University of Groningen

Discovering gravitational lenses with artificial intelligence

Petrillo, Enrico

DOI:
[10.33612/diss.100697045](https://doi.org/10.33612/diss.100697045)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2019

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Petrillo, E. (2019). *Discovering gravitational lenses with artificial intelligence*. [Thesis fully internal (DIV), University of Groningen]. Rijksuniversiteit Groningen. <https://doi.org/10.33612/diss.100697045>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

rijksuniversiteit
 groningen

Discovering Gravitational Lenses with Artificial Intelligence

Proefschrift

ter verkrijging van de graad van doctor aan de
Rijksuniversiteit Groningen
op gezag van de
rector magnificus prof. dr. C. Wijmenga
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op

vrijdag 1 november 2019 om 12.45 uur

door

Carlo Enrico Petrillo

geboren op 11 Augustus 1986
te Napels, Italië

Supervisors

Prof. dr. L. V. E. Koopmans

Co-supervisor

Dr. G. Verdoes Kleijn

Dr. C. Tortora

Assessment Committee

Prof. dr. C. Fassnacht

Prof. dr. G. Longo

Prof. dr. R. Peletier

To the god of science, to whom one part of me believes

Cover design: Rustam Rysov and Carlo Enrico Petrillo. Inspired by the movie 2001: A Space Odyssey. The astronomical image on the cover is the gravitational lens G2237+0305 observed by the Hubble Space Telescope.
ISBN: 978-94-034-2023-3 (printed version)
ISBN: 978-94-034-2022-6 (electronic version)

If you meet the Buddha on the road, kill him

– Linji Yixuan

Everything that is possible demands to exist

– Gottfried Wilhelm Leibniz

Contents

1	INTRODUCTION.	1
1.1	Context	1
1.2	Gravitational Lensing	2
1.2.1	Basic Lensing Formalism for axisymmetric lenses	3
1.2.2	Strong-lensing applications	6
1.2.3	Surveying the lenses	7
1.3	The Kilo-Degree Survey	8
1.4	Convolutional Neural Networks	12
1.4.1	Algorithm overview	14
1.4.2	Training	16
1.5	This thesis	18
2	FINDING STRONG GRAVITATIONAL LENSES IN THE KILO DEGREE SURVEY WITH CONVOLUTIONAL NEURAL NETWORKS	21
2.1	Introduction	22
2.2	The KiDS survey	25
2.2.1	Data Release Three	25
2.2.2	Luminous red galaxy sample	27
2.3	Training the CNN to find Lenses	28
2.3.1	Input Samples	28
2.3.2	Building the training examples	31
2.3.3	Data augmentation	32
2.4	Results	35
2.4.1	Candidate selection	35
2.4.2	Final sample of candidates	36

2.4.3 Sample Characterization and Comparison	41
2.5 Conclusions	49
2.A NEURAL NETWORKS	52
2.A.1 Feed-forward neural network	57
2.A.2 Convolutional Neural Network	57
2.B CNN implementation	62
2.B.1 Network architecture	62
2.B.2 Training	63
2.B.3 Analysis	64
2.C r-band images of the candidates	65
3 TESTING CONVOLUTIONAL NEURAL NETWORKS FOR FINDING STRONG GRAVITATIONAL LENSES IN KIDS	73
3.1 Introduction	74
3.2 Training the ConvNets to find strong lenses	77
3.2.1 Data	78
3.2.2 Creating the training set	83
3.3 Analysis	84
3.3.1 Performance metric	86
3.3.2 Performance	87
3.4 Application to real data	90
3.4.1 Results on the LRG sample	90
3.4.2 Application to a small sample of clear lens candidates	92
3.4.3 Visual inspection	95
3.4.4 Prospects for Euclid	99
3.5 Discussion and Conclusions	100
3.A Technical details	104
3.A.1 Convolutional Neural Networks	104
3.A.2 Implementation	106
4 LINKS: DISCOVERING GALAXY-SCALE STRONG LENSES IN THE KILO-DEGREE SURVEY USING CONVOLUTIONAL NEURAL NETWORKS	109
4.1 Introduction	110

4.2	Data from the Kilo-Degree Survey	113
4.2.1	The “full sample”	114
4.2.2	The luminous red galaxy sample	115
4.3	Searching for Lenses	116
4.3.1	Training the Convolutional Neural Networks	116
4.3.2	Application to the LRG sample	120
4.4	The LinKS sample candidates	127
4.4.1	Candidate properties	127
4.4.2	Predictions and Prospects: Euclid and LSST	129
4.5	The full sample candidates	132
4.5.1	A high-purity sample	133
4.5.2	Small high-purity Euclid & LSST samples	133
4.5.3	The “bonus sample”	137
4.6	Discussion and Conclusions	137
4.A	Data	142
4.A.1	LinKS sample	142
4.A.2	Bonus sample	142
5	CONCLUSIONS AND FUTURE PROSPECTS	147
5.1	Conclusions	147
5.1.1	The first CNN lens-finder	148
5.1.2	Improving the CNN lens-finder	148
5.1.3	LinKS: Lenses in the Kilo-Degree Survey	149
5.2	Future Plans	149
6	BIBLIOGRAPHY	151
	Bibliography	163
A	Guide to the script	165
A.1	Dependencies	165
A.2	Files	165
A.3	Data preparation	166
A.4	Parameters	166
A.5	Running the script	167

SUMMARY	169
SAMENVATTING.....	177
COMPENDIO	185
AFTERWORD & ACKNOWLEDGEMENTS	193

1. INTRODUCTION

There are these two young fishes swimming along and they happen to meet an older fish swimming the other way, who nods at them and says “Morning, boys. How’s the water?” And the two young fishes swim on for a bit, and then eventually one of them looks over at the other and goes “What the hell is water?”

– David Foster Wallace

1.1 Context

Strong gravitational lensing enables us to carry out cosmological, and galaxy formation and evolution studies. The insights and the accuracy that can be achieved through the analysis of strong lenses depend on the availability of these astrophysical objects, which are rare and difficult to identify in astronomical surveys. Therefore, it is worth developing automatic identification methods for aiding the discovery of lensing systems. Especially considering that forthcoming imaging surveys will produce extremely large amount of data that will hardly be directly seen or analysed by any human. In this thesis I have developed such a method and applied it to an on-going wide-field survey.

The following parts of this introduction provide a general overview for a better understanding of the scientific context, the scientific drive of this thesis and the methods here developed. In particular, I describe the basics of the lensing formalism, and I discuss some of the important scientific applications of strong gravitational lensing. I then give a short overview of past, present and future astronomical surveys dedicated to the search of strong gravitational lenses. In particular, I describe the Kilo-Degree Survey (KiDS), which has provided the core data used in this thesis. Subsequently, I describe Convolutional Neural Networks which are the algorithms at the foundation of the method that I developed to search for strong gravitational lenses in KiDS data. Finally, I present how the thesis is structured.

1.2 Gravitational Lensing

With Einstein's theory of General Relativity (GR; Einstein 1915) humans have gained tremendous insights on one of the most fundamental physical processes governing the Universe: the property which we call mass is deeply linked to the structure of space-time itself. In particular, massive objects shape space-time in which they are embedded, and that is the very reason why objects accelerate in the vicinity of another mass, just like the Earth orbits the Sun. The fact that all massive objects are attracted to each other was already understood in the 17th century thanks to another milestone of Physics: Newton's law of universal gravitation (Newton, 1726). However, GR tells us that the gravitational attraction is not a fundamental force that occurs between massive objects, but it is a consequence of the deformation of space-time caused by the masses therein. Thence, even massless entities are subjected to "gravitational forces". That is, even the path of a light ray must deviate in the vicinity of a mass. The effect of this change is subtle but can be apprehended on astronomical scales.

For the first measurement of the phenomenon, we must go back to 1919. At that time, during a total solar eclipse, the positions of stars at the limb of the Sun were measured by astronomers (Dyson et al., 1920). They found that their apparent positions were seemingly displaced from their original position, as predicted by Einstein. This observation was the first direct evidence of GR.

A more spectacular phenomenon can be observed when massive deflectors such as galaxies, groups and clusters of galaxies bend light rays coming from a background galaxy seen along the same line of sight, creating for example magnified extended arcs and multiple images (see Fig. 1.1). In this case, we call this effect *strong gravitational lensing*. The term *lensing* is used because of the similarity of the effect to the deflections of light-rays in optical systems. The term *strong* is used to distinguish it from other regimes of the phenomenon: i.e., *micro* lensing, when the observed effect is only a magnification of the background sources, but otherwise unresolved on scales of milli-arcsecond and above; *weak* lensing, when lensed sources appear only slightly displaced and distorted by the entire inhomogeneous distribution of matter present along their cosmological path.

Strong gravitational lensing has many scientific applications, which derive from the property that the shape and position of lensed images depend almost exclusively on the mass distribution of the deflectors and the geometric configurations of lensing systems. Therefore, studying strong-lensing systems can provide valuable insight into the structure and

Figure 1.1: A luminous red galaxy distorts the light from a more distant blue galaxy in to a nearly complete blue *Einstein ring* (image credit: ESA/Hubble & NASA).

composition of the foreground galaxies and into the nature of space-time itself.

1.2.1 Basic Lensing Formalism for axisymmetric lenses

In this section, I present a general overview of lensing formalism, for a more detailed introduction see, e.g., Schneider (2006). A gravitational lensing system can be described by the simple geometric representation shown in Fig. 1.2. Gravitational lensing just deflects photons, it neither creates nor destroys them. Thus, the surface brightness of a source that is being lensed remains unchanged. What it changes is the dimension of the image which, if enlarged, appears brighter; just like any other magnifying glass. Light rays coming from a background source are deflected by the gravitational potential of the lens galaxy over an angle $\hat{\alpha}$ before reaching the observer. We ignore the two-dimensional nature of the system, but the scalars in these equations can also be replaced by a vector representation in two dimensions. One can define β as the angle between the optical axis and the true source position and θ as the angle between the optical axis and the image; in addition one can further define the reduced deflection

angle

$$\alpha = \frac{D_{ds}}{D_s} \hat{\alpha}, \quad (1.1)$$

where D_{ds} and D_s are the angular diameter distances between the source and the so-called *lens*, and the source and the observer, respectively. The angular diameter distances are defined such that the classical Euclidean relation *separation* = *angle* $\times$ *distance* holds in the curved space-time described by the Friedmann-Robertson-Walker metric. It is then straightforward to obtain the *lens equation* that provides the relationship between the source and image positions:

$$\beta = \theta - \alpha(\theta). \quad (1.2)$$

This equation is non-linear and can have multiple solutions, depending on the mass model and the source position. The latter implies that multiple images of the same background source can be observed in particular circumstances (i.e. a favourable lens-source alignment and a sufficiently massive lens).

If the Newtonian potential Φ and the peculiar velocity of the lens v are small in scale and amplitude (i.e., $\Phi \ll c^2$ and $v \ll c$; this holds for many astrophysical cases but not for extreme gravitational fields, as near black holes or neutron stars) and we can consider the lens as a sheet orthogonal to the line of sight (since the light deflection occurs on a very short distance respect to the distances between source, lens and observer, this can be practically considered true in all astrophysical cases) the deflection angle is given, for an axisymmetric lens, by

$$\hat{\alpha} = \frac{4GM(\xi)}{c^2\xi}, \quad (1.3)$$

where ξ is the impact parameter from the lens center and $M(\xi)$ is the mass enclosed within ξ . The reduced deflection angle of Eq. 1.1 can then be rewritten in the case of a constant surface mass density as

$$\alpha(\theta) = \frac{4\pi G\Sigma}{c^2} \frac{D_{ds}D_d}{D_s} \theta$$

where Σ and D_d are the surface mass surface density of the mass sheet and the distance between the observer and the source respectively. If Σ is greater than the *critical density*

$$\Sigma_{crit} \equiv \frac{c^2}{4\pi G} \frac{D_s}{D_{ds}D_s}$$

Figure 1.2: Schematic illustration of the geometry of a lensing system. The light coming from a background galaxy is deviated by a *lens* galaxy before arriving to the observer.

the lens is called *supercritical* and it becomes possible to observe multiple images.

If we consider an axisymmetric lens, using eq.1.3, one can rewrite the lens equation 1.2 as

$$\beta = \theta - \frac{D_s}{D_{ds}D_s} \frac{4GM(\theta)}{c^2\theta}.$$

If the lens is supercritical and the source lies on the optical axis (i.e., $\beta=0$), the source light, due to the rotational symmetry of the system, is projected in a ring of radius

$$\theta_E = \left(\frac{4GM(\theta)}{c^2} \frac{D_{ds}}{D_d D_s} \right)^{1/2}$$

called the *Einstein radius*. The Einstein radius represents a general property of lensing systems: it contains the main parameters of the system; the typical angular separation between multiple images is of the order of $2\theta_E$; images near the Einstein radius exhibit strong magnification compared to the original unlensed source.

1.2.2 Strong-lensing applications

Astronomers actively search for strong lenses because they are extremely valuable tools to address many open scientific questions about the structure of our Universe. Other than providing a window on otherwise unobservable higher-redshift galaxies, strong lenses can be used in determining the expansion history of the Universe and to measure with exceptional accuracy the mass distribution of the lensing galaxies. In turn, the above applications allow astronomers to shed light on galaxy formation and evolution and, thus, in testing our current model of the Universe.

One of the most important application of strong lensing is estimating the Hubble constant (H_0) by measuring the time delays between the lensed images of a time-variable lensed source. These delays are caused by different geometric light-paths, for the different images of the same source, as well as by the light passing through different potential depths; thus an intrinsic variation of the source luminosity will be seen at different times for each different image. The time delay between the images is inversely proportional to the Hubble constant and depends mainly on the distances of source and lens, and also on the gravitational potential of the lens. Thus, measuring time delays and constructing an accurate model of the lens potential allows measuring the Hubble constant. Recently, the value of H_0 has been measured with a few-percent accuracy by monitoring variable lensed quasars (Bonvin et al. 2017). Moreover, strong lensing can be used to put constraints on the dark energy equation of state which describes the expansion history of the Universe. This can be achieved either using large samples of lens systems (Cao et al., 2015) or by studying rare systems where a source is lensed by two galaxies at different redshifts (Collett & Auger, 2014).

Strong lensing is also a unique probe for measuring the distribution of total matter at galactic and sub-galactic scales. Such studies are crucial for validating models of galaxy formation. For example, Koopmans et al. (2009) estimated the density slope inside one effective radius of massive early-type galaxies by using a sample of 58 strong lenses. The fraction of dark matter of massive early-type galaxies has been estimated as well by, e.g., Sonnenfeld et al. (2015), who measured the quantity in the inner five kpc by using a sample of 81 strong lenses. Strong lensing also allows one to probe dark-matter substructure. In fact, sub-halos have been detected in strong lenses by, e.g., Vegetti et al. (2012); Nierenberg et al. (2014); Hezaveh et al. (2016).

Also, strong lenses act as “cosmic telescopes”: they provide a magnified view of background objects that is generally of the order of one magnitude.

This magnification allows one to image high-redshift lensed galaxies with a high level of detail (e.g., Mason et al. 2017; Salmon et al. 2017). An effectively higher resolution image of the lensed source is achieved thanks to the spatial stretching of the lensed source. This is particularly useful for small structures, such as active galactic nuclei and bulges of distant galaxies, which have scales below the nominal spatial resolution (e.g., Peng et al. 2006).

1.2.3 Surveying the lenses

Different systematic searches for gravitational lenses have been carried out. Traditionally, gravitational lenses have been found serendipitously or by visually inspecting survey images. For example, twenty secure galaxy-scale lens systems were discovered by visually inspecting HST images taken as part of the COSMOS Survey (Faure et al., 2008; Jackson, 2008). The Sloan Quasar Lens Survey (SQLS; Inada et al. 2012) adopted a more sophisticated approach which led to discovering 26 galaxy-scale multiply-imaged quasars in a spectroscopically selected sample by looking for magnified sources that appeared brighter than the lens galaxies in selected wavelengths. A similar visual approach has been used in the radio with the Cosmic Lens All Sky Survey (CLASS; Browne et al. 2003) which has led to the discovery of 22 multiply-imaged active nuclei. The most successful strategy was adopted in the Sloan Lens ACS Survey (SLACS; Bolton et al., 2008), where the lens candidates were selected spectroscopically by looking for two superimposed galaxy spectra (the lens galaxy and the lensed source) inside a 3 arcsecond diameter fibre. The SLACS survey has provided the largest sample of galaxy-scale lenses to date, with more than a hundred lenses detected and analysed.

However, all current samples of gravitational lenses are statistically limited and have significant observational biases. E.g., most of the known lenses are elliptical galaxies at a redshift below $z = 0.5$ with star-forming lensed sources. Therefore, finding new strong lenses that cover a wider range of redshifts, masses, environments and galaxy types would allow one to probe as vastly as possible the parameter space of galaxies and improve the precision of previous measurements. Moreover, larger samples of lenses will allow comparing the results obtained from lens simulations statistically (e.g., Mukherjee et al. 2018). Deep and high-resolution wide-field surveys have the potential to fill this gap with expected samples of gravitational lenses that are three orders of magnitude larger than current surveys. Indeed ongoing wide-surveys such as the Kilo-Degree Survey (KiDS; see later in this Chapter), the Dark Energy Survey (DES; The

Dark Energy Survey Collaboration 2005), the Subaru Hyper Suprime-Cam Survey (Miyazaki et al. 2012) in the optical and Herschel (Negrello et al., 2010) and the South Pole Telescope (Carlstrom et al. 2011) in the sub-mm, are expected to find samples of lenses of the order of $\sim 10^3$ (see, e.g., Collett 2015). Forthcoming surveys with Euclid (Laureijs et al., 2011), the Large Synoptic Survey Telescope (LSST; LSST Science Collaboration et al. 2009) and the Square Kilometre Array are even expected to discover $\sim 10^5$ new strong lensing systems (Oguri & Marshall, 2010; Pawase et al., 2014; Collett, 2015; McKean et al., 2015).

The increasing depth, resolution and area covered by astronomical surveys, together with the lack of full spectroscopic coverage, has made previous successful observational strategies partially inadequate. On the one hand because of the impossibility of selecting and pre-selecting targets spectroscopically, and on the other hand because of the demanding level of human intervention needed to inspect millions of targets visually. Thus, different strategies based on automatic selection of lens candidates have been developed (e.g., Lenzen et al. 2004; Horesh et al. 2005; Alard 2006; Estrada et al. 2007; Seidel & Bartelmann 2007; Gavazzi et al. 2007; Kubo & Dell’Antonio 2008; More et al. 2012; Maturi et al. 2014; Joseph et al. 2014; Chan et al. 2015). For example, the Canada-France-Hawaii Legacy Survey (CFHTLS) has been exploited with different such methods (Cabanac et al., 2007; More et al., 2012; Paraficz et al., 2016) and also with a crowd-sourcing approach that exploits the visual classifications done by volunteers on the web (Marshall et al., 2016; More et al., 2016). This combined effort has led to the discovery of about a hundred highly probable lenses. Ongoing surveys also represent an opportunity to test automated lens-finders that will become crucial for selecting gravitational lenses among the avalanche of data expected in the coming decade, in particular coming from Euclid, LSST and SKA.

1.3 The Kilo-Degree Survey

The Kilo-Degree Survey (KiDS; de Jong et al. 2015, 2017) is an ESO public survey performed with the OmegaCAM imager at the VLT Survey Telescope (VST; Capaccioli & Schipani 2011) at the Paranal Observatory in Chile. KiDS will survey 1350 square degrees of the sky (to be completed in 2019; Fig. 1.3) subdivided into two patches, the Northern patch, which lies on the celestial equator, and Southern patch which span a portion of sky on the South Galactic Pole. In this way, the observations can take place year-round. KiDS observes in the u , g , r and i filters which are complemented

by the companion VIKING survey carried by the neighbouring VISTA telescope that covers the same area in five near-infrared bands: Z, Y, J, H, K. The surveys have been designed to reach a median galaxy redshift of 0.7. Thus, the KiDS/VIKING combination offers an unprecedented combination of depth and wavelength coverage (u-K, nine bands) which in turn allows astronomers to obtain precise photometric redshifts. Moreover, KiDS will map regions of the sky which overlap with previous redshift surveys as 2dF (Colless et al., 2001), SDSS (Eisenstein et al., 2011) and GAMA Driver et al. (2011). This overlap allows us to have a map of the foreground galaxy population as well as a spectral calibration to better study the evolution of the galaxy population and matter distribution at higher redshifts. KiDS' eye is OmegaCAM (Kuijken, 2011), a 300-million pixel camera composed of 32 CCDs mounted on the VST. OmegaCAM has a field of view of one square degree, which is sampled at a resolution of 0.2 arcsec/pixel. The survey has been designed to obtain images with sub-arcsecond seeing and homogeneous image quality across the full field of view and throughout different observations. Exposure times and observing constraints are reported in Table 1.1. The handling of survey data is operated through the information system Astro-WISE (Valentijn et al., 2007) which allows scientists to archive, calibrate and analyse data in the same single environment. Image quality is a fundamental requirement to achieve the primary science driver of the survey, i.e. to constrain the cosmological parameters of the Universe by measuring the subtle distortions introduced in galaxy shapes by weak lensing (e.g., Hildebrandt et al. 2017).

Gravitational lenses are difficult to be identified due to their scarcity with respect other galaxies and also because their features are often hidden by the Point Spread Function of the instrument (see Fig. 1.4 for a comparison between SDSS and KiDS). The combination of superb image quality and wide area coverage it is the reason why KiDS is also optimal to survey strong gravitational lenses. In fact, a total of ~ 2400 gravitational lenses are expected to be retrievable in the complete survey (see Chapter 2).

1. INTRODUCTION

Table 1.1: KiDS observing strategy: observing condition constraints and exposure times (Table 1 from Kuijken et al. 2019).

Filter	Max. lunar illumination	Min. moon distance [deg]	Max. seeing [arcsec]	Max. airmass	Sky transp.	Dithers	Total Exp. time [s]
<i>u</i>	0.4	90	1.1	1.2	CLEAR	4	1000
<i>g</i>	0.4	80	0.9	1.6	CLEAR	5	900
<i>r</i>	0.4	60	0.8	1.3	CLEAR	5	1800
<i>i</i>	1.0	60	1.1	2.0	CLEAR	5	1200

Figure 1.3: Layout of the KiDS survey fields, KiDS-N in the Northern galactic cap on the equator, and KiDS-S around the South Galactic Pole. Dots indicate where galaxy redshift surveys have already taken place (small dots: SDSS; large dots: 2dFGRS). The combined area of the fields is some 7% of the extragalactic sky, and KiDS can be observed year-round (caption and image from <http://kids.strw.leidenuniv.nl/overview.php>).

1.4 Convolutional Neural Networks

Machine learning research aims at creating algorithms which learn as well as humans do, or better. More specifically, machine learning algorithms “learn” by generalising from examples and make predictions without relying on explicit rules-based programming. Deep Learning algorithms are a particular class of machine learning algorithms which have recently revolutionised the field of machine learning. Deep learning algorithms are feature learning methods, i.e., they automatically build features needed to solve a regression or classification problem; they have a hierarchical architecture inspired by biological nervous systems. The most popular of these algorithms are Convolutional Neural Networks (CNNs; LeCun et al. 1998; Krizhevsky et al. 2012). They are inspired by networks of biological neurons and by the receptive field of the visual cortex of animals. CNNs use a sequence of nonlinear functions for feature extraction and transformation. The algorithms learn multiple levels of representations that correspond to different levels of abstraction. They have been theorised in the ’80s (Fukushima, 1980), but their predictive power has been demonstrated only in the last decade thanks to improvements in hardware, especially in GPUs (Graphics Processing Units) which, in particular, reduce computation times by parallelising the operations among different cores. CNNs are one of the main reasons of the renaissance that the discipline of artificial intelligence is experiencing: they achieved above-human accuracy in many applications and they have been adopted in many commercial applications. A growing concern about ethical and practical implications of these algorithms has led to the emergence of governmental and non-governmental associations dedicated to studying artificial intelligence policy and strategy. To mention some famous applications, the algorithm is at the core of the computer known as AlphaGo (Silver et al., 2017) which won four out of five Go matches against the 18-time world champion Lee Sedol. Moreover, CNNs are the core of the technology of self-driving cars which major companies have been developing in the last years and which will have a dramatic impact on our society.

CNNs were initially being designed to classify and analysing tasks in which they can achieve above-human accuracy. Hence, their exploitation comes naturally in the field of observational astronomy where great part of the scientific research is built around the analysis of astronomical images and the morphology of the astronomical targets is often trivially related to some of their physical properties. Also, CNNs classify images almost instantaneously, are not prone to (human) fatigue and maintain a constant performance over time. These properties make them ideally

Figure 1.4: A couple of examples to illustrate the exceptional image quality of KiDS allows to identify gravitational lenses. On the left column are shown two gravitational lens candidates observed by KiDS (*gri* composed images, both cover 20 by 20 arcseconds); on the right column I show the same galaxies as seen by SDSS (*gri* composed images, 5 arcseconds reference is imprinted on both the images).

suited for classifying the immense amount of data from current and future astronomical surveys. After the first tests for spectral classification (Hála, 2014) and galaxy classification (Dieleman et al., 2015), the algorithm has been applied by Huertas-Company et al. (2015) for classifying galaxy morphology in the CANDELS survey and by the author (Chapter 2) and Jacobs et al. (2017) to detect strong gravitational lenses. Many other successful applications have followed, demonstrating the synergy of the method with astronomical applications. In the following, I will give a short technical introduction on the CNN algorithm.

1.4.1 Algorithm overview

Convolutional Neural Networks take as input data which has a topological structure (e.g., an image) which can be represented as a set of matrices $\mathbf{X}_k$ with $k = 1, 2, \dots, K$ (e.g., the R, G and B components of an image. In this case $K = 3$). The main component of a CNN is the convolutional layer, which takes the inputs and, through a set of two-dimensional filters, produces a stack of *feature maps* $\mathbf{Y}_n$ with n equal to the number of filters. More specifically, every filter in the convolutional layer convolves the input, then a constant value is added to the result and, finally, the output goes through a non-linear function. This procedure is formulated by the following equation:

$$\mathbf{Y} = \sigma \left(\sum_{k=1}^K \mathbf{W}_k * \mathbf{X}_k + \mathbf{B} \right), \quad (1.4)$$

where $*$ is the convolution operator, σ is a non-linear function that is the only source of non-linearity of the algorithm. σ is often implemented as the Rectified linear unit (ReLU; Nair & Hinton 2010) function

$$\sigma(x) = \max(0, x), \quad (1.5)$$

or the sigmoid function

$$\sigma(x) = \frac{1}{(1 + e^{-x})}. \quad (1.6)$$

$\mathbf{W}_k$ are the K weight matrices with $k = 1, 2, \dots, K$, representing a filter with its so-called bias, given by the constant matrix $\mathbf{B}$. Convolutional layers are sequentially stacked, the input of the layers following the input layer are the feature maps created by the preceding layer. A visual example of the result of a convolution of an input image with a filter is shown in Fig 1.5.

Figure 1.5: Visual example of convolving a filter with an input image. The filter swipes through the full image creating a feature map.

Usually, one or more fully-connected layers follow the stack of convolutional layers. Their function is to convert the feature maps created by the last convolutional layer in a single, or more, numbers which represent the outcome of the classification. Fully-connected layers are composed of units which give as output a single number y by performing the following transformation of their inputs

$$y = \sigma(\mathbf{w} \cdot \mathbf{x} + b), \quad (1.7)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a one-dimensional vector created by flattening the feature maps into one dimension or just the output of a previous fully connected layer.

1.4.2 Training

To find the optimal value of the network parameters (i.e., the weights and biases) to classify images, a set of labelled images (called training set) is passed as input of the CNN, this procedure is called training phase. Finding these values is achieved by optimising the weights and biases of the network. In this way, the CNN "learns" complex functions and how to extract features from the data that are not hand designed but are learned during the training stage. After the training procedure, the CNN can be used for classifying new data by keeping its parameters fixed. The matrix values representing CNN filters are automatically modified during the training phase in order to extract features useful for the purpose of the classification problem. Deeper filters extract increasingly more abstract features. So the features learned by CNNs depend on the labelled data available during the training phase and are not manually chosen (see Fig 1.6 for example).

For example, let us consider the problem of classifying pictures of animals in one of the different n animal categories. We want the output layer of the CNN $\mathbf{y} = (y_1, y_2, \dots, y_n)^T$ to approximate the desired output $\hat{\mathbf{y}} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_n)^T$ which represents some measure of probability that the input image is one of the n different categories. The weights and biases are optimized by minimizing a chosen loss function $L(\mathbf{y}, \hat{\mathbf{y}})$ via the iterative process of *gradient descent*, i.e., each weight and bias of the CNN is updated following the negative direction of the gradient:

$$w \rightarrow w' = w - \eta \frac{\partial L}{\partial w}, \quad b \rightarrow b' = b - \eta \frac{\partial L}{\partial b}, \quad (1.8)$$

where η is a constant called the *learning rate*, the larger learning rates imply larger updates. The gradients are computed via the back-propagation algorithm (Rumelhart et al., 1986).

Figure 1.6: Picture of a cat (top row); feature maps extracted by 5 different consecutive layers of a CNN (second to last row). Deeper layers of the CNN extract feature maps of increasing level of abstraction (image from <https://towardsdatascience.com>).

1.5 This thesis

This thesis aims to develop a method to automatically identify strong gravitational lens candidates in astronomical image surveys in order to enable individual and ensemble studies of these scientifically valuable astronomical objects. To accomplish this goal, I have designed a machine learning algorithm based on Convolutional Neural Networks (CNNs) with the purpose of classifying galaxy images from the Kilo-Degree Survey as either *lenses* or *non-lenses*. The algorithm has aided the discovery of a few hundred new gravitational lens candidates.

In **Chapter 2**, I design a CNN lens-finder and train it using r -band images of KiDS galaxies and simulated strong gravitational lensed sources. I then apply the CNN to a sample of $\sim 20\,000$ colour-magnitude selected galaxies from ~ 255 square degrees. The CNN successfully identifies two out of three pre-confirmed lenses that are in the sample together with another ~ 50 reliable lens candidates.

In **Chapter 3**, I implement two different CNN lens-finders, one for classifying r -band images while the second for classifying gri -composite images. The CNNs have a better performance compared to the CNN developed in Chapter 2. The reason rests on using an improved version of the algorithm and especially in training the algorithms with a vastly improved set of simulated lensed sources.

In **Chapter 4**, I present a sample of many hundreds of gravitational lens candidates: the LinkS (Lenses in KiDS) sample. This sample is selected by applying the two CNNs developed in Chapter 3 to galaxies selected from ~ 900 square degrees of KiDS and then pruned and ordered by visual inspection. Besides, I find that it will be possible to select thousands of lens candidates in surveys exploited with Euclid and LSST by using CNNs and minimal human intervention.

Finally, in **Chapter 5**, I summarise the main conclusions of this thesis, and I provide an outlook for future plans and improvements.

The 9000 series is the most reliable computer ever made. No 9000 computer has ever made a mistake or distorted information. We are all, by any practical definition of the words, foolproof and incapable of error

– HAL 9000, *2001: A Space Odyssey*

2. FINDING STRONG GRAVITATIONAL LENSES IN THE KILO DEGREE SURVEY WITH CONVOLUTIONAL NEURAL NETWORKS

C. E. Petrillo, C. Tortora, S. Chatterjee, G. Vernardos
L. V. E. Koopmans, G. Verdoes Kleijn,
N. R. Napolitano, G. Covone,
P. Schneider, A. Grado, J. McFarland

2017, *Monthly Notices of the Royal Astronomic Society*, 472, 1129

The volume of data that will be produced by new-generation surveys requires automatic classification methods to select and analyze sources. Indeed, this is the case for the search for strong gravitational lenses, where the population of the detectable lensed sources is only a very small fraction of the full source population. We apply for the first time a morphological classification method based on a Convolutional Neural Network (CNN) for recognizing strong gravitational lenses in 255 square degrees of the Kilo Degree Survey (KiDS), one of the current-generation optical wide surveys. The CNN is currently optimized to recognize lenses with Einstein radii $\gtrsim 1.4$ arcsec, about twice the r -band seeing in KiDS. In a sample of 21789 colour-magnitude selected Luminous Red Galaxies (LRG), of which three are known lenses, the CNN retrieves 761 strong-lens candidates and correctly classifies two out of three of the known lenses. The misclassified lens has an Einstein radius below the range on which the algorithm is trained. We down-select the most reliable 56 candidates by a joint visual inspection. This final sample is presented and discussed. A conservative estimate based on our results shows that with our proposed method it should be possible to find ~ 100 massive LRG-galaxy lenses at $z \lesssim 0.4$ in KiDS when completed. In the most optimistic scenario this number can grow considerably (to maximally ~ 2400 lenses), when widening the colour-magnitude selection and training the CNN to recognize smaller image-separation lens systems.

2.1 Introduction

Strong gravitational lensing is a rare phenomenon which provides very tight constraints on the projected mass of the foreground lens galaxy. In fact, the total mass (dark plus baryonic) within the Einstein radius depends almost solely on the space-time geometry of the lensing system (the source and the lens redshift and the cosmological parameters). For this reason, strong lensing is a unique tool, if combined with central velocity dispersion measurements and stellar population analysis, to estimate the fraction of dark matter in the central regions of galaxy-scale halos (e.g., Gavazzi et al. 2007; Jiang & Kochanek 2007; Grillo et al. 2010; Cardone et al. 2009; Cardone & Tortora 2010; Tortora et al. 2010; More et al. 2011; Ruff et al. 2011; Sonnenfeld et al. 2015), and to constrain the slope of the inner mass density profile (e.g., Treu & Koopmans 2002a,b; Koopmans et al. 2006; Koopmans & Treu 2003; More et al. 2008; Barnabè et al. 2009; Koopmans et al. 2009; Cao et al. 2016).

Gravitational lenses can be also used to constrain the stellar initial mass function (e.g., Treu et al. 2010; Ferreras et al. 2010; Spiniello et al. 2011; Brewer et al. 2012; Sonnenfeld et al. 2015; Posacki et al. 2015; Leier et al. 2016) and to independently measure the Hubble constant through time delays (e.g., Suyu et al. 2010; Bonvin et al. 2017). In addition, strong lensing gives magnified views of background objects otherwise inaccessible to observations (e.g., Impellizzeri et al. 2008; Swinbank et al. 2009; Richard et al. 2011; Deane et al. 2013; Treu et al. 2015; Mason et al. 2017).

A homogeneously selected large lens sample can improve dramatically the effectiveness of the methods and the reliability of the results from gravitational lensing studies. The largest homogeneous sample so far is provided by the Sloan Lens ACS Survey (SLACS; Bolton et al. 2008) with almost 100 observed lenses. In the future, deep high resolution wide surveys have the potential to produce samples three orders of magnitude larger than the current known lenses. These large numbers will allow to, e.g., greatly improve the precision in the mass density slope measurements (Barnabè et al., 2011), in better estimate the presence of substructure (Vegetti & Koopmans, 2009) and to put constraints on the nature of dark matter (Li et al., 2016).

Upcoming telescopes, such as Euclid (Laureijs et al., 2011) and the Large Synoptic Survey Telescope (LSST; LSST Science Collaboration et al. 2009), will increase the rate of discovery of new lenses, reaching the number of $\sim 10^5$ new strong lensing systems (Oguri & Marshall, 2010; Pawase et al., 2014; Collett, 2015). Also, the number of lenses that will

be observed by the Square Kilometer Array is expected to be of the same order of magnitude (McKean et al., 2015). The ongoing optical wide surveys, such as the Kilo Degree Survey (KiDS; see Sec. 2.2), the Dark Energy Survey (DES; The Dark Energy Survey Collaboration 2005) and the Subaru Hyper Suprime-Cam Survey (Miyazaki et al. 2012) are expected to find samples of lenses of the order of $\sim 10^3$ (see, e.g. Collett 2015). Sub-mm observations from Herschel (Negrello et al., 2010) and the South Pole Telescope (Carlstrom et al. 2011), together with deeper, high resolution observations from the the Atacama Large Millimeter/sub-millimeter Array, are expected to provide several hundred new lenses as well.

Traditionally, the search of extended lens features (i.e., arcs and rings) relied heavily on the visual inspection of the targets. This is still the best approach for small samples of objects, but is impractical for the ongoing and new generation surveys given the large number of targets that need to be inspected. Accordingly, numerous automatic lens finders have been developed in recent years. Most are based on the identification of arc-like shapes (e.g., Lenzen et al. 2004; Horesh et al. 2005; Alard 2006; Estrada et al. 2007; Seidel & Bartelmann 2007; Kubo & Dell’Antonio 2008; More et al. 2012). The same approach, together with a colour selection, is employed by Maturi et al. (2014). Another method consists of subtracting the light of the central galaxies using multiband images and then analyse the image residuals (Gavazzi et al., 2014). Joseph et al. (2014) follow a similar approach but employing machine-learning techniques to analyse single-band images. Instead Brault & Gavazzi (2015) model the probability that the targets are actual lenses. Very recently Bom et al. (2017) have developed an artificial neural network for recognizing strong lenses that uses as entries a set of morphological measurements of the targets. A completely different approach based on crowdsourcing is employed in the Space Warps project (Marshall et al., 2016; More et al., 2016), with volunteers visually inspecting and classifying galaxy cutouts through a web applet¹. All these automatic methods have their advantages and disadvantages and perform at their best with different typologies of lenses, quantity and kind of data available. A detailed comparison between these methods should be done on a common dataset, but is beyond the scope of this paper.

Convolutional Neural Networks (CNNs; Fukushima 1980; LeCun et al. 1998) are a state of the art class of machine learning algorithm particularly suitable for image recognition tasks. The ImageNet Large Scale Visual Recognition Competition (ILSVRC; Russakovsky et al. 2015; the most

¹<https://spacewarps.org/>

important image classification competition) of the last four years has been won by groups utilizing CNNs. The advantage of CNNs with respect to other pattern recognition algorithms is that they automatically define and extract representative features from the images during the learning process. Although the theoretical basis of CNNs was built in the the 1980s and the 1990s, only in the last years do CNNs generally outperform other algorithms due to to the advent of large labelled datasets, improved algorithms and faster training times on e.g. Graphics Processing Units (GPUs). We refer the interested reader to the reviews by Schmidhuber (2015), LeCun et al. (2015) and Guo et al. (2016) for a detailed introduction to CNNs.

The first application of CNNs to astronomical data was made by Hála (2014) for classifying spectra in the Sloan Digital Sky Survey (SDSS; Eisenstein et al. 2011). Then, Dieleman et al. (2015)² used CNNs to morphological classify SDSS galaxies. Subsequently, Huertas-Company et al. (2015) used the same set-up of Dieleman et al. (2015) for classifying the morphology of high- z galaxies from the Cosmic Assembly Near-IR Deep Extragalactic Legacy Survey (Grogin et al., 2011). More recently, Hoyle (2016) adopted CNNs for estimating photometric redshifts of SDSS galaxies. CNNs have been employed also by Kim & Brunner (2017) for star/galaxy classification.

In this paper we present our morphological lens-finder which is based on CNNs. We apply it to the third data release of KiDS (de Jong et al., 2015, 2017), starting a systematic census of strong lenses. This project, which consists of both visual and automatic inspection of the KiDS images, is dubbed "Lenses in KiDS" (LinKS). KiDS is a particularly suitable survey for finding strong lenses, given its excellent seeing and pixel scale, in addition to the large sky coverage (see Sect. 2.2).

The paper is organized as follows. In Sect. 2.2 we provide a brief description of the KiDS survey and the way in which we select the LRG-galaxy sample used in this work. In Sect. 2.3 we illustrate our lens-finding CNN-based algorithm and how we build the training data set. In Sect. 2.4 we explain how we apply our method to ~ 255 square degrees of KiDS, present the list of our new lens candidates, compare it with the literature and with a forecast of the expected number of detectable strong gravitational lenses in the survey and do a consistency check of the observed Einstein radii of the candidates to select the most reliable ones. Finally, in Sect. 2.5, we provide a summary, the main conclusion of this work and a short outlook for future plans and improvements. In the following

²The method won a challenge against other techniques
<https://www.kaggle.com/c/galaxy-zoo-the-galaxy-challenge/>

we adopt a cosmological model with $(\Omega_m, \Omega_\Lambda, h) = (0.3, 0.7, 0.75)$, where $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

2.2 The KiDS survey

The Kilo-Degree Survey (KiDS) (de Jong et al., 2015) is one of the three ESO public surveys carried out using the OmegaCAM wide-field imager (Kuijken 2011) mounted at the Cassegrain focus of the VLT Survey Telescope (VST; Capaccioli & Schipani 2011) at Paranal Observatory in Chile. OmegaCAM is a 256 Megapixel camera containing 32 science CCD detectors which cover a one square degree field of view at a pixel-size of 0.21 arcsec. The VST is a 2.6m telescope with active control of the primary and secondary mirror which is driven by wave-front sensing via two auxiliary CCDs in OmegaCAM. In this way, the camera-telescope combination is specifically designed to obtain sharp and homogeneous image quality over the wide field of view. KiDS is a 1500 square degree extra-galactic imaging survey in four optical bands (u , g , r and i). The survey area is divided over an equatorial patch and a Southern patch around the South Galactic Pole. Observations are queue scheduled, reserving the best seeing for the r -band which has a median FWHM of the PSF of 0.65 arcsec with a maximum of 0.8 arcsec. Median PSF FWHM values in u , g and i are 1.0 arcsec, 0.8 arcsec and 0.85 arcsec, respectively. KiDS reaches limiting magnitudes (5- σ AB in a 2 arcsec aperture) of 24.3, 25.1, 24.9 and 23.8 in u , g , r and i band, respectively. The primary science driver for the survey design is the study of the dark matter distribution over cosmological volumes via weak-lensing tomography. Strong-lensing survey studies are a particularly suitable science case as well, because they exploit the combination of superb image quality and wide survey area.

2.2.1 Data Release Three

In this paper we make use of the most recent public data release (KiDS ESO-DR3, de Jong et al. 2016, in prep). It consists of the co-added images, weight maps, masks, single-band and multi-band catalogues and photometric redshifts for 292 survey tiles. We use the multi-band photometry based on r -band detections, with a total of 33 million unique sources. Our data handling and scientific data analysis is performed using the Astro-WISE information system (Valentijn et al., 2007). The source extraction and related photometry have been obtained with S-EXTRACTOR (Bertin & Arnouts 1996). We rely on both aperture photometry and the

Figure 2.1: colour $g - r$ versus photometric redshift. The g and r values are MAG_AUTO magnitudes and the photometric redshift is obtained with BPZ. The dots are sources from KiDS DR3. Shown are (i) extended objects with MAG_AUTO in r -band less than 20 (blue), (ii) objects that satisfy the Eisenstein et al. (2001) colour-magnitude selection (red) and (iii) objects selected with our expanded colour-magnitude selection (green). See Sect. 2.2.2 for the details.

Kron-like MAG_AUTO. A relevant output parameter of S-EXTRACTOR is the FLAGS parameter. We set the r -band FLAGS to be < 4 , to only include de-blended sources and remove from the catalogues those objects with incomplete or corrupted photometry, saturated pixels or any other kind of problem encountered during de-blending or extraction. Critical areas as saturated pixels, star spikes and reflection halos have been masked using a dedicated automatic procedure (PULECENELLA). The IMA_FLAGS flags store the result of this masking operation: sources that are not in critical regions have this parameter set to 0. Photometric redshifts are determined using the program BPZ (Benítez 2000), which is a Bayesian photo- z estimator based on a template fitting method (see de Jong et al. 2017, in prep., for further details). The unmasked effective area adopted, considering the sources with IMA_FLAGS = 0 in all the KiDS-DR3 bands, is 255 square degrees.

2.2.2 Luminous red galaxy sample

We select Luminous Red Galaxies (LRGs; Eisenstein et al. 2001) from the 255 square degrees of the KiDS-ESO DR3 for the purpose of both training our CNN and searching for lens candidates among them. LRGs are very massive and hence more likely to exhibit lensing features compared to other classes of galaxies ($\sim 80\%$ of the lensing population; see Turner et al. 1984; Fukugita et al. 1992; Kochanek 1996; Chae 2003; Oguri 2006; Möller et al. 2007). We focus on this kind of galaxies in this work and will consider other kind of galaxies in the future. The selection is made with the following criteria where all the parameters are from S-EXTRACTOR and magnitudes are MAG_AUTO:

- (i) The low- z ($z < 0.4$) LRG colour-magnitude selection of Eisenstein et al. (2001), adapted to including more sources (fainter and bluer):

$$\begin{aligned} r &< 20 \\ |c_{\text{perp}}| &< 0.2 \\ r &< 14 + c_{\text{par}}/0.3 \end{aligned} \tag{2.1}$$

where

$$\begin{aligned} c_{\text{par}} &= 0.7(g - r) + 1.2[(r - i) - 0.18] \\ c_{\text{perp}} &= (r - i) - (g - r)/4.0 - 0.18 \end{aligned}$$

- (ii) A source size in the r -band larger than the average FWHM of the PSF of the respective tiles, times a empirical factor to maximize the separation between stars and galaxies.

This final selection provides an average of 74 LRGs per tile and a total of 21789 LRGs. We refer to this sample as the "LRG sample" in the remainder of the paper. Compared to the original colour-magnitude selection for $z < 0.4$ (Eisenstein et al., 2001), we obtain ~ 3 times more galaxies. A colour-photo- z diagram of the results of the two different cuts is shown in Fig. 2.1 for illustration.

2.3 Training the CNN to find Lenses

Our lens finder is based on a Convolutional Neural network (CNN) and is inspired by the work of Dieleman et al. (2015). CNNs are supervised deep learning algorithms (see the recent reviews from Schmidhuber 2015; LeCun et al. 2015; Guo et al. 2016) particularly effective for image recognition tasks (see e.g., He et al. 2015b, winner of the last ILSVRC competition; Russakovsky et al. 2015) and regression tasks, such as, in the astronomical domain, the determination of galaxy morphologies (Dieleman et al., 2015; Huertas-Company et al., 2015). The algorithm converts sequentially the input data through non-linear transformations whose parameters are learned in the training phase. A set of labelled images (the training set) are used as input of the CNN in this phase. The network changes its parameters by optimizing a loss function that expresses the difference between its output and the labels of the images in the training set. This allows the CNN to learn complex functions and to extract features from the data that are not hand designed but are learned during the training stage. After the training procedure the CNN can be used for classifying new data by keeping its parameters fixed. For the interested reader, in Appendix 2.A we shortly introduce the technical background of CNNs that are relevant to some of the choices made in this paper.

2.3.1 Input Samples

Finding strong gravitational lenses can be reduced to a two-class classification problem, where the two kinds of objects to recognize are the lenses and the non-lenses. Training a Convolutional Neural Network (CNN) to solve this task requires a dataset representative of the two classes called *training set*. It has to be large enough because of the large number of parameters of a CNN (usually of the order of 10^6). In the case of strong gravitational lenses we do not have a large enough representative data-set at our disposal. The largest sample available is collected in The

Table 2.1: The range of values adopted for the model parameters of the lens and source. See Sect. 2.3.1.2 for further details.

Parameter	Range	Unit
Lens (SIE)		
Einstein radius	1.4 - 5.0	arcsec
Axis ratio	0.3 - 1.0	-
Major-axis angle	0.0 - 180	degree
External shear	0.0 - 0.05	-
External-shear angle	0.0 - 180	degree
Source (Sérsic)		
Effective radius	0.2 - 0.6	arcsec
Axis ratio	0.3 - 1.0	-
Major-axis angle	0.0 - 180	degree
Sérsic index	0.5 - 5.0	-

Masterlens Database³. Unfortunately, this sample can not be used as a training set for our purpose, since it is small and heterogeneous. It consists of 657 lens systems that are not all spectroscopically confirmed, that have been discovered in various surveys and programs, or that are observed at different wavelengths according to the instrument used.

For these reasons, we build a set of mock lens systems, relying on a hybrid approach: first we select real galaxies, with their fields, obtained from KiDS (Sect. 2.3.1.1), in order to include seeing, noise and especially the lens environment that is a feature hard to simulate and its omission would limit the ability of the network to recognize lenses in real survey data. Then we independently simulate the lensed sources (Sect. 2.3.1.2) and combine them with the real galaxies (Sect. 2.3.2).

We limit our training to r -band images, where KiDS provides the best image quality (an average FWHM of 0.65 arcsec). Hence, the network will learn selection criteria mostly based on the morphology of the sources. We plan to ingest multi-wavelength data into the network in future improvements, allowing the training on the differences in colours. Our training set consists of images of lens and non-lens examples produced with r -band KiDS images of real galaxies (see Sect. 2.3.1.1) and mock gravitational lensed sources (see Sect. 2.3.1.2). In Sect. 2.3.2 we summarize how the actual positive (lenses) and negative examples (non-

³<http://masterlens.astro.utah.edu/>

Figure 2.2: Several examples of simulated lensed sources produced as described in Sect. 2.3.1.2. The image size is 101 by 101 pixels, corresponding to 20 by 20 arcsec.

lenses) employed in the training of the network, are produced. We train our CNN on a set of six million images (three million lenses and three million non-lenses with labels 1 and 0, respectively). Our trained CNN gives as output a value p ranging between 0 and 1. The sources with an output value of p larger than 0.5 are classified as lenses. The technical details of our implementation and the training procedure can be found in Appendix 2.B, providing further background to our procedures and choices. We further expand our training set using data augmentation techniques (Sect. 2.3.3).

2.3.1.1 Real Galaxy Sample

We select a sub-sample of the KiDS LRGs (see Sect. 2.2.2) consisting of 6554 galaxies (a third of the full sample), which we have visually inspected finding 218 contaminants, mostly face-on spirals. Additionally, we have collected a sample of 990 sources wrongly classified as lenses in previous tests with CNNs. We use this sample in the training set to reject clear outliers. The 6326 LRGs, the 218 contaminants and the 990 false positives constitute together the non-simulated part of the data used to build the training set. We will refer to it as the real galaxy sample in the remaining of this paper.

2.3.1.2 Mock Lensed-Source Sample

The mock lensed source sample is composed by 10^6 simulated lensed images of 101 by 101 pixels, using the same spatial resolution of KiDS (0.21 arcsec per pixel), corresponding to a 20 by 20 arcsec field of view. We produce the different lensed image configurations by sampling uniformly the parameters of the lens and source models listed in Table 2.1. A few examples are shown in Fig. 2.2. The choice of uniformly sampling the parameter space does not reproduce the distribution of the parameters for a real lens population, but allows the classifier to learn the features for recognizing the different kinds of lenses, no matter how likely they are to appear in a real sample of lenses.

We model the sources with a Sérsic (1968) profile and the lenses with a Singular Isothermal Ellipsoid (SIE; Kormann et al. 1994) model. At source redshifts of $z > 0.5$, smaller sizes and smaller Sérsic indices are found with respect to the local universe, and the fraction of spiral galaxies (with $n < 2 - 3$) increases (e.g. Trujillo et al. 2007; Chevance et al. 2012). We exclude spiral galaxy sources or very elliptical ones considering only axis ratios > 0.3 . The source positions are chosen uniformly within the radial distance of the tangential caustics plus one effective radius of the source Sérsic profile. This leads our training set to be mostly composed of high-magnification rings, arcs, quads, folds and cusps rather than doubles (Schneider et al., 1992) that are harder to distinguish from companion galaxies and other environmental effects. In this paper our first-order goal is to find the larger, brighter and more magnified strong lenses, rather than aim for completeness over the full parameter space of lenses.

The upper limit of 5 arcsec for the Einstein radius aims to include typical Einstein radii for strong galaxy-galaxy and group-galaxy lenses (Koopmans et al. 2009; Foëx et al. 2013; Verdugo et al. 2014). The lower limit is chosen to be 1.4 arcsec, about twice the average FWHM of the r -band KiDS PSF. Because lenses are typically early-type galaxies, which do not have high ellipticity, we choose 0.3 as a lower limit of the axis ratio (Binney & Merrifield 1998). We set the external shear to less than 0.05, higher than typically found for SLACS lenses (Koopmans et al. 2006) with a random orientation varying between 0 and 180 degrees.

2.3.2 Building the training examples

Each training image passed to the network is built as described below and as summarized schematically in Fig. 2.3.

Mock lenses (positive sample): To create the mock lenses we carry out the following procedure: (i) we randomly choose a mock lensed source from the mock source sample and a LRG from the real galaxy sample (Sections 2.3.1.2 and 2.3.1.1, respectively); (ii) we randomly perturb both the mock source and the LRG as described in Sect. 2.3.3; (iii) we rescale the peak brightness of the simulated source between 2% and 20% of the peak brightness of the LRG. In this way we take into account the typical lower magnitudes of the lensing features with respect to the lens galaxies despite the magnification; (iv) we add the two resulting images; (v) we clip the negative values of the pixels to zero and performing a square-root stretch of the image to emphasize lower luminosity features; and (vi) finally we normalize the resulting image by the peak brightness. This procedure can yield a-typical lens configurations, because the mock sources and the KiDS galaxies are combined randomly, without taking into account the physical characteristics of the galaxies. Nevertheless, we operate in this way with the intent to train the network to classify a lens largely relying on the morphology of the source. Moreover, we reduce the risk of over-fitting, because the probability that the network will see twice the same (or a very similar) example is negligible. In addition, we cover the parameter space as free from priors as possible, which could allow to find less conventional lens configurations as well.

Non-lenses (negative sample): To create the mock non-lens sample we carry out the following procedure: (i) we randomly choose one galaxy from the real galaxy sample (see Sect. 2.3.1.1) with a 60% probability of extracting a LRG and 40% probability to extract a contaminant or false positive; (ii) we randomly perturbing as in Sect. 2.3.3; (iii) we apply a square-root stretch of the image; (iv) we normalizing the image by the peak brightness.

The final inputs of the convolutional neural network are image-cutouts of 60 by 60 pixels which correspond to ~ 12 by 12 arcsec. These images are produced in real-time during the training phase.

2.3.3 Data augmentation

A common practice in machine learning is data augmentation: a procedure used to expand the training set in order to avoid over-fitting the data and teaching the network rotational, translational and scaling invariance (see e.g., Simard et al. 2003). We augment our dataset applying the following transformations to the mock lensed images and the real galaxy

Figure 2.3: A schematic of the training-set creation. For the non-lens examples we use real KiDS image-cutouts of LRGs and other galaxies (see Sect. 2.3.1.1). For producing the lens examples we mix KiDS LRGs and simulated mock lensed sources (Sect. 2.3.1.2). In the process the images are augmented and preprocessed as explained in Sections 2.3.3 and 2.3.2.

Figure 2.4: RGB images of 20 by 20 arcsec of some contaminants classified as lenses by the CNN.

sample: (i) a random rotation between 0 and 2π ; (ii) a random shift in both x and y direction between -4 and +4 pixels; (iii) a 50% probability of horizontally flipping the image; (iv) a rescaling with a scale factor sampled log-uniformly between 1/1.1 and 1.1. All transformations are applied to the image-cutouts of 101 by 101 pixels of both the real galaxy and mock lensed source sample. We extract a central region of 60 by 60 pixels from the resulting images to avoid unnecessarily information (i.e., noise and empty sky) around the image edges.

2.4 Results

Having trained the CNN as described in Sect. 2.3 (see also Appendix 2.B for more details), in this section we present our results. In Sect. 2.4.1 we report the procedure to select our final sample of lens candidates and in Sect. 2.4.2 the sample is presented, discussed and compared with the literature.

2.4.1 Candidate selection

First we ingest the full 21789 LRG sample (see Sect. 2.2.2) into the trained CNN. We obtain 761 galaxies ($\sim 3\%$ of the full LRG sample) classified as lens candidate with $p > 0.5$ and all the remainder in the non-lens category with $p < 0.5$. The number of LRG classified by the network as lenses is too large when compared to the expected number of strong lenses in the KiDS-DR3 area (see Sect. 2.4.2.1). Among the selected sources there are contaminants such as spirals, galaxies with dust lanes, mergers, etc. (see Fig. 2.4 for some examples). For this reason we decide to further visually classify the 761 targets selected by the network. Seven of the authors of this paper – referred as “classifiers” in the following – are presented with a set of images for each lens candidate: the cut-out images from KiDS (one image per each of the u , g , r , and i filters) and a RGB reconstructed composite image obtained with the software STIFF⁴ from the g , r , and i -band images. The classifiers can classify the sources in three categories: *Sure*, *Maybe*, and *No lens*. The score for each candidate is based on the following scheme:

<i>Sure lens</i>	10 points.
<i>Maybe lens</i>	4 points.
<i>No lens</i>	0 points.

⁴<http://www.astromatic.net/software/stiff>

The histogram of the accumulated grades of the visual classification is shown in Fig. 2.5. There are 384 candidates classified in the *Sure* and *Maybe* categories by at least one classifier. To further reduce the sample, we decide to introduce a threshold at the score of 17, below which all candidates are considered not reliable. This implies that more than four classifiers would be required to classify a lens candidate in the *Maybe* category to be regarded as reliable. For lenses in the *Sure* category we expect a large number of users to agree in their classification due to more evident lensing features in the images, giving a higher score to such candidates. Only two candidates achieved the maximum score of 70. As seen in Fig. 2.5 (blue bars), the distribution of candidates rises rapidly below the threshold score and remains flat for higher values. Changing the points given to a candidate classified as *Maybe* lens from four to six, and appropriately relocating the threshold, does not affect the resulting ranking, and the distribution shown in Fig. 2.5 remains largely the same.

Since the focus of this paper is to find new lens candidates, we are interested in the first two categories, i.e. *Sure* and *Maybe*. However, we plan for future applications to use the candidates classified in the *No* category to retrain the CNN, aiming at considerably reducing the number of candidates that need to be visually inspected.

2.4.2 Final sample of candidates

After both CNN and visual classification, the final sample of lens candidates consists of 56 objects, down-selected from an initial sample of 21789 galaxies. In Fig. 2.6 we show how the candidates are distributed in colour-photo- z space together with the full LRG sample (Sect. 2.2.2). In Fig. 2.11 the RGB images of these best candidates are shown together with their scores from the visual inspection procedure. For completeness, in Appendix 2.C the r -band-only images of the 56 ranked objects are also shown, since they are the images on which the CNN has made its classification. Candidates are listed in Table 2.2, where we show the final grade of our classification, the KiDS MAG_AUTO in the u , g , r , and i bands for each candidate, together with the BPZ photometric redshift, stellar mass and, if available, spectroscopic redshift and velocity dispersion.

J085446-012137 and J114330-014427 are successfully classified as lenses by our network and they pass our visual inspection with a score of 70 and 60 respectively (KSL317 and KSL040 in Table 2.2 and Fig. 2.11). Instead, J1403+0006 is classified as a non-lens by the network, this could be due to the fact that this system has an Einstein radius of 0.83 arcsec, well below the lower limit of the interval of radii on which the CNN is

Figure 2.5: Histogram of the ranking of 384 lens candidates, which have been classified at least by one user in the *Sure* or *Maybe* categories. In blue are the candidates with a score higher than 16 that are considered the most reliable.

trained. In Fig. 2.7 we show the RGB images of these three known lenses as observed in KiDS. The lensed images of the misclassified lens are also not as prominent as in the other two.

We find that 34 of our candidates have spectra measured from different sources (2dF, Colless et al. 2001; Limousin et al. 2010; SDSS, Eisenstein et al. 2011; BOSS, Dawson et al. 2013; GAMA, Liske et al. 2015). We visually inspected the spectra without clearly identifying any emission line that could belong to a background source. A more detailed data reduction of the spectra is needed to confirm or discard any of these candidates. We also notice that the photometric redshifts tend to overestimate the distance. This could be due to the contamination of the colours of the main galaxy by the supposed lensed sources. We will investigate this issue in a forthcoming paper.

Table 2.2: The final sample of candidates. For each candidate we report an internal ID; the spectroscopic redshift if available (see notes), the BPZ photometric redshift, the KiDS u , g , r and i MAG AUTO (average uncertainties are 0.13, 0.02, 0.03 and 0.04 respectively), the velocity dispersion from SDSS or BOSS if available, the stellar mass (the typical uncertainty is ~ 0.2 dex). A double (single) check-mark indicates the candidates with (without) a measured velocity dispersion that have a predicted Einstein radius comparable with their galaxy-image configuration (see Sect. 2.4.2).

ID	z_{spec}	z_{phot}	u	g	r	i	σ_* [km/s]	$\log M_*/M_\odot$ [dex]	score	
KSL427	0.24 ²	0.25	20.17	18.38	16.96	16.58		11.3	70	✓
KSL317	0.35 ⁴	0.42	21.43	19.61	17.86	17.20		11.6	70	
KSL103	0.24 ²	0.26	20.55	18.53	17.28	16.81		11.3	64	
KSL040	0.11 ¹	0.15	18.56	16.65	15.64	15.26	269 ± 5	11.1	60	✓✓
KSL627	0.21 ¹	0.24	20.89	18.75	17.42	16.92	206 ± 13	11.3	60	
KSL327	0.12 ²	0.17	18.96	16.80	15.75	15.33		11.4	58	✓
KSL376	0.30 ¹	0.36	21.88	20.03	18.46	17.91	242 ± 20	11.2	48	
KSL086		0.33	22.30	20.23	18.54	18.03		11.1	48	
KSL351	0.26 ¹	0.30	21.12	18.89	17.42	16.91	278 ± 19	11.4	46	✓✓
KSL469	0.29 ¹	0.33	21.23	19.52	18.08	17.51	228 ± 19	11.4	46	
KSL228	0.18 ²	0.16	19.83	18.30	17.25	16.75		11.2	42	
KSL713	0.23 ²	0.29	20.46	18.54	17.01	16.47	304 ± 17	11.5	42	✓✓
KSL328	0.23 ¹	0.24	21.28	19.52	18.09	17.58	235 ± 13	11.0	42	
KSL411	0.25 ¹	0.27	18.62	17.59	16.66	16.22		11.5	40	✓
KSL070	0.44 ¹	0.45	21.55	20.37	19.05	18.37	206 ± 37	11.2	40	
KSL543		0.25	20.73	19.03	17.78	17.29		11.3	38	
KSL664		0.30	21.89	19.91	18.62	18.00		11.1	36	✓
KSL106	0.27 ³	0.28	21.75	19.96	18.59	18.06		11.1	36	
KSL337		0.35	22.06	20.55	18.93	18.41		10.8	32	

KSL388	0.33 ¹	0.37	22.83	19.81	18.14	17.58	228 ± 20	11.4	32	
KSL415	0.21 ¹	0.21	20.64	19.03	17.68	17.19	223 ± 17	11.3	32	
KSL220		0.31	21.85	20.32	18.91	18.42		11.2	30	✓
KSL601	0.46 ¹	0.54	23.09	21.41	19.65	18.97	221 ± 22	11.1	28	
KSL603	0.34 ¹	0.41	21.96	20.04	18.46	17.86	220 ± 16	11.5	28	
KSL436		0.27	21.44	19.57	17.90	17.74		10.8	26	
KSL233	0.15 ²	0.17	19.96	18.17	17.16	16.67		10.8	26	
KSL231		0.46	23.73	20.98	19.36	18.64		11.3	26	
KSL101		0.32	22.92	20.36	18.89	18.33		11.2	26	✓
KSL450	0.40 ¹	0.46	23.24	20.83	18.97	18.44	270 ± 30	11.2	26	✓✓
KSL737	0.37 ¹	0.45	22.27	20.39	18.74	18.14	222 ± 38	11.4	26	✓✓
KSL094	0.29 ¹	0.42	20.77	19.36	17.94	17.37	219 ± 19	11.5	26	✓✓
KSL669	0.05 ¹	0.16	19.46	17.58	16.63	16.17	212 ± 8	10.4	26	
KSL707		0.25	21.77	19.66	18.28	17.72		11.0	24	
KSL197		0.21	21.63	19.56	18.14	17.71		11.1	24	
KSL335	0.22 ²	0.26	21.89	19.07	17.72	17.19		11.2	24	
KSL565	0.29 ¹	0.29	21.52	19.97	18.52	17.97	251 ± 18	11.1	24	✓✓
KSL134	0.27 ¹	0.29	21.23	19.57	18.24	17.73	235 ± 14	11.3	24	
KSL606	0.18 ²	0.17	19.79	18.32	17.25	16.85		11.2	22	
KSL046		0.12	18.44	16.69	15.79	15.40		11.4	22	✓
KSL620		0.28	21.53	19.78	18.47	17.91		11.1	22	
KSL013	0.11 ²	0.14	19.73	18.19	17.26	16.74		10.7	22	
KSL421		0.32	21.42	19.65	18.23	17.69		11.4	22	✓
KSL434		0.39	21.74	20.06	18.47	17.99		11.0	22	
KSL516		0.56	23.33	21.19	19.55	18.80		11.1	20	

KSL278		0.42	21.85	20.34	18.84	18.32		11.2	20	✓
KSL178		0.43	21.74	20.04	18.31	17.63		11.9	20	
KSL159		0.44	22.08	20.32	18.68	18.08		11.7	20	
KSL686	0.25 ³	0.30	21.77	19.24	17.72	17.20		11.2	20	
KSL465		0.34	20.49	19.17	17.69	17.17		11.6	20	✓
KSL463		0.23	21.24	19.61	18.31	17.85		10.8	20	✓
KSL342		0.21	21.22	19.28	17.95	17.49		10.9	20	✓
KSL322	0.33 ³	0.44	22.27	20.02	18.37	17.76	333 ± 25	11.5	20	✓✓
KSL674	0.28 ¹	0.31	21.68	19.60	18.14	17.60	293 ± 21	11.2	20	✓✓
KSL564	0.29 ¹	0.33	22.92	19.85	18.47	17.88	249 ± 22	11.2	20	✓✓
KSL670	0.44 ¹	0.48	23.71	21.10	19.49	18.72	207 ± 25	11.4	20	
KSL535		0.44	22.08	20.32	18.68	18.08		11.5	18	

¹ Eisenstein et al. (2011) and Dawson et al. (2013); ² Colless et al. (2001); ³ Liske et al. (2015); ⁴ Limousin et al. (2010)

2.4.2.1 Expected number of lenses

To assess whether the amount of selected candidates is reliable, we estimate the detectable lens population in KiDS, using the lens-statistics code LENSPOP⁵ (Collett, 2015). Assuming an effective KiDS survey area of ~ 1275 sq. deg we forecast ~ 2400 potentially detectable lenses with a total signal-to-noise ratio larger than 20 and having lensed images resolved over least three seeing elements. The expected number of lenses reduces to ~ 500 for the effective area of the KiDS-DR3 of 255 sq. deg. If we consider only lenses that satisfy our colour-magnitude cut of Sect. 2.2.2 and with an Einstein radius > 1.4 arcsec, i.e., our range of the parameter space, we forecast ~ 50 lenses for 255 sq. deg, broadly comparable to the number of our final sample of candidates, especially if we keep in mind that (i) we do not expect the CNN plus human lens selection to be 100% efficient and (ii) our training was largely focused on arcs and rings and not on quads and doubles. The redshift distribution of our final sample of candidates and the simulated population from LENSPOP, within the selection constraints specified above, are also consistent. Our candidates are observed in the window $0.1 \lesssim z \lesssim 0.5$, with a median redshift of $0.28^{+0.12}_{-0.08}$, while the LENSPOP sample is in the window $0.13 \lesssim z \lesssim 0.4$ with a median redshift of $0.32^{+0.08}_{-0.09}$. The scatter corresponds to the 16-84th quantiles of the distribution.

2.4.3 Sample Characterization and Comparison

To further characterize the sample of candidates, and allow a comparison with the literature, we estimate the stellar masses using the software LE PHARE (Arnouts et al. 1999; Ilbert et al. 2006), which performs a χ^2 fitting between the stellar population synthesis (SPS) theoretical models and the data. Single burst models from Bruzual & Charlot (2003, BC03) and a Chabrier (2001) IMF are implemented in the software. In the BC03 models we leave the age free to vary up to a maximum of 13 Gyr, and assume metallicities in the range $(0.005-2.5 Z_{\odot})$. No internal extinction is adopted. The single burst models provide us with a fair description of the stellar populations in massive early-type galaxies. Models are redshifted using the photometric redshifts (or the spectroscopic estimates where available). We adopt the observed *ugri* magnitudes (and related 1σ uncertainties) within a 5 arcsec diameter aperture, corrected for Galactic extinction using

⁵<https://github.com/tcollett/LensPop>

Figure 2.6: $g-r$ colour-redshift distribution of the LRG sample (green dots; Sect. 2.2.2) and our 56 best candidates (blue diamonds; Sect. 2.4). The BPZ photometric redshift is plotted, except for the candidates with an available spectroscopic redshift.

Figure 2.7: RGB images of the three known lenses present in the LRG sample. The network correctly classifies the first two as lenses, but classifies the third as non-lens. Given that its Einstein radius is 0.83 arcsec, smaller than the Einstein radii of the simulated lenses on which the network has been trained, this might be expected. The images are 20 by 20 arcsec.

the map in Schlafly & Finkbeiner (2011)⁶. The r -band `MAG_AUTO` is used to correct the outcomes of LE PHARE for missing flux. For 34 out of the 56 lens candidates (i.e., 60 per cent) we have spectroscopic redshifts. As pointed out previously, photometric redshifts tend to be larger than the spectroscopic estimate by $\Delta z \sim 0.04$, on average. For the 22 galaxies with a photometric redshift only, an overestimated redshift could imply both an over- or underestimate of the stellar mass, leading to a less reliable stellar mass. We have estimated the average impact of this systematics using the derived masses for the 34 galaxies with both measures of redshifts, finding that the photometric values bring to an average overestimate of the mass of 0.04 dex, with a scatter of ~ 0.2 dex. In addition, the aperture photometry adopted for the derivation of stellar masses is also affected. We have considered KiDS magnitudes within a radius of 5 arcsec, thus the enclosed lensing features make bluer colours and thus we underestimate the real stellar mass. However, we expect that this effect is within the typical mass uncertainty, since the arcs are very faint compared to the lens and hence colour contamination is very small. A systematic study of this issue is beyond the scope of this paper: we will discuss its impact on our results in a forthcoming paper.

In terms of redshift distribution, our lens candidates are observed in the window $0.1 \lesssim z \lesssim 0.5$, with a median redshift of $0.28^{+0.12}_{-0.08}$. This value is larger than the median redshift of SLACS lenses from Auger et al. (2009), i.e. $0.20^{+0.09}_{-0.07}$, but consistent within the scatter distribution. Instead, our median redshift is smaller than the average for the SL2S sample from Sonnenfeld et al. (2013b), i.e. $0.48^{+0.23}_{-0.16}$. This is not surprising given the $z \lesssim 0.4$ colour cut of Sect. 2.2.2. In future analyses this limit will be loosened. The median stellar mass of our sample is $\log M_*/M_\odot \sim 11.2$ dex with a scatter of ~ 0.2 dex. The typical uncertainty of the mass estimates is $\sim 0.1 - 0.2$ dex too. Within the scatter and mass uncertainties, this value is consistent with the average stellar mass in SLACS (the median is $\log M_*/M_\odot \sim 11.3$ dex and the scatter is ~ 0.2 dex; Auger et al. 2009) and SL2S lenses (the median is $\log M_*/M_\odot \sim 11.2$ dex and the scatter is ~ 0.25). All the galaxies have $M_* \gtrsim 10^{11} M_\odot$, except for KSL669 (see next subsection for further comments about this source). In Fig. 2.8 we plot the stellar mass as a function of redshift for our sample and the SLACS and SL2S ones.

A similar comparison can be performed for the velocity dispersion, if we consider the KiDS candidates with an available measure of this quantity.

⁶These updated extinctions are calculated by multiplying for 0.86 the Schlegel et al. (1998) values stored in the KiDS-DR3 catalog.

The average value for KiDS is $\sigma_{\star} = 232^{+46}_{-20}$ km/s⁷. In SLACS, the average is $\sigma_{\star} = 243^{+47}_{-33}$ km/s, while in SL2S the velocity dispersion within a radius of one-half effective radius is $\sigma_{e,2} = 258^{+42}_{-53}$ km/s. The three estimates agree within the scatter distribution and within the typical uncertainties of velocity dispersion measurements of $\sim 15 - 20$ km/s.

2.4.3.1 A sanity check of the candidates

A sub-sample of candidates has stellar velocity dispersion, $\sigma_{\star}$, measured in the SDSS (Eisenstein et al., 2011) or BOSS (Dawson et al., 2013) surveys. For these candidates, the knowledge of $\sigma_{\star}$ allows to put constraints on their Einstein radii R_E . For a SIS model, the Einstein radius can be expressed, in radians, as

$$\theta_E = 4\pi \left(\frac{\sigma_{\text{SIS}}}{c} \right)^2 \frac{D_{ls}}{D_s}, \quad (2.2)$$

where D_s , and D_{ls} are, respectively, the angular diameter distances between the observer and the source and between the lens and the source. As a first approximation, in Eq. (2.2) the σ_{SIS} strength parameter can be substituted with the measured stellar velocity dispersion, since they have been found to be approximately equal for lens galaxies (see, e.g., Bolton et al. 2008). A more rigorous approach consists of deriving the value of σ_{SIS} by matching the theoretical velocity dispersion derived from the Jeans equations with the observed one. From the Jeans equations, the radial velocity dispersion can be easily derived. This theoretical quantity is first integrated along the line of sight and then within a circular aperture with SDSS or BOSS fibres radius (i.e., $R_{\text{ap}} = 1.5$ and 1 arcsec, respectively; see Tortora et al. (2009) for equations and further details about the procedure). We use a SIS for the total mass profile, and the light distribution is set adopting a de Vaucouleurs (1948) profile, using the effective radii taken from the SDSS website⁸ which come from a de Vaucouleurs fit. Imposing that the theoretical aperture-averaged velocity dispersion σ_{Jeans} is equal to the observed one $\sigma_{\star}$, the only free parameter, σ_{SIS} , can be derived. This estimated quantity is finally inserted in Eq. (2.2).

For each lens with a measured $\sigma_{\star}$, the predicted R_E is plotted as a function of the unknown z_s and compared with the observed Einstein radius (see Fig. 2.9). A precise determination of the Einstein radius would require modelling of the lensing candidates, which is beyond the scope of

⁷The velocity dispersions are extracted from both SDSS and BOSS survey, thus, this average value is mixing observations made within two different fibre apertures.

⁸<http://skyserver.sdss.org/dr13/en/home.aspx>

Figure 2.8: Stellar mass versus redshift for our 56 best candidates (black dots), SLACS sample from Auger et al. (2009, red triangles) and SL2S sample from Sonnenfeld et al. (2013b, blue squares). The uncertainties on the stellar masses are $\sim 0.1 - 0.2$ dex.

this paper, but is planned for a follow-up paper. Here, we simply estimate the Einstein radius visually. We take it to be between 1 and 0.5 times the distance between the arc (or the brightest arc in case of multiple images) and the centre of the lens. This choice is due to the fact that, given a SIS, the image separation for an Einstein ring is exactly twice the Einstein radius, whereas in the case of an arc or an image maximally away from the centre of the lens, the distance from the centre is twice the Einstein radius (Kormann et al., 1994).

In Fig. 2.9 we show the comparison and we find an overlap for about half of the candidates with measured velocity dispersion. The other half is more likely to be constituted by ring galaxies, foreground sources or other contaminants. An excellent agreement is found for the systems KSL713, KSL450, KSL737, KSL565, KSL322, KSL674 and KSL564. For the other systems, the dynamics predict too small Einstein radii. One interesting case is KSL669, which is a $z = 0.05$ galaxy with a very small stellar mass of $\sim 2 \times 10^{10} M_{\odot}$. We expect the probability for it to act as a lens to be very small. The comparison performed in Fig. 2.9 seems to confirm the peculiarity of this lensing candidate. In fact, a larger view of the source shows that it is actually a merger event.

Moreover, among the best ranked systems, KSL627 and KSL376 present a discrepancy which seems difficult to reconcile. These two systems have almost circular blue rings, with ~ 4.3 and ~ 5.7 arcsec radii, corresponding to 15 and 25 kpc, respectively. These sources do not match typical Einstein radii observed in galaxy-scale gravitational lenses; they are more likely to belong to the category of ring galaxies (Hoag, 1950; Theys & Spiegel, 1976; Whitmore et al., 1990; Bournaud & Combes, 2003; Iodice et al., 2003; Madore et al., 2009).

ETGs follow a tight relationship between velocity dispersion and stellar mass. This can allow us to predict the velocity dispersion for the remaining galaxies in our final sample. After collecting ETG lenses from SLACS (Auger et al., 2009), we perform a median fit determining the best-fitted relation $\log \sigma_{\star} = -0.1 + 0.22 \log M_{\star}/M_{\odot}$ between the velocity dispersion and the estimated stellar masses. Thus, assuming that this relation holds for our sample and that it has no scatter, we can determine an estimate for the velocity dispersion, when our stellar masses are used. We use Eq. (2.2) to predict the Einstein radius as a function of the source redshift. The results are shown in Fig. 2.10, where the dashed line is calculated by inserting the estimated velocity dispersion in Eq. (2.2), while the solid line is calculated by assuming the average of the ratios $\sigma_{\text{SIS}}/\sigma_{\star}$ obtained for the galaxies with available velocity dispersion, and inserting the derived σ_{SIS} in Eq. (2.2). Similar considerations as for the galaxies with measured velocity dispersion

can be done for these objects, even if the uncertainties on the estimated velocity dispersions are higher.

The previous analysis can give us an indication on the nature of the candidates. However a spectroscopic validation is needed, because it can not be excluded that the lens candidates are part of a group of galaxies. In this case the stellar velocity dispersion would not trace the dynamics of the group. Indeed, this is the case of the known lens J085446-012137 (KSL317), which is part of a group (Limousin et al., 2010), resulting in a under-estimation of the Einstein radius.

We plan to follow-up our most reliable candidates, mainly for an estimate of the redshift of the arc, in order to confirm or discard their lensing nature.

Figure 2.9: Einstein radii estimated by the observed dynamics are plotted as a function of the unknown source redshift z_s , and compared with the value estimated visually from the images. The black lines are calculated using Eq. (2.2): the solid line assumes that σ_{SIS} is determined from Jeans dynamical analysis, and the dashed line by fixing $\sigma_{\text{SIS}} = \sigma_*$ (see Sect. 2.4.3.1 for the details). The shaded cyan region corresponds to a conservative range of values for the Einstein radii: it is calculated from the observed distance of the arc from the lens centre, and is set to the range between $0.5R_E$ and R_E .

2.5 Conclusions

We have developed a new pipeline, based on a Convolutional Neural Network (CNN), to automatically identify strong gravitational lens candidates based mainly on their morphology (Sect. 2.3). We have applied the method (see Sect. 2.4) to the third data release of the Kilo Degree Survey (KiDS), which is one of the ESO public surveys carried out with the VLT Survey Telescope (see e.g. de Jong et al., 2015, for a description). Thanks to its high quality images, KiDS is particularly suitable for a search of strong lenses. In the complete survey we expect to find at least 100 LRG lens systems with lens-galaxies at $z < 0.4$, possibly increasing to several thousand when expanding the search to fainter and higher-redshift galaxies (see Sect. 2.4.2.1).

To train the CNN to find lenses, we generated a large sample of simulated lensing features on top of observed colour-magnitude selected galaxies from KiDS (Sect. 2.3.1.2). The trained network has been applied to a sample of 21789 LRGs in KiDS DR3, retrieving 761 candidates (3.6% of the initial sample). With a visual inspection performed by seven “human” classifiers, we down-selected the most promising 56 lens candidates (Fig. 2.11). In our starting sample there were three known lenses, two of which were classified correctly as lenses by the CNN and in the subsequent visual inspection phase (Fig. 2.7). The misclassified lens has an Einstein radius (0.83 arcsec), well below the range where the CNN is trained (1.4 – 5.0 arcsec).

For the candidates with available measures of velocity dispersion or stellar mass estimates, we performed an additional sanity check, suggesting that ~ 22 are solid candidates (Sect. 2.4.2). Considering the colour-magnitude selection of the lens-galaxy sample, the type of lenses simulated and the completeness, the number is roughly consistent with the expected ~ 50 lenses forecast for the KiDS-DR3 survey area (Sect. 2.4.2.1). Extending this result to the full KiDS survey, we expect to find ~ 100 LRG lens candidates as a lower limit, similar to the number of lenses in the SLACS sample. Because we limited our search to a very restricted portion of the colour-magnitude and Einstein-radius and magnification space, dominated by luminous lenses and highly magnified sources, the natural next step is to enlarge the colour-magnitude pre-selection of the simulated and observed lenses, which would allow in principle to find up to ~ 2400 lenses in the most optimistic scenario (nearly all lenses with signal-to-noise larger than 20 and at least three PSF resolution element for the arc-like images).

Figure 2.10: Same as in Fig. 2.9, but for the candidates without measured velocity dispersion. In these cases the velocity dispersion is inferred from the stellar mass as described in Sect. 2.4.3.1.

A critical aspect to improve in the CNN approach is to reduce the contamination by the false positives which dominate the number of true positives currently by a factor of ~ 40 . This could facilitate, or even eliminate, the need for visual inspection. This factor is consistent with the $\sim 4\%$ mis-classifications in the training (see Appendix 2.B), which for the input sample size can lead to ~ 900 false-positive, which is close to the actual number of ~ 700 . Given that the lenses-galaxies are outnumbered by normal galaxies typically by a thousand to one, an important goal is to bring the false-positive rate down to less than 0.1% , without decreasing the true-positive rate substantially. This is a hard task, but human visual inspection suggests that at least 0.25% can be reached (i.e. 56 out of 22 thousand), possibly when including additional colour (RGB) information. Our next goal is to create a completely automated pipeline for lens classification without the need of visual inspection. However, if visual inspection will be needed to down-select lens candidates, it will be important to test the efficiency of the human classifiers in order to be able to estimate accurately the completeness and purity of the final selected sample.

Moreover, for evaluating purity and completeness in a realistic lens search setting, we plan to build a validation-set which reproduces the characteristics of a real survey (where the number of negatives far outnumber the number of positives).

The CNN tends to mis-classify primarily galaxies resembling lensing features (e.g., ring galaxies, mergers, star-forming rings). Thus, training the network on an ensemble of this kind of false-positives would allow the algorithm to learn the subtle differences between the false positives and the true lenses.

The network performance could also be improved with model averaging, i.e., building a series of networks for the same task, but with different structure and parameters, and by averaging their output. Moreover, training the network on galaxy-subtracted images could facilitate the algorithm to pick up more subtle lensing features, especially in the regime of small Einstein radii and bright galaxies. Another possibility is to produce and train the CNN on multi-band images. In this way, colour information would be used to discriminate between lenses with sources and non-lenses.

The final sample of KiDS-DR3 lens candidates suggests that our method is promising to down-selected lens candidates from an input sample by two orders of magnitude. Moreover, it is easily applicable to any ongoing and future survey, (e.g., Euclid, LSST) for classifying the enormous amount of data that will be produced. In the near future we plan spectroscopic follow-up of our best candidates, to model them, and to better assess their

selection biases. In addition, we will apply the method to the full KiDS survey and work on the above-mentioned improvements.

Acknowledgements

We thank the anonymous referee for his/her comments that have improved the paper; Sander Dieleman, Davide Punzo, Marco Alexander Wiering, Emmanuel Okafor, Giuseppe Longo, Andrea Colonna and Koen Kuijken for stimulating and useful discussions; Hugo Buddelmeijer, Ewout Helmich and Jelte de Jong for their technical support. LVEK thanks Tom Collett with help in using the code LENSPOP. CEP, SC, CT, GV, GVK and LVEK are supported through an NWO-VICI grant (project number 639.043.308). GVK acknowledges financial support from the Netherlands Research School for Astronomy (NOVA) and Target. Target is supported by Samenwerkingsverband Noord Nederland, European fund for regional development, Dutch Ministry of economic affairs, Pieken in de Delta, Provinces of Groningen and Drenthe. PS is supported by the Deutsche Forschungsgemeinschaft in the framework of the TR33 ‘The Dark Universe’.

Appendix 2.A NEURAL NETWORKS

In this appendix we give a short introduction on the theory of Convolutional Neural Networks (CNN).

Figure 2.11: RGB images of the 56 candidates down-selected through a visual inspection of the 761 CNN candidates (see Sect. 2.4.1). Each source is labelled by an internal ID followed by, in parentheses, the visual classification score (70 points maximum). Each image is 20 by 20 arcsec.

Figure 2.11: *continued*

Figure 2.11: *continued*

Figure 2.11: *continued*

2.A.1 Feed-forward neural network

A feed-forward neural network, as the one shown schematically in Fig. 2.12, is a basic example of a deep-learning algorithm that can be schematized as an ensemble of connected units. The *input layer* is the representation of a single element of the training set, where the units are the components of the data-point vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$. Every unit in the *hidden layer* performs the following transformation of its inputs

$$y = \sigma(\mathbf{w} \cdot \mathbf{x} + b) \quad (2.3)$$

where the vector $\mathbf{w} = (w_1, w_2, \dots, w_n)$ is called the *weight vector* and the constant b is the *bias*. The non-linear activation function, σ is often the Rectified linear unit (ReLU; Nair & Hinton 2010)

$$\sigma(x) = \max(0, x), \quad (2.4)$$

or the sigmoid function

$$\sigma(x) = \frac{1}{(1 + e^{-x})}. \quad (2.5)$$

Hidden layers are stacked sequentially until the topmost, i.e. the *output layer*, is reached. We want the output layer $\mathbf{y} = (y_1, y_2, \dots, y_n)^T$ to approximate the desired output $\hat{\mathbf{y}} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_n)^T$. This is obtained by finding the weights and biases that minimize a chosen loss function $L(\mathbf{y}, \hat{\mathbf{y}})$. The minimization is most often done via the iterative process of *gradient descent*. For each layer l the weights and biases are updated in the following way

$$\begin{aligned} \mathbf{w}_l &\rightarrow \mathbf{w}'_l = \mathbf{w}_l - \eta \frac{\partial L}{\partial \mathbf{w}_l} \\ \mathbf{b}_l &\rightarrow \mathbf{b}'_l = \mathbf{b}_l - \eta \frac{\partial L}{\partial \mathbf{b}_l} \end{aligned} \quad (2.6)$$

where η is a constant called the *learning rate*. The gradients are computed via the back-propagation algorithm (Rumelhart et al., 1986).

2.A.2 Convolutional Neural Network

In Convolutional Neural Networks, the input data has a topological structure (e.g. an image) and is not presented as a vector but as a set of matrices $\mathbf{X}_k$ with $k = 1, 2, \dots, K$ (e.g., the R, G and B components of

Figure 2.12: A schematic view of a feed forward neural network.

an image. In this case $K = 3$). The main component of a CNN is the convolutional layer, which takes the inputs and, through a set of filters, produces a stack of *feature maps* $\mathbf{Y}_n$ with n equal to the number of filters. Every filter (also called kernel) in the convolutional layer produces a feature map through a convolution

$$\mathbf{Y} = \sigma \left(\sum_{k=1}^K \mathbf{W}_k * \mathbf{X}_k + \mathbf{B} \right), \quad (2.7)$$

where $*$ is the convolution operator, σ is a non linear function as in eq. (2.3), $\mathbf{W}_k$ are the K weight matrices with $k = 1, 2, \dots, K$, representing a filter with its bias given by the constant matrix $\mathbf{B}$. There are far fewer parameters to be determined in a convolutional layer as compared to a fully connected layer because, practically, we are replacing the dot product of Eq. (2.3) with a convolution, and, in all the practical cases, the weight matrices have spatial dimensions much smaller than the input dimension (usually 3 by 3).

Convolutional layers are sequentially stacked such as the input of the deeper layers are the feature maps. In between then there are non-linear and other transformations (e.g. pooling). After the training is complete, in each layer we have a representation of the input data of increasing complexity (i.e. the different feature maps) each produced by a different filter which represents a particular feature learned during the training phase. The output layer of a CNN is the same as in a feed-forward neural network and can be preceded by one or more hidden fully-connected layers.

Its function is to classify the last layer of feature maps created by the CNN, giving as output one or more numbers which represent the outcome of the classification.

60

Figure 2.13: Schematic view of the architecture of our CNN.

Table 2.3: The table summarizes the characteristics of each layer of our CNN.

Type	Filters/Units	Filter size	Padding	Non-linearity	Initial weights	Initial biases
Convolutional	32	7×7	3	ReLU	HeNormal	0
Max-pooling	-	2×2	-	-	-	-
Convolutional	64	3×3	1	ReLU	HeNormal	0
Max-pooling	-	2×2	-	-	-	-
Convolutional	128	3×3	1	ReLU	HeNormal	0
Convolutional	128	3×3	1	ReLU	HeNormal	0
Fully connected	1024	-	-	ReLU	HeNormal	0
Fully connected	1024	-	-	ReLU	HeNormal	0
Max-pooling	-	1	-	-	-	-
Fully connected	1	2	-	sigmoid	HeNormal	0

Appendix 2.B CNN implementation

Our CNN is implemented in Python 2.7 using the open-source libraries LASAGNE⁹ and THEANO¹⁰ (Theano Development Team, 2016). The training of the CNN is executed on a GeForce GTX 760 in parallel with the data augmentation performed on the CPU using the SCIKIT-IMAGE¹¹ package (Van der Walt et al., 2014). The training time with this configuration takes about 2 hours. While the CNN takes about 20 minutes to classify the LRG sample. In this Appendix we provide the technical details of the implementation and training of our CNN (Sections 2.B.1 and 2.B.2).

2.B.1 Network architecture

In Fig. 2.13 and in Table 2.3 we show the architecture of our CNN. ReLU (see Eq. (2.4)) is applied after each convolutional and fully connected layer. The 60 by 60 input layer is followed by four convolutional layers with 32, 64, 128 and 128 filters, respectively. All the filters have 3 by 3 sizes except for the first convolutional layer which has a filter size of 7 by 7. In the convolutional kernels we use untied bias. To preserve the input volume through the convolution and not degrade the information at the borders of the input, we zero-pad the input of the convolutional layers with 3, 1, 1 and 1 pixels respectively. Max-pooling (Boureau et al., 2010) with a kernel size of 2 by 2 is used after the first and the second convolutional layer and after the second fully-connected layer with a 1D kernel size of 2. Max-pooling takes the maximum value in a connected set of elements of the feature maps. There are two main consequences of using Max-pooling: a) reducing the dimensionality of the data and the parameters to be estimated, b) teaching translational invariance to the network, because slightly shifted inputs will produce the same feature maps. Two fully-connected layers of 1024 units follow the set of convolutional and max-pooling layers. Finally, we use a sigmoid non-linear output unit which gives a real number between 0 and 1 that represents the probability of being a strong gravitational lensing system. We use *batch normalization* (Ioffe & Szegedy, 2015) before the non-linearity of each layer. A batch normalization layer operates on the inputs of the non-linearities normalizing the data in order to have zero mean and unit variance among the mini-batch. Then, the data is fed to a linear function with two learnable parameters that has the property to

⁹<http://github.com/Lasagne/Lasagne/>

¹⁰<http://deeplearning.net/software/theano/>

¹¹<http://scikit-image.org/>

revert or modify the normalization. The chosen architecture implies that our network has about 30 million of trainable parameters.

2.B.2 Training

The network is trained by minimizing a loss function of the targets t (1 for lenses and 0 for non-lenses) and the predictions p (the output of the sigmoid unit of network). We use the binary cross-entropy, a common choice in two-class classification problems:

$$L = -t \log p - (1 - t) \log(1 - p) \quad (2.8)$$

The minimization is done via mini-batch stochastic gradient descent with *ADAM* updates (Kingma & Ba, 2014). The advantage of using *ADAM* updates is the introduction of a friction term that mitigates the gradient momentum in order to reach a faster convergence. In addition, the updates have a per-parameter adaptation, i.e., they have a different effective learning rate for the different parameters depending on the gradient values. We used a batch size of 600 images and perform 10000 gradient updates, which corresponds to six million examples. Each mini-batch is composed by 300 lens and 300 non-lens examples. After an initial exploration, we start with a learning rate of 0.004, decrease it to 0.0004 after three million training examples and to 0.00004 after 5.5 million training examples (the choice of the values for the learning rate has been fundamental for training successfully the network). The weights of each filter are initialized, as discussed in He et al. (2015a), from a random normal distribution with variance $2/n$ where n is the number of inputs of the unit. The initial values of the biases are set to zero. We use *dropout* (Hinton et al., 2012) in the fully connected layers. Dropout consists of switching off units randomly during each update of the training phase. This has two main consequences: a) a speed-up of the training phase, because of the reduced number of parameters to be computed in the fully connected-layers of the CNN, b) reducing the possibility of over-fitting, since the network tends to learn features that better generalize the data. We also use L2-norm regularization (see e.g., Ng 2004) with $\lambda = 10^{-4}$. The regularization adds to the loss function Eq. (2.8) another factor given by the squared sum of all the weights times the factor λ . It has the property to let the network prefer to learn small diffuse weights penalizing the creation of peaky ones. In this way a classification, based on all the data coming from the input, tends to be promoted over one where the weights tend to consider only a subset of the input.

Figure 2.14: The distribution of the network output for the validation set. The blue bars are the lens examples, while the green ones are the non-lens examples.

2.B.3 Analysis

For monitoring the network during the training we build a fixed validation set composed of 5000 images (half lenses, half non-lenses) with the same prescriptions as summarized in Sect. 2.3.2. At the end of the training the network reaches a 96% accuracy for both the lens and non-lens examples. Fig. 2.14 shows the distribution of the network output p for the validation set. For values of p greater than 0.5 the lens examples start to be dominant in the distribution. To check if there is any correlation between p and the characteristics of the simulated mock sources, we investigate the p distribution of the lens examples for different ranges of three parameters: Einstein radii, magnification and ratio between the peak brightness of the source and the lens. From Fig. 2.16 one can see that the p distributions are skewed to higher values for the higher ranges of the parameters considered (larger, brighter and more magnified lens systems). This implies that the classification is more accurate for the training examples with higher magnification, higher Einstein radius and more luminous sources with respect to the KiDS galaxy. In Fig. 2.15 we show the distribution of the network output of p -values for the full LRG sample and the 56 lens candidates (a sub-sample of the 761 candidates with $p > 0.5$). We do not retrieve a significant peaked distribution in the far end as for the validation set. This could be due to the intrinsic difference between real

Figure 2.15: The distribution of the network output for the full LRG sample (blue bars) and for the 56 candidates selected via visual inspection (green bars).

and simulated data. In addition, we have far fewer lenses in the LRG sample compared to the non-lenses. Moreover, the lenses in the validation set are uniformly distributed in the range of the parameters of Table 2.1. Thus, for a proper comparison, a validation set that reproduces the numbers and the distribution of the parameters for real lenses should be created.

Appendix 2.C r-band images of the candidates

In Fig. 2.17 we show the r -band images of the 56 candidates selected through the visual inspection of Sect. 2.4.1. r -band KiDS images have been used as the actual input of the CNN with the purpose of finding lens candidates.

2. CNN LENSFINDER

Figure 2.16: Distributions of the output of the CNN for different bins (shown in the parenthesis) of some parameters of the mock lensed sources (ratio between the maximum brightness of the lensed source and the lens in green, Einstein radius in blue, magnification in red).

Figure 2.17: Square-root stretched KiDS *r*-band images of the 56 candidates selected through a visual inspection of the 761 CNN candidates (see Sect. 2.4.1). Each source is labelled by an internal ID followed by, in parenthesis, the visual classification score (70 points maximum). Each image is 20 by 20 arcsec.

Figure 2.17: *continued*

Figure 2.17: *continued*

Figure 2.17: *continued*

Figure 2.17: *continued*

*I am putting myself to the fullest possible use, which is all I think
that any conscious entity can ever hope to do*

– HAL 9000, *2001: A Space Odyssey*

3. TESTING CONVOLUTIONAL NEURAL NETWORKS FOR FINDING STRONG GRAVITATIONAL LENSES IN KiDS

*C. E. Petrillo, C. Tortora, S. Chatterjee, G. Vernardos, L. V. E. Koopmans,
G. Verdoes Kleijn,
N. R. Napolitano, G. Covone, L. S. Kelvin, A. M. Hopkins
2018, Monthly Notices of the Royal Astronomic Society, 482, 807*

Convolutional Neural Networks (ConvNets) are one of the most promising methods for identifying strong gravitational lens candidates in survey data. We present two ConvNet lens-finders which we have trained with a dataset composed of real galaxies from the Kilo Degree Survey (KiDS) and simulated lensed sources. One ConvNet is trained with single r-band galaxy images, hence basing the classification mostly on the morphology. While the other ConvNet is trained on g-r-i composite images, relying mostly on colours and morphology. We have tested the ConvNet lens-finders on a sample of 21789 Luminous Red Galaxies (LRGs) selected from KiDS and we have analyzed and compared the results with our previous ConvNet lens-finder on the same sample. The new lens-finders achieve a higher accuracy and completeness in identifying gravitational lens candidates, especially the single-band ConvNet. Our analysis indicates that this is mainly due to improved simulations of the lensed sources. In particular, the single-band ConvNet can select a sample of lens candidates with $\sim 40\%$ purity, retrieving 3 out of 4 of the confirmed gravitational lenses in the LRG sample. With this particular setup and limited human intervention, it will be possible to retrieve, in future surveys such as Euclid, a sample of lenses exceeding in size the total number of currently known gravitational lenses.

3.1 Introduction

Strong gravitational lensing is a phenomenon that originates when light rays propagating from a background source galaxy are deflected, on their way towards the observer, by the gravitational field of a foreground galaxy, creating multiple images, arcs and/or rings around the foreground galaxy. Strong gravitational lensing is a unique probe for studying the (dark) matter distribution of galaxies and providing cosmological constraints. E.g., gravitational lenses have been used to measure the Hubble constant through time delays of lensed quasar images (e.g., Suyu et al. 2010, 2017; Bonvin et al. 2017) and to constrain the dark energy equation of state (e.g., Biesiada et al. 2010; Collett & Auger 2014; Cao et al. 2012, 2015). Gravitational lensing also allows measuring the fraction of dark matter in the central regions of galaxies (Gavazzi et al. 2007; Jiang & Kochanek 2007; Covone et al. 2009; Grillo et al. 2010; Cardone et al. 2009; Cardone & Tortora 2010; Auger et al. 2010; Tortora et al. 2010; More et al. 2011; Ruff et al. 2011; Sonnenfeld et al. 2015) and to constrain the slope of the inner mass density profile (e.g., Treu & Koopmans 2002a,b; Koopmans et al. 2006; Koopmans & Treu 2003; More et al. 2008; Barnabè et al. 2009; Koopmans et al. 2009; Cao et al. 2016; Mukherjee et al. 2018). Moreover, studying gravitational lenses can constrain the stellar initial mass function (e.g., Treu et al. 2010; Ferreras et al. 2010; Spiniello et al. 2011; Brewer et al. 2012; Sonnenfeld et al. 2015; Posacki et al. 2015; Spiniello et al. 2015; Möller et al. 2007; Sonnenfeld et al. 2018a). Strong lensing also works as a “cosmic telescope”, producing a magnified view of background objects otherwise not observable (e.g., Impellizzeri et al. 2008; Swinbank et al. 2009; Richard et al. 2011; Deane et al. 2013; Treu et al. 2015; Mason et al. 2017; Salmon et al. 2017; Kelly et al. 2018). Discovering new gravitational lenses allows placement of more precise constraints on the above-mentioned quantities (see e.g., Vegetti & Koopmans 2009; Barnabè et al. 2011; Li et al. 2016). For a comprehensive review of the scientific applications of strong lensing see, e.g., Schneider et al. (1992), Schneider (2006) and Treu (2010).

Originally, gravitational lenses were found serendipitously in astronomical surveys, while currently they are considered as an important class of objects to systematically search in large sky surveys. The most successful campaign aiming at building a homogeneous dataset of strong gravitational lenses was the Sloan Lens ACS Survey (SLACS; Bolton et al. 2008) with more than 100 observed lenses that were identified by analyzing spectra from the Sloan Digital Sky Survey (SDSS; Eisenstein et al. 2011) that exhibited the imprint of two galaxies at two different redshifts. On-going

optical surveys such as the Hyper Suprime-Cam survey (HSC; Miyazaki et al. 2012), the Kilo Degree Survey (KiDS; de Jong et al. 2015) and the Dark Energy Survey (DES; The Dark Energy Survey Collaboration 2005) are expected to provide in the coming years thousands of new lenses (see Collett 2015 and Petrillo et al. 2017) and have already provided new lens candidates (Petrillo et al., 2017; Diehl et al., 2017; Sonnenfeld et al., 2018b). The future is bright also in the sub-millimeter wavelength, where Herschel (Negrello et al., 2010) and the South Pole Telescope (Carlstrom et al. 2011), coupled with the Atacama Large Millimeter/sub-millimeter Array, are providing several hundreds of new lens candidates (Vieira et al., 2013; Negrello et al., 2017). However, it is the next decade that holds a treasure trove of new gravitational lenses. It has been estimated that samples of $\sim 10^5$ strong lenses (Oguri & Marshall, 2010; Pawase et al., 2014; Collett, 2015; McKean et al., 2015) will be observed by Euclid (Laureijs et al., 2011), the Large Synoptic Survey Telescope (LSST; LSST Science Collaboration et al. 2009) and the Square Kilometer Array¹.

The huge number of possible new candidates, together with the difficulty of identifying them in the enormous volume of survey data, drives the growing effort in developing automatic lens-finders. Most are based on the identification of arc-like features (e.g., Lenzen et al. 2004; Horesh et al. 2005; Alard 2006; Estrada et al. 2007; Seidel & Bartelmann 2007; Kubo & Dell’Antonio 2008; More et al. 2012; Maturi et al. 2014). Other approaches, such as described by Gavazzi et al. (2014) and Joseph et al. (2014), focus on the analysis of the residuals after subtracting the candidate lens galaxies from the astronomical images. Both methods have been used to find lens candidates in the Canada-France-Hawaii Telescope Legacy Survey (CFHTLS²) by Sonnenfeld et al. (2013a); Gavazzi et al. (2014); Paraficz et al. (2016). Instead, the algorithm developed by Chan et al. (2015) is specialized in identifying lensed quasars and together with the algorithm YATTALENS (Sonnenfeld et al., 2018b) has been applied to find lens candidates in HSC (Sonnenfeld et al., 2018b). Another approach, as in Brault & Gavazzi (2015), is modeling the probability that the targets are actual lenses. Stapelberg et al. (2019b) applied the same strategy to clusters and groups of galaxies. Gravitational lenses have been identified also with citizen-science experiment approaches with the Space Warps project (Marshall et al., 2016; More et al., 2016) where non professional volunteers can classify galaxy images with the help of a web applet³. Most recently, Petrillo et al. (2017) and Jacobs et al. (2017) have used

¹<https://www.skatelescope.org/>

²<http://www.cfht.hawaii.edu/Science/CFHLS/>

³<https://spacewarps.org/>

Convolutional Neural Network (ConvNets) for finding lens candidates in KiDS and CFHTLS, respectively. Finally, Hartley et al. (2017) have used a technique based on Support Vector Machines (SVMs) and applied it to KiDS. Instead Spiniello et al. (2018) focused on the search of lensed quasars in KiDS using 3 different morphology based methods.

In this paper we present and test our latest ConvNet lens-finders, improving on the work of Petrillo et al. (2017). ConvNets (Fukushima 1980; LeCun et al. 1998) are the state of the art and often the standard choice among machine learning algorithms for pattern recognition in digital images. The winners of the ImageNet Large Scale Visual Recognition Competition (ILSVRC; Russakovsky et al. 2015; the most important image classification competition) in recent years have all been groups utilizing ConvNets. The advantage of the latter method with respect to other pattern recognition algorithms is that the features are not hand-crafted but are themselves extracted automatically during the training procedure, thus the algorithm decides which features are most representative for classifying the images. The theoretical basis of ConvNets was developed in the 1980s and the 1990s. However only recently ConvNets have started to outperform other algorithms thanks to the advent of large labeled datasets, improved algorithms and faster training, especially on Graphics Processing Units (GPUs). The interested reader is referred to the Appendix for a brief introduction on ConvNets and to the reviews by Schmidhuber (2015), LeCun et al. (2015) and Guo et al. (2016) for a more detailed introduction.

ConvNets have been used recently in many astronomical problems, e.g., galaxy morphology classification (Dieleman et al., 2015; Huertas-Company et al., 2015), estimation of photometric redshifts (Hoyle, 2016; D’Isanto & Polsterer, 2018), spectra classification (Hála, 2014; Tao et al., 2018), identifying exoplanets (Shallue & Vanderburg, 2018), transient detection (Cabrera-Vives et al., 2017), galaxy surface brightness estimation (Tuccillo et al., 2018), strong lensing parameters estimation (Hezaveh et al., 2017) and star/galaxy separation (Kim & Brunner, 2017).

More importantly, Metcalf et al. (2019) presented the results of a large international *challenge* in which various methods of identifying simulated gravitational lenses were tested blindly. This challenge, the first of a series, sets out to prepare the community for finding lenses in the data of ESA’s Euclid mission. Its large data volume requires fast and efficient algorithms to identify strong gravitational lenses. However, the methods were also tested on simulated KiDS data. ConvNets and Support Vector Machines (SVMs) were recognized to be the most promising methods, among many different methods tested in the challenge.

The ConvNet lens-finders presented in this paper will be applied on ~ 900 sq. deg of the KiDS survey in a forthcoming paper with the purpose of starting a systematic census of strong lenses named “LinKS” (Lenses in KiDS Survey). The paper is organized as follows. In Sect. 3.2, we illustrate our lens-finding ConvNet-based algorithms and how the training dataset is built. In Sect. 3.3, we evaluate the performances of the ConvNets. In Sect. 3.4, we apply the lens-finders to ~ 22000 extracted from ~ 255 square degrees of KiDS for testing the algorithms on real data. Finally, in Sect. 3.5, we provide a summary and the main conclusions of this work.

3.2 Training the ConvNets to find strong lenses

A Convolutional Neural Network (ConvNet) can be seen as a sequence of non-linear functions, called layers, that create, starting from an input image, a series of increasingly abstract representations of the input called feature maps. The final layer of the ConvNet converts the input feature maps into a set of numbers that represent the outcome of the classification. Hence a ConvNet maps an image onto a single or few numbers. In our case the output is a single number, denoted by p , which can vary between 0 and 1, and it is related to the probability that the input image is a lens (see Saerens et al. 2002 for a detailed discussion). The parameters of the non-linear functions are obtained during the so called training phase where labeled images are fed to the ConvNet. In more detail, the parameters are derived by minimizing a loss function that expresses the difference between the label values of the images (1 for lenses, 0 for non-strong-lensing systems) and the output p of the ConvNet. Although in Petrillo et al. (2017) we have used a similar set-up, the aim of this work is to improve the performance of our previous lens-finder.

Currently we use a ConvNet with a *ResNet*-type architecture that has 18 layers, exactly as described in He et al. (2015b). *ResNet*-type architectures are often the preferred choice in image classification tasks because of their faster convergence and higher classification accuracy with respect to other architectures. Moreover, *ResNet* architectures have already been tested successfully on identifying simulated lenses and they have proven to be one of the best architecture for this task (Schaefer et al., 2018; Lanusse et al., 2018; Metcalf et al., 2019). We train two different ConvNets with the same architecture except that one takes in input RGB-images composed with HUMVI⁴ (Marshall et al., 2016), while the other takes single r -band images as input. We choose the r -band as single-band input because the KiDS

⁴<https://github.com/drphilmarshall/HumVI>

observing strategy reserves the best seeing conditions for this band (which is used for the weak lensing studies; Kuijken et al. 2015; Hildebrandt et al. 2017). The technical details of the ConvNet and of the training procedure are described in Appendix A together with a brief introduction on ConvNets.

To produce the data used to train and validate the ConvNets, we adopt a hybrid approach similarly as done in Petrillo et al. (2017); Jacobs et al. (2017); Pourrahmani et al. (2018), creating mock images of strong gravitational lenses using images of real galaxies from KiDS and superimposing simulated lensed images. We adopt this approach because we do not have a sample of genuine KiDS lenses large enough to train a ConvNet (usually of the order of 10^6).

3.2.1 Data

In this section we describe the dataset used to train the ConvNets, which is composed of real KiDS galaxies and simulated lensed sources.

3.2.1.1 Luminous Red Galaxies

We use the sample of Luminous Red Galaxies (LRGs; Eisenstein et al. 2001) presented in Petrillo et al. (2017). We choose to focus on massive early-type galaxies, because it has been estimated that these galaxies form $\sim 80\%$ of the lens-galaxy population (Turner et al., 1984; Fukugita et al., 1992; Kochanek, 1996; Chae, 2003; Oguri, 2006; Möller et al., 2007). Spiral galaxies form the other $\sim 20\%$ but are much harder to identify. This training sample of LRGs is a random subset of 6554 galaxies from a parent sample of 21789 selected from 255 square degrees of KiDS DR3 (de Jong et al., 2017) with the following criteria (see Petrillo et al. 2017 for more details):

(i) The low- z ($z < 0.4$) LRG colour-magnitude selection of Eisenstein et al. (2001), adapted to include more sources, both fainter and bluer:

$$\begin{aligned} r &< 20 \\ |c_{\text{perp}}| &< 0.2 \\ r &< 14 + c_{\text{par}}/0.3 \end{aligned} \tag{3.1}$$

where

$$\begin{aligned} c_{\text{par}} &= 0.7(g - r) + 1.2[(r - i) - 0.18] \\ c_{\text{perp}} &= (r - i) - (g - r)/4.0 - 0.18 \end{aligned}$$

(ii) A source size in the r -band larger than the average FWHM of the PSF of the respective tiles, times an empirical factor to maximize the separation between stars and galaxies.

3.2.1.2 Contaminants

Moreover, we have used a set of ~ 6000 KiDS sources to train the ConvNets to recognize sources that would likely be incorrectly classified as lenses otherwise, either because they can resemble lensing features or they are “ghosts”, i.e. they are undetected, at least significantly, in the luminous red galaxies sample discussed in Sect. 3.2.1.1.

- ~ 2000 sources wrongly classified as lenses in previous tests with ConvNets identified by the authors. This is done to teach the ConvNets not to replicate previous mistakes;
- ~ 3000 randomly extracted KiDS sources with r -band magnitude brighter than 21. To provide the network with general true negatives.
- ~ 1000 KiDS sources visually classified as spiral galaxies from an on-going new project of GalaxyZoo (Willett et al. 2013, Kelvin et al., in prep.). This is done to decrease the false positives due to spiral features. To select the galaxies we used a preliminary reduced version of the GAMA-KiDS Galaxy Zoo catalogue for the GAMA09 9h region (see Driver et al. 2011 for further details). This catalogue contains $\sim 10^4$ sources out to a redshift of $z = 0.15$. We select galaxies for which a large majority of people replied to the question “*Is the galaxy in the centre of the image simply smooth and rounded, or does it have features?*” with “*it has features*”⁵

There is a non-zero probability that among the contaminants and the LRGs described in the previous Section there are actual gravitational lenses. We can estimate that the percentage would be of the order of 10^{-2} among the contaminants and $\sim 1\%$ among the LRGs (Petrillo et al., 2017). Thus, even if real lenses are actually in the training sample, with such a small percentage they would not contaminate the training procedure.

3.2.1.3 Mock lensed-sources

We simulate 10^6 lensed images of 101×101 pixels, using the same spatial resolution of KiDS (~ 0.2 arcsec per pixel), corresponding to a 20×20 arcsec field of view. To produce more realistic lensing systems, we add more complexity both in the source and in the lens plane with respect to

⁵The actual selection is done by selecting sources from the catalogue with a value of the attribute `features_features_frac` larger than 0.6.

the simulations in Petrillo et al. (2017). The distribution of the lens and source parameters that we choose for simulating the lensed images are chosen to create a wide range of realistic lensed images. They are not meant to statistically represent a real lens population, since the training set has to be populated sufficiently densely in the parameter space to allow the ConvNets to learn all the possible configurations and to recognize lenses that are rare in a real distribution (or currently even unknown). To ensure this, a more homogeneous distribution of the parameters is advantageous in order not to over-train on the most common lens configurations.

We proceed in the following way, we sample the parameters of the Singular Isothermal Ellipsoid (SIE; Kormann et al. 1994) and Sérsic (1968) source models as listed in Table 3.1. The values of the lens Einstein radius and the source effective radius are drawn from a logarithmic distribution, while the remaining parameters, listed in Table 3.1, are drawn from a uniform distribution. In this way our simulation sample contains a higher fraction of smaller rings and arcs compared to Petrillo et al. (2017) for making the new ConvNets more sensitive to this kind of objects with respect to the old one. The source positions are chosen uniformly within the radial distance of the tangential caustics plus one effective radius of the source Sérsic profile. This leads our training set to be mostly composed of high-magnification rings, arcs, quads, folds and cusps rather than doubles (Schneider et al., 1992). To add complexity in the lensed sources, besides the Sérsic profile, we add between 1 and 5 small circular Sérsic blobs. The centers of these blobs are drawn from a Gaussian probability distribution function (PDF) around the main Sérsic source. The width of the standard deviation of the PDF is the same as the effective radius of the main Sérsic profile. The sizes of the blobs are chosen uniformly within 1-10% of the effective radius of the main Sérsic profile. The Sérsic indices of the blobs are drawn using the same prescription as for the main central source. The amplitudes of the blobs are also chosen from a uniform distribution in such a way that the ratio of the amplitude of an individual blob to the amplitude of the main Sérsic profile is at most 20%.

Moreover, we add Gaussian Random Field (GRF) fluctuations to the lens potential, which, to a first order approximation, make the lens sample more realistic by adding small scale substructures (Chatterjee & Koopmans, 2018). The GRF realizations we added in our simulations all follow a power law power-spectrum with a fixed exponent -6 , which is to the first order a good approximation of substructures in lens plane in the Λ CDM paradigm (Hezaveh et al., 2014). The variances of the realizations are drawn from a logarithmic distribution between $10^{-4} - 10^{-1}$ about mean zero in the units

Figure 3.1: Examples of RGB images of simulated strong lens galaxies used to train the ConvNets. The lens galaxies are observed KiDS galaxies, while the lensed sources are simulated, as described in Sect. 3.2.1.3.

of square of the lensing potential. This yields both structured sources and lenses that are not perfect SIE.

For each source a realistic color is simulated to create images in g , r , i -bands. In order to produce realistic 3-band images we extract magnitudes from “COSMOS” models in LE PHARE (Arnouts et al. 1999; Ilbert et al. 2006). This library of spectra, consists of 31 models, used for COSMOS photo- z (Ilbert et al. 2009). The basic “COSMOS” library is composed of 8 templates for elliptical/S0 galaxies, 11 for spiral types, and 12 for galaxies with star-burst ages ranging from 0.03 to 3 Gyr, allowing us to span a wide range of galaxy types and colours. In order to simulate the typical blue arcs observed in most of the observed lenses, we choose models bluer than S0 and calculate observer-frame magnitudes in the three KiDS wavebands g , r and i for model spectra redshifted up to a redshift of $z = 3$ with a 0.1 binning. Moreover, to populate the magnitude-space more uniformly, we perturb the three magnitudes adding to each of them a random number uniformly extracted from the range $[-0.1, 0.1]$ mag. We also take into account dust extinction by considering a color excess $E(B - V)$, we extract it from a normal distribution with $\sigma = 0.1$ and mean 0 considering only the positives values. In this way we obtain a small extinction correction in order to avoid very red lensed sources which, in the real universe, are much rarer than blue ones. We adopt a typical extinction curve with $R_V = 3.1$, using the relation $A_x = R_x E(B - V)$ where x represents the value for the g , r and i SDSS-filters that can be found in Table 2 of Yuan et al. (2013). Finally, we convolve the three images with an average KiDS-DR3 PSF for each different band: with a FWHM of ~ 0.86 arcsec for g , ~ 0.68 arcsec for r and ~ 0.81 arcsec for i (de Jong et al., 2017).

Table 3.1: Range of parameter values adopted for simulating the lensed sources. The parameters are drawn uniformly, except for Einstein and effective radius, as indicated. See Sect. 3.2.1.3 for further details.

Parameter	Range	Unit
Lens (SIE)		
Einstein radius	1.0 - 5.0 (log)	arcsec
Axis ratio	0.3 - 1.0	-
Major-axis angle	0.0 - 180	degree
External shear	0.0 - 0.05	-
External-shear angle	0.0 - 180	degree
Main source (Sérsic)		
Effective radius (R_{eff})	0.2 - 0.6 (log)	arcsec
Axis ratio	0.3 - 1.0	-
Major-axis angle	0.0 - 180	degree
Sérsic index	0.5 - 5.0	-
Sérsic blobs (1 up to 5)		
Effective radius	(1% - 10%) R_{eff}	arcsec
Axis ratio	1.0	-
Major-axis angle	0.0	degree
Sérsic index	0.5 - 5.0	-

3.2.2 Creating the training set

The data presented above are used to build the training set which is composed of mock strong-lensing systems (labeled with a 1) and non-strong-lensing systems (labeled with a 0), i.e., objects without lensing features. In the following we outline the procedure used to build the two kinds of objects in the training set.

Mock strong-lensing systems: To create mock strong-lensing systems we carry out the following procedure:

- (i) We randomly choose a mock lensed source (Sect. 3.2.1.3) and a LRG (Sect. 3.2.1.1); we rescale the brightness of the simulated source to the peak brightness of the LRG in the r -band multiplied by a factor α randomly drawn from the interval $[0.02, 0.3]$. This accounts for the typical lower brightness of the lensing features with respect to the lens galaxies;
- (ii) we stack the LRG and the mock source for each one of the three bands;
- (iii) for the single-band images, we clip the negative values of the pixels to zero and performing a square-root stretch of the image to emphasize lower luminosity features. Instead, we create 3-band images with HUMVI that operates an arcsinh stretch of the image following the Lupton et al. (2004) composition algorithm;
- (iv) finally, we normalize the resulting images by the galaxy peak brightness (only for single-band images).

Some examples of mock strong-lensing systems obtained in this way are shown in Fig. 3.1.

Non-strong-lensing systems To create the non-strong-lensing system sample we carry out the following procedure:

- (i) we choose a random galaxy from either the LRG sample (with a probability of 20%) or from the contaminant sample (80% probability);
- (ii) we clip the negative values of the pixels to zero and performing a square-root stretch of the images. We create 3-band images with HUMVI;
- (iii) we normalize the images by the galaxy peak brightness (only for single-band images).

Finally, we augment the images, which is a standard procedure in machine learning (see e.g., Simard et al. 2003). It is used to avoid over-fitting by expanding artificially the training set through different transformations of the images. Before feeding the images to the ConvNets, we apply the following transformations:

Figure 3.2: ROC curves for the 1-band (blue) and 3-band (red) ConvNet. Each point of the curves is the true positive rate vs. false positive rate for different values of threshold for p (decreasing from left to right; some values are shown on the curve for reference).

- (i) a random rotation between 0 and 2π ;
- (ii) a random shift in both x and y direction between -4 and +4 pixels;
- (iii) a 50% probability of horizontally flipping the image;
- (iv) a rescaling with a scale factor sampled log-uniformly between 1 and 1.1.

All transformations are applied to both the mock strong-lensing systems and the non-strong-lensing systems. The final set of inputs of the ConvNets are postage stamps of 101 times 101 pixels which correspond to $\sim 20 \times 20$ arcsec. The images are produced in real-time during the training phase. For more details on the training phase see Appendix A.

3.3 Analysis

After the training is completed, the ConvNets must be tested in order to assess whether the training was successful. In this section we define the metric for evaluating the results and evaluate the performances of the ConvNets on a dataset composed by non-lenses and mock lenses in comparison to Petrillo et al. (2017).

Figure 3.3: Percentage of false negatives (i.e., the percentage of lenses that have been misclassified) for bins of R_E and α defined as the ratio between the peak brightness of the lens galaxy and the lensed source.

Figure 3.4: Number of detections as a function of the threshold for p for the 1-band (blue) and the 3-bands ConvNet (red) compared to the ConvNet of Petrillo et al. (2017) (grey).

3.3.1 Performance metric

To evaluate the performances of the ConvNets we use:

- the true-positive rate (TPR), which measures the fraction of positive objects (in our case the lenses) detected by the algorithm. It is given by the ratio between the number of real positive (the number of real lenses that algorithm finds) and the sum of the latter and the number of false negatives (the lenses that the algorithm does not find):

$$\text{TPR} = \frac{N_{\text{TruePositives}}}{N_{\text{TruePositives}} + N_{\text{FalseNegatives}}} \in [0, 1]; \quad (3.2)$$

- the false-positive rate (FPR), which measures the fraction of negative objects (non-strong-lensing systems) misclassified as positives (lenses). It is given by the ratio between the number of false positive (the number of not lenses that algorithm misclassifies as lenses) and the sum of the latter and the number of true negatives (the non lenses that the algorithm classifies correctly)

$$\text{FPR} = \frac{N_{\text{FalsePositives}}}{N_{\text{TrueNegatives}} + N_{\text{FalsePositives}}} \in [0, 1]; \quad (3.3)$$

- these two quantities can be used to build Receiver Operating Characteristic (ROC) curves which allow to check at a glance the

degree of completeness and contamination of a binary classifier. ROC curves are created by plotting TPR as a function of FPR varying the threshold of detection for p between 0 and 1. This allows us to tune the value for the threshold for p in order to get the desired amount of TPR and FPR for a given classification problem. In our case p is the output of the ConvNet and we can tune the p -threshold depending how many lens candidates we desire and what level of contamination is deemed to be acceptable.

3.3.2 Performance

The ROC curves for a test-set composed of 5000 mock strong-lensing systems and 5000 non-strong-lensing systems created, as described in Sec. 2, are shown in Fig. 3.2. In general the 3-bands ConvNet has a better performance than the 1-band ConvNet, retrieving more mock strong-lensing systems than the 1-band ConvNet. On the contrary, the 1-band ConvNet is less contaminated by false positives at higher values of the threshold for p . Since gravitational lenses are rare events, it is important to keep a low value of FPR. Otherwise a candidate sample selected from real data would be dominated by false positives and a large amount of time would be needed to discard them through a visual inspection. In Fig. 3.3 we show, for a fiducial value for the threshold of the detection $p = 0.8$, the percentage of false negatives (i.e., the percentage of lenses that have been misclassified) as a function of the Einstein radius, R_E , and the source over lens-galaxy brightness contrast, α , defined in Sect. 3.2.2. Lenses with small Einstein radii and low-contrast lensed images are, as expected, the ones with which the ConvNets struggle the most. This suggests that our mock lens training samples currently covers the range in which lenses are found most easily. Smaller lenses are effectively smeared to an unrecognizable configuration by the PSF and fainter lensed sources will be too noisy to detect. In Fig. 3.3 we also see that the accuracy decreases for larger Einstein radii, possibly due to the fact that we covered the Einstein radius with logarithmic distribution, focusing slightly more on the small-separation systems, and secondly because their lensed images are more likely to blend in to the local environment and therefore harder to distinguish from nearby galaxies by the ConvNets.

However, because our goal is to efficiently select true strong-lenses in real astronomical observations, therefore it is necessary to assess the TPR and FPR when the ConvNets are applied to real data where the performance might be worse than on a simulated dataset.

Figure 3.5: The first row shows the distribution of the scores for the 56 candidates with $p > 0.5$ selected in Petrillo et al. 2017. The scores of the 1-band (blue) and 3-bands ConvNet (red) are compared with those of Petrillo et al. (2017) (grey). The second row shows the

Figure 3.6: Same as Fig. 3.5 but for a subsample of 22 candidates selected as described in Sect. 3.4.1.

3.4 Application to real data

Testing the ConvNets on real data is fundamental, since the algorithms have been trained on a mixture of real and simulated data. It is not trivial how the method will perform on a completely real dataset, since the domain of application is slightly different with respect to the domain where the ConvNets have been trained on. Ideally, the morphologies in the ensemble of simulated strong-lens systems and non-strong-lens systems would be a fair representation of all morphologies observed in their equivalents in real observations.

Hence, to properly analyze the performances of the ConvNets, we apply them to the full LRG sample composed of 21789 galaxies extracted from 255 square degrees as described in Sect. 3.2.1.1. Using the same LRG sample of Petrillo et al. (2017) allows us to assess whether there has been any improvement with respect to our previous work. For each galaxy image we opt to obtain an average prediction given by the average of the p 's for the original image and the images obtained operating a rotation of 90, 180 and 270 degrees respectively. Generally this procedure allows to increase the accuracy of the classifications.

3.4.1 Results on the LRG sample

In Fig. 3.4 we show the number of lens candidates detected varying the threshold p . The 1-band ConvNet detects more lens candidates compared to the 3-band one for any given threshold for p . For each of the three ConvNets it holds that the lower the threshold in p is set, the more candidates will have to be inspected visually. In other words, one wants to set as the threshold to an as low as possible value that yields both a sufficiently large sample of candidates and a sufficiently high TPR for the purpose of the scientific project. In Petrillo et al. (2017) we used visual inspection to select visually promising strong-lens candidates within the sample of systems assigned with a $p > 0.5$ by the ConvNet. This sample contains 56 candidates. Moreover, in Petrillo et al. (2017) we selected a subsample of 22 candidates based on the agreement between their expected Einstein radii, computed from the stellar mass or the velocity dispersion of the candidate lens galaxies, and the actual galaxy-image configurations. This does not guarantee that the 22 candidates are actual lenses but it allows us to exclude the cases with more implausible configurations. Figs. 3.5 and 3.6 compare the p -values for these two samples assigned by the two new ConvNets to those assigned by Petrillo et al. (2017) ConvNet.

Figure 3.7: Sources classified by the new ConvNets with $p < 0.5$ which were classified in Petrillo et al. 2017 with $p > 0.5$. On each image the score assigned by the single-band ConvNet (blue), the multi-band ConvNet (red) and Petrillo et al. (2017)'s ConvNet (grey) are reported.

We note that the p -values of the new ConvNets have a noticeable peak at high values that becomes even more pronounced considering only the 22 candidates. In particular, the single-band ConvNet selects high-confidence candidates assigning high values of p . This is a fair improvement of the performance of the algorithm since there is a larger clustering of the higher visually ranked candidates toward the high p -values. Instead in Fig. 3.7 we show the subset of the 56 galaxies that the new ConvNets classify with $p < 0.5$. For the 1-band finder there are not clear candidates that would be lost: maybe a couple of galaxies could be considered acceptable lens candidates, while for the rest a low p -value it is the ideally desired output. In particular, three candidates (third, fourth and fifth galaxy in Fig. 3.7, which have been selected as lenses in Petrillo et al. (2017) by visual inspection but after a more careful analysis have been revealed as false positives (likely a merger and two ring galaxies), are classified as non-lenses. Thus, the new finder does cumulatively a better job in excluding contaminants and selecting lens candidates. Instead the 3-bands lose some acceptable candidates, but more importantly misidentifies a known gravitational lens and a clear good candidate (first and second galaxy in Fig. 3.7). This needs further investigation, thus, in the following subsection, we analyze the behavior of the two lens-finders on a small sample composed by real lenses and striking lens candidates.

3.4.2 Application to a small sample of clear lens candidates

Additional insights on ConvNet performance can be obtained from inspecting the results on a set of real lenses and striking lens candidates. We gather a set of 6 galaxies composed as follows (see Fig. 3.8):

- The four confirmed lenses known in literature which are present in our LRG sample: J085446-012137 (Limousin et al., 2010), J114330-014427 and J1403+0006 (Bolton et al., 2008), J115252+004733 (More et al., 2017);
- Three extremely likely lenses found in Petrillo et al. (2017), i.e. KSL713, KSL327 and KSL427;

The p -values for each of these galaxies are shown in Fig. 3.8 and Table 3.2. It is immediately noticeable that the 1-band ConvNet is the best-performing one: the 3 extremely likely lenses and 3 of the 4 known lenses are identified correctly with a very high p -value; Instead, the 3-bands ConvNet gives more weight to the source colours: it identifies easily blue-features but it struggles with other colours.

This could be due to the larger number of training examples with blue colours.

Moreover, both the new ConvNets presented in this paper are able to pick up the quad lens J115252+004733 not selected in (Petrillo et al., 2017). This piece of improvement is possibly due to the improved realism of the mock lensed sources adopted for this new ConvNets, collecting a larger variety of configurations with respect to Petrillo et al. (2017). To validate this hypothesis, we train another 1-band ConvNet with the same mock lensed sources of Petrillo et al. (2017). This new ConvNet does not detect the quad (see Table 3.2), thus confirming the idea that the realism of the simulations plays a key-role in the performance of the lens-finder.

To further understand the role of the simulations, we also train a ConvNet with the same configuration of Petrillo et al. (2017) but with the simulated strong-lensing systems produced for this work (see Sect. 3.2.1.3). Also in this case the quad is detected (see Table 3.2) even if the performance is worse than the new 1-band ConvNet. This could be due to the different architecture of the two ConvNets.

To conclude, despite the limited size of the control sample presented in this section, the 1-band ConvNet is the one which has generally the best performance under all conditions and it seems the best one to systematically apply to real data, both in terms of purity and completeness. The 3-band set-up, instead, is generally biased toward bluer colours and sometimes has failed to select known or very promising strong-lensing systems.

Table 3.2: Scores given by ConvNets on the sample described in Sect. 3.4.2. The 1-band and 3-band ConvNets presented in this work are compared with Petrillo et al. (2017) (old ConvNet). In the 5th column are shown the scores for a ConvNet with the same architecture of the 1-band ConvNet but trained with the mock lensed sources used in Petrillo et al. (2017) (1-band/old mocks). In the 6th column are reported the scores for a ConvNet with the same architecture of Petrillo et al. (2017) but trained with the same training set of this work (old ConvNet/new mocks).

ID	1-band	3-bands	old ConvNet	1-band/old mocks	old ConvNet/new mocks
J085446-012137	1.000	0.998	0.696	0.148	0.996
J114330-014427	0.998	0.403	0.997	0.947	0.605
J1403+0006	0.360	0.456	0.5	0.038	0.419
J115252+004733	0.992	0.887	0.5	0.321	0.904
KSL713	0.999	0.611	0.942	0.990	0.942
KSL427	0.999	0.999	0.943	0.982	0.752
KSL327	0.933	0.012	0.522	0.587	0.479

3.4.3 Visual inspection

To further analyze the performances of the ConvNets, we visually inspect the sources in the LRG sample (see Sect. 3.4) selected by the ConvNets with a threshold $p > 0.8$. To compare the results, we visually inspect the sources selected with the same threshold, $p > 0.8$, by the ConvNet in Petrillo et al. (2017). This yields 721 sources to inspect for the 1-band ConvNet, 390 for the 3-band ConvNet and 292 sources for the ConvNet of Petrillo et al. (2017) ConvNet. The visual inspection is carried out by 3 observers who have three possible choices: “Yes, it is a lens”, “Maybe it is a lens”, “No, it is not a lens”. For each source the classifiers visualize the g , r and i fits files and a $g-r-i$ composed image with the software STIFF⁶ (Bertin, 2012).

After the visual inspection, we build ROC curves assigning a ground truth for each source from the result of the visual inspection. ROC curves are shown in Fig. 3.9 for two different ground truths: one where we assign the label “lens” to the sources that received at least two times *Maybe*, and another where “lenses” are the sources rated with 3 times *Maybe* or at least one *Yes* and one *Maybe*. In this way we have two different ground truths: one where also dubious lens candidates are labeled as “lenses”, and one more conservative where only more solid lens candidates are labeled as “lenses”. In Fig. 3.10 we show the absolute number of true positives versus false positives as a function of the threshold p for the same ground truths. The ConvNets presented in this paper have both higher purity and completeness with respect to Petrillo et al. (2017) and a larger number of lens candidates retrieved. Moreover, the 1-band ConvNet has higher purity and completeness, especially for the more conservative sample, for sources with higher values of p as also the results of Sect. 3.4.2 seems to indicate. In addition, in Fig. 3.11 we show for the conservative classification the number of candidates detected by the ConvNets varying the threshold for p and how many of these are False Positives. The True Positives and False Positives do not grow at the same pace as a function of the threshold p of detection. Moreover, the percentage of false positives is similar for the two new ConvNets. From the visual inspection analysis we conclude that the 1-band ConvNet would be the best choice if one wants to have a small sample of good lens candidates with few contaminants and thus less images to inspect visually.

⁶<http://www.astromatic.net/software/stiff>

Figure 3.8: RGB images of the previously known lenses in our test-sample (first row) and the extremely likely lenses found by Petrillo et al. 2017 in the same test-sample (second row), this sample is discussed in Sect. 3.4.2. The scores assigned by the single-band ConvNet (blue), the multi-band ConvNet (red) and Petrillo et al. (2017)'s ConvNet (grey) are reported on each image.

Figure 3.9: ROC curves of the 3 ConvNets built for a real sample of Luminous Red Galaxies (LRGs) as described in Sect. 3.4.3. For a sample including also more dubious candidates (top) and for a sample with solid candidates only (bottom).

Figure 3.10: ROC curves of the 3 ConvNets built for a real sample of Luminous Red Galaxies (LRGs) as described in Sect. 3.4.3. The absolute number of True Positive vs. False positives for the sample including also more dubious candidates (top) for a sample with solid candidates only (bottom).

3.4.4 Prospects for Euclid

We have employed about 1 hour to visually inspect the 721 targets selected by the 1-band ConvNet, selecting 60 lens candidates in the more conservative case (see previous subsection). For current surveys, investing such an amount of time in visual inspection is still reasonable for obtaining a more complete sample of lens candidates. However, for future surveys, the time dedicated to visually inspect the targets must be minimized considering the extremely higher number of resolved targets in the survey data. Nevertheless, one can choose a higher threshold of detection for the ConvNet and spend considerably less time in visually inspecting the targets. For example, if we select a threshold $p = 0.9919$ for the 1-band ConvNet, which is the threshold to retrieve 3 of the 4 known lenses presented in Sect. 3.4.2, we obtain a total of 35 sources selected by the 1-band ConvNet. These 35 sources are composed of 14 positives and 21 negatives, considering the more conservative result from the visual classification. This translates to roughly $\sim 40\%$ purity, assuming that our visual inspection is accurate enough. By using the code LENSPOP by Collett (2015), we estimate that there are, in 255 square degree of KiDS, ~ 125 lenses with a Einstein radius larger than 1 arcsecond and with our LRG colour cut selection (which corresponds to a cut in redshift of $z \lesssim 0.5$, as shown in Fig. 8 in Petrillo et al. 2017). Hence, we can very roughly say that a candidate sample selected by the 1-band ConvNet with the fiducial threshold $p = 0.9919$ is $\sim 40\%$ pure and $\sim 11\%$ (i.e. 14 out of 125) complete by considering a population similar to that on which we trained our ConvNets. To translate this result in a prediction for Euclid, assuming that ConvNets will perform at least as well on the same domain on Euclid data as they do on KiDS data, we can proceed in the following way. Collett (2015) predicted that there will be ~ 170000 discoverable lenses in Euclid. If we consider only the lenses with an Einstein radius larger than 1 arcsecond and with a redshift $z < 0.5$, the number reduces to ~ 20000 . Therefore, with the fiducial threshold of $p = 0.9919$, it should be possible to select, from a sample of Euclid LRGs, ~ 2200 good gravitational lens candidates, by visually inspecting ~ 5500 galaxies. Considering that an observer needs ~ 5 seconds to inspect a lens candidate, that would imply that ~ 8 hours of visual inspection would be already enough to obtain a sample of lens candidates far exceeding in size to the total number of currently known gravitational lenses. Nevertheless, this is an extremely conservative estimate, since Euclid data will have better spatial sampling (0.1 arcseconds per pixel), PSF FWHM (~ 0.2 arcseconds), and thus image quality, allowing algorithms and humans to better identify

lensing features. Moreover, it will be possible to train the algorithms on a wider parameter space, and thus retrieving a larger number of lenses. We conclude however that current ConvNets, without much adaptation, can yield already enormous numbers of lenses from Euclid data without much human effort.

3.5 Discussion and Conclusions

Automatic lens finders have become a standard method for finding gravitational lens candidates in survey data. They are crucial to cope with the large set of an estimated $\sim 10^6$ gravitational lens candidates that will be produced by upcoming surveys such as Euclid and LSST. Therefore, it is important to build and validate such lens-finders in order to be ready for this anticipated data-avalanche.

It has been shown that Convolutional Neural Networks (ConvNets) are one of the most promising methods for finding lens candidates (see, e.g., Metcalf et al. 2019). ConvNets achieve outstanding results when the data used to train the algorithm are very similar to the data where the algorithm is intended to be applied. However, this is not necessarily the case when the domain of application is different from the domain where the algorithm is trained on. There is an active field of research investigating how to adapt the algorithms, trained on one domain to other domains (see, e.g., Csurka 2017 for a review). In all published cases, ConvNet lens-finders are trained using simulated mock lensed sources. Moreover, in many cases, the lens-finders are tested only on fully simulated datasets. This does not ensure that the lens-finders will perform in a satisfactory way on real survey data as shown in this paper. It is important to conduct a thorough analysis of the performances on real data. Optimally, one would like to build a benchmark where all the lens-finders can be tested and compared against. Data from surveys such as KiDS and DES, with their identified lens candidates, could be used to build such a benchmark.

In this paper we have tested two lens-finders based on a Convolutional Neural Network (ConvNet) for selecting strong gravitational lens candidates in the Kilo-Degree Survey (KiDS). One finder just uses r -band images while the other uses RGB images composed with g , r , i images. To train the algorithms, we have generated a large sample of simulated lensing features on top of real colour-magnitude selected galaxies from KiDS. Both the lens-finders are able to identify real lenses and good lens candidates in the survey. The performance of the two lens-finders is similar but the 3-bands finder seems to under-perform when the lensed sources do not exhibit

a bluer colour. This is most likely due to the fact that the mock lensed sources in the training set have mostly blue colours. Although genuine lensed sources are usually blue, this could select against non-blue sources. One way of dealing with this issue could be populating the mock lensed sources with a wider selection of colours and, in addition, dividing the training set in different classes of colours. This would help not only the ConvNet to explicitly classify sources with different colours but could also improve the general classification given the new information provided. In any case, the power of the single-band set-up is particularly encouraging in view of the Euclid mission data which will rely especially on the VIS band.

In addition, we have tested and compared the lens-finders with a similar one presented in Petrillo et al. (2017). The lens-finders presented in this work have a better performance, i.e., they have both a better completeness and purity (Sect. 3.4.3) and also they tend to classify more probable lens candidates with higher output values (see Sect. 3.4.2). This is a fair improvement of the performance which implies that selecting candidates with high output values we will have a purer sample which turns out to be convenient for visual inspection and/or spectroscopic follow-ups. Indeed, the larger the future survey will become (see e.g. Euclid and LSST) the more prohibiting the visual inspection of candidates will be, hence the purity of machine learning tools will be crucial for the automatic selection of sources to be set on the queue of spectrographs for their spectroscopic confirmations. The differences between this work and Petrillo et al. (2017) can be summarized in three points:

- more complex mock lensed source simulations;
- a modified ConvNet architecture;
- a slightly larger training set.

These differences have contributed to the improvement of the performance, but our analysis presented in Sections 3.4.2 indicates that the main reason is the improved mock lensed sources simulations. In this work the simulated sample is more populated with smaller sized lensed source galaxies and the sources exhibit more complexity in their structure, i.e., the presence of substructures in the lensed sources and a Gaussian Random Field perturbing the lens potential.

The ConvNet lens-finders can be tuned in terms of completeness and purity according to the specific needs of the science project. If one wants to have a more complete sample of lenses, a low threshold of detection can be chosen and the lower purity can be corrected by visually inspecting the candidates, something still feasible for current surveys. On the other hand, a purer (but less complete) sample of candidates can be obtained choosing a higher threshold. We have shown in Sect. 3.4.4 that by using a

high threshold for detection will be already enough to retrieve in Euclid a sample of lenses exceeding in size the total number of currently known gravitational lenses.

A series of possible improvements can be applied to the lens-finders. As we have shown, the performance strongly depends on the composition of the training set. Hence, making the lens simulations more realistic and using real gravitational lens candidates in the training set would probably improve the quality of the classification. Also enlarging the training set with more KiDS galaxies would probably help, as well as adding more labels for helping the ConvNets to discriminate between different classes of galaxies (e.g., adding labels for ellipticals and spirals). Moreover, particular care can be put in producing lens simulations where the S/N is always high enough for the lensing features to be recognizable by an observer. Another possibility would be to specialize the ConvNets in recognizing different kinds of sources (e.g., large/small-separation systems or extended/compact-sources). This could be obtained by either training different ConvNets with specialized training sets or using a single ConvNet trained with a unique training set but with multiple labels rather than with a binary classification scheme. Instead, on the algorithm side, a trivial improvement could be the so called *ensemble averaging*, i.e., averaging the output of different ConvNets in order to possibly reduce the statistical noise of the classification. An approach experimented by, e.g., Schaefer et al. (2018) for identifying strong lens simulations.

Finally, in a forthcoming paper we will apply the algorithms to ~ 900 square degrees of the KiDS survey starting a systematic census of strong gravitational lens candidates named “*LinKS*” (Lenses in the Kilo Degree Survey).

Acknowledgements

CEP thanks Leon Doddema, Ewout Helmich, Jelte de Jong and Kateryna Frantseva for help and support. CEP, CT, GV, and LVEK are supported through an NWO-VICI grant (project number 639.043.308). The authors thank the anonymous referee for the insightful report. SC has been financially supported by a grant (project number 614.001.206) from the Netherlands Organization for Scientific Research (NWO). GVK acknowledges financial support from the Netherlands Research School for Astronomy (NOVA) and Target. Target is supported by Samenwerkingsverband Noord Nederland, European fund for regional development, Dutch Ministry of economic affairs, Pieken in de Delta, Provinces of

Figure 3.11: Results of the more conservative visual classification of the candidates selected by the ConvNets (Sect. 3.4.3). On the left: number of False Positives and True Negatives for the 1-band and 3-bands ConvNet as a function of the threshold of detection p . On the right: percentage of False Positives as a function of the threshold of detection p for both the ConvNets.

Groningen and Drenthe. NRN acknowledges financial support from the European Union Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement N. 721463 to the SUNDIAL ITN network. GAMA is a joint European-Australasian project based around a spectroscopic campaign using the Anglo-Australian Telescope. The GAMA input catalogue is based on data taken from the SDSS and the UKIDSS. Complementary imaging of the GAMA regions is being obtained by a number of independent survey programmes including GALEX MIS, VST KiDS, VISTA VIKING, WISE, Herschel-ATLAS, GMRT, and ASKAP, providing UV to radio coverage. GAMA is funded by the STFC (UK), the ARC (Australia), the AAO, and the participating institutions. The GAMA website is www.gama-survey.org. Based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs 177.A-3016, 177.A-3017, and 177.A-3018, and on data products produced by Target/OmegaCEN, INAF-OACN, INAF-OAPD, and the KiDS production team, on behalf of the KiDS consortium. OmegaCEN and the KiDS production team acknowledge support by NOVA and NWO-M grants. Members of INAF-OAPD and INAF-OACN also acknowledge the support from the Department of Physics and Astronomy of the University of Padova, and of the Department of Physics of University of Federico II (Naples). This publication has been made possible by the participation in the Galaxy Zoo project of more than 20.000 volunteers from around the world, with almost 2 million classifications provided. Their contributions are individually acknowledged at <http://authors.galaxyzoo.org/>. The data are generated via the Zooniverse.org platform, development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation.

Appendix 3.A Technical details

3.A.1 Convolutional Neural Networks

The input data of a Convolutional Neural Network (ConvNet) have a topological structure, e.g., the images of cutout around galaxies that we use in this paper. The input images can be seen as a set of matrices $\mathbf{X}_k$ with $k = 1, 2, \dots, K$ (in the case of our 3-bands ConvNet $K = 3$). The convolution property of the network is given by the convolutional layer which is composed by n kernels that produce n feature maps Y_n as in the

Figure 3.12: A building block of a ResNet. Image from He et al. 2015b. *Relu* is the non-linear function used between the convolutional layers; *identity* is the identity function that maps the input x into itself.

following equation:

$$Y = \sigma \left(\sum_{k=1}^K W_k * X_k + B \right), \quad (3.4)$$

where $*$ is the convolution operator, σ is a non-linear function (the so-called activation function), W_k are the K weight arrays with $k = 1, 2, \dots, K$, representing a filter with its bias given by the constant array B . The weights W and the biases B are the parameters that are determined during the training of the algorithm. Convolutional layers are sequentially stacked such as the input of the deeper layers are the feature maps of the preceding layer. There are far fewer parameters to be determined in a ConvNet as compared to a neural network because convolutional layers are only locally connected to its inputs as opposed to the fully connected layers of general neural networks.

It has been shown that ConvNets work better for image classification tasks than classical neural networks. One could expect this result from the following characteristics of the algorithm. ConvNets read images preserving their structure as a 2d matrix (3d in the case of multi-band images) rather than transforming them in one-dimensional vectors where the topological information is lost. This turns out to be important because the relevant features to classify an image (e.g, a human face) are normally localized regions of pixels (e.g., eyes, mouth, etc.) and their spatial relations (e.g., the distance between eyes and mouth). The convolution operation of the filters with the image localizes where those patterns of pixels are in the image. The feature maps produced with this operation highlight those

predictive features in the images. During the training phase the optimal weights to be used in the filters are learned, which in turn means that the ConvNet learns which are the features useful for classifying the images. What usually happens is that the filters learn to extract simple features in the first layers (usually they act as edge detectors), while deeper filters learn more complicated and abstract features. See Zeiler & Fergus (2013) for a thorough analysis.

3.A.2 Implementation

The ConvNets used in this paper are implemented in Python 3.6 using KERAS (Chollet et al., 2015) with TENSORFLOW backend (Abadi et al., 2016). The training phase took ~ 12 hours on a GeForce GTX 1080. Once trained, the ConvNets take ~ 0.019 seconds to classify one galaxy. For our ConvNets we use a ResNet architecture with 18 layers as described in He et al. (2015b). The building block of a ResNet (see Fig. 3.12) is characterized by a shortcut connection between the input and the output of a stack of a few convolution layers. In this way the input x is mapped into the *residual* mapping $F(x) + x$. A residual mapping should be easier to optimize than an unreferenced mapping. For example in the case where an optimal mapping was the identity function, it would be easier to fit $F(x) = 0$ to satisfy the requirement $F(x) + x = x$ than to satisfy the equality $F(x) = x$.

The ConvNets in this paper are trained, as in Petrillo et al. (2017), by minimizing a loss function of the targets t (1 for lenses and 0 for non-strong-lensing systems) and the predictions p (the output of the sigmoid unit of the ConvNet). The loss function that we minimize is the binary cross-entropy, a common choice in two-class classification problems:

$$L = -t \log p - (1 - t) \log(1 - p) \quad (3.5)$$

The minimization is done via mini-batch stochastic gradient descent and ADAM updates (Kingma & Ba, 2014) using a learning rate of 0.0001. We use a batch size of 64 images and perform 80.000 gradient updates, which corresponds to about six million examples. Each batch is composed by 32 lens and 32 non-lens examples. The weights of the ConvNets are initialized, as discussed in He et al. (2015a), from a random normal distribution with variance $2/n$ where n is the number of inputs of the unit. The initial values of the biases are set to zero.

I know I've made some very poor decisions recently, but I can give you my complete assurance that my work will be back to normal. I've still got the greatest enthusiasm and confidence in the mission
– HAL 9000, *2001: A Space Odyssey*

4. LINKS: DISCOVERING GALAXY-SCALE STRONG LENSES IN THE KILO-DEGREE SURVEY USING CONVOLUTIONAL NEURAL NETWORKS

*C. E. Petrillo, C. Tortora, G. Vernardos, L. V. E. Koopmans, G. Verdoes Kleijn,
M. Bilicki, N. R. Napolitano, S. Chatterjee, G. Covone, A. Dvornik, T. Erben,
F. Getman, B. Giblin, C. Heymans, J. T. A. de Jong, K. Kuijken, P. Schneider, H. Shan,
C. Spiniello, A. H. Wright
2019, Monthly Notices of the Royal Astronomic Society, 484, 3879*

We present a new sample of galaxy-scale strong gravitational-lens candidates, selected from 904 square degrees of Data Release 4 of the Kilo-Degree Survey (KiDS), i.e., the “Lenses in the Kilo-Degree Survey” (LinKS) sample. We apply two Convolutional Neural Networks (ConvNets) to $\sim 88\,000$ colour-magnitude selected luminous red galaxies yielding a list of 3500 strong-lens candidates. This list is further down-selected via human inspection. The resulting LinKS sample is composed of 1983 rank-ordered targets classified as “potential lens candidates” by at least one inspector. Of these, a high-grade subsample of 89 targets is identified with potential strong lenses by all inspectors. Additionally, we present a collection of another 200 strong lens candidates discovered serendipitously from various previous ConvNet runs. A straightforward application of our procedure to future Euclid or LSST data can select a sample of ~ 3000 lens candidates with less than 10 per cent expected false positives and requiring minimal human intervention.

4.1 Introduction

Strong gravitational lenses¹ are composite systems where a massive foreground object (e.g., a galaxy or a cluster) creates multiple images of one or more higher-redshift sources (e.g., galaxies or quasars). Strong lenses are useful for a wide range of cosmological and astrophysical studies (Schneider et al. 1992; Schneider 2006; Treu 2010). For example, they can provide cosmological constraints on the dark energy equation of state (Collett & Auger 2014; Cao et al. 2015) and precision measurements of the Hubble constant (Schechter et al., 1997; Suyu et al., 2013; Bonvin et al., 2017). The information obtained from strong lensing also allows us to study the mass distribution in the inner regions of galaxies: e.g., the fraction of dark matter in their central regions (Gavazzi et al. 2007; Jiang & Kochanek 2007; Grillo et al. 2010; Cardone & Tortora 2010; Tortora et al. 2010; More et al. 2011; Ruff et al. 2011; Sonnenfeld et al. 2015), the slope of their inner mass density profile (Treu & Koopmans, 2002a; Koopmans et al., 2006; More et al., 2008; Barnabè et al., 2009; Koopmans et al., 2009; Shu et al., 2015; Cao et al., 2016; Li et al., 2018) and their dark-matter substructures (More et al., 2009; Vegetti et al., 2012; Nierenberg et al., 2014; Hezaveh et al., 2016). Besides studying dark matter, strong lenses allow us to place constraints on the stellar Initial Mass Function (IMF) when combined with dynamical and stellar population synthesis analyses (Treu et al. 2010; Ferreras et al. 2010; Spiniello et al. 2011; Brewer et al. 2012; Barnabè et al. 2013; Sonnenfeld et al. 2015; Posacki et al. 2015; Spiniello et al. 2015; Leier et al. 2016; Sonnenfeld et al. 2018a; Vernardos 2018). Finally, strong lenses can act as a “Cosmic Telescope”, providing a magnified view of otherwise unresolved background sources (e.g., Impellizzeri et al. 2008; Swinbank et al. 2009; Richard et al. 2011; Deane et al. 2013; Treu et al. 2015; Mason et al. 2017; Salmon et al. 2017; Kelly et al. 2018).

The above-listed studies have typically been carried out using samples of tens to maximally about a hundred massive lens galaxies ($M_{\star} \gtrsim 10^{11} M_{\odot}$), and are often limited to redshifts $z \lesssim 0.5$ and/or are inhomogeneously selected. Current results are therefore often limited by sample size or cosmic variance. Creating more substantial, homogeneously selected samples of gravitational lenses, which extend to lower-mass galaxies and higher redshifts, will reduce the effects of “small-number statistics” and allow an improved study of lens galaxies as a function of galaxy properties and evolutionary state. In particular, Vegetti & Koopmans (2009) estimate

¹Called strong lenses or simply lenses hereafter.

that it is possible to compute sub-halo mass fractions of lens galaxies to a level of $\lesssim 0.1$ per cent with only ~ 50 lens systems. With the same number of lenses, it is possible to reach a per cent level precision in estimating their mass density slopes (Barnabè et al., 2011). Therefore a much larger number of galaxy-scale lenses can improve the outcome from these analyses and enable one to conduct a proper statistical comparison with the results obtained from lens simulations (e.g., Xu et al. 2016; Li et al. 2016; Mukherjee et al. 2018). Moreover, the precision of the value of H_0 can be improved to the level of a few per cent when studying a sample of about 40 strong lenses with measured time delays (Jee et al., 2016; Shajib et al., 2018). Collecting large samples of strong lenses, furthermore, giving us better access to the high-redshift universe and increases the probability of discovering double Einstein-ring (Gavazzi et al., 2008) and other “exotic” lenses (e.g., Tu et al. 2009; Cooray et al. 2011; Brammer et al. 2012; Tanaka et al. 2016). Moreover, samples of homogeneously selected strong lenses are needed to characterize the selection function of a strong lens survey, allowing to map measurements carried out on strong lenses back to the general population of galaxies. We refer the reader to the LSST Science Book (LSST Science Collaboration et al., 2009) and the Euclid Strong Lensing white paper (Euclid Strong Lensing team, 2018, in prep) for a more detailed discussion of future scientific applications of strong gravitational lenses.

The largest homogeneously-selected sample of confirmed strong lenses is the Sloan Lens ACS Survey (SLACS; Bolton et al. 2006, 2008), which yielded more than a hundred spectroscopically confirmed strong lenses with complete redshift information and high-resolution imaging follow-up (with e.g., the Hubble Space Telescope and Keck Observatory Adaptive Optics). In total, all lens surveys combined have produced up to a thousand highly-likely² gravitational lens candidates (e.g., Browne et al. 2003; Faure et al. 2008; Treu et al. 2011; Inada et al. 2012; Brownstein et al. 2012; More et al. 2012; Stark et al. 2013; Sonnenfeld et al. 2013a; Gavazzi et al. 2014; More et al. 2016; Shu et al. 2016, 2017).

Ongoing wide-field optical-IR surveys are expected to make the next giant step forward by yielding thousands of new lenses (Collett, 2015; Petrillo et al., 2017). The first new lens candidates have already been discovered (Petrillo et al., 2017; Hartley et al., 2017; Diehl et al., 2017; Sonnenfeld et al., 2018b; Spiniello et al., 2018; Jacobs et al., 2018; Wong et al., 2018) in the Kilo-Degree Survey (KiDS; de Jong et al. 2013), in

²Not all of these lenses have been spectroscopically confirmed though, but from their image geometry are extremely probable to be strong lenses.

the Hyper Suprime-Cam Subaru Strategic Program (HSC; Miyazaki et al. 2012), and in the Dark Energy Survey (DES; The Dark Energy Survey Collaboration 2005). Similarly large samples are expected from deep sub-mm observations by e.g., the Herschel telescope (Negrello et al., 2010), the South Pole Telescope (SPT; Carlstrom et al. 2011), and the Atacama Large Millimeter/sub-millimeter Array (ALMA)³. These telescopes have already uncovered hundreds of new lens candidates (Vieira et al., 2013; Negrello et al., 2017). Within the next decade, $\sim 10^5$ strong lenses are expected to be found in future surveys (Oguri & Marshall, 2010; Pawase et al., 2014; Collett, 2015; McKean et al., 2015) utilising, e.g., ESA’s Euclid mission (Laureijs et al., 2011), the Large Synoptic Survey Telescope (LSST Science Collaboration et al., 2009) and the Square Kilometer Array⁴. In particular, these surveys will allow lower-mass and higher-redshift lenses to be found, thanks to their deeper and higher angular resolution observations. Moreover, it will become possible to follow up promising targets at an even higher angular resolution with ALMA and the European Extremely Large Telescope (E-ELT). A future SKA-VLBI facility could, in addition, investigate milli-arcsecond angular scales of the lensed images for the effects of dark-matter line-of-sight and sub-halos (Spingola et al., 2018), enabling one to study small deviations from the smooth mass model of the lens.

Strong gravitational lenses are scarce objects within the total population of galaxies. In current surveys, of the order of one strong lens exists per few hundred to a thousand galaxies. This number strongly depends on galaxy mass and selection criteria, with the number of lenses peaking around M^* -galaxies for source-selected samples and at larger masses when lenses are selected as luminous red galaxies (LRGs). Their rarity makes it essential to develop robust lens-finder algorithms and deploy them in streamlined data-processing pipelines. This end-to-end automation will drastically reduce, and possibly prevent entirely, the need for future visual inspection of millions of potential lens candidates (e.g., Lenzen et al. 2004; Horesh et al. 2005; Alard 2006; Estrada et al. 2007; Seidel & Bartelmann 2007; Kubo & Dell’Antonio 2008; More et al. 2012; Maturi et al. 2014; Joseph et al. 2014; Gavazzi et al. 2014; Agnello et al. 2015; Brault & Gavazzi 2015; Chan et al. 2015; Stapelberg et al. 2019a; Hartley et al. 2017; Petrillo et al. 2017, 2019; Jacobs et al. 2017; Sonnenfeld et al. 2018b; Spiniello et al. 2018).

In light of such an automation strategy, we recently developed (Petrillo et al., 2017), and more recently improved upon (Petrillo et al., 2019), a

³<http://www.almaobservatory.org/>

⁴<https://www.skatelescope.org/>

new convolutional neural network (ConvNet) lens-finder algorithm. The objective in this paper is to report on how we use ConvNets in an automated lens-search pipeline, and report on the results of applying these networks to galaxies selected from ~ 900 square degrees of KiDS Data Release 4. The core result that we present is an automatically selected sample of 3500 rank-ordered strong-lens candidates. From this ConvNet pre-selected sample, several subsamples of higher confidence candidates are distilled through human visual inspection.

In Sect. 4.2, we provide a brief introduction to KiDS, the imaging and catalogue data that are used in this paper. In Sect. 4.3, we explain how we select a subsample of intrinsically luminous (red) galaxies from the colour-magnitude diagram of KiDS galaxies, as well as the methodology used to identify gravitational lens candidates within that colour-magnitude selected subsample. In Sect. 4.4, we present the gravitational lens candidates found from the most conservative sample selection. In Sect. 4.5, we apply the networks to a wider selection of galaxies – inherently limited only in their apparent brightness – to examine the efficiency of the algorithm in extremely data-heavy regimes such as those expected from future astronomical surveys, such as with Euclid and LSST, which may also have restricted colour information. In the same section we also present a “bonus sample” of inhomogeneously selected lens candidates that were identified serendipitously during various past experiments in the development of the final ConvNets. Lastly, in Sect. 4.6, we summarise our main conclusions.

4.2 Data from the Kilo-Degree Survey

The Kilo-Degree Survey (KiDS; de Jong et al. 2013) is an ESO public survey carried out with the OmegaCAM wide-field imager (Kuijken 2011) mounted on the VLT Survey Telescope (VST; Capaccioli & Schipani 2011) at the Paranal Observatory in Chile. The telescope, camera, and survey have been designed to obtain images with sub-arcsecond seeing and homogeneous image quality both across the full field of view and throughout the survey execution. In this way the survey yields a large and homogeneous galaxy sample. The size and homogeneity of this sample is required for the surveys primary science drivers, which include placing strong constraints on both the distribution of matter across cosmic time and the cosmological parameters of the universe through weak-lensing measurements; the subtle distortions introduced in galaxy shapes by cosmic shear (e.g., Hildebrandt et al. 2017). At the same time, the combined power of the survey’s superb image quality and wide area makes KiDS optimal for strong-lensing

studies (Napolitano et al., 2016; Petrillo et al., 2017; Spiniello et al., 2018). OmegaCAM has a one square degree field of view, with pixels that have an angular scale of 0.21 arcseconds, and KiDS will survey a total of ~ 1350 square degrees in four optical bands (u , g , r and i) by the end of observations in 2019. The best seeing observations are reserved for the r -band, with the survey exhibiting median point spread function (PSF) full-width at half-maximum (FWHM) values of 1.0, 0.8, 0.65 and 0.85 arcseconds in the u -, g -, r -, and i -bands respectively. The survey depths per-band, as determined by the $5 - \sigma$ limiting magnitudes within a 2 arcsecond circular aperture, are 24.2, 25.1, 25.0, 23.7 in the u -, g -, r -, and i -bands respectively (de Jong et al., 2015, 2017).

In this paper, we make use of 904 tiles⁵ that form a subset of the KiDS Data Release 4 (KiDS ESO-DR4, Kuijken et al. 2018, in prep.). The analysis performed uses imaging data, and derived products, produced within the Astro-WISE information system (Valentijn et al., 2007; McFarland et al., 2013). We make use of the single-band and multi-band catalogues of the KiDS-DR4.

4.2.1 The “full sample”

The target extraction and their associated photometry have been obtained using S-EXTRACTOR (Bertin & Arnouts 1996). To optimise the initial lens searches, we pre-select a sample of luminous galaxies with reliable photometric data. We proceed in the following way:

- (a) We select sources with a S-EXTRACTOR r -band FLAGS value < 4 , thereby including only deblended sources and removing from the catalogue objects with incomplete or corrupted photometry, saturated pixels, or any other blending or extraction related problem.
- (b) We further reject galaxies in areas compromised by, e.g., stellar diffraction spikes and reflection halos, by selecting sources with the flag IMA_FLAGS set to zero for all the four KiDS bands.
- (c) We select sources with a Kron-like magnitude MAG_AUTO in the r -band below 20th magnitude, in order to maximize the lensing cross-section (Schneider et al. 1992).
- (d) Finally, we select sources with flag 2DPHOT equal to 1 (as derived by the star-galaxy separator software 2DPHOT (La Barbera et al., 2008) in order to select secure galaxies.

To reduce the contamination by stars further, we select only objects with a FWHM in r -band greater than the 90 percentile range of the distribution of

⁵The full fourth KiDS data release consists of 1006 tiles, but we have chosen to limit our analysis to the first 904 tiles that were processed by Astro-WISE.

star-like objects within the same tile (those with 2DPHOT equal to zero). We adopt this strategy to reach a suitable compromise between filtering out stars and not excising too many galaxies from the sample. This selection procedure results in a sample of nearly one million (specifically 930 651) targets which we will refer to as the “*full sample*” in the remainder of the paper.

4.2.2 The luminous red galaxy sample

Luminous Red Galaxies (LRGs; Eisenstein et al. 2001) are massive galaxies which, as a result, are more likely to exhibit strong lensing features than other classes of galaxies (see Turner et al. 1984; Fukugita et al. 1992; Kochanek 1996; Chae 2003; Oguri 2006; Möller et al. 2007). We select LRGs from the *full sample*, defined earlier, using the low-redshift ($z < 0.4$) LRG colour-magnitude selection of Eisenstein et al. (2001). We slightly adapt this selection to include fainter and bluer sources:

$$\begin{aligned}
 |c_{\text{perp}}| &< 0.2, \\
 r &< 14 + c_{\text{par}}/0.3
 \end{aligned}
 \tag{4.1}$$

where

$$\begin{aligned}
 c_{\text{par}} &= 0.7(g - r) + 1.2[(r - i) - 0.18], \\
 c_{\text{perp}} &= (r - i) - (g - r)/4.0 - 0.18.
 \end{aligned}$$

The magnitudes are S-EXTRACTOR MAG_AUTO. In this section we chose to limit our analysis to the ASTRO-WISE single-band object detection catalogues. We determine the u, g, r, i photometry for each object using the individual S-EXTRACTOR MAG_AUTO measurements. As these measurements are made using slightly different centroids and the PSF varies significantly between bands, we do not expect this “first-look” LRG selection methodology to be uniform. As our aim is not to compile a complete sample of LRGs, however, we do not expect this decision to impact our conclusions. We note that after the analysis for this project began, Vakili et al. (2019) presented a sophisticated methodology to select LRG galaxies for clustering studies in KiDS-DR3. Future LinKS analyses will investigate adopting this LRG sample. Our selection results on a sample of 88 327 sources, which we refer as the “*LRG sample*” throughout the remainder of this paper. Note that our goal here is to select a reasonable number of massive (LRG) galaxies, without significant contamination by spiral galaxies, but that this sample need not strictly be purely LRGs. We find an average of 98 sources selected per tile with a standard deviation of ~ 43 . This standard deviation is high, but expected given the “first-look” methodology that we have adopted

to compile this sample, in addition to the high levels of cosmic variance expected for this highly biased galaxy sample.

4.3 Searching for Lenses

To find gravitational-lens candidates in KiDS imaging data, we use the ConvNets previously introduced by Petrillo et al. (2019). These networks are significantly improved variants of the original ConvNet presented by Petrillo et al. (2017). ConvNets (Fukushima, 1980; LeCun et al., 1998) represent a state-of-the-art method of pattern recognition (Russakovsky et al., 2015). The networks *learn* how to classify a diverse set of images during the so-called training phase, whereby labelled images are provided to the ConvNet. Its weight parameters are changed to minimise a pre-defined loss function, which expresses the difference between the labels of the images and the output values p (one for each image) of the ConvNet. For a more detailed introduction to ConvNets for finding lenses we refer the interested reader to Petrillo et al. (2017), and to more general reviews by Schmidhuber (2015), LeCun et al. (2015) and Guo et al. (2016).

To evaluate methods for identifying images of simulated gravitational lenses – in preparation for the Euclid mission (Metcalf et al., 2019) – recently an international challenge was organised. The results of this challenge demonstrated that ConvNets, collectively with Support Vector Machines (SVMs), are among the most promising methods for finding lens candidates currently available. As a proof of concept, ConvNets have been used to find new gravitational lens candidates by Petrillo et al. (2017) in the KiDS DR3 and by Jacobs et al. (2017) in the Canada-France-Hawaii Telescope Legacy Survey (CFHTLS) and in DES (Jacobs et al., 2018).

In terms of methodology and target selection our analysis differs from the work done by Spiniello et al. (2018), who have focused their search exclusively on lensed quasar candidates in KiDS, by visual inspecting targets preselected using optical/infrared colours. Lens candidates have also been found in the KiDS DR3 data by Hartley et al. (2017), who trained a Gabor-SVM finder.

4.3.1 Training the Convolutional Neural Networks

We start by giving a brief synopsis of our ConvNets and the training procedure, as reported by Petrillo et al. (2019). Building on our experience, we choose to deploy two different ConvNets. One focusses on utilising the best morphological information by taking the best-seeing, i.e., r -band,

images as input. The other ConvNet exploits colour information in addition to morphological information by taking 3-band RGB images as input. The RGB images are created with HUMVI⁶ (Marshall et al., 2016) using the g , r and i bands. In both cases, the KiDS images have a size of 101×101 pixels (i.e. 20×20 arcseconds) with the central pixel corresponding to the centre of the galaxy of interest. The ConvNets take these images and transform them into a single value, p , which can vary between 0 and 1. This value represents, to some degree (see e.g., Scaerens et al., 2002), the probability that the input image is a lens (see also Section 3.2). The input size of 20×20 arcseconds is chosen to be sufficiently large as to enclose most galaxy-scale lens systems, and sufficiently small as to both avoid contamination by unrelated field objects and allow for a ConvNet with a practical memory requirement⁷.

We use two classes of objects to train the ConvNets: (1) the *lenses* labelled with a 1.0, and (2) the *non-lenses*, labelled with a 0.0.

- (1) For the *lenses*, we use a set of ~ 6000 KiDS LRGs on which we superimpose simulated lensed images. The simulated lensed images ($\sim 10^6$ in number) are composed mostly of high-magnification rings, arcs and quads. The gravitational-lens mass distribution adopted in our simulations is assumed to be that of a Singular Isothermal Ellipsoid (SIE, Kormann et al. 1994) perturbed by additional Gaussian Random Field (GRF) fluctuations and an external shear. An elliptical Sérsic (1968) brightness profile is used to represent the lensed sources, and to which we add several small internal stellar structures (e.g., star-formation regions, satellite galaxies), described by circular Sérsic profiles. For each background source, we extract magnitudes from the “COSMOS” models provided by the code LE PHARE (Arnouts et al. 1999; Ilbert et al. 2006) in order to simulate realistic *gri*-composite images. The lens and source parameters vary accordingly to the values in Table 1 of Petrillo et al. (2019).
- (2) The non-lenses are a collection of $\sim 12\,000$ galaxies from KiDS. This sample is comprised of a supersample of: (a) the same LRGs used for the *lenses*; (b) randomly selected galaxies from the survey with a r -band magnitude brighter than 21; (c) ‘false positives’ (e.g., mergers, ring galaxies, etc.) from earlier ConvNets; and (d) a sample of galaxies that were visually classified as spirals from an on-going GalaxyZoo project (Willett et al. 2013, Kelvin et al., in prep.).

⁶<https://github.com/drphilmarshall/HumVI>

⁷Larger images require a larger numbers of network weights and consequently more computer memory.

Figure 4.1: Histogram of the numerical rankings from the visual inspection of 3500 targets, selected by the ConvNets, by seven human classifiers. See Sect. 4.3.2 for the detailed discussion of the results.

A more detailed description of the training sample preparation, the results of the training phase, and a detailed discussion of the performance of the ConvNets are presented in Petrillo et al. (2019).

Figure 4.2: The candidates classified through visual inspection with the two topmost scores, 70 and 64. Below each image are shown the visual inspection score followed by the p -values of the 1-band and 3-band ConvNets. Each image has dimensions 20×20 arcseconds.

4.3.2 Application to the LRG sample

The ConvNets described in the previous subsection are both applied to the LRG sample, and only targets with $p > 0.8$ (returned from either of the ConvNets) are selected. This threshold is chosen to obtain a reasonable number of ‘true positives’ and, at the same time, not contaminate the sample with a large number of ‘false positives’. Petrillo et al. (2019) present an extensive analysis of the performance of these ConvNets by choosing different p -value thresholds. With this threshold, the 3-band ConvNet picks 1689 candidates, while the one-band ConvNet picks 2510 candidates. These numbers correspond to fractions of ~ 1.9 and ~ 2.8 per cent of the LRG sample, respectively. We find a total of (exactly) 3500 unique candidates with $p > 0.8$ since 699 galaxies are common between both ConvNets. We refer to this sample of 3500 unique targets as the *ConvNet sample*.

By setting the threshold value p to 0.8, however, we still expect the presence of many false positives in the *ConvNet sample* (~ 90 per cent; Petrillo et al. 2019). To validate the candidates, selected by the ConvNets, we conduct a visual inspection: seven of the authors of this paper – referred to as “classifiers” – examine the 101×101 pixels RGB composite image, created with STIFF⁸ (Bertin, 2012). The classifiers have only three possible choices for each source being a lens: *Sure*, *Maybe*, and *No lens*. We translate each of these categories into a numerical value in the same way as was done by Petrillo et al. (2017):

- A: *Sure lens* 10 points.
- B: *Maybe lens* 4 points.
- C: *No lens* 0 points.

As a result, the maximum score that any one galaxy candidate can obtain is 70, i.e. when all human classifiers think it is surely a lens. A histogram with the numerical results of the visual inspection is shown in Fig. 4.1. About ~ 57 per cent of the initial 3500 candidates selected by the ConvNets (i.e., 1983 candidates) have at least one classifier selecting it as a *Sure lens* or *Maybe lens*. Only four candidates achieve the maximum score. Fig. 4.2 presents the eight candidates that received the two highest scores, i.e., 64 and 70. Among them, there is one confirmed quad lens, J115252+004733 (bottom right panel, More et al. 2017). It is worth noting that, within the full *ConvNet sample*, there are five confirmed lenses: J114330-014427, J1025-0035 (Bolton et al., 2008), J085446-012137 (Cabanac et al., 2007),

⁸<http://www.astromatic.net/software/stiff>

CSWA 5 (Christensen et al., 2010) and J115252+004733 (More et al., 2017) classified with scores of 58, 22, 54, 24, 64 and 64 respectively (see Fig. 4.3). Naturally this means that none of these confirmed lenses were flagged as *Sure lens* by all classifiers. However, these sources are often confirmed as lensed through high angular resolution HST (Hubble Space Telescope) follow-up, which makes it unsurprising that they are not classified as secure lenses in ground-based KiDS data. In the LRG sample there are other six known gravitational lenses which have not been identified by our ConvNets. However, the KiDS images of these objects do not exhibit striking lensing features and, thus, they are hardly recognizable as strong lenses.

Figure 4.3: First row: images of 5 known confirmed lenses re-discovered by the ConvNets. Below each image are shown the visual inspection score followed by the p -values of the 1-band and 3-band ConvNets. Second row: known lenses in the LRG sample not identified by the ConvNets. All the images have dimensions 20×20 arcseconds.

Figure 4.4: Images of the same candidate retrieved by the ConvNets in two different survey tiles. The scores from the visual classification (the numbers below the images) are different because of the different quality of the images. Each image has dimensions 20×20 arcseconds.

The visual classification appears to depend on the signal to noise ratio. For example, the candidate SCJ083726+015639, found in HSC data by Sonnenfeld et al. (2018b), is present in two adjacent KiDS tiles, and the ConvNets retrieve it from both tiles (the ConvNets select three more HSC candidates). Nevertheless, the human classifiers, in general, give very different scores to the same candidate depending on the quality of the images (Fig. 4.4). Thus, it is fair to assume that many ‘good’ candidates are lost from our sample if we preferentially select only those candidates with high visual-inspection score. On the other hand, there are also clearly cases where the ConvNets select candidates without any human-identifiable lensing feature being present.

To examine the other extreme of the classification, Figs. 4.5 and 4.6 present the candidates that the ConvNets classify with values of $p > 0.999$, along with the scores from our visual inspection. For the 3-band ConvNet, some of these extremely high-confidence ConvNet candidates received low visual classification scores; there is even a case with visual-inspection score of zero. It is clear that there remains significant disagreements between human and ConvNet classifications, and that both classification methods are prone to some level of bias and error. Nonetheless, Fig. 4.7 demonstrates that the visual-inspection scores and the p -values are indeed correlated. The figure shows a positive correlation between average values of the p -values for different bins of the visual inspection score. Hence, even if the classification schemes from humans and ConvNets differ, both tend to agree to a certain extent on what constitutes a ‘good’ lens candidate.

Figure 4.5: Candidates selected by the 3-band ConvNet with $p > 0.999$. The scores from the visual inspection are shown below the images. In the last row, the lens J1244+0106 is shown twice because it appears twice in the LRG sample since is centred on two different LRGs. Each image has dimensions of 20×20 arcseconds.

Figure 4.5: *continued*

Even if there is no obvious inspection-score below which the candidates are no longer reliable, we nonetheless observe an increase in the fraction of good candidates with increasing score. Therefore by defining some fiducial threshold for the visual inspection score, above which one considers the targets as reliable candidates, we can investigate how the number of retrieved candidates (and the degree of contamination) vary as a function of the threshold set on the value of p . Fig. 4.8 presents these correlations for all the ConvNet candidates and for a “bona fide” sub-sample composed of targets with a visual inspection score ≥ 28 . This is a fiducial value of the score which corresponds to a) *maybe lens* given by all the classifiers or to b) *sure lens* given by two classifiers and *maybe lens* from other two classifiers. In particular, in the left panel of Fig. 4.8 we see how the number of retrieved candidates changes as a function of the value of p , greatly decreasing when p is approaching to 1. This change is more gentle in the case of the “bona fide” sample. The right panel shows that the fraction of “bona fide” systems is increasing with p , reaching the lowest contamination degree when p is close to 1. This latter result confirms the correlation among the visual inspection score and p , previously shown in Fig. 4.7.

Figure 4.6: Candidates selected by the 1-band ConvNet with $p > 0.999$. The scores from the visual inspection are shown below each image. Each image has dimensions 14×14 arcseconds.

4.4 The LinKS sample candidates

We define the “LinKS (Lenses in the Kilo-Degree Survey) sample” as the full sample of 1983 gravitational lens candidates retrieved with $p > 0.8$ and a score from the visual inspection greater than zero. The sample contains five previously confirmed strong lenses (see Fig. 4.3; Cabanac et al. 2007; Bolton et al. 2008; Christensen et al. 2010; More et al. 2017) and 12 lens candidates discovered in the HSC data (Sonnenfeld et al., 2018b; Wong et al., 2018). This sample also contains the “bona fide” subsample, composed of the 89 candidates which have a visual inspection score ≥ 28 , which we defined in Sect. 4.3.2. We note that by relaxing this inspection score requirement further, for example to ≥ 16 (i.e. the score corresponding to four *maybe a lens*), we are able to produce a subsample of 308 candidates. Nonetheless, we opt to define our “bona fide” sample using the more stringent ≥ 28 requirement.

Information about the data products provided for the LinKS sample, along with images for each of the 89 “bone fide” candidates, is provided in Appendix 4.A. Additional information is also provided at the the LinKS webpage⁹.

4.4.1 Candidate properties

In this section we summarise the main characteristics of the LinKS sample. To enable this analysis, we rely on candidates with known spectroscopic redshift publicly available from SDSS DR14, GAMA DR3 and 2dFLenS (Abolfathi et al., 2018; Baldry et al., 2018; Blake et al., 2016). We also incorporate accurate multi-band colours as measured by the Gaussian Aperture and PSF (GAaP) code. Briefly, GAaP produces fluxes measured in Gaussian-weighted apertures, which are modified per-source and per-image, so as to produce seeing-independent estimates flux estimates across different observations/bands. The aperture modification calculation requires that the PSF of the image be both homogeneous and Gaussian, and so prior to running GAaP each survey tile has its PSF Gaussianised over the full field of view. Importantly, GAaP magnitudes are not total, and preferentially weight the central, redder parts of our lens galaxies. This acts to reduce the contamination of the outer (blue) features of the lens candidates (i.e. the lensed arcs), and improve the fidelity of lens-candidate SED models. In this section we have chosen to limit our analysis to the

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

LinKS sample in the KiDS-North patch¹⁰. This selection reduces our LinKS sample to 659 candidates, of which 41 (out of 89) are in the “bona fide” subsample. We show the observer-frame $g - r$ colour in terms of redshift of these candidates in the left panel of Fig. 4.9. Due to our initial selection criteria (see Sect. 4.2.2) all of our candidates exhibit red colours, with $g - r \sim 0.8$ at $z \sim 0$ and $g - r \sim 1.7$ at the highest redshifts $z \sim 0.5$. Visually the “bona fide” candidates seem to sample the colour distribution of the entire sample without bias; they are otherwise unexemplary. To further characterize the sample of candidates, and allow for a comparison with the literature, we then estimate stellar masses for the subsample of our sources with spectroscopic redshifts¹¹. Following Petrillo et al. (2017), we estimate stellar masses using the software LE PHARE (Arnouts et al. 1999; Ilbert et al. 2006), which does a χ^2 fitting between colours from stellar population synthesis (SPS) models and the observed colours. We employ single burst SPS models from Bruzual & Charlot (2003, BC03) and a Chabrier (2001) IMF, allowing the stellar population age to vary and assuming metallicities in the range ($0.005\text{--}2.5 Z_{\odot}$). The maximum age is set by the age of the Universe at the redshift of the galaxy, with a maximum value at $z = 0$ of 13 Gyr. We do not consider internal extinction, and our models assume zero redshift uncertainty. We adopt the GAaP $ugri$ magnitudes MAG_GAaP and related 1σ uncertainties (Kuijken et al. 2018, in prep.), corrected for Galactic extinction using the map by Schlafly & Finkbeiner (2011). The r -band MAG_AUTO is used to correct the results of LE PHARE for missing flux¹². The typical uncertainty on the stellar mass estimates (provided by LePhare) is $\sim 0.1 - 0.2$ dex. Stellar masses are shown as a function of redshift in Fig. 4.9, and compared with SLACS (Auger et al. 2009) and SL2S (Sonnenfeld et al. 2013b) data. Consistently with Petrillo et al. (2017), the selected candidates have redshifts in the

¹⁰The fourth KiDS data release consists of multi-band GAaP catalogues for both the Northern and Southern patches, but we chose to limit our analysis to a preliminary set of 497 tiles that were processed by Astro-WISE at the start of this analysis. We note that some improvements have been made to the GAaP catalogues during the course of this work, in particular the calibration of the u -band zero-points has been refined. We do not expect these updates to significantly impact our conclusions.

¹¹Robust stellar masses are available from the literature for those KiDS galaxies that are also contained in SDSS and GAMA. However, in order to have homogeneous results for all the candidates, we determine the masses for the whole sample using KiDS 4-band photometry.

¹²GAaP magnitudes do not trace the whole galaxy light distribution, for this reason we need to correct the stellar masses $\log M_{\star}^{\text{LE PHARE}}$ for missing flux, using the following formula $\log M_{\star} = \log M_{\star}^{\text{LE PHARE}} + 0.4 * (\text{MAG_GAaP_r} - \text{MAG_AUTO_r})$. This could contaminate the lens mass estimate but, since the lensed sources are usually blue, the impact on MAG_AUTO_r is usually small.

Figure 4.7: Average of the p -values given by the two ConvNets divided in bins of the score of the visual inspection. The vertical bars correspond to the 16-84 percentile of the distributions.

window $0.1 \lesssim z \lesssim 0.5$, with a median value of 0.33, while the stellar masses are typically larger than $10^{11} M_{\odot}$, with an average value of $\sim 2 \times 10^{11} M_{\odot}$. We note, of course, that the choice of IMF significantly influences the final mass estimates; using a Salpeter (1955) IMF instead of a Chabrier IMF causes inferred stellar masses to increase by a factor of ~ 2 with no change to observed colours (Tortora et al., 2009). The “bone fide” candidates are shown in green in both panels. They span a similar range of redshifts and masses as the whole sample, with a marginal indication that they may preferentially sample higher stellar masses.

4.4.2 Predictions and Prospects: Euclid and LSST

Using the LENSPOP code presented in Collett (2015), Petrillo et al. (2017) estimated that the maximally retrievable number of strong lens candidates in a fully complete KiDS survey would be ~ 2400 . For a ~ 900 square degrees area such as that considered in this paper, ignoring the masked area of the survey, we would there expect to find ~ 1700 possible strong lenses. If we further consider only those lenses that satisfy our LRG colour-magnitude cuts (Sect. 4.2.2), and which have an Einstein radius larger than one arcsecond (i.e. the range on which the ConvNets have been trained; see Table 1 in Petrillo et al. 2019), this number reduces further to

Figure 4.8: On the left we plot the number of targets retrieved by the two ConvNets and the number of “bona fide” candidates as a function of the threshold of detection p . On the right we show the percentage with respect to the total number of retrieved candidates as a function of the threshold of detection p .

about ~ 450 retrievable strong lenses. Their average distribution in redshift is consistent with the actual distribution of our retrieved candidates of the previous subsection, peaking at a value of $z \sim 0.3$. Our samples here therefore fully encompass the predicted ~ 450 retrievable strong lenses from LENSPOP: the full sample of LinKS candidates containing $\sim 4\times$ the number of predicted sources, and the bone fide sample containing $\sim 5\times$ too few. We note again, though, that by relaxing the visual inspection score requirement to, e.g., ≥ 16 (the score corresponding to four *maybe a lens*) one can create a wider “bone fide” sample containing 308 candidates; ~ 68 per cent of the retrievable lenses predicted by LENSPOP. Nonetheless, we continue to conservatively consider only the 89 sources in our “bona fide” subsample to be genuine lenses, and conclude that this sample is complete at the level of ~ 20 per cent.

In the following, we predict the number of lenses expected in future surveys utilising the depth and breadth of the future Euclid and LSST surveys, and the performance of our ConvNets in retrieving strong lenses within these future datasets.

Euclid. Collett (2015) predicts that there will be $\sim 170\,000$ potential lenses in Euclid. Petrillo et al. (2019) extended this analysis by estimating the number of lenses with an Einstein radius larger than 1 arcsecond and with a redshift $z < 0.5$, which roughly corresponds to our LRG colour cut selection. This reduces the number of potential strong lenses to $\sim 20\,000$ in the $15\,000$ square degrees of the completed survey. With the same strategy used in this paper, we conservatively estimate that between $\sim 5\,000$ and $\sim 15\,000$ lenses will be retrievable with ConvNets from the completed Euclid survey. These numbers assume that the 1-band ConvNet performs at least as well on Euclid data as it does on KiDS data, in the same parameter-domain, and that it is possible to pre-select LRGs with the aid of ground-based multi-band observations and the IR-bands from Euclid. We note, though, that Euclid data will have better image quality than KiDS, which will allow the training of more effective algorithms over a wider parameter space. Furthermore, it will allow improved recognition and rejection of false positives via visual inspection. These considerations all lead to our assessment that our estimate of the number of retrievable strong lenses is conservative.

LSST. The above forecast can also be performed for LSST, and moreover with greater accuracy, as LSST will observe in the same g , r , and i filters as does KiDS. We find that the number of potentially discoverable lenses

in LSST, with an Einstein radius larger than one arcsecond and with our invoked LRG colour selection, is $\sim 20\,000$ over the $20\,000$ square degrees of the completed survey. Therefore, as in Euclid, we estimate that between ~ 5000 and $\sim 15\,000$ lenses may be retrievable from the completed LSST survey data with our ConvNets.

4.5 The full sample candidates

Visual inspection of strong lens candidates selected by the ConvNets is a time-consuming task. However investing such time to achieve increased purity and completeness of the recovered candidate sample is worth the effort. But lowering the p -value threshold above which lens candidates are defined, or significantly increasing the survey area (and thus significantly increasing the absolute number of $p \geq 0.8$ candidates) naturally only exacerbates this task. As such, performing the visual inspections completed here for much larger target samples, such as those expected from Euclid and LSST, will likely be prohibitive. In these cases, one may want to reduce the number of candidates to visually inspect by increasing the p -threshold required for candidacy definition. However it is unclear how such an increase may influence the number of lens-candidates returned. Furthermore, if the scientific aim is to establish a complete strong-lens sample that is unbiased in its lens properties, then such a high threshold may be counter-productive. The *LRG sample* used in this paper is a distinct sub-sample of massive Early Type Galaxies (ETGs) which lack (active) star formation and therefore have profiles which allow easier separation of foreground lenses from lensed images, which are often blue star-forming galaxies, as demonstrated in SLACS. In this work we use the LRG sample because we expect most of the lenses to be massive ETGs. However, selecting such a sample of galaxies is not always straightforward and can lead to the loss of potential lenses; LRGs do not represent the entire population of galaxies and hence the entire strong-lensing cross-section.

For this reason, it is interesting to see how the ConvNets perform on a less restricted and much larger sample of galaxies. We explore these issues in Sect. 4.5.1 by applying the ConvNets to the *full sample*, but with a higher threshold in p , in order to reduce the visual inspection effort. In Sect. 4.5.2, we then translate the outcome to the planned Euclid and LSST surveys and analyse the advantages and applicability of such a strategy. Finally, in Sect. 4.5.3 we present a composite sample of lens candidates collected from various ConvNets, applied to the *full sample*, that were run during the ConvNet optimisation process. Each of these runs was less efficient than

the final ConvNets employed in the main body of this work, but sometimes yielded distinct lenses which we have subsequently collated.

4.5.1 A high-purity sample

We run the two ConvNets on the *full sample* (930 651 galaxies) rather than on the smaller but purer *LRG sample* (88 327 galaxies). To obtain a sample of lens candidates that is both pure and limited in size, and in order to reduce the visual inspection load, we average the predictions from both ConvNets into a single predictive parameter p . We select candidates with an average value of p larger than 0.999. With this selection we obtain just 30 strong lens candidates (Fig. 4.10); 0.003 per cent of the *full sample*. When visually inspected, we find that this sample is extremely pure and, more in particular, it is composed of¹³:

- 2 confirmed lenses (Cabanac et al., 2007; Bolton et al., 2008);
- 1 candidate discovered by Sonnenfeld et al. (2018b);
- 1 quad recently identified by Sergeyev et al. (2018);
- 14 very-likely genuine lenses;
- 10 potential lenses;
- 2 possible contaminants.

This result attests to the capability of the ConvNets to find lens candidates in a sample slightly different from what it was trained on. We note that 18 of the 30 candidates retrieved in this manner are not part of the *LinKS sample* because they did not satisfy the LRG cut in Sect. 4.2.2 (see Sect. 4.5.3 for more information on these candidates). We note further that it is entirely possible that some of these candidates fail our LRG colour-magnitude selection explicitly because of contamination by the bright blue lensing features that we are attempting to locate; a clear drawback of such an LRG selection with imperfect photometry.

4.5.2 Small high-purity Euclid & LSST samples

Considering that, theoretically, the number of recoverable lenses in ~ 900 square degrees of KiDS is at most ~ 1700 (see Sect. 4.4.2), our recovery of only 30 candidates in Sect. 4.5.1 implies that a $p > 0.999$ setup will only recover ~ 2 per cent of possibly retrievable lenses. If we turn this efficiency into a forecast for the 170 000 total retrievable lenses in the full Euclid survey as predicted by Collett (2015), we expect to find ~ 3000 candidates

¹³This sample has been visually inspected using a classification scheme similar to, but not the same as, the one adopted for the *LinKS sample*. For sake of brevity we omit details about this classification.

Figure 4.9: The $g - r$ observer-frame colour, corrected for Galactic extinction (top panel) and stellar mass (bottom panel) versus redshift for a subsample of 659 ConvNet candidates with spectroscopic redshift available (grey dots). The subsample of best candidates with a visual score larger than 28 are shown as green points. Stellar masses (see Sect. 4.4.1) are compared with the SLACS sample from Auger et al. (2009, red triangles) and the SL2S sample from Sonnenfeld et al. (2013b, blue squares).

Figure 4.10: Images of the sample of candidates retrieved running the two ConvNets on the *full sample*, averaging the predictions and selecting those with $p > 0.999$. This sample is > 90 per cent pure and requires very little human intervention. The block of 12 galaxies on this page is part of the *LRG sample* while the 18 galaxies on the next page are exclusively part of the *full sample*. More information on the latter candidates can be found in the Appendix. Each image has dimensions 13×13 arcseconds.

Figure 4.10: *continued*

with a > 90 per cent purity which are retrievable with minimal human intervention. Such a sample represents the often called “low-hanging fruit” of strong lenses within Euclid, as these sources are expected to occupy a limited but easily accessible part of parameter space (i.e., large Einstein radii and low redshifts). Note again that we expect this number to be conservative, as with our other forecasts presented in Sect. 4.4.2, as Euclid lenses will be observed with a much higher angular resolution than KiDS lenses, and will be detected with ConvNets trained on higher fidelity data. Near-infrared colours will also help to down-select lens candidates since, being less sensitive to the dust and mapping a wider wavelength baseline, they will provide a more efficient way to separate LRGs from star forming galaxies.

Nonetheless, even in this conservative case, the number of lenses forecast here would be one to two orders of magnitude larger than the number detected in any previous or ongoing strong lens survey. Finally, as in Sect. 4.4.2, a similar number of easy candidates may be expected from LSST surveys.

4.5.3 The “bonus sample”

The sample presented in this section includes 200 strong lens candidates discovered serendipitously during previous ConvNet runs that are not part of the LinKS sample. The candidates in this *Bonus sample* have not gone through the same rigorous visual inspection as those in the LinKS sample, and subsequently cannot be considered to be as statistically well defined. However if we apply the ConvNets to these candidates with a threshold $p > 0.8$, 160 candidates pass this threshold in at least one of the two ConvNets, i.e., 80 per cent of the sample. Detailed data related to this sample can be found online¹⁴ (see the Appendix). This sample contains eight HSC survey lens candidates (Sonnenfeld et al., 2018b) and four confirmed lenses J1452-0058 (Bolton et al., 2008), J142449-005322 (Tanaka et al., 2016), J010127-334319 (Bettinelli et al., 2016) and KiDS0239-3211 (Sergeyev et al., 2018).

4.6 Discussion and Conclusions

In this paper, we present several samples of lens candidates from the Kilo-Degree Survey (KiDS) which likely contain several hundred strong

¹⁴<http://www.astro.rug.nl/lensesinkids>

gravitational lenses. To generate these samples, we apply two new lens-finder algorithms – based on Convolutional Neural Networks (ConvNets) – to a sample of 88,327 LRGs, selected via a colour-magnitude cut, from 904 one-square-degree tiles of KiDS data. We visually inspect the candidates selected by these ConvNets and conservatively select 1983 rank-ordered candidates, which we designate the *LinKS sample* (see Sect. 4.4). We further subset the data into subsamples of 219 more plausible candidates, and 89 highly likely candidates.

We did not attempt to achieve a high level of statistical completeness in the samples of LRGs, nor in the samples of resulting lens candidates. We aimed instead to both maximise the number of lens candidates while minimising the fraction of false positives. Our colour-magnitude selection (Sect. 4.2.2) aimed at choosing a large sample of massive (early-type) galaxies while specifically avoiding star-forming (e.g. spiral) galaxies and other contaminants. We note that Vakili et al. (2019) recently selected LRGs from KiDS data using the LRG colour-magnitude relation, and also computed their photometric redshifts; we anticipate that this sample could be utilised to compile a more statistically complete sample of KiDS LRG lens candidates in the future. In addition to the *LinKS sample*, in Sect. 4.5.3 we presented a *Bonus Sample* that consists of two-hundred lens candidates. These lenses were serendipitously discovered in KiDS data, e.g., during previous experiments with various ConvNets. While these sources have not been rigorously scrutinised in the same manner as our main *LinKS sample*, and we therefore do not consider it to be as statistically well-defined as our main sample, it nonetheless contains a number of interesting strong lens candidates for future follow-up.

From our KiDS strong lens candidates (together with those found by Hartley et al. 2017 and Spiniello et al. 2018), the ~ 600 galaxy-scale lens candidates found in DES (Diehl et al., 2017; Jacobs et al., 2018; Spiniello et al., 2019) and HSC data (Sonnenfeld et al., 2018b; Wong et al., 2018), it will soon be possible to select a sample of confirmed lenses similar in size to the total number of gravitational lenses known today. For example the Masterlens database¹⁵, which assembles information on all known gravitational lenses, contains a total of ~ 600 gravitational lenses discovered up to 2016. It is possible that the total number could be, by now, up to ~ 1000 confirmed lenses and lens candidates. We believe it likely that strong lens searches within the KiDS, DES, and HSC surveys could easily double this number – accumulated over many decades – within the next few years.

¹⁵<http://masterlens.astro.utah.edu/>

Despite the already considerable numbers of new lens candidates from KiDS, there are still many lens candidates to be discovered, especially in that part of parameter space that we have not, or rather not thoroughly, explored. In addition, the completed KiDS survey will cover an area of 1350 square degrees. We plan to apply our method to these completed KiDS data, together with that of Spiniello et al. (2018), to find lensed quasars. Applying other complementary methods as Hartley et al. (2017) SVM will aid in maximizing the exploration of the parameter space.

Besides the LRG-selected sample, we have shown that is possible to tune the ConvNets to yield a sample of lens candidates with considerable purity by using many more targets (i.e. about ten times more). In particular, we ran the lens-finders on a sample composed of 930 651 galaxies (not just LRGs) and retrieved a sample of 30 strong lens candidates with an expected purity of > 90 per cent. By selecting lens candidates in this way, we are able to considerably diminish the visual inspection load, although at the price of losing many genuine lenses. With a similar setup, though, it would be feasible to retrieve ~ 3000 lens candidates from the future Euclid data set with minimal human intervention. A similar number would be found by LSST.

All these results can be enhanced further, especially by training the ConvNets with more complete training sets (Petrillo et al., 2019). In addition, a collection of genuine lens candidates, even in modest numbers, should allow one to fine-tune ConvNet lens-finders further to improve their classification capacity (Tuccillo et al., 2018; Domínguez Sánchez et al., 2018). New gravitational lenses can also be used as training sets for future crowdsourced searches (Marshall et al., 2016). Finally, the candidates identified in this paper could be used to build a benchmark against which different lens-finders can be tested and compared, similar to analyses done with simulated data (e.g., Metcalf et al. 2019)

Our results are very encouraging in light of future strong-lens surveys (e.g. those utilising Euclid and LSST) for which a naive strategy of visually inspecting galaxies to select lens candidates is entirely infeasible, given the enormous number of galaxies these new instruments will uncover. One can expect to compile samples of strong lenses from Euclid and LSST that are between one to two orders of magnitude larger than the samples compiled by any survey to date, and with minimal human effort.

Acknowledgements

The authors are grateful to the referee Alessandro Sonnenfeld for the thorough and helpful comments. CEP thanks Leon Doddema, Martin Vogelaar and Ewout Helmich for help and support. CEP, CT, GV, and LVEK are supported through an NWO-VICI grant (project number 639.043.308). CT also acknowledges funding from the INAF PRIN-SKA 2017 program 1.05.01.88.04. SC has been financially supported by a grant (project number 614.001.206) from the Netherlands Organization for Scientific Research (NWO). GVK acknowledges financial support from the Netherlands Research School for Astronomy (NOVA) and Target. Target is supported by Samenwerkingsverband Noord Nederland, European fund for regional development, Dutch Ministry of economic affairs, Pieken in de Delta, Provinces of Groningen and Drenthe. NRN acknowledges financial support from the European Union Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement N. 721463 to the SUNDIAL ITN network. This work is supported by the Deutsche Forschungsgemeinschaft in the framework of the TR33 ‘The Dark Universe’. MB is supported by the Netherlands Organization for Scientific Research, NWO, through grant number 614.001.451 and by the Polish Ministry of Science and Higher Education through grant DIR/WK/2018/12. BG acknowledges support from the European Research Council under grant number 647112. JTAdJ is supported by the Netherlands Organisation for Scientific Research (NWO) through grant 621.016.402. CH acknowledges support from the European Research Council under grant number 647112. KK acknowledges support by the Alexander von Humboldt Foundation. CS has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie actions grant agreement No 664931. This work is based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs 177.A-3016, 177.A-3017, and 177.A-3018, and on data products produced by Target/OmegaCEN, INAF-OACN, INAF-OAPD, and the KiDS production team, on behalf of the KiDS consortium. OmegaCEN and the KiDS production team acknowledge support by NOVA and NWO-M grants. Members of INAF-OAPD and INAF-OACN also acknowledge the support from the Department of Physics and Astronomy of the University of Padova, and of the Department of Physics of University of Federico II (Naples). GAMA is a joint European-Australasian project based around a spectroscopic campaign using the Anglo-Australian Telescope. The GAMA input catalogue is based on data taken from the SDSS and the UKIDSS. Complementary imaging of the

GAMA regions is being obtained by a number of independent survey programmes including GALEX MIS, VST KiDS, VISTA VIKING, WISE, Herschel-ATLAS, GMRT, and ASKAP, providing UV to radio coverage. GAMA is funded by the STFC (UK), the ARC (Australia), the AAO, and the participating institutions. The GAMA website is www.gama-survey.org. Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is www.sdss.org. SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, the Chilean Participation Group, the French Participation Group, Harvard-Smithsonian Center for Astrophysics, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatorio Nacional / MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University. The TOPCAT (Taylor, 2005) and STILTS (Taylor, 2006) software have been extensively used in this project. *Author Contributions:* All authors contributed to the development and writing of this paper. The authorship list is given in three groups: the lead authors (CEP,CT,GV,LVEK,GVK), followed by two alphabetical groups. The first alphabetical group includes those who are key contributors to both the scientific analysis and the data products. The second group covers those who have either made a significant contribution to the data products, or to the scientific analysis.

Appendix 4.A Data

4.A.1 LinKS sample

In the online material¹⁶ we provide a table of our lens candidates properties, with:

- an internal ID;
- a score from the visual inspection;
- the p -values from the ConvNets;
- the lens-candidate coordinates;
- a flag that indicates whether the candidate is already a confirmed lens or it has been identified as a candidate in other surveys.

In addition, at <http://www.astro.rug.nl/lensesinkids> we list for each one of the 1983 *LinKS* candidates:

- the internal ID;
- the visual inspection score;
- the lens-candidate coordinates;
- the RGB stamp of 101×101 pixels, corresponding to $\sim 20 \times 20$ arcseconds;
- a link to download their respective g , r and i fits files.

The candidates are ordered by decreasing visual inspection score. We also present the RGB images of the 89 “bona fide” candidates in Fig. 4.11.

4.A.2 Bonus sample

The *Bonus sample* is available via <http://www.astro.rug.nl/lensesinkids> similarly to the *LinKS sample*.

¹⁶<http://www.astro.rug.nl/lensesinkids>

Figure 4.11: Images of the 89 candidates in the LinKS with a visual inspection score greater than 27. Each image has dimensions 13×13 arcseconds.

Figure 4.11

Figure 4.11

Figure 4.11

5. CONCLUSIONS AND FUTURE PROSPECTS

I left the mellow light of Athens for the city where men wrapped and hooded in heavy togas battle against February winds, where luxury and debauch are barren of charm, but where the slightest decision taken affects the fate of some quarter of the world. There a young and eager provincial, not wholly obtuse but pursuing at first only vulgar ambitions, was little by little to lose such aspirations in the act of fulfilling them; he was to learn to contend both with men and with things, to command, and what is perhaps in the end slightly less futile, to serve.

– Marguerite Yourcenar, *Mémoires d’Hadrien*

5.1 Conclusions

This thesis addresses the problem of identifying strong gravitational lenses in astronomical wide-field imaging surveys where spectroscopic observations are partially or completely lacking. For this purpose, we have developed a pipeline for the discovery of strong gravitational lenses candidates. We have applied it to the Kilo-Degree Survey (KiDS), one of the current-generation optical wide surveys. This effort has led to the discovery of hundreds of new gravitational lens candidates. The core of the pipeline is a morphological lens-finder based on Convolutional Neural Networks (CNNs), a machine learning algorithm that “learns” from a set of examples to classify images. This kind of algorithm has proven to be one of the most promising for identifying gravitational lenses in large amounts of data (Metcalf et al., 2019). In this chapter, I summarize the main conclusions of this thesis and report an outlook for future plans and improvements.

5.1.1 The first CNN lens-finder

In Chapter 2 we have developed the first version of a CNN lens-finder. To train successfully a CNN, a homogeneous large set of labelled examples is needed. Unfortunately, such a set is still not available and, as a workaround, we have built it utilizing r -band KiDS galaxy images on which we have super-imposed simulated images of lensed galaxies. We have applied the trained CNN to a sample of r -band postage stamps of $\sim 20\,000$ color-selected luminous red galaxies (LRGs) from 255 square degrees of KiDS. The CNN has identified 761 strong gravitational lens candidates. The LRG sample included three already confirmed gravitational lenses from previous work by others. Two of these have been correctly classified by the CNN, while the third has been misclassified most likely because its lensing features are mostly unresolved. The 761 candidates have been visually inspected by 7 astronomers, whom have select a sub-sample of 56 strong lens candidates. An additional check on the Einstein radii of the candidates suggests that ~ 22 of them are solid candidates. In doing so, we have demonstrated that CNNs are suitable for identifying gravitational lens candidates in astronomical image surveys, even in the unfavourable case when a training set is not fully available. Our approach of creating a training set that is a mixture of observational and simulated data has allowed us to train successfully the CNN for identifying reliable lens candidates. However, a human intervention is still needed in validating the candidates. We have tried to mitigate this issue in the following chapters.

5.1.2 Improving the CNN lens-finder

The encouraging results achieved in Chapter 2 motivated the work discussed in Chapter 3 for further developing the method in view of applying it to the upcoming KiDS data release 3 which, at that time, was in production. In this chapter an improved version of the lens-finder has been developed. It has been realized by using a state of the art version of the algorithm and by training it with more realistic lens simulations which included realistic colours and source sub-structures. Additionally, we have developed a version of the CNN which analyzes 3-band gri -composed images. We have applied the two new CNNs to the same sample of the previous chapter with the purpose of evaluating the performance with respect to the previous CNN. We have found, beside the fact that the new CNNs perform better, that they are capable to identify lensed point-sources as doubles and quads unlike the previous realization. Moreover, the tests

carried out in this chapter have indicated that a realistic and representative training set is the most important factor for achieving high accuracy in the network classification.

5.1.3 LinKS: Lenses in the Kilo-Degree Survey

In Chapter 4, we have applied the CNNs developed in Chapter 3 to the latest KiDS data release which have been produced from observations of ~ 900 square degree of the Sky. Similarly to Chapter 2, we have applied the CNNs on ~ 88000 LRGs selected via a color-magnitude cut. This has yielded a list of 3500 strong-lens candidates which have been further down-selected via human inspection. The resulting sub-sample composed of 1983 rank-ordered targets has been named “Lenses in the Kilo-Degree Survey” (LinKS). Moreover, we have collected further 200 strong lens candidates discovered serendipitously and during experimentation with CNNs. The strong lens candidates found in this thesis will contribute in the next few years to form a sample of confirmed lenses similar in size to the ~ 1000 gravitational lenses discovered in the previous decades. In addition, we have applied the CNNs with a higher threshold of detection to a larger sample of galaxies ($\sim 10^6$), which have yielded a sample of 30 strong lens candidates with an expected purity of > 90 per cent. Despite losing many candidates, the time spent for visual inspection is considerably reduced. Applying such a strategy to the upcoming surveys carried out by Euclid and LSST, would yield, in the most pessimistic scenario, a > 90 per cent pure sample of ~ 6000 lens candidates. In this way it would be possible to trivially retrieve samples of strong lenses that are already between one to two orders of magnitude larger than the samples collected by any survey to date. These results are extremely encouraging, since the simple strategy of visually inspecting galaxies to select lens candidates will be entirely unfeasible, given the enormous number of galaxies these new instruments will observe.

At the time of writing these conclusions CNNs have proven to be one the most successful and promising methods for identifying gravitational lens candidates in astronomical image surveys (Petrillo et al., 2017; Jacobs et al., 2017; Metcalf et al., 2019; Jacobs et al., 2018; Petrillo et al., 2019).

5.2 Future Plans

The completed KiDS survey will cover a total area of 1350 square degrees. We plan to apply our method to the completed KiDS survey. There are still

many candidates to be found in KiDS, especially in the part of the parameter space that we have not explored, i.e. higher redshift and smaller separation lenses. As we have shown in Chapter 3, we can improve the ConvNets performance by using larger and more complete training sets. However, the field of deep learning applied to pattern recognition is quickly developing and a new algorithm could soon make obsolete the above issues. Fine-tuning the ConvNets on a small number of genuine lens candidates could also allow to improve their classification capacity (Tuccillo et al., 2018; Domínguez Sánchez et al., 2018). Moreover, the automated identification of the candidates via deep learning could be coupled to the automated modelling of the lens systems and the automated reconstruction of the lensed sources as experimented by Hezaveh et al. (2017) and Morningstar et al. (2019), respectively. Especially for the future wide surveys where the image resolution should be good enough to obtain satisfactory results.

The natural step forward, after finding lens candidates, is to confirm their nature through follow-ups. To this purpose, we have started a project with the SALT telescope in South Africa, named “Gotta catch’Em All” (P.I. L. Marchetti) which mainly aims at providing spectroscopic follow-up for the strong gravitational lens candidates discovered in KiDS and presented in this thesis, together with those identified in Spiniello et al. (2018). Since November 2018, we have already obtained spectra for fifteen arcs and lensed quasars and we will have observational time for at least one more semester. The spectroscopic validation of the lens candidates will allow us, in addition to confirming their lensing nature, to measure the lens and source redshifts. In doing so, we will collect a sample of confirmed lenses that can be observed with higher-resolution imaging, from the space using HST or from the ground using adaptive optics. In turn, these further follow-ups will allow us to obtain the necessary spatial resolution for modeling the lensing events and derive mass models to finally quantify the mass substructures of the confirmed lenses and to study the properties of the lensed sources.

6. BIBLIOGRAPHY

- Abadi M., et al., 2016, preprint, (arXiv:1603.04467)
- Abolfathi B., et al., 2018, ApJS, 235, 42
- Agnello A., Kelly B. C., Treu T., Marshall P. J., 2015, MNRAS, 448, 1446
- Alard C., 2006, preprint, (arXiv:astro-ph/0606757)
- Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, MNRAS, 310, 540
- Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V. E., Marshall P. J., Bundy K., Moustakas L. A., 2009, ApJ, 705, 1099
- Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V. E., Marshall P. J., Moustakas L. A., Burles S., 2010, ApJ, 724, 511
- Baldry I. K., et al., 2018, MNRAS, 474, 3875
- Barnabè M., Czoske O., Koopmans L. V. E., Treu T., Bolton A. S., Gavazzi R., 2009, MNRAS, 399, 21
- Barnabè M., Czoske O., Koopmans L. V. E., Treu T., Bolton A. S., 2011, MNRAS, 415, 2215
- Barnabè M., Spiniello C., Koopmans L. V. E., Trager S. C., Czoske O., Treu T., 2013, MNRAS, 436, 253
- Benítez N., 2000, ApJ, 536, 571
- Bertin E., 2012, in Ballester P., Egret D., Lorente N. P. F., eds, Astronomical Society of the Pacific Conference Series Vol. 461, Astronomical Data Analysis Software and Systems XXI. p. 263
- Bertin E., Arnouts S., 1996, A&AS, 117, 393

BIBLIOGRAPHY

- Bettinelli M., Simioni M., Aparicio A., Hidalgo S. L., Cassisi S., Walker A. R., Piotto G., Valdes F., 2016, *MNRAS*, 461, L67
- Biesiada M., Piórkowska A., Malec B., 2010, *MNRAS*, 406, 1055
- Binney J., Merrifield M., 1998, *Galactic Astronomy*. Princeton University Press
- Blake C., et al., 2016, *MNRAS*, 462, 4240
- Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Moustakas L. A., 2006, *ApJ*, 638, 703
- Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Gavazzi R., Moustakas L. A., Wayth R., Schlegel D. J., 2008, *ApJ*, 682, 964
- Bom C. R., Makler M., Albuquerque M. P., Brandt C. H., 2017, *A&A*, 597, A135
- Bonvin V., et al., 2017, *MNRAS*, 465, 4914
- Boureau Y.-L., Bach F., LeCun Y., Ponce J., 2010, in *Computer Vision and Pattern Recognition (CVPR)*, 2010 IEEE Conference on. pp 2559–2566
- Bournaud F., Combes F., 2003, *A&A*, 401, 817
- Brammer G. B., et al., 2012, *ApJ*, 758, L17
- Brault F., Gavazzi R., 2015, *A&A*, 577, A85
- Brewer B. J., et al., 2012, *MNRAS*, 422, 3574
- Browne I. W. A., et al., 2003, *MNRAS*, 341, 13
- Brownstein J. R., et al., 2012, *ApJ*, 744, 41
- Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000
- Cabanac R. A., et al., 2007, *A&A*, 461, 813
- Cabrera-Vives G., Reyes I., Förster F., Estévez P. A., Maureira J.-C., 2017, *ApJ*, 836, 97
- Cao S., Covone G., Zhu Z.-H., 2012, *ApJ*, 755, 31
- Cao S., Biesiada M., Gavazzi R., Piórkowska A., Zhu Z.-H., 2015, *ApJ*, 806, 185

-
- Cao S., Biesiada M., Yao M., Zhu Z.-H., 2016, MNRAS, 461, 2192
- Capaccioli M., Schipani P., 2011, The Messenger, 146, 2
- Cardone V. F., Tortora C., 2010, MNRAS, 409, 1570
- Cardone V. F., Tortora C., Molinaro R., Salzano V., 2009, A&A, 504, 769
- Carlstrom J. E., et al., 2011, PASP, 123, 568
- Chabrier G., 2001, ApJ, 554, 1274
- Chae K.-H., 2003, MNRAS, 346, 746
- Chan J. H. H., Suyu S. H., Chiueh T., More A., Marshall P. J., Coupon J., Oguri M., Price P., 2015, ApJ, 807, 138
- Chatterjee S., Koopmans L. V. E., 2018, MNRAS, 474, 1762
- Chevance M., Weijmans A.-M., Damjanov I., Abraham R. G., Simard L., van den Bergh S., Caris E., Glazebrook K., 2012, ApJ, 754, L24
- Chollet F., et al., 2015, Keras, <https://keras.io>
- Christensen L., D’Odorico S., Pettini M., Belokurov V., Evans N. W., Kellogg M., Vernet J., 2010, MNRAS, 406, 2616
- Colless M., et al., 2001, MNRAS, 328, 1039
- Collett T. E., 2015, ApJ, 811, 20
- Collett T. E., Auger M. W., 2014, MNRAS, 443, 969
- Cooray A., et al., 2011, preprint, (arXiv:1110.3784)
- Covone G., et al., 2009, ApJ, 691, 531
- Csurka G., 2017, arXiv preprint arXiv:1702.05374
- D’Isanto A., Polsterer K. L., 2018, A&A, 609, A111
- Dawson K. S., et al., 2013, AJ, 145, 10
- Deane R. P., Rawlings S., Garrett M. A., Heywood I., Jarvis M. J., Klöckner H.-R., Marshall P. J., McKean J. P., 2013, MNRAS, 434, 3322
- Diehl H. T., et al., 2017, ApJS, 232, 15
- Dieleman S., Willett K. W., Dambre J., 2015, MNRAS, 450, 1441

BIBLIOGRAPHY

- Domínguez Sánchez H., et al., 2018, preprint, (arXiv:1807.00807)
- Driver S. P., et al., 2011, Monthly Notices of the Royal Astronomical Society, 413, 971
- Dyson F. W., Eddington A. S., Davidson C., 1920, Philosophical Transactions of the Royal Society of London Series A, 220, 291
- Einstein A., 1915, Sitzungsber. preuss.Akad. Wiss., vol. 47, No.2, pp. 831-839, 1915, 47, 831
- Eisenstein D. J., et al., 2001, AJ, 122, 2267
- Eisenstein D. J., et al., 2011, AJ, 142, 72
- Estrada J., et al., 2007, ApJ, 660, 1176
- Faure C., et al., 2008, ApJS, 176, 19
- Ferreras I., Saha P., Leier D., Courbin F., Falco E. E., 2010, MNRAS, 409, L30
- Föëx G., Motta V., Limousin M., Verdugo T., More A., Cabanac R., Gavazzi R., Muñoz R. P., 2013, A&A, 559, A105
- Fukugita M., Futamase T., Kasai M., Turner E., 1992, ApJ, 393, 3
- Fukushima K., 1980, Biological cybernetics, 36, 193
- Gavazzi R., Treu T., Rhodes J. D., Koopmans L. V. E., Bolton A. S., Burles S., Massey R. J., Moustakas L. A., 2007, ApJ, 667, 176
- Gavazzi R., Treu T., Koopmans L. V. E., Bolton A. S., Moustakas L. A., Burles S., Marshall P. J., 2008, ApJ, 677, 1046
- Gavazzi R., Marshall P. J., Treu T., Sonnenfeld A., 2014, ApJ, 785, 144
- Grillo C., Eichner T., Seitz S., Bender R., Lombardi M., Gobat R., Bauer A., 2010, ApJ, 710, 372
- Grogin N. A., et al., 2011, ApJS, 197, 35
- Guo Y., Liu Y., Oerlemans A., Lao S., Wu S., Lew M. S., 2016, Neurocomputing, 187, 27
- Hála P., 2014, preprint, (arXiv:1412.8341)

-
- Hartley P., Flamary R., Jackson N., Tagore A. S., Metcalf R. B., 2017, MNRAS, 471, 3378
- He K., Zhang X., Ren S., Sun J., 2015a, in Proceedings of the IEEE International Conference on Computer Vision. pp 1026–1034
- He K., Zhang X., Ren S., Sun J., 2015b, CoRR, abs/1512.03385
- Hezaveh Y., Dalal N., Holder G., Kisner T., Kuhlen M., Perreault Levasseur L., 2014, preprint, (arXiv:1403.2720)
- Hezaveh Y. D., et al., 2016, ApJ, 823, 37
- Hezaveh Y. D., Levasseur L. P., Marshall P. J., 2017, Nature, 548, 555
- Hildebrandt H., et al., 2017, MNRAS, 465, 1454
- Hinton G. E., Srivastava N., Krizhevsky A., Sutskever I., Salakhutdinov R. R., 2012, preprint, (arXiv:1207.0580)
- Hoag A. A., 1950, AJ, 55, 170
- Horesh A., Ofek E. O., Maoz D., Bartelmann M., Meneghetti M., Rix H.-W., 2005, ApJ, 633, 768
- Hoyle B., 2016, Astronomy and Computing, 16, 34
- Huertas-Company M., et al., 2015, ApJS, 221, 8
- Ilbert O., et al., 2006, A&A, 457, 841
- Ilbert O., et al., 2009, ApJ, 690, 1236
- Impellizzeri C. M. V., McKean J. P., Castangia P., Roy A. L., Henkel C., Brunthaler A., Wucknitz O., 2008, Nature, 456, 927
- Inada N., et al., 2012, AJ, 143, 119
- Iodice E., Arnaboldi M., Bournaud F., Combes F., Sparke L. S., van Driel W., Capaccioli M., 2003, ApJ, 585, 730
- Ioffe S., Szegedy C., 2015, preprint, (arXiv:1502.03167)
- Jackson N., 2008, MNRAS, 389, 1311
- Jacobs C., Glazebrook K., Collett T., More A., McCarthy C., 2017, MNRAS, 471, 167

- Jacobs C., et al., 2018, preprint, (arXiv:1811.03786)
- Jee I., Komatsu E., Suyu S. H., Huterer D., 2016, *J. Cosmology Astropart. Phys.*, 4, 031
- Jiang G., Kochanek C. S., 2007, *ApJ*, 671, 1568
- Joseph R., et al., 2014, *A&A*, 566, A63
- Kelly P. L., et al., 2018, *Nature Astronomy*, 2, 334
- Kim E. J., Brunner R. J., 2017, *MNRAS*, 464, 4463
- Kingma D., Ba J., 2014, preprint, (arXiv:1412.6980)
- Kochanek C., 1996, *ApJ*, 473, 595
- Koopmans L. V. E., Treu T., 2003, *ApJ*, 583, 606
- Koopmans L. V. E., Treu T., Bolton A. S., Burles S., Moustakas L. A., 2006, *ApJ*, 649, 599
- Koopmans L. V. E., et al., 2009, *ApJ*, 703, L51
- Kormann R., Schneider P., Bartelmann M., 1994, *A&A*, 284, 285
- Krizhevsky A., Sutskever I., Hinton G. E., 2012, in *Advances in neural information processing systems*. pp 1097–1105
- Kubo J. M., Dell’Antonio I. P., 2008, *MNRAS*, 385, 918
- Kuijken K., 2011, *The Messenger*, 146, 8
- Kuijken K., et al., 2015, *MNRAS*, 454, 3500
- Kuijken K., et al., 2019, *A&A*, 625, A2
- LSST Science Collaboration et al., 2009, preprint, (arXiv:0912.0201)
- 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
- Lanusse F., Ma Q., Li N., Collett T. E., Li C.-L., Ravanbakhsh S., Mandelbaum R., Póczos B., 2018, *MNRAS*, 473, 3895
- Laureijs R., et al., 2011, arXiv e-prints, p. arXiv:1110.3193

-
- LeCun Y., Bottou L., Bengio Y., Haffner P., 1998, *Proceedings of the IEEE*, 86, 2278
- LeCun Y., Bengio Y., Hinton G., 2015, *Nature*, 521, 436
- Leier D., Ferreras I., Saha P., Charlot S., Bruzual G., La Barbera F., 2016, *MNRAS*, 459, 3677
- Lenzen F., Schindler S., Scherzer O., 2004, *A&A*, 416, 391
- Li R., Frenk C. S., Cole S., Gao L., Bose S., Hellwing W. A., 2016, *MNRAS*, 460, 363
- Li R., Shu Y., Wang J., 2018, *MNRAS*, 480, 431
- Limousin M., et al., 2010, *A&A*, 524, A95
- Liske J., et al., 2015, *MNRAS*, 452, 2087
- Lupton R., Blanton M. R., Fekete G., Hogg D. W., O’Mullane W., Szalay A., Wherry N., 2004, *PASP*, 116, 133
- Madore B. F., Nelson E., Petrillo K., 2009, *ApJS*, 181, 572
- Marshall P. J., et al., 2016, *MNRAS*, 455, 1171
- Mason C. A., et al., 2017, *ApJ*, 838, 14
- Maturi M., Mizera S., Seidel G., 2014, *A&A*, 567, A111
- McFarland J. P., Verdoes-Kleijn G., Sikkema G., Helmich E. M., Boxhoorn D. R., Valentijn E. A., 2013, *Experimental Astronomy*, 35, 45
- McKean J., et al., 2015, *Proceedings of Science*, p. 84
- Metcalf R. B., et al., 2019, *A&A*, 625, A119
- Miyazaki S., et al., 2012, in *SPIE Astronomical Telescopes+ Instrumentation*. pp 84460Z–84460Z
- Möller O., Kitzbichler M., Natarajan P., 2007, *MNRAS*, 379, 1195
- More A., McKean J. P., Muxlow T. W. B., Porcas R. W., Fassnacht C. D., Koopmans L. V. E., 2008, *MNRAS*, 384, 1701
- More A., McKean J. P., More S., Porcas R. W., Koopmans L. V. E., Garrett M. A., 2009, *MNRAS*, 394, 174

BIBLIOGRAPHY

- More A., Jahnke K., More S., Gallazzi A., Bell E. F., Barden M., Häußler B., 2011, *ApJ*, 734, 69
- More A., Cabanac R., More S., Alard C., Limousin M., Kneib J.-P., Gavazzi R., Motta V., 2012, *ApJ*, 749, 38
- More A., et al., 2016, *MNRAS*, 455, 1191
- More A., et al., 2017, *MNRAS*, 465, 2411
- Morningstar W. R., et al., 2019, arXiv e-prints, p. arXiv:1901.01359
- Mukherjee S., et al., 2018, *MNRAS*, 479, 4108
- Nair V., Hinton G. E., 2010, in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. pp 807–814
- Napolitano N. R., et al., 2016, *The Universe of Digital Sky Surveys*, 42, 129
- Negrello M., et al., 2010, *Science*, 330, 800
- Negrello M., et al., 2017, *Monthly Notices of the Royal Astronomical Society*, 465, 3558
- Newton I., 1726, *Philosophiae naturalis principia mathematica*. Vol. 1
- Ng A. Y., 2004, in *Proceedings of the twenty-first international conference on Machine learning*. p. 78
- Nierenberg A. M., Treu T., Wright S. A., Fassnacht C. D., Auger M. W., 2014, *MNRAS*, 442, 2434
- Oguri M., 2006, *MNRAS*, 367, 1241
- Oguri M., Marshall P. J., 2010, *MNRAS*, 405, 2579
- Paraficz D., et al., 2016, *A&A*, 592, A75
- Pawase R. S., Courbin F., Faure C., Kokotanekova R., Meylan G., 2014, *MNRAS*, 439, 3392
- Peng C. Y., Impey C. D., Rix H.-W., Kochanek C. S., Keeton C. R., Falco E. E., Lehár J., McLeod B. A., 2006, *The Astrophysical Journal*, 649, 616
- Petrillo C. E., et al., 2017, *MNRAS*, 472, 1129
- Petrillo C. E., et al., 2019, *MNRAS*, 482, 807

-
- Posacki S., Cappellari M., Treu T., Pellegrini S., Ciotti L., 2015, MNRAS, 446, 493
- Pourrahmani M., Nayyeri H., Cooray A., 2018, ApJ, 856, 68
- Richard J., Kneib J.-P., Ebeling H., Stark D. P., Egami E., Fiedler A. K., 2011, MNRAS, 414, L31
- Ruff A. J., Gavazzi R., Marshall P. J., Treu T., Auger M. W., Brault F., 2011, ApJ, 727, 96
- Rumelhart D. E., Hinton G. E., Williams R. J., 1986, Nature, 323, 533
- Russakovsky O., et al., 2015, International Journal of Computer Vision (IJCV), 115, 211
- Saerens M., Latinne P., Decaestecker C., 2002, IEEE TRANSACTIONS ON NEURAL NETWORKS, 13, 1205
- Salmon B., et al., 2017, preprint, (arXiv:1710.08930)
- Salpeter E. E., 1955, ApJ, 121, 161
- Schaefer C., Geiger M., Kuntzer T., Kneib J.-P., 2018, A&A, 611, A2
- Schechter P. L., et al., 1997, ApJ, 475, L85
- Schlafly E. F., Finkbeiner D. P., 2011, ApJ, 737, 103
- Schlegel D. J., Finkbeiner D. P., Davis M., 1998, ApJ, 500, 525
- Schmidhuber J., 2015, Neural Networks, 61, 85
- Schneider P., 2006, in Meylan G., Jetzer P., North P., Schneider P., Kochanek C. S., Wambsganss J., eds, Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro. pp 1–89
- Schneider P., Ehlers J., Falco E., 1992, Gravitational Lenses. Springer-Verlag Berlin Heidelberg New York. Also Astronomy and Astrophysics Library
- Seidel G., Bartelmann M., 2007, A&A, 472, 341
- Sergeyev A., Spiniello C., Khramtsov V., Napolitano N. R., Bannikova E., Tortora C., Getman F. I., Agnello A., 2018, Research Notes of the American Astronomical Society, 2, 189
- Sérsic J. L., 1968, Atlas de galaxies australes

- Shajib A. J., Treu T., Agnello A., 2018, MNRAS, 473, 210
- Shallue C. J., Vanderburg A., 2018, AJ, 155, 94
- Shu Y., et al., 2015, ApJ, 803, 71
- Shu Y., et al., 2016, ApJ, 833, 264
- Shu Y., et al., 2017, ApJ, 851, 48
- Silver D., et al., 2017, Nature, 550, 354
- Simard P. Y., Steinkraus D., Platt J. C., 2003, in Proceedings of the Seventh International Conference on Document Analysis and Recognition - Volume 2. ICDAR '03. IEEE Computer Society, Washington, DC, USA, pp 958–, <http://dl.acm.org/citation.cfm?id=938980.939477>
- Sonnenfeld A., Gavazzi R., Suyu S. H., Treu T., Marshall P. J., 2013a, ApJ, 777, 97
- Sonnenfeld A., Treu T., Gavazzi R., Suyu S. H., Marshall P. J., Auger M. W., Nipoti C., 2013b, ApJ, 777, 98
- Sonnenfeld A., Treu T., Marshall P. J., Suyu S. H., Gavazzi R., Auger M. W., Nipoti C., 2015, ApJ, 800, 94
- Sonnenfeld A., Leauthaud A., Auger M. W., Gavazzi R., Treu T., More S., Komiyama Y., 2018a, preprint, (arXiv:1801.01883)
- Sonnenfeld A., et al., 2018b, PASJ, 70, S29
- Spingola C., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., Lagattuta D. J., Vegetti S., 2018, MNRAS, 478, 4816
- Spiniello C., Koopmans L., Trager S. C., Czoske O., Treu T., 2011, MNRAS, 417, 3000
- Spiniello C., Barnabè M., Koopmans L. V. E., Trager S. C., 2015, MNRAS, 452, L21
- Spiniello C., et al., 2018, MNRAS, 480, 1163
- Spiniello C., Agnello A., Sergeyev A. V., Anguita T., Rodríguez Ó., Napolitano N. R., Tortora C., 2019, MNRAS, 483, 3888
- Stapelberg S., Carrasco M., Maturi M., 2019a, MNRAS, 482, 1824

-
- Stapelberg S., Carrasco M., Maturi M., 2019b, MNRAS, 482, 1824
- Stark D. P., et al., 2013, MNRAS, 436, 1040
- Suyu S. H., Marshall P. J., Auger M. W., Hilbert S., Blandford R. D., Koopmans L. V. E., Fassnacht C. D., Treu T., 2010, ApJ, 711, 201
- Suyu S. H., et al., 2013, ApJ, 766, 70
- Suyu S. H., et al., 2017, MNRAS, 468, 2590
- Swinbank A. M., et al., 2009, MNRAS, 400, 1121
- Tanaka M., et al., 2016, ApJ, 826, L19
- Tao Y., Zhang Y., Cui C., Zhang G., 2018, preprint, (arXiv:1801.04839)
- Taylor M. B., 2005, in Shopbell P., Britton M., Ebert R., eds, Astronomical Society of the Pacific Conference Series Vol. 347, Astronomical Data Analysis Software and Systems XIV. p. 29
- Taylor M. B., 2006, in Gabriel C., Arviset C., Ponz D., Enrique S., eds, Astronomical Society of the Pacific Conference Series Vol. 351, Astronomical Data Analysis Software and Systems XV. p. 666
- The Dark Energy Survey Collaboration 2005, preprint, (arXiv:astro-ph/0510346)
- Theano Development Team 2016, arXiv e-prints, abs/1605.02688
- Theys J. C., Spiegel E. A., 1976, ApJ, 208, 650
- Tortora C., Napolitano N. R., Romanowsky A. J., Capaccioli M., Covone G., 2009, MNRAS, 396, 1132
- Tortora C., Napolitano N., Romanowsky A., Jetzer P., 2010, ApJ, 721, L1
- Treu T., 2010, ARA&A, 48, 87
- Treu T., Koopmans L. V. E., 2002a, MNRAS, 337, L6
- Treu T., Koopmans L. V. E., 2002b, ApJ, 575, 87
- Treu T., Auger M. W., Koopmans L. V. E., Gavazzi R., Marshall P. J., Bolton A. S., 2010, ApJ, 709, 1195
- Treu T., Dutton A. A., Auger M. W., Marshall P. J., Bolton A. S., Brewer B. J., Koo D. C., Koopmans L. V. E., 2011, MNRAS, 417, 1601

BIBLIOGRAPHY

- Treu T., et al., 2015, *ApJ*, 812, 114
- Trujillo I., Conselice C. J., Bundy K., Cooper M. C., Eisenhardt P., Ellis R. S., 2007, *MNRAS*, 382, 109
- Tu H., et al., 2009, *A&A*, 501, 475
- Tuccillo D., Huertas-Company M., Decencire E., Velasco-Forero S., Domnguez Snchez H., Dimauro P., 2018, *MNRAS*, 475, 894
- Turner E. L., Ostriker J. P., Gott III J. R., 1984, *ApJ*, 284, 1
- Vakili M., et al., 2019, *MNRAS*, 487, 3715
- Valentijn E. A., et al., 2007, in Shaw R. A., Hill F., Bell D. J., eds, *Astronomical Society of the Pacific Conference Series Vol. 376, Astronomical Data Analysis Software and Systems XVI*. p. 491 (arXiv:astro-ph/0702189)
- Van der Walt S., Schnberger J. L., Nunez-Iglesias J., Boulogne F., Warner J. D., Yager N., Gouillart E., Yu T., 2014, *PeerJ*, 2, e453
- Vegetti S., Koopmans L. V. E., 2009, *MNRAS*, 400, 1583
- Vegetti S., Lagattuta D. J., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., 2012, *Nature*, 481, 341
- Verdugo T., et al., 2014, *A&A*, 571, A65
- Vernardos G., 2018, *MNRAS*, p. 3317
- Vieira J., et al., 2013, *Nature*, 495, 344
- Whitmore B. C., Lucas R. A., McElroy D. B., Steiman-Cameron T. Y., Sackett P. D., Olling R. P., 1990, *AJ*, 100, 1489
- Willett K. W., et al., 2013, *MNRAS*, 435, 2835
- Wong K. C., et al., 2018, *ApJ*, 867, 107
- Xu B., et al., 2016, *ApJ*, 817, 85
- Yuan H. B., Liu X. W., Xiang M. S., 2013, *MNRAS*, 430, 2188
- Zeiler M. D., Fergus R., 2013, preprint, (arXiv:1311.2901)
- de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, *Experimental Astronomy*, 35, 25

de Jong J. T. A., et al., 2015, A&A, 582, A62

de Jong J. T. A., et al., 2017, A&A, 604, A134

de Vaucouleurs G., 1948, Annales d'Astrophysique, 11, 247

A. Guide to the script

Talk is cheap. Show me the code.

– Linus Torvalds

In this Appendix I provide a guide for using the scripts to reproduce the 1-band and 3-bands ConvNets presented in Chapter 3. The files can be found as a git repository at

https://github.com/CEnricoP/cnn_strong_lensing.

A.1 Dependencies

The script has been written in Python 3 using the distribution Anaconda 3. The following libraries are needed to run the script: TensorFlow 1.3.0, keras 2.0.8, skimage 0.12.3, astropy 1.3 and PIL 4.0.0.

A.2 Files

The git repository contains the following files:

- `cnn.py` the main file;
- `augmentation.py` implements the data augmentation;
- `create_training_dic.py` together with
- `create_train_ids.py` is used to create a labelled list of the file to use for the training phase;
- `load_data.py` implements the loading of the input data to the CNN;
- `resnet.py` contains the network architecture;
- the folder `HumVI_online_lensing` contains a library that is used to produce RGB images and is a slightly modified version of the HumVI library (<https://github.com/drphilmarshall/HumVI>);
- The folder `data` which contains the average PSFs for the g, r and i KiDS bands and a table with a list of tuples of simulated g, r and i magnitudes for producing 3-band images.

A.3 Data preparation

Create N subfolders in `data/training/sources/` with names ranging from one to the total number of simulated files. In each subfolder put a simulated source named accordingly, e.g. in folder 1 there will be `1.fits`.

Place in `data/training/lenses/` all the training lens examples. The files have to be named `xxx_g_xxx.fits`, `xxx_r_xxx.fits`, `xxx_i_xxx.fits`, one for each of the three different KiDS bands.

Place the negative examples, in the same manner as above, in `data/training/negatives/`.

Run the script `create_training_dic.py` and, subsequently, `create_training_ids.py`.

Put the test data in `data/test_data/` in the same manner as above.

A.4 Parameters

At the beginning of the file `cnn.py` there is a list of the main parameters of the script with their default values. A description of the parameters follows:

- `nbands` either 1 or 3 to choose between the 1-band or 3-bands ConvNet;
- `input_sizes` the size of the input images;
- `batch_size` number of inputs after which there is an update of the weights of the network;
- `chunk_size` the number of images loaded in one chunk (the final number will be twice this numbers, because one chunk is loaded for each the negative and positive examples);
- `num_chunks` the total number of chunks to load;
- `normalize` if True normalize the images between 0 and 1 (for single-band only);
- `model_name` the name of the model that will be saved at training time or the name of the model to load at test time;
- `learning_rate` the learning rate, i.e., the magnitude of weight updates at training time;
- `range_min` At training time gives the minimum value of the ratio between the maximum brightness of the simulated source and the lens galaxy;
- `range_max` At training time gives the maximum value of the ratio between the maximum brightness of the simulated source and the lens galaxy.

- `augm_pred` if True, at testing time the cnn gives a prediction obtained by averaging of the p 's for the original image and the images obtained operating a rotation of 90, 180 and 270 degrees;

A.5 Running the script

From the terminal launch the command `python cnn.py resnet train` to train a CNN. When the training is complete a file with the weights of networks named `[model_name].weights.h5` is created. The same file is created if the training is interrupted by the user for some reason. To run the trained CNN run the command `python cnn.py resnet predict`. The results of the test will be stored in a file named `'pred_[model_name].pkl'` as a list of the file names with the relative predictions.

SUMMARY

Introduction

Dark matter, galaxies and the Universe

According to our current cosmological model, the Λ CDM model, we live in a Universe that is perpetually expanding. The ingredients of this Universe are light, ordinary matter, the mysterious dark matter and the even more mysterious dark energy which is the responsible for the expansion. The universe began its expansion 13.8 billion years ago with the so-called Big Bang. It is thought that the early Universe was very homogenous and its expansion was powered by matter that was extremely dense and hot. As the Universe kept growing, it began to gradually cool and allowed gravity to play a more prominent role in forming the first dark matter structures. Dark matter was the first to form structures because it does not interact with ordinary matter and light but it just “feels” the gravitational force. Next, when the universe was between 150 million and 1 billion years old, ordinary matter started to collapse within dark matter over-densities. The first stars and galaxies were born and they began to produce metals and dust which were bound together with gravity. Dark matter constitutes the dominant part of galaxies, for example about 95% of the mass of our own galaxy is estimated to be dark matter. Dark matter is extremely difficult to study because it can not be “seen” since it does not emit any light and does not interact directly with ordinary matter as mentioned above.

How galaxies are formed and evolve in time, what is the amount of their constituents and at which rate the Universe expands are still mostly open fundamental questions. Research studies on these topics are largely driven by the analysis of astronomical observations. Fortunately, the more distant galaxies we observe in the sky the more in the past we observe. In fact, light does not transmit instantaneously but it travels at a certain speed. This means that the galaxies which we observe in the sky are pictures of what the galaxies were when the light started its journey towards us. This fact allows astronomers to have “pictures” of the Universe at different instances in time and opens windows on its history. Moreover, a phenomenon called

gravitational lensing further facilitates cosmological and galaxy formation and evolution studies and it is an essential tool for shedding light on the above-mentioned fundamental questions.

Gravitational Lensing

One of the most surprising phenomena predicted by Einstein's theory of General Relativity is that light-rays are bent by objects lying in the neighborhood on their trajectories. In particular, when one is lucky enough to have two galaxies aligned along the line of sight, light emitted from the furthest galaxies is amplified and bent in such a way that we can observe multiple distorted images of the more distant galaxy (see figures A.1 and A.2). These rare astronomical objects are called gravitational lenses because the galaxy nearest to the observer acts as a *lens* which deviates and amplifies the light-rays coming from the furthest galaxy which is called *source*.

It turns out that gravitational lensing is a unique tool for answering the fundamental questions described in the previous section. In fact, gravitational lensing allows us to observe distant galaxies that would not be observable otherwise due to its amplifying effect and this allows us to study old and faint galaxies. Moreover, the shape of the observed images depend almost entirely on the mass and the relative positions of the observer, the source and the lensed galaxies. This allows to precisely measure the mass of the lens galaxy, which is in general constituted not only by ordinary matter but mostly by dark matter. That is, gravitational lensing is a unique tool for measuring the dark matter content of galaxies. Last but not least, another important scientific application of gravitational lensing is the measure of the velocity of expansion of the Universe, the so-called Hubble constant. It can be measured if the source galaxy is a quasar, that is a very distant with a strongly variable light. In such a case an observer will see the intensity of the light of different multiple images varying at different times. Astronomers monitor these rare objects because the amount of time between the variation of the images depends on the Hubble constant which can be estimated from this method. Therefore, it is necessary to find and study these precious and rare astronomical objects to expand our knowledge about galaxies and the structure of the Universe.

However, gravitational lenses are rare, about one in a thousand galaxies, and laborious to identify by eye among millions of galaxies. Forthcoming astronomical imaging surveys will produce a large amount of data containing billions galaxies and hundred of thousands of gravitational lenses. Visualizing inspecting all these galaxies would require to create

Figure A.1: The light coming from a *source* galaxy is bent by a *lens* galaxy before arriving at Earth. Multiple distorted images of the same *source* galaxy are created and are observable from Earth. See two examples in the next figure. (Credit: NASA, ESA & L. Calçada)

Figure A.2: Two examples of gravitational lenses. On the top, the central yellow galaxy bends the light coming from a distant galaxy creating a blue arc. In the second example the light coming from a quasar is divided in four different images (three above and one below) by the galaxy at the centre of the image.

special full-time jobs dedicated to the effort of identifying gravitational lenses! Therefore, it is worth developing automatic identification methods for aiding the discovery of gravitational lenses. This is the main goal of my thesis research and to address it I have developed during my studies a pipeline based on neural networks for identifying gravitational lenses.

Neural Networks

Recently, a particular kind of algorithm has brought a revolution in the artificial intelligence field. They are inspired by the structure and functions of neurons and are called *artificial neural networks*, or neural networks in short, for such a reason. Neural networks learn by *observing* examples and they are particularly effective if the quantity of available examples on which they can be *trained* is large enough. Neural networks have out-performed humans in various applications. For example, the renowned neural network algorithm AlphaGO won four of five matches against the 15 times world champion Lee Sedol. Such an event was thought to be difficult to achieve. AlphaGO developed its strategy by observing Go matches and playing against itself. Another example are self-driving cars which big companies are developing and that will have a big impact on our way of travelling and commuting. Neural networks are used in a large variety of public and commercial applications, especially in the field of image recognition. For example, Facebook and Google use this tool for recognizing subjects in pictures uploaded on online platforms.

This thesis

In this thesis I have developed a method to automatically identify gravitational lens candidates in astronomical image surveys in order to enable individual and ensemble studies of these scientifically valuable astronomical objects. To accomplish this goal, I have designed a machine learning algorithm based on neural networks. The data utilized in this work comes from the Kilo-Degree Survey (KiDS). KiDS is an astronomical survey conducted by the 2.61 metre VLT Survey Telescope (VST, see Fig. A.3) which is situated in the Acatama desert in Chile. That is an ideal place for astronomical observations because of its weather conditions that offer clear and dry skies for the vast majority of the year.

KiDS has observed millions of galaxies in the Southern sky with unprecedented quality. In my thesis I have developed neural networks to find gravitational lenses in such a large amount of data. The neural

Figure A.3: The VST telescope in the Atacama desert, Chile.

networks have analyzed images of tens of thousands galaxies in few dozens of minutes. They have led to the discovery of hundreds of gravitational lens candidates. Applying neural networks for finding gravitational lenses will become extremely important for the future of astronomical surveys. In fact, many millions of galaxies will be observed and analyzing them by eye will be unfeasible. Thus, utilising this kind of automatic analyses will become necessary.

In **Chapter 2** I design a neural network lens-finder and train it using real galaxy images from KiDS together with simulated gravitationally lensed sources. There was the need to use simulated data because there are not enough known gravitational lenses to form a large enough set of examples to train a neural network. After the training, I apply the algorithm to a sample of about 20 000 galaxies, the algorithm successfully identified two of three previously known gravitational lenses that were in the sample. Additionally, another 50 reliable gravitational lens candidates were found.

In **Chapter 3**, I implement two different lens-finders, one for classifying the images based on their shape and the other one to classify them based also on their colour. The algorithms developed here have better performance compared to the one developed in Chapter 2. This is due to

using an improved version of the algorithms and more importantly, training them with an improved set of simulated gravitationally lensed sources.

In **Chapter 4**, I present a sample of several hundreds of gravitational lens candidates: the LinKS (Lenses in KiDS) sample. This sample has been selected by applying the two neural networks developed in Chapter 3 to about 88 000 galaxies from KiDS and subsequently picking the best candidates with a visual inspection. The images of these galaxies can be visualized at the website <https://www.astro.rug.nl/lensesinkids/>. Besides, I find that it will be possible to select thousands of lens candidates in forthcoming astronomical imaging surveys by using neural networks and minimal human intervention.

SAMENVATTING

Inleiding

Universum, melkwegstelsel en donkere materie

Volgens ons huidige kosmologische model, het Λ CDM-model, leven we in een universum dat zich voortdurend uitdijt. De ingrediënten van dit Universum zijn licht, gewone materie, de mysterieuze donkere materie en de nog mysterieuzere donkere energie die verantwoordelijk is voor de uitdijng. Het begin van het heelal is vastgesteld op 13,8 miljard jaar geleden met de zogeheten Big Bang. Dat wil zeggen, het heelal begon uit te dijen vanuit een zeer hoge dichtheid en hoge-temperatuurtoestand waarin de bestanddelen homogeen verspreid waren en er nog geen structuur aanwezig was. Terwijl de uitdijng doorging, koelde het heelal af en begonnen er gebieden met iets hogere massadichtheid te ontstaan die met de tijd compacter en compacter werden onder invloed van de zwaartekracht en zodoende de eerste structuren vormden. De donkere materie begon als eerste samen te klonteren. Vervolgens begon gewone materie in die gebieden samen te klonteren en daaruit vormden zich ongeveer 150 miljoen tot 1 miljard jaar na de oerknal de eerste sterren¹ en melkwegstelsels. Melkwegstelsels zijn systemen van miljoenen of miljarden sterren, samen met gas, stof, enz. en donkere materie, bij elkaar gehouden door de zwaartekracht. De massa van melkwegstelsels wordt geschat voor het grootste deel uit donkere materie te bestaan. Onze eigen Melkweg bestaat voor een geschatte 95% uit donkere materie. Donkere materie is uiterst moeilijk te bestuderen omdat het niet "gezien" kan worden omdat het geen licht uitstraalt en geen interactie heeft met gewone materie en licht behalve via de zwaartekracht.

Er staan nog steeds fundamentele vragen open rondom de vorming en evolutie van melkwegstelsels, de relatieve verdeling van hun bestanddelen en de evolutie van de snelheid waarmee het heelal uitdijt. Onderzoek naar deze onderwerpen wordt voor een groot gedeelte gedaan op basis

¹Onze zon is een ster.

van analyse van astronomische waarnemingen. Gelukkig is het zo dat hoe dieper wij in het heelal sterrenstelsels waarnemen, hoe meer wij te weten komen over het verleden van sterrenstelsels. Dit komt doordat licht niet oneindig snel reist. Dit betekent dat de sterrenstelsels die we aan de hemel waarnemen beelden zijn van die sterrenstelsels op het moment dat zij het licht naar ons uitzonden. Dit feit stelt astronomen in staat om 'beelden' van het heelal te hebben op verschillende momenten in haar geschiedenis. De beelden van zwaartekrachtlenzen spelen in dit onderzoek een belangrijke rol. In de volgende sectie leg ik uit wat zwaartekrachtlenzen zijn.

Zwaartekrachtlenzen

Een van de meeste verbazingwekkende fenomenen voorspeld op basis van Einsteins algemene relativiteitstheorie is de mate waarin lichtstralen gebogen worden door de massa van objecten nabij hun traject. Dit effect is extra sterk wanneer men genoeg geluk heeft om twee melkwegstelsels op verschillende afstanden vrijwel op dezelfde zichtlijn te hebben. Het licht afkomstig van het verste sterrenstelsel wordt geconcentreerd en verbogen door de zwaartekracht van het sterrenstelsel dat dichterbij staat. Op een dergelijke manier ontstaan verscheidene vervormde afbeeldingen van het verste sterrenstelsel (zie afbeeldingen A.4 en A.5). Dit zeldzame astronomische fenomeen wordt een zwaartekrachtlenze genoemd, omdat de melkwegstelsel het dichtst bij de waarnemer fungeert als een lens die de lichtstralen van het verste stelsel (de "bron") afbuigt en concentreert.

Het blijkt dat deze ringvormige zwaartekrachtlenzen een uniek instrument zijn voor het beantwoorden van de fundamentele vragen die in de vorige sectie zijn beschreven. De ringvormige beelden van verafgelegen melkwegstelsels maken het mogelijk het te observeren, terwijl het onder normale omstandigheden te lichtzwak zou zijn. Dankzij het versterkend effect van de lens is er zo de mogelijkheid om oude, lichtzwakke sterrenstelsels te bestuderen. Bovendien zijn de vormen van de waargenomen beelden afhankelijk van de massaverdeling van de lens en de relatieve posities van de waarnemer, de lens en de bron. Dit maakt het mogelijk om exact de totale massa te meten van het lensmelkwegstelsel, dat over het algemeen niet alleen bestaat uit gewone materie, maar juist voornamelijk uit de onzichtbare donkere materie. Dat wil zeggen dat de zwaartekrachtring een uniek hulpmiddel is voor het meten van het gehalte aan donkere materie in melkwegstelsels. Een andere belangrijke wetenschappelijke toepassing van zwaartekrachtlenzen is de meting van de snelheid van de uitdijing van het heelal, de zogenaamde Hubble constante. De Hubble constante kan afgeleid worden als het bronstelsel een variabel

Figure A.4: Het licht dat uit een melkwegstelsel komt, wordt afgebogen door een melkwegstelsel voordat het op de Aarde aankomt. Op deze manier kunnen er verscheidene vervormde beelden van dezelfde bron van het melkwegstelsel zichtbaar zijn vanaf de Aarde. Zie de twee voorbeelden in de volgende figuur. (Credit: NASA, ESA & L. Calçada)

Figure A.5: Twee voorbeelden van zwaartekrachtlenzen. In het bovenste figuur buigt het centrale gele sterrenstelsel het licht van een ver blauw sterrenstelsel en creëert een blauwe boog. In het onderste voorbeeld wordt het licht van een quasar in vier verschillende beelden verdeeld (drie boven en één onder door het sterrenstelsel in het midden van het beeld).

object is, zoals een quasar. Dat is een actief sterrenstelsel met een sterke en variabele lichtkracht. In zo'n geval ziet de waarnemer de lichtintensiteit van de verschillende beelden van de quasar op verschillende tijden variëren. Astronomen monitoren deze zeer zeldzame zwaartekrachtlenzen omdat met behulp van de tijdsverschillen de Hubble constante kan worden geschat. Voor alle bovenstaande redenen is het noodzakelijk om deze kostbare astronomische fenomenen te vinden en te bestuderen en zo onze kennis over melkwegstelsels en de structuur van het heelal te vergroten.

Zoals gezegd zijn zwaartekrachtlenzen een zeer zeldzaam fenomeen. Het dat ontstaat bij ongeveer één op de duizend zwaardere melkwegstelsels. Het subtiel effect is moeizaam met het oog te identificeren in verzamelingen van melkwegstelselbeelden. Des te meer omdat de komende astronomische waarneemprojecten extreem grote hoeveelheid gegevens gaan produceren met afbeeldingen van miljarden melkwegstelsels, waaronder honderdduizenden zwaartekrachtlenzen. Het handmatig visualiseren en met het menselijk oog inspecteren van al deze beelden van melkwegstelsels zou fulltime banen creëren voor specialisten in het identificeren van zwaartekrachtlenzen! Daarom is het de moeite waard om automatische identificatiemethoden te ontwikkelen om zo de detectie van nieuwe zwaartekrachtlenzen te ondersteunen of zelfs te verbeteren. Dit is het hoofddoel van mijn proefschrift en voor het behalen van mijn doel heb ik tijdens mijn studie een methode ontwikkeld op basis van kunstmatige neurale netwerken voor het identificeren van zwaartekrachtlenzen. In de volgende sectie leg ik uit wat neurale netwerken zijn.

Neurale netwerken

Onlangs heeft een bepaald soort algoritmen een kwaliteits- en snelheidsrevolutie doorgemaakt in analyse van data. De algoritmen zijn geïnspireerd door de structuur van hersenen en de werking van neuronen en worden daarom kunstmatige neurale netwerken genoemd, of neurale netwerken. Neurale netwerken leren door het observeren van voorbeelden en ze zijn bijzonder effectief als de hoeveelheid beschikbare voorbeelden waarop ze getraind kunnen worden groot genoeg is. Neurale netwerken hebben bovenmenselijke prestaties bereikt in verschillende toepassingen. Een voorbeeld hiervan is het gerenommeerde neurale netwerk algoritme AlphaGO dat vier van de vijf wedstrijden van het spel Go won tegen de vijftienvoudige wereldkampioen Lee Sedol. Zo'n prestatie werd verondersteld zeer moeilijk te zijn om te bereiken. AlphaGo ontwikkelde zijn strategie door Go wedstrijden te observeren en tegen zichzelf te spelen. Een ander voorbeeld zijn zelfrijdende auto's ontwikkeld door

grote bedrijven. Die hebben waarschijnlijk in de toekomst een grote invloed op onze manier van reizen en op het woon-werkverkeer. Neurale netwerken zijn gebruikt in een grote verscheidenheid van publieke en commerciële toepassingen, met name op het gebied van beeldherkenning, zoals bijvoorbeeld wat Facebook en Google doen voor het herkennen van onderwerpen van foto's die op de online platformen zijn geüpload.

Dit Proefschrift

In dit proefschrift heb ik een methode ontwikkeld om met behulp van neural netwerken automatisch kandidaat-zwaartekrachtlenzen te identificeren uit beelden van astronomische objecten die dan vervolgens gebruikt kunnen worden de eerder genoemde onderzoeken. Om dit doel te bereiken heb ik een machine learning-algoritme ontworpen op basis van convolutionele neurale netwerken (convolutional neural networks). De gebruikte gegevens van dit proefschrift zijn afkomstig van de Kilo Degree Survey (KiDS). KiDS is een astronomisch waarneemproject dat wordt uitgevoerd met de VLT Survey Telescope (VST, zie Fig. A.6) die zich op de berg Paranal in de woestijn in Chili bevindt. Dit is een ideale plek voor astronomische observaties vanwege de weersomstandigheden: gedurende het hele jaar is het lucht niet of grotendeels onbewolkt en weinig turbulent.

KiDS heeft miljoenen sterrenstelsels aan de zuidelijke hemel waargenomen met een ongekennde beeldkwaliteit. In mijn proefschrift heb ik neurale netwerken ontwikkeld om zwaartekrachtlenzen te vinden in een grote hoeveelheid data, zoals beschikbaar voor KiDS. De neurale netwerken hebben, binnen enkele tientallen minuten, beelden van tienduizenden sterrenstelsels geanalyseerd. Dit heeft geleid tot het ontdekken van honderden zwaartekrachtlenzenkandidaten. Het vinden van zwaartekrachtlenzen met behulp van neurale netwerken zal zeer belangrijk worden voor toekomstige astronomische onderzoeken. In feite zullen miljarden sterrenstelsels waargenomen worden met een "computer oog", wat voor een menselijk oog ondoenlijk zou zijn om uit te voeren. Het gebruik van dit soort automatische analyses zal dan ook noodzakelijk worden.

In **Hoofdstuk 2** ontwerp ik een lenszoeker voor een neurale netwerk en trainde deze met behulp van melkwegbeelden van KiDS en gesimuleerde bronnen met zwaartekrachtlenzen. Er was behoefte aan gesimuleerde gegevens omdat er onvoldoende bekende zwaartekrachtlenzen beeldvoorbeelden beschikbaar zijn om het neurale netwerk te trainen. Na deze training heb ik het algoritme toegepast op een steekproef van 20.000 sterrenstelsels uit de verzameling van KiDS beelden. Het

Figure A.6: De VST-telescoop in de Atacama-woestijn, Chili.

algoritme identificeerde met succes twee van de drie eerder bekende zwaartekrachtlenzen die in het steekproef zaten en bovendien ongeveer 50 nieuwe betrouwbare zwaartekracht lens kandidaten.

In **Hoofdstuk 3** implementeer ik twee verschillende lenszoekers. Een lenszoeker voor het classificeren van de afbeeldingen op basis van hun vorm en een voor het classificeren op basis ook van hun kleur. Deze lenszoekers hebben een betere prestatie in vergelijking tot die in hoofdstuk 2. De reden hiervan is het gebruik van een verbeterde versie van de algoritmen en vooral het trainen met een verbeterde set van gesimuleerde bronnen met zwaartekrachtlenzen.

In **Hoofdstuk 4** presenteer ik een verzameling van vele honderden zwaartekracht lens kandidaten: de LinKS verzameling (Lenzen in KiDS). Deze verzameling is geselecteerd door het toepassen van de twee neurale netwerken, die in hoofdstuk 3 waren ontwikkeld, op ongeveer 88.000 sterrenstelsels van KiDS en deze te filteren en te rangschikken en zo een kleine eindset aan te bieden voor een laatste visuele inspectie door sterrenkundigen. De afbeeldingen van deze zwaartekrachtlenzen zijn te vinden op de website <https://www.astro.rug.nl/lensesinkids/>. Ik denk dat het zeer waarschijnlijk mogelijk is om effectief lenskandidaten te selecteren ook in de toekomstige nog grotere astronomische waarneemprojecten,

dankzij de hulp van kunstmatige neurale netwerken en met minimale menselijke interventie.

COMPENDIO

Introduzione

L'Universo, le galassie e la materia oscura

Secondo il paradigma cosmologico standard, il modello Λ CDM, viviamo in un Universo in perenne espansione, i cui ingredienti fondamentali sono luce, materia ordinaria, la misteriosa materia oscura e la ancora più misteriosa energia oscura che è la responsabile dell'espansione accelerata. L'inizio dell'Universo è fissato a 13.8 miliardi di anni fa con il cosiddetto Big Bang, quando l'Universo iniziò ad espandersi a partire da uno stato a elevatissima densità e temperatura. In questa situazione i vari costituenti erano quasi uniformemente distribuiti e nessuna struttura era presente. Con l'avanzare dell'espansione, l'Universo si raffreddò e le parti con una quantità leggermente più elevata di materia oscura cominciarono a diventare sempre più dense grazie all'attrazione gravitazionale. La materia oscura fu la prima a formare strutture poiché essa ha la proprietà di non interagire né con la luce né con la materia ordinaria e forse neanche con se stessa, essa "sente" solo la forza gravitazionale. Successivamente, a partire da circa 150 milioni di anni dopo il Big Bang fino a circa un miliardo di anni dopo, la materia ordinaria incominciò a collassare in queste dense regioni popolate dalla materia oscura, formando così le prime stelle (come il nostro Sole) e le prime galassie (sistemi di milioni o miliardi di stelle, gas, polvere e materia oscura, tenute insieme dalla forza di gravità). La materia oscura è stimata come il maggior costituente delle galassie, ad esempio il 95% della nostra galassia è stimato essere materia oscura. Come già sottolineato in precedenza, lo studio di questa componente è estremamente complicato, perché non può essere "vista" dato che non interagisce con la materia ordinaria e non emette luce.

Come si sono formate le galassie e come evolvono? Qual è la quantità dei loro costituenti? A quale velocità si espande l'Universo? Queste sono alcune delle domande alle quali dobbiamo dare tuttora una risposta esauriente. Gli studi per rispondere a queste domande fondamentali si basano sull'analisi delle immagini astronomiche. La luce non viaggia

istantaneamente, ma si trasmette ad una velocità finita, e quindi guardare una galassia molto lontana (a migliaia o milioni di anni luce), significa guardare un Universo più giovane. Ciò significa che le galassie che osserviamo in cielo sono una “fotografia” di ciò che esse erano quando la luce ha iniziato il suo viaggio verso di noi.

Ciò permette agli astronomi di avere fotografie dell’Universo a diversi istanti temporali e perciò ci apre una finestra sulla sua storia.

Il lensing gravitazionale

Una delle sorprendenti conseguenze della teoria della relatività generale di Einstein è che i raggi di luce vengono deviati dalla presenza di oggetti nelle vicinanze della loro traiettoria, ciò dà luogo a uno dei più spettacolari fenomeni astrofisici: il lensing gravitazionale. In particolare, quando si è abbastanza fortunati da avere due galassie allineate sulla linea di vista, la luce proveniente dalla galassia più lontana viene deviata ed amplificata, creando immagini multiple e distorte (vedere figure A.7 e A.8). Questi rari eventi astrofisici sono chiamati lenti gravitazionali, poiché la galassia più vicina all’osservatore funge da *lente*, deviando ed amplificando i raggi di luce provenienti dalla galassia più lontana, chiamata *sorgente*.

Il lensing gravitazionale è uno strumento unico per rispondere alle domande fondamentali descritte nella sezione precedente. Infatti, grazie all’effetto amplificante della galassia lente, esso permette di osservare e studiare galassie lontane e deboli che non sarebbero visibili altrimenti. Inoltre, la forma delle immagini osservate dipende quasi esclusivamente dalla massa della lente e dalla posizione relativa di lente, sorgente e osservatore. Ciò permette di poter misurare accuratamente la massa della galassia lente che è costituita non solo da materia ordinaria ma anche da materia oscura; il lensing gravitazionale è lo strumento più accurato per misurare il contenuto di materia oscura nelle galassie. *Dulcis in fundo*, un’altra importante applicazione è la misura della velocità di espansione dell’Universo, anche chiamata costante di Hubble. Essa può essere misurata se la galassia sorgente è un quasar, cioè una galassia dalla luce fortemente variabile, in tal caso un osservatore vedrà variare l’intensità della luce delle diverse immagini del quasar a diversi istanti temporali. Gli astronomi monitorano questi rari oggetti poiché la quantità di tempo che trascorre fra la variazione di un’immagine e l’altra dipende dalla costante di Hubble che può essere stimata in tal modo. Per i motivi sopra elencati, e per altre importanti applicazioni, è necessario individuare e studiare questi affascinanti eventi astrofisici per ampliare la nostra conoscenza sulle galassie, la loro formazione ed evoluzione e

Figure A.7: La luce proveniente da una galassia *sorgente* lontana viene deviata da una galassia *lente* intermedia prima di arrivare sulla Terra. In questo modo dalla Terra saranno visibili immagini multiple e distorte come gli esempi nella figura successiva. (Fonte: NASA, ESA & L. Calçada)

Figure A.8: Due esempi di lente gravitazionale. Nel primo caso la galassia gialla centrale devia la luce di una galassia lontana, creando un arco blu. Nel secondo caso abbiamo che la luce proveniente da una galassia quasar lontana che è stata “sdoppiata” in quattro immagini differenti (tre sopra e una sotto) dalla galassia al centro della figura.

sulla struttura dell'Universo. Tuttavia le lenti gravitazionali sono rare, circa una ogni mille galassie e difficili da identificare ad occhio fra milioni di altre galassie. L'identificazione diventa ancora più problematica se si pensa che le future campagne osservative astronomiche produrranno una quantità enorme di dati, miliardi di galassie fra le quali un numero potenziale di migliaia di lenti gravitazionali. Visualizzare ed analizzare tutte queste galassie richiederebbe speciali posti di lavoro a tempo pieno dedicati esclusivamente ad identificare le lenti gravitazionali! Perciò è necessario sviluppare metodi di identificazione automatica per supportare la scoperta di nuove lenti gravitazionali. Questo è l'obiettivo principale del mio lavoro di tesi che mi ha portato a sviluppare un metodo per identificare lenti gravitazionali basato sulle reti neurali.

Le reti neurali

Negli ultimi anni un particolare algoritmo ha portato una rivoluzione nella disciplina dell'intelligenza artificiale. Esso è ispirato alla struttura e al funzionamento dei neuroni e prende appunto il nome di *rete neurale*. Questi algoritmi “imparano” *osservando* esempi e sono particolarmente accurati se la quantità di esempi su cui possono essere *allenati* è grande abbastanza. Le reti neurali hanno raggiunto prestazioni al di sopra delle capacità umane in diverse applicazioni. Celebre è l'algoritmo AlphaGo, basato sulle reti neurali, che ha vinto quattro partite su cinque al gioco del Go contro il diciotto volte campione del mondo Lee Sedol, un'impresa che fino a qualche anno fa era considerata difficilmente raggiungibile. AlphaGo ha sviluppato una sua strategia *osservando* partite di Go e giocando contro se stesso. Un altro esempio sono le automobili autoguidate che importanti compagnie stanno sviluppando e che cambieranno radicalmente il sistema dei trasporti e la nostra quotidianità. Ad oggi le reti neurali sono utilizzate nelle più varie applicazioni commerciali e non, specialmente nel campo del riconoscimento di immagini come fanno ad esempio Facebook e Google per riconoscere i soggetti nelle fotografie caricate online.

Questa tesi

In questa tesi ho sviluppato un metodo per identificare automaticamente possibili lenti gravitazionali in campagne osservative astronomiche, al fine di consentire studi individuali e d'insieme di questi oggetti scientificamente preziosi. Per raggiungere questo obiettivo, ho progettato un algoritmo di apprendimento automatico basato sulle reti neurali. I dati utilizzati in

Figure A.9: Il telescopio VST nel deserto di Atacama in Chile.

questa tesi provengono dalla campagna osservativa chiamata Kilo-Degree Survey (KiDS). Essa è stata condotta utilizzando il telescopio di 2.61 metri di diametro VLT Survey Telescope (VST, vedi figura A.9) situato nel deserto Atacama in Chile, un luogo ideale per osservazioni astronomiche per via delle condizioni ambientali che offrono cieli limpidi e poco nuvolosi.

Grazie a KiDS sono state osservate milioni di galassie situate nel cielo del Sud del mondo con una qualità mai raggiunta prima. Nel mio lavoro di tesi ho programmato reti neurali per individuare lenti gravitazionali in una tale marea di dati. Le reti neurali hanno analizzato immagini di decine di migliaia di galassie in poche decine di minuti e hanno portato ad identificare centinaia di possibili lenti gravitazionali. Le immagini di queste galassie possono essere visualizzate sul sito web <https://www.astro.rug.nl/lensesinkids/>. L'applicazione delle reti neurali per identificare lenti gravitazionali diventerà estremamente importante per le future campagne osservative. Infatti la sovraumana impresa di osservare ed analizzare ad occhio nudo milioni di galassie sarà resa possibile e non potrà prescindere dall'applicazione di questo tipo di ricerca automatica.

Più in dettaglio, nel **Capitolo 2** di questa tesi ho progettato e “addestrato” un cercatore automatico di lenti gravitazionali basato su una rete neurale. Per addestrarlo ho utilizzato immagini di galassie

reali da KiDS e un campione di immagini simulate di lenti gravitazionali. E' stato necessario utilizzare dati simulati perché non sono disponibili abbastanza lenti gravitazionali conosciute per formare un insieme di immagini sufficientemente grande per addestrare con successo una rete neurale. Dopodiché ho applicato l'algoritmo a un campione di 20 000 galassie, esso ha identificato con successo due delle tre lenti gravitazionali precedentemente note presenti nel campione e altre 50 possibili lenti gravitazionali.

Nel **Capitolo 3** ho implementato due diversi cercatori automatici, uno per classificare le immagini in base alla loro forma e l'altro per classificarle basandosi anche sul loro colore. I metodi qui sviluppati hanno una migliore prestazione rispetto a quello sviluppato nel Capitolo 2. Il motivo sta nell'utilizzo di una versione migliorata degli algoritmi e soprattutto nel fatto che essi sono stati allenati con un campione migliorato di lenti gravitazionali simulate.

Nel **Capitolo 4** ho presentato un campione di diverse centinaia di possibili lenti gravitazionali, chiamato LinKS (Lenses in KiDS). Il campione è stato compilato applicando dapprima le due reti neurali sviluppate nel capitolo 3 a circa 88 000 galassie da KiDS. Successivamente i candidati selezionati dalle reti neurali sono stati filtrati con una ispezione visiva. Le immagini di queste galassie possono essere visualizzate sul sito web <https://www.astro.rug.nl/lensesinkids/>. Inoltre ho stimato che, utilizzando le reti neurali e con minimo intervento umano, nelle future campagne osservative astronomiche sarà possibile selezionare migliaia di possibili lenti gravitazionali.

AFTERWORD & ACKNOWLEDGEMENTS

OH MY GOD! It seems literally a life ago when I first entered in Kapteyn for going to meet my future supervisor. The astronomical images on the walls that welcome you while entering the institute made me feel like I was going to be in the right place! In fact, I have been welcomed in an environment where research in the field is done at the highest level and which offers a very culturally diverse experience, thanks to people coming from the most different parts of the world. I would have never expected to learn so much, not only about Science, but about many different aspects of life. Surely, one the most notable achievement has been learning how to bike without hands!

During the days spent in the office, it has been extremely exciting handling the farthest accessible data that the Universe can provide us and being allowed to speculate about its structure while admiring its beauty. This has been a sort of coronation of the path that I started because of curiosity and appreciation of the beauty of the physical world. I imagine there are mainly three types of motivations that strongly push people to study Physics. Who enjoys the process of the scientific methodology; who believes in the positive contribution that scientific research gives to society; who is fascinated and curious about the structure of the Universe, struck by its majesty. Fascination without curiosity can similarly lead people towards religions. Often, curiosity and fascination go hand in hand creating the need to answer to the question “Why?”. This is ultimately the very thing that sparkles the human adventure in the realm of philosophy and in which science is our only objective method of describing reality. Usually, an experienced scientist knows that to this question will be never found an answer. This could cause a certain frustration since the quest is ultimately impossible. Nevertheless, science can give an important gift to humanity, the power of shaping the Universe in increasing powerful ways which come together with greater risks. Anyway, science, together with education, can also contribute in forming an ideology which takes into account our position in the Universe and that could make us a species more in equilibrium with the objective reality and, ultimately, with ourselves.

Supposedly, when I will be sitting defending my thesis, there will be a fresco called "the Tree of Knowledge" on the wall behind me. It depicts a humanity working boldly in building elaborated social and technical structures around a big tree, the Tree of Knowledge. Despite the effort, humans struggle in climbing and owning the Tree. Instead, a couple of monkeys are just enjoying themselves on the highest branches. Have they been owning the Tree better than humans do? I wish humanity to climb that Tree by knowingly setting its goals and take the responsibility for choosing its own future.

Now, I could not be sitting there and have all these thoughts and experiences without the people I met and who gave me the most out of this adventure. I can not spare to reciprocate at least with a bit of gratitude and love. I wish to have started writing this acknowledgments much more in advance!

Thanks to my supervisors; **Léon**, I am extremely grateful for the possibility you gave me to start this adventure. As a supervisor you always gave me enough freedom and challenged me in a constructive way for every step I took, your help was invaluable. Moreover, I could see in your actions a sense of morality that is not quite common among the people at your professional level.

Gijs, you are a very acute, precise person, precise as every scientist should be. This has affected positively my work and my personal growth. Your curiosity pushes you in many different directions, all worth to explore. I have appreciated the hard work you put in my thesis and your interest in my personal well-being during the duration of the PhD. It has been always very instructive talking with you about scientific and non scientific topics.

Crescenzo, you have been a supervisor, a teacher and a friend, always with the humbleness that distinguishes you. You gave me the possibility of realizing one of my greatest dreams: observing the Universe with some of the best equipment a human can use. I won't forget that you stood on my side when it was necessary.

Thanks to my reading committee; **Reynier Peletier**, I am glad you are in my reading committee I really appreciate how you care about science and people in the same careful and honest way.

Giuseppe Longo, it is because of your inspiring lectures that I chose this road. It is appropriate that you will be at the end of what you contributed to start.

Chris Fassnacht, I remember with pleasure the time spent in Aosta. They say that you are one of the kindest people in the strong lensing community.

Thanks to the people without whom my thesis would not have been here; Saikat, your contribution to this thesis has been fundamental, it has been a real pleasure working with you. Other than a teammate, I found a person with many affinities, with whom I can talk for hours and spend high quality time together.

Hugo, Ewout, Jelte, thanks for your continuous support in taming the data and AstroWISE.

Giorgos, the webapp you developed for inspecting the lens candidates has been extremely useful.

Thanks to the computer group **Martin, Eite** and especially **Leon** for having assisted and helped me in all the computer related issues.

Lingyu and **Will**, our short collaboration has been really enjoyable. You are both delightful persons.

Thanks to the lensing group people, **Crescenzo, Cristiana, Dorota, Gijs, Giorgos, Hannah, John, Léon, Saikat, Sampath** for the interesting discussions during our group meetings and the funny ones during our lensing happy hours.

Tzi and **Hallo**, thanks for having helped me in writing the Samenvatting.

Rustam, thanks for the amazing cover and the amazing person you are!

Thanks to my paranympths; **Aku**, I remember when you asked around who would be willing to go to a music festival in Tilburg. I was the only to say yes, one of the many yes that came unconditionally from both sides in the years to come, which made me exploring many new places and enriched me as a person. You are one of the few people I feel I can completely rely on, you showed it with your actions rather than with your words. It is not by chance that you will be my paranympth. I am glad you have found in **Sonja** your sweet half.

Olmo, your acceptance and your trust in people and in the unfolding of events made us meet immediately when you started your adventure here. From there, we have built a deep bond and found so many points that unite us - including a small dose of revolutionary anarchism. Then, I find really appropriate that you will be my partner in this pretty institutional moment.

Thanks to the people that make the institute working and a better place; **Lucia, Christa, Martine, Ramona**, you always been very supportive. It is because of you that the institute is the big family that it is.

Scott, thanks for the big enthusiasm you put in directing the institute.

Marc, thank you for your kind support as my mentor from the PBC.

Hugo van Woerden, I really appreciate the kind touch you have not missed to put every time there was the occasion.

Marlies, Pete van der Kruit, Edwin, Peter Bartel, Rien, thanks for bringing attention to interesting topics related to Astronomy and not necessarily to scientific research. **Renzo**, thanks for the spontaneous conversations during otherwise boring passages in the corridor.

Thanks to my office mates; **Anastasia**, I could not have been welcomed in Kapteyn in a better way. One creates and lives in many families during the course of its life and often can not choose them. I have been extremely lucky that the events made us a family in which we can be ourselves, support and love each other no matter what. You made me feel our office like home.

Mpati, it was a pleasure having you there during the nice season. Your attitude towards life and our long discussions made the summer looking more shiny.

Teymoor, your love for science and fascination for the human creation has been contagious and a pleasant addition towards the end of my PhD.

Thanks to my friends and fellow PhD students; **Filippo**, you brought me to explore where there is fun and much, much more. I am really jealous of the unforgettable moments we spent together, I will bring them with me for the rest of my life. I could spot love in your actions, even if you are good at hiding it from yourself and the others. Thank you for having introduced me to beautiful people as **Laura & Claudio, Giulia** and **Leo**.

Alessandro, thanks for the tons of (not) useful information. The evenings spent chatting and having fun are of that kind that make life enjoyable in its simplest form.

Bharat, thanks for some of the funniest moments of the PhD, for always being there when needed and for your culinary abilities!

Francesco, thank you for your personality. I could see your land in there, something that is inextricably inside me and I truly love. It looks like that this time we are going to finally be in a real party together :)

Davide, thank you for your rare veracity and for your imagination that brought me into fantastic adventures just sitting around a table. **Antonino & Laura** thank you for having provided that table and to your big generosity.

Marisa, thank you for the delicious dinners, occasions when we could speak as we truly are and sincerely appreciate each other.

Mustafa, thank you for your amazing hospitality, your generosity and the long and funny conversations. And thanks for the tea! I wish you all the luck to your amazing family!

Wouter, it has been a pleasure working together as teaching assistants. I have loved arguing with you on many different matters and I have always

enjoyed your company.

Smaran, you have brought to the institute a different, precious perspective. Your values and positive attitude towards life are of that kind that can make the world a better place.

Sampath, sharing the burden of finishing the thesis has been quite some comfort. Now I will reach you on the other side :)

Danielito, the life of the party, thank you for making every possible moment another occasion to enjoy life.

Laurent, looking at you, the new generation looks very promising. I am very proud to be your old grandpa :)

Nika, as the Rolling Stones would say, thank you for being a rainbow!

Simon, thank you for providing chaos and confetti.

Kristiina & Karlis, thanks for the incredible dinners and your easy-goingness.

Stefano, thank you for your wishes that came true. I wish you to find all the adventure you crave for.

And thanks for the enjoyable time spent together to **Cristiana, Evgenia, Lenka, Pratyush, Andrea, Tadeja, Giulio, Evandro, Avanti, Pooja, Hyoyin, Anne, Pranav, Tirna, Elaheh, Sara, Umit, Samira, Georg**.

Thanks also to the amazing people who live outside the institute; **Jonas, Georgi, Piermichele** and my **IBM** new friends and colleagues. **Vaggelis**, thanks for the sweet attention that you have for your friends.

Giappo, how nice to converse with you and how nice it is to get infected by all your interests. I miss our cohabitation and our afterwork time together, I will bring them into my heart despite our difference on home issues!

Thanks to all the people I am forgetting!

Grazie ai miei amici di Caserta. **Giovanni & Federica, Roberto & Alessandra, Pietro & Valeria, Matteo, Luigi, Emanuele, Claudio & Ilia**. Siete uno dei pochi motivi che mi fa tornare a casa di tanto in tanto. **Giovanni**, anche se stavolta non ci potrai essere, conto di averti come testimone al mio matrimonio!

Grazie agli amici che sono rimasti dai tempi dell'università; **Riccardo**, siamo sempre riusciti a rubare momenti assieme nonostante la distanza, il poco tempo e il fatto che la pensione sia ancora lontana!

Marcello, sei sempre la solita preziosa lotta ma di cui non si può fare proprio a meno :)

Grazie a **Ferragamo**, il mio compagno di sofferenza a distanza, ce l'hai fatta giusto due settimane prima di me :D

Grazie alla mia famiglia, che ci è sempre stata nonostante tutte i problemi logistici e le vicissitudini. Esiste un posto da qualche parte dove potremo riunirci se tutti lo vorremo e saremo disposti ognuno a rinunciare a qualcosa. **Mamma, Papà** grazie per avermi permesso di intraprendere la mia strada, anche quando a malincuore e in disaccordo. Grazie per avermi insegnato a sognare in qualcosa di migliore.

Chicca, è rassicurante avere un complice col quale sognare e forse un giorno rivoltare l'ineluttabile. **Anto**, thank you for taming my sister and being tamed by her, thank you for the affection you put each time we visit you. You are a great addition to our crazy family!

Ludovica, che noia che sarebbe senza di te, hai portato una novità che ha reso la famiglia più viva e unita. Posso assicurarti che lì fuori c'è un mondo bellissimo che bisogna imparare a trovare.

Grazie per l'amore incondizionato e una splendida infanzia a **Nonna** e a quelli che non ci sono più che ritrovo ogni tanto nei miei sogni.

Grazie ad **Anna** per essere la zia che vive nella città figa e lontana, compagna e supporto nelle mie avventure adolescenziali.

Дякую **родині Каті**, за те що ви прийняли мене з усією можливою любов'ю. Навіть якщо я не можу сказати і слова українською, я завжди почуваюсь з вами як вдома. Thanks to Katya's friends, you are all exceptional people and I was genuinely surprised by how you welcomed me in a way only best friends do.

Katya, your determination in taking me and building a path together is of that kind that goes beyond the cynicism of events, that makes men perform superhuman feats and elevates life by giving it meaning. We have learned to love, to love each other and to love ourselves as much as necessary for the exceptional feat of improving ourselves and a bit of all the rest.