



Universiteit
Leiden
The Netherlands

Reducing false positives in strong lens detection through effective augmentation and ensemble learning

Rezaei Badafshani S.; Chegeni, A.; Nagam, B.C.; McKean, J.P.; Baratchi, M.; Kuijken, K.H.; Koopmans, L.V.E.

Citation

Chegeni, A., Nagam, B. C., McKean, J. P., Baratchi, M., Kuijken, K. H., & Koopmans, L. V. E. (2025). Reducing false positives in strong lens detection through effective augmentation and ensemble learning. *Monthly Notices Of The Royal Astronomical Society*, 538(2), 1081-1095. doi:10.1093/mnras/staf327




Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/4292223>

Note: To cite this publication please use the final published version (if applicable).

Reducing false positives in strong lens detection through effective augmentation and ensemble learning

Samira Rezaei ^{1,2★}, Amirmohammad Chegeni,^{1,3,4} Bharath Chowdhary Nagam,⁵ J. P. McKean ^{5,6,7},
Mitra Baratchi,² Koen Kuijken¹ and Léon V. E. Koopmans ⁵

¹*Leiden Observatory, Leiden University, NL-2333 CC Leiden, the Netherlands*

²*Leiden Institute of Advanced Computer Science (LIACS), Leiden University, NL-2333 CC Leiden, the Netherlands*

³*Dipartimento di Fisica e Astronomia ‘G. Galilei’, Università di Padova, Via Marzolo 8, I-35131 Padova, Italy*

⁴*INFN-Padova, Via Marzolo 8, I-35131 Padova, Italy*

⁵*Kapteyn Astronomical Institute, University of Groningen, Postbus 800, NL-9700 AV Groningen, the Netherlands*

⁶*South African Radio Astronomy Observatory (SARAO), P.O. Box 443, Krugersdorp 1740, South Africa*

⁷*Department of Physics, University of Pretoria, Lynnwood Road, Hatfield, Pretoria, 0083, South Africa*

Accepted 2025 February 19. Received 2025 February 10; in original form 2024 September 2

ABSTRACT

This research studies the impact of high-quality training data sets on the performance of Convolutional Neural Networks (CNNs) in detecting strong gravitational lenses. We stress the importance of data diversity and representativeness, demonstrating how variations in sample populations influence CNN performance. In addition to the quality of training data, our results highlight the effectiveness of various techniques, such as data augmentation and ensemble learning, in reducing false positives while maintaining model completeness at an acceptable level. This enhances the robustness of gravitational lens detection models and advancing capabilities in this field. Our experiments, employing variations of DenseNet and EfficientNet, achieved a best false positive rate (FP rate) of 10^{-4} , while successfully identifying over 88 per cent of genuine gravitational lenses in the test data set. This represents an 11-fold reduction in the FP rate compared to the original training data set. Notably, this substantial enhancement in the FP rate is accompanied by only a 2.3 per cent decrease in the number of true positive samples. Validated on the Kilo Degree Survey data set, our findings offer insights applicable to ongoing missions, like *Euclid*.

Key words: gravitational lensing; strong – methods: data analysis – techniques: image processing.

1 INTRODUCTION

Strong gravitational lensing occurs when light from a distant galaxy gets bent by the curvature of space–time that is caused by another galaxy along our line of sight. This phenomenon significantly impacts our understanding of the Universe. Strong gravitational lensing, as predicted by the theory of general relativity, occurs when the foreground galaxy acts as a lens, creating multiple magnified and distorted images of the background galaxy (see Treu 2010 for a review). Strong lensing serves as a unique tool for testing models for galaxy formation (Renneby et al. 2020) and cosmology (Cao et al. 2015). Measuring the mass components of early-type galaxies, constraining their stellar initial mass function, and determining their inner mass density profiles (Wucknitz, Biggs & Browne 2004; Koopmans et al. 2006; Treu et al. 2006; Bolton et al. 2008; Auger et al. 2009, 2010; Spiniello et al. 2012, 2014; Spingola et al. 2018), probing the nature of dark matter through detailed modelling of the surface brightness distribution of lensed images (Vegetti et al. 2012, 2014; Ritondale et al. 2019; Gilman et al. 2020; Hsueh et al. 2020), testing models for the expansion of the Universe and dark energy

(Suyu et al. 2010, 2013; Bonvin et al. 2017; Wong et al. 2020), and measuring the Hubble constant (Birrer & Treu 2021) are some applications of studying this phenomenon.

Despite the profound insights gained by strong lensing, identifying gravitational lenses remains a challenge, as these events are rare, with the probability of one gravitational lens being found in about a thousand observed galaxies (Chae et al. 2002; Wardlow et al. 2013; Amante et al. 2020). Visual inspection of the extensive parent population is both time-consuming and susceptible to incompleteness (Jackson 2008; Marshall et al. 2016; More et al. 2016). Consequently, the discovery of gravitational lenses often relies on the application of selection criteria in catalogue space, considering factors such as optical colour, radio spectral index, total flux density, and the morphology of candidate lensed images (see Spingola et al. 2019 for an illustrative example). Despite the use of these criteria, some degree of visual inspection remains necessary to validate potential lens candidates. As we anticipate parent samples to surpass 10^7 galaxies in size, driven by the data volumes from wide-field surveys conducted by the Vera C. Rubin Observatory (Bianco et al. 2022), the Nancy Grace Roman Space Telescope (Mosby et al. 2020), and *Euclid* (Laureijs et al. 2011), the imperative for sophisticated automated search techniques becomes evident. Developing such techniques is crucial to efficiently navigate and analyse data sets of this scale. The

* E-mail: rezaei@strw.leidenuniv.nl

literature has seen the emergence of several intelligent approaches tailored to diverse imaging surveys, reflecting the need for innovative and effective solutions in the face of increasingly massive data sets. Metcalf et al. (2019) conducted a lens finding challenge focused on optical/infrared data sets, revealing that automated approaches, including machine learning algorithms, outperformed traditional methods.

Convolutional Neural Networks (CNNs), a subset of machine learning models, have emerged as significant tools for cosmological and astronomical applications. These models excel at identifying patterns in complex data sets, making them highly effective for a range of scientific tasks. CNN methods have been applied in the study of exoplanet detection, where they enhance the identification and analysis of potential exoplanets from large data sets (Shallue & Vanderburg 2018; Priyadarshini & Puri 2021). In the field of radio astronomy, CNNs have been instrumental in classifying radio galaxies by analysing their morphological features (Aniyam & Thorat 2017; Becker et al. 2021). They have also been used in the detection and analysis of gravitational waves, helping to filter noise and improve signal detection from data collected by observatories such as Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo (Gebhard et al. 2019; Baltus et al. 2021). Furthermore, CNNs have contributed significantly to distinguishing alternative dark energy (Chegeni et al. 2024) and dark matter (Khosa et al. 2020) scenarios compared to the standard model of cosmology. In the study of the cosmic microwave background (CMB) maps, CNNs have been employed to fill the masked regions of the CMB (Sadr & Farsian 2021), extract cosmological parameters and identify subtle features in the CMB data (Krachmalnicoff & Tomasi 2019; Sadr & Farsian 2021). These diverse applications underscore the versatility and power of CNNs in enhancing our understanding of the Universe. A comprehensive review of CNN applications across various astronomical problems was presented by Rezaei et al. (2025).

Numerous studies have investigated the application of CNNs for identifying strong gravitational lenses in vast data sets, automating a process that traditionally required extensive manual effort. For instance, Petrillo et al. (2019b) and Rojas et al. (2022) showcased the efficiency of CNNs in detecting gravitational lenses in the Kilo Degree Survey (KiDS) data, while Nagam et al. (2023) recently proposed a new pipeline-ensemble model using Densely Connected Neural Networks (DenseNets) to reduce FPs in the identification of strong gravitational lenses, making it more suitable for large-scale astronomical surveys like *Euclid*. An additional example of using ensemble techniques for identifying gravitationally lensed quasars was presented by Andika et al. (2023), where an ensemble averaging approach combines state-of-the-art convolutional and transformer-based neural networks applied to multiband images.

Utilizing a CNN classifier on both single-band and multiband KiDS images, Li et al. (2021) identified a sample of 97 new high-quality strong lensing candidates. This effort contributes to a total of 268 high-quality candidates from KiDS, optimizing classifier efficiency for future large-scale surveys. In addition, Rezaei et al. (2022) developed and validated CNNs for detecting galaxy-scale gravitational lenses in interferometric data from the International Low-Frequency Array (LOFAR) Telescope, aiming for predicting a pure selection of lens candidates. The effectiveness of CNNs in these diverse settings highlights their robustness and adaptability, making them indispensable tools for the automated detection of strong gravitational lenses across various astronomical surveys.

Current research on CNNs has largely focused on refining model architectures and optimizing hyperparameters to improve perfor-

mance metrics, as highlighted in recent reviews (e.g. Rahman Minar & Naher 2018; Menghani 2023). While these technical improvements are essential, the quality of training data is equally crucial for CNN effectiveness. High-quality, diverse, and representative data sets significantly boost a model's generalization ability and robustness. For instance, Canameras et al. (2024) demonstrates that tailoring the training data set in both the lensed and non-lensed samples can lead to improved results. That study shows how CNNs perform best when the data set includes mock lenses closely resembling real, detectable systems with multiple lensed images that are bright and deblended. Such samples help the model to clearly differentiate between lensed and non-lensed objects. Likewise, including a high proportion of non-lensed contaminants (images resembling lenses without strong lensing features), improves the model's performance by learning to filter out such non-lensed objects.

By exploring the relationship between training data quality and the performance of a CNN, with a particular architecture performance, we aim to uncover insights that go beyond conventional optimization strategies. Specifically, we seek to understand how variations in training data characteristics, such as data distribution, labelling accuracy, and data set size, can influence a CNN's ability to learn and generalize from the provided data. Additionally, we investigate whether certain architectures exhibit greater resilience to variations in training data.

The structure of this paper is organized as follows. Section 2 outlines the methodology for generating the training, validation, and test data sets, focusing on the creation of two primary classes: lensed and non-lensed data. In Section 3, we delve into the exploration of different training strategies and their respective definitions. Moreover, we describe the CNN architectures utilized in this study. We conduct an analysis of several CNN architectures to demonstrate that the dependence on training data is not specific to any particular CNN architecture. Section 4 presents a comparative analysis of the results obtained from each training strategy, considering both the training data and the CNN architecture employed. Lastly, in Section 5, we engage in a thorough discussion of the results obtained in this study, providing insights and reflections on the methodology employed and suggesting potential avenues for future research and development.

2 TRAINING AND TEST DATA SET

In order to assemble the training, validation, and test data sets for this research, it is necessary to establish two main categories of data: 'lensed' and 'non-lensed' samples. A lens system is created by pairing an elliptical galaxy, specifically a Luminous Red Galaxy (LRG) from the actual KiDS data release 4 (DR4; Kuijken et al. 2019), denoted as the 'lensing galaxy' or 'foreground galaxy', with a simulated lens configuration (mock lens). The same LRGs can also be used in the non-lensed class together with a selection of spiral galaxies and contaminants collected from the KiDS catalogue. A visual representation of both classes is provided in Fig. 1. The top row exhibits a diverse array of lensed samples, showcasing various morphologies and configurations, offering a representative glimpse into the spectrum of lensed phenomena within our data set. The middle panel displays samples of LRGs, serving as both foreground lensing galaxies and non-lensed instances. Furthermore, the non-lensed category encompasses not only LRGs, but also spiral galaxies depicted in the bottom panel of Fig. 1.

The lensed and non-lensed samples within our data set display a wide range of surface brightness distributions. To ensure consistency




Figure 1. These images illustrate the data set’s diversity, presenting lensed phenomena and non-lensed galaxies. The top row showcases various lensed samples, offering insights into their morphology and configurations. The middle panel displays a selection of LRGs, serving as both foreground and non-lensed instances, while the bottom panel includes spiral galaxies. More details about the properties of these galaxies are available in Section 2 and Table 1. Each image has a size of 101×101 pixels, which corresponds to an area of 20×20 arcsec.

in pixel value ranges across both classes of data, we implement a *MinMax* normalization. The objective of this normalization is to standardize the model by aligning the distribution of inputs. Although this normalization method results in scaling the absolute loss amplitude in training, it does not affect our capability to identify lens candidates. This is because our analysis focuses on the relative surface brightness of the lensed images and non-lensed source emission within each simulated sample. Mathematically, this normalization can be expressed as,

$$x_{\text{normalized}} = \frac{x - \min(x_d)}{\max(x_d) - \min(x_d)}, \quad (1)$$

where x represents the value of a specific pixel and x_d encompasses all pixels in an image. Through this process, all pixel values within a given image are mapped to the range of (0, 1). A square-root stretch is applied to this normalized image to enhance features with lower surface brightness.

In the following subsections, we provide detailed information on how both lensed and non-lensed classes of data are generated.

2.1 Non-lensed samples

The task of distinguishing genuine gravitational lenses from various celestial objects in surveys like KiDS is challenging due to factors such as the diverse range of objects observed and variations in colour. Objects spanning different morphologies, including galaxies and artefacts, can potentially mimic lensing features, while limitations in survey resolution and depth may obscure faint lensed signals. To effectively train CNNs for lens detection, it is crucial to compile a comprehensive data set that encompasses a wide variety of objects with diverse characteristics. By exposing the CNNs to this diverse data set during training, they can learn to distinguish between genuine lensing events and contaminants, thereby improving the

accuracy and reliability of lens detection methods in astronomical surveys.

Our approach in selecting non-lensed samples aligns with previous studies, such as by Nagam et al. (2023) and Petrillo et al. (2019b). It comprises three distinct classes of data: (a) 3000 LRGs with an r magnitude below 21; (b) 2000 sources that were previously misidentified as mock lenses in earlier tests conducted by Petrillo et al. (2017); and (c) 1000 galaxies visually classified as spiral galaxies through the GalaxyZoo project (Willett et al. 2013; Melvin et al. 2014). This diverse selection introduces a wide range of objects with different characteristics, enhancing the robustness of our training data set for the CNNs.

We use the same backbone of data as in previous studies (Petrillo et al. 2019b; Nagam et al. 2023), but we make changes to improve the CNN performance. In order to understand the data, we need to investigate deeper into the properties of the training data set. For those 3000 LRGs (see above), we calculate the Sérsic profile (Sérsic 1968), which is a mathematical function used to describe the distribution of light in galaxies. The parameters of this profile provide insights into their structural properties such as size, magnitude/flux and morphology. By fitting the observed intensity profiles of LRGs, cropped to 20×20 pixel cutouts, with the Sérsic function, we can extract parameters such as the effective radius and Sérsic index. We use these cutouts to eliminate the effects of contamination in our Sérsic profile calculations. Our aim is to shed light on the nature and distribution of galaxies in our sample. Moreover, we introduced an additional parameter referred to as *galaxy complexity*, which is defined as,

$$\text{Complexity} = \frac{S_I}{S_P}, \quad (2)$$

where S_I denotes the galaxy’s integrated surface brightness, and S_P represents the galaxy’s peak surface brightness. This metric provides insights into the extent of an elliptical galaxy, helping us assess its




Figure 2. The distribution of the lensing galaxy model parameters used for the training data set; lensing galaxy effective radius (top) and the lensing galaxy complexity (bottom).

size. In our experiments, we considered this property to enhance our understanding of the lensing galaxy’s dimensions, a critical factor when selecting a suitable mock lensed system to be added to a lensing galaxy.

Fig. 2 shows the distribution of the effective radius (top) and compactness (bottom) for the LRG samples to be considered as both non-lensed samples and foreground lensing galaxies. Notably, the selection of foreground lensing galaxies exhibits a significant bias toward effective radii in the range of 0.65–0.85 arcsec and complexity between 70 and 80. This observation proves invaluable, particularly in understanding the training data and its impact on the performance of trained CNNs. Later in Sections 4.2 and 4.3 we show how the properties of training data affect the completeness and purity of the generated lens candidates.

2.2 Lensed samples

Lensed samples are created by combining simulated gravitational arcs, rings, quads, and doubles and a selection of foreground elliptical galaxies. These elliptical galaxies, which also act as the foreground lensing galaxies, are collected from the KiDS DR4, and therefore, provide a varied and representative collection of instances for both training and evaluation purposes. The process entails producing artificial distortions, such as arcs, rings, quads, and doubles, around selected foreground galaxies to replicate the appearance of gravitational lensing. The goal of these simulated lens configurations is to mimic the gravitational lensing characteristics observed in actual astronomical data, thereby providing a comprehensive and diverse data set for training and evaluation. In the following, we explain

Table 1. The parameter ranges for the Singular Isothermal Ellipsoid (SIE) representing the lens and the Sérsic model representing the source used in our simulations. These parameters contribute to the diversity of lensed images generated for training and testing purposes. The units and the range of values for each parameter are also indicated for reference.

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

how the ‘lensing galaxies’ and the ‘mock lenses’ are selected for the purpose of this work.

2.2.1 Foreground galaxies

Similar to previous work, such as by Petrillo et al. (2017, 2019a) and Nagam et al. (2023), we focus on low redshift ($z \leq 0.4$), massive early-type galaxies (i.e. LRGs), which have been established as the predominant contributors to the lensing galaxy population (e.g. Oguri 2006; Möller, Kitzbichler & Natarajan 2007). The training data set comprises 4411 unique LRGs, with an additional 511 samples allocated for validation and another 511 for samples for testing. Further details on the selection process of LRGs to be considered as foreground galaxies is presented by Nagam et al. (2023).

In contrast to previous studies (e.g. Petrillo et al. 2019b; Li et al. 2021; Nagam et al. 2023), which randomly select a mock lensed source and a foreground galaxy to form a lens system, we consider the properties of both components. This approach enables us to generate realistic gravitational lens samples, but to also provide our deep learning model with clear samples that minimize the risk of misinterpretation, which should improve the accuracy of our analysis. This would also enhance the model’s ability to generalize to unseen data. More details on how we have adjusted the selection of LRGs to be used in combination with a mock lens is provided in Section 4.3.

2.2.2 Mock lenses

Developing a robust training data set for strong gravitational lensing detection poses unique challenges, demanding representation across a diverse set of lensing configurations, while maintaining the spatial sampling of KiDS (~ 0.2 arcsec pixel $^{-1}$). The selection of lens and source parameters is detailed in Table 1.

We model the sources using a Sérsic profile and we exclude highly elliptical sources by restricting to axis ratios greater than 0.3 (Petrillo et al. 2017). The effective radius is randomly assigned within the range of 0.2 to 0.6 arcsec, while the Sérsic index is randomly selected between 0.5 and 5, which extends lower when compared to the range considered by Petrillo et al. (2017). This is done in order to consider a wider range of source morphologies in our samples. The range of Sérsic indices and effective radii for the source galaxies shows a slight bias toward early-type galaxies (He et al. 2020). The major-axis

position angle is also randomly assigned across the entire range of 0 to 180 deg, ensuring a robust and realistic data set for source modelling. As evident from the mock lens parameter distribution shown in Fig. 3, the objective is not to replicate a statistically accurate representation of the real lens population. Instead, the emphasis is on densely populating the training data set within the considered parameter space. This strategic approach, which has been previously employed by other studies (Petrillo et al. 2019b; Rezaei et al. 2022), empowers developed model architectures to learn various configurations, even those that may be rare or currently unknown in real distributions. This approach enhances the model’s ability to generalize beyond common scenarios.

In total, 10^6 mock lenses are generated, each covering a 20×20 arcsec field, introducing heightened complexity in both source and lens planes. Among these, 800 000 are randomly selected for the training phase, while the validation and test data sets each encompass 100 000 unique lens configurations.

2.2.3 Creating real-looking lens systems

To create realistic lens systems, we employ the method described by Petrillo et al. (2017), which combines a chosen mock lens (detailed in Section 2.2.2) with a potential lensing galaxy (LRGs; as explained in Section 2.2.1). When combining a LRG and mock lensed emission, we adjust their peak brightness using a scaling factor ($0.03 \leq K \leq 0.5$), ensuring the lower magnitudes typically observed in lensing features relative to LRGs are preserved. Specifically, we scaled the brightness of the mock lens to K times the peak brightness of the selected LRG, allowing it to resemble the lensing galaxy in the system. Additional steps to create realistic lens systems include clipping negative pixel values to zero, which eliminates non-physical intensities, and applying a square-root stretch to emphasize low-brightness features, such as extended gravitational arcs. Finally, the images are normalized to a pixel value range of 0 and 1, ensuring uniformity across the data set and optimizing the training process for the CNN.

Previous studies, such as by Petrillo et al. (2017), have often employed a random selection strategy to pair mock lenses and foreground galaxies; however, in our investigation of the influence of training data on CNN performance, we identified potential limitations in this approach. Fig. 4 illustrates an example in which such a random strategy could cause potential problems by generating confusing training samples. Two scenarios are depicted: in one, a mock lens with a small Einstein radius is added to a LRG, resulting in a lens sample where the ring configuration of the lens is entirely hidden by the LRG emission. Conversely, incorporating a mock lens with a larger Einstein radius offers a more suitable match for the same LRG. This observation aligns with the equation $M = \pi\rho\theta^2$, where ρ is the average surface mass density inside the Einstein radius (θ). It is proportional to the square root of the enclosed mass (M). Assuming the enclosed mass correlates with the total integrated surface brightness, the Einstein radius is linked to the total flux of a galaxy.

Avoiding the generation of samples similar to the one on the left side of Fig. 4 enhances quality assurance for the CNN. This sample, which would be labelled as lensed despite presenting no clear lensing emission, can be considered a non-lensed sample. In other words, incorporating nearly identical samples with differing labels into the training data set leads to confusion for the model. This underscores the importance of carefully selecting the pair of foreground galaxy and mock lensed emission, as it significantly

influences the composition of the generated lens population in the training data set.

By avoiding such confusing samples, the model trains more effectively and produces the expected output, translating to a lower false positive (FP) rate. Reducing the FP rate is particularly important in large-scale surveys such as KiDS and *Euclid*, where the volume of data is immense, and manual inspection of each candidate lens is impractical. High FP rates make it challenging to accurately identify true lenses. This not only wastes valuable time and resources, but also introduces uncertainties that can propagate into subsequent analyses, leading to erroneous conclusions about the properties and distribution of galaxies and dark matter.

3 METHOD

In this section, we provide an overview of the architecture employed in our deep learning algorithm designed for the detection and ranking of gravitational lensing candidates. CNNs recognized for their adeptness in processing input imaging data with a topological structure, stand out as the primary approach for object detection and classification. The strength of CNNs lies in their capacity to utilize multiple layers, each serving a distinct function, and the arrangement of these layers can produce various convolutional components. As a consequence, the efficacy of a CNN is closely tied to the specific components implemented, and the performance may vary based on the application’s requirements.

To quantify the performance of our network in terms of dissimilarity between estimated and true class labels, we employ loss functions. In our task of binary classification to distinguish between lensed and non-lensed objects, we have chosen the Binary Cross Entropy (BCE) form of the loss function for our evaluations. The BCE loss function, represented as,

$$\text{BCE} = -\frac{1}{N} \sum_{i=1}^N [y_i \log p_i + (1 - y_i) \log(1 - p_i)], \quad (3)$$

is employed where y_i denotes the given class label for the i th sample in our data set of N training samples and p_i represents the estimated probability of the model indicating the i th sample as a strong gravitational lens system.

While some studies have developed custom architectures for strong lens detection (Rezaei et al. 2022), leveraging existing network structures is a common practice in the field. Among the widely utilized architectures, *ResNet* (He et al. 2016) stands out as one of the most popular choices in the strong lensing literature, as evidenced by studies such as from Petrillo et al. (2017) and Lanusse et al. (2018). The *ResNet* model is built on the concept of training deeper CNNs by incorporating shortcuts or by skipping connections between the front and back layers. This strategy helps in facilitating the backpropagation of gradients during training, allowing for better optimization of the model. In a comparative analysis conducted by Nagam et al. (2023), the performance of *DenseNet* (Huang et al. 2017) in comparison to *ResNet* was assessed. The findings revealed that *DenseNet* achieved comparable true positive (TP) rates while exhibiting lower FP rates. The *DenseNet* model builds upon the skipped connections concept, but introduces dense connections between all previous and subsequent layers.

This unique characteristic allows *DenseNet* to achieve superior performance compared to *ResNet*, all while requiring fewer parameters and incurring less computational cost. Motivated by those results, we have opted for *DenseNet* in our further analysis.




Figure 3. The distribution of the lens model parameters used for the training data set; these are (left panel) the lens axis ratio (b/a), (middle panel) the lens external shear (γ_{ext}), and (right panel) the lens Einstein radius (θ_E). Notably, the Einstein radii follow a logarithmic distribution, while the other parameters adhere to a flat distribution. The position angles of the ellipsoidal mass distribution and the external shear were set randomly between ± 90 deg.




Figure 4. An example of a potential issue arising when the lensed emission is faint with respect to the brightness of the foreground lensing galaxy and has an Einstein radius that is significantly lower than the effective radius of the foreground lensing galaxy. Two scenarios employing the same LRG as the foreground galaxy are shown, but with different lens configurations. As depicted, when lensed emission with a smaller Einstein radius is introduced to the selected LRG, the ring configuration of the lensed emission is entirely obscured by the LRG emission. The inclusion of lensed samples, such as the example on the left (labelled as 1 or lensed), may confuse the CNN model, as this sample resembles a non-lensed sample (labelled as 0) in our training data set.

Specifically, we explore multiple variants, including *DenseNet-121* and *DenseNet-169*.

Furthermore, we investigate various architectures based on the understanding that the selection of model architecture plays a crucial role in determining the overall performance in tasks such as strong lens detection. Taking this direction further, we consider another branch of CNN architectures, called *EfficientNet* (Tan & Le 2019), which achieve state-of-the-art performance on image classification tasks while also being computationally efficient. In particular, we investigate various versions, such as *EfficientNet-B3* and *EfficientNet-B4*. This comprehensive analysis aims to uncover insights into how different neural network architectures influence the accuracy and reliability of strong lens detection models. In the following, we provide an overview of these selected architectures.

3.1 DenseNet

DenseNet, short for Densely Connected Convolutional Networks, is a type of neural network architecture that emphasizes dense connectivity within ‘dense blocks’. In each dense block, every layer receives the feature maps generated by all preceding layers as input, while also passing on its own feature maps to every subsequent layer in the block. This dense connectivity structure is distinct from traditional CNNs, where each layer is only connected to the immediately following layer. By establishing direct connections between all layers in a block, *DenseNet* allows each layer to directly access the features of all previous layers, promoting both efficient information flow and rich feature representation. *DenseNet* comes in various versions, such as *DenseNet-121*, *DenseNet-169*, and *DenseNet-201*, where the numbers represent the total number of layers in each network. The choice of model variant depends on the complexity of the task and the available computational resources. For tasks that demand a balance between model complexity and computational efficiency, *DenseNet-121* and *DenseNet-169* are widely adopted due to their favourable trade-off between performance and resource consumption. We refer the interested reader to the review by Huang et al. (2017) for a detailed discussion on the architecture of *DenseNet*.

3.2 EfficientNet

EfficientNet is a family of CNNs with the key innovation of compound scaling, which enables an optimal trade-off between model size and performance by uniformly scaling the network’s depth, width, and resolution. *EfficientNet* has been widely used in a variety of scientific applications, such as galaxy morphology classification (Kalvankar, Pandit & Parwate 2020), spectral classification of astronomical objects (Wu et al. 2023), skin cancer detection (Venugopal et al. 2023), brain tumor detection (Curci & Esposito 2024), and lung cancer detection (Raza et al. 2023).

The key innovation of *EfficientNet* lies in the balance between model depth, width, and resolution, as governed by a compound scaling method. This approach ensures that the network scales efficiently across these dimensions, making it well suited for diverse tasks. Each variant of *EfficientNet* is denoted by a scaling factor (e.g. *EfficientNet-B3*, *EfficientNet-B4*), reflecting its capacity for increased depth and width. These scaling factors allow users to choose a model that aligns with the specific requirements of their task and computational resources. This concept can be mathematically expressed as,

$$\alpha \cdot \beta^2 \cdot \gamma^2 \approx 2 \quad (4)$$

for $\alpha \geq 1$, $\beta \geq 1$, and $\gamma \geq 1$. Here, the depth is α^ϕ , the width is β^ϕ and the resolution is γ^ϕ , where ϕ denotes the scaling coefficient

that uniformly scales the network. The choice of scaling coefficients impacts the trade-off between model complexity and computational efficiency, making *EfficientNet* a versatile architecture that is adaptable to different resource constraints and task requirements. Further details on the structure of *EfficientNet* are given by Tan & Le (2019).

In our investigation, we incorporate *EfficientNet-B3* and *EfficientNet-B4*, considering their widespread adoption and ability to strike a suitable balance between model complexity and computational efficiency in various applications. The initial weights were randomly set using a uniform distribution. While different weight initializations can impact the training process during the early epochs, our experiments show that the network converges shortly after this period. Thus, the observed performance remains largely unaffected by the initial weight settings.

3.3 Evaluation criteria

To assess and compare the effectiveness of different methodologies and training data strategies on the test data set, it is essential to establish appropriate evaluation criteria. Given that lens detection is treated as a classification problem, where samples are categorized as lensed or non-lensed, the following representation can be utilized for evaluation purposes:

		True data	
		Lens	Not lens
Test results	Lens	TP	FP
	Not lens	FN	TN

In this representation, TP indicates correctly identified gravitational lens systems, True Negative (TN) corresponds to accurately recognized non-lensed sources, FP denotes mis-classification of non-lensed sources as gravitational lenses, and False Negative (FN) refers to gravitational lensing events missed by the algorithm and classified as non-lensed sources.

Based on these terms, several evaluation criteria can be defined. Accuracy measures the proportion of correctly identified samples (TP and TN) out of the total number of samples. Precision quantifies the ratio of TPs to the sum of TPs and FPs, indicating the reliability of positive predictions. Recall assesses the fraction of TPs correctly identified by the algorithm, reflecting the model's completeness. Fall-out, also known as the FP rate, calculates the proportion of negative samples incorrectly classified as positive, providing insight into the purity of detected lens candidates. Mathematically, these metrics are represented as:

$$\begin{aligned}
 \text{Accuracy} &= \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FN} + \text{TN} + \text{FP}}, \\
 \text{Precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}}, \\
 \text{Recall} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\
 \text{Fall-out or FP rate} &= \frac{\text{FP}}{\text{FP} + \text{TN}}. \tag{5}
 \end{aligned}$$

In assessing gravitational lens search algorithms, the emphasis extends beyond completeness alone, especially given the anticipation of a large number of gravitational lenses to be detected with upcoming all-sky surveys. Rather, the focus often centres on achieving a low FP rate, with the aim of identifying a high number of genuine lens candidates in the ranked list. Additionally, the Receiver Operating Characteristic (ROC) curve serves as another valuable metric. The ROC plot visually represents the trade-off between the TP rate and

the FP rate for each model. Each point on the ROC curve signifies a different threshold for classifying samples as positive or negative based on their predicted probabilities. By examining the ROC plot, we can discern how well each model discriminates between positive (lensed) and negative (non-lensed) samples. A model with better performance will exhibit a curve that closely approaches the top-left corner of the plot, indicating higher TP rate and lower FP rate across various threshold values. The comparison of ROC curves for different models provides insights into their relative effectiveness in identifying lensed samples. This analysis aids in selecting the most suitable model for the task at hand, considering both sensitivity to TPs and robustness against FPs.

4 RESULTS

The success of any machine learning model relies on the quality of the training data set. This section examines various properties of training data sets, focusing on their impact in strong lens detection projects. The relevance and representation of the training data set are crucial. A well-curated data set must encompass a diverse and representative set of examples that reflect the variety and complexity of real-world data. In the context of strong lens detection, this means including images with different types of strong lensing phenomena, as well as a high variety in the non-lensed samples. Also, a representative data set helps the model learn to generalize from the training data to unseen data, improving its robustness and accuracy. Another important factor is accurate labelling as the model relies on these labels to learn the correct associations between input data and the desired output. An example of the importance of accurate labelling in the context of strong lens detection is provided in Fig. 4.

Considering the significant impact of training data on the performance of CNNs, we have provided two novel strategies to handle the training data set, complementing the conventional approach typically adopted in the literature (e.g. Petrillo et al. 2019b; Nagam et al. 2023), which we refer to as the Vanilla strategy. Our aim is to craft a data set that enhances the purity of detected candidates. While completeness is still a consideration, in this study, we prioritize the necessity of mitigating FPs.

Our first approach, termed 'Applied1', prioritizes the treatment of non-lensed samples within the training data set. This strategy is tailored to address specific challenges associated with the characterization of non-lensed objects. Conversely, our second approach, 'Applied2', focuses on optimizing the representation of the lens population within the training data set, thereby enhancing the CNN's ability to accurately identify and classify gravitational lensing events. Further elaboration on these strategies is provided below, including their respective methodologies and rationales for effectively training CNNs in gravitational lens detection tasks.

Although we have tested CNN architectures on various training data sets, we ensured that each round incorporated the same number of samples to provide a fair analysis of model performance. Each CNN architecture was trained on a total of 500 000 samples. This training data set was balanced with 250 000 samples labelled as lensed and 250 000 as non-lensed.

4.1 Vanilla

The Vanilla setting, pioneered by Petrillo et al. (2017) and recently utilized by Nagam et al. (2023) with an expanded repertoire of non-lensed samples, stands as a benchmark methodology for gravitational lens detection tasks using KiDS data. As described above, the generation of lensed samples entails a random pairing of a lens and




Figure 5. The ROC plot demonstrates how different machine learning models, trained with the Vanilla setting, perform in distinguishing between positive (lensed) and negative (non-lensed) samples across various thresholds balancing TP and FP rates. For presentation purposes, the range of the ROC plot has changed from $[0, 1]$ to the current display. By comparing these curves, we can identify the most efficient model for detecting lensed samples, while minimizing FPs. The plotted results show the improvement in FP rate as we use an ensemble technique, by averaging the predicted lens probability of individual models. The best performing ensemble belongs to averaging the output of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-121* with the FP rate of 1.1×10^{-3} and a TP rate of 0.906.

an LRG sample to create a realistic-looking strong lensing sample. Within this framework, non-lensed samples are predominantly drawn from a pool of observed spiral (contaminations) and elliptical galaxies. These samples, originating from authentic KiDS observations, undergo augmentation to mitigate overfitting and enhance model performance. While Nagam et al. (2023) maintains a ratio of 20 per cent LRGs to 80 per cent contamination samples in their non-lensed selection process, we have opted for a balanced approach, incorporating an equal number of LRGs and contaminations in our data set. This deliberate choice aims to provide the CNNs with a more comprehensive representation of LRGs, thereby strengthening their capacity to differentiate between LRGs acting as foreground galaxies in gravitational lens systems and those exhibiting typical LRG characteristics devoid of lensing effects.

Fig. 5 shows the performance comparison of the four different models: *DenseNet-121*, *DenseNet-169*, *EfficientNet-B3*, and *EfficientNet-B4*. Since the model output represents the probability of an object being a lens, ranging between 0 and 1, setting a threshold is necessary to classify objects into two classes: lensed and non-lensed. To ensure a stringent selection of strong lenses, a threshold of 0.99 is set for further analysis. This decision is driven by our primary goal of reducing the FP rate. As illustrated in Fig. 5, increasing the threshold does not significantly impact the TP rate. In other words, it is acceptable to sacrifice a small percentage of the TP rate to achieve a better FP rate.

Building upon findings of Rezaei et al. (2022) and Andika et al. (2023), it has been observed that averaging the output probabilities of multiple models can effectively mitigate FP rates. This improvement stems from the models' tendency to disagree on suspicious cases, thus allowing for a voting scheme to be applied. As depicted in Fig. 5, averaging the model outputs indeed yields enhanced results.

Table 2 provides a detailed overview of each method's performance when the threshold is set at 0.99. Notably, the average of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-121* demonstrates

Table 2. The main evaluation metrics, such as TP and FP rates, derived from training the featured models, using the Vanilla setting on the training data set. The predicted lensing probability from each model is a value between 0 and 1. However, for calculating the evaluation metrics, we have established a threshold of 0.99 to differentiate between lensed and non-lensed samples. The test data set consists of 96 072 samples equally distributed between lensed and non-lensed categories.

Model	TP	FP
DenseNet-121	0.925	5.3×10^{-3}
DenseNet-169	0.925	6.4×10^{-3}
EfficientNet-B3	0.929	3.6×10^{-3}
EfficientNet-B4	0.937	7.1×10^{-3}
DenseNet-121, EfficientNet-B3	0.908	1.3×10^{-3}
DenseNet-169, EfficientNet-B3	0.910	1.8×10^{-3}
DenseNet-121, EfficientNet-B4	0.916	2.7×10^{-3}
DenseNet-169, EfficientNet-B4	0.915	3.1×10^{-3}
EfficientNet-B3, EfficientNet-B4	0.921	2.1×10^{-3}
DenseNet-121, DenseNet-169	0.910	2.7×10^{-3}
DenseNet-121, EfficientNet-B3, EfficientNet-B4	0.906	1.1×10^{-3}
DenseNet-169, EfficientNet-B3, EfficientNet-B4	0.906	1.2×10^{-3}

the most promising outcomes, with a FP rate of 1.1×10^{-3} and an acceptable TP rate of 0.90. Following closely, the average of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-169* achieves a similar TP rate and a slightly higher FP rate of 1.2×10^{-3} . These results underscore the effectiveness of leveraging ensemble methods to enhance the performance of gravitational lens detection algorithms.

In order to comprehensively understand the model behaviour and analyse its performance, we conducted a thorough investigation into the samples contributing to the FP rate of 1.1×10^{-3} by our best model. Several examples are provided in Fig. 6 and their properties are given in Table 3. Our analysis revealed that a significant portion of these FPs stem from the misclassification of extended LRGs




Figure 6. A selection of LRGs falsely detected as strong lens systems by the Vanilla setting. These samples (FPs) are having a high complexity value of 200 or more, which are not well represented in the Vanilla setting. More details regarding these samples are provided in Table 3.

Table 3. Properties of FPs detected by the best performing model using the Vanilla setting. The indices correspond to the images depicted in Fig. 6. The lens probability and uncertainty metrics were derived from the averaged predictions of the lensing probability across the *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-121* models. These findings indicate that all three models consistently predict a probability exceeding 99 per cent for this selection of LRGs, even though they are non-lensed samples.

Index	Lens probability	Complexity	Effective radii
a	0.997	271	21.28
b	0.999	282	24.73
c	0.995	283	32.41
d	0.999	297	21.17
e	0.999	285	25.38
f	0.997	400	50.23

as strong gravitational lensed samples. This observation, coupled with the complexity distribution shown in the lower panel of Fig. 2, indicates that our training data for LRG samples lacks a sufficient number of instances representing LRGs with extended emission or those with nearby contaminants. This realization prompted us to introduce the first variation of this study, which is detailed in the next sub-section.

This investigation highlights the importance of ensuring the diversity and representativeness of training data not only in lensed samples, but also in non-lensed samples. By addressing the imbalance in the representation of extended LRGs in our training data set, we aim to enhance the robustness of our model in distinguishing between genuine gravitational lenses and other astronomical objects.

4.2 Applied1

We now demonstrate how the properties of the training data set can significantly influence the achieved results and the overall performance of the model. This is done by changing the characteristics and sample population within the training data set, while maintaining consistency in the test data set. Drawing insight from our analysis of the FP samples presented in Section 4.1, here we focus on improving the representation of the training data set for non-lensed samples. By incorporating this solution, we have achieved compelling results, which we detail below.




Figure 7. A comparison of the complexity between observational KiDS LRGs data provided by Li et al. (2020) and the LRG samples within our non-lensed population from the training data set. It highlights the adjustments made to the LRG population, achieved through strategic augmentation of LRG samples. These modifications aim to enhance the visibility of sources with less representative samples, particularly those with complexity values exceeding 100. By augmenting these samples, the CNN is exposed more to the complex LRGs and, thereby enhancing its ability in identifying and analysing LRGs as non-lensed samples.

Fig. 7 shows the distribution of complexity values among the LRG samples within our training data set. Notably, this distribution exhibits non-uniformity, which is particularly evident in the complexity range of 120 and higher, where less than 1.4 per cent of the LRG samples are represented. To validate these findings, we compare them with the complexity distribution of a larger LRG data set produced by Li et al. (2020). This data set comprises low-redshift samples ($z \leq 0.4$) selected based on their r -band magnitude, limited to values below 20. The complexity distribution of these LRG samples is also shown in Fig. 7. Remarkably, the comparison reveals that the distribution of LRG complexity in our training data set does not accurately reflect the distribution observed in real observational data.

Even when augmentation techniques are applied to the current LRG samples, the resulting distribution tends to mirror the original distribution, which worsens the imbalance between compactness categories. Therefore, augmentation alone does not offer a viable solution to this issue. To address this challenge, we propose a strategic augmentation approach aimed at achieving a more balanced distribution of samples across various complexity bins. Specifically, we advocate for augmenting the less frequent sources with greater variations compared to the more densely populated regions in the complexity distribution. This strategic augmentation aims to ensure a semi-uniform distribution of samples across various complexity categories, thereby enhancing the representativeness of our training data set. The distribution of the generated LRG population with this strategy is also shown in Fig. 7.

Following this, we adapted our training data set to align with the adjusted LRG population. However, we have kept other training parameters, such as the total number of training data points, the distribution of lensed samples, the learning rate, and other hyperparameters consistent with the Vanilla setting. The results obtained from this training data configuration, referred to as ‘Applied1’, are presented in Fig. 8. It is evident that the best achieved FP rate comes from the ensemble of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-169*, with a value of 3.5×10^{-4} , marking a notable decrease from the FP rate of 1.1×10^{-3} observed for the Vanilla setting. With a total of 48 036 non-lensed samples in our test data set, the number of FP detections has decreased from 53 to 17, which




Figure 8. The ROC plot comparing the performance of trained CNN models under the Applied1 scenario in distinguishing between lensed and non-lensed samples. Given the significance of minimizing FPs, the most effective model is one that can provide a purer selection of candidates. Consequently, the ensemble comprising *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-169*, with an FP rate of 3.5×10^{-3} and a TP rate of 0.906, emerges as the top-performing model.

Table 4. The primary evaluation metrics obtained from training the CNN models using the Applied1 setting on the training data set. The models predict lensing probability values ranging between 0 and 1. However, to calculate the evaluation metrics, we have set a threshold of 0.99 to distinguish between lensed and non-lensed samples. The test data set comprises 96 072 samples evenly split between lensed and non-lensed categories.

Model	TP	FP
DenseNet121	0.899	1.4×10^{-3}
DenseNet169	0.914	2.9×10^{-3}
EfficientNetB3	0.946	2.2×10^{-3}
EfficientNetB4	0.951	2.4×10^{-3}
DenseNet121, EfficientNetB3	0.898	6.8×10^{-4}
DenseNet169, EfficientNetB3	0.900	5.4×10^{-4}
DenseNet121, EfficientNetB4	0.898	6.2×10^{-4}
DenseNet169, EfficientNetB4	0.910	4.9×10^{-4}
EfficientNetB3, EfficientNetB4	0.938	9.5×10^{-4}
DenseNet121, DenseNet169	0.886	6.2×10^{-4}
DenseNet121, EfficientNetB3, EfficientNetB4	0.897	4.7×10^{-4}
DenseNet169, EfficientNetB3, EfficientNetB4	0.906	3.5×10^{-4}

represents more than a threefold reduction. This effect will likely be further magnified when dealing with a larger test data set, such as the 126 884 LRGs from the KiDS DR4, as described by Li et al. (2020).

Moreover, this reduction in the number of FPs has not compromised the TP rate. The specific details of the FP and TP rates are presented in Table 4. In comparison to the Vanilla setting, where the ensemble of *DenseNet-121*, *EfficientNet-B3*, and *EfficientNet-B4* yielded the best results, in this scenario, *DenseNet-169* has replaced *DenseNet-121* in the ensemble. The combination of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-169* achieves a TP rate that is 0.9 percent better, while exhibiting a slightly higher FP rate of 1.2×10^{-4} , when compared to the ensemble of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-121*.

Table 5. The lens probability of the FPs detected using the Vanilla setting that were not detected as FPs using the Applied1 scenario. Alongside the ensemble predicted lens probability, the individual lens probabilities for *EfficientNet-B3* (ENet-B3), *EfficientNet-B4* (ENet-B4), and *DenseNet-169* (DNet-169) are provided. All of these predictions are based on the Applied1 setting. A visual representation of these samples can be found in Fig. 6, while additional details regarding the Vanilla setting predictions are presented in Table 3.

Index	Ensemble	ENet-B3	ENet-B4	DNet-169
a	0.65	0.950	0.997	3.3×10^{-6}
b	0.96	0.999	1	0.89
c	0.64	0.950	0.98	1.5×10^{-3} ,
d	0.98	0.999	1	0.95
e	0.67	0.999	0.999	1.1×10^{-2} ,
f	0.70	0.991	0.998	0.12

Table 5 offers some insight to the variability of individual model predictions and their influence on the ensemble probability. Also, the recorded samples are shown in Fig. 6 and their predicted lensing probabilities under the Vanilla setting are given in Table 3. A comparison between Tables 3 and 5 highlights the better performance of the training strategy implemented in Applied1, when compared to the results using the Vanilla setting.

4.3 Applied2

We now introduce our second modification to the training data set. While the previous section addressed issues concerning the non-lensed population, the focus here is on potential challenges within the lensed population. As discussed previously in Fig. 4, it is critical to consider the morphology of the foreground galaxy when pairing it with a mock lens. Our adjustments target how the lens population is constructed within the training data set. As outlined in Section 2, this process involves pairing mock lenses with foreground galaxies to simulate realistic strong gravitational lens systems. Through visual

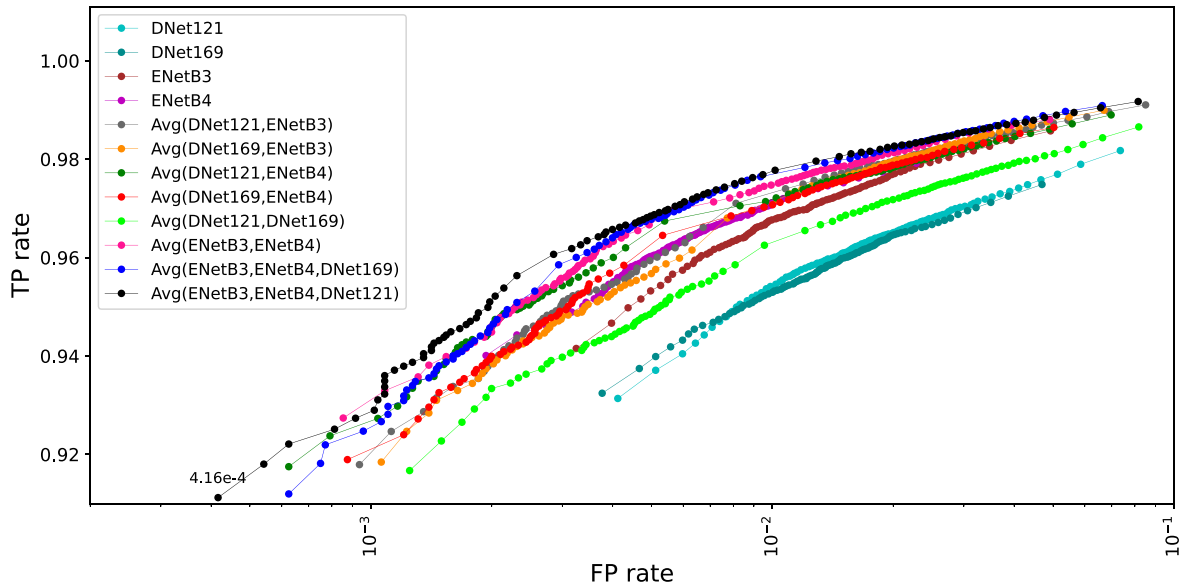


Figure 9. The ROC plot comparing the performance of trained CNN models under the Applied1 scenario in distinguishing between lensed and non-lensed samples. In this scenario, both the lensed and non-lensed populations have been altered compared to the Vanilla setting. The ensemble of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-121* achieves a TP rate of 0.9112 with a FP rate of 4.16×10^{-4} . The achieved TP rate surpasses that of both the Vanilla and Applied1 settings.

Table 6. A comparison of the TP and FP rates for different CNN architectures, when the training data follows the Applied2 scenario. The best results are obtained using an ensemble of *DenseNet-121*, *EfficientNet-B3*, and *EfficientNet-B4* with a TP rate of 0.9112 and a FP rate of 4.16×10^{-4} .

Model	TP	FP
DenseNet-121	0.931	4.1×10^{-3}
DenseNet-169	0.932	3.7×10^{-3}
EfficientNet-B3	0.941	3.2×10^{-3}
EfficientNet-B4	0.940	1.9×10^{-3}
DenseNet-121, EfficientNet-B3	0.917	9.3×10^{-4}
DenseNet-169, EfficientNet-B3	0.918	1.1×10^{-3}
DenseNet-121, EfficientNet-B4	0.817	6.2×10^{-4}
DenseNet-169, EfficientNet-B4	0.918	8.7×10^{-4}
EfficientNet-B3, EfficientNet-B4	0.927	8.5×10^{-4}
DenseNet-121, DenseNet-169	0.916	1.2×10^{-3}
DenseNet-121, EfficientNet-B3, EfficientNet-B4	0.911	4.2×10^{-4}
DenseNet-169, EfficientNet-B3, EfficientNet-B4	0.919	6.2×10^{-4}

examination, we identified potential issues, particularly with small-separation lensed emission ($\theta_E < 0.85$ arcsec). An analysis of the mock lens parameters (see the right panel of Fig. 3) reveals that a considerable portion of existing samples in the training data set belong to the small Einstein radii population, with approximately 23 per cent having Einstein radii below 0.85 arcsec.

Our primary focus is on selecting appropriate lensing galaxies, specifically tailored for these smaller Einstein radii lenses. An example illustrating the potential issue is presented in Fig. 4, where an unsuitable choice of lensing galaxy results in a perplexing sample that lacks distinct lensed emission. Given the considerable likelihood (approximately 23 per cent) of such mock lenses being incorporated into the training data set, it is crucial to address this challenge. Therefore, we propose selecting only a subset of LRGs as potential foreground galaxies for small Einstein radii lenses. This approach ensures that the distribution of source parameters remains consistent

with the previous analysis. The variable aspect here is the selection process for the corresponding lensing galaxy associated with the mock lensed emission that have smaller Einstein radii. In Section 2, we discuss the assumption that the Einstein radius (θ_E) is proportional to the total integrated flux of a foreground galaxy. Here, we use the LRGs’ effective radii as a proxy for their integrated flux. Under the ‘Applied2’ setting, the training data is modified to include only LRGs with effective radii below 0.5 arcsec, paired with mock lensed emission with Einstein radii below 0.85 arcsec. However, it is important to note that the same test data set has been utilized here as in the Vanilla and Applied1 settings.

Fig. 9 and Table 6 show the results obtained from employing this strategy. We see that the most optimal performance is achieved by the ensemble of *EfficientNet-B3*, *EfficientNet-B4*, and *DenseNet-121*, with a FP rate of 4.16×10^{-4} when a detection threshold of 0.99 is used. This translates to the detection of 20 FPs within the 48,036 total non-lensed samples in our data set, which is three more than the number of FPs detected using the Applied1 setting (see Fig. 8 and Table 4). Interestingly, this variation in the training data set, by solely modifying the lens class of the training data, has influenced the number of FP detections by the model. This indicates that the challenges of strong lensing detection are complex, relying on multiple parameters. Another notable point is the 222 additional TP samples achieved using the Applied2 setting, when compared to the Applied1 and Vanilla settings. This corresponds to approximately a 0.46 per cent improvement in TPs. This presents a trade-off between selecting a model with a higher TP rate or a lower FP rate, considering that both models significantly outperform the Vanilla setting in terms of the FP rate.

4.4 Combined model

Our findings have thus far highlighted the impact of modifications in the training data set on the model’s performance. We have observed that adjustments to either the lensed or non-lensed classes of data can




Figure 10. The 8 FP samples within the test data set of 48 036 non-lensed samples, when considering the prediction of the Applied1 and Applied2 settings. All of the samples belong to the LRG population, which indicates that the models have learned to distinguish between spiral galaxies and lensed samples. The remaining issue is in separating the FP samples belonging to non-lensed LRGs with those that exhibit lensed emission.

influence the TP and FP rates, thus affecting the overall quality and reliability of the results obtained. A comparison between Tables 4 and 6 reveals that the Applied2 setting exhibits a higher FP rate, resulting in three more FP samples compared to the Applied1 setting, within our non-lensed test data set of 48 036 samples. Upon examination, it becomes apparent that the two models only agree on eight FP samples, which are shown in Fig. 10. Among these, the predicted lensing probability of only four samples exceeds 0.99 for the Vanilla setting, indicating that if a voting strategy were employed across all models, we would identify only four FP samples within the entire non-lensed sample; these are labelled as (b), (c), (f), and (h) in Fig. 10. However, incorporating a strategy that averages out the predicted lensing probability for the Vanilla, Applied1 and Applied2 settings results in five FP samples, which now includes sample (a) in Fig. 10. This intriguing result also underscores how

Table 7. A comparison of TP and FP rates for different training data set settings (Vanilla, Applied1, and Applied2), alongside the combined model. These results highlight the superior performance of the combined model in detecting fewer FPs compared to each of the investigated training settings. The selected threshold is 0.99.

Model	TP	FP
Vanilla	0.906	1.1×10^{-3}
Applied1	0.906	3.5×10^{-4}
Applied2	0.911	4.2×10^{-4}
Combined	0.883	1.0×10^{-4}

different ensemble approaches can impact the actual number of FPs encountered. Another interesting finding is that all of the FPs shown in Fig. 10 are LRGs, which do not exhibit any spiral emission. This suggests that the primary challenge in reducing the number of FPs from the KiDS data lies not in spiral emission, but rather in distinguishing between the LRG population and contaminants that may mislead the model into interpreting them as lensed emission.

Fig. 11 presents a comparison of different settings in the training data set, namely the Vanilla, Applied1 and Applied2 settings. As previously discussed, averaging through all of these predictions yields the combined lens probability prediction, which demonstrates a better FP rate, as the models may not agree on identifying challenging samples as lensed. A more detailed view of these results is provided in Table 7, which shows the efficacy of ensemble techniques. According to our results, the combined model exhibits a FP rate of 10^{-4} , representing an 11-fold decrease in the number of FPs compared to the Vanilla setting, a 3.5-fold improvement compared to Applied1, and a 4.1-fold improvement compared to Applied2. This significant reduction in FPs comes at the cost of a 2.3 per cent decrease in TPs when compared to the Vanilla and Applied1 settings, while compared to the Applied2 setting, the TP rate has decreased




Figure 11. The ROC plot comparing the performance of each training data set setting (Vanilla, Applied1, and Applied2). This illustrates the dependency of the TP and FP rates on the chosen threshold that separates lensed and non-lensed samples. The combined model incorporates the average lensing probability of all three training settings, demonstrating an improved FP rate, albeit with a slight reduction of the TP rate by a few per cent.

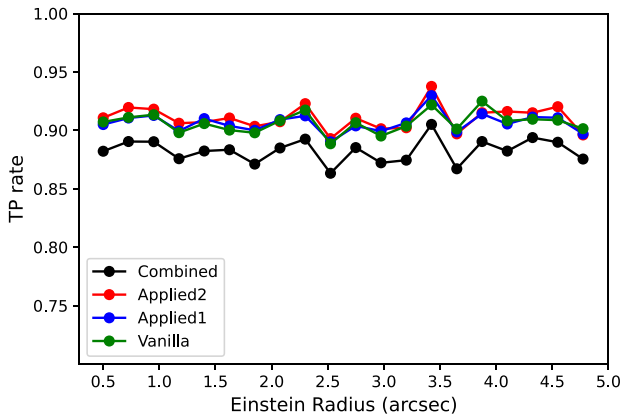


Figure 12. The distribution of TP rates as a function of the Einstein radius of the mock lenses. This reveals that the effectiveness of the detection method does not seem to be directly influenced by the size of the Einstein radius of the mock lenses. As expected from the results presented in Table 7 and Fig. 11, the Combined model has the lowest TP rate, when compared to the individual models.

by 2.8 per cent. In the following, we provide details on how these TP and FP rates are related to the population of lensed and non-lensed samples and their underlying properties.

4.5 Parameter-space analysis

Utilizing the results obtained from Fig. 11 and Table 7, we now investigate the parameter-space for detection, specifically examining how the Einstein radius of a lensed object or the complexity of the LRG influences the TP and FP rates. Through these experiments, our aim is to understand how different training data sets affect the types of lenses to which the trained CNN architectures are sensitive, as well as the parameters that may impact the model’s ability to accurately label samples in the test data set.

Fig. 12 shows the TP rate as a function of the lens Einstein radius. This result indicates that although the Applied2 setting shows a slightly higher TP rate when compared to the Vanilla and Applied1 settings, this difference is consistent across all Einstein radius bins, as opposed to belonging to any specific range. On the other hand, when examining the behaviour of the combined model, it becomes evident that this model exhibits a lower TP rate, when compared to the individual models. This outcome is expected because in the combined model, a sample is labelled as a lens if all three settings of Vanilla, Applied1, and Applied2 agree. This ensemble technique, achieved by averaging the predicted lens probabilities of each setting, results in a smaller set of final candidates, as is also demonstrated in Table 7.

Another noteworthy comparison lies in examining the FP rate alongside the complexity of the foreground lens galaxies. As previously demonstrated, the adjustment made in the training data set for Applied1 aims to balance the non-lensed population, ensuring that extended, complex sources are adequately represented in that class of data. A comparison between the Vanilla, Applied1, and Applied2 settings clearly highlights the impact of such adjustments on the achieved results. Fig. 13 illustrates a significant difference in the FP rate obtained by the Vanilla, Applied1 and Applied2 settings. An intriguing observation from Fig. 13 is the disparity between the FP rates of the Applied1 and Applied2 settings, despite both utilizing the exact same non-lensed samples. The sole difference

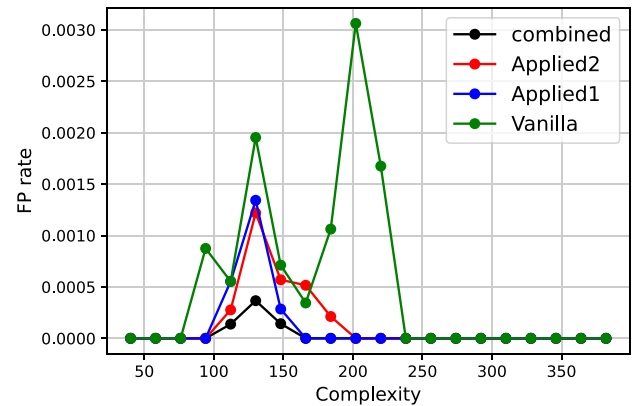


Figure 13. The FP rate as a function of the LRG compactness employed within the non-lensed population. This illustrates the influence of the distribution of the implemented training data set on the behaviour of the Vanilla, Applied1, and Applied2 settings.

between these two settings lies in the distribution of lensed samples. This observation underscores the complexity of the lens detection problem, where the behaviour of the model can be influenced by numerous parameters that may initially seem unrelated.

4.6 Evaluation on real KiDS data

To evaluate the performance of our methodology on real KiDS data, we have applied the Vanilla, Applied1, Applied2, and Combined strategies to 126 000 LRGs from the KiDS DR4 (Li et al. 2020). The predictions for each lensing probability bin are presented in Fig. 14. We find that the results align well with our previous analyses, such as those shown in Fig. 11. The Vanilla setting predicts the highest number of lens candidates for the probability range of [0.9–1], whereas the Applied2, Applied1, and Combined settings predict fewer samples with a high lensing probability. When using a threshold of 0.99 to identify potential lenses, the Combined strategy identifies 347 samples, whereas the Vanilla setting identifies 997 samples, indicating a threefold reduction in the number of lens candidates. This reduction is advantageous as it minimizes the need for expert visual inspection and reduces the time required for follow-up observations to confirm the strong gravitational lensing nature of these candidates. The Applied1 and Applied2 settings detect 551 and 710 samples, respectively, with a lensing probability higher than 0.99. It is crucial to assess how many genuine lensed samples are being missed among the detected candidates and whether some lenses are being overlooked in this selection process. This aspect will be addressed in our future work, where we will verify the predicted lens candidates through visual inspection.

5 CONCLUSIONS

In this study, we have investigated the complicated task of detecting gravitational lens systems through analysing the intricate relationship between the composition of the training data set and the performance of the detection models. Through a comprehensive analysis and experimentation, we uncovered vital insights that shed light on the multifaceted challenges and opportunities within this domain. Our findings underscored the importance of data diversity and representativeness, revealing how variations in the sample populations can




Figure 14. A comparison of the predicted lens probability for 126,000 KiDS LRG samples, as a real test data set. As expected, the Combined model has the lowest number of predicted lens candidates, when compared to the other scenarios that have been tested in this study.

cause significant influence on the behaviour and efficacy of the detection models. Our study highlighted the importance of understanding the underlying reasons for model output, beyond merely assessing its performance. This iterative process often necessitates a return to the foundational step of data collection and analysis. For instance, our research revealed that the underlying distribution of LRGs in our non-lensed sample affected the number of FPs. By continuously refining the data collection methods, reassessing the data set properties, and fine-tuning the model architectures, based on the insights gleaned from the model behaviours, we iteratively enhanced the accuracy and robustness of the detection models.

One pivotal discovery from our research revolves around the critical need to address imbalances within the training data set. In particular, distinguishing between extended, complex LRGs and genuine gravitational lenses posed a significant challenge. We noticed that the original test data set, which has been widely used in the literature, has an un-balanced population of LRGs as non-lensed samples in terms of complexity. The number of complex LRGs are significantly lower in the original data set, when compared to the compact ones. By strategically mitigating these imbalances, through techniques such as data augmentation and ensemble learning approaches, we were able to achieve notable reductions in the number of FPs, which enhanced the overall reliability of our detection model.

Beside the population of non-lensed samples in the training data set, we also made modifications that focused on the lensed population. The adjustments changed how the lensed population is constructed, particularly concerning lenses with Einstein radii below 0.85 arcsec. We proposed to only use a subset of LRGs with a small effective radius, to act as potential foreground galaxies for these small Einstein radii lenses. The effectiveness of this modification was evaluated through experimentation, which showed an improved performance in terms of the TP and FP rates, when compared to the standard Vanilla setting.

Our examination of the FP samples obtained from the Applied1 and Applied2 settings revealed that within our test data set of 48 036 non-lensed samples, Applied1 incorrectly labelled 17 samples as strong lenses, while Applied2 identified 20 samples as strong lenses. Further analysis showed that among those samples, Applied1 and Applied2 shared only eight common FP samples. Interestingly, when incorporating the lensing probability predicted

by the Vanilla setting into the ensemble average, only five FP samples remained, all of which were found to be LRGs without any evidence of spiral emission. This underscored the effectiveness of ensemble methods and highlighted the challenge of distinguishing LRG populations from potential contaminants in reducing the FP rate.

The FP rate achieved when we used the ensemble method of the Vanilla, Applied1 and Applied2 (Combined) setting was found to be 10^{-4} , which represented an 11-fold improvement, when compared to the Vanilla setting. While this reduction in FPs came with a 2.3 per cent decrease in TPs, when compared to the Vanilla and Applied1 settings, it signified an advancement considering our primary goal of minimizing FP rates in strong gravitational lens detection algorithms.

We then evaluated our methodology using real KiDS data by applying the various strategies (Vanilla, Applied1, Applied2, and Combined) to a data set of 126 000 LRGs. The results showed that different strategies yielded varying numbers of high-probability lens candidates, with the Combined strategy significantly reducing the number of candidates, when compared to the Vanilla setting. This reduction is beneficial for minimizing the effort required for expert visual inspection and follow-up observations. However, it remains essential to investigate how many genuine lenses may be missed and whether some lenses are overlooked from this process. In a future work, we will focus on addressing these concerns through detailed visual inspection of the predicted lens candidates.

In conclusion, our research contributes significantly to advancing the field of gravitational lens detection by offering insights into the interplay between the training data set, model performance, and the iterative refinement process. By emphasizing the importance of data diversity, imbalance mitigation, and continuous refinement through iterative analysis, our study provides a road-map for developing accurate and reliable detection models that are capable of unraveling the mysteries of the Universe’s most intriguing phenomena.

ACKNOWLEDGEMENTS

This work was performed using the compute resources from the Academic Leiden Interdisciplinary Cluster Environment (ALICE)

provided by Leiden University. AC was supported by the MUR PRIN2022 project 20222JBKCN with title ‘LaScaLa’ – funded by the European Union – NextGenerationEU. This work is based on the research supported in part by the National Research Foundation of South Africa (grant number: 128943).

DATA AVAILABILITY

Upon reasonable request, the underlying data used for this article will be shared by the corresponding author.

REFERENCES

- Amante M. H., Magaña J., Motta V., García-Aspeitia M. A., Verdugo T., 2020, *MNRAS*, 498, 6013
- Andika I. T. et al., 2023, *A&A*, 678, A103
- Aniyan A. K., Thorat K., 2017, *ApJS*, 230, 20
- 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
- Baltus G., Janquart J., Lopez M., Reza A., Caudill S., Cudell J.-R., 2021, *Phys. Rev. D*, 103, 102003
- Becker B., Vaccari M., Prescott M., Grobler T., 2021, *MNRAS*, 503, 1828
- Bianco F. B. et al., 2022, *ApJS*, 258, 1
- Birrer S., Treu T., 2021, *A&A*, 649, A61
- 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
- Bonvin V. et al., 2017, *MNRAS*, 465, 4914
- Canameras R. et al., 2024, *A&A*, 692, A72
- Cao S., Biesiada M., Gavazzi R., Piórkowska A., Zhu Z.-H., 2015, *ApJ*, 806, 185
- Chae K. H. et al., 2002, *Phys. Rev. Lett.*, 89, 151301
- Chegeni A., Hassani F., Vafaei Sadr A., Khosravi N., Kunz M., 2024, *MNRAS*, 531, 1534
- Curci A., et al. 2024, Proceedings of the 13th International Conference on Pattern Recognition Applications and Methods, Vol. 1. p. 995–1000
- Gebhard T. D., Kilbertus N., Harry I., Schölkopf B., 2019, *Phys. Rev. D*, 100, 063015
- Gilman D., Birrer S., Nierenberg A., Treu T., Du X., Benson A., 2020, *MNRAS*, 491, 6077
- He K., Zhang X., Ren S., Sun J., 2016, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778 Las Vegas,
- He Z. et al., 2020, *MNRAS*, 497, 556
- Hsueh J. W., Enzi W., Vegetti S., Auger M. W., Fassnacht C. D., Despali G., Koopmans L. V. E., McKean J. P., 2020, *MNRAS*, 492, 3047
- Huang G., Liu Z., Van Der Maaten L., Weinberger K. Q., 2017, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, p. 4700–4708,
- Jackson N., 2008, *MNRAS*, 389, 1311
- Kalvankar S., Pandit H., Parwate P., 2020, preprint (arXiv:2008.13611)
- Khosa C. K., Mars L., Richards J., Sanz V., 2020, *J. Phys. G: Nucl. Part. Phys.*, 47, 095201
- Koopmans L. V. E., Treu T., Bolton A. S., Burles S., Moustakas L. A., 2006, *ApJ*, 649, 599
- Krachmalnicoff N., Tomasi M., 2019, *A&A*, 628, A129
- Kuijken K. et al., 2019, *A&A*, 625, A2
- 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, preprint (arXiv:1110.3193)
- Li R. et al., 2020, *ApJ*, 899, 30
- Li R. et al., 2021, *ApJ*, 923, 16
- Marshall P. J. et al., 2016, *MNRAS*, 455, 1171
- Melvin T. et al., 2014, *MNRAS*, 438, 2882
- Menghani G., 2023, *ACM Computing Surveys*, 55, 12
- Metcalfe R. B. et al., 2019, *A&A*, 625, A119
- Möller O., Kitzbichler M., Natarajan P., 2007, *MNRAS*, 379, 1195
- More A. et al., 2016, *MNRAS*, 455, 1191
- Mosby G. et al., 2020, *J. Astron. Telesc. Instrum. Syst.*, 6, 046001
- Nagam B. C., Koopmans L. V. E., Valentijn E. A., Kleijn G. V., de Jong J. T. A., Napolitano N., Li R., Tortora C., 2023, *MNRAS*, 523, 4188
- Oguri M., 2006, *MNRAS*, 367, 1241
- Petrillo C. E. et al., 2017, *MNRAS*, 472, 1129
- Petrillo C. E. et al., 2019a, *MNRAS*, 482, 807
- Petrillo C. E. et al., 2019b, *MNRAS*, 484, 3879
- Priyadarshini I., Puri V., 2021, *Earth Sci. Inform.*, 14, 735
- Rahman Minar M., Naher J., 2018, preprint (arXiv:1807.08169)
- Raza R., Zulfikar F., Khan M. O., Arif M., Alvi A., Iftikhar M. A., Alam T., 2023, *Eng. Appl. Artif. Intell.*, 126, 106902
- Renneby M., Henriques B. M. B., Hilbert S., Nelson D., Vogelsberger M., Angulo R. E., Springel V., Hernquist L., 2020, *MNRAS*, 498, 5804
- Rezaei S., McKean J. P., Biehl M., de Roo W., Lafontaine A., 2022, *MNRAS*, 517, 1156
- Rezaei S., Chegeni A., Javadpour A., Vafaei Sadr A., Cao L., Röttgering H., Staring M., 2025, *Astron. Comput.*, 51, 100921
- Ritondale E., Vegetti S., Despali G., Auger M. W., Koopmans L. V. E., McKean J. P., 2019, *MNRAS*, 485, 2179
- Rojas K. et al., 2022, *Astronomy and Astrophysics*, 668
- Sadr A. V., Farsian F., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 012
- Sérsic J. L., 1968, Atlas de Galaxias Australes. Observatorio Astronomico, Cordoba, Argentina
- Shallue C. J., Vanderburg A., 2018, *AJ*, 155, 94
- Spingola C., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., Lagattuta D. J., Vegetti S., 2018, *MNRAS*, 478, 4816
- Spingola C., McKean J. P., Lee M., Deller A., Moldon J., 2019, *MNRAS*, 483, 2125
- Spiniello C., Trager S. C., Koopmans L. V. E., Chen Y. P., 2012, *ApJ*, 753, L32
- Spiniello C., Trager S., Koopmans L. V. E., Conroy C., 2014, *MNRAS*, 438, 1483
- 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
- Tan M., Le Q., 2019, Proceedings of Machine Learning Research. International conference on machine learning, Long Beach, p. 6105–6114
- Treu T., 2010, *ARA&A*, 48, 87
- Treu T., Koopmans L. V., Bolton A. S., Burles S., Moustakas L. A., 2006, *ApJ*, 640, 662
- Vegetti S., Lagattuta D. J., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., 2012, *Nature*, 481, 341
- Vegetti S., Koopmans L. V. E., Auger M. W., Treu T., Bolton A. S., 2014, *MNRAS*, 442, 2017
- Venugopal V., Raj N. I., Nath M. K., Stephen N., 2023, *Decis. Analytics J.*, 8, 100278
- Wardlow J. L. et al., 2013, *ApJ*, 762, 59
- Willett K. W. et al., 2013, *MNRAS*, 435, 2835
- Wong K. C. et al., 2020, *MNRAS*, 498, 1420
- Wu J., Zhang Y., Qu M., Jiang B., Wang W., 2023, *Universe*, 9, 477
- Wucknitz O., Biggs A. D., Browne I. W. A., 2004, *MNRAS*, 349, 14

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