DeepZipper: A Novel Deep Learning Architecture for Lensed Supernovae Identification

Robert Morgan
Physics Department
University of Wisconsin-Madison
Madison, WI 53706
robert.morgan@wisc.edu

Brian Nord
Fermi National Accelerator Laboratory
Batavia, IL 60510
nord@fnal.gov

Keith Bechtol
Physics Department
University of Wisconsin-Madison
Madison, WI 53706
kbechtol@wisc.edu

Jimena González
Physics Department
University of Wisconsin-Madison
Madison, WI 53706
sgonzalezloz@wisc.edu

Abstract

The identification of gravitationally lensed supernovae in modern astronomical datasets is a needle-in-a-haystack problem with dramatic scientific implications: discovered systems can be used to directly measure and resolve the current tension on the value of the expansion rate of the Universe today. We hypothesize that the image-based features of the gravitational lensing and the temporal-based features of the time-varying brightness are equally important in classifications. We therefore develop a deep learning technique that utilizes long short-term memory cells for the time-varying brightness of astronomical systems and convolutional layers for the raw images of astronomical systems simultaneously, and then concatenates the feature maps with multiple fully connected layers. This novel approach achieves a receiver operating characteristic area under curve of 0.97 on simulated astronomical data and more importantly outperforms standalone versions of its recurrent and convolutional constituents. We find that combining recurrent and convolutional layers within one coherent network architecture allows the network to optimally weight and aggregate the temporal and image features to yield a promising tool for lensed supernovae identification.

1 Introduction

Astronomical systems in which two galaxies are aligned along the line of sight for an Earthly observer and a star explodes into a supernova in the background galaxy are both rare and powerful probes of the underlying laws of our Universe (Treu & Marshall, 2016). To date, only a handful of lensed supernovae (LSNe) have been observed, however in the upcoming Rubin Observatory's Legacy Survey of Space and Time (LSST), thousands are expected to be detectable (Oguri, 2019). The LSST will produce time sequences of 4-channel color images over 10 years of observations of approximately 37 billion objects (Ivezić et al., 2019), necessitating an automated detection strategy for LSNe. In this letter, we propose a novel deep learning architecture tailored to the problem of LSNe identification.

Fourth Workshop on Machine Learning and the Physical Sciences (NeurIPS 2021).

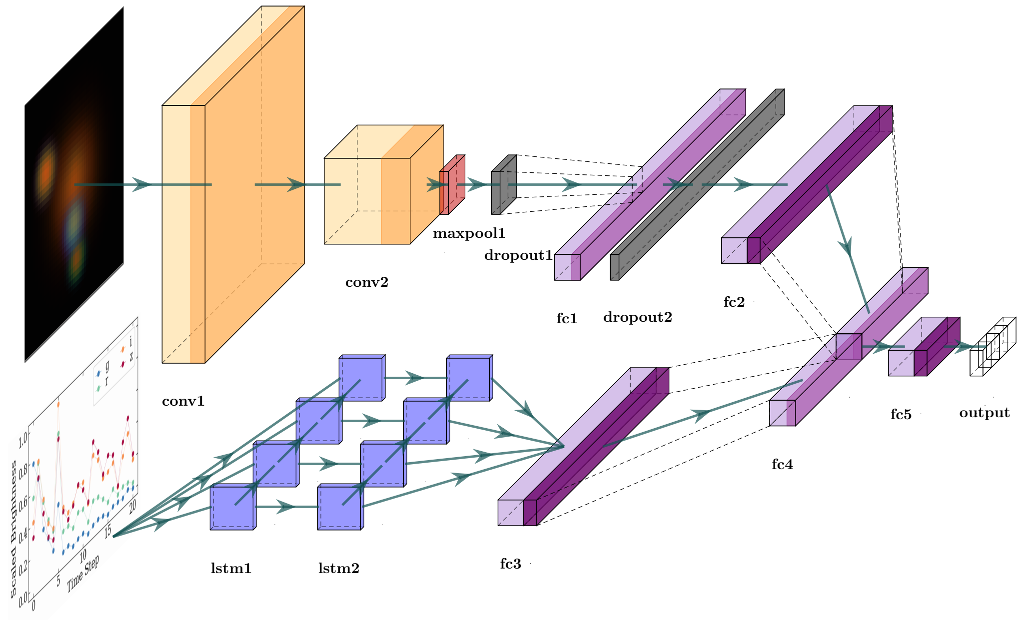

Figure 1: A diagram of the network architecture. Two convolutional layers (orange) receive 4-channel color images while two LSTM layers (blue) receive extracted 4-channel light curves. Fully-connected layers (purple) classify the flattened and concatenated outputs. Activation functions are indicated by a darkened band at the right edge of a layer. Table 1 displays the layer specifications. This visualization was made with the `PlotNeuralNet` library (Iqbal, 2018).

2 Architecture

The network architecture was implemented with `PyTorch` (Paszke et al., 2019) and contains a convolutional branch, a recurrent branch, and a fully connected classifier to weight the extracted features from the two branches. The convolutional branch contains two convolutional layers (LeCun et al., 1989), followed by max pooling. Then, prior to entering the classifier, the featurized data representation is passed through two sets of dropout and fully connected layers to reduce the size of the feature representation to be comparable with the recurrent branch. The recurrent branch contains two long short-term memory (LSTM; Hochreiter & Schmidhuber, 1997) layers followed by a fully connected layer that yields the same number of output features as the convolutional branch. Prior to the classifier, a log softmax activation function is applied to the outputs of the two branches. The data representations are then concatenated and passed through two more fully connected layers with ReLU and softmax activation functions, respectively, which reduce the data representation size to an array equal in size to the number of classes. Figure 1 illustrates the architecture and Table 1 presents the specifications for each layer within the network.

3 Experiment

To test the network architecture, we simulate a realistic astronomical dataset and train a standalone CNN, a standalone RNN, and our experimental network architecture. The main metric of interest is the receiver operating characteristic area under curve (ROC AUC) for the two class problem of general LSNe identification. This experiment has the limitation of using simulated data as opposed to real data, so performance on real astronomical datasets cannot be directly inferred without retraining, but we are nonetheless able to evaluate the experimental architecture.

Data We simulate astronomical data using the exact characteristics of the upcoming Vera C. Rubin Observatory's Legacy Survey of Space and Time. The simulated dataset contains 17 sub-classes: nine types of common astronomical objects for the "No Lens" class, four types of non-time-varying lensing systems for the "Lens" class, two types of lensed type-Ia supernovae for the "LSNe-Ia"

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

Table 1: ZipperNet layer specifications. A shorthand notation has been utilized for kernel size (k), padding (p), stride (s), dropout fraction (f), and hidden units (h).
† Uses a ReLU activation function
‡ Uses a LogSoftmax activation function

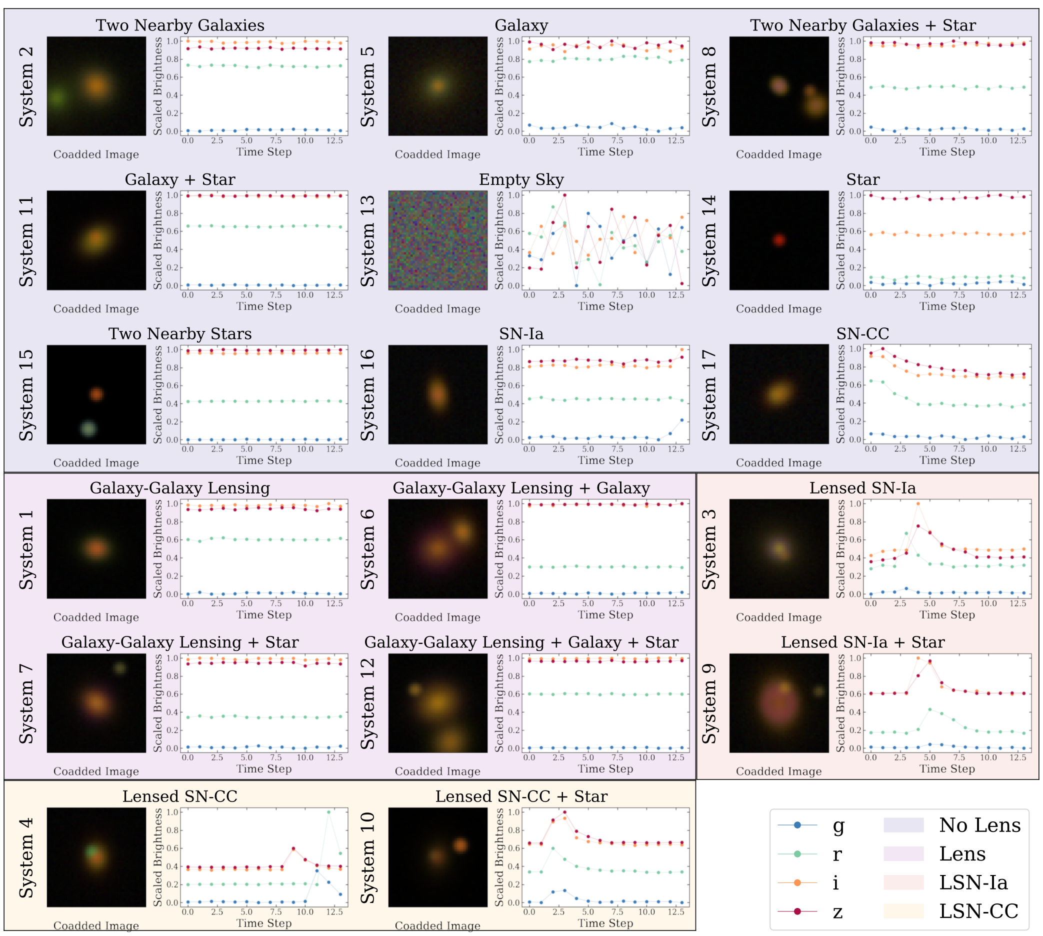

Figure 2: Examples of the 17 simulated systems from the dataset grouped into No Lens (blue), Lens (magenta), LSN-Ia (orange), and LSN-CC (yellow) classes. Each example displays an average image of the 4 color channels and the extracted light curve from our processing.

class, and two types of lensed core-collapse supernovae for the "LSNe-CC" class. Each of the four classes accounts for 25 percent of the dataset. Object colors, shapes, distances, time variability, and all known physical correlations are sampled from catalogs of astronomical observations to ensure realistic simulations (Abbott et al., 2021). We simulate a modest dataset of 10,000 examples using the `deeplenstronomy` library (Morgan et al., 2021) and apply rotations and mirroring for preprocessing. Each dataset example is a sequence of 14 images from different time points in a regular observing sequence. Prior to entering the network, the sum of the pixel values in a circular aperture is extracted from each image and stored in an array and the sequence of images is averaged into a single image. Both the light curve extraction and image averaging are performed in each of the four color channels, so that network inputs are a $45 \times 45 \times 4$ color image and a 14×4 color light curve. Configuration files to reproduce the dataset are available online (Morgan, 2021). Figure 2 shows examples of each of the systems that comprise the four-class dataset as well as their extracted time-variability.

Training We partition our simulated data into 80 percent training, 10 percent validation, and 10 percent testing splits. All networks are trained using the Adam (Kingma & Ba, 2017) optimizer, a constant learning rate of 0.001, and categorical cross entropy loss. We utilize a batch size of 5 since the recurrent cells were observed to not perform well at larger batch sizes. The learning algorithm hyperparameters and the network hyperparameters were determined through iterative tests on the validation dataset, though an exhaustive hyperparameter optimization was not performed. The results presented in this letter are for the 10 percent testing split after training on the training plus validation sets using the selected hyperparameters. We train the networks for 40 epochs in total, and the validation accuracies were observed to plateau after approximately 20 epochs. Because of the modest network and dataset size, we were able to train each network locally on a MacBookPro with a 2.3 GHz Dual-Core Intel Core i5 processor and 16GB of RAM in a wall-clock time of 7 hours.

Results Testing dataset ROC curves for the standalone RNN, standalone CNN, and our experimental ZIPPER architecture are presented in Figure 3. We also present a fourth classifier called COMBO, which requires both the RNN and CNN to predict that an object is a LSN. We note the presence of sharp features in the RNN (and consequently the COMBO) ROC curves, and interpret them as the influence of incorrectly classified unlensed SNe (systems 16 and 17 in Figure 2). Without the image-based information from the CNN feature representation, the RNN is unable to differentiate between unlensed and lensed SNe.

While it is interesting (and validating for our approach) that the smooth ROC curve features are recovered with the addition of image-based information, the key result is that the ZIPPER architecture outperforms all other networks in terms of ROC AUC. The ZIPPER's leveraging of image-based (temporal) information allows it to perform well in areas where the RNN (CNN) has low performance. Furthermore, the ZIPPER outperforming the COMBO classifier evidences an advantage in combining the convolutional and recurrent structures within a unified framework: the network is able to optimally weight the image and temporal features before using them to make a classification.

4 Conclusion

This analysis presents a novel deep learning architecture tailored to the problem of lensed supernovae identification in large astronomical datasets. The architecture processes images with convolutional layers, simultaneously processes brightness as a function of time with recurrent layers, concatenates the feature representation, and classifies the data with a series of fully connected layers. This multi-branch approach is ideal for putting image information of gravitational lensing and temporal information of variability on the same footing, and outperforms comparable feed-forward networks. With the performance demonstrated here, we expect this architecture to be a useful tool in the identification of thousands of lensed supernovae in the upcoming Rubin Observatory's Legacy Survey of Space and Time and in the resolution of the tension on the value of the expansion rate of the Universe.

5 Broader Impacts

The broader impacts of this work lie in its accessibility, reproducibility, simplicity, and generalizability. We make the neural network code and all dataset simulation inputs public (Morgan, 2021).

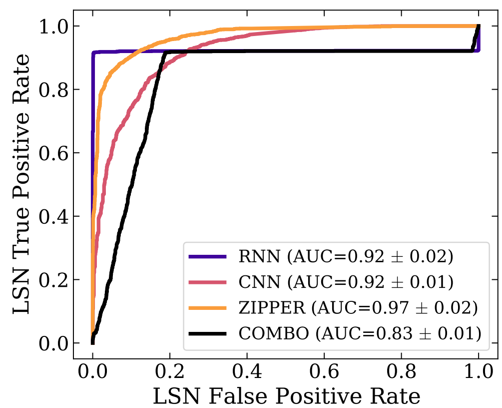

Figure 3: Receiver Operating Characteristic (ROC) curves using different networks. The ROC curves are calculated for the two-class problem of LSNIa and LSNCC versus No Lens and Lens. An Area Under Curve (AUC) of 1.0 indicates perfect performance, while an AUC of 0.5 indicates random guessing. The mean and standard deviation for 3 random seeds is presented.

Furthermore, the data from the Rubin Observatory's Legacy Survey of Space and Time will be public 1 year after it is collected, meaning any interested scientist can take this technique, apply it to real data, and perform their own search for gravitationally lensed supernovae. Beyond data and code access, we have intentionally designed this first iteration of our ZipperNet architecture as simply as possible. We anticipate that several applications could be found outside of astronomy where multiple types of information (in our case temporal and image-based), can be utilized within a unified framework to achieve higher performance. By itself, the architecture presented here has no direct negative impact on society; however, it is important to be mindful of how improved performance by leveraging multiple data formats could one day produce ethical dilemmas, e.g. invasions of privacy. Keeping ethical concerns in mind with use cases, we can look on the bright side of the broader impacts of expansions on this multi-branch network architecture. Problems in medical physics, self-driving cars, and several other applications routinely have to combine multiple types of information into artificial intelligence frameworks and could improve many features of society as a whole. We encourage other researchers to explore the multi-branch approach presented here, expand on it, and apply it to problems across several fields while cautiously practicing mindfulness regarding ethical concerns.

Acknowledgments and Disclosure of Funding

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, idea-generation, and collaboration. This community was important for the development of this project.

References

Abbott, T. M. C., Adamow, M., Aguena, M., et al. 2021, The Dark Energy Survey Data Release 2, arXiv:2101.05765

Hochreiter, S. & Schmidhuber, J. 1997, Neural Comput., 9, 1735–1780

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

Kingma, D. P. & Ba, J. 2017, Adam: A Method for Stochastic Optimization, arXiv:1412.6980

LeCun, Y., Boser, B., Denker, J. S., et al. 1989, Neural Computation, 1, 541

Morgan, R. 2021, DeepZipper: Code and Simulation Inputs, DOI: 10.5281/zenodo.5425727

Morgan, R., Nord, B., Birrer, S., Lin, J. Y.-Y., & Poh, J. 2021, Journal of Open Source Software, 6, 2854

Oguri, M. 2019, Reports on Progress in Physics, 82, 126901

Paszke, A., Gross, S., Massa, F., et al. 2019, PyTorch: An Imperative Style, High-Performance Deep Learning Library, arXiv:1912.01703

Treu, T. & Marshall, P. J. 2016, The Astronomy and Astrophysics Review, 24, 11

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See the beginning of Section 3.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See (Morgan, 2021).
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See the Training paragraph of Section 3.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Figure 3.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See the Training paragraph of Section 3.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] See (Morgan, 2021).
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]