



SCONE: Supernova Classification with a Convolutional Neural Network

Helen Qu¹ , Masao Sako¹ , Anais Möller² , and Cyrille Doux¹ 

¹ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA; helenqu@sas.upenn.edu

² LPC, Université Clermont Auvergne, CNRS/IN2P3, F-63000 Clermont-Ferrand, France

Received 2021 February 19; revised 2021 June 2; accepted 2021 June 2; published 2021 July 20

Abstract

We present a novel method of classifying Type Ia supernovae using convolutional neural networks, a neural network framework typically used for image recognition. Our model is trained on photometric information only, eliminating the need for accurate redshift data. Photometric data is preprocessed via 2D Gaussian process regression into two-dimensional images created from flux values at each location in wavelength-time space. These “flux heatmaps” of each supernova detection, along with “uncertainty heatmaps” of the Gaussian process uncertainty, constitute the data set for our model. This preprocessing step not only smooths over irregular sampling rates between filters but also allows SCONE to be independent of the filter set on which it was trained. Our model has achieved impressive performance without redshift on the in-distribution SNIa classification problem: $99.73 \pm 0.26\%$ test accuracy with no over/underfitting on a subset of supernovae from PLAsTiCC’s unblinded test data set. We have also achieved $98.18 \pm 0.3\%$ test accuracy performing six-way classification of supernovae by type. The out-of-distribution performance does not fully match the in-distribution results, suggesting that the detailed characteristics of the training sample in comparison to the test sample have a big impact on the performance. We discuss the implication and directions for future work. All of the data processing and model code developed for this paper can be found in the [SCONE software package](#) located at github.com/helenqu/scone.

Unified Astronomy Thesaurus concepts: Type Ia supernovae (1728); Classification (1907); Convolutional neural networks (1938); Supernovae (1668); Gaussian Processes regression (1930)

1. Introduction

The discovery of the accelerating expansion of the universe (Riess et al. 1998; Perlmutter et al. 1999) has led to an era of sky surveys designed to probe the nature of dark energy. Type Ia supernovae (SNe Ia) have been instrumental to this effort due to their standard brightness and light-curve profiles. Building a robust data set of SNe Ia across a wide range of redshifts will allow for the construction of an accurate Hubble diagram that will enrich our understanding of the expansion history of the universe as well as place constraints on the dark energy equation of state.

Modern-scale sky surveys, including SDSS, Pan-STARRS, and the Dark Energy Survey, have identified thousands of supernovae throughout their operational lifetimes (Frieman et al. 2008; Chambers et al. 2016; Smith et al. 2020). However, it has been logistically challenging to follow up on most of these detections spectroscopically. The result is a low number of spectroscopically confirmed SNe Ia and a large photometric data set of SNe Ia candidates. The upcoming Rubin Observatory Legacy Survey of Space and Time (LSST) is projected to discover 10^7 supernovae (LSST Science Collaboration et al. 2009), with millions of transient alerts each observing night. As spectroscopic resources are not expected to scale with the size of these surveys, the ratio of spectroscopically confirmed SNe to total detections will continue to shrink. With only photometric data, distinguishing between SNe Ia and other types can be difficult. A reliable photometric SNe Ia classification algorithm will allow us to tap into the vast potential of the photometric data set and pave the way for confident classification and analysis of the ever-growing library of transients from current and future sky surveys.

Significant progress has been made in the past decade in the development of such an algorithm. Most approaches involve light-curve template matching (Sako et al. 2011), feature

extraction paired with either sequential cuts (Bazin et al. 2011; Campbell et al. 2013), or machine learning algorithms (Lochner et al. 2016; Möller et al. 2016; Dai et al. 2018; Boone 2019). Most recently, the spotlight has been on deep learning techniques since it has been shown that classification based on handcrafted features is not only more time intensive for the researcher but is outperformed by neural networks trained on raw data (Charnock & Moss 2017; Kimura et al. 2017; Moss 2018). Since then, many neural network architectures have been explored for SN photometric classification, such as PELICAN’s CNN architecture (Pasquet et al. 2019) and SuperNNova’s deep recurrent network (Möller & de Boissière 2020).

Several photometric classification competitions have been hosted, including the Supernova Photometric Classification Challenge (Kessler et al. 2010) and the Photometric LSST Astronomical Time Series Classification Challenge (PLAsTiCC; The PLAsTiCC team et al. 2018). These have not only resulted in the development of new techniques, such as PSNID (Sako et al. 2011) and Avocado (Boone 2019), but have also provided representative data sets available to researchers during and after the competition, such as the PLAsTiCC unblinded data set used in this paper.

In this paper we present SCONE, a novel application of deep learning to the photometric classification problem. SCONE is a convolutional neural network (CNN), an architecture prized in the deep learning community for its state-of-the-art image recognition capabilities (LeCun et al. 1989, 1998; Krizhevsky et al. 2012; Simonyan & Zisserman 2014; Zeiler & Fergus 2014). Our model Requires raw photometric data only, precluding the necessity for accurate redshift approximations. The data set is preprocessed using a light-curve modeling technique via Gaussian processes described in Boone (2019), which alleviates the issue of irregular sampling between filters

but also allows the CNN to learn from information in all filters simultaneously. The model also has relatively low computational and data set size requirements without compromising on performance—400 epochs of training on our $\sim 10^4$ data set requires around 15 minutes on a GPU.

We will introduce the data sets used to train and evaluate SCONE its computational requirements, as well as the algorithm itself in Section 2. Section 3 will focus on the performance of SCONE on both binary and categorical classification. Section 4 presents an analysis of misclassified light curves and heatmaps for both modes of classification.

2. Methods

2.1. Data Sets

The PLAsTiCC training and test data sets were originally created for the 2018 Photometric LSST Astronomical Time Series Classification Challenge.

The PLAsTiCC training set includes ~ 8000 simulated observations of low-redshift, bright astronomical sources, representing objects that are good candidates for spectroscopic follow-up observation. This data set will be referred to in future sections as the “spectroscopic data set.” We use this data set to evaluate the out-of-distribution performance of SCONE in Section 3.1.2.

The PLAsTiCC test set consists of 453 million simulated observations of 3.5 million transient and variable sources, representing 3 yr of expected LSST output (Kessler et al. 2019). The objects in this data set are generally fainter, higher redshift, and do not have associated spectroscopy. Note that most of the results presented in this paper are produced from this data set alone and will be referred to as “the data set,” “the main data set,” or “the PLAsTiCC data set” in future sections.

All observations in both data sets were made in LSST’s *ugrizY* bands and realistic observing conditions were simulated using the LSST Operations Simulator (Delgado & Schumacher 2014). While PLAsTiCC includes data from many other transient sources, we are using only the supernovae in the data sets. We selected all of the type II, Iax, Ibc, Ia-91bg, Ia, and SLSN-1 sources (corresponding to `true_target` values of 42, 52, 62, 67, 90, and 95, respectively) from this data set and chose only well-sampled light curves by restricting ourselves to observations simulating LSST’s deep drilling fields (`ddf = 1`). `peak_mjd` values, the modified Julian date of peak flux for each object, were calculated for both the main data set and the spectroscopic data set by taking the signal-to-noise weighted average of all observation dates. `peak_mjd` is referred to as t_{peak} in future sections. The total source count for the spectroscopic set is 4556 and the total source count for the main data set is 32,087. A detailed breakdown by type is provided in Table 1.

The categorical data set was created using SNANA (Kessler et al. 2009) with the same models and redshift distribution as the main data set in order to perform categorical classification with balanced classes. 2000 examples of each type were randomly selected to constitute a class-balanced data set of 12,000 examples.

2.2. Quality Cuts

In order to ensure that the model is learning only from high-quality information, we have instituted some additional quality-

Table 1
Makeup of the PLAsTiCC Data Set by Type

SN Type	Number of Sources	
	Spectroscopic	Main
SNIa	2313	12,640
SNI	1193	15,986
SNIbc	484	2194
SNIa-91bg	208	362
SNIax	183	807
SLSN-1	175	98

Table 2
Makeup of the PLAsTiCC Data Set by Type After Applying Quality Cuts

SN Type	Number of Sources	
	Spectroscopic	Main
SNIa	654	6128
SNI	262	5252
SNIbc	97	779
SNIa-91bg	41	113
SNIax	59	281
SLSN-1	37	58

based cuts on all data sets. These cuts are based on light-curve quality, so all metrics are defined for a single source. The metrics evaluated for these cuts are as follows:

1. Number of detection data points (n_{detected}): number of observations where the source was detected. We chose a detection threshold of $S/N > 5$, based on Figure 8 of Kessler et al. (2015).
2. Cumulative signal-to-noise ratio (CS/N): cumulative S/N for all points in the light curve

$$\text{CSNR} = \sqrt{\sum \frac{f^2}{\sigma_f^2}}.$$

3. Duration: timespan of detection data points

$$t_{\text{active}} = t_{\text{last}} - t_{\text{first}},$$

where f represents the flux measurements from all observations of a given source, σ_f represents the corresponding uncertainties, t represents the timestamps of all observations of a given source, t_{first} is the time of initial detection, and t_{last} is the time of final detection. Our established quality thresholds require that:

1. $n_{\text{detected}} \geq 5$
2. $\text{CS/N} > 10$
3. $t_{\text{active}} \geq 30$ days

for light-curve points in the range $t_{\text{peak}} - 50 \leq t \leq t_{\text{peak}} + 130$. 1150 out of 4556 sources passed these cuts in the spectroscopic data set and 12,611 out of 32,087 sources passed these cuts in the main data set. The makeup of these data sets is detailed in Table 2. The categorical data set was created from sources that already passed the cuts, so the makeup is unchanged.

2.3. Class Balancing

Maintaining an equal number of examples of each class, or a balanced class distribution, is important for machine learning data sets. Balanced data sets allow for an intuitive interpretation

Table 3
Central Wavelength of Each Filter

Filter	Central Wavelength (Å)
u	3670.69
g	4826.85
r	6223.24
i	7545.98
z	8590.90
Y	9710.28

of the accuracy metric as well as provide ample examples of each class for the machine learning model to learn from.

As shown in Table 2, the natural distribution of the spectroscopic data set is more abundant in Ia sources than non-Ia sources. Thus, all non-Ia sources were retained in the class balancing process for binary classification and an equivalent number of Ia sources were randomly chosen. SNIax and SNIa-91bg sources were labeled as non-Ia sources for binary classification. The class-balanced spectroscopic data set has 496 sources of each class for a total of 992 sources.

In contrast, the natural distribution of the main data set is more abundant in non-Ia sources than Ia sources. Thus, all Ia sources were retained in the class balancing process for binary classification and an equivalent number of non-Ia sources were randomly chosen. The random selection process does not necessarily preserve the original distribution of non-Ia types. The class-balanced data set has 6128 sources of each class for a total of 12,256 sources.

The categorical data set of 2000 sources for each of the six types was created explicitly for the purpose of retaining balanced classes in categorical classification, as mentioned in Section 2.1.

All data sets were split by class into 80% training, 10% validation, and 10% testing. Splitting by class ensures balanced classes in each of the training, validation, and test sets.

2.4. Heatmap Creation

Prior to training, we preprocess our light-curve data into heatmap form. First, all observations are labeled with the central wavelength of the observing filter according to Table 3, which was calculated from the filter functions used by The PLAsTiCC team et al. (2018). We then use the approach described by Boone (2019) to apply two-dimensional Gaussian process regression to the raw light-curve data to model the event in the wavelength (λ) and time (t) dimensions. We use the Matern 32 kernel with a fixed 6000 Å characteristic length scale in λ and fit for the length scale in t . Once the Gaussian process model has been trained, we obtain its predictions on a λ , t grid and call this our “heatmap.” Our choice for the λ , t grid was $t_{\text{peak}} - 50 \leq t \leq t_{\text{peak}} + 130$ with a one-day interval and $3000 < \lambda < 10,100$ Å with a 221.875 Å interval. The significance of this choice is explored further in Section 3.

The result of this is a $n_\lambda \times n_t$ image-like matrix, where n_λ and n_t are the lengths of the wavelengths array and times array, respectively, given to the Gaussian process. We also take into account the uncertainties on the Gaussian process predictions at each time and wavelength, producing a second image-like matrix. We stack these two matrices depthwise and divide by the maximum flux value to constrain all entries to [0,1]. This matrix is our input to the convolutional neural network.

Figure 1 shows the raw light-curve data in the $t_{\text{peak}} - 50 \leq t \leq t_{\text{peak}} + 130$ range for each filter in blue, the Gaussian process model in gray, and the resulting flux and uncertainty heatmaps at the bottom. Figure 2 shows a representative example of a heatmap for each supernova type.

2.5. Convolutional Neural Networks

Convolutional neural networks (CNNs) are a type of artificial neural network that makes use of the convolution operation to learn local, small-scale structures in an image. This is paired with an averaging or subsampling layer, often called a *pooling layer*, that reduces the resolution of the image and allows the subsequent convolutional layers to learn hierarchically more complex and less localized structures.

In a *convolutional layer*, each unit receives input from only a small neighborhood of the input image. This use of a restricted receptive field, or *kernel size*, resembles the neural architecture of the animal visual cortex and allows for extraction of local, elementary features such as edges, endpoints, or corners. All units, each corresponding to a different small neighborhood of the image, share the same set of learned weights. This allows them to detect the presence of the same feature in each neighborhood of the image. Each convolutional layer often has several layers of these units, each of which is called a *filter* and extracts a different feature. The output of a convolutional layer is called a *feature map*.

Pooling layers reduce the local precision of a detected feature by subsampling the each unit’s receptive field, often 2×2 pixels, based on some rule. Average pooling, for example, extracts the average of the four pixels and max pooling extracts the maximum value. Assuming no overlap in the receptive fields, the spatial dimensions of the resulting feature maps will be reduced by half.

Dropout (Hinton et al. 2012) is a commonly used regularization technique in fully-connected layers. A *dropout layer* chooses a random user-defined percentage of the input weights to set to zero, improving the robustness of the learning process.

A convolutional neural network typically consists of alternating convolutional layers and pooling layers, followed by a series of fully-connected layers that learn a mapping between the result of the convolutions and the desired output.

2.6. SCONE Architecture

The relatively simple architecture of SCONE, shown in Figure 3, allows for a minimal number of trainable parameters, speeding up the training process significantly without compromising on performance. It has a total of 22,606 trainable parameters for categorical classification and 22,441 trainable parameters for binary classification when trained on heatmaps of size $32 \times 180 \times 2$ ($h \times w \times d$).

As mentioned in Section 2.4, each heatmap is divided by its maximum flux value for normalization. After receiving the normalized heatmap as input, the network pads the heatmap with a column of zeros on both sides, bringing the heatmap size to $32 \times 182 \times 2$. Then, a convolutional layer is applied with h filters and a kernel size of $h \times 3$, which in this case is 32 filters and a 32×3 kernel, resulting in a feature map of size $1 \times 180 \times 32$. We reshape this feature map to be $32 \times 180 \times 1$, apply batch normalization, and repeat the above process one more time. We have now processed our heatmap through two

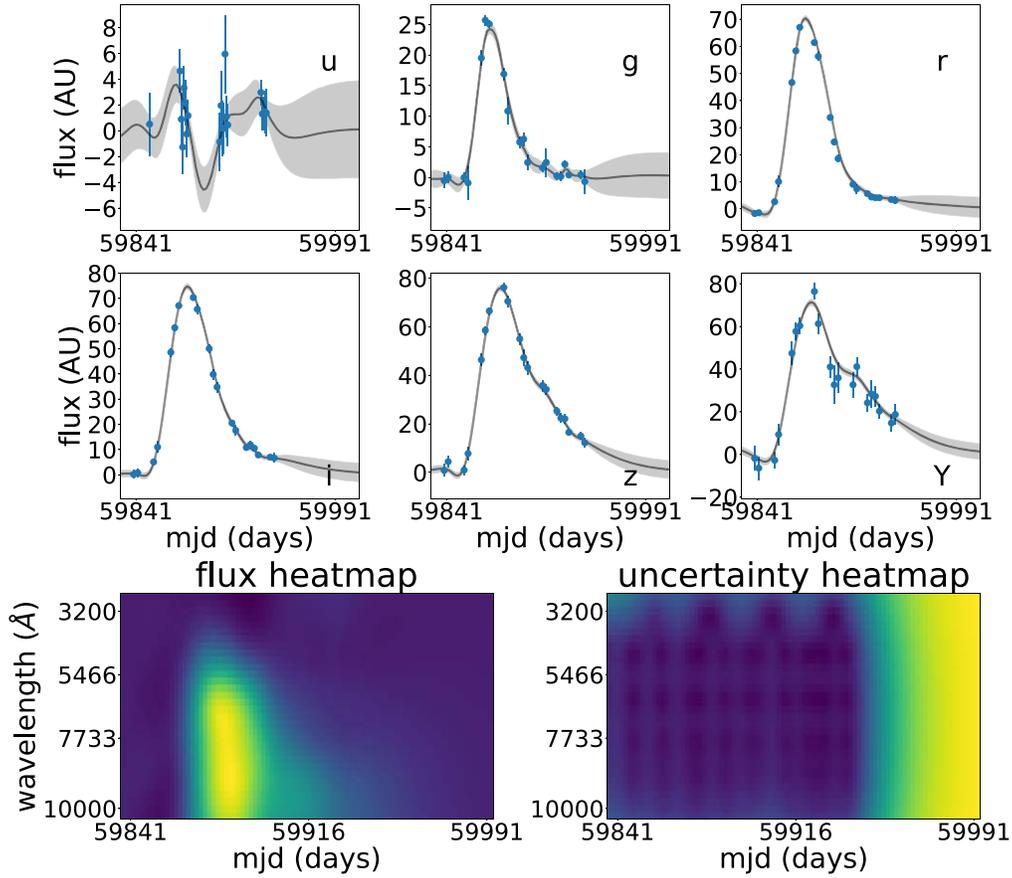


Figure 1. Raw *ugrizY* light-curve data with the Gaussian process model at corresponding wavelengths and resulting heatmap and error heatmap for a type Ia SN. The shaded regions in the Gaussian process plots represent the Gaussian process error.

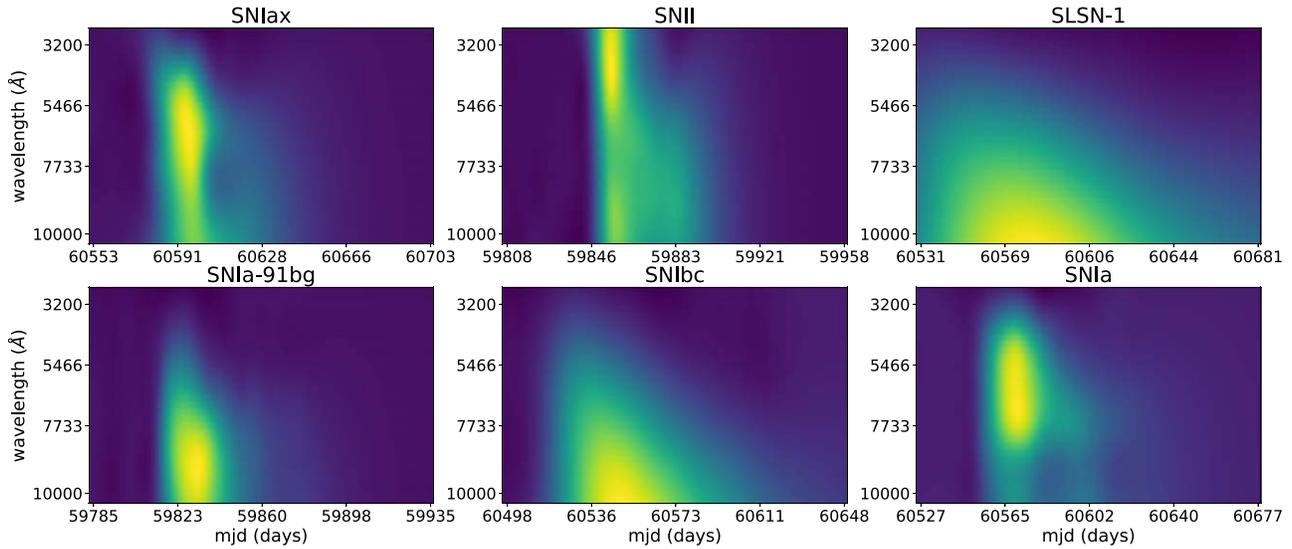


Figure 2. Example flux heatmaps for each supernova type.

“convolutional blocks” with an output feature map of size $32 \times 180 \times 1$.

We apply 2×2 max pooling to our output, reducing its dimensions to $16 \times 90 \times 1$, and pass it through two more convolutional blocks; but this time $h = 16$.

We pass our output through a final 2×2 max pooling layer, resulting in an $8 \times 45 \times 1$ feature map. This is subsequently flattened into a 360-element array and passed through a 50%

dropout layer. A 32-unit fully-connected layer followed by a 30% dropout layer feeds into the final layer. For binary classification, this is a node with a sigmoid activation that returns the model’s predicted Ia probability. For categorical classification, the final layer contains six nodes with softmax activations that return the respective probabilities of each of the six SN types. Both of these versions of SCONE are shown in Figure 3.

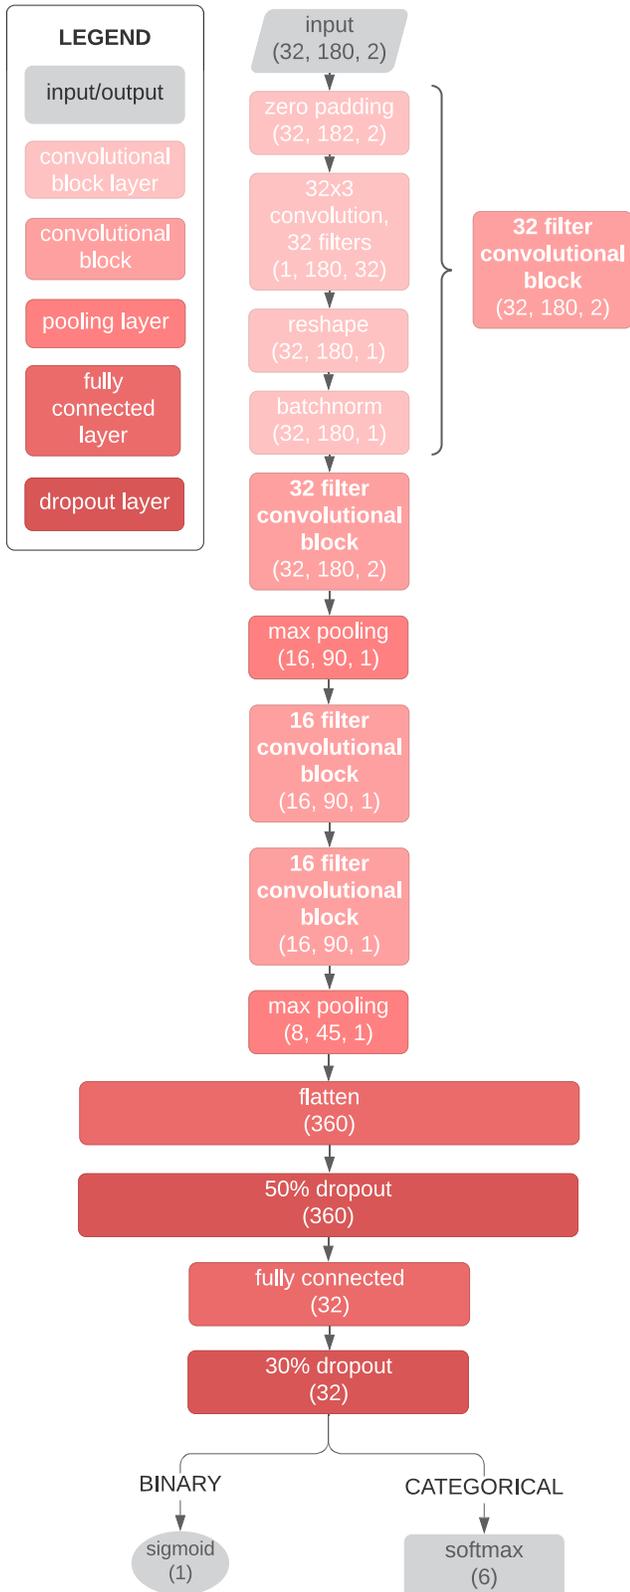


Figure 3. SCONE architecture for binary and categorical classification.

The model is trained with the binary cross-entropy loss function for binary classification and the sparse categorical cross-entropy loss function for categorical classification. Both classification modes use the Adam optimizer (Kingma & Ba 2017) at a constant $1e-3$ learning rate for 400 epochs.

2.7. Evaluation Metrics

Accuracy is defined as the number of correct predictions divided by the number of total predictions. We also evaluated the model on a number of other performance metrics: purity, efficiency, and AUC.

Purity and *efficiency* are defined as:

$$\text{purity} = \frac{\text{TP}}{\text{TP} + \text{FP}}; \text{efficiency} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

where TP is true positive, FP is false positive, and FN is false negative.

AUC, or *area under the receiver operating characteristic (ROC) curve*, is a common metric used to evaluate binary classifiers. The ROC curve is created by plotting the false positive rate against the true positive rate at various discrimination thresholds, showing the sensitivity of the classifier to the chosen threshold. A perfect classifier would score a 1.0 AUC value while a random classifier would score a 0.5.

2.8. Computational Requirements

Due to the minimal number of trainable parameters and data-set size, the time and hardware requirements for training and evaluating with SCONE are relatively low. The first training epoch on one NVIDIA V100 Volta GPU takes approximately 2 seconds, and subsequent training epochs take approximately 1 second each with TensorFlow’s data-set caching. Training epochs on one Haswell node (with Intel Xeon Processor E5-2698 v3), which has 32 cores, take approximately 26 seconds each.

3. Results

3.1. Binary Classification

Our model achieved $99.93 \pm 0.06\%$ training accuracy, $99.71 \pm 0.2\%$ validation accuracy, and $99.73 \pm 0.26\%$ test accuracy on the Ia versus non-Ia binary classification problem performed on the class-balanced data set of 12,256 sources.

Figure 4 shows the confusion matrices for binary classification and Table 4 shows the model’s performance on all the evaluation metrics described in Section 2.7. Both Figure 4 and Table 4 were created with data from five independent training, validation, and test runs of the classifier. Unless otherwise noted, the default threshold for binary classification is 0.5, where classifier confidence equal to or exceeding 0.5 counts as an SNIa classification and vice versa. Altering this threshold produces Figure 5, the ROC curve of one of these test runs.

3.1.1. Heatmap Dimensions

We explored the binary classification performance of different heatmap dimensions in both the time (width) and wavelength (height) axes. For our 7100 Å wavelength range ($3000 \text{ Å} < \lambda < 10,100 \text{ Å}$), we chose intervals of 221.875 Å, 443.8 Å, 887.5 Å, and 1775 Å, resulting in heatmaps with 32, 16, 8, and 4 wavelength “bins”, respectively. For our 180 day range ($t_{\text{peak}} - 50 \leq t \leq t_{\text{peak}} + 130$), we chose intervals of 1, 2, 6, and 18 days, resulting in heatmaps with 180, 90, 30, and 10 time bins.

Figure 6 shows the training, validation, and test accuracies for each choice of wavelength and time dimensions. Our classifier seems relatively robust to these changes, showing minimal performance impacts for wavelength bins ≥ 16 and impressive performance for the smaller sizes as well. Test accuracy drops a

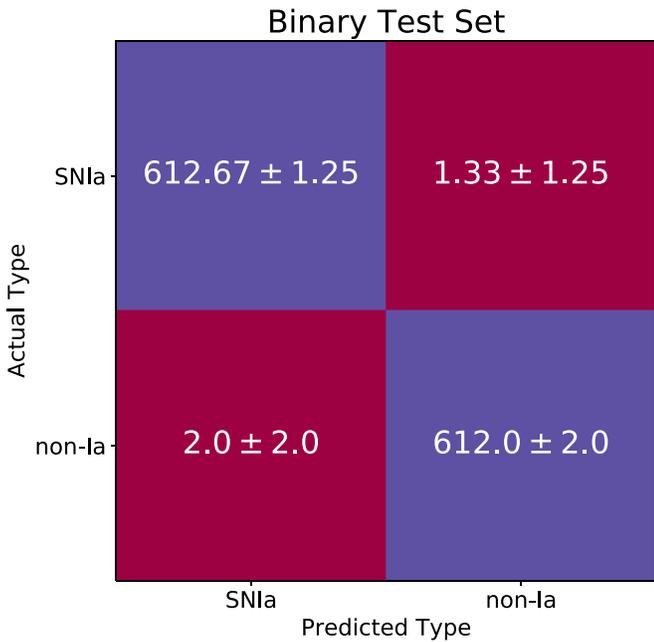


Figure 4. Confusion matrix showing average and standard deviation over five runs for binary classification on the test set.

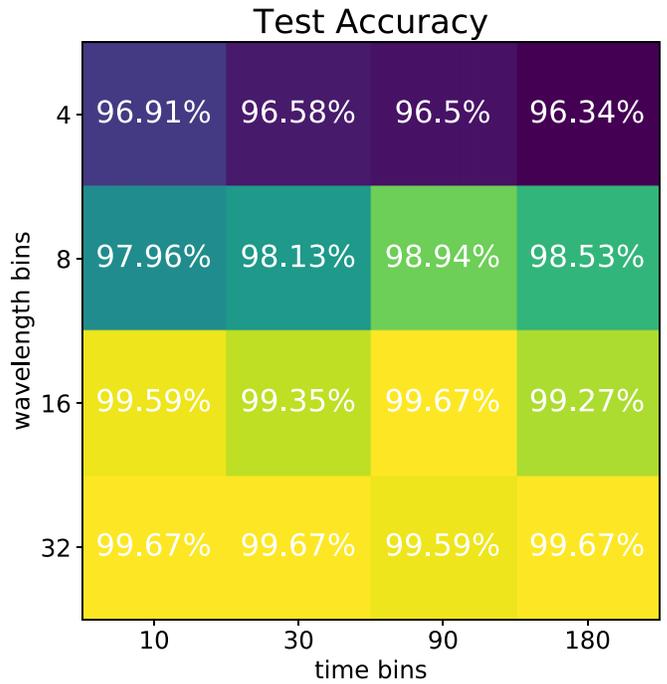


Figure 6. Test set accuracies for each choice of wavelength and time bins.

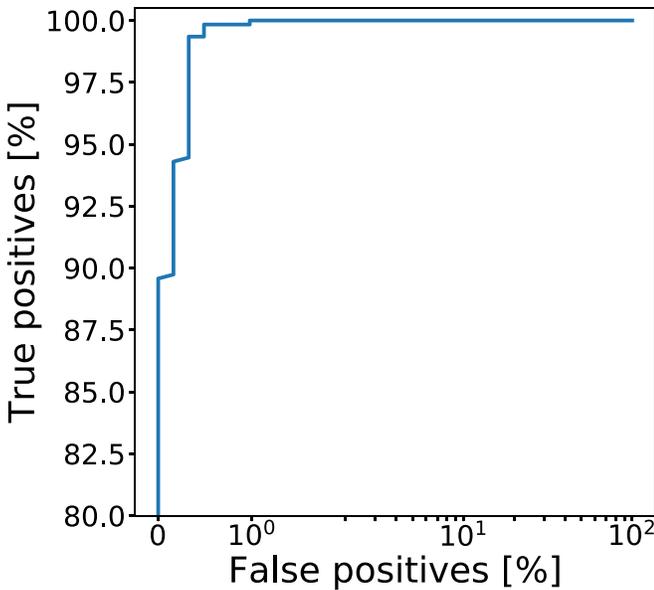


Figure 5. Semilog plot of the ROC curve for binary classification on the test set.

Table 4
Evaluation Metrics for Ia versus Non-Ia Classification on Cut Data Set

Metric	Training	Validation	Test
Accuracy	$99.93 \pm 0.06\%$	$99.71 \pm 0.2\%$	$99.73 \pm 0.26\%$
Purity	$99.93 \pm 0.06\%$	$99.76 \pm 0.25\%$	$99.68 \pm 0.35\%$
Efficiency	$99.93 \pm 0.05\%$	$99.67 \pm 0.23\%$	$99.78 \pm 0.22\%$
AUC	$1.0 \pm 4.1e-5$	$0.9991 \pm 1.6e-3$	$0.9994 \pm 1e-3$

maximum of 2.67% between the best- and worst-performing variants, 32×180 and 4×10 . This is noteworthy as the 32×180 heatmaps contain 144 times the number of pixels of the 4×10 heatmaps.

Table 5
Evaluation Metrics for Out-of-distribution Ia versus Non-Ia Classification

Metric	Training	Validation	Test
Accuracy	$99.65 \pm 0.36\%$	$98.84 \pm 1.1\%$	$80.61 \pm 1.75\%$
Purity	$99.86 \pm 0.31\%$	$98.74 \pm 1.31\%$	$81.86 \pm 2.22\%$
Efficiency	$99.67 \pm 0.46\%$	$98.75 \pm 1.27\%$	$80.06 \pm 1.31\%$
AUC	$1.0 \pm 8.9e-5$	$0.9939 \pm 9.5e-3$	$0.8552 \pm 1.4e-2$

The performance seems to unilaterally improve as expected as the number of wavelength bins increases, but increasing the number of time bins seems to yield varying, though likely not statistically significant, results.

We have reported all of our SCONE results using one of the best performing variants, the 32×180 heatmaps.

3.1.2. Out-of-distribution Results

Preliminary exploration into the out-of-distribution task of training on the spectroscopic data set and testing on the main data set yielded 80.6% test accuracy. The full results of training on 85% of the spectroscopic data set, validating on the remaining 15%, and testing on the full main data set are shown in Table 5. Since the redshift distribution of the spectroscopic data set is skewed toward lower redshifts, testing on class-balanced low-redshift subsets of the main data set yielded 83% test accuracy for $z < 0.4$ and 87% test accuracy for $z < 0.3$. Boone (2019) introduced redshift augmentation to mitigate this effect. The mismatch between the test and training sets, however, comes in other forms. For example, if the training data is generated using models rather than real data, differences in the characteristics of the spectral surfaces can have an impact on classification. The unknown relative rates of each type of event also affect the overall performance. Since out-of-distribution robustness is an integral part of the challenge of photometric SNe classification, improving the performance of SCONE on these metrics will be the topic of a future paper.

Table 6
Misclassified Test set Heatmaps by True and Predicted Type for Binary Classification

True Type	Predicted Type	>90% Confidence	90%–70% Confidence	70%–50% Confidence	Total	Percentage
SNIa	non-Ia	1	0	0	1	25%
SNIa	SNIa	2	0	0	2	50%
SNIax		1	0	0	1	25%

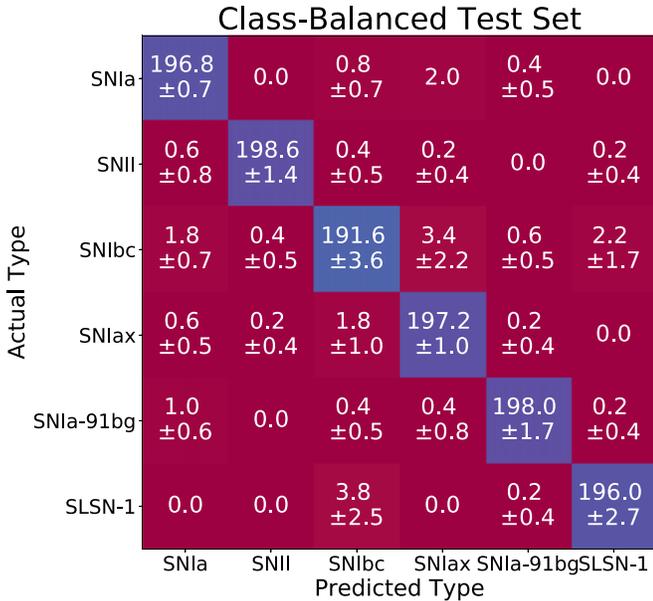


Figure 7. Confusion matrix showing average and standard deviation over five runs for categorical classification on the balanced test set.

3.2. Categorical Classification

In addition to binary classification, SCONE is able to perform categorical classification and discriminate between different types of SNe. We performed six-way categorical classification with the same PLASrTiCC data set used for binary classification as well as the class-balanced data set described in Section 2.1. Our model differentiated between SN types Ia, II, Ibc, Iax, SN-91bg, and SLSN-1. On the PLASrTiCC data set (not class balanced), it achieved $99.26 \pm 0.18\%$ training accuracy, $99.13 \pm 0.34\%$ validation accuracy, and $99.18 \pm 0.18\%$ test accuracy. The confusion matrices in Figure 8 show the average by-type breakdown for five independent runs.

On the balanced data set, it achieved $97.8 \pm 0.32\%$ training accuracy, $98.52 \pm 0.28\%$ validation accuracy, and $98.18 \pm 0.3\%$ test accuracy. The confusion matrices in Figure 7 show the average and standard deviations of the by-type breakdown for five independent runs.

It is worth noting that we trained and tested with the PLASrTiCC data set even though it is not class balanced for this task to try to evaluate the model’s performance on a data set emulating the relative frequencies of these events in nature.

4. Discussion

Analysis of misclassified heatmaps was performed for both binary and class-balanced categorical classification. No clear evidence of the effect of redshift on accuracy was found for either mode of classification. The quantity of misclassified SNIa examples per run is not sufficient for us to draw conclusions about the accuracy evolution as a function of redshift.

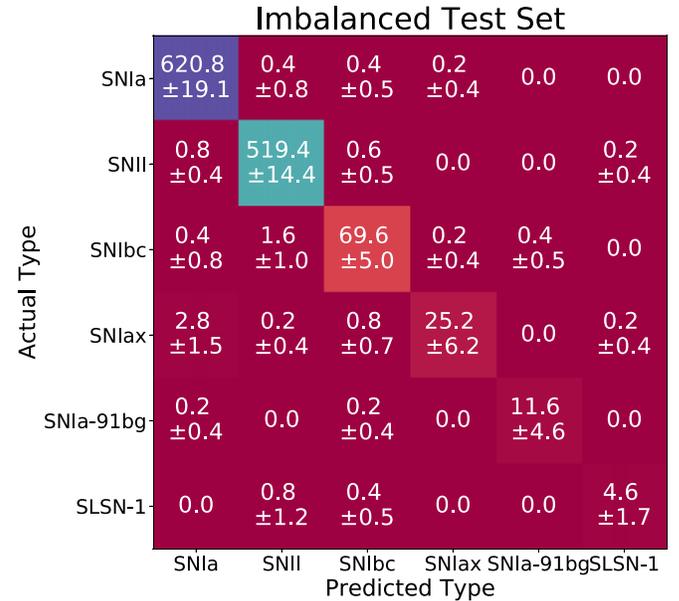


Figure 8. Confusion matrix showing average and standard deviation over five runs for categorical classification on the PLASrTiCC (imbalanced) test set.

4.1. Binary Classification

According to the data presented in Figure 4, the model seems to mispredict about the same number of Ias as non-Ias. An average of 3.33 Ia’s are mispredicted in the training set compared with 3.5 non-Ias. For the validation set, two Ias are mispredicted compared to 1.5 and 1.33 Ias and two non-Ias are mislabeled for the test set.

The misclassified set summarized in Table 6 is the result of one of the five runs represented by the data in Figure 4. In this example, the model missed five examples total during testing—1 SNIa, 2 SNIi, and 1 SNIax. It was >90% confident about all of these misclassifications, which is certainly not the case for categorical classification. This could be due to the fact that there are more examples of each type for binary classification than categorical (~6000 and 1200 per type for training).

4.2. Categorical Classification

The data in Figures 7 and 8 show some level of symmetry between misclassifications. SNIax and SLSN-1 seem to be easily distinguishable across the board, for example, with 0s in all relevant cells in both figures except one. In Figure 7, SNIbc and SLSN-1 are seemingly very similar, as they are misclassified as one another at similarly high rates.

There are notable differences between Figures 7 and 8, however. In Figure 7, representing the classifier’s performance on class-balanced categorical classification, the model mispredicts SNIas as other types at a similar rate as non-Ias mispredicted as Ias. An average of 3.2 Ias are mispredicted, whereas an average of four non-Ias are misclassified as Ia. In the confusion matrix shown

Table 7
Misclassified Test set Heatmaps by True and Predicted Type for Categorical Class-balanced Classification

True Type	Predicted Type	>90% Confidence	90%–70% Confidence	70%–50% Confidence	Total	Percentage
SNIa	SNIax	1	0	0	1	4.8%
SNI	SNIa	1	0	0	1	4.8%
	SNIbc	0	0	1	1	4.8%
SNIbc	SNIa	0	0	1	1	4.8%
	SNIax	1	0	0	1	4.8%
	SLSN-1	0	0	1	1	4.8%
SNIax	SNIa	1	2	0	3	14.3%
	SNIbc	1	2	0	3	14.3%
	SNIa-91bg	1	0	0	1	4.9%
SLSN-1	SNIbc	3	1	3	7	33.3%
SNIa-91bg	SNIbc	0	1	0	1	4.8%

in Figure 8, significantly more non-Ias were mispredicted as Ia. An average of one Ia was mispredicted compared to an average of 4.2 non-Ias misclassified as Ia. The rate of SNIbcs mispredicted as SLSN-1 is also significantly lower for the PLAsTiCC data set than for the balanced data set. These observations further reinforce the impact of imbalanced classes in classification tasks.

The misclassified set summarized in Table 7 is the result of one of the five class-balanced categorical classification runs represented by the data in Figure 7.

One point of interest is the lack of symmetry between misclassifications, in contrast with the analysis of Figures 7 and 8. This is clear in the significantly larger number of SLSN-1 misclassified as SNIbc (7) compared with the number of SNIbc misclassified as SLSN-1 (1). SNIax is also more often misclassified as other types (three as SNIa, three as SNIbc, and one as SNIa-91bg) than non-Iax misclassified as Iax (one SNIa and one SNIbc). The more symmetric Figure 7 suggests that the asymmetry of this table is due to randomness and would be corrected with data from other runs.

The distribution of misclassified examples across the confidence spectrum is nonuniform. In this table, *confidence* refers to the probability assigned to the predicted type by the classifier. Confidence near 100% for a misclassified example is potentially more insightful than one near 50%. Nine out of the 21 misclassified heatmaps were misclassified at >90% confidence, six at 90%–70% confidence, and six at 70%–50%. Surprisingly, the classifier is confidently wrong almost half the time. One particularly interesting example is a SNIbc “misclassified” as SLSN-1, but the classification probabilities for both SLSN-1 and SNIbc were 50%.

4.3. Limitations and Future Work

As stated in Section 2.1, it is important to note that the metrics reported in this paper are in-distribution results since the training, validation, and test sets are mutually exclusive segments of the main data set. The out-of-distribution performance of SCONE, as evaluated in Section 3.1.2, is noticeably diminished from the >99% in-distribution test accuracy. The high in-distribution test accuracy shows that SCONE is robust to previously unseen data, but the lower out-of-distribution test accuracy demonstrates SCONE’s sensitivity to variations in the parameters of the data set, such as the redshift distribution, relative rates of different types of SNe, small variations in the SNIa model, as well as telescope characteristics. Generalizing SCONE to become robust to these variations will be the subject of a future paper.

5. Conclusions

In this paper we have presented SCONE, a novel application of deep learning to the photometric supernova classification problem. We have shown that SCONE has achieved unprecedented performance on the in-distribution Ia versus non-Ia classification problem and impressive performance on classifying SNe by type without the need for accurate redshift approximations or handcrafted features.

Using the wavelength-time flux and error heatmaps from the Gaussian process for image recognition also allows the convolutional neural network to learn about the development of the supernova over time in all filter bands simultaneously. This provides the network with far more information than a photograph taken at one moment in time. Our choice of an $h \times 3$ convolutional kernel, where h is the number of wavelength bins, supplements these benefits by allowing the network to learn from data on the full spectrum of wavelengths in a sliding window of three days.

As future large-scale sky surveys continue to add to our ever-expanding transients library, we will need an accurate and computationally inexpensive photometric classification algorithm. Such a model can inform the best choice for allocation of our limited spectroscopic resources as well as allow researchers to further cosmological science using minimally contaminated SNIa data sets. SCONE can be trained on tens of thousands of light curves in minutes and confidently classify thousands of light curves every second at >99% accuracy.

Although SCONE was formulated with supernovae in mind, it can easily be applied to classification problems with other transient sources. The documented source code has been released on Github (github.com/helenqu/scone) to ensure reproducibility and encourage the discovery of new applications.

The authors would like to thank Rick Kessler for the help with SNANA simulations and Michael Xie for guidance on the model architecture. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. This work was supported by DOE grant DE-FOA-0001781 and NASA grant NNH15ZDA001N-WFIRST.

ORCID iDs

Helen Qu  <https://orcid.org/0000-0003-1899-9791>
Masao Sako  <https://orcid.org/0000-0003-2764-7093>

Anais Möller  <https://orcid.org/0000-0001-8211-8608>
 Cyrille Doux  <https://orcid.org/0000-0003-4480-0096>

References

- Bazin, G., Ruhlmann-Kleider, V., Palanque-Delabrouille, N., et al. 2011, *A&A*, **534**, A43
- Boone, K. 2019, *AJ*, **158**, 257
- Campbell, H., D’Andrea, C. B., Nichol, R. C., et al. 2013, *ApJ*, **763**, 88
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
- Charnock, T., & Moss, A. 2017, *ApJL*, **837**, L28
- Dai, M., Kuhlmann, S., Wang, Y., & Kovacs, E. 2018, *MNRAS*, **477**, 4142
- Delgado, F., & Schumacher, G. 2014, *Proc. SPIE*, **9149**, 91490G
- Frieman, J. A., Bassett, B., Becker, A., et al. 2008, *AJ*, **135**, 338
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. 2012, arXiv:1207.0580
- Kessler, R., Bassett, B., Belov, P., et al. 2010, *PASP*, **122**, 1415
- Kessler, R., Bernstein, J. P., Cinabro, D., et al. 2009, *PASP*, **121**, 1028
- Kessler, R., Marriner, J., Childress, M., et al. 2015, *AJ*, **150**, 172
- Kessler, R., Narayan, G., Avelino, A., et al. 2019, *PASP*, **131**, 094501
- Kimura, A., Takahashi, I., Tanaka, M., et al. 2017, in IEEE 37th Int. Conf. on Distributed Computing Systems Workshops (ICDCSW)
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, *Commun. ACM*, **60**, 84
- LeCun, Y., Boser, B., Denker, J. S., et al. 1989, *Neural Comput.*, **1**, 541
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proc. IEEE*, **86**, 2278
- Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K. 2016, *ApJS*, **225**, 31
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
- Möller, A., & de Boissière, T. 2020, *MNRAS*, **491**, 4277
- Möller, A., Ruhlmann-Kleider, V., Leloup, C., et al. 2016, *JCAP*, **2016**, 008
- Moss, A. 2018, arXiv:1810.06441
- Pasquet, J., Pasquet, J., Chaumont, M., & Fouchez, D. 2019, *A&A*, **627**, A21
- Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, *ApJ*, **517**, 565
- Qu, H. 2021, helenqu/scone: first release, v1.0.0, Zenodo, doi:10.5281/zenodo.4660288
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, **116**, 1009
- Sako, M., Bassett, B., Connolly, B., et al. 2011, *ApJ*, **738**, 162
- Simonyan, K., & Zisserman, A. 2014, arXiv:1409.1556
- Smith, M., D’Andrea, C. B., Sullivan, M., et al. 2020, *AJ*, **160**, 267
- The PLAsTiCC team, Allam, T. J., Bahmanyar, A., et al. 2018, arXiv:1810.00001
- Zeiler, M. D., & Fergus, R. 2014, in European Conf. on Computer Vision (Berlin: Springer), **818**