

DeepZipper: A Novel Deep Learning Architecture for Lensed Supernovae Identification

R. MORGAN,^{1,2,3} B. NORD,^{2,4,5} K. BECHTOL,^{1,6} S. J. GONZÁLEZ,¹ E. BUCKLEY-GEER,^{2,4} A. MÖLLER,⁷ J. W. PARK,^{8,9}
A. G. KIM,¹⁰ S. BIRRER,^{8,9} M. AGUENA,¹¹ J. ANNIS,² S. BOCQUET,¹² D. BROOKS,¹³ A. CARNERO ROSELL,¹¹
M. CARRASCO KIND,^{14,15} J. CARRETERO,¹⁶ R. CAWTHON,¹⁷ L. N. DA COSTA,^{11,18} T. M. DAVIS,¹⁹ J. DE VICENTE,²⁰
P. DOEL,¹³ I. FERRERO,²¹ D. FRIEDEL,¹⁴ J. GARCÍA-BELLIDO,²² M. GATTI,²³ E. GAZTANAGA,^{24,25} G. GIANNINI,¹⁶
D. GRUEN,¹² R. A. GRUENDL,^{14,15} G. GUTIERREZ,² D. L. HOLLOWOOD,²⁶ K. HONSCHIED,^{27,28} D. J. JAMES,²⁹
K. KUEHN,^{30,31} N. KUROPATKIN,² M. A. G. MAIA,^{11,18} R. MIQUEL,^{32,16} A. PALMESE,³³ F. PAZ-CHINCHÓN,^{14,34}
M. E. S. PEREIRA,^{35,36} A. PIERES,^{11,18} A. A. PLAZAS MALAGÓN,³⁷ K. REIL,⁹ A. ROODMAN,^{38,9} E. SANCHEZ,²⁰ M. SMITH,³⁹
E. SUCHYTA,⁴⁰ M. E. C. SWANSON,³⁹ G. TARLE,³⁵ AND C. TO²⁷

¹Physics Department, University of Wisconsin-Madison, 1150 University Avenue Madison, WI 53706, USA

²Fermi National Accelerator Laboratory, P. O. Box 500, Batavia, IL 60510, USA

³Legacy Survey of Space and Time Corporation Data Science Fellowship Program, USA

⁴Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

⁵Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁶Legacy Survey of Space and Time, 933 North Cherry Avenue, Tucson, AZ 85721, USA

⁷Centre for Astrophysics & Supercomputing, Swinburne University of Technology, Victoria 3122, Australia

⁸Kavli Institute for Particle Astrophysics and Cosmology, Department of Physics, Stanford University, Stanford, CA 94305, USA

⁹SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

¹⁰Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

¹¹Laboratório Interinstitucional de e-Astronomia - LIneA, Rua Gal. José Cristino 77, Rio de Janeiro, RJ - 20921-400, Brazil

¹²Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr. 1, 81679 Munich, Germany

¹³Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK

¹⁴Center for Astrophysical Surveys, National Center for Supercomputing Applications, 1205 West Clark St., Urbana, IL 61801, USA

¹⁵Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, Urbana, IL 61801, USA

¹⁶Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona) Spain

¹⁷Physics Department, William Jewell College, Liberty, MO 64068 USA

¹⁸Observatório Nacional, Rua Gal. José Cristino 77, Rio de Janeiro, RJ - 20921-400, Brazil

¹⁹School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia

²⁰Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

²¹Institute of Theoretical Astrophysics, University of Oslo. P.O. Box 1029 Blindern, NO-0315 Oslo, Norway

²²Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, 28049 Madrid, Spain

²³Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA

²⁴Institut d'Estudis Espacials de Catalunya (IEEC), 08034 Barcelona, Spain

²⁵Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain

²⁶Santa Cruz Institute for Particle Physics, Santa Cruz, CA 95064, USA

²⁷Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210, USA

²⁸Department of Physics, The Ohio State University, Columbus, OH 43210, USA

²⁹Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA

³⁰Australian Astronomical Optics, Macquarie University, North Ryde, NSW 2113, Australia

³¹Lowell Observatory, 1400 Mars Hill Rd, Flagstaff, AZ 86001, USA

³²Institució Catalana de Recerca i Estudis Avançats, E-08010 Barcelona, Spain

³³Department of Astronomy, University of California, Berkeley, 501 Campbell Hall, Berkeley, CA 94720, USA

³⁴Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

³⁵Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

³⁶Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, 21029 Hamburg, Germany

³⁷Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

³⁸Kavli Institute for Particle Astrophysics & Cosmology, P. O. Box 2450, Stanford University, Stanford, CA 94305, USA

³⁹School of Physics and Astronomy, University of Southampton, Southampton, SO17 1BJ, UK

⁴⁰Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831

Corresponding author: Robert Morgan
robert.morgan@wisc.edu

ABSTRACT

Large-scale astronomical surveys have the potential to capture data on large numbers of strongly gravitationally lensed supernovae (LSNe). To facilitate timely analysis and spectroscopic follow-up before the supernova fades, an LSN needs to be identified soon after it begins. To quickly identify LSNe in optical survey datasets, we designed ZipperNet, a multi-branch deep neural network that combines convolutional layers (traditionally used for images) with long short-term memory (LSTM) layers (traditionally used for time series). We tested ZipperNet on the task of classifying objects from four categories – no lens, galaxy-galaxy lens, lensed type Ia supernova, lensed core-collapse supernova – within high-fidelity simulations of three cosmic survey data sets – the Dark Energy Survey (DES), Rubin Observatory’s Legacy Survey of Space and Time (LSST), and a Dark Energy Spectroscopic Instrument (DESI) imaging survey. Among our results, we find that for the LSST-like dataset, ZipperNet classifies LSNe with a receiver operating characteristic area under the curve of 0.97, predicts the spectroscopic type of the lensed supernovae with 79% accuracy, and demonstrates similarly high performance for LSNe 1-2 epochs after first detection. We anticipate that a model like ZipperNet, which simultaneously incorporates spatial and temporal information, can play a significant role in the rapid identification of lensed transient systems in cosmic survey experiments.

Keywords: Optical astronomy – Machine learning – Transient sources – Gravitational lensing

1. INTRODUCTION

Strong gravitational lensing is an expansive probe of both astrophysics and cosmology. In systems with strong lensing, light from a background object (the source) is deflected by the gravitational potential of an interposed foreground object (the lens), producing characteristic features such as arcs, Einstein rings, and/or multiple images (Treu 2010). A key subclass of lenses, strongly lensed transients, have time-variable brightness in the source galaxy. Due to the cosmological distances involved in strong lensing, the most common transient objects to observe lensed are quasars, which vary in brightness on the time scales of several years (Hook et al. 1994; Helfand et al. 2001), and supernovae, which can reach peak brightness within days and then dim over the course of weeks to months (Mihalas 1963).

One of the principle features of lensed transients is the time delay between arrival of photons that take different paths around the lenses. Photons from multiple images of source objects travel different distances to Earth and experience different magnitudes of the gravitational potential due to the geometry of the lensing system. Given the constant speed of light, the differences in the path lengths and gravitational potentials traversed by the photons produce an offset in the arrival times for photons that are emitted at the same point in the source’s lightcurve. Therefore, sources with time-varying brightness in these systems exhibit an observable time delay between the individual images of the lensed source (Refsdal 1964).

Strongly lensed transients are particularly useful for a wide range of investigations. For example, the magni-

fication of these background sources and their environments can reveal new information on the eruption and prevalence of these objects at earlier times in the Universe (Treu & Marshall 2016). Time-delay cosmography (TDC) is a technique that uses the time-varying brightness of compact systems that have undergone strong gravitational lensing (SL) to perform a geometric measurement of H_0 (Refsdal 1964). TDC entails a measurement of this time delay and modeling of the full strong lensing system. H_0 is then inversely proportional to the time delay between photons. This technique has been utilized with quasars (persistent variable objects) to measure H_0 to better than five percent precision in time-varying SL systems (Shajib et al. 2020; Wong et al. 2019).

The cosmic expansion rate today H_0 is a critical parameter for understanding the evolution of the Universe. There are multiple probes of H_0 , including the cosmic distance ladder, extrapolation from the cosmic microwave background (CMB), and strongly lensed variable sources, like supernovae and quasars. There is currently a significant tension amongst the probes (Freedman 2021), particularly between early-Universe probes like the CMB (Planck Collaboration et al. 2020) and late-Universe probes like the type-Ia SNe distance ladder anchored with cepheid variable stars (Riess et al. 2021). Time-delay cosmography, as probe of H_0 , does not rely on anchoring measurements to late-Universe objects or extrapolating from early-Universe physics, so it offers a new perspective on the expansion rate of the Universe today.

Aside from quasars, supernovae (SNe) are the other major class of common time-varying sources that are

Dataset	Gain [e ⁻ /count]	Read Noise [e ⁻]	Pixel Size [arcsec]	Exp. Time [seconds]	Seeing [FWHM arcsec]	Cadence [days]	Reference
DES-wide	6.083	7.0	0.263	90	0.91 ± 0.12	220 ± 191	Abbott et al. (2018)
LSST-wide	2.3	10.0	0.2	30	0.71	12 ± 5	Marshall et al. (2017)
DES-deep	6.083	7.0	0.263	200	1.04 ± 0.44	6 ± 1	Abbott et al. (2019)
DESI-DOT	6.083	7.0	0.263	60	0.90 ± 0.12	3	Abbott et al. (2018)

Table 1. Summary of the instruments and observational procedures emulated in the simulated data sets: camera properties, observing conditions, and survey cadence. We show only the i band for each property, because it illustrates the key discerning features between the datasets. The cadence displays the mean and standard deviation of the intra-band time separation. The full cadence information and data quality properties are available in the `deepLenstronomy` input files, which accompany this work (Morgan et al. 2021a).

bright enough to be detected at cosmological distances. SNe present an experimental challenge in TDC because they become bright and fade on the scale of months, therefore requiring rapid identification and analysis to obtain a measurement of the time delay. However, the larger variability on shorter timescales compared to quasars offers a competitive advantage in measuring the time delay. Another advantage of LSNe is a common subclass of SNe (SNe-Ia) can facilitate highly accurate modeling of the lensing gravitational potential – and therefore a more precise H_0 measurement (Foxley-Marrable et al. 2018a; Birrer et al. 2021a; Kolatt & Bartelmann 1998; Oguri & Kawano 2003) – as a result of their standardizable brightness (Tripp 1998). To date, only a handful of LSNe have been detected (Kelly et al. 2015; Rodney et al. 2021), and only one LSNe-Ia has been discovered (Goobar et al. 2017).

Large optical surveys are ideal datasets to search for LSNe, since high area coverage and returning to the same field multiple times increase the chances of observing a LSN, though the rarity of LSNe still makes their detection a challenging problem. For example, based on the area covered, imaging depth, and length of observations, only 0.5-2 LSNe are expected to be in DES data (Oguri 2019). Looking forward to the next era of optical survey astronomy, the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) plans to cover the entire southern sky to greater depth and with higher cadence than any predecessor survey — e.g. the Dark Energy Survey (DES; Diehl et al. 2018). Preliminary forecasts indicate that the Rubin Observatory will detect hundreds to thousands of these systems (Goldstein & Nugent 2016; Oguri 2019; Wojtak et al. 2019). Each detected time-varying SL system has the potential to produce an independent measurement of H_0 , meaning that the measured statistical precision on H_0 using this technique will scale with $1/\sqrt{N}$, where N is the number of detected systems. Therefore, one of the main goals in the LSST-era is the identification

and characterization of as many LSNe-Ia as possible to precisely measure H_0 (The LSST Dark Energy Science Collaboration et al. 2018).

Fast and robust algorithms for detecting LSNe-Ia are essential to keep pace with the data stream of LSST and surveys with comparable data size. Furthermore, because supernovae fade after their explosion, they must be identified rapidly to facilitate follow-up observations and more detailed characterization of the system for lensing analyses. One approach to detecting LSNe is to observe known SL systems and wait for SN. This approach leverages existing SNe detection strategies and infrastructure and is expected to detect all SNe in known SL systems in the southern hemisphere. However, the Rubin Observatory’s LSST will probe deeper than any previous optical survey. It is expected that a large population ($\sim 10\%$) of all SL systems will have a source galaxy that is too faint to have been detected previously, and they will be missing from the list of target systems (Ryczanowski et al. 2020). Another approach leverages the standardizable brightness of SNe-Ia and proposes the search for brighter than expected (due to lensing magnification) SNe near elliptical galaxies (Goldstein & Nugent 2016). This magnitude threshold technique is expected to detect ~ 500 LSNe-Ia, but is specific to elliptical lens galaxies and will miss LSNe-Ia whose date of peak brightness do not align with the LSST cadence.

In this work, we introduce a deep neural network architecture designed to identify LSNe-Ia without the requirements of targeting known SL systems, elliptical galaxies, or observing the peak brightness of the SN. Deep learning algorithms have been highly successful, both in terms of accuracy and speed, in the fields of image-based SL system detection (Jacobs et al. 2019, among several others) and light-curve-based SNe classification (Möller & de Boissière 2019, among several others). A convolutional neural network (CNN; LeCun et al. 1989) is a kind of deep algorithm that slides learnable matrix operators along images to emphasize

or de-emphasize characteristic shapes, such as lensing arcs. Recurrent neural networks (RNNs; Elman 1990) can model sequences of data, such as lightcurves, to make classifications based on how the data vary in time. Within the RNN architecture, long short-term memory (LSTM) networks (Hochreiter & Schmidhuber 1997), which introduce cells with learnable pathways for information to be passed along, have demonstrated improved performance in sequence characterization. In our approach, we combined a CNN and an LSTM as two branches and optimally utilize the information from both branches — an architecture that leverages the data structures of each input and is tailored to the challenge of immediate LSNe-Ia identification.

We present this work as follows: In Section 2, we describe the simulations used for training and testing ZipperNet, the data processing procedure for utilizing both spatial and time series information, and the architecture of ZipperNet. In Section 3, we describe the results from applying ZipperNet to four simulated datasets that emulate modern survey data products. In Section 4, we discuss the performance of ZipperNet with respect to the dataset properties and the network architecture. We conclude in Section 5.

2. METHODS

2.1. Data Simulation

We simulated images of astronomical strong lensing systems with the open-source software package, `deeplens` (Morgan et al. 2021b), which is built around `lenstronomy` (Birrer & Amara 2018; Birrer et al. 2021b), a widely used package that performs gravitational lensing calculations, modeling, and simulations in a variety of contexts. `deeplens` provides additional features that are important for the era of large-scale surveys and deep learning studies of strong lenses – e.g., image and SN injection, probability distribution sampling, and realistic observing conditions.

2.1.1. Survey Emulation

We simulated four datasets, distinguished by their camera specifications, observing conditions, and cadence, each emulating a distinct modern or next-generation cosmic survey: one wide- and one deep-field dataset for DES – DES-wide and DES-deep, respectively; a wide-field LSST dataset (LSST-wide), and a three-day cadence Dark Energy Camera (DECam; Flaugher et al. 2015) dataset similar to the Dark Energy Spectroscopic Instrument DECam Observation of Transients (DESI-DOT) program (Palmese & Wang; DECam Proposal 2021A-0148). All datasets utilize the g , r , i , and z optical filters and include 45-by-45-pixel

images. The DES-wide, DES-deep, and DESI-DOT datasets simulate images from the DECam, which connects the pixel size, gain, and read noise across those datasets. The LSST-wide dataset simulates LSSTCam (Stalder et al. 2020) images, which have similar but slightly adjusted values for the camera properties.

The DES-wide and DESI-DOT datasets both use the real observing conditions (seeing and sky brightness) from the DES wide-field survey (Abbott et al. 2018). For DES-wide and DESI-DOT, the exposure times are 90 and 60 seconds, respectively. The DES-deep dataset has different seeing, sky brightness, zeropoint, and exposure times chosen for the DES SN program (Abbott et al. 2019). In general, these exposure times are on the order of 200 seconds, but the acceptable seeing criteria can be worse than the DES wide-field survey. The LSST-wide observing conditions are estimated from simulations of the first year of the survey and utilize 30-second exposures (Marshall et al. 2017).

Each dataset has a specific and distinct cadence, and the density of observations significantly affects the analysis in this work. The LSST-wide, DESI-DOT, and DES-deep datasets contain a baseline of 14 epochs per band. While the LSST main survey cadence is still being designed at the time of this writing, our fiducial 14-epoch data sequences are sufficiently short that they are obtainable from both the “baseline” and “rolling” cadences that are under consideration for the survey. The DES-wide dataset contains seven exposures in each band spread over 5.5 years to match the real survey. We sampled the observation times of several fields from the DES footprint to generate the `deeplens` simulations. The LSST-wide cadence is estimated using several realizations of an intra-band spacing of 12 ± 5 days over a three-month period (Marshall et al. 2017). The DES-deep cadence is estimated using several realizations of an intra-band spacing of 6 ± 1 days over a one-month period (Abbott et al. 2019). Lastly, the DESI-DOT cadence is an exposure in each band every three nights over a one-month period.

We seek to avoid avoid jargon confusion between astronomy and machine learning contexts with respect to the term, “epoch.” In this work, “epoch” refers to one astronomical exposure or data-collection period. When discussing neural network training steps, we use the term “training iteration” instead of the traditional “epoch” that is used in machine learning. The data sets are summarized in Table 1. All `deeplens` input files for this analysis are accessible for reproduction of the datasets in Morgan et al. (2021a).

2.1.2. Object, System, and Population Simulation

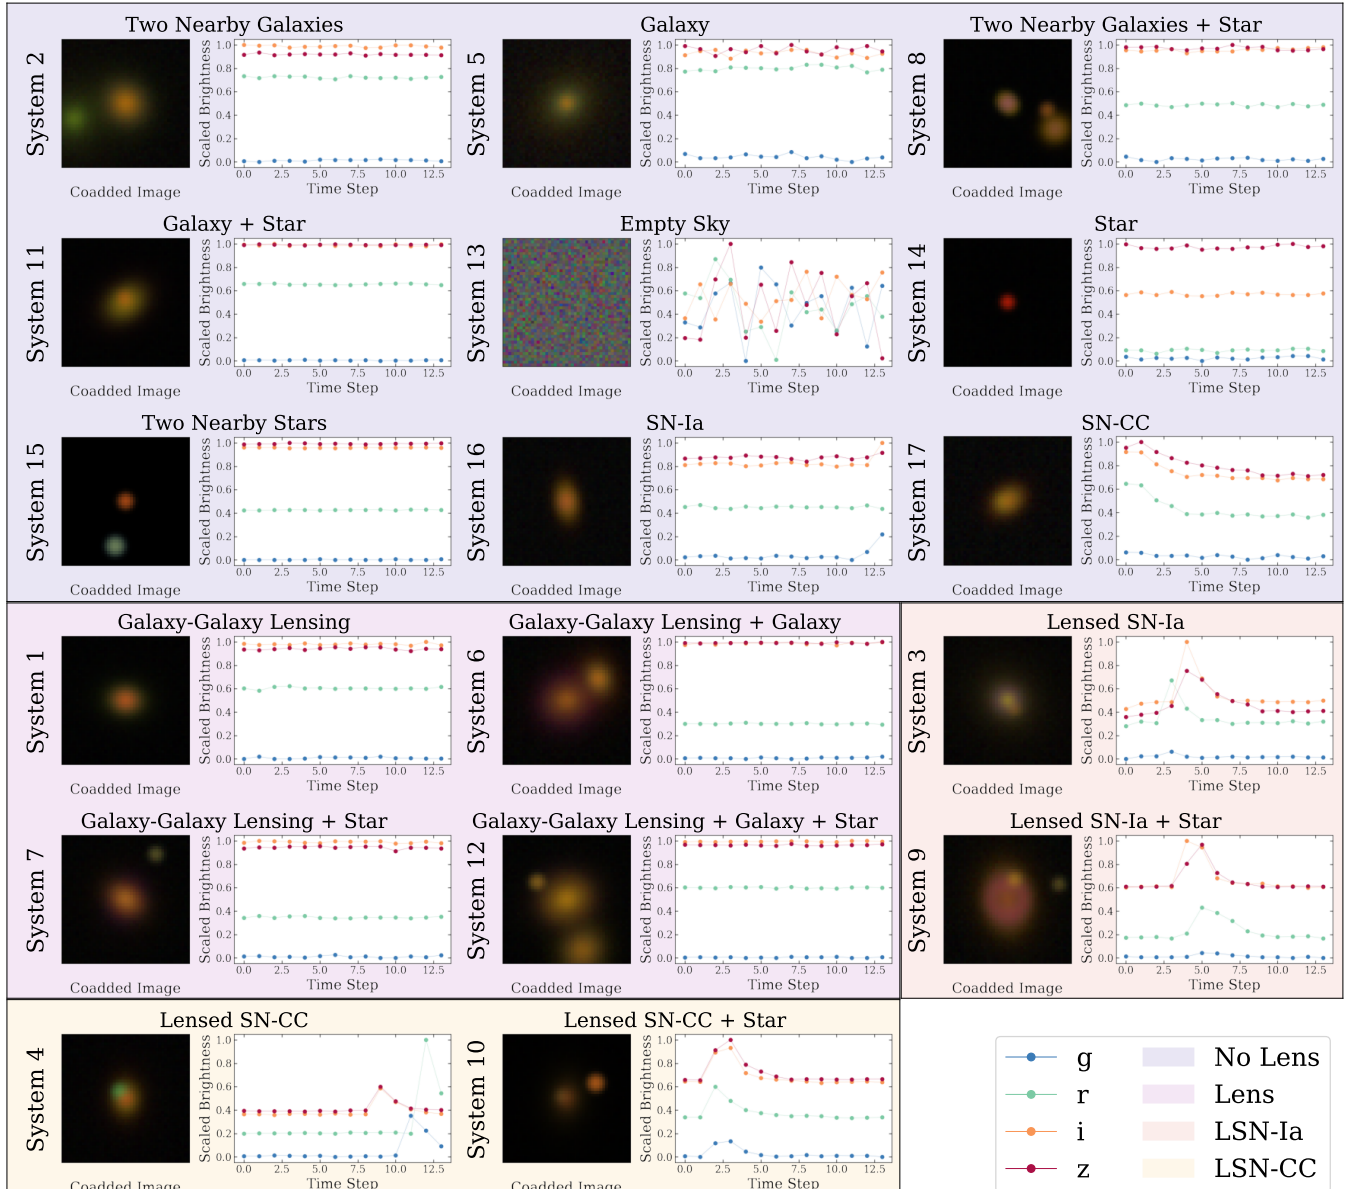


Figure 1. Examples of the 17 simulated systems from the LSST-wide dataset grouped into No Lens (blue), Lens (magenta), LSN-Ia (orange), and LSN-CC (yellow) classes. Each example displays a coadded composite image of the gri bands and the extracted lightcurve from our processing in Section 2.2. Each time step in the LSST-wide dataset is approximately 12 days.

In total, we simulate 17 different types of astronomical systems to reflect the diversity of systems that classifiers are likely to encounter when applied to observed optical survey data: (1) one galaxy behind one foreground galaxy; (2) two galaxies at the same redshift and with small angular separation (1 to 4 arcseconds); (3) one galaxy (with one SN-Ia) behind one foreground galaxy; (4) one galaxy (with one SN-CC) behind one foreground galaxy; (5) one galaxy; (6) two galaxies (with small angular separation and at the same redshift) in front of one background galaxy that is at a higher redshift; (7) one galaxy behind one foreground galaxy that has one

star from the Milky Way in the image cutout; (8) two galaxies with small angular separation and at similar redshifts that has one star from the Milky Way in the image cutout; (9) one galaxy (with one SN-Ia) behind one foreground galaxy that has one star from the Milky Way in the image cutout; (10) one galaxy (with one SN-CC) behind a foreground galaxy that has one star from the Milky Way in the image cutout; (11) one galaxy that has one star from the Milky Way in the image cutout; (12) two galaxies with small angular separation and at similar redshifts in front of one background galaxy at a higher redshift and one star from the Milky Way in the

image cutout; (13) empty sky; (14) one Milky Way star; (15) two Milky Way stars; (16) one galaxy with one SN-Ia; and (17) one galaxy with one SN-CC. Figure 1 shows sample images and time series of the 17 systems.

To further enhance realism, the properties of all simulated objects are drawn from real data: all inherent physical correlations of these parameters are included in our dataset. First, a galaxy that enters the simulations as the lens has properties drawn from a population of $\sim 2,000$ observed galaxies. The velocity dispersion and spectroscopic redshift were measured by the Sloan Digital Sky Survey (SDSS; York et al. 2000). We obtain a color-independent ellipticity, as well as a band-wise half-light radius, Sersic profile index, magnitude from DES Year 1 data (Tarsitano et al. 2018; Abbott et al. 2018). Next, a putative source galaxy, whether used in a lensing or non-lensing system, draws its properties from a population of $\sim 500,000$ galaxies measured by DES (Abbott et al. 2021). The band-wise magnitudes of foreground Milky Way stars were also drawn from DES data (Abbott et al. 2021). Finally, the SNe were injected using public SN spectral energy distributions (Kessler et al. 2010) available in `deeplens`, which redshifts the distribution and calculates the observed magnitude in each band. The injected SN reaches peak brightness anytime between 20 days before the first observation and 20 days after the final observation, so the dataset contains falling lightcurves, rising lightcurves, and complete lightcurves. We do not include the effects of microlensing in our simulated dataset, because it is expected to be small compared to the change in brightness observed from a SN (Foxley-Marrable et al. 2018b).

For all four survey emulation datasets, we simulate the same strong lensing systems: all strong lensing systems are emulated in all four of the cosmic survey contexts. While our simulated datasets are subject to selection biases from the detection limits of DES and SDSS for source and galaxies, respectively, the data nonetheless contain realistic collections of object properties, which enables the validation of our deep learning detection method on realistic survey data.

2.2. Data Processing

`deeplens` emulates observational surveys by producing a time series of images. With the exception of small effects from observing conditions, images in a series will be approximately identical, because astronomical objects are approximately stationary on month-long timescales. Even in the case of a SN, the primary difference is the presence of one or more point sources in some of the images in the time series. We condensed the image information to single-image input for ZipperNet

by averaging all images in the time series on a pixel-by-pixel basis within each band. This processing reduces noise fluctuations from the observing conditions while preserving the presence of SNe, thus increasing the overall signal-to-noise ratio of the image and making faint objects more visible. After averaging the images, the pixel values of the mean images are scaled to range from 0 to 1 on a per-example basis to preserve color relationships.

To concisely characterize the temporal behavior of a time series of images and to avoid relying on source identification or deblending algorithms, we follow a process that reflects a standardized background-subtracted aperture flux measurement in astronomy. We measure the signal (S) and background (B) to extract a background-subtracted brightness (S – B) of individual images within a predefined circular aperture at the center of each image. In equation form, this process can be expressed as

$$S = \sum_{i,j}^N W_{i,j} X_{i,j}, \quad (1)$$

$$B = \text{median}_{W_{i,j}=0} [X_{i,j}] \times \sum_{i,j}^N W_{i,j}, \quad (2)$$

$$\text{Brightness} = S - B, \quad (3)$$

where i and j index the row and column of the image pixels, N gives the number of pixels along one dimension of the images, X is an image, and W is the aperture. $W_{i,j}$ is zero outside the aperture and one inside the aperture. In this work, the circular aperture has a 20-pixel radius, which corresponds to 5.26 arcseconds for DECam and 4.0 arcseconds for LSST-Cam – both much larger than typical galaxy-scale lens Einstein radii, which are approximately in the range [0.5, 1.2] arcsec. For processing within the neural networks, we again scale the extracted brightness to between 0 and 1 on a per-example basis.

A byproduct of this process is the significant increase in noise in the photometry measurements; we find, however, that this effect does not hinder the deep learning methods. After the averaging images and extracting lightcurves, each example input to ZipperNet is a 45-pixel-by-45-pixel image in each band and a lightcurve of the extracted brightness at each time step in each band. The operations required to extract the photometry of the systems have little computational cost and can be easily broadcasted, which is a great benefit considering the scale of modern astronomical datasets.

Lastly, we define a classification scheme for our 17 simulated systems. We construct a four-class problem, where the classes are “No Lens,” “Lens,” “LSNe-Ia,”

and “LSNe-CC,” as shown in Figure 1. The No Lens class collects the cases where there is no gravitational lensing present – labeled as 2, 5, 8, 11, 13, 14, 15, 16, and 17. The Lens class collects cases where there is gravitational lensing but no SNe in the background galaxy – labeled as 1, 6, 7, and 12. The LSNe-Ia and LSNe-CC classes collect cases with gravitational lensing and a SN in the background galaxy – labeled as 3 and 9 and as 4 and 10, respectively. Each of the four classes contain 1,250 examples, with equal representation of the individual constituent cases. To augment the datasets and increase the size of each class eightfold to 10,000, we rotated and mirrored the images; the lightcurve was unaffected due to the circular aperture extraction method. When we split the datasets into smaller training and validation datasets, none of the examples in the validation dataset are rotated or mirrored versions of objects in the training dataset: the two datasets are rotated and mirrored independently.

2.3. ZipperNet

Our ZipperNet architecture is designed to treat the image-based information and the lightcurve-based information on equal footing. On one branch, the images are passed through convolutional layers, flattened into one dimensional arrays, and condensed in size. On another branch, the lightcurves are passed through recurrent layers composed of LSTM cells and flattened into one dimensional arrays. The flattened images and lightcurves output by each branch are condensed to equal sizes, concatenated, and then reduced to a single classification. By zipping convolutional layers and recurrent layers into one coherent deep learning architecture (in joining the branches), the training of the network will optimize weights in both types of layers simultaneously. The architecture is illustrated in Figure 2, and the specifications of each layer are presented in Table 2. All deep learning code in this analysis utilizes the PyTorch (Paszke et al. 2019) library.

For each of the four datasets in our analysis, we trained an individual ZipperNet: we trained on 90% (9,000 samples) of the simulated data and used the remaining 10% (1,000 samples) for validation. We chose a batch size of five for the training because LSTM cells generally perform better when processing smaller amounts of information at the same time. We also utilized categorical cross entropy as the loss function of the network and a learning rate of 0.001 with the Adam (Kingma & Ba 2017) optimizer. As the network trained, we monitored the accuracy (the number of correct classifications divided by the total number of samples) for the training and validation datasets. In each case, the

Layer	Specifications
conv1 [†]	Conv2D — (k : 15, p : 2, s : 3) — (4 → 48)
conv2 [†]	Conv2D — (k : 5, p : 2, s : 1) — (48 → 96)
maxpool1	MaxPool2D (k : 2)
dropout1	Dropout2D (f : 0.25)
fc1 [†]	Fully-Connected (3456 → 408)
dropout2	Dropout (f : 0.5)
fc2 [‡]	Fully-Connected (408 → 25)
lstm1	LSTM (h : 128)
lstm2	LSTM (h :128)
fc3 [‡]	Fully-Connected (128 → 25)
fc4 [†]	Fully-Connected (50 → 8)
fc5 [‡]	Fully-Connected (8 → 4)

Table 2. ZipperNet layer specifications. We use the following shorthand: kernel size (k), padding (p), stride (s), dropout fraction (f), and hidden units (h). Arrows indicate the change in the size of the data representation as it is passed through the layer. “[†]” indicates a Rectified Linear Unit (ReLU) activation function. “[‡]” indicates a LogSoftmax activation function

training and validation accuracy plateaued after ~ 10 training iterations, but we allowed the training to continue for 40 training iterations. The fully trained network is chosen as the point during training with the highest validation accuracy. We use the term “training iteration” in place of the traditional “epoch” to avoid confusion with the astronomy term “epoch” utilized in other parts of this analysis. The accuracy for the training and validation sets for each of the four datasets is presented in Table 3.

3. RESULTS

In general, we find that the ZipperNet model is capable of identifying SL systems, identifying the presence of an SN within SL systems, and distinguishing between LSN-Ia LSN-CC. Furthermore, we demonstrate that ZipperNet can perform these classifications on simulated datasets across wide ranges of depth, observing conditions, and cadence. We use a four-class confusion matrix to compare the predicted and true labels resulting from ZipperNet’s execution (see Figure 3). In all the matrices, the strong representation along the main diagonal indicates correct classifications. ZipperNet had the lowest performance on the DES-wide data, and we discuss this result in Section 4.

There are multiple primary sources of confusion. Confusion between the No Lens and Lens classes is likely due to pixel-based features being difficult to distinguish from the images. This confusion is the strongest within




Figure 2. A diagram of the ZipperNet architecture. Two convolutional layers (orange) receive *griz* images while two LSTM layers (blue) receive extracted *griz* lightcurves. Fully-connected layers (purple) process the flattened and concatenated outputs. Nonlinear activation functions are indicated by a darkened band at the right edge of a layer. Table 2 displays the layer specifications. The final output is a length-4 array, with each element representing the score of one of the four classes in the dataset. This visualization was made with the `PlotNeuralNet` library (Iqbal 2018).

the DES-wide dataset, so we attribute this behavior to the optical depth of the images, since the DES-wide dataset is the shallowest dataset simulated and faint source galaxies would become more difficult to identify. In the deeper datasets, this confusion is caused by exposures with high seeing or systems with small Einstein radii, where in both cases objects blur together. The confusion between the LSNe-Ia and LSNe-CC classes is likely due to the difference in cadence and seeing in the surveys. DES-deep, LSST-wide, and DESI-DOT are much higher cadence datasets than DES-wide (See Table 1), indicating that more densely sampled SN lightcurves are easier for the LSTM cells to classify. A comparison of the DES-deep and DES-wide confusion matrices indicates that there are situations where cadence can be more important than seeing (specifically noting the LSNe-Ia versus LSNe-CC confusion), since the DES-wide dataset had better seeing than DES-deep, but that these situations require dramatic differences in cadence density. Furthermore, the LSST-wide and DESI-DOT datasets have much better seeing than the DES-deep dataset, which shows the importance of being able to resolve spatial features when making classifications. We

Dataset	Training Acc.		Validation Acc.	
	Iter. 10	Iter. 40	Iter. 10	Iter. 40
DES-wide	0.496	0.503	0.473	0.487
LSST-wide	0.750	0.838	0.709	0.785
DES-deep	0.550	0.613	0.535	0.573
DESI-DOT	0.737	0.813	0.680	0.735

Table 3. Training and validation accuracy for ZipperNet on each of the four datasets (rows) at training iterations 10 and 40 (columns).

discuss the importance of the cadence and seeing in more detail in Section 4.

In practice, a general LSN identifier is itself a useful tool: LSNe-CC can be utilized for time-delay cosmography measurements, though they offer less precision on the final H_0 measurement. If we reframe this classification scheme from a four-class problem to a two-class problem (No Lens and Lens in one class and LSNe-Ia and LSNe-CC in second class), the performance is boosted. Figure 4 shows a Receiver Operating Char-

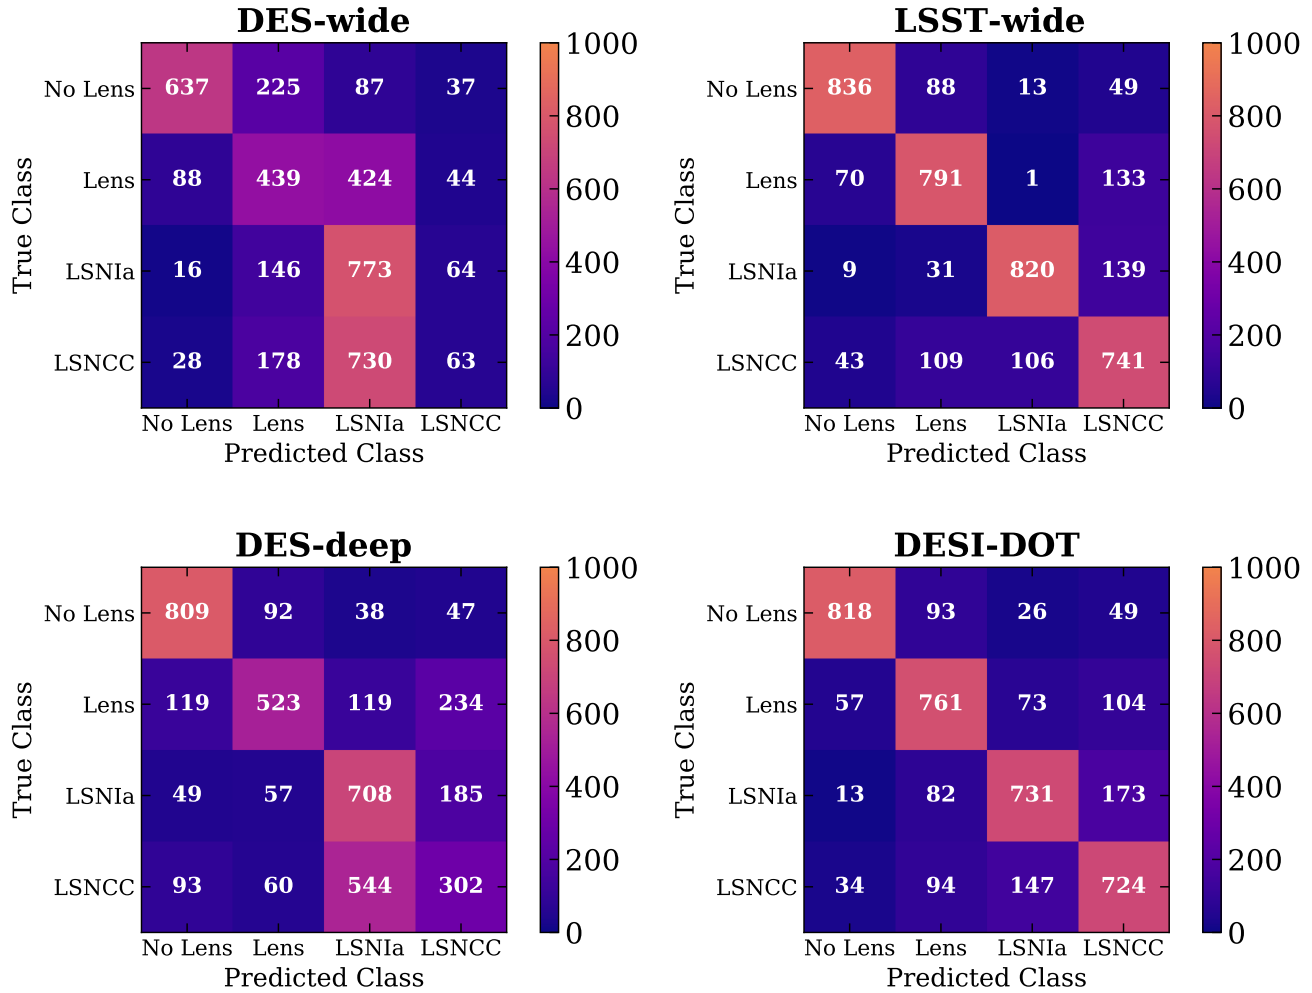


Figure 3. Confusion matrices from applying a trained ZipperNet to the DES-wide, LSST-wide, DES-deep, and DESI-DOT datasets. The matrices display the predictions of ZipperNet as a function of the true classes of the examples for the validation datasets.

acteristic (ROC) curve for the DES-wide, LSST-wide, DES-deep, and DESI-DOT datasets. ROC curves are standard tools for assessing the predictive power of a classifier by calculating the false positive rate and true positive rate at all possible probabilities output by the classifier. A perfect classifier will have an Area Under Curve (AUC) of 1.0 while a classifier that guesses randomly will have an AUC of 0.5. ZipperNet shows high performance when classifying LSNe versus everything else, and this high performance extends across the LSST-wide, DES-deep, and DESI-DOT datasets. This result can also be interpreted from the confusion matrix (Figure 3), where there is little confusion between the LSNe classes and non-LSNe classes. We comment on this high performance in the context of ZipperNet’s architecture in Section 4.

We also estimate the baseline true-positive rate and false-positive rate for LSNe using this technique in the different datasets. When applying the ZipperNet technique to real data, we would expect real data to be used in the training and validation, which would produce more accurate estimates. That being said, we can initially report a LSNe true (false) positive rate of 90.2 % (29.6 %) for DES-wide, 87.0 % (21.9 %) for DES-deep, 91.5 % (9.8 %) for LSST-wide, and 89.0 % (12.6 %) for DESI-DOT. The true-positive rate of approximately 91.5 % for LSST-wide is the highest recovery rate of LSNe predicted to date. Furthermore, with relatively low false-positive rates, we do not expect the data stream of the LSST to be overwhelmed by other astronomical systems being incorrectly labeled as LSNe.

In the two-class problem, we performed additional analysis to interpret the features ZipperNet identified for

making classifications. Figure 5 displays examples from the LSST-wide dataset arranged into groups of correctly classified LSNe (true positives), other astronomical systems correctly classified as not LSNe (true negatives), other astronomical systems erroneously labeled as LSNe (false positives), and LSNe that were missed (false negatives). These examples were found to be representative of the general relationship between the properties of objects and the predictions made by ZipperNet. ZipperNet is able to correctly classify LSNe when the Einstein radius is small ($0.5 - 1.0$ arcsec) and with foreground stars in the image, both of which would trouble a standalone CNN. Upon close inspection, the true-positive images display galaxies with non-uniform light profiles, hinting at the presence of lensing, but the dominating feature is clearly the large fluctuation in brightness detected by the LSTM. The true negatives show systems with evidence of lensing as well, but this time there is no temporal behavior to indicate the presence of a SN. The false positives contain images of lensing or crowded fields with small, but non-negligible coherent time-varying behavior; they in general contain both spatial and temporal features similar to the LSNe class. Lastly, the false negatives are the most important group to understand due to the rareness of LSNe. In some cases, the presence of a star in the aperture used to extract the scaled brightnesses can be bright enough to obscure the change in brightness of the LSNe. Similarly, if the source galaxy is distant and the alignment of the lensing system does not produce sufficient magnification, the imaging may not be deep enough to see a lensed source or a background SNe. Both these cases of false negatives demonstrate difficult to detect systems with LSNe and point toward a physically-motivated selection as opposed to inaccurate feature representations learned by ZipperNet. In general, we find that ZipperNet learns features we would *a priori* expect.

4. DISCUSSION

In general, ZipperNet can identify SL systems, identify LSNe within those systems, and classify the LSNe as LSNe-Ia or LSNe-CC. The variance in the performance for across different datasets indicates a correspondence between data quality and LSN identification power. The DES-deep and DESI-DOT datasets have slightly higher cadences (more samples within a time series) than the LSST-wide dataset, but the ZipperNet accuracy was considerably higher for the LSST-wide dataset, which has higher seeing (lower image quality). The DES-deep dataset emulates that of the DES SNIa observing program, in which exposures were collected on nights with slightly poorer (higher) seeing to optimize




Figure 4. Receiver Operating Characteristic (ROC) curves for the LSST-wide, DES-deep, and DESI-DOT datasets. The ROC curves are calculated for the two-class problem of LSNIa and LSNCC versus No Lens and Lens. An Area Under Curve (AUC) of 1.0 indicates perfect performance, while an AUC of 0.5 indicates random guessing.

the seeing of the exposures used for weak lensing measurements. In DES data processing and SNIa analysis, the difference-imaging (Kessler et al. 2015) and scene-modeling (Brout et al. 2019) techniques can nevertheless detect and measure SNe when seeing is up to $2''$. We bypassed those time-consuming techniques with our circular aperture extraction method for the lightcurve brightnesses. Without those techniques, the performance of ZipperNet is degraded, because the higher seeing in the DES-deep dataset obscures image patterns that would otherwise be detectable in LSST-wide and DESI-DOT datasets.

A second data quality factor in ZipperNet’s predictive power is the cadence of the observations. The DES-wide dataset has excellent depth and seeing, but a low-density time sampling: ZipperNet fails to perform at similar levels to the other datasets. With DESI-DOT, which has high cadence and similar depth and seeing as DES-wide, ZipperNet was able to learn the underlying features of the four classes extremely well. Overall, we find that together seeing of $\lesssim 1.2''$ (corresponding to typical upper limits on Einstein radii of galaxy-scale lenses) and a cadence with intra-band spacing of $\lesssim 15$ nights (corresponding to roughly the timescale for SNe evolution) can also improve performance.

Nevertheless, even when ZipperNet had confusion between LSNIa and LSN-CC (primarily caused by deficient cadence of the dataset) or confusion between the No Lens and Lens classes (primarily caused by seeing

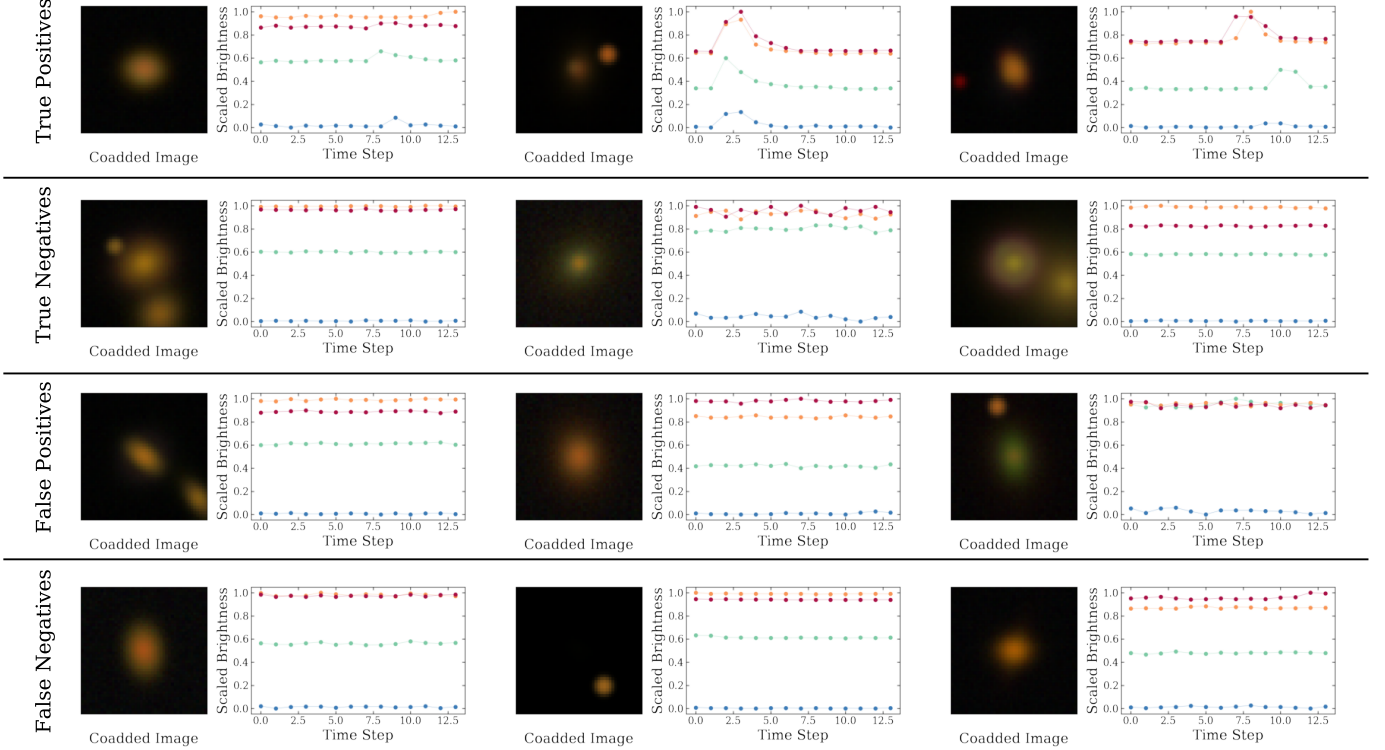


Figure 5. Examples of classifications made by ZipperNet from the LSST-wide dataset, grouped by true positives, true negatives, false positives, and false negatives. The color scheme for the bands used in the lightcurves is the same as Figure 1.

greater than typical Einstein radii), ZipperNet performs extremely well as a LSNe finder. Reducing the classification to two classes – LSN and everything else – shows high performance for the LSST-wide, DES-deep, and DESI-DOT datasets (Figure 4). The ZipperNet architecture as a LSNe finder in this setting does not suffer from the dependence on seeing observed in the four-class problem and is slightly less dependent on the cadence. By balancing the image and temporal inputs, ZipperNet finds weightings and combinations of the two data products optimal for LSNe detection. We interpret this result as a demonstration of a key strength of the ZipperNet architecture.

Finally, we compare the ZipperNet architecture to a standalone “RNN,” a standalone “CNN,” and a combination of those two standalone networks (“COMBO”) in the context of the two-class problem and the LSST-wide dataset. The standalone networks are identical in structure to the corresponding constituents of ZipperNet and trained under the same conditions. The combination classifier does not connect the feature representations of the standalone networks internally: it merely uses the outputs from each of the standalone classifiers, requiring them to individually identify an object as a LSN. We expect that the standalone CNN will identify systems with lensing, while the standalone RNN will identify

systems with SNe: by requiring both a lensing classification and a SN classification, we establish a baseline for the performance of deep learning architectures that are not connected internally in ways similar to ZipperNet’s design. ROC curves for the different classifiers are shown in Figure 6. The CNN performs well and mostly identifies all systems with lensing or point sources in galaxies, but this performance can still result in high numbers of false positives (due to the rarity of LSNe) in practical applications. The RNN is effectively a SN identifier in this context: the vast majority of the false positives come from classifying unlensed SNe (labels 16 and 17) as LSNe. While the ROC AUC is high for the RNN, it alone could not be used as a LSN finder due to the much larger volumetric rate of SNe compared to LSNe. As shown by this test, ZipperNet outperforms both constituent networks (CNN and RNN) and more importantly the combination of its constituent networks (COMBO), indicating that connecting the feature representations of the RNN and CNN internally yields better overall performance and is a worthwhile deep learning strategy for the problem of LSNe detection. The ZipperNet architecture shows the highest performance.

The motivating challenge for ZipperNet is to facilitate the identification of LSNe during the Rubin Observatory’s LSST, which is expected to have a very high data




Figure 6. Receiver Operating Characteristic (ROC) curves for the LSST-wide datasets using different networks. The “COMBO” label is an optimized weighting of the “CNN” and “RNN” network predictions. The ROC curves are calculated for the two-class problem of LSNIa and LSNCc versus No Lens and Lens. An Area Under Curve (AUC) of 1.0 indicates perfect performance, while an AUC of 0.5 indicates random guessing.

stream rate compared to previous large-scale surveys. The high data stream rate is not challenge for ZipperNet because classifications can be parallelized and preprocessing is economical (because we directly extract band-wise lightcurves from images without the need deblending or expensive photometric analysis). The remaining test of ZipperNet for Rubin Observatory main survey operations is detection at various epochs into the SN light curves for community alerts: How early into the light curve can ZipperNet find a LSN? We present in Figure 7 the LSN true-positive rates as functions of the number of lightcurve epochs after the first detection. In this context, an r -band magnitude brighter than 24.4 mag serves as the first detection, which is motivated by the LSST science requirements (The LSST Dark Energy Science Collaboration et al. 2018). We find that even when only one or two epochs are present in the lightcurve after the first detection, the true-positive rate is $> 75\%$, and it improves as more epochs are added. After five post-detection observing epochs in the light curve, the LSN identification has a true-positive rate $> 95\%$. Furthermore, even LSNe fainter than the detection threshold are identifiable with a true-positive rate of $\sim 80\%$, indicating high performance where other detection methods would lose sensitivity: both proposed methods discussed in Section 1 rely on SNe being brighter than the detection threshold and realized as in-




Figure 7. The true-positive rates for LSNe-Ia and LSNe-CC as functions of lightcurve phase for the LSST-wide dataset. The threshold used for the first detection is 24.4 mag in the r band based on the single-epoch survey requirements (The LSST Dark Energy Science Collaboration et al. 2018). The “All LSNe” label refers to the two-class problem of general LSNe detection.

dividual objects. For these faint LSNe, ZipperNet likely finds success by identifying systems with image-based strong lensing features and then noticing small changes in brightness. This result supports the expectation that ZipperNet will perform well as a real-time LSN identifier in Rubin Observatory main survey data.

In summary, ZipperNet introduces a new deep learning architecture for lensed transient detection where spatial and temporal features are treated on the same footing in a single framework. This balanced framework produces a ROC curve AUC of 0.97 when identifying LSNe and 79% accuracy at outright identification of LSNe-Ia in LSST wide-field data even in the early phases of the lightcurve. Therefore, we expect ZipperNet to play a large role in the rapid identification of LSNe for spectroscopic characterization and time-delay cosmography during the main survey operations of the Rubin Observatory.

5. CONCLUSION

Detecting LSNe soon after explosion will be an important goal for the Vera C. Rubin Observatory and other high-cadence optical surveys. In this work, we introduced ZipperNet a deep learning tool for LSNe identification. ZipperNet combines a convolutional neural network with a recurrent neural network to simultaneously process spatial and temporal data. We utilized `deeplens` to simulate four distinct optical survey datasets for testing. ZipperNet performed well when the cadence and seeing were LSST-like or better. Specifically, ZipperNet was able to identify LSNe in LSST-like data with a ROC AUC of 0.97 and distinguish LSNe-Ia

from LSNe-CC in LSST-like data with 79% accuracy. With ZipperNet, high-cadence optical surveys can accurately identify both LSNe-CC and LSNe-Ia early into their light curves, and furthermore distinguish between the two transient classes. Thus, we expect ZipperNet to be a powerful tool in identifying lensing systems for time-delay cosmography measurements during the Rubin Observatory main survey operations.

ACKNOWLEDGMENTS

According to the CRediT system we acknowledge the contributions from authors in detail. R. Morgan designed the analysis and deep learning methodologies, prepared simulations, trained networks, and prepared this publication. B. Nord and K. Bechtol offered feedback and guidance on the analysis and publication throughout the project. S. J. Gonzalez prepared merged SDSS / DES catalogs essential for realistic dataset simulation. L. Buckley-Geer, A. Möller, J. W. Park, and A. G. Kim served as internal reviewers within DES. S. Birrer provided useful comments on the draft and analysis. All other authors contributed to DES infrastructure upon which this project was based.

R. Morgan thanks the Universities Research Association Fermilab Visiting Scholars Program for funding his work on this project. R. Morgan also thanks the LSSTC Data Science Fellowship Program, which is funded by LSSTC, NSF Cybertraining Grant #1829740, the Brinson Foundation, and the Moore Foundation; his participation in the program has benefited this work.

We acknowledge the Deep Skies Lab as a community of multi-domain experts and collaborators who've facilitated an environment of open discussion, ideageneration, and collaboration. This community was important for the development of this project.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1744555. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

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, NSF's NOIRLab, the University of Nottingham, The Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, Texas A&M University, and the OzDES Membership Consortium.

Based in part on observations at Cerro Tololo Inter-American Observatory at NSF's NOIRLab (NOIRLab Prop. ID 2012B-0001; PI: J. Frieman), which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

The DES data management system is supported by the National Science Foundation under Grant Numbers AST-1138766 and AST-1536171. The DES participants from Spanish institutions are partially supported by MICINN under grants ESP2017-89838, PGC2018-094773, PGC2018-102021, SEV-2016-0588, SEV-2016-0597, and MDM-2015-0509, some of which include ERDF funds from the European Union. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. Research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) including ERC grant agreements 240672, 291329, and 306478. We acknowledge support from the Brazilian Instituto Nacional de Ciência e Tecnologia (INCT) e-Universe (CNPq grant 465376/2014-2).

This paper has gone through internal review by the DES collaboration.

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Software: `astropy` (Astropy Collaboration et al. 2013), `deeplens Astronomy` (Morgan et al. 2021b), `lenstronomy` (Birrer & Amara 2018), `matplotlib` (Hunter 2007), `numpy` (Harris et al. 2020), `pandas` (McKinney et al. 2010), `PlotNeuralNet` (Iqbal 2018), `PyTorch` (Paszke et al. 2019), `Scikit-Learn` (Pedregosa et al. 2011), `scipy` (Jones et al. 2001).

REFERENCES

- Abbott, T. M. C., Abdalla, F. B., Allam, S., et al. 2018, *The Astrophysical Journal Supplement Series*, 239, 18
- Abbott, T. M. C., Allam, S., Andersen, P., et al. 2019, *The Astrophysical Journal*, 872, L30
- Abbott, T. M. C., Adamow, M., Aguena, M., et al. 2021, *The Dark Energy Survey Data Release 2*, arXiv:2101.05765
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, arXiv:1307.6212
- Birrer, S. & Amara, A. 2018, *Physics of the Dark Universe*, 22, 189
- Birrer, S., Dhawan, S., & Shajib, A. J. 2021a, *The Hubble constant from strongly lensed supernovae with standardizable magnifications*, arXiv:2107.12385
- Birrer, S., Shajib, A. J., Gilman, D., et al. 2021b, *Journal of Open Source Software*, 6, 3283
- Brout, D., Sako, M., Scolnic, D., et al. 2019, *The Astrophysical Journal*, 874, 106
- Diehl, H. T., Neilsen, E., Gruendl, R. A., et al. 2018, in *Observatory Operations: Strategies, Processes, and Systems VII*, ed. A. B. Peck, R. L. Seaman, & C. R. Benn, Vol. 10704, International Society for Optics and Photonics (SPIE), 138
- Elman, J. L. 1990, *Cognitive Science*, 14, 179
- Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, *The Astronomical Journal*, 150, 150
- Foxley-Marrable, M., Collett, T. E., Vernardos, G., Goldstein, D. A., & Bacon, D. 2018a, *Monthly Notices of the Royal Astronomical Society*, 478, 5081, <https://academic.oup.com/mnras/article-pdf/478/4/5081/25105265/sty1346.pdf>
- Foxley-Marrable, M., Collett, T. E., Vernardos, G., Goldstein, D. A., & Bacon, D. 2018b, *Monthly Notices of the Royal Astronomical Society*, 478, 5081, <https://academic.oup.com/mnras/article-pdf/478/4/5081/25105265/sty1346.pdf>
- Freedman, W. L. 2021, *Measurements of the Hubble Constant: Tensions in Perspective*, arXiv:2106.15656
- Goldstein, D. A. & Nugent, P. E. 2016, *The Astrophysical Journal*, 834, L5
- Goobar, A., Amanullah, R., Kulkarni, S. R., et al. 2017, *Science*, 356
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357–362
- Helfand, D. J., Stone, R. P. S., Willman, B., et al. 2001, *AJ*, 121, 1872, arXiv:astro-ph/0012442
- Hochreiter, S. & Schmidhuber, J. 1997, *Neural Comput.*, 9, 1735–1780
- Hook, I. M., McMahon, R. G., Boyle, B. J., & Irwin, M. J. 1994, *MNRAS*, 268, 305
- Hunter, J. D. 2007, *Computing in science and engineering*, 9, 90
- Iqbal, H. 2018, *HarisIqbal88/PlotNeuralNet v1.0.0*
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *The Astrophysical Journal*, 873, 111, arXiv:0805.2366
- Jacobs, C., Collett, T., Glazebrook, K., et al. 2019, *The Astrophysical Journal Supplement Series*, 243, 17
- Jones, E. D., Oliphant, T. E., & Peterson, P. 2001, *SciPy: Open Source Scientific Tools for Python*
- Kelly, P. L., Rodney, S. A., Treu, T., et al. 2015, *Science*, 347, 1123, <https://www.science.org/doi/pdf/10.1126/science.aaa3350>
- Kessler, R., Bassett, B., Belov, P., et al. 2010, *Publications of the Astronomical Society of the Pacific*, 122, 1415
- Kessler, R., Marriner, J., Childress, M., et al. 2015, *The Astronomical Journal*, 150, 172
- Kingma, D. P. & Ba, J. 2017, *Adam: A Method for Stochastic Optimization*, arXiv:1412.6980
- Kolatt, T. S. & Bartelmann, M. 1998, *Monthly Notices of the Royal Astronomical Society*, 296, 763, <https://academic.oup.com/mnras/article-pdf/296/3/763/18540536/296-3-763.pdf>
- LeCun, Y., Boser, B., Denker, J. S., et al. 1989, *Neural Computation*, 1, 541
- Marshall, P., Clarkson, W., Shemmer, O., et al. 2017, *LSST Science Collaborations Observing Strategy White Paper: "Science-driven Optimization of the LSST Observing Strategy"*
- McKinney, W. et al. 2010, in *Proceedings of the 9th Python in Science Conference*, Vol. 445, Austin, TX, 51
- Mihalas, D. 1963, *PASP*, 75, 256
- Morgan, R., Nord, B., Bechtol, K., & Jonzález, J. 2021a, *DeepZipper: Code and Simulation Inputs*, DOI: 10.5281/zenodo.5425727
- Morgan, R., Nord, B., Birrer, S., Lin, J. Y.-Y., & Poh, J. 2021b, *Journal of Open Source Software*, 6, 2854
- Möller, A. & de Boissière, T. 2019, *Monthly Notices of the Royal Astronomical Society*, 491, 4277, <https://academic.oup.com/mnras/article-pdf/491/3/4277/31555240/stz3312.pdf>
- Oguri, M. 2019, *Reports on Progress in Physics*, 82, 126901
- Oguri, M. & Kawano, Y. 2003, *Monthly Notices of the Royal Astronomical Society*, 338, L25, <https://academic.oup.com/mnras/article-pdf/338/4/L25/3954312/338-4-L25.pdf>

- Paszke, A., Gross, S., Massa, F., et al. 2019, in *Advances in Neural Information Processing Systems* 32, ed. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Curran Associates, Inc.), 8024
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, 12, 2825–2830
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, 641, A6
- Refsdal, S. 1964, *Monthly Notices of the Royal Astronomical Society*, 128, 307
- Riess, A. G., Casertano, S., Yuan, W., et al. 2021, *The Astrophysical Journal*, 908, L6
- Rodney, S. A., Brammer, G. B., Pierel, J. D. R., et al. 2021, *Nature Astronomy*, 5, 1118
- Ryczanowski, D., Smith, G. P., Bianconi, M., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 495, 1666
- Shajib, A. J., Birrer, S., Treu, T., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 6072, <https://academic.oup.com/mnras/article-pdf/494/4/6072/33212512/staa828.pdf>
- Stalder, B., Reil, K., Claver, C., et al. 2020, in *Ground-based and Airborne Instrumentation for Astronomy VIII*, ed. C. J. Evans, J. J. Bryant, & K. Motohara, Vol. 11447, International Society for Optics and Photonics (SPIE), 86
- Tarsitano, F., Hartley, W. G., Amara, A., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 481, 2018, <https://academic.oup.com/mnras/article-pdf/481/2/2018/25729998/sty1970.pdf>
- The LSST Dark Energy Science Collaboration, Mandelbaum, R., Eifler, T., et al. 2018, The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document, arXiv:1809.01669
- Treu, T. 2010, *ARA&A*, 48, 87, arXiv:1003.5567
- Treu, T. & Marshall, P. J. 2016, *The Astronomy and Astrophysics Review*, 24, 11
- Tripp, R. 1998, *A&A*, 331, 815
- Wojtak, R., Hjorth, J., & Gall, C. 2019, *Monthly Notices of the Royal Astronomical Society*, 487, 3342, <https://academic.oup.com/mnras/article-pdf/487/3/3342/28840832/stz1516.pdf>
- Wong, K. C., Suyu, S. H., Chen, G. C.-F., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 498, 1420, <https://academic.oup.com/mnras/article-pdf/498/1/1420/33755111/stz3094.pdf>
- York, D. G., Adelman, J., Anderson, John E., J., et al. 2000, *The Astronomical Journal*, 120, 1579, arXiv:astro-ph/0006396