

Euclid preparation

LIII. LensMC, weak lensing cosmic shear measurement with forward modelling and Markov Chain Monte Carlo sampling

Euclid Collaboration: G. Congedo^{1,*}, L. Miller², A. N. Taylor¹, N. Cross¹, C. A. J. Duncan^{3,2}, T. Kitching⁴, N. Martinet⁵, S. Matthew¹, T. Schrabback⁶, M. Tewes⁷, N. Welikala¹, N. Aghanim⁸, A. Amara⁹, S. Andreon¹⁰, N. Auricchio¹¹, M. Baldi^{12,11,13}, S. Bardelli¹¹, R. Bender^{14,15}, C. Bodendorf¹⁴, D. Bonino¹⁶, E. Branchini^{17,18,10}, M. Brescia^{19,20,21}, J. Brinchmann²², S. Camera^{23,24,16}, V. Capobianco¹⁶, C. Carbone²⁵, V. F. Cardone^{26,27}, J. Carretero^{28,29}, S. Casas³⁰, F. J. Castander^{31,32}, M. Castellano²⁶, S. Cavuoti^{20,21}, A. Cimatti³³, C. J. Conselice³, L. Conversi^{34,35}, Y. Copin³⁶, F. Courbin³⁷, H. M. Courtois³⁸, M. Cropper⁴, A. Da Silva^{39,40}, H. Degaudenzi⁴¹, A. M. Di Giorgio⁴², J. Dinis^{40,39}, F. Dubath⁴¹, X. Dupac³⁵, M. Farina⁴², S. Farrens⁴³, S. Ferriol³⁶, P. Fosalba^{32,44}, M. Frailis⁴⁵, E. Franceschi¹¹, S. Galeotta⁴⁵, B. Garilli²⁵, B. Gillis¹, C. Giocoli^{11,46}, A. Grazian⁴⁷, F. Grupp^{14,15}, S. V. H. Haugan⁴⁸, M. S. Holliman⁴⁹, W. Holmes⁵⁰, F. Hormuth⁵¹, A. Hornstrup^{52,53}, P. Hudelot⁵⁴, K. Jahnke⁵⁵, E. Keihänen⁵⁶, S. Kermiche⁵⁷, A. Kiessling⁵⁰, M. Kilbinger⁵⁸, B. Kubik³⁶, K. Kuijken⁵⁹, M. Kümmel¹⁵, M. Kunz⁶⁰, H. Kurki-Suonio^{61,62}, S. Ligorì¹⁶, P. B. Lilje⁴⁸, V. Lindholm^{61,62}, I. Lloro⁶³, D. Maino^{64,25,65}, E. Maiorano¹¹, O. Mansutti⁴⁵, O. Marggraf⁷, K. Markovic⁵⁰, F. Marulli^{66,11,13}, R. Massey⁶⁷, S. Maurogordato⁶⁸, H. J. McCracken⁵⁴, E. Medinaceli¹¹, S. Mei⁶⁹, M. Melchior⁷⁰, M. Meneghetti^{11,13}, E. Merlin²⁶, G. Meylan³⁷, M. Moresco^{66,11}, B. Morin⁵⁸, L. Moscardini^{66,11,13}, E. Munari⁴⁵, S.-M. Niemi⁷¹, J. W. Nightingale^{72,67}, C. Padilla²⁸, S. Paltani⁴¹, F. Pasian⁴⁵, K. Pedersen⁷³, W. J. Percival^{74,75,76}, V. Pettorino⁷⁷, S. Pires⁴³, G. Polenta⁷⁸, M. Poncet⁷⁹, L. A. Popa⁸⁰, L. Pozzetti¹¹, F. Raison¹⁴, R. Rebolo^{81,82}, A. Renzi^{83,84}, J. Rhodes⁵⁰, G. Riccio²⁰, E. Romelli⁴⁵, M. Roncarelli¹¹, E. Rossetti¹², R. Saglia^{15,14}, D. Sapone⁸⁵, B. Sartoris^{15,45}, P. Schneider⁷, A. Secroun⁵⁷, G. Seidel⁵⁵, S. Serrano^{32,31,86}, C. Sirignano^{83,84}, G. Sirri¹³, L. Stanco⁸⁴, P. Tallada-Crespi^{87,29}, D. Tavagnacco⁴⁵, I. Tereno^{39,88}, R. Toledo-Moreo⁸⁹, F. Torradeflot^{29,87}, I. Tutusaus⁹⁰, E. A. Valentijn⁹¹, L. Valenziano^{11,92}, T. Vassallo^{15,45}, A. Veropalumbo^{10,18}, Y. Wang⁹³, J. Weller^{15,14}, G. Zamorani¹¹, J. Zoubian⁵⁷, E. Zucca¹¹, A. Biviano^{45,94}, M. Bolzonella¹¹, A. Boucaud⁶⁹, E. Bozzo⁴¹, C. Burigana^{95,92}, C. Colodro-Conde⁸¹, D. Di Ferdinando¹³, J. Graciá-Carpio¹⁴, N. Mauri^{33,13}, C. Neissner^{28,29}, A. A. Nucita^{96,97,98}, Z. Sakr^{99,90,100}, V. Scottez^{101,102}, M. Tenti¹³, M. Viel^{94,45,103,104,105}, M. Wiesmann⁴⁸, Y. Akrami^{106,107}, V. Allevato²⁰, S. Anselmi^{83,84,108}, C. Baccigalupi^{103,45,104,94}, M. Ballardini^{109,110,11}, S. Borgani^{111,94,45,104}, A. S. Borlaff^{112,113,114}, S. Bruton¹¹⁵, R. Cabanac⁹⁰, A. Cappi^{11,68}, C. S. Carvalho⁸⁸, G. Castignani^{66,11}, T. Castro^{45,104,94,105}, G. Cañas-Herrera^{71,116}, K. C. Chambers¹¹⁷, A. R. Cooray¹¹⁸, J. Coupon⁴¹, S. Davini¹⁸, G. De Lucia⁴⁵, G. Desprez¹¹⁹, S. Di Domizio^{17,18}, H. Dole⁸, A. Díaz-Sánchez¹²⁰, J. A. Escartin Vigo¹⁴, S. Escoffier⁵⁷, I. Ferrero⁴⁸, F. Finelli^{11,92}, L. Gabarra^{83,84}, J. García-Bellido¹⁰⁶, E. Gaztanaga^{31,32,9}, F. Giacomini¹³, G. Gozalias^{121,61}, D. Guinet³⁶, A. Hall¹, H. Hildebrandt¹²², S. Ilić^{123,79,90}, A. Jimenez Muñoz¹²⁴, S. Joudaki^{9,74,75}, J. J. E. Kajava^{125,126}, V. Kansal^{127,128,129}, D. Karagiannis¹³⁰, C. C. Kirkpatrick⁵⁶, L. Legrand⁶⁰, J. Macias-Perez¹²⁴, G. Maggio⁴⁵, M. Magliocchetti⁴², R. Maoli^{131,26}, M. Martinelli^{26,27}, C. J. A. P. Martins^{132,22}, M. Maturi^{99,133}, L. Maurin⁸, R. B. Metcalf^{66,11}, M. Migliaccio^{134,135}, P. Monaco^{111,45,104,94}, G. Morgante¹¹, S. Nadathur⁹, L. Patrizzii¹³, A. Peel³⁷, A. Pezzotta¹⁴, V. Popa⁸⁰, C. Porciani⁷, D. Potter¹³⁶, M. Pöntinen⁶¹, P. Reimberg¹⁰¹, P.-F. Rocci⁸, A. G. Sánchez¹⁴, J. A. Schewtschenko¹, A. Schneider¹³⁶, E. Sefusatti^{45,104,94}, M. Sereno^{11,13}, P. Simon⁷, A. Spurio Mancini⁴, J. Stadel¹³⁶, J. Steinwagner¹⁴, G. Testera¹⁸, R. Teyssier¹³⁷, S. Toft^{53,138}, S. Tosi^{17,18,10}, A. Troja^{83,84}, M. Tucci⁴¹, C. Valieri¹³, J. Valiviita^{61,62}, and D. Vergani¹¹

(Affiliations can be found after the references)

Received 6 May 2024 / Accepted 17 September 2024

* Corresponding author; giuseppe.congedo@ed.ac.uk

ABSTRACT

LENSMC is a weak lensing shear measurement method developed for *Euclid* and Stage-IV surveys. It is based on forward modelling in order to deal with convolution by a point spread function (PSF) with comparable size to many galaxies, sampling the posterior distribution of galaxy parameters via Markov chain Monte Carlo, and marginalisation over nuisance parameters for each of the 1.5 billion galaxies observed by *Euclid*. We quantified the scientific performance through high-fidelity images based on the *Euclid* Flagship simulations and emulation of the *Euclid* VIS images, realistic clustering with a mean surface number density of 250 arcmin^{-2} ($I_E < 29.5$) for galaxies, and 6 arcmin^{-2} ($I_E < 26$) for stars, and a diffraction-limited chromatic PSF with a full width at half maximum of $0''.2$ and spatial variation across the field of view. LENSMC measured objects with a density of 90 arcmin^{-2} ($I_E < 26.5$) in 4500 deg^2 . The total shear bias was broken down into measurement (our main focus here) and selection effects (which will be addressed in future work). We found measurement multiplicative and additive biases of $m_1 = (-3.6 \pm 0.2) \times 10^{-3}$, $m_2 = (-4.3 \pm 0.2) \times 10^{-3}$, $c_1 = (-1.78 \pm 0.03) \times 10^{-4}$, and $c_2 = (0.09 \pm 0.03) \times 10^{-4}$; a large detection bias with a multiplicative component of 1.2×10^{-2} and an additive component of -3×10^{-4} ; and a measurement PSF leakage of $\alpha_1 = (-9 \pm 3) \times 10^{-4}$ and $\alpha_2 = (2 \pm 3) \times 10^{-4}$. When model bias is suppressed, the obtained measurement biases are close to *Euclid* requirement and largely dominated by undetected faint galaxies (-5×10^{-3}). Although significant, model bias will be straightforward to calibrate given its weak sensitivity on galaxy morphology parameters. LENSMC is publicly available at gitlab.com/gcongedo/LENSMC.

Key words. gravitational lensing: weak – methods: data analysis – cosmology: observations

1. Introduction

Weak gravitational lensing by large-scale structure is a mature cosmological tool to measure the distribution of dark matter and study dark energy through its evolution with redshift (Schneider 2006; Kilbinger 2015; Mandelbaum 2018). Weak lensing is particularly sensitive to modifications of the theory of gravity and the emergence of physics beyond the concordance Λ -cold dark matter model, which affect the clustering of dark matter (Amendola et al. 2018).

Galaxy surveys from the ground, such as the Dark Energy Survey (DES; Abbott et al. 2018b), the Kilo Degree Survey (KiDS; Kuijken et al. 2015), and the Hyper Suprime-Cam survey (HSC; Aihara et al. 2017), are now achieving constraints on the dark matter sector (primarily in the Ω_m - σ_8 parameter space and their combination S_8) to a few percent (Abbott et al. 2018a; Hildebrandt et al. 2017; Hikage et al. 2019; Amon et al. 2022a; Secco et al. 2022; Dalal et al. 2023; Li et al. 2023d). After extensive consistency checks and sensitivity studies, recent lensing measurements from galaxy surveys are shown to be broadly in agreement with each other but in mild tension with the *Planck* satellite at the 2σ level or greater (Planck Collaboration VI 2020; Joudaki et al. 2020; Amon et al. 2022b; Heymans et al. 2021; Asgari et al. 2021; Loureiro et al. 2022) – this includes the latest joint analysis of DES and KiDS (Abbott et al. 2023).

In the coming years, galaxy surveys will enter a new regime of area, depth, and image quality. The space-based *Euclid* telescope (survey area of $14\,000 \text{ deg}^2$, full width at half maximum (FWHM) resolution of $0''.2$, depth of 24.5, $I_E + Y_E + J_E + H_E$ filters, see Laureijs et al. 2011; Cropper et al. 2016; Euclid Collaboration 2022b, 2024d); the planned space-based *Roman* telescope (1700 deg^2 , $0''.2$, 26.5, $YJH+F184$, see Spergel et al. 2015; Akeson et al. 2019); and the ground-based *Rubin* Observatory Legacy Survey of Space and Time (LSST; $18\,000 \text{ deg}^2$, $0''.7$, 27.5, $ugrizy$, see Ivezić et al. 2019) will substantially increase the number of observable galaxies compared to current surveys. The systematic observation of a billion galaxies or more across one-third of the visible sky will then be possible for the first time. The combined effect of improved survey area and angular resolution will be an enhanced ability to probe both the large and small scales via weak lensing and galaxy clustering, allowing us to constrain cosmological models and dark energy to percent-level precision (Mandelbaum et al. 2018; Euclid Collaboration 2020a), or even an order of magnitude better when combined with data from *Planck* (Euclid Collaboration 2022a).

With the dramatic improvement in precision that will be achieved in the coming years, experiments are now focussing on understanding the accuracy of their analyses (see, e.g., Euclid Collaboration 2020b). Along with theory uncertainties, the cosmic shear measurement and redshift estimation are the most challenging aspects of any large-scale weak lensing surveys. The concern of this paper is on cosmic shear measurement, which provides the necessary data for the weak lensing cosmological analysis (Kilbinger 2015). In order to achieve an order of one-percent precision on the dark energy equation of state, a billion galaxies or more with median redshift around one need to be observed. This observation has to be carried out consistently so the same shape measurement procedure is applied to all objects. This measurement has to be conducted with outstanding accuracy to satisfy the stringent requirement of 2×10^{-3} and 3×10^{-4} on the measured multiplicative and additive shear biases that were set in the early development phase of weak lensing space telescopes (Massey et al. 2012; Cropper et al. 2013).

Throughout the past years, a number of shear measurement methods were developed, tested on data challenges, and applied to real data. These can be categorised into two main classes: non-parametric and parametric. Among the non-parametric is Kaiser-Squires-Broadhurst (KSB; Kaiser et al. 1995; Hoekstra et al. 1998), which is based on weighted moments of image data. Because of its simplicity, these estimators were used for the very first attempts at measuring cosmic shear in the early 2000s. These methods are fast and can be quickly calibrated, but they are sensitive to effects that need to be characterised. With better precision, it later became clear that more effects needed to be factored in, particularly a realistic PSF and the sensitivity of bias on the PSF ellipticity, known as leakage. Parametric methods, based on forward modelling and model fitting, soon appeared to be better suited to accurately incorporating such real data features building on solid statistical grounds. In recent years, a more systematic use of model fitting techniques was observed across all major lensing surveys. The Bayesian-inspired shape method *lensfit* (Miller et al. 2007; Kitching et al. 2008; Miller et al. 2013) was extremely successful, first in the Canada-France Hawaii Telescope Lensing Survey (CFHTLenS; Heymans et al. 2012) and more recently in KiDS. This method is based on forward modelling and marginalisation over galaxy nuisance parameters. A similar method is IM3SHAPE (Zuntz et al. 2013), which is a maximum likelihood estimator based again on analytic forward modelling that was applied to DES. While many real-data effects, including the PSF, are accounted for and can be directly built in parametric methods, any in-built correction clearly

introduces extra computational overhead, as the parameter probability distribution needs to be sampled accurately.

While lensing measurements became more precise over time, accuracy also needed to be examined more carefully. Methods were compared in data challenges and were run on common simulations with increasing realism, such as in the Shear TEsting Programme (STEP; Heymans et al. 2006; Massey et al. 2007) and the Gravitational LEnsing Accuracy Testing (GREAT; Bridle et al. 2010; Kitching et al. 2012; Mandelbaum et al. 2015). With no method outshining in absolute terms, and methods being better at some aspects of the measurement but worse at others, it became evident that some form of calibration was still necessary. More than ever before, the field has become reliant on galaxy simulations. Sophisticated high-fidelity simulations now need to reproduce the realism of actual observations as close as possible so all biases from detection, measurement, and selection can be fully captured (Fenech Conti et al. 2017; Kannawadi et al. 2019; Euclid Collaboration 2019; MacCrann et al. 2021; Li et al. 2023a,b). Calibration naturally raises the question about how sensitive results are to the assumptions that are made in simulations (Hoekstra et al. 2017), or how large these simulations need to be to meet the desired precision (Pujol et al. 2019; Jansen et al. 2024). Other methods now rely on some form of calibration that is directly built in the measurement process. Galaxy images used in the calibration were simulated internally, inferred from real data, or obtained through a combination of the two methods. Metacalibration (Sheldon & Huff 2017; Huff & Mandelbaum 2017) derives internal estimates of the sensitivity of the ellipticity estimator to input shears and was extremely successful on DES Year 3 (Gatti et al. 2021); Bayesian Fourier Domain (BFD; Bernstein & Armstrong 2014; Bernstein et al. 2016) estimates the Taylor coefficients of the galaxy likelihood expanded over shear with information about moments measured from calibration fields; a similar implementation to BFD uses forward modelling (Sheldon 2014); the KiDS self-calibration (Fenech Conti et al. 2017) derives internal estimates of the ellipticity bias from noise-free galaxy images; MomentsML relies on simulated images to train shear-predicting artificial neural networks (Tewes et al. 2019). Because many selection biases happen before the shear measurement introduces its own bias, the field has gradually become more aware that those will probably be the limiting factor in future lensing surveys. Further work around Metacalibration led to Metadetection to address the issue (Sheldon et al. 2020). An application of the method to *Euclid*-like simulations showed that while selection biases may exceed requirements, the outlook is still positive, with demonstrated success at handling detection and blending biases (Hoekstra 2021; Hoekstra et al. 2021; Melchior et al. 2021). Furthermore, the recent Anacal method aims to correct measurement and selection biases via analytic differentiation (Li et al. 2023c).

The impact of neighbours to lensing measurements has also become one of the most important issues that current and future surveys will need to address. In space, the large number density of detected galaxies of about 30 arcmin^{-2} ($I_E < 24.5$) is compensated by a good image resolution, so the impact of neighbours may not be as bad as from the ground. In fact, due to the worse resolution on the ground, the impact of neighbours is serious, affecting 60% of the sample (Bosch et al. 2017; Arcelin et al. 2021). DES Year 1 results (Samuroff et al. 2017) showed that the total neighbour bias can be as large as 9% (reaching 80% at a close distance to the neighbour, if uncorrected). Moreover, the impact of ‘unrecognised’ blends (undetected neighbours) on the S_8 parameter from simulated *Rubin* Year 1 data can be as large as 15% (Nourbakhsh et al. 2022). Cuts to the catalogue to remove

those objects can reduce the total bias to 1% (reaching 30% at a close distance to the neighbour), however at the cost of reducing the effective number density by 30% and leaving a residual bias on S_8 of 2σ . While Metacalibration is extremely successful in a few idealised cases (e.g., isolated galaxies) and Metadetection in the handling of blending and detection bias, a suite of advanced simulations were required for the tomographic calibration of DES Year 3 (MacCrann et al. 2021). These simulations show that an external calibration is still required, as the unresolved neighbour introduces a correlation between the two galaxies at different redshift, plus the Metacalibration shear responses will be biased by the presence of the neighbour itself. Therefore, calibration simulations are necessary to correctly capture neighbour bias and the interplay between shear and redshift. However, these simulations assume uniform random distributions of galaxies that were re-weighted to mimic clustering, which raises the question as to whether the inferred bias is likely underestimated. The most recent simulations by KiDS have realistic clustering of N -body simulations, mimic a number of measurement effects, and address the shear-redshift interplay (Li et al. 2023a). With a larger number density, it is expected that the situation may be more serious in future surveys.

In this paper, we present our advanced shear measurement method LENSEMC specifically developed for *Euclid* that builds on the knowledge and success of ground-based measurement at handling real data effects. Similarly to *lensfit*, it adopts a mean estimator. Contrary to *lensfit*, it does not marginalise over nuisance parameters with numerical approximations. With full flexibility in the choice of the prior, all the marginalisation in LENSEMC is performed by Markov chain Monte Carlo (MCMC) sampling for individual galaxies or jointly across groups of neighbouring galaxies. While IM3SHAPE returns the maximum of the likelihood estimated via the Levenberg-Marquardt algorithm, LENSEMC employs a combination of large-scale and small-scale algorithms (such as conjugate gradient and simplex methods) to estimate a suitable initial guess for the MCMC sampling, thus requiring no information about the model derivatives and dramatically reducing the sensitivity on the initial guess (which is assumed to always be the same for all galaxies). The galaxy models in LENSEMC are rendered directly in the Fourier space; hence only a single Fourier transform is required. A recent profile-fitting method, The Farmer, has also drawn attention (Weaver et al. 2023). This method is a maximum likelihood estimator whose initial guesses of position and shape are provided by the detection method. It includes a decision tree based on χ^2 values to classify objects on their likely type and provides error estimates via Cramer-Rao bounds. Preliminary results are encouraging; however, the method has not been tested on full space-based cosmic shear accuracy yet.

Accurate cosmic shear measurement requires controlling the bias from a number of sources. Key examples include source clustering, faint objects, neighbours, PSF leakage, astrometry, image artefacts, and cosmic rays. Additionally, any forward-modelling method is plagued by potential model bias. One of the main sources of model bias was addressed in this work. In summary, LENSEMC employs: (i) forward modelling to deal with *Euclid* image undersampling and convolution by a PSF with comparable size to many galaxies; (ii) joint measuring object groups to correctly handle bias due to neighbours; (iii) masking out objects belonging to different groups; (iv) MCMC sampling to sample the posterior in a multi-dimensional parameter space, calculate weights, and correctly marginalise ellipticity over nuisance parameters and other objects in the same group. We particularly focussed on the realism of our simulations,

including clustering, stars, object detection, the handling of neighbours due to high number density, and the use of realistic undersampled chromatic PSF images with spatial variation across the field of view. We did not include further real data effects such as non-linearities or cosmic rays, as these will be addressed separately. Also we assumed the same broadband PSF as obtained from a spectral energy distribution (SED) of an SBc-type galaxy at a redshift of one in both simulations and measurements.

Sect. 2 introduces our method and the practical advantages in addressing real data problems. Our forward-modelling method is sufficiently fast to analyse the typical data volume of Stage-IV surveys and can be applied to the complexity of *Euclid* measurements, including undersampled data and a complex PSF while accounting for the full degrees of freedom in the galaxy modelling. Additionally, it allows for the proper handling of resolved neighbours by joint measurement and masking of more distant galaxies, stars, and artefacts in the images. Sect. 3 describes the simulations used for our intensive testing of the method. The images are fiducial realisations of the *Euclid* VIS detector (Euclid Collaboration 2024b), and galaxies were rendered based on N -body simulations with full variability of the morphological properties (Potter et al. 2017; Euclid Collaboration 2024a). All galaxies were convolved with a realistic pre-flight PSF model with full spatial variation, but the chromatic variability was ignored. Sect. 4 presents the main results of this testing. When model bias, chromaticity, and selection biases are suppressed, the obtained biases are close to the *Euclid* requirement. This measurement bias is largely dominated by undetected faint galaxies in the images. The bias was also found to be stable and mostly insensitive to the many effects in the simulations. As the *Euclid* analysis will also need to correct for other artefacts in the images, the residual bias will be straightforward to calibrate through image simulations. Once we included the model bias by allowing the full variability in the galaxy models, the overall bias became significant. However, since the sensitivity is weak (the derivative of the bias with respect to the assumptions made in the simulations appears negligible), it will then be straightforward to also calibrate the model bias through image simulations. Sect. 5 discusses the main findings and draws the conclusions of our work.

2. Method

The main physical quantity of interest in weak lensing is the reduced cosmic shear (Kilbinger 2015),

$$g = \frac{\gamma}{1 - \kappa}, \quad (1)$$

where κ and γ are convergence and shear (both related to the gravitational potential), and $g \approx \gamma$ in the weak lensing regime. The related observable in weak lensing is the ellipticity of a galaxy,

$$\epsilon = \frac{a - b}{a + b} e^{2i\varphi}, \quad (2)$$

where a and b are, respectively, the semi-major and semi-minor axes,¹ $\varphi \in [0, \pi)$ is the orientation angle of the galaxy, and $|\epsilon| \leq 1$. The effect of weak lensing is to distort the ellipticity

¹ This holds true only for elliptical isophotes, but the ellipticity remains well-defined if one specifies how it is measured, that is, it becomes method dependent.

of a source galaxy, ϵ_s , by the canonical transformation (Seitz & Schneider 1997),

$$\epsilon = \frac{\epsilon_s + g}{1 + \epsilon_s g^*}, \quad (3)$$

where all spin-2 quantities are expressed in complex notation (e.g., $\epsilon = \epsilon_1 + i\epsilon_2$, where ϵ_1 quantifies the distortion along x and y , and ϵ_2 along the coordinate axes rotated by $\pi/4$). As it is customary in weak lensing, we will refer to ϵ_s as the intrinsic ellipticity of the source galaxy, and ϵ as the lensed or observed ellipticity. The ellipticity in Eq. (3) is a point estimate for shear in that information on the underlying cosmic shear can be derived in a statistical sense as a sample average, $g = \langle \epsilon \rangle$, which holds whenever the distribution of orientation angles is uniform, for example, when there are no astrophysical intrinsic alignments (Joachimi et al. 2015) or shear dependent selection effects.

In weak lensing measurements we infer the reduced shear through sample averages. In this work, we use the ellipticity as a point estimator for shear and the problem of measuring ellipticity can be formulated fully in Bayesian terms. Suppose we have a pixel data vector, \mathbf{D} , and a model for the galaxy brightness distribution, $\mathbf{I} = \mathbf{I}(\epsilon, \theta, \phi)$, as a function of ellipticity, ϵ , intrinsic nuisance parameters, θ , and linear nuisance parameters, ϕ .² Thanks to Bayes' theorem, we can define a joint posterior as follows:

$$p(\epsilon, \theta, \phi | \mathbf{D}) = \frac{p(\mathbf{D} | \epsilon, \theta, \phi) p(\epsilon, \theta, \phi)}{p(\mathbf{D})}, \quad (4)$$

where $p(\mathbf{D} | \epsilon, \theta, \phi)$ is the likelihood, $p(\epsilon, \theta, \phi)$ is the prior on ellipticity and nuisance parameters, and $p(\mathbf{D})$ is the evidence or marginal likelihood,

$$p(\mathbf{D}) = \int p(\mathbf{D} | \epsilon, \theta, \phi) p(\epsilon, \theta, \phi) d\epsilon d\theta d\phi. \quad (5)$$

We can then construct the ellipticity marginal posterior:

$$p(\epsilon | \mathbf{D}) = \frac{1}{p(\mathbf{D})} \int p(\mathbf{D} | \epsilon, \theta, \phi) p(\epsilon, \theta, \phi) d\theta d\phi, \quad (6)$$

marginalising out the nuisance parameters. Common choices of estimators are the maximum likelihood or maximum posterior probability, but these are usually biased if the distribution is not Gaussian. However, the bias can be predicted in simple cases of low dimensionality or when the probability function is fully analytic (Cox & Snell 1968; Hall & Taylor 2017). Another option that was successful in ground-based surveys (Miller et al. 2013) is to set the estimator to the mean of the posterior distribution,

$$\hat{\epsilon} = \int \epsilon p(\epsilon | \mathbf{D}) d\epsilon. \quad (7)$$

We adopt this definition, as it has some useful properties:

1. as the nuisance parameters are marginalised out, their impact on the ellipticity estimator via their correlation is mitigated;

² The parameter vectors θ and ϕ are not to be confused with angular coordinates. Here θ represents non-linear parameters (such as object size and position offsets) and ϕ represents linear parameters (such as component fluxes) that can be analytically marginalised out. The modelling is linear in the ϕ parameters as long as the flux components are co-centred and do not depend on the positions.

2. overfitting³ is inherently reduced as we pick an average representative of all possible realisations that are statistically equivalent;
3. following on from the previous point, we expect the mean estimator to be, in general, less biased than the maximum estimators;
4. the mean of the distribution can be estimated through MCMC sampling techniques (see Sect. 2.3); such estimators satisfy the central limit theorem and therefore converge to the true mean.

We will discuss more about those points later in this section. Whatever choice is made, any estimator can be seen as a non-linear mapping between \mathbf{D} and ϵ . Therefore even if \mathbf{D} were to be Gaussian distributed, the estimator will not, hence leading to a fundamental bias in the measurement, which is commonly referred to as noise bias (Melchior & Viola 2012; Refregier et al. 2012; Viola et al. 2014). Moreover, as the shear is estimated through a sample average over a population of galaxies with varying morphological properties and complex selections, the properties of the shear bias will be different from that of galaxy ellipticity. Assuming shear is small, it is customary in the field to model the shear bias on each component with a linear model (Guzik & Bernstein 2005; Huterer et al. 2006; Heymans et al. 2006),

$$\hat{g}_i = (1 + m_i)g_i + c_i + n_i, \quad (8)$$

where m_i and c_i are the multiplicative and additive biases for the i -th component, g_i is the true shear, \hat{g}_i is an estimate of it, and n_i is the corresponding statistical noise. The transformation in Eq. (8) should in principle have m_i replaced by a matrix m_{ij} to model any potential cross-talk between shear components. Alternatively, it could be rewritten as a spin-2 equation (Kitching & Deshpande 2022),

$$\hat{g} = (1 + m_0)g + m_4 g^* + c + n. \quad (9)$$

However, generalising upon previous work, m_0 and m_4 are now spin 0 and 4 complex numbers acting on spin-2 shear fields. Physically, this added flexibility allows for complete mode-mixing: m_0 models a dilation and rotation of the true shear, whereas m_4 models a reflection around the axis determined by its phase. We defer the application of this approach to future work. Multiplicative terms can be induced by, for example, non-Gaussianities in the posterior (skewness at first order), caused by, for example, pixel noise and a small galaxy size relative to the PSF. Additive terms are due to anisotropies induced by, for example, the PSF and its spatial variability across the field of view. This effect is referred to as PSF leakage and is defined by the dependence of c_i on the PSF ellipticity $\epsilon_{\text{PSF},i}$ (see, e.g. Gatti et al. 2021, and references therein),

$$\alpha_i = \frac{dc_i}{d\epsilon_{\text{PSF},i}}. \quad (10)$$

We focussed primarily on the c dependence on ϵ_{PSF} , as the m dependence is negligible as long as the PSF is stable (i.e., the variation in PSF size is within a percent level). Earlier studies (Massey et al. 2012; Cropper et al. 2013) set out requirements for space-based missions on m and c based on a top-down error breakdown from cosmology to two-point statistics. For *Euclid*,

³ This is the tendency of some estimators, in particular the maximum of the probability distribution function, to interpret random fluctuations in the noise as actual signal in the data.

the requirement is on the statistical error on bias, $\sigma_m < 2 \times 10^{-3}$ and $\sigma_c < 3 \times 10^{-4}$. That is roughly equivalent to saying that a shear of 1% should be measured with a fractional accuracy and precision of 0.2%. We note that the requirement is an order of magnitude more stringent than current ground-based experiments (Hildebrandt et al. 2017). The detailed breakdown of the total budget on m and c into various error terms (Cropper et al. 2013) suggests we can set the required statistical error on the bias due to the measurement alone to $\sigma_m < 5 \times 10^{-4}$ and $\sigma_c < 5 \times 10^{-5}$. Therefore, in order to measure $|g| \approx 0.03$ with a residual post-calibration multiplicative bias smaller than σ_m , one will need at least $N = \sigma_\epsilon^2 |g|^{-2} \sigma_m^{-2} \approx 4 \times 10^8$ galaxies⁴, where $\sigma_\epsilon \approx 0.3$ is the ‘shape noise’, that is, the standard deviation of the per-component intrinsic ellipticity distribution. Obviously measurement noise and intrinsic scatter in the morphological properties will also need to be factored in. A ballpark estimate for the *Euclid* requirement on PSF leakage that we will be using as benchmark in our analysis is $\sigma_\alpha \lesssim \sigma_c / |\delta\epsilon_{\text{PSF}}|$, where $|\delta\epsilon_{\text{PSF}}| \approx 0.1$ is the order of magnitude (absolute) variation in PSF ellipticity across the field of view, which yields $\sigma_\alpha < 1 \times 10^{-3}$ if we assume a budget of $\sigma_c < 1 \times 10^{-4}$. This derivation may be too conservative, as a full propagation of the errors and biases through to cosmological parameters is demonstrated to be able to capture the spatial pattern imprinted by the PSF and other effects (Euclid Collaboration 2020b). Other surveys implemented other solutions such a first-order expansion on PSF ellipticity and PSF model residuals in KiDS (Heymans et al. 2006; Giblin et al. 2021) or the angular correlations between PSF ellipticity and size implemented in the rho statistics in DES (Jarvis et al. 2020).

In the next subsections we address the key elements of the LENSMC measurement method: galaxy modelling, PSF convolution, likelihood, sampling, and a further discussion about handling real data effects. We emphasise the role of joint measurement of objects to address neighbour bias, which is a concern for current and upcoming surveys, and also our MCMC strategy to sample a multi-dimensional parameter space and marginalise each lensing target over all nuisance parameters and other objects.

2.1. PSF-convolved galaxy models

We assumed 2D-projected galaxy models as a mixture of two circular Sérsic profiles (Sérsic 1963). The disc component is

$$I_d(r) \propto \exp\left(-\frac{r}{r_e}\right), \quad (11)$$

and the bulge component is

$$I_b(r) \propto \exp\left[-a_b \left(\frac{r}{r_h}\right)^{\frac{1}{n_b}}\right], \quad (12)$$

where r is the distance from the centre, r_e is the exponential scale length of the disc, r_h is the bulge half-light radius, $n_b = 1$, and $a_b \approx 2 n_b - 0.331$ (Peng et al. 2002). The bulge Sérsic index was fixed to 1 based on recent multi-wavelength observations of the *Hubble* CANDELS/GOODS-South field (Welikala et al., in prep.)⁵, while bulge profiles with $n_b = 4$ (de Vaucouleurs) are

⁴ This assumes we measure shear with accuracy given by $\hat{g} = (1 + m)g$. The standard deviation of the measured shear scales as $\sigma_\epsilon / \sqrt{N}$ and therefore we required that $\sigma_\epsilon / \sqrt{N} \lesssim \sigma_m |g|$.

⁵ Their work highlights that both the dust in the disc and the 3D modelling of the galaxy influence the inferred bulge structural parameters.

instead more typical for early-type galaxies at low redshift. Both profiles were normalised so that their integral is 1. In the measurement, r_e plays the role of object size parameter, and we fixed the bulge-to-disc scale length ratio to $r_h/r_e = 0.15$ based on the same *Hubble* Space Telescope measurements (Welikala et al., in prep.). Finally, we imposed a hard cut-off on the surface brightness profile at $r_{\max}/r_e = 4.5$ since observations indicate that galaxies have truncated surface brightness distributions (Van der Kruit & Searle 1981, 1982). The parameters n_b , r_{\max}/r_e , and r_h/r_e are assumptions made in the modelling that can potentially lead to large biases in presence of a mismatch in the assumed Sérsic index compared to simulations (Simon & Schneider 2017). We stress that our choice of fixed values are based on recent observations, and the model bias due to incorrect assumptions are often intertwined with the particular simulation setup and its complexity. A detailed investigation of sensitivity of the calibration to bulge parameters is presented later in this paper.

The Sérsic model introduced above is an isotropic profile with zero ellipticity. To make it anisotropic (i.e., elliptical with ellipticity $\epsilon = \epsilon_1 + i\epsilon_2$), we used the following distortion matrix:

$$S = \frac{\bar{r}_0}{q_\epsilon r_0} \begin{pmatrix} 1 - \epsilon_1 & -\epsilon_2 \\ -\epsilon_2 & 1 + \epsilon_1 \end{pmatrix}, \quad (13)$$

where r_0/\bar{r}_0 is the scale factor necessary to scale a model of size \bar{r}_0 to the desired size r_0 . Because observed galaxy shapes are a 2D visual projection of an intrinsically 3D distribution, we introduce the additional scale factor, $q_\epsilon = 1 - |\epsilon|$, to make the profile semi-major axis invariant under ellipticity transformation.⁶ Discs and bulges typically show different intrinsic ellipticity. As discs will be observed more elliptical if edge-on, their ellipticity is primarily driven by inclination angle. In contrast, bulges are spheroids that are almost invariant under inclination, so they will appear more circular. In the measurement, we still applied the same ellipticity to both components as part of our 2D modelling, but we are aware that a positive ellipticity gradient from the intrinsically 3D distribution would induce a bias if not fully captured (Bernstein 2010). Our ellipticity estimate will therefore be a proxy of the inclination angle, especially for disc-dominated galaxies. Any residual ellipticity gradient, if significant, will have to be addressed separately as part of the calibration.

The *Euclid* telescope, optical elements, and detector introduce distortions of the input galaxy brightness distribution, which must be corrected. The effect is mostly convolutive, which tends to blur the galaxy image further. An example of a typical PSF image for a space-based telescope like *Euclid* is given in Fig. 1. The PSF is: (i) close to being diffraction limited; (ii) undersampled due to the half width being comparable with the pixel size; (iii) chromatic due to the integration over a wide range of wavelengths in the VIS filter; (iv) SED dependent due to integration being weighted by a combination of galaxy bulge and disc SEDs⁷; (v) spatially variant across the field of view due to optical distortions at the exit pupil; and (vi) epoch variant due to varying Solar aspect angle throughout the mission inducing thermal distortion on the optics. A comprehensive study of the modelling will be presented elsewhere (Duncan et al., in prep.). A smaller contribution also comes from the CCD pixel response function, which models the response of the detector pixel as an

⁶ Instead, $q_\epsilon = 1$ would make the profile *area* invariant under an ellipticity transformation. The choice of q_ϵ leads to different shear bias properties that can have significant impact on the final calibration (Fenech Conti et al. 2017).

⁷ With galaxy bulges being, on average, redder than discs. However, PSF images will be generated from the total galaxy SED.




Fig. 1. Example of a *Euclid* chromatic PSF for an assumed SED of a SBC-type galaxy at a redshift of one. The flux in the image was rescaled to its maximum value and oversampled by a factor of three with respect to the native VIS pixel size of $0''.1$. Diffraction spikes are clearly visible at a significant distance from the centre despite the blurring due to the chromaticity. The full width at half maximum, including its variation across the field of view, is $0''.1564^{+0''.0040}_{-0''.0019}$, comparable with the *Euclid* pixel size, so images will be undersampled at the Nyquist limit. The ellipticity is $\epsilon_{1,\text{PSF}} = 0.017^{+0.038}_{-0.024}$ and $\epsilon_{2,\text{PSF}} = 0.001^{+0.042}_{-0.020}$, with the superscript and subscript denoting absolute ranges.

integrated measure of the incoming flux illuminating individual pixels. In forward modelling, we included all the convolutive effects due to the telescope PSF and CCD, individually for bulge, disc, and for each image exposure. Colour gradients originate from incorrectly using the total galaxy SED when generating the PSF image, while the bulge and disc will have naturally different SEDs (Semboloni et al. 2013; Er et al. 2018). Using individually PSF-convolved model components may help control colour gradients, if individual SEDs were available. However, the impact of colour dependence and gradients on our analysis was not addressed here since we assumed the same broadband PSF as obtained from an SED of an SBC-type galaxy at a redshift of one in both simulations and measurements. Further non-linear distortions, such as in the case of charge transfer inefficiency (CTI, Rhodes et al. 2010) and the brighter-fatter effect (BFE, Antilogus et al. 2014), are typically corrected for at the image pre-processing stage, but residuals could still affect the shear measurement (Massey et al. 2014; Israel et al. 2015), which we did not include here.

Galaxy modelling for large-volume surveys like *Euclid* requires fast and efficient rendering of the images. All operations described so far are best suited to work in the Fourier space. We adopted a similar approach to galsim (Rowe et al. 2015). Consider the generic profile $I(r)$, which could be either Eq. (11) or (12). Because of its isotropy, the 2D Fourier transform is the 1D Hankel transform,

$$\tilde{I}(k) = 2\pi \int_0^\infty I(r) J_0(kr) r dr, \quad (14)$$

where k is the Fourier frequency (sampled on an oversampled grid) and J_0 is the Bessel function of the first kind. We call $\tilde{I}(k)$ the template model which any profile with arbitrary choice of parameter values (ellipticity, size, and position offset) can

be derived from⁸. To render a galaxy profile with parameters ϵ_1 , ϵ_2 , and r_e , we applied the inverse distortion matrix \mathbf{S}^{-1} to coordinates in Fourier space, so anisotropic coordinates are now defined by $\mathbf{k}' = \mathbf{S}^{-1}\mathbf{k}$. A position shift by $\delta\mathbf{r} = (\delta x, \delta y)$ from the centre⁹ is equivalent to the phase $\mathbf{k}' \cdot \delta\mathbf{r}$. The sheared-stretched-shifted model becomes

$$\tilde{I}(\mathbf{k}') = \frac{1}{\det \mathbf{S}} e^{-2\pi i \mathbf{k}' \cdot \delta\mathbf{r}} \tilde{I}(\mathbf{k}'), \quad (15)$$

where $\tilde{I}(\mathbf{k}')$ is the template calculated at the sheared-stretched Fourier mode \mathbf{k}' . Since isotropy is lost through the operation above, $\tilde{I}(\mathbf{k}')$ is no longer a Hankel transform but a full Fourier transform, which is a function of the vector \mathbf{k}' . We alleviate the problem of undersampling by calculating the PSF and galaxy model on coordinates with a common oversampling factor of three. Finally, the oversampled models for PSF and galaxy is multiplied together, the convolved model is downsampled by the same factor to the actual pixel scale¹⁰, and the downsampled convolved model is inverse fast Fourier transformed to real space. We applied the operations of shear and stretch to the template bank described in Appendix A to get any object of arbitrary ellipticity and size before the convolution with the PSF takes place as previously discussed.

The final galaxy model is a linear mixture of PSF-convolved co-centred components. We label the profiles with a subscript ‘d’ for disc and ‘b’ bulge:

$$I(r) = F_d I_d(r) + F_b I_b(r), \quad (16)$$

where F_d and F_b are disc and bulge fluxes, and $F = F_d + F_b$ is the total flux. If B/T is the bulge fraction, then the fluxes are also defined by $F_d = F(1 - B/T)$ and $F_b = F B/T$.

To summarise, given pre-computed template models for disc and bulge, we can generate a galaxy model with a desired ellipticity, size, and position by carrying out all the operations with simple algebra in Fourier space on oversampled coordinates, and then take one Fourier transform each time at the end. This is sufficiently fast for an intensive, repeated calculation of the same model with varying realisations of galaxy parameters (ellipticity, size, position offset, and fluxes). However, we kept n_b , r_{\max}/r_e , and r_b/r_e fixed in our modelling as allowing too much freedom would induce strong degeneracies between parameters and complicate the measurement substantially. We address the model bias sensitivity in Sect. 4.4.

2.2. Likelihood

Suppose we have multi-exposure image data vectors $\mathbf{D} = \{\mathbf{D}_{\text{exp}}\}$ ¹¹. We wish to estimate the model, $\mathbf{I} = \{\mathbf{I}_{\text{exp}}\}$, that best represents the available exposures. The model $\mathbf{I} = \mathbf{I}(\epsilon, \theta, \phi)$ is a function of ellipticity $\epsilon = (\epsilon_1, \epsilon_2)$, nuisance parameters $\theta = (r_e, \delta x, \delta y)$,¹² and linear flux nuisance parameters $\phi = (F_d, F_b)$.

⁸ For computational purposes, the transform is calculated once and saved in a cached 1D array to minimise computing time.

⁹ As it will be explained later in the section, in reality the modelling actually includes position shifts from right ascension, α , and declination, δ . Therefore, the position parameters will be $\Delta\alpha$ and $\Delta\delta$.

¹⁰ Downsampling in real space corresponds to aliasing in Fourier space, that is, n -folding the transform and summing up.

¹¹ On average 4 exposures for the Euclid Wide Survey and 64 for the Deep Survey fields.

¹² As explained, though, we model positions in world coordinates.

Assuming Gaussian data¹³, the likelihood can be written as

$$\ln p(\mathbf{D}|\epsilon, \theta, \phi) = -\frac{1}{2} [\mathbf{D} - \mathbf{I}(\epsilon, \theta, \phi)]^T \mathbf{C}^{-1} [\mathbf{D} - \mathbf{I}(\epsilon, \theta, \phi)] + \text{const}, \quad (17)$$

where \mathbf{C} is the noise covariance matrix usually estimated from the data as a block diagonal matrix, and the normalisation constant, $1/2 \ln[(2\pi)^{\lambda} \det \mathbf{C}]$ (λ is the dimensionality), is ignored. The noise is intrinsically non-stationary since various noise sources (such as the Poisson noise¹⁴ from the objects in the image) vary spatially. Because the model is linear in the component fluxes, $\mathbf{I} = F_d \mathbf{I}_d + F_b \mathbf{I}_b$, it is straightforward to integrate over the fluxes, $\phi = (F_d, F_b)$, and we have the following marginalised likelihood,

$$\ln p(\mathbf{D}|\epsilon, \theta) = \frac{1}{2} \mathcal{F}_{ij}^{-1}(\epsilon, \theta) \rho_i(\epsilon, \theta) \rho_j(\epsilon, \theta) + \text{const}, \quad (18)$$

where i indexes the model component, $\rho_i(\epsilon, \theta) = \mathbf{D}^T \mathbf{C}^{-1} \mathbf{I}_i(\epsilon, \theta)$ is a 2×1 vector, $\mathbf{I}_i = \partial \mathbf{I} / \partial F_i$ ($i = d, b$), and \mathcal{F}_{ij} is the 2×2 Fisher matrix,

$$\mathcal{F}_{ij}(\epsilon, \theta) = \mathbf{I}_i(\epsilon, \theta)^T \mathbf{C}^{-1} \mathbf{I}_j(\epsilon, \theta). \quad (19)$$

We note that because the right-hand side of Eq. (18) is quadratic in ρ_i and \mathcal{F}_{ij} is positive definite, we find $\ln p(\mathbf{D}|\epsilon, \theta) > 0$. A full derivation of the marginalised likelihood, including edge cases and implementation, can be found in Appendix B. The dimensionality of the problem has now been reduced from 7 free parameters to 5: ellipticity, size, and position offsets¹⁵.

Forward modelling provides solid grounds for a further generalisation to measuring multiple objects jointly, especially if they are observed within a short angular separation such as for neighbours. We label each likelihood with the index ω running through the objects being jointly measured, $\ln p_{\omega}(\mathbf{D}|\epsilon, \theta, \phi)$. The joint likelihood is then

$$\ln p(\mathbf{D}|\{\epsilon, \theta, \phi\}_{\omega}) = \sum_{\omega} \ln p_{\omega}(\mathbf{D}|\epsilon, \theta, \phi), \quad (20)$$

where $\{\epsilon, \theta, \phi\}_{\omega}$ is the set of all parameters for all the objects being measured. In the above equation we assumed the independence of the individual likelihoods. This is a fair assumption since close neighbours will very often be so due to random visual alignment. Consequently, those galaxies will be at different redshift and will have different shear. A much smaller fraction will include tidal interaction. In this case, the galaxies will be at the same redshift and have the same shear. It is then expected to have some degree of correlation between the individual likelihoods. In more extreme but much rarer cases, the galaxies will be tidally interacting, and therefore our Sersic modelling would break down entirely as we did not include any extra correlation term. Despite affecting a very small fraction of objects, dedicated simulations would be required to assess the impact on shear bias. Also, it is worth noting that we need to be careful with

¹³ The Gaussian approximation holds true in the limit of large counts in the image.

¹⁴ Poisson noise is a significant noise source especially for bright objects. This term is included in the simulations, but not in the measurement as it would require prior knowledge of the distribution profile that is being measured.

¹⁵ Having assumed that the two components are co-centred and the bulge size is locked to the disc size by a fixed rescaling.

Table 1. Summary of free, fixed, and derived parameters in the modelling.

	Free	Fixed	Derived	Initial value	Bounds	Unit	Description
ϵ_1	✓			0	$ \epsilon < 1$		First component of ellipticity in the tangent plane to $(-\alpha, \delta)$ as defined in Eq. (13)
ϵ_2	✓			0	$ \epsilon < 1$		Second component of ellipticity in the tangent plane to $(-\alpha, \delta)$ as defined in Eq. (13)
r_e	✓			0.3	$[0, 2]$	arcsec	Effective radius: disc scale length as defined in Eq. (13)
$\Delta\alpha$	✓			0	$[-0.3, 0.3]$	arcsec	Offset in right ascension as a phase prefactor in Eq. (15)
$\Delta\delta$	✓			0	$[-0.3, 0.3]$	arcsec	Offset in declination as a phase prefactor in Eq. (15)
α			✓			degree	Right ascension
δ			✓			degree	Declination
F_d	✓				$F_d \geq 0$		Disc flux: marginalised over or free as defined in Eq. (16)
F_b	✓				$F_b \geq 0$		Bulge flux: marginalised over or free as defined in Eq. (16)
F			✓		$F \geq 0$		Total flux
B/T			✓		$[0, 1]$		Bulge fraction as defined in text after Eq. (16)
S/N			✓		$S/N \geq 0$		Signal-to-noise ratio
I_E			✓				Magnitude: depends on the assumed zero-point, exposure gain and integration time
n_b		✓		1			Sérsic index of the bulge as defined in Eq. (12)
a_b		✓		depends on n_b			Sérsic coefficient of the bulge as defined in Eq. (12), see Peng et al. (2002)
r_h/r_e		✓		0.15			Bulge half-light radius to effective radius ratio as defined in text after Eq. (12), see Welikala et al., in prep.
r_{\max}/r_e		✓		4.5			Truncation radius to effective radius ratio as defined in text after Eq. (12)

Notes. The value should be interpreted as an initial guess or constant depending on whether the parameter was allowed to vary or remain constant.

the marginalisation of the individual likelihoods. The main issue lies in the marginalised likelihood of Eq. (18). This relies on calculating $\rho_i(\epsilon, \theta)$ for the various model components. However, when multiple objects are present in the same neighbourhood, this quantity will effectively introduce correlation between the likelihoods of the two objects. Therefore, the statistical independence required to multiply likelihoods together will not be ensured. We verified in testing that not marginalising individual likelihoods is indeed the correct approach to the problem. The joint likelihood is defined in a $7 \times N$ -dimensional parameter space, where N is the number of objects being measured jointly, with $N = 2$ being a typical number found in testing. For increased stability, we first optimised the likelihood for fluxes and positions offsets, and then also for ellipticity and sizes. This proved to be very robust as opposed to iterating over individual objects after masking neighbours to achieve a reliable initial guess (Drlica-Wagner et al. 2018). One key benefit of MCMC is that it marginalises the ellipticity of an object over all remaining nuisance parameters, which include object nuisance parameters as well as the parameters of the other objects included in the joint sampling (see Sect. 2.3).

Our prior is based on enforcing hard bounds on all parameters: $0 \leq |\epsilon| < 1$ given that the modulus of ellipticity in Eq. (2) cannot exceed 1; $0 \leq r_e \leq 2''$, where the upper bound is based on observations made in the *Hubble* Deep Fields; position offsets are restricted to ± 0.3 since the accuracy to which detections are made is typically sub-pixel; and fluxes are positive. A more

informative prior could be derived from real observations in the future. A summary of all parameters, being free, fixed or derived, is presented in Table 1.

2.3. Massive Markov chain Monte Carlo sampling

Shear measurement poses serious difficulties in identifying the best strategy to sample the posterior probability distribution of Eq. (4), assuming the likelihood of Eq. (18) or (20).

1. Since the lensing sample is very broad in morphological properties, it will contain both low and high S/N objects, whose posterior probability distribution can be either very broad or very narrow; hence any sampling strategy must be robust to this variability.
2. If N is the number of objects being measured jointly to mitigate the neighbour bias, the dimensionality of the problem is $7 \times N$ (with N being typically 2 and rarely 3 or 4); sampling must then be resilient to the large dimensionality, and provide marginalisation and error estimations with minimum overheads.
3. The shape of the distribution is a strong function of object parameters, such as ellipticity and size, and therefore it varies significantly across the sample; without prior knowledge of the physical properties of each galaxy, any sampling method must run in a consistent, robust way.
4. Given the large sample size of order 10^9 galaxies, sampling the posterior is a computational challenge, so a trade-off

between method complexity, runtime, and access to computing resources must be identified.

5. The sampling must be completely automated, without human supervision, and no fine tuning of sampling parameters and options is allowed.

Considering all these challenges, our priority must be the average convergence property of the sampler. The best strategy identified is MCMC, which allowed us to sample the posterior generated from the marginalised likelihood of Eq. (18) for an individual object, or the joint likelihood of Eq. (20) for a group of objects in an efficient and consistent way, for all 10^9 objects in the sample (hence the adjective ‘massive’). More importantly, MCMC seems to be the best choice to tackle neighbours, particularly as an estimate can be found in a high dimensional parameter space. It is worth noting that another key benefit of MCMC sampling is that it is both a maximisation and sampling procedure. The maximisation happens during the burn-in phase where the sampler tries to reach the region of higher probability. The actual sampling happens in the later stage of the chain after the burn-in phase. The marginalisation over nuisance and error estimation are then natural by-products with no extra overhead. This implies that not only can ellipticity estimates be marginalised over object nuisance parameters, but also over other object parameters in the joint group, hence minimising the impact of neighbours in the final ellipticity estimate.

When searching for such an algorithm that could potentially suit our needs, we considered a number of potential candidates that are widely used in cosmology and other fields (MacKay 2002). The development of various sampling methods is primarily driven by the quest to achieve lower auto-correlation and higher acceptance rate (Hastings 1970; Swendsen & Wang 1986; Skilling 2006; Goodman & Weare 2010; Foreman-Mackey et al. 2013; Betancourt 2017; Karamanis & Beutler 2021; Lemos et al. 2022). Although appealing, all these methods do suffer from increased complexity, which is the limiting factor in large-scale applications, where ‘large’ in this context implies runs repeated over a billion times. Even for the most sophisticated methods, it is often realised that a good initial guess is the key for good sampling of the posterior.

For shear measurement on the scale of large galaxy surveys, there is not much room for sampling complexity. The method has to be light enough and yet robust to all the posteriors that need to be sampled. Furthermore, the likelihood runtime limits the maximum number of samples that can be drawn for each galaxy without having an overall impact on the survey analysis runtime. The likelihood runtime is mostly dominated by the model component generation. The measurement is dominated by sampling the posterior, plus some additional pre-processing, therefore to limit the galaxy runtime to around a few seconds per galaxy, the MCMC sampling must be sparse. This is considered acceptable as we are interested in accurate shear estimates, which are found by averaging over ellipticity measurements. We chose an improved version of the Metropolis-Hastings algorithm, which was modified in two ways: (i) it incorporates some of the ideas of parallel tempering (Swendsen & Wang 1986; Sambridge 2013), so the parameter space can be sampled on a larger scale initially, and then on the smaller scale; (ii) it updates the proposal distribution during the burn-in phase, automatically tuning it to find a good match with the target distribution. Let ϑ be the generic vector of all parameters. As explained in Sect. 2.2, this is $\vartheta = (\epsilon, \theta)$ for individual galaxy measurement or $\vartheta = (\epsilon, \theta, \phi)$ when jointly measuring groups of galaxies. Ignoring the evidence since our

method is invariant to it, the Bayes’ theorem gives us

$$p(\vartheta|\mathbf{D}) \propto p(\mathbf{D}|\vartheta) p(\vartheta). \quad (21)$$

The parameter vector at the current iteration step is denoted with ϑ_t , where $t = 0, \dots$ is the MCMC sample index. We also define a tempering function, T_t , as a function of t . This acts as a Boltzmann temperature and its expression will be defined later in the text. When the temperature is high, $T_t \gg 1$, sampling from the posterior is equivalent to sampling globally from the prior. When the temperature is gradually reduced, as in annealing, $T_t \rightarrow 1$, we begin sampling directly from the posterior. We define such a tempering function for application during the burn-in phase only, and make sure $T_t = 1$ for the final part of the chain where we will take sample from. The method goes as follows.

1. At t , draw a new sample ϑ'_t from the proposal distribution $q(\vartheta'_t|\vartheta_t)$; here we assume a symmetric Gaussian proposal with mean ϑ_t and a pre-defined diagonal covariance of 10^{-4} on all parameters (in units of arcsec for size and position offsets).
2. Calculate the logarithm of the acceptance ratio

$$\ln A = \frac{\ln p(\mathbf{D}|\vartheta'_t) - \ln p(\mathbf{D}|\vartheta_t)}{T_t} + \ln p(\vartheta'_t) - \ln p(\vartheta_t). \quad (22)$$

3. Accept or reject ϑ'_t with probability A , that is, draw u from the uniform distribution on $[0, 1]$ and accept ϑ'_t if $u < \min(1, A)$; to speed things up and avoid calculating the likelihood outside the prior, we immediately reject ϑ'_t if $p(\vartheta'_t) = 0$.

For consistency, all posteriors are sampled from an initial guess that is a circular galaxy of mean size, $r_e = 0''.3$, and zero offset from the nominal position in the sky. If we were to run any MCMC method from this point onward we would end up with varying autocorrelation time depending on how far the actual galaxy is from the initial guess; therefore, we would need to wait longer for very elliptical, small, or large galaxies. To improve the convergence of the chains within a smaller number of iterations, we achieved a better initial guess by running an initial maximisation of the posterior before the actual MCMC. We ran the conjugate-gradient search method (Powell 1964) restricted to only 100 function evaluations, and then a downhill simplex search (Nelder & Mead 1965). The burn-in phase of the MCMC starts right afterwards. During this phase the temperature is gradually lowered to one. We adopted the following power law cooling scheme (Cornish & Porter 2007),

$$T_t = \begin{cases} 10^{f_{\text{heat}}(1-t/t_{\text{cool}})}, & \text{if } t < t_{\text{cool}} \\ 1, & \text{otherwise} \end{cases}, \quad (23)$$

where $f_{\text{heat}} = 10$ is the heat factor and $t_{\text{cool}} = 100$ is cooling-down sample index. The parameter t_{cool} represents the number of samples it takes for the tempering function to become one. We note that $T_0 = 10^{f_{\text{heat}}}$ and $T_{t_{\text{cool}}} = 1$. We begin with a diagonal Gaussian proposal of width 0.01 on all parameters, which is then recalculated from the chains every 100 samples and rescaled by the factor $2.4 \lambda^{-1/2}$, with λ being the number of parameters (Dunkley et al. 2005). The burn-in phase lasts for a total of 500 samples, which is long enough for the tempering function to become 1, the proposal covariance to be recalculated 5 times, and the chain to stabilise and reach the high probability region (well before we start accumulating the final chain samples). The

final phase lasts for an additional $N_{\text{MC}} = 200$ samples. Again, this is plenty to estimate both the mean and covariance of the chains with enough precision, but we recognise that sampling noise may still be non-negligible.

A good quantitative way to test the convergence of the chains is to investigate their auto-correlation. We do so for a variety of galaxies and results are shown in Appendix C. A less quantitative way is to verify that the acceptance rate is within the expected range. We also increased the final 200 samples up by a factor of five, without noticing any significant difference in the shear results. For further verification, we compared the method with our implementation of affine invariant (Goodman & Weare 2010) and parallel tempering (Swendsen & Wang 1986; Sambridge 2013)¹⁶. While these methods produce better ellipticity chains, they did not show any significant advantage in recovering shear, but increased complexity and therefore runtime overhead.

Once the samples are drawn from the distribution function, calculating the mean and width of the marginalised distribution becomes straightforward. Our point estimate for the ellipticity component marginalised over nuisance is the mean of the chain after the burn-in phase,

$$\hat{\epsilon}_i = \frac{1}{N_{\text{MC}}} \sum_t \epsilon_{i,t}, \quad (24)$$

where $i = 1, 2$, and $\epsilon_{1,t}$ and $\epsilon_{2,t}$ are the two ellipticity chains. The marginalised ellipticity covariance matrix is

$$C_{ij} = \frac{1}{N_{\text{MC}} - 1} \sum_t (\epsilon_{i,t} - \hat{\epsilon}_i)(\epsilon_{j,t} - \hat{\epsilon}_j). \quad (25)$$

We calculated the averaged per-component variance (ignoring negligible covariance between components),

$$C_\epsilon = \frac{1}{2}(C_{11} + C_{22}), \quad (26)$$

and chose to define the galaxy shear weight by

$$w = \frac{1}{C_\epsilon + \sigma_\epsilon^2}, \quad (27)$$

where σ_ϵ is the assumed shape noise (i.e., the width of the 1D intrinsic ellipticity distribution), ideally estimated in tomographic bins from deeper measurements. In this modelling, as C_ϵ quantifies the noise level in the data, faint galaxies will be automatically downweighted as opposed to very bright galaxies that will always have a maximum finite weight. Additionally, we note the negligible sensitivity to the choice of the $1/2$ factor in C_ϵ ¹⁷. The MCMC provides a convenient and efficient way to calculate both the mean and width of the ellipticity posterior at no extra computational cost. The weights can then be used to define sample averages, such as in one-point estimates:

$$\hat{g}_i = \frac{1}{\sum_k w_k} \sum_k w_k \hat{\epsilon}_{i,k}, \quad (28)$$

¹⁶ These alternative methods can be controlled by specific parameters in the code.

¹⁷ In fact, ignoring the $1/2$ factor would lead to a redefinition of the weight, $w = 1/[(C_{11} + C_{22})/2 + \sigma_\epsilon^2] \propto 1/[(C_{11} + C_{22}) + \sigma_\epsilon'^2]$ with $\sigma_\epsilon' = \sqrt{2} \sigma_\epsilon$, but results show weak sensitivity to the value assumed for σ_ϵ , as it will be demonstrated at the end of Sect. 4.2.

where k indexes the galaxies in the lensing catalogue. The generalisation to two-point estimates is straightforward. Please note that any choice of weight leads to shear bias due to correlation with the measured ellipticity, and this is tested in Sect. 4.2.

We also implemented the self-calibration of ellipticity proposed by Fenech Conti et al. (2017)¹⁸ via importance sampling and likelihood ratio while checking the quality of the sampling weights (Wraith et al. 2009) without finding strong evidence of improvement. As results will be dominated by other larger effects, we leave out further discussion from this paper.

2.4. Handling real data

Handling real data requires being careful with additional aspects of the measurement. For instance, throughout our discussion we proposed that our sampling strategy is best suited to handle the presence of neighbours, that is, resolved objects¹⁹ close to the lensing targets. However, the situation is complicated by the fact that there is more variety in real data as the brightness distribution of an object can be contaminated in different ways depending on the nature of the nearby objects:

1. neighbours (resolved galaxies or stars);
 2. blends (unresolved galaxies or stars);
 3. any other contamination (cosmic rays, transients, or ghosts).
- Each case leads to a particular type of bias, and therefore we deal with close objects in two ways. First, neighbours are grouped with a friend-of-friend algorithm to a maximum angular separation of $r_{\text{friend}} = 1''$. If the separation is too small, the objects will be mostly measured in isolation; therefore, they will still be affected by the neighbours due to improper masking. If the separation is too large, the groups will begin to be unmanageable in size, but the benefit in controlling the neighbour bias will be negligible. We found that $1''$ is a good trade-off between measuring N close neighbours jointly within a default postage stamp size of $38''.4$, and the non-negligible overhead in sampling a $7 \times N$ dimensional posterior. The joint analysis of object groups also gave us a good control of neighbour bias, leading to a correction of $m \approx -7 \times 10^{-4}$ as it will be shown later in the paper²⁰. Second, the segmentation maps and masks that are usually provided with the data (Bertin et al. 2020; Kümmel et al. 2020) were combined in a binary map and passed on to the likelihood to mask out objects in different groups. Detector artefacts or cosmic rays were also masked out in a similar way. In this case, to be even more conservative, we further dilated the masks by one pixel so most of contamination bias should be taken care of. But masking also helps partially with blends because objects that are false negatives according to the detection strategy may sometimes be true positives according to the masking procedure and would therefore be masked out. Blending with faint galaxies were demonstrated to be relevant when trying to calibrate methods that are particularly sensitive to the problem (Euclid Collaboration 2019). We demonstrate that, to some extent, this is also the case in our simulations where we measured objects deeper than the *Euclid* nominal survey depth, as we will show in the next section.

¹⁸ The same correction can also be proved to map, within some approximations, to other studies (Cox & Snell 1968; Refregier et al. 2012; Hall & Taylor 2017).

¹⁹ In this context ‘resolved’ implies that the object has been detected and at least partially deblended so that our measurement can be applied to all reported object positions.

²⁰ We do not attempt to optimise our choice of r_{friend} due to a number of other effects being more substantial than this.

Real images have a background level that needs to be subtracted. LENS MC uses the background estimates and noise maps that the *Euclid* processing provides, but residual local background variations are subtracted at the individual object group level. This is implemented by post-masking median subtraction. Likewise, the standard deviation of the background noise is estimated after masking. We measured galaxies and stars with the same model described earlier in this section. We found that good star-galaxy separation is based on selecting objects whose measured size is greater than the PSF size. This method fits well with the joint measurement of groups, however at the price of rejecting faint galaxies that would nevertheless have a negligible weight or be hard to distinguish from faint stars. More details will be given in Sect. 4.

The measurement was made in sky coordinates using the supplied world coordinate system (WCS) solution, which includes both linear and non-linear distortions (Greisen & Calabretta 2002). We assumed the default coordinate system $(-\alpha, \delta)$, where α is the right ascension and δ is the declination. We measured position offsets from the provided nominal position in arcsec. The resulting α and δ were reported in degrees, and r_e in arcsec. Likewise, the measured ellipticity is defined in the tangent plane to the $(-\alpha, \delta)$ coordinate system centred at the object position. We used the WCS to estimate a local linear approximation of the mapping from sky coordinates to tangent plane coordinates at the nominal position. We defined 9 points in a square grid of size $0''.3$ in pixel coordinates centred at the nominal position, mapped them to sky coordinates, and finally mapped the sky coordinates back to the undistorted tangent plane. The Jacobian matrix, which models the local linear approximation of the mapping, is the least-square solution to the mapping from sky coordinates to tangent coordinates. As part of this procedure, we also calculated the astrometric offsets due to the exposures being dithered differently. The brightness model was then correctly generated taking into account both the local distortion and astrometric offsets so all the observables were measured uniquely in tangent plane to sky coordinates.

When reporting our measurement we always compute $\chi^2 = -2 \ln p/\nu$, where p is the likelihood of Eq. (18) or (20) calculated at the mean estimate and ν is the number of degrees of freedom. The χ^2 will not in general follow the theoretical distribution for a number of reasons. The noise is only approximately Gaussian and non-Gaussianities will always be present in the data. For instance, key examples are the Poisson noise from the background and the object, digitalisation noise, non-linear artefacts, modelling mismatches, or failures in the sampling. Nonetheless, the χ^2 metric defined in this way is still a good statistical measure of the quality of the measurement. We also compared the χ^2 calculated above with the same quantity, which we call χ^2_{bkg} , after having masked out all the objects, which is expected to be just noise. Objects will be flagged up if the χ^2 is not consistent with the background. Following an F-test procedure, we calculated the test statistic $(\chi^2/\nu) (\chi^2_{\text{bkg}}/\nu_{\text{bkg}})^{-1}$ and rejected the null hypothesis (the measured χ^2 is consistent with the background) if the p-value was less than 0.01. Nonetheless we found that the impact of flagged objects is negligible, so we usually included them in our results. However, that may not be true for real data where the contamination from data artefacts will be more important.

The measurement includes estimation of the object magnitude based on the supplied zero-point, gain, and exposure time. Each exposure may come with its own values, as these vary both spatially and temporally. Therefore it is important to

normalise the data to common units. As the data is measured in analogue-to-digital units (ADU), we multiply each exposure by $G 10^{-(I_{E,0} - \bar{I}_{E,0})/2.5} / \tau$, where G is the gain in e^-/ADU , $I_{E,0}$ is the magnitude zero-point, $\bar{I}_{E,0}$ is the average magnitude zero-point across the exposures, τ is the exposure time, and the data is now in normalised photoelectron count rate of e^-/s . The flux, F , is then measured in the same units, and we can estimate the magnitude as follows:

$$I_E = -2.5 \log_{10} \frac{F}{e^-/\text{s}} + \bar{I}_{E,0}. \quad (29)$$

The specific values for zero-point, gain and exposure time assumed in our simulations will be provided in Sect. 3.

Analysing a volume of about 1.5 billion galaxies for *Euclid* will be a massive computational challenge, especially if employing MCMC to sample the posterior. Our measurement on highly realistic images ran at about 5 seconds per galaxy per exposure per computing core, including joint measurement of groups²¹. We discussed the benefits of using a fast, efficient implementation of MCMC in the previous section. Here we want to highlight the fact that all the pre- and post-processing described above adds very little overhead to the measurement. We found that the maximum posterior does suffer from a large bias of $m \approx -1 \times 10^{-2}$, which is about twice the bias obtained when using the mean of the MCMC samples. Since the bias tends to increase with magnitude, we interpret it as the maximum posterior estimate of the ellipticity being more prone to noise bias. This is further evidence that the MCMC can mitigate noise bias by consistent sampling and marginalisation of a multi-dimensional posterior, in particular when jointly measuring groups of objects, with the full complexity of real data and at the modest cost of slightly more overhead²².

3. Simulations

In order to validate our measurement method in a realistic setup, we designed a suite of simulations that incorporate most of the real data effects that future lensing surveys like *Euclid* will need to account for. It is essential then to bring in as much realism as possible. One problem that all shear methods have to deal with is clustering that leads to close neighbours, which is a concern for *Euclid*, Rubin, and present surveys as well. Because the inferred bias depends on the details about the realism of clustering of faint galaxies, this has to be incorporated in simulations particularly for calibration purposes (Kannawadi et al. 2019; Euclid Collaboration 2019). To make our custom simulations realistic and bring in all those effects we are most concerned about, we utilised the exquisite, state-of-the-art Flagship simulation mock galaxy catalogue (Potter et al. 2017; Euclid Collaboration 2024a),²³ developed specifically for *Euclid*. The same Flagship simulation is also used for the Euclid Science Ground Segment simulations (Euclid Collaboration 2024e). Flagship provides, in particular, a realistic distribution of galaxy morphologies, and clustering of galaxies obtained through a full N -body dark matter simulation.

²¹ The overhead of the joint measurement is about half of the quoted total.

²² Compared to the maximum estimate, the MCMC adds only 40% to the total runtime.

²³ We ingested the catalogue version 2.1.10 retrieved from the official website cosmohub.pic.es (Carretero et al. 2017; Tallada et al. 2020).




Fig. 2. Input magnitude-size distribution of galaxies. The data points are the mean r_e as a function of I_E . Also shown are the standard deviation of r_e and I_E in each bin. The horizontal band denotes the PSF full width at half maximum and its variation across the field of view. A significant fraction of the galaxies have intrinsic effective radii similar to the PSF, especially at the Euclid wide field (WF) depth.

The morphological parameters and spatial distributions are provided over an octant of the full sky, which is just less than 40% of the Euclid Wide Survey. Here we used values for the provided disc ellipticity and orientation angle, disc scale length, VIS flux, bulge fraction, and position over a region defined by $150^\circ < \alpha < 230^\circ$ and $15^\circ < \delta < 85^\circ$. We also selected all galaxies that are classified in the catalogue as being either central or satellite in the halo, and excluded quasars or high-redshift galaxies. Figure 2 shows the joint size-magnitude distribution of galaxies. A significant fraction of the galaxies have intrinsic effective radii similar to the PSF, which has a full width at half maximum of $0''.1564^{+0''.0040}_{-0''.0019}$, and therefore appear only marginally resolved in the PSF-convolved images. It is worth highlighting that because of the very faint magnitude limit ($I_E < 29.5$, but complete to $I_E < 27$) a significant fraction of the objects will be too faint to be detected, but these will be still included in the background noise. In addition to galaxies provided by Flagship, we also simulated a uniform spatial distribution of stars up to $I_E < 26$. Figure 3 shows the number count of galaxies and stars. The galaxy count was obtained from Flagship and compared against a polynomial model to VIS-corrected magnitudes in the GOODS-South field up to 26 (Giavalisco et al. 2004) and Ultra Deep Field beyond 26 (Beckwith et al. 2006). Stars were drawn from a polynomial model of i magnitudes generated with the Besançon model in the north ecliptic pole. Overall, we obtained a number density of 250 arcmin^{-2} ($I_E < 29.5$) for galaxies and 6 arcmin^{-2} ($I_E < 26$) for stars.

We defined sky patches of size $380''$ (about 40 arcmin^2), broadly corresponding to the size of a single *Euclid* CCD, but also included an adjacent area of extra 10% (called buffer/guard region) all around the patch to draw objects in the image that will not be part of the measurement. When selecting the morphological properties from the Flagship catalogue (ellipticity, disc scale length, bulge fraction, position, and flux), their AB flux was first converted to AB magnitude, then further converted to VIS photoelectron count rate via the magnitude-flux relation of Eq. (29). We assumed a constant magnitude zero-point of $I_{E,0} = 25.719$, which was calculated using *Euclid* as-designed




Fig. 3. Input differential number count for galaxies and stars. The solid line is the measured galaxy count from Flagship. The dashed line is a polynomial model of VIS-corrected magnitudes in the GOODS South ($I_E < 26$) and Ultra Deep Field ($I_E > 26$). Stars were drawn from a polynomial model of i magnitudes generated with the Besançon model in the north ecliptic pole. The cumulative number counts are 250 arcmin^{-2} ($I_E < 29.5$) for galaxies and 6 arcmin^{-2} ($I_E < 26$) for stars.

system throughput data. Star positions are drawn from the count model uniformly in each patch.

All galaxies in each patch had ellipticity assigned by Flagship. In principle we could use the cosmic shear from Flagship. However, in this work we applied the same constant shear to all galaxies in a patch, with the shear varying from one patch to another according to a uniform distribution on a circle of radius $|q| = 0.02$ and random orientation. This choice mimics the typical shear expected for a real survey and also minimises the error on multiplicative bias. We assumed an (infinitely thin) annular distribution as opposed to a disc distribution or an even more realistic log-normal distribution because we wanted to minimise the statistical error on multiplicative bias, and obviously be as cosmology agnostic as possible. On the other hand, a variable shear field might in principle introduce an additional correlation with neighbour bias, particularly if neighbours at different redshifts are considered (MacCrann et al. 2021). However, capturing realistic clustering is the most important aspect of the simulations, which is what we focussed on in this work. Similarly to previous work (Bridle et al. 2010; Kannawadi et al. 2019), we applied shape noise mitigation by making, in total, 4 clones of each patch with all ellipticities rotated by 45° while maintaining the same overall constant shear, which gave us significant leverage on the required simulation volume. It is worth noting, though, that a varying shear could also be dealt with in a shear response approach, leading to an increased effective sample size and reduce simulation volume in calibration (Pujol et al. 2019; Jansen et al. 2024).

We set up a suite of simulations for each of 9 realisations of the PSF image drawn at different positions in the field of view. While varying the PSF image, we kept the objects at the same positions. We assumed a *Euclid* PSF model for a fixed SBC-type galaxy SED at a redshift of one, the median of the distribution. The mean ellipticity and its variation across the field of view was: $\epsilon_{1,\text{PSF}} = 0.017^{+0.038}_{-0.024}$ and $\epsilon_{2,\text{PSF}} = 0.001^{+0.042}_{-0.020}$, with the superscript and subscript denoting absolute ranges. We note that this variation, if not included in the modelling, would be responsible

for an error in the shear measurement that would far exceed science requirements. We did not include PSF mismodelling in our simulations as the current *Euclid* requirement on PSF ellipticity error is already quite stringent, but will be addressed elsewhere. The Euclid Wide Survey is designed to take 4 dithered exposures (pointings), plus two extra short exposures, of the same sky area. Most often these will be taken in the same observation. Hence the PSF model is not expected to vary too much across the exposures, but the images will be different because taken at different positions in the field of view.

We generated *Euclid* detector images containing galaxies rendered with the brightness model of Sect. 2.1 with varying ellipticity, r_e , position, and fluxes. In our initial tests, we made our results insensitive to model bias by construction, and therefore we fixed n_b , r_h/r_e and r_{\max}/r_e . Later on, we address model bias sensitivity by allowing n_b , and r_h/r_e to vary. For stars, we used a restricted model with zero ellipticity and r_e , so we effectively rendered point-like PSF images. For the measurement, we used the same galaxy profile (with fixed bulge parameters) for all detected objects.

The pixel photoelectron noise is given by

$$\sigma_{\text{px}}^2(x, y) = (R_{\text{bkg}} + R_{\text{dark}}) \tau + \lambda_{\text{obj}}(x, y) + \sigma_{\text{read}}^2. \quad (30)$$

The first term is Poisson noise from a constant zodiacal light background and dark current, with rates $R_{\text{bkg}} = 0.225 \text{ e}^-/\text{s}$ and $R_{\text{dark}} = 0.001 \text{ e}^-/\text{s}$, and exposure time $\tau = 565 \text{ s}$. The second term, $\lambda_{\text{obj}}(x, y)$, is spatially varying Poisson noise from all the objects in the image, which is non-negligible in the *Euclid* VIS images. The third term is Gaussian noise from the CCD read-out with a constant $\sigma_{\text{read}} = 4.5 \text{ e}^-$. For the purpose of this work, we assumed that all noise sources are uncorrelated²⁴. In generating the images, we also applied a bias of 2000 e^- (about as expected for *Euclid*), and finally digitised the data. Digitalisation corresponds to dividing the image by a gain of $3.1 \text{ e}^-/\text{ADU}$ and floor truncating it to nearest integer, which itself adds uniform noise of variance $1/12 \text{ ADU}$ ²⁵. We set a tangent projection as our WCS at the centre of the patch, drew 4 undithered exposures and stacked them up by taking their average. These images were used by the object detection for the main results presented here, but we will also include a discussion about the dithering. An example of stacked CCD image is shown in Fig. 4.

4. Results

To quantify the performance of LENS MC on our realistic LENS MC-Flagship simulations, we ran the measurement on about 45 000 random patches, which is equivalent to an area of about 500 deg^2 , with mean number density, according to Fig. 3, of 250 arcmin^{-2} ($I_E < 29.5$) for galaxies and 6 arcmin^{-2} ($I_E < 26$) for stars. We measured the same area (with the objects at the same positions) 9 times for varying noise realisations and PSF across the field of view, totalling an equivalent, effective area of 4500 deg^2 . We ran all our simulations across the GridPP UK network (Faulkner et al. 2005; Britton et al. 2009)²⁶. A qualitative test of the measurement performance is shown in Fig. 5.

²⁴ Pixel noise is also uncorrelated with good approximation in real data especially when working with individual exposures, except for residual detector artefacts.

²⁵ The in-orbit detectors will have slightly larger gain, likely around $3.4 \text{ e}^-/\text{ADU}$, but this is not expected to change any of our results.

²⁶ Testing took 6 months, with our final run averaging 15 000 cores/day for two weeks.




Fig. 4. Example of LENS MC-Flagship image. The input galaxy distribution was provided by Flagship and stars were drawn from a model. We emulated the VIS detector by including realistic image properties and noise, but we did not include non-linearities, CTI, BFE, or cosmic rays. To aid the visualisation of the faint objects, the image was clipped between the tenth and ninetieth percentiles, illustrating the sheer number of objects and their clustering. The image size is 4096 pixels.

After the galaxy models was subtracted from the image data, the residuals looked consistent with noise, for galaxies measured individually or jointly in groups, despite the presence of neighbours. More quantitative tests will be discussed as part of the validation presented in Appendix D.

4.1. Selection

For our baseline test, we ran SourceExtractor++ (Bertin et al. 2020; Kümmel et al. 2020)²⁷ to detect galaxies and stars in each of the undithered stacked images. The code attempts to deblend detected objects and produces a detection catalogue with a total number density of 88 arcmin^{-2} ($I_E < 26.5$), and 34 arcmin^{-2} ($I_E < 24.5$). Figure 6 contains the galaxy magnitude selection applied by SourceExtractor++. This shows the number count of the objects in the simulation and after the detection by SourceExtractor++. The detection catalogue is complete to the magnitudes of interest, apart from false positives of about 6 arcmin^{-2} ($I_E < 24.5$), probably due to a combination of sub-optimal detection and mismatching with the true input catalogue in presence of neighbours at those magnitudes.

The detections were grouped (with $r_{\text{friend}} = 1''$) according to their reported SourceExtractor++ positions. LENS MC went through each object group and measured the object parameters starting from an initial guess at the provided SourceExtractor++ positions. If the size of the group is 1, LENS MC will measure the object in isolation and masks out neighbours through the supplied segmentation maps. Instead, if it is greater than 1, LENS MC will measure the objects jointly, while masking out neighbours belonging to other groups. We

²⁷ Version 0.19.2 with default settings as used in *Euclid*. We do not test the sensitivity to changes in these settings and that will be the focus of future work.




Fig. 5. Examples of measurement performance. The target galaxies are denoted with a cross. The image residuals look consistent with noise, for galaxies measured individually or jointly in groups, despite the presence of neighbours. All images have a size of 128 pixels.

matched the input catalogue with the measurement catalogue and within a maximum angular distance of $0''.3$ from the measured object (which also corresponds to the LENS MC maximum search region around the detected object). The few measured objects that did not get a useful match to within that distance were then flagged up and removed from the analysis. We tested the sensitivity to the maximum match distance without noticing any appreciable change to the bias.

A key selection applied to the measurement catalogue is the star-galaxy separation. As found in applications to real data, the object size is an excellent statistic to discriminate between galaxies and stars (Sevilla-Noarbe et al. 2018). Therefore we classified objects according to measured $r_e > r_{s/g}$ where $r_{s/g} = 0''.15$, which is slightly larger than the pixel size and image resolution. We note that we applied our star-galaxy separation to broadband data simulated with a fixed choice of SED representative of a typical galaxy at a redshift of one. However, this does not test how well the star-galaxy separation works with a broad range of galaxy SEDs and with a clear distinction between galaxy and star SEDs.




Fig. 6. Differential number count for all objects in the simulation and after the detection by SourceExtractor++. The detection catalogue is mostly complete to $I_E < 24.5$, apart from false positives of about 6 arcmin^{-2} ($I_E < 24.5$). The distribution starts rolling off at 26, so a large fraction of faint objects is not detected.

We quantified the performance of our separation by calculating the following: *i*) N_g , the number of true positives – galaxies correctly identified as such; *ii*) N_s , the number of true negatives – stars correctly identified as such; *iii*) N_{-g} , the number of false positives – stars wrongly identified as galaxies; *iv*) N_{-s} , the number of false negatives – galaxies wrongly identified as stars. The above numbers are always defined in the measurement catalogue. The true positive rate (TPR) and false positive rate (FPR) are

$$\text{TPR} = \frac{N_g}{N_g + N_{-s}}, \quad (31)$$

$$\text{FPR} = \frac{N_{-g}}{N_{-g} + N_s}. \quad (32)$$

Realistic values of $\text{FPR} > 0$ and $\text{TPR} < 1$ are always linked to type I and II errors in the shear analysis. Type I is the inclusion of stars in the lensing sample, hence leading to potentially large negative multiplicative bias. Type II is the omission of galaxies (with potentially large shear signal) from the lensing sample which introduces selection bias and also a dilation in statistical error.

For the sample of detected objects to the detection limit ($I_E < 26.5$) we found $\text{TPR} = 93.3\%$, $\text{FPR} = 4.6\%$, purity of 99.8% ²⁸, and a star fraction of 6.6% ²⁹. The TPR gives us the frequentist probability of a positive being a galaxy, so $\text{TPR} = p(+|g)$. Similarly, $\text{FPR} = p(+|s)$. Bayesian posterior probabilities provide a more meaningful interpretation of those numbers. The prior probability of an object being a galaxy is $p(g)$ and a star is $p(s) = 1 - p(g)$ (i.e., the star fraction). Applying Bayes' theorem, we get the probability of a galaxy given a positive detection,

$$p(g|+) = \frac{p(+|g)p(g)}{p(+|g)p(g) + p(+|s)p(s)}, \quad (33)$$

and similarly for $p(s|+)$. With the numbers above, we obtained $p(g|+) = 99.7\%$ and $p(s|+) = 0.3\%$ for all objects in the detection catalogue. A more relaxed FPR of about 20% would still

²⁸ Astronomical completeness coincides with TPR, but purity = $N_g/(N_g + N_{-g})$.

²⁹ Star fraction = $N_s/(N_g + N_s)$.




Fig. 7. Star-galaxy separation of detected objects. (Top) Observed size-magnitude distribution of true galaxies and stars. The data points are the mean of r_e as a function of I_E . Also shown are the standard deviation of r_e and I_E in each bin. $r_{s/g} = 0''.15$ provides a good separation threshold working up to $I_E < 24$. (Bottom) Operating curve showing where our threshold (indicated with a cross) sits in terms of true positives (92.9%) and false positives (4.3%). The area under the operating curve is large, and the curve itself is reasonably flat for a wide range of false positive rate suggesting excellent discrimination and weak sensitivity on the threshold (the shear bias does not appreciably change). In real measurements this could be further optimised through access to external data or simulations.

give us $p(g|+) = 99\%$ and $p(s|+) = 1\%$, given the strong imbalance between the galaxy and star samples. These numbers give us reassurance that once an object is classified as a galaxy, there will be an average 3σ confidence that it will indeed be a galaxy for the entire sample up to the detection limit ($I_E < 26.5$).

Figure 7 shows the size distribution of true galaxies and true stars and the operating curve (TPR versus FPR for varying threshold) of our classification with either a horizontal line or a cross to denote the default threshold, $r_{s/g} = 0''.15$. Both plots provide solid justification for our choice of $r_{s/g}$, but confusion is evident around $I_E = 24$, which might explain most of the false positives. The area under the operating curve at the bottom of Fig. 7 is large, and the curve itself is reasonably flat for a wide range of FPRs, suggesting excellent discrimination and weak sensitivity on the threshold (in that the shear bias does not appreciably change for a wide range of threshold values around the




Fig. 8. Shear weight after star-galaxy separation of detected objects as a function of magnitude separately for true galaxies and stars. As the star weight is systematically lower than the galaxy weight, this drastically reduces the impact of those residual stars in the catalogue up to the faint magnitudes.

nominal one). However, in real measurements, an optimal value could be inferred from external data or simulations, hence allowing for a dramatic reduction of the false positives at the expense of a modest reduction in the true positives. Our TPR, FPR, and operating curve of Fig. 7 are consistent or better than the best estimators presented in Sevilla-Noarbe et al. (2018), although a key caveat in our work is likely to be that we did not include a full colour variation of galaxy and star SEDs, which would lead to variable PSFs and potentially harder separation. Moreover, we did not investigate any effect due to star density variation, which might well change by a factor of two or three going from the high to the low latitudes.

We defined our final shear weight by multiplying Eq. (27) by the step function, $H(r_e - r_{s/g})$ and show this as a function of magnitude for detected true galaxies and stars in Fig. 8. As the star weight is systematically lower than the galaxy weight, this drastically reduces the impact of those residual stars (false positives) in the catalogue up to the faint magnitudes.

The quality of our star-galaxy separation can only be tested by fully propagating results through shear bias. We calculated the shear bias for perfect star-galaxy separation (where we enforce knowledge about the truth; that is, we do not use our classification but exclude true stars from the galaxy catalogue), and compared it with our nominal analysis. We did not see any statistically significant difference in shear bias between the two cases. Additionally, we varied the value of $r_{s/g}$ and again found that the bias does not change appreciably, but this may not hold true for more realistic SED variation.

4.2. Shear bias

As a preliminary validation, Appendix D contains a few distributions and correlations. We tested the bias as a function of input true magnitude to avoid large selection biases due to binning by observed quantities, such as S/N or r_e , which strongly correlate with shear (Fenech Conti et al. 2017). We defined 12 bins in $20 < I_E < 24.5$ and, in each bin, we regressed the measured ellipticity against input true shear via weighted least square as described in detail in Appendix E. We also wanted to clearly separate measurement (shear measurement method)




Fig. 9. Multiplicative and additive biases averaged over the magnitude selection $20 < I_E < 24.5$ and PSF variation across the field of view for the measurement and selection. The triangle and the circle denote each of the two components. The light shaded area is the *Euclid* requirement on knowledge of m and c , respectively $\sigma_m < 2 \times 10^{-3}$ and $\sigma_c < 3 \times 10^{-4}$. The dark shaded area is the ideal target for measurement alone, respectively $\sigma_m < 5 \times 10^{-4}$ and $\sigma_c < 5 \times 10^{-5}$. Reference values can be found in Table 2.

from purely selection (detection, catalogue matching, weights, and star-galaxy separation) effects. We let g_i be the input true shear, $\hat{\epsilon}_i$ the measured ellipticity, and $\epsilon_{i,\text{sel}}$ the input true sheared ellipticity on the same selection. Similarly to Eq. (8), where we regress estimates of shear with input true shear, we can define the same regression of our estimate of shear (i.e., ellipticity)³⁰,

$$\hat{\epsilon}_i = (1 + m_i) g_i + c_i + n_i, \quad (34)$$

$$\epsilon_{i,\text{sel}} = (1 + m_{i,\text{sel}}) g_i + c_{i,\text{sel}} + n_{i,\text{sel}}, \quad (35)$$

where $i = 1, 2$ indexes the ellipticity or shear component, n_i and $n_{i,\text{sel}}$ are the statistical noise components for measurement and selection. The first equation estimates the total of measurement and selection bias, whereas the second one estimates the selection bias. Therefore, if we take the difference between the two and regress $\hat{\epsilon}_i - \epsilon_{i,\text{sel}}$ with g_i ,

$$\hat{\epsilon}_i - \epsilon_{i,\text{sel}} = m_{i,\text{meas}} g_i + c_{i,\text{meas}} + n_i - n_{i,\text{sel}}, \quad (36)$$

we can then directly estimate the measurement bias, having coherently subtracted the selection bias, and reduce the statistical noise thanks to the correlation between n_i and $n_{i,\text{sel}}$. The measurement-only bias is then defined as $m_{i,\text{meas}} = m_i - m_{i,\text{sel}}$ and $c_{i,\text{meas}} = c_i - c_{i,\text{sel}}$.

The main performance metric that we present here is the total bias computed as a weighted average over magnitude bins and PSF variation across the field of view as shown in Fig. 9. The performance metric is defined in an m - c plane for each of the two components fitted independently. We indicate the *Euclid* requirements σ_m and σ_c in the shaded areas, both the total one and the desired for measurement alone. From looking at the summary figures of Table 2, selection effects are dominating the error budget, with a pronounced asymmetry between the two components. We tested that this is not due to the star-galaxy separation, weights, or the particular PSF used here by

³⁰ It is worth noting that a least-square regression of ellipticity is a linear operation that corresponds to calculating the mean ellipticity, that is, estimating shear.

Table 2. Multiplicative and additive biases.

	Measurement	Selection
$m_1/10^{-3}$	-3.58 ± 0.18	11.8 ± 1.0
$c_1/10^{-4}$	-1.797 ± 0.025	-3.40 ± 0.15
$m_2/10^{-3}$	-4.30 ± 0.18	0.0 ± 1.0
$c_2/10^{-4}$	0.088 ± 0.025	-0.21 ± 0.15

Notes. Biases averaged over the magnitude selection $20 < I_E < 24.5$ and PSF variation across the field of view, for measurement and selection (detection, catalogue matching, star-galaxy separation, and weights).

varying each parameter and checking that results remain consistent with the default analysis. To investigate if the origin of this asymmetry could be due to the input distributions, we estimated the bias of the input ellipticity (before detection) in exactly the same way as we did for our measurements, finding no bias up until the detection is run. We defer the investigation of sensitivity to the SourceXtractor++ configuration to future work. For the time being, we highlight that the multiplicative bias owing to measurement alone (i.e., the shear measurement method) is about -4×10^{-3} with an uncertainty of 2×10^{-4} , and a small residual asymmetry in additive bias. All statistical errors for selection and measurement were estimated following the modelling of Eq. (34) and the analytical solution presented in Appendix E. The modelling effectively uses the high correlation between measurement and selection estimates to reduce the error on bias. Once selection and measurement biases are calculated, the total bias is just the sum of the two individual values. The total error is then the sum in quadrature, given that the individual values have had their correlation removed.

A more in-depth investigation can be carried out when looking at bias as a function of true input magnitude as shown in Fig. 10. As already mentioned, not only is the correlation between magnitude (or flux) with shear negligible (apart from magnification, not included here), but defining true input bins is also essential to minimise the impact of selection bias and not to misinterpret results. Curves were averaged over the PSF variations across the field of view. We note that m and c show a negative trend at the faint magnitudes. This suggests that the total bias shown in Fig. 9 is mostly dominated by those bins, which happen to have the largest relative weight due to the number count increasing with magnitude. The requirement, σ_m and σ_c , on each bin was derived from the total requirement by increasing the per-bin variance by the decrease in number count in each bin. The error bars are consistent with this.

We tested the sensitivity to the faint undetected objects by calculating the total bias as a function of the intrinsic limiting magnitude in the simulations. For this we repeatedly rendered images excluding the faintest objects, with a varying magnitude limit. We expect that the brighter the magnitude limit in the simulation, the weaker the impact of faint objects on the measurement (as their relative fraction becomes small). Figure 11 shows different trends in bias with the magnitude limit. For multiplicative bias, the selection bias seems to be insensitive to the magnitude limit; instead, the measurement bias seems to be symmetric (components are consistent with one another) and shows a trend with the magnitude limit, with a slight hint of flattening out at the faintest end. This effect on the measurement bias indicates a circularisation bias due to the faint objects. Overall, we estimated $m_{\text{faint}} \approx -5 \times 10^{-3}$ just due to the presence of faint




Fig. 10. Measurement multiplicative and additive biases in bins of I_E and averaged over the PSF variations across the field of view. The triangles and circles denote each of the two components. Shaded areas are the *Euclid* requirements relaxed by the increase in variance in each bin. Except for $c_{1,\text{meas}}$, which is fully consistent with requirements, all other biases show a slight trend in the faintest bins.

undetected objects. This is fully consistent with earlier predictions of a few 10^{-3} up to 10^{-2} depending on the clustering of the faint objects (Euclid Collaboration 2019). We think the flattening of the measurement curves with the magnitude limit might be due to faint galaxies having less impact the fainter they are or perhaps a lack of an ultra-faint population. However, this indicates that any calibration strategy relying on external images should render galaxies with a magnitude limit of at least 27.5, which is 3 deeper than the Euclid Wide Survey (in agreement with Hoekstra et al. 2017), and the sensitivity to that limit should be investigated as well.

We estimated the bias due to close detected neighbours by comparing results from running the measurement in two modes: $r_{\text{friend}} = 0$ or $r_{\text{friend}} = 1''$. As discussed in Sect. 2.4, the first mode corresponds to no grouping of detections at all, so LENS MC measures all the objects individually, hence relying on the supplied maps to implement the masking of neighbours. In fact, masking provide limited help when the objects are too close to each other: the final ellipticity estimate of the target object will be slightly biased towards the neighbour. Because of the random orientation around the target object, the net effect is a circularisation of the average ellipticity if not corrected for.




Fig. 11. Multiplicative and additive biases for measurement and selection as a function of the intrinsic limiting magnitude in the simulations for a PSF chosen at the centre of the field of view. The triangles and circles denote each of the two components. Shaded areas are the *Euclid* requirements. The measurement bias shows varying trends with the limiting magnitude and asymmetries between components in some cases. (See text for discussion.)

Instead, the second mode makes groups of objects that are measured jointly. In total, almost 96% of the objects were still measured in isolation, 4% in pairs, and 0.2% in groups that include triplets, quadruplets, and some rare quintuplets. The measurement of the groups increases the robustness to neighbours and mitigates the reliance on the accuracy of the maps at short angular separations between objects. We found a differential multiplicative bias of $m_{\text{neighbour}} \approx -7 \times 10^{-4}$ due to the masking of close neighbours (4.2% of the sample) when measuring objects individually ($r_{\text{friend}} = 0$), and when joint measuring them in groups ($r_{\text{friend}} = 1''$). As the neighbour bias predominantly effects the short separation between objects and the fraction of multiplets compared to the pairs is very small, we expect that increasing r_{friend} further would provide little benefit to the bias, but at a much increased computational expense.

We also checked the dependence of bias on the weight definition of Eq. (27) or the one employed by KiDS (Miller et al. 2013),

$$w_{\text{KiDS}} = \left(\frac{C_\epsilon \epsilon_{\text{max}}^2}{\epsilon_{\text{max}}^2 - 2C_\epsilon} + \sigma_\epsilon^2 \right)^{-1}, \quad (37)$$

as a function of the assumed shape noise, σ_ϵ . We estimated the first-order sensitivity as a linear regression to the bias for varying σ_ϵ . For either definition, we found weak sensitivity: $dm/d\sigma_\epsilon \approx 4 \times 10^{-4}$ and $dc/d\sigma_\epsilon \lesssim 4 \times 10^{-6}$. The implication is such that a change in the assumed σ_ϵ of, for example, 20% would be responsible for an additional bias of order of 10^{-5} .

Finally, we carried out a test on exposures dithered randomly between $[-0.05, 0.05]''$. In reality, the dithering could be as large as a few arcminutes but sampling is always affected by the random sub-pixel shifts. In a real survey, any stacking procedure, even if applied to nominally undithered exposures, will be affected by the random shifts in the telescope pointing at the sub-pixel level. The combination of such exposures will inevitably introduce pixel correlations in the stacked images and PSFs, a problem that would be exacerbated by combining exposures at different epochs. The aim of this test is to verify that the results between dithered exposures and our perfectly undithered exposures (as presented throughout the text) are fully consistent, so to demonstrate that the method will perform well on dithered exposures on real data. However, we did not quantify the impact of stacking in our tests, nor directly assess any benefit from analysing dithered exposures over stacked exposures. For dithered exposures, we resample-coadded the images with SWarp (Bertin et al. 2002), ran SourceXtractor++, and remapped the segmentation maps back to individual exposures with SWarp. This data was then passed on to LENSMC as usual, with the only difference that it now carried out the measurement jointly across exposures. We ran the measurement on one of the PSF images at the centre of the field of view and processed the catalogues as usual. We found no statistically significant difference in the estimated bias between the two cases of random dithering and perfectly undithered exposures (as used in the main simulations of this paper). This gives us reassurance that the joint measurement of individual exposures would perform well on real data, and also better than stacking because data interpolation would be avoided. However, it also suggests that any undersampling bias was probably below the statistical uncertainty to be seen in these simulations.

4.3. PSF leakage

Having multiple realisations for the spatially varying PSF model allows us to investigate the dependence of bias on the PSF. Figure 12 shows the calculated leakage terms, α_1 and α_2 , obtained from a linear fit to measured c against input PSF ellipticity. We found $\alpha_{1,\text{sel}} = (-2 \pm 3) \times 10^{-3}$ and $\alpha_{2,\text{sel}} = (-8 \pm 3) \times 10^{-3}$ for selection and $\alpha_{1,\text{meas}} = (-9 \pm 3) \times 10^{-4}$ and $\alpha_{2,\text{meas}} = (2 \pm 3) \times 10^{-4}$ for measurement, averaged over all magnitude bins. We note that the measurement leakage is within the empirical requirement derived in Sect. 2, but $\alpha_{2,\text{sel}}$ is clearly not. Combined with the results from the previous section, in particular Table 2, we can observe that the asymmetry in c is most likely due to the PSF having a factor of ten larger first component of ellipticity in detector frame (which in our setup is anti-aligned with world coordinates along right ascension). It is worth adding that the negligible measurement leakage ensures that this residual additive bias is constant across the field of view and hence potentially straightforward to calibrate. However, the same statement would not entirely apply to the selection leakage where a residual term would still complicate the calibration.

Finally, we tested the consistency between the measurement curves when assuming perfect star-galaxy separation or not. In fact, the more elliptical the PSF, the larger the residual additive bias, and the harder separating galaxies from stars will be.




Fig. 12. Measurement PSF leakage in bins of I_E . The triangles and circles denote each of the two components. Shaded areas are the projected *Euclid* requirements forward propagated by the corresponding ones on c , relaxed by the increase in variance in each bin. The measurement terms are consistent with requirements, except for the first component in the faintest bins.

Due to the intertwining of chromaticity and variation across the field of view of the PSF, we would expect some leakage due to imperfect star-galaxy separation. Luckily, we did not see that the star-galaxy separation could be appreciably impacting our results in our simulations, but we realise this may not be the case with the full chromaticity of real data.

4.4. Model bias calibration

Having the same galaxy brightness model in both simulations and measurements allows us to isolate the various contributions to the total shear bias, while separating them from the issue of model bias. However, this raises the reasonable concern about whether model bias could be the dominant source of error. Bearing in mind that model bias is usually addressed as part of shear calibration (which is outside of the scope of this paper), we still want to provide reassurance by providing results after we relax some key assumptions made in the previous sections. Then the question of calibratability ties in closely with the sensitivity to the assumptions that are made in the calibration simulations. To address this, the KiDS calibration relies on applying the measurement method to reference simulations and new realisations of the same after having scaled key parameter distributions up and down, with residual biases present in some cases (Li et al. 2023a,b).

Here we chose a similar approach by relaxing some key assumptions in the simulations and then investigated the sensitivity of the calibration. By looking at the list of model parameters in Table 1, we can immediately recognise that the parameters that were fixed and matched in measurement and simulations take priority. Key parameters are those of the bulges, namely n_b (hence a_b) and r_b (hence r_b/r_e).³¹ Investigating the impact of a distribution of variable bulges to shear bias also makes sense as these are more compact than the discs and not fully captured by LENSMC, which always assumes a fixed r_b/r_e in the measurement which could lead to bias. In these tests, we allowed

³¹ We leave out r_{max}/r_e from this work for now as we are more concerned about the impact of the bulges on the calibration of the model bias and the uncertainty in the knowledge of their distribution.

for a broad variation of bulge parameters, shifted their distributions up and down as in the KiDS calibration, and saw what the impact was on the calibrated shear bias. As part of this tests, we also made some changes to the reference dataset. We selected the patches randomly within the available simulation area, simulated objects within a large model array size, and included bulge-only galaxies (accounting for an additional 1% of the galaxies). The last change is required because when allowing the full realism of the bulges, this sub-population of bulge-only galaxies could play a role in model bias. We therefore have the following datasets:

1. a revised reference dataset where the bulge parameters were still fixed and bulge-only galaxies were artificially rendered as two-component galaxies;
2. a new dataset where the bulge parameters were now allowed to vary following a broad distribution of n_b and r_h ;
3. calibration datasets where the bulge parameters were scaled up and down.

In all cases, during the measurement, LENSMC was applied with the same assumptions described earlier in this paper, which allowed us to study the differential bias from any assumptions made in the simulations.

While the revised dataset in 1 contains slightly more galaxies, the measurement bias is consistent with results presented in the previous sections. In particular, in this case we found $m_{1,\text{meas}} \approx m_{2,\text{meas}} \approx -3 \times 10^{-3}$, $c_{1,\text{meas}} \approx -1 \times 10^{-4}$, and $c_{2,\text{meas}} \approx -2 \times 10^{-5}$ (cf. Table 2). However, the selection bias now shows a less pronounced asymmetry between the two components: $m_{1,\text{sel}} \approx 7 \times 10^{-3}$, and $m_{2,\text{sel}} \approx 1 \times 10^{-2}$, suggesting a possible sensitivity to the changes introduces in the revised dataset in 1. We hypothesise that one of the three changes included in the revised dataset (random patches, larger model array size, and included sub-population) could be playing a role, but we defer the study to a future work since the property of the selection bias seems to depend on the details of the image simulation. After running the measurement on dataset 2 and by comparing it with 1, we saw the emergence of a bias of $\approx -8 \times 10^{-3}$ due to the sub-population of bulge-only galaxies and the full variability of bulges (n_b and r_h), now included in the simulations but not captured in the measurement³².

This additional bias would be corrected through direct calibration via realistic image simulations. However, it remains to be seen how sensitive the calibration would be to choices made about the bulges in the calibration datasets. This is the main idea behind setting up the new simulations in 3, with distributions that were scaled up and down to quantify the sensitivity. The top row of Fig. 13 shows the distributions of simulations in 2 and scaled distributions of n_b and r_h in 3. As clarified above, the bulges were always allowed to vary in the simulations, but were not fully captured by the measurement (since LENSMC fixes both parameters to their nominal values as in Table 1). These distributions were scaled up and down by $\pm 20\%$, which is about 30σ away from the mean, comfortably outside the statistical uncertainty of the mean. The following rows of Fig. 13 show the sensitivity of the corrected multiplicative and additive biases to the variation in the scaled calibration dataset. These are shown separately for selection and measurement biases. While the evidence of sensitivity of selection bias is weak due to the large error bars, we do see some sensitivity particularly of measurement multiplicative bias on r_h . This is important as knowing the distribution of bulge sizes is essential for an accurate calibration of the bias. We estimated: $dm_{1,\text{meas}}/d[\Delta n_b/n_b] \approx -7 \times 10^{-4}$,




Fig. 13. Bias calibration sensitivity. (Top row) Distributions of bulge parameters scaled up and down by $\pm 20\%$ relative to the nominal one at the centre. (Following rows) Calibrated bias using the simulations with scaled bulge parameters. Semi-transparent triangles are selection biases; circles are measurement biases.

$dm_{2,\text{meas}}/d[\Delta n_b/n_b] \approx 4 \times 10^{-6}$, $dm_{1,\text{meas}}/d[\Delta r_h/r_h] \approx -8 \times 10^{-3}$, $dm_{2,\text{meas}}/d[\Delta r_h/r_h] \approx -9 \times 10^{-3}$. For instance, that would imply a level of bias of $\approx -1 \times 10^{-4}$ for every (positive) percent variation on r_h assumed in the calibration dataset. Conversely, a bias requirement of, for example, 1×10^{-3} would imply a calibration requirement on r_h of 10%.

While the results presented in this section provides some reassurance that the bias is stable and the sensitivity of the calibration is under control, it is worth mentioning that the variability in bulge distributions included here did not fully capture the variability observed in real data, with realistic morphologies of faint galaxies potentially playing a major role in the *Euclid* analysis (Euclid Collaboration 2024c).

³² However, we did not see any degradation in runtime due to model bias, suggesting that the measurement is robust.

5. Summary and discussion

LENSMC is our advanced cosmic shear measurement method based on galaxy forward modelling and MCMC sampling that is being developed for *Euclid* and Stage-IV surveys. We discussed the key components of the measurement and how to handle real data problems robustly. We demonstrated its performance on a suite of suitably complex images, our LENS-Flagship simulations, which take the Flagship catalogue as input to produce full *Euclid*-like detector images. These images include realistic galaxy properties and clustering to $I_E < 29.5$, and stars to $I_E < 26$. We made emulations of the VIS images to include realistic pixel noise and a broadband chromatic PSF model with spatial variation across the field of view. SourceExtractor++ was run to detect objects down to $I_E < 26$, and LENS-Flagship used segmentation maps to mask out objects, if not belonging to the same group, or jointly measured all objects within the same group.

The bias can be broken down into measurement (from running the method on the detected objects) and selection (detection, catalogue matching, star-galaxy separation, and weights). From Table 2, the selection accounts for a bias of $m_{1,\text{sel}} = (12 \pm 1) \times 10^{-3}$, $m_{2,\text{sel}} = (0 \pm 1) \times 10^{-3}$, $c_{1,\text{sel}} = (-3.4 \pm 0.2) \times 10^{-4}$, and $c_{2,\text{sel}} = (-0.2 \pm 0.2) \times 10^{-4}$. The measurement instead accounts for a bias of $m_{1,\text{meas}} = (-3.6 \pm 0.2) \times 10^{-3}$, $m_{2,\text{meas}} = (-4.3 \pm 0.2) \times 10^{-3}$, $c_{1,\text{meas}} = (-1.80 \pm 0.03) \times 10^{-4}$, and $c_{2,\text{meas}} = (0.09 \pm 0.03) \times 10^{-4}$. We tested that the measurement bias would be larger by an additional multiplicative bias of -7×10^{-4} if detected objects were not measured jointly. This alleviates the need of having extremely accurate deblended segmentation maps that are usually needed for masking out detected neighbours at a close angular separation. Undetected faint objects remain buried in the noise, and we estimated their multiplicative bias to -5×10^{-3} . Therefore, the measurement bias is dominated by the faint objects and to some extent by the measurement method. We tested the sensitivity to other effects, including the star-galaxy separation or the weight definition, but we did not find any statistical significance. The leakage due to the PSF variation across the field of view was found to be limited by selection, with the measurement contribution mostly consistent with requirements. Our total m and c biases are to a large extent limited by selection and secondarily by the presence of faint objects. The detection bias might first be improved with an optimisation of detection parameters, for example, by a better choice of the SourceExtractor++ convolution filter that should match the PSF size as closely as possible. Second, as the faint objects account for the other largest source of bias after selection, we aim to study the sensitivity on the choice of that distribution by adding an ultra-faint population. Since the bias can be measured by varying the magnitude limit in the input simulations, the differential bias between a shallow and deep limit yielded the calibration coefficient required for the correction of the effect due to the faint objects, which is $m_{\text{faint}} = (-5.0 \pm 0.2) \times 10^{-3}$. Once the simulation complexity was increased after the inclusion of full variability in bulge parameters and bulge-only galaxies, we observed the emergence of a bias of $\approx -8 \times 10^{-3}$. As this is not captured in the measurement, it is usually corrected as part of the shear bias calibration. In this case, we studied the sensitivity of our calibration on assumptions made about the bulges in the calibration datasets, finding that it was modest overall (with a bias of $\approx 1 \times 10^{-4}$ for every percent variation in bulge sizes assumed in the calibration dataset). As part of this investigation, we revisited the datasets and found a change in the selection bias with a less pronounced asymmetry and smaller magnitude,

suggesting that the details of the simulation setup may play a key role. We defer further study of the selection bias to future work following a further upgrade of the simulations. We found that breaking down the bias into leading effects as shown in this work proves itself as a useful tool when deriving the calibration corrections required for real data applications.

We recognise that our simulations will need further elements of realism. First, we showed initial results for nominal simulation settings and studied the sensitivity to some key parameters. This proved essential in order to break down the bias in its main contributions, but it did not give an exhaustive answer to the full calibratability of our method. While we tested the sensitivity to some key effects, we will need to carry out a more thorough study of the sensitivity to choices made in SourceExtractor++ (e.g., detection thresholds and convolution filter) and the distribution of faint galaxies. Furthermore, we showed results for models that match those in the simulations for three parameters (see fixed ones in Table 1) and then also after relaxing assumptions made for bulges. The residual bias is still under control, and its sensitivity is weak, but more parameters will need to be varied in a follow-up work. Additionally, we assumed the same shear magnitude for all galaxies, bright and faint, in our simulations. In reality, bright galaxies (affected by model bias) have a small shear, while faint galaxies have a large shear. The implication is that we might have overestimated the model bias while underestimating faint bias. We defer investigation into this effect with more realistic shear fields to a follow-up publication. Furthermore, we did not include the study of bias due to galaxy substructure. However, as this effect is observed for bright galaxies at relatively lower redshift and with smaller shear compared to faint galaxies, we anticipate that the substructure bias will not only be calibrated with *Hubble* or *Euclid* deep fields but will also be easily distinguishable from faint bias. Having ignored any redshift dependence, we used a constant SED for all the objects, resulting in no SED variation in the PSF model and perfect knowledge of the resulting PSF. There are a number of implications: (i) Stars are bluer than galaxies, and therefore their PSF appears more compact, which leads to a more difficult star-galaxy separation in real data. (ii) Any SED mismodelling (e.g., in the slope of the spectrum) will lead to multiplicative bias due to the differing PSF size. (iii) The measured galaxy SEDs are based on ground-based photometry and are therefore noisy, which causes shear bias once averaged over many galaxies. (iv) The quantification of the shear bias due to the SED estimation depends on the particular implementation details of the SED estimation method, which is currently being finalised for the *Euclid* analysis. An accurate investigation of SED variation will be included in a follow-up paper, but the impact of SED mismodelling depends on wider aspects of the *Euclid* pipeline, which are outside the scope of this paper. Another point of future work is the inclusion of cosmic rays, which are identified as one of the main causes of concern for space-based lensing surveys. In fact, while these are usually masked out at the detector level, residuals and undetected cosmic rays may still impact the shear measurement significantly. Similarly, detector non-linearities, CTI, and BFE were not included, and while again there are already strategies to correct for those at the detector level, residuals should still be fully propagated through. A complete study of redshift-dependent biases is a further essential step, as it will also be necessary to account for the colour dependence (leading to an effective redshift dependence) of the PSF modelling. As the PSF modelling is a strong function of both colour and redshift, future realistic simulations that include tomographic binning will also have to include the full spatial and colour variability of the

galaxies. Closely related galaxy colour gradients³³ might lead to additional redshift-dependent biases (Semboloni et al. 2013; Er et al. 2018).

Furthermore, our work concentrated on weak lensing shear bias for the cosmic shear using $|g| = 0.02$. *Euclid* will also provide lensing measurements for galaxy clusters, where very massive systems feature reduced shears $|g| > 0.1$ even outside the core region (e.g., Schrabback et al. 2018). In this regime, non-linear shear responses and increased blending can affect shear calibrations at the percent level (e.g., Hernández-Martín et al. 2020). This must be accounted for in order to reach the accuracy requirements of next-generation cluster cosmology analyses (e.g., Grandis et al. 2019). We therefore additionally plan to conduct dedicated analyses of LENS MC using cluster field image simulations, anticipating that ongoing preliminary tests already demonstrate linearity of the method up to $|g| = 0.1$.

While our simulations are already up to the standard of the most recent shear measurement simulations, the goal of this paper was to set up suitably complex simulations that can prove the robustness and performance of LENS MC in real-data scenarios, so we are confident about its application to real *Euclid* data. The added benefit of our simulations is the relative flexibility in the possibility to incorporate and study new effects individually and for the shear bias to be broken down into individual effects. This is something that has not been possible with the fully fledged simulations implemented in the *Euclid* science ground segment.

We showed results that were derived without resorting to external calibration and will demonstrate the full calibratability in a separate paper investigating further realism and sensitivity in much more detail. To summarise, the key points of our measurement and why we think this is best for real data are as follows: (i) forward modelling to deal with *Euclid* image undersampling and convolution by a PSF with comparable size to the many galaxies; (ii) joint measuring object groups to correctly handle neighbour bias; (iii) masking out objects belonging to different groups; (iv) MCMC sampling of the posterior in a multi-dimensional parameter space, which provides shear weights and correct marginalisation of ellipticity over nuisance parameters and other objects in the same group.

The main findings and takeaways can be summarised as follows. When model bias, chromaticity, and selection biases are suppressed, the obtained biases are close to the *Euclid* requirement. This measurement bias is largely dominated by undetected faint galaxies in the images. The bias was also found to be stable and mostly insensitive to the many effects in the simulations, which we explored in detail. As the *Euclid* analysis will also need to correct for other artefacts in the images, because of its stability, the residual bias will be straightforward to calibrate through image simulations. Once we included the model bias in the simulations, the overall bias was found to be significant. However, since sensitivity of model bias on galaxy morphological parameters is weak, it will be straightforward to also calibrate it through the same image simulations.

Acknowledgements. GC acknowledges support provided by the United Kingdom Space Agency with grant numbers ST/X00189X/1, ST/W002655/1, ST/V001701/1, and ST/N001761/1. GC thanks the University of Oxford for support, where this work was started. GC thanks D. Bauer, S. Fayer (Imperial College London), P. Clarke, R. Currie (University of Edinburgh), and the GridPP Collaboration for support on getting access to and use of GridPP where the final

simulations were run. GC acknowledges the use of the *Euclid* clusters in Edinburgh, and thanks J. Patterson for support on the Astrophysics cluster in Oxford where the initial simulations were run. TS acknowledges support provided by the Austrian Research Promotion Agency (FFG) and the Federal Ministry of the Republic of Austria for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK) via the Austrian Space Applications Programme with grant numbers 899537 and 900565. MT acknowledges support from the German Federal Ministry for Economic Affairs and Climate Action (BMWK) provided by DLR under projects 50QE1103, 50QE2002, and 50QE2302. The Euclid Consortium acknowledges the European Space Agency and a number of agencies and institutes that have supported the development of *Euclid*, in particular the Agenzia Spaziale Italiana, the Austrian Forschungsförderungsgesellschaft funded through BMK, the Belgian Science Policy, the Canadian Euclid Consortium, the Deutsches Zentrum für Luft- und Raumfahrt, the DTU Space and the Niels Bohr Institute in Denmark, the French Centre National d'Etudes Spatiales, the Fundação para a Ciência e a Tecnologia, the Hungarian Academy of Sciences, the Ministerio de Ciencia, Innovación y Universidades, the National Aeronautics and Space Administration, the National Astronomical Observatory of Japan, the Nederlandse Onderzoekschool Voor Astronomie, the Norwegian Space Agency, the Research Council of Finland, the Romanian Space Agency, the State Secretariat for Education, Research, and Innovation (SERI) at the Swiss Space Office (SSO), and the United Kingdom Space Agency. A complete and detailed list is available on the *Euclid* web site (<https://www.euclid-ec.org>). Python (Van Rossum & Drake 2009, and also python.org); astropy (Price-Whelan et al. 2022), cython (Behnel et al. 2011), numpy (Harris et al. 2020), pyfftw (Frigo & Johnson 2005, see also github.com/pyfftw/pyfftw), and scipy (Virtanen et al. 2020) for the core measurement code; DIRAC (Bauer et al. 2015) for the submission to GridPP; dask (Rocklin 2015) and h5py (Collette 2013, see also github.com/h5py/h5py) for the final analysis. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

References

- Abbott, T. M. C., Abdalla, F. B., Alarcon, A., et al. 2018a, *Phys. Rev. D*, **98**, 043526
- Abbott, T. M. C., Abdalla, F. B., Allam, S., et al. 2018b, *ApJS*, **239**, 18
- Abbott, T. M. C., Aguena, M., Alarcon, A., et al. 2023, *Open J. Astrophys.*, **6**, 36
- Aihara, H., Armstrong, R., Bickerton, S., et al. 2017, *PASJ*, **70**, S8
- Akeson, R., Armus, L., Bachelet, E., et al. 2019, arXiv e-prints [arXiv:1902.05569]
- Amendola, L., Appleby, S., Avgoustidis, A., et al. 2018, *Living Rev. Relativ.*, **21**, 2
- Amon, A., Gruen, D., Troxel, M. A., et al. 2022a, *Phys. Rev. D*, **105**, 023514
- Amon, A., Robertson, N. C., Miyatake, H., et al. 2022b, *MNRAS*, **518**, 477
- Antilogus, P., Astier, P., Doherty, P., Guyonnet, A., & Regnault, N. 2014, *JINST*, **9**, C03048
- Arcelin, B., Doux, C., Aubourg, E., Roucelle, C., & LSST Dark Energy Science Collaboration 2021, *MNRAS*, **500**, 531
- Asgari, M., Lin, C.-A., Joachimi, B., et al. 2021, *A&A*, **645**, A104
- Bauer, D., Colling, D., Currie, R., et al. 2015, *J. Phys. Conf. Ser.*, **664**, 062036
- Beckwith, S. V. W., Stiavelli, M., Koekemoer, A. M., et al. 2006, *ApJ*, **132**, 1729
- Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *Comput. Sci. Eng.*, **13**, 31
- Bernstein, G. M. 2010, *MNRAS*, **406**, 2793
- Bernstein, G. M., & Armstrong, R. 2014, *MNRAS*, **438**, 1880
- Bernstein, G. M., Armstrong, R., Krawiec, C., & March, M. C. 2016, *MNRAS*, **459**, 4467
- Bertin, E., Mellier, Y., Radovich, M., et al. 2002, in *ASPCS*, **281**, Astronomical Data Analysis Software and Systems XI, eds. D. A. Bohlender, D. Durand, & T. H. Handley, 228
- Bertin, E., Schefer, M., Apostolakis, N., et al. 2020, *ASPCS*, **527**, 461
- Betancourt, M. 2017, arXiv e-prints [arXiv:1701.02434]
- Bosch, J., Armstrong, R., Bickerton, S., et al. 2017, *PASJ*, **70**, S5
- Bridle, S., Balan, S. T., Bethge, M., et al. 2010, *MNRAS*, **405**, 2044
- Britton, D., Cass, A. J., Clarke, P. E. L., et al. 2009, *Phil. Trans. R. Soc. A*, **367**, 2447
- Carretero, J., Tallada, P., Casals, J., et al. 2017, in *Proceedings of the European Physical Society Conference on High Energy Physics*, 488
- Collette, A. 2013, *Python and HDF5* (O'Reilly)
- Congedo, G. 2015, *Phys. Rev. D*, **91**, 062006
- Cornish, N. J., & Porter, E. K. 2007, *Phys. Rev. D*, **75**, 021301
- Cox, D. R., & Snell, E. J. 1968, *J. Roy. Stat. Soc. Ser. B (Method)*, **30**, 248
- Cropper, M., Hoekstra, H., Kitching, T., et al. 2013, *MNRAS*, **431**, 3103
- Cropper, M., Pottinger, S., Niemi, S., et al. 2016, *SPIE*, **9904**, 99040Q
- Czekaj, M. A., Robin, A. C., Figueras, F., Luri, X., & Haywood, M. 2014, *A&A*, **564**, A102
- Dalal, R., Li, X., Nicola, A., et al. 2023, *Phys. Rev. D*, **108**, 123519

³³ Bulges and discs can have different colors, and in such cases, the two components are convolved with a slightly wrong PSF in the shear measurement.

- Drlica-Wagner, A., Sevilla-Noarbe, I., Rykoff, E. S., et al. 2018, *ApJS*, **235**, 33
- Dunkley, J., Bucher, M., Ferreira, P. G., Moodley, K., & Skordis, C. 2005, *MNRAS*, **356**, 925
- Er, X., Hoekstra, H., Schrabback, T., et al. 2018, *MNRAS*, **476**, 5645
- Euclid Collaboration (Martinet, N., et al.) 2019, *A&A*, **627**, A59
- Euclid Collaboration (Blanchard, A., et al.) 2020a, *A&A*, **642**, A191
- Euclid Collaboration (Paykari, P., et al.) 2020b, *A&A*, **635**, A139
- Euclid Collaboration (Ilić, S., et al.) 2022a, *A&A*, **657**, A91
- Euclid Collaboration (Scaramella, R., et al.) 2022b, *A&A*, **662**, A112
- Euclid Collaboration (Castander, F. J., et al.) 2024a, *A&A*, submitted [arXiv:2405.13495]
- Euclid Collaboration (Cropper, M. S., et al.) 2024b, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202450996>
- Euclid Collaboration (Csizi, B., et al.) 2024c, *A&A*, submitted [arXiv:2409.07528]
- Euclid Collaboration (Mellier, Y., et al.) 2024d, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202450810>
- Euclid Collaboration (Serrano, S., et al.) 2024e, *A&A*, **690**, A103
- Faulkner, P. J. W., Lowe, L. S., Tan, C. L. A., et al. 2005, *J. Phys. G*, **32**, N1
- Fenech Conti, I., Herbonnet, R., Hoekstra, H., et al. 2017, *MNRAS*, **467**, 1627
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Frigo, M., & Johnson, S. G. 2005, *Proc. IEEE*, **93**, 216
- Gatti, M., Sheldon, E., Amon, A., et al. 2021, *MNRAS*, **504**, 4312
- Geyer, C. J. 1992, *Statist. Sci.*, **7**, 473
- Giavalisco, M., Ferguson, H. C., Koekemoer, A. M., et al. 2004, *ApJ*, **600**, L93
- Giblin, B., Heymans, C., Asgari, M., et al. 2021, *A&A*, **645**, A105
- Goodman, J., & Weare, J. 2010, *Commun. Appl. Math. Comput. Sci.*, **5**, 65
- Grandis, S., Mohr, J. J., Dietrich, J. P., et al. 2019, *MNRAS*, **488**, 2041
- Greisen, E. W., & Calabretta, M. R. 2002, *A&A*, **395**, 1061
- Guzik, J., & Bernstein, G. 2005, *Phys. Rev. D*, **72**, 043503
- Hall, A., & Taylor, A. 2017, *MNRAS*, **468**, 346
- Harris, C. R., Millman, K. J., der Walt, S. J. V., et al. 2020, *Nature*, **585**, 357
- Hastings, W. K. 1970, *Biometrika*, **57**, 97
- Hernández-Martín, B., Schrabback, T., Hoekstra, H., et al. 2020, *A&A*, **640**, A117
- Heymans, C., Van Waerbeke, L., Bacon, D., et al. 2006, *MNRAS*, **368**, 1323
- Heymans, C., Van Waerbeke, L., Miller, L., et al. 2012, *MNRAS*, **427**, 146
- Heymans, C., Tröster, T., Asgari, M., et al. 2021, *A&A*, **646**, A140
- Hikage, C., Oguri, M., Hamana, T., et al. 2019, *PASJ*, **71**, 43
- Hildebrandt, H., Viola, M., Heymans, C., et al. 2017, *MNRAS*, **465**, 1454
- Hoekstra, H. 2021, *A&A*, **656**, A135
- Hoekstra, H., Franx, M., Kuijken, K., & Squires, G. 1998, *ApJ*, **504**, 636
- Hoekstra, H., Viola, M., & Herbonnet, R. 2017, *MNRAS*, **468**, 3295
- Hoekstra, H., Kannawadi, A., & Kitching, T. D. 2021, *A&A*, **646**, A124
- Huff, E., & Mandelbaum, R. 2017, arXiv e-prints [arXiv:1702.02600]
- Huterer, D., Takada, M., Bernstein, G., & Jain, B. 2006, *MNRAS*, **366**, 101
- Israel, H., Massey, R., Prod'homme, T., et al. 2015, *MNRAS*, **453**, 561
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Jansen, H., Tewes, M., Schrabback, T., et al. 2024, *A&A*, **683**, A240
- Jarvis, M., Bernstein, G. M., Amon, A., et al. 2020, *MNRAS*, **501**, 1282
- Joachimi, B., Cacciato, M., Kitching, T. D., et al. 2015, *SSRv*, **193**, 1
- Joudaki, S., Hildebrandt, H., Traykova, D., et al. 2020, *A&A*, **638**, L1
- Kaiser, N., Squires, G., & Broadhurst, T. 1995, *ApJ*, **449**, 460
- Kannawadi, A., Hoekstra, H., Miller, L., et al. 2019, *A&A*, **624**, A92
- Karamanis, M., & Beutler, F. 2021, *Stat. Comput.*, **31**, 61
- Kilbinger, M. 2015, *Rep. Prog. Phys.*, **78**, 086901
- Kitching, T. D., & Deshpande, A. C. 2022, *Open J. Astrophys.*, **5**, 6
- Kitching, T. D., Miller, L., Heymans, C. E., Van Waerbeke, L., & Heavens, A. F. 2008, *MNRAS*, **390**, 149
- Kitching, T. D., Balan, S. T., Bridle, S., et al. 2012, *MNRAS*, **423**, 3163
- Kuijken, K., Heymans, C., Hildebrandt, H., et al. 2015, *MNRAS*, **454**, 3500
- Kümmel, M., Bertin, E., Schefer, M., et al. 2020, *ASPCS*, **527**, 29
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv e-prints [arXiv:1110.3193]
- Lawson, C. L., & Hanson, R. J. 1995, *Solving Least Squares Problems* (SIAM)
- Lemos, P., Weaverdyck, N., Rollins, R. P., et al. 2022, *MNRAS*, **521**, 1184
- Li, S.-S., Kuijken, K., Hoekstra, H., et al. 2023a, *A&A*, **670**, A100
- Li, S.-S., Hoekstra, H., Kuijken, K., et al. 2023b, *A&A*, **679**, A133
- Li, X., Mandelbaum, R., Jarvis, M., et al. 2023c, *MNRAS*, **527**, 10388
- Li, X., Zhang, T., Sugiyama, S., et al. 2023d, *Phys. Rev. D*, **108**, 123518
- Loureiro, A., Whittaker, L., Spurio Mancini, A., et al. 2022, *A&A*, **665**, A56
- MacCrann, N., Becker, M. R., McCullough, J., et al. 2021, *MNRAS*, **509**, 3371
- MacKay, D. J. C. 2002, *Information Theory, Inference & Learning Algorithms* (New York, NY, USA: Cambridge University Press)
- Mandelbaum, R. 2018, *ARA&A*, **56**, 393
- Mandelbaum, R., Rowe, B., Armstrong, R., et al. 2015, *MNRAS*, **450**, 2963
- Mandelbaum, R., Eifler, T., Hložek, R., et al. 2018, arXiv e-prints [arXiv:1809.01669]
- Massey, R., Heymans, C., Bergé, J., et al. 2007, *MNRAS*, **376**, 13
- Massey, R., Hoekstra, H., Kitching, T., et al. 2012, *MNRAS*, **429**, 661
- Massey, R., Schrabback, T., Cordes, O., et al. 2014, *MNRAS*, **439**, 887
- Melchior, P., & Viola, M. 2012, *MNRAS*, **424**, 2757
- Melchior, P., Joseph, R., Sanchez, J., MacCrann, N., & Gruen, D. 2021, *Nat. Rev. Phys.*, **3**, 712
- Miller, L., Kitching, T. D., Heymans, C., Heavens, A. F., & Van Waerbeke, L. 2007, *MNRAS*, **382**, 315
- Miller, L., Heymans, C., Kitching, T. D., et al. 2013, *MNRAS*, **429**, 2858
- Nelder, J. A., & Mead, R. 1965, *CompJ*, **7**, 308
- Nourbakhsh, E., Tyson, J. A., Schmidt, S. J., & The LSST Dark Energy Science Collaboration 2022, *MNRAS*, **514**, 5905
- Peng, C. Y., Ho, L. C., Impey, C. D., & Rix, H.-W. 2002, *ApJ*, **124**, 266
- Planck Collaboration VI. 2020, *A&A*, **641**, A6
- Potter, D., Stadel, J., & Teyssier, R. 2017, *ComAC*, **4**, 2
- Powell, M. J. D. 1964, *CompJ*, **7**, 155
- Price-Whelan, A. M., Lim, P. L., Earl, N., et al. 2022, *ApJ*, **935**, 167
- Pujol, A., Kilbinger, M., Sureau, F., & Bobin, J. 2019, *A&A*, **621**, A2
- Refregier, A., Kacprzak, T., Amara, A., Bridle, S., & Rowe, B. 2012, *MNRAS*, **425**, 1951
- Rhodes, J., Leauthaud, A., Stoughton, C., et al. 2010, *PASP*, **122**, 439
- Rocklin, M. 2015, *Proc. of the 14th Python in Science Conference*, 130
- Rowe, B., Jarvis, M., Mandelbaum, R., et al. 2015, *Astron. Comput.*, **10**, 121
- Sambridge, M. 2013, *Geophys. J. Int.*, **196**, 357
- Samuroff, S., Bridle, S. L., Zuntz, J., et al. 2017, *MNRAS*, **475**, 4524
- Schneider, P. 2006, in *Gravitational Lensing: Strong, Weak and Micro*, eds. Meylan, G., Jetzer, P., & North, P. (Springer Berlin Heidelberg), 269
- Schrabback, T., Schirmer, M., Van der Burg, R. F. J., et al. 2018, *A&A*, **610**, A85
- Secco, L. F., Samuroff, S., Krause, E., et al. 2022, *Phys. Rev. D*, **105**, 023515
- Seitz, C., & Schneider, P. 1997, *A&A*, **318**, 687
- Semboloni, E., Hoekstra, H., Huang, Z., et al. 2013, *MNRAS*, **432**, 2385
- Sérsic, J. L. 1963, *Bol. AAA*, **6**, 41
- Sevilla-Noarbe, I., Hoyle, B., Marchă, M. J., et al. 2018, *MNRAS*, **481**, 5451
- Sheldon, E. S. 2014, *MNRAS*, **444**, L25
- Sheldon, E. S., & Huff, E. M. 2017, *ApJ*, **841**, 24
- Sheldon, E. S., Becker, M. R., MacCrann, N., & Jarvis, M. 2020, *ApJ*, **902**, 138
- Simon, P., & Schneider, P. 2017, *A&A*, **604**, A109
- Skilling, J. 2006, *Bayesian Anal.*, **1**, 833
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv e-prints [arXiv:1503.03757]
- Swendsen, R. H., & Wang, J.-S. 1986, *Phys. Rev. Lett.*, **57**, 2607
- Tallada, P., Carretero, J., Casals, J., et al. 2020, *Astron. Comput.*, **32**, 100391
- Taylor, A. N., & Kitching, T. D. 2010, *MNRAS*, **408**, 865
- Tewes, M., Kuntzer, T., Nakajima, R., et al. 2019, *A&A*, **621**, A36
- Van der Kruit, P. C., & Searle, L. 1981, *A&A*, **95**, 105
- Van der Kruit, P. C., & Searle, L. 1982, *A&A*, **110**, 61
- Van Rossum, G., & Drake, F. L. 2009, *Python 3 Reference Manual* (Scotts Valley, CA: CreateSpace)
- Viola, M., Kitching, T. D., & Joachimi, B. 2014, *MNRAS*, **439**, 1909
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, **17**, 261
- Weaver, J. R., Zalesky, L., Kokorev, V., et al. 2023, *ApJS*, **269**, 20
- Wraith, D., Kilbinger, M., Benabed, K., et al. 2009, *Phys. Rev. D*, **80**, 023507
- Zuntz, J., Kacprzak, T., Voigt, L., et al. 2013, *MNRAS*, **434**, 1604

¹ Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK

² Department of Physics, Oxford University, Keble Road, Oxford OX1 3RH, UK

³ Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK

⁴ Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK

⁵ Aix-Marseille Université, CNRS, CNES, LAM, Marseille, France

⁶ Universität Innsbruck, Institut für Astro- und Teilchenphysik, Technikerstr. 25/8, 6020 Innsbruck, Austria

⁷ Universität Bonn, Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany

⁸ Université Paris-Saclay, CNRS, Institut d'astrophysique spatiale, 91405 Orsay, France

⁹ Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK

- ¹⁰ INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy
- ¹¹ INAF-Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Via Piero Gobetti 93/3, 40129 Bologna, Italy
- ¹² Dipartimento di Fisica e Astronomia, Università di Bologna, Via Gobetti 93/2, 40129 Bologna, Italy
- ¹³ INFN-Sezione di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
- ¹⁴ Max Planck Institute for Extraterrestrial Physics, Giessenbachstr. 1, 85748 Garching, Germany
- ¹⁵ Universitäts-Sternwarte München, Fakultät für Physik, Ludwig-Maximilians-Universität München, Scheinerstrasse 1, 81679 München, Germany
- ¹⁶ INAF – Osservatorio Astrofisico di Torino, Via Osservatorio 20, 10025 Pino Torinese (TO), Italy
- ¹⁷ Dipartimento di Fisica, Università di Genova, Via Dodecaneso 33, 16146, Genova, Italy
- ¹⁸ INFN-Sezione di Genova, Via Dodecaneso 33, 16146 Genova, Italy
- ¹⁹ Department of Physics “E. Pancini”, University Federico II, Via Cinthia 6, 80126 Napoli, Italy
- ²⁰ INAF – Osservatorio Astronomico di Capodimonte, Via Moiarillo 16, 80131 Napoli, Italy
- ²¹ INFN section of Naples, Via Cinthia 6, 80126 Napoli, Italy
- ²² Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Rua das Estrelas, 4150-762 Porto, Portugal
- ²³ Dipartimento di Fisica, Università degli Studi di Torino, Via P. Giuria 1, 10125 Torino, Italy
- ²⁴ INFN-Sezione di Torino, Via P. Giuria 1, 10125 Torino, Italy
- ²⁵ INAF – IASF Milano, Via Alfonso Corti 12, 20133 Milano, Italy
- ²⁶ INAF-Osservatorio Astronomico di Roma, Via Frascati 33, 00078 Monteporzio Catone, Italy
- ²⁷ INFN – Sezione di Roma, Piazzale Aldo Moro 2, c/o Dipartimento di Fisica, Edificio G. Marconi, 00185 Roma, Italy
- ²⁸ Institut de Física d’Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain
- ²⁹ Port d’Informació Científica, Campus UAB, C. Albareda s/n, 08193 Bellaterra (Barcelona), Spain
- ³⁰ Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, 52056 Aachen, Germany
- ³¹ Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain
- ³² Institut d’Estudis Espacials de Catalunya (IEEC), Edifici RDIT, Campus UPC, 08860 Castelldefels, Barcelona, Spain
- ³³ Dipartimento di Fisica e Astronomia “Augusto Righi” – Alma Mater Studiorum Università di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
- ³⁴ European Space Agency/ESRIN, Largo Galileo Galilei 1, 00044 Frascati, Roma, Italy
- ³⁵ ESAC/ESA, Camino Bajo del Castillo, s/n., Urb. Villafranca del Castillo, 28692 Villanueva de la Cañada, Madrid, Spain
- ³⁶ Université Claude Bernard Lyon 1, CNRS/IN2P3, IP2I Lyon, UMR 5822, Villeurbanne, 69100 France
- ³⁷ Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, 1290 Versoix, Switzerland
- ³⁸ UCB Lyon 1, CNRS/IN2P3, IUF, IP2I Lyon, 4 rue Enrico Fermi, 69622 Villeurbanne, France
- ³⁹ Departamento de Física, Faculdade de Ciências, Universidade de Lisboa, Edifício C8, Campo Grande, 1749-016 Lisboa, Portugal
- ⁴⁰ Instituto de Astrofísica e Ciências do Espaço, Faculdade de Ciências, Universidade de Lisboa, Campo Grande, 1749-016 Lisboa, Portugal
- ⁴¹ Department of Astronomy, University of Geneva, ch. d’Ecogia 16, 1290 Versoix, Switzerland
- ⁴² INAF-Istituto di Astrofisica e Planetologia Spaziali, via del Fosso del Cavaliere, 100, 00100 Roma, Italy
- ⁴³ Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France
- ⁴⁴ Institut de Ciències de l’Espai (IEEC-CSIC), Campus UAB, Carrer de Can Magrans, s/n Cerdanyola del Vallès, 08193 Barcelona, Spain
- ⁴⁵ INAF – Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, 34143 Trieste, Italy
- ⁴⁶ Istituto Nazionale di Fisica Nucleare, Sezione di Bologna, Via Irnerio 46, 40126 Bologna, Italy
- ⁴⁷ INAF – Osservatorio Astronomico di Padova, Via dell’Osservatorio 5, 35122 Padova, Italy
- ⁴⁸ Institute of Theoretical Astrophysics, University of Oslo, PO Box 1029 Blindern, 0315 Oslo, Norway
- ⁴⁹ Higgs Centre for Theoretical Physics, School of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3FD, UK
- ⁵⁰ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
- ⁵¹ von Hoerner & Sulger GmbH, Schlossplatz 8, 68723 Schwetzingen, Germany
- ⁵² Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark
- ⁵³ Cosmic Dawn Center (DAWN), Denmark
- ⁵⁴ Institut d’Astrophysique de Paris, UMR 7095, CNRS, and Sorbonne Université, 98 bis boulevard Arago, 75014 Paris, France
- ⁵⁵ Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
- ⁵⁶ Department of Physics and Helsinki Institute of Physics, Gustaf Hållströmin katu 2, 00014 University of Helsinki, Finland
- ⁵⁷ Aix-Marseille Université, CNRS/IN2P3, CPPM, Marseille, France
- ⁵⁸ AIM, CEA, CNRS, Université Paris-Saclay, Université de Paris, 91191 Gif-sur-Yvette, France
- ⁵⁹ Leiden Observatory, Leiden University, Einsteinweg 55, 2333 CC Leiden, The Netherlands
- ⁶⁰ Université de Genève, Département de Physique Théorique and Centre for Astroparticle Physics, 24 quai Ernest-Ansermet, 1211 Genève 4, Switzerland
- ⁶¹ Department of Physics, PO Box 64, 00014 University of Helsinki, Finland
- ⁶² Helsinki Institute of Physics, Gustaf Hållströmin katu 2, University of Helsinki, Helsinki, Finland
- ⁶³ NOVA optical infrared instrumentation group at ASTRON, Oude Hoogeveensedijk 4, 7991PD, Dwingeloo, The Netherlands
- ⁶⁴ Dipartimento di Fisica “Aldo Pontremoli”, Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy
- ⁶⁵ INFN – Sezione di Milano, Via Celoria 16, 20133 Milano, Italy
- ⁶⁶ Dipartimento di Fisica e Astronomia “Augusto Righi” – Alma Mater Studiorum Università di Bologna, via Piero Gobetti 93/2, 40129 Bologna, Italy
- ⁶⁷ Department of Physics, Institute for Computational Cosmology, Durham University, South Road, DH1 3LE, UK
- ⁶⁸ Université Côte d’Azur, Observatoire de la Côte d’Azur, CNRS, Laboratoire Lagrange, Bd de l’Observatoire, CS 34229, 06304 Nice cedex 4, France
- ⁶⁹ Université Paris Cité, CNRS, Astroparticule et Cosmologie, 75013 Paris, France
- ⁷⁰ University of Applied Sciences and Arts of Northwestern Switzerland, School of Engineering, 5210 Windisch, Switzerland
- ⁷¹ European Space Agency/ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
- ⁷² School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, Newcastle-upon-Tyne NE1 7RU, UK
- ⁷³ Department of Physics and Astronomy, University of Aarhus, Ny Munkegade 120, 8000 Aarhus C, Denmark
- ⁷⁴ Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada
- ⁷⁵ Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada
- ⁷⁶ Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada
- ⁷⁷ Université Paris-Saclay, Université Paris Cité, CEA, CNRS, Astrophysique, Instrumentation et Modélisation Paris-Saclay, 91191 Gif-sur-Yvette, France

- ⁷⁸ Space Science Data Center, Italian Space Agency, via del Politecnico snc, 00133 Roma, Italy
- ⁷⁹ Centre National d'Etudes Spatiales – Centre spatial de Toulouse, 18 avenue Edouard Belin, 31401 Toulouse Cedex 9, France
- ⁸⁰ Institute of Space Science, Str. Atomistilor, nr. 409 Măgurele, Ilfov 077125, Romania
- ⁸¹ Instituto de Astrofísica de Canarias, Calle Vía Láctea s/n, 38204, San Cristóbal de La Laguna, Tenerife, Spain
- ⁸² Departamento de Astrofísica, Universidad de La Laguna, 38206, La Laguna, Tenerife, Spain
- ⁸³ Dipartimento di Fisica e Astronomia “G. Galilei”, Università di Padova, Via Marzolo 8, 35131 Padova, Italy
- ⁸⁴ INFN – Padova, Via Marzolo 8, 35131 Padova, Italy
- ⁸⁵ Departamento de Física, FCFM, Universidad de Chile, Blanco Encalada 2008, Santiago, Chile
- ⁸⁶ Satlantís, University Science Park, Sede Bld 48940, Leioa-Bilbao, Spain
- ⁸⁷ Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Avenida Complutense 40, 28040 Madrid, Spain
- ⁸⁸ Instituto de Astrofísica e Ciências do Espaço, Faculdade de Ciências, Universidade de Lisboa, Tapada da Ajuda, 1349-018 Lisboa, Portugal
- ⁸⁹ Universidad Politécnica de Cartagena, Departamento de Electrónica y Tecnología de Computadoras, Plaza del Hospital 1, 30202 Cartagena, Spain
- ⁹⁰ Institut de Recherche en Astrophysique et Planétologie (IRAP), Université de Toulouse, CNRS, UPS, CNES, 14 Av. Edouard Belin, 31400 Toulouse, France
- ⁹¹ Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
- ⁹² INFN – Bologna, Via Irnerio 46, 40126 Bologna, Italy
- ⁹³ Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA
- ⁹⁴ IFPU, Institute for Fundamental Physics of the Universe, via Beirut 2, 34151 Trieste, Italy
- ⁹⁵ INAF, Istituto di Radioastronomia, Via Piero Gobetti 101, 40129 Bologna, Italy
- ⁹⁶ Department of Mathematics and Physics E. De Giorgi, University of Salento, Via per Arnesano, CP-193, 73100 Lecce, Italy
- ⁹⁷ INAF-Sezione di Lecce, c/o Dipartimento Matematica e Fisica, Via per Arnesano, 73100, Lecce, Italy
- ⁹⁸ INFN, Sezione di Lecce, Via per Arnesano, CP-193, 73100 Lecce, Italy
- ⁹⁹ Institut für Theoretische Physik, University of Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany
- ¹⁰⁰ Université St Joseph; Faculty of Sciences, Beirut, Lebanon
- ¹⁰¹ Institut d'Astrophysique de Paris, 98bis Boulevard Arago, 75014 Paris, France
- ¹⁰² Junia, EPA department, 41 Bd Vauban, 59800 Lille, France
- ¹⁰³ SISSA, International School for Advanced Studies, Via Bonomea 265, 34136 Trieste TS, Italy
- ¹⁰⁴ INFN, Sezione di Trieste, Via Valerio 2, 34127 Trieste TS, Italy
- ¹⁰⁵ ICSC – Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing, Via Magnanelli 2, Bologna, Italy
- ¹⁰⁶ Instituto de Física Teórica UAM-CSIC, Campus de Cantoblanco, 28049 Madrid, Spain
- ¹⁰⁷ CERCA/ISO, Department of Physics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, USA
- ¹⁰⁸ Laboratoire Univers et Théorie, Observatoire de Paris, Université PSL, Université Paris Cité, CNRS, 92190 Meudon, France
- ¹⁰⁹ Dipartimento di Fisica e Scienze della Terra, Università degli Studi di Ferrara, Via Giuseppe Saragat 1, 44122 Ferrara, Italy
- ¹¹⁰ Istituto Nazionale di Fisica Nucleare, Sezione di Ferrara, Via Giuseppe Saragat 1, 44122 Ferrara, Italy
- ¹¹¹ Dipartimento di Fisica – Sezione di Astronomia, Università di Trieste, Via Tiepolo 11, 34131 Trieste, Italy
- ¹¹² NASA Ames Research Center, Moffett Field, CA 94035, USA
- ¹¹³ Kavli Institute for Particle Astrophysics & Cosmology (KIPAC), Stanford University, Stanford, CA 94305, USA
- ¹¹⁴ Bay Area Environmental Research Institute, Moffett Field, CA 94035, USA
- ¹¹⁵ Minnesota Institute for Astrophysics, University of Minnesota, 116 Church St SE, Minneapolis, MN 55455, USA
- ¹¹⁶ Institute Lorentz, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands
- ¹¹⁷ Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA
- ¹¹⁸ Department of Physics & Astronomy, University of California Irvine, Irvine CA 92697, USA
- ¹¹⁹ Department of Astronomy & Physics and Institute for Computational Astrophysics, Saint Mary's University, 923 Robie Street, Halifax, Nova Scotia B3H 3C3, Canada
- ¹²⁰ Departamento Física Aplicada, Universidad Politécnica de Cartagena, Campus Muralla del Mar, 30202 Cartagena, Murcia, Spain
- ¹²¹ Department of Computer Science, Aalto University, PO Box 15400, Espoo, 00 076, Finland
- ¹²² Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute (AIRUB), German Centre for Cosmological Lensing (GCCL), 44780 Bochum, Germany
- ¹²³ Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France
- ¹²⁴ Univ. Grenoble Alpes, CNRS, Grenoble INP, LPSC-IN2P3, 53, Avenue des Martyrs, 38000 Grenoble, France
- ¹²⁵ Department of Physics and Astronomy, Vesilinnantie 5, 20014 University of Turku, Finland
- ¹²⁶ Serco for European Space Agency (ESA), Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain
- ¹²⁷ ARC Centre of Excellence for Dark Matter Particle Physics, Melbourne, Australia
- ¹²⁸ Centre for Astrophysics & Supercomputing, Swinburne University of Technology, Victoria 3122, Australia
- ¹²⁹ W.M. Keck Observatory, 65-1120 Mamalahoa Hwy, Kamuela, HI, USA
- ¹³⁰ Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town, 7535, South Africa
- ¹³¹ Dipartimento di Fisica, Sapienza Università di Roma, Piazzale Aldo Moro 2, 00185 Roma, Italy
- ¹³² Centro de Astrofísica da Universidade do Porto, Rua das Estrelas, 4150-762 Porto, Portugal
- ¹³³ Zentrum für Astronomie, Universität Heidelberg, Philosophenweg 12, 69120 Heidelberg, Germany
- ¹³⁴ Dipartimento di Fisica, Università di Roma Tor Vergata, Via della Ricerca Scientifica 1, Roma, Italy
- ¹³⁵ INFN, Sezione di Roma 2, Via della Ricerca Scientifica 1, Roma, Italy
- ¹³⁶ Department of Astrophysics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland
- ¹³⁷ Department of Astrophysical Sciences, Peyton Hall, Princeton University, Princeton, NJ 08544, USA
- ¹³⁸ Niels Bohr Institute, University of Copenhagen, Jagtvej 128, 2200 Copenhagen, Denmark

Appendix A: Template bank

The Fourier model of Eq (14), $\tilde{I}(k)$, was precomputed and stored as a template bank in a cached array. Subsequently, these templates were interpolated at the new coordinates, k' . There is a residual interpolation error in the convolved model at a large radius from the centre, but is around the same level of the precision used to store the model itself. It is worth noting that, contrarily to the analytic approach adopted by *galsim*, our template models are numerical arrays obtained by a Fourier transform of the isotropic model arrays. The main reason for this approach is that the theoretical definition of Hankel transform of Eq. (14) is invalidated by the finite limit of integration and the oversampling of the model, which make our models a bit more realistic.

We further optimised the calculation of the models for speed by drawing on square images of size proportional to the galaxy size being rendered, so larger galaxies require larger arrays. Given a cut-off at $r_{\max} = 4.5 r_e$ from the centre, a minimum image half-size of $2 r_{\max}$ was required to avoid aliasing from the Fourier transform. Therefore the minimum image size will always be larger than $18 r_e$. As galaxies are expected to have mean size of $0''.3$, but spanning the whole range from just below resolution, $0''.1$, to the largest (although rare) galaxies, $1''$, we defined a template bank of pre-calculated Hankel transforms of circular objects of different scale lengths $\bar{r}_0 = \{0.2, 0.4, 0.8, 1.6, 3.2\}''$ and model arrays of different sizes $\{3.2, 4.8, 6.4, 9.6, 12.8, 19.2, 25.6, 38.4\}''$. To avoid aliasing from the interpolation, the template scale length \bar{r}_0 is required to be slightly larger than the galaxy size being rendered; we required $\bar{r}_0 > \sqrt{2} r_e$.

Appendix B: Likelihood

We wish to linearly marginalise the likelihood in Eq. (17) over flux parameters, ϕ_i . Taking the derivative and equating to zero, we have

$$\frac{\partial}{\partial \phi_i} \ln p(\mathbf{D}|\epsilon, \theta, \phi) = \mathbf{I}_i(\epsilon, \theta)^\top \mathbf{C}^{-1} [\mathbf{D} - \phi_j \mathbf{I}_j(\epsilon, \theta)] = 0, \quad (\text{B.1})$$

where $\mathbf{I}_i = \partial \mathbf{I} / \partial \phi_i$. For a linear model, marginalising is equivalent to solving the equation in a least-square sense. The Fisher matrix of the problem is defined by the derivatives of the likelihood and is given by Eq. (19). The least-square solution is therefore

$$\hat{\phi}_i(\epsilon, \theta) = \mathcal{F}_{ii}^{-1}(\epsilon, \theta) \rho_j(\epsilon, \theta). \quad (\text{B.2})$$

where $\rho_j(\epsilon, \theta) = \mathbf{D}^\top \mathbf{C}^{-1} \mathbf{I}_j(\epsilon, \theta)$. The marginal likelihood is found by plugging the solution above into the original likelihood of Eq. (17). This partial marginalisation technique is customary in many fields including cosmology (Taylor & Kitching 2010), and gravitational wave analysis (Congedo 2015).

The problem is invertible whenever the Fisher matrix is full rank, $|\mathcal{F}_{ij}(\epsilon, \theta)| > 0$. The following two conditions must be satisfied: (i) bulge and disc components must not be the same to avoid degeneracy built in the modelling; (ii) both components must not go to zero. If we were to naively apply the linear solution of Eq. (B.2) to very faint galaxies, we would often get negative fluxes. This is undesirable; hence we adopted the non-negative least-square implementation of Lawson & Hanson (1995). Effectively this is equivalent to the standard least square while enforcing a hard constraint on fluxes, $\phi_i > 0$. Whenever one of the two fluxes is zero, we make the likelihood collapse to single




Fig. C.1. Distribution of shifts of ellipticity components from the truth marginalised over nuisance for a variety of galaxies. The $\epsilon_{i,t}$ is a chain for an ellipticity component with $i = 1, 2$ and MCMC sample t ; ϵ_i is the corresponding true value. The distribution peaks at zero but shows large random values that are usually washed out by taking averages over large samples.

component modelling. The Fisher matrix of the i -th component is now a scalar,

$$\mathcal{F}_{ii}(\epsilon, \theta) = \mathbf{I}_i(\epsilon, \theta)^\top \mathbf{C}^{-1} \mathbf{I}_i(\epsilon, \theta). \quad (\text{B.3})$$

The linear solution is

$$\hat{\phi}_i(\epsilon, \theta) = \frac{\rho_i(\epsilon, \theta)}{\mathcal{F}_{ii}(\epsilon, \theta)}, \quad (\text{B.4})$$

and the marginalised likelihood is given by

$$\ln p(\mathbf{D}|\epsilon, \theta) = \frac{1}{2} \frac{\rho_i^2(\epsilon, \theta)}{\mathcal{F}_{ii}(\epsilon, \theta)} + \text{const}. \quad (\text{B.5})$$

We note that this is just a particular case of the more general Eq. (18), which is valid for more than one model component.

Appendix C: MCMC convergence

We quantified the average convergence property by calculating the shifts of every MCMC sample from the truth. Figure C.1 shows the distribution of ellipticity component shifts, marginalised over nuisance, for a variety of galaxies of broad morphological properties as described in Sect. 3. This distribution peaks at zero but also shows some large random shifts from the truth. Any such large shift is usually washed out by taking averages over large samples, as it is the case for cosmic shear analyses.

A complementary test of convergence is through the analysis of the autocorrelation function.³⁴ Suppose we have a chain for a generic parameter, θ_t , with mean $\bar{\theta}$ and standard deviation $\hat{\sigma}_{\bar{\theta}}$. The sample autocorrelation function is defined as

$$R_k = \frac{1}{\sigma_{\bar{\theta}}^2 (N_{\text{MC}} - k)} \sum_t (\theta_{t+k} - \bar{\theta})(\theta_t - \bar{\theta}), \quad (\text{C.1})$$

³⁴ Alternatively, for longer chains, the power spectrum may be better suited (Dunkley et al. 2005).




Fig. C.2. (Top) Sample autocorrelation function of the same ellipticity chains as shown in Fig. C.1. The function dies off rapidly approximately as an exponential decay (shown as the analytic curve without error bands) with the same autocorrelation time. A small negative correlation at large lags is an indicator of fast convergence. (Bottom) Distribution of effective sample size for the same chains, whose mean is at about 14 (compare with $N_{\text{MC}} = 200$).

where $t = 0, \dots, N_{\text{MC}} - k - 1$ and N_{MC} is the chain length. The autocorrelation time quantifies how long the samples will be correlated for, and is given by

$$\tau = 1 + 2 \sum_k R_k, \quad (\text{C.2})$$

for $k = 1, \dots, M$ and M is a cut-off point set by $R_{k-1} + R_k < 0$ (Geyer 1992). This truncation is usually required to avoid the inclusion of too much noise at large lags. The top of Fig. C.2 shows the mean autocorrelation function of the same ellipticity chains. The $1\text{-}\sigma$ error bars are due to the random variation in the sample. At small lags, the mean autocorrelation function approximately follows an $\exp(-2k/\tau)$ trend, and becomes slightly negative at large lags. Positive autocorrelation at large lags is an unwanted feature of any MCMC method as this suggests poor convergence. In contrast, a small negative autocorrelation as shown here suggests faster convergence. This can be seen from, for example, the estimator of Eq. (C.2): if $R_k < 0$ consistently for some $k > 0$, the final estimate of τ would be reduced compared to the case of positive autocorrelation. However, as discussed above, at large lags,




Fig. D.1. Distribution of χ^2/ν values. A shift of about 0.15 as well as a slightly positive skewness can be seen in the distribution.

the impact of noise on the estimator will be important. To account for autocorrelation in the chain, the sample variance needs to be rescaled by τ , so it is worth introducing the effective sample size as follows,

$$N_{\text{eff}} = \frac{N_{\text{MC}}}{\tau}. \quad (\text{C.3})$$

The distribution of this quantity is shown at the bottom of Fig. C.2, and its mean is at about 14 (compare with $N_{\text{MC}} = 200$) with a positive skewness towards the larger values. In this context, a large effective sample size is a good indicator of the quality of the chains. In fact, since the MC noise scales as $N_{\text{eff}}^{-1/2}$, a galaxy with typical measurement ellipticity noise of 0.3 will be affected, on average, by a sampling noise of 0.08. Both the intrinsic ellipticity dispersion and sampling noise are diluted by the average over large samples. As the averaged shear will include both noise components, the sampling noise will always be a factor of four smaller than the intrinsic dispersion. However, in general, one would expect that many more sources of noise may be present in the measurement from image detection to shear, so the sampling term would be expected to be significantly smaller than the total noise.

Appendix D: Validation

Figure D.1 shows the distribution of χ^2/ν for all galaxies in the measurement, where $\chi^2 = -2 \ln p$, p is the likelihood calculated at the mean estimate, and ν is the number of degrees of freedom. This distribution is broadly consistent with 1, except for a small shift of about 0.15 and a slightly positive skewness. There are likely two reasons for this small shift. The first one could be due to residual features in the data due to inaccurate masking and the presence of undetected faint objects. The second one, more general, is due to the fact that we calculate the χ^2 at the mean, not the maximum, and therefore a positive shift should always be expected. In our simulations, we found $\Delta\chi^2/\nu \approx 0.015$.

Figure D.2 shows the input-output ellipticity correlation for all the selected galaxies in the catalogue (see discussion about selection in Sect. 4.1), split in three magnitude bins: relatively bright, faint, and very faint. The measured ellipticity correlates very accurately with the true input value. However, as the galaxies become fainter, we observe an increase of noise and a negative bias (which is a reflection of what is noted about shear

in Sect. 4.2). We quantified this bias as the deviation from the perfect correlation line in each of the three magnitude bins: $(1 \pm 7) \times 10^{-4}$, $(-2 \pm 1) \times 10^{-3}$, and $(-13.0 \pm 0.6) \times 10^{-2}$.

Finally, Fig. D.3 shows input-output magnitude and size correlation. The measured magnitude correlates very accurately with the true input value, except at the very faint end. However, the interpretation of the size correlation is a bit more difficult. The small sizes ($r_e < 0.25$), comparable with the PSF, are biased high due to the faintness of the galaxies and the poor constraint from the posterior which does not incorporate a realistic size prior. Residual PSF errors are also expected to bias high the sizes. Additionally, a possible leakage from stars might make the situation worse (as discussed in Sect. 4.1). The issue is reversed for the large sizes ($r_e > 1.8$), which are biased low. In this case, the galaxies are very bright and their brightness profile extends to very large radii, with a large variation in brightness from the peak to the tail of the distribution. Indeed it is generally hard to measure these objects due to under-modelling of the faint tails. Either way, the impact on lensing is negligible because small, faint galaxies tend to have systematically smaller shear weight (see Fig. 8), while large, bright galaxies are small in number and carry negligible shear signal.

Appendix E: Shear bias estimate

We wish to derive a maximum likelihood estimator for the bias model of Eq. (8). We aim to regress values for measured ellipticity, $\hat{\epsilon}$, against input shear, g , with weights w (not necessarily inverse variance). The corresponding data vectors are denoted as $\hat{\epsilon}$ and g , and the weights are assumed to be uncorrelated as a diagonal matrix w . All data vectors and matrices are of the same size N_{data} . The solution $\mu = (m, c)$ is found in least-square sense,

$$\hat{\mu} = F^{-1} [(\hat{\epsilon} - g)^T w J], \quad (\text{E.1})$$

where $J = (g, 1)$ is the Jacobian matrix of size $N_{\text{data}} \times 2$. We assume matrix multiplication throughout and diagonal weights. The Fisher matrix is given by

$$F = J^T w J, \quad (\text{E.2})$$

which leads to the variance on our estimate,

$$C_{\mu} = \frac{\chi^2}{\nu} F^{-1}, \quad (\text{E.3})$$

where χ^2/ν is the rms of the fit residuals.

The explicit solution (in data index α) is as follows

$$\hat{m} = \frac{\sum_{\alpha} w_{\alpha} \delta g_{\alpha} g_{\alpha} \sum_{\alpha} w_{\alpha} - \sum_{\alpha} w_{\alpha} \delta g_{\alpha} \sum_{\alpha} w_{\alpha} g_{\alpha}}{\delta}, \quad (\text{E.4})$$

$$\hat{c} = \frac{\sum_{\alpha} w_{\alpha} \delta g_{\alpha} \sum_{\alpha} w_{\alpha} g_{\alpha}^2 - \sum_{\alpha} w_{\alpha} \delta g_{\alpha} g_{\alpha} \sum_{\alpha} w_{\alpha} g_{\alpha}}{\delta}, \quad (\text{E.5})$$

where $\delta = \sum_{\alpha} w_{\alpha} g_{\alpha}^2 \sum_{\alpha} w_{\alpha} - (\sum_{\alpha} w_{\alpha} g_{\alpha})^2$, $\delta g_{\alpha} = \hat{\epsilon}_{\alpha} - g_{\alpha}$. The variance on m and c are given by

$$\sigma_m^2 = \frac{\chi^2}{\nu} \frac{\sum_{\alpha} w_{\alpha}}{\delta}, \quad (\text{E.6})$$

$$\sigma_c^2 = \frac{\chi^2}{\nu} \frac{\sum_{\alpha} w_{\alpha} g_{\alpha}^2}{\delta}, \quad (\text{E.7})$$

where $\chi^2 = \sum_{\alpha} w_{\alpha} (\delta g_{\alpha} - \hat{m} g_{\alpha} - \hat{c})^2$ and $\nu = N_{\text{data}} - 2$.




Fig. D.2. Input-output ellipticity correlation. The correlation has been calculated for the selected sample of galaxies for a relatively bright (top), faint (middle), and very faint (bottom) magnitude bins. The measured ellipticity shows increased noise and a negative bias at the faint end, highlighted as the deviation from the perfect correlation line.




Fig. D.3. Input-output correlation for magnitude and size. (Top) Magnitude correlation showing a slightly negative bias for very faint galaxies. (Bottom) Size correlation showing a positive bias for small sizes and a negative bias for large sizes.