









This article has been accepted for publication in Monthly Notices of the Royal Astronomical Society ©: 2021 The Authors. Published by Oxford University Press on behalf of the Royal Astronomical Society. All rights reserved.

Lenses In VoicE (LIVE): searching for strong gravitational lenses in the VOICE@VST survey using convolutional neural networks

Fabrizio Gentile ^{1,2,3★}, Crescenzo Tortora ⁴, Giovanni Covone ^{3,4,5★}, Léon V. E. Koopmans,⁶
Chiara Spiniello ^{4,7}, Zuhui Fan,⁸ Rui Li,⁹ Dezi Liu ⁸, Nicola R. Napolitano ^{3,9}, Mattia Vaccari ^{10,11}
and Liping Fu ¹²

¹Department of Physics and Astronomy (DIFA), University of Bologna, Via Gobetti 93/2, I-40129 Bologna, Italy

²INAF – Osservatorio di Astrofisica e Scienza dello Spazio, via Gobetti 93/3, I-40129 Bologna, Italy

³Dipartimento di Fisica ‘Ettore Pancini’, Università di Napoli Federico II, Compl. Univ. Monte S. Angelo, I-80126 Napoli, Italy

⁴INAF – Osservatorio Astronomico di Capodimonte, Salita Moiariello, 16, I-80131 Napoli, Italy

⁵INFN, Sezione di Napoli, C.U. Monte S. Angelo, Via Cinthia, I-80126 Napoli, Italy

⁶Kapteyn Astronomical Institute, University of Groningen, PO Box 800, NL-9700 AV Groningen, The Netherlands

⁷Sub-Dep. of Astrophysics, Dep. of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

⁸South-Western Institute for Astronomy Research, Yunnan University, Kunming 650500, China

⁹School of Physics and Astronomy, Sun Yat-sen University Zhuhai Campus, Daxue Road 2, 519082-Tangjia, Zhuhai, Guangdong, China

¹⁰Inter-University Institute for Data Intensive Astronomy, Department of Physics and Astronomy, University of the Western Cape, Robert Sobukwe Road, 7535 Bellville, Cape Town, South Africa

¹¹INAF – Istituto di Radioastronomia, via Gobetti 101, I-40129 Bologna, Italy

¹²The Shanghai Key Lab for Astrophysics, Shanghai Normal University, 100 Guilin Road, Shanghai 200234, China

Accepted 2021 November 18. Received 2021 November 18; in original form 2021 May 10

ABSTRACT

We present a sample of 16 likely strong gravitational lenses identified in the *VST Optical Imaging of the CDFS and ESI fields* (VOICE survey) using convolutional neural networks (CNNs). We train two different CNNs on composite images produced by superimposing simulated gravitational arcs on real Luminous Red Galaxies observed in VOICE. Specifically, the first CNN is trained on single-band images and more easily identifies systems with large Einstein radii, while the second one, trained on composite RGB images, is more accurate in retrieving systems with smaller Einstein radii. We apply both networks to real data from the VOICE survey, taking advantage of the high limiting magnitude (26.1 in the *r* band) and low PSF FWHM (0.8 arcsec in the *r* band) of this deep survey. We analyse $\sim 21\,200$ images with $mag_r < 21.5$, identifying 257 lens candidates. To retrieve a high-confidence sample and to assess the accuracy of our technique, nine of the authors perform a visual inspection. Roughly 75 per cent of the systems are classified as likely lenses by at least one of the authors. Finally, we assemble the LIVE sample (Lenses In VoicE) composed by the 16 systems passing the chosen grading threshold. Three of these candidates show likely lensing features when observed by the *Hubble Space Telescope*. This work represents a further confirmation of the ability of CNNs to inspect large samples of galaxies searching for gravitational lenses. These algorithms will be crucial to exploit the full scientific potential of forthcoming surveys with the *Euclid* satellite and the *Vera Rubin Observatory*.

Key words: gravitational lensing; strong – galaxies; elliptical and lenticular, cD.

1 INTRODUCTION

Gravitational lenses are astrophysical systems created when the space–time warp around foreground astrophysical objects (the lenses) deflects light rays from distant background sources. In the presence of a massive object (e.g. a galaxy or a galaxy cluster), one defines *strong gravitational lensing* that produces multiple images of a distant source (when the source is a point-like object) or gravitational arcs (when the background is an extended object, such as high-*z* galaxy), as predicted by Zwicky (1937). The main observables of strong lensing (i.e. position and shape of the lensed

images) strongly rely on two factors: the angular diameter distances involving the observer, lens, and source, and the mass distribution (baryons plus dark matter) of the lens (Schneider, Ehlers & Falco 1992; Bartelmann 2010). For these reasons, strong gravitational lensing is suitable for a wide range of astrophysical and cosmological studies, among which the estimation of the Hubble constant (see e.g. Refsdal 1964; Wong et al. 2020) and the measure of the dark matter fraction in early-type galaxies (e.g. Treu & Koopmans 2004; Covone et al. 2009; Tortora et al. 2010; Auger et al. 2010b; Spiniello et al. 2011). Strong lensing has also been used to constrain the initial mass function in early-type galaxies (e.g. Treu et al. 2010; Auger et al. 2010a; Barnabè et al. 2013; Sonnenfeld et al. 2019), to identify dark matter substructures (e.g. Mao & Schneider 1998; Dalal & Kochanek 2002; Koopmans 2005; Vegetti et al. 2014), and

* E-mail: fabrizio.gentile3@unibo.it (FG); giovanni.covone@unina.it (GC)

to constrain cosmological models (e.g. Chae 2003; Cao et al. 2012). For a more detailed review of strong lensing applications, please refer to Treu (2010) and Blandford & Narayan (1992). All these analyses, however, require large samples of observed and modelled strong lenses. Unfortunately, due to the limited cross-section, strong lensing is a rare phenomenon (Schneider, Ehlers & Falco 1992). Traditionally, visual inspection used to be the main approach to lens finding (see e.g. Le Fevre & Hammer 1988; Sygnet et al. 2010), often preceded by a spectroscopic or photometric selection of the most promising candidates (e.g. Browne et al. 2003; Bolton et al. 2006; Faure et al. 2008). However, next-generation surveys with forthcoming facilities such as the *Euclid* satellite (Laureijs et al. 2011), the *Vera Rubin Observatory* (LSST Science Collaboration et al. 2009), and the *Chinese Space Station* (Gong et al. 2019) are expected to retrieve $\sim 10^5$ strong lenses in $\sim 10^9$ observed galaxies (Collett 2015). The high number of strong lenses identified in these surveys will allow new statistical studies about the strong lensing phenomenon (see e.g. Oguri, Rusu & Falco 2014; Sonnenfeld & Cautun 2021). A complete review of the possible applications of large samples of strong lenses can be found in the White Papers provided by the LSST collaboration (LSST Science Collaboration et al. 2009) and the *Euclid* Collaboration (Bergamini et al. in preparation).

It is then clear that we need more efficient methods to analyse the large amounts of data produced by these facilities, reducing the need for visual inspection (a time-consuming procedure and prone to multiple biases). In the last years, several alternative methods have been developed. These spanned from crowd science (e.g. Marshall et al. 2016) to automated source extraction (e.g. More et al. 2012). Among these, machine learning-based algorithms appeared to be the most efficient and reliable (see, e.g. the results of the first *strong lens finding challenge*; Metcalf et al. 2019). Convolutional neural networks (CNNs; LeCun et al. 1998; LeCun, Bengio & Hinton 2015) represent a special class of these algorithms. These networks are designed to resemble animal and human visual cortex and are currently the state of the art in image recognition and classification (see e.g. Russakovsky et al. 2015). CNN-based lens-finders have already been employed to search for galaxy–galaxy strong lenses in several wide sky surveys such as the *Kilo-Degree Survey* (KiDS; Petrillo et al. 2017, 2019a,b; Li et al. 2020, 2021), the *Dark Energy Survey* (DES; Jacobs et al. 2019a,b), the *Pan-STARRS* survey (Canameras et al. 2020), the *DESI* survey (Huang et al. 2020), and the *Canada–France–Hawaii Telescope Legacy Survey* (CFHTLS; Jacobs et al. 2017).

While a large amount of work has been done in analysing wide and shallow surveys, little interest has been devoted to smaller and deeper surveys. Surveys with longer exposure times and fainter limiting magnitudes are expected to retrieve more easily lenses with faint lensing features, increasing the number of identified strong lenses per square degree (Collett 2015). The samples of systems retrieved in these deep surveys will have higher mean redshifts (for both lenses and lensed sources). This will allow several applications to be extended to higher redshifts (see e.g. Treu & Koopmans 2004; Koopmans 2005; Treu et al. 2010; Vegetti et al. 2014). Both the *Euclid* satellite and the *Vera Rubin Observatory* will have, in fact, a deep survey besides their wide surveys (LSST Science Collaboration et al. 2009; Laureijs et al. 2011). Testing machine learning techniques on data from deep surveys is therefore crucial to exploit the full scientific potential of these forthcoming facilities.

In this paper, we employ the two CNNs developed in Petrillo et al. (2017, 2019a) to search for strong gravitational lenses in the *VST Optical Imaging of the CDFS and ES1 fields* (VOICE survey; Vaccari et al. 2016). Both networks were already successfully employed to

search for gravitational lenses in the KiDS survey (Petrillo et al. 2017, 2019a,b) and in the *Fornax Deep Survey* (FDS; see the preliminary results in Cantiello et al. 2020). Applying these CNNs to a smaller but deeper survey than KiDS, as VOICE, which has a *r*-band limiting magnitude at 5σ for point-like sources of 26.1 (i.e. 1 mag deeper than KiDS; Kuijken et al. 2019), we expect to identify a larger number of lenses per square degree than in the KiDS survey (Petrillo et al. 2019b; He et al. 2020; Li et al. 2020). Furthermore, since the limiting magnitude of VOICE in the *r* band is comparable with the one expected for the *Euclid* deep survey (~ 26.4 at 10σ for extended sources in the VIS band; Laureijs et al. 2011), our results will be useful to predict the performances of machine learning algorithms like CNNs on these future observations.

This paper is organized as follows. In Section 2, we briefly introduce the VOICE survey and describe the data employed to train the CNNs and to search for strong lenses. In Section 3, we describe the two lens finding algorithms and the procedure followed to create the training set. In Section 4, we assess the performances of the CNNs by applying the networks to a validation set. In Section 5, we describe the application of the algorithms to real data from the VOICE survey. In Section 6, we present and analyse the LIVE sample (*Lenses In VoicE*), comparing its size and properties with the expected number of lenses estimated with LENSPOP (Collett 2015) and with the results found by Petrillo et al. (2019b). Finally, we summarize our conclusions in Section 7.

2 DATA FROM THE VOICE SURVEY

The VOICE survey (PIs: Giovanni Covone and Mattia Vaccari; Vaccari et al. 2016) is a deep optical survey performed with the *VLT Survey Telescope* (VST) during the INAF Guaranteed Time of Observation. The VST (Capaccioli & Schipani 2011) is a 2.6-m optical telescope located at the ESO Paranal Observatory (Chile). Its main scientific instrument is a wide-field imager called OmegaCAM (Kuijken 2011), which consists of a 32 CCDs grid, each made up of $4k \times 2k$ pixels, with a field of view of about 1 deg^2 and a pixel size of $0.214 \text{ arcsec pixel}^{-1}$. The VOICE survey, once completed, will observe in the four photometric bands *ugri* a sky area of $\sim 8 \text{ deg}^2$ evenly split between the *Chandra Deep Field South* (CDFS; Giacconi et al. 2001; Tozzi et al. 2001) and the *European ISO Field 1* (ES1; Oliver et al. 2000; Rowan-Robinson et al. 2004). Several facilities already observed these regions, collecting data in different wavelengths from radio to X-rays, providing a unique set of ancillary data for these two fields (Vaccari et al. 2016). This paper focuses on a 4.9 deg^2 area in the CDFS (RA: $3^{\text{h}}32^{\text{m}}32^{\text{s}}$, Dec.: $-27^{\circ}48'30''$) whose VST observations took place between 2011 and 2015 and are now concluded.

The survey observing strategy consists of splitting each field in four tiles of about 1 deg^2 . Each tile is observed several times (more than 100 exposures were taken for the *r*-band observations, ~ 50 for the other bands), reserving best observing conditions (lower seeing and darker moon phases) for the *r*-band imaging. Single exposure times are 360 s for the *r* and *g* bands, and 400 s for the *i* band, respectively. Since observations covered about four years, image quality is not constant throughout the exposures. The PSF FWHM spans from 0.4 to 1.5 arcsec with a median value of 0.85 arcsec. Images analysed in this work are obtained stacking selected exposures with PSF FWHM < 1.1 arcsec. The averaged PSF FWHMs in the final images are 0.8 arcsec for the *r* and *i* bands, and 0.6 arcsec for the *g* band, respectively. The total exposure time of the coadds in the *r* band spans from 11.3 to 14.2 h (Table 1). Such long exposure times allowed us to reach a 5σ limiting magnitude for

Table 1. Total exposure times (in hours) of the four VOICE-CDFS fields in the three photometric bands *gri* selecting only the best exposures with PSF FWHM < 1.1 arcsec.

	<i>g</i>	<i>r</i>	<i>i</i>
CDFS-1	2.4	12.0	6.3
CDFS-2	2.8	11.3	3.7
CDFS-3	2.3	14.2	6.0
CDFS-4	2.4	12.5	6.1
Mean seeing (arcsec)	0.6	0.8	0.8
Limiting magnitude	25.4	26.1	25.2

Note. The mean seeing and the mean limiting magnitude at 5σ for point-like sources are reported in the last rows for each band (Section 2).

point-like sources of 26.1 in the *r* band, 25.4 in the *g* band, and 25.2 in the *i* band. These deep observations were used for weak-lensing studies (Fu et al. 2018; Liu et al. 2018), while the multi-epoch imaging of the CDFS allowed variability selection of supernovae (Cappellaro et al. 2015) and AGN (De Cicco et al. 2015; Poulain et al. 2020).

In this work, we use VOICE data to search for strong gravitational lenses in the CDFS. There are two main reasons why these data are particularly suitable for this research. First, the faint limiting magnitude makes it easier to identify strong lensing features (which are generally faint). Secondly, the low value of the PSF FWHM makes it possible to resolve lenses with small values of the Einstein Radius (i.e. the typical angular separation between arcs and deflectors) of the order of the arcsecond. This kind of lenses is generally harder to identify, but is also the most common (Collett 2015). Furthermore, the CDFS will be covered by the forthcoming LSST deep survey (LSST Science Collaboration et al. 2009), providing a multi-epoch, high-resolution follow-up for our lens candidates. Finally, VOICE data are similar to the data from the *KiDS* (Kuijken et al. 2019), which have been already analysed by the CNNs employed in this work (Petrillo et al. 2017, 2019a,b). This will allow an interesting comparison (Section 6.2).

2.1 Sample selection

The full VOICE catalogue contains 736 518 detected sources. Many of these are stars and low-mass galaxies that have a negligible strong lensing cross-section (Schneider et al. 1992). Furthermore, spirals represent a small portion of the strong lenses population (~20 percent; Möller, Kitzbichler & Natarajan 2007; Oguri & Marshall 2010) and, due to their morphology, are harder to identify since the spiral arms can easily be mistaken for strong lensing features. Supplying these images to the CNN lens-finders could produce a highly contaminated candidate sample and increase network confusion during the training. We thus select galaxies with a higher probability of being strong lenses. Before operating any cut in magnitude or colour, we cross-match the VOICE catalogue with ancillary data from the VIDEO survey (Jarvis et al. 2013), obtaining photometry in the two NIR photometric bands *J* and *K*. We then remove from the VOICE catalogue all the objects with corrupted photometry in the photometric bands *griJK* and exclude stars employing the 2DPHOT SG index (La Barbera et al. 2008). This index is particularly efficient in performing the star–galaxy separation, as can be seen in Fig. 1. The full catalogue is thus reduced to 172 316 objects. Then, we assemble two subsets as follows:

(i) *Bright galaxies sample*: Lensing cross-section increases with the square of the mass of a galaxy, and therefore with luminosity




Figure 1. The $g - J$ versus $J - K$ diagram for the objects with correct photometry and SEXTRACTOR’s $MAG_AUTO_r < 21.5$. The star/galaxy separation is performed through the 2DPHOT SG index. In the highlighted area, there are the selected Luminous Red Galaxies. To generate this plot, we employ NIR photometry from the VIDEO survey (Jarvis et al. 2013, Section 2.1).

(Schneider et al. 1992). To select galaxies likely to be strong lenses, we select all the objects with Kron-like magnitude MAG_AUTO_r provided by SEXTRACTOR (Bertin & Arnouts 1996) brighter than 21.5. The final BG sample consists of 21 216 galaxies.

(ii) *Luminous red galaxies sample*: LRGs are thought to represent most of the strong lens population (Eisenstein et al. 2001; Oguri et al. 2006). These galaxies are generally selected using the criteria of Eisenstein et al. (2001). However, such selection would limit our sample to too few objects to successfully train the CNN. We thus employ a slightly modified version of the colour cut presented in Tortora et al. (2018). Starting from the BG sample ($r < 21.5$), we select galaxies in the colour range,

$$\begin{cases} g - J > 2.6 \\ J - K > 0.2 \end{cases} ,$$

shown in Fig. 1. We choose this $g - J$ threshold through a visual inspection of the sample to qualitatively assess the fraction of blue and star-forming galaxies. The chosen threshold results to be higher than the one adopted in Tortora et al. (2018). The final LRG sample consists of 3450 galaxies.

3 METHODS

In this section, we briefly introduce the CNNs employed to search for strong gravitational lenses in the VOICE survey. We describe the procedure followed to create the training set, to simulate mock gravitational lenses, and to train the CNNs.

3.1 Convolutional neural networks

Artificial neural networks (ANNs; e.g. LeCun et al. 1998) are among the most popular supervised machine learning algorithms. Their architecture reflects the natural neural networks, centre of animal (and human) learning process. ANNs look for the highly complex relationship between input data (e.g. galaxy images) and the target value (e.g. the probability of being a strong gravitational lens). According to the Universal Approximation Theorem (Hornik 1991), ANNs try to approximate this relationship applying several non-linear functions to the input data. In classic ANNs, the input data pass through different layers. Each layer is made up of multiple

neurons, each of which takes as input a vector x_i from the previous layer and returns a scalar y given by

$$y = f \left(\sum_{i=0}^N x_i \cdot w_i + b \right). \quad (1)$$

The non-linear function f is called *activation function*, w_i are free parameters called *weights*, and b is the *bias*. During the training phase, an ANN inspects labelled (i.e. pre-classified) examples and *learns* the classification scheme. Learning is achieved gradually adjusting the weights w_i and the bias b to minimize the difference between predicted and actual target value (Rumelhart, Hinton & Williams 1986). This difference is measured, for example, through a *loss function* such as binary cross-entropy (e.g. Goodfellow, Bengio & Courville 2016):

$$H = -t \log(p) - (1 - t) \log(1 - p), \quad (2)$$

where t is the target value and p is the predicted one. Convolutional neural networks (CNNs; e.g. LeCun, Bengio & Hinton 2015) are a noteworthy subclass of ANNs. These algorithms use convolutional kernels to extract features, maintaining the 2D topology of input data. Thanks to this property, CNNs are usually employed in images classification problems where they can achieve even higher accuracy than humans (He et al. 2015a; Russakovsky et al. 2015; see Metcalf et al. 2019 and Becker et al. 2021 for some astrophysical examples).

In this work, we use the two CNNs developed in Petrillo et al. (2017, 2019a) that implement a *ResNet*-like architecture¹ with four residual blocks of two convolutional layers (He et al. 2015b). Further details on the architecture of the CNNs can be found in Petrillo et al. (2019a). The first CNN (*single-band CNN* hereafter) takes as input only r -band images. We choose this photometric band because of its better image quality (see Section 2) that simplifies identifying strong lensing morphological features. The second CNN (*three-band CNN*) takes as input composite images obtained combining *gri* data through the HUMVI opensource library² (Marshall et al. 2016). The *three-band CNN*, analysing RGB images, can recognize gravitational lenses through the colour gradient between the redder deflecting galaxy and the bluer alleged gravitational arc.

Both CNNs take as input 101×101 pixels² stamps (equivalent to 20×20 arcsec²) and give as output a single value p (the *score*, hereafter) in the range $[0,1]$, related to the probability that the object in the image is a strong gravitational lens (Saerens, Latinne & Decaestecker 2002). As already done in Petrillo et al. (2017), we choose the size of the stamp to be small enough to speed up training phase, to exclude environment galaxies that could confuse the network, but large enough to include the largest Einstein radius expected for galaxy–galaxy lensing (Collett 2015). The CNNs are implemented in PYTHON 3.7 using the opensource libraries KERAS³ (Chollet et al. 2015) and TENSORFLOW⁴ (Abadi et al. 2015). Both networks minimize the binary cross-entropy (equation 2) using the ADAM optimizer (Kingma & Ba 2014).

3.2 Creating the training set

From a machine learning perspective, identifying strong gravitational lenses is a two-classes classification problem. We can success-

fully address such issue using ANNs through appropriate training. Training these algorithms requires feeding examples from the two classes (i.e. lenses and non-lenses) to the ANNs. To successfully train our CNNs (each having $\sim 10^7$ free parameters to estimate) we need a vast pre-classified training set. However, strong gravitational lensing is a rare phenomenon. Currently, the *Sloan Lens ACS Survey* (SLACS; Bolton et al. 2006) provides the largest catalogue of confirmed strong lenses comprising just 118 objects (Shu et al. 2017). Larger data bases (e.g. the MasterLens project⁵) reach up to ~ 700 lenses, but many of them still require high-resolution follow-up or spectroscopic confirmation. Furthermore, all these samples do not cover homogeneously the lensing parameter space, resulting thus unsuitable for training a CNN-based lens finder to detect all possible strong lensing configurations. With a few exceptions (e.g. Huang et al. 2020), training this kind of classifiers requires strong lensing simulations.

3.2.1 Simulating strong lenses

To simulate strong gravitational lenses, we can follow two different strategies: We can simulate both deflectors and gravitational arcs (e.g. Pourrahmani, Nayyeri & Cooray 2018; Metcalf et al. 2019) or we can simulate the arcs and superimpose them on real galaxy images (e.g. Petrillo et al. 2017; Li et al. 2020). In this work, we follow the second strategy. By doing so, we obtain realistic images (Fig. 2) without having to simulate sky and instrument noise nor the nearby environment or line-of-sight structures around the lens galaxy. We produce simulations using the software described in Chatterjee (2019). We model the mass distribution of the lens galaxies (deflectors) using a Singular Isothermal Ellipsoid model (SIE; Kormann, Schneider & Bartelmann 1994; Gavazzi et al. 2007) with external shear (Keeton, Kochanek & Seljak 1997). The parameters of the model are sampled in the range used in Petrillo et al. (2019a,b) and summarized in Table 2.

We choose a uniform sampling for the axial ratio, inclination, shear strength, and shear angle, while we employ logarithmic sampling for the Einstein radius. By doing so, we train our CNN to identify more easily systems with small Einstein radii that are generally harder to detect but also more common (Collett 2015). We also simulate background lensed galaxies: we use a Sérsic brightness profile (Sérsic 1963) with parameters sampled from the range in Table 2. Similarly, we choose uniform sampling for the axial ratio, inclination, and Sérsic index, while we employ logarithmic sampling for effective radius. To add additional structures to the matter distribution, as in Petrillo et al. (2019a), we add a Gaussian Random Field in the lens plane (Hezaveh et al. 2016) and from one to five Sérsic components in the brightness distribution of the lensed source, to crudely mimic star-forming regions (Chatterjee & Koopmans 2018). These perturbations were shown to increase the accuracy of CNN-based lensfinders (Petrillo et al. 2019a,b). Further details on our simulation strategy can be found in Petrillo et al. (2019a) and Chatterjee (2019).

3.2.2 Positive training set

For the *single-band CNN*, we produce mock strong lenses following a slightly modified version of Petrillo et al. (2019a) strategy:

(i) We randomly select deflectors from the LRG sample (see Section 2.1).

¹https://github.com/CENricoP/cnn_strong_lensing

²<https://github.com/drphilmarshall/HumVI>

³<https://keras.io/>

⁴<https://www.tensorflow.org/>

⁵masterlens.astro.utah.edu/




Figure 2. Examples of mock strong gravitational lenses simulated to train the *three-band CNN*. All images are created superimposing a simulated gravitational arc on a real LRG observed in the VOICE survey. All images have a 20-arcsec side. Further details in Section 3.2.1.

Table 2. Range of parameters used to simulate mock gravitational arcs according to Petrillo et al. (2019a).

Parameter	Range	Units
Lens (SIE)		
Einstein radius	1.0–5.0 (log)	arcsec
Axis ratio	0 to 3–1.0	–
Major-axis angle	0–180	degrees
External shear	0–0.05	–
External shear angle	0–180	degrees
Source (Sérsic)		
Effective Radius R_e	0.2–0.6 (log)	arcsec
Axis ratio	0 to 3–1.0	–
Major-axis angle	0–180	degree
Sérsic Index	0.5–5.0	–
Sérsic Blobs (1 up to 5)		
Effective radius	1–10 R_e	arcsec
Axis ratio	1.0	–
Major-axis angle	0	degrees
Sérsic Index	0.5–5.0	–

Note. We perform uniform sampling for all parameters except for Einstein radius and source effective radius that are sampled logarithmically (Section 3.2.1).

(ii) We simulate 101×101 pixels² stamps of gravitational arcs with the same pixel scale as the VST. We convolve them with an averaged PSF obtained by applying the PSFEX software (Bertin 2011) to the *r*-band VOICE tiles. Differently from Petrillo et al. (2019a), we directly simulate gravitational arcs during the training phase to increase the number of strong lensing configuration examined by the CNN.

(iii) We normalize gravitational arcs to the deflector maximum brightness multiplied by an α -factor in the range [0.02,0.3]. This factor accounts for the expected luminosity gradient between deflector and arc.

(iv) We coadd the two images, applying a square root stretching to enhance lensing features,

(v) Finally, we normalize all pixel values to the maximum brightness in the image.

We create images to supply to the *three-band CNN* through the same procedure, with a few differences:

(i) We simulate three copies of each arc, one for each photometric band. We convolve each arc with the corresponding averaged PSF.

(ii) We ‘colour’ gravitational arcs using synthetic photometries of late-type galaxies (LTG) from the COSMOS templates in the LEPHARE library (Arnouts et al. 1999). These are synthetic models used to estimate photometric redshifts of galaxies in the COSMOS fields (Ilbert et al. 2006). The full library contains 31 templates in

total, for elliptical/S0 galaxies (8 models), spirals (11 models), and star-bursting galaxies (12 models). We select photometries of LTG and star-bursting galaxies (template index > 19) and redshift them to different values of z up to $z = 3$. We employ later-type templates than Petrillo et al. (2019a) to increase the colour gradient between deflectors (i.e. LRGs) and lensed sources. This choice is shown to decrease the number of environment galaxies erroneously classified as arcs.

(iii) To homogeneously sample colour space and to account for possible errors in the photometry, we add a random term in the range [−0.1, 0.1] to LEPHARE magnitudes. We also add a colour-excess term $A_x = R_x E(B - V)$ to account for extinction. In this relation, x is the SDSS filter considered and R_x factors are taken from Yuan, Liu & Xiang (2013).

(iv) We combine the three images using the HUMVI open-source code (Marshall et al. 2016) that applies the Lupton’s algorithm (Lupton et al. 2004), performing a \sinh stretching instead of a more standard square-root one.

3.2.3 Negative training set

A good lens-finder is required to produce a pure candidate sample. We thus need to teach the CNN how to recognize and exclude contaminants. Several studies (e.g. Petrillo et al. 2017, 2019a,b; Li et al. 2020) reported how some objects (e.g. spirals, merging galaxies, polar rings) can easily confuse CNN-based classifiers because of their morphology and colour gradient. To limit such effects, we populate our negative training set (~40 per cent of the full set) with the bluest sources in the BG sample with $g - J < 2.6$ (i.e. the ones with a higher probability of being spirals or star-forming galaxies). We populate the remaining 60 per cent with other random galaxies in the BG sample (30 per cent) and LRGs from the homonymous sample (30 per cent). As highlighted by Petrillo et al. (2017, 2019a), we cannot exclude that a few real lenses are present in our negative sample, but their expected low number (less than 1 in a 1000) should not strongly affect our training.

3.3 Training phase

Once the training set has been created, we train our CNNs using the *mini-batch stochastic gradient descent* technique. Each mini-batch is made up of 64 images (32 strong gravitational lenses and 32 contaminants). Our CNNs minimize the binary cross-entropy (equation 2) using the ADAM optimizer (Kingma & Ba 2014) (see Section 3.1). We initially set ADAM’s learning rate to 10^{-2} , gradually lowering it up to 10^{-5} during the training phase to fine-tune the weights. As done in He et al. (2015b), we initialize the CNNs weights w_i following a normal distribution with $\mu = 0$ and $\sigma = 1/n$, where n is




Figure 3. Distribution of *single-band CNN* (panel a) and *three-band CNN* (panel b) scores for mock gravitational lenses and contaminants in the validation set. Further details are in Section 4.

the number of inputs of each unit. To increase the training set size and to teach the CNNs rotational, scaling, and translational invariance, we employ *data augmentation* (Simard, Steinkraus & Platt 2003). This is a common strategy in machine learning, consisting of feeding several copies of the same image to the CNNs. Each copy is

- (i) rotated by an angle between 0° and 360° ;
- (ii) translated in the horizontal and vertical direction by N pixels, with N in the range $[-4, 4]$;
- (iii) reflected on the vertical and horizontal axis with a 50 per cent probability;
- (iv) rescaled by a factor in the range $[1/1.1, 1.1]$.

We directly perform data augmentation during the training phase using the opensource PYTHON library SCIKIT-IMAGE (van der Walt et al. 2014). Comparing to other analogous experiments, we employ a more limited-size training set. To prevent overfitting, we use cross-validation on the 25 per cent of the training set, constantly monitoring validation loss and accuracy. We stop the training when the validation accuracy reaches its maximum at ~ 90 per cent. The CNN performs overall $\sim 40\,000$ weights updates, examining $\sim 10^6$ examples in total.

4 TESTING THE PERFORMANCE OF THE CNNs

Before applying the CNNs to real data searching for strong gravitational lenses, we need to assess their performances. We thus apply both CNNs to a validation set made up of mock gravitational lenses and contaminants. We produce mock lenses following the same procedure as discussed in Section 3.2.2, while we select contaminants through the same distribution as described in Section 3.2.3. Each CNN assigns a score p between 0 and 1 to all images, related to the probability of being a strong gravitational lens. We show the score distributions for mock lenses and contaminants, where the ground truth is known, in Fig. 3. An ideal classifier would assign $p = 1$ to all lenses and $p = 0$ to all contaminants. We thus need statistical indicators (i.e. ‘metrics’) to measure the difference between our lens finder and an ideal one.

4.1 Confusion matrix

A *confusion matrix* is a table containing four values: *true positive rate* (TPR), *false positive rate* (FPR), *true negative rate*

(TNR), and *false negative rate* (FNR). They are defined as follows:

$$TPR = \frac{TP}{TP + FN}, \quad (3)$$

$$TNR = \frac{TN}{TN + FP}, \quad (4)$$

$$FPR = \frac{FP}{TN + FP} = 1 - TNR, \quad (5)$$

$$FNR = \frac{FN}{FN + TP} = 1 - TPR, \quad (6)$$

where FN, FP, TP, and TN are, respectively, the number of false negatives, false positives, true positives, and true negatives. All these values are computed once a threshold value (p_{Th}) is chosen, and considering all objects with $p \geq p_{Th}$ as valid lens candidates. An ideal classifier would have $TPR = TNR = 1$ or, equivalently, $FNR = FPR = 0$ for all possible threshold values that are not exactly 0 or 1. Fig. 4 represents our CNN’s confusion matrices for different values of p_{Th} . As expected, for both networks the fraction of false positives decreases towards higher p_{Th} , while the fraction of false negatives increases. Following Petrillo et al. (2019a,b), we choose an intermediate threshold value of 0.8 to get a fair trade-off between purity (i.e. a low number of false positives) and completeness (i.e. a low number of false negatives) for the resulting candidate sample.

4.2 Receiver operating characteristic

A *receiver operating characteristic curve* (or ROC curve) is produced computing TPR and FPR for all possible threshold values and plotting them against each other. An ideal classifier would provide an ROC curve passing by the point ($TPR = 1, FPR = 0$), while a random one would produce a ROC curve lying on the bisector of the TPR–FPR plane. Fig. 5 represents our CNNs’ ROC curves. To quantitatively measure the performances of a classifier, we can compute the AUROC (*Area under the ROC curve*) that is equal to 1 for an ideal classifier and to 0.5 for an untrained one. Our *single-band CNN* produces an AUROC = 0.98, while the *three-band CNN* produces an AUROC = 0.96. Both metrics are similar to other analogous CNN-based lens finders (see, e.g. the results of the first *strong lens finding challenge*; Metcalf et al. 2019).

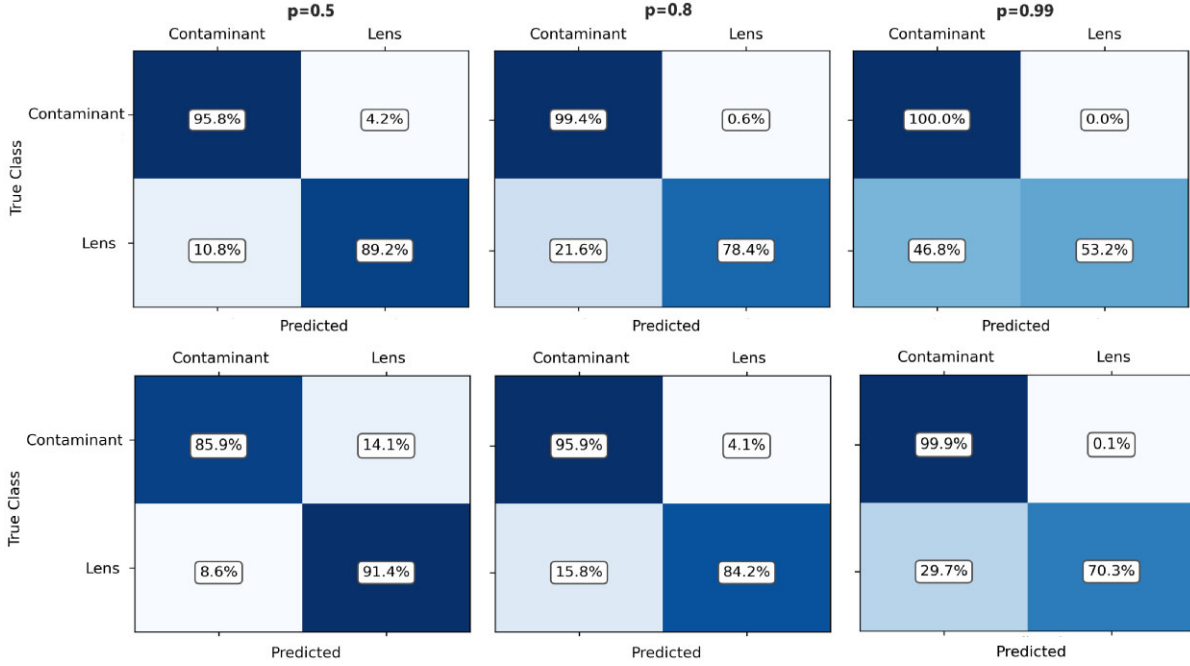


Figure 4. Confusion matrices for the *single-band CNN* (top row) and *three-band CNN* (bottom row). All the metrics are computed applying the lens-finders to the validation set and choosing different values of the threshold. Analysing the different matrices, we choose a threshold value of 0.8 (Section 4).




Figure 5. ROC curve for the two CNNs built during cross-validation. The plot is in a semi-logarithmic scale to better show low FPR values and to show the little difference between the two curves. On the plot are reported different values of the threshold value p_{Tresh} (Section 4.2).

4.3 F_β

We employ a third metric called F_β (Baeza-Yates & Ribeiro-Neto 2000). This metric, commonly employed to measure performances of classification algorithms, was also used to rank the entries in the second edition of the *strong lens finding challenge* (Metcalf et al. in preparation). It is defined as a weighted geometric average of the precision and recall of the CNN:

$$F_\beta = (1 + \beta^2) \frac{P \times R}{\beta^2 P + R}, \quad (7)$$

where

$$P = \text{precision} = \frac{TP}{TP + FP}, \quad (8)$$

$$R = \text{recall} = \frac{TP}{TP + FN} = \text{TPR}. \quad (9)$$

Varying the β -factor, we can differently weight precision and recall. Since in real data non-lenses are much more abundant than lenses, we prefer having a highly pure candidate sample rather than a highly complete one. We thus use a $\beta^2 = 0.001$, as in Metcalf et al. (in preparation). An ideal classifier would have a maximum $F_\beta = 1$. Our *single-band CNN* reached a maximum $F_\beta = 0.9994$, while the *three-band CNN* reached a maximum $F_\beta = 0.9993$. As before, these values are similar to other analogous CNN-based lens finders (Metcalf et al. in preparation).

4.4 Further Analyses

It is interesting to measure performances as a function of lens parameters such as the α -factor (which describes the brightness of the source versus the lens; see Section 3.2.2) or Einstein radius. Fig. 6 shows our results. As expected, the FNR decreases towards larger α -factors and thus towards gravitational arcs with higher brightness. It is worth noting that the two CNNs react differently to different Einstein radii. Lenses with smaller Einstein radius often have unresolved gravitational arcs. These are harder to detect using only r -band images. On the contrary, the *three-band CNN* can more easily recognize the colour gradient between the deflector and the gravitational arc, producing a lower FNR. Conversely, lenses with larger Einstein radii more easily confuse the *three-band CNN*: Distant gravitational arcs are often mistaken for blue galaxies in the lens environment. *Single-band CNN*, thanks to better image quality, can more easily detect the arc because of its morphology.

4.5 Final considerations

Analysing the different metrics, we conclude that the two CNNs are complementary (an analogous result was found by Petrillo et al.

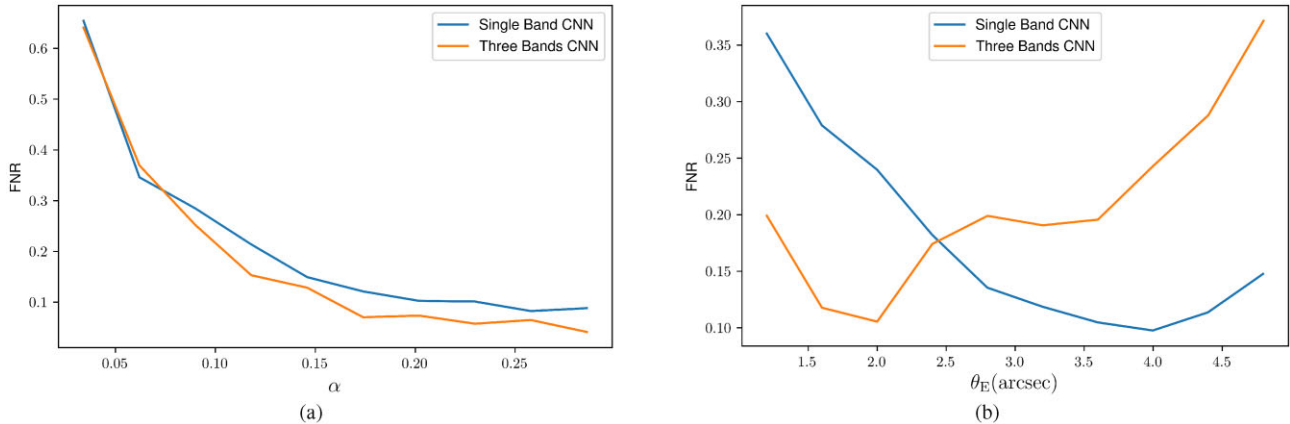


Figure 6. FNR as a function of the luminosity ratio between arc and deflector (panel a) and as a function of Einstein radius (panel b). FNRs are computed applying both CNNs to the validation set (Section 4.4).

2019a,b). In particular, the *single-band CNN* performs slightly better when global metrics (i.e. the ROC curve and the F_β index) are considered. This result can be quite counter-intuitive, since the colour information is generally useful for a human inspector in recognizing a gravitational lens. However, analysing the confusion matrices (Fig. 4), it can be seen how the lower performances of the *three-band CNN* can be explained by the larger fraction of false positives reported by the network. As discussed in the previous paragraph, blue galaxies in the environment of non-lensing objects can be easily mistaken for gravitational arcs. Never the less, the *three-band CNN* is shown to produce a more complete candidate sample for the chosen threshold score (i.e. a lower number of real lenses wrongly classified as contaminants, see Fig. 4). Finally, analysing Fig. 6(b), it can be seen that the *three-band CNN* attains lower FNR for smaller Einstein radii (which are the most common; see Collett 2015). We thus decide to use both CNNs to search for strong gravitational lenses in real data produced by the VOICE survey.

5 RESULTS

Having assessed the performances of the two CNNs, we apply both algorithms to real data from the VOICE survey. This step has a double importance. On one hand, it allows us to assemble a sample of likely strong gravitational lenses in the CDFS. On the other hand, it represents a further confirmation of the ability of the CNNs to identify strong gravitational lenses in real astronomical images. Since we trained the networks only on simulated arcs, applying the CNNs to real data helps us to exclude any possible bias in the simulation procedure.

5.1 Application to real data

Differently from analogous experiments (e.g. Petrillo et al. 2019a,b), applying the CNNs to a smaller survey, we are able to search for strong lenses in a larger fraction of observed galaxies than just in the LRG sample (Li et al. 2020). We analyse all the $\sim 21\,200$ galaxies in the Bright Galaxies Sample (see Section 2.1) passing their 101×101 pixels² stamps in the *gri* photometric bands, or only in *r*, to the two CNNs. Both algorithms give as output two values (p_1 for the *single-band CNN* and p_3 for the *three-band CNN*) in the range $[0,1]$. Choosing a threshold value $p_{\text{Th}} = 0.8$ (Section 4.1) and considering all the images with $p > p_{\text{Th}}$ as lens candidates, we assemble two samples: the *single-band lens-candidate sample* (103 systems with

$p_1 > 0.8$) and the *three-band lens-candidate sample* (161 systems with $p_3 > 0.8$), which we finally join in a combined candidate sample (*CNNs sample* hereafter). The full *CNNs sample* consists of 257 galaxies with at least one score above the chosen threshold (7 of which have both scores above the threshold), ~ 1 per cent of the full BG sample.

5.2 Visual Inspection

We do not expect the CNNs to retrieve a completely pure candidate sample (see Section 4 and Fig. 4). We expect slightly lower performances passing from the validation set (made up of simulations) to real data. To further clean the final sample from false positives, we perform a visual inspection of the images retrieved by the CNNs. Nine of the authors (the *graders* hereafter) inspected all the 257 images in the *CNNs sample*. Each grader could choose three different quality values for each image:

- (i) A : *sure lens*;
- (ii) B : *maybe lens*;
- (iii) C : *not lens*.

To combine the different rankings, as in Petrillo et al. (2017, 2019a,b), we assign a numerical value to each grading (10 to A, 4 to B, and 0 to C). This choice allows us to weight more the *sure lens* grade than the *maybe lens* when combining the rankings. Although a visual inspection is still necessary to identify false positives, it is still prone to several biases. The first is the subjectivity of the visual inspection: As in other studies (see e.g. Petrillo et al. 2017, 2019a,b), also in our sample, there are objects graded from different authors as ‘*sure lens*’ and ‘*no lens*’. To mitigate the effects of subjectivity, we involved more than one grader and choose a threshold value smaller than 90 (i.e. nine classifications as ‘*sure lens*’). In doing so, we include in the final sample also candidates without unanimous ranking as ‘*sure lens*’. The second possible bias is the inter-dependence of the graders. To mitigate this effect, all graders independently rank the images using a specially-designed grading software. This also accelerates the inspection phase and avoids any accuracy loss due to a time-consuming, tedious procedure. The results of the visual inspection are summarized in Fig. 7, where the sum of the visual scores of all the graders is considered for each candidate.

At the end of the visual inspection, 194 of the 257 images (~ 75 per cent) attain at least one classification as ‘*maybe lens*’. Among the most common objects with unanimous classification as




Figure 7. Results of the visual inspection performed on the 257 candidates accepted by the CNNs. We choose a threshold score of 36 to consider a system as a lens candidate. Further details are in Section 5.2.

‘not lens’ there are spirals, merging galaxies, polar rings, and galaxies with close companions (see some examples in Fig. 8). Although the number of false positives is low in our candidate sample, these objects could represent a problem for the future applications of the CNNs. Including a higher percentage of these objects in the negative training set could improve the quality of future trainings.

Analysing the results, we decide to accept as lens candidates objects with a total visual score ≥ 36 . This value corresponds to an unanimous classification of ‘maybe lens’, but includes in the final sample systems with some ‘not lens’ grades balanced by some ‘sure lens’ grades. We thus assemble the ‘LIVE’ sample, made up of 16 likely strong gravitational lenses. RGB stamps of the systems in the sample are shown in Fig. 9 and listed in Table 3.

6 THE LIVE SAMPLE

In this section, we analyse the gravitational lens candidates in the LIVE sample, assembled in the previous section. Among those systems, seven candidates were identified by the *single-band CNN* and ten by the *three-band CNN* (Table 3). Only one object (LIVE-1, attaining the highest score in the visual inspection) passed the score threshold for both CNNs. This result confirms the expected performances of the CNNs discussed in Section 4. The *three-band CNN* was indeed expected to retrieve a more complete candidate sample (i.e. to identify a larger fraction of real lenses), while the *single-band CNN* was expected to retrieve a purer one (i.e. with a smaller fraction of contaminants). However, we highlight that, in the final candidate sample, there are objects with one score nearly close to zero (LIVE-5 for the *single-band CNN* and LIVE-11 for the *three-band CNN*). This represents a further confirmation of the complementarity of the two algorithms (Section 4).

Among the systems in the LIVE sample, there is a previously discovered gravitational lens (LIVE-5; Blakeslee et al. 2004). It is worth noting that this object did not attain a unanimous classification as ‘sure lens’ during the visual inspection, albeit it did not obtain any ‘not lens’ classification. This confirms the possible biases in the visual inspection (Section 5.2) and exemplifies the dependence of the classification on the image quality and the signal-to-noise ratio (S/N; Fig. 11). Furthermore, it is interesting to note that this real lens has p_1 below the chosen threshold. This can be explained by the higher FNR of the *single-band CNN* at lower values of the Einstein radius (Section 4 and Fig. 6). Through research in the current literature,

we retrieve one more lens previously identified in the CDFS (DES J0329–2820; Nord et al. 2016). However, since this system has a *r*-band magnitude of 22.4, it is too faint to be part of the BG sample (Section 2.1), and thus it is not analysed by the CNNs. Never the less, we manually pass its image to the CNNs, obtaining both scores below the chosen threshold ($p_1 = 0.6$ and $p_3 = 0.5$). This result can be explained by the large value of the Einstein estimated by Nord et al. (2016) for this system ($\theta_E = 7.8 \text{ arcsec} \pm 1.4 \text{ arcsec}$; see Fig. 10). This value is well outside the range of Einstein radii on which we trained our algorithms and in which we expect the CNNs to be accurate. To further support this hypothesis, we re-train the CNNs employing only simulated lenses with $5 < \theta_E < 10 \text{ arcsec}$. By passing the image of DES J0329–2820 to the algorithms, we obtain $p_1 = 0.85$ and $p_3 = 0.79$. Hence, with a different training the lens would be part of the candidate sample produced by the algorithms. Moreover, when applied to the whole BG sample, the re-trained CNNs manage to detect 8 out of 16 candidates in the LIVE sample (i.e. with $\theta_E < 5 \text{ arcsec}$), even though in a candidate sample characterized by a larger fraction of false positives. This demonstrates that these specialized CNNs have some abilities to extrapolate their knowledge beyond the parameter space they are trained on, even though these are not sufficient to provide a complete and pure candidate sample.

Furthermore, for four of the objects in the LIVE sample, we retrieve high-resolution imaging from the *Hubble Space Telescope* Legacy Archive (HLA⁶). One of the objects is LIVE-5 (previously discovered, Fig. 11). Other two objects, LIVE-11 and LIVE-12, show likely lensing features when observed with *HST*. On the contrary, high-resolution data for LIVE-3 makes it possible to identify a likely spiral structure, revealing a non-lens nature for this candidate.

Continuing the analysis, we emphasize that six of the candidates in the LIVE sample satisfy the colour cut we introduced in Section 2.1 and are thus part of the LRG sample (see Table 3 and Section 2.1). In contrast, none of the systems satisfies the criteria espoused by Eisenstein et al. (2001) to select LRGs. Hence, using these criteria to select the input sample for the CNNs (as done in analogous studies; e.g. Petrillo et al. 2017, 2019a,b) we would therefore have missed all these candidates.

Finally, to fully characterize our set of candidates, we retrieve spectroscopic and photometric redshifts for most of the systems in the LIVE sample. These values are summarized in Table 3. The photometric redshifts are computed using the METAPHOR algorithm (Cavuoti et al. 2017), previously applied to the galaxies in the VOICE survey. The spectroscopic redshifts are retrieved from the VizieR archive⁷ (Ochsenbein, Bauer & Marcout 2000) querying the catalogues from previous spectroscopic surveys of the CDFS (Eales et al. 2009; Cowie, Barger & Hu 2011; Cooper et al. 2012; Cool et al. 2013).

6.1 Comparison with LENSPOP

We employ the lens-statistics code LENSPOP to assess the reliability of the LIVE sample. LENSPOP is a software introduced by Collett (2015) and able to simulate a realistic population of strong gravitational lenses. By opportunely tuning its parameters, LENSPOP can predict the number of lenses observable in a given survey, and their global properties. LENSPOP defines as ‘observable’ all the lenses satisfying

⁶<https://hla.stsci.edu/>

⁷<https://vizier.u-strasbg.fr/>



Figure 8. Some examples of contaminants wrongly classified as lenses by the CNNs and with visual score = 0. Among the most common misclassified objects there are spirals, merging, galaxies with close companions, and peculiar objects. All images have a 20-arcsec side; the north corresponds to the upper side and the east to the left-hand side of the stamps. Further details are in Section 5.2.

these criteria:

$$\begin{cases} \theta_E^2 \geq x_s^2 + y_s^2 \\ \theta_E^2 \geq r_s^2 + (s/2)^2 \\ \mu_{\text{TOT}} r_s > s, \quad \mu_{\text{TOT}} > 3 \\ S/N \geq 20 \end{cases} \quad (10)$$

where x_s , y_s , and r_s represent, respectively, the coordinates and the size of the unlensed source. θ_E represents the Einstein radius, μ the magnification, S/N the signal-to-noise ratio, and s the mean seeing of the image (Collett 2015). Besides these properties, we require that

$$\begin{cases} \text{mag}_r < 21.5 \\ 1 < \theta_E < 5 \text{ arcsec} \end{cases} \quad (11)$$

The first property requires the galaxy being in the BG sample (and, thus, being analysed by the CNNs; Section 2.1). The second property considers that, since we trained our algorithms on lenses with Einstein radii between 1 and 5 arcsec (Table 2), we do not expect our CNN to be accurate outside this range (Petrillo et al. 2019a,b). LENSPOP predicts that 10 strong lenses are observable in the 4.9 deg² of the CDFS covered by the VOICE survey. Assuming a Poissonian noise on the code prediction, we estimate a confidence interval of 10 ± 3 . Comparing this value with the size of the LIVE sample, we expect the latter to be nearly complete, but not entirely pure. This result agrees with the likely spiral galaxy identified in the sample using the *HST* data (Section 6). The lens population simulated by LENSPOP is predicted to have a mean redshift of 0.4 with a standard deviation of 0.2. Using the spectroscopic and photometric redshifts retrieved in the previous section for the LIVE sample, we estimate a mean redshift of 0.5 with a standard deviation of 0.2. Finally, using LENSPOP, we predict a mean value of the Einstein radii for the lenses observable in the VOICE survey equals to 1.4 arcsec with a standard deviation of 0.3 arcsec. Visually estimating the Einstein radii for the lenses in the LIVE sample as half the distance between the alleged multiple lensed images,⁸ we obtain a mean value of 1.5 arcsec with a standard deviation of 0.4 arcsec. Both the redshifts and the Einstein radii are consistent, within 1σ error, with the predictions made by LENSPOP, representing a further confirmation of the reliability of the candidates in the LIVE sample.

6.2 Comparison with KiDS

The CNNs employed in this work were previously applied to data from the *KiDS* (Kuijken et al. 2019) by Petrillo et al. (2017, 2019a,b). It is interesting to compare those results with ours, to investigate what

performances the same CNN architectures can achieve when applied to different data, although both produced by the same telescope, and thus similar. The images passed to the CNNs for VOICE and KiDS have the same pixel size (0.2 arcsec pixel⁻¹) and a comparable mean value of the PSF FWHM (0.8 arcsec for VOICE r band, and 0.7 arcsec for the same band in KiDS; section 2 in Kuijken et al. 2019). However, KiDS is a wide and shallow survey (about 900 deg² observed in the fourth data release with a 5σ limiting magnitude in the r band of 25.0; Kuijken et al. 2019). VOICE, on the contrary, is a smaller but deeper survey (4.9 deg² observed with a 5σ limiting magnitude in the same band of 26.1; see Section 2). In Petrillo et al. (2019b), the CNNs analysed 88 327 LRGs selected using the criteria espoused by Eisenstein et al. (2001). The LRG sample covered less than 0.01 per cent of the full KiDS catalogue. Adopting the same threshold used in this paper ($p_{\text{Th}} = 0.8$), the CNNs retrieved a sample of 3500 systems (about 4 per cent of the LRG sample) with at least one score above the threshold. Finally, performing a visual inspection of the selected candidates similar to the one presented in this paper (with seven graders and three possible grades), the authors assembled the LinKS (Lenses In KiDS) sample⁹ composed by 1983 likely strong gravitational lenses with at least one score above the threshold and at least one classification as ‘*maybe lens*’ or ‘*sure lens*’. Of those candidates (the ‘*bona fide*’ sample), 89 attained a visual score ≥ 28 ,¹⁰ equivalent to our threshold of unanimous classification as ‘*maybe lens*’ (considering seven inspectors instead of our nine, see Section 5.2). In comparing the results from the two studies, we must consider two main differences between the surveys: the higher number of galaxies observed in KiDS and the higher S/N (and fainter limiting magnitude) of VOICE. Since we worked on a smaller survey, we could relax the criteria of Eisenstein et al. (2001) for the selection of the galaxies to analyse. The BG sample inspected by the CNNs covered a larger fraction of the galaxies observed in the VOICE survey (about 3 per cent against the 0.01 per cent of KiDS). However, the fractions of systems retrieved by the CNNs are quite similar (1 per cent of the BG sample in VOICE, 3 per cent of the LRG sample in KiDS). Performing the visual inspection, we could independently assess the contamination rate of the candidate samples. In Petrillo et al. (2019b), ~ 57 per cent of the candidates attained at least one classification as ‘*maybe lens*’ or ‘*sure lens*’. This fraction is higher in our case (~ 75 per cent, see Section 5.2). This result can be explained by the fainter limiting magnitude reached by the VOICE survey. This allows the CNNs to identify more easily faint characteristics (e.g. spiral structures), revealing the contaminant nature of some lens candidates. This point can be studied in detail,

⁸This estimate is not completely accurate for strongly asymmetric lenses; a complete modelling would be required for a better accuracy.

⁹<https://www.astro.rug.nl/lensesinkids/>

¹⁰Petrillo et al. (2019b) used the same numerical values given in Section 5.2.




Figure 9. RGB stamps of lens candidates in the LIVE sample. Each image contains the two scores from the *single-band CNN* (red) and *three-band CNN* (blue). The systems are ordered from the top to bottom according to the final score obtained from visual inspection (grey box). All stamps have a 20-arcsec side and are produced using the HUMVI library. The north corresponds to the upper side and the east to the left-hand side of the stamps. Further details are in Section 5.2.

since the KiDS and VOICE fields overlap in a region of $\sim 2 \text{ deg}^2$. In particular, six lens candidates in the LinKS sample (but with visual score < 28 , i.e. not part of the ‘bona fide sample’) are also in the BG sample analysed by our CNNs. KiDS and VOICE cutouts for these systems are shown in Fig. 12. The higher S/N and better image quality of VOICE reveal the contaminant nature of all these systems. Five of these objects obtained both scores under the chosen threshold from our CNNs. Only the object with LinKS ID = 68 (the first system in Fig. 12) obtained $p_1 = 0.9$ from our *single-*

band CNN. This candidate, however, attained a visual score of 22 during our visual inspection, well below the chosen threshold of 36.

The fainter limiting magnitude (and the consequent higher S/N), combined with more flexible criteria to select the galaxies to analyse, is also responsible for the higher number density of lenses found by our CNNs. In fact, according to equation (10), the higher S/N reached by a deep survey augments the number of strong lenses retrievable in a given area (Collett 2015). To assess quantitatively this property, we

Table 3. Results of the visual inspection performed on the systems in the CNN sample (Section 5.1).

LIVE ID	Name	RA	Dec.	r	$photo-z^a$	$spec-z$	p_1	p_3	Score	LRG
1	VOICE J520934–282157	52.1595	–28.3658	21.24	0.56	0.48 ^(b)	0.99	0.87	72	
2	VOICE J522422–284602	52.4060	–28.7672	20.92	0.56		0.92	0.69	66	
3	VOICE J531403–273933	53.2342	–27.6592	20.33	0.57	0.62 ^c	0.52	0.99	66	
4	VOICE J534529–270652	53.7581	–27.1145	21.47			0.57	0.97	66	✓
5	VOICE J530933–275654	53.1592	–27.9482	20.51	0.61	0.61 ^c	0.07	0.8	60	
6	VOICE J524537–275130	52.7603	–27.8585	19.14	0.43	0.34 ^d	0.79	0.89	60	✓
7	VOICE J532836–274052	53.4766	–27.6811	21.38	0.83		0.37	0.92	58	
8	VOICE J532701–270801	53.4504	–27.1337	20.38			0.27	0.98	56	
9	VOICE J533716–275300	53.6210	–27.8833	20.64	0.57		0.94	0.11	56	✓
10	VOICE J514601–284848	51.7670	–28.8133	20.84	0.75		0.25	0.85	54	
11	VOICE J525404–274159	52.9012	–27.6998	18.94	0.14	0.07 ^c	0.87	0.04	48	
12	VOICE J530605–280115	53.1013	–28.0207	21.39	0.63	0.62 ^c	0.99	0.07	42	✓
13	VOICE J521812–280115	52.3033	–28.0375	20.27	0.55	0.54 ^e	0.29	0.9	40	✓
14	VOICE J520548–282538	52.0968	–28.4273	21.38	0.64	0.65 ^b	0.91	0.28	40	✓
15	VOICE J515150–272926	51.8638	–27.4905	21.09	0.55		0.85	0.23	40	
16	VOICE J534638–271022	53.7771	–27.1727	20.76			0.29	0.98	36	

Notes. For all the candidates with score ≥ 36 , the LIVE ID, the coordinates, the r -band magnitude, the redshift, the scores from the CNNs, and the visual score are provided. Moreover, we report the IAU designation of each object employing the standard $J\pm DDMSS$ convention. The last column reports if the candidate is part of the LRG sample or not (Section 2.1). Further details are in Section 5.2. ^aPhotometric redshift estimated with METAPHOR (Cavuoti et al. 2017).

^bSpectroscopic redshift retrieved from the PRIMUS survey (Cool et al. 2013).

^cSpectroscopic redshift retrieved from the ACES survey (Cooper et al. 2012).

^dSpectroscopic redshift retrieved from Cowie, Barger & Hu (2010).

^eSpectroscopic redshift retrieved from the BLAST survey (Eales et al. 2009).




Figure 10. RGB stamp of the gravitational lens DES J0329–2820 previously discovered by Nord et al. (2016) observed in the VOICE survey. The system has $r > 21.5$; thus, it is not part of the BG sample. We manually pass its image to the CNNs, receiving both scores below the chosen threshold. The high value of the Einstein radius ($7.8 \text{ arcsec} \pm 1.4 \text{ arcsec}$) can explain this result. Further details in Section 6.

can consider the $\sim 2 \text{ deg}^2$ area observed in both surveys. While the ‘bona fide’ LinKS sample contains no lens candidates in this area, the LIVE sample contains six candidates (LIVE IDs: 1, 5, 6, 12, 13, and 14) retrieved by our CNNs in this region.

Finally, a further interesting comparison between KiDS and VOICE concerns the mean redshift of the retrieved lens candidates. The systems in the LinKS sample have a mean redshift of 0.3 (Petrillo et al. 2019b), while the systems in the LIVE sample have a mean redshift of 0.5 (Section 6). Identifying strong lenses at higher redshift is crucial for extending many analyses on a larger scale (see e.g. Treu & Koopmans 2004; Koopmans 2005; Treu et al. 2010; Vegetti et al. 2014). All these results justify the increasing interest in the forthcoming deep surveys conducted with the *Euclid* satellite (Laureijs et al. 2011) and the *Vera Rubin Observatory* (LSST Science Collaboration et al. 2009).

7 CONCLUSION

In this paper, we presented a sample of 16 likely strong gravitational lenses identified in the CDFS. We analysed the data from the VOICE survey (Vaccari et al. 2016) using two CNNs.

Both algorithms were previously developed by Petrillo et al. (2017, 2019a) and employed to search for strong lenses in the *KiDS* (Kuijken et al. 2019) by Petrillo et al. (2017, 2019a,b) and the *FDS* (Iodice et al. 2016) by Cantiello et al. (2020). We trained the CNNs on composite images obtained by superimposing simulated gravitational arcs on real LRGs observed in VOICE (Section 3.2.2). The first CNN, *single-band CNN*, analysed images in the r photometric band, while the second one, *three-band CNN*, inspected composite RGB images obtained combining the data in the *gri* bands with the HUMVI library. Once the algorithms have been trained, we assessed their performances by applying them to a validation set consisting of both simulated lenses and real contaminants (Section 4). The performances of both networks (i.e. the FPR and the TNR) are comparable to the previous applications in Petrillo et al. (2019a,b). Moreover, we found that the *three-band CNN* can identify more easily systems with smaller Einstein radii, where the colour gradient can help to recognize unresolved gravitational arcs. On the contrary, the *single-band CNN* shows a better accuracy in identifying systems with larger Einstein radii. In this case, however, high- z groups of star-forming galaxies can be more easily mistaken for distant gravitational arcs.

Concluding that the two CNNs are complementary, we applied both networks to real data from the VOICE survey. The CNNs analysed in total $\sim 21\,200$ galaxies with $mag_r < 21.5$, retrieving a sample of 257 lens candidates with at least one score above the chosen threshold of 0.8 (Section 5.1). To improve the purity of the candidate sample, we performed a visual inspection with nine graders judging the systems in a blind way (Section 5.2). About 75 per cent of the candidates attained at least one classification as ‘*maybe lens*’ or ‘*sure lens*’.




Figure 11. Systems LIVE-5, LIVE-12, LIVE-11, and LIVE-3 observed by the *Hubble Space Telescope* (top row) and in the VOICE survey (bottom row). The *HST* images are part of the GEMS survey (*HST* observing program 9500; Rix et al. 2004) and are retrieved from the Hubble Legacy Archive. All the images have a 20-arcsec side. LIVE-5 is a previously discovered strong gravitational lens (Blakeslee et al. 2004). LIVE-12 and LIVE-11 show likely lensing features, while LIVE-3 shows a likely spiral structure. Further details in Section 6.




Figure 12. Six lens candidates found in the KiDS survey (top row) by Petrillo et al. (2019b) and observed in the VOICE survey (bottom row). The higher *S/N* of VOICE reveals the non-lensing nature of all these candidates and makes it easier for the CNNs to reject these systems. For all the systems, we show the LinKS ID (white box) and the two scores obtained by the *single-band CNN* (red box) and the *three-band CNN* (blue box). All the stamps have a 20-arcsec side; the north corresponds to the upper side and the east to the left-hand side of the stamps. Further details are in Section 6.2.

Finally, we assembled the ‘LIVE sample’ consisting of 16 likely strong gravitational lenses with at least one score above the threshold and a visual score ≥ 36 . To fully characterize the final set, we retrieved spectroscopic and photometric redshifts for most of the lens candidates. We also retrieved high-resolution data from the Hubble Legacy Archive for four of the systems (Section 6). The entire process described here allowed us to identify a gravitational lens previously discovered in the CDFS (Blakeslee et al. 2004) and at least two very high-probability candidates when observed by *HST* (Fig. 11). To assess the reliability of the LIVE sample, we compared its global properties with the ones predicted by the lens-statistics software LENSPOP (Collett 2015). We concluded that our sample is likely to be complete albeit not totally pure, while its global properties fully encompass the code predictions (Section 6.1). Finally, we compared our results with the ones presented in Petrillo et al. (2019b), obtained using the same CNNs applied to the KiDS survey.

Since we applied the algorithms to a smaller but deeper survey, we were able to retrieve a less contaminated candidate sample, with a higher number density of lens candidates and a higher mean redshift (Section 6.2).

Although the probability to be confirmed as lens is high for most of the objects in the LIVE sample, we stress that an unambiguous validation requires a high-resolution and/or a spectroscopic follow-up (see e.g. Bolton et al. 2006; Anguita et al. 2018; Lemon et al. 2020; Spiniello et al. 2019a, b), which will be provided by the *Vera Rubin Observatory* deep survey that will observe the CDFS in the near future (LSST Science Collaboration et al. 2009).

In conclusion, this work represents a further confirmation of the ability of machine learning algorithms like CNNs to analyse efficiently large amounts of data searching for strong gravitational lenses. These algorithms will reach their full scientific potential in the analysis of forthcoming large sky surveys such as the one performed

with the ESA's *Euclid* satellite (Laureijs et al. 2011), the *Vera Rubin Observatory* (LSST Science Collaboration et al. 2009), and the *Chinese Space Station* (Gong et al. 2019). These surveys are indeed expected to retrieve $\sim 10^5$ strong gravitational lenses in a data set of $\sim 10^9$ observed galaxies. Solely visually inspecting all the galaxies retrieved by these forthcoming facilities would require several years and would be prone to several biases (Section 5.2), even applying some a priori cut to select only galaxies with a high lensing cross-section. However, even applying CNNs to select the most promising lens candidates, a low contamination rate is still crucial to reduce the need for a visual inspection. This goal can be achieved, on one hand, by employing the latest CNN architectures available (see e.g. Chollet 2016; Szegedy et al. 2016), and thus taking advantage of the latest results in machine learning and computer vision. On the other hand, training these algorithms requires reliable strong lensing simulations to avoid possible biases in the training phase. This is the reason why, in the last few years, some collaborations started to investigate possible alternatives to the supervised-learning paradigm. Unsupervised learning (requiring no training set or a small one just for labelling (see e.g. Cheng et al. 2020) or self-supervised learning (requiring smaller data sets, e.g. Abul Hayat et al. 2020) can allow training based only on real observed strong lenses.

Finally, it is worth noting that the large amount of lenses retrieved from these forthcoming large surveys will pose the non-trivial problem of how efficiently one can then analyse and model these systems to constrain structural parameters of the lens to be used for scientific purposes. The computational cost of classical bayesian techniques (e.g. Jullo et al. 2007; Birrer & Amara 2018; Nightingale, Dye & Massey 2018) and their need for a human intervention, make these algorithms unfit to model large samples of lenses. Machine learning algorithms like CNNs have already been applied to the fast automated analysis of strong gravitational lenses (Hezaveh, Perreault Levasseur & Marshall 2017; Madireddy et al. 2019; Pearson, Li & Dye 2019; Schuldt et al. 2020). This represents a future perspective of this work, towards a full exploitation of the scientific potential of forthcoming facilities like *Euclid*, the *Vera Rubin Observatory*, and the *Chinese Space Station*.

ACKNOWLEDGEMENTS

We thank the anonymous referee for useful suggestions to improve this paper. FG acknowledges the support from grant PRIN MIUR 2017-20173ML3WW.001. CS is supported by a Hintze Fellowship at the Oxford Centre for Astrophysical Surveys, which is funded through generous support from the Hintze Family Charitable Foundation. MV acknowledges support from the Italian Ministry of Foreign Affairs and International Cooperation (MAECI: Grant Number ZA18GR02) and the South African Department of Science and Innovation's National Research Foundation (DSI-NRF Grant Number 113121) as part of the ISARP RADIOSKY2020 Joint Research Scheme. LPF acknowledges the support from NSFC grants 11933002, STCSM grant 18590780100, and the Dawn Program 19SG41 and the Innovation Program 2019-01-07-00-02-E00032 of SMEC. Based on observations made with the NASA/ESA *Hubble Space Telescope*, and obtained from the Hubble Legacy Archive, which is a collaboration between the Space Telescope Science Institute (STScI/NASA), the Space Telescope European Coordinating Facility (ST-ECF/ESA), and the Canadian Astronomy Data Centre (CADN/CRC/CSA). This research has made use of the VizieR catalogue access tool, CDS, Strasbourg, France (DOI: 10.26093/cds/vizieR).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author, GC, upon reasonable request.

REFERENCES

- Abadi M. et al., 2015, preprint (arXiv:1603.04467)
 Abul Hayat M., Stein G., Harrington P., Lukić Z., Mustafa M., 2020, preprint (arXiv:2012.13083)
 Anguita T. et al., 2018, *MNRAS*, 480, 5017
 Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, *MNRAS*, 310, 540
 Auger M. W., Treu T., Gavazzi R., Bolton A. S., Koopmans L. V. E., Marshall P. J., 2010a, *ApJ*, 721, L163
 Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V. E., Marshall P. J., Moustakas L. A., Burles S., 2010b, *ApJ*, 724, 511
 Baeza-Yates R., Ribeiro-Neto B., 1999, *Modern Information Retrieval*, Vol. 463. ACM press, New York
 Barnabè M., Spiniello C., Koopmans L. V. E., Trager S. C., Czoske O., Treu T., 2013, *MNRAS*, 436, 253
 Bartelmann M., 2010, *Class. Quantum Gravity*, 27, 233001
 Becker B., Vaccari M., Prescott M., Grobler T., 2021, *MNRAS*, 503, 1828
 Bertin E., 2011, in Evans I. N., Accomazzi A., Mink D. J., Rots A. H., eds, *ASP Conf. Ser. Vol. 442, Astronomical Data Analysis Software and Systems XX*. Astron. Soc. Pac., San Francisco, p. 435
 Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
 Birrer S., Amara A., 2018, *Phys. Dark Univ.*, 22, 189
 Blakeslee J. P. et al., 2004, *ApJ*, 602, L9
 Blandford R. D., Narayan R., 1992, *ARA&A*, 30, 311
 Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Moustakas L. A., 2006, *ApJ*, 638, 703
 Browne I. W. A. et al., 2003, *MNRAS*, 341, 13
 Canameras R. et al., 2020, *A&A*, 644, A163
 Cantiello M. et al., 2020, *A&A*, 639, A136
 Cao S., Pan Y., Biesiada M., Godlowski W., Zhu Z.-H., 2012, *J. Cosmol. Astropart. Phys.*, 2012, 016
 Capaccioli M., Schipani P., 2011, *The Messenger*, 146, 2
 Cappellaro E. et al., 2015, *A&A*, 584, A62
 Cavuoti S., Amaro V., Brescia M., Vellucci C., Tortora C., Longo G., 2017, *MNRAS*, 465, 1959
 Chae K.-H., 2003, *MNRAS*, 346, 746
 Chatterjee S., 2019, PhD thesis, Univ. Groningen
 Chatterjee S., Koopmans L. V. E., 2018, *MNRAS*, 474, 1762
 Cheng T.-Y., Li N., Conselice C. J., Aragón-Salamanca A., Dye S., Metcalf R. B., 2020, *MNRAS*, 494, 3750
 Chollet F. et al., 2015, *Keras*
 Chollet F., 2016, preprint (arXiv:1610.02357)
 Collett T. E., 2015, *ApJ*, 811, 20
 Cool R. J. et al., 2013, *ApJ*, 767, 118
 Cooper M. C. et al., 2012, *MNRAS*, 425, 2116
 Covone G. et al., 2009, *ApJ*, 691, 531
 Cowie L. L., Barger A. J., Hu E. M., 2010, *ApJ*, 711, 928
 Cowie L. L., Barger A. J., Hu E. M., 2011, *ApJ*, 738, 136
 Dalal N., Kochanek C. S., 2002, *ApJ*, 572, 25
 De Cicco D. et al., 2015, *A&A*, 574, A112
 Eales S. et al., 2009, *ApJ*, 707, 1779
 Eisenstein D. J. et al., 2001, *AJ*, 122, 2267
 Faure C. et al., 2008, *ApJS*, 176, 19
 Fu L. et al., 2018, *MNRAS*, 479, 3858
 Gavazzi R., Treu T., Rhodes J. D., Koopmans L. V. E., Bolton A. S., Burles S., Massey R. J., Moustakas L. A., 2007, *ApJ*, 667, 176
 Giacomini R. et al., 2001, *ApJ*, 551, 624
 Gong Y. et al., 2019, *ApJ*, 883, 203
 Goodfellow I., Bengio Y., Courville A., 2016, *Deep Learning*. MIT Press, Cambridge
 He Z. et al., 2020, *MNRAS*, 497, 556
 He K., Zhang X., Ren S., Sun J., 2015a, preprint (arXiv:1502.01852)

- He K., Zhang X., Ren S., Sun J., 2015b, preprint ([arXiv:1512.03385](https://arxiv.org/abs/1512.03385))
- Hezaveh Y., Dalal N., Holder G., Kisner T., Kuhlen M., Perreault Levasseur L., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 048
- Hezaveh Y. D., Perreault Levasseur L., Marshall P. J., 2017, *Nature*, 548, 555
- Hornik K., 1991, *Neural Netw.*, 4, 251
- Huang X. et al., 2020, *ApJ*, 894, 78
- Ilbert O. et al., 2006, *A&A*, 457, 841
- Iodice E. et al., 2016, *ApJ*, 820, 42
- Jacobs C. et al., 2019a, *ApJS*, 243, 17
- Jacobs C. et al., 2019b, *MNRAS*, 484, 5330
- Jacobs C., Glazebrook K., Collett T., More A., McCarthy C., 2017, *MNRAS*, 471, 167
- Jarvis M. J. et al., 2013, *MNRAS*, 428, 1281
- Jullo E., Kneib J.-P., Limousin M., Elíasdóttir Á., Marshall P. J., Verdugo T., 2007, *New J. Phys.*, 9, 447
- Keeton C. R., Kochanek C. S., Seljak U., 1997, *ApJ*, 482, 604
- Kingma D. P., Ba J., 2014, preprint ([arXiv:1412.6980](https://arxiv.org/abs/1412.6980))
- Koopmans L. V. E., 2005, *MNRAS*, 363, 1136
- Kormann R., Schneider P., Bartelmann M., 1994, *A&A*, 284, 285
- Kuijken K. et al., 2019, *A&A*, 625, A2
- Kuijken K., 2011, *The Messenger*, 146, 8
- La Barbera F., de Carvalho R. R., Kohl-Moreira J. L., Gal R. R., Soares-Santos M., Capaccioli M., Santos R., Sant'Anna N., 2008, *PASP*, 120, 681
- Laureijs R. et al., 2011, preprint ([arXiv:1110.3193](https://arxiv.org/abs/1110.3193))
- LeCun Y., Bottou L., Bengio Y., Haffner P., 1998, *Proc. IEEE*, 86, 2278
- LeCun Y., Bengio Y., Hinton G., 2015, *Nature*, 521, 436
- Lemon C. et al., 2020, *MNRAS*, 494, 3491
- Le Fevre O., Hammer F., 1988, *ApJ*, 333, L37
- Li R. et al., 2020, *ApJ*, 899, 30
- Li R. et al., 2021, preprint ([arXiv:2110.01905](https://arxiv.org/abs/2110.01905))
- Liu D. et al., 2018, *MNRAS*, 478, 2388
- LSST Science Collaboration et al., 2009, preprint ([arXiv:0912.0201](https://arxiv.org/abs/0912.0201))
- Lupton R., Blanton M. R., Fekete G., Hogg D. W., O'Mullane W., Szalay A., Wherry N., 2004, *PASP*, 116, 133
- Madireddy S., Li N., Ramachandra N., Butler J., Balaprakash P., Habib S., Heitmann K., 2019, preprint ([arXiv:1911.03867](https://arxiv.org/abs/1911.03867))
- Mao S., Schneider P., 1998, *MNRAS*, 295, 587
- Marshall P. J. et al., 2016, *MNRAS*, 455, 1171
- Metcalf R. B. et al., 2019, *A&A*, 625, A119
- Möller O., Kitzbichler M., Natarajan P., 2007, *MNRAS*, 379, 1195
- More A., Cabanac R., More S., Alard C., Limousin M., Kneib J.-P., Gavazzi R., Motta V., 2012, *ApJ*, 749, 38
- Nightingale J. W., Dye S., Massey R. J., 2018, *MNRAS*, 478, 4738
- Nord B. et al., 2016, *ApJ*, 827, 51
- Ochsenbein F., Bauer P., Marcout J., 2000, *A&AS*, 143, 23
- Oguri M. et al., 2006, *AJ*, 132, 999
- Oguri M., Marshall P. J., 2010, *MNRAS*, 405, 2579
- Oguri M., Rusu C. E., Falco E. E., 2014, *MNRAS*, 439, 2494
- Oliver S. et al., 2000, *MNRAS*, 316, 749
- Pearson J., Li N., Dye S., 2019, *MNRAS*, 488, 991
- Petrillo C. E. et al., 2017, *MNRAS*, 472, 1129
- Petrillo C. E. et al., 2019a, *MNRAS*, 482, 807
- Petrillo C. E. et al., 2019b, *MNRAS*, 484, 3879
- Poulain M. et al., 2020, *A&A*, 634, A50
- Pourrahmani M., Nayyeri H., Cooray A., 2018, *ApJ*, 856, 68
- Refsdal S., 1964, *MNRAS*, 128, 307
- Rix H.-W. et al., 2004, *ApJS*, 152, 163
- Rowan-Robinson M. et al., 2004, *MNRAS*, 351, 1290
- Rumelhart D. E., Hinton G. E., Williams R. J., 1986, *Learning Internal Representations by Error Propagation*. MIT Press, Cambridge, MA
- Russakovsky O. et al., 2015, *Int. J. Comput. Vision*, 115, 211
- Saerens M., Latinne P., Decaestecker C., 2002, *Neural Comput.*, 14, 21
- Schneider P., Ehlers J., Falco E. E., 1992, *Gravitational Lenses*. Springer-Verlag, New York
- Schuldt S., Suyu S. H., Meinhardt T., Leal-Taixé L., Cañameras R., Taubenberger S., Halkola A., 2020, *A&A*, 646, A126
- Sérsic J. L., 1963, *Bol. Asoc. Astron. La Plata Argentina*, 6, 41
- Shu Y. et al., 2017, *ApJ*, 851, 48
- Simard P., Steinkraus D., Platt J., 2003, *Seventh International Conference on Document Analysis and Recognition*, Vol. 2. p. 958
- Sonnenfeld A., Cautun M., 2021, *A&A*, 651, A18
- Sonnenfeld A., Jaelani A. T., Chan J., More A., Suyu S. H., Wong K. C., Oguri M., Lee C.-H., 2019, *A&A*, 630, A71
- Spiniello C., Koopmans L. V. E., Trager S. C., Czoske O., Treu T., 2011, *MNRAS*, 417, 3000
- Spiniello C., Agnello A., Sergeev A. V., Anguita T., Rodríguez Ó., Napolitano N. R., Tortora C., 2019a, *MNRAS*, 483, 3888
- Spiniello C. et al., 2019b, *MNRAS*, 485, 5086
- Syget J. F., Tu H., Fort B., Gavazzi R., 2010, *A&A*, 517, A25
- Szegedy C., Ioffe S., Vanhoucke V., Alemi A., 2016, preprint ([arXiv:1602.07261](https://arxiv.org/abs/1602.07261))
- Tortora C. et al., 2018, *MNRAS*, 481, 4728
- Tortora C., Napolitano N. R., Romanowsky A. J., Jetzer P., 2010, *ApJ*, 721, L1
- Tozzi P. et al., 2001, *ApJ*, 562, 42
- Treu T., 2010, *ARA&A*, 48, 87
- Treu T., Koopmans L. V. E., 2004, *ApJ*, 611, 739
- Treu T., Auger M. W., Koopmans L. V. E., Gavazzi R., Marshall P. J., Bolton A. S., 2010, *ApJ*, 709, 1195
- Vaccari M. et al., 2016, *The 4th Annual Conference on High Energy Astrophysics in Southern Africa (HEASA 2016)*. p. 26
- van der Walt S. et al., 2014, *PeerJ*, 2, e453
- Vegetti S., Koopmans L. V. E., Auger M. W., Treu T., Bolton A. S., 2014, *MNRAS*, 442, 2017
- Wong K. C. et al., 2020, *MNRAS*, 498, 1420
- Yuan H. B., Liu X. W., Xiang M. S., 2013, *MNRAS*, 430, 2188
- Zwicky F., 1937, *Phys. Rev.*, 51, 290

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