

# NOMAI: A real-time photometric classifier for superluminous supernova identification

## A science module for the Fink broker

E. Russeil<sup>1,\*</sup>, R. Lunnan<sup>1</sup>, J. Peloton<sup>2</sup>, S. Schulze<sup>3</sup>, P. J. Pessi<sup>4</sup>, D. Perley<sup>5</sup>, J. Sollerman<sup>1</sup>, A. Gkini<sup>1</sup>,  
Y. Hu<sup>1</sup>, T.-W. Chen<sup>6</sup>, E. C. Bellm<sup>7</sup>, T. X. Chen<sup>8</sup>, and B. Rusholme<sup>8</sup>

<sup>1</sup> Department of Astronomy, Oskar Klein Center, Stockholm University, 106 91 Stockholm, Sweden

<sup>2</sup> Université Paris-Saclay, CNRS/IN2P3, IJCLab, Orsay, France

<sup>3</sup> Department of Particle Physics and Astrophysics, Weizmann Institute of Science, 234 Herzl St, 76100 Rehovot, Israel

<sup>4</sup> Astrophysics Division, National Centre for Nuclear Research, Pasteura 7, 02-093 Warsaw, Poland

<sup>5</sup> Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK

<sup>6</sup> Graduate Institute of Astronomy, National Central University, 300 Jhongda Road, 32001 Jhongli, Taiwan

<sup>7</sup> DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA

<sup>8</sup> IPAC, California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125, USA

Received 13 April 2026 / Accepted 12 July 2026

### ABSTRACT

Superluminous supernovae (SLSNe) are one of the most luminous stellar explosions known, yet they remain poorly understood. Because they are intrinsically rare, efficiently identifying them in the large alert streams produced by modern time-domain surveys is essential for building larger observational samples and enabling spectroscopic follow-up. We present NOMAI, a machine-learning classifier designed to identify SLSN candidates directly from photometric alerts in the ZTF stream, using light curves accumulated over at least 30 days. It does not require any spectroscopic redshift and is running in real time within the Fink broker. ZTF light curves were transformed into a set of physically motivated features derived primarily from model-fitting procedures using SALT2 and Rainbow, a blackbody-based multiband fitting framework. These features were used to train an XGBoost classifier on a curated dataset of labeled ZTF sources constructed using literature samples of SLSNe, along with TNS and internal ZTF labeled sources. The final dataset contains 5280 unique sources, including 225 spectroscopically classified SLSNe. Averaged over 100 bootstrap realizations, the classifier reaches 66% completeness and 58% purity. Deployed within the Fink broker, NOMAI has been running continuously since 18 December 2025 on the ZTF alert stream and publicly reports SLSN candidates every night by automatically posting them to dedicated communication channels. Based on this, we also report the first two-month evaluation period, where the classifier successfully recovered 22 of the 24 active SLSNe reported on the Transient Name Server, providing an upper limit on its real-time completeness. The achieved performances, particularly the high completeness in real-time operations, demonstrate that the classifier provides a valuable tool for experts to efficiently scan the alert stream and identify promising candidates. In the near future, NOMAI is intended to be adapted to operate on the Legacy Survey of Space and Time conducted by the Vera C. Rubin Observatory, which is expected to uncover an unprecedented number of transients, making machine-learning based photometric classification essential.

**Key words.** methods: data analysis – methods: numerical – techniques: photometric – astronomical databases: miscellaneous – surveys – supernovae: general

## 1. Introduction

Superluminous supernovae (SLSNe) are rare stellar explosions whose peak luminosities exceed those of standard SNe by one to two orders of magnitude (Quimby et al. 2011; Gal-Yam 2012, 2019). Since their identification as a distinct population (Quimby et al. 2011), SLSNe have attracted considerable attention due to their extreme energetics. Although they display very diverse photometric properties, their light curves typically evolve more slowly than those of normal SNe and often display blue colors (e.g., Lunnan et al. 2018; Angus et al. 2019; De Cia et al. 2018; Chen et al. 2023; Gomez et al. 2024; Pessi et al. 2025). Because of their high luminosities, they are detectable at large cosmological distances, making them valuable probes for the high-redshift universe.

Despite significant observational and theoretical progress over the past decade, the physical mechanisms powering these explosions remain debated. Several scenarios have been proposed, including magnetar spin-down (Ostriker & Gunn 1971; Kasen & Bildsten 2010; Woosley 2010; Vurm & Metzger 2021; Mukhopadhyay et al. 2026), interaction with dense circumstellar material (Chatzopoulos et al. 2012; Sorokina et al. 2016; Wheeler et al. 2017; Chen et al. 2023), long-term fallback accretion of material onto a black hole (Dexter & Kasen 2013; Moriya et al. 2018), or pair-instability explosions (Fowler & Hoyle 1964; Barkat et al. 1967; Rakavy & Shaviv 1967; Woosley et al. 2002; Heger & Woosley 2002). Further constraining these models requires larger samples of well-observed events spanning a broad range of environments and redshifts. However, SLSNe are intrinsically rare (e.g., Prajs et al. 2017; Perley et al. 2020; Frohmaier et al. 2020) and only represent a small fraction of the overall SN

\* Corresponding author: [etienne.russeil@astro.su.se](mailto:etienne.russeil@astro.su.se)

population. Identifying them efficiently in wide-field surveys is therefore a key step toward improving our understanding of these extreme transients.

Modern time-domain surveys are now producing unprecedented volumes of transient alerts. The Zwicky Transient Facility (ZTF, [Bellm et al. 2019](#); [Graham et al. 2019](#); [Dekany et al. 2020](#); [Masci et al. 2019](#)) currently delivers hundreds of thousands of alerts per night, and the recent Vera C. Rubin Observatory will increase this number to ten million through the Legacy Survey of Space and Time (LSST, [Ivezić et al. 2019](#)). In this context, the classical approach of spectroscopic follow-up for every candidate becomes impossible. Only a small fraction of transient discoveries can realistically be spectroscopically classified, making automated photometric classification an essential part of modern astronomy. The alert streams produced by modern surveys require dedicated infrastructures capable of processing and distributing transient classifications in real time. Alert brokers ([Nordin, J. et al. 2019](#); [Möller et al. 2021](#); [Sánchez-Sáez et al. 2021](#); [Matheson et al. 2021](#); [Williams et al. 2024](#); [Jegou du Laz et al. 2026](#))<sup>1</sup> have emerged to fulfill this role by ingesting survey alerts, enriching them with additional information, and providing value-added products to the community. In this context, machine-learning classifiers can naturally be integrated within alert brokers to operate directly on the incoming data stream, and multiple works have explored the specific task of classifying SLSN light curves, although some classification pipelines are still developed and evaluated outside broker infrastructures.

Some approaches rely on multi-class classifiers in which SLSNe constitute one of several transient classes, such as the ALerCE data broker ([Sánchez-Sáez et al. 2021](#)), or Superphot+ ([de Soto et al. 2024](#)) running on the ANTARES broker ([Matheson et al. 2021](#)), both of which classify alerts from the ZTF data stream nightly. Other pipelines are specifically designed to identify SLSN events among the broader population of transients. The “NEural Engine for Discovering Luminous Events” (NEEDLE, [Sheng et al. 2024](#)), hosted within the Lasair broker ([Williams et al. 2024](#)), uses a hybrid neural network to process both the photometry and the telescope images to classify SLSN and TDE sources. “Finding Luminous and Exotic Extragalactic Transients” (FLEET, [Gomez et al. 2020, 2023](#)) attempts to classify SLSNe by relying on host galaxy association and photometric feature extraction to train a random forest classifier. Finally, by combining SuperRAENN ([Villar et al. 2020](#)), a recurrent autoencoder architecture, with Superphot ([Villar et al. 2019](#)), a versatile feature-extraction pipeline based on light-curve model fitting, [Hsu et al. \(2022\)](#) developed a classifier designed to identify SLSNe in Pan-STARRS1 data ([Tonry et al. 2012](#)). Furthermore, other works have focused on anticipating the LSST era using simulated datasets such as the Photometric LSST Astronomical Time-series Classification Challenge (PLAsTiCC, [Hložek et al. 2023](#)) to train SLSN classification models ([Qu & Sako 2022](#); [Burhanudin & Maund 2023](#); [Li et al. 2024](#)). Similarly, more recent works using the Extended LSST Astronomical Time-Series Classification Challenge (ELAsTiCC, [Knop & ELAsTiCC Team 2023](#)) have explored analogous problems using a more realistic LSST-like data stream format ([Fraga et al. 2024](#); [Cabrera-Vives et al. 2024](#)).

Among the brokers, Fink ([Möller et al. 2021](#)) has been operating on the ZTF public stream since late 2019, while also helping prepare for the upcoming Rubin era. It provides an infrastructure that allows the scientific community to deploy dedicated “science modules,” which operate on the alert stream to

compute additional features and classifications. This framework makes it possible to run machine-learning models continuously on the incoming alerts and distribute their outputs in almost real time.

In this work, we present NOMAI, a machine-learning classifier designed to identify SLSN candidates directly from ZTF photometric alerts and deploy it within the Fink broker as a science module. The classifier relies on physically motivated features extracted from transient light curves without requiring a spectroscopic redshift, and it is optimized to operate in a real-time streaming environment. Its integration within the broker enables the automatic identification and public sharing of SLSN candidates from the ZTF alert stream. In what follows, we describe the methodology used to build the classifier, in particular the construction of an informative dataset (Section 2), the extraction of the features, and the training of the model (Section 3). We then present the performance obtained using bootstrapping (Section 4), as well as the practical metrics achieved on the ZTF stream during the first months of operation (Section 5). Finally, we discuss the model’s limitations and its performance within the current landscape (Section 6), before concluding (Section 7).

## 2. Labeled dataset

In this section, we describe the construction of the labeled dataset. Because the ZTF alert stream contains a large fraction of non-transient sources, an initial filtering step is required before performing any feature extraction or classification. We therefore apply a sequence of selection criteria to isolate likely transient events and ensure sufficient photometric quality for the analysis. We then describe the origin of the data and the procedure used to assign labels to the sources.

### 2.1. Transient cut

The real-time ZTF data stream processed by the Fink broker consists of compact data packets called alerts ([Patterson et al. 2019](#)), generated whenever a source exhibits a significant brightness variation relative to the reference template image. Hence, a single transient source will trigger multiple alerts. Photometry-wise, they do not contain the full light curve generated by the events, but only their last detection along with 30 days of history. However, given the long-lived nature of SLSNe, it is crucial that the complete light curve is used to provide the most informed classification possible. Consequently, the Fink API must be queried nightly to retrieve the complete public photometric history of each source. But this operation is computationally expensive, and given the large stream of alerts provided by ZTF, which averages ~250,000 alerts per night, it is infeasible to perform this step on all alerts every night. Therefore, simple cuts must be applied in order to remove as much non-transient contamination as possible upstream. They constitute highly computationally efficient solutions that can be applied to process even the largest datasets.

Hence, the *transient\_complete* alert filter was integrated to the Fink broker. It is based on the Bright Transient Survey filtering pipeline<sup>2</sup> ([Fremming et al. 2020](#); [Perley et al. 2020](#)), which enables high transient completeness (recall) while considerably

<sup>1</sup> <https://pitt-broker.readthedocs.io/en/latest/#>

<sup>2</sup> A public implementation can be found here: <https://zenodo.org/records/4054129>. It is available on Fritz under the name Redshift Completeness Factor (RCF).

reducing contamination. The full *transient\_complete* pipeline consists of eight conditions:

- *not\_faint: magpsf* (PSF-fit difference magnitude) is currently brighter than 19.8, or the source had a very recent detection brighter than 19.
- *no\_pointunderneath*: is likely not sitting on top of or blended with a star in Pan-STARRS.
- *positivesubtraction*: is brighter than the template image.
- *real*: is likely a genuine astrophysical transient and not an artifact.
- *stationary*: is not a moving source.
- *no\_brightstar*: is likely not contaminated by a nearby bright star.
- *not\_variablesource*: is likely not a variable star.
- *not\_roid*: is likely not a solar system object.

The exact implementation is publicly available on Fink<sup>3</sup>. On average, this filter alone removes more than 99% of the alert stream, reducing it to a few hundred alerts per night and making it possible to retrieve the photometric history for the remaining events.

## 2.2. Quality cuts

The performance of any alert-classification system strongly depends on the minimal amount of information it requires from a light curve to produce a reliable output. Because SLSNe are generally slowly evolving events, it is possible to accumulate a substantial number of observations before classifying a source, thereby improving the purity (precision). However, since one of the primary goals of the classifier is to direct scientists toward promising candidates for spectroscopic follow-up, the classification should ideally be produced while the source is still bright and young. We balance early detection and improved purity by imposing minimal requirements on the photometry. To ensure that a substantial portion of the transient’s evolution is sampled, we require that each source has at least 30 days of photometric history before we can provide a classification. Given the characteristic timescale of SLSNe, this still allows the classifier to return a prediction near peak brightness for sources first detected close to explosion; however, faster-evolving SLSNe may be missed due to this requirement (Section 6.1). On top of a duration criterion, we also impose that the sources have more than 7 photometric observations in total, and more than 3 observations in each particular filter (*ZTF-g* and *ZTF-r*). These requirements guarantee sufficiently constrained fits during the feature-extraction stage (Section 3.2).

## 2.3. Data source

The dataset is built around the labels available on the Fritz<sup>4</sup> platform. Fritz is an extension of Kowalski<sup>5</sup> and SkyPortal (Coughlin et al. 2023; van der Walt et al. 2019), serving as an alert broker, multi-survey data archive, marshal, and tool for target management and follow-up. It also allows the creation of working groups, that receive a sub-stream of alerts based on customizable filters. Over the years, experts across multiple areas of transient astronomy have used this interface to assign spectroscopic classifications to transient sources. Using these high-quality labels for our dataset, the machine-learning

classifier learns from numerous confidently labeled events. We proceeded by collecting the *objectId* of every source that passed the filter presented in Sections 2.1 and 2.2.

As Fritz is a private platform not fully accessible to the broader scientific community, we instead relied on the public Fink broker to retrieve the corresponding public light-curve data. Hence, the list of *objectId* was passed to the Fink API to collect all associated ZTF light curves (*ZTF-r* and *ZTF-g* filters only), including detections ( $S/N > 5$ ) and bad quality points ( $S/N > 3$ ) but excluding upper-limits. In order to gather as many positive events as possible, we complemented this dataset by adding potentially missing SLSN sources reported in the most recent and complete sample papers. Hence, we gathered all available ZTF SLSN-I light curves reported by Gomez et al. (2024), which regroups all SLSN-I from previous sample papers and even provides a python package to access the data<sup>6</sup>. We also added all SLSN-II reported by Pessi et al. (2025). Fritz automatically filters out most known AGN and stellar contaminants. However, excluding these sources removes light-curve behaviors that would otherwise be encountered in real-time classification and could lead to incorrect predictions by the classifier. To account for this, we manually included ~200 such non-transient sources that we found could be misclassified as SLSNe during preliminary tests of NOMAI on the real-time alert stream. These contaminants are mainly non-variable stars produced by faulty difference imaging, high-proper-motion stars, and AGNs with only a few detections spread over several years. Such events may appear either as flat light curves or as very slowly evolving ones, and can therefore be mistaken for SLSNe.

The final intent of NOMAI is to provide real-time classification of transients, and therefore the light curves that are processed rarely display a complete photometric picture of the sources. Indeed, classification is performed as soon as the previously stated criteria are met, even if the source is barely fading, is at peak, or is even still rising. Therefore, the dataset should be representative of this reality and learn all the possible light curve configurations. Each light curve is therefore decomposed into multiple pseudo-alerts by iteratively removing the most recent observation, thereby generating truncated light curves that represent progressively earlier evolutionary stages of the transient. Following this method, a transient with 100 data points can be split into 100 pseudo-alerts (among which a fraction would not pass the quality cuts). We refer to these truncated light curves as pseudo-alerts to distinguish them from native ZTF alerts, which are restricted to a maximum of 30 days of photometric history. Pseudo-alerts, however, are generated from the fully aggregated photometric evolution of each transient (see Section 2.1) and are therefore not restricted in terms of duration.

## 2.4. Labeling

After applying the quality cuts presented in Sections 2.2 and 2.1, we were left with 8849 unique sources. We then proceeded to label them by combining Fritz and TNS labels, which are not always simultaneously available and do not necessarily match. Given that we use a classification scheme from three distinct sources (Fritz, TNS, and sample papers), the attribution of a definitive label is not always trivial. We proceeded using the following priorities:


1. If there exists a label from a published SLSN sample study, we prioritize it. Given that such analysis proposes highly curated samples, we consider this as the most reliable source.

<sup>3</sup> [https://github.com/astrolabsoftware/fink-filters/blob/master/fink\\_filters/ztf/filter\\_transient\\_complete/filter.py](https://github.com/astrolabsoftware/fink-filters/blob/master/fink_filters/ztf/filter_transient_complete/filter.py)

<sup>4</sup> <https://fritz.science/>

<sup>5</sup> <https://github.com/skyportal/kowalski>

<sup>6</sup> <https://github.com/gmzsebastian/SLSNe>



**Fig. 1.** Class distribution of the labeled dataset for the unique sources (left panel) and for the pseudo-alerts generated from them (right panel). The “other” class mostly includes variable stars, novae, and the non-transient contaminants added.

2. If only a TNS or a Fritz label is available, or if both labels are available and in agreement, we use that label.
3. If both TNS and Fritz labels are available but are in minor disagreement, we use the closest common denomination as a final label (e.g., if one is SLSN-I and the other is SLSN-II, the final label will be SLSN).
4. If both TNS and Fritz labels are available and are in major disagreement (e.g., TDE and SLSN-I), the object is labeled as “Ambiguous”
5. Finally, if no label is available, the source is marked as unclassified.

In addition to this labeling procedure, we must account for the fact that some SLSNe may be missed by the original classifications, assuming a peak magnitude definition. Following the method proposed by [Pessi et al. \(2025\)](#), we set an absolute magnitude threshold of  $-19.9$  in any passband, below which a source is labeled as superluminous (note that other similar threshold values could also be used, such as [Gomez et al. 2022](#), using  $-20$ ). Based on this limit, we further refined the labels by gathering from Fritz all available spectroscopic redshifts in order to compute the peak absolute magnitude of the light curves. We computed the distance modulus using the `astropy.cosmology` package ([Astropy Collaboration 2022](#)) with cosmological parameters coming from the Planck 2015 results ([Planck Collaboration XXII 2016](#)):  $H_0 = 67.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ,  $\Omega_0 = 0.308$ ,  $\Omega_\Lambda = 0.692$ . Following the procedure from [Chen et al. \(2023\)](#), we used only the cosmological term  $-2.5 \times \log(1 + z)$  for the K-correction ([Hogg et al. 2002](#)). Finally we corrected for the Milky Way extinction using the `dust_extinction` python package ([Gordon 2024](#)) assuming a [Fitzpatrick \(1999\)](#) reddening law with  $R_v = 3.1$ .

Any object brighter than the threshold that was previously labeled as core-collapse supernova was upgraded to the SLSN category. In many cases this procedure is essential. For example, SLSN-II sources are often initially labeled as SN IIn based on spectroscopic classification, but the labels are not always updated even when experts recognize that the source is photometrically superluminous. In the end, after applying all cuts and labeling procedures, we obtain a definitive label for 5280 unique ZTF sources, among which 225 unique SLSNe. The left panel of Figure 1 shows the class distribution of the different types of sources. It shows the high imbalance of the dataset with SLSNe representing  $\sim 4\%$  of all sources, with a large majority of SNe Ia or core-collapse supernovae. The right panel shows the dataset after converting the full light curves to pseudo-alerts (Section 2.3). Since each individual source produces multiple pseudo-alerts, the total number of events increases by roughly

a factor of twenty. However, longer-duration transients, such as SLSNe, generate more pseudo-alerts than shorter-lived events. It leads to an improvement regarding the dataset imbalance, with SLSN alerts representing  $\sim 12\%$  of all pseudo-alerts.

### 3. Methodology

In this section, we describe the methodology used to construct the classification pipeline. Because machine-learning algorithms require homogeneous inputs, the irregular and noisy nature of astronomical light curves must first be translated into a set of numerical features. Therefore, we begin by describing the feature-extraction procedure applied to each pseudo-alert, before presenting the implementation and optimization of the machine-learning classifier used to identify SLSN candidates.

#### 3.1. Feature extraction

Light curves are noisy, multidimensional, and unevenly sampled objects, making the task to learn their patterns challenging, given that most machine-learning algorithms require homogeneous input for training. Although some neural network architectures such as Recurrent Neural Networks (e.g., [Möller & de Boissière 2020](#); [Burhanudin et al. 2021](#)) or transformers (e.g., [Cabrera-Vives et al. 2024](#); [Moreno-Cartagena et al. 2025](#)) can handle sequential data without major transformation, most light curve classification pipelines have relied on the extraction of physical features (e.g., [Ishida et al. 2018](#); [Gomez et al. 2020](#); [Sánchez-Sáez et al. 2021](#); [Leoni et al. 2022](#); [Llamas Lanza et al. 2026](#)). This method relies on the automatic computation of pre-defined set of numerical values carrying physical information able to separate different types of transients, effectively converting inhomogeneous light curves into fixed sized arrays. The extraction should be constructed such that as much information as possible about the transient light curve is encompassed into a small amount of features.

An effective way to characterize a transient light curve is to fit it with an appropriate mathematical model and use the optimized parameters of that fit as compact, descriptive features. Following this principle, each light curve is fitted using the Rainbow framework ([Russeil et al. 2024](#)), as implemented in the `light-curve` Python package<sup>7</sup>. This approach performs a single simultaneous surface fit across all fluxes and passbands, rather than treating them independently. Features extracted using

<sup>7</sup> <https://github.com/light-curve/light-curve-python>

the Rainbow method have been shown to outperform the standard approach in classification tasks, particularly for SLSNe (Russeil et al. 2024), and it has already been adopted in several classification pipelines (Fraga et al. 2024; Llamas Lanza et al. 2026; Bommireddy, Hemanth et al. 2026).

Rainbow relies on the assumption that the transient emission can be approximated by a blackbody, thereby introducing temperature parameters to be fitted in addition to the standard flux-related parameters. Therefore, two models should be chosen, one to describe the bolometric flux evolution and one to describe the temperature evolution through time. We chose to model the bolometric evolution using the Bazin (Bazin et al. 2009) function, defined as

$$F(t; t_0, A, \tau_{\text{rise}}, \tau_{\text{fall}}) = A \times \frac{e^{-\frac{(t-t_0)}{\tau_{\text{fall}}}}}{1 + e^{\frac{(t-t_0)}{\tau_{\text{rise}}}}}, \quad (1)$$

where  $t_0$  is a reference time parameter,  $A$  is the amplitude,  $\tau_{\text{rise}}$  is the characteristic rise time, and  $\tau_{\text{fall}}$  is the characteristic decay time. The model describes the flux evolution as an exponential rise followed by an exponential decay. It enables the proper description of most transients, but is insufficient to characterize more complex behaviors such as double peaks or a plateau phase during the decay. However, preliminary tests with more complex functions (e.g., additional shape parameters) did not improve the performance of the classifier, while increasing the computational cost.

We model the temperature using a decaying sigmoid function. It provides a first-order approximation of the temperature evolution of a supernova, whose effective temperature decreases with time as the ejecta expand and cool. It is expressed as

$$T(t; t_0, T_{\text{min}}, T_{\text{max}}, t_{\text{color}}) = T_{\text{min}} + \frac{(T_{\text{max}} - T_{\text{min}})}{1 + \exp \frac{(t-t_0)}{t_{\text{color}}}}, \quad (2)$$

where  $t_0$  is a reference time parameter,  $T_{\text{max}}$  is the initial temperature,  $T_{\text{min}}$  is the final temperature, and  $t_{\text{color}}$  is the exponent driving the speed of the cool down phase. Following Russeil et al. (2024), we assume a common  $t_0$  parameter for both the flux and temperature models, thereby reducing the number of free parameters by one. Combining the flux and the temperature component, we obtain a final model consisting of 7 parameters that we can extract as features. The light-curve package used to perform the fit also provides uncertainties associated with each optimized parameter. Hence, we compute additional S/N features by dividing each best-fit parameter by its corresponding uncertainty. Finally, we extract the  $\chi^2$  of the fit, representing how well the model matches the data.

Given the high number of SNe Ia in the dataset, we also perform a SALT2 (Guy et al. 2007) fit using the `sncosmo`<sup>8</sup> python package. SALT2 is an empirical parametric model designed to describe Type Ia supernovae, which are known for their relative homogeneity. The model is characterized by three physical parameters:  $x_0$  an amplitude parameter,  $x_1$ , the light-curve stretch, and  $c$  the color. In addition, `sncosmo` can also treat the redshift  $z$  as a free parameter which gets simultaneously fitted. Finally, the  $\chi^2$  of the fit is also collected and constitutes an additional feature from the SALT2 model.

The numerical convergence of both the Rainbow and SALT2 fits is ensured by the minimum photometric-point requirements described in Section 2.2. However, note that, in the regime where

**Table 1.** Summary of complementary features computed without a data-fitting procedure.

| Features              | Description                                                                                        |
|-----------------------|----------------------------------------------------------------------------------------------------|
| <i>flux_amplitude</i> | Difference between the <i>max</i> and <i>min</i> flux.                                             |
| <i>max_slope</i>      | Maximum flux slope between two consecutive observations.                                           |
| <i>std_flux</i>       | Standard deviation of the flux.                                                                    |
| <i>duration</i>       | Time difference between the last and the first observation.                                        |
| <i>q15</i>            | 15th percentile of the observation epochs.                                                         |
| <i>q85</i>            | 85th percentile of the observation epochs.                                                         |
| <i>skew</i>           | Skewness of the flux. Measures the asymmetry of a distribution.                                    |
| <i>distnr</i>         | Mean angular distance to the nearest object in the reference image. Available in ZTF alert packets |
| <i>E(B - V)</i>       | Milky Way extinction. Retrieved using SkyCoord from the <code>astropy</code> package.              |


the number of observations is close to the number of fitted parameters, the physical significance of the recovered parameters may not always be guaranteed and depends on the information content and temporal distribution of the available observations. Nevertheless, these parameters are used in combination with the other extracted features, allowing the classifier to learn their practical predictive value within the overall classification framework.

In addition to the model-fitting procedures, we compute a set of features directly from the raw light curves, or from the alert packet. These features capture valuable information about the data while remaining computationally inexpensive to derive. A summary of all these additional features is provided in Table 1. Most of them are meant to provide additional description of the data. The features *q15* and *q85* are intended to identify and mitigate cases where a recent transient is associated with a single spurious observation from several years ago—causing it to appear artificially long-lived—or, conversely, where an older transient produces a single detection long after its main event. The *distnr* constitute one way for the model to separate SLSNe from Tidal Disruption Events (TDEs) which occur at the center of their host galaxy.

We derive statistical features, apply both Rainbow and SALT2 fitting procedures, and extract the physically meaningful parameters from these models for each alert. Altogether, this process yields a set of 27 features per alert, as detailed in Appendix (Table A.2). These features were selected from a broader pool of candidate features following an iterative curation process, during which redundant or non-informative quantities were removed. We find that further reducing the feature set degrades the classifier performance. While some features exhibit only limited discriminating power when considered individually, tree-based classifiers are able to exploit nonlinear interactions between features to improve the classification.

Figure 2 displays a projected view of the parameter space for pairs of the 8 most discriminative features. It illustrates that SLSNe occupy characteristic regions of the projected feature space, although substantial overlap with the heterogeneous population of non-SLSN transients remains. This demonstrates that the extracted features capture physically meaningful information that can be exploited by the machine-learning classifier.

<sup>8</sup> <https://sncosmo.readthedocs.io/en/stable/index.html>



**Fig. 2.** Projection of the parameter space along four pairs of the eight most important features, for SLSNe (orange) and the other alerts (blue). Normalized kernel density estimated feature (KDE) distributions are also displayed.

### 3.2. Classifier implementation

We chose to train an XGBoost<sup>9</sup> (Chen & Guestrin 2016) classifier to separate SLSNe from the other sources. It is a highly efficient and scalable machine-learning algorithm that combines multiple weak learners, implemented as decision trees, to form a stronger, more accurate model. To mitigate dataset imbalance, we weighted each label inversely to its frequency, so that SLSNe have a similar impact on the learning process as the other sources.

The output probabilities of the classifier are subsequently refined using an additional physically motivated criterion based on the absolute magnitude of the sources, computed from the publicly available photometric redshifts. This step aims to leverage external information that is readily accessible for a large fraction of sources in wide-field surveys, even for newly discovered transients.

Therefore, for each source, we query the Sloan Digital Sky Survey (SDSS, Blanton et al. 2017) for a photometric redshift,  $z$ , and its associated uncertainty  $\Delta z$ . Photometric redshifts are available for approximately half of the transients in our labeled dataset. If no photometric redshift is available, the probability is left unchanged. Otherwise, an upper limit on the peak absolute magnitude is computed empirically from the brightest observed photometric point of the light curve, following the methodology of Section 2.4, and using  $z - \Delta z$  as the photometric redshift to obtain a conservative estimate. If the upper limit of the Milky Way corrected absolute magnitude computed from the photometric redshift is less bright than  $-19.75$  mag, the probability of the source being a SLSN is set to 0. Note that this threshold is independent from the  $-19.9$  mag used for the labeling procedure. Using the bootstrap evaluation described in Section 4, we find that, among alerts classified as SLSNe by NOMAI, this

**Table 2.** Summary of the hyperparameter search space used for the random grid search.

| Hyperparameter        | Description         | Optimal value |
|-----------------------|---------------------|---------------|
| <i>reg_alpha</i>      | Uniform(0, 100)     | 68.3          |
| <i>reg_lambda</i>     | Uniform(0, 10)      | 0.34          |
| <i>max_depth</i>      | UniformInt(3, 5)    | 5             |
| <i>max_delta_step</i> | UniformInt(0, 2)    | 2             |
| <i>learning_rate</i>  | Uniform(0.01, 0.31) | 0.26          |
| <i>n_estimators</i>   | Fixed               | 100           |

**Notes.** The right column shows the optimized values used to train the final model.

empirical cut removes 5% of the non-SLSN contaminants while rejecting only 0.3% of true SLSN alerts. Therefore, this empirical cut reduces the false-positive rate with a negligible impact on SLSN completeness.

We then performed a hyperparameter optimization procedure to select the most effective model. We performed 1000 training iterations, sampling the XGBoost hyperparameters from a random grid, and we selected the set that maximized the F1-score defined as

$$F_1 = 2 \times \frac{\text{purity} \times \text{completeness}}{\text{purity} + \text{completeness}}. \quad (3)$$


The hyperparameters optimized are *reg\_alpha* the L1 regularization term, *reg\_lambda* the L2 regularization term, *max\_depth* the maximum depth of a tree, *max\_delta\_step* the maximum delta step for each tree’s weight, and the learning rate. The precise sampling distributions and the final optimal values are presented in Table 2.

The number of trees, *n\_estimators*, was manually set to 100, as increasing the value did not lead to any meaningful improvement but reduced the computation performances. We also emphasize that the *max\_depth* sampling distribution is intentionally kept to small values in order to improve the generalization performances of the final classifier. Since the SLSN class is not yet fully characterized, we favor simpler and more general models to reduce the risk of excluding atypical or previously unseen behaviors that could still correspond to SLSNe.


## 4. Model performances

We assessed the performances of the model by performing 100 iterations of bootstrapping (Efron 1979), which provided more stable results than K-folding methods given the data imbalance. For each of the 100 bootstrap realizations, unique transient identifiers were sampled with replacement from the labeled dataset to form a training sample. Objects not selected in the bootstrap realization constituted the corresponding testing (out-of-bag) sample. Performance metrics were computed on the testing samples and averaged over the 100 bootstrap realizations. Since multiple pseudo-alerts are generated from the same transient, care must be taken to avoid information leakage between the training and testing samples. To this end, the bootstrap sampling was performed at the level of unique transient identifiers rather than individual alerts. Consequently, all pseudo-alerts originating from a given transient were assigned exclusively to either the training or testing sample in a given bootstrap realization, ensuring complete separation between training and evaluation data.

<sup>9</sup> <https://xgboost.readthedocs.io/>



**Fig. 3.** Evolution of the purity and completeness scores as a function of the classification threshold. The uncertainties correspond to the standard deviation of the bootstrapping iterations.




**Fig. 4.** Confusion matrices of 100 iterations of bootstrapping. The left (right) matrix is normalized on completeness (purity). The main value of each cell corresponds to the mean performance of the bootstrapping procedure, while the value displayed below corresponds to the standard deviation. The original unnormalized confusion matrix is presented in Appendix (Table A.1).

The final metrics however, were computed on every pseudo-alert independently, even when multiple alerts originate from the same transient.

Figure 3 displays the evolution of the scoring metrics depending on the classification threshold. As it increases from 0.1 to 0.9, the completeness decreases from 80 to 25%, whereas the purity increases from 44 to 83%, illustrating the trade-off between sensitivity and contamination. A threshold value of  $\sim 0.35$  maximizes the final F1-score (Equation (3)), therefore any alert receiving a score higher than this is classified as a SLSN.

Figure 4 presents the completeness and purity performances of NOMAI using the optimal threshold. Despite the high imbalance, it reaches 58% purity and 66% completeness in the task of classifying SLSNe. Identifying two out of three SLSNe using only light curves, and without access to spectroscopic redshifts, represents a significant achievement.

This is particularly noteworthy given that the distinction between SLSNe-II and SNe IIn remains debated, with both classes potentially forming a continuum rather than being intrinsically distinct (Hiramatsu et al. 2026). This is reflected in the false positives, since approximately one third of the misclassifications belong to the SN IIn class. Another third arise from SNe II, which are known for their extended plateau phase during decline and can therefore mimic the characteristically slow




**Fig. 5.** Collection of the 15 objects with a score closest to the classification threshold. The Rainbow model fitted to the data is also displayed.

evolution of SLSNe. The remaining major source of contamination is AGNs, which account for 13% of the alerts incorrectly classified as SLSNe. Their stochastic variability can randomly produce luminous, slowly evolving light curves resembling those of SLSNe. Although the pre-filtering procedure (Section 2.1) removes most AGN contamination, it is statistically unavoidable that some brightening AGN sources temporarily behave like transients and pass the selection cuts.

In order to better understand the limitations of the classifier, it is useful to study alerts with a score close to the classification threshold, which highlights events for which the classifier is the least confident. Figure 5 presents the 15 objects closest to the threshold. The only SLSN located at the decision boundary is a relatively fast evolving but luminous SN IIn, that was reclassified as an SLSN IIn during the labeling procedure (Section 2.4). Besides this event, approximately a third of the borderline cases are SNe Ia, partly reflecting their prevalence in the dataset (Figure 1). However, they also correspond to atypical SNe Ia, characterized by poorly sampled light curves, high Milky Way extinction, a missing rising phase, or luminosities at the upper end of the SN Ia distribution. The remaining events near the decision threshold are predominantly slow-evolving core-collapse supernovae, in particular SNe II (1/3 of the events).

Figure 6 presents the feature importance analysis. It shows the 10 most informative features, thereby providing insight into




**Fig. 6.** Normalized feature importance of the top ten most informative features of the classifier, averaged over the 100 bootstrapping iterations. The uncertainty corresponding to the standard deviation of the feature importance over the bootstrapping procedure.

the classifier’s decision-making process. By far the most influential feature is  $x_1$ , the stretch parameter of the SALT2 model. This dominance reflects the dataset imbalance, with SNe Ia being by far the most common unique sources. Nevertheless,  $x_1$ , together with the  $\chi^2$  of the SALT2 fit (5th feature), highlights the effectiveness of SALT2 in rapidly identifying SN Ia light curves. The Rainbow fit also contributes essential information to the classification, particularly through the *rise\_time* and the fit  $\chi^2$  (2nd and 3rd features), but also to a lesser extent with the temperature and its evolution (7th and 8th features).

Finally, an important property of a real-time classifier is its stability across predictions. Users must be confident that a source identified as a SLSN candidate is unlikely to rapidly change classification as new observations become available. To quantify this behavior, we define the stability of a SLSN classification as the fraction of subsequent alerts that remain classified as SLSNe after the first positive NOMAI classification. Figure 7 shows the cumulative distribution of this quantity for all SLSNe across the 100 bootstrap realizations. The median stability reaches 100%, meaning that more than half of the SLSNe always remain correctly classified once the first positive classification is obtained. The mean stability is lower, at 81%, due to a small number of borderline cases located near the classification threshold that occasionally switch labels as additional observations become available. Overall, these results demonstrate that NOMAI is stable, and that classification reversals are limited to a small minority of marginal cases.

## 5. Fink implementation

NOMAI is efficient at identifying SLSNe in the labeled dataset. However, for it to be useful for the scientific community, it should be able to run continuously on the ZTF data stream, and produce public classifications that can be easily retrieved. In order to meet these requirements, we incorporated the classifier into the Fink ecosystem in the form of a science module. Science modules are pipelines, built directly by the scientific community, able to process ZTF alerts and enrich them with extra information. In our case, the SLSN classification pipeline can ingest the data stream nightly and add NOMAI’s score as extra information to the alerts. The module is integrated into a



**Fig. 7.** Cumulative distribution of the SLSN label stability after the first positive NOMAI classification. The stability is defined as the percentage of all subsequent alerts that remained classified as SLSNe.

distributed computing infrastructure relying on Apache Spark (Zaharia et al. 2016) to efficiently perform all steps (transient filtering, light curve aggregation, feature extraction, and classification) on distributed chunks of data coming in real-time. The complete implementation is publicly available on Fink<sup>10</sup>. It should be stated that the performance scores on the labeled sample may not directly translate to the live ZTF stream. Indeed, the labeled dataset is inherently biased toward specific classes and brighter sources, and therefore does not fully represent the statistical properties of the raw alert stream. Nevertheless, the real-time evaluation presented in the following section demonstrates that the classifier remains effective when applied to the live ZTF alert stream, despite the distributional differences with the labeled dataset.


Since the completion of this work, NOMAI predictions have been integrated into the Fink science portal<sup>11</sup> and can therefore be accessed directly using the class “SLSN candidate.” This allows users to quickly query and inspect SLSN candidates from the public ZTF data stream directly from a web interface in the same way as other Fink science modules. For operational follow-up, we additionally maintain a dedicated Slack bot that automatically posts all alerts classified as SLSN candidates each night to dedicated channels. As shown in Figure 8, candidates are reported with multiple contextual information, facilitating fast decision making for experts looking to the sub-stream daily.

NOMAI has been running live since 18 December 2025, and in what follows we discuss its real-time performance up to 18 February 2026, corresponding to two months of operation. During this period, minor corrections were applied, leading to slight adjustments toward the final version presented in this paper. However, these modifications only marginally improved the performance metrics on the labeled sample, and we therefore consider the results to be consistent throughout this interval.

Evaluating performance metrics on the live alert stream is challenging, as most sources remain unclassified. Despite this limitation, we attempt to provide informative indicators to assess the classifier’s behavior. To estimate the model purity, we collected all sources classified by the model as SLSNe during the two months of operation. A total of 83 unique sources were identified, among which 55 have received a TNS classification. For the remaining objects, we inspected their photometry and adopted a simple classification scheme: “True positives”

<sup>10</sup> [https://github.com/astrolabsoftware/fink-science/tree/master/fink\\_science/ztf/superluminous](https://github.com/astrolabsoftware/fink-science/tree/master/fink_science/ztf/superluminous)


<sup>11</sup> <https://ztf.fink-portal.org/>



**Fig. 8.** Screenshot of the Slack bot nightly reporting SLSN candidates. It includes, from top to bottom: the TNS classification; a link to the Fink webpage (public); a link to the Fritz webpage (private); the SLSN score from NOMAI; the Milky Way extinction; the SDSS photometric redshift, if available; the peak absolute magnitude based on the SDSS photometric redshift, if available; the alert stamp; and the full ZTF light curve of the source.

correspond to TNS-classified SLSNe, or CCSNe exceeding the absolute-magnitude threshold; “Potential-SLSN” are sources displaying clear SLSN characteristics, such as long duration, blue color, a faint host, or a high absolute magnitude inferred from the photometric redshift; “Compatible” refers to transient sources that could be SLSNe but do not exhibit most of their typical features; and “False positives” correspond to non-SLSN TNS-classified sources, or events for which the photometry alone clearly rules out the SLSN scenario. Figure 9 summarizes the classification of the 83 candidates. Although a definitive purity cannot be established, the results broadly match the performance observed on the labeled dataset. A total of 39 sources can be considered correct identifications, including 22 confirmed in TNS. However, a significant fraction of misclassifications remains, with 32 false positives. These mainly correspond to AGNs or cases of poor photometry leading to artificially flat light curves. The model also identified four SNe IIn which, although they do not reach the absolute magnitude threshold, remain of interest for the SLSN community, as they could provide insight into the possible continuum between SNe IIn and SLSNe-II.

We also attempted to estimate the completeness of NOMAI by collecting all active TNS-classified sources that triggered a detection during the two-month period (excluding events used for training the model). After verifying that their light curves satisfy the minimal requirements of the classifier (Sections 2.1 and 2.2), a total of 279 labeled sources were gathered. Among these, 24 objects were classified as SLSNe, or as CCSNe with absolute magnitudes exceeding the adopted threshold. As shown in Figure 9, the classifier correctly identifies 22 (92%) of them, suggesting a high level of completeness. However, as indicated by the “potential-SLSN” candidates, TNS-classified events likely represent only a fraction of all SLSNe present in the stream, and it is impossible to determine how many may have been missed by human classification. Consequently, this value should be regarded as an upper limit on the classifier’s completeness.



**Fig. 9.** Classification of the 83 SLSN candidates reported by the model during two months of live operation. Candidates are divided into categories based on TNS classifications (True positives, SNe IIn, False positives) and photometric inspection (Potential-SLSNe, Compatible, False positives). The names of the sources are detailed in Appendix (Table A.1).

## 6. Discussion


NOMAI has demonstrated robust performance both on the labeled dataset and under real-time operating conditions. It has been running continuously since 18 December 2025 and is already used by experts as a recommendation system for further analysis. The daily scans and automated summaries have contributed to the identification of previously unclassified SLSN candidates. Some examples are shown in Figure 10, where several sources were flagged by NOMAI as promising SLSN candidates prior to attracting expert attention and before their eventual spectroscopic confirmation, illustrating its ability to spread noteworthy events at early stages. The figure also includes a compelling candidate (bottom-right panel) that, despite being identified and shared by NOMAI, did not receive spectroscopic follow-up, highlighting a limitation in current workflows where promising events can be missed. Such cases are common and emphasize the importance of classification tools in prioritizing and increasing the visibility of candidates that might otherwise go unnoticed.

While the current methodology provides a solid foundation, it also presents limitations that could be improved in the future. In this section we discuss these limitations and possible developments, and we compare it to the current landscape of existing SLSN classifiers.

### 6.1. Limitation and future development

As with any real-time broker-based classification system, NOMAI is expected to quickly evolve as additional operational experience is accumulated. The results presented in this work therefore provide a reference characterization of the classifier at the time of deployment, upon which future improvements can be built.

One of the main limitations of the present approach concerns the minimum light-curve duration required for reliable classification. In the current implementation, the classifier uses the first 30 days of photometric evolution, which is sufficient to achieve the high completeness reported in this work and therefore does not represent a fundamental limitation of the method itself. In practice, however, earlier classifications would be desirable for survey operations. Rapid identification is important to trigger follow-up observations, in particular spectroscopy obtained close to maximum light. Even if SLSNe are generally long lived sources, requiring a full 30-day baseline may reduce the fraction of objects that can be spectroscopically observed during



**Fig. 10.** Examples SLSN candidate light curves identified by NOMAI. The two top panels and the bottom-left panel show sources that were flagged as promising candidates prior to attracting expert attention and were later spectroscopically confirmed as SLSNe. The bottom-right panel shows a promising candidate identified and shared by NOMAI that did not receive spectroscopic follow-up. The dashed vertical line indicates the first alert classified as SLSN by NOMAI.

their most informative phases. This requirement also potentially biases the classifier toward slowly evolving SLSNe, neglecting faster ones. Future work could focus on reducing this minimum duration while maintaining satisfactory performances.

A second limitation arises from the limited use of host-galaxy information in the current feature set. The classifier relies primarily on the transient light curve and only indirectly includes environmental information through the alert feature describing the angular distance to the nearest reference object, and the SDSS photometric redshift if available. While these numbers are very informative, they provide no information about the host properties, meaning that morphology, brightness, or local environment are not exploited. Incorporating image-level information (as shown in [Kisley et al. 2023](#); [Sheng et al. 2024](#)) or host-galaxy features (such as [Gomez et al. 2020](#)) derived directly from survey images could therefore provide additional discriminatory power and improve classification performance.

These developments are particularly relevant to process the LSST public alert stream, which delivers an unprecedented volume of transient detections. In practice, adapting NOMAI from ZTF to LSST should remain straightforward from a technical perspective, since the feature extraction relies on light-curve fits that require few observations and uses Rainbow, which is intrinsically designed to handle multiband photometry. However, the main challenge lies in constructing an informative and representative training dataset.

## 6.2. Comparison to other classifiers

As discussed in the introduction, NOMAI complements the various SLSN classification methods available in the literature. It is therefore natural to compare its performance with similar implementations, such as NEEDLE or FLEET, which both aim to identify SLSNe directly from the ZTF data stream without relying on spectroscopic redshifts. In this context, FLEET

2.0 ([Gomez et al. 2023](#)), the most recent version of the model, reports approximately 50 (40)% purity and 50 (30)% completeness when classifying light curves at late (early) phases using a classification threshold of  $P(\text{SLSN}) > 0.5$ . These values are broadly similar to the metrics obtained in this work, namely 58% purity and 66% completeness. In contrast, NEEDLE ([Sheng et al. 2024](#)) reports higher performance, with an average purity of 84% and completeness of 75%, suggesting that the detection and reference image cutouts are highly informative for classification. However, it remains difficult to determine to what extent these differences reflect intrinsically better performance or simply variations in the evaluation conditions.

More generally, direct comparisons between different classifiers should be interpreted with a lot of caution. Reported performance metrics depend strongly on the characteristics of the datasets used for training and evaluation, including the selection criteria, preprocessing steps, class distributions, and the survey from which the data are drawn. Consequently, without a dedicated benchmarking framework, any strict quantitative comparison between models is impossible. At the same time, the diversity of methodologies used by existing classifiers suggests that they may capture complementary information from the data. In this context, maintaining multiple independent classification approaches may be beneficial, as combining their outputs could ultimately provide a more robust identification of SLSN candidates.

## 7. Conclusion

In this work, we presented the methodology of NOMAI, running on the Fink broker. It relies on a strong filtering step that removes most non-transient events, allowing the efficient retrieval of full light curves for the sources. These light curves are subsequently processed through feature extraction using simple statistical quantities as well as model-fitting procedures based

on SALT2 and Rainbow. These features are highly informative and summarize the inhomogeneous light curves into fixed-size arrays.

The dataset is composed of 5280 unique sources, among which 225 are SLSNe, carefully labeled by combining public classifications from TNS and private labels from Fritz, along with a reclassification of some sources based on their absolute magnitudes. After converting the dataset to alert format, an XGBoost classifier is trained on these features to separate SLSNe from the rest of the ZTF sources. It reaches 66% completeness and 58% purity on the labeled sample, making it, a priori, a powerful tool for analyzing the ZTF alert stream. However, these metrics are only valid for the labeled dataset, which, by the nature of the labeling process, is subject to various biases.

NOMAI has been integrated into the Fink ecosystem and has been processing public alerts every night for the past months. We therefore attempted to estimate its real-world performance by evaluating two months of SLSN classifications. Although exact metrics cannot be computed because of the lack of labels, the purity predicted from the labeled sample appears broadly consistent with the results obtained under real conditions. Regarding completeness, the classifier performs very well, successfully identifying 22 of the 24 active TNS SLSNe during this period. This value, however, does not represent a true completeness, since the total number of SLSNe that triggered alerts during this time interval is unknown; it should therefore be regarded as an upper limit.

Beyond these quantitative metrics, the classifier has been continuously operating within the Fink broker and providing nightly recommendations to the community. The results are automatically shared through dedicated public Slack channels, allowing experts to quickly inspect promising candidates without the need to interact directly with the alert stream. This user-friendly interface has proven effective for scanning the large volume of incoming alerts and highlighting potentially interesting SLSN candidates. In practice, NOMAI has already facilitated the identification and discussion of several events within the community, demonstrating its usefulness as a real-time recommendation tool for transient follow-up.

Many questions remain about SLSNe, since their progenitors, powering mechanisms, and environments are still debated. Increasing the total number of events is essential to improve our understanding of these extreme transients. Tools such as the classifier presented in this work contribute to this effort by enabling the live identification of promising SLSN candidates within large alert streams, facilitating targeted follow-up observations. This need will become even more pressing with the advent of the Legacy Survey of Space and Time conducted by the Vera C. Rubin Observatory, which is expected to discover superluminous supernovae in unprecedented numbers. In this context, automated pipelines capable of identifying these events in real time are not only useful, but necessary for fully exploiting the scientific potential of next-generation time-domain surveys.

**Acknowledgments.** This project is funded by the European Union (ERC, project number 101042299, TransPIre). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. Based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under Grants No. AST-1440341 and AST-2034437 and a collaboration including current partners Caltech, IPAC, the Weizmann Institute of Science, the Oskar Klein Center

at Stockholm University, the University of Maryland, Deutsches Elektronen-Synchrotron and Humboldt University, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, Trinity College Dublin, Lawrence Livermore National Laboratories, IN2P3, University of Warwick, Ruhr University Bochum, Northwestern University and former partners the University of Washington, Los Alamos National Laboratories, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW. The Gordon and Betty Moore Foundation, through both the Data-Driven Investigator Program and a dedicated grant, provided critical funding for SkyPortal.

## References

- Angus, C. R., Smith, M., Sullivan, M., et al. 2019, *MNRAS*, **487**, 2215  
 Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, *ApJ*, **935**, 167  
 Barkat, Z., Rakavy, G., & Sack, N. 1967, *Phys. Rev. Lett.*, **18**, 379  
 Bazin, G., Palanque-Delabrouille, N., Rich, J., et al. 2009, *A&A*, **499**, 653  
 Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002  
 Blanton, M. R., Bershad, M. A., Abolfathi, B., et al. 2017, *AJ*, **154**, 28  
 Bommireddy, H., Förster, F., McMahon, I., et al. 2026, *A&A*, **708**, A325  
 Burhanudin, U. F., & Maund, J. R. 2023, *MNRAS*, **521**, 1601  
 Burhanudin, U. F., Maund, J. R., Killestein, T., et al. 2021, *MNRAS*, **505**, 4345  
 Cabrera-Vives, G., Moreno-Cartagena, D., Astorga, N., et al. 2024, *A&A*, **689**, A289  
 Chatzopoulos, E., Wheeler, J. C., & Vinko, J. 2012, *ApJ*, **746**, 121  
 Chen, T., & Guestrin, C. 2016, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16* (New York, NY, USA: Association for Computing Machinery), 785  
 Chen, Z. H., Yan, L., Kangas, T., et al. 2023, *ApJ*, **943**, 41  
 Chen, Z. H., Yan, L., Kangas, T., et al. 2023, *ApJ*, **943**, 42  
 Coughlin, M. W., Bloom, J. S., Nir, G., et al. 2023, *ApJSS*, **267**, 31  
 De Cia, A., Gal-Yam, A., Rubin, A., et al. 2018, *ApJ*, **860**, 100  
 de Soto, K. M., Villar, V. A., Berger, E., et al. 2024, *ApJ*, **974**, 169  
 Dekany, R., Smith, R. M., Riddle, R., et al. 2020, *PASP*, **132**, 038001  
 Dexter, J., & Kasen, D. 2013, *ApJ*, **772**, 30  
 Efron, B. 1979, *Ann. Statist.*, **7**, 1  
 Fitzpatrick, E. L. 1999, *PASP*, **111**, 63  
 Fowler, W. A., & Hoyle, F. 1964, *ApJ*, **9**, 201  
 Fraga, B. M. O., Bom, C. R., Santos, A., et al. 2024, *A&A*, **692**, A208  
 Fremling, C., Miller, A. A., Sharma, Y., et al. 2020, *ApJ*, **895**, 32  
 Frohmaier, C., Angus, C. R., Vincenzi, M., et al. 2020, *MNRAS*, **500**, 5142  
 Gal-Yam, A. 2012, *Science*, **337**, 927  
 Gal-Yam, A. 2019, *ARA&A*, **57**, 305  
 Gomez, S., Berger, E., Blanchard, P. K., et al. 2020, *ApJ*, **904**, 74  
 Gomez, S., Berger, E., Nicholl, M., Blanchard, P. K., & Hosseinzadeh, G. 2022, *ApJ*, **941**, 107  
 Gomez, S., Berger, E., Blanchard, P. K., et al. 2023, *ApJ*, **949**, 114  
 Gomez, S., Nicholl, M., Berger, E., et al. 2024, *MNRAS*, **535**, 471  
 Gomez, S., Nicholl, M., Berger, E., et al. 2024, *MNRAS*, **535**, 471  
 Gordon, K. D. 2024, *J. Open Source Softw.*, **9**, 7023  
 Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, **131**, 078001  
 Guy, J., Astier, P., Baumont, S., et al. 2007, *A&A*, **466**, 11  
 Heger, A., & Woosley, S. E. 2002, *ApJ*, **567**, 532  
 Hiramatsu, D., Berger, E., Gomez, S., et al. 2026, *ApJ*, **1005**, 82  
 Hložek, R., Malz, A. I., Ponder, K. A., et al. 2023, *ApJ*, **267**, 25  
 Hogg, D. W., Baldry, I. K., Blanton, M. R., & Eisenstein, D. J. 2002, arXiv e-prints [arXiv:astro-ph/0210394]  
 Hsu, B., Hosseinzadeh, G., Villar, V. A., & Berger, E. 2022, *ApJ*, **937**, 13  
 Ishida, E. E. O., Beck, R., González-Gaitán, S., et al. 2018, *MNRAS*, **483**, 2  
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111  
 Jegou du Laz, T., Coughlin, M. W., Bachant, P., et al. 2026, *PASP*, **138**, 064503  
 Kasen, D., & Bildsten, L. 2010, *ApJ*, **717**, 245  
 Kisley, M., Qin, Y.-J., Zabludoff, A., Barnard, K., & Ko, C.-L. 2023, *ApJ*, **942**, 29  
 Knop, R., & ELASiTIC Team 2023, in *American Astronomical Society Meeting Abstracts*, **55**, American Astronomical Society Meeting Abstracts, 117.02  
 Leoni, M., Ishida, E. E. O., Peloton, J., & Möller, A. 2022, *A&A*, **663**, A13  
 Li, M., Wu, C., Kang, Z., et al. 2024, *Universe*, **10**, 31  
 Llamas Lanza, M., Karpov, S., Russeil, E., et al. 2026, *A&A*, **705**, A144  
 Lunnan, R., Chornock, R., Berger, E., et al. 2018, *ApJ*, **852**, 81  
 Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003  
 Matheson, T., Stubbs, C., Wolf, N., et al. 2021, *ApJ*, **161**, 107  
 Möller, A., & de Boissière, T. 2020, *MNRAS*, **491**, 4277  
 Möller, A., Peloton, J., Ishida, E. E. O., et al. 2021, *MNRAS*, **501**, 3272  
 Moreno-Cartagena, D., Protopapas, P., Cabrera-Vives, G., et al. 2025, *A&A*, **703**, A41

- Moriya, T. J., Nicholl, M., & Guillochon, J. 2018, [ApJ](#), **867**, 113
- Mukhopadhyay, M., Kimura, S. S., Vurm, I., & Metzger, B. D. 2026, [ApJ](#), **1006**, 24
- Nordin, J., Brinnel, V., van Santen, J., et al. 2019, [A&A](#), **631**, A147
- Ostriker, J. P., & Gunn, J. E. 1971, [ApJ](#), **164**, L95
- Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, [PASP](#), **131**, 018001
- Perley, D. A., Fremling, C., Sollerman, J., et al. 2020, [ApJ](#), **904**, 35
- Pessi, P. J., Lunnan, R., Sollerman, J., et al. 2025, [A&A](#), **695**, A142
- Planck Collaboration XXII. 2016, [A&A](#), **594**, A1
- Prajs, S., Sullivan, M., Smith, M., et al. 2017, [MNRAS](#), **464**, 3568
- Qu, H., & Sako, M. 2022, [AJ](#), **163**, 57
- Quimby, R. M., Kulkarni, S. R., Kasliwal, M. M., et al. 2011, [Nature](#), **474**, 487
- Rakavy, G., & Shaviv, G. 1967, [ApJ](#), **148**, 803
- Russeil, E., Malanchev, K. L., Aleo, P. D., et al. 2024, [A&A](#), **683**, A251
- Sánchez-Sáez, P., Reyes, I., Valenzuela, C., et al. 2021, [AJ](#), **161**, 141
- Sheng, X., Nicholl, M., Smith, K. W., et al. 2024, [MNRAS](#), **531**, 2474
- Sorokina, E., Blinnikov, S., Nomoto, K., Quimby, R., & Tolstov, A. 2016, [ApJ](#), **829**, 17
- Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, [ApJ](#), **750**, 99
- van der Walt, S. J., Crellin-Quick, A., & Bloom, J. S. 2019, [J. Open Source Softw.](#), **4**
- Villar, V. A., Berger, E., Miller, G., et al. 2019, [ApJ](#), **884**, 83
- Villar, V. A., Hosseinzadeh, G., Berger, E., et al. 2020, [ApJ](#), **905**, 94
- Vurm, I., & Metzger, B. D. 2021, [ApJ](#), **917**, 77
- Wheeler, J. C., Chatzopoulos, E., Vinkó, J., & Tuminello, R. 2017, [ApJ](#), **851**, L14
- Williams, R. D., Francis, G. P., Lawrence, A., et al. 2024, [RAS Tech. Instrum.](#), **3**, 362
- Woosley, S. E. 2010, [ApJ](#), **719**, L204
- Woosley, S. E., Heger, A., & Weaver, T. A. 2002, [Rev. Mod. Phys.](#), **74**, 1015
- Zaharia, M., Xin, R. S., Wendell, P., et al. 2016, [Commun. ACM](#), **59**, 56

## Appendix A: Supplementary material

Table A.1: List of inspected ZTF objects grouped by classification outcome.

| Class           | ZTF objectIds                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| True positives  | ZTF24abyaowk, ZTF25abxjtkv, ZTF25aaggkxk, ZTF25acibvls, ZTF25aauczue, ZTF25aatmfzy, ZTF25achwsgj, ZTF25abwvblg, ZTF25acemaph, ZTF25abxlcab, ZTF25acfybcc, ZTF25abyrwhz, ZTF25abzyzac, ZTF25acgvbov, ZTF25accmopr, ZTF25achwzws, ZTF25abmhchv, ZTF25abxfsza, ZTF25acfbplr, ZTF25acgorke, ZTF25acelgwl, ZTF26aaakmmc                                                                                    |
| Potential-SLSN  | ZTF25acelkdo, ZTF25acgtgdl, ZTF25achvfus, ZTF25acicrjr, ZTF18aaweguw, ZTF25acjsejg, ZTF25acasutk, ZTF25acekwdl, ZTF25achctne, ZTF25acehskf, ZTF25acgojmh, ZTF25aclldhis, ZTF25acbjqvq, ZTF25achbrei, ZTF25acghjar, ZTF25abxpson, ZTF26aaadhkv                                                                                                                                                         |
| Compatible      | ZTF25accfued, ZTF25aceeehm, ZTF25acendyn, ZTF25acjedts, ZTF25acgaku, ZTF25ackkmlt, ZTF25achqjdq, ZTF25acktbjh, ZTF26aaaqepm, ZTF26aabdhbl, ZTF25acfpqje, ZTF26aaaqdbi                                                                                                                                                                                                                                 |
| SNIIn           | ZTF25aadhpuz, ZTF25abxqslz, ZTF25acgggyl, ZTF25acggpbd                                                                                                                                                                                                                                                                                                                                                |
| False positives | ZTF18aaiuwqh, ZTF25abubydw, ZTF25abuovxv, ZTF18aczazkn, ZTF23aagapcz, ZTF25abxqrvf, ZTF25aahmacs, ZTF25abzivnk, ZTF25acgtjkh, ZTF24aarwhrh, ZTF18aaxqitv, ZTF25acinfidh, ZTF25acfqjb, ZTF25acjqzbs, ZTF25abwtoay, ZTF25aadoqxb, ZTF18aaisaqo, ZTF25acemsej, ZTF20aahbzae, ZTF25abwbkdq, ZTF25abgviek, ZTF25abtgdoi, ZTF25abkeuac, ZTF25aceqoyb, ZTF25acefgbd, ZTF21acqhbdp, ZTF25ackarzf, ZTF26aaarpj |

**Notes.** Potential-SLSN were purely classified based on present public photometric data, therefore they might be confirmed or rejected in the future.




Fig. A.1: Unnormalized confusion matrix corresponding to the classifier predictions pooled across all bootstrapping iterations. Normalized completeness and purity matrices are shown in Figure 4.

Table A.2: Summary of all features used by the model.

| Rainbow features             |                                                                                                     |
|------------------------------|-----------------------------------------------------------------------------------------------------|
| <i>amplitude</i>             | Amplitude parameter of the Bazin flux model.                                                        |
| <i>rise_time</i>             | Characteristic rise timescale ( $\tau_{\text{rise}}$ ) of the Bazin model.                          |
| <i>fall_time</i>             | Characteristic decay timescale ( $\tau_{\text{fall}}$ ) of the Bazin model.                         |
| $T_{\text{min}}$             | Late-time temperature from the temperature model.                                                   |
| $T_{\text{max}}$             | Initial temperature from the temperature model.                                                     |
| $t_{\text{color}}$           | Cooling timescale.                                                                                  |
| <i>chi2_rainbow</i>          | $\chi^2$ value of the Rainbow model fit.                                                            |
| <i>snr_amplitude</i>         | S/N of the amplitude parameter.                                                                     |
| <i>snr_rise_time</i>         | S/N of the rise time parameter.                                                                     |
| <i>snr_fall_time</i>         | S/N of the fall time parameter.                                                                     |
| <i>snr_T<sub>min</sub></i>   | S/N of the minimum temperature.                                                                     |
| <i>snr_T<sub>max</sub></i>   | S/N of the maximum temperature.                                                                     |
| <i>snr_t<sub>color</sub></i> | S/N of the cooling timescale.                                                                       |
| SALT2 features               |                                                                                                     |
| $x_0$                        | Amplitude parameter of the SALT2 model.                                                             |
| $x_1$                        | Light-curve stretch parameter from the SALT2 model.                                                 |
| $c$                          | Color parameter from the SALT2 model.                                                               |
| $z$                          | Redshift estimated from the SALT2 fit.                                                              |
| <i>chi2_salt</i>             | $\chi^2$ value of the SALT2 model fit.                                                              |
| Statistical / Alert features |                                                                                                     |
| <i>flux_amplitude</i>        | Difference between the maximum and minimum flux.                                                    |
| <i>max_slope</i>             | Maximum flux slope between two consecutive observations.                                            |
| <i>std_flux</i>              | Standard deviation of the flux.                                                                     |
| <i>duration</i>              | Time difference between the last and the first observation.                                         |
| $q_{15}$                     | 15th percentile of the observation epochs.                                                          |
| $q_{85}$                     | 85th percentile of the observation epochs.                                                          |
| <i>skew</i>                  | Skewness of the flux distribution, measuring asymmetry.                                             |
| <i>distnr</i>                | Mean angular distance to the nearest object in the reference image. Available in ZTF alert packets. |
| $E(B - V)$                   | Milky Way extinction estimated using SkyCoord from <i>astropy</i> .                                 |