



<i>Publication Year</i>	2023
<i>Acceptance in OA</i>	2023-11-03T16:44:14Z
<i>Title</i>	DeepGraviLens: a multi-modal architecture for classifying gravitational lensing data
<i>Authors</i>	PINCIROLI VAGO, Nicolo Oreste, FRATERNALI, Piero
<i>Publisher's version (DOI)</i>	https://doi.org/10.1007/s00521-023-08766-9
<i>Handle</i>	http://hdl.handle.net/20.500.12386/34465
<i>Journal</i>	NEURAL COMPUTING & APPLICATIONS



DeepGraviLens: a multi-modal architecture for classifying gravitational lensing data

Nicolò Oreste Pincirolì Vago^{1,2} · Piero Fraternali¹

Received: 14 March 2023 / Accepted: 12 June 2023 / Published online: 23 June 2023
© The Author(s) 2023

Abstract

Gravitational lensing is the relativistic effect generated by massive bodies, which bend the space-time surrounding them. It is a deeply investigated topic in astrophysics and allows validating theoretical relativistic results and studying faint astrophysical objects that would not be visible otherwise. In recent years, machine learning methods have been applied to support the analysis of the gravitational lensing phenomena by detecting lensing effects in datasets consisting of images associated with brightness variation time series. However, the state-of-the-art approaches either consider only images and neglect time-series data or achieve relatively low accuracy on the most difficult datasets. This paper introduces DeepGraviLens, a novel multi-modal network that classifies spatio-temporal data belonging to one non-lensed system type and three lensed system types. It surpasses the current state-of-the-art accuracy results by $\approx 3\%$ to $\approx 11\%$, depending on the considered data set. Such an improvement will enable the acceleration of the analysis of lensed objects in upcoming astrophysical surveys, which will exploit the petabytes of data collected, e.g., from the Vera C. Rubin Observatory.

Keywords Multi-modal deep learning · Fusion · Gravitational lensing · Time series

1 Introduction

In astrophysics, a gravitational lens is a matter distribution (e.g., a black hole) able to bend the trajectory of transiting light, similar to an optical lens. Such apparent distortion is caused by the curvature of the geometry of space-time around the massive body acting as a lens, a phenomenon that forces the light to travel along the geodesics (i.e., the shortest paths in the curved space-time). Strong and weak gravitational lensing focus on the effects produced by particularly massive bodies (e.g., galaxies and black holes), while microlensing addresses the consequences produced by lighter entities (e.g., stars). This research proposes an

approach to automatically classify strong gravitational lenses with respect to the lensed objects and to their evolution through time.

Automatically finding and classifying gravitational lenses is a major challenge in astrophysics. As [1–4] show, gravitational lensing systems can be complex, ubiquitous and hard to detect without computer-aided data processing. The volumes of data gathered by contemporary instruments make manual inspection unfeasible. As an example, the Vera C. Rubin Observatory is expected to collect petabytes of data [5].

Moreover, strong lensing is involved in major astrophysical problems: studying massive bodies that are too faint to be analyzed with current instrumentation; characterizing the geometry, content and kinematics of the universe; and investigating mass distribution in the galaxy formation process [1]. Discovery is only the first step, yet a fundamental one, in the study of gravitational lenses. Finding evidence of strong gravitational lensing enables the validation and the advancement of existing astrophysical theories, such as the theory of general relativity [2], and supports specialized studies aimed at modeling the effects of gravitational lensing on specific entities, such as

✉ Nicolò Oreste Pincirolì Vago
nicolooreste.pincirolì@polimi.it

Piero Fraternali
piero.fraternali@polimi.it

¹ Department of Electronics, Information and Bioengineering, Politecnico di Milano, Via Giuseppe Ponzio, 34, 20133 Milan, MI, Italy

² Osservatorio Astronomico di Roma, INAF, Via Frascati 33, 00078 Monte Porzio Catone, RM, Italy

wormholes [2], Simpson–Visser black holes [3], and Einstein–Gauss–Bonnet black holes [4].

The gravitational lenses discovery task takes as input spatiotemporal observations consisting of images and time series and associates each observation with one class (e.g., “Lens,” “No lens,” “Lensed galaxy”...). Images are obtained from specific regions of the electromagnetic field (e.g., visible and infrared [6], ultra-violet [7], and green, red, and near-infrared [8]), depending on the specific experiment. Time series are also collected in specific electromagnetic field regions. They typically describe brightness variation through time (e.g., [8, 9]), and their sampling frequency depends on the technological constraints of the acquisition instrument. In general, they can be multivariate time series [10, 11]. Observations can be either real (i.e., collected by actual instruments) or simulated (i.e., generated by a software system that replicates the characteristics of real instruments).

Several gravitational lenses discovery approaches and tools have been introduced in the past. Originally, observations were analyzed without the aid of computers [12]. Even after the advent of computer science, observations were initially processed without automated classification systems [13–15]. More recently, machine learning (ML) methods have been exploited. The works [16, 17] use convolutional neural networks (CNNs) to classify gravitational lensing images, [18] exploits a Bayesian approach to categorize image data, and [8] applies a multi-modal approach to classify spatio-temporal data in four simulated datasets generated by the *deeplens* simulator [19]. In particular, [8] classifies gravitational lensing data by applying a CNN to the image and a long short-term memory (LSTM) network to the brightness time series and then fusing the outputs of the two branches, achieving a test accuracy ranging from 48.7 to 78.5%.

State-of-the-art lens detection systems, however, still present several limitations. Some of them (e.g., [16–18]) rely on images only and neglect time-domain data and thus cannot detect transient phenomena such as supernovae explosions, which are of great importance for estimating the rate of expansion of the Universe [20]. The work [8] considers spatio-temporal data, but the proposed DeepZipper multi-modal (image + time series) multi-class (“no lens,” “lensed galaxy,” “lensed type-Ia supernova,” “lensed core-collapse supernova”) classification architecture shows relatively low accuracy on the most challenging simulated datasets. Moreover, the simulated data set, as presented in [8], contains ≈ 4000 samples in each test set, obtained after an eightfold data augmentation. The unique samples, before augmentation, then, amount to 500, making the number of test samples of some sub-classes (e.g., “SN-Ia”) low (≈ 14 samples) and yielding high uncertainty on the test set accuracy results.

The authors of [8] have recently proposed the DeepZipper II architecture [20], which exploits a multi-modal (image + time series) binary (“lensed supernova” vs “other”) classification architecture similar to that of [8], achieving an accuracy of 93% over a mix of real and simulated data. The work [21] applies to image time series (i.e., sequences of images), but the classifier works only on the observations where a supernova is known to be present to infer if lensing has occurred or not.

Multi-modal classification architectures have been exploited in many fields other than astrophysics (e.g., remote sensing and medicine) [22–24]. Only a few approaches consider the combination of a single image and one or more time series [25–27], and most approaches are similar to the architecture proposed in [8, 20]. Other modalities have been also considered (e.g., videos and texts), but such inputs differ from those relevant to astrophysical observations and thus such architectures do not carry over the gravitational lensing discovery task. Section 2 briefly surveys them.

Finally, the evaluation of an automatic system for gravitational lens classification poses specific challenges due to the very nature of the task. In real astrophysical observations, gravitational lenses, especially lensed supernovae, are extremely rare and only a few discoveries have already been validated by the scientific community. The extreme scarcity of ground truth data (i.e., verified discoveries) challenges both training and testing of classification algorithms and motivates the use of simulators for creating synthetic datasets. Such datasets can be used for training, validating and testing a classifier in the usual way. However, when it comes to real data, evaluation can only be done a posteriori by submitting the candidate lensing phenomena to the expert judgement for verification.

This paper presents DeepGraviLens, a novel architecture for the classification of strong gravitational lensing multi-modal data. The considered classes concern both transient and non-transient phenomena, and this research shows the superiority of DeepGraviLens over other spatio-temporal networks not only at finding gravitational lenses, but also at finding gravitationally-lensed supernovae, rare objects of particular interest to the astrophysical community. The contributions can be summarized as follows:

- We introduce the architecture of DeepGraviLens, which takes in input spatio-temporal data of real or simulated astrophysical observations and produces in output a multi-class single-label classification of each spatio-temporal sample. DeepGraviLens exploits three complementary sub-networks trained independently and combines their outputs by means of an SVM final stage. The three sub-networks apply different and complementary ways of combining image and time-series data,

taking advantage of both the local and the global features of the input data.

- We evaluate the designed architecture on four simulated datasets formed by $\approx 20,000$ unique examples, split into a train set with $\approx 14,000$ samples (70% of the data set), a validation set with ≈ 3000 samples (15% of the data set), and a test set with ≈ 3000 samples (15% of the data set). We compare the predictions of DeepGraviLens to the results obtained by the DeepZipper network [8] and by a version of DeepZipper II [20] extended from 2 to 4 classes. DeepGraviLens yields accuracy improvements ranging from ≈ 10 to $\approx 36\%$ with respect to the best version of DeepZipper on each test set and significantly reduces the confusion between similar classes, one of the major issues of gravitational lenses classification.
- We have also compared DeepGraviLens with STNet [25], a spatio-temporal multi-modal neural network recently proposed in remote sensing applications, with an improvement in accuracy ranging from ≈ 3 to 11%.
- Finally, we demonstrate that DeepGraviLens is able to detect the presence of gravitational lenses, and specifically gravitationally-lensed supernovae, in real Dark Energy Survey (DES) data [28].

The obtained improvements in the classification of lensing phenomena will enable a faster and more accurate characterization of future real observations, such as those of the Vera C. Rubin Observatory, and will open the way to the discovery of lensed supernovae, which are among the hardest bodies to detect due to their rarity, scattered spatial distribution and relatively short observable life [28–32].

The rest of this paper is organized as follows: Sect. 2 surveys the related work; Sect. 3 describes the data set and the architecture of DeepGraviLens; Sect. 4 describes the adopted evaluation protocol and presents quantitative and qualitative results; finally, Sect. 5 draws the conclusions and outlines our future work.

2 Related work

This section surveys the previous research in the fields of automated gravitational lensing analysis and multi-modal Deep Learning, which are the foundations of this work.

2.1 Automated gravitational lensing analysis

Classifying gravitational lensing phenomena is a challenging task and the subject of many studies. This section concentrates on data-driven techniques, as opposed to the analytical methods that focus on the design of mathematical models capable of explaining the observed data. It

considers the specific case of lensed supernovae, as representatives of transient phenomena, as they are particularly interesting for the astrophysics community. Some of the most recent and promising approaches are listed in Table 1.

In gravitational lens search, finding lensed supernovae (LSNe) is challenging, as they are rare and fast transient phenomena. The main challenges connected with rarity have been thoroughly analyzed in [48]. A common problem across several lens-finding approaches is the lack of large datasets comprising a sufficient number of real gravitational lens observations. The work [40], then, proposes a training set with mock lenses and real non-lensed data, which is a widespread strategy. Several works [20, 21, 35] also test their trained models on real data and propose some candidate gravitational lenses.

The second major challenge is considering the transient nature of supernovae. The explosion of a supernova leads to a peak in its brightness, which first increases and can then decline at a slower rate in a few months [49]. The benefits of considering brightness time series in the LSNe case have been illustrated by [8, 20], and [21] uses image time series to consider brightness variability. [8] justifies the extraction of brightness time series from image time series noticing that the differences between images in a series are negligible in 17 representative sub-classes of lensed and non-lensed astrophysical objects. For this reason, their input is formed by a representative image and a normalized brightness time series. The work [21], instead, uses image time series for finding lensed supernovae and shows promising results on simulated data. However, it considers only two classes: non-lensed supernovae and lensed supernovae, while [8] considers also other astrophysical objects, both lensed and non-lensed, making the input used by [21] a particular case of theirs.

The work described in [18] applies a Bayesian approach to classify high-resolution images of non-transient phenomena to reproduce the categorization performed by human experts. However, high-resolution images are not always available and the human classification (“Definitely not a lens,” “Possibly a lens,” “Probably a lens,” “Definitely a lens”) is intrinsically imprecise and prone to bias depending on the human classifier.

An alternative to Bayesian methods [46] relies on domain-specific features and separates lensed and non-lensed systems using an SVM, whose output is assessed by human experts. The classifier obtains, in the best case, an AUC of 0.95 on simulated data, but the presence of manually-defined features makes this approach less general than deep learning methods. In particular, it exploits specific hard-coded characteristics of lenses, such as the prevalence of a specific color, which can hardly generalize

Table 1 This table summarizes the main approaches for finding gravitational lenses using data-driven techniques

Paper	Year	Algorithm	Survey	Metric	Metric result	Real data	LSNe discoveries	Trans	LSNe class	Class. type
[20]	2022	Multi-modal NN	DES	Accuracy	0.930	Y	Y	Y	Y	B
[21]	2022	Spatiotemporal NN	YSE	Accuracy	0.950	Y	Y	Y	Y	B
			LSST	Accuracy	0.990	N	None			
[33]	2022	CNN committee	CFIS	Precision*	0.014	Y	N/A	N	N	B
[8]	2022	Multi-modal NN	DES (DES-wide)	Accuracy	0.487	N	None	Y	Y	M
			DES (DES-deep)	Accuracy	0.573	N	None			
			DES (DESI-DOT)	Accuracy	0.735	N	None			
			LSST	Accuracy	0.785	N	None			
[34]	2021	Tree-based	Gaia	Found lenses*	14	Y	N/A	Y	N	B
[35]	2020	CNN	Pan-STARRS 3 π	Accuracy	0.942	Y	N/A	N	N	B
[36]	2020	Rule-based	HSC	FPR	2.30%	Y	N/A	Y	N	B
[37]	2020	HSC	Rule-based	Found lenses*	6	Y	N/A	Y	N	B
[38]	2020	CNN	KiDS	FPR	< 0.4%	Y	N/A	N	N	B
[39]	2020	ConvAE and BGM	Euclid	Accuracy	0.773	N	N/A	N	N	B
[40]	2019	CNN	KiDS	Recall*	0.750	Y	N/A	N	N	B
[41]	2019	Tree-based	Gaia	AUC*	0.997	Y	N/A	N	N	B
[42]	2019	CNN	KiDS	Precision*	0.025	Y	N/A	N	N	B
[43]	2019	Tree-based	KiDS	Precision*	0.013	Y	N/A	N	N	M
[44]	2018	CNN	Various	Accuracy	0.982	N	N/A	N	N	B
[45]	2018	CNN committee	GGSLC	AUC	0.988	N	N/A	N	N	B
[46]	2017	SVM	Euclid	AUC	0.89	N	N/A	N	N	B
			KiDS	AUC	0.95	Y				
[47]	2017	CNN	KiDS	Precision*	0.029	Y	N/A	N	N	B
[18]	2009	Bayes	HSC	Completeness	0.900	Y	N/A	N	N	M

In the “Metric” column, “*” indicates that the metric was computed on real data. The “Real data” column indicates whether the algorithm was tested also on real data, the “Trans” column indicates whether transient phenomena are considered, “LSNe class” indicates whether the “LSNe” class is present in the data set, and “Class. type” is the classification type, which can be either binary (B) or multi-class (M)

to multi-label classification tasks or to scenarios where transient phenomena are relevant.

Deep learning-based methods rely mostly on Convolutional Neural Networks (CNNs), as in the binary classifiers illustrated in [16, 40] and in [50], which do not consider time-domain information nor support the fine-grain classification of lensed systems. The work [35] exploits a CNN architecture and tests it also on real data, reporting good results on a binary classification problem that does not focus on LSNe. The authors observe that in their experiments CNN performances relied “heavily on the design of lens simulations and on the choice of negative examples for training, but little on the network architecture.” The works

[8, 20] argue instead that architecture design can lead to great improvements in the results, reporting that multi-modal architectures outperform single-modality CNNs on transient phenomena data. The work [44] describes a CNN-based algorithm trained and tested only on simulated data, which achieves an accuracy of 98% and finds the position of the gravitational lens in the input image. However, the classifier is binary and does not consider LSNe. An interesting approach has been proposed in [45], which focuses on the binary classification of simulated data and proposes a committee of networks, yielding an improvement with respect to individual networks.

As an alternative to supervised methods, [39] defines an unsupervised method for binary classification, which first uses an autoencoder to denoise the image (reducing its resolution), then applies a second autoencoder to extract features from the denoised image, and finally exploits a Bayesian Gaussian Mixture (BGM) to cluster the extracted features. This approach, however, requires human intervention for associating labels to clusters corresponding to the lensed objects.

Several works focused on finding other gravitationally lensed transient phenomena, such as quasars [8], shows that, compared to supernovae, the brightness of quasars changes in a timescale of several years, because they are not explosive phenomena. For this reason, many studies targeting lensed quasars do not use time series information. [43] exploits the image magnitudes in different bands, which is an ad-hoc method that would need adaptation to be applied to the LSNe search. [36] also focuses on finding lensed quasars, but aims at finding quadruply-lensed quasars using an essentially rule-based pipeline. While this method can be effective for the specific application, it should also be modified to tackle more generic and complex cases.

Differently from binary approaches, DeepZipper [8] casts the problem as a multi-class single-label classification task for datasets consisting of images associated with time series of brightness variation. To analyze both images and time-series data, the authors propose a multi-modal network, formed by a CNN and an LSTM, whose outputs are then fused. The resulting system is applied to four simulated datasets corresponding to different astronomical surveys (DES-wide, LSST-wide, DES-deep, and DESI-DOT). This approach, although relatively simple, achieves relatively good results on all four datasets, with accuracies ranging from 48.7% to 78.5%. DeepZipper II [20], an evolution of DeepZipper, introduces minor changes to the network, casts the problem as a binary classification task (“LSNe” vs “other”) instead of a multi-class one, and performs testing on a new dataset partially based on real data. It reaches an accuracy of 93% on DES data and a false positive rate of 0.02%. Three new candidate lensed supernovae found in the DES survey are offered to the astrophysical scientific community for confirmation.

DeepGraviLens, similar to DeepZipper, casts the problem as a multi-class single-label classification task, on the same types of classes and datasets. Compared to previous approaches, it employs more effective unimodal networks and more advanced fusion techniques, which improve the effectiveness in dealing with shared information between the two modalities.

2.2 Multi-modal deep learning and fusion

Several phenomena in the most varied disciplines are characterized by heterogeneous data that give complementary information about the subject under investigation. Multi-modal DL has proved its effectiveness in those domains that require the integrated analysis of multiple data types (e.g., images, videos, and time series). The survey [22] overviews the advances and the trends in multi-modal DL until 2017 and documents usage in such areas as medicine [51–53], human–computer interaction [54] and autonomous driving [55, 56]. The recent survey [57] discusses several applications combining image and text [58, 59], video and text [60, 61], and text and audio [62, 63]. Some applications rely on physiological signals for behavioral studies, such as face recognition [64–66]. In the medical field, [67] overviews the use of AI in oncology and shows the benefits of multi-modal DL. The work [68] diagnoses cervical dysplasia with the integrated analysis of images and numerical data. [69] employs multi-modal DL for classifying malware using textual data from different sources. [70] exploits images and texts to detect hate speech in memes. [71] uses multiple robotic sensors (e.g., cameras, tactile and force sensors) for object manipulation.

From the architecture viewpoint, the processing of heterogeneous inputs can be performed by analyzing the individual data types separately and then fusing the outcome of the different branches to produce an output (late fusion), by stacking the inputs, which are processed together (early fusion), or by introducing fusion at a middle stage (intermediate fusion) [22, 72]. The survey [23] overviews DL methods for multi-modal data fusion in general whereas [72] focuses on biomedical data fusion. The work [57] broadens the comparison beyond DL and contrasts alternative methods employed in multi-modal classification tasks, including SVMs [73], RNNs [66, 74], CNNs [75, 76] and even GANs [77]. The combination of a single image and a time series has been considered by a few works, mainly in the remote sensing [78] and medicine [79] fields. It is apparently similar to the problem of classifying data formed by a video and a time series [80–82]. However, the combination of a single image and a time series, differently from the case of videos, does not require addressing the time-dependent synchronization, connection and interaction between modalities [83]. Another similar case is the joint analysis of image and text. However, text processing poses different challenges and adopts different methods with respect to numeric signals [84]. Another correlated problem is classifying image time series (i.e., sequences of images), as done in several remote sensing applications (e.g., [85–87]). This task, addressed also by [21] for gravitational lensing data, is best applied when

images in the time series vary noticeably. In gravitational lensing data applications such as the one addressed in this paper, instead, the images in the series have small variations. In such a scenario, the use of time series is preferred to the use of image sequences and can be regarded as the extraction of the relevant features from the image sequence [8, 20].

3 Datasets and methods

3.1 Datasets

An input to the lensed object classification task consists of four images and four brightness variation time series, which together represent an astrophysical observation. One image and one time series are provided for each band of the *griz* photometric system, widely used in CCD cameras [88]. In this system, the *g* band is centered on green, the *r* band is centered on red, the *i* band is the near-infrared one, and the *z* band is the infrared one.

Each input is labeled with one of four classes: “No Lens” (no lensed system), “Lens” (Galaxy-Galaxy lensing), “LSNIa” (the lensed object is a Type-Ia supernova), and “LSNCC” (the lensed object is a core-collapse supernova). Section 4.2 shows various examples of input samples and of their classification by DeepGraviLens.

Four distinct datasets (DESI-DOT, LSST-wide, DES-wide, and DES-deep) are built via simulation and are used for training and evaluating DeepGraviLens. The details of their construction are similar to the ones presented in [8, 19, 89]. Each dataset simulates a current or next-generation cosmic survey and is characterized by different specifications of the images and of the associated time series.

The DESI-DOT dataset simulates the observations made by the Dark Energy Camera (DECam) [90] and mirrors the real observing conditions of the DES wide-field survey reported in [91]. The exposure time, a simulation parameter that affects the image quality (higher is better), was set to 60 s. The LSST-wide dataset simulates the LSST survey images acquired using the LSSTCam camera [92]. The simulation parameters were estimated from the conditions of the first year of the survey and the exposure time was set to 30 s [93]. The DES-wide dataset emulates the images from the DECam and uses the real observing conditions from the DES wide-field survey, but the exposure time is 90 s. The DES-deep dataset also reproduces the images from DECam but its characteristics are simulated according to the DES SN program [94] with the exposure time set to 200 s.

Due to the use of the four-bands *griz* photometric system, each image has 4 layers. The image size is $45 \times 45 \times$

4 pixels for all the four datasets. The length of the time series depends on the technical limitations of the simulated instruments. DESI-DOT, LSST-wide, and DES-deep time series contain 14 samples for each band, while DES-wide contains 7 samples for each band.

For each dataset, 17 astrophysical systems were defined and grouped into the four classes “No Lens,” “Lens,” “LSNIa,” and “LSNCC” as proposed in [8]. The examples of the four classes were generated randomly: each class covers $\approx 25\%$ of each dataset and the distribution of the 17 subsystems is the same in all the datasets. Each dataset comprises $\approx 20,000$ elements, split into the train set, containing $\approx 14,000$ samples ($\approx 70\%$), the validation set, containing ≈ 3000 samples ($\approx 15\%$), and the test set, containing ≈ 3000 samples ($\approx 15\%$).

3.2 Extraction of statistical quantities

Two statistical quantities (mean μ and standard deviation σ) are extracted from the brightness time series and used as inputs. Such derived data have a physical meaning. For example, an empty sky is expected to have approximately the same mean value for the four bands and a high standard deviation (because the fluctuations are random). A non-lensed star is expected to be characterized by a low standard deviation, as the means are approximately constant. Even when they manifest a transient behavior (e.g., the explosion of a supernova), the brightness variation is attenuated by the distance. Lensed bodies instead are expected to have a higher standard deviation, because when they display a transient behavior their brightness is amplified by the lens. Other statistical quantities, instead, are not considered as they could not be associated with explainable physical behaviors. The contribution of such derived inputs is quantified in the ablation study described in Sect. 4.

3.3 Overall architecture

Figure 1 illustrates the multi-stage multi-modal inference pipeline of DeepGraviLens. It is formed by three sub-networks (LoNet, GloNet, and MuNet), whose outputs (i.e., the 4 features preceding the final activation function) are concatenated and ensembled using SVM. In this case, concatenation is the process of joining the outputs of length 4 from the three sub-networks end-to-end to obtain a vector of length 12. LoNet and MuNet, in turn, rely on unimodal sub-networks focusing on local or global features in the images and time series. Table 2 summarizes the characteristics of the three networks. GloNet exploits the combination of the image and time-series data, which are merged using early fusion. This approach emphasizes the global features of the multi-modal inputs. LoNet focuses on

Fig. 1 The DeepGraviLens pipeline comprises four steps: (1) the inputs are fed into three independent networks (LoNet, GloNet, and MuNet); (2) the outputs of the three networks are concatenated; (3) the ReFuse network receives the concatenated outputs and (4) outputs a predicted class




Table 2 The three sub-networks pursue different goals: GloNet emphasizes global features and applies early fusion; LoNet accentuates local features and employs intermediate fusion; MuNet extracts both global and local image features

Fusion type	Feature extraction	
	Global	Local
Early	GloNet	
Intermediate	MuNet	LoNet

the local features of the distinct data types: the image and the time series pass through two separate sub-networks and then intermediate fusion is applied. Finally, MuNet extracts both local and global features from the image, using an FC sub-network and a CNN in parallel, and then applies intermediate fusion. The next sections present the three proposed multi-modal networks.


3.4 LoNet, a network focusing on local features

Figure 2 shows the architecture of the LoNet sub-network and Table 3 summarizes its features. It comprises two branches, one for the image (processed through a CNN) and one for the time series (processed by a Gated Recurrent Unit recurrent neural network, from now on referred to as GRU). The network is composed of several layers. After processing the image and the time series, a transformer (similar to [25]) takes in input the concatenation of the

Table 3 Summary of the LoNet neural network architecture showing its layers, output shape, and number of parameters

	Output shape	Parameters #
GRU (features)	[128, 64]	4,068
CNN (features)	[128, 64]	2,337,284
Transformer1d: 1-3	[128, 136]	–
– TransformerEncoder: 2-15	[128, 136]	–
– ModuleList: 3-2	–	2,537,248
Sequential: 1-4	[128, 32]	–
– Linear: 2-16	[128, 64]	8,768
– ReLU: 2-17	[128, 64]	–
– BatchNorm1d: 2-18	[128, 64]	128
– Dropout: 2-19	[128, 64]	–
– Linear: 2-20	[128, 32]	2,080
– ReLU: 2-21	[128, 32]	–
– BatchNorm1d: 2-22	[128, 32]	64
Sequential: 1-5	[128, 4]	–
– Linear: 2-23	[128, 8]	264
– ReLU: 2-24	[128, 8]	–
– BatchNorm1d: 2-25	[128, 8]	16
– Dropout: 2-26	[128, 8]	–
– Linear: 2-27	[128, 4]	36
Total params: 4,889,956		
Trainable params: 4,889,956		
Non-trainable params: 0		
Total mult-adds (G): 3.40		

Fig. 2 LoNet architecture. The time series is processed by the GRU module and the image by a CNN. The two outputs together with the statistics are fused and fed as input to a final transformer module



feature vectors from the GRU and CNN, the means and the standard deviations of the time series. Finally, two sequential components incorporate fully connected layers, dropout, batch normalization, and ReLU activation functions. The last sequential component generates the output of the model. This structure is similar to the one of ZipperNet [8] but replaces the LSTM [95] module with a GRU module with a smaller hidden unit size [96] and batch normalization. The benefits of GRU over LSTM have been shown in several applications [97–101]. In the considered datasets, the short length of the time series makes GRU advantageous over LSTM because the former has fewer training parameters and thus better generalization abilities.

The use of CNN for extracting features from images privileges the focus on contiguous pixels (i.e., small regions of the image), as shown in several studies [102–104]. The use of transformers privileges the extraction of the most important features and contextual information from the CNN and GRU sub-networks.

3.5 GloNet, a network focusing on global features

Figure 3 shows the architecture of the GloNet sub-network and Table 4 summarizes its features. GloNet, differently from LoNet, applies early fusion and relies on a Fully Connected sub-network applied to the flattened inputs. This network comprises multiple linear layers with ReLU activation functions, batch normalization, and dropout. The network comprises two parts (Sequential: 1-1 and Sequential: 1-2). The first part outputs a feature vector, used as input during SVM ensembling, while the second part generates the model output. In this architecture, the image, the time series, the means and the standard deviations are concatenated and given in input to the first linear layer (Linear: 2-1). This approach is complementary to the one of LoNet: it combines the original time series and the original image up-front, rather than merging the features derived from their pre-processing by the GRU and CNN modules. Table 4 also shows that the number of parameters

is higher than in LoNet. Having more parameters allows learning from more complex patterns, which compensates for the absence of convolutional layers.

3.6 MuNet, a network focusing on local and global features

Figure 4 shows the architecture of the MuNet sub-network and Table 5 summarizes its features. It processes the image using two parallel branches: a CNN and an FC sub-network. The time series is processed in the same way as in LoNet. In this network, unlike the case of LoNet, sub-networks results are fused using only a fully connected network, which relies on the ReLU activation function, batch normalization and dropout layers. Compared to LoNet, MuNet adds the FC module applied to the image, to extract local and global features simultaneously. The latter may provide a relevant contribution due to the small size of the images. To avoid overfitting, the number of parameters in the FC sub-network is smaller than in GloNet. In total, the number of parameters is similar to the one of LoNet.

3.7 Ensembling

The three multi-modal networks introduced in this study extract distinct information from the data, emphasizing local features, global features, or a combination of both. To fully leverage the complementary information provided by these networks, ensemble methods can be employed. Table 6 details the ensemble methods used in this study and their associated experimental parameters. For each parameter combination of every method, accuracy is computed on both the train and validation sets. The best parameter combination is then selected based on the highest validation set result, and the accuracy is finally computed on the test set. Moreover, an ablation study is conducted to assess the performance of the best ensemble method when using only two out of the three networks.

Fig. 3 GloNet architecture. The input data are (1) flattened, (2) concatenated, and (3) fed to an FC module




Table 4 Summary of the GloNet neural network architecture showing its layers, output shape, and number of parameters

	Output shape	Parameters #
Sequential: 1-1	[128, 32]	–
– Linear: 2-1	[128, 4096]	33,443,840
– ReLU: 2-2	[128, 4096]	–
– BatchNorm1d: 2-3	[128, 4096]	8,192
– Dropout: 2-4	[128, 4096]	–
– Linear: 2-5	[128, 2048]	8,390,656
– ReLU: 2-6	[128, 2048]	–
– BatchNorm1d: 2-7	[128, 2048]	4,096
– Dropout: 2-8	[128, 2048]	–
– Linear: 2-9	[128, 1024]	2,098,176
– ReLU: 2-10	[128, 1024]	–
– BatchNorm1d: 2-11	[128, 1024]	2,048
– Dropout: 2-12	[128, 1024]	–
– Linear: 2-13	[128, 512]	524,800
– ReLU: 2-14	[128, 512]	–
– BatchNorm1d: 2-15	[128, 512]	1,024
– Dropout: 2-16	[128, 512]	–
– Linear: 2-17	[128, 256]	131,328
– ReLU: 2-18	[128, 256]	–
– BatchNorm1d: 2-19	[128, 256]	512
– Linear: 2-20	[128, 128]	32,896
– ReLU: 2-21	[128, 128]	–
– BatchNorm1d: 2-22	[128, 128]	256
– Dropout: 2-23	[128, 128]	–
– Linear: 2-24	[128, 64]	8,256
– ReLU: 2-25	[128, 64]	–
– BatchNorm1d: 2-26	[128, 64]	128
– Dropout: 2-27	[128, 64]	–
– Linear: 2-28	[128, 32]	2,080
Sequential: 1-2	[128, 4]	–
– Linear: 2-29	[128, 8]	264
– ReLU: 2-30	[128, 8]	–
– Linear: 2-31	[128, 4]	36
Total params: 44,648,588		
Trainable params: 44,648,588		
Non-trainable params: 0		
Total mult-adds (G): 5.72		

In this case, a time series of 14 steps is considered

3.8 Training and validation

The training process of DeepGraviLens is divided into two stages. In the first step, LoNet, GloNet, and MuNet are trained separately, using the same inputs. The second stage consists of training the SVM, which exploits as inputs the values obtained before the application of the final

activation function of the LoNet, GloNet, and MuNet sub-networks. LoNet, GloNet, and MuNet are trained for a maximum of 500 epochs, and the Early Stopping patience is set to 20 epochs. In both stages, the best model is the one with the highest accuracy obtained with the validation process. The validation process aims at selecting the best configurations of the components and of the overall architecture by evaluating the performances on the validation set. It consists of several phases. For unimodal networks, the best-performing model is the one that obtains the highest accuracy value on the validation set. Then, the parameters of the unimodal components of the multi-modal networks are initialized using the previously selected pre-trained unimodal models, and each multi-modal network is trained. The best-performing model is the one that obtains the highest accuracy value on the validation set. Finally, the SVM ensembling phase relies on the previously selected models, and the best ensembling model is the one achieving the highest validation accuracy.

4 Evaluation

This section reports the quantitative and qualitative evaluation of DeepGraviLens on the datasets introduced in 3.1.

For each accuracy result, a confidence interval amounting to 1 standard deviation is calculated to take the limited size of the test set into account. C.R. represents the radius of the confidence interval [106]:

$$\text{C.R.} = \sqrt{\frac{a \cdot (1 - a)}{n}} \quad (1)$$

where a is the mean accuracy (scaled to $[0, 1]$) on the test set and n is the number of samples in the test set.

4.1 Quantitative results

This section presents the outcome of the performance analysis of DeepGraviLens on the four datasets described in Sect. 3.1. For assessing the improvement induced by the proposed architecture, the approach of [8] is used as a baseline, since it is the only research which used a dataset with the same classes as ours. Accuracy is used as the performance metrics because the dataset is balanced. In addition, results were compared with other two multi-modal networks using the time and image modalities, presented in Table 7, and with seven unimodal networks, presented in Table 8. Both DeepZipper II [20] and STNet [25] have been adapted to use four classes rather than the original two.

Ablation experiments with respect to the sub-networks preceding the final ensembling stage are also performed to verify their contribution.

Fig. 4 MuNet architecture. While LoNet processes the image using only a CNN, MuNet employs both a CNN and an FC component




Table 5 Summary of the MuNet neural network architecture showing its layers, output shape, and number of parameters

	Output shape	Parameters #
FC: 1-1	[128, 32]	2,337,284
GRU: 1-2	[128, 64]	4,068
CNN: 1-3	[128, 64]	2,118,804
Linear: 1-4	[128, 32]	2,080
Sequential: 1-5	[128, 32]	–
– Linear: 2-23	[128, 64]	8,768
– ReLU: 2-24	[128, 64]	–
– BatchNorm1d: 2-25	[128, 64]	128
– Dropout: 2-26	[128, 64]	–
– Linear: 2-27	[128, 32]	2,080
– ReLU: 2-28	[128, 32]	–
– BatchNorm1d: 2-29	[128, 32]	64
Sequential: 1-6	[128, 4]	–
– Linear: 2-30	[128, 16]	528
– ReLU: 2-31	[128, 16]	–
– BatchNorm1d: 2-32	[128, 16]	32
– Dropout: 2-33	[128, 16]	–
– Linear: 2-34	[128, 8]	136
– ReLU: 2-35	[128, 8]	–
– BatchNorm1d: 2-36	[128, 8]	16
– Dropout: 2-37	[128, 8]	–
– Linear: 2-38	[128, 4]	36
Total params: 4,474,024		
Trainable params: 4,474,024		
Non-trainable params: 0		
Total mult-adds (G): 3.39		

4.1.1 Prediction performance

Table 7 presents the accuracy results on the four considered test sets. The test set accuracy is similar for the DESI-DOT, LSST-wide and DES-deep datasets and decreases for the more complex DES-wide dataset. In all cases, the accuracy shows an improvement with respect to both the DeepZipper baseline and the best method in the state of the art. Such improvement is observed not only in the case of DeepGraviLens, but also for LoNet and GloNet, making

them viable alternatives to state-of-the-art approaches. Moreover, the performances of GloNet, a simple network, are similar to the ones of DeepZipper and DeepZipper II.

In addition to LoNet and MuNet, the networks EvidentialLoNet and EvidentialMuNet were also implemented and tested. These networks exploit the evidence-based late fusion approach proposed in [107], which dynamically weights the contribution of each modality based on the degree of uncertainty associated with its predictions. Our experiments show that the proposed intermediate fusion approach outperforms the evidence-based fusion approach, with an average improvement of $\approx 4.5\%$.

Figure 5 illustrates the confusion matrices for the four datasets. For the DES-deep dataset, the greatest confusion is observed between “LSNIa” and “LSNCC.” A similar, yet more accentuated pattern, was found in [8] too.

For the DES-wide dataset, the confusions between classes are similar, different from [8], in which the greatest confusion is between “LSNIa” and “LSNCC.” This demonstrates that DeepGraviLens is more effective at discerning between different gravitationally-lensed transient phenomena, reducing the confusion with respect to the baseline [8] significantly.

For the DESI-DOT dataset, the confusion between classes is lower than the one presented in [8]. The greatest confusion is between the “No Lens” and the “Lens” classes, which can be justified by the similarity of the brightness time series of some systems. An example is the “Galaxy + Star” system, in which a galaxy and a star appear close together but without the lensing effect, and the “Galaxy-Galaxy Lensing + Star” system, in which a galaxy stands in front of another galaxy producing the lensing effect and a star appears close to the lensed galaxy from the point of view of the observer.

For the LSST-wide dataset, the greatest confusion is between the “LSNIa” and the “LSNCC” classes as in DES-deep, similar to the pattern observed in [8].

The reported results prove that DeepGraviLens can classify the samples of all the datasets accurately and with

Table 6 Experimental parameters of ensemble methods for aggregating decisions of LoNet, GloNet, and MuNet

Methods	Parametrization
AdaBoost	Estimators: 200
Random forest	Estimators: [10, 50, 100, 200, 500, 1000, 2000]
Extra trees	Estimators: 100
Fuzzy ranking [105]	None
Average	None
MLP	Hidden layer sizes: 100 Activation: 'relu' Solver: 'adam' Alpha: 0.0001
KNN	Neighbors: [2, 4, 6, 8, 16, 32]
FCNN	Early stop
Max	None
SVM	C: [10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 10^0 , 10^1] Kernel: ['poly', 'linear', 'rbf', 'sigmoid']

Table 7 Accuracy—comparison of the accuracy of DeepGraviLens and of the best result obtained using state-of-the-art multi-modal methods

	DESI-DOT	DES-deep	DES-wide	LSST-wide
DeepZipper [8]	77.1	58.6	51.7	74.3
DeepZipper II [20]	78.9	57.4	49.8	70.7
STNet [25]	85.1	58.4	82.5	84.3
EvidentialLoNet (Ours)	81.6	65.6	79.9	84.5
EvidentialMuNet (Ours)	81.1	65.6	79.4	84.2
LoNet (ours)	87.0	67.5	85.8	87.2
GloNet (ours)	77.2	62.3	76.8	76.8
MuNet (ours)	87.9	67.9	86.5	88.5
DeepGraviLens (ours)	<u>88.7</u>	<u>69.6</u>	<u>87.7</u>	<u>88.8</u>
Improvement	3.6	11.0	5.2	4.5

An improvement of $\approx 10\%$ to $\approx 36\%$ is achieved with respect to DeepZipper [8], the only work using a dataset with the same classes as DeepGraviLens. When compared to the best result obtained by reproducing state-of-the-art approaches, the improvement ranges between $\approx 3\%$ and $\approx 11\%$

Table 8 Comparison of the unimodal networks and DeepGraviLens—the table shows the performance of different unimodal networks on image and time modalities, used in Deep Zipper, STNet, and DeepGraviLens

Modality	Unimodal network	Multi-modal networks	DESI-DOT	DES-deep	DES-wide	LSST-wide	Average
Image	ResMixer	STNet	81.4	65.1	82.5	82.1	77.8
	CNN (ours)	LoNet, MuNet	78.4	65.9	79.7	81.4	76.4
	CNN (DZ)	DeepZipper	74.3	61.9	70.0	74.3	70.1
	FCNN (ours)	MuNet	67.7	57.7	62.3	59.3	61.8
Time	GRU (DZ)	DeepZipper	70.9	28.5	39.1	67.0	51.4
	GRU (Ours)	LoNet, MuNet	69.2	32.4	38.9	61.0	50.4
	PDNet	STNet	63.8	28.5	32.9	60.9	46.5
Multi-modal	DeepGraviLens (ours)		<u>88.7</u>	<u>69.6</u>	<u>87.7</u>	<u>88.8</u>	<u>83.7</u>

The best unimodal results are highlighted in bold, and the proposed network's performance is underlined

a significant performance improvement with respect to the compared methods. The results on DES-wide show a

significant improvement, reducing the confusion between lensed supernovae classes. This dataset is particularly

Fig. 5 Confusion matrices of the **a** DES-deep, **b** DES-wide, **c** DESI-DOT, and **d** LSST-wide datasets. In general, the greatest confusion is observed between “Lens” and “No Lens,” and in the case of the DES-wide dataset, between “LSNCC” and “LSNIa,” due to the low sampling rate




Table 9 Comparison of 10 ensemble methods accuracies

Ensemble method	DESI-DOT	DES-deep	DES-wide	LSST-wide	Average
AdaBoost	87.2 ± 0.6	66.7 ± 0.9	86.1 ± 0.6	87.9 ± 0.6	82.0
Random Forest	88.3 ± 0.6	68.6 ± 0.8	87.0 ± 0.6	88.8 ± 0.6	83.2
Extra Trees	88.3 ± 0.6	68.7 ± 0.8	87.1 ± 0.6	88.8 ± 0.6	83.2
Fuzzy ranking [105]	88.3 ± 0.6	68.8 ± 0.8	87.2 ± 0.6	88.6 ± 0.6	83.2
Average	88.1 ± 0.6	68.8 ± 0.8	87.3 ± 0.6	88.7 ± 0.6	83.2
MLP	88.0 ± 0.6	68.9 ± 0.8	87.6 ± 0.6	88.5 ± 0.6	83.3
KNN	88.3 ± 0.6	68.7 ± 0.8	87.1 ± 0.6	88.9 ± 0.6	83.3
FCNN	88.4 ± 0.6	69.2 ± 0.8	87.6 ± 0.6	88.4 ± 0.6	83.4
Max	88.6 ± 0.6	68.7 ± 0.8	87.3 ± 0.6	89.1 ± 0.6	83.4
SVM	88.7 ± 0.6	69.6 ± 0.8	87.7 ± 0.6	88.8 ± 0.6	83.7
Improvement w.r.t. MuNet	0.8	1.7	1.2	0.3	1.0

The underlined results are the best ones for each dataset. The values in bold are the ones comprised in the 1σ confidence interval of the best results. The best performances are obtained using SVM on DESI-DOT, DES-deep, and DES-wide, while Max is the best on LSST-wide

Table 10 Ablation studies on SVM ensemble—when a single network is considered, accuracy refers to the results obtained by applying it without any additional decision-level algorithm

Dataset	LoNet	GloNet	MuNet	Accuracy $\pm 1\sigma$
DESI-DOT	✓			87.0 ± 0.6
		✓		77.2 ± 0.8
			✓	87.9 ± 0.6
	✓	✓		87.0 ± 0.6
	✓		✓	<u>88.7 ± 0.6</u>
		✓	✓	87.9 ± 0.6
	✓	✓	✓	<u>88.7 ± 0.6</u>
	✓			67.5 ± 0.9
DES-deep		✓		62.3 ± 0.9
			✓	67.9 ± 0.9
	✓	✓		68.4 ± 0.8
	✓		✓	68.7 ± 0.8
		✓	✓	68.7 ± 0.8
	✓	✓	✓	<u>69.6 ± 0.8</u>
	✓			85.8 ± 0.6
		✓		76.8 ± 0.8
DES-wide			✓	86.5 ± 0.6
	✓	✓		<u>87.2 ± 0.6</u>
	✓		✓	<u>87.3 ± 0.6</u>
		✓	✓	86.9 ± 0.6
	✓	✓	✓	<u>87.7 ± 0.6</u>
	✓			87.2 ± 0.6
		✓		76.8 ± 0.8
			✓	<u>88.5 ± 0.6</u>
LSST-wide	✓	✓		87.4 ± 0.6
	✓		✓	<u>88.5 ± 0.6</u>
		✓	✓	<u>88.5 ± 0.6</u>
	✓	✓	✓	<u>88.5 ± 0.6</u>
	✓	✓	✓	<u>88.8 ± 0.6</u>

The underlined results are the best mean accuracy results for every dataset, and results in bold are contained within the confidence intervals of the best results. All the values are expressed in %

challenging because lensed galaxies are fainter due to the simulated optical depth of the images, which depends on the technical characteristics of the simulated instrumentation. Moreover, the time series are shorter than in the other datasets and thus contain less information.

4.1.2 Ablation studies

Table 9 compares SVM with other ensemble methods. The use of SVM brings an average 1% improvement over the best multi-modal network (MuNet) and surpasses the performances of other ensemble methods in three datasets out of four. Considering the LSST-wide dataset, Max performs better than SVM, but the SVM result is inside Max's

confidence interval. Moreover, Max's accuracy on DES-deep is outside the SVM confidence interval. Considering the analyzed ensemble methods, only SVM, Fuzzy Ranking [105] and Average are inside the confidence interval of the best ensemble approach for all the datasets. However, both Fuzzy Ranking and Average have an accuracy significantly inferior to that of SVM.

Table 10 presents the results of the ablation experiments with respect to the multi-modal sub-networks. The presence of the three sub-networks guarantees the highest accuracy, with the results obtained ensembling one or two networks being often outside the confidence interval of the result obtained by ensembling three networks. In particular, combining three networks yields an improvement ranging from +0.3% to +12.0% with respect to single networks, and a change ranging from 0.0% to +1.7% with respect to the combination of two networks.

In DESI-DOT, the contribution of GloNet is dominated by that of the other two sub-networks and thus eliminating GloNet does not affect accuracy. This can be explained by the use of early fusion in GloNet which does not preserve the information of the image, which is immediately fused with the time series.

The introduction of the means μ and standard deviations σ of the time series yields an additional modest average improvement of 0.5% in accuracy consistently across the datasets. Compared to the predictions made using a random forest with inputs μ and σ , DeepGraviLens accuracy improves from 18 to 49%.

4.1.3 Execution time

DeepGraviLens has been trained using an NVIDIA GeForce GTX 1080 Ti for GloNet, MuNet and LoNet. On average, the network training requires less than 3 h for a single dataset. SVM training time is negligible with respect to the other networks.

4.2 Qualitative results

Section 4.2.1 illustrates some representative examples of the results obtained by DeepGraviLens on the four test sets. All the images are obtained by coadding the *griz* layers, as done in [8]. In the plots, the g band is displayed in green, the r band in red, the i band in blue and the z band in grey. The brightness time series have been obtained from each initial sequence of images by calculating, for each band, the difference between the sum of the central pixels of the image and the image background, reproducing the approach presented in [8].

Section 4.2.2 shows how the application of DeepGraviLens to real data recognizes the presence of




Fig. 6 A positive example on the LSST-wide dataset—This datum belongs to the “No Lens” class. The image shows two separate stars that have a spherical geometry, which suggests they are not lensed.

Moreover, the curves on the right show no consistent brightness variation through time, which indicates the absence of transient phenomena




Fig. 7 A positive example on the DESI-DOT dataset—This datum belongs to the “Lens” class. The lensing effect is visible in the ring pattern around the central body. The flatness of the brightness time

series, instead, indicates the absence of transient phenomena (e.g., explosions), which is expected because the involved entities are galaxies

gravitational lensing phenomena, also confirming the three lensed supernovae candidate systems, a very rare occurrence, reported in [20].

4.2.1 Simulated data

Figure 6 presents a true positive example belonging to the “No Lens” class in the LSST-wide dataset. It shows two stars close to each other, which exhibit a spherical

symmetry, which suggests the absence of lensing. In addition, the brightness curves do not show consistent variations, which indicates the absence of transient phenomena.

Figure 7 presents a true positive example belonging to the “Lens” class, in the DESI-DOT dataset. In this system, the lensing effect is manifested by the ring pattern on the

¹ <https://datalab.noirlab.edu/> (As of April 2023).




Fig. 8 A positive example on the DESI-DOT dataset—This datum belongs to the “LSNIa” class. The lensing effect is visible from the elliptical shape of the central body, while the presence of a supernova can be observed by the peaks in the brightness time series, which

indicates the presence of explosive transient phenomena. The supernova type can be inferred from the flatness of the g band time series




Fig. 9 A positive example on the DES-wide dataset—This datum belongs to the “LSNIa” class. The lensing effect is visible because of the elliptical shape of the central body. Even if the peaks that indicate

the presence of transient phenomena are absent, the network is still able to correctly classify the datum

central body. The flatness of the brightness curves indicates the absence of transient phenomena, as expected, because the system is formed by galaxies, which are not characterized by explosive events.

Figure 8 presents a true positive example belonging to the “LSNIa” class in the DESI-DOT dataset. The peak in the time series indicates the presence of an exploding supernova and the image shows an elliptical shape, which

signals the presence of lensing. The brightness in the g band is almost flat, which is distinctive of Type Ia supernovae. Type Ia and core-collapse supernovae release chemical elements during the explosion and produce photons at different wavelengths, which are detected by sensors in specific bands. During explosions, the emission of an element with a certain wavelength produces a temporary brightness peak in the corresponding band. Both types of




Fig. 10 A positive example on the DESI-DOT dataset—This datum belongs to the “LSNCC” class. In this case, the lensing effect is suggested both by the presence of varying time curves (indicating the presence of a supernova) and the green body lensed by the galaxy




Fig. 11 A negative example on the LSST-wide dataset—This datum belongs to the “LSNCC” class, but has been classified as “Lens.” The lensing effect is alluded by the halo surrounding the star, while

the flat time series suggests the absence of a transient phenomenon, which induces the wrong classification

supernovae release chemical elements whose detection can be observed in the g band, but Type-Ia supernovae emit less materials than core-collapse supernovae, which makes the latter exhibit a more pronounced peak in the g band. The absence of such a peak in Fig. 8 justifies the “LSNIa” classification.

The same type of system is shown in Fig. 9, from the DES-wide dataset. In this case, the peaks are not detected because of the lower sampling rate, which misses rapid transient events. However, the network correctly classifies

this example thanks to the information contained in the image.

Figure 10 presents a true positive example belonging to the “LSNCC” class, in the DESI-DOT dataset. In this case, the presence of a supernova is indicated by the rapid variation in the brightness time series. Since also the g band exhibits a peak, the input is classified as a core-collapse supernova. The lensing effect is manifested in the image by the supernova (the green body), lensed by the galaxy in front of it. The green color confirms the presence




Fig. 12 A negative example on the DES-deep dataset—This datum belongs to the “Lens” class, but has been classified as “No Lens.” The lensing effect is suggested by the halo surrounding the central body




Fig. 13 A negative example on the DESI-DOT dataset—This datum belongs to the “No Lens” class, but it has been classified as belonging to the “Lens” class. The lensing effect is suggested by the elliptical shape, but such shape may suggest also the presence of a non-lensing

elliptical galaxy. The flatness of the time series, in addition, does not allow to discern “Lens” and “No Lens” systems, as some “No Lens” systems also have flat time series

of elements emitting photons in the g band and the body itself is visible because of the magnifying effect induced by the galaxy.

Figure 11 presents a negative example in the LSST-wide dataset. The datum belongs to the “LSNCC” class, but is classified as “Lens,” which means that the model was not able to detect the presence of a supernova and interpreted the example as a lensed system without evident transient phenomena. The wrong classification is caused by

the low-quality time series and the ambiguous image. The lensing effect is visible thanks to the faint halo surrounding the star in the background, but the time series (wrongly) suggest the absence of a transient phenomenon. The apparent lack of the transient phenomenon can be explained by considering that supernovae explosions can happen in a short time and the brightness variation may not be recorded by the camera. Soon after the explosion, the

Table 11 Summary of results on the considered real data, including system ID, coordinates, number of observations, predicted class, and the proportion of observations in which that class was observed

System ID [20]	Coordinates [°]		# Observations	Predicted class	
	RA	DEC		Class	Proportion
691022126	53.898910	−28.912293	77	LSNCC	65%
701263907	40.969218	−0.619054	69	Lens	100%
699919273	10.155917	−44.437515	76	Lens	93%

Here, RA indicates the right ascension, and DEC indicates the declination




Fig. 14 The detection of a real gravitationally-lensed supernova—This system is formed by four objects, whose boundaries are not well-defined. The time series shows the presence of peaks in the four

bands. The presence of a peak in the g band suggests the presence of an LSNCC, as predicted by DeepGraviLens

brightness returns to the original value, which explains the flatness of the curves.

Figure 12 presents a negative example from the DES-deep dataset, belonging to the “Lens” class, but classified as “No Lens.” The lensing effect is visible on the central body, which has a halo. However, because of the low image resolution, this effect is not as clear in most of the positive examples. In addition, the presence of multiple peaks is not frequently associated with the “Lens” class and induces the wrong classification.

As a final example, Fig. 13 shows an ambiguous image in the DESI-DOT dataset, incorrectly classified. The sample belongs to the “No Lens” class, but is classified as “Lens.” The confusion is generated chiefly by the elliptical object, which is confused with a lensing effect, while it can represent, e.g., a non-lensed elliptical galaxy. The time series are flat, so they do not help discern “Lens” and “No Lens” systems, because some “No Lens” systems also have flat time series.

4.2.2 Real data

The authors of [20] analyze real data from the Dark Energy Survey over a five-year period (Y1–Y5) with the aim of detecting gravitationally-lensed supernovae. They identify three potential lensed supernova systems (identified as 691022126, 701263907, and 699919273), two of which were detected using only Y5 data, indicating that the supernovae likely exploded during that year. Our research tries to reproduce such results using public data provided by NoirLab,¹ which currently only includes data up to Y4, using the network trained on the DES-deep dataset.

DeepGraviLens successfully identified the lensed supernova with ID 691022126 and also detected the presence of a gravitational lens for the other two systems. To extract brightness time series, we followed a methodology similar to the one employed in [20], using 14 time steps

¹ <https://datalab.noirlab.edu/> (As of April 2023).




Fig. 15 The missed detection of a real gravitationally-lensed supernova—The system presented in this figure is the same as the one in Fig. 14, but the time series, for this time interval, does not show

significant peaks, suggesting the absence of a transient phenomenon. The clearer separation between the four bodies in the image is not enough to suggest the presence of a supernova




Fig. 16 A real gravitational lens—The system presented in this figure has been classified as a gravitationally-lensed supernova by [20]. However, the detection was performed on the fifth year of the

observation, which is not publicly available. At the time of the observation, the lens is already present, but the supernova explosion is not visible yet. The time series, indeed, are almost flat or noisy

with a 6-day interval between each step, resulting in a 78-day period. The corresponding image was obtained by averaging the images captured during this period. Each system has been observed for more than 78 days, and as such, multiple observations are associated with each system. Finally, images bigger than 45×45 pixels are resized to such dimension.

Table 11 presents a summary of our results on the real data. The number of observations associated with each

system may differ slightly due to missing observations in the database. Our results confirm the findings of [20]. The systems in which a lensed supernova was discovered only in Y5 have a prevalence of “Lens” prediction.

The object with ID 691022126 is shown in Fig. 14. It has been classified as “LSNCC” in 65% of the observations. The presence of a gravitational lens is signaled by the multiple objects visible in the image. Additionally, the peaks in the four bands indicate the presence of a




Fig. 17 A real gravitational lens—The system presented in this figure has been indicated as a gravitationally-lensed supernova by [20]. However, the detection was performed on the fifth year of the

supernova and the peak in the g band suggests it belongs to the “LSNCC” class, similar to the case shown in Fig. 10. Figure 15 shows the same system at a different time. Although the four objects are more clearly visible in the image, the time series appears more flat and does not exhibit the typical peaks of exploding supernovae. There are several possible explanations for this. One hypothesis is that the supernova has already exploded and the brightness change is no longer detectable. Another possibility is that real data are inherently more variable than simulated data and noise makes peaks difficult to detect.

Figure 16 presents the system identified with ID 699919273, which exhibits a clear gravitational lens. Additionally, this system contains multiple objects, which are likely to be lensed versions of the same astrophysical object. The authors of [20] classify this system as a gravitationally-lensed supernova, based on Y5 data (not publicly available). With the available data up to Y4, the system is classified as a “Lens,” which confirms the category assigned by [20] with the public data up to Y4.

Figure 17 presents the more complex system with ID 701263907, in which the identification of individual objects is challenging due to their blurred boundaries. The presence of halos around the central bodies and in the bottom-right corner of the image suggests the existence of a gravitationally-lensed object. It is possible that the lens extends beyond the boundaries of the image, further complicating its identification. The absence of evident peaks in the time series data suggests the absence of transient phenomena. Specifically, the peaks observed in the g band do not correspond with significant peaks in other

observation, which is not publicly available. Before, the lens is already present, but the supernova explosion is not visible yet. The time series, indeed, are almost flat or noisy

bands, indicating the absence of relevant transient effects. Similar to system 699919273, data up to Y4 hint at the presence of a lens, which DeepGraviLens correctly identifies.

5 Conclusions and future work

This work has introduced DeepGraviLens, a neural architecture for the classification of simulated and real gravitational lensing phenomena that processes multi-modal inputs by means of sub-networks focusing on complementary data aspects. DeepGraviLens surpasses the state-of-the-art accuracy results by $\approx 3\%$ to $\approx 11\%$ on four simulated datasets with different data quality. In particular, it attains a 4.5% performance increase on the LSST-wide dataset, which simulates the acquisitions of the Vera C. Rubin Observatory, whose operations are scheduled to start in 2023. The Vera C. Rubin Observatory is expected to detect hundreds to thousands of lensed supernovae systems, which represents a breakthrough with respect to the capacity of previous instruments. The enormous amount of data that will be acquired demands highly accurate and fast computer-aided classification tools, such as DeepGraviLens.

Despite the promising results, some limitations should be considered. The presented approach should be validated on additional real data for a more robust evaluation. In addition, a generalization study should be conducted to assess the accuracy of the proposed methods when applied to other types of gravitationally-lensed objects such as

quasars and to datasets collected by astronomical surveys characterized by different acquisition parameters.

Future work will concentrate on the application of DeepGraviLens to real observations as soon as they become available. The presented approach could be applied to automatic and extensive sky searches. First, sequences of images can be extracted from survey data in the positions with a previously detected astrophysical object. Then, they can be processed to extract a single image and a time series in the same format as the ones accepted by DeepGraviLens. Finally, predictions can be made. However, such an extensive search would require the analysis of a large number of sources before finding a gravitational lens because of the rarity of gravitational lenses with respect to other objects.

The envisioned research work will also pursue the objective of creating a scientist-friendly system that allows experts to import and manually classify data from real observations to create a non-simulated dataset and compute relevant classification and object detection metrics for automated data analysis, following an approach similar to the one implemented in [108–110]. Finally, we plan to employ the multi-modal architecture designed for DeepGraviLens for the analysis of other (possibly non-astrophysical) datasets characterized by images and time series.

Acknowledgements The authors thank Robert Morgan, author of [8], for having provided guidance in the use of the *deeplens* astronomy simulator and the generation and use of the dataset, and professors Hans Georg Schaathun and Ben David Normann for proofreading the initial version of the article. This project used public archival data from the Dark Energy Survey (DES). Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft, and the Collaborating Institutions in the Dark Energy Survey. The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, the National Optical Astronomy Observatory, the University of Nottingham, The Ohio State University, the

OzDES Membership Consortium, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, and Texas A&M University. Based in part on observations at Cerro Tololo Inter-American Observatory, National Optical Astronomy Observatory, which is operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

Author contributions All authors contributed to the study conception, design, material preparation, data collection and analysis. All authors read and approved the final manuscript.

Funding Open access funding provided by Politecnico di Milano within the CRUI-CARE Agreement.

Availability of data and materials The datasets and best models generated during the current study are available in the DeepGraviLens repository, <https://github.com/nicolopinci/deepgravilens>.

Code availability The code used in the current study is available in the DeepGraviLens repository, <https://github.com/nicolopinci/deepgravilens>.

Declarations

Conflict of interest None.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Treu T (2010) Strong lensing by galaxies. *Ann Rev Astron Astrophys* 48(1):87–125. <https://doi.org/10.1146/annurev-astro-081309-130924>
2. Shaikh R, Banerjee P, Paul S, Sarkar T (2019) Strong gravitational lensing by wormholes. *J Cosmol Astropart Phys* 2019(07):028
3. Islam SU, Kumar J, Ghosh SG (2021) Strong gravitational lensing by rotating Simpson–Visser black holes. *J Cosmol Astropart Phys* 2021(10):013
4. Jin X-H, Gao Y-X, Liu D-J (2020) Strong gravitational lensing of a 4-dimensional Einstein–Gauss–Bonnet black hole in homogeneous plasma. *Int J Mod Phys D* 29(09):2050065

5. Vicedomini M, Brescia M, Cavuoti S, Riccio G, Longo G (2021) Statistical characterization and classification of astronomical transients with machine learning in the era of the Vera C. Rubin observatory. Springer, Cham, pp 81–113. https://doi.org/10.1007/978-3-030-65867-0_4
6. Dahle H, Kaiser N, Irgens RJ, Lilje PB, Maddox SJ (2002) Weak gravitational lensing by a sample of X-ray luminous clusters of galaxies. I. the data set. *Astrophys J Suppl Ser* 139(2):313
7. Quider AM, Pettini M, Shapley AE, Steidel CC (2009) The ultraviolet spectrum of the gravitationally lensed galaxy ‘the Cosmic Horseshoe’: a close-up of a star-forming galaxy at $z = 2$. *Mon Not R Astron Soc* 398(3):1263–1278
8. Morgan R, Nord B, Bechtol K, González S, Buckley-Geer E, Möller A, Park J, Kim A, Birrer S, Aguena M et al (2022) DeepZipper: a novel deep-learning architecture for lensed supernovae identification. *Astrophys J* 927(1):109
9. Vakulik V, Schild R, Dudinov V, Nuritdinov S, Tsvetkova V, Burkhonov O, Akhunov T (2006) Observational determination of the time delays in gravitational lens system Q2237+ 0305. *Astron Astrophys* 447(3):905–913
10. Park JW, Villar A, Li Y, Jiang Y-F, Ho S, Lin JY-Y, Marshall PJ, Roodman A (2021) Inferring black hole properties from astronomical multivariate time series with Bayesian attentive neural processes. [arXiv:2106.01450](https://arxiv.org/abs/2106.01450)
11. Park JW (2018) Strongly-lensed quasar selection based on both multi-band tabular data. Project report of the “CS230 Deep Learning” (2018 edition) course at Stanford
12. Zwicky F (1937) On the probability of detecting nebulae which act as gravitational lenses. *Phys Rev* 51(8):679
13. Gorenstein M, Shapiro I, Cohen N, Corey B, Falco E, Marcaide J, Rogers A, Whitney A, Porcas R, Preston R et al (1983) Detection of a compact radio source near the center of a gravitational lens: quasar image or galactic core? *Science* 219(4580):54–56
14. Lawrence C, Schneider D, Schmidt M, Bennett C, Hewitt J, Burke B, Turner E, Gunn J (1984) Discovery of a new gravitational lens system. *Science* 223(4631):46–49
15. Tyson JA, Valdes F, Wenk R (1990) Detection of systematic gravitational lens galaxy image alignments-mapping dark matter in galaxy clusters. *Astrophys J* 349:1–4
16. Davies A, Serjeant S, Bromley JM (2019) Using convolutional neural networks to identify gravitational lenses in astronomical images. *Mon Not Roy Astron Soc* 487(4): 5263–5271. <https://academic.oup.com/mnras/article-pdf/487/4/5263/28893573/stz1288.pdf>. <https://doi.org/10.1093/mnras/stz1288>
17. Teimoorinia H, Toyonaga RD, Fabbro S, Bottrell C (2020) Comparison of multi-class and binary classification machine learning models in identifying strong gravitational lenses. *Publ Astron Soc Pac* 132(1010):044501
18. Marshall PJ, Hogg DW, Moustakas LA, Fassnacht CD, Bradač M, Schrabback T, Blandford RD (2009) Automated detection of galaxy-scale gravitational lenses in high-resolution imaging data. *Astrophys J* 694(2):924
19. Morgan R, Nord B, Birrer S, Lin JY-Y, Poh J (2021) Deepenstronomy: a dataset simulation package for strong gravitational lensing. [arXiv:2102.02830](https://arxiv.org/abs/2102.02830)
20. Morgan R, Nord B, Bechtol K, Möller A, Hartley W, Birrer S, González S, Martínez M, Gruendl R, Buckley-Geer E et al (2022) Deepzipper ii: searching for lensed supernovae in dark energy survey data with deep learning. [arXiv:2204.05924](https://arxiv.org/abs/2204.05924)
21. Kodi Ramanah D, Arendse N, Wojtak R (2022) AI-driven spatio-temporal engine for finding gravitationally lensed type Ia supernovae. *Mon Not Roy Astron Soc* 512(4):5404–5417. <https://academic.oup.com/mnras/article-pdf/512/4/5404/43377895/stac838.pdf>. <https://doi.org/10.1093/mnras/stac838>
22. Ramachandram D, Taylor GW (2017) Deep multimodal learning: a survey on recent advances and trends. *IEEE Signal Process Mag* 34(6):96–108
23. Gao J, Li P, Chen Z, Zhang J (2020) A survey on deep learning for multimodal data fusion. *Neural Comput* 32(5):829–864
24. Gao J, Li P, Chen Z, Zhang J (2020) A survey on deep learning for multimodal data fusion. *Neural Comput* 32(5): 829–864. https://direct.mit.edu/neco/article-pdf/32/5/829/1865303/neco_a_01273.pdf. https://doi.org/10.1162/neco_a_01273
25. Fan R, Li J, Song W, Han W, Yan J, Wang L (2022) Urban informal settlements classification via a transformer-based spatial-temporal fusion network using multimodal remote sensing and time-series human activity data. *Int J Appl Earth Obs Geoinf* 111:102831
26. Jayachitra V, Nivetha S, Nivetha R, Harini R (2021) A cognitive iot-based framework for effective diagnosis of covid-19 using multimodal data. *Biomed Signal Process Control* 70:102960
27. Gadiraju KK, Ramachandra B, Chen Z, Vatsavai RR (2020) Multimodal deep learning based crop classification using multi-spectral and multitemporal satellite imagery. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery and data mining, pp 3234–3242
28. Diehl HT, Neilsen E, Gruendl RA, Abbott TMC, Allam S, Alvarez O, Annis J, Balbinot E, Bhargava S, Bechtol K, Bernstein GM, Bhatawdekar R, Bocquet S, Brout D, Capasso R, Cawthon R, Chang C, Cook E, Conselice CJ, Cruz J, D’Andrea C, da Costa L, Das R, DePoy DL, Drlica-Wagner A, Elliott A, Everett SW, Frieman J, Neto AF, Ferté A, Friswell I, Furnell KE, Gelman L, Gerdes DW, Gill MSS, Goldstein DA, Gruen D, Gullledge DJ, Hamilton S, Hollowood D, Honscheid K, James DJ, Johnson MD, Johnson MWG, Kent S, Kessler RS, Khullar G, Kovacs E, Kremin A, Kron R, Kuropatkin N, Lasker J, Lathrop A, Li TS, Manera M, March M, Marshall JL, Medford M, Menanteau F, Mohammed I, Monroy M, Moraes B, Morganson E, Muir J, Murphy M, Nord B, Pace AB, Palmese A, Park Y, Paz-Chinchón F, Pereira MES, Petravick D, Plazas AA, Poh J, Prochaska T, Romer AK, Reil K, Roodman A, Sako M, Sauseda M, Scolnic D, Secco LF, Sevilla-Noarbe I, Shipp N, Smith JA, Soares-Santos M, Soergel B, Stebbins A, Story KT, Stringer K, Tarsitano F, Thomas B, Tucker DL, Vivas K, Walker AR, Wang M-Y, Weaverdyck C, Weaverdyck N, Wester W, Wethers CF, Wilkenson R, Wu H-Y, Yanny B, Zenteno A, Zhang Y (2018) Dark energy survey operations: years 4 and 5. In: Peck AB, Seaman RL, Benn CR (eds) *Observatory Operations: strategies, Processes, and Systems VII*, vol 10704, pp 138–155. SPIE. <https://doi.org/10.1117/12.2312113>. International Society for Optics and Photonics
29. Ivezić Ž, Kahn SM, Tyson JA, Abel B, Acosta E, Allsman R, Alonso D, AlSayyad Y, Anderson SF, Andrew J et al (2019) LSST: from science drivers to reference design and anticipated data products. *Astrophys J* 873(2):111
30. Oguri M (2019) Strong gravitational lensing of explosive transients. *Rep Prog Phys* 82(12):126901
31. Goldstein DA, Nugent PE (2016) How to find gravitationally lensed type Ia supernovae. *Astrophys J Lett* 834(1):5
32. Wojtak R, Hjorth J, Gall C (2019) Magnified or multiply imaged?—Search strategies for gravitationally lensed supernovae in wide-field surveys. *Mon Not Roy Astron Soc* 487(3):3342–3355. <https://academic.oup.com/mnras/article-pdf/487/3/3342/28840832/stz1516.pdf>. <https://doi.org/10.1093/mnras/stz1516>
33. Savary E, Rojas K, Maus M, Clément B, Courbin F, Gavazzi R, Chan J, Lemon C, Vernardos G, Cañameras R et al (2022) Strong lensing in unions: toward a pipeline from discovery to modeling. *Astron Astrophys* 666(ARTICLE), 1

34. Stern D, Djorgovski S, Krone-Martins A, Sluse D, Delchambre L, Ducourant C, Teixeira R, Surdej J, Boehm C, Den Brok J et al (2021) Gaia gral: Gaia dr2 gravitational lens systems. vi. spectroscopic confirmation and modeling of quadruply imaged lensed quasars. *Astrophys J* 921(1):42
35. Cañameras R, Schuldt S, Suyu S, Taubenberger S, Meinhardt T, Leal-Taixé L, Lemon C, Rojas K, Savary E (2020) Holismokes-ii. identifying galaxy-scale strong gravitational lenses in panstarrs using convolutional neural networks. *Astron Astrophys* 644:163
36. Chao DC-Y, Chan JH-H, Suyu SH, Yasuda N, More A, Oguri M, Morokuma T, Jaelani AT (2020) Lensed quasar search via time variability with the hsc transient survey. *Astron Astrophys* 640:88
37. Chan JH, Suyu SH, Sonnenfeld A, Jaelani AT, More A, Yonehara A, Kubota Y, Coupon J, Lee C-H, Oguri M et al (2020) Survey of gravitationally lensed objects in hsc imaging (sugohi)-iv. Lensed quasar search in the hsc survey. *Astron Astrophys* 636:87
38. Li R, Napolitano N, Tortora C, Spiniello C, Koopmans L, Huang Z, Roy N, Vernardos G, Chatterjee S, Giblin B et al (2020) New high-quality strong lens candidates with deep learning in the kilo-degree survey. *Astrophys J* 899(1):30
39. Cheng T-Y, Li N, Conselice CJ, Aragón-Salamanca A, Dye S, Metcalf RB (2020) Identifying strong lenses with unsupervised machine learning using convolutional autoencoder. *Mon Not R Astron Soc* 494(3):3750–3765
40. Petrillo C, Tortora C, Chatterjee S, Vernardos G, Koopmans L, Verdoes Kleijn G, Napolitano NR, Covone G, Kelvin L, Hopkins A (2019) Testing convolutional neural networks for finding strong gravitational lenses in kids. *Mon Not R Astron Soc* 482(1):807–820
41. Delchambre L, Krone-Martins A, Wertz O, Ducourant C, Galluccio L, Klüter J, Mignard F, Teixeira R, Djorgovski S, Stern D et al (2019) Gaia gral: Gaia dr2 gravitational lens systems-iii. a systematic blind search for new lensed systems. *Astron Astrophys* 622:165
42. Petrillo C, Tortora C, Vernardos G, Koopmans L, Verdoes Kleijn G, Bilicki M, Napolitano NR, Chatterjee S, Covone G, Dvornik A et al (2019) Links: discovering galaxy-scale strong lenses in the kilo-degree survey using convolutional neural networks. *Mon Not R Astron Soc* 484(3):3879–3896
43. Khramtsov V, Sergeyev A, Spiniello C, Tortora C, Napolitano NR, Agnello A, Getman F, De Jong JT, Kuijken K, Radovich M et al (2019) Kids-squad-ii. machine learning selection of bright extragalactic objects to search for new gravitationally lensed quasars. *Astron Astrophys* 632:56
44. Pearson J, Pennock C, Robinson T (2018) Auto-detection of strong gravitational lenses using convolutional neural networks. *Emerg Sci* 2:1
45. Schaefer C, Geiger M, Kuntzer T, Kneib J-P (2018) Deep convolutional neural networks as strong gravitational lens detectors. *Astron Astrophys* 611:2
46. Hartley P, Flamary R, Jackson N, Tagore A, Metcalf R (2017) Support vector machine classification of strong gravitational lenses. *Mon Not R Astron Soc* 471(3):3378–3397
47. Petrillo C, Tortora C, Chatterjee S, Vernardos G, Koopmans L, Verdoes Kleijn G, Napolitano NR, Covone G, Schneider P, Grado A et al (2017) Finding strong gravitational lenses in the kilo degree survey with convolutional neural networks. *Mon Not R Astron Soc* 472(1):1129–1150
48. Savary EMC (2022) Teaching machines how to find strongly lensed galaxies in cosmological sky surveys. Technical report, EPFL
49. Hoeflich P, Khokhlov A, Wheeler JC, Phillips MM, Suntzeff NB, Hamuy M (1996) Maximum brightness and postmaximum decline of light curves of type supernovae ia: a comparison of theory and observations. *Astrophys J* 472(2):81
50. Pourrahmani M, Nayyeri H, Cooray A (2018) LensFlow: a convolutional neural network in search of strong gravitational lenses. *Astrophys J* 856(1):68. <https://doi.org/10.3847/1538-4357/aaae6a>
51. Banos O, Villalonga C, Garcia R, Saez A, Damas M, Holgado-Terriza JA, Lee S, Pomares H, Rojas I (2015) Design, implementation and validation of a novel open framework for agile development of mobile health applications. *Biomed Eng Online* 14(2):1–20
52. Cao Y, Steffey S, He J, Xiao D, Tao C, Chen P, Müller H (2014) Medical image retrieval: a multimodal approach. *Cancer Inf* 13:14053
53. Liang M, Li Z, Chen T, Zeng J (2014) Integrative data analysis of multi-platform cancer data with a multimodal deep learning approach. *IEEE/ACM Trans Comput Biol Bioinf* 12(4):928–937
54. Azagra P, Mollard Y, Golemo F, Murillo AC, Lopes M, Civera J (2016) A multimodal human-robot interaction dataset. In: *NIPS 2016, workshop future of interactive learning machines*
55. Geiger A, Lenz P, Stiller C, Urtasun R (2013) Vision meets robotics: the kitti dataset. *Int J Robot Res* 32(11):1231–1237
56. Maddern W, Pascoe G, Linegar C, Newman P (2017) 1 year, 1000 km: the Oxford RobotCar dataset. *Int J Robot Res* 36(1):3–15
57. Summaira J, Li X, Shoib AM, Li S, Abdul J (2021) Recent advances and trends in multimodal deep learning: a review. [arXiv:2105.11087](https://arxiv.org/abs/2105.11087)
58. Liu M, Hu H, Li L, Yu Y, Guan W (2020) Chinese image caption generation via visual attention and topic modeling. *IEEE Trans Cybern*
59. Liu M, Li L, Hu H, Guan W, Tian J (2020) Image caption generation with dual attention mechanism. *Inf Process Manag* 57(2):102178
60. Liu S, Ren Z, Yuan J (2020) Sibnet: sibling convolutional encoder for video captioning. *IEEE Trans Pattern Anal Mach Intell* 43(9):3259–3272
61. Rahman M, Abedin T, Prottoy KS, Moshruha A, Siddiqui FH et al (2020) Semantically sensible video captioning (SSVC). [arXiv:2009.07335](https://arxiv.org/abs/2009.07335)
62. Elias I, Zen H, Shen J, Zhang Y, Jia Y, Weiss RJ, Wu Y (2021) Parallel tacotron: non-autoregressive and controllable tts. In: *ICASSP 2021–2021 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, pp 5709–5713
63. Shen J, Pang R, Weiss RJ, Schuster M, Jaitly N, Yang Z, Chen Z, Zhang Y, Wang Y, Skerrv-Ryan R et al (2018) Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In: *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, pp 4779–4783
64. Cimtay Y, Ekmekcioglu E, Caglar-Ozhan S (2020) Cross-subject multimodal emotion recognition based on hybrid fusion. *IEEE Access* 8:168865–168878
65. Li M, Xie L, Lv Z, Li J, Wang Z (2020) Multistep deep system for multimodal emotion detection with invalid data in the internet of things. *IEEE Access* 8:187208–187221
66. Jaiswal M, Aldeneh Z, Mower Provost E (2019) Controlling for confounders in multimodal emotion classification via adversarial learning. In: *2019 International conference on multimodal interaction*, pp 174–184
67. Shimizu H, Nakayama KI (2020) Artificial intelligence in oncology. *Cancer Sci* 111(5):1452–1460
68. Xu T, Zhang H, Huang X, Zhang S, Metaxas DN (2016) Multimodal deep learning for cervical dysplasia diagnosis. In: *Ourselin S, Joskowicz L, Sabuncu MR, Unal G, Wells W (eds) Medical image computing and computer-assisted intervention-MICCAI 2016*. Springer, Cham, pp 115–123

69. Gibert D, Mateu C, Planes J (2020) HYDRA: a multimodal deep learning framework for malware classification. *Comput Secur* 95:101873. <https://doi.org/10.1016/j.cose.2020.101873>
70. Velioglu R, Rose J (2020) Detecting hate speech in memes using multimodal deep learning approaches: prize-winning solution to hateful memes challenge. *arXiv:2012.12975*
71. Saito N, Ogata T, Funabashi S, Mori H, Sugano S (2021) How to select and use tools?: Active perception of target objects using multimodal deep learning. *IEEE Robot Autom Lett* 6(2):2517–2524
72. Stahlschmidt SR, Ulfenborg B, Synnergren J (2022) Multimodal deep learning for biomedical data fusion: a review. *Briefings Bioinf* 23(2) <https://academic.oup.com/bib/article-pdf/23/2/bbab569/42805085/bbab569.pdf>. <https://doi.org/10.1093/bib/bbab569>
73. Huang Y, Yang J, Liao P, Pan J (2017) Fusion of facial expressions and EEG for multimodal emotion recognition. In: *Computational intelligence and neuroscience 2017*
74. Wei R, Mi L, Hu Y, Chen Z (2020) Exploiting the local temporal information for video captioning. *J Vis Commun Image Represent* 67:102751
75. Hazarika D, Poria S, Mihalcea R, Cambria E, Zimmermann R (2018) Icon: interactive conversational memory network for multimodal emotion detection. In: *Proceedings of the 2018 conference on empirical methods in natural language processing*, pp 2594–2604
76. Oord A, Li Y, Babuschkin I, Simonyan K, Vinyals O, Kavukcuoglu K, Driessche G, Lockhart E, Cobo L, Stimberg F et al (2018) Parallel wavenet: fast high-fidelity speech synthesis. In: *International conference on machine learning*. PMLR, pp 3918–3926
77. Wei Y, Wang L, Cao H, Shao M, Wu C (2020) Multi-attention generative adversarial network for image captioning. *Neurocomputing* 387:91–99
78. Gómez-Chova L, Tuia D, Moser G, Camps-Valls G (2015) Multimodal classification of remote sensing images: a review and future directions. *Proc IEEE* 103(9):1560–1584
79. Arioiz U, Smrke U, Plohl N, Mlakar I (2022) Scoping review on the multimodal classification of depression and experimental study on existing multimodal models. *Diagnostics* 12(11):2683
80. Salekin MS, Zamzmi G, Goldgof D, Kasturi R, Ho T, Sun Y (2021) Multimodal spatio-temporal deep learning approach for neonatal postoperative pain assessment. *Comput Biol Med* 129:104150
81. Cai T, Ni H, Yu M, Huang X, Wong K, Volpi J, Wang JZ, Wong ST (2022) Deepstroke: an efficient stroke screening framework for emergency rooms with multimodal adversarial deep learning. *Med Image Anal* 80:102522
82. Pouyanfar S, Tao Y, Tian H, Chen S-C, Shyu M-L (2019) Multimodal deep learning based on multiple correspondence analysis for disaster management. *World Wide Web* 22:1893–1911
83. Liang PP, Zadeh A, Morency L-P (2022) Foundations and recent trends in multimodal machine learning: principles, challenges, and open questions. *arXiv:2209.03430*
84. Feghali J, Jimenez AE, Schilling AT, Azad TD (2022) Overview of algorithms for natural language processing and time series analyses. In: *Machine learning in clinical neuroscience: foundations and applications*. Springer, pp 221–242
85. Suel E, Bhatt S, Brauer M, Flaxman S, Ezzati M (2021) Multimodal deep learning from satellite and street-level imagery for measuring income, overcrowding, and environmental deprivation in urban areas. *Remote Sens Environ* 257:112339
86. Pelletier C, Webb GI, Petitjean F (2019) Temporal convolutional neural network for the classification of satellite image time series. *Remote Sens* 11(5):523
87. do Nascimento Bendini H, Fonseca LMG, Schwieder M, Körting TS, Rufin P, Sanches IDA, Leitao PJ, Hostert P (2019) Detailed agricultural land classification in the Brazilian Cerrado based on phenological information from dense satellite image time series. *Int J Appl Earth Observ Geoinf* 82:101872
88. Schneider D, Gunn J, Hoessel J (1983) CCD photometry of Abell clusters. I-magnitudes and redshifts for 84 brightest cluster galaxies. *Astrophys J* 264:337–355
89. Birrer S, Shajib AJ, Gilman D, Galan A, Aalbers J, Millon M, Morgan R, Pagano G, Park JW, Teodori L et al (2021) lenstronomy ii: a gravitational lensing software ecosystem. *arXiv:2106.05976*
90. Flaugher B, Diehl HT, Honscheid K, Abbott TMC, Alvarez O, Angstadt R, Annis JT, Antonik M, Ballester O, Beaufore L, Bernstein GM, Bernstein RA, Bigelow B, Bonati M, Boprie D, Brooks D, Buckley-Geer EJ, Campa J, Cardiel-Sas L, Castander FJ, Castilla J, Cease H, Cela-Ruiz JM, Chappa S, Chi E, Cooper C, da Costa LN, Dede E, Derylo G, DePoy DL, de Vicente J, Doel P, Drlica-Wagner A, Eiting J, Elliott AE, Emes J, Estrada J, Neto AF, Finley DA, Flores R, Frieman J, Gerdes D, Gladders MD, Gregory B, Gutierrez GR, Hao J, Holland SE, Holm S, Huffman D, Jackson C, James DJ, Jonas M, Karcher A, Karliner I, Kent S, Kessler R, Kozlovsky M, Kron RG, Kubik D, Kuehn K, Kuhlmann S, Kuk K, Lahav O, Lathrop A, Lee J, Levi ME, Lewis P, Li TS, Mandrichenko I, Marshall JL, Martinez G, Merritt KW, Miquel R, Muñoz F, Neilsen EH, Nichol RC, Nord B, Ogando R, Olsen J, Palaio N, Patton K, Peoples J, Plazas AA, Rauch J, Reil K, Rheault J-P, Roe NA, Rogers H, Roodman A, Sanchez E, Scarpine V, Schindler RH, Schmidt R, Schmitt R, Schubnell M, Schultz K, Schurter P, Scott L, Serrano S, Shaw TM, Smith RC, Soares-Santos M, Stefanik A, Stuermer W, Suchyta E, Sypniewski A, Tarle G, Thaler J, Tighe R, Tran C, Tucker D, Walker AR, Wang G, Watson M, Weaverdyck C, Wester W, Woods R, BY (2015) The dark energy camera. *Astron J* 150(5):150. <https://doi.org/10.1088/0004-6256/150/5/150>
91. Abbott TMC, Abdalla FB, Allam S, Amara A, Annis J, Asorey J, Avila S, Ballester O, Banerji M, Barkhouse W, Baruah L, Baumer M, Bechtol K, Becker MR, Benoit-Lévy A, Bernstein GM, Bertin E, Blazek J, Bocquet S, Brooks D, Brout D, Buckley-Geer E, Burke DL, Busti V, Campisano R, Cardiel-Sas L, Rosell AC, Kind MC, Carretero J, Castander FJ, Cawthon R, Chang C, Chen X, Conselice C, Costa G, Crocce M, Cunha CE, D'Andrea CB, da Costa LN, Das R, Daues G, Davis TM, Davis C, Vicente JD, DePoy DL, DeRose J, Desai S, Diehl HT, Dietrich JP, Dodelson S, Doel P, Drlica-Wagner A, Eifler TF, Elliott AE, Evrard AE, Farahi A, Neto AF, Fernandez E, Finley DA, Flaugher B, Foley RJ, Fosalba P, Friedel DN, Frieman J, García-Bellido J, Gaztanaga E, Gerdes DW, Giannantonio T, Gill MSS, Glazebrook K, Goldstein DA, Gower M, Gruen D, Gruendl RA, Gschwend J, Gupta RR, Gutierrez G, Hamilton S, Hartley WG, Hinton SR, Hislop JM, Hollowood D, Honscheid K, Hoyle B, Huterer D, Jain B, James DJ, Jeltema T, Johnson MWG, Johnson MD, Kacprzak T, Kent S, Khullar G, Klein M, Kovacs A, Koziol AMG, Krause E, Kremin A, Kron R, Kuehn K, Kuhlmann S, Kuropatkin N, Lahav O, Lasker J, Li TS, Li RT, Liddle AR, Lima M, Lin H, López-Reyes P, MacCrann N, Maia MAG, Maloney JD, Manera M, March M, Marriner J, Marshall JL, Martini P, McClintock T, McKay T, McMahon RG, Melchior P, Menanteau F, Miller CJ, Miquel R, Mohr JJ, Morganson E, Mould J, Neilsen E, Nichol RC, Nogueira F, Nord B, Nugent P, Nunes L, Ogando RLC, Old L, Pace AB, Palmese A, Paz-Chinchón F, Peiris HV, Percival WJ, Petravick D, Plazas AA, Poh J, Pond C, Porredon A, Pujol A, Refregier A, Reil K, Ricker PM, Rollins RP, Romer AK, Roodman A, Rooney P, Ross AJ, Rykoff ES, Sako M, Sanchez ML, Sanchez E, Santiago B, Saro

- A, Scarpine V, Scolnic D, Serrano S, Sevilla-Noarbe I, Sheldon E, Shipp N, Silveira ML, Smith M, Smith RC, Smith JA, Soares-Santos M, Sobreira F, Song J, Stebbins A, Suchyta E, Sullivan M, Swanson MEC, Tarle G, Thaler J, Thomas D, Thomas RC, Troxel MA, Tucker DL, Vikram V, Vivas AK, Walker AR, Wechsler RH, Weller J, Wester W, Wolf RC, Wu H, Yanny B, Zenteno A, Zhang Y, Zuntz J, Juneau S, Fitzpatrick M, Nikutta R, Nidever D, Olsen K, Scott A (2018) The dark energy survey: data release 1. *Astrophys J Suppl Ser* 239(2):18. <https://doi.org/10.3847/1538-4365/aac9f0>
92. Stalder B, Reil K, Claver C, Liang M, Tsai TW, Lange T, Haupt J, Wiecha O, Lopez M, Poculp G et al (2020) Rubin commissioning camera: integration, functional testing, and lab performance. In: *Ground-based and airborne instrumentation for astronomy VIII*, vol 11447. SPIE, pp 86–98
93. Marshall P, Clarkson W, Shemmer O, Biswas R, de Val-Borro M, Rho J, Jones L, Anguita T, Ridgway S, Bianco F, Ivezić Z, Lochner M, Meyers J, Vivas K, Graham M, Claver C, Digel S, Kasliwal V, McGehee PM, Gawiser E, Bellm E, Walkowicz L, Olsen K, Yoachim P, Bell K, Nidever D, Lund M, Connolly A, Arcavi I, Awan H (2017) LSST science collaborations observing strategy white paper: “science-driven optimization of the LSST observing strategy”. Zenodo. <https://doi.org/10.5281/zenodo.842713>
94. Abbott TMC, Allam S, Andersen P, Angus C, Asorey J, Avelino A, Avila S, Bassett BA, Bechtol K, Bernstein GM, Bertin E, Brooks D, Brout D, Brown P, Burke DL, Calcino J, Rosell AC, Carollo D, Kind MC, Carretero J, Casas R, Castander FJ, Cawthon R, Challis P, Childress M, Clocchiatti A, Cunha CE, D’Andrea CB, da Costa LN, Davis C, Davis TM, Vicente JD, DePoy DL, Desai S, Diehl HT, Doel P, Drlica-Wagner A, Eifler TF, Evrard AE, Fernandez E, Filippenko AV, Finley DA, Flaugh B, Foley RJ, Fosalba P, Frieman J, Galbany L, García-Bellido J, Gaztanaga E, Giannantonio T, Glazebrook K, Goldstein DA, González-Gaitán S, Gruen D, Gruendl RA, Gschwend J, Gupta RR, Gutierrez G, Hartley WG, Hinton SR, Hollowood DL, Honscheid K, Hoormann JK, Hoyle B, James DJ, Jeltema T, Johnson MWG, Johnson MD, Kasai E, Kent S, Kessler R, Kim AG, Kirshner RP, Kovacs E, Krause E, Kron R, Kuehn K, Kuhlmann S, Kuropatkin N, Lahav O, Lasker J, Lewis GF, Li TS, Lidman C, Lima M, Lin H, Macaulay E, Maia MAG, Mandel KS, March M, Marriner J, Marshall JL, Martini P, Menanteau F, Miller CJ, Miquel R, Miranda V, Mohr JJ, Morganson E, Muthukrishna D, Möller A, Neilsen E, Nichol RC, Nord B, Nugent P, Ogando RLC, Palmese A, Pan Y-C, Plazas AA, Pursiainen M, Romer AK, Roodman A, Roza E, Rykoff ES, Sako M, Sanchez E, Scarpine V, Schindler R, Schubnell M, Scolnic D, Serrano S, Sevilla-Noarbe I, Sharp R, Smith M, Soares-Santos M, Sobreira F, Sommer NE, Spinka H, Suchyta E, Sullivan M, Swann E, Tarle G, Thomas D, Thomas RC, Troxel MA, Tucker BE, Uddin SA, Walker AR, Wester W, Wiseman P, Wolf RC, Yanny B, Zhang B, YZ (2019) First cosmology results using type Ia supernovae from the dark energy survey: constraints on cosmological parameters. *Astrophys J* 872(2):30. <https://doi.org/10.3847/2041-8213/ab04fa>
95. Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9(8):1735–1780
96. Cho K, Van Merriënboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, Bengio Y (2014) Learning phrase representations using rnn encoder-decoder for statistical machine translation. [arXiv:1406.1078](https://arxiv.org/abs/1406.1078)
97. Tang Z, Shi Y, Wang D, Feng Y, Zhang S (2017) Memory visualization for gated recurrent neural networks in speech recognition. In: 2017 IEEE international conference on acoustics, speech and signal processing (ICASSP). IEEE, pp 2736–2740
98. Chung J, Gulcehre C, Cho K, Bengio Y (2014) Empirical evaluation of gated recurrent neural networks on sequence modeling. [arXiv:1412.3555](https://arxiv.org/abs/1412.3555)
99. Collins J, Sohl-Dickstein J, Sussillo D (2016) Capacity and trainability in recurrent neural networks. [arXiv:1611.09913](https://arxiv.org/abs/1611.09913)
100. Jozefowicz R, Zaremba W, Sutskever I (2015) An empirical exploration of recurrent network architectures. In: *International conference on machine learning*. PMLR, pp 2342–2350
101. Yang S, Yu X, Zhou Y (2020) LSTM and GRU neural network performance comparison study: taking yelp review dataset as an example. In: 2020 International workshop on electronic communication and artificial intelligence (IWECAL), pp. 98–101. <https://doi.org/10.1109/IWECAL50956.2020.00027>
102. Luo W, Li Y, Urtasun R, Zemel R (2016) Understanding the effective receptive field in deep convolutional neural networks. In: *Advances in neural information processing systems* 29
103. Pincioli Vago NO, Milani F, Fraternali P, da Silva Torres R (2021) Comparing cam algorithms for the identification of salient image features in iconography artwork analysis. *J Imaging* 7(7):106
104. Milani F, Pincioli Vago NO, Fraternali P (2022) Proposals generation for weakly supervised object detection in artwork images. *J Imaging* 8(8):215
105. Manna A, Kundu R, Kaplun D, Sinitca A, Sarkar R (2021) A fuzzy rank-based ensemble of cnn models for classification of cervical cytology. *Sci Rep* 11(1):14538
106. Witten IH, Frank E (2002) Data mining: practical machine learning tools and techniques with java implementations. *ACM SIGMOD Rec* 31(1):76–77
107. Zhao K, Gao Q, Hao S, Sun J, Zhou L (2023) Credible remote sensing scene classification using evidential fusion on aerial-ground dual-view images. [arXiv:2301.00622](https://arxiv.org/abs/2301.00622)
108. Torres RN, Fraternali P, Romero J (2020) ODIN: An object detection and instance segmentation diagnosis framework. In: *European conference on computer vision*. Springer, pp 19–31
109. Torres RN, Milani F, Fraternali P (2021) ODIN: pluggable meta-annotations and metrics for the diagnosis of classification and localization. In: *International conference on machine learning, optimization, and data science*. Springer, pp 383–398
110. Zangrando N (2021) The ODIN framework, a tool for image classification diagnosis