

Using wavelets to capture deviations from smoothness in galaxy-scale strong lenses^{★,★★}

A. Galan¹, G. Vernardos¹, A. Peel¹, F. Courbin¹, and J.-L. Starck²

¹ Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, 51 chemin Pegasi, 1290 Versoix, Switzerland
e-mail: aymeric.galan@epfl.ch

² AIM, CEA, CNRS, Université Paris-Saclay, Université de Paris, Orme des Merisier, 91191 Gif-sur-Yvette, France

Received 11 July 2022 / Accepted 5 September 2022

ABSTRACT

Modeling the mass distribution of galaxy-scale strong gravitational lenses is a task of increasing difficulty. The high-resolution and depth of imaging data now available render simple analytical forms ineffective at capturing lens structures spanning a large range in spatial scale, mass scale, and morphology. In this work, we address the problem with a novel multiscale method based on wavelets. We tested our method on simulated *Hubble* Space Telescope (HST) imaging data of strong lenses containing the following different types of mass substructures making them deviate from smooth models: (1) a localized small dark matter subhalo, (2) a Gaussian random field (GRF) that mimics a nonlocalized population of subhalos along the line of sight, and (3) galaxy-scale multipoles that break elliptical symmetry. We show that wavelets are able to recover all of these structures accurately. This is made technically possible by using gradient-informed optimization based on automatic differentiation over thousands of parameters, which also allow us to sample the posterior distributions of all model parameters simultaneously. By construction, our method merges the two main modeling paradigms – analytical and pixelated – with machine-learning optimization techniques into a single modular framework. It is also well-suited for the fast modeling of large samples of lenses.

Key words. galaxies: structure – dark matter – methods: data analysis – gravitation – gravitational lensing: strong

1. Introduction

The Lambda cold dark matter cosmological model encapsulates our current understanding of the Universe, accurately explaining a number of observations on large scales (>10 Mpc), such as the cosmic microwave background temperature and polarization anisotropies (Planck Collaboration VI 2020), baryon acoustic oscillations (e.g., Raichoor et al. 2021), and the accelerated expansion of the Universe (e.g., Riess 2020). The dark matter (DM) component of this model plays a major role in the hierarchical collapse of matter due to gravitational instability, which eventually produced the galaxies populated by stars that we observe today (e.g., Toomre & Toomre 1972; Dubinski 1994; Springel et al. 2006). Despite its successes on large scales, the effect of DM on galactic and subgalactic scales (<10 Mpc), where its interplay with baryons is mediated by nonlinear and poorly understood mechanisms (e.g., Blumenthal et al. 1986; Zubovas & King 2012), poses a challenge in explaining observations (e.g., Vogelsberger et al. 2020). For instance, predictions from the cold DM model lead to overly cuspy central density profiles (the so-called cusp-core problem, Moore 1994; de Blok 2010), too many low-mass ($\lesssim 10^8 M_\odot$) dwarf galaxies (the missing satellites problem, Moore et al. 1999; Klypin et al. 1999) and intermediate-mass ($\sim 10^{10} M_\odot$) galaxies populated by too few stars (the too-big-to-fail problem, Boylan-Kolchin et al. 2011; Papastergis et al. 2015).

[★] Movies associated to Fig. 4 are available at <https://www.aanda.org>

^{★★} All methods presented here are publicly available in our new HERCULENS package, <https://github.com/austinpeel/herculens>

One path to reconcile theory with observations is to examine alternative DM models that assume different properties of the DM particle. For example, warm DM simulations (e.g., Li et al. 2017) have been able to produce far fewer dwarf galaxy satellites, which in line with observations. Another path is to improve our ability to detect observational signatures of DM through its gravitational influence on baryonic matter, such as those observed in the local Universe within the Milky Way (e.g., with stellar streams, Erkal et al. 2016) and its neighborhood (e.g., within dwarf galaxy satellites, Nadler et al. 2021). At cosmological distances, gravitational lensing is a direct and more efficient probe of the mass in galaxies, and has the potential to measure DM effects down to kpc in size and $10^7 M_\odot$ in mass (e.g., Hezaveh et al. 2016a; Chatterjee & Koopmans 2018; Gilman et al. 2019). This is a crucial range of mass where most DM models tend to disagree in their predictions.

Galaxy-galaxy strong lensing occurs when two galaxies at different redshifts align along the line of sight, with the foreground galaxy (the lens) deflecting incoming light rays of the background one (the source). From the observer's perspective, the source appears magnified, distorted, and split into multiple images. By modeling the observed lensed features through the well-established physics of the lens equation, one can constrain the total mass distribution in galaxies and measure their DM content and formation history. However, lens modeling is an under-constrained problem, mostly because both the lens mass and the source light distributions are a priori unknown, and the lensed source is often blended with the lens light. Assumptions are therefore needed to reconstruct the lensing mass, the lens light, and the source light distributions by inverting the lens equation.

These assumptions are often priors on the shape of mass and light profiles, or on their higher order statistical properties. Currently used analytical profiles are sufficient to capture first-order properties of the lens mass distribution (e.g., power-law profiles), but they lack the degrees of freedom to capture those small-scale features that are critical for determining DM properties (e.g., [He et al. 2022](#)).

Increasing the complexity of a lens model is not trivial due to degeneracies between the lens potential and the source light: a complex structure in the observed lensed features could be attributed to either local perturbations in the potential or to an intrinsically complex source light distribution (e.g., clumpy star-forming galaxies or galaxy mergers). [Koopmans \(2005\)](#) used a Taylor expansion to address this, deriving a perturbative correction for the lens equation that combines spatial derivatives of both the source and the potential. This approach extends the semilinear inversion technique of [Warren & Dye \(2003\)](#) by adding a perturbing field to the smooth lens potential, which does not assume any specific shape and can be solved for in the same way as the source. The lens potential and the source light were discretized and cast on two grids of pixels in the lens and source planes respectively, on which the two fields were reconstructed (see also [Vegetti & Koopmans 2009](#); [Suyu et al. 2009](#); [Vernardos & Koopmans 2022](#), for extensions of the technique). For the source, various priors have been explored, including analytical functions ranging from the Sérsic profile ([Sérsic 1963](#)) to shapelet basis sets ([Tagore & Keeton 2014](#); [Birrer et al. 2015](#)), Gaussian processes ([Karchev et al. 2022](#)), or deep generative models ([Chianese et al. 2020](#)). More complex sources can be modeled using grids of pixels, combined with curvature-based ([Suyu et al. 2006](#); [Nightingale & Dye 2015](#)), Gaussian process ([Vernardos & Koopmans 2022](#)), or wavelet-based ([Joseph et al. 2019](#); [Galan et al. 2021](#)) regularizations. Recently, [Vernardos & Koopmans \(2022\)](#) studied the effect of specific prior assumptions on both the potential perturbations and the source light distribution, confirming their degeneracy and that the particular choice of priors plays an important role in recovering the underlying potential perturbations.

In a real-world application, we do not know a priori the dominating type of perturbations, such that we need a method that is both flexible yet robust in recovering their main properties. These perturbations may be due to isolated subhalos, or subhalo populations along the line of sight. Both cases allow us to probe the low-mass end of the dark matter mass function, and possibly the shape of DM halos, for which different DM model predictions disagree. The gravitational imaging technique has been used for the detection of subhalos based on the reconstructed potential perturbations. If the pixelated reconstruction contains a well-localized mass over-density, it is replaced by an analytical profile such as a Navarro-Frenk-White (NFW) halo ([Navarro et al. 1996](#)) whose parameters are further optimized ([Vegetti & Koopmans 2009](#)). Using this method, [Vegetti et al. \(2012\)](#) reported the detection of a $10^{8.28} M_{\odot}$ dark halo, and [Vegetti et al. \(2014\)](#) used nondetections in eleven systems of the Sloan Lens ACS Survey (SLACS) sample ([Bolton et al. 2006](#)) observed with the *Hubble* Space Telescope (HST) to constrain the mean projected substructure mass fraction in the context of cold DM. Using a different method based on comparing model likelihoods between a model that explicitly includes a subhalo at a given position with a model that does not, [Hezaveh et al. \(2016b\)](#) reported the detection of a substructure of mass $10^{8.96} M_{\odot}$ from ALMA observations. In general, constraints from individual subhalo detections can mainly be improved with the analysis of larger samples of lenses

([Vegetti et al. 2014](#)), and depend on the shape of the profile used to describe the subhalo mass distribution ([Despali et al. 2022](#)).

These direct detection methods quickly become too computationally expensive for detecting $\gtrsim 2$ subhalos. Considering populations of subhalos instead allows one to probe even lower subhalo masses ($\lesssim 10^7 M_{\odot}$) through their collective effects, since low-mass subhalos are predicted to be more numerous ([Hezaveh et al. 2016a](#)). A population of subhalos is described in a statistical way, but its properties are directly related to those of DM models, which can thus be constrained (e.g., the thermal relic mass of the warm DM particle, [Birrer et al. 2017](#)). One such statistical description of the perturbed lensing potential has been introduced in [Chatterjee & Koopmans \(2018\)](#), using a power-spectrum analysis of the (lensed) source surface brightness fluctuations, reaching a sensitivity down to a few kpc in the subhalo mass power-spectrum (see also [Bayer et al. 2018](#)). Similarly, [Cyr-Racine et al. \(2019\)](#) decomposed model residuals into Fourier modes and linked them to the substructure power-spectrum. More recently, [Vernardos & Koopmans \(2022\)](#) used the gravitational imaging technique to reconstruct the perturbing field of a population of subhalos and recovered its power-spectrum properties, especially the slope, remarkably well. Several studies have also employed deep learning techniques to infer the presence of subhalo populations (e.g., [Brehmer et al. 2019](#); [Diaz Rivero & Dvorkin 2020](#); [Varma et al. 2020](#); [Coogan et al. 2020](#); [Vernardos et al. 2020](#); [Ostdiek et al. 2022](#); [Adam et al. 2022](#)). While it is still unclear if these methods are strongly limited by the simplifying assumptions on their training data (e.g., the absence of lens light, fixed instrumental properties, etc.), they are a promising path forward to speed up computations for application to large samples of lenses, possibly in combination with more classical approaches.

Independently of the presence of subhalos, deviations from smoothness can occur within the lens galaxy mass distribution itself, manifesting as additional radial or azimuthal structures. For instance, deviations along the radial direction can be mass-to-light radial gradients ([Oldham & Auger 2018](#); [Sonnenfeld et al. 2018](#); [Shajib et al. 2021](#)), while angular structures can be the consequence of ellipticity gradients and twists ([Keeton et al. 2000](#); [Van de Vyvere et al. 2022a](#)). These cannot be captured by the typically employed elliptical power-law mass models. Free-form techniques have been one way to address these limitations by dismissing the smooth component of the potential and relying solely on a grid of mass pixels governed by a set of priors (either physically or mathematically motivated) to prevent the appearance of un-physical mass distributions ([Saha & Williams 1997](#); [Coles et al. 2014](#)). However, retaining the smooth component already provides a reliable first-order model, from which to explore higher order deviations. Azimuthal structures can be described by higher-order multipoles that have been identified in stellar populations from both real observations and cosmological simulations ([Trotter et al. 2000](#); [Claeskens et al. 2006](#)). These can be explained by AGN feedback that suppresses the formation of disks in massive galaxies, as they tend to evolve from disk to elliptical or even boxy shapes (i.e., quadrupoles, [Frigo et al. 2019](#)). Recently, [Van de Vyvere et al. \(2022b\)](#) demonstrated that quadrupole moments of low amplitude, based on results of [Hao et al. \(2006\)](#), can be detected in current HST observations of strong lenses, although their detectability depends on numerous factors such as their alignment with the smooth potential or the degrees of freedom in the source model. Using very long baseline interferometric observations of a strong lens, [Powell et al. \(2022\)](#) recently reported the detection of multipole

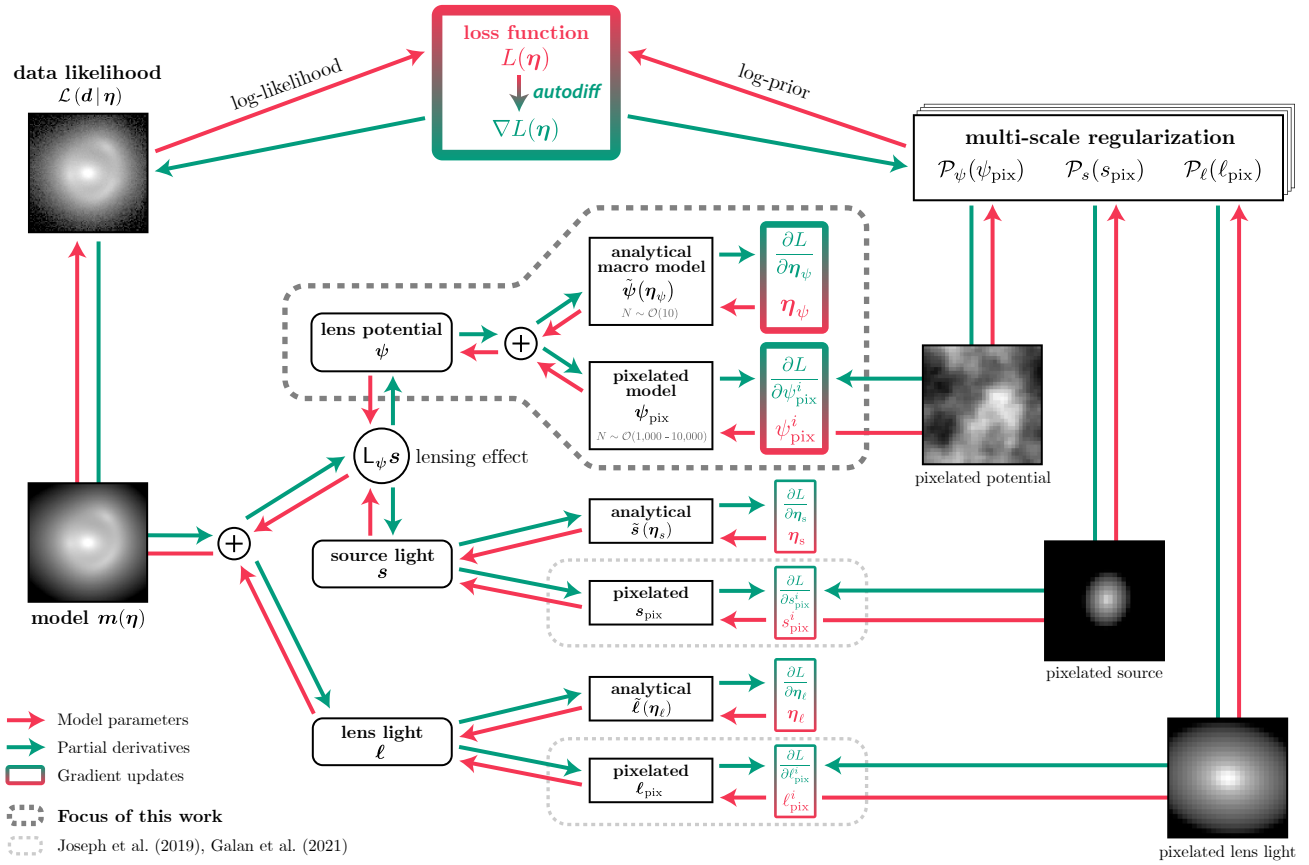


Fig. 1. Flow chart of the fully differentiable method used in this work and implemented in HERCULENS, which merges analytical profiles with pixelated components. We indicate the typical number of parameters N for the analytical potential η_ψ and the pixelated potential ψ_{pix} . Partial derivatives are computed and propagated for all parameters. These derivatives are then used to update parameter values in the direction indicated by the gradient of the loss function L . We note that the blurring operator \mathbf{B} (see Eq. (4)) is also used to generate the model \mathbf{m} but is not shown in the diagram to avoid clutter. This work focuses on modeling the lensing potential ψ , and we refer to Joseph et al. (2019) and Galan et al. (2021) for the modeling of surface brightness distributions.

structures beyond ellipticity in the deflector mass distribution. While not accounting for those multipoles in the model can bias the measurement of the lens mass profile up to a few percent (Van de Vyvere et al. 2022b; Powell et al. 2022), their detection most importantly holds valuable information on the formation history of galaxies, as signatures of their past evolution.

Our goal in this paper is to unify the modeling of generic lens potential perturbations in a robust framework that includes, but is not limited to, the three categories presented above: individual or populations of subhalos and higher order moments in the lens mass distribution. To achieve this, we extend our previous work in Galan et al. (2021) by including a reconstruction of lens potential perturbations regularized with wavelets, and demonstrate that our method can successfully reconstruct perturbed lens potentials of different origin. Our technique benefits from the multiscale properties of the wavelet transform, along with well-motivated sparsity constraints to reconstruct the various spatial scales over which perturbations to the smooth potential can occur. We implement our method using a fully differentiable algorithm based on automatic differentiation, first introduced by Wengert (1964). This programming framework gives direct access to all the derivatives of the highly nonlinear loss function at a negligible computational cost, which enables the use of robust gradient-informed algorithms to explore the parameter space and optimize the model parameters. As a result, convergence to the maximum-a-posteriori (MAP) solution is fast

and first-order error estimates are obtained through Fisher matrix analysis. An efficient exploration of the parameter space for estimating posterior distributions is then performed via gradient-informed Hamiltonian Monte-Carlo (HMC) sampling.

We present our methodology in Sect. 2. The simulated examples used to demonstrate the capabilities of our method are presented in Sect. 3. We perform the reconstruction of perturbations to the lens potential by uniformly modeling these examples and present our results in Sect. 4. We then evaluate the reconstructions of the perturbations in Sect. 5. We conclude this work and discuss its future prospects in Sect. 7.

2. Methodology

In this section we introduce the gravitational lensing formalism and describe in detail the various aspects of our method. Figure 1 summarizes the main components of the method and can be referred to throughout this section. Some important mathematical notations are also summarized Table 1.

2.1. Discrete formulation of lensing

From the observation of a strongly lensed source, or data \mathbf{d} , we aim to recover the lens potential ψ and the (unlensed) source light \mathbf{s} . Additionally, the lens light \mathbf{l} , which is often partially blended with the lensed features, must also be modeled, either

Table 1. List of the main notations used in the paper.

Notation	Definition
<i>Observation</i>	
\mathbf{d}	Imaging data
\mathbf{C}_d	Data covariance matrix
<i>Model components</i>	
\mathbf{m}	Image model
\mathbf{B}	Blurring operator
\mathbf{L}_ψ	Lensing operator
ψ, s, ℓ	Lens potential, source light, lens light
<i>Model parametrization</i>	
η	Vector containing all model parameters
$\eta_\psi, \eta_s, \eta_\ell$	Parameters for analytical profiles $\tilde{\psi}, \tilde{s}, \tilde{\ell}$
ψ_{pix}	Pixel values of the pixelated potential
$\lambda_{\psi, \text{bl}}, \lambda_{s, \text{bl}}$	Regularization strengths for ψ_{pix}
$\mathbf{W}_{\psi, \text{bl}}, \mathbf{W}_{s, \text{bl}}$	Regularization weights for ψ_{pix}
<i>Optimization & Inference</i>	
$L, \nabla L$	Loss function and its gradient
η_{MAP}	MAP solution
\mathbf{C}_η	Parameter covariance matrix

jointly with the lens potential and the source, or carefully subtracted from the data beforehand.

The mapping between an observed angular position θ on the sky and a corresponding position on the source plane β before lensing is given by the lens equation

$$\beta = \theta - \alpha(\theta), \quad (1)$$

where the reduced deflection angle α is given by the first derivatives of the lens potential

$$\alpha(\theta) = \nabla \psi(\theta). \quad (2)$$

It is possible to infer the (dimensionless) surface mass density of the lens from the lens potential, called the convergence κ , which can be computed from the second derivatives of the lens potential, as

$$\kappa(\theta) \equiv \frac{\Sigma(\theta)}{\Sigma_{\text{crit}}} = \frac{1}{2} \nabla^2 \psi(\theta), \quad (3)$$

where Σ is the surface mass density in physical units and Σ_{crit} the critical density that depends on the cosmology and the redshifts of the lens and source galaxies. For quoting quantities in physical units (e.g., masses), we assume a flat Λ CDM model with Hubble constant $H_0 = 70$ and matter density at present time $\Omega_{m,0} = 0.3$.

The lens equation holds for any light ray from the source to the observer, and as such describes the lensing effect in a continuous way. However, the data is pixelated and includes additional nuisance effects due to limiting seeing conditions and instrumental noise. We thus discretize the problem and write \mathbf{d} as

$$\mathbf{d} = \mathbf{B} \mathbf{L}_\psi \mathbf{s} + \mathbf{B} \ell + \mathbf{n}, \quad (4)$$

where \mathbf{s} and ℓ are vectors holding the true (discretized) surface brightness values of the unlensed source and the lens, respectively. The lensing operator \mathbf{L}_ψ encodes a discrete version of Eq. (1) that models the lensing effect by mapping source surface brightness values onto the lens plane, based on the lens potential ψ (which can also be discretized) through Eq. (2). The

blurring operator \mathbf{B} models the seeing conditions by convolving the images of the lens and the lensed source with the point spread function (PSF) of the instrument. Throughout this paper, we assume that the PSF has been modeled beforehand (e.g., from stars in the field) and that it remains constant in Eq. (4). The last term, \mathbf{n} , represents additive noise, and is usually a combination of instrumental read-out noise (Gaussian noise) and shot noise (Poisson noise that depends on \mathbf{s} and ℓ). We note that while the model depends linearly on \mathbf{s} and ℓ , the lensing operator \mathbf{L}_ψ depends non-linearly on ψ through the lens equation (Eq. (1)).

2.2. Lens modeling and inversion of the lens equation

The model defined in Eq. (4) cannot easily be inverted to retrieve the lens potential, the unlensed source light, and the lens light. The problem is ill-posed and subject to degeneracies that cannot be mitigated based only on the data: (1) some locations in the source plane do not map to any data pixel (thus are unconstrained), (2) the lens potential is directly constrained only at locations where there are lensed features, (3) the lens light is often blended with these same lensed features. Therefore additional constraints, based on priors on ψ , \mathbf{s} and ℓ , are needed to solve Eq. (4) and obtain physically motivated solutions. This is particularly important in situations where the number of model parameters is comparable to the number of pixels in the data, which is a common situation when modeling complex sources or perturbations in the lens potential.

Before specifying the choice of priors, we first simplify the notation by defining the model \mathbf{m} and the corresponding full set of parameters η that describe ψ , \mathbf{s} and ℓ as

$$\mathbf{m}(\eta) \equiv \mathbf{m}(\psi, \mathbf{s}, \ell) = \mathbf{B} \mathbf{L}_\psi \mathbf{s} + \mathbf{B} \ell. \quad (5)$$

Inverting Eq. (4) comes down to obtaining the set of maximum a posteriori (MAP) parameters η_{MAP} that maximizes the posterior probability distribution $P(\eta | \mathbf{d}, \mathbf{m})$ of the parameters given the data. From Bayes theorem we have

$$P(\eta | \mathbf{d}, \mathbf{m}) = \frac{P(\mathbf{d} | \eta, \mathbf{m}) P(\eta | \mathbf{m})}{P(\mathbf{d} | \mathbf{m})}, \quad (6)$$

where the first term of the numerator is the data likelihood, the second term is the prior, and the denominator is the Bayesian evidence. The evidence does not change η_{MAP} , but is particularly relevant for comparing different models (through evidence ratios).

In practice, we obtain η_{MAP} by minimizing a loss function $L(\eta)$ defined as

$$\begin{aligned} L(\eta) &= -\log P(\mathbf{d} | \eta, \mathbf{m}) - \log P(\eta | \mathbf{m}) \\ &\equiv \mathcal{L}(\mathbf{d} | \eta, \mathbf{m}) + \mathcal{P}(\eta), \end{aligned} \quad (7)$$

where the data-fidelity term \mathcal{L} is the negative log-likelihood, and the regularization term \mathcal{P} is the negative log-prior. The MAP solution can then be obtained as

$$\eta_{\text{MAP}} = \arg \min_{\eta} L(\eta). \quad (8)$$

The choice of the data-fidelity term is tightly linked to the statistical properties of the data noise \mathbf{n} , characterized by the covariance matrix of the data \mathbf{C}_d . The noise is composed of instrumental readout noise and shot noise due to both the flux from the observed target and the sky brightness, which we assume follow Gaussian distributions. The target shot noise is estimated from the “modeled” flux \mathbf{m} , as estimating it from the data itself can introduce biases (Horne 1986); therefore we formally have $\mathbf{C}_d \equiv \mathbf{C}_d(\eta)$ that we implicitly assume throughout

the following equations to avoid clutter. Moreover, we assume uncorrelated noise as is usually true for charge-coupled device images, hence \mathbf{C}_d is a diagonal matrix. A diagonal element of \mathbf{C}_d for a given data pixel p is given by

$$\sigma_{d,p}^2 = \sigma_{\text{bkg}}^2 + \frac{m_p}{t_{\text{exp}}}, \quad (9)$$

where σ_{bkg}^2 is the variance of the background noise (readout and sky brightness), m_p is the modeled flux \mathbf{m} at pixel p (in electrons per second) and t_{exp} is the exposure time. The last term of the equation is the Gaussian approximation of the shot noise variance (Poisson noise) due to the lens and source flux.

Under the assumption of Gaussian noise, the data-fidelity term is the χ^2 of the data given the model

$$\begin{aligned} \mathcal{L}(d|\boldsymbol{\eta}, \mathbf{m}) &= \frac{1}{2} \chi^2 \\ &= \frac{1}{2} \left[\mathbf{d} - \mathbf{m}(\boldsymbol{\eta}) \right]^\top \mathbf{C}_d^{-1} \left[\mathbf{d} - \mathbf{m}(\boldsymbol{\eta}) \right]. \end{aligned} \quad (10)$$

The priors and corresponding regularization terms depend on the specific choices of parametrization of model components $\boldsymbol{\psi}$, \mathbf{s} and $\boldsymbol{\ell}$. In general, each of these model components can be described by a set of analytical profiles, pixelated profiles, or a combination of both.

2.3. Analytical and pixelated components

We model smooth mass and light distributions with a set of analytical profiles. In this case, regularization terms in Eq. (7) are not explicitly defined, but rather directly encoded in the parametrization of the model \mathbf{m} , which we write as

$$\boldsymbol{\psi} \equiv \tilde{\boldsymbol{\psi}}(\boldsymbol{\eta}_\psi), \quad \mathbf{s} \equiv \tilde{\mathbf{s}}(\boldsymbol{\eta}_s), \quad \boldsymbol{\ell} \equiv \tilde{\boldsymbol{\ell}}(\boldsymbol{\eta}_\ell), \quad (11)$$

where $\boldsymbol{\eta}_\psi$, $\boldsymbol{\eta}_s$, and $\boldsymbol{\eta}_\ell$ are the set of parameters for the analytical profiles $\tilde{\boldsymbol{\psi}}$, $\tilde{\mathbf{s}}$ and $\tilde{\boldsymbol{\ell}}$ which describe the lens potential, the source and the lens light, respectively.

To describe more complex mass and light distributions that cannot be captured by analytical functions, we rely on pixelated components, where pixel values represent the lens potential, the lens light, or source light (in source plane) at each pixel position. We adopt the following notation for those components

$$\boldsymbol{\psi} \equiv \boldsymbol{\psi}_{\text{pix}}, \quad \mathbf{s} \equiv \mathbf{s}_{\text{pix}}, \quad \boldsymbol{\ell} \equiv \boldsymbol{\ell}_{\text{pix}}. \quad (12)$$

Such pixelated components typically imply a much larger number of parameters (i.e., each single pixel value). The model inversion is then highly under-constrained, and the choice of regularization plays a central role in the success of the method in recovering the underlying lens potential and light distributions (see e.g., [Vernardos & Koopmans 2022](#)).

2.4. Multiscale regularization of pixelated components

We employ a multiscale strategy based on wavelet transforms and sparsity constraints to regularize the pixelated components of the model. While in [Galan et al. \(2021\)](#) we focused on the source model, in this work we consider a pixelated component only in the lens potential. To clarify the relationship between the two works, we first recall the principle of the technique in the context of the source reconstruction, then we apply it to the case of the lens potential.

Our regularization strategy is based on wavelet transforms, which decompose an image into a set of wavelet coefficients

organized by spatial scale. Each of these scales is a filtered version of the signal (i.e., same number of pixels) that contains emphasized features at a given spatial scale, similar to a frequency decomposition using the Fourier transform. In [Galan et al. \(2021\)](#), we use the following regularization term (Eq. (7)) for the pixelated source component \mathbf{s}_{pix} :

$$\mathcal{P}_s(\mathbf{s}_{\text{pix}}) = i_{\geq 0}(\mathbf{s}_{\text{pix}}) + \lambda_s \left\| \mathbf{W}_s \circ \boldsymbol{\Phi}^\top \mathbf{s}_{\text{pix}} \right\|_1. \quad (13)$$

The first term in the above equation is a positivity constraint¹ that enforces pixel values to be nonnegative. The second term combines the ℓ_1 -norm $\|\cdot\|_1$ with the wavelet transform operator $\boldsymbol{\Phi}^\top$. The effect of the ℓ_1 -norm is to impose a sparsity constraint on wavelet coefficients, which is effectively equivalent to a soft-thresholding² of the coefficients ([Starck et al. 2015](#)). The threshold level depends directly on the hyper-parameter λ_s , which is further adapted to each wavelet scale through the weight matrix \mathbf{W}_s to efficiently regularize features that span different spatial scales in the source plane (the operation \circ represents the element-wise product). We note that a similar regularization term $\mathcal{P}_\ell(\boldsymbol{\ell}_{\text{pix}})$ can also be written for reconstructing the lens light ([Joseph et al. 2019](#)).

The success of this regularization strategy thus relies on our ability to correctly estimate the regularization weights \mathbf{W}_s . We follow a data-driven approach that relies on the noise, in the goal to control the statistical significance of the reconstructed source light distribution via λ_s ([Paykari et al. 2014](#)). We achieve this by estimating the standard deviation of the noise in the source plane for each wavelet scale. The details of this procedure are given in [Joseph et al. \(2019\)](#). The only remaining hyper-parameter is the overall strength of the regularization λ_s , which is a scalar and usually set between 3 and 5 to ensure high enough statistical significance of the reconstructed source light distribution (e.g., [Starck et al. 2007](#)).

In this work, we aim to regularize the pixelated potential component $\boldsymbol{\psi}_{\text{pix}}$. While in principle, this component can be used to model either the full lens potential or only perturbations to the underlying smooth potential, we focus on the latter case and use $\boldsymbol{\psi}_{\text{pix}}$ to reconstruct various types of perturbations. The regularization needs to be flexible enough to allow for a variety of different features to be reconstructed, ranging at least from localized sub-halos to populations of subhalos and multipolar moments. Therefore, we expect that a multiscale regularization strategy similar to \mathbf{s}_{pix} can be applied on $\boldsymbol{\psi}_{\text{pix}}$ as well. However we do not impose a positivity constraint on the values of $\boldsymbol{\psi}_{\text{pix}}$, as negative potential pixels correspond to a local decrease of the potential, relative the smooth potential component.

The full regularization term for $\boldsymbol{\psi}_{\text{pix}}$ is

$$\begin{aligned} \mathcal{P}_\psi(\boldsymbol{\psi}_{\text{pix}}) &= \lambda_{\psi,\text{st}} \left\| \mathbf{W}_{\psi,\text{st}} \circ \boldsymbol{\Phi}_{\text{st}}^\top \boldsymbol{\psi}_{\text{pix}} \right\|_1 \\ &\quad + \lambda_{\psi,\text{bl}} \left\| \mathbf{W}_{\psi,\text{bl}} \circ \boldsymbol{\Phi}_{\text{bl}}^\top \boldsymbol{\psi}_{\text{pix}} \right\|_1, \end{aligned} \quad (14)$$

where we use two different wavelet transforms: the starlet transform ($\boldsymbol{\Phi}_{\text{st}}^\top$) and the Battle-Lemarié wavelet transform of order 3 ($\boldsymbol{\Phi}_{\text{bl}}^\top$). As in Eq. (13), the scalars $\lambda_{\psi,\text{st}}$ and $\lambda_{\psi,\text{bl}}$ are the regularization strengths, and the elements of the matrices $\mathbf{W}_{\psi,\text{st}}$ and $\mathbf{W}_{\psi,\text{bl}}$ are the regularization weights. The starlet transform is the transform used for modeling the source light distribution in [Galan et al. \(2021\)](#), and is well-suited for the reconstruction of multiple locally isotropic features at different spatial scales.

¹ The indicator function $i_{\geq 0}(\cdot)$ is formally equal to 0 if its argument contains only nonnegative values, and $+\infty$ otherwise.

² Soft-thresholding is the proximal operator of the ℓ_1 -norm.




Fig. 2. Example maps illustrating the computation of the regularization weights for Eq. (14), to properly scale the multiscale regularization of the pixelated lens potential model ψ_{pix} . *Top-left panel:* noise standard deviation estimated from a preliminary smooth model of the source, that is the square root of the diagonal elements of \mathbf{C}_d , without the contribution from the lens light. The remaining panels show the noise standard deviation propagated to each potential pixel for the first four wavelet scales of the starlet transform, that we used as the regularization weights $\mathbf{W}_{\psi,\text{st}}$. We see that the standard deviation strongly decreases as a function of the wavelet scale. The resulting maps are similar for the Battle-Lemarié wavelet transform. More details are given Sect. 2.4.

The Battle-Lemarié transform is introduced to reduce the appearance of spurious isolated potential pixels, inspired by a similar use for mass-mapping studies from weak lensing observations (Lanusse et al. 2016).

Similarly to the regularization of the source, we aim to find the regularization weights $\mathbf{W}_{\psi,\text{st}}$ and $\mathbf{W}_{\psi,\text{bl}}$ based on the noise in the data. The situation is however more complicated than for the pixelated source reconstruction discussed in Sect. 2.4, because the relation between the lens potential and the lensed source light distribution is highly nonlinear due to the lens equation. Nevertheless, in the limit of small perturbations $\delta\psi$ to the smooth potential ψ , it is possible to linearize the lens equation and relate changes in the lens potential to changes in the lensed light distribution of the source. We use this first-order treatment for estimating the covariance – more specifically the standard deviation – in the lens potential, and scaling the regularization strengths accordingly.

A first-order Taylor expansion of the lens equation around β leads to the following equation (Blandford et al. 2001; Koopmans 2005)

$$\delta d(\theta) \approx -\nabla_{\beta} s(\beta) \cdot \nabla_{\theta} \delta\psi(\theta), \quad (15)$$

where residuals $\delta d \equiv d - \tilde{m}(\psi, s)$ are based on a preliminary model \tilde{m} that does not include any perturbations $\delta\psi$ in the lens potential. Gradients are computed with respect to the coordinates indicated by the subscripts β and θ (defined in Eq. (1)). The above equation provides a linear relation that connects individual pixels in potential space to individual pixels in data space, which we can use to propagate noise levels for tuning the regularization strengths. We give all the remaining details of the computation in Appendix A. The resulting weight matrices $\mathbf{W}_{\psi,\text{st}}$ and $\mathbf{W}_{\psi,\text{bl}}$ contain the standard deviation of the noise for each wavelet scale in the lens potential. We show an example of weights with respect to the starlet transform (i.e., $\mathbf{W}_{\psi,\text{st}}$) in Fig. 2.

As for the source, the remaining hyper-parameters are the two regularization strengths $\lambda_{\psi,\text{st}}$ and $\lambda_{\psi,\text{bl}}$, which directly control the statistical significance of the starlet and Battle-Lemarié wavelet coefficients, respectively. These scalars are set in practice between 3 and 5, depending on how strongly certain features need to be regularized.

2.5. Optimization with differentiable programming

So far we have only defined how the different model components (ψ , s , ℓ) can be parametrized and cast into an optimization problem. However, combining components which are described either with analytical profiles or on pixelated grids

is a challenging task, as the standard methods to minimize the loss function can be fundamentally different in each case. With analytical model components, the MAP solution η_{MAP} can be approached via stochastic algorithms such as particle swarm optimization (PSO, Kennedy & Eberhart 2001). With wavelet-regularized pixelated components (Eqs. (13) and (14)), convergence to η_{MAP} usually requires the use of carefully chosen iterative algorithms relying on the formalism of proximal operators to apply constraints such as sparsity and positivity to the solution (so-called proximal splitting algorithms, see e.g., Starck et al. 2015).

In Galan et al. (2021), we used a hybrid scheme that first optimizes the smooth lens potential described analytically using a PSO, followed by a source reconstruction step (at fixed lens potential parameters) using an iterative proximal algorithm. However, this strategy does not scale well with model complexity. Each additional pixelated model component would require similar hybrid schemes, which rapidly become inefficient at converging to the MAP solution. In addition, estimating the joint posterior distribution of model parameters using traditional sampling techniques (e.g., Markov chain Monte Carlo, MCMC) becomes computationally too expensive. In this work, we overcome this issue by implementing a fully differentiable loss function, meaning that we can obtain its full gradient and higher order derivatives analytically. This allows us to simultaneously optimize all parameters, both analytical parameters and individual pixel values, using robust gradient descent algorithms which remain efficient even in a large parameter space. This naturally replaces proximal algorithms that are usually necessary to solve for pixelated components, and leads to a self-consistent combination of both analytical and pixelated model components.

Gradient descent optimization guarantees convergence to a minimum of the loss function, which is typically not the case for stochastic optimization algorithms that do not use the gradient (or higher order derivatives) of the loss function. However there is still no definitive guarantee to converge to the global minimum (i.e., η_{MAP}), which can depend on the parameter initial values. It is possible to address this limitation by using a multi-start gradient descent optimization, which runs the same minimization multiple times for different parameter initializations (e.g., Gu et al. 2022). We leave this optimization improvement to future work.

We use the automatic differentiation Python library JAX (Bradbury et al. 2018) to construct a fully differentiable model and loss function. Wavelet transforms are implemented using convolutions, making them straightforwardly differentiable (as in convolutional neural networks). We also take advantage of efficient compilation features of JAX to speed up computations

during evaluation of the loss function and its derivatives. All of our modeling methods and algorithms are implemented in the Python software package HERCULENS, which we make publicly available³. The code structure and part of the modeling routines of HERCULENS are based on the open-source modeling software package LENSTRONOMY (Birrer & Amara 2018; Birrer et al. 2021). All modeling and analysis scripts are also publicly available⁴.

2.6. Estimating the parameter covariance matrix

Estimating parameter uncertainties and covariances is crucial to reliably interpret the model that corresponds to the MAP solution. Sampling techniques such as MCMC, which draw samples from the full posterior distribution function (Eq. (6)), are often used for this purpose. However, for large parameter spaces, such stochastic techniques become inefficient, especially when parameters depend non-linearly on each other, as is the case in lensing.

In this regime, Hamiltonian Monte Carlo sampling (HMC, introduced by Duane et al. 1987; Neal 2011, for a review) is particularly efficient as each new sample is drawn based on the gradient of the loss function, resulting in high acceptance rates. While individual HMC steps might be more expensive to compute, the number of samples required for a reliable estimation of the parameters' posterior distributions is largely reduced compared to stochastic techniques such as MCMC, which tend to deliver noisier distributions. Moreover, we note that gradient-informed nested sampling for calculating the Bayesian evidence also exists (see e.g., Albert 2020), which is an important tool for model comparison (although not explored in this work).

In addition to HMC sampling, we also explore the possibility of using the Fisher information matrix (FIM) and its relation to the second-order derivatives of the loss function to estimate the parameters' covariance matrix. This is usually referred to as a Fisher analysis, which has been used in various studies including the forecast of constraints on cosmological parameters (e.g., Philcox et al. 2021), the mapping of large-scale structures (e.g., Abramo 2012), the analysis of gravitational waves (e.g., Belgacem et al. 2019), and the substructure power-spectrum from lensing (Cyr-Racine et al. 2019).

Second-order partial derivatives of the loss function L define the Hessian matrix, which reflects the local shape of the loss function at any point of the parameter space. Each entry of the Hessian matrix \mathbf{H}_L is defined as

$$\mathbf{H}_L(\boldsymbol{\eta}) = \left\{ \frac{\partial^2 L(\boldsymbol{\eta})}{\partial \eta_i \partial \eta_j} \right\}, \quad (16)$$

where the indices i and j indicate different model parameters from the entire parameter set $\boldsymbol{\eta}$. The “expected” Fisher information matrix, denoted as \mathcal{I} , is defined as

$$\mathcal{I}(\boldsymbol{\eta}) = \mathbb{E}[\mathbf{H}_L(\boldsymbol{\eta})], \quad (17)$$

where \mathbb{E} is the expectation operator over the distribution of all the realizations of the data for a fixed set of parameters $\boldsymbol{\eta}$ (i.e., $P(\mathbf{d}|\boldsymbol{\eta})$). In practice, however, we cannot compute the expected value $\mathcal{I}(\boldsymbol{\eta})$, because we only have access to a single observation \mathbf{d} , with a specific realization of the noise. Instead, we can compute the “observed” FIM, $\mathcal{I}(\boldsymbol{\eta})$, evaluated at the MAP solution

$$\mathcal{I}(\boldsymbol{\eta}_{\text{MAP}}) = \mathbf{H}_L(\boldsymbol{\eta}_{\text{MAP}}). \quad (18)$$

The FIM can then be used to approximate the covariance matrix of the MAP solution through matrix inversion:

$$\mathbf{C}_\eta \approx \mathcal{I}(\boldsymbol{\eta}_{\text{MAP}})^{-1}. \quad (19)$$

This covariance directly gives a lower bound on the uncertainty for each parameter⁵. We note that this (first-order) approximation of \mathbf{C}_η becomes exact if the loss function locally behaves quadratically or, equivalently, follows a Gaussian distribution. We expect this to be the case for simple, smooth models described with analytical profiles, as shown in Vernardos & Koopmans (2022). We find that while for fully analytical models we can rely on the above first-order approximation of parameter covariance matrix, a sampling-based exploration of the highly nonlinear parameter space using HMC is warranted.

3. Experimental setup

We test our method on mock *Hubble* Space Telescope (HST) observations that include realizations of the three perturbation categories introduced in Sect. 1: a localized DM subhalo, a population of DM subhalos, and higher-order multipoles. We simulate these perturbing fields based on physical considerations and results of previous works (detailed in the sections below), which result in each case in very different levels of perturbations when compared to the underlying smooth lens potential.

We focus on the strong lensing of a source with smooth surface brightness, resembling an early-type galaxy lensed by a foreground early-type galaxy (so-called EEL, Oldham et al. 2017). We simulate both the source and the lens light as single elliptical Sérsic profiles, as early-type galaxies are often well-fitted by these profiles (e.g., Shu et al. 2017). The smooth lens potential is described by a singular isothermal ellipsoid (SIE) for the main deflector, embedded in an external shear to simulate the net influence of neighboring galaxies. All analytical profiles used in this work are defined in Appendix B.

To create the mock data, we use an algorithm that is entirely independent of our modeling code, in order to prevent the occurrence of artificially advantageous minima during optimization. This also has the advantage to closely mimic a real-world situation, as the data never exactly corresponds to any model generated by the modeling code itself. We use the software package MOLET (Vernardos 2022) to simulate typical observations of strongly lensed galaxies as observed with HST and the Wide Field Camera 3 (WFC3) instrument, in the near infrared ($F160W$ filter). The pixel size is $0''.08$, and the field of view is $8'' \times 8''$. In this work we do not explore effects due to incorrect PSF modeling, hence for simplicity we use a Gaussian PSF with $0''.3$ FWHM. This results in a simulated data set that is sufficiently realistic to evaluate our method. The first column of Fig. 3 shows the simulated image without perturbations to the smooth potential. The remaining three columns show the different perturbation cases to which we apply our method. Full details of all instrumental settings and input model parameters are listed in Appendix C.

3.1. Localized subhalo (LS)

We simulate a single localized DM subhalo, with a spherical isothermal profile (SIS) and Einstein radius $\theta_{\text{E,halo}} = 0''.07$ as

⁵ This is more formally called the Cramér-Rao bound, which states that the variance of an unbiased parameter estimator is at least as high as the inverse of the Fisher information (e.g., Cramér 1999).

³ <https://github.com/austinpeel/herculens>

⁴ <https://github.com/aymgal/wavelet-lensing-papers>

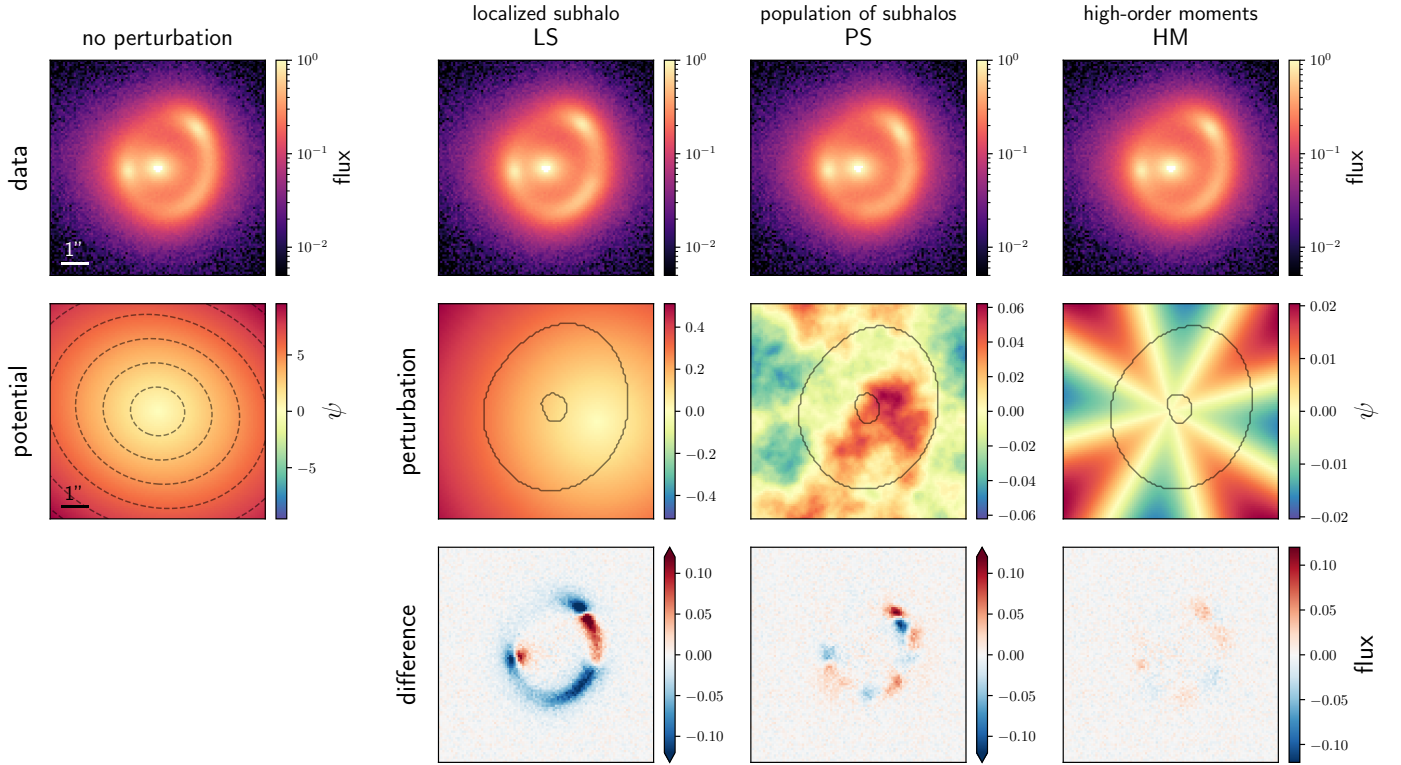


Fig. 3. Overview of the simulated HST observations. The left-most column shows the simulated data without perturbations to the smooth lens potential. The dashed lines in the *bottom-left panel* are isopotential contours. The remaining columns show the mock data used in this work (*top row*), the three different input potential perturbations (*middle row*), and the difference between the unperturbed and perturbed lenses (*bottom row*). The range (min and max) of the perturbations varies according to the different nature of the perturbers. The solid lines enclose the region where the S/N of the lensed source is higher than 5. The isopotential contours are almost indistinguishable between the unperturbed and perturbed cases, therefore we omit them.

input perturbing potential (for reference, the Einstein radius of the main deflector is $1''.6$). For simplicity, we assume that the subhalo is in the same redshift plane as the main lens, although it can be located at any redshift (through a simple scaling of its mass). We assume typical lens and source redshifts of EEL lenses $z_d = 0.3$ and $z_s = 0.7$ respectively (Oldham & Auger 2018), from which we get a mass of $10^9 M_\odot$ within the SIS Einstein radius, which is comparable to previous dark halo detections based on HST observations (Vegetti et al. 2010). The resulting simulated observation is shown in the second column of Fig. 3. The mean perturbation level, computed within the region containing the lensed arcs, is 6.4% relative to the smooth potential.

3.2. Population of subhalos (PS)

We follow recent work and simulate the net effect of a population of DM subhalos along the line-of-sight with a Gaussian random field (GRF; Chatterjee & Koopmans 2018; Bayer et al. 2018; Vernardos et al. 2020). GRF perturbations are random fluctuations that have a Fourier power-spectrum following a power-law. We parametrize this power-law relation as

$$\text{PS}(\psi_{\text{GRF}})(k) = C_\beta \sigma_{\text{GRF}}^2 k^{-\beta_{\text{GRF}}}, \quad (20)$$

where k is the wavenumber, β_{GRF} is the power-law slope, and C_β is a normalizing factor that depends on β_{GRF} and the size of the field of view, which ensures that σ_{GRF}^2 is the variance of the GRF (for the exact formula, see Chatterjee 2019). The power-spectrum is then converted to a specific random realization of

direct-space perturbations using the inverse Fourier transform. The value of β_{GRF} determines the distribution of power at each length scale: a large value leads to extended and smooth variations, whereas a small value creates a large number of localized and grainy structures.

Typical ranges for σ_{GRF}^2 and β_{GRF} have been explored in the literature and are justified in Chatterjee (2019) and Vernardos et al. (2020). In these works, the authors explored ranges $\sigma_{\text{GRF}}^2 \in [10^{-5}, 10^{-2}]$ and $\beta_{\text{GRF}} \in [3, 8]$. Additionally, Bayer et al. (2018) excluded GRF variance larger than $\sigma_{\text{GRF}}^2 = 10^{-2.5}$, based on HST observations of the strong lens system SDSS J0252+0039. Hence, we set $\sigma_{\text{GRF}}^2 = 10^{-3}$ and $\beta_{\text{GRF}} = 4$, such that it leads to a GRF that is not unphysically large, and contains both small and large-scale features. The specific GRF realization and the corresponding simulated observation are shown in the third column of Fig. 3. The corresponding relative mean perturbation level is 0.83%.

3.3. Higher-order multipoles (HM)

We introduce higher-order deviations to the smooth potential as a multipole of order 4 (octupole). We use the same definition of multipole as in Van de Vyvere et al. (2022b)

$$\psi_{\text{multipole}}(\theta) = \frac{r}{1 - m^2} a_m \cos(m\phi - m\phi_m), \quad (21)$$

where r and ϕ are polar coordinates transformed from θ . In our case, we fix the multipole order to $m = 4$ (octupole), and the orientation ϕ_4 to be aligned with the main axis of the SIE component of the smooth potential (in this case the full observation

becomes more disk-like). The octupole strength is set to $a_4 = 0.06$, which corresponds to the high end of the distribution found by Hao et al. (2006), from isophote measurements over a large sample of elliptical and lenticular galaxies in SDSS data. The resulting simulated observation is shown in the last column of Fig. 3. The corresponding relative mean perturbation level is 0.16%.

4. Modeling the full lens potential

4.1. Baseline model, parameter optimization and sampling

We model the simulated data set presented above, with the goal of retrieving the full lens potential, including the perturbations. The lens potential is modeled as $\psi = \tilde{\psi}(\eta_\psi) + \psi_{\text{pix}}$. The smooth component $\tilde{\psi}$ is parameterized as a SIE and external shear, and the a priori unknown perturbations are captured in the pixelated component ψ_{pix} , regularized with sparsity and wavelets as discussed in Sect. 2.4. The deflection angles at each position in the image plane are computed based on bicubic interpolation of ψ_{pix} . We use bicubic instead of bilinear interpolation in order to compute the surface mass density (Eq. (3)) corresponding to the pixelated model, as it requires second-order spatial derivatives of the potential (these derivatives are always zero with bilinear interpolation). We model the surface brightness of the source using a Sérsic profile, that is $s = \tilde{s}(\eta_s)$. This modeling choice means that we assume an accurate knowledge of the underlying shape of the source galaxy.

The lens light profile is modeled only once with a Sérsic profile for one of the system, assuming the other model components are known, then it is then fixed during the rest of the modeling. Fixing the lens light is not identical to subtracting it from the observation, as it still contributes to the noise model and reduces the contrast of the lensed features. We assume the lens light centroid is a good tracer of the mass centroid, which is a realistic scenario for fairly isolated lens galaxies (see e.g., Shajib et al. 2019). Therefore, we join the center of the SIE profile to that of the lens light profile. We note that it is however outside the scope of this work to assess the impact of inaccurate lens light modeling.

We model instrumental and seeing effects by assuming perfect knowledge of the PSF, background noise level (read-out noise and shot noise from sky brightness), however we estimate the shot noise from the modeled lens and source light distributions to estimate the diagonal of the data covariance matrix following Eq. (9).

When optimizing only analytical profile parameters (η_ψ, η_s) – which we do before including pixelated components in the model –, we find that the quasi-Newton optimization method BFGS⁶ (Nocedal & Wright 2006) is sufficient to reach convergence to the MAP solution. However, the optimization of both analytical and pixelated model components is more challenging, as the dynamic range of analytical and pixel parameters can vary significantly during optimization according to their impact on the loss function and the different regularization terms. Therefore, in this case, parameter updates are performed using the adaptive gradient descent algorithm ADABELIEF (Zhuang et al. 2020), which is extremely efficient for optimizing a large number parameters (typically as large as for convolutional neural networks). The initial learning rate is set in order to obtain a smooth decrease of the loss function until convergence, coupled with an exponential decay of the learning rate. We use the opti-

mization library OPTAX (Hessel et al. 2020) that implements the ADABELIEF algorithm.

For fully analytical models, we find that using the fast estimation from the FIM leads to a parameter covariance matrix almost indistinguishable from the one obtained via sampling methods. Therefore we only rely on the FIM for these simple models. However, for the more complex models that include a pixelated component, we find that HMC sampling of the parameter space is warranted to obtain reliable estimates of the posterior distributions. We use the python package BLACKJAX⁷ that provides an implementation of HMC well-integrated with JAX, that we run using the “No U-Turn Sampler” algorithm (NUTS, Hoffman & Gelman 2011) to dynamically adapt the step size, and their “Stan’s adaptation window” feature to improve sampling efficiency.

Our fiducial model is defined with a ψ_{pix} pixel size set to three times the data pixel size, leading to a resolution of $0''.24$. The choice of pixel size is based on preliminary models comparing the best-fit reconstructions and model residuals for different ψ_{pix} resolutions. While the residuals do not vary significantly for a pixel scale between 4 and 1.5, best-fit reconstructions obtained with pixel scales larger than 3 display artifacts at the scale of individual pixels. Although most of these artifacts can be efficiently reduced using our multiscale regularization strategy by increasing further the regularization strength for small wavelet scales, we find that using a larger number of model parameters for only marginal improvements in terms of residuals is not necessary for this work. We refer to Appendix D for further discussion on the choice of pixel scale for ψ_{pix} . The total number of parameters for our fiducial model is thus 1101 (1089 for ψ_{pix} , 5 for η_ψ and 7 for η_s), jointly optimized and constrained by 10^4 data pixels.

The regularization strengths for ψ_{pix} are set to $\lambda_{\psi, \text{st}} = 3$ and $\lambda_{\psi, \text{bl}} = 4$, that is in the range of values discussed in Sect. 2.4. We use a 1σ higher strength for the Battle-Lemarié regularization in order to penalize more the appearance of spurious pixels in the solution.

4.2. Modeling the perturbations only

We first reconstruct the perturbations in the idealized case where all smooth components are perfectly known and fixed. This unrealistic scenario allows us to assess the best level of perturbations that can be recovered, given the quality of the data set. As the source and the smooth potential are fixed, the regularization weights $\mathbf{W}_{\psi, \text{st}}$ and $\mathbf{W}_{\psi, \text{bl}}$ can be precomputed and kept constant throughout the gradient descent optimization.

The resulting models, corresponding to the MAP parameters η_{MAP} , are shown in Fig. 4. We also note that the pixelated model ψ_{pix} is expected to differ from the input by a uniform offset, as a uniform value in the potential corresponds to zero deflection of light rays (Eq. (2)). This constant offset is thus never constrained by the data alone (and depends on the initialization). Hence, for better visualizing the reconstruction for lens LS (for which we know that the input potential is positive), we shift the ψ_{pix} model such that the minimum pixel value displayed on the figure is zero. We do not add such an offset to our reconstructions for lenses PS and HM shown in Fig. 4, as the underlying perturbations have roughly zero-mean.

Overall the characteristic features of each type of perturbations are well recovered. For lens LS, both the subhalo position and the shape of the underlying SIS profile are captured by the

⁶ <https://docs.scipy.org/doc/scipy/tutorial/optimize.html>

⁷ <https://github.com/blackjax-devs/blackjax>

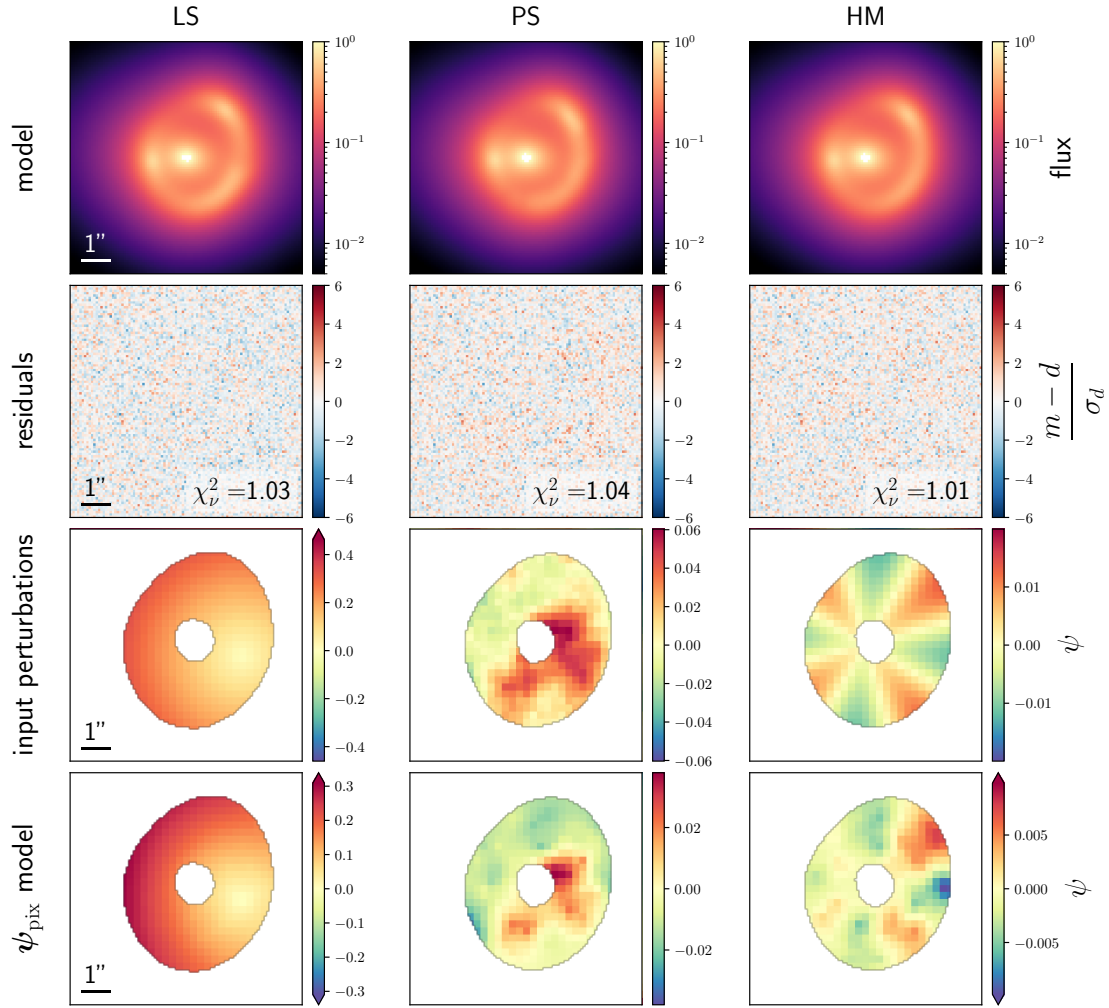


Fig. 4. Best-fit pixelated potential reconstruction of our simulated data set, assuming all other model components are perfectly known. Each column represents the different types of perturbations considered in this work. *From top to bottom:* simulated data, image model, normalized residuals, input perturbations, ψ_{pix} model. The outlined annular region on the potential panels corresponds to the solid lines in the panels of Fig. 3, and is where the S/N of the lensed source is higher than 5. Perturbations outside this area are cropped only to ease the visual comparison with the input perturbations (there is no cropping or masking during modeling). Additionally, for lens LS only, the minimum value of the perturbations is subtracted from the model, again to ease the comparison with the input SIS profile (this does not affect the lensing observables). An animated visualization of the gradient descent optimization is available [online](#).

model. For lens PS, the clear over-density region on the bottom right part of the arc is recovered, although the correspondence with the input perturbations is less than for lens LS. For lens HM, the reconstruction displays imprints of azimuthal periodicity between over- and under-density regions, despite the low level (0.2%) of perturbation. Lastly, we notice that the amplitude of the modeled perturbations is systematically lower, by a factor of ~ 2 , than the input perturbations. We note that a similar result can be seen in some of the pixelated reconstructions of Vernardos & Koopmans (2022, see their Fig. 8).

4.3. Modeling the full potential and the source

Our method is then applied to the more realistic situation in which the full lens potential and the source light are unknown as well. We do so by uniformly modeling the simulated data set using a series of steps involving gradient descent optimization. After each step, the model degrees of freedom are gradually increased, in order to prevent the optimizer to be trapped in local minima. These optimization steps are as follows.

Firstly, the pixelated component ψ_{pix} of the lens potential is initialized to zero (in practice to 10^{-8} , to prevent gradients to be evaluated at zero, see also Appendix E) and kept fixed. The initial model is thus fully smooth, with corresponding parameters η_{ψ} and η_s . We use a multistart gradient descent (as advocated by Gu et al. 2022) with 30 runs, for which we verified that it leads to convergence to the global minimum.

Secondly, the pixelated potential component is released and regularization strengths are deliberately set to large values, such that only the most significant potential pixels enter the solution. The regularization weights are computed based on the smooth model from the previous steps, and fixed throughout the gradient descent. This prevents the pixelated component from fitting all model residuals from the previous step, which can strongly bias the recovered perturbations. We find that setting $\lambda_{\psi, \text{st}} = 10$ and $\lambda_{\psi, \text{bl}} = 20$ leads to an intermediate solution that contains the main features of the perturbations, reducing slightly model residuals but preventing the model from getting trapped in a spurious minimum. All model parameters (η_{ψ} , η_s and ψ_{pix}) are simultaneously optimized.

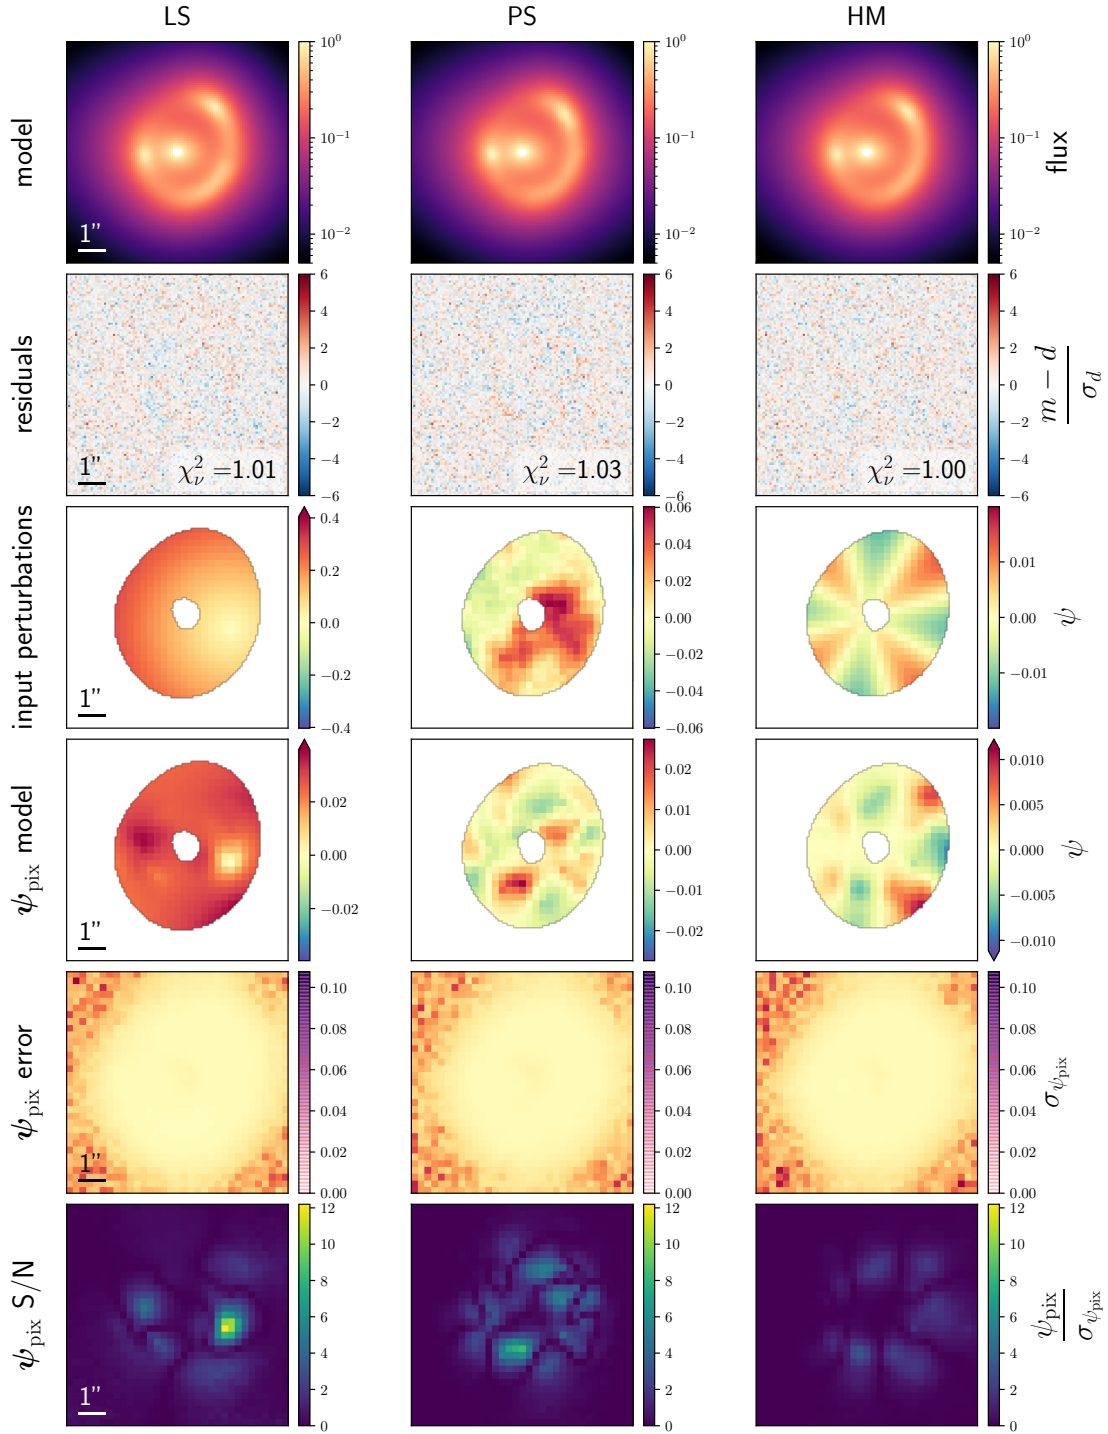


Fig. 5. Best-fit models on the simulated data set, after full modeling of the lens potential and the source. The first four rows are as in Fig. 4. The last two rows show, in order: the standard deviation for each ψ_{pix} pixel obtained from HMC sampling of the posterior, and the S/N maps that correspond to the absolute value of the modeled potential divided by the standard deviation.

Thirdly, the regularization strengths are set to their fiducial values ($\lambda_{\psi, \text{st}} = 3$, $\lambda_{\psi, \text{bl}} = 4$), and regularization weights are recomputed based on the previous models. All model parameters are simultaneously further optimized.

Lastly, we perform HMC sampling to estimate the error and posterior distributions for all parameters. We find that, thanks to the gradient information, only 400 samples (after a warm-up phase of 100 samples) are necessary to obtain well-sampled posteriors. For posterior distributions that have multiple modes,

which is not the case in our models, the number samples should be increased.

In Fig. 5 we summarize, for each lens, the models corresponding to the best-fit parameters η_{MAP} . The panels are identical to those of the ideal case (Fig. 4), although we added two additional panels to help the interpretation of the results. The first of these panels is the standard deviation for each ψ_{pix} pixel, based on the HMC samples, which can be interpreted as an error map for the best-fit model. These error maps are almost identical




Fig. 6. Reconstructed potential perturbations for different values of the regularization strength. *First row:* corresponds to a low regularization strength with $\lambda_{\psi,\text{st}} = 1$. *Second row:* shows our fiducial models of Fig. 5, with $\lambda_{\psi,\text{st}} = 3$. *Bottom row:* corresponds to a high strength with $\lambda_{\psi,\text{st}} = 5$. In all cases, the strength for the Battle-Lemarié wavelet ($\lambda_{\psi,\text{bl}}$) is set 1σ higher than for the starlet ($\lambda_{\psi,\text{st}}$) to penalize isolated spurious features in the reconstructed. The main features of the perturbations are recovered in all panels, so the model is robust against reasonable changes in regularization strengths. A too low regularization strength leads to over-fitting ($\chi^2_\nu < 1$) whereas a higher regularization strength leads to a smoother model, as expected.

for each modeled system, but are necessary to compute the second additional panel, which shows a measure of the signal-to-noise (S/N) of the reconstruction. We define the S/N for ψ_{pix} as the absolute value of the best-fit model divided pixel-wise by the standard deviation. These maps reveal the regions where the reconstruction is the most statistically significant.

For lens LS, the position of the subhalo is clearly and accurately recovered. However, the overall shape of the underlying smooth subhalo profile is not recovered equally well as in the ideal case. We also observe some features on the other side of the Einstein ring, although those are less significant than at the subhalo position. The S/N map clearly shows that the feature at the position of the subhalo has a higher likelihood compared to other features present in the solution.

For lens PS, we see a relatively good agreement with the input field, where both over- and under-dense regions are recovered. However, again, the amplitude of the perturbation is well below the input value, by a factor between 1.5 and 3. Contrary to the reconstruction of the ideal case, the over-dense region on the lower part of the ring is now the most prominent. This feature is well aligned with a lower-intensity region of the arc, which can partly explain why the model favors a correction to the smooth potential at this location. The S/N map confirms the significance of this feature. All the reconstructed features have a S/N of approximately 3 and above. This system is arguably the most complex to model, as the perturbations affect the lensed features at many different locations with different strengths and orientations, which is likely to translate to a larger set of possible solutions.

For lens HM, the reconstruction reveals a clear azimuthal periodicity, with no regions significantly different from the input perturbations. Interestingly, the reconstruction is very similar to

the ideal case, despite the input level of perturbations (relative to the smooth potential) being lower than in the other systems. Similar to lens PS, we notice a region where the model over-predicts a correction to the smooth potential. This region is better revealed in the S/N map and corresponds to the same low-flux region of the arc. This is not surprising as lower S/N in the data leads to looser constraints on the model parameters. For this lens we also notice that the model does not predict any perturbation at the position of left-most image of the source. At this location, which is closer to the lens center compared to other parts of the arc, the input perturbation level is lower while the data noise is higher (shot noise), which leads to almost no detection of perturbations.

4.4. Effect of the regularization strength

The models presented in the previous section correspond to well-motivated but fixed strengths for the regularization of the pixelated potential component. We investigate the effect of the regularization strength on the reconstructed perturbations by running the same modeling procedure with different values of $\lambda_{\psi,\text{st}}$ and $\lambda_{\psi,\text{bl}}$. Compared to our fiducial models with $\{\lambda_{\psi,\text{st}} = 3, \lambda_{\psi,\text{bl}} = 4\}$, we define a low regularization case $\{\lambda_{\psi,\text{st}} = 1, \lambda_{\psi,\text{bl}} = 2\}$ and a high regularization case $\{\lambda_{\psi,\text{st}} = 5, \lambda_{\psi,\text{bl}} = 6\}$. The resulting ψ_{pix} models and corresponding residuals are shown in Fig. 6.

Low-regularization models lead to a slight over-fitting as seen from the reduced χ^2 below unity. The reconstructed perturbations display higher frequency features, some being present in the input perturbations as in lenses LS and HM. Increasing the regularization strength filters out the high-frequencies, leading to smoother variations in the ψ_{pix} maps. For highest regularization strengths,

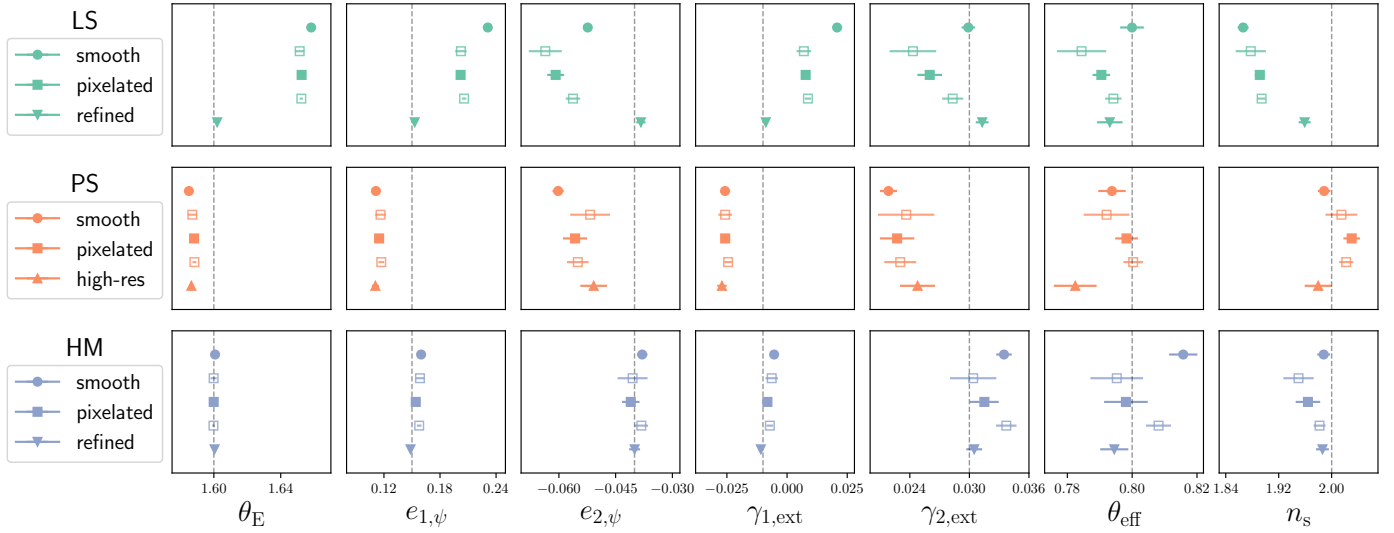


Fig. 7. MAP values and estimated uncertainties for the smooth potential parameters η_ψ and a subset of the source parameters η_s . From left to right: Einstein radius, complex ellipticities of the SIE, external shear components, source effective radius, and Sérsic index. The dashed lines indicate the input values. The “smooth” models correspond to models without including pixelated perturbations; “pixelated” models correspond to our baseline models shown in Fig. 5; “high-res” (lens PS) is similar to the fiducial model but with a smaller ψ_{pix} pixel size; “refined” corresponds to models where the pixelated perturbations have been replaced with an analytical profile (see text for more details). For “pixelated” models we also show model parameters obtained with different regularization strengths with empty square symbols, above (low strength) and below (high strength) the fiducial model (see Fig. 6). For several parameters, error bars are smaller than the marker size.

more features are visible, which is translated in higher χ^2 values. Moreover, the amplitude of the modeled perturbations decreases as the regularization strength increases, which is also expected since high-frequencies, which are suppressed, have the highest amplitudes. Overall, the main features of the perturbations are well-recovered for these three reasonable choices of hyper-parameters $\lambda_{\psi,\text{st}}$ and $\lambda_{\psi,\text{bl}}$. These results also confirm that our fiducial setting is close to be optimal since it provides a good compromise between over-fitting and correctly fitting the data.

4.5. Smooth potential parameters

In the previous section, we have seen that the main features of the perturbations are overall well recovered, but the reconstructions are not perfect despite model residuals almost at the noise level. Therefore, we expect that some of the inaccuracies in the pixelated potential model are absorbed by the smooth analytical component of the lens potential, or vice versa.

We check this hypothesis by comparing different models of the smooth potential, with and without including potential perturbations. As discussed earlier in this section, parameter uncertainties are either computed using the fast approximation of the FIM (fully smooth models), or based on HMC sampling of the parameter space (models including pixelated perturbations). The MAP values and uncertainties of a subset of analytical parameters are shown in Fig. 7, and compared against the input values. We discuss here only lens potential parameters, but the conclusions are similar for the source parameters as well, that we present in Appendix F.

We observe that fully smooth models display strong biases in almost all parameters as expected. Interestingly, models including a pixelated model in the potential, while having larger error bars, still lead to statistically significant biases. One particularly informative parameter is the inferred Einstein radius, as the value of θ_E is slightly more accurate after including the potential perturbations, but a substantial bias still remains. We attribute those

biases to a manifestation of the degeneracies that exist between the smooth component of the lens potential and the pixelated perturbations, where some adjustments of one component can compensate for the other, still leading to comparable model residuals (see [Vegetti et al. 2014](#), who also observed reabsorption of residuals in the macro potential model). Additionally, we see that the regularization strengths of the pixelated potential marginally impact the results. Parameter uncertainties are increased with lower regularization strength, which is expected as the model has effectively more degrees of freedom that are not fully constrained.

4.6. Analytically refined models

To assess if the biases discussed above can be mitigated, we reduce the parameter space by replacing the pixelated component with analytical profiles. In a real-world scenario, where the underlying type of perturbations is unknown, this would correspond to imposing stronger priors on the model, motivated by the characteristic features observed in the pixelated model. This strategy is similar to previous studies based on the gravitational imaging technique (e.g., [Vegetti & Koopmans 2009](#); [Vegetti et al. 2012](#)), where an analytical profile is optimized at the position of a tentative detection of a subhalo, in order to better characterize its properties (position, radial profile, mass).

Among the different systems modeled in this work, the most characteristic features that can be noticed in the pixelated models are for the LS and HM cases, with either a very localized decrease of the potential, or azimuthal periodicity centered on the lens galaxy, respectively. Therefore, for these two lenses, we start from the MAP solution obtained with our fiducial model and include a SIS profile (lens LS) or an multipole component (lens HM) instead of the pixelated component. We refer to these additional fully analytical models as “refined” in the remaining of the text.

The smooth potential parameters inferred from these refined models are compared to the smooth models in Fig. 7. We see

Table 2. Recovered properties of potential perturbations, from the MAP ψ_{pix} models shown on Fig. 5.

Lens	Parameters	Input values	Model	Measured values
LS	$(\theta_{x,\text{sub}}, \theta_{y,\text{sub}})$	$(1.90, -0.40)$	Pixelated	$(1.94 \pm 0.12, -0.49 \pm 0.12)$
			Refined	$(1.899 \pm 0.008, -0.408 \pm 0.005)$
	$\log_{10} M_{\text{sub}}$	8.91	Pixelated	8.71 ± 0.05
			Refined	8.90 ± 0.01
PS	$(\log_{10} \sigma_{\text{GRF}}^2, \beta_{\text{GRF}})$	$(-3.11 \pm 0.15, 3.23 \pm 0.05)$	Pixelated, ideal	$(-3.35 \pm 0.20, 3.40 \pm 0.15)$
			Pixelated, fiducial	$(-3.89 \pm 0.25, 4.15 \pm 0.19)$
			Pixelated, high-res	$(-3.49 \pm 0.15, 3.26 \pm 0.09)$
HM	(a_m, ϕ_m)	$(0.06, 82.50)$	Pixelated	$(0.024 \pm 0.002, 83.78 \pm 0.88)$
			Refined	$(0.061 \pm 0.002, 83.15 \pm 0.39)$
	m	4	Refined	4.01 ± 0.02

Notes. The quoted values for the subhalo mass is given in solar units and computed within the region with $S/N(\psi_{\text{pix}}) > 3$ (see Fig. 8 and text for more details). The input GRF parameters are given with uncertainties as they were fitted on the power-spectrum of the input perturbing field.

that the significant biases observed with models with too few or too many degrees of freedom have been correctly mitigated. This result is not surprising as the refined model is now parametrized identically to the simulated data. Nevertheless, this allows us to confirm that even after optimizing a more complex and possibly inaccurate model, it does not prevent us from accessing the optimal global solution after correctly identifying the underlying type of perturbations. In Sect. 5 we use these refined models to retrieve the properties of the underlying perturbations.

4.7. Higher resolution pixelated model

Contrary to the localized subhalo and multipolar structures, the underlying perturbations for lens PS are described by a GRF, which does not have a specific analytical profile as it is a random realization. Instead, we use this system to test if using a higher resolution grid for the pixelated component ψ_{pix} (i.e., more parameters) allows us to reduce the biases on smooth model parameters. The resulting model, named “high-res”, is compared to the fully smooth and fiducial models in Fig. 7. While this model is insufficient to recover unbiased smooth potential parameters, we subsequently see in Sect. 5 that the recovered power-spectrum of the perturbations is closer to the input one.

Nevertheless, we note that recovering the smooth potential parameters in the presence of perturbations such as a subhalo population modeled via a GRF is achievable by imposing informed priors on the pixelated potential model. This has been shown in the recent work of Vernardos & Koopmans (2022), which extended the gravitational imaging technique using a covariance-based regularization of the pixelated potential model. The covariance matrix governing the regularization term can be specifically adapted to GRF-like perturbations, leading to an effective regularization of the solution if the assumption matches the underlying perturbations. We plan to implement the strategy presented in Vernardos & Koopmans (2022) in the HERCULENS package, and leave for future works its comparison with the method presented here.

5. Constraints on the underlying perturbations

In the previous section we discussed in a qualitative manner the reconstructed perturbations. Here we seek to quantify the prop-

erties that can be recovered from these models, and discuss the robustness of those measurements as well as their applicability to real data sets. All the inferred quantities discussed in the next subsections are summarized in Table 2, and based on the fiducial models shown in Fig. 5.

5.1. Subhalo mass and position

We consider the two models of lens LS to quantify the properties of the underlying DM subhalo: our fiducial model including a pixelated component in the lens potential, and the refined model assuming the detected subhalo mass distribution follows a SIS profile.

For the pixelated model we assign the position of the detected subhalo to the minimum of the pixelated potential ψ_{pix} . We show the location of the pixel in the top left panel of Fig. 8. For related uncertainties, we compute the minimum of each HMC sample of ψ_{pix} , but we note that the minimum remains in the same potential pixel, leading to error bars smaller than its size. We thus turn to a more conservative estimate of the uncertainty and simply set it to half the pixel size (i.e., $0''.12$). For the refined model, we take the optimized position of the SIS as the position of the subhalo, with uncertainties estimated from the FIM. The resulting positions and error bars are listed in the top row of Table 2.

The mass of the subhalo is more difficult to estimate from our pixelated model. Nevertheless, achieving this would be powerful, as it does not require the choice of a specific shape for the subhalo mass distribution. We start by computing the surface mass density (i.e., the convergence) corresponding to each potential pixel using Eq. (3). This results in the pixelated convergence model κ_{pix} shown in the top right panel of Fig. 8. Next, we need to define a region in which to integrate κ_{pix} pixels before converting the surface mass density to proper solar mass units using Σ_{crit} . As discussed in Sect. 3.1, input parameters of the subhalo correspond to a subhalo mass of about 10^9 within the Einstein radius of $0''.07$. As this scale is much smaller than a pixel of our ψ_{pix} model, we cannot rely on summing the convergence pixels.

We address this issue by considering a larger region within which the subhalo mass can be inferred for both the analytical and pixelated models. We select a high-significance region of

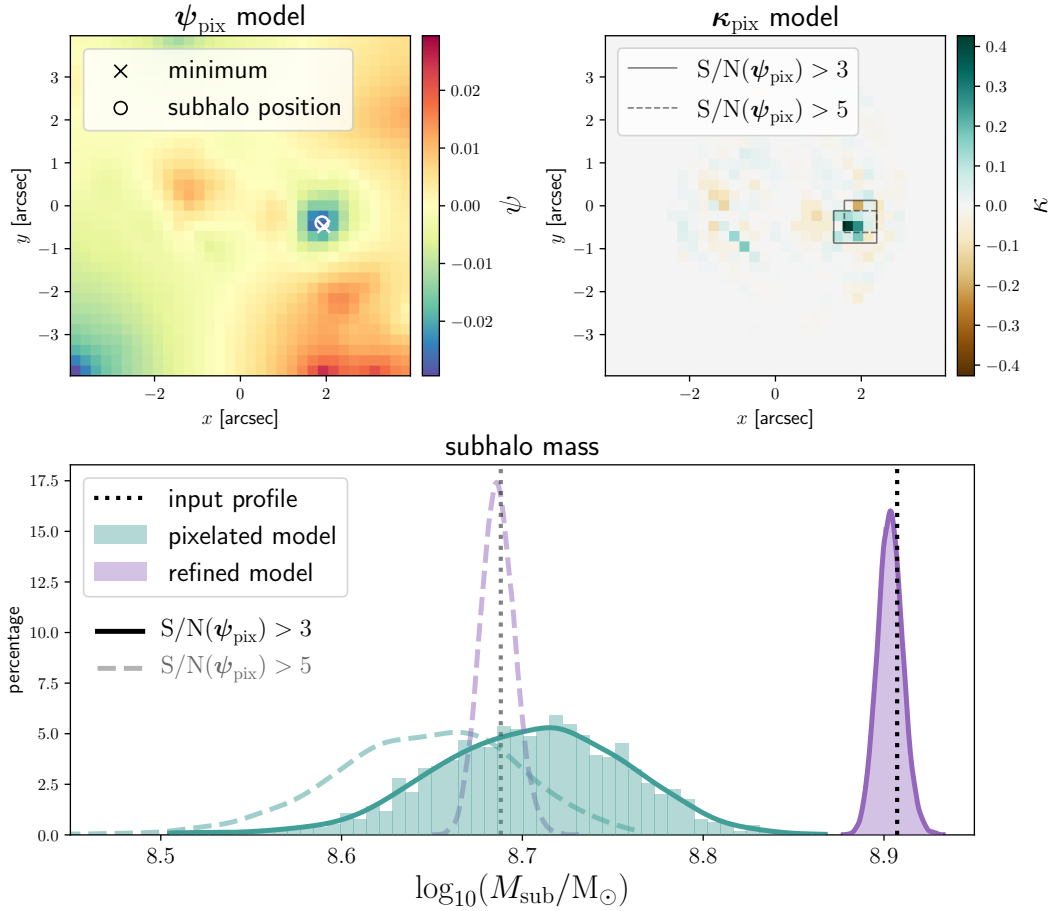


Fig. 8. Characterization of the position and mass of the localized subhalo based on our model of lens LS. *Top left:* ψ_{pix} model, from which we assign the detected subhalo position to the minimum of the clear negative feature, in good agreement with the input subhalo position. *Top right:* pixelated convergence model κ_{pix} obtained from the ψ_{pix} model. We indicate the fiducial pixelated region within which we compute the subhalo mass, defined as the pixels with $S/N(\psi_{\text{pix}}) > 3$ (see also the bottom panel of Fig. 5). A smaller region corresponding to $S/N(\psi_{\text{pix}}) > 5$ is also shown, in dashed lines. *Bottom:* posterior distributions of the subhalo mass, as estimated from the pixelated model, and from the refined model that replaces ψ_{pix} with a SIS profile for the subhalo. The two pairs of posterior distributions corresponding to each pixelated region are shown in continuous ($S/N(\psi_{\text{pix}}) > 3$) or dashed lines ($S/N(\psi_{\text{pix}}) > 5$). For each pixelated region, the black dotted line indicates the mass computed from the input convergence profile of the subhalo.

the reconstruction that contains all pixels with $S/N(\psi_{\text{pix}}) > 3$ (see bottom left panel of Fig. 5). This region contains 15 convergence pixels, that we sum and convert to proper units to estimate the subhalo mass. We repeat the same procedure for each sample of the joint posterior distribution (~ 1300 samples) and find the distribution shown in the bottom panel of Fig. 8. To compare the inferred value with that of the input subhalo, we discretize the input SIS profile by evaluating it on the same grid of pixels, and compute the mass as for the pixelated model. We note that the resulting “input” mass is lower than the one computed analytically within Einstein radius, because of the discretization of the SIS profile that diverges in the center.

We find that the inferred mean value of the subhalo is lower than the measured mass on the input perturbations. In addition, we find that the amplitude of the disagreement is significantly affected by the choice of the region in which we integrate the convergence. For instance, considering only κ_{pix} pixels with $S/N(\psi_{\text{pix}}) > 5$ instead of 3 (6 pixels instead of 15) leads in fact to a very nice agreement with the input value. This assumption corresponds to the dashed-line histograms in Fig. 8. The main reason of the better agreement is that this smaller region essentially excludes the few pixels with negative convergence, seen

in brown color in Fig. 8 (negative convergence pixels do have a physical meaning, as they indicate a local decrease of the lens mass relative to the smooth component). On one hand, this leads to a higher mass inferred from the model; on the other hand, the input value used as a reference is smaller, because it is computed within a smaller region. These two effects combined lead to an overall better agreement between the model and the input. We note that a similar behavior is observed when the ψ_{pix} regularization strength is too low (see Sect. 4.4). However, in this case, the pixelated convergence map is very noisy due to over-fitting the imaging data and the region inside which the subhalo mass is measured cannot be reliably defined.

Measuring the mass directly from the pixelated model is therefore challenging and possibly depends on additional assumptions. Currently, a more robust approach is to infer the subhalo mass from our refined model, which is based on an analytical profile for the subhalo. After applying the same procedure, we obtain the resulting posterior distribution shown in purple in Fig. 8, which is in perfect agreement with the input. Again, this is not surprising, as the SIS profile reflects well the underlying shape of the subhalo. Nevertheless, these results showcase the requirement of stronger priors on the shape of the




Fig. 9. Azimuthally averaged power spectra of input and modeled (ψ_{pix}) perturbations for lens PS (computed within the same mask as in Fig. 5). Linear relations used to estimate the variance σ_{GRF}^2 and slope β_{GRF} are overlaid (dashed lines), assuming the underlying perturbing field is a GRF (uncertainties are estimated from the least-squares fit). The crosses indicate the bin (2×10^{-1}) that is ignored for the linear fit, as it corresponds to the size of the mask.

subhalo in order to infer unbiased properties. Additionally, for proper inference on real data, several works have demonstrated the need to carefully compare different choices of subhalo profiles (e.g., Çağan Şengül et al. 2022; Despali et al. 2022).

5.2. Statistics of the subhalo population

We analyze our pixelated reconstruction for lens PS following a Fourier power-spectrum analysis, motivated by the assumption of GRF perturbations (Eq. (20)). We show in Fig. 9 the azimuthally averaged power-spectra from the three different pixelated models explored in this work: the “ideal” (Fig. 4), the “fiducial” (Fig. 5), and the “high-res” models (i.e., finer ψ_{pix} pixels, twice the data pixels). We compute the power-spectra inside the region of Fig. 5, in order to consider only features in the region of interest. We then compare the obtained power-spectra with those from the input perturbations by fitting linear relations in log-space and list the resulting best-fit values for σ_{GRF}^2 and β_{GRF} in Table 2. We note that the first bin is excluded from the linear fit because it corresponds to a wavenumber that translates to the roughly the size of the region used for computing the power-spectra, hence it is no informative. The quoted uncertainties are estimated from the least-square fit.

As expected, the model in the ideal case (i.e., with fixed smooth potential and source light) agrees very well with the input power-spectra, which translates in a good agreement for GRF parameters as well. Regarding our fiducial model, the amplitude is overall lower than the input, consistent with what we discussed in Sect. 4. At intermediate wavenumbers, the recovered power-spectrum is close to the input one, however it is strongly attenuated at large wavenumbers. This leads to an overall steeper slope, and translates to a $\sim 5\sigma$ difference in β_{GRF} with respect to the input. This attenuation of small spatial scales is fully mitigated by modeling the perturbation on a higher-resolution grid, that allows us to better model small scale features. Indeed, the power-spectrum of the high-res model exhibits an excellent agreement with the input for all wavenum-




Fig. 10. Measurement of the multipole orientation from the ψ_{pix} model of lens HM. An analytical octupole profile (plus a constant offset) is fitted directly on the reconstructed perturbations, and compared to the input one. Each pair of lines shows two of the main octupole axes, to ease the visualization. The gray dotted lines are corresponds to $\pm 3\sigma$ fits, where σ is estimated from the FIM.

bers $k > 4 \times 10^{-1} \text{ arcsec}^{-1}$. The inferred slope β_{GRF} is within 1σ with respect to the input value (see Table 2).

Overall, our pixelated method correctly retrieves the locations of the main perturbations that mimic a subhalo population, and allows us to obtain a first-order estimate of its statistics. Our results suggest that using a higher-resolution grid for the pixelated component allows to better recover the full power-spectrum. However, the precise characterization of the power-spectrum of the perturbations under the assumption of GRF is a challenging task that requires additional priors in the model (e.g., Bayer et al. 2018; Vernardos & Koopmans 2022).

5.3. Properties of multipolar structures

Based on our models of lens HM, we can recover the underlying octupole using two different methods: fitting an octupole directly on the ψ_{pix} model, or using the refined model that includes explicitly an octupole profile in addition to the other components of the lens potential (see Fig. 10).

We perform the octupole fitting (i.e., we fix $m = 4$) on the ψ_{pix} model via gradient descent with three free parameters, namely the amplitude, orientation and an additional constant offset (remember that this offset in the potential is not constrained by the data). After converging to the MAP solution, parameter uncertainties are estimated from the FIM. The recovered octupole orientation, reported in Table 2, agrees extremely well with the input. However, the recovered amplitude is lower, as expected from the pixelated reconstructions and discussed in Sect. 4.

With the refined model it is possible to measure the constraints on the multipole order m , in addition to the orientation and amplitude. The MAP values of the multipole component obtained with this refined model are reported in Table 2, which all agree very well with the input. The parameter m is expected to be more challenging to optimize, as it has nonlinear effects on the profile shape. We tested the robustness of the optimization to different initial values m , and found that initializing m to a value of 3 or higher leads to the correct MAP value, but setting it closer to 2 drives the model toward a quadrupole ($m = 2$), which is degenerate with the shear and ellipticity of the smooth potential.

6. Computation time

HERCULENS uses JAX to exactly differentiate the loss function and significantly decrease runtime (Sect. 2.5). The entire analysis of this work, including parameters optimization and sampling, was performed on a personal computer. We give average timings of the main modeling steps for a personal computer⁸:

- Optimizing the smooth analytical models (12 parameters), takes one minute for a multistart gradient descent with 30 starts (i.e., ~ 2 s for a single gradient descent).
- Computing the regularization weights (Eq. (14)) takes about 20 s. This step can be accelerated at least by a factor of two, by precomputing some of the operators involved in Eq. (15); we leave this for future improvements.
- Optimizing the idealized models of Fig. 4 in which only the pixelated potential component is optimized (1089 parameters) takes 30 s. This is for 10^3 iterations, which is sufficient to reach convergence.
- Compared to these idealized models, the run time is virtually identical for optimizing the full models of Fig. 5, due to the marginal increase in the number parameters (1101 parameters).
- Computing the FIM and its inverse typically takes 20 s.
- Performing HMC sampling for a total of 500 samples takes about 1.1 h for a single chain.

These numbers can be extended to the modeling of a typical HST observation of a strong lens. For instance, modeling the lens SDSS J1630+4520 (from the SLACS sample, Bolton et al. 2006), with a smooth lens potential and a pixelated source regularized with wavelets (as in Galan et al. 2021) showed that convergence to the MAP takes about 1.5 min for a single gradient descent (still on a personal computer). This includes preoptimization steps with a smooth model for source whose complexity is progressively improved with a pixelated source. Fitting the lens light with analytical profiles would only lead to a marginal increase of the total run time (~ 10 s). Next, modeling deviations to the smooth lens potential assuming the initial model fits reasonably well the data would require about 30 s for a single gradient descent, similar to the models presented in this work. Finally, sampling the full parameter space using HMC would take from one to two hours for $\sim O(10^3)$ samples, which we expect to be sufficient to ensure well-sampled posterior distributions.

While the timings quoted above demonstrate that our code can be conveniently run on a single CPU, they do not reveal the full potential of the method. HERCULENS is fully based on JAX so it supports large scale parallelization and GPU capabilities, which can lead to dramatic improvements in terms of computation time (see e.g., Gu et al. 2022). This will be exploited in future analyses of real data sets.

7. Summary and conclusion

In this work we develop and apply a novel lens modeling method that is able to recover perturbations to a mostly smooth lens potential with minimal assumptions about their nature. This is made possible by modeling the perturbations on a grid of pixels regularized using a well-established multiscale technique based on sparsity and wavelets. This grid of pixels is superimposed on other analytical profiles for the joint modeling of the full potential and the source light. We show that merging the two main state-of-the-art modeling paradigms (analytical and pixelated) is

possible within the framework of differentiable programming. This enables us to seamlessly optimize lens models with over one thousand individual parameters and obtain their uncertainties, either via Fisher information analysis or gradient-informed HMC sampling.

We summarize the key results of this work as follows:

- We extend our previous work in Joseph et al. (2019) and Galan et al. (2021), by introducing a pixelated mass component in addition to the smooth lens potential, and recover three different types of perturbations: a localized DM subhalo, a population of such subhalos, and high-order moments in the lens potential.
- Sparse regularization is usually performed via iterative algorithms to converge to the solution, which makes it challenging to incorporate within standard lens modeling codes. In this work we demonstrate that the solution can also be obtained via gradient descent, capitalizing on the access to derivatives of the loss function with automatic differentiation.
- Differentiable programming enables the simultaneous optimization of analytical and pixelated components. One benefit is that the perturbative approach of Koopmans (2005) for pixelated perturbations to the lens potential is no longer warranted, although we can still use it to estimate regularization weights. This is because the full inverse problem can now be solved from the explicit superposition of smooth analytical and pixelated components, jointly optimized via gradient descent.
- For each type of perturbation explored in this work, the main features are correctly recovered by the pixelated potential component. In particular, the signature of a localized subhalo and octupolar structures are accurately captured. The subhalo population, simulated here as a GRF, clearly represents a more challenging situation, although the main over- and under-density regions can still be recovered.
- We test for a model-independent recovery of the DM subhalo mass, directly from the reconstructed pixelated potential model. We find that the resulting measurement of the mass is sensitive to the region where the surface density is integrated, leading to either an under-estimation of the mass, or a value in agreement with the true subhalo mass. Nevertheless, switching to an analytical profile for the subhalo, as is standard practice, allows to robustly infer its mass.
- The statistical properties of the subhalo population can be recovered via the power spectrum of the pixelated model. The underlying GRF parameters, which effectively act as parameters of the subhalo mass function, are recovered with our higher resolution model. Nevertheless, in a real-world scenario, we advocate for a comparison of different model variants, typically with different grid resolutions, and a cross-checking with other methods relying on more informative priors (e.g., as in Vernardos & Koopmans 2022) to improve the robustness of the inference.
- High-order moments in the lens potential (here as an octupole) are remarkably well recovered, despite being small in amplitude compared to lower-order moments such as quadrupoles (e.g., external shear). While the amplitude is biased low in the pixelated reconstruction, the octupole orientation is accurately measured, either from the pixelated model, or using a refined model including a multipole component in the lens potential explicitly.

Our method is readily applicable to real HST observations of EELs, such as the systems presented in Oldham & Auger (2018), as the source surface brightness is smooth. This assumption of

⁸ We note that most of the timings quoted here also include an overhead time of about 2 to 4 s, due to JAX “just-in-time” compilation feature.

smoothness, although well motivated by real observations, is arguably the strongest assumption of this work. Indeed, there are many situations in which the source galaxy is more complex, featuring for instance spiral arms and localized clumpy star forming regions. This requires the joint modeling of deviations from smoothness both in the source and the lens potential, which is challenging due to degeneracies between those two components, as some of the lensing features may be equally well explained by underlying features in the source or in the potential.

Nevertheless, the modeling methods presented in this work and implemented in HERCULENS enable the joint modeling of more complex sources on a grid of pixels in the source plane, which we will apply in future analyses. Moreover, the flexible framework provided by differentiable programming allows one to implement a large fraction (if not all) of the modeling methods explored in the literature so far. The careful comparison between their different assumptions is absolutely critical to mitigate degeneracies and robustly infer key physical quantities, including constraints on the subhalo mass function and the mass of DM particles. Recently, there have been a few works along these lines, including the reanalysis of a subhalo detection based on different assumptions such as the shape of the subhalo mass profile and its redshift (Çağan Şengül et al. 2022), or the thorough comparison of different modeling codes on the same data set (Shajib et al. 2022).

In this work we limit ourselves to three categories of perturbations in the lens potential, however there exist others. Another departure from the simple elliptical profiles is the presence of a disk component in the lens galaxy. In particular, such a disk structure can have observational effects that are similar to flux ratio anomalies in multiply imaged lensed quasars. These anomalies are usually considered as a tracer for DM subhalos, hence unique probes of the subhalo mass function. In a series of papers based on cosmological simulations and real observations, Hsueh et al. (2018, and references therein) showed that not taking into account these baryonic effects on the lens potential can bias the inferred constraints on DM properties. The detection and modeling of a disk component in the lens potential is therefore a situation where our pixelated reconstruction method could be successful.

HERCULENS is built on top of differentiable programming libraries that readily enable massive parallelization and GPU acceleration. This important aspect has been extensively discussed in Gu et al. (2022), where authors demonstrated the efficiency of gradient-informed methods for modeling large samples of strong lenses. Their results remain relevant and are applicable in the context of our work as well. This is a promising path towards the modeling of the tens of thousands strong lenses that will be discovered in the era of *Vera Rubin* and *Euclid* survey telescopes (Collett 2015). Moreover, highly flexible yet computationally efficient models will be essential to handle the even higher resolution data sets soon delivered by the *James Webb* Space Telescope and future extremely large telescopes.

Acknowledgements. We thank the anonymous referee for helpful comments that improved this manuscript. We thank Luca Biggio, Martin Millon, Eric Paic and Simona Vegetti for useful feedback and discussion on the present work. We thank Simon Birrer for useful discussion, and for making the modeling software package LENSTRONOMY open-source, which made the development of HERCULENS much easier. This programme is supported by the Swiss National Science Foundation (SNSF) and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (COSMICLENS: grant agreement No. 787886). G.V. has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement

No. 897124. This research has also made use of SCIPY (Virtanen et al. 2020), NUMPY (Oliphant 2006; Van Der Walt et al. 2011), MATPLOTLIB (Hunter 2007), ASTROPY (Astropy Collaboration 2013, 2018), JUPYTER (Kluyver et al. 2016) and GETDIST (Lewis 2019).

References

- Abramo, L. R. 2012, *MNRAS*, **420**, 2042
- Adam, A., Perreault-Levasseur, L., & Hezaveh, Y. 2022, ArXiv e-prints [arXiv:2207.01073]
- Albert, J. G. 2020, ApJ, submitted [arXiv:2012.15286]
- Astropy Collaboration (Robitaille, T. P., et al.) 2013, *A&A*, **558**, A33
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, **156**, 123
- Bayer, D., Chatterjee, S., Koopmans, L. V. E., et al. 2018, ArXiv e-prints [arXiv:1803.05952]
- Belgacem, E., Calcagni, G., Crisostomi, M., et al. 2019, *JCAP*, **2019**, 024
- Birrer, S., & Amara, A. 2018, *Phys. Dark Univ.*, **22**, 189
- Birrer, S., Amara, A., & Refregier, A. 2015, *ApJ*, **813**, 102
- Birrer, S., Amara, A., & Refregier, A. 2017, *JCAP*, **2017**, 037
- Birrer, S., Shajib, A., Gilman, D., et al. 2021, *J. Open Sour. Softw.*, **6**, 3283
- Blandford, R., Surpi, G., & Kundić, T. 2001, in *Gravitational Lensing: Recent Progress and Future Go*, eds. T. G. Brainerd, & C. S. Kochanek, *ASP Conf. Ser.*, **237**, 65
- Blumenthal, G. R., Faber, S. M., Flores, R., & Primack, J. R. 1986, *ApJ*, **301**, 27
- Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, *ApJ*, **638**, 703
- Boylan-Kolchin, M., Bullock, J. S., & Kaplinghat, M. 2011, *MNRAS*, **415**, L40
- Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, *JAX: Composable Transformations of Python+NumPy Programs*, <http://github.com/google/jax>
- Brehmer, J., Mishra-Sharma, S., Hermans, J., Louppe, G., & Cranmer, K. 2019, *ApJ*, **886**, 49
- Çağan Şengül, A., Dvorkin, C., Ostdieck, B., & Tsang, A. 2022, *MNRAS*, **515**, 4391
- Chatterjee, S. 2019, PhD Thesis, University of Groningen, The Netherlands
- Chatterjee, S., & Koopmans, L. V. E. 2018, *MNRAS*, **474**, 1762
- Chianese, M., Coogan, A., Hofma, P., Otten, S., & Weniger, C. 2020, *MNRAS*, **496**, 381
- Claeskens, J. F., Sluse, D., Riaud, P., & Surdej, J. 2006, *A&A*, **451**, 865
- Coles, J. P., Read, J. I., & Saha, P. 2014, *MNRAS*, **445**, 2181
- Collett, T. E. 2015, *ApJ*, **811**, 20
- Coogan, A., Karchev, K., & Weniger, C. 2020, ArXiv e-prints [arXiv:2010.07032]
- Cramér, H. 1999, *Mathematical Methods of Statistics, Princeton Landmarks in Mathematics and Physics* (Princeton: Princeton University Press)
- Cyr-Racine, F.-Y., Keeton, C. R., & Moustakas, L. A. 2019, *Phys. Rev. D*, **100**, 023013
- de Blok, W. J. G. 2010, *Adv. Astron.*, **2010**, 789293
- Despali, G., Vegetti, S., White, S. D. M., et al. 2022, *MNRAS*, **510**, 2480
- Diaz Rivero, A., & Dvorkin, C. 2020, *Phys. Rev. D*, **101**, 023515
- Duane, S., Kennedy, A., Pendleton, B. J., & Roweth, D. 1987, *Phys. Lett. B*, **195**, 216
- Dubinski, J. 1994, *ApJ*, **431**, 617
- Erkal, D., Belokurov, V., Bovy, J., & Sanders, J. L. 2016, *MNRAS*, **463**, 102
- Frigo, M., Naab, T., Hirschmann, M., et al. 2019, *MNRAS*, **489**, 2702
- Galan, A., Peel, A., Joseph, R., Courbin, F., & Starck, J. L. 2021, *A&A*, **647**, A176
- Gilman, D., Birrer, S., Treu, T., Nierenberg, A., & Benson, A. 2019, *MNRAS*, **487**, 5721
- Gu, A., Huang, X., Sheu, W., et al. 2022, *ApJ*, **935**, 49
- Hao, C. N., Mao, S., Deng, Z. G., Xia, X. Y., & Wu, H. 2006, *MNRAS*, **370**, 1339
- He, Q., Nightingale, J., Massey, R., et al. 2023, *MNRAS*, **518**, 220
- Hessel, M., Budden, D., Viola, F., et al. 2020, *Optax: Composable Gradient Transformation and Optimisation*, in *JAX*, <http://github.com/deepmind/optax>
- Hezaveh, Y., Dalal, N., Holder, G., et al. 2016a, *JCAP*, **2016**, 048
- Hezaveh, Y. D., Dalal, N., Marrone, D. P., et al. 2016b, *ApJ*, **823**, 37
- Hoffman, M. D., & Gelman, A. 2011, ArXiv e-prints [arXiv:1111.4246]
- Home, K. 1986, *PASP*, **98**, 609
- Hsueh, J.-W., Despali, G., Vegetti, S., et al. 2018, *MNRAS*, **475**, 2438
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Joseph, R., Courbin, F., Starck, J. L., & Birrer, S. 2019, *A&A*, **623**, A14
- Karchev, K., Coogan, A., & Weniger, C. 2022, *MNRAS*, **512**, 661
- Keeton, C. R. 2001, ArXiv e-prints [arXiv:astro-ph/0102341]
- Keeton, C. R., Falco, E. E., Impey, C. D., et al. 2000, *ApJ*, **542**, 74

- Kennedy, J., & Eberhart, R. C. 2001, *Swarm Intelligence* (San Francisco: Morgan Kaufmann Publishers Inc.)
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, eds. F. Loizides, & B. Schmidt (IOS Press), 87
- Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, *ApJ*, **522**, 82
- Koopmans, L. V. E. 2005, *MNRAS*, **363**, 1136
- Lanusse, F., Starck, J. L., Leonard, A., & Pires, S. 2016, *A&A*, **591**, A2
- Lee, W., Yu, H., Rival, X., & Yang, H. 2020, ArXiv e-prints [arXiv:2006.06903]
- Lewis, A. 2019, <https://getdist.readthedocs.io/en/latest/> [arXiv:1910.13970]
- Li, R., Frenk, C. S., Cole, S., Wang, Q., & Gao, L. 2017, *MNRAS*, **468**, 1426
- Moore, B. 1994, *Nature*, **370**, 629
- Moore, B., Ghigna, S., Governato, F., et al. 1999, *ApJ*, **524**, L19
- Nadler, E. O., Birrer, S., Gilman, D., et al. 2021, *ApJ*, **917**, 7
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, *ApJ*, **462**, 563
- Neal, R. 2011, *Handbook of Markov Chain Monte Carlo*, 113
- Nightingale, J. W., & Dye, S. 2015, *MNRAS*, **452**, 2940
- Nocedal, J., & Wright, S. J. 2006, *Numerical Optimization*, 2nd edn. (New York: Springer)
- Oldham, L. J., & Auger, M. W. 2018, *MNRAS*, **476**, 133
- Oldham, L., Auger, M. W., Fassnacht, C. D., et al. 2017, *MNRAS*, **465**, 3185
- Oliphant, T. E. 2006, *A Guide to NumPy* (USA: Trelgol Publishing), 1
- Ostdiek, B., Diaz Rivero, A., & Dvorkin, C. 2022, *A&A*, **657**, L14
- Papastergis, E., Giovanelli, R., Haynes, M. P., & Shankar, F. 2015, *A&A*, **574**, A113
- Paykari, P., Lanusse, F., Starck, J. L., Sureau, F., & Bobin, J. 2014, *A&A*, **566**, A77
- Philcox, O. H. E., Sherwin, B. D., Farren, G. S., & Baxter, E. J. 2021, *Phys. Rev. D*, **103**, 023538
- Planck Collaboration VI. 2020, *A&A*, **641**, A6
- Powell, D. M., Vegetti, S., McKean, J. P., et al. 2022, *MNRAS*, **516**, 1808
- Raichoor, A., de Mattia, A., Ross, A. J., et al. 2021, *MNRAS*, **500**, 3254
- Riess, A. G. 2020, *Nat. Rev. Phys.*, **2**, 10
- Saha, P., & Williams, L. L. R. 1997, *MNRAS*, **292**, 148
- Sérsic, J. L. 1963, *Boletín de la Asociación Argentina de Astronomía La Plata Argentina*, **6**, 41
- Shajib, A. J., Birrer, S., Treu, T., et al. 2019, *MNRAS*, **483**, 5649
- 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
- Shu, Y., Brownstein, J. R., Bolton, A. S., et al. 2017, *ApJ*, **851**, 48
- Sonnenfeld, A., Leauthaud, A., Auger, M. W., et al. 2018, *MNRAS*, **481**, 164
- Springel, V., Frenk, C. S., & White, S. D. M. 2006, *Nature*, **440**, 1137
- Starck, J.-L., Fadili, J., & Murtagh, F. 2007, *IEEE Trans. Image Process.*, **16**, 297
- Starck, J., Murtagh, F., & Fadili, J. 2015, *Sparse Image and Signal Processing: Wavelets and Related Geometric Multiscale Analysis* (Cambridge: Cambridge University Press)
- Suyu, S. H., Marshall, P. J., Hobson, M. P., & Blandford, R. D. 2006, *MNRAS*, **371**, 983
- Suyu, S. H., Marshall, P. J., Blandford, R. D., et al. 2009, *ApJ*, **691**, 277
- Tagore, A. S., & Keeton, C. R. 2014, *MNRAS*, **445**, 694
- Toomre, A., & Toomre, J. 1972, *ApJ*, **178**, 623
- Trotter, C. S., Winn, J. N., & Hewitt, J. N. 2000, *ApJ*, **535**, 671
- Van de Vyvere, L., Sluse, D., Gomer, M. R., & Mukherjee, S. 2022a, *A&A*, **663**, A179
- Van de Vyvere, L., Gomer, M. R., Sluse, D., et al. 2022b, *A&A*, **659**, A127
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, **13**, 22
- Varma, S., Fairbairn, M., Figueroa, J. 2020, *ApJ*, submitted [arXiv:2005.05353]
- Vegetti, S., & Koopmans, L. V. E. 2009, *MNRAS*, **392**, 945
- Vegetti, S., Koopmans, L. V. E., Bolton, A., Treu, T., & Gavazzi, R. 2010, *MNRAS*, **408**, 1969
- Vegetti, S., Lagattuta, D. J., McKean, J. P., et al. 2012, *Nature*, **481**, 341
- Vegetti, S., Koopmans, L. V. E., Auger, M. W., Treu, T., & Bolton, A. S. 2014, *MNRAS*, **442**, 2017
- Vernardos, G. 2022, *MNRAS*, **511**, 4417
- Vernardos, G., & Koopmans, L. V. E. 2022, *MNRAS*, **516**, 1347
- Vernardos, G., Tsagkatakis, G., & Pantazis, Y. 2020, *MNRAS*, **499**, 5641
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Vogelsberger, M., Marinacci, F., Torrey, P., & Puchwein, E. 2020, *Nat. Rev. Phys.*, **2**, 42
- Warren, S. J., & Dye, S. 2003, *ApJ*, **590**, 673
- Wengert, R. E. 1964, *Commun. ACM*, **7**, 463
- Zhuang, J., Tang, T., Ding, Y., et al. 2020, ArXiv e-prints [arXiv:2010.07468]
- Zubovas, K., & King, A. 2012, *ApJ*, **745**, L34

Appendix A: Multiscale regularization weights

In this section we give details about the estimation of the weights in Eq. 14 used to scale the regularization strength of the pixelated potential ψ_{pix} . Let us assume that we have in hand a model fitted to the data without ψ_{pix} , such that we can assume that a sufficiently good model of the source s and the smooth lens potential are known.

The minimization problem corresponding to the full lens potential, parametrized as $\psi = \tilde{\psi}(\eta_\psi) + \psi_{\text{pix}}$ can be written as

$$\arg \min_{\eta_\psi, \psi_{\text{pix}}} \frac{1}{2} \left[d - \mathbf{B}\mathbf{L}_\psi s \right]^\top \mathbf{C}_d^{-1} \left[d - \mathbf{B}\mathbf{L}_\psi s \right] + \lambda_\psi \left\| \mathbf{W}_\psi \circ \Phi^\top \psi_{\text{pix}} \right\|_1, \quad (\text{A.1})$$

where \mathbf{L}_ψ depends on the smooth potential parameters η_ψ and the pixelated component ψ_{pix} . The regularization term is identical to Eq. 14.

In contrast to the source reconstruction problem, namely where the variable of interest is s , the model we optimize here does not depend linearly on η_ψ and ψ_{pix} . Consequently, the gradient of the data-fidelity term in Eq. A.1 does not have a simple closed-form expression, which prevents us from propagating the noise from the data to the lens potential as easily as to the source plane.

We address this issue by considering the perturbative approach of Koopmans (2005) as a means to relate the data space to the potential space. We argue that we can use the weights obtained with this model for the minimization problem above.

Based on a Taylor expansion of the lens equation with respect to small deviations $\delta\psi$ to an underlying smooth potential $\tilde{\psi}(\eta_\psi)$, Koopmans (2005) established the following linear relation

$$\delta d = \underbrace{-\mathbf{B}\mathbf{D}_s(s)}_{\mathbf{D}} \mathbf{D}_{\delta\psi} \delta\psi, \quad (\text{A.2})$$

which corresponds to Eq. 15 rewritten with linear operators. The term $\delta d \equiv d - \tilde{m}$ represents residuals between the data and a model $\tilde{m} = \mathbf{B}\mathbf{L}_\psi s$ without perturbations to the smooth potential (i.e., $\delta\psi = 0$). The operator \mathbf{D}_s contains spatial derivatives of the source light model with respect to source plane coordinates mapped from the data grid via \mathbf{L}_ψ . The operator $\mathbf{D}_{\delta\psi}$ combines bilinear interpolation and finite difference coefficients to compute the spatial derivatives of $\delta\psi$ on the data grid. We refer the reader to Koopmans (2005) for more details about the exact structure of these operators, that we implement as matrices.

From the linear relation of Eq. A.2, we can formulate yet a new minimization problem for $\delta\psi$ as

$$\arg \min_{\delta\psi} \frac{1}{2} \left[\delta d - \mathbf{D}\delta\psi \right]^\top \mathbf{C}_{\delta d}^{-1} \left[\delta d - \mathbf{D}\delta\psi \right] + \lambda_{\delta\psi} \left\| \mathbf{W}_{\delta\psi} \circ \Phi^\top \delta\psi \right\|_1, \quad (\text{A.3})$$

where $\mathbf{C}_{\delta d}$ is the covariance matrix associated to δd . From the definition of model residuals, we have

$$\begin{aligned} \mathbf{C}_{\delta d} &\equiv \text{cov}(\delta d, \delta d) \\ &= \text{cov}(d - \tilde{m}, d - \tilde{m}) \\ &= \mathbf{C}_d + \mathbf{C}_{\tilde{m}} - 2 \text{cov}(d, \tilde{m}) \\ &\approx \mathbf{C}_d. \end{aligned} \quad (\text{A.4})$$

The cross-covariance matrix $\text{cov}(d, \tilde{m})$ is zero, since the data and model \tilde{m} are uncorrelated (nature does not correlate with our

model). Additionally, the covariance term $\mathbf{C}_{\tilde{m}}$ is much smaller than \mathbf{C}_d in the case of smooth analytical profiles for the potential and source (we checked this numerically). This is because the fully smooth model \tilde{m} is strongly limited by the shape of the analytical profiles described by a small number of degrees of freedom, constrained by a much larger number of data pixels. Although not explicitly specified for conciseness, the elements of \mathbf{C}_d depend on the model \tilde{m} to estimate the shot noise (Poisson noise) from the lensed features.

We now need to establish which operators are necessary to propagate the noise from δd space to wavelet coefficients of $\delta\psi$, in order to compute the coefficients of the matrix $\mathbf{W}_{\delta\psi}$. We do so by considering a gradient descent update based on the data-fidelity term of Eq. A.3 (e.g., Lanusse et al. 2016), which reads

$$\delta\psi^{(n+1)} = \delta\psi^{(n)} + \overbrace{\mathbf{D}^\top \mathbf{C}_{\delta d}^{-1} (\delta d - \mathbf{D}\delta\psi^{(n)})}^{\text{negative data-fidelity gradient}}. \quad (\text{A.5})$$

The above gradient translates model residuals (the term in parenthesis) into the coefficients of $\delta\psi$. Therefore, we see that the propagation of a given noise realization δd_N to the potential pixels is $\mathbf{D}^\top \mathbf{C}_{\delta d}^{-1} \delta d_N$. We simply apply the wavelet transform operator to obtain wavelet coefficients instead, which leads to $\Phi^\top \mathbf{D}^\top \mathbf{C}_{\delta d}^{-1} \delta d_N$. Using Monte-Carlo simulations of δd_N , we can compute the standard deviation $\sigma_{\delta\psi}^{(n,m,j)}$ for each potential pixel (n, m) in each wavelet scale j , and populate the matrix $\mathbf{W}_{\delta\psi}$ accordingly.

As a final step, we return to the original problem described by Eq. A.1. In this problem, the pixelated potential parameters are ψ_{pix} , instead of $\delta\psi$. First, let us assume that both ψ_{pix} and $\delta\psi$ are defined on the same grid of potential pixels. Second, ψ_{pix} is used in combination with a smooth lens potential component $\tilde{\psi}(\eta_\psi)$, such that it captures deviations to the smooth potential in a similar way as $\delta\psi$. Third, from Eq. A.4, we see that the noise in the data (characterized by \mathbf{C}_d) is a good approximation of the noise in the model residuals δd (characterized by $\mathbf{C}_{\delta d}$). Therefore, the standard deviation of the noise in potential space, at the location of potential pixels and in each wavelet scale, can be approximated by the values $\sigma_{\delta\psi}^{(i,j,k)}$ computed above, namely $\mathbf{W}_\psi \approx \mathbf{W}_{\delta\psi}$.

In summary, we have used the linear expression of Eq. A.2 as a way to set the weights of our regularization term (Eq. A.1) based on the noise levels from the data. In this work, we use two different wavelet transforms, namely the starlet and Battle-Lemarié wavelet transforms. Thus we compute the corresponding weights following the above procedure to obtain $\mathbf{W}_{\psi, \text{st}}$ and $\mathbf{W}_{\psi, \text{bl}}$ respectively (only the wavelet operator Φ^\top is different). We show in Fig. 2 the corresponding weights for each scale of the starlet transform.

Appendix B: Analytical profiles

The smooth models used throughout this work are based on analytical profiles, which we define in detail here. Elliptical profiles are described with coordinates (θ_1, θ_2) , obtained by rotating the original coordinates $\theta \equiv (\theta_x, \theta_y)$ along the major axis of the ellipse with position angle ϕ .

To describe the smooth component of the lens potential, we use the singular isothermal ellipsoid (SIE) profile. This profile is originally defined in surface mass density (convergence) but also has analytical formulae for the potential and deflection angles.

The SIE potential is given by (e.g., Keeton 2001)

$$\psi_{\text{SIE}}(\theta_1, \theta_2) = \frac{\theta_E \sqrt{q_m} \theta_1}{\sqrt{1 - q_m^2}} \arctan \left(\frac{\sqrt{1 - q_m^2} \theta_1}{\sqrt{q_m^2 \theta_1^2 + \theta_2^2}} \right) + \frac{\theta_E \sqrt{q_m} \theta_2}{\sqrt{1 - q_m^2}} \operatorname{arctanh} \left(\frac{\sqrt{1 - q_m^2} \theta_2}{\sqrt{q_m^2 \theta_1^2 + \theta_2^2}} \right), \quad (\text{B.1})$$

where θ_E is the Einstein radius, q_m is the axis ratio of the elliptical profile, and the corresponding position angle is ϕ_m . In practice, we do not optimize the axis ratio and position angle, but rather ellipticity components $\{e_{1,\psi}, e_{2,\psi}\}$ to prevent sampling issues with the periodicity of ϕ_m (particularly for small values of q_m). These components are defined as

$$\begin{cases} e_{1,\psi} &= \frac{1-q_m}{1+q_m} \cos(2\phi_m) \\ e_{2,\psi} &= \frac{1-q_m}{1+q_m} \sin(2\phi_m). \end{cases} \quad (\text{B.2})$$

The singular isothermal sphere (SIS) is a particular case of an SIE with no ellipticity ($q_m = 1$).

The influence of other galaxies in the vicinity of the main deflector—whose lensing effects remain linear with respect to the lensed source—is modeled as a uniform external shear. It has amplitude γ_{ext} and orientation ϕ_{ext} . The corresponding lensing potential in polar coordinates (r, ϕ) is

$$\psi_{\text{ext}}(r, \phi) = \frac{1}{2} \gamma_{\text{ext}} r^2 \cos(2\phi - 2\phi_{\text{ext}}). \quad (\text{B.3})$$

The corresponding shear ellipticity parameters that we used for optimization are defined as

$$\begin{cases} \gamma_{1,\text{ext}} &= \gamma_{\text{ext}} \cos(2\phi_{\text{ext}}) \\ \gamma_{2,\text{ext}} &= \gamma_{\text{ext}} \sin(2\phi_{\text{ext}}). \end{cases} \quad (\text{B.4})$$

The elliptical Sérsic profile, suitable for modeling smooth galaxy light distributions, is defined as (Sérsic 1963)

$$I_{\text{Sérsic}}(\theta_1, \theta_2) = I_{\text{eff}} \exp \left[-b_n \left(\frac{\sqrt{\theta_1^2 + \theta_2^2 / q_1^2}}{\theta_{\text{eff}}} \right)^{1/n_s} + b_n \right], \quad (\text{B.5})$$

where I_{eff} is the amplitude of the profile at the effective radius θ_{eff} , and n_s is the Sérsic index which defines the slope of the profile. The axis ratio is q_1 , and the corresponding position angle is ϕ_1 . Ellipticities are defined as for the potential (see Eq. B.4). The term b_n is not a free parameter and is computed such that θ_{eff} is always equal to the half-light radius of the profile.

Appendix C: Simulated HST observations

In Table C.1 we summarize all instrumental settings and specific model assumptions that were used to generate HST observations of EELs systems. We used the simulation software package MOLET (Vernardos 2022) to perform high-resolution ray-tracing and to add instrumental effects (instrumental noise and PSF convolutions).

Table C.1. Instrument settings and model assumptions used for generating HST/WFC3/F160W mock observations of EELs. The coordinates are oriented such as north is up and east is right. The orientation angle is east-of-north.

Observation	
Pixel size [arcsec]	0.08
Exposure time [s]	$4 \times 2400 = 9600$
Zero-point [mag]	25.9463
Sky brightness [mag/arcsec ²]	22
Read noise [e ⁻]	4
PSF (gaussian FWHM) [arcsec]	0.3
Noise	readout + shot noise
Deflector , $z_d = 0.3$	
Singular isothermal ellipsoid (SIE)	
Einstein radius, θ_E [arcsec]	1.6
Ellipticity, q_m	0.73
Orientation, ϕ_m	82.5
Position, (x, y) [arcsec]	(0, 0)
External shear	
Strength, γ_{ext}	0.032
Orientation, ϕ_{ext}	144.2
Sérsic (light)	
Ellipticity, q_l	0.73 (= q_m)
Orientation, ϕ_l	82.5 (= ϕ_m)
Half-light radius, θ_{eff}	2
Sérsic index, $n_{\text{Sérsic}}$	2
Position, (x, y) [arcsec]	(0, 0)
Magnitude	19
Perturbations to the potential	
LS: Localized subhalo (SIS)	
Einstein radius, $\theta_{\text{E,halo}}$ [arcsec]	0.07
Position, (x, y) [arcsec]	(1.9, -0.4)
PS: Population of subhalos (GRF)	
Variance, σ_{GRF}^2	10^{-3}
Slope, β_{GRF}	4
HM: High-order moments (octupole)	
Strength, a_4	0.06
Orientation, ϕ_4	82.5 = ϕ_m
Source , $z_s = 0.7$	
Sérsic	
Ellipticity, q_l	0.82
Orientation, ϕ_l	170.8
Half-light radius, θ_{eff} [arcsec]	0.8
Sérsic index, $n_{\text{Sérsic}}$	2
Position (x, y) [arcsec]	(0.4, 0.15)
Magnitude	21

Appendix D: Choice of pixel size for ψ_{pix}

The pixelated model ψ_{pix} used to capture deviations from the smooth lens potential requires a choice of pixel size. In HERCULENS, this can be set to any multiplicative factor (not necessarily an integer) of the data pixel size. In this work we use a pixel scale factor of 3 as our fiducial model, meaning the grid on which the ψ_{pix} component is defined has a pixel size of $3 \times 0''.08 = 0''.24$. This resolution is sufficient to accurately model the simulated data set as well as to characterize the reconstructed features in the lens potential.

The multiscale aspect of the regularization strategy detailed in Sect. 2.4 enables the independent treatment of different spatial scales in the lens potential. This avoids the need to choose a

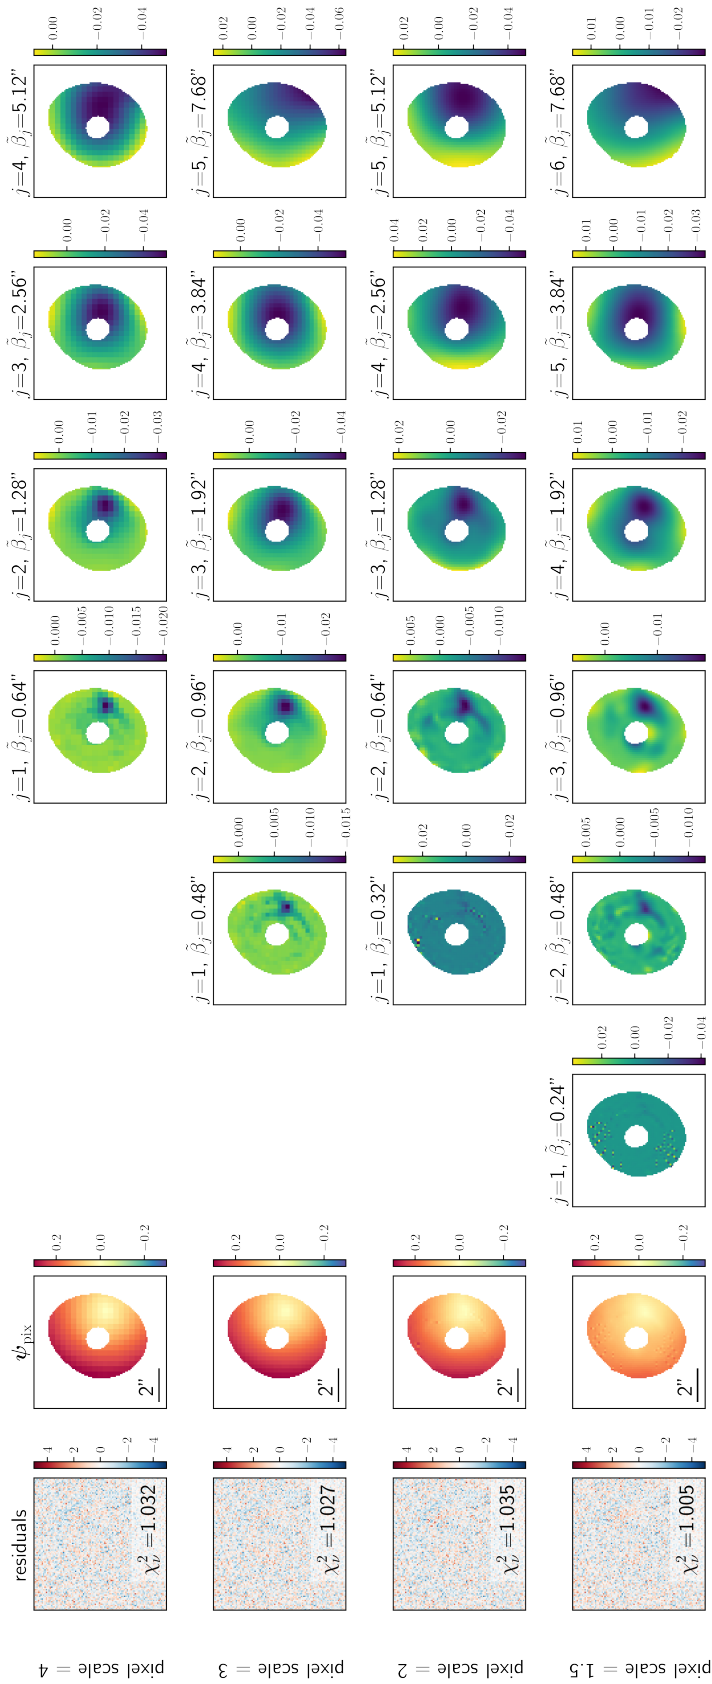


Fig. D.1. Best-fit models for different choices of pixel size of the ψ_{pix} component, corresponding to 4, 3, 2, and 1.5 times the data pixel size. The modeling assumptions are otherwise identical to the results presented in Fig. 5 (the second row is exactly the same). From left to right, the columns show: model residuals with reduced χ^2 value, best-fit ψ_{pix} model, followed by each scale of the starlet transform of ψ_{pix} . Each of these scales are images of the same size as ψ_{pix} , above which we indicate the scale index j and the characteristic scale β_j (in arcsec) of the features. We see that a pixel scale of 3 is sufficient to model the full dynamical range of features supported by the data. Finer pixels are not required, as the smaller starlet scales reveal mainly low-significance artifacts over the level of individual pixels.

specific pixel size, as long as it is small compared to the lensed features (otherwise it would lead to a poor fit to the data). Let us take the example of a pixel size much smaller than that of the data; if there are features below a given spatial scale that are not supported by the data (typically below the smallest detectable deflection angle), the multiscale wavelet decomposition combined with sparsity constraints will suppress all wavelet coefficients below this scale. Therefore, these features should not be visible in the reconstruction. This means that in comparing two models that fit the data equally well, we should select the one with the smaller pixel size. In practice, however, and from the perspective of parameter optimization, a smaller pixel size translates into a larger number of parameters, for which the regularization becomes even more crucial to balance the lack of constraints provided by the data alone. Moreover, the principle of Occam's razor, at the basis of Bayesian approaches, advocates for fewer parameters if it does not bring significant improvements to the fit.

We investigate the impact of the pixelated model resolution further by modeling lens LS (see Fig. 3) using four different pixel sizes for ψ_{pix} : 4, 3, 2 and 1.5 times the data pixel size (leading to 625, 1089, 2500 and 4489 parameters, respectively). For this simple exercise, we fix all other parameters to their input values. The resulting model residuals, best-fit ψ_{pix} solution, and the full starlet decomposition of the solution are shown in Fig. D.1. For each of the starlet scales j , we quote the characteristic scales β_j (in arcsec) captured by that scale, which is 2^j times the pixel size. Among the different versions of the model, the reduced χ^2 values are very similar, the smallest value being achieved by the highest resolution model⁹. An interesting aspect is visible on the finest starlet scales: for a pixel scale factor below 3, spurious isolated features appear in the solution, which is a consequence of poor regularization at small spatial scales. It is possible to address this issue by locally increasing the regularization strengths (or, alternatively, by modifying the regularization weights). Nevertheless, we find that the results of this work do not strongly depend on pixel size, and we leave such improvements for future versions of the method. Along these lines, we will investigate fully differentiable methods to optimize regularization strengths in the future.

Appendix E: Differentiation of the ℓ_1 -norm terms

The ℓ_1 -norm used to regularize s_{pix} and ψ_{pix} is not differentiable when its argument is strictly zero. However, Lee et al. (2020) validates the use of the many nondifferentiable functions used in machine learning (e.g., activation functions such as ReLU), by observing that they are part of a special class of functions that

allow their partial derivatives to be computed. Such functions are said to have piecewise analyticity under analytic partition.

The absolute value function at the basis of the ℓ_1 -norm belongs to this class of functions. Moreover, in practice, the function is never evaluated exactly at zero: the source or lens potential pixel values are never all exactly zero due to, for example, the presence of noise in the data. In addition, we never initialize pixel values to exactly zero on order to avoid undefined gradients at the start of the optimization.

Appendix F: Posterior distributions for analytical profile parameters

In Sect. 4.5, we discussed the MAP parameters of the smooth lens potential. Here we complete this discussion by showing in Fig. F.1 the full (smooth) parameter space, including the source parameters (Sérsic profile). For fully smooth models (upper triangle of the figure), the covariance matrix is estimated via the FIM (Sect. 2.6). For the perturbed models, HMC sampling was performed to obtain the full posterior distribution. For readability, we only plot ellipses corresponding to a multivariate normal distribution based on the parameter covariance matrices.

Overall, we observe two main features from these posterior distributions. First, fully smooth models of lenses LS and PS display strong biases on almost all parameters. This is expected as the model is insufficient to capture the complexity of the lens potential. Second, for many of the parameters, biases are reduced by including the pixelated potential in the model, although the reduction is mainly due to the larger error bars. As discussed in Sect. 4.5, these full models do not recover all smooth properties of the lens mass and source light distributions within uncertainties, despite the model fitting the data down to the noise and recovering the correct features in the pixelated potential component.

We interpret these biases as due to absorption of model residuals into the pixelated component ψ_{pix} , beyond the simple deviations to the smooth potential. This hypothesis is supported by our refined models, in which we replace the pixelated component by an analytical prescription, that show no remaining bias in smooth model parameters (see e.g., Fig. 7). We expect that changing to a regularization strategy that is motivated by the specific type of perturbations is expected to mitigate these biases. For instance, based on the assumption of GRF perturbations as considered in this work, we expect the pixelated reconstruction method presented in Vernardos & Koopmans (2022) to be less impacted by a slightly inaccurate absorption of model residuals. In a future work, we plan to explore in more detail the effect of different regularization methods on model parameters.

⁹ Seemingly small differences in χ^2 can still translate to potentially large differences in terms of Bayesian evidence. However, it is outside the scope of this work to use nested sampling to compute the Bayesian evidence, as it is too computationally expensive on a personal computer. Nevertheless, we note that nested sampling also benefits from a gradient-informed exploration of the parameter space (see e.g., Albert 2020).

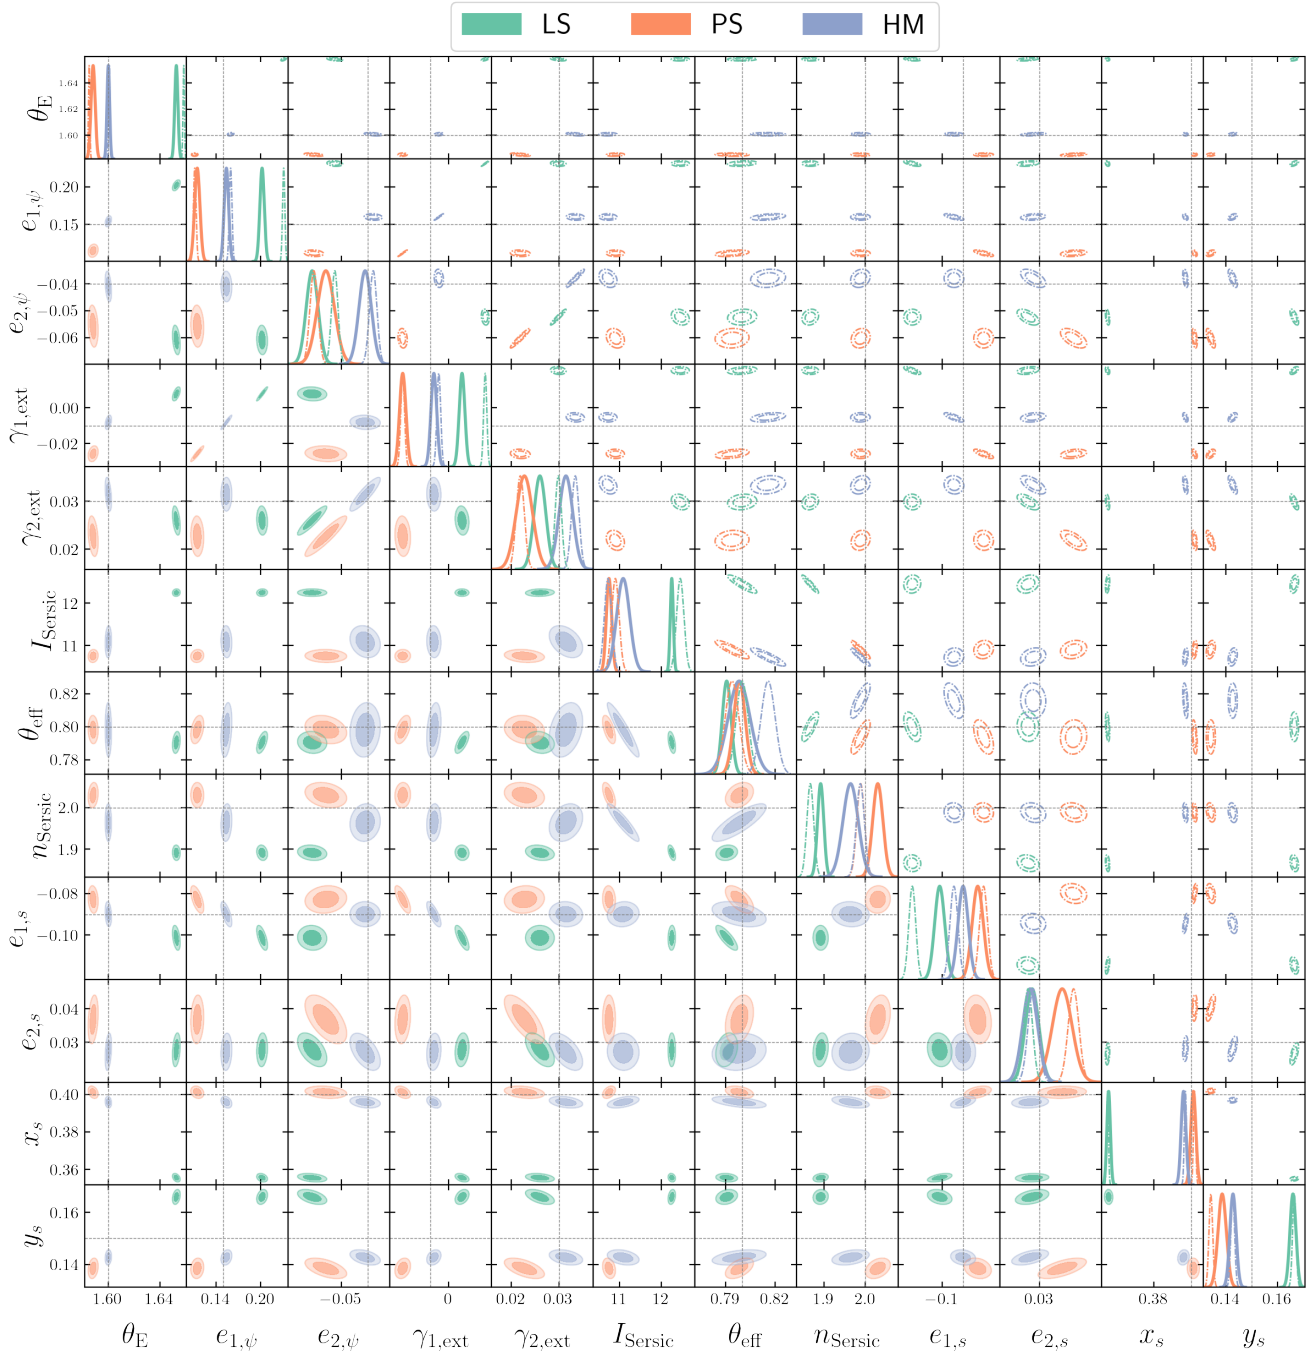


Fig. F.1. Posterior distributions for some analytical model parameters (smooth potential and smooth source), based on parameter covariance matrices computed either from the FIM or via HMC sampling (see Sect. 2.6). The three colors correspond to the three mocks of Fig. 3: localized subhalo (LS), population of subhalos (PS) and high-order moments (HM). Dot-dashed distributions in the upper right triangle correspond to fully smooth models, while distributions from lower left triangle correspond to models including a pixelated component to account for perturbations of the smooth potential. Parameters are, from left to right: Einstein radius of the lens θ_E , ellipticity of the smooth lens potential $\{e_{1,\psi}, e_{2,\psi}\}$, external shear components $\{\gamma_{1,\text{ext}}, \gamma_{2,\text{ext}}\}$, central intensity of the source I_{Sersic} , half-light radius r_{eff} , Sérsic index of the source n_{Sersic} , ellipticity of the source $\{e_{1,s}, e_{2,s}\}$, position of the source $\{x_s, y_s\}$. Overall, parameter biases are reduced after including a pixelated component in the potential to model deviations from smoothness; however, some still remain, in particular for lens LS, for which turning to an analytical model informed by the pixelated model is warranted.