ_1 norm and element-wise product, respectively. The second term in Eq. (14) is a non-negativity constraint on the source since we reconstruct surface brightness values. The operator Φ^\top represents an hybrid wavelet transform composed of all scales from of the starlet transform (Starck et al. 2007) and the first scale of the Battle-Lemarié wavelet transform (for more details see e.g., Lanusse et al. 2016; Galan et al. 2022). The matrix elements of \mathbf{W}_{ms} are computed by propagating the data noise from the source plane to the wavelet domain (hence the dependence on η_ψ, η_s), allowing us to attach a clear meaning to the global regularization strength λ_{ms} , interpreted as the statistical significance of the

regularized source model and given in units of the noise (e.g., $\lambda_{\text{ms}} = 3\sigma$).

We dynamically adapt the extent of the source plane regular grid based on the lens mass by defining the annular region in image plane (the ‘‘arc mask’’) within which the pixelated source light is evaluated. This adaptive scheme ensures that the effective source pixel size covers the same angular scale of the source for any realization of the mass model². This treatment is different from the fixed source plane grid used in Galan et al. (2021), and provides more stability when optimizing lens mass models with a free density slope.

We used the strong lens modeling code HERCULENS (Galan et al. 2022) that implements the multi-scale regularization strategy described in Eq. (14) using JAX (Bradbury et al. 2018), such that the full model can be pre-compiled and is differentiable. We used the probabilistic programming library NUMPYRO (Phan et al. 2019; Bingham et al. 2019) to implement prior distributions and constraints on model parameters, and to perform the inference of posterior distributions. Additional technical details regarding the Sparsity+Wavelets models are given in Appendix B.4.

3.5. Non-parametric Gaussian processes

Finally, we applied a recently introduced source reconstruction method that relies on Gaussian processes and information field theory (IFT Enßlin 2019). The first applications of Gaussian processes for modeling strongly lensed sources are the recent works of Karchev et al. (2022), Rüstig et al. (2024) and Galan et al. (2024), but we describe below the main principles for completeness. In the IFT framework, such Gaussian processes are often referred to as correlated fields, which is the term we use in the remaining.

We use IFT to represent the source light distribution as a two-dimensional non-negative correlated field. Such a field is based on two main components: (1) an analytical parametrization of its power spectrum (mainly its amplitude and slope) that accounts for correlated structures in the source; (2) a discretization onto a regular square grid, with elements following standardized Gaussian distributions (we call the latter an excitation field). More formally, we describe the source field, s , as

$$s = \exp \left[\mathbf{F}^{-1} (A_0 \odot \xi) + \delta \right], \quad (15)$$

where A_0 is a zero-mode spectral field generated from the parametrized power spectrum, ξ is the excitation field, and \odot is the point-wise multiplication. The resulting field in harmonic space is then transformed to real space by applying the inverse Fourier transform operator, \mathbf{F}^{-1} , to which a constant offset, δ , is added. Finally, we take the exponential of the resulting field to enforce the positivity of source pixel values, since they represent surface brightness. We note that no extra regularization term is added to the loss function, as Eq. (15) describes a generative model that already incorporates smoothness conditions through its power spectrum.

We use the Python library NIFTY³ (Selig et al. 2013; Steininger et al. 2019; Arras et al. 2019) to implement the above field model and variational inference samplers to get the joint posterior distribution over the parameter space. As in Galan et al.

(2024), we use the JAX interface of NIFTY (nifty.re, Edenhofer et al. 2024) and combine it with HERCULENS to evaluate the forward model (Eq. (7)). For further technical details regarding the Correlated field model, see Appendix B.5.

3.6. Additional choices left to the modelers

There are some additional choices that are left to the modelers. In particular, they are free to choose if and how they mask out some regions in the imaging data and exclude those from the data likelihood evaluation. Similarly, super-sampling of the coordinates grid when performing ray-tracing evaluations and surface brightness convolutions with the PSF are optional. Modelers are free to run variations in their fiducial models, varying, for example, hyper-parameters, random number generator seeds, or inference algorithms. This may allow for robustness tests, and any eventual marginalization over families of models leading to a joint posterior distribution as the final result.

4. Independent simulation of imaging data

We produced a simulated mock system to model with all methods presented above and compare the results. We simulated the mock using MOLET⁴ (Vernardos 2022), a simulator code that is independent of any of the codes used to fit the mock. The simulated mock is constructed using a power-law mass profile with external shear and a relatively structured source, as is expected from most lensed galaxies. The true source used in the simulation is kept hidden (blind) from the modelers, as well as all other input parameters.

For the source, we used an HST image of the galaxy NGC 1084 that was previously used for simulated lenses in the context of the TDLMC (Ding et al. 2021). This source is a local galaxy with a detailed structure that is resolved without any prominent PSF spikes, which could introduce nonphysical features if lensed. We further avoided introducing unphysical features in the resulting lensed source due to edge effects (e.g., sky background) in the cutout that we used by forcing the values in the source pixels to decay exponentially to zero away from the brightness peak.

The simulated lensing data is then created by providing the mass model and source to MOLET, which performs ray-tracing on a high-resolution grid, convolution with the PSF, downsampling, and adding noise to the final mock observation. We use a simulated HST PSF using TINYTIM (Krist et al. 2011) based on the WFPC2 instrument with the F814W filter (we do not consider the more recent WFC3 instrument as TINYTIM does not support it). Since ray-tracing is performed on a 10 times higher resolution grid compared to the final data, we use a simulated PSF at that resolution for more accurate surface brightness convolutions (although modelers are given a PSF at the data resolution, as is mentioned in Sect. 3.1). Our settings for the noise correspond to 2200 s of exposure in the chosen instrument setup. The angular size of the source cutout is set to 4 arcsec, and we scale its total flux such that it has an apparent (unlensed) AB magnitude of 23.2, which is in the range of observed source galaxies from the Sloan Lens ACS Survey (SLACS) sample (Bolton et al. 2006; Newton et al. 2011). We show in the top panel of Fig. 1 the simulated lens image, while the bottom row shows the supersampled and data-resolution PSFs (only the latter is provided to the modelers).

² We note that a similar strategy is used in the GLEE modeling code to adapt the extent of the source plane regular grid (Suyu & Halkola 2010; Suyu et al. 2012).

³ <https://gitlab.mpcdf.mpg.de/ift/nifty>

⁴ <https://github.com/gvernard/molet>




Fig. 1. Simulated HST imaging data used for the blind lens modeling experiment (see Sect. 4 for details). The top left panel shows a zoom-in cutout of the data. The top right panels shows the true (unlensed) source surface brightness. The bottom left panel shows the true PSF kernel downsampled to the data resolution (and given to the modelers). The bottom right panel shows normalized residuals in image plane obtained with a too simplistic source model (a single Sérsic profile).

We note that the lensed source galaxy and the data S/N are such that a simple source model is not able to fit the data. We visualize this in the top right panel of Fig. 1, which shows normalized residuals between the data and a model based on a single Sérsic profile for the source. Such residuals are strong evidence for the necessity of more flexible source models as the ones we employ in this work (described in Sect. 3).

In Appendix A we give useful details about a previous version of the mock we attempted to model, for which we detected issues related to the input source light distribution. In a nutshell, the original source did not have an accurate background subtraction and displayed sharp edges (visually unnoticeable after lensing the addition of noise), which led to biases in the lens models. For this reason, a second mock with different input parameters and source light had to be re-created (and the subsequent modeling re-done).

5. Standardized comparison framework

Our work relies on several collections of modeling methods and software packages, that we systematically apply on the same data. These codes have been developed following different conventions (e.g., angles, units, profile definitions), are written in different programming languages (e.g., Python, C++), and differ in their final modeling products. Therefore, we must ensure that we can both simulate and model strong lensing data in a consistent way, in order to mitigate problems arising from the heterogeneous collection of methods we consider.

We used the recently released strong gravitational lensing standard COOLEST (for CODE-independent Organized LENSing STandard) as a framework unifying the different components of our analysis⁵. Below we briefly describe COOLEST and its specific features that we used in this work, but refer the reader

⁵ COOLEST is an open source Python package publicly available at <https://github.com/aymgal/COOLEST>. We used the released version 0.1.9.

to Galan et al. (2023) and the online documentation⁶ for more details.

The foundation of COOLEST is a set of conventions to serve as a reference point for modeling assumptions and codes, for example, coordinate systems, units and profile definitions. Given these conventions, any lens model – together with the data being modeled and other modeling components such as the PSF – can be concisely described in a single file following the JSON format. We refer to the latter as a COOLEST template file, which we use both to describe an instance of a strong lens to be simulated and to store lens modeling results (e.g., best-fit parameters and uncertainties). This standard way of storing lens modeling information allows us to straightforwardly compare any modeling results to each other as well as to an existing groundtruth (in the case of simulated data). In practice, it requires each modeling or simulation code to have an interface with COOLEST to create or update such template files⁷. Lastly, we use the analysis features of COOLEST to read the content of template files, compute key lensing quantities, and produce comparison plots. In particular, we use this interface to plot all lens models side-by-side and compare them to the groundtruth, compute morphological features of reconstructed source galaxies, and plot joint posterior distributions over the parameter space.

6. Results from blind modeling with six methods

After agreeing on the properties of the simulated data (instrumental setup, S/N, type of mass model), one of the authors (M.G.) first simulated the imaging data with a first software (Sect. 4). The rest of the authors proceeded with the blind modeling of the data, whose results are presented here. The modeling workload was split between different modelers: L.V.V. (Sérsic+Shapelets models), G.V. (Adaptive+Matérn model), Q.M. (Cluster+Exp and Cluster+Exp+Lumweight models) and A.G. (Sparsity+Wavelets and Correlated Field models). In total, we used four distinct software packages and six different modeling methods. Each modeler was free to perform their analysis to the best of their judgment. Once confident that no significant improvements to the models could be made by further fine-tuning, the modelers converted their results to the COOLEST format (Sect. 5) and submitted them to A.G. (this may have been a single or a marginalization of model instances). During the simulation and the subsequent modeling phase, there was minimum amount of information shared between the authors. In particular, the shared information between the simulating author and the modelers was restricted to the only data pixels, the corresponding noise map as well as the data-resolution PSF. Afterwards, the modelers did not share their lens modeling results before all authors agree to do so. Keeping this part of the analysis blind ensured unbiased and independent models. Unblinding took place when M.G. also submitted the true model, and the final figures presented in this section were produced. The specific steps necessary to obtain these models, as well as their estimated computation times, are detailed in Appendix B. All models along with the simulated data products are publicly available⁸.

⁶ <https://coolest.readthedocs.io>

⁷ This task is made easier by using the dedicated COOLEST Python interface.

⁸ https://github.com/aymgal/LensSourceDegeneracy_public. Upon acceptance of the manuscript, analysis and visualization notebooks will also be released.

6.1. Overall fit to the imaging data

We show in Fig. 2 all the six models that have been blindly submitted, in direct comparison with the true model from the simulation. We compare side-to-side the image model, relative error on the convergence, model normalized residuals (in unit of the noise), and the reconstructed sources. The reconstructed sources should be directly compared to the top right panel of Fig. 1. The residuals show clear improvements over the residuals shown in the bottom right panel of Fig. 1 with a too simplistic source model. We also quote the reduced chi-square value χ^2_{ν} computed within the likelihood mask chosen by the modeler, in order to quantitatively compare the quality of fits. The largest χ^2_{ν} value is 1.06 while the lowest value is 0.88, and four models have a χ^2_{ν} below unity which indicate slight overfitting. In several models, residuals at the $\sim 3\sigma$ level remain where the arcs are the brightest (multiple images of the bright, cuspy central region of the source). We note that overall, all models fit the imaging data to very close to noise level.

While the Sérsic+Shapelets model captures less small-scale structures in the source compared to other models, the resulting fit in the image plane still achieves noise-level residuals (except maybe for a few pixels in the outskirts of the lensed source). Finer structures in the source like spiral arms and star forming regions are overall better modeled by semi-linear inversion methods, than by the wavelets and the correlated field.

The second column of Fig. 2 shows the relative error in the lens convergence, throughout the field-of-view. Offsets in the lens centroid and in the ellipticity are visible, although strongest errors remain at the very center of the lens. At the position of the multiple images (roughly traced by the critical lines), the relative error in convergence remains below 5%.

6.2. Recovery of source properties

To visualize better which source features are captured by the models, we show in Fig. 3, maps of source plane residuals computed as the difference between the true and reconstructed sources. Far from the optical axis, models are not well constrained by the data and significantly deviates from the true light distribution; thus, we darken these areas for better visualization. Pixelated models defined on irregular grids capture similarly well most of the spiral features as well as star forming region located in the left spiral arm. The center of the source galaxy seems to be retrieved equally well by all models, despite the persistent image plane residuals (see Fig. 2).

We show in Fig. 4 the recovery of several properties of the source galaxy: the two-point correlation function, the effective radius and the axis ratio. Bottom parts of each panel show the relative error computed as $(\text{truth} - \text{model})/\text{truth}$ (i.e., negative values are over-estimates). We measure these properties within a square field of view of size $2'' \times 2''$, after projecting (using bi-cubic interpolation) each source model on a regular grid with 10 times higher resolution than the data. As is shown in Fig. 2, such a field of view contains the entire flux of the source galaxy.

The left panel of Fig. 4 shows the two-point correlation function $\xi(r)$, which gives the azimuthally averaged correlation between the source intensity at two positions in source plane, as a function of their angular separation. All reconstructed sources exhibit two-point correlations close to the one of the input galaxy. Over all the models, the maximum error remains small, except for some models for which it exceeds 15% error on two-point correlations on arcsecond scales. We note that the Sérsic+Shapelets model under-estimates two-point correlations

at all scales, while other models over-estimate these correlations. The Cluster+Exp+Lumweight reaches minimal error on the smallest scales ($\lesssim 0''.4$), which shows that the more detailed reconstruction obtained with this model, visible in Fig. 2, is accurate over these small spatial scales.

We investigate the recovery of the size of the source galaxy through its effective radius, r_{eff} , which we define as the radius that encloses half of the total light within a circular aperture of radius $2''/2$. The middle panel of Fig. 4 shows r_{eff} and its relative error with respect to the true value for all source models. The effective radius is well recovered by all models with a maximum error of 2.5%. In addition we observe a tendency to over-estimate the effective radius, as only the Cluster+Exp model slightly under-estimates it, which is also the model with the lowest error. However, a better quantification of these errors should involve posterior distributions over the source models, which we do explore in this work. A first order uncertainty quantification can be obtained through the scatter among the different models, shown as the shaded gray area in Fig. 4. Given this scatter, the average source half-light radius remains lies close to 1σ from the true value.

Over the different source models considered here, none explicitly parametrizes the ellipticity of the source galaxy, in particular its axis ratio. Therefore, we use central moments of source model images (projected onto the same coordinates), in order to empirically measure an axis ratio q_s . More specifically, we compute the second order central moments of the source image, and use its eigenvalues to estimate the major and minor axes, from which we obtain the axis ratio q_s . The rightmost panel of Fig. 4 shows the resulting values, along with the true value measured on the true source image. The relative error is overall larger than for the effective radius although it remains below 6%. Taking the mean over the ensemble of modeling methods lies very close to the true value, within the 1σ scatter among the models.

Figure 5 shows the histogram of source pixel intensities for each model, compared to the true intensities (after interpolation, i.e., rightmost column of Fig. 2). Interestingly, we clearly see that the Sérsic+Shapelets and Cluster+Exp+Lumweight models reach higher intensity values. This is expected for the former (Sérsic+Shapelets) as the Sérsic profile diverges (with a best-fit Sérsic index ≈ 1.6) in its center and thus can predict large flux values. The luminosity-weighted regularization of the latter (Cluster+Exp+Lumweight) is behaving similarly by decreasing the regularization strength in regions with high observed flux. Consequently, these two models are best at capturing the three most magnified images of the source (see third column of Fig. 2). We also note from Fig. 5 that some of the models exhibit slightly negative values, as those are not penalized in their underlying prior, although these value are not statistically significant when compared to the noise level, and located mainly on the outskirts of the reconstructed source.

6.3. Recovery of lens properties

We investigate the different constraints on the mass distribution of our simulated strong lens system obtained from the different modeling methods. Already on Fig. 3, showing the predicted tangential caustics, we can visualize slight differences among the models. The predicted caustics have different sizes and positions, overall slightly larger than the true caustics. These differences correlate with the corresponding biases we discuss below. In particular, larger predicted density slopes are responsible for

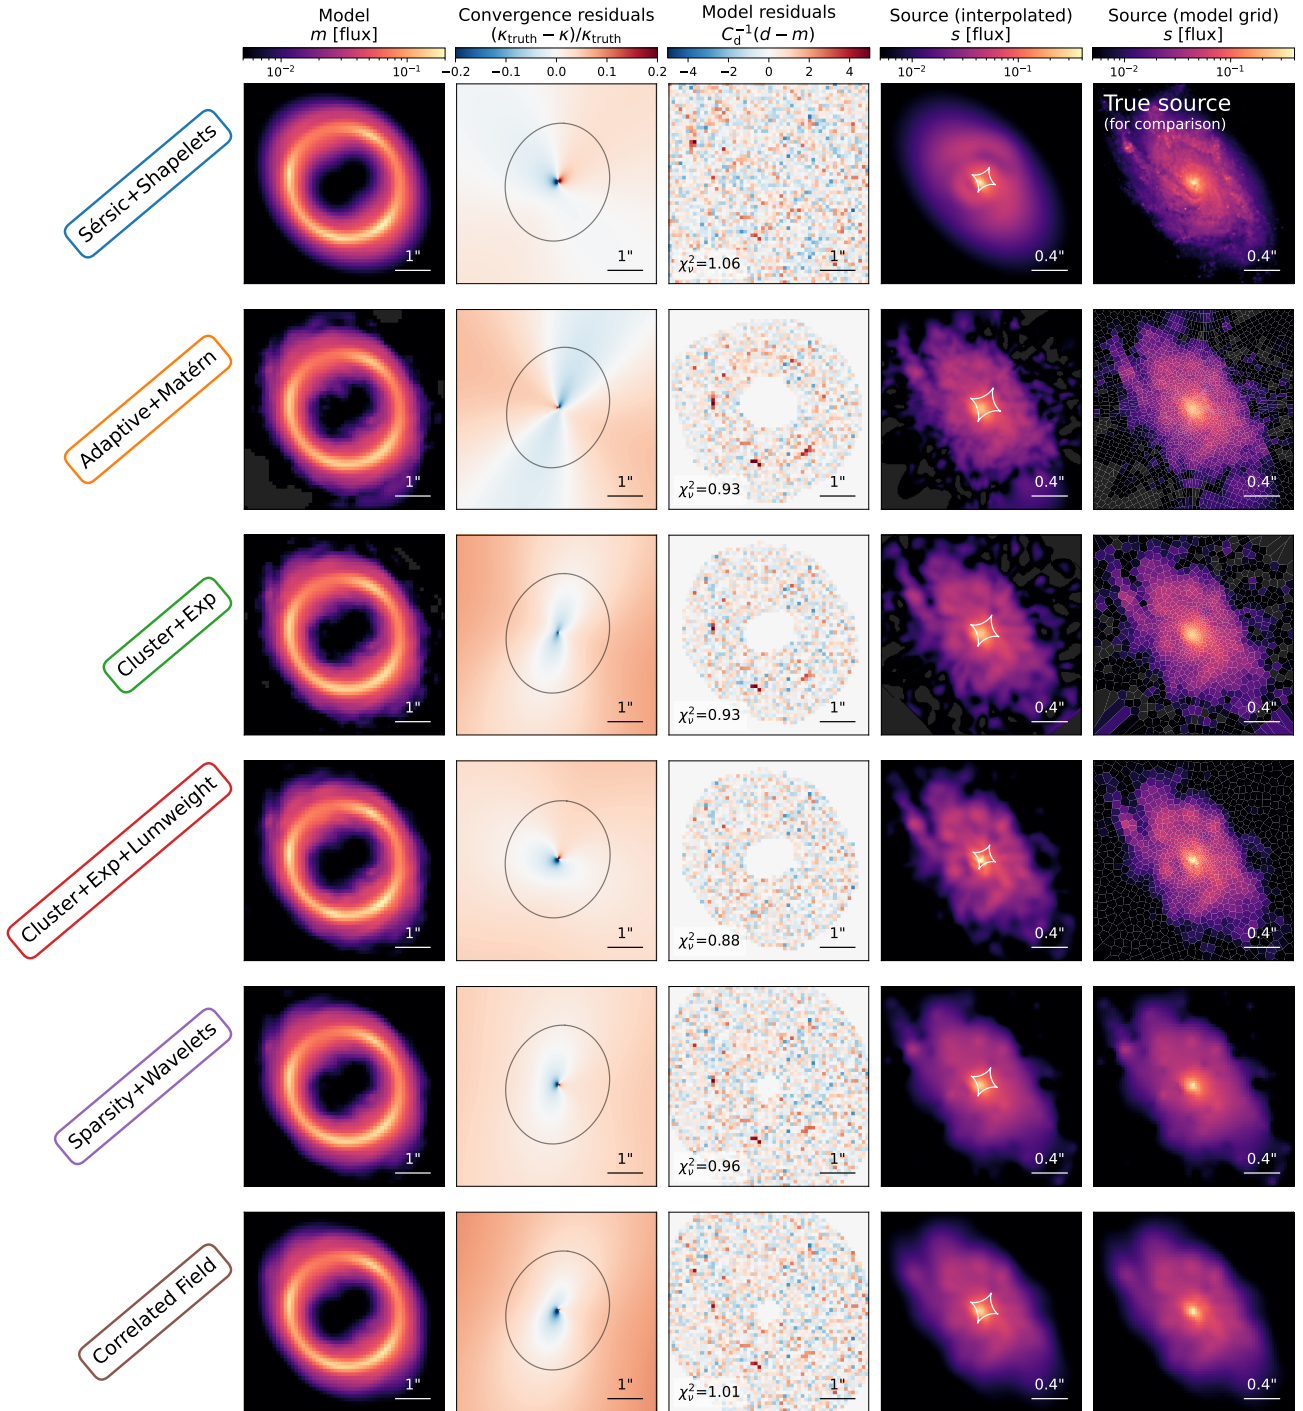


Fig. 2. Comparison between all blindly submitted models of the data shown in Fig. 1. The leftmost area gives the labels associated with each modeling method (see also Table 1), as well as their associated color used for the subsequent figures of this paper. First column: Image model. Second column: Relative difference between the true and modeled convergence maps, with the predicted tangential critical line shown as a black line. Third column: Normalized model residuals, with associated reduced chi-squared χ_v^2 indicated in the bottom left. The white areas are outside the likelihood mask chosen by the modelers and are thus excluded during modeling. Fourth column: reconstructed source models, all interpolated onto the same (regular) high-resolution grid to ease visual comparison. The predicted tangential caustic is also indicated as a white line. Last column, first panel: True source as in Fig. 1, to ease comparison with the models. Last column, remaining panels: Reconstructed source models on their original discretization grid, which can be regular or irregular, when applicable (the Sérsic+Shapelets model is not defined on a grid). All panels have been generated using COOLEST routines from the standardized storage of each model.

increasing the caustic size, and lens ellipticity and position offsets both impact the position and orientation of the astroid.

To quantitatively compare the constraints on lens potential parameters among the six models, we show in Fig. 6 the joint

posterior distributions for all mass model parameters, as well as true values from the data. Overall we find that the posterior distributions are within $\sim 3\sigma$ from the truth. The parameter with the smallest scatter relative to the posterior width is the logarithmic

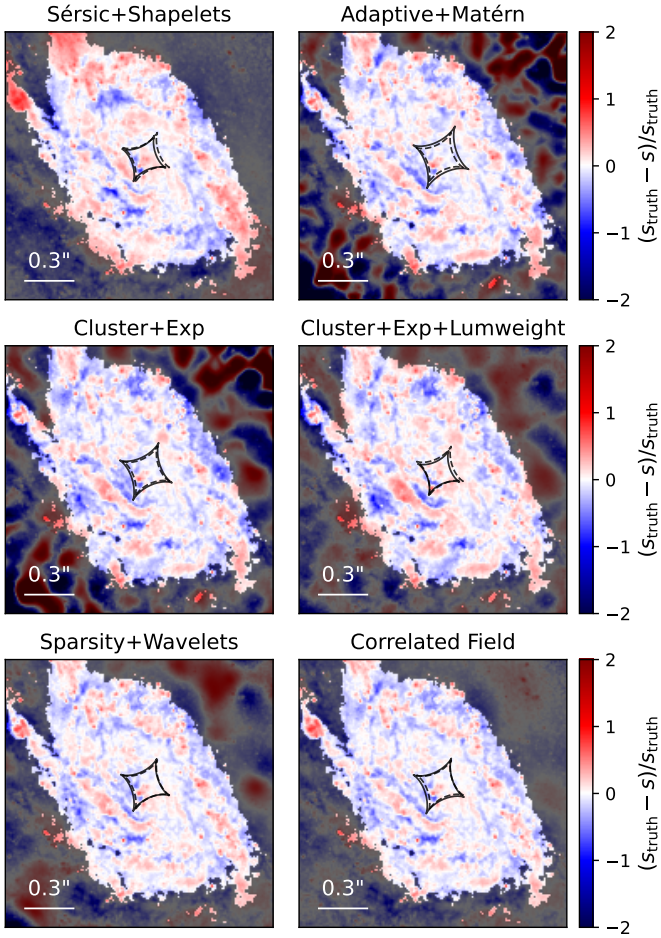


Fig. 3. Source plane residuals between the true source and the models shown in the fourth row of Fig. 2. Away from the center, the true source intensity is close to zero and models are less accurate, thus these areas are darkened for better visualization. Caustics from the best-fit (solid lines) and true (dashed lines) lens models are shown as well. Within the caustics, all models overall recover the source structure.

density slope γ . While all models are slightly biased towards values larger than the true value by approximately 2%, they are all compatible with it at $\leq 1\sigma$. Such a low scatter in the slope may be perhaps surprising, as it has been shown that the density slope may differ significantly between models (e.g., Etherington et al. 2022; Tan et al. 2024). However, we are in a regime where both the data and all models are parametrized by a single power-law density profile, hence the data is by construction a realization of the true model (modulo our inexact knowledge of the PSF and different numerics settings). Besides the density slope, the Einstein radius, θ_E , also has low scatter among the models, which is expected as it is the primary quantity constrained by strong lensing observables. Models whose median values are further away from the truth also tend to have broader posteriors (larger uncertainties), which contribute to reduce the systematic bias (see also Fig. 7 below). The remaining mass model parameters all show visible scatter around the true values, while no systematic shift can be associated with a specific modeling method or model parameter.

From the full joint distributions over lens potential parameters, we investigate further the difference in uncertainties between the models. In Fig. 7 we plot posterior standard deviations for four parameters, θ_E , γ , q , and γ_{ext} . As is expected

from imaging lensing data, the uncertainty on θ_E is the smallest, with a relative precision on the order of 0.1%. For comparison, the relative precision on the mass density slope γ is around 1.2%, and around to 1.5% for the lens axis ratio q (the relative error for γ_{ext} is inconclusive since it is close to zero). We notice that models with largest dynamic ranges (see Sect. 6.2) in their reconstructed source (Cluster+Exp+Lumweight) have smaller uncertainties on lens potential parameters. This trend is particularly clear for θ_E , γ , q (first column in Fig. 7). Over the parameters shown in Fig. 7, the difference in uncertainties between the models remains relatively small, and amounts to a factor of approximately 1.6 between the least and most precise models.

Ensemble models – namely, the combination of the posteriors of multiple models – can help improve modeling accuracy by correcting for the observed systematic biases of individual models. Depending on the assumptions about the original individual models, in particular regarding their statistical independence, there exist different approaches to combine their posteriors together. Here we follow a conservative approach and simply combine individual posteriors with equal weights. We show with dash-dotted black lines and contours in Fig. 6 the resulting combined posteriors. We find that these combined distributions are all within 1σ from the true values, except for the lens center along the x direction which is at 1.7σ . The marginalized statistics of the combined posterior are reported in the last row of Table 2 and compared to the simple average and standard deviation among individual models.

We quantify the improvement in systematic bias between individual models and the combined model. In the fourth row of Table 2, we list the largest bias (in units of standard deviation) that arises among the six lens models, for each lens potential parameter. As already seen in Fig. 6, the most biased parameters appear to be the coordinates of the center of the power-law profile, (x_0, y_0) , while the least biased is the density slope γ . As a comparison, the last row of Table 2 lists the corresponding bias values of the combined posterior, showing as expected a substantial decrease for all mass model parameters. We discuss further these results in Sect. 8.

6.4. Correlations between the recovery of lens and source properties

Having in hand multiple lens models of the same data, we have the opportunity to explore how the accuracy of inferred lens and source properties correlate among the models. In particular, it is interesting to understand if certain biases observed in lens potential model parameters have an origin in biases in the reconstructed source light distribution, and vice versa. If such correlations exist, they could be used to design better parametrizations to jointly model the lens and source components that specifically break these degeneracies. Such degeneracies may also be broken using non-lensing observations to further improve the accuracy of inferred lens and source properties (e.g., stellar kinematics of the source galaxy to place complementary constraints on its morphology). We emphasize that such correlations may be system- and data-dependent, which would warrant additional analyses complementary to ours.

We show in Figs. C.1–C.3 a series of scatter plots that correlate the relative error on lens potential parameters with the relative error on a given source property. We compute uncertainties on lens potential parameters from their posterior standard deviation. As we do not have such posterior distributions for all source models, we assume a fiducial uncertainty based on the Corre-

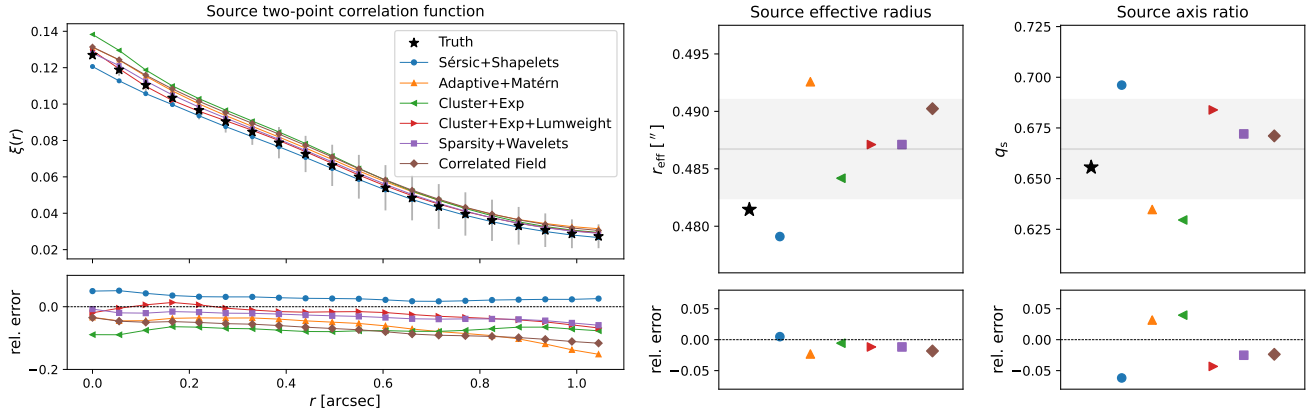


Fig. 4. Recovery of source properties among the six blind lens models shown in Fig. 2 (the legend in the right panel is also valid of other panels). Left panels: Azimuthally averaged two-point correlation functions, $\xi(r)$, of the recovered sources, compared to the truth, which measures the correlation between pixels separated by a given distance, r (in arcsec). Middle panels: Source effective radius r_{eff} , computed within a circular aperture. Right panels: Source axis ratio, q_s , estimated from the central moments of the source model. The bottom panel in each column indicates, the relative error with respect to the true quantity, computed as $(\text{truth} - \text{model})/\text{truth}$. For r_{eff} and q_s , the shaded gray region indicates the mean and 1σ scatter among the models.




Fig. 5. Histogram of pixel intensities for each reconstructed source shown in Fig. 2. Some models allow for slightly negative intensities, while some can capture very high intensity and compact source features.

lated Field model, for which we have posterior samples (see Sect. 3.5).

To quantify possible correlations, we compute the biweight mid-correlation coefficient r , indicated on each panel in Figs. C.1–C.3. The uncertainty on r is estimated by drawing 1000 random samples from bivariate uncorrelated Gaussian distributions centered on each data point. Based on the biweight mid-correlation coefficients, the largest (anti) correlation arises between the x -coordinate of the lens centroid x_0 and the axis ratio of the source q_s , with $r = -1.0 \pm 0.2$. On the other hand, the strongest absence of correlation is seen between the external shear strength γ_{ext} , and the total source magnitude m_s , with $r = 0.1 \pm 0.2$. We discuss and interpret the observed correlations in Sect. 8.

7. Investigating sources of systematics

While a thorough investigation of all possible sources of systematics is beyond the scope of this study, we nevertheless attempt to assess the impact of some key modeling assumptions and data properties. The results of this effort will be useful in guiding future in-depth investigations. Specifically, we explore the role

of the knowledge of the lens position, the presence of small-scale high-contrast regions (cusps, or point-like features) in the light profile of the source galaxy, and imperfect knowledge of the PSF.

7.1. Intrinsic source morphology

The spiral galaxy light profile that was used as the lensed source in producing the mock data presented in Sect. 4, has a prominent bright spot its center. The best-fit Sérsic index from the Sérsic+Shapelets model is approximately 1.6, indicative of a cuspy radial profile. This feature, which we shall refer to as the source cusp in the following paragraphs, consists of only a handful of pixels (<20) that contain a significant amount of light ($\sim 5\%$). This region is clearly hard to model, as is shown by the reconstructions in the last two columns of Fig. 2, where only two models are able to capture it (Sérsic+Shapelets and Cluster+Exp+Lumweight). The remaining models fail to do so and leave behind characteristic residual flux at the data pixels where this compact region is multiply imaged.

Driven by this observation, we argue that this cusp is closer to a point-like flux component than to an extended source. Some algorithms, like the plain semi-linear inversion (even on an adaptive grid), have been explicitly designed to model the latter and are known to have a poor performance with the former. This can be understood in terms of the regularization, which tries to impose smoothness on the source and inevitably suppresses such cusps. In Fig. 5, we see that only a few source pixels with the highest flux (>0.3), corresponding to the central cusp, are considerably extending the dynamic range of the source light profile. The two aforementioned models that successfully model the cusp have a similar dynamic range, while the rest of the models fall short, barely reaching a flux of 0.3.

Here, we examine in more details how the performance of the model with the lowest brightness range (smoothest), Adaptive+Matérn, is affected by the prominence of the cusp. We select the central bright region with flux >0.3 and reduce the flux⁹ in each pixel by 50 and 95%, creating two new mocks that

⁹ The noise map that is used as the covariance matrix in Eq. (5) is also changed accordingly.

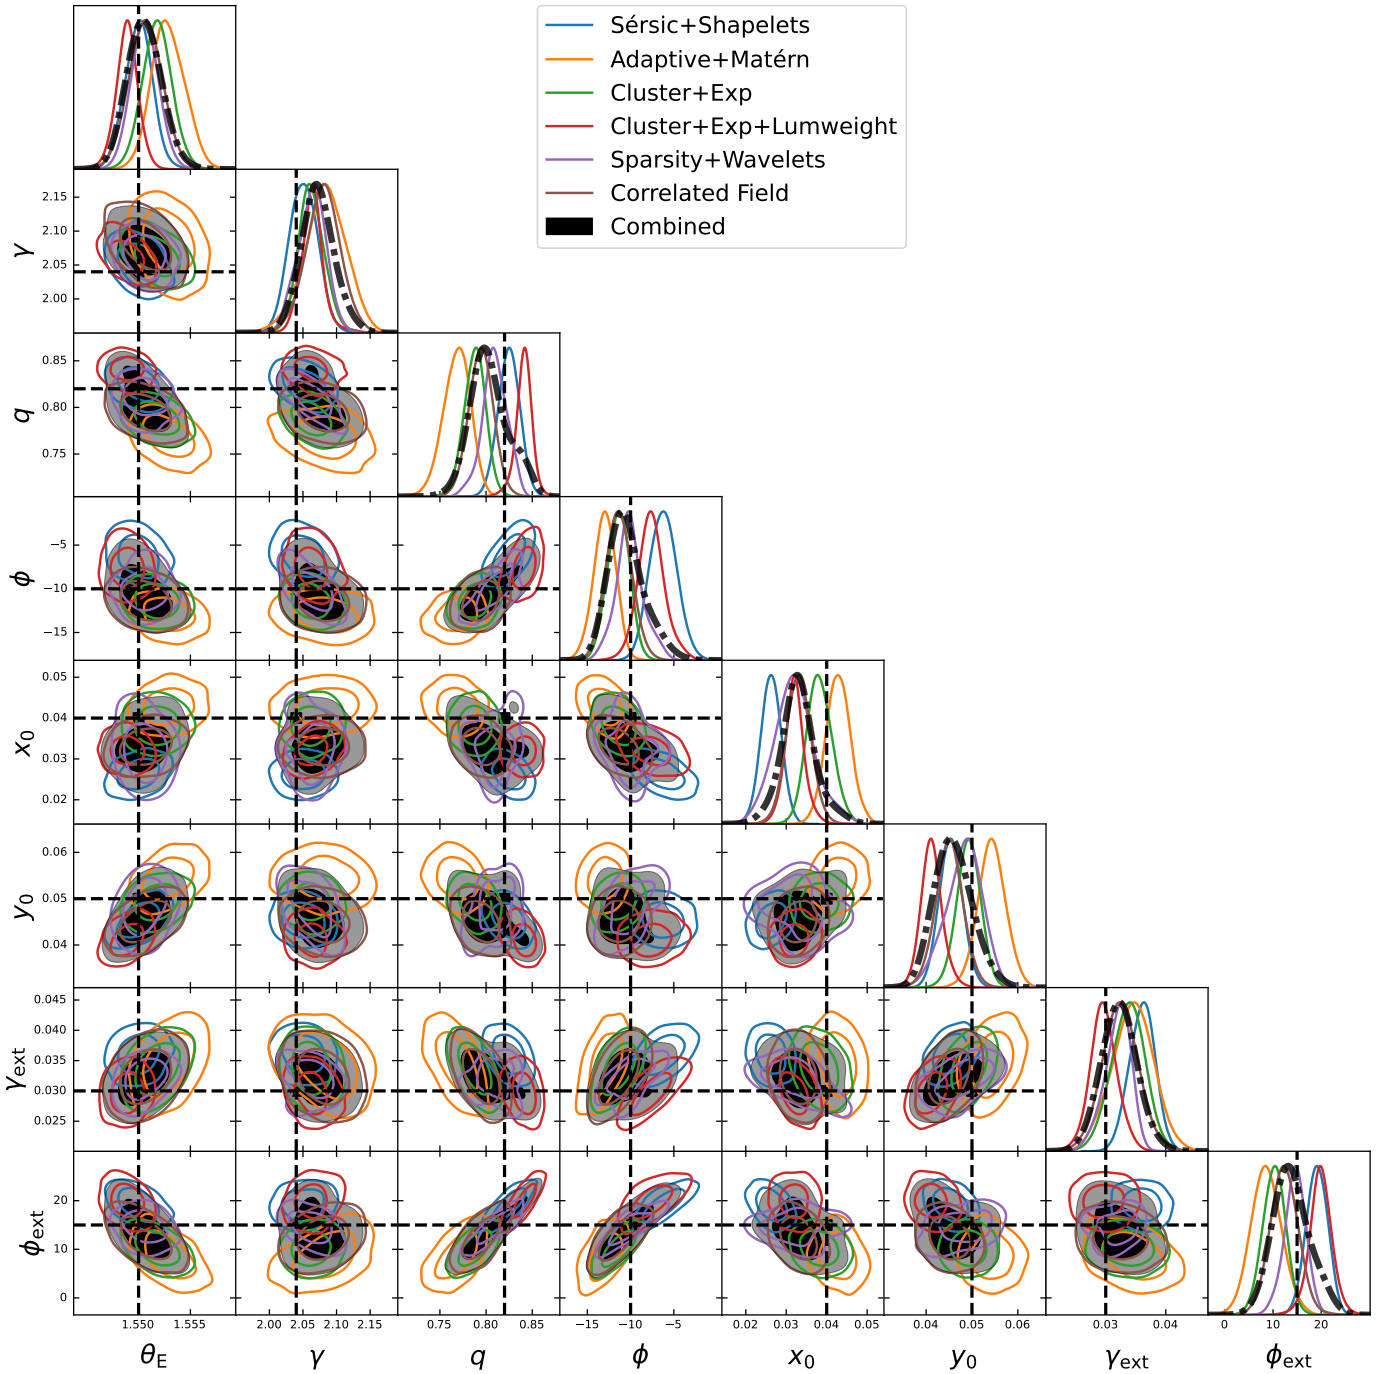


Fig. 6. Posterior distributions of the lens potential parameters, inferred blindly from the simulated data shown in Fig. 1, with the true values indicated by the dashed lines. The parameters from left to right are: the Einstein radius, mass density logarithmic slope, axis ratio, position angle, lens center coordinates, external shear strength, and external shear orientation. Note that we use the COOLEST definitions of these parameters.

we then model with exactly the same setup as the model shown in Fig. 2. The resulting (true) source and model residuals are shown in Fig. 8, while Fig. 9 shows the posterior distributions of the lens potential parameters. It can be seen that the more we suppress the cusp the better the model; that is, we get less residual flux and less biased parameters.

We conclude that a plain semi-linear inversion approach is much better suited for modeling smoother sources, without cuspy, point-like features in their light profile. A regularization scheme imposed just by Eq. (10) leads to reconstructions that are too smooth, and additional constraints, like the luminosity-

weighted scheme presented in Sect. 3.3, give better results. The cuspy nature of the source can thus explain, at least partially, the systematic errors in lens potential parameters, in particular for the case of the Adaptive+Matérn model.

7.2. Supersampling and inexact PSF model

We have seen in the previous section that a cuspy source can lead to biases in the lens potential parameters even in models that fit the central cusp well, in particular the Cluster+Exp+Lumweight model which achieved noise-level residuals. We now investigate

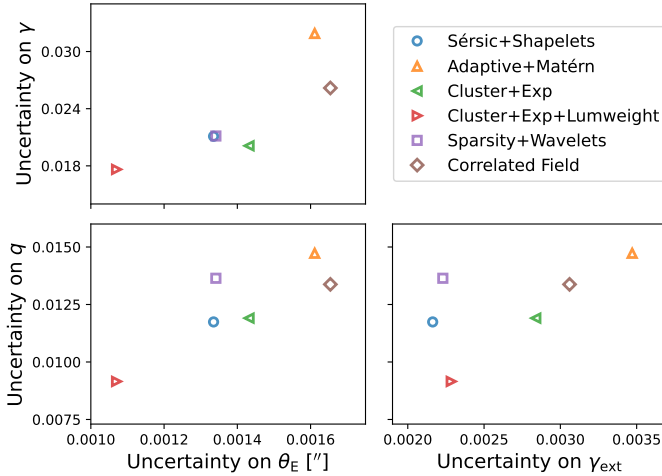


Fig. 7. Uncertainty on a subset of key lens potential parameters, defined as the standard deviation of their marginalized posterior distribution (shown in Fig. 6): the Einstein radius, θ_E , the mass density slope, γ , the axis ratio q , and the external shear strength γ_{ext} . The Einstein radius and the mass density slope are central quantities to many analyses, while the lens axis ratio and external shear are known to share degeneracies.

whether using a supersampled PSF can further eliminate this bias. We implement PSF supersampling in two different ways: first, by interpolating in the pixel-level PSF (which is often implemented in practice when point sources are present in a lensed source); and second, by using the original supersampled PSF that was used to produce the mock data. Although the original PSF corresponds to a supersampling factor of 10 (i.e., each pixel is split into 10×10 subpixels), this would be much too computationally expensive to employ directly. We therefore choose a supersampling factor of $f = 5$ and downsample the original PSF with $f = 10$ accordingly; this is what we shall refer to as the “true supersampled PSF.” In the first model, we use bicubic interpolation in the pixel-level PSF to generate our interpolated $f = 5$ supersampled PSF.

We first check whether PSF supersampling can eliminate bias in models without a luminosity-weighted regularization, by applying the above supersampling procedure with the “true supersampled PSF” to the Cluster+Exp model. The resulting posteriors are plotted as the dot-dashed green curves in Fig. 10. In this case we find a similar bias in the lens parameters as in the case where no supersampling is performed, which is perhaps explained by the fact that the best-fit model produces similar residuals in the regions where the lensed images are brightest. This is due to the fact that the reconstructed source is not significantly better resolved without the luminosity-weighted regularization prior; the regularization strength is too high to allow for a cuspy source in the central high-intensity region of the source.

Next, we apply PSF supersampling to the corresponding model with luminosity-weighted regularization (Cluster+Exp+Lumweight). The results are shown as the red curves in Fig. 10, with the solid red curve showing the original luminosity-weighted model (i.e., same as Fig. 6) and the dashed and dot-dashed curves showing the models with supersampled PSF and “true supersampled PSF”, respectively. Note that both supersampled models have significantly reduced bias compared to the original luminosity-weighted model that did not use PSF supersampling. Some bias is still present in the lens model parameters for the interpolated PSF model, particularly in the slope γ and ϕ_{ext} parameters, whereas in the “true PSF” model,

bias is largely eliminated in all lens parameters except for the center coordinates. In addition, the parameter uncertainties are significantly reduced when the true PSF is used. The essential difference can be seen by comparing the two supersampled PSF’s directly in Fig. 11, where the interpolated PSF is quite poorly resolved. We conclude that for sufficiently cuspy sources such as this one, supersampling can significantly reduce bias in the lens model parameters, but may not entirely eliminate bias in the lens parameters if one generates a supersampled PSF by interpolating in the observed pixel-level PSF.

It is interesting that without a luminosity-weighted source prior, the lens center coordinates are *more* accurately recovered (regardless of whether supersampling is used), despite all the other lens parameters being significantly biased. The Sérsic+Shapelet model produced a similar bias in the lens center coordinates as the luminosity-weighted models did, which is noteworthy since these are the only models that were able to reproduce the central cusp in the source well. While it is unclear exactly why this is the case, in real applications this may be ameliorated by the fact that the foreground lens light can furnish a prior in the lens center coordinates. Aside from this caveat, the bias in these parameters is at least somewhat reduced by supersampling with the true rather than interpolated PSF. We conclude that when fitting cuspy lensed sources, PSF supersampling can significantly reduce parameter biases only if it is accompanied by a regularization scheme (e.g., luminosity-weighted) that allows the source pixels to have steeper variations in the central bright region of the source galaxy where PSF supersampling is of the greatest benefit.

8. Discussion

8.1. Quantifying model complexity and its impact on posterior uncertainties

In lens modeling there is inherently a trade-off between model complexity and tractability of the final inference. On the one hand, a more complex model – namely with more model parameters, or degrees of freedom – is likely to provide a better fit to the data (commonly measured as a smaller χ^2 value) with the risk of over-fitting. On the other hand, a simpler model is usually faster to optimize and often leads to a more robust inference (lower risk of local minima and multi-modal posteriors) but may not fit the data well. Moreover, models with different complexity can still fit the same data seemingly equally well; in such a case, one typically invokes the principle of Occam’s razor, and prefer the one being the least complex a priori (which, when feasible, takes the form of the Bayesian evidence or other proxies such as the Bayesian information criterion).

In our work, as is described in Sect. 3, one of the main differences between the modeling methods we consider is the way the source galaxy is reconstructed. In the past, several works focused on consistently comparing a set of lens models with different source reconstruction techniques developed under the same general formalism (i.e., the semi-linear inversion formalism, see Sect. 3.3, and Warren & Dye 2003), but differing in their regularization terms (e.g., Suyu et al. 2006; Tagore & Keeton 2014). Unfortunately, in the present work, there is no clear way to quantitatively and unambiguously rank the complexity of the source models considered here, given their fundamental differences in terms of mathematical formalism and underlying assumptions about the morphology of galaxies. One possibility would be to count their number of degrees of freedom, but this quantity is not readily accessible for all models. In regularized pixelated source

Table 2. Posterior statistics of source and lens properties inferred from the six lens models shown in Fig. 1.

Quantity	$r_{\text{eff}} ["]^{(\dagger)}$	$q_s^{(\ddagger)}$	$m_s [\text{mag}]$	$\theta_E ["]$	γ	q_m	ϕ_m	$x_0 ["]$	$y_0 ["]$	γ_{ext}	$\phi_{\text{ext}} [^\circ]$
Truth	0.481	0.66	23.2	1.550	2.040	0.82	-10.0	0.040	0.050	0.030	15.0
μ_{models}	0.487	0.67	23.1	1.551	2.066	0.80	-9.9	0.034	0.047	0.033	13.9
σ_{models}	0.004	0.02	0.04	0.001	0.009	0.03	2.3	0.005	0.004	0.002	4.6
Max. bias $(\dagger\dagger)$	×	×	×	1.7σ	1.6σ	-3.3σ	2.3σ	-5.5σ	-3.7σ	2.7σ	2.2σ
Combined	×	×	×	$1.55^{+0.002}_{-0.002}$	$2.07^{+0.03}_{-0.02}$	$0.80^{+0.03}_{-0.03}$	$-10.5^{+1.4}_{-2.6}$	$0.033^{+0.004}_{-0.004}$	$0.046^{+0.004}_{-0.005}$	$0.032^{+0.003}_{-0.003}$	$13.8^{+3.3}_{-4.5}$
Bias $(\dagger\dagger)$	×	×	×	0.3σ	1.2σ	-0.6σ	-0.2σ	-1.6σ	-0.9σ	0.7σ	-0.3σ

Notes. In the second and third row μ_{model} and σ_{model} show the average and standard deviation of the posterior means of all models, respectively. The fourth row gives the maximal bias among all models for each lens potential parameter. The two last rows correspond to the combined posterior distribution shown in black in Fig. 6. (\dagger) The source properties correspond to best-fit models, whereas lens properties relate to the corresponding marginalized posterior distributions shown in Fig. 6 (individual and combined models). Note that we do not provide combined estimates for the source properties as we do not have formal posterior distributions not uncertainty estimations for all source models. $(\dagger\dagger)$ The maximum bias among all models is based on the difference between the mean posterior value and the true value, in units of the posterior standard deviation (negative and positive biases correspond to under- or over-estimates, respectively). The last row of the table shows the corresponding bias value from the combined posterior distribution.




Fig. 8. Model residuals from the Adaptive+Matérn model and central region of the true source (top row, as in Fig. 2). Suppressing the flux in the central most bright pixels (>0.3 in Fig. 5) by 50 and 95% (middle and bottom rows), creating new mocks, and modeling them with the same setup, leads to improved residuals. The plain semi-linear inversion technique using just a regularization of the form given in Eq. (10) cannot adequately capture point-like, cuspy features in the source light profile. More advanced schemes, like a luminosity-weighted regularization scheme (see Sect. 3.3), perform better in this case.

models, the effective number of degrees of freedom is lower than them total number of source pixels, as regularization correlates source pixels over different spatial scales (e.g., Suyu et al. 2006; Nightingale & Dye 2015). As a concrete example, the regularization based on sparsity and wavelets do not allow one to unambiguously estimate the number of degrees of freedom, because

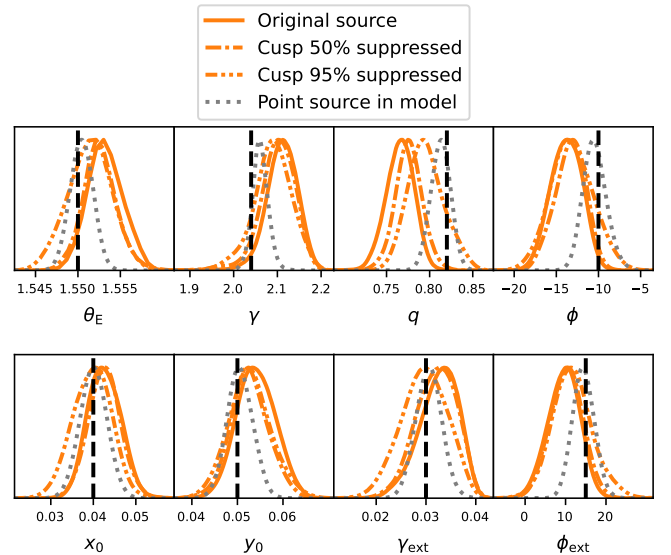


Fig. 9. Posterior distributions of the lens potential parameters for the three models shown in Fig. 8. The fourth model (bottom in the legend) corresponds to a model of the original mock after including a point source component in the source model. When the cuspy central region of the source is suppressed from the data, or if a point source feature is added in the source model, the resulting distributions are less biased and shift closer to the true values.

such regularization imposes sparsity simultaneously over various spatial scales.

Nevertheless, our ensemble of six models allow us to qualitatively discuss how differences in model complexity can affect the resulting inference. In Sect. 6.3 we compare the blindly submitted posterior distributions, and notice some differences in the inferred parameters uncertainties (in terms of the posterior standard deviation). Based solely on the reconstructed sources (see Fig. 2), the model that displays the least complex features is Sérsic+Shapelets. Oppositely, one of the models that capture the most complex features is Cluster+Exp+Lumweight. One may naively expect the Cluster+Exp+Lumweight model to be more affected by issues related to over-fitting and local minima, as it is defined on a locally fine grid and has more flexibility compared to the Sérsic+Shapelets. In case of over-fitting, the uncertainties on model parameters can be significantly under-estimated, partly

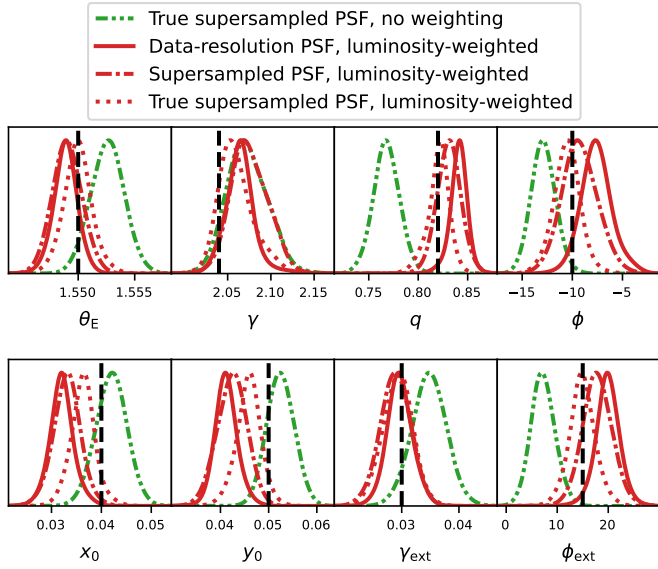


Fig. 10. Posterior distributions of lens potential parameters obtained with models exploring the role of the PSF. The dash-dotted green line distributions correspond to the Cluster+Exp model (i.e., similar to the green model in Fig. 6) using the true supersampled PSF used for simulating the data (top left panel of Fig. 11). The solid red line distributions are showing, for reference, the blindly submitted model Cluster+Exp+Lumweight (i.e., the same model as in Fig. 6). The dash-dotted and dotted red distributions, also obtained with the Cluster+Exp+Lumweight model, use a supersampled (interpolated) version of the data-resolution PSF (bottom left panel in Fig. 11) or the true supersampled PSF, respectively. The main result of this comparison is that biases in lens model parameters are most reduced only with a combination of a more accurate PSF and a model that can capture magnified cuspy features in the source.

due to discretization biases that artificially narrows the likelihood profile (e.g., [Nightingale & Dye 2015](#); [Etherington et al. 2022](#)). Interestingly, while Cluster+Exp+Lumweight indeed lead to overall small uncertainties on lens potential parameters (compared to other models, see Fig. 7), we see that they do not significantly differ from the Sérsic+Shapelets model. Therefore, while we observe variations among the models regarding their posterior uncertainties (up to a factor of ~ 1.6), these variations are not solely driven by differences in model complexity. This result also highlights the difficulty in measuring the complexity of model simply based on the modeling assumptions and forms of regularization, or on the visual impression of the reconstructed source. However, this kind of analysis offers directions to explore for better understanding the origin of systematic biases in lens potential parameters.

8.2. Model degeneracies between the lens and the source

As is presented in Sect. 6.4, we observed correlations between the level of recovery (the relative error) of lens and source properties. In particular, we find a correlation between the lens power-law slope γ and the source effective r_{eff} , with correlation coefficient $r = 0.9 \pm 0.4$, namely a positive correlation with a statistical significance of 2.3σ . It is well-known that correlations between the source scale and the mass density profile of the lens can arise, which are often seen as a manifestation of the MSD ([Falco et al. 1985](#)). The MSD originates from the mass sheet transformation (MST), and its consequence is as follows: the addition or subtraction of a infinitely thin mass

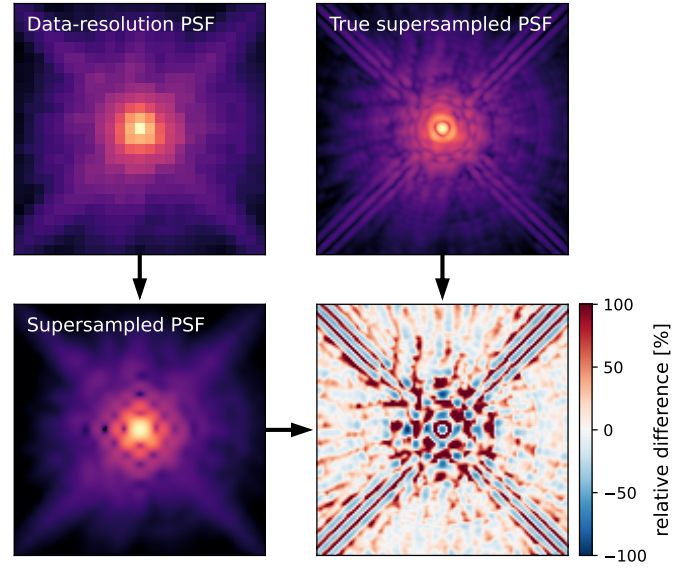


Fig. 11. Comparison between the different PSF kernels used in this work (each panel shows a zoom on the central $2''5$ of the PSF kernel). The top left panel shows the original true PSF used for simulating the data, while the top right panel shows the PSF given to the modelers (same as in Fig. 1). The bottom left panel shows the bicubic-interpolated (5 times supersampled) PSF used in the model shown with dotted lines in Fig. 10, and the bottom right panel shows the relative difference with the true PSF. Except for the bottom right panel, all color scales are arbitrary and chosen to help the visual comparison between the PSF kernels.

sheet from a power-law mass density profile changes, to first order, the slope of the density profile at the Einstein radius (e.g., [Schneider & Sluse 2013](#); [Blum et al. 2020](#); [Birrer et al. 2020](#)), while rescaling the source proportionally. A positive mass sheet locally increases the density slope, which in turns induces an increase in the source size through the MSD. This effect has also been empirically explored in [Birrer et al. \(2016\)](#) by using the explicit scale encoded in a source model based on shapelets. The positive correlation we find between the biases in γ and r_{eff} may thus be the signature of the MSD: all the lens models we consider infer slightly too large density slopes and source effective radii.

There is a general trend of correlations between source light shape and lens mass shape that we observe in our results. For instance, we note correlations between the recovery of q_s and q_m or ϕ_{ext} , r_{eff} , and anti-correlations with ϕ and the lens centroid. We also observe strong correlations between the error on q_s and that of the lens centroid. These correlations are more challenging to unambiguously interpret compared to those associated with the MST. We argue that they may be related to the source position transformation (SPT) outlined in [Schneider & Sluse \(2014\)](#) and further developed in [Unruh et al. \(2017\)](#). The latter work effectively shows that transformation of the source profile can be compensated by changes in ellipticity and radial profile of the lens mass distribution. The biases that we observe in our models may be manifestations of this transformation and qualitatively match the expectation for a SPT. Because the SPT leads to an approximate degeneracy (unlike the MSD), it is expected to be broken with higher quality data. It may therefore be interesting to see how these correlations change with higher S/N data or even noise-free mock data.

More generally, approaches that are independent of a specific choice of mass model offer a complementary route to

further study the interplay between lens and source structure. As reviewed in [Wagner \(2019\)](#), and references therein), instead of the direct modeling of imaging data pixels, a careful extraction of specific lensing observables can provide constraints onto local lens properties, without assuming global functional form of the lens potential. While such an approach is scale-independent by construction – it does not rely on a priori assumptions on the nature of the lensing object (isolated galaxy, cluster member, etc.) – it requires the ability to unambiguously extract lensing features like well-defined multiple images, their shape and orientation, or individual clumps within lensed arcs (e.g., [Wagner et al. 2018](#)). In galaxy-scale strong lenses such as the simulated one of [Fig. 1](#), the observed images take the form of relatively smooth highly distorted arcs (i.e., rings), for which it is error-prone to attempt extracting local features of multiply imaged regions of the source without performing full lens modeling ([Galan et al. 2024](#)). Nevertheless, the local lensing formalism can be expanded to lensed arcs with globally smooth surface brightness, as developed in [Birrer \(2021\)](#). Such an approach is well-suited to make use of the multiple elongated arcs observed around massive galaxy clusters, as a way to locally correct the global mass model (e.g., [Yang et al. 2020](#)).

8.3. Combining methods to mitigate systematic biases

We quantify the scatter associated with the choice of modeling method in [Table 2](#), which records the mean and scatter of relevant source light and lens mass quantities among all the models. Although dependent on the type of data considered in this work, these numbers provide an estimation of systematic errors that one can expect from modeling imaging data with different methods. The systematic errors we quote here can be interpreted as lower bounds to those of a real-case scenario, since we placed ourselves in an idealized setting – no contamination by the lens light and perfect knowledge of noise properties and mass model family – which removes a subset of the known sources of biases or degeneracies. Ideally, the addition of such complicating factors may broaden the posterior distributions from individual models, thus making them statistically compatible (underlying systematic biases would then be unnoticeable). However, more realistically, the scatter between models is likely to increase as a consequence of more complex models potentially subject to different sources of biases. Performing similar analyses as the one presented here to a wider variety of strong lensing data (different resolution, S/N, lensing configuration, nature of the source, etc.) will help understand the generalization of our results. For this purpose, the framework we have developed should help and encourage the multiplication of such analyses on both simulated and real data sets.

As we show in [Fig. 6](#) and [Table 2](#), combining together the results from an ensemble of methods using uniform weights removes the observed biases. While individual methods display systematic errors in different parameters – and not necessarily always for the same parameters and in the same directions – it is reassuring that overall, we do not observe a significant residual offset after combination. Only for the mass density slope (γ) one can observe a global trend towards larger values. We quantify in [Table 2](#) the bias reduction between the biases from individual models and the one from the combined model. Among the eight mass model parameters, we find that an average bias reduction of 5.4, which is a substantial improvement on inferences from individual models alone. This result is reassuring and shows that analyzing a given data set using independent modeling methods is an efficient way to mitigate systematic errors.

In this work, we conservatively assumed equal weights when combining individual methods. This is similar to the recent work of [Wong et al. \(2024\)](#), where the authors combined measurements of H_0 obtained from by modeling the same lensed quasar with two different modeling codes, assuming equal weights. Nevertheless, we note that other approaches exist. The least conservative approach would be to multiply individual posteriors together, which is equivalent to assume that all methods are independent from each other and do not share systematic errors. The six modeling methods are only partially independent, as some of the models were performed by the same modelers using the same modeling code (see [Appendix B](#)) albeit with different source reconstruction techniques. Additionally, we modeled a single imaging data realization of our simulated strong lens, which further introduces some degree of statistical dependence. The most Bayesian approach would be to combine individual models based on their Bayesian evidence. However, for reasons we detailed in [Sect. 8.1](#), objectively ranking models obtained from methods based on fundamentally different assumptions is still an unresolved issue. Nonetheless, better quantifying how different strong lens modeling methods perform on identical or similar datasets with respect to their complexity will be important to design better statistical combination procedures.

Beyond the accuracy improvement of a combined posterior distribution over the lens model parameters, the comparison of multiple models – and if possible, truly independent models – is extremely valuable for detecting unknown sources of systematics, updating our current modeling assumptions and techniques and testing additional models with different levels of complexity. Similarly, in source plane, comparing the morphology of different versions of the unlensed object (which is never directly observable) is necessarily valuable for straightening the confidence in a given feature, which in turn improves the robustness of its interpretation. Moreover, different models may predict different observables that may otherwise be overlooked, enabling the potential detection of anomalous lens systems (such as missing images, [Ertl et al. 2024](#)) or anticipate the detection of additional images unobserved in current observations (e.g., similar to the geometric confirmation of multiple images in cluster-scale lenses, [Diego et al. 2023](#)).

8.4. Extrapolating to time-delay applications

As lens modeling is a key ingredient in time-delay cosmography applications, one may be interested in the propagation of our results to the recovery of the Fermat potential difference (defined in [Eq. \(4\)](#)). The Fermat potential difference between two images i and j is proportional to the time delay measured if the source is varying in time. Such a source can be a quasar centered on its host galaxy and orders of magnitude brighter than the lensed arcs, or supernova (SN) which can appear almost anywhere in its host and fades away after the explosion, leaving behind only the arcs as in our simulated data.

We show in [Fig. 12](#) the Fermat potential differences for three hypothetical pairs of lensed images. While the simulated data we analyze in this work ([Fig. 1](#)) mimics the case of a lensed SN that faded away, we assume for simplicity that its location in source plane coincides with its host galaxy (although one could also proceed similarly for any other source position that leads to at least two lensed images). We label the lensed images ABCD, order them by ascending Fermat potential (i.e., $\Phi(\theta_A)$ is the lowest) and consider the three independent pairs $ij \in \{AB, AC, AD\}$. In addition to the posterior mass parameters uncertainties, we add in quadrature an additional uncertainty to mimic the

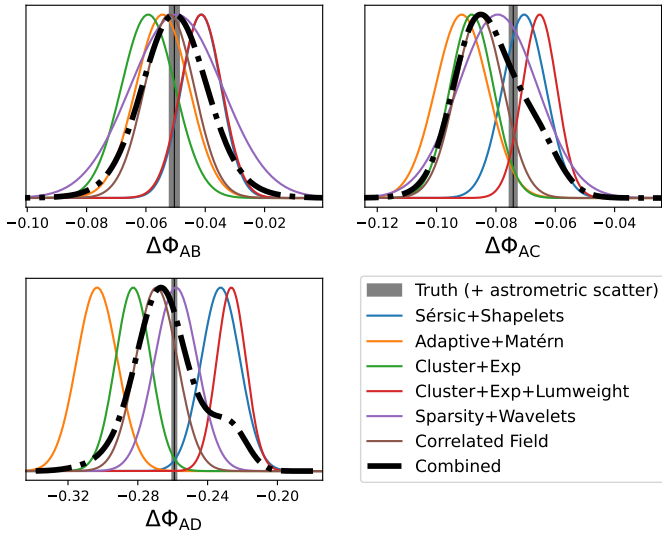


Fig. 12. Fermat potential differences evaluated at image positions of a multiply point source located at the center of the source galaxy. Given that the modeled data (Fig. 1) does not contain the multiply imaged point source, this can be seen as an intermediate scenario between that of a lensed quasar (i.e., centered on its host) and a lensed SN (i.e., has faded away).

limited astrometric precision (for details, see Birrer & Treu 2019), assuming a conservative precision of 10 mas in image plane. The shaded gray area in each panels of Fig. 12 around the true values isolates the typical contribution of the astrometric scatter term, which is in our case small compared to the uncertainties modeling the extended lensed source. As is expected from the constraints on mass profile parameters (Fig. 6), we observe a similar scatter around the true values among the models.

Interestingly, we do not observe a clear correlation between model biases on individual lens potential parameters and their resulting biases on Fermat potential differences. For example, the Sparsity+Wavelets model, which display slight biases on some source properties and lens potential parameters (e.g., the lens position), give Fermat potential posteriors well-centered on the expected values. While a robust generalization of these results require further work (e.g., on a larger sample of simulations), the general trend we previously found still remains: the combination of individual posteriors assuming equal weights results in posteriors that are free of biases. These combined distributions are shown with dash-dotted black lines in Fig. 12, and encapsulate well the Fermat potential difference values. This further gives motivation for time delay cosmography analyses to compare and combine lens models together, in particular those obtained using independent modeling methods (see e.g., Shajib et al. 2022; Kelly et al. 2023). In a follow-up paper, we shall explore the dependency of systematic lens modeling errors on the source plane position of the time-delay background object.

8.5. Application to real data sets

For the purpose of demonstration, we focus on simulated imaging data, but the framework and ideas we present is intended to be applied to real data. The COOLEST standard, which provides a common ground to express lens models from different origins, already supports real data and their corresponding mod-

els (see e.g., Galan et al. 2024). Moreover, extending the standard to other types of models (for both the lens and the source) is straightforward. One complication for applying the same framework presented here to real data arises from the human time and computation time required to apply multiple codes and methods to a given data set (see also Appendix B).

Although COOLEST significantly reduces the burden to express a lens model in a format that can be readily compared to other models, the acquisition of combined posterior distributions over lens model parameters still requires that at least two distinct lens models are in hand. Here we model the same data using four software packages and six modeling methods, which is likely unrealistic in most real scenarios, especially given the large amount of lenses that are still to be modeled (e.g., in the archival HST data, see Garvin et al. 2022) and that will be discovered in the near future (e.g., Collett 2015). Nevertheless, it is reasonable to assume that it is feasible to apply two or three methods to the same data (not necessarily within the same analysis), since the amount of lens modeling experts naturally grow over time and inference methods are becoming less and less time consuming. Deep learning methods, which we do not explore in the present work, offer promising avenues to accelerate the overall procedure either by proposing preliminary models to be refined with classical methods, or providing additional posterior distributions for final combination, at a negligible cost (ignoring the training phases).

When modeling real data, the chosen mass model parametrization is only an approximation of the true mass distribution. Extending the present analysis to account for differences between the truth and the model is obviously conceivable. However, it is important to note that such an extension inherently requires key assumptions in order to properly define “how different” is the truth compared to the model. A concrete example is the work of Gomer et al. (2022, Sect. 3.2) who simulated a population of SLACS-like strong lenses following a composite (baryons + dark matter) mass model, for which they had to assume a series of additional assumptions (ellipticity, halo scale radius, etc.), which they subsequently modeled with a power-law profile radially modified with a mass sheet parameter. Cao et al. (2022, Sect. 2.4) simulated strong lenses with mass distribution based on stellar kinematics models and performed modeling using single power-law profiles with shear. Van de Vyvere et al. (2022) quantified the biases caused by a lack of freedom along the azimuthal direction, when the true mass distribution contains multipoles or ellipticity gradients. The Rung 3 of the TDLMC, although focused on the recovery of H_0 instead of mass model parameters, has been a blind experiment where independent teams modeled lenses extracted from cosmological simulations (Ding et al. 2021). On real and simulated systems, the Time-Delay Cosmography collaboration (Millon et al. 2020) systematically uses both power-law and composite mass model families and marginalizes over the resulting posteriors, which are further constrained along the radial direction with stellar kinematics information (e.g., Shajib et al. 2022). Given the diversity of deviations to the commonly assumed sheared power-law elliptical mass profile in real galaxies – such as truncations (e.g., Limousin et al. 2007, in dense environments), cores (e.g., Collett et al. 2017) or multipoles (e.g., Stacey et al. 2024) – the generalization of the above-cited results to all strong lens systems is still unclear. It is possible that for a given real system, a set of dedicated simulations based on a set of likely mass models, followed by lens modeling as for the real data, may be the best way to properly quantify systematic errors caused by the choices of mass model.

9. Conclusion

We have conducted a fully blind modeling experiment on strong lensing data simulated with a dedicated software package by one author, while four other authors used four independent software packages to model it. In total, six modeling methods – which differ in their source reconstruction techniques and inference pipelines – have been applied on that same data. We have made a series of simplifying assumptions to keep this novel kind of analysis tractable, in particular regarding the lens light, the mass model parametrization and the noise properties. In contrast, the true PSF, lens parameters and source morphology were hidden and the optimization and inference strategies left free to the modelers. We used the lensing standard COOLEST to overcome the challenges that arise when comparing results obtained with different modeling codes. The resulting image and source plane models, as well as model residuals are given in Fig. 2, and constraints on lens potential parameters are shown in Fig. 6. Below we summarize our main results:

- While no modeling method resulted in strong statistical biases systematically for all lens and source properties, we observed a measurable scatter among the models. Strongest biases arise for the lens centroid, while the mass density slope at Einstein radius is only mildly biased with a small inter-model scatter. We also observe differences in the dynamic range of the reconstructed source intensities.
- Combining results from different modeling techniques enables to mitigate systematic uncertainties arising for individual models. For the particular data we consider, the systematic error on lens potential parameters is reduced by a factor 5 on average. The reason is that models tend to scatter around the true parameters values but stay statistically compatible, such that the combined posterior distributions effectively broaden and include the true values. This results also holds regarding the Fermat potential differences between hypothetical lensed images of a point source component, which is relevant for time delay cosmography applications. While the amount of bias reduction is evidently data-dependent, we argue that model combination is generally beneficial to remove some biases from strong lensing analyses.
- Towards the goal of better understanding the origin of model biases, we used our ensemble of models to investigate possible correlations between lens and source properties. We observed correlations between errors on lens potential parameters (e.g., the mass density slope) and on the morphology of the source (e.g., the effective radius or axis ratio). We argue that such correlations can be manifestations of the well-known mass-sheet transformation (MST), but also more generally of the source-position transformation (SPT). Better handling these degeneracies in current and future modeling methods will be key to further minimize model biases.
- We investigated how certain model assumptions affect the recovery of lens potential parameters. In particular, we explore (1) how the cuspy nature of the lensed galaxy can lead to systematic errors if the source model is not flexible enough to capture large intensity variations, and (2) how the accuracy of the PSF (the true PSF being unavailable to the modelers) plays a role even for extended source modeling. We find that both an accurately sampled PSF and a source model with large dynamic range (e.g., using a luminosity-weighted source prior) are warranted to reconstruct cuspy lensed sources while minimizing systematic errors on lens potential parameters.

Over the past years, numerous lens modeling methods have been proposed and implemented in different software packages. Here, we selected a subset of those with the goal of using them together, instead of only opposing them. Typically, we refrain from explicitly ranking the modeling methods, which would only be meaningful over an extremely large sample of strong lenses with different data quality and modeling assumptions to ensure proper generalization. In a real-case scenario, we do not have access to the truth; therefore, combining the results from different methods is a pragmatic and efficient way to detect and mitigate systematic errors. As is shown, for example, in [Suyu et al. \(2006\)](#) by testing three types of source regularizations, there exists an inherent dependence on the data properties, the lens configuration and the (unobservable) intrinsic source morphology, such that it is unlikely that a single source reconstruction method gives unbiased results in all cases. Our work strengthens this idea and goes further by combining a large collection of models, while investigating their specific effects on inferred parameters. Moreover, we purposefully allowed some level of freedom for the modelers (e.g., masks, PSF, posterior marginalization, etc.), such that our work also illustrates the role of specific modelers' choices. These choices play a role in the observed scatter between models, and can be marginalized over by combining the methods together.

As stated in the introduction, we have not used lens modeling methods based on deep learning, as those would require many additional assumptions, in particular regarding their training phase. Nevertheless, the comparison framework presented here is very general and does not depend on the nature (classical, deep learning, etc.) of the underlying methods. Therefore, we encourage future studies to complement and combine classical methods with those based on deep learning, as the latter have clear advantages such as fast computation time after training and a large flexibility through different network architectures. Examples of using neural networks to complement classical techniques have been proposed in [Maresca et al. \(2021\)](#) and [Pearson et al. \(2021\)](#). Our publicly released simulated data and lens modeling products can be directly used to test and improve such deep learning (or any other) approaches.

The framework and ideas presented here is designed to be applied on real data sets and expanded beyond our initial simplifying assumptions. In particular, the role of the lens surface brightness model should be investigated further (see e.g., [Nightingale et al. 2024](#), in the context of subhalo detection). Similarly, the standard assumptions of uncorrelated Gaussian noise used in lens modeling analyses should be re-assessed (e.g., recent JWST imaging data show strongly correlated noise patterns, see [Rigby et al. 2023](#)) to ensure that analyses of the many future observations of strong lenses remain fully accurate.

Acknowledgements. The authors thank the anonymous referee for reviewing the original manuscript and providing useful comments. This work originated in the Lensing Odyssey 2021 workshop¹⁰, and so we would like to acknowledge the organizers and attendees for the fruitful discussions. AG acknowledges the Swiss National Science Foundation (SNSF) for supporting this work. This work was also supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (COSMICLENS: grant agreement No 787886). GV and QM were both supported by the generosity of Eric and Wendy Schmidt by recommendation of the Schmidt Futures program. QM gratefully acknowledges a grant of computer time from ACCESS allocation TG-AST130007. This research has made use of SCIPY ([Virtanen et al. 2020](#)), NUMPY ([Oliphant 2006](#); [Van Der Walt et al. 2011](#)), MATPLOTLIB ([Hunter 2007](#)), ASTROPY ([Astropy Collaboration 2013](#); [Price-Whelan et al. 2018](#)), JUPYTER ([Kluyver et al. 2016](#)) and GETDIST ([Lewis 2019](#)).

¹⁰ <https://lensingodyssey.com>

References

- Adam, A., Perreault-Levasseur, L., Hezaveh, Y., & Welling, M. 2023, *ApJ*, **951**, 6
- Arras, P., Baltac, M., Ensslin, T. A., et al. 2019, *Astrophysics Source Code Library* [record ascl:1903.008]
- Astropy Collaboration (Robitaille, T. P., et al.) 2013, *A&A*, **558**, A33
- Barkana, R. 1998, *ApJ*, **502**, 531
- Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, *J. Mach. Learn. Res.*, **20**, 28:1
- Birrer, S. 2021, *ApJ*, **919**, 38
- Birrer, S., & Amara, A. 2018, *Phys. Dark Univ.*, **22**, 189
- Birrer, S., & Treu, T. 2019, *MNRAS*, **489**, 2097
- Birrer, S., Amara, A., & Refregier, A. 2015, *ApJ*, **813**, 102
- Birrer, S., Amara, A., & Refregier, A. 2016, *JCAP*, **2016**, 020
- Birrer, S., Shajib, A. J., Galan, A., et al. 2020, *A&A*, **643**, A165
- Birrer, S., Shajib, A., Gilman, D., et al. 2021, *J. Open Source Softw.*, **6**, 3283
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. 2016, *ArXiv e-prints* [arXiv:1601.00670]
- Blum, K., Castorina, E., & Simonović, M. 2020, *ApJ*, **892**, L27
- Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, *ApJ*, **638**, 703
- Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, *JAX: Composable Transformations of Python+NumPy Programs*, <http://github.com/google/jax>
- Cao, X., Li, R., Nightingale, J. W., et al. 2022, *Res. Astron. Astrophys.*, **22**, 025014
- Collett, T. E. 2015, *ApJ*, **811**, 20
- Collett, T. E., & Auger, M. W. 2014, *MNRAS*, **443**, 969
- Collett, T. E., Buckley-Geer, E., Lin, H., et al. 2017, *ApJ*, **843**, 148
- Diego, J. M., Meena, A. K., Adams, N. J., et al. 2023, *A&A*, **672**, A3
- Ding, X., Treu, T., Birrer, S., et al. 2021, *MNRAS*, **503**, 1096
- Edenhofer, G., Frank, P., Roth, J., et al. 2024, *J. Open Source Softw.*, **9**, 6593
- Enßlin, T. A. 2019, *Ann. Phys.*, **531**, 1800127
- Ertl, S., Schuldt, S., Suyu, S. H., et al. 2023, *A&A*, **672**, A2
- Ertl, S., Schuldt, S., Suyu, S. H., et al. 2024, *A&A*, **685**, A15
- Etherington, A., Nightingale, J. W., Massey, R., et al. 2022, *MNRAS*, **517**, 3275
- Etherington, A., Nightingale, J. W., Massey, R., et al. 2024, *MNRAS*, **531**, 3684
- Falco, E. E., Gorenstein, M. V., & Shapiro, I. I. 1985, *ApJ*, **289**, L1
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Frank, P., Leike, R., & Enßlin, T. A. 2021, *Entropy*, **23**, 853
- Galan, A., Peel, A., Joseph, R., Courbin, F., & Starck, J. L. 2021, *A&A*, **647**, A176
- Galan, A., Vernardos, G., Peel, A., Courbin, F., & Starck, J. L. 2022, *A&A*, **668**, A155
- Galan, A., Van de Vyvere, L., Gomer, M., Vernardos, G., & Sluse, D. 2023, *J. Open Source Softw.*, **8**, 5567
- Galan, A., Caminha, G. B., Knollmüller, J., Roth, J., & Suyu, S. H. 2024, *A&A*, **689**, A304
- Garvin, E. O., Kruk, S., Cornen, C., et al. 2022, *A&A*, **667**, A141
- Gawade, P., More, A., More, S., et al. 2024, *ArXiv e-prints* [arXiv:2404.18897]
- Gentile, F., Tortora, C., Covone, G., et al. 2023, *MNRAS*, **522**, 5442
- Gomer, M., & Williams, L. L. R. 2020, *JCAP*, **2020**, 045
- Gomer, M. R., & Williams, L. L. R. 2021, *MNRAS*, **504**, 1340
- Gomer, M. R., Sluse, D., Van de Vyvere, L., Birrer, S., & Courbin, F. 2022, *A&A*, **667**, A86
- Gu, A., Huang, X., Sheu, W., et al. 2022, *ApJ*, **935**, 49
- Hessel, M., Budden, D., Viola, F., et al. 2020, *The DeepMind JAX Ecosystem*, <http://github.com/google-deepmind>
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Joseph, R., Courbin, F., Starck, J. L., & Birrer, S. 2019, *A&A*, **623**, A14
- Karчев, K., Coogan, A., & Weniger, C. 2022, *MNRAS*, submitted [arXiv:2105.09465]
- Kelly, P. L., Rodney, S., Treu, T., et al. 2023, *Science*, **380**, abh1322
- Kennedy, J., & Eberhart, R. 2002, *Proc. ICNN'95 - Int. Conf. Neural Networks*, **4**, 1942
- Khadka, N., Birrer, S., Leauthaud, A., & Nix, H. 2024, *MNRAS*, **533**, 795
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, 87
- Knollmüller, J., & Enßlin, T. A. 2019, *ArXiv e-prints* [arXiv:1901.11033]
- Koopmans, L. V. E., Treu, T., Bolton, A. S., Burles, S., & Moustakas, L. A. 2006, *ApJ*, **649**, 599
- Krist, J. E., Hook, R. N., & Stoehr, F. 2011, *Proc. SPIE*, **8127**, 81270J
- Lanusse, F., Starck, J. L., Leonard, A., & Pires, S. 2016, *A&A*, **591**, A2
- Lefor, A. T., & Futamase, T. 2015, *ArXiv e-prints* [arXiv:1505.00502]
- Lewis, A. 2019, *ArXiv e-prints* [arXiv:1910.13970]
- Limousin, M., Kneib, J. P., Bardeau, S., et al. 2007, *A&A*, **461**, 881
- Maresca, J., Dye, S., & Li, N. 2021, *MNRAS*, **503**, 2229
- Meneghetti, M., Natarajan, P., Coe, D., et al. 2017, *MNRAS*, **472**, 3177
- Millon, M., Galan, A., Courbin, F., et al. 2020, *A&A*, **639**, A101
- Newton, E. R., Marshall, P. J., Treu, T., et al. 2011, *ApJ*, **734**, 104
- Nightingale, J. W., & Dye, S. 2015, *MNRAS*, **452**, 2940
- Nightingale, J. W., Dye, S., & Massey, R. J. 2018, *MNRAS*, **478**, 4738
- Nightingale, J. W., Hayes, R. G., Kelly, A., et al. 2021, *J. Open Source Softw.*, **6**, 2825
- Nightingale, J. W., He, Q., Cao, X., et al. 2024, *MNRAS*, **527**, 10480
- Oguri, M. 2010, *PASJ*, **62**, 1017
- Oguri, M., & Marshall, P. J. 2010, *MNRAS*, **405**, 2579
- Oliphant, T. E. 2006, *A Guide to NumPy* (Trelgol Publishing USA)
- Pascale, M., Frye, B. L., Pierel, J. D. R., et al. 2024, *ArXiv e-prints* [arXiv:2403.18902]
- Pearson, J., Maresca, J., Li, N., & Dye, S. 2021, *MNRAS*, **505**, 4362
- Peel, A., Lanusse, F., & Starck, J.-L. 2017, *ApJ*, **847**, 23
- Phan, D., Pradhan, N., & Jankowiak, M. 2019, *ArXiv e-prints* [arXiv:1912.11554]
- Powell, D. M., Vegetti, S., McKean, J. P., et al. 2022, *MNRAS*, **516**, 1808
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, **156**, 123
- Rigby, J. R., Vieira, J. D., Phadke, K. A., et al. 2023, *ArXiv e-prints* [arXiv:2312.10465]
- Ruan, D., & Keeton, C. R. 2023, *ArXiv e-prints* [arXiv:2309.16529]
- Rüstig, J., Guardiani, M., Roth, J., Frank, P., & Enßlin, T. 2024, *A&A*, **682**, A146
- Saha, P., Sluse, D., Wagner, J., & Williams, L. L. R. 2024, *Space Sci. Rev.*, **220**, 12
- Schneider, P., & Sluse, D. 2013, *A&A*, **559**, A37
- Schneider, P., & Sluse, D. 2014, *A&A*, **564**, A103
- Schuldt, S., Suyu, S. H., Cañameras, R., et al. 2023a, *A&A*, **673**, A33
- Schuldt, S., Cañameras, R., Shu, Y., et al. 2023b, *A&A*, **671**, A147
- Selig, M., Bell, M. R., Junklewitz, H., et al. 2013, *A&A*, **554**, A26
- Sengül, A. Ç., Dvorkin, C., Ostdiek, B., & Tsang, A. 2022, *MNRAS*, **515**, 4391
- Shajib, A. J. 2024, *IAU Symp.*, **381**, 3
- Shajib, A. J., Treu, T., Birrer, S., & Sonnenfeld, A. 2021, *MNRAS*, **503**, 2380
- Shajib, A. J., Wong, K. C., Birrer, S., et al. 2022, *A&A*, **667**, A123
- Sonnenfeld, A. 2018, *MNRAS*, **474**, 4648
- Stacey, H. R., Powell, D. M., Vegetti, S., et al. 2024, *A&A*, **688**, A110
- Starck, J.-L., Fadili, J., & Murtagh, F. 2007, *IEEE Trans. Image Process.*, **16**, 297
- Steininger, T., Dixit, J., Frank, P., et al. 2019, *Ann. Phys.*, **531**, 1800290
- Suyu, S. H., & Halkola, A. 2010, *A&A*, **524**, A94
- Suyu, S. H., Marshall, P. J., Hobson, M. P., & Blandford, R. D. 2006, *MNRAS*, **371**, 983
- Suyu, S. H., Hensel, S. W., McKean, J. P., et al. 2012, *ApJ*, **750**, 10
- Suyu, S. H., Auger, M. W., Hilbert, S., et al. 2013, *ApJ*, **766**, 70
- Tagore, A. S., & Keeton, C. R. 2014, *MNRAS*, **445**, 694
- Tan, C. Y., Shajib, A. J., Birrer, S., et al. 2024, *MNRAS*, **530**, 1474
- Tessore, N., & Metcalf, R. B. 2015, *A&A*, **580**, A79
- Unruh, S., Schneider, P., & Sluse, D. 2017, *A&A*, **601**, A77
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, **13**, 22
- Van de Vyvere, L., Sluse, D., Gomer, M. R., & Mukherjee, S. 2022, *A&A*, **663**, A179
- Vegetti, S., & Koopmans, L. V. E. 2009, *MNRAS*, **392**, 945
- Vegetti, S., Czoske, O., & Koopmans, L. V. E. 2010, *MNRAS*, **407**, 225
- Vegetti, S., Birrer, S., Despali, G., et al. 2023, *ArXiv e-prints* [arXiv:2306.11781]
- Vernardos, G. 2022, *MNRAS*, **511**, 4417
- Vernardos, G., & Koopmans, L. V. E. 2022, *MNRAS*, submitted [arXiv:2202.09378]
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Wagner, J. 2018, *A&A*, **620**, A86
- Wagner, J. 2019, *Universe*, **5**, 177
- Wagner, J., Liesenborgs, J., & Tessore, N. 2018, *A&A*, **612**, A17
- Wang, H., Cañameras, R., Caminha, G. B., et al. 2022, *A&A*, **668**, A162
- Warren, S. J., & Dye, S. 2003, *ApJ*, **590**, 673
- Wong, K. C., Suyu, S. H., Chen, G. C. F., et al. 2020, *MNRAS*, **498**, 1420
- Wong, K. C., Dux, F., Shajib, A. J., et al. 2024, *A&A*, **689**, A168
- Yang, L., Birrer, S., & Treu, T. 2020, *MNRAS*, **496**, 2648
- Yıldırım, A., Suyu, S. H., Chen, G. C. F., & Komatsu, E. 2023, *A&A*, **675**, A21
- Zhuang, J., Tang, T., Ding, Y., et al. 2020, *ArXiv e-prints* [arXiv:2010.07468]




Fig. A.1. Old simulated lensing data based on another source galaxy. For that mock the source image background was insufficiently subtracted, which left a faint imprints in the data after being lensed. While unnoticeable to the eye in the data (left panel), the square boundaries of the source are clearly noticeable in the input source (middle panel) and even on the reconstructed source (right panel, using a correlated field model as an example). These artificial boundaries led to biases in lens potential parameters and source size. Note that a logarithmic scale is used in all panels.

Appendix A: Realistic source surface brightness in mock data

The mock data shown in Fig. 1 is not the first mock we built for this work. Initially, an other prescription for the background source was used, which caused large biases on many model parameters, which was noticed only after unblinding. Therefore, it was decided to repeat the entire procedure—namely, the data simulation and blind submissions of models from independent modelers—which lead to the mock used throughout this work. As it can be beneficial to some readers, we give more details about the old mock data below and why it caused large parameter biases.

We show in the left panel of Fig. A.1 the old simulated data. For the source, shown in the middle panel, we used a B filter image of M31 available through the ESO Online Digitized Sky Survey. This image was selected because a local galaxy has high-detail resolved structure with negligible PSF spikes, which could introduce nonphysical features if lensed. After ray-tracing and the addition of noise, the data visually resembles to genuine lensing data, despite the input source boundaries being visible at some locations. Therefore, any potential issue related to these boundaries could not be detected before modeling the data.

After the blind modeling stage, it was noticed by the modelers that their reconstructed source models were displaying strong boxy features surrounding the ellipsoidal shape of the source galaxy. Among the six models, the Sparsity+Wavelets model was strongly affected by the square and boxy nature of the source, for which an extremely low resolution grid was systematically preferred as it represented better these sharp boundaries. This artificial bias towards a low resolution grid (i.e., few source pixels) was in turn strongly biasing the lens potential parameters, in particular the mass density slope due to its degeneracy with the source size (through this MST). The right panel of Fig. A.1 shows, as an example, the Correlated Field model of that source, which also displays a boxy shape. After comparing the different models together, the mass density slope was found to be biased in almost all models. After extensive checks, the origin of these biases was attributed to the sharp boundaries and large background flux of the input source. The source used in the new simulated data (shown in Fig. 1) was carefully processed, which solved these issues.

Appendix B: Technical modeling details

This section gives technical details on the specific choices made by modelers to analyze the simulated data. We complement it with Table B.1 which gives the approximate computation time necessary to obtain the various models used in this work.

B.1. Sérsic+Shapelets

For the analytical model using LENSSTRONOMY (Sect. 3.2), a $4.5''$ -radius circular mask around the lens is used. The source light is composed of a single elliptical Sérsic profile. Shapelets components are added sequentially to that profile. The significance of those components to the model is evaluated by calculating the Bayesian Information Criterion (BIC), which balances the likelihood with the number of parameters following:

$$\text{BIC} = n_{\text{par}} \ln(n_{\text{data}}) - 2 \ln L(\tilde{\eta}), \quad (\text{B.1})$$

where n_{par} is the number of parameter, n_{data} is the number of data points (i.e., data pixels used as constraints), and $L(\tilde{\eta})$ is the loss function evaluated at the best fit position (Eq. 6). For this specific mock, the BIC favors a maximum order of shapelets $n_{\text{max}} = 5$ (Table B.2).

The point spread function (PSF) is treated in LENSSTRONOMY as follows. The surface brightness of the lensing system, which is simulated at each iteration, can be sampled on a grid with higher resolution than the observed image before being averaged to the data resolution. If the modeled surface brightness is supersampled, the user can perform the PSF convolution on the finer grid. In that case, a supersampled PSF is calculated by interpolation, and an iterative process allowing for perturbation of individual PSF pixels, ensures that the downsampling of the supersampled PSF recovers the input PSF. We elected a supersampling factor of 5 and performed the PSF convolution on the supersampled grid.

B.2. Adaptive+Matérn

We begin with the broadest possible range of the lens potential parameters of the model (see Eqs. 8 and 9), except for the slope, which we fix to isothermal, and the Einstein radius, whose range is fixed by rough estimation of the radius of the Einstein ring from the data. The former serves simply to speed up the calculation, while the latter is required to avoid over- or under-focused, non-physical solutions. The main choices at this first-pass stage are the type of regularization and the resolution of the adaptive grid. For the regularization we choose curvature, which has only one free parameter, the regularization strength λ , and a fixed covariance matrix C_s (see Eq. 10). The adaptive grid is created simply by using the deflected positions of every third pixel of the data image in the x and y directions. We create a custom mask that we use in all of our subsequent models, and we do not perform any supersampling of the PSF. This setup is very fast to run, taking only a few minutes on a standard multi-core laptop.

After this first pass, we switch to using the Matérn kernel regularization given in Eq. 11, which has 3 free parameters, and increase the resolution of our adaptive grid by using every second pixel in both directions in the data image to construct it. Based on our first crude model, we restrict each lens potential parameter to about 20 per cent of its full range, centered roughly on the bulk of the posterior probability. This model takes about an hour to run on a standard multi-core laptop.

Finally, we further restrict the parameters to 10 per cent of their full range and run a final model with the same regulariza-

Table B.1. Computation times for the six modeling methods, including best-fit optimization and posterior distribution sampling.

Modeling method label	Software package	Involved steps	Longest instance	Total	Hardware configuration
Sérsic+Shapelets	LENSTRONOMY	PSO + MCMC	17 hours	38 hours	Intel Core i7 CPU (2 cores)
Adaptive+Matérn	VKL	Nested sampling	4 days	4.2 days	Intel Xeon Platinum 8260 CPU (96 cores)
Cluster+Exp	QLENS	Nested sampling	1.5 hours	1.5 hours	Intel Xeon Platinum 8160 CPU (192 cores)
Cluster+Exp+Lumweight	QLENS	Nested sampling	2 hours	2 hours	Intel Xeon Platinum 8160 CPU (192 cores)
Sparsity+Wavelets	HERCULENS	Grad. descent + SVI	6 minutes	3.4 hours	NVIDIA A10 GPU
Correlated Field	HERCULENS	Grad. descent + geoVI	45 minutes	4.5 hours	NVIDIA A10 GPU

Notes. Some of the final models were obtained by combining several model instances, so we quote the runtime for a single instance (the most expensive) in addition to the total runtime (assuming single instances were not performed in parallel). For additional details and software package references, see Sect. 3 and Appendix B.

Table B.2. Comparison of best-fit log-likelihood ($\ln L(\tilde{\eta})$) and corresponding BIC values (Eqs. 6 and B.1) considering different maximum order of the shapelets basis set n_{\max} for the Sérsic+Shapelets modeling method (and corresponding number of parameters n_{par}). The number of data pixels is 6376 used to constrain the model (within a circular mask of 4'5 radius).

n_{\max}	n_{par}	$\ln L(\tilde{\eta})$	BIC
3	28	-3581.4	7408.0
4	33	-3562.8	7414.7
5	39	-3485.1	7311.9
6	46	-3502.2	7407.1
7	54	-3434.2	7341.4

tion but even higher resolution; that is, shooting back every pixel in the data image to create the adaptive grid. We combine the last two models, which differ only in the resolution of the reconstructed source (although the first served to initialize the second in order to save on computations), by merging their posterior samples in such a way that the probability mass of each one is the same independently of its actual size.

B.3. Cluster+Exp and Cluster+Exp+Lumweight

In QLENS, the image plane is supersampled by splitting each image pixel into 3×3 subpixels; each subpixel is ray-traced to the source plane, and the source grid is generated using a k -means clustering algorithm exactly as is done in Nightingale et al. (2018), but with each ray-traced point receiving equal weight. We choose the number of source pixels to be equal to half the number of image pixels within the mask. The surface brightness of each ray-traced subpixel is determined by interpolating in the three source pixels whose Delaunay triangle the ray-traced point is in (or closest to); the surface brightness for all the subpixels within a given image pixel are then averaged to obtain the surface brightness for the image pixel. The pixel surface brightness values obtained this way are then convolved with the pixel-level PSF (note that although we are supersampling the image plane, we do not supersample the PSF here; the effect of PSF supersampling will be explored in Section 7.2). The regularization is performed using an exponential kernel (equivalent to a Matérn kernel with $\nu = 1/2$).

To encourage convergence, we make use of two additional priors on the source: first, there is a prior that discourages producing lensed images outside the mask. We accomplish this by temporarily unmasking after the source pixels are solved for, generating the lensed images without a mask, and imposing a

steep penalty if surface brightness is found outside the mask whose value is greater than 0.2 times the maximum surface brightness of the images. Second, we place a prior on the number of lensed images produced. This is accomplished by creating a Cartesian grid in the source plane and finding the overlap area of all the ray-traced image pixels for each Cartesian grid cell; by dividing the total overlap area by the area of each grid cell, we obtain the number of images produced by that cell. We can then take the average number of images over all the cells. We impose a steep penalty if the average number of images is less than 1.5. This discourages solutions that are not multiply imaged, where the source looks identical to the observed configuration of lensed images. With these priors in place, we can obtain a good solution with a single nested sampling run, provided the parameter priors are broad enough.

The Cluster+Exp+Lumweight model uses all of the methods described above, but in addition it uses a luminosity-weighted regularization, as is described in Sect. 3.3. Thus we include the additional parameter ρ (Eq. 13), which controls the steepness of the luminosity weighting, as an additional nonlinear parameter to be varied.

B.4. Sparsity+Wavelets

Before modeling the source on a regular grid with multi-scale regularization, we start with an approximate lens mass model obtained by modeling the source with a single Sérsic profile. Since at this stage of the modeling process the mass model may be rather inaccurate, using spatially varying regularization weights (\mathbf{W}_{ms} in Eq. 14) could bias the source reconstruction. Therefore we approximate the weights by their median value within each wavelet decomposition scale (i.e., we use spatially uniform weights within each frequency range). We set the global regularization strength $\lambda_{\text{ms}} = 3\sigma$ for the first wavelet scale (highest frequency features), and $\lambda_{\text{ms}} = 1\sigma$ for the remaining scales. We choose a lower threshold for low frequency features as advocated in various works relying on similar multi-scale regularization strategies (e.g., Lanusse et al. 2016; Peel et al. 2017; Galan et al. 2022), since high frequencies are more impacted by the presence of noise in the data.

We obtain a first approximate model of the pixelated source by jointly optimizing all parameters except for the fixed lens center, and we impose a strong isothermal prior (i.e., $\gamma \sim \mathcal{N}(2, 10^{-3})$) on the mass density slope to avoid introducing degeneracies early in the modeling sequence. We use the gradient descent optimizer ADABELIEF (Zhuang et al. 2020) implemented in the OPTAX (Hessel et al. 2020) library to obtain best-fit parameters. We then re-optimize model parameters by

re-computing regularization weights and releasing priors on the lens center and density slope.

The last step is to estimate the posterior distribution of lens mass parameters while further refining the source model. At this stage, the lens model is very close to the best-fit model so we do not rely anymore on the approximation of uniform regularization weights per wavelet scale, and properly propagate the noise to source plane. These more accurate regularization weights significantly help eliminating remaining artifacts located on the outskirts of the source galaxy, when the data is the least constraining. The resulting χ^2 being below unity, we further boost the regularization strength of high-frequencies (for this particular data, by 5σ) in order to obtain a χ^2 of the order of unity. We note that the precise value of this boost does not significantly impact the final posterior distribution as we marginalize over many model variations. In particular, we vary the number of source pixels from 80×80 to 160×160 pixels with steps of 5. We also run the same ensemble of models by globally increasing the regularization strength by 1σ . In total, we consider 34 model variations.

The joint posterior distribution for lens mass parameters is estimated using stochastic variational inference (SVI, see review by Blei et al. 2016), which directly makes use of known gradient of the loss function. SVI is also less computationally expensive than other sampling methods such as Markov Chain Monte Carlo or Hamiltonian Monte Carlo, which allows us to run a larger number of model variations. Since in this work, we are mainly interested in the joint posterior distribution of the lens mass parameters which are not expected to exhibit strongly non-Gaussian correlations¹¹, we find that a multi-variate Gaussian distribution is a sufficient surrogate posterior model. However, we acknowledge that SVI can underestimate uncertainties (as in Gu et al. 2022), a limitation that we address by marginalizing over the 34 model variations assuming equal weights. In addition to the full posterior distributions and first-order posterior statistics, we save in COOLEST format a point-estimate model, that corresponds to the mean model as obtained from SVI, with 120^2 pixels and fiducial regularization strength.

B.5. Correlated Field

Similar to the strategy used with the multi-scale source model, we first model the imaging data with a single Sérsic source profile, which we then replace with the correlated field model defined in Eq. 15. For our baseline model, we set the shape of the Gaussian excitation field to $90^2 = 8'100$ pixels. The field power spectrum is modeled as power law parametrized with an amplitude and a slope. These two parameters are themselves sampled from a log-normal distribution, described in turn by a mean and scale parameters. The additive offset in real space is modeled by a single scalar (initialized to the mean flux values from the Sérsic model), whereas variations in this offset are sampled from a log-normal distribution with additional mean and scale parameters. Lens potential model parameters (PEMD and external shear) are sampled from Gaussian distribution, for which we check that the prior widths are large enough. We optimize the full model—lens mass parameters and source field parameters—using the set of minimizers and samplers implemented in NIFTY. More specifically, we converge to the maximum a posteriori solution and estimate the joint posterior distribution using metric-Gaussian variational inference (MGVI, Knollmüller & Enßlin 2019).

We ran several variations of the above fiducial models. In particular, we increased the field resolution to $120^2 = 14'400$ pixels, alter the sampler random seeds, initialized the model with a worst lens model (obtained from a different Sérsic source model), and used geometric variational inference (geoVI, Frank et al. 2021) instead of MGVI. All these model variations result in almost identical posterior distributions for lens mass parameters and consistent source models. Nevertheless, we conservatively marginalize over these models with equal weights. For the point-estimate parameters, we set those to the mean values of the VI samples of the model with the most resolved source (i.e., 120^2 pixels), although the fiducial model is virtually indistinguishable.

Appendix C: Correlations between lens and source properties

Given the six independent models we gather in this work, we can investigate of the errors on the lens and source properties of interest correlate with each other. In other, we are interested in signs of degeneracies between the lens and source properties, that can be revealed over the ensemble of models. We show in Figs. C.1, C.2 and C.3 a series of plots that correlate all lens potential parameters with the three main source properties we investigate (r_{eff} , q_s and m_s respectively). These results are presented and discussed in the main text, in particular in Sect. 6.4 and Sect. 8.2.

¹¹ This assumption is validated with the correlated field model by using more flexible SVI methods (see Sect. B.5).

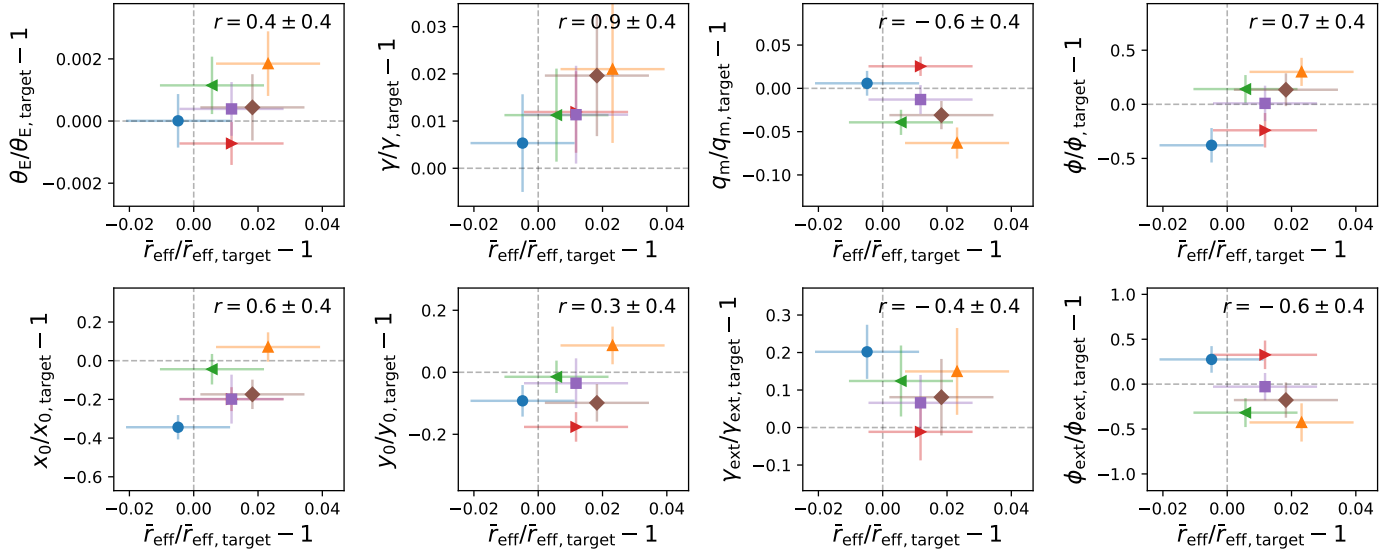


Fig. C.1. Relative error with respect to the true values. x axis: source effective radius, y axis: mass model parameters. The legend markers and colors are the same as in Fig. 4. In the top right part of each panel, the biweight mid-correlation r is indicated (0 means no correlation).

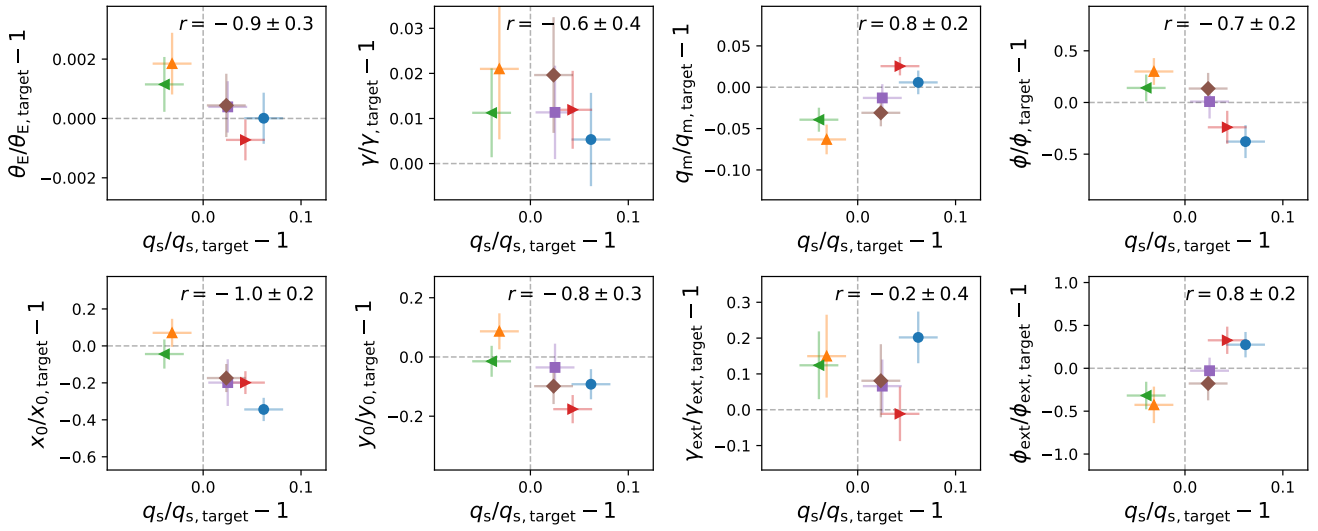


Fig. C.2. Relative error with respect to the true values. x axis: source axis ratio (estimated from central moments), y axis: mass model parameters. The legend markers and colors are the same as in Fig. 4. In the top right part of each panel, the biweight mid-correlation r is indicated (0 means no correlation).

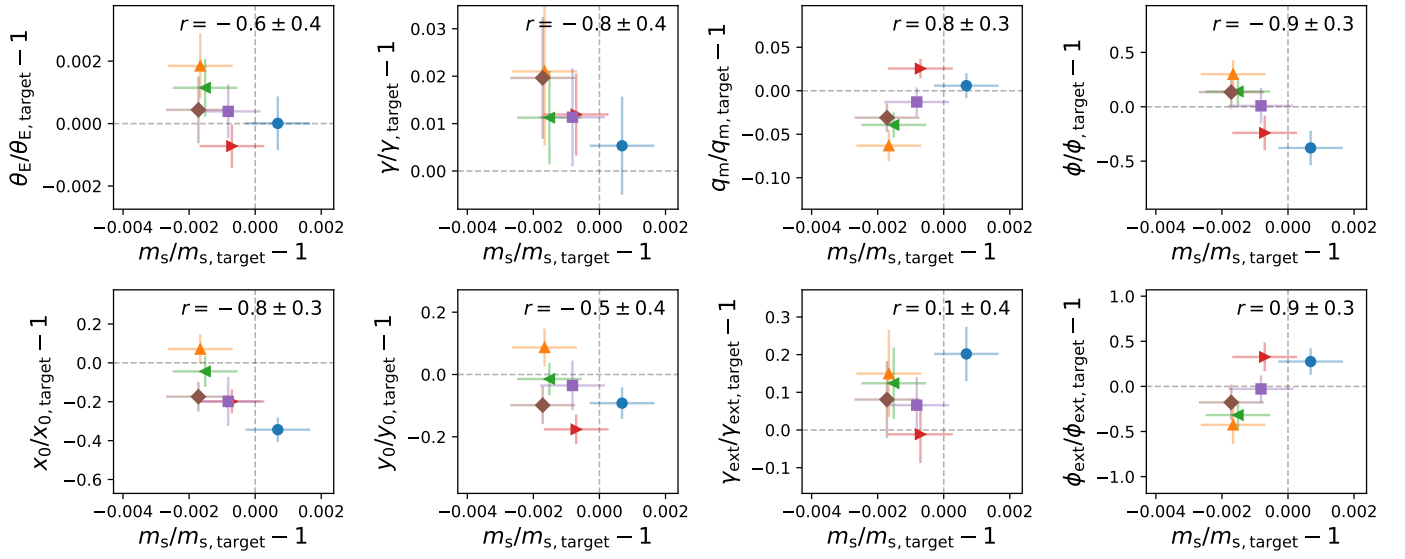


Fig. C.3. Relative error with respect to the true values. x axis: source total magnitude, y axis: mass model parameters. The legend markers and colors are the same as in Fig. 4. In the top right part of each panel, the biweight mid-correlation r is indicated (0 means no correlation).